Decidability in extensions of $\mathbb{F}_p((t))$

Benjamin Rigler

St. Peter's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2021

Contents

Acknowledgements

This project was generously supported by EPSRC research studentship award number MATH1314.

There are innumerable friends and family who have supported me throughout this long process, and I am greatly indebted to them, as I am to the many mathematicians who have provided discussion, advice, or inspiration in one form or another. It is not possible to include them all, but in particular I would like to single out the following:

First and foremost is my wonderful supervisor Professor Jochen Koenigsmann, without whose wisdom, patience, understanding, and kind encouragement, the present work would have been impossible. A workshop on definability and decidability problems in number theory, held in October 2016 at Oberwolfach, was especially helpful and allowed me to meet many leading mathematicians in the field, for which opportunity I am truly grateful. I should specifically credit talks by Itay Kaplan, Franz-Viktor Kuhlmann and Raf Cluckers, a brief conversation with Tom Scanlon, and others which I am no doubt forgetting, as providing new perspectives, new ideas or simply moral support. A conference at the University of Mons in January 2017 was similarly encouraging. I'd also like to extend my thanks to the model theory community as a whole, for the warm-hearted and generous atmosphere that I have always encountered within it.

Although care has been taken over proofreading this thesis, undoubtedly some errors will still be present; all such errors remain my own.

Abstract

In this thesis we primarily consider the first-order theory of the local field $\mathbb{F}_p((t))$ and the question of whether it is decidable. To this end we expand the language of valued fields together with a constant symbol for t by adding predicates R_f representing the existence of a root of the multivariable polynomial $f \in \mathbb{F}_p(t)[X_1, ..., X_n]$, hoping in this way to control the behaviour of purely wild extensions of valued fields (which has historically caused problems in this area), and seek to prove a quantifier elimination result for a proposed axiomatisation of $\mathbb{F}_p((t))$ (originally, we believe, due to F-V. Kuhlmann). This proceeds along classical lines, *à la* Shoenfield (used in Macintyre's celebrated proof of quantifier elimination for the p-adic numbers $\mathbb{Q}_p$). We obtain a conditional result, sketching a proof that if a particular technical condition is satisfied by an arbitrary $\aleph_1$-saturated model of our theory, then we will indeed obtain quantifier elimination and thereby (via the existence of an algebraically prime model) decidability of the first-order theory of $\mathbb{F}_p((t))$. However, this technical condition is still largely mysterious. We sketch an (unsuccessful) attempt at an unconditional proof and highlight where difficulties arise.

In Chapter 1 we introduce the area in general, and our question and approach more specifically. Chapter 2 contains necessary preliminaries for the rest of the thesis, on valuation theory, model theory, Galois theory, and a smidgen of algebraic geometry; we have intended to keep these preliminaries as brief as possible. Chapter 3 starts with an explanation of our

approach and contains sections (among others) introducing and analysing our candidate axiomatisation, proceeding to a conditional Shoenfield-style quantifier elimination proof (Theorem 3.4.2) and presenting (in §3.5) a sketch of where difficulties arise when trying an unconditional proof. Chapter 4 introduces several notions which we feel may be helpful in future study of this area, as well as discussing ultraproducts of generalised Laurent series fields in relation to extremality. Chapter 5 includes a self-contained examination of the structure and first-order theory of some distinguished fields extending $\mathbb{F}_p((t))$, this time with p-divisible value group. We conclude with some brief remarks in Chapter 6.

Chapter 1

Introduction

The question of decidability in $\mathbb{F}_p((t))$ has been open since at least the 1960s, and the seminal work of Ax–Kochen and Ershov on the closely analogous field $\mathbb{Q}_p$ of p-adic numbers. The problem, of whether or not there exists a complete recursive axiomatisation of the theory of the valued field $\mathbb{F}_p((t))$ (the field of Laurent series over a finite field with p elements) has remained annoyingly intractable since then, despite many advances in modern model theory and many important results in the area of decidability/undecidability of fields in particular. (For an overview of the area as a whole, see [Koe14]. [Poo08] provides a brief overview focusing mainly on Hilbert's 10th Problem over rings.)

Why is this gap so frustrating? In modern mathematics there is an important link between *local fields*, those which are locally compact with respect to a non-discrete topology, and *global fields*, finite extensions of $\mathbb{Q}$ or $\mathbb{F}_q(t)$. Local-global principles such as the Hasse–Minkowski principle can sometimes reduce problems about global fields — those of particular number-theoretic interest — to local fields, with which it is often more practical to work. (This is touched on in the introduction to [EP05].) A local field is precisely a field which is locally compact with respect to an absolute value, and any such field is isomorphic to $\mathbb{R}$, $\mathbb{C}$, or a finite extension of $\mathbb{Q}_p$ or $\mathbb{F}_p((t))$ for some prime p. (See [Ser79] for a fairly definitive guide to the theory of local fields.) Decidability of $\mathbb{R}$ and $\mathbb{C}$ was already known in the 1930s from the work of Tarski, and Ax–Kochen and independently Ershov answered the question positively for $\mathbb{Q}_p$ and its finite extensions. (See [PR84] for a comprehensive overview on decidability of finite extensions of $\mathbb{Q}_p$.) On the other hand, in terms of global fields, we can go even further — not only do these fields have undecidable first-order theories, but in most cases we even have *existential* undecidability (equivalent to a negative solution to Hilbert's 10th Problem), with the only remaining unknown cases being $\mathbb{Q}$ and (some of) its finite extensions. (See [Koe14] in particular on this question.)

We might therefore hope for a clear distinction, with global fields being existentially undecidable on the one hand, and local fields being decidable on the other — this would give us, from a model-theoretic perspective, a compelling narrative for how local-global principles can turn difficult problems in number theory, formulated over global fields, into much easier and in principle answerable problems over local fields. The one major holdout is finite extensions of $\mathbb{F}_p((t))$ (each of which is isomorphic to $\mathbb{F}_q((t))$ for some q a power of p), so that a proof of decidability for $\mathbb{F}_p((t))$ would largely complete the picture. This has been the motivation behind the current project.

There is a significant amount of previous work on the first-order theory of $\mathbb{F}_p((t))$ on which we hope to build. It has long been known that the existential part of the theory is decidable conditional on resolution of singularities in positive characteristic (see [DS03]); in the past few years this has been proved unconditionally by Anscombe and Fehm [AF16], albeit in a language without a symbol for t, building on several years of progress in the theory of definable henselian valuations since the publication of the seminal [CDLM13] (a useful survey of this progress is given in [FJ17]). Unfortunately the approach taken there cannot extend to the language with t, and to our knowledge the result of Denef and Schoutens has not been improved upon in that particular case.

The model theory of $\mathbb{F}_p((t))$ as a ring has been examined by T. Rohwer in his doctoral thesis [Roh03], in which model-completeness is proved in a certain sense as a module, but unfortunately the techniques used there (largely dealing with the model theory of modules and rings) does not seem to relate to the question of decidability as a field. We know from earlier work that if certain structures are definable in $\mathbb{F}_p((t))$ then we will get an *un*decidable theory: for example, in [DR88] it is proved that a definable section of the valuation map would imply undecidability, and it is clear that if undecidability results for global fields (such as [ES09] or, more recently, [ES17]) could be extended to $\mathbb{F}_p((t))$ then we would get undecidability. And clearly if any undecidable structure is definable or interpretable inside $\mathbb{F}_p((t))$ then we will get an undecidable first-order theory, so that a definite answer to our question might come from e.g. being able to interpret $\langle \mathbb{Z}, +, P \rangle$ for certain sets of primes P, under certain hypotheses, following [BJW93] and subsequent developments, or via interpreting any number of other known undecidable structures.

Another important development in recent years has been the notion of *extremal valued fields*; these are fields for which certain images in the value group of multivariable polynomials over the valuation ring satisfy a sort of extreme value theorem. Such a strong condition happens only rarely, but is satisfied by both $\mathbb{F}_p((t))$ and $\mathbb{Q}_p$, leading us to suspect that there is something important at work here. First introduced by Ershov in [Ers04], a critical mistake in the definition was spotted by Starchenko and subsequently corrected (cf. [AKP12, Remark 3.2]); further developments have included [AKP12] and [AK16], but progress here has been unexpectedly difficult, and there are very obvious questions about extremal fields which despite our best efforts are still unresolved.

One barrier to decidability results in this area is the non-uniqueness of maximal immediate extensions of valued fields of characteristic $p > 0$; this can prove a decisive stumbling block when it comes to standard model-theoretic proof methods along the lines of quantifier elimination or model-completeness. Certain algebraic extensions of positive characteristic valued fields, so-called 'purely wild' extensions, are also poorly understood and pose a definite challenge for model theory in this area. When working in this area one can obtain a surprising amount of model-theoretic knowledge from a good understanding of the Galois theory, in particular how specific subgroups (such as the ramification group) of a valued field's absolute Galois group relate to the valuation on the field. We started by asking the question, what would happen if we were to add new predicates into our language to control these problematic purely wild extensions, and pin down the 'right' complement of the ramification group? The question then becomes one of 'what *else* is needed to guarantee decidability/quantifier elimination/etc?'. After our early efforts proved insufficient, we expanded the range of predicates considered, eventually opening it up past 'predicates for all purely wild extensions' to 'predicates for all $\varnothing$-definable multivariable polynomials' and considering the assumption of existential decidability in this enriched language to see if we could at least arrive at a conditional decidability result.

This has been a partial success — in Chapter 3 we introduce a language $\mathcal{L}_{\mathcal{F}}$ and a recursively axiomatised $\mathcal{L}_{\mathcal{F}}$-theory T, satisfied by $\mathbb{F}_p((t))$, which has algebraically prime model $\mathbb{F}_p(t)^h$ (the henselization of $\mathbb{F}_p(t)$), and we present some results about the models of this theory. §3.3 contains a definition of the *algebraic configuration* of an element $a \in K$ where $K \models T_{\mathcal{F}}$, which we introduce as a way of focusing on the relevant part of the quantifier-free type of a over K for a proof of quantifier elimination. In §3.4 we carry out a conditional Shoenfield-type proof of quantifier elimination for T, culminating in Theorem 3.4.2, giving quantifier elimination for T conditional on the truth of a technical statement about algebraic configurations. §3.5 is a sketch of some of the difficulties in reaching an unconditional proof. The remainder of the chapter concerns models of T with value group precisely $\mathbb{Z}$ and on considerations about existential decidability.

Chapter 4 introduces several other notions which we believe may be fruitfully employed when studying this particular area of valuation theory and model theory, attempting to deploy them to yield further results about models of T, in particular the question of whether the generalised Laurent series fields $\mathbb{F}_p((\Gamma))$ (for $\Gamma \equiv \mathbb{Z}$) are extremal and therefore genuinely models of T (which sadly we have not been able to resolve). We make a number of conjectures in this chapter, all of which have remained annoyingly out of reach.

Chapter 5 is a relatively self-contained discussion on certain extensions of $\mathbb{F}_p((t))$ with dense value group. Several such extensions of $\mathbb{F}_p((t))$ are of mathematical interest, such as its perfect closure and the completion thereof, but in many such cases we have no hope of extremality holding without massively expanding the value group and residue field. As such, we have examined several of these fields via the theory of tame valued fields, which has been developed over several years primarily by F-V. Kuhlmann (e.g. [KPR86], [Kuh90], [Kuh16]). This allows us to say something about the model theory of certain fields with p-divisible value group, as well as some analysis of the structure of certain fields which to our knowledge is not contained elsewhere in the literature.

We conclude with some brief remarks in Chapter 6.

Chapter 2

Preliminaries

In this chapter we present some necessary preliminaries for the rest of the thesis. Familiarity with basic concepts from algebra (such as field theory, finite Galois theory, etc.) are assumed, and we direct the reader to [Lan03] for all such background knowledge.

2.1 Preliminaries on valuation theory

We present the following basics on valuation theory, derived primarily from [EP05] which is our preferred reference on the topic. This section can be skipped over by those familiar with the subject. The interested reader may also consult [Kuh] (currently a work in progress) or [Efr06], which takes a relativised approach.

A ***valuation*** on a field K is a surjective map $v : K \to \Gamma \cup \{\infty\}$, where Γ is an ordered abelian group (written additively), satisfying the rules (for all $x, y \in K$)

(i) $v(x) = \infty \iff x = 0$

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

(iii) $v(x + y) \geq \min\{v(x), v(y)\}$.

We call (K, v) (sometimes written just K when unambiguous) a ***valued field*** and Γ the ***value group*** (usually written vK). The set

$$\mathcal{O}_v := \{x \in K \mid v(x) \geq 0\}$$

(or $\mathcal{O}_K$ or just $\mathcal{O}$ if unambiguous) forms a local ring called the ***valuation ring***, with unique maximal ideal $\mathcal{M}_v = \{x \in K \mid v(x) > 0\}$ and group of units $\mathcal{O}^\times = \{x \in K \mid v(x) = 0\}$. Two valuations on the same field sharing the same valuation ring are called ***equivalent***, and the two structures differ only by an order-preserving isomorphism of the value group, so we will often speak interchangeably of a valuation and its valuation ring. The field $\mathcal{O}/\mathcal{M}$ is called the ***residue field*** and written Kv or $\overline{K}_v$ (dropping the subscript v when unambiguous).

Where L/K is a field extension and K admits a valuation v, there is automatically a valuation w on L which ***extends*** v to L (i.e. $\mathcal{O}_w \cap K = \mathcal{O}_v$), inducing embeddings $vK \to wL$ and $Kv \to Lw$ (and therefore implying $vK \leq wL$ and $Kv \subseteq Lw$). In this case we call $e = e(L/K) := (wL : vK)$ the ***ramification index*** of the extension and $f = f(L/K) := [Lw : Kv]$ the ***inertia degree***. In the event that we are working with multiple valuations, say a valuation ring $\mathcal{O}_K$ on K and $\mathcal{O}_L$ on L, we also use the notation $e(\mathcal{O}_L/\mathcal{O}_K), f(\mathcal{O}_L/\mathcal{O}_K)$. An extension is called ***immediate*** if $e = f = 1$. A field is called ***maximal*** or ***maximal valued*** if it admits no proper immediate extension, and ***algebraically maximal*** if it admits no proper immediate algebraic extension.

Generally an extension (or ***prolongation***) of a valuation v_K to a field extension L will not be unique. For L/K a finite extension, the number of distinct prolongations is bounded by the degree of the field extension $[L : K]$, and where $\mathcal{O}_1, ..., \mathcal{O}_n$ are all the distinct prolongations of v to L, we have the ***fundamental inequality*** $\sum_{i=1}^n e(\mathcal{O}_i/\mathcal{O}_K) f(\mathcal{O}_i/\mathcal{O}_K) \leq [L : K]$. A case of particular interest is where there is only one such prolongation, in which case we can say even more, namely that $[L : K] = e \cdot f \cdot p^\delta$, where p is the characteristic of the residue field (if greater than 0; otherwise set $p = 1$) and δ is a natural number, known as the ***defect*** of the extension (with the extension being called ***defectless*** if $\delta = 0$). A valued field is called ***defectless*** if all its finite extensions are.

A valuation v on K that extends uniquely to the algebraic closure of K is called ***henselian***. We also speak of the valued field itself being henselian and of a pure field being henselian if it admits a nontrivial henselian valuation. We note the following equivalent definitions of a valuation being henselian:

Theorem 2.1.1. [EP05, §4.1] *The following properties of a valued field $(K, \mathcal{O})$ are equivalent:*

(i) *$\mathcal{O}$ extends uniquely to every algebraic extension of K.*

(ii) *$\mathcal{O}$ extends uniquely to the separable closure K^{sep} of K.*

(iii) *Every polynomial $X^n + a_{n-1}X^{n-1} + \ldots + a_0 \in \mathcal{O}[X]$ with $a_{n-1} \notin \mathcal{M}$ and $a_{n-2}, \ldots, a_0 \in \mathcal{M}$ has a root in K.*

(iv) *Every polynomial $X^n + X^{n-1} + a_{n-2}X^{n-2} + \ldots + a_0 \in \mathcal{O}[X]$ with $a_{n-2}, \ldots, a_0 \in \mathcal{M}$ has a root in K.*

(v) *For each $f \in \mathcal{O}[X]$ and $a \in \mathcal{O}$ with $\overline{f}(\overline{a}) = 0$ and $\overline{f}'(\overline{a}) \neq 0$, there exists $\alpha \in \mathcal{O}$ with $f(\alpha) = 0$ and $\overline{\alpha} = \overline{a}$.*

(vi) *For each $f \in \mathcal{O}[X]$ and $a \in \mathcal{O}$ such that $v(f(a)) > 2v(f'(a))$, there exists $\alpha \in \mathcal{O}$ with $f(\alpha) = 0$ and $v(a - \alpha) > v(f'(a))$.*

Here $\overline{*}$ denotes the residue homomorphism $\mathcal{O} \to Kv$ and its canonical extension $\mathcal{O}[X] \to Kv[X]$. Note that our list is not exhaustive and that other equivalent properties are standardly used. We also see that separably closed fields are automatically henselian, and that any algebraic extension of a henselian valued field is also henselian.

Due to the characterisations in terms of polynomials, henselianity can be expressed in first-order terms in any suitable language of valued fields (e.g. one with signature $\{+, ., 0, 1, \mathcal{O}\}$ where $\mathcal{O}$ is a unary predicate interpreted as $\mathcal{O}(x) \Leftrightarrow x \in \mathcal{O}$; see e.g. [Kuh16, p. 21] for such an expression). Given an arbitrary valued field (K, v) we can, by adjoining roots of Hensel polynomials (those appearing in the above lemma), form a smallest henselian field K^h extending K; this is called a ***henselization*** of K and is unique up to isomorphism of valued fields over K. It is "smallest" in the sense that whenever L is a henselian field extending K we have that L contains an isomorphic copy of K^h. This can be expressed as the following universal property:

Theorem 2.1.2. [EP05, Theorem 5.2.2] *Let $(K, \mathcal{O})$ be a valued field and $(K^h, \mathcal{O}^h)$ its henselization. Then $(K^h, \mathcal{O}^h)$ is henselian and, given any valued field extension $(K_1, \mathcal{O}_1)$ of $(K, \mathcal{O})$ which is henselian, there is a uniquely determined K-embedding of valued fields $\iota : (K^h, \mathcal{O}^h) \to (K_1, \mathcal{O}_1)$, i.e. $\iota(\mathcal{O}^h) = \mathcal{O}_1 \cap \iota(K^h)$ and $\iota|_K = \mathrm{id}_K$.*

The henselization of a field is an immediate algebraic extension of valued fields [EP05, Theorem 5.2.5]. In the specific case where vK is an ordered subgroup of $\mathbb{R}$ and v is induced by an absolute value, we have that the standard notion of *completion* also always yields an immediate extension [EP05, Theorem 1.3.4].

We note the main application of the key theorem proved in [Pop10], that "henselian implies large".

Theorem 2.1.3. [Pop10, Theorem 1.1] *Suppose (K, v) is a henselian valued field. Then K is a **large field**; that is, for every smooth K-curve C, if $C(K)$ is nonempty then $C(K)$ is infinite.*

See [Pop96] for more information about large fields and their connection to embedding problems in inverse Galois theory. The notion of ***large field*** will reappear later in connection with extremal valued fields.

2.2 Preliminaries on model theory

Familiarity with basic first-order model theory is assumed. See Marker [Mar06] or [Poi12] for a good introduction, or [Hod93] for a more extensive treatment from which much of this section is taken). We will throughout be working with first-order theories T, not assumed complete, admitting no finite models, in some recursive language $\mathcal{L}$. (We require that $\mathcal{L}$ is recursive as we are considering questions of decidability; it would be unhelpful working in a language with e.g. predicate symbols P_ϕ for every true sentence of arithmetic ϕ if we wished to obtain a decidability result for a theory in $\mathcal{L}$.)

A first-order $\mathcal{L}$-theory T is called ***recursively enumerable*** if there is an algorithm which enumerates its axioms (and nothing else), and ***decidable*** if there is an algorithm that, on input φ a sentence in $\mathcal{L}$, outputs YES if $T \models \varphi$ holds, and outputs NO otherwise. A complete recursively enumerable theory is decidable.

A standard approach to proving the completeness (and therefore decidability) of a countable recursively enumerable theory T is to try and reduce the complexity of arbitrary formulas to something more manageable:

Definition 2.2.1. *A first-order $\mathcal{L}$-theory T has* ***quantifier elimination*** *if, for every $\mathcal{L}$-formula φ with free variables $\overline{x} = (x_1, ..., x_n)$, there exists a quantifier-free $\mathcal{L}$-formula ψ such that $T \models \forall \overline{x}(\varphi(\overline{x})) \leftrightarrow \psi(\overline{x})))$.*

A slightly weaker condition is that of model-completeness:

Definition 2.2.2. *A first-order $\mathcal{L}$-theory T is called* ***model-complete*** *if whenever $\mathcal{A}$, $\mathcal{B}$*

are models of T, every embedding $\mathcal{A} \to \mathcal{B}$ is in fact an elementary embedding $\mathcal{A} \preceq \mathcal{B}$.

We have that model-completeness is equivalent to the property that 'every formula is modulo T equivalent to an $\exists$-formula' (or, equivalently, 'every formula is modulo T equivalent to a $\forall$-formula'), so that quantifier elimination implies model-completeness.

Of especial interest is the case where T has a unique 'smallest' model:

Definition 2.2.3. *$\mathcal{A}$ is called an* ***algebraically prime model*** *of T if $\mathcal{A} \models T$ and, for every model $\mathcal{B}$ of T, we have that $\mathcal{A}$ embeds as a substructure of $\mathcal{B}$.*

(Note that these definitions date back to the 1950s or earlier, some in particular to the work of Abraham Robinson, who simply called these 'prime models'; however, model theorists since then have adopted the convention that a prime model is one which *elementarily* embeds in any model, which is much stronger.)

It will be readily seen that if a theory T admits an algebraically prime model $\mathcal{A}$ and is model-complete, then any model of T is a model of the complete theory $\mathrm{Th}(\mathcal{A})$, and T is therefore complete.

We also note a different way of conceptualising model-completeness. First we need:

Definition 2.2.4. *A model $\mathcal{A}$ of an $\mathcal{L}$-theory T is called* ***existentially closed*** *if, whenever $\exists \overline{x}\ \varphi(\overline{x}, \overline{y})$ is an existential $\mathcal{L}$-formula, $\overline{a}$ is a tuple of elements from A, $\mathcal{B}$ is a model of T such that $\mathcal{A} \subseteq \mathcal{B}$, and we have that $\mathcal{B} \models \exists \overline{x}\ \varphi(\overline{x}, \overline{a})$, then also we have that $\mathcal{A} \models \exists \overline{x}\ \varphi(\overline{x}, \overline{a})$.*

Note that this definition is always relative to an ambient theory T (or, more generally, a class of $\mathcal{L}$-structures $\mathbf{K}$). In general, if $\mathcal{A}$ and $\mathcal{B}$ are arbitrary $\mathcal{L}$-structures satisfying the above condition, we say that $\mathcal{A}$ is ***existentially closed*** in $\mathcal{B}$ (denoted $\mathcal{A} \leq_\exists \mathcal{B}$). Now we can characterise model-completeness in the following terms:

Proposition 2.2.5. *A theory T is model-complete if and only if every model of T is existentially closed.*

We also have the following useful theorem: [Hod93, Theorem 2.5.1]

Proposition 2.2.6. [Tarski–Vaught test.] *Suppose we are given structures $\mathcal{N} \subseteq \mathcal{M}$ with domains N, M respectively, such that for every formula $\phi(x, \overline{y})$, and every tuple $\overline{a}$ from N, if $\mathcal{M} \models \exists x\, \phi(x, \overline{a})$ then there is $b \in N$ such that $\mathcal{M} \models \phi(b, \overline{a})$. Then the embedding $\mathcal{N} \to \mathcal{M}$ is elementary.*

That is to say, if there is some $a \in M$ satisfying some formula ϕ over N, then there must already be some $b \in N$ satisfying ϕ. Note however that here our formulas may be *arbitrary*, not just existential.

Before we are in a position to characterise quantifier elimination more helpfully, we need some additional pure model theory. Our strategy for better understanding the first-order theory of $\mathbb{F}_p((t))$ is however already clear: write down a recursively enumerable set of axioms for T, show that T is model-complete (perhaps via quantifier elimination), and exhibit an algebraically prime model. Conversely, if our theory T is not in fact complete, then pursuing this strategy should guide us towards the place in the proof where existential closedness is not guaranteed, thus shedding light on the problem and on what would remain to be done in order to prove decidability.

2.2.1 Preliminaries on saturated models

We will need to be delicate with definitions here as our theory cannot be assumed complete. First we need to introduce the notion of a type, fundamental in model theory: let $\mathcal{A}$ be a model of T and C some subset of the underlying domain A.

A ***partial*** n***-type over*** C is a set of formulas p in the language $\mathcal{L}_C$, extending $\mathcal{L}$ by adding constant symbols for every element in C, such that there are at most n free variables in p, and for every finite collection of formulas $\Phi_1(\overline{x}), ..., \Phi_m(\overline{x})$ from p we have that $T \not\models \forall\overline{x}\ (\neg\Phi_1(\overline{x}) \vee ... \vee \neg\Phi_m(\overline{x}))$. By Compactness we are thereby guaranteed that for no subset Σ of p, even infinite, do we have that $T \cup \Sigma$ is inconsistent. In other words, p is a collection of formulas which could simultaneously hold of a tuple $\overline{x}$ in a model of T which contains C. We often denote it $p(\overline{x}/C)$ or $p(\overline{x})$.

A ***complete*** n***-type over*** C is a partial n-type p for which, given any $\mathcal{L}_C$-formula $\Phi(\overline{x})$ with at most n free variables, we have that either $\Phi(\overline{x}) \in p$ or $\neg\Phi(\overline{x}) \in p$. Any partial type may be extended to a complete type by Zorn's Lemma.

If $\mathcal{B}$ is a model of T containing C and $p(\overline{x}/C)$ is a type, we say that $p(\overline{x}/C)$ is ***realised*** in $\mathcal{B}$ if there is a tuple $\overline{b}$ of elements from $\mathcal{B}$ such that $\mathcal{B} \models \Phi(\overline{b})$ for every $\Phi \in p(\overline{x}/C)$. We call $\overline{b}$ a ***realisation*** of the type; if there is no realisation of $p(\overline{x}/C)$ in $\mathcal{B}$ then we say that $\mathcal{B}$ ***omits*** the type. If $\overline{b}$ is a tuple from $\mathcal{B}$ then we denote by $\mathrm{tp}(\overline{b}/C)$ the complete type consisting of all $\mathcal{L}_C$-formulas Φ such that $\mathcal{B} \models \Phi(\overline{b})$.

It is easy to see how one may formulate definitions e.g. of partial/complete quantifier-free types, etc., if required, by restricting the sorts of $\mathcal{L}_C$-formula appearing in p.

Our next notion is highly important in model theory, and here we need to take care as it is usually used in the context of complete theories. A κ***-saturated model*** of a first-order theory T, for κ an infinite cardinal, is a model $\mathcal{B}$ of T such that, for every $C \subseteq B$ with $|C| < \kappa$, every complete n-type over C, *with respect to the complete* $\mathcal{L}_C$*-theory* $\mathrm{Th}(\mathcal{B})$, is realised in $\mathcal{B}$. That is, whenever $p(\overline{x}/C)$ is a (complete) collection of $\mathcal{L}_C$-formulas such that

$\mathcal{B} \not\models \forall \overline{x}\ (\neg\Phi_i(\overline{x}) \vee ... \vee \neg\Phi_m(\overline{x}))$ with $\Phi_1, ..., \Phi_m$ in p, then p has a realisation in $\mathcal{B}$.

Remark: In effect, a saturated model is one where 'anything that *can* happen *does* happen'; if there is no inconsistency in a set of formulas simultaneously holding of an element of $\mathcal{B}$, then there must be *something somewhere* in $\mathcal{B}$ which simultaneously satisfies them all. However, $\mathcal{B}$ need not realise all complete types *over* T, if T is not complete. For example, take DLO$^-$ to be the theory of densely linearly ordered sets without specifying whether or not there are endpoints. Then if the $\aleph_0$-saturated structure $\mathcal{B}$ is a model of DLO$^-$ and has an upper endpoint e, then by taking $C = \{e\}$ we see that $\mathcal{B} \models \forall x\,(x = e \vee x < e)$ and so the type $p(x/C) = \{x > e\}$ is not realised. However, there are other saturated models of DLO$^-$ which do not have upper endpoints, e.g. $\mathbb{R}$ is an $\aleph_0$-saturated model.

One may also conceptualise saturated models in the following way: $\mathcal{B}$ is κ-saturated if and only if whenever an n-type over C (with $|C| < \kappa$) is realisable in some *elementary extension* $\mathcal{B}^+ \succeq \mathcal{B}$, it is already realised in $\mathcal{B}$.

Remark: In model theory we sometimes speak loosely of 'saturated' models instead of κ-saturated models for a given κ; this is because we often wish to have a model where 'everything that can happen does happen', and the restriction to $|C| < \kappa$ is only necessary to ensure that the domain of our model does not become a proper class instead of a set. It is also obvious that a model $\mathcal{A}$ of cardinality κ cannot be κ^+-saturated by considering the type $\{x \neq a \mid a \in \mathcal{A}\}$. We cannot assume the existence of saturated models of arbitrary theories T without making assumptions beyond ZFC in our underlying set theory, but common practice in model theory is usually to assume the existence of saturated models and worry about the set-theoretic implications later, if at all. However, as the main part of this thesis requires quite a technical understanding of saturation, we will henceforth adopt the convention that for a model $\mathcal{M}$ to be saturated, this means that it is $||\mathcal{M}||$-saturated.

If T is a **complete** theory with a κ-saturated model of cardinality κ then this model is unique up to isomorphism in that cardinality (via a back-and-forth construction). A κ-saturated model will always be κ-homogeneous (partial embeddings whose domain has cardinality less than κ can be extended to automorphisms of the whole space, and if two tuples share the same complete n-type then we can construct an automorphism sending one to the other) and κ^+-universal (any model of T of cardinality $\leq \kappa$ elementarily embeds inside our saturated model). See [Mar06, §4.3] for details.

We are now in a position to introduce the following test for quantifier elimination, as quoted in [Mac76] (originally due to Shoenfield [Sho71]).

Theorem 2.2.7. *Let T be an $\mathcal{L}$-theory. Then T has quantifier elimination if and only if the following holds: whenever we have $\mathcal{L}$-structures $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{M}_1$, $\mathcal{M}_2$ such that $\mathcal{A}_i \subseteq \mathcal{M}_i \models T$ (for $i = 1, 2$), $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{M}_1$ are countable, $\mathcal{M}_2$ is $\aleph_1$-saturated, and we are provided with an isomorphism $f : \mathcal{A}_1 \cong \mathcal{A}_2$, we can always extend f to an embedding of $\mathcal{M}_1$ into $\mathcal{M}_2$.*

This is a special case of the more general theorem proved by Shoenfield in [Sho71], that quantifier elimination is equivalent to the statement that a partial embedding of any model $\mathcal{M}_1$ into an $||\mathcal{M}_1||^+$-saturated model $\mathcal{M}_2$ can be extended to a full embedding of $\mathcal{M}_1$ into $\mathcal{M}_2$, but we can reduce to the more specific case via standard model-theoretic techniques (namely, using Compactness).

Quick remark: Without wishing to go deeper into pure model theory than necessary, we note that if T is, as hoped, complete and model-complete, then it will also satisfy what is known as the ***strong amalgamation property***, which roughly states that if a model $\mathcal{A}$ can be embedded into two models $\mathcal{B}$, $\mathcal{C}$ then we can find a model $\mathcal{D}$ into which $\mathcal{B}$ and

$\mathcal{C}$ embed in a sufficiently nice way (namely, such that the composition of the embeddings $\mathcal{A} \to \mathcal{B} \to \mathcal{D}$ is the same as the composition of the embeddings $\mathcal{A} \to \mathcal{C} \to \mathcal{D}$, and such that the intersection of the images of the embeddings $\mathcal{B} \to \mathcal{D}$ and $\mathcal{C} \to \mathcal{D}$ is precisely the image of the induced embedding $\mathcal{A} \to \mathcal{D}$). See [Hod93, §7.1] on amalgamation properties. We mention this purely as general background to provide some indication of how models of T might behave in practice.

Remarks on stability, cardinality, etc: Much modern model theory revolves around nice model-theoretic situations (such as stable theories, simple theories, etc.) which are unfortunately not present when considering valued fields. For instance, in any suitable language of valued fields, the first-order theory of a (nontrivially) valued field will interpret an infinite ordered set (namely the value group), therefore will be unstable and have 2^κ models of cardinality κ for every uncountably infinite cardinal κ. (See [Mar06, §5.3], originally due to Shelah; the formula $x^{-1}y \in \mathcal{O}^\times$ will suffice to define a linear order $x < y$ on any set $\{a_0, a_1, ..., a_n, ...\}$ such that $v(a_i) = i$.) Proving that there are $2^{\aleph_0}$ models of cardinality $\aleph_0$ is more involved but can often be done by explicit construction.

Before finishing this section, we note that we also have the following alternative test for quantifier elimination, due, we believe, to van den Dries (see [VDD85, "EQ-Test"] and [VDD88, p. 11]), following A. Robinson and Shoenfield.

Theorem 2.2.8. *An $\mathcal{L}$-theory T has quantifier elimination if and only if it satisfies the following two properties:*

*(i) Every substructure of a model admits a T-**closure**: if $\mathcal{A} \subseteq \mathcal{B}$ and $\mathcal{B} \models T$ then there is a model $\mathcal{A}^*$ of T containing $\mathcal{A}$ which embeds over $\mathcal{A}$ in any model of T containing $\mathcal{A}$.*

(ii) Suppose $\mathcal{M}$, $\mathcal{N}$ are two models of T such that $\mathcal{M} \subsetneq \mathcal{N}$. Then there is $a \in \mathcal{N} \setminus \mathcal{M}$ and

a set of quantifier-free $\mathcal{L}_{\mathcal{M}}$-formulas $\{\phi_i(x) \mid i \in I\}$ realised by a and determining its isomorphism type over $\mathcal{M}$, such that every finite subset can be realised in $\mathcal{M}$.

By general model-theoretic considerations it suffices to prove (i) for finitely-generated structures and (ii) for T-closures of finitely-generated substructures. The canonical example of (ii) ("specializability of selected elements") is the fact that a simple ordered field extension $R(x)$ of a real closed field R is determined up to isomorphism over R by the set $\{r \in R : r < x\}$.

2.2.2 Preliminaries on ultraproducts

This material is primarily drawn from [Poi12, §4.1], but should be contained in any good textbook on model theory (see e.g. [Hod93, §9.5], [Mar06, pp. 63-65]).

Let I be some infinite indexing set. Suppose we are given an ***ultrafilter*** $\mathcal{U}$ on I: this is a proper subset of the powerset of I which is closed under intersection and taking supersets, satisfying (for every $A \subseteq I$) either $A \in \mathcal{U}$ or $I \setminus A \in \mathcal{U}$.

(Spelling this out further, $\mathcal{U} \subseteq \mathcal{P}(I)$ is a family of subsets of I satisfying:

- $\varnothing \notin \mathcal{U}$, $I \in \mathcal{U}$,
- if $X \in \mathcal{U}$ and $Y \in \mathcal{U}$ then $X \cap Y \in \mathcal{U}$,
- if $X \in \mathcal{U}$ and $X \subseteq Y \subseteq I$ then $Y \in \mathcal{U}$,
- for every $X \subseteq I$ either $X \in \mathcal{U}$ or $I \setminus X \in \mathcal{U}$.)

If an ultrafilter is generated by a singleton in $\mathcal{P}(I)$ then we call it ***principal***, otherwise we call it ***non-principal***. It is straightforward to see that if an ultrafilter contains any finite set then it must be principal. We may always assume the existence of non-principal ultrafilters, implied by assuming Axiom of Choice.

Now given an indexing set I and an ultrafilter $\mathcal{U}$ on I, suppose $\mathcal{L}$ is a first-order language and that we are given, for each $i \in I$, some $\mathcal{L}$-structure $\mathcal{A}_i$. Then we define the ***ultraproduct*** $\prod_I \mathcal{A}_i/\mathcal{U}$ (or $\prod_\mathcal{U} \mathcal{A}_i$ if I is unambiguous) in the following way:

- Let A_i be the domain of $\mathcal{A}_i$ for each $i \in I$, and consider the product $A := \prod_I A_i$. We define an equivalence relation $\sim$ on A by stipulating that $(a_i)_{i\in I} \sim (b_i)_{i\in I}$ if and only if the set $\{i \in I \mid \mathcal{A}_i \models a_i \doteq b_i\}$ is in $\mathcal{U}$. The domain of the ultraproduct $\prod_\mathcal{U} \mathcal{A}_i$ is defined to be the quotient $A/\sim$. We will abuse notation and write $(a_i)_{i\in I}$ to mean the equivalence class $[(a_i)_{i\in I}]$, and will write $a(i)$ to mean the image of the sequence $a \in \prod_I \mathcal{A}_i$ under the ith projection map $\pi_i : \prod_I \mathcal{A}_i \to \mathcal{A}_i$

- The interpretation of an n-ary predicate symbol P in $\mathcal{L}$ is given by stipulating that $P(a_1, a_2, ..., a_n)$ holds if and only if $\{i \in I \mid \mathcal{A}_i \models P(a_1(i), a_2(i), ..., a_n(i))\}$ is in $\mathcal{U}$.

- The interpretation of a constant symbol c in $\mathcal{L}$ is simply given by $c^{(\prod_\mathcal{U} \mathcal{A}_i)} := (c^{\mathcal{A}_i})_{i\in I}$.

- Let $f(x_1, x_2, ..., x_n)$ be an n-ary function symbol in $\mathcal{L}$. We define

$$f^{(\prod_\mathcal{U} \mathcal{A}_i)}(a_1, a_2, ..., a_n) := (f^{\mathcal{A}_i}(a_1(i), a_2(i), .., a_n(i)))_{i\in I}.$$

(These are well-defined and independent of the choice of representative of a class $[(a_i)_{i\in I}]$.) A good way to conceptualise this is to think of sets in $\mathcal{U}$ as being 'big sets' or as being 'almost all' of I; viewed in this way the ultraproduct construction is essentially just a product of structures, where we identify elements if they are the same 'almost everywhere', predicates hold of an element if they hold 'almost everywhere', etc. Note that this is entirely uninteresting in the case of principal ultrafilters (where the ultraproduct ends up isomorphic to one of the $\mathcal{A}_i$s), so we will always assume that $\mathcal{U}$ is non-principal.

Ultraproducts are useful for the purposes of model theory in light of the following theorem.

Theorem 2.2.9. [Łoś's Theorem] *Let $\mathcal{U}$ be an ultrafilter on I, $(\mathcal{A}_i)_{i\in I}$ a family of $\mathcal{L}$-structures, and $\phi(x_1, ..., x_n)$ an $\mathcal{L}$-formula. Then for every tuple $\overline{a} := (a_1, ..., a_n) \in \prod_{\mathcal{U}} \mathcal{A}_i$ we have $\prod_{\mathcal{U}} \mathcal{A}_i \models \phi(\overline{a})$ if and only if the set $\{i \in I : \mathcal{A}_i \models \phi(a_1(i), a_2(i), ..., a_n(i))\}$ is in $\mathcal{U}$.*

Thus satisfaction of a formula in an ultraproduct is the same thing as satisfaction 'almost everywhere'. This gives us a construction with a lot of useful applications.

2.2.3 Preliminaries on model theory of ordered abelian groups

Here we outline some necessary model theory on ordered abelian groups, originally treated by A. Robinson & Zakon in [RZ60], which is fairly comprehensive. See e.g. [Con62] for a discussion on the algebra of ordered abelian groups.

An ***ordered abelian group*** is an abelian group $(G, +, 0)$ (usually written additively in this context) equipped with an ordering $<$ compatible with the group operation, i.e. so that for all $x, y, z \in G$ we have that $x < y$ implies $x + z < y + z$. It follows from this that every nonzero element must be of infinite order. We call such a group ***discretely ordered*** (or just ***discrete***) if it has a smallest (nonzero) positive element, and ***densely ordered*** (or just ***dense***) otherwise. We should note that in some classical valuation theory the term "discrete" is sometimes used to mean "isomorphic to $\mathbb{Z}$", but here we are using it in a much wider sense.

If an ordered abelian group is densely ordered then, between any two elements $x < x'$ in G, we can find some y such that $x < y < x'$. We call an element $x \in G$ n-***divisible*** if there is some $y \in G$ such that $y + y + ... + y$ (n times) is equal to x. We call G itself n-

divisible if all its elements are, and we call G ***divisible*** if it is n-divisible for all n. Letting 1 be the smallest positive nonzero element of G (if G is discrete), or some arbitrary positive nonzero element of G (if G is dense), we call G ***regular*** if for any $x \in G$ we have that at least one of the elements x, $x+1$, $x+2$, ..., $x+n-1$ is n-divisible. (This is expressible in first-order terms in the language of ordered groups.)

The regular ordered abelian groups are those which are elementarily equivalent to subgroups of $\mathbb{R}$ in the language $\{+, <, 0\}$ of ordered abelian groups. Concretely, call $H \leq G$ a ***convex subgroup*** if whenever $h_1 < h_2$ are two elements of H and $g \in G$ satisfies $h_1 < g < h_2$, then $g \in H$. An ordered abelian group G is called ***archimedean*** if the only convex subgroups are $\{0\}$ and G; every archimedean group is isomorphic to an ordered subgroup of $\mathbb{R}$, and the regular ordered groups are precisely those which are elementarily equivalent to an archimedean group.

A regular discrete ordered abelian group is called a $\mathbb{Z}$***-group***. The first-order theory of $\mathbb{Z}$-groups is complete, and model-complete in the language of ordered abelian groups supplemented with divisibility predicates D_n (such that $D_n(x)$ holds if and only if x is n-divisible) and a constant symbol for 1. This is precisely the class of ordered abelian groups elementarily equivalent to $\mathbb{Z}$, and any such group G must contain a convex subgroup isomorphic to $\mathbb{Z}$; in this case $G/\mathbb{Z}$ is a divisible ordered abelian group. In the densely ordered case, the theory of a given regular densely ordered abelian group is again model-complete once supplemented with divisibility predicates, and is axiomatised by adding (for each prime p) axioms asserting the index $(G : pG)$ of the subgroup of p-divisible elements inside G.

If $H \leq G$ is a subgroup then we call H ***pure*** in G if, whenever we have $\gamma \in G$ such that $n\gamma \in H$ for some $n \in \mathbb{N}$, there is some $\delta \in H$ such that $n\delta = n\gamma$. That is, H is already 'divisibly closed' inside G, in that if $h \in H$ is n-divisible inside G it is already n-divisible inside H. One sees that in the language of ordered abelian groups with divisibility predicates added, where G and H are considered as models of the theory of ordered abelian groups, H is a pure subgroup of G if and only if there is an embedding of H into G as first-order structures.

We have the following characterisation of $\mathbb{Z}$-groups (see [Con62]): suppose G is an abelian ordered group with Z any convex subgroup of G. Then G is regularly ordered if and only if G/Z is divisible. Thus a $\mathbb{Z}$-group is an extension of $(\mathbb{Z}, +, <, 0)$ by a $\mathbb{Q}$-vector space.

2.2.4 Ax–Kochen–Ershov principles

While implicit elsewhere, in this section we borrow Franz-Viktor Kuhlmann's formulation of what it means for a class of valued fields to satisfy an Ax–Kochen–Ershov principle (see e.g. [Kuh06, p. 8], [Kuh16, pp. 3-4]).

Let $\mathfrak{G}$ be a class of valued fields (K, v). We say that $\mathfrak{G}$ satisfies an ***Ax–Kochen–Ershov principle*** (or an ***AKE principle***, or an ***AKE$^{\equiv}$ principle***) if for all $(K_1, v_1), (K_2, v_2) \in \mathfrak{G}$ we have:

$$K_1v_1 \equiv K_2v_2 \quad \& \quad v_1K_1 \equiv v_2K_2 \quad \Rightarrow \quad (K_1, v_1) \equiv (K_2, v_2)$$

(in the language of pure fields, ordered abelian groups, and valued fields, respectively — for these definitions one may take the language of valued fields, as [Kuh16, pp. 3-4] does, to be the language of pure fields together with a binary predicate $\mathcal{O}(a, b)$ interpreted as $v(a) \geq v(b)$.)

We also have the following condition: $\mathfrak{G}$ is said to satisfy an ***AKE$^{\preceq}$ principle*** if for all $(K_1, v_1) \subseteq (K_2, v_2)$ with $(K_1, v_1), (K_2, v_2) \in \mathfrak{G}$ we have:

$$K_1v_1 \preceq K_2v_2 \quad \& \quad v_1K_1 \preceq v_2K_2 \quad \Rightarrow \quad (K_1, v_1) \preceq (K_2, v_2).$$

There is also a version of this condition for existential closedness: we say $\mathfrak{G}$ satisfies an ***AKE*** $^{\exists}$ ***principle*** if for all $(K_1, v_1) \subseteq (K_2, v_2)$ with $(K_1, v_1), (K_2, v_2) \in \mathfrak{G}$ we have:

$$K_1 v_1 \leq_\exists K_2 v_2 \quad \& \quad v_1 K_1 \leq_\exists v_2 K_2 \quad \Rightarrow \quad (K_1, v_1) \leq_\exists (K_2, v_2).$$

While these definitions make sense in any class of valued fields $\mathfrak{G}$, they are most useful in *elementary* classes, as this allows us to reduce questions of decidability in valued fields to (generally far easier) questions of decidability in ordered abelian groups and residue fields.

Some examples of elementary classes satisfying AKE principles are:

- (importantly) the class of all tame valued fields of a given characteristic $p > 0$,
- algebraically closed valued fields,
- henselian fields of residue characteristic 0,
- p-adically closed fields.

See [Kuh16] for details, especially in the first case. See Chapter 5 for some applications of this.

2.3 Brief preliminaries on infinite Galois theory

We outline as briefly as possible the necessary background on infinite Galois theory as it relates to valued fields. The Galois-theoretic material in this section is drawn from [EP05] and [Jar96]. See [RZ00] for a comprehensive overview of the theory of profinite groups, and specifically §2.11 on profinite groups as Galois groups.

Let (K, v) be a valued field and (L, v') an infinite Galois extension. The ***Galois group*** $\mathrm{Gal}(L/K)$ is defined, as in the finite case, to be the group of all automorphisms of L which fix every element of K. This is a *profinite* group, obtained as an inverse limit over all Galois groups $\mathrm{Gal}(K'/K)$ of finite Galois extensions K' of K contained in L. As a profinite group it is a compact Hausdorff totally disconnected topological group, and naturally comes equipped with a topology (the ***Krull topology***) induced by the product topology. Conceretely, a basis of neighbourhoods of $1 \in \mathrm{Gal}(L/K)$ is given by the collection of subgroups $\{\, \mathrm{Gal}(L/N) \mid N$ is a finite Galois extension of K contained in $L\}$. Many of the theorems of finite Galois theory carry over to the infinite case if the word 'subgroup' is replaced by 'closed subgroup', 'normal subgroup' by 'closed normal subgroup', etc. In particular we have a correspondence between intermediate field extensions $K \subseteq K' \subseteq L$ and closed subgroups of $\mathrm{Gal}(L/K)$.

In the special case where K^{sep} is the separable-algebraic closure of K, we write $\mathrm{Gal}(K) := \mathrm{Aut}(K^{\mathrm{sep}}/K)$ for the ***absolute Galois group*** over K, or simply G_K where this is unambiguous. This is the same as considering the Galois group of the full algebraic closure $\overline{K}$ of K. For the next few definitions, where $\mathcal{O} = \mathcal{O}_v$ is the valuation ring of the valued field (K, v), let $\mathcal{O}'$ be a prolongation of $\mathcal{O}$ to K^{sep}. (It is possible to formulate these definitions relative to a normal extension N/K rather than relative to the separable closure; this is the approach taken in [Efr06].)

The ***decomposition group*** of $\mathcal{O}'$ is given by

$$G^h(\mathcal{O}') = \{\sigma \in \mathrm{Gal}(K^{\mathrm{sep}}/K) \mid \sigma(\mathcal{O}') = \mathcal{O}'\},$$

which is a closed subgroup of $\mathrm{Gal}(K)$ and whose fixed field is a henselization of K. If $\mathcal{O}^*$ is another prolongation of $\mathcal{O}$ to K^{sep} then $G^h(\mathcal{O}^*)$ is conjugate to $G^h(\mathcal{O}')$ in $\mathrm{Gal}(K^{\mathrm{sep}}/K)$.

Assume for the rest of this section that K is henselian. The ***inertia group*** of $\mathcal{O}'$ is given by

$$G^I(\mathcal{O}') = \left\{\sigma \in \mathrm{Gal}(K^{\mathrm{sep}}/K) \mid \sigma(x) - x \in \mathcal{M}' \text{ for all } x \in \mathcal{O}'\right\},$$

(where $\mathcal{M}'$ is the maximal ideal of $\mathcal{O}'$) and is a normal subgroup of G^h. We denote the fixed field of the inertia group by K^I when unambiguous. We have that $vK^I = vK$ and $K^Iv = (Kv)^{\mathrm{sep}}$, with an inclusion-preserving correspondence between intermediate fields $K \subseteq L \subseteq K^I$ and separable extensions $Kv \subseteq Lv \subseteq K^Iv$.

The ***ramification group*** of $\mathcal{O}'$ is given by

$$G^r(\mathcal{O}') = \left\{\sigma \in \mathrm{Gal}(K^{\mathrm{sep}}/K) \;\middle|\; \frac{\sigma(x)}{x} - 1 \in \mathcal{M}' \text{ for all } x \in (K^{\mathrm{sep}})^\times\right\}$$

and is a closed subgroup of the inertia group. It can be characterised as the unique (up to conjugation) p-Sylow subgroup of G^I, where $p > 0$ is the residue characteristic of K (with G^r trivial in residue characteristic 0), and it is a normal subgroup of G^h (and therefore G^I). Moreover, $G^I/G^r \cong \prod_{q \neq p} \mathbb{Z}_q^{r_q}$ as profinite groups, where for each prime number $q \neq p$ we denote by r_q the $\mathbb{F}_q$-dimension of the group $vK/q(vK)$. We denote the fixed field of the ramification group by K^r when unambiguous; its residue field is the same as that of K^I, and its value group satisfies $vK^r/vK = (vK^{\mathrm{sep}}/vK)_{p}$ (the subgroup consisting of all elements whose order is prime to p). There is an inclusion-preserving correspondence between all subextensions of K^r/K^I and all totally ordered groups lying between vK and vK^r.

The ramification group G^r and its fixed field K^r have intimate connections to the theory of tame and wild algebraic extensions of valued fields, as outlined below. The understanding of such extensions is crucical to understanding the model theory of the fields we are considering.

2.4 Preliminaries on tame, wild and purely wild extensions of valued fields

The best overview for this material can be found in [Kuh16].

Suppose throughout this section that (K, v) is a henselian valued field of residue characteristic $p > 0$. An algebraic extension $(L, w)/(K, v)$ is called ***tame*** if every finite subextension $(L_0, w_0)/(K, v)$ satisfies the following:

(i) $(w_0L_0 : vK)$ is prime to p,

(ii) the induced residue field extension L_0w_0/Kv is separable,

(iii) the extension is defectless; recall this means $[L_0 : K] = (w_0L_0 : vK)[L_0w_0 : Kv]$.

An extension of valued fields satisfying only the first two conditions is sometimes referred to as ***tamely ramified***, so we are using "tame" to mean "defectless tamely ramified". The definition can be extended to the residue characteristic 0 case by setting $p = 1$. General ramification theory gives us that K^r is the unique maximal tame extension of K.

An algebraic extension which is not tame is called ***wild***. We call an algebraic extension L/K ***purely wild*** if it is linearly disjoint from every tame extension; this will be the case if and only if L is linearly disjoint from K^r, if and only if vL/vK is a p-group and Lv/Kv is purely inseparable. Every immediate extension of valued fields is purely wild.

The key theorem of [KPR86] gives us that, where (K, v) is a henselian valued field of positive residue characteristic $p > 0$, there exist field complements L of K^r, such that $L.K^r = K^{\mathrm{alg}}$ and $L \cap K^r = K$. (Almost) equivalently, on the group side, the absolute Galois group $G = \mathrm{Gal}(K)$ splits over its ramification subgroup G^r, so that there exist group-theoretic complements H (closed subgroups of G) of G^r, such that $H \cap G^r = \{e\}$ and $H.G^r = G$. These two cases are in correspondence via the usual Galois correspondence between fixed fields of closed subgroups and intermediate fields inside K^{alg}. We note that the complements will in general not be unique. Note also that this allows us to talk of maximal purely wild extensions, to look at algebraic extensions of K as being made up of a tame part and a purely wild part, and so on.

A valued field is called a ***tame valued field*** (or simply ***tame***) if it is henselian and all its algebraic extensions are tame extensions. The class of all tame fields in a given characteristic satisfies an Ax–Kochen–Ershov principle, so that if the theories of Kv and vK are both decidable, then the theory of (K, v) is decidable. Moreover, the class of tame valued fields in a given characteristic satisfies an AKE$^{\exists}$ principle, so that where $(K_1, v_1) \subseteq (K_2, v_2)$ are two valued fields in this class, we have $(K_1, v_1) \leq_{\exists} (K_2, v_2)$ if and only if $K_1 v_1 \leq_{\exists} K_2 v_2$ and $v_1 K_1 \leq_{\exists} v_2 K_2$. We have the following characterisation of tame fields [Kuh16, Theorem 3.2]:

Proposition 2.4.1. *The following properties of a valued field (K, v) with residue characteristic $p > 0$ (or by setting $p = 1$ if char $Kv = 0$) are equivalent:*

- *(K, v) is a tame field,*
- *K^r is algebraically closed,*
- *(K, v) admits no proper purely wild extensions,*
- *(K, v) is algebraically maximal and closed under purely wild extensions by pth roots,*
- *(K, v) is algebraically maximal, vK is p-divisible and Kv is perfect.*

In the case where (K, v) is an equal characteristic valued field, then the property '(K, v) is a tame field' is also equivalent to '(K, v) is algebraically maximal and perfect'.

The last bullet point is the most useful for our purposes. It implies, for example, that every extension of $\mathbb{F}_p((t))$ with p-divisible value group (and residue field $\mathbb{F}_p$) admits a tame field as an algebraically maximal immediate extension. See Chapter 5 for some applications of this.

2.5 Fields of generalised Laurent series

We outline a notable class of fields which will be useful to study; see [EP05, § 3.5] or [Efr06, § 2.8] for details. It is a generalisation of the standard notion of $K((t))$ for a field K.

Definition 2.5.1. *Suppose Γ is an ordered abelian group and K a field. We define the* ***formal Laurent series field*** *$K((\Gamma))$ whose elements are the formal expressions $\sum_{\gamma\in\Gamma} a_\gamma t^\gamma$, with coefficients $a_\gamma \in K$ and indices $\gamma \in \Gamma$, such that the* ***support*** *$\{\gamma \in \Gamma \mid a_\gamma \neq 0\}$ is well-ordered. We define addition componentwise and multiplication is defined following the obvious rule*

$$\left(\sum_{\gamma\in\Gamma} a_\gamma t^\gamma\right) \cdot \left(\sum_{\gamma\in\Gamma} b_\gamma t^\gamma\right) = \sum_{\gamma} \left(\sum_{\delta+\epsilon=\gamma} a_\delta b_\epsilon\right) t^\gamma.$$

We define a valuation v on $K((\Gamma))$ by setting $v\left(\sum_{\gamma\in\Gamma} a_\gamma t^\gamma\right)$ to be the minimal element of the suppport, i.e. the first $\gamma \in \Gamma$ for which a_γ is nonzero. It is easily seen that if $\Gamma \cong \mathbb{Z}$ then we recover the standard Laurent series field $K((t))$.

These are often also known as ***Hahn series***, ***formal/generalised power series***, etc.

These fields have the important property that they do not admit proper immediate extensions — in other words, they are ***maximal***. In particular this is known to be equivalent to the property of being pseudo-complete: let λ be a limit ordinal and a_i $(i < \lambda)$ be elements of a valued field (K, v), satisfying $v(a_\gamma - a_\beta) > v(a_\beta - a_\alpha)$ whenever $\alpha < \beta < \gamma < \lambda$. (Such sequences are called ***pseudo-Cauchy***.) Then K is ***pseudo-complete*** if every pseudo-Cauchy sequence in K has a ***pseudo-limit*** $b \in K$, i.e., if there is some $b \in K$ such that

$v(b - a_\alpha) = v(a_{\alpha+1} - a_\alpha)$ for all $\alpha < \lambda$. (See [Kap42] for some applications and fleshing out of this theory.)

Note that if v is induced by an absolute value on K, with vK an ordered subgroup of $\mathbb{R}$, then the usual notion of *completion* will yield an immediate extension, so that a *maximal* valued field satisfying such conditions will be **complete** and not just pseudo-complete. The model theory of the fields $K((\Gamma))$ is relatively well-understood if $\mathrm{char}K = 0$ or if Γ is sufficiently divisible, as here we may treat it via the classification of tame valued fields (see below). However, in our situation much less is known, and we do not even know if e.g. the field $\mathbb{F}_p((\mathbb{Q} \oplus \mathbb{Z})) \cong \mathbb{F}_p((t))((\mathbb{Q}))$ is extremal, let alone if it is elementarily equivalent to $\mathbb{F}_p((t))$.

There are a few important unsolved questions we can ask about the class of fields $\mathfrak{G} := \{\mathbb{F}_p((\Gamma)) \mid \Gamma \equiv \mathbb{Z}\}$, none of which have been definitively ruled out (to the best of our knowledge):

- Are the fields $\mathbb{F}_p((\Gamma))$ extremal, for $\Gamma \equiv \mathbb{Z}$? Are there any maximal fields with value group elementarily equivalent to $\mathbb{Z}$ which are *not* extremal? (See preliminaries on extremal fields below.)

- If $\Gamma_1 \equiv \Gamma_2$ as ordered abelian groups, do we have that $\mathbb{F}_p((\Gamma_1)) \equiv \mathbb{F}_p((\Gamma_2))$? If not, what would a counterexample look like?

- An Ax–Kochen–Ershov principle in the (non-elementary) class $\mathfrak{G}$ would reduce the model theory of $\mathfrak{G}$ to the (already understood) model theory of $(\mathbb{Z}, +, <)$. How far away from this are we? For example, could we manage to prove that if $G_1 \leq_\exists G_2$ then $\mathbb{F}_p((G_1)) \leq_\exists \mathbb{F}_p((G_2))$?

2.6 Kaplansky's Hypothesis A and maximal immediate extensions of valued fields

We may always assume the existence of a maximal immediate extension of any given valued field by Zorn's Lemma. The question of whether such an extension is *unique* was treated in-depth by Kaplansky (see [Kap42], [Kap45]), where it was proved that in certain restricted situations we do in fact have uniqueness, but that in general it may not be assumed (with explicit counterexamples). Non-uniqueness of maximal immediate extensions has long been identified as one of the stumbling blocks on the road towards proving decidability for $\mathbb{F}_p((t))$.

In [Kap42] Kaplansky introduced a set of conditions, labelled 'Hypothesis A', that were sufficient (but *not* necessary) to guarantee that a valued field (K, v) admits a unique maximal immediate extension. These conditions were later shown to be equivalent to the following conditions in residue characteristic $p > 0$ (see [Wha57], [Del82]):

(i) (K, v) is perfect,

(ii) vK is p-divisible,

(iii) Kv admits no finite extensions of degree divisible by p.

(In residue characteristic 0 Hypothesis A is trivially satisfied.)

Kaplansky demonstrated that in the equicharacteristic $p > 0$ case, a maximal field (K, v) which satisfied Hypothesis A is isomorphic to a generalised Laurent series field $Kv((vK))$, and therefore that every equicharacteristic valued field can be identified with a subfield of some (possibly much larger) generalised Laurent series field. However, in the absence of Hypothesis A, there are maximal valued fields (K, v) which are not isomorphic to a generalised Laurent series field.

We have already seen that a valued field (K, v) being maximal is equivalent to the property that every pseudo-Cauchy sequence in K has a pseudo-limit in K. Now suppose $(a_\gamma)_{\gamma<\lambda}$ is a pseudo-Cauchy sequence in K with no pseudo-limit in K: define the ***breadth*** of the sequence $(a_\gamma)_{\gamma<\lambda}$ to be the (possibly fractional — see [Lan03, p. 88]) ideal in $\mathcal{O}_K$ consisting of precisely those elements $x \in K$ such that $v(x) > v(a_{\gamma+1} - a_\gamma)$ for all $\gamma < \lambda$. A pseudo-Cauchy sequence thus has breadth 0 if and only if the distance between elements of the sequence genuinely approaches 0. In [Kap45] it was proved that, if a valued field (K, v) satisfies the property that all of its pseudo-Cauchy sequences with no limit in K have breadth 0, then (K, v) admits a unique maximal immediate extension up to valuation-preserving isomorphism. Unfortunately in the non-Archimedean case it is easy to construct pseudo-convergent sequences with nonzero breadth, so that we should expect the existence of non-isomorphic maximal immediate extensions.

Can we at least hope that for a given valued field (K, v), the maximal immediate *algebraic* extension of (K, v) is unique up to isomorphism? Unfortunately, again here the answer is "no", as proved in [KPR86, Prop. 8.5]. The unfortunate reality in the equicharacteristic $p > 0$ case is that a valued field may have many non-isomorphic maximal immediate extensions, each of infinite transcendence degree over K (proved in [BK17]), and that we cannot even rely on the algebraic part being fixed.

Uniqueness of maximal immediate extensions is (thankfully) not a necessary condition for the existence of an AKE principle or otherwise having a 'good' model theory; for example, tame fields do not in general have unique maximal immediate extensions. (See [Kuh90].)

2.7 Preliminaries on extremal fields

We call a valued field (K, v) ***extremal*** if, for every $n \in \mathbb{N}$ and for every multivariable polynomial $f \in K[X_1, ..., X_n]$, we have that the set of values taken by f at points of $\mathcal{O}$, i.e.

$$\{v(f(a_1, ..., a_n)) \mid a_1, ..., a_n \in \mathcal{O}\} \subseteq vK \cup \{\infty\},$$

has a maximal element. That is, either f has a root in $\mathcal{O}$, or there is some tuple in $\mathcal{O}$ whose image under f has the maximal possible value realisable over $\mathcal{O}$.

Remark: It is crucial in the definition of extremality that the coefficients may range over K but we look only at values achieved on $\mathcal{O}$ — this corrects the original incorrect definition given by Ershov in [Ers04], after a simple counterexample due to Starchenko, namely the polynomial $X^2 + (XY - 1)^2$ over $\mathbb{R}((t))$. (See [AKP12, Remark 3.2] for details.) Observe that this counterexample uses only two variables. The current state of the art on extremal fields is outlined in [AK16], since which there have been no significant developments to the best of our knowledge. We present an outline of the theory of extremal fields as contained there.

Theorem 2.7.1. *Let (K, v) be a nontrivially valued field. If (K, v) is extremal then it is henselian, defectless and algebraically maximal. Furthermore it falls into one of the two following cases.*

(i) vK is a $\mathbb{Z}$-group.

(ii) vK is divisible and Kv is a large field.

As a partial converse we get the following.

(i) If a nontrivially valued field (K, v) is henselian, defectless and algebraically maximal, and either $\mathrm{char} Kv = 0$ with $vK \equiv \mathbb{Z}$, or $\mathrm{char} Kv = p > 0$ and $vK \cong \mathbb{Z}$, then (K, v) is

extremal.

(ii) If a nontrivially valued field (K, v) is henselian, defectless and algebraically maximal, and furthermore vK is divisible and Kv is large and ***perfect****, then (K, v) is extremal.*

However, this converse is the best we can do so far:

- *if (K, v) is henselian, defectless and algebraically maximal of residue characteristic $p > 0$ and $vK \equiv \mathbb{Z}$ then (K, v) may or may not be extremal,*
- *if (K, v) is henselian, defectless and algebraically maximal with vK divisible and Kv large but not perfect, then (K, v) may or may not be extremal.*

The situation is much better in residue characteristic 0, where we have a precise classification, but unfortunately that is not the situation we are in.

For clarity we will distinguish between the case where vK is a $\mathbb{Z}$-group, where we will call K ***discretely extremal***, and the case where vK is divisible and Kv is large, which we will call ***densely extremal***.

A few worked examples. As an example, consider the polynomial $f(X, Y) = X^2 + (XY - 1)^2$ over $\mathbb{F}_p((t))$, with $p = 13$ say. Then we have $f(0, 0) = 1$. To wipe out the constant term we can take e.g. $X = 3 + t \cdot \gamma$, $Y = 1 + t \cdot \delta$ (γ, $\delta \in \mathbb{F}_p[[t]]$) so that $f(X, Y) = (3^2 + 2^2) + 9t^2\gamma^2 + 3t^2\gamma^2\delta^2$, so that by appropriately selecting γ and δ we can achieve an arbitrarily high value of $f(X, Y)$. In this case we should expect a root of the polynomial in the extremal field $\mathbb{F}_p((t))$, and indeed if $\gamma = \delta = 0$ then we have $f(X, Y) = 0$.

If instead we set $p = 3$ and take the polynomial $f(X) = X^2 + 1$ we see that there is no root of $f(X)$ in $\mathbb{F}_p$, and that therefore over $\mathbb{F}_p[[t]]$ we will always have $vf(X) = 0$. Working in a more arbitrary field where x is any element ($\in \mathcal{O}$) which is not a square, we can similarly use the polynomial $X^2 - x$, whose value will not exceed $v(x)$.

We see that in these examples the calculations involved tend to hinge on algebra over the residue field $\mathbb{F}_p$. Whether a given value can be achieved or not — and, therefore, whether or not our field is extremal — will reduce to a question of whether a set of simultaneous equations has a solution over the residue field or not. In practice it is difficult to see how to turn this style of argument into something concrete, but the principle of extremality reducing to algebra over the residue field should extend to the power series fields $K((\Gamma))$ for (regular) discrete Γ, although it is unclear if the argument can go any further than that. Indeed we can formulate the following conjecture:

Conjecture 2.7.2. *Suppose K is a field of characteristic $p > 0$ and Γ a $\mathbb{Z}$-group such that $K((\Gamma))$ is extremal. Then for all $\mathbb{Z}$-groups Δ, $K((\Delta))$ is extremal.*

In many ways a counterexample to this conjecture might turn out to be more interesting than a proof.

Remark: Suppose (K, v) is a discretely extremal valued field with $Kv = \mathbb{F}_p$. If $\text{char}K = p$ then we end up with some extension of $\mathbb{F}_p(t)^h$, the henselization of $\mathbb{F}_p(t)$. If $\text{char}K = 0$, however, we need only stipulate that $v(p)$ is the smallest element of the value group and we already recover that our field is p-adically closed, and therefore elementarily equivalent to the field $\mathbb{Q}_p$ of p-adic numbers. It at least seems plausible that in the residue characteristic p case, stipulating $v(t)$ to be minimal positive (as well as any necessary information on the degree of imperfection) may well be enough by itself to recover the elementary theory of $\mathbb{F}_p((t))$.

To the best of our knowledge, it is currently still open as to whether one can find extremal fields of arbitrary finite Archimedean rank, and how to construct them if so. There remain many unanswered questions about extremal fields in general, and we hope to see more work in this area in coming years.

2.8 Brief preliminaries on algebraic geometry

We present some very brief preliminaries on a few concepts from algebraic geometry which are necessary to include. These are all taken from Chapter 10 of [FJ08], occasionally simplified for our purposes.

- Let K be a field, $n \in \mathbb{N}$, Ω an algebraically closed field of infinite transcendence degree over K, and define the ***affine space*** $\mathbb{A}^n := (x_1, ..., x_n)$ to be the set of all points with coordinates in Ω. A K***-variety*** (or simply ***variety***) $V(\mathfrak{a}) \subseteq \mathbb{A}^n$ is the zero locus of an ideal $\mathfrak{a}$ of the polynomial ring $K[X_1, ..., X_n]$. (These form the closed sets of the ***Zariski topology*** on $\mathbb{A}^n$.) Such a variety is ***irreducible*** if it cannot be written nontrivially as the union $V(\mathfrak{a}) = V(\mathfrak{b}) \cup V(\mathfrak{c})$ of two varieties.

- We will sometimes speak of the intersection of such a variety with the base field K, i.e. its set of K***-rational points***. If $X \subseteq \mathbb{A}^n$ is a variety then the set of its K-rational points is denoted $X(K) \subseteq K^n$.

- Let $A \subseteq \mathbb{A}^n$. Take $I(A)$ to be the set of polynomials in $K[X_1, ..., X_n]$ which vanish everywhere on A. This is an ideal of $K[X_1, ..., X_n]$. If A is a variety then $I(A)$ is prime iff A is irreducible. Given an irreducible variety V, there is a K-embedding of $K[X_1, ..., X_n]/I(V)$ in Ω; the image $x = (x_1, ..., x_n) = (X_1, ..., X_n) + I(V)$ is called a ***generic point*** of V over K; $K[x]$ is the ***coordinate ring*** of V and $K(x)$ its ***function field***. The transcendence degree of $K(x)$ over K is an invariant, called the ***dimension*** of V.

- An irreducible K-variety of dimension 1 is called a K***-curve***.

- If L/K is an extension of fields then a variety over K may also be considered as a variety over L. A K-variety which remains irreducible over every extension of K is called ***absolutely irreducible***. (For a non-example, consider $V(X^2+Y^2)$ over $\mathbb{Q}$, which decomposes as $V(X+iY)\cup V(X-iY)$ over the extension $\mathbb{Q}(i)$.)

- Let $V\subseteq\mathbb{A}^n$, $W\subseteq\mathbb{A}^m$ be (irreducible) varieties. A ***K-rational map*** $\phi:V\to W$ is a morphism of varieties which is locally a rational function: suppose x and y are generic points of V and W, respectively, with $y=(y_1,...,y_m)$ such that each y_i is an element in the function field $K(x)$. Then $I:=\{g\in K[X_1,...,X_n]\mid g(x)y_i\in K[x]\quad 1\le i\le m\}$ is a nonzero ideal of $K[X_1,...,X_n]$ and $U_0:=V\setminus V(I)$ is a Zariski-open subset of V. For each $a\in U_0$ there exist $f_1,...,f_m,g\in K[X_1,...,X_n]$ such that $g(a)\neq 0$ and $g(x)y_i=f_i(x)$ for $1\le i\le m$. Hence $b=\left(\frac{f_1(a)}{g(a)},\frac{f_2(a)}{g(a)},...,\frac{f_m(a)}{g(a)}\right)$ is a well-defined point of W which depends only on a and y, and not on any particular choice of f_i, g. We can then form $\phi:U_0\to W$ as the map sending a to b in this way; the Zariski-open set U_0 is the ***domain of definition*** of ϕ. (We speak of $\phi:V\to W$ as being a rational map when, strictly speaking, ϕ is only defined on an open subset of V.)

- Suppose $\phi:V\to W$ is a K-rational map and suppose further that there is some K-rational map $\psi:W\to V$ such that $\phi\circ\psi$ and $\psi\circ\phi$ are each the identity map on some open subset of the domain of definition of ϕ, ψ (respectively). Then ϕ and ψ are called ***K-birational equivalences*** and we say that the varieties V and W are ***K-birationally equivalent***.

Chapter 3

Towards a quantifier elimination result

At the top of this chapter we seek to outline our approach, the necessary preliminaries now being in place.

Decidability of the first-order theory of $\mathbb{F}_p((t))$ would be implied by a quantifier elimination result in a suitable language, combined with the existence of an algebraically prime model. We have an obvious candidate for an algebraically prime model — the henselization $\mathbb{F}_p(t)^h$ of $\mathbb{F}_p((t))$ — and we have what appears to be a relatively tractable way to approach quantifier elimination, i.e. the Shoenfield criterion as used by Macintyre in [Mac76]. Recall from preliminaries:

Theorem 2.2.7. *Let T be an $\mathcal{L}$-theory. Then T has quantifier elimination if and only if the following holds: whenever we have $\mathcal{L}$-structures $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{M}_1$, $\mathcal{M}_2$ such that $\mathcal{A}_i \subseteq \mathcal{M}_i \models T$ (for i = 1, 2), $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{M}_1$ are countable, $\mathcal{M}_2$ is $\aleph_1$-saturated, and we are provided with an isomorphism $f : \mathcal{A}_1 \cong \mathcal{A}_2$, we can always extend f to an embedding of $\mathcal{M}_1$ into $\mathcal{M}_2$.*

Historically problems have arisen due to the non-uniqueness of maximal immediate (algebraic) extensions of valued fields in positive characteristic, so that an initial approach might be to add predicates to our language to control the purely wild extensions that appear over the course of the proof. Although we somewhat understand the structure of purely wild extensions (see §3.1), our attempts in this direction were not successful, and even if we were to add predicates for all purely wild extensions into our language we would still run into obstacles: suppose for example that we have an arbitrary wild extension L/K that can be decomposed as a tower $K \subseteq F \subseteq L$ of extensions with L/F purely wild and F/K solvable. How would we go about guaranteeing that all the new predicates introduced to handle purely wild extensions would survive the step from K to F? It seemed to us that this got us bogged down in intractable technical details when what was desired was an overarching approach, so we decided to broaden out the class of predicates to add to the language. To that end we consider things in fairly wide generality — we take $\mathcal{F}$ to be the set of all multivariable polynomials with coefficients in $\mathbb{F}_p(t)$, and for each $f \in \mathcal{F}$ we add a predicate R_f to represent the existence of a root of f in K.

After some discussion of relevant algebraic considerations in §3.1, we introduce our language and theory in §3.2. Our main proof of quantifier elimination in our theory $T_{\mathcal{F}}$, conditional on the truth of a technical statement about certain subtypes of existential types, is carried out over §3.3 and §3.4, with the culmination being Theorem 3.4.2; §3.6 sketches an attempted proof of an *unconditional* result, and highlights two areas where difficulties arise. Additional results about models of $T_{\mathcal{F}}$ are obtained in §3.6.

As all our newly introduced predicates are $\varnothing$-definable in the language $\mathcal{L}_t$ of valued fields together with t, our enriched language is merely a reduct of our original language (in the standard modern model-theoretic sense), so that the structures considered do not fundamentally change. However, when it comes to results about quantifier complexity etc., more care is needed; there is some discussion of this in §3.7.

3.1 Purely wild extensions and additive polynomials

The background theory in this section can be found in [Kuh06]; we are abbreviating it somewhat for our purposes.

Let K be an infinite field of characteristic $p > 0$. An ***additive polynomial*** $f \in K[X]$ is a polynomial such that $f(a+b) = f(a) + f(b)$ for all $a, b \in K$. This will happen if and only if $f(X)$ is of the form $\sum_{i=0}^{n} a_i X^{p^i}$ for some $n \in \mathbb{N}, a_i \in K$. We call $g \in K[X]$ a p***-polynomial*** if it is of the form $f(X) - c$ for some additive polynomial $f \in K[X]$ and $c \in K$. One can generalise these notions to the multivariable polynomial case fairly easily, and in fact a polynomial $f \in K[X_1, ..., X_n]$ is additive if and only if we have $f(X_1, ..., X_n) = \sum_{i=1}^{n} f_i(X_i)$ for additive polynomials $f_i \in K[X_i]$ $(1 \leq i \leq n)$. If a finite algebraic extension L/K admits no proper subextension we call it ***minimal***. We have the following theorem ([Kuh06, Theorem 13], originally due to Pop):

Proposition 3.1.1. *Suppose (K, v) is a henselian valued field of characteristic $p > 0$ and L/K is a minimal purely wild extension. Then there exists an additive polynomial $f \in \mathcal{O}_K[X]$ and an element $\vartheta \in L$ such that $L = K(\vartheta)$ and the p-polynomial $f(X) - f(\vartheta)$ is the minimal polynomial of ϑ over K.*

One can therefore see any finite purely wild extension of a henselian valued field (K, v) as being generated by a tower of p-polynomials. As such, understanding such extensions will be important when considering immediate purely wild extensions of valued fields (which tend to cause the most trouble for the model theory of such fields).

For technical reasons we also note the following ([Kuh06, Theorem 17], originally [Kuh11]):

Definition 3.1.2. *We call a valued field* (K, v) ***inseparably defectless*** *if every finite purely inseparable extension* L/K *is defectless.*

Proposition 3.1.3. *Let* (K, v) *be a valued field of characteristic* $p > 0$*. Then* K *is inseparably defectlesss if and only if it is extremal with respect to every* p*-polynomial of the form* $b - \sum_{i=1}^{n} b_i X_i^p$*, for* $n \in \mathbb{N}$ *and* $b, b_1, ..., b_n \in K$*.*

An example: iterated Artin–Schreier extensions

The simplest case of an immediate purely wild algebraic extension in equicharacteristic $p > 0$ is an Artin–Schreier extension generated by a polynomial of the form $X^p - X - \alpha = 0$, where $v(\alpha) < 0$ (by Hensel's Lemma). The next simplest case would be a tower of such extensions; let K be a characteristic p valued field and suppose L/K is a field extension which decomposes as a finite tower of Artin–Schreier extensions $K = L_0 \subseteq L_1 \subseteq L_2 \subseteq ... \subseteq L_n = L$ such that each L_{i+1}/L_i is generated by the Artin–Schreier polynomial $X^p - X - \alpha_i = 0$, with $\alpha_0 \in K$ and $\alpha_1, ..., \alpha_n \in L$. That is, we have $\alpha_1^p - \alpha_1 = \alpha_0$, $\alpha_2^p - \alpha_2 = \alpha_1$, etc.

One sees that $(\alpha_2^p - \alpha_2)^p - (\alpha_2^p - \alpha_2) = \alpha_1^p - \alpha_1 = \alpha_0$, so that by direct computation we get that

- L_2/K is generated by the polynomial $X^{p^2} - 2X + X = \alpha_0$
- L_3/K is generated by the polynomial $X^{p^3} - 3X^{p^2} + 3X^p - X = \alpha_0$
- L_4/K is generated by the polynomial $X^{p^4} - 4X^{p^3} + 6X^{p^2} - 4X^p + X = \alpha_0$

and a straightforward inductive argument gives us that L_i/K is generated by the polynomial

$$\alpha_0 = \sum_{i=0}^{n} (-1)^i \binom{n}{i} X^{p^i}$$

so that this would give us examples of the sort of additive polynomials we would need to consider. Note also that in this example we're only using coefficients from $\mathbb{F}_p$, so that this is quite a restricted case. Even more restrictive is our assumption on the inductive definitions of α_i, so that where $a \in K \setminus \wp(K)$ is a non-Artin–Schreier root of K, the polynomial $\sum_{i=0}^{n}(-1)^i\binom{n}{i}X^{p^i} - a = 0$ might not be the only polynomial which generates the extension L_n/K.

One might hope that this sort of example is enough to understand the problematic behaviour of purely wild extensions for model-theoretic purposes, in the light of a theorem (see e.g. [Kuh06, p. 11]) that given a finite purely wild extension L/K and where K^r is the ramification field of K, the extension $L.K^r/K^r$ (where . denotes the compositum) can be realised as a tower of Artin–Schreier extensions. That is, once ramification is taken into account, any finite purely wild extension is 'morally speaking' a tower of Artin–Schreiers. However, this does not suffice for our purposes: to be sufficient we would need that fixing the 'maximal immediate iterated Artin–Schreier extension of K' also pins down the precise structure of purely wild immediate algebraic extensions of K, corresponding to choosing a specific complement of the ramification group of K in its absolute Galois group.

Translating via Galois theory, we would need the following group-theoretic statement to be true:

Statement 3.1.4. *Suppose G is a finite soluble group containing a normal subgroup N of order p^k and a complement H (so $G = N \rtimes H$), such that there is no normal subgroup $U \leq G$ with $H \leq U$ and $[G : U] = p$. Then H is unique up to conjugation in G.*

Unfortunately the statement is false. We present the following counterexample:

Take $G = N \rtimes H$, where $N = C_3^5 = \langle a\rangle \times \langle b\rangle \times \langle c\rangle \times \langle d\rangle \times \langle e\rangle$, $H = C_3 \times C_5 = \langle\gamma\rangle \times \langle\delta\rangle$, such that γ acts trivially on N via conjugation, and δ acts on N via the 5-cycle $(a\,b\,c\,d\,e)$. The only nontrivial proper H-invariant subgroup of N is $\langle a\,b\,c\,d\,e\rangle$, so the only subgroup U

with $H \subsetneq U \subsetneq G$ has $[G:U] = 3^5$. But $a\,b\,c\,d\,e \in Z(G)$ and $\gamma \in Z(G)$, so $H' = \langle \gamma' \rangle \times \langle \delta \rangle$ with $\gamma' := (a\,b\,c\,d\,e)\gamma$ is another complement of N in G satisfying the same conditions on $H' \subsetneq U \subsetneq G$. However, H' is not conjugate to H in G, as the only conjugate of γ is $\gamma \notin H'$.

Moreover we can realise this group over $K = \mathbb{F}_3((t))((u))$ by some L/K with $\mathrm{Gal}(L/K) = G$ such that N is the ramification subgroup with respect to the u-adic valuation on K: take $L = \mathbb{F}_{3^{15}}(\alpha_1, ..., \alpha_5)$ where α_i is a zero of $X^3 - X - \zeta^{i-1}u^{-1}$ and ζ is a primitive 11th root of unity, as $3^5 - 1 = 242 = 2 \cdot (11)^2$.

The upshot of all this is that there are examples of fields K such that the absolute Galois group G_K, which splits as $H \cdot G^r$ for some complement H of the ramification subgroup G^r, cannot be made to split *uniquely* (up to conjugation) solely by fixing a maximal Artin–Schreier extension. Therefore in order to deal with the problem we will need to tackle the case of purely wild extensions in greater generality. This would involve at the very least including all polynomials which might represent e.g. the purely wild Artin–Schreier extensions of some finite subextension L of K inside K^r; as this seems like a difficult question to answer, we take the approach of encoding as much of the algebra as we can via predicates.

3.2 Setting up our language $\mathcal{L}_{\mathcal{F}}$ and theory $T_{\mathcal{F}}$

Let $\mathcal{F}$ be a given set of multivariable polynomials with coefficients in $\mathbb{F}_p(t)$. Define the expansion $\mathcal{L}_{\mathcal{F}}$ of the language $\mathcal{L}_t = \{+, -, \cdot, *^{-1}, 0, 1, \mathcal{O}, |_v, t\}$ (of valued fields with a constant symbol for t) by adding an n-ary predicate $R_f(x_1, ..., x_n)$ for each $f \in \mathbb{F}_p(t)[X_1, ..., X_n, Y]$ that appears in $\mathcal{F}$. (We include certain $\varnothing$-definable symbols e.g. $*^{-1}$, $|_v$ purely to make life easier regarding quantifiers; $x|_v y$ has the standard interpretation $v(x) \leq v(y)$.)

Define the theory $T_{\mathcal{F}}$ to be the first-order $\mathcal{L}_{\mathcal{F}}$-theory axiomatised by the following, for a structure $\mathcal{K}$.

(i) $\mathcal{K}$ is an extremal valued field with valuation ring $\mathcal{O}$, of characteristic $p > 0$, with $+, -, \cdot, *^{-1}, 0, 1, \mathcal{O}$ and $|_v$ having their usual interpretations,

(ii) The value group vK is a $\mathbb{Z}$-group and the residue field Kv is precisely $\mathbb{F}_p$,

(iii) $v(t)$ is minimal positive in vK,

(iv) K has degree of imperfection p, with p-basis $\{1, t, t^2, ..., t^{p-1}\}$,

(v) $\mathcal{K} \models R_f(\overline{a})$ if and only if $\mathcal{K} \models \exists y\, f(\overline{a}, y) = 0$, for every $f \in \mathcal{F}$.

(See [Cha, §1.16] or [Bou13, Ch. V §13] on p-bases and degree of imperfection.)

From now on we take $\mathcal{F}$ to be the set of all multivariable polynomials over $\mathbb{F}_p(t)$, and write T to mean $T_{\mathcal{F}}$ when unambiguous. We leave open the question of which recursive sets $\mathcal{F}$ would suffice to prove a Shoenfield-style result.

A key ingredient in our proof attempt is the following.

Proposition 3.2.1. *The henselization $\mathbb{F}_p(t)^h$ of $\mathbb{F}_p(t)$ is an algebraically prime model of $T_{\mathcal{F}}$.*

Proof. Suppose $K \models T_{\mathcal{F}}$; we wish to show there is an embedding of $\mathbb{F}_p(t)^h$ into K. By extremality we know that K is henselian, and as it is an $\mathcal{L}_{\mathcal{F}}$-structure it contains a copy of t, so we know that K must contain a copy of the henselization of $\mathbb{F}_p(t)$. Indeed, $\mathbb{F}_p(t)^h$ itself is a model of $T_{\mathcal{F}}$, and it embeds as a valued field (though not necessarily elementarily) in every model of our theory by construction. Moreover we know (see e.g. [Kuh16, Theorem 5.12]) that $\mathbb{F}_p(t)^h \leq_\exists \mathbb{F}_p((t))$ so that given any new predicate R_f, both fields will agree on $R_f \cap \mathbb{F}_p(t)^h$. We deduce that $\mathbb{F}_p(t)^h$ is an algebraically prime model of our theory and therefore if $T_{\mathcal{F}}$ is model-complete, it is complete. □

We note also the following. (It is trivial if we include (v).)

Proposition 3.2.2. *Let $K_1 \subseteq K_2$ be two models of $(i) - (iv)$ of T. Then K_1 is relatively algebraically closed inside K_2.*

Proof. As K_1 is extremal it is defectless and algebraically maximal, and so any proper algebraic extension L/K_1 contained in K_2 must satisfy $[L : K_1] = (vL : vK_1) \cdot [Lv : K_1v]$. As L is contained within K_2 it must have residue field $\mathbb{F}_p$, leaving only extensions arising from adding a new element (α, say) to the value group. Now as $v(t)$ remains minimal positive in K_2 this new element α must fall outside the isomorphic copy of $\mathbb{Z}$ in vL. Without loss of generality suppose $(vL : vK_1) = q$ is prime, with α a generator of vL/vK_1, and let $a \in L$ be such that $v(a) = \alpha$. Then $q \cdot \alpha$ is an element of vK_1, say β, with $b := a^q$ having value β. Now $b \in L$ so adjoining b to K_1 does not alter the residue field $\mathbb{F}_p = Lv = K_1v$, and by construction b does not alter the value group vK_1, meaning that $K_1(b)/K_1$ is an immediate extension. By assumption $b \in L$ is algebraic, therefore $K_1(b)/K_1$ is an immediate algebraic extension, therefore (as K_1 is algebraically maximal) it is trivial. Therefore K_2 contains no proper algebraic extension of K_1, i.e. K_1 is relatively algebraically closed inside K_2.

(One may also deduce this from the fact that, as vK_1 and vK_2 are both $\mathbb{Z}$-groups, vK_1 is pure in vK_2 and so any algebraic intermediate extension must be immediate, by [EP05, Lemma 6.2.2].) □

As such, the new predicates we introduce are not so much to control algebraic extensions between two different models of our theory; instead, we introduce them to control algebraic extensions of *subfields* which may generally be much smaller.

3.3 Algebraic configurations

We introduce the following notion for technical reasons.

Definition 3.3.1. *Let $\mathcal{A}$ be a model of $T_{\mathcal{F}}$ with domain A with C a subfield of A. Define the* ***algebraic configuration*** $\Phi_{\mathcal{A}}(a/C)$ *of an element $a \in A$ to be the set*

$$\{f(X, \overline{c}, Y) \mid f(X, \overline{Z}, Y) \in \mathcal{F},\ \overline{c} \text{ is a tuple in } C, \text{ and } \mathcal{A} \models \exists y\ f(a, \overline{c}, y) = 0\ \}.$$

This set $\Phi_{\mathcal{A}}(a/C)$ is a set of polynomials from $\mathcal{F}$ (which, recall, we take to be all multivariable polynomials over $\mathbb{F}_p(t)$), with parameters from C. We see that $f(X, \overline{c}, \overline{Z}) \in \Phi_{\mathcal{A}}(a/C)$ if and only if $\mathcal{A} \models R_f(a, \overline{c})$. What we are doing here is picking out a certain subtype of the quantifier-free $\mathcal{L}_{\mathcal{F}}$-type / existential $\mathcal{L}_t$-type of a, so as to capture the behaviour at a of all polynomials over C. (Compare in particular the notion of 'specializability of selected elements' in the van den Dries test for quantifier elimination.)

Although the definition of the set references $\mathcal{A}$, the elements of the set are polynomials from $\mathcal{F}$ with parameters from C, so that if $\mathcal{B}$ is another model of T which contains C as a subfield, and b is an element from B, we can ask the question whether $\Phi_{\mathcal{A}}(a/C) = \Phi_{\mathcal{B}}(b/C)$. Our purpose is the following result.

Lemma 3.3.2. *Suppose $\mathcal{A}$, $\mathcal{B}$ are two models of $T_{\mathcal{F}}$ with common subfield C, $a \in A \setminus C$, $b \in B \setminus C$. Then $\Phi_{\mathcal{A}}(a/C) = \Phi_{\mathcal{B}}(b/C)$ if and only if $C(a) \cong C(b)$ as $\mathcal{L}_{\mathcal{F}}$-structures.*

Proof. Clearly if $C(a) \cong C(b)$ as $\mathcal{L}_{\mathcal{F}}$-structures then $\mathcal{A}$ and $\mathcal{B}$ agree on which polynomials over $C \cup \{a\}$ (resp. $C \cup \{b\}$) have roots and which do not, by virtue of the predicates R_f, whence $\Phi_{\mathcal{A}}(a/C) = \Phi_{\mathcal{B}}(b/C)$. For the other direction, suppose $\Phi_{\mathcal{A}}(a/C) = \Phi_{\mathcal{B}}(b/C)$. Define the partial embedding $\phi_0 : \mathcal{A} \to \mathcal{B}$ as the identity map $C \to C$ together with the mapping $\phi_0(a) = b$. This can be extended uniquely to a valued field homomorphism $\phi : C(a) \to C(b)$. As $\Phi_{\mathcal{A}}(a/C) = \Phi_{\mathcal{B}}(b/C)$ we have that a and b will satisfy the same R_f predicates, as will all rational functions in a (resp. b) over C. Moreover as $\mathcal{O}$ is existentially $\varnothing$-definable in the language of rings then membership of $\mathcal{O}$ will be represented by some polynomial, so will be captured by one of the R_f predicates. We conclude that ϕ is an $\mathcal{L}_{\mathcal{F}}$-isomorphism

$C(a) \to C(b)$. $\square$

3.4 A conditional quantifier elimination result

Recall again our test for quantifier elimination.

Theorem 2.2.7. *Let T be an $\mathcal{L}$-theory. Then T has quantifier elimination if and only if the following holds: whenever we have $\mathcal{L}$-structures $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{M}_1$, $\mathcal{M}_2$ such that $\mathcal{A}_i \subseteq \mathcal{M}_i \models T$ (for $i = 1, 2$), $\mathcal{A}_1$, $\mathcal{A}_2$ and $\mathcal{M}_1$ are countable, $\mathcal{M}_2$ is $\aleph_1$-saturated, and we are provided with an isomorphism $f : \mathcal{A}_1 \cong \mathcal{A}_2$, we can always extend f to an embedding of $\mathcal{M}_1$ into $\mathcal{M}_2$.*

We first prove an important lemma.

Lemma 3.4.1. *Let $\mathcal{L}_{\mathcal{F}}$ and $T_{\mathcal{F}}$ be as defined in §3.2, (K, v) a countable model of $T_{\mathcal{F}}$, (K^*, v^*) an $\aleph_1$-saturated model of $T_{\mathcal{F}}$, with $C_1 \subseteq K$ and $C_2 \subseteq K^*$ two $\mathcal{L}_{\mathcal{F}}$-structures together with an isomorphism $\phi : C_1 \to C_2$. Then ϕ can be extended to an isomorphism $\phi' : C_1^{\mathrm{alg}} \cap K \to C_2^{\mathrm{alg}} \cap K^*$.*

Proof. As C_1 and C_2 are isomorphic we will treat them as a common substructure C where appropriate. An $\mathcal{L}_{\mathcal{F}}$-substructure C of a model of $T_{\mathcal{F}}$ will automatically inherit the structure of a valued field as it contains 0, 1, is closed under $+$, $-$, $\cdot$ and $*^{-1}$, and for all $x \in C$ we have one of $x \in \mathcal{O}$ or $x^{-1} \in \mathcal{O}$ holding.

We may recursively extend our embedding $\phi : C_1 \to C_2$ to an embedding $C_1^{\mathrm{alg}} \cap K \to C_2^{\mathrm{alg}} \cap K^*$ in the following way, with the inductive hypothesis that ϕ_β is a partial $\mathcal{L}_{\mathcal{F}}$-embedding (and thus an isomorphism onto its image).

Let $\phi_0 := \phi$ and well-order the elements of $C_1^{\mathrm{alg}} \cap K$ by some well-ordering $<_{\text{w.o.}}$ with order-type $\alpha \in \mathbf{ON}$. Given a partial embedding ϕ_β, define $\phi_{\beta+1}$ by setting $\phi_{\beta+1} = \phi_\beta$ if the first β elements of $C_1^{\mathrm{alg}} \cap K$ are already in $D := \mathrm{dom}\ \phi_\beta$, else pick x the first element under $<_{\text{w.o.}}$ not in D such that the extension $D(x)/D$ is purely inseparable (if this class is nonempty). Then where $x^{p^n} \in D$ and $f(X,Y)$ is the polynomial $Y^{p^n} - X$, we have $K \models R_f(x^{p^n})$ so that $K^* \models R_f(\phi_\beta(x^{p^n}))$, whence there is a p^n-th root x' of $\phi_\beta(x^{p^n})$ in K^*. Take $\phi_{\beta+1}$ to be the field homomorphism given by taking the identity on D and setting $\phi_{\beta+1}(x) = x'$. In this way we handle all purely inseparable extensions.

Else, take x the first element under $<_{\text{w.o.}}$ not in D; then the extension $D(x)/D$ is separable. Now by inductive hypothesis ϕ_β is an $\mathcal{L}_{\mathcal{F}}$-embedding and so D is a common $\mathcal{L}_{\mathcal{F}}$-substructure of both K and K^*, so that for every algebraic element $x \in K$ over D, having minimal polynomial $m_x(X)$ over D with coefficients $a_1, ..., a_n \in D$, and where $f(Y_1, ..., Y_n, X)$ is the polynomial obtained from $m_x(X)$ by substituting every coefficient a_i by the variable Y_i, we have that $K^* \models R_f(a_1, ..., a_n)$; as K^* is a model of $T_{\mathcal{F}}$ then we have $K^* \models \exists z\, f(a_1, ..., a_n, z) = m_x(z) = 0$. As $D(x)/D$ is separable then $m_x(X)$ is separable and so it has roots $x_1, ..., x_m \in K$ and $x'_1, .., x'_m \in K^*$. So take $\phi_{\beta+1}$ to be the field homomorphism given by taking the identity on D together with the mappings $x_i \mapsto x'_i$, $1 \leq i \leq m$. (If the extension $D(x)/D$ is not normal then we may need to take more care here, but in such a situation both K and K^* will have the same number of roots of m_x due to both satisfying the same R_f predicates over D, so that we will be able to extend $\phi_{\beta+1}$ to some ϕ_β.)

If λ is a limit ordinal then take $\Phi_\lambda = \bigcup_{\beta<\lambda} \Phi_\beta$, and let $\phi' := \phi_\alpha$. Thus ϕ_α is well-defined if we can show that each ϕ_β is indeed a $\mathcal{L}_{\mathcal{F}}$-embedding and not just a field homomorphism.

We need to show that ϕ_β preserves $\mathcal{O}$ and our new predicates R_f to show that it is an $\mathcal{L}_{\mathcal{F}}$-embedding. For $\mathcal{O}$, one may assume C henselian (as K, K^* are both henselian fields containing C so that both will contain a henselization of C, and this is unique up to isomorphism of valued fields: alternatively one may use our R_f predicates to adjoin roots of Hensel polynomials directly), so that where $x \in K$ is algebraic over C, the algebraic extension $C(x)/C$ admits a unique prolongation of the valuation on K, so that where x' is the image of x in K^* under ϕ_β, we have that $C(x) \cong C(x')$ as valued fields. Now if $f \in F$ is a multivariable polynomial then we need to show that, where $a_1, ..., a_n \in C^{\text{alg}} \cap K$ have images $a'_1, a'_2, ..., a'_n \in C^{\text{alg}} \cap K^*$ under ϕ_β, we have $K^* \models R_f(a'_1, ..., a'_n)$ if and only if $K \models R_f(a_1, ..., a_n)$; but as each a_i is algebraic over C we may replace each a_i with a variable y_i and stipulate that each y_i is a root of the minimal polynomial $m_{a_i}(X)$ over C, the upshot of all of which is that we can replace $f(X_1, ..., X_n)$ with a polynomial $g(X_1, ..., X_n, \overline{Y})$ such that $K \models R_f(a_1, ..., a_n)$ if and only if there are $c_1, ..., c_n \in C$ and some tuple $\overline{z}$ (representing the coefficients of all minimal polynomials involved) such that $K \models R_g(c_1, ..., c_n, \overline{z})$. And as C is a common $\mathcal{L}_\mathcal{F}$-substructure of both K and K^*, this will automatically be satisfied in K^*. Hence ϕ_β preserves $\mathcal{O}$ and all our R_f predicates, and is an $\mathcal{L}_\mathcal{F}$-embedding as claimed; therefore ϕ_α is a well-defined $\mathcal{L}_\mathcal{F}$-embedding.

We have thus recursively defined an embedding $C_1^{\text{alg}} \to C_2^{\text{alg}}$, and moreover, for every element $x \in C_2^{\text{alg}} \cap K^*$, x must satisfy some minimal polynomial m_x over C_2, whose roots will already have been mapped to in the above recursion, so that ϕ_α is surjective. We conclude that $\phi_\alpha : C_1^{\text{alg}} \cap K \to C_2^{\text{alg}} \cap K^*$ is an isomorphism of $\mathcal{L}_\mathcal{F}$-structures. □

We are now ready to prove our main theorem.

Theorem 3.4.2. *Suppose the following statement holds: whenever K, K^* are two models of $T_\mathcal{F}$ with common subfield D, such that K, D are countable, $D \subsetneq K$, and K^* is $\aleph_1$-saturated; then there is at least one $a \in K \setminus D^{\text{alg}}$ and one $b \in K^* \setminus D^{\text{alg}}$ with $\Phi_K(a/D) = \Phi_{K^*}(b/D)$. Then $T_\mathcal{F}$ admits quantifier elimination. Consequently, it is model-complete, complete, and decidable.*

Proof. Let $K = \mathcal{M}_1$, $K^* = \mathcal{M}_2$ and $C = \mathcal{A}_1 = \mathcal{A}_2$ be as in Theorem 2.2.7, i.e. K is countable, K^* is $\aleph_1$-saturated, both K and K^* are models of $T_{\mathcal{F}}$ with C as common subfield. Enumerate the elements of K via a well-order $<_{\text{w.o.}}$ with order-type $\alpha \in \mathbf{ON}$ and define the partial embedding $\phi_0 : K \to K^*$ as the identity function on C. Given a partial embedding $\phi_\beta : K \to K^*$, we define $\phi_{\beta+1} := \phi_\beta$ if the first i elements of K are all in dom ϕ_β.

Else, extend $\phi_\beta : \text{dom } \phi_\beta \to \text{im } \phi_\beta$ to an isomorphism $\phi'_\beta : (\text{dom } \phi_\beta)^{\text{alg}} \cap K \to (\text{im } \phi_\beta)^{\text{alg}} \cap K^*$ by Lemma 3.4.1, and set $D := (\text{dom } \phi_\beta)^{\text{alg}} \cap K$. Take a the first element in K under the well-ordering $<_{\text{w.o.}}$ such that there is $b \in K^*$ with $\Phi_K(a/D) = \Phi_{K^*}(b/D)$, which exists via our assumption, and take $\phi_{\beta+1}$ to be the field homomorphism given by taking the identity on D and setting $\phi_{\beta+1}(a) = b$. By Lemma 3.3.2 we are given that this yields an isomorphism of $\mathcal{L}_{\mathcal{F}}$-structures $D(a) \cong D(b)$.

As before, set $\phi_\lambda = \bigcup_{\beta<\lambda} \phi_\beta$ for limit ordinals λ, and again we conclude that ϕ_α is an $\mathcal{L}_{\mathcal{F}}$-embedding of K into K^*. Thus we have extended our embedding $C \to K^*$ to an embedding $K \to K^*$, and by the Shoenfield test for quantifier elimination we have that $T_{\mathcal{F}}$ admits quantifier elimination as claimed.

Model-completeness follows automatically from quantifier elimination, completeness follows from the existence of an algebraically prime model (Proposition 3.2.1), and decidability follows as our axioms for $T_{\mathcal{F}}$ are recursively enumerable. □

3.5 Towards an unconditional result

Obviously the condition in the above theorem is unsatisfactory: we attempt to sketch here a proof that given a relatively algebraically closed subfield C and an arbitrary $a \in K$ transcendental over C we can find some $b \in K^*$ such that $\Phi_K(a/C) = \Phi_{K^*}(b/C)$, highlighting where problems arise. It is clear that if we could prove such a statement, we would not need the assumption in the above proof.

Sketch. Let $\alpha \in K$ be transcendental over C. To extend to an embedding $C \cup \{a\} \to K^*$ it is sufficient to find some $b \in K^*$ satisfying the same *atomic* formulas over C, or in other words, a realisation of the quantifier-free type $\mathrm{tp}^{\mathrm{QF}}(a/C)$. As C has cardinality at most $\aleph_0$ and K^* is $\aleph_1$-saturated, it suffices to show that every finite subset in $\mathrm{tp}^{\mathrm{QF}}(a/C)$ is consistent with K^* by Compactness. So take Σ a finite subset of this set, and let $\Psi(x)$ be the conjunction of all the formulas in Σ, so that we need to establish $K^* \models \exists x \, \Psi(x)$.

A one-variable quantifier-free sentence in $\mathcal{L}_{\mathcal{F}}$ will be equivalent to a sentence of the form $\bigwedge_{i=1}^{n} R_{f_i}(x, \overline{a}_i) \wedge \bigwedge_{j=1}^{m} \neg R_{f_j}(x, \overline{b}_j) \wedge \bigwedge_{k=1}^{s} v(F_k(x, \overline{c}_k)) \leq v(G_k(x, \overline{d}_k)) \wedge \bigwedge_{l=1}^{r} H_l(x, \overline{e}_l) = 0 \wedge \bigwedge_{q=1}^{N} J_q(x, \overline{f}_q) \neq 0$, with $\overline{a}_i, \overline{b}_i, \overline{c}_i$ etc tuples from C and polynomials f_i, F_k, H_l, J_q etc with coefficients from $\mathbb{F}_p(t)$; i.e., we get a boolean combination of polynomial equalities and inequalities, valuation equalities and inequalities, and asserting the existence or non-existence of roots of certain polynomials.

We distinguish cases according to whether x appears inside the scope of a predicate of the form R_f or not. If it does not then K^* inherits information about the R_f predicates from C and so here we need only check that we can guarantee $K^* \models \exists x \bigwedge_{k=1}^{s} v(F_k(x, \overline{c}_k)) \leq v(G_k(x, \overline{d}_k)) \wedge \bigwedge_{l=1}^{r} H_l(x, \overline{e}_l) = 0 \wedge \bigwedge_{q=1}^{N} J_q(x, \overline{f}_q) \neq 0$. As by assumption C is relatively algebraically closed inside K^* and α is transcendental over C, the H_l and J_q will give only trivial information. So we are done if we can show that K^* will always contain an element satisfying a finite collection of statements $v(g(x, \overline{d})) - v(f(x, \overline{c})) \geq 0$ whenever these are jointly satisfiable inside K.

As x is transcendental over C we cannot have $f(x, \overline{c}) = 0$ so that we will end up trying to satisfy some finite collection of statements of the form:

$$v\left(\frac{d_0 + d_1 x + d_2 x^2 + ... + d_n x^n}{c_0 + c_1 x + ... + c_m x^m}\right) \geq 0,$$

where all c_i, d_i are coefficients from C, and we know that there is a solution inside K.

First problem. This may happen two different ways inside K — first, the value of a may be outside $v(C^\times)$ and e.g. end up dominating due to its size (e.g. if $n > m$ and $v(a) >> v(C^\times)$). This should always be possible to realise inside K^* as the cofinality of $v(K^*)$ must necessarily be uncountable due to $\aleph_1$-saturation of K^*, and by assumption C is countable. So we should always be able to find large enough elements of the value group to guarantee satisfaction here. Similar arguments should deal with any other cases where the inequality is guaranteed by the value $v(a)$ (outside $v(C^\times)$). In general in such cases we will have $v(d_0 + d_1 a + ... + d_n a^n) = \min\{v(d_0), v(a) + v(d_1), 2v(a) + v(d_2), ..., nv(a) + v(d_n)\}$, etc. The second way that these statements may be satisfied is if $v(a) \in v(C^\times)$ and a satisfies some relation with the coefficients c_i, d_i to ensure that certain cancellations happen and therefore $v(d_0 + ... + d_n a^n) \geq \min\{v(d_0), ..., nv(a) + v(d_n)\}$. While it seems we can find a suitable $b \in K^*$ in practice, and such an approach seems to work concretely, the whole problem of finding $b \in K^*$ satisfying the given set of valuation inequalities requires further argument to ensure a genuine *proof.*

For the harder case, if x occurs in the scope of some R_f, further to the above, we must guarantee that for some given polynomials $f_1, ..., f_n, f_{n+1}$ with coefficients in $C \cup \{X\}$ there exist roots in K^* of $f_1, ..., f_n$ and that f_{n+1} has no root in K^*. First let us consider the case of a single polynomial — consider $R_{f_1}(X, \bar{c})$, asserting that there is some y such that $f_1(X, \bar{c}, y) = 0$, and reorder the variables of f_1 so we get a multivariable polynomial g_1 such that $R_{g_1}(\bar{c})$ if $\exists x, y \ g_1(\bar{c}, x, y) = 0$. Then $R_{g_1}(\bar{c})$ will be contained in the quantifier-free $\mathcal{L}_{\mathcal{F}}$-theory of C, so K^* will satisfy it. Thus if we take the first element of a witnessing tuple (x_0, y_0) (i.e. such that $g_1(x_0, y_0) = 0$), we will have a $b \in K^*$ such that $R_{f_1}(\overline{c}, b)$ holds. Finitely many polynomials can be reduced to this case as follows: similar to the trick $f_1 = f_2 = ... = f_n = 0$ iff $f_1^2 + ... + f_n^2 = 0$ in characteristic 0, we may let a be a non-square in $\mathbb{F}_p$ (if $p > 2$) and note that $f_1 = f_2 = 0$ iff $f_1^2 - af_2^2 = 0$: iteratively we may reduce to the case of a single polynomial. (For $p = 2$, or for a uniform approach, we may also consider the polynomial $f_1^2 - tf_2^2$, exploiting the fact that t has minimal positive value and thus cannot be a square.)

Second problem. How do we ensure that $f_{n+1}(\overline{c}, b)$ DOES NOT have a root in K^*? It may be that $g_{n+1}(\overline{c})$ does have one but if we restrict to the line $x = b$ it does not. This seems like an 'intersection of varieties' question (once we have somehow extracted irreducible varieties X_g from the relevant polynomials g), but so far we cannot see a way of guaranteeing this. (Arguments about henselian fields being large, and about the dimension of certain varieties, could provide a way forward here.)

If we can solve the above two issues, we will have that $K^* \models \exists x \Psi(x)$, and so by Compactness we have that $\mathrm{Th}(K^*) \cup \mathrm{tp}^{\mathrm{QF}}(a/C)$ is consistent, whence the quantifier-free type of a over C is realised in K^* by some b, implying $\Phi_K(a/C) = \Phi_{K^*}(b/C)$ as desired. We would then be able to extend the embedding in Theorem 3.4.2 unconditionally.

3.6 The case $\Gamma \cong \mathbb{Z}$

Suppose (K, v) is a model of our theory T with value group precisely $\mathbb{Z}$. In this special case we have some very nice theorems available.

Proposition 3.6.1. *(K, v) has unique maximal immediate extension $\mathbb{F}_p((t))$.*

Proof. If (K, v) has value group $\mathbb{Z}$ then it admits a completion $\hat{K}$, which is an immediate extension (see §2.1). We easily see that $\hat{K}$ must contain a copy of $\mathbb{F}_p((t))$, either by general theory ($\hat{K}$, as a field in characteristic $p > 0$ which is complete with respect to a henselian valuation with value group $\mathbb{Z}$ and residue field $\mathbb{F}_p$, is thereby a local field and must be isomorphic to $\mathbb{F}_p((t))$ as a pure valued field) or by explicitly constructing elements of $\mathbb{F}_p((t))$ as limits of Cauchy sequences over $\mathbb{F}_p(t)^h$, and thus $\hat{K}$ is an immediate extension of $\mathbb{F}_p((t))$. As $\mathbb{F}_p((t))$ admits no proper immediate extensions, we have $\hat{K} \cong \mathbb{F}_p((t))$ as pure valued fields, and by construction we can ensure this isomorphism preserves t. Because the diagrams of addition and multiplication in the $\mathcal{L}$-structure $\hat{K}$ are precisely the same as those in the $\mathcal{L}$-structure $\mathbb{F}_p((t))$, and by construction of T, we have that $\hat{K}$ and $\mathbb{F}_p((t))$ must agree on the interpretations of all our new predicates R_f. Therefore $\hat{K} \cong \mathbb{F}_p((t))$ as $\mathcal{L}_\mathcal{F}$-structures and we obtain that, again as $\mathcal{L}_\mathcal{F}$-structures, $(\mathbb{F}_p(t)^h, v) \leq (K, v) \leq (\mathbb{F}_p((t)), v)$. □

Proposition 3.6.2. *The class of all models of T with value group precisely $\mathbb{Z}$ is precisely the class of all intermediate fields between $\mathbb{F}_p(t)^h$ and $\mathbb{F}_p((t))$ which are relatively algebraically closed inside $\mathbb{F}_p((t))$ and which have degree of imperfection p.*

Proof. If K is a model of T with value group precisely $\mathbb{Z}$ then by the above it is a subfield of $\mathbb{F}_p((t))$, and as it is a model of T it is relatively algebraically closed inside the model $\mathbb{F}_p((t))$ of T by §3.2.

Now suppose that (K, v) is a field with $\mathbb{F}_p(t)^h \subseteq K \subseteq \mathbb{F}_p((t))$ which is relatively algebraically closed inside $\mathbb{F}_p((t))$ and which has degree of imperfection p. Then $vK = \mathbb{Z}$ and $Kv = \mathbb{F}_p$, and clearly $v(t)$ will have minimal positive value in vK. We need only show that K is extremal and that $\{1, t, t^2, ..., t^{p-1}\}$ forms a p-basis. Extremality will follow if we can show that (K, v) is henselian, algebraically maximal and defectless, as it has value group precisely $\mathbb{Z}$. Now (K, v) is closed under all Hensel polynomials as it is relatively algebraically closed inside the henselian field $\mathbb{F}_p((t))$, so that K is henselian. Suppose E/K is an immediate algebraic extension of valued fields; then E is contained in some maximal immediate extension of K, which must be isomorphic to $\mathbb{F}_p((t))$. As E is immediate over K we cannot have that this isomorphism alters the residue field in any way or makes the interpretation of t in K no longer minimal positive, so that $K \subseteq E \subseteq \mathbb{F}_p((t))$ and therefore, by assumption that K is relatively algebraically closed inside $\mathbb{F}_p((t))$, we have that E/K is trivial. Hence K is algebraically maximal.

By [Kuh01, Theorem 21] we need only show that K is inseparably defectless to get that it is defectless and therefore extremal. From [Kuh01, Lemma 17] (originally Delon [Del82]), we have that K is inseparably defectless if and only if $[K : K^p] = p$, i.e. if it has degree of imperfection p, which we have already assumed. As K is inseparably defectless, [Kuh01, Lemma 18] gives us that $\{1, t, ..., t^{p-1}\}$ forms a p-basis for K. Thus (once the interpretations of the R_f predicates are fixed in the obvious way), we see that K is indeed a model of T as claimed. □

3.7 Remark on existential decidability in $\mathcal{L}_{\mathcal{F}}$

Language is important — for instance, we know (unconditionally) existential decidability of $\mathbb{F}_p((t))$ holds in the language of pure valued fields [AF16], and (conditional on Resolution of Singularities in positive characteristic) existential decidability holds in the language $\mathcal{L}_t$ of valued fields together with a constant symbol t [DS03]. We claim the following.

Proposition 3.7.1. *Suppose the existential $\mathcal{L}_t$-theory of $T_{\mathcal{F}}$ is decidable — that is, there is a decision procedure which, given ϕ an existential $\mathcal{L}_t$-sentence, outputs YES if $T_{\mathcal{F}} \models \phi$ and which outputs NO otherwise. Then $T_{\mathcal{F}}$ is existentially decidable in $\mathcal{L}_{\mathcal{F}}$.*

Proof. It is possible to rewrite existential $\mathcal{L}_{\mathcal{F}}$-sentences in terms of $\mathcal{L}_t$, but we may end up with universal quantifiers, so that this result wouldn't be strong enough to give us existential decidability in $\mathcal{L}_{\mathcal{F}}$ without an extra argument.

Concretely, suppose $\exists \overline{x}\, \phi(\overline{x})$ is an existential $\mathcal{L}_{\mathcal{F}}$-sentence, so that ϕ is quantifier-free. Then ϕ is logically equivalent to a finite boolean combination of atomic formulas and negated atomic formulas, so that $\phi(\overline{x})$ is equivalent to a combination of formulas of the form $\tau_1 = \tau_2$, $v(\tau_3) \leq v(\tau_4)$, $R(\tau_5)$, and negations of these, where τ_i are $\mathcal{L}_{\mathcal{F}}$-terms. (It is possible to interpret all other predicates etc in this way modulo T.) It is easy to see that an $\mathcal{L}_{\mathcal{F}}$-term is simply going to be a rational function over $\mathbb{F}_p(t)$.

Now, among models of T, $K \models R_f(\tau)$ iff $K \models \exists \overline{x}\, f(\tau, \overline{x}) = 0$ so that, as every polynomial in F is defined over $\mathbb{F}_p(t)$, every *positive* existential $\mathcal{L}_{\mathcal{F}}$-formula is equivalent to an existential $\mathcal{L}_t$-formula. To translate the formula $\neg R_f(\tau)$ into an existential formula in $\mathcal{L}_t$, first take ψ the $\mathcal{L}_t$-formula $\forall x\, f(\tau, x) \neq 0$ and let T_t be the $\mathcal{L}_t$-theory obtained by taking only axioms (i)-(iv) of $T_{\mathcal{F}}$ in §3.2 (so that every model of $T_{\mathcal{F}}$ is a model of T_t). Note that if $\mathcal{M}_1 \subseteq \mathcal{M}_2$ are two models of $T_{\mathcal{F}}$ with $\mathcal{M}_1 \models \neg R_f(\tau)$, then $\mathcal{M}_1$ does not contain a root of the polynomial $f(\tau, X)$; by Proposition 3.2.2, which is proved for models of axioms (i)-(iv) only (i.e. models of T_t), then $\mathcal{M}_2$ cannot contain a root of this polynomial either, so that for every two models $\mathcal{M}_1$ and $\mathcal{M}_2$ of T_t with $\mathcal{M}_1 \subseteq \mathcal{M}_2$ and $\mathcal{M}_1 \models \psi$ we have $\mathcal{M}_2 \models \psi$. By [Hod93, Corollary 6.5.5] (a corollary of the Łoś–Tarski theorem [Hod93, Theorem 6.5.4]), we have that $\neg R_f(\tau)$ is modulo $T_{\mathcal{F}}$ equivalent to an existential $\mathcal{L}_t$-formula. □

We can thereby reduce the question of existential decidability in $\mathcal{L}_{\mathcal{F}}$ to the question of existential decidability in $\mathcal{L}_t$, and thereby reduce existential decidability in $\mathcal{L}_{\mathcal{F}}$ down to an assumption of Resolution of Singularities in positive characteristic.

Remark. To what extent would the assumption of existential decidability in $\mathcal{L}_{\mathcal{F}}$ help us to approach an unconditional proof of quantifier elimination in $T_{\mathcal{F}}$? Unfortunately it does not quite seem to provide the argument we need: we would be able to prove, where K^* is our $\aleph_1$-saturated model and C' is some subfield of K^* elementarily equivalent to C in $\mathcal{L}_{\mathcal{F}}$, that K^* contains some realisation b' of $\mathrm{tp}^{\mathrm{QF}}(a/C)$, interpreting constant symbols for elements of C by corresponding elements of C', but because we do not know (and certainly cannot assume!) that this type is *complete*, we cannot use κ^+-homogeneity of K^* to yield an automorphism of K^* sending C' to C and b' to some $b \in K^*$ genuinely satisfying the algebraic configuration of a *over* C. This may seem like a subtle point but it appears to be of critical importance.

Chapter 4

Tools for further study

4.1 Pseudomaximal fields

As we see with large fields, sometimes in the model theory of fields we can fruitfully formulate and use definitions which involve a field K containing rational points of a K-variety if some larger field contains rational points of that K-variety. Large fields turning up in theorems around extremality provide one example of this, but there are many others, e.g. the theories of pseudo algebraically closed (PAC) fields, pseudo real closed (PRC) fields, pseudo p-adically closed (PpC) fields, etc. See the excellent [FJ08] for a comprehensive overview of many such examples.

We introduce a notion which to our knowledge has not previously been presented in this form, although it has been previously studied.

Definition 4.1.1. *Let (K, v) be a nontrivially valued field. We will call (K, v)* ***pseudomaximal*** *if for every absolutely irreducible variety X defined over K, if X has a K'-rational point in every maximal immediate extension (K', v') of (K, v), then X has a K-rational point.*

Standard techniques give us the following; the proof should work *mutatis mutandis* similarly to the case for PRC fields shown in [Pre81]. (We credit [Mon17] as the paper which first gave us the idea of reformulating this concept in model-theoretic terms.)

Theorem 4.1.2. *Let (K, v) be a valued field. Then (K, v) is pseudomaximal if and only if (K, v) is existentially closed in each of its maximal immediate extensions, considered as first-order structures in the language of valued fields.*

Remark: Existential closedness in maximal immediate extensions has been previously touched on in the literature; see [Kuh90, §10], particularly with reference to **(Mod 2)** (K is existentially closed in every immediate extension), **(Mod 1)** (K is existentially closed in at least one maximal immediate extension) and Lemma 10.2 (where it is shown that the property **(Mod 1)** implies algebraic completeness).

We have already found our first nontrivial example of a pseudomaximal field, as $\mathbb{F}_p(t)^h$ is existentially closed in its unique (up to isomorphism) maximal immediate extension $\mathbb{F}_p((t))$.

Conjecture 4.1.3. *Suppose (K, v) is a pseudomaximal valued field which admits a maximal immediate extension (M, v) which is extremal. Then (K, v) is extremal.*

Our attempt to prove this statement ran aground. Our strategy was to let $f \in K[X_1, ..., X_n]$ be a polynomial and try to show that $\{v(f(\overline{x})) \mid \overline{x} \in \mathcal{O}_K^n\}$ contains a maximal element, by somehow extracting an absolutely irreducible variety from f; unfortunately, f may not have a root in all maximal immediate extensions and so we cannot necessarily take a witnessing element from K of the maximal value achieved in a maximal immediate extension M and show that this implies f has that maximal value over $\mathcal{O}_K$, as we would not be guaranteed that our variety has a K-rational point.

Despite these setbacks, pseudomaximality seems like a natural notion which would allow us to 'pull results down' from maximal valued fields to more general valued fields, and we hope that it may yet bear fruit. In particular we feel that Conjecture 4.1.3 should be within reach once extremality is a little better understood.

4.2 n-extremality

Recall that a field (K, v) with valuation ring $\mathcal{O}$ is said to be **extremal** if, for every $n \in \mathbb{N}$ and every $f \in K[X_1, ..., X_n]$, the image of $v \circ f$ on $\mathcal{O}^n$ has a maximal element in $vK \cup \{\infty\}$. Let us form the following definition:

Definition 4.2.1. Let $n \in \mathbb{N}$ (≥ 1). We call a valued field (K, v) with valuation ring $\mathcal{O}$ ***n-extremal*** if, for every $f \in K[X_1, ..., X_n]$, the image of $v \circ f$ on $\mathcal{O}^n$ has a maximal element in $vK \cup \{\infty\}$.

Thus a valued field is extremal iff it is n-extremal for all $n \in \mathbb{N}$.

This definition has several advantages: first, it may be much easier to prove that a field is e.g. 2-extremal than it is to prove extremality as a whole. Indeed, the behaviour of curves over a valued field (K, v) with respect to its valuation seems like a natural question for those interested in valuation theory and algebraic geometry.

Secondly, once we have the concept of n-extremality in mind, a natural question presents itself: how much extremality do we actually require to guarantee 'the whole lot'? Concretely, and again by analogy with the cases of PAC, PRC and PpC fields:

Conjecture 4.2.2. *Let (K, v) be a valued field. Then (K, v) is extremal if and only if (K, v) is 2-extremal.*

We pursued this idea but had difficulties reconciling the different topologies involved (the Zariski topology and the topology induced by the ordering on vK). Our idea was as follows:

one direction of the implication is trivial, so for the other direction, suppose (K, v) is *not* extremal, and let $f \in K[X_1, ..., X_n]$ be a polynomial witnessing this fact. Consider now the family of polynomials $\{f(X_1, ..., X_n) - c \mid c \in K\}$; the failure of extremality means that there will be a subfamily $\{g(X_1, ..., X_n) \mid g \in \mathfrak{G}\}$, where $\mathfrak{G}$ consists precisely of those polynomials $f - c$ which have roots in $\mathcal{O}$, and the $v(c)$s will witness the failure of extremality, i.e. there will not be an attained upper bound. If necessary we can multiply through by constants at this stage to ensure that every g is defined over $\mathcal{O}$ and not just over K. From here one attempts to obtain an absolutely irreducible variety X_g from each g, and, following the analogous result for PAC fields (see [FJ08, §11.2]), show that each such variety must contain an absolutely irreducible curve; the fact that henselian fields are large should provide the required arguments on the dimension of X_g and the existence of rational points, and we would thus obtain a family of curves (say, $C_g \subseteq X_g$) whose K-rational points witness the failure of extremality.

The problems with this approach are as follows: first, it requires us to reconstruct a fair amount of algebraic geometry over $\mathcal{O}$ rather than over a general field K. In particular it seems that the key statement required is an analogue of the Matsusaka–Zariski theorem (Proposition 10.5.2 in [FJ08]), and that things should be manageable if one can prove that over some suitable family of valuation rings. At this point the theory of schemes became involved and we felt that it was probably beyond the scope of this project.

Secondly, and most crucially to our mind, rational points on our curve C_g are still going to be elements $(x_1, ..., x_n) \in \mathcal{O}^n$; one also needs a way to go from C_g to a curve writable in two variables, defined over $\mathcal{O}$, so that a rational point will genuinely be an element $(x, y) \in \mathcal{O}^2$. We know that C_g will be birationally equivalent to such a curve (assuming we can recreate the relevant algebraic geometry over $\mathcal{O}$), but we do not know if, in practice, this birational equivalence will preserve the property that the $v(c)$s witness the counterexample to extremality. Therefore we may have no idea, when considering the family of all curves C_g over $\mathcal{O}$, whether the structure of $\mathcal{O}$-rational points of C_g actually informs us about the values of the polynomial f over $\mathcal{O}$.

The problem in a nutshell is that our maps will preserve the Zariski topology but we have no reason to believe they preserve the topology induced by the order topology on vK, and the existence of two different topologies here is making things much more complicated than we felt able to tackle as part of the present work.

(There is also the matter of obtaining an *absolutely irreducible* variety from each g, but we felt this was less significant than the problems outlined above.)

4.3 Ultraproducts of generalised Laurent series fields

Write $\mathbb{P}$ for the set of all prime numbers, and let $\mathcal{U}$ be a non-principal ultrafilter on $\mathbb{P}$. Now let Γ be a $\mathbb{Z}$-group and consider the family of generalised Laurent series fields $\{\mathbb{F}_p((\Gamma)) : p \in \mathbb{P}\}$. This is a family of $\mathcal{L}_{\mathcal{F}}$-structures and we can thus form the ultraproduct $\Pi(\mathcal{U}, \Gamma) := \prod_{\mathcal{U}} \mathbb{F}_p((\Gamma))$. The following can easily be seen to hold true of $\Pi(\mathcal{U}, \Gamma)$:

- $\Pi(\mathcal{U}, \Gamma)$ is a henselian valued field of residue characteristic 0:

  As we have a symbol $\mathcal{O}$ for the valuation in our language and all our generalised Laurent series fields satisfy $\forall x\, (x \in \mathcal{O} \vee 1/x \in \mathcal{O})$ we have that $\Pi(\mathcal{U}, \Gamma)$ is a valued field. Each $\mathbb{F}_p((\Gamma))$ is henselian so (via Hensel's Lemma) satisfies $\forall a_0, ..., a_{n-2} \in \mathcal{M}\ \exists x\, (x^n + x^{n-1} + a_{n-2}x^{n-2} + ... + a_0 = 0)$ for every $n \in \mathbb{N}$, from which we see that $\Pi(\mathcal{U}, \Gamma)$ does likewise. And as we can write the residue field uniformly as $\mathcal{O}/\mathcal{M}$ without reference to p, and we have $1 + ... + 1$ (p times) $\notin \mathcal{M}$ for all but one $\mathbb{F}_p((\Gamma))$, we see that $\Pi(\mathcal{U}, \Gamma)$

has residue characteristic 0.

- $v(\Pi(\mathcal{U},\Gamma))$ is a $\mathbb{Z}$-group with $v(t) = 1$:

  As we have a symbol for $\mathcal{O}$ in our language we can express the sentence "$v(t)$ is minimal positive" without reference to p via any sentence interpreting $\neg\exists x\ v(1) < v(x) < v(t)$. Similarly, one can express the fact that Γ is a regularly ordered group via the same infinite schema of sentences in all models without reference to p. This means we have these sentences holding in all $\mathbb{F}_p((\Gamma))$, whence they necessarily hold in the ultraproduct, giving us our result directly.

- $\Pi(\mathcal{U},\Gamma)$ is extremal:

  $\Pi(\mathcal{U},\Gamma)$ is a henselian field of residue characteristic 0. It is thus automatically defectless, and so cannot admit any proper immediate algebraic extensions, whence it is henselian, defectless and algebraically maximal. As its value group is a $\mathbb{Z}$-group we have immediately that $\Pi(\mathcal{U},\Gamma)$ is extremal.

Unfortunately as there are infinitely many formulas involved in defining extremality in first-order terms, we cannot directly conclude that $\mathbb{F}_p((\Gamma))$ is extremal for almost all p. However, we can obtain some results nonetheless:

Proposition 4.3.1. *Fix $m, n \in \mathbb{N}$, and let Γ be any $\mathbb{Z}$-group. Then for almost all primes p, the following statement holds: $\mathbb{F}_p((\Gamma))$ is n-extremal with respect to all multivariable polynomials of order $\leq m$.*

Proof. If $f \in K[X_1, ..., X_n]$ has order $\leq m$ then it has only finitely many terms and finitely many coefficients from K, so that where $\overline{a}$ is a tuple of all the coefficients appearing in f and we set $g(x_1, ..., x_n, \overline{y}) := f(x_1, ..., x_n)[\overline{y}/\overline{a}]$ (i.e., replacing each coefficient a_i appearing in f with a variable y_i), one can simply use the sentence

$$\forall\overline{y}\ \exists z_1, ..., z_n \in \mathcal{O}\ \forall x_1, ..., x_n \in \mathcal{O}\ v(g(x_1, ..., x_n, \overline{y})) \leq v(g(z_1, ..., z_n, \overline{y})).$$

This is satisfied in the ultraproduct and therefore must be satisfied by $\mathbb{F}_p((\Gamma))$ for almost all p. □

This allows us to conclude extremality with respect to certain classes of polynomials: for instance, we obtain that $\mathbb{F}_p((\Gamma))$ is extremal with respect to all elliptic curves, for almost all p. Similarly, we can say that $\mathbb{F}_p((\Gamma))$ is $\mathbb{F}_p((\Gamma))$-extremal with respect to all p-polynomials of order $\leq m$ (in the terminology of [AK16, §3]) for almost all p. Unfortunately we cannot seem to dispense with the finite m here, and we have not succeeded in our attempts at a direct proof of n-extremality even for the field $\mathbb{F}_p((t))((\mathbb{Q}))$.

We note that we have not used many particular properties of the fields $\mathbb{F}_p((\Gamma))$ here, so that this argument applies much more generally: indeed our argument holds for any collection of henselian fields such that the residue characteristic of the ultraproduct ends up being 0, and the value group ends up a $\mathbb{Z}$-group.

While we have not succeeded in proving that each $\mathbb{F}_p((\Gamma))$ is a model of T where Γ is an arbitrary $\mathbb{Z}$-group, a simple Compactness argument yields the following result:

Proposition 4.3.2. *Let Γ be a $\mathbb{Z}$-group. Then $\mathbb{F}_p((\Gamma))$ is contained in some model of T.*

Proof. It suffices to show that we can embed $\mathbb{F}_p((\Gamma))$ as an $\mathcal{L}_t$-structure in some model of T. Extend $\mathcal{L}_t$ by adding constant symbols c_s for every sequence $s \subseteq \Gamma \times \mathbb{F}_p$ such that the support of s, i.e. the set $\{\gamma \in \Gamma \mid \exists n \in \mathbb{F}_p^{\times}\ (\gamma, n) \in s\}$, is well-ordered. The intention is to have every element of $\mathbb{F}_p((\Gamma))$ encoded by such a sequence s and interpreted by c_s. Add to T the following axioms:

- $c_s = \sum_s a_i t^i$ for all finite sums $\sum_s a_i t^i$ in $\mathbb{F}_p((t))$, where $s \subseteq \mathbb{Z} \times \mathbb{F}_p$ and $a_i \in \mathbb{F}_p$ such that a nonzero coefficient a_i appears in the sum only if $(i, a_i) \in s$.

- $c_{s_1} + c_{s_2} = c_{s_1+s_2}$ and $c_{s_1} \cdot c_{s_2} = c_M$ for all sequences s_1, s_2, where M is the sequence associated with multiplying together the elements represented by c_{s_1} and c_{s_2} in $\mathbb{F}_p((\Gamma))$.

- $c_s \in \mathcal{O}$ whenever the first element of the support of s is non-negative, and $v(c_s) = v(t^n)$ whenever the first element of the support of s is $n \in \mathbb{Z}$.

Denote this expanded set of axioms by T^+. We claim that $\mathbb{F}_p((t))$ is a model for any finite subset Σ_0 of T^+. Indeed, let $s_1, ..., s_n$ be all sequences appearing in constant symbols in sentences of Σ_0; map c_{s_i} to the relevant finite sum in $\mathbb{F}_p((t))$ wherever $\sum_s a_i t^i$ is a finite sum in $\mathbb{F}_p((t))$, and let N be the highest power of t appearing in any of these sums or any other term appearing in Σ_0. By inductively choosing, as an interpretation for c_{s_j} (where s_j is the first sequence from $s_1, ..., s_n$ not already mapped to an element of $\mathbb{F}_p((t))$) the sum $\sum_{k=1}^{N} a_k t^k$ plus some arbitrarily large power of t (which we then take as our new N), it is straightforward to pick a finite set of elements of $\mathbb{F}_p((t))$ satisfying all (quantifier-free first-order in $\mathcal{L}_t$) constraints in Σ_0, by simply choosing larger and larger powers of t. (Indeed, by choice of axioms, we need only really concern ourselves with making sure that every relation of the form $c_{s_1} + c_{s_2} = c_{s_1+s_2}$ and $c_{s_1} \cdot c_{s_2} = c_M$ is preserved.)

Therefore T^+ is satisfied by some model $\mathcal{M}$. But our additional axioms guarantee that the set $\{x \in \mathcal{M} \mid \exists s \; x = c_s\}$ is isomorphic as a valued field to $\mathbb{F}_p((\Gamma))$. This proves the statement. □

4.4 Extremal closures

Let Σ be an $\mathcal{L}$-theory (we use Σ instead of T to avoid confusion with $T_{\mathcal{F}}$). Recall the notion of Σ-closure introduced in §2.2.1 in van den Dries' test for quantifier elimination (Theorem 2.2.8): if $\mathcal{A}$ is a substructure of a model $\mathcal{B}$ of Σ, then a Σ-***closure*** of $\mathcal{A}$ is a model $\mathcal{A}^*$ of Σ which embeds over $\mathcal{A}$ in every model of Σ containing $\mathcal{A}$. In every first-order theory of valued fields which admits quantifier elimination (which, we hope, includes $T_{\mathcal{F}}$) then we should expect some analogous notion to make sense. To that end we formulate the following attempted definition.

Definition 4.4.1. *Fix $\mathcal{L}$ a first-order language extending the language of valued fields. Let Γ be an ordered abelian group, k a field, $p = \mathrm{char}(k)$ either 0 or a positive prime, and q equal either to 0 or p, such that the following class of fields is nonempty:*

$$\boldsymbol{E} = \{(K, v) \mid (K, v) \text{ is an extremal valued field with } \mathrm{char}K = q,\ vK \equiv \Gamma, \text{ and } Kv \equiv k\}.$$

*This is an elementary class of fields; let $\mathcal{B} \in \boldsymbol{E}$ be an $\mathcal{L}$-structure and let $\mathcal{A} \subseteq \mathcal{B}$ be some substructure. We will call $\mathcal{A}^*$ an **extremal closure** of $\mathcal{A}$ if we have both:*

(i) $\mathcal{A} \subseteq \mathcal{A}^* \in \boldsymbol{E}$,

(ii) for every $\mathcal{K}$ in $\boldsymbol{E}$ containing $\mathcal{A}$, there is an embedding of $\mathcal{A}^*$ into $\mathcal{K}$ over $\mathcal{A}$.

The first thing to note is that we say *an* extremal closure rather than *the* extremal closure because we do not *a priori* require that this embedding be unique, which would be needed to imply a universal property of extremal closures and thereby guaranteeing uniqueness up to valuation-preserving isomorphism. In particular, we suspect that the non-uniqueness of algebraically maximal immediate extensions of fields in positive characteristic may provide examples where such a closure, if it exists, will not be unique.

In which classes **E** does this provide a useful notion, and furthermore, in which classes **E** are extremal closures unique up to valuation-preserving isomorphism? It seems to us that this question is wide open, and other than a few easy cases (e.g. if $vK \equiv \mathbb{Z}$ and Kv is of characteristic 0 then the extremal closure of K will be unique and will coincide with the henselization of K, as fields of residue characteristic 0 are automatically defectless and so

(K^h, v^h) will be extremal directly), it is not even clear how one would concretely obtain an extremal closure of a given field.

If we can prove that $T_{\mathcal{F}}$ admits extremal closures, then this is close to a proof of quantifier elimination via the Van den Dries-style test (Theorem 2.2.8) — the notion of 'extremal closure' would work towards satisfying condition (i), and our notion of 'algebraic configuration' can be seen as an attempt to capture condition (ii).

Unfortunately we cannot see a clear path towards constructing an extremal closure of a given valued field. One might seek to approach the problem by attempting to construct new elements witnessing extremality, by taking pseudolimits of some pseudo-Cauchy sequence of elements $a_0, a_1, a_2, ...$ such that the values $f(a_0), f(a_1), f(a_2), ...$ approach the maximum value achieved on $\mathcal{O}$; however, this approach could only work for polynomials in a single variable and thus will not be sufficient for extremality, which involves multivariable polynomials. Even in the case of 2-extremality this would cause issues, because for a 2-variable polynomial $f(X, Y)$ there would be no *a priori* guarantee that the pseudoconvergence of the sequence $f(a_0, b_0), f(a_1, b_1), f(a_2, b_2), ...$ would imply the pseudoconvergence of the sequences $a_0, a_1, a_2, ...$ and $b_0, b_1, b_2, ...$. As such, we cannot see a clear way to begin constructing an extremal closure of a given valued field.

Chapter 5

Extensions with dense value group

We turn now to extensions of $\mathbb{F}_p((t))$ whose value group is no longer a $\mathbb{Z}$-group. These fields will no longer be discretely extremal and thus we should expect their model theory to be altogether different from that of $\mathbb{F}_p((t))$. For all of the extensions we consider, we will keep the residue field as $\mathbb{F}_p$ throughout.

Let us first consider the perfect hull $\mathbb{F}_p((t))^{\text{perf}}$ of $\mathbb{F}_p((t))$ — this is the smallest perfect field containing $\mathbb{F}_p((t))$, or, equivalently, the maximal purely inseparable extension of $\mathbb{F}_p((t))$. We have that $\mathbb{F}_p((t))^{\text{perf}}$ consists of precisely those elements x in the algebraic closure of $\mathbb{F}_p((t))$ such that for some $k \in \mathbb{N}$ we have $x^{p^k} \in \mathbb{F}_p((t))$; one sees easily that such an element will be a Laurent series over $\mathbb{F}_p$ in some t^{1/p^k}, and therefore that $\mathbb{F}_p((t))^{\text{perf}} = \bigcup_{k=0}^{\infty} \mathbb{F}_p((t^{1/p^k}))$. This is an algebraic extension of $\mathbb{F}_p((t))$ and so the henselian valuation there extends uniquely to a henselian valuation on $\mathbb{F}_p((t))^{\text{perf}}$ by setting $v(t^{1/p^k}) = 1/p^k$ in the obvious way. We thus end up with a valued field of characteristic $p > 0$ with value group $p^{-\infty}\mathbb{Z} := \bigcup_{k=0}^{\infty} p^{-k}\,\mathbb{Z}$, the p-divisible hull of $\mathbb{Z}$, and residue field $\mathbb{F}_p$.

Whereas $\mathbb{F}_p((t))$ is complete (in an analytic sense), $\mathbb{F}_p((t))^{\text{perf}}$ is not: consider for example the Cauchy sequence given by setting $a_n := \sum_{k=1}^{n} t^{k-(1/p^k)}$ for each $n \in \mathbb{N}$. By construction this sequence can have no limit in any $\mathbb{F}_p((t^{1/p^k}))$ and therefore has no limit in $\mathbb{F}_p((t))^{\mathrm{perf}}$, and the sequence is clearly Cauchy because $v(a_{n+1}-a_n) \to \infty$ as $n \to \infty$. (This translates to $|a_{n+1}-a_n| \to 0$ as $n \to \infty$ via the identification $|x| = e^{-v(x)}$.)

We know that $\mathbb{F}_p((t))^{\mathrm{perf}}$ admits a completion, and that this completion will be an immediate extension of valued fields (see preliminaries). Furthermore we know that $\mathbb{F}_p((t))^{\mathrm{perf}} \subseteq \mathbb{F}_p((p^{-\infty}\mathbb{Z}))$ and that this latter field is maximal, admitting no proper immediate extensions; it is therefore complete, and is a suitable ambient space to consider the question of the completion of $\mathbb{F}_p((t))^{\mathrm{perf}}$.

Let us denote $K := \mathbb{F}_p((t))^{\mathrm{perf}}$, $\widehat{K}$ the completion of K, and $F := \mathbb{F}_p((p^{-\infty}\mathbb{Z}))$: then we have $K \subseteq \widehat{K} \subseteq F$, with F/K an immediate extension of valued fields (thereby implying that the extension $F/\widehat{K}$ is also immediate). We first note a simple result.

Proposition 5.1. *The fields K, $\widehat{K}$ and F are all henselian.*

Proof. K is an algebraic extension of the henselian field $\mathbb{F}_p((t))$ and is therefore henselian. Henselianity of $\widehat{K}$ follows from [EP05, Theorem 1.3.1]. As a generalised Laurent series field, F is automatically henselian, as per preliminaries. □

We next seek to characterise $\widehat{K}$ inside F.

Given an element α of any generalised Laurent series field, the support $\mathrm{supp}(\alpha)$ is a well-ordered set; define the ***order-type*** of an element $\alpha \in F$ to be the unique ordinal number that is order-isomorphic to $\mathrm{supp}(\alpha)$, giving us a function $\mathrm{ot} : F \to \mathbf{ON}$.

Further define $N : F \to \mathbb{N} \cup \{\infty\}$ by setting $N(\alpha)$ to be the least $n \in \mathbb{N}$ such that $\mathrm{supp}(\alpha) \subseteq p^{-n}\mathbb{Z}$, with $N(\alpha) = \infty$ if there is no such $n \in \mathbb{N}$. (This is well-defined as we know that $\operatorname{supp}(\alpha) \subseteq \bigcup_{n=0}^{\infty} p^{-n}\mathbb{Z}$.)

We can now provide the following characterisation.

Proposition 5.2. *Let $\alpha \in F$. We distinguish the following cases:*

(i) If $\operatorname{ot}(\alpha) < \omega$ then $\alpha \in K$.

(ii) If $\operatorname{ot}(\alpha) > \omega$ then $\alpha \in F \setminus \widehat{K}$.

(iii) If $\operatorname{ot}(\alpha) = \omega$ and $\operatorname{supp}(\alpha)$ is bounded above in $p^{-\infty}\mathbb{Z}$, then $\alpha \in F \setminus \widehat{K}$.

(iv) If $\operatorname{ot}(\alpha) = \omega$, $\operatorname{supp}(\alpha)$ is unbounded above in $p^{-\infty}\mathbb{Z}$, and $N(\alpha) < \infty$, then $\alpha \in K$.

(v) Else, $\alpha \in \widehat{K} \setminus K$.

Proof. Case (i) is clear.

For case (ii), we have $\operatorname{ot}(\alpha) \geq \omega + 1$, so where γ_1 is the first and $\gamma_{\omega+1}$ is the $\omega + 1$th element of $\operatorname{supp}(\alpha)$, we have that there are infinitely many elements of $\operatorname{supp}(\alpha)$ between γ_1 and $\gamma_{\omega+1}$. But for every finite n there are only finitely many elements of $p^{-n}\mathbb{Z}$ between γ_1 and $\gamma_{\omega+1}$, so if α were the limit of a Cauchy sequence in $K = \bigcup_{n\in\mathbb{N}} \mathbb{F}_p((t^{1/p^n}))$, no element of the Cauchy sequence would agree with α up to $\gamma_{\omega+1}$: that is, if $a_n \to \alpha$ as $n \to \infty$, we have $v(a_n - a) < \gamma_{\omega+1}$ for all n. So certainly $v(a_n - \alpha) \not\to \infty$ as $n \to \infty$. But then $|a_n - \alpha| \not\to 0$ as $n \to \infty$, contradicting our assumption that α is a limit of the Cauchy sequence (a_n).

This generalises to give us case (iii) by replacing $\gamma_{\omega+1}$ with any upper bound for $\operatorname{supp}(\alpha)$.

In case (iv) we have $\alpha \in \mathbb{F}_p((t^{1/p^n}))$ directly for some n, by definition of $N(\alpha)$, whence $\alpha \in K$.

Finally, we need to show that if $\mathrm{ot}(\alpha) = \omega$, $\mathrm{supp}(\alpha)$ is unbounded above and $N(\alpha) = \infty$ then α is genuinely an element of $\widehat{K} \backslash K$. Now $N(\alpha) = \infty$ implies that α cannot be contained in any $\mathbb{F}_p((t^{1/p^n}))$ and therefore $\alpha \notin K$. Let $\{\gamma_n : n \in \mathbb{N}\}$ be the canonical enumeration of $\mathrm{supp}(\alpha)$ compatible with the ordering of $p^{-\infty}\mathbb{Z}$ (i.e. $\gamma_n < \gamma_m$ iff $n < m$), so that $\alpha = \sum_n a_n t^{\gamma_n}$. We claim that the partial sums $s_k := \sum_{n=0}^{k} a_n t^{\gamma_n}$ form a Cauchy sequence in K converging to α.

Clearly every s_k, being a finite sum, must already be contained in some $\mathbb{F}_p((t^{1/p^n}))$ directly, so that each element of the sequence is already in K. We have $v(s_{k+1} - s_k) = \gamma_{k+1} \to \infty$ as $k \to \infty$, this being equivalent to the statement that $\mathrm{supp}(\alpha)$ is unbounded above, so our sequence is indeed Cauchy. And clearly $v(\alpha - s_k) = \gamma_{k+1}$ for all k, whence α is indeed a limit of the Cauchy sequence $(s_k)_{k \in \mathbb{N}}$, as desired.

This proves the result for all cases. $\square$

We can see immediately that $\widehat{K}$ is perfect, via the rule $\sqrt[p]{\sum_\gamma a_\gamma t^\gamma} = \sum_\gamma \sqrt[p]{a_\gamma}\, t^{\gamma/p}$. We see easily that K is relatively algebraically closed inside $\widehat{K}$, as every algebraic extension of a henselian field within its completion is purely inseparable (see e.g. [Kuh90, Lemma 7.1]), and as K is perfect it admits no (proper) purely inseparable algebraic extensions.

We note that there are proper algebraic extensions of $\widehat{K}$ inside F: indeed, consider $s := \sum_{k=1}^{\infty} t^{-1/p^k} \in F$. We see that $s \notin \widehat{K}$ as its support has order-type ω and is bounded above by 0. Calculating $s^p = \sum_{k=1}^{\infty} (t^{-1/p^k})^p = \sum_{k=1}^{\infty} t^{-1/p^{k-1}}$, we see that $s^p - s = t^{-1}$ and so this gives an algebraic purely wild extension of $\widehat{K}$ inside F. Let us therefore denote by $F_1 := F \cap \mathbb{F}_p((t))^{\mathrm{alg}}$ the relative algebraic closure of $\mathbb{F}_p((t))$ inside F, and $F_2 := F \cap \widehat{K}^{\mathrm{alg}}$ the relative algebraic closure of $\widehat{K}$ inside F.

We next turn to the Galois theory of these fields. Recall that the absolute Galois group of a field is the same as that of its perfect hull (see [Jar96, §1]), so that $G_K \cong G_{\mathbb{F}_p}((t))$ via the canonical restriction map. (For a concrete description of the absolute Galois group of $\mathbb{F}_p((t))$, see [EF99].)

Proposition 5.3. *The absolute Galois groups $G_{\widehat{K}}$ and G_K are isomorphic via the canonical restriction map* $\mathrm{res} : G_{\widehat{K}} \to G_K$.

Proof. As K is relatively algebraically closed inside $\widehat{K}$, we have that $K^{\mathrm{sep}} \cap \widehat{K} = K$. As K^{sep} is a Galois extension of K, [Cha, §1.3] gives us immediately that the restriction map $\mathrm{Gal}(K^{\mathrm{sep}}\widehat{K}/\widehat{K}) \to \mathrm{Gal}(K^{\mathrm{sep}}/K)$ is an isomorphism of profinite groups. But the group on the right-hand side is just G_K, and the field composite $K^{\mathrm{sep}}\widehat{K}$ is equal to $(\widehat{K})^{\mathrm{sep}}$ via a simple corollary of Krasner's Lemma (see [Efr06, Corollary 18.5.3]). □

Proposition 5.4. *The canonical restriction maps* $\mathrm{res} : G_{F_1} \to G_K$, $\mathrm{res} : G_{F_2} \to \widehat{K}$ *are* $1-1$ *but not onto.*

Proof. By the definition of 'absolute Galois group' and the fact that F_1/K, $F_2/\widehat{K}$ are both algebraic extensions, the restriction maps are necessarily injective: two automorphisms of $F_1^{\mathrm{sep}} = K^{\mathrm{sep}}$ (resp. $F_2^{\mathrm{sep}} = \widehat{K}^{\mathrm{sep}}$) which are fixed pointwise on F_1 (resp. F_2) have identical restrictions in G_K (resp. $G_{\widehat{K}}$) if they take precisely the same values as automorphisms of K^{sep} fixed pointwise on K (resp. automorphisms of $\widehat{K}^{\mathrm{sep}}$ fixed pointwise on $\widehat{K}$), but this happens just when they are identical. Noting that F_1/K (resp. $F_2/\widehat{K}$) is a separable extension, [Cha, § 1.12] gives that the restriction map is onto iff $F_1 \cap K^{\mathrm{sep}} = K$ (resp. iff $F_2 \cap \widehat{K}^{\mathrm{sep}} = \widehat{K}$), but this is clearly not the case due to the existence of algebraic elements in $F \setminus K$ (resp. $F \setminus \widehat{K}$). □

We know that F is a tame valued field (see preliminaries), and that these fields have particularly nice model-theoretic properties. In particular their first-order theory is decidable if and only if the first-order theories of their value group and residue field are decidable.

Furthermore we know that a valued field in positive characteristic is tame if and only if it is perfect and algebraically maximal; thus if we demonstrate that F_1 and F_2 are algebraically maximal we will have a number of useful results immediately available to us.

Let L/F_1 be an immediate algebraic extension. (The case F_2 works identically.) Now as F_1 is relatively algebraically closed in F, either $L \subseteq F$ already (and therefore L/F_1 is trivial), or else L is linearly disjoint from F. Then $[L : F_1] = [L.F : F] = p^n$ for some n, as L/F_1 is nontrivial immediate and must therefore be a defect extension. But F is tame and admits no nontrivial defect extensions, so that $L.F/F$ must be a residue field extension (F does not admit any value group extension of order p^n as its value group is p-divisible). But then $(L.F)v = \mathbb{F}_{p^n}$, and therefore $L.F \setminus F$ contains some $(p^n - 1)$th root of unity. But by assumption L is linearly disjoint from F, so that this root of unity must already exist in L, contradicting our assumption that L/F_1 is immediate.

From this, taking our preliminaries on the model theory of tame valued fields into account, we obtain:

- F_1, F_2 are tame fields, and F_1 is the smallest tame field containing $\mathbb{F}_p((t))$,
- the fields F_1, F_2 and F are all elementarily equivalent,
- F_1 and F_2 are both existentially closed in F, and F_1 is existentially closed in F_2,
- the first-order theory of F_1 is decidable.

Therefore we see that decidability is obtainable in extensions of $\mathbb{F}_p((t))$ once we've overcome the issues posed by the fact that $\mathbb{F}_p((t))$ is not perfect and that there are defect extensions lying between $\mathbb{F}_p((t))^{\text{perf}}$ and $\mathbb{F}_p((p^{-\infty}\mathbb{Z}))$.

Chapter 6

Brief concluding remarks

We have, in Theorem 3.4.2, a conditional proof of quantifier elimination of $T_{\mathcal{F}}$, and the existence of an algebraically prime model gives us a conditional decidability result for $\mathbb{F}_p((t))$. The required condition is, however, something that appears difficult to tackle, and we have highlighted a couple of areas where an unconditional proof seems to run aground. Nevertheless, the theory $T_{\mathcal{F}}$ certainly seems like a good candidate for further investigation, and we have a complete characterisation of its models with value group precisely $\mathbb{Z}$. Furthermore we have that existential decidability of $T_{\mathcal{F}}$ would be implied by Resolution of Singularities in positive characteristic.

We note that our notion of *algebraic configuration* would exactly correspond to the second part of van den Dries' quantifier elimination test, and note that the first part, the existence of some form of extremal closure for T, seems to be readily supplanted by being able to extend to the relative algebraic closure of a subfield. We suggest that this is because extremality is an *intrinsic* property of a valued field which can be 'read off' from the graph of the map $v : f(\mathcal{O}^n) \to vK$ where f is a multivariable polynomial; in essence, because we know that K is extremal, we must have that its image in K^* *becomes* extremal once we have added enough elements.

We have introduced some notions to the study of this area and done our best to obtain results involving them, with varying degrees of success. The formulation in terms of algebraic configurations, in particular, allows us to specify precisely what is needed to guarantee decidability.

There have been several developments around this area recently, not all of which have been possible to incorporate or internalise in the time available, such as the notion of Hensel minimality developed in [CHRK19], of which we first became aware at a workshop in Oberwolfach in October 2016, where it was presented under the name 'resplendent minimality' (which we prefer as being more poetic).

We hope to see the model theory of valued fields, particularly in this area, continue its recent spell of significant progress over the coming years. One way or another we are optimistic that the question of whether or not the first-order theory of $\mathbb{F}_p((t))$ is decidable is a question which will soon be answered.

Bibliography

[AF16] Sylvy Anscombe and Arno Fehm. The existential theory of equicharacteristic henselian valued fields. *Algebra & Number Theory*, 10(3):665–683, 2016.

[AK16] Sylvy Anscombe and Franz-Viktor Kuhlmann. Notes on extremal and tame valued fields. *The Journal of Symbolic Logic*, 81(2):400–416, 2016.

[AKP12] Salih Azgin, Franz-Viktor Kuhlmann, and Florian Pop. Characterization of extremal valued fields. In *Proc. Amer. Math. Soc*, volume 43, pages 1535–1547, 2012.

[BJW93] P Bateman, C Jockush, and A Woods. Decidability and undecidability of theories with a predicate for the primes. *Journal of Symbolic Logic*, 58(2):672, 1993.

[BK17] Anna Blaszczok and Franz-Viktor Kuhlmann. On maximal immediate extensions of valued fields. *Mathematische Nachrichten*, 290(1):7–18, 2017.

[Bou13] Nicolas Bourbaki. *Algebra II: Chapters 4-7*. Springer Science & Business Media, 2013. reprint of 1990 translation of 1981 original.

[CDLM13] Raf Cluckers, Jamshid Derakhshan, Eva Leenknegt, and Angus Macintyre. Uniformly defining valuation rings in henselian valued fields with finite or pseudo-finite residue fields. *Annals of Pure and Applied Logic*, 164(12):1236–1246, 2013.

[Cha] Zoé Chatzidakis. Preliminaries on fields. `http://www.logique.jussieu.fr/~zoe/papiers/3Fields.pdf`.

[CHRK19] Raf Cluckers, Immanuel Halupczok, and Silvain Rideau-Kikuchi. Hensel minimality. *arXiv preprint arXiv:1909.13792*, 2019.

[Con62] Paul F Conrad. Regularly ordered groups. *Proceedings of the American Mathematical Society*, 13(5):726–731, 1962.

[Del82] Françoise Delon. *Quelques propriétés des corps valués en théorie des modèles.* PhD thesis, 1982. doctoral thesis.

[DR88] Françoise Delon and Yamina Rouani. Indécidabilité de corps de séries formelles. *Journal of Symbolic Logic*, 53(4):1227–1234, 1988.

[DS03] Jan Denef and Hans Schoutens. On the decidability of the existential theory of $\mathbb{F}_p[[t]]$. *Valuation theory and its applications*, 2:43–60, 2003.

[EF99] Ido Efrat and Ivan Fesenko. Fields Galois-equivalent to a local field of positive characteristic. *Mathematical Research Letters*, 6(3):345–356, 1999.

[Efr06] Ido Efrat. *Valuations, orderings, and Milnor K-theory.* Number 124 in Mathematical Surveys and Monographs. American Mathematical Soc., 2006.

[EP05] Antonio J Engler and Alexander Prestel. *Valued Fields.* Springer Berlin, 2005.

[Ers04] Yu L Ershov. Extremal valued fields. *Algebra and Logic*, 43(5):327–330, 2004.

[ES09] Kirsten Eisenträger and Alexandra Shlapentokh. Undecidability in function fields of positive characteristic. *International Mathematics Research Notices*, 2009(21):4051–4086, 2009.

[ES17] Kirsten Eisenträger and Alexandra Shlapentokh. Hilbert's tenth problem over function fields of positive characteristic not containing the algebraic closure of a finite field. *Journal of the European Mathematical Society*, 19(7):2103–2138, 2017.

[FJ08] Michael D Fried and Moshe Jarden. *Field Arithmetic, Vol. 11 3rd ed., A Series of Modern Surveys in Mathematics*. Berlin, Heidelberg: Springer-Verlag, 2008.

[FJ17] Arno Fehm and Franziska Jahnke. Recent progress on definability of henselian valuations. In Fabrizio Broglia, Françoise Delon, Max Dickmann, Danielle Gondard-Cozette, and Victoria Powers, editors, *Ordered algebraic structures and related topics*, volume 697, pages 135–143. American Mathematical Society, 2017.

[Hod93] Wilfrid Hodges. *Model theory*. Cambridge University Press, 1993.

[Jar96] Moshe Jarden. Infinite Galois theory. In Michiel Hazewinkel, editor, *Handbook of algebra*, volume 1, pages 271–319. North-Holland, 1996.

[Kap42] Irving Kaplansky. Maximal fields with valuations. *Duke Math. J.*, 9(2):303–321, 06 1942.

[Kap45] Irving Kaplansky. Maximal fields with valuations, II. *Duke Math. J.*, 12(2):243–248, 06 1945.

[Koe14] Jochen Koenigsmann. Undecidability in number theory. In *Model Theory in Algebra, Analysis and Arithmetic*, pages 159–195. Springer, 2014.

[KPR86] Franz-Viktor Kuhlmann, Matthias Pank, and Peter Roquette. Immediate and purely wild extensions of valued fields. *manuscripta mathematica*, 55(1):39–67, 1986.

[Kuh] Franz-Viktor Kuhlmann. Valuation theory. Work in progress available at http://math.usask.ca/∼fvk/Fvkbook.htm.

[Kuh90] Franz-Viktor Kuhlmann. *Henselian function fields and tame fields.* Heidelberg, 1990. (extended doctoral thesis).

[Kuh01] Franz-Viktor Kuhlmann. Elementary properties of power series fields over finite fields. *The Journal of Symbolic Logic*, 66(2):771–791, 2001.

[Kuh06] Franz-Viktor Kuhlmann. Additive polynomials and their role in the model theory of valued fields. In *Logic in Tehran*, pages 160–203, 2006.

[Kuh11] Franz-Viktor Kuhlmann. The defect. In *Commutative Algebra*, pages 277–318. Springer, 2011.

[Kuh16] Franz-Viktor Kuhlmann. The algebra and model theory of tame valued fields. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2016(719):1–43, 2016.

[Lan03] S Lang. *Algebra. Revised third edition.* Number 211 in Graduate Texts in Mathematics. 2003.

[Mac76] Angus Macintyre. On definable subsets of p-adic fields. *The Journal of Symbolic Logic*, 41(3):605–610, 1976.

[Mar06] David Marker. *Model theory: an introduction*, volume 217. Springer Science & Business Media, 2006.

[Mon17] Samaria Montenegro. Pseudo real closed fields, pseudo p-adically closed fields and NTP2. *Annals of Pure and Applied Logic*, 168(1):191–232, 2017.

[Poi12] Bruno Poizat. *A course in model theory: an introduction to contemporary mathematical logic.* Springer Science & Business Media, 2012.

[Poo08] Bjorn Poonen. Undecidability in number theory. *Notices of the AMS*, 55(3), 2008.

[Pop96] Florian Pop. Embedding problems over large fields. *Annals of Mathematics*, 144(1):1–34, 1996.

[Pop10] Florian Pop. Henselian implies large. *Annals of Mathematics*, 172(3):2183–2195, 2010.

[PR84] Alexander Prestel and Peter Roquette. Formally p-adic fields, 1984.

[Pre81] Alexander Prestel. Pseudo real closed fields. In *Set theory and model theory*, pages 127–156. Springer, 1981.

[Roh03] Thomas Rohwer. *Valued difference fields as modules over twisted polynomial rings*. PhD thesis, University of Illinois at Urbana-Champaign, 2003.

[RZ60] Abraham Robinson and Elias Zakon. Elementary properties of ordered abelian groups. *Transactions of the American Mathematical Society*, 96(2):222–236, 1960.

[RZ00] Luis Ribes and Pavel Zalesskii. *Profinite groups*. Springer, 2000.

[Ser79] Jean-Pierre Serre. *Local fields*. Number 67 in Graduate Texts in Mathematics. Springer-Verlag, New York, 1979.

[Sho71] Joseph R Shoenfield. A theorem on quantifier elimination. In *Symposia Mathematica*, volume 5, pages 173–176, 1971.

[VDD85] Lou Van Den Dries. The field of reals with a predicate for the powers of two. *Manuscripta mathematica*, 54(1):187–195, 1985.

[VDD88] Lou Van Den Dries. Alfred Tarski's elimination theory for real closed fields. *The Journal of Symbolic Logic*, 53(1):7–19, 1988.

[Wha57] G Whaples. Galois cohomology of additive polynomial and n-th power mappings of fields. *Duke Mathematical Journal*, 24(2):143–150, 1957.