MATHEMATIK

Towards a model theory of Frobenius lifts

VALUED DIFFERENCE FIELDS AND DEFINABLE HENSELIAN VALUATIONS

Inauguraldissertation
zur Erlangung des Doktorgrades der Naturwissenschaften
im Fachbereich Mathematik und Informatik
der Mathematisch-Naturwissenschaftlichen Fakultät
der Universität Münster

vorgelegt von

SIMONE RAMELLO
aus ALBA (ITALIEN)

– SEPTEMBER 2024 –

DEKAN: *Prof. Dr. Arthur Bartels*
ERSTGUTACHTER: *Prof. Dr. Franziska Jahnke*
ZWEITGUTACHTER: *Prof. Dr. Matthias Aschenbrenner*

TAG DER MÜNDLICHEN PRÜFUNG: 11.04.2025
TAG DER PROMOTION: 11.04.2025

The crucial idea behind the main work of this thesis came to me in conversation, in front of a blackboard on the dusty corridor on the fourth floor of the ENS, in Paris. That insight only came to light thanks to the guidance of one of the greatest model theorists and mentors I have had the honour of meeting. This thesis is dedicated to the memory of Zoé Chatzidakis.

Contents

Summary

We tackle two different questions in the model theory of valued fields.

- CHAPTER 3, based on [KRS24]: In joint work with Margarete Ketelsen and Piotr Szewczyk, we study the question of $\mathcal{L}_{\text{ring}}$-definability of non-trivial henselian valuation rings. Building on work of Jahnke and Koenigsmann, we provide a characterization of henselian fields that admit a non-trivial definable henselian valuation. In particular, we treat cases where the canonical henselian valuation has positive residue characteristic, using techniques from the model theory and algebra of tame fields.
- CHAPTER 4, based on [Ram24]: We establish relative quantifier elimination for valued fields of residue characteristic zero enriched with a non-surjective endomorphism, building on recent work of Dor and Halevi. In particular, we deduce relative quantifier elimination for the ultralimit of the Frobenius action on separably closed valued fields of positive characteristic.

Zusammenfassung

Wir befassen uns mit zwei verschiedenen Fragen in der Modelltheorie bewerteter Körper.

- CHAPTER 3, basierend auf [KRS24]: In Kollaboration mit Margarete Ketelsen und Piotr Szewczyk untersuchen wir die Frage der $\mathcal{L}_{\text{ring}}$-Definierbarkeit von nicht-trivialen henselschen Bewertungsringen. Aufbauend auf Ergebnissen von Jahnke und Koenigsmann liefern wir eine Charakterisierung von henselschen Körpern, die eine nicht-triviale definierbare henselsche Bewertung besitzen. Insbesondere behandeln wir Fälle, in denen die kanonische henselsche Bewertung positive Restklassencharakteristik besitzt. Wir verwenden dazu Techniken aus Modelltheorie und Algebra von zahmen Körper.
- CHAPTER 4, basierend auf [Ram24]: Wir beweisen relative Quantorenelimination für bewertete Körper der Restklassencharakteristik null, die mit einem nicht-surjektiven Endomorphismus ausgestattet sind, und bauen dabei auf der jüngsten Arbeit von Dor und Halevi auf. Insbesondere folgern wir daraus relative Quantorenelimination für den Ultralimes der Frobenius-Wirkung auf separabel abgeschlossenen bewerteten Körpern positiver Charakteristik.

I. *Fighting dragons with you:* ACKNOWLEDGEMENTS

> Too much joy, I swear, is lost in our desperation to keep it.
>
> – Ocean Vuong, *On Earth We're Briefly Gorgeous*

My most heartfelt thanks must go to my supervisors, Franzi and Martin. The world is full of clever mathematicians, human beings that somehow transpire insights and knowledge, and they are certainly two of them. However, I am also sure that this PhD would have been very different had I not been blessed with two caring and compassionate supervisors. Throughout these years, I was allowed to carve my own path, make my own choices and mistakes, and I was always able to rely on them for support, understanding, grace, and advice. Some of my fondest memories involve them, climbing cherry trees, long nights somewhere on the French coast, and perhaps a bit too much wine. Simply put: thanks for everything.

For the past three minus epsilon years, the Cluster has been a home, with all its joys and idiosyncrasies. I hope I have managed to leave it a little bit better than how I found it (the campsite rule: "*Leave the campsite better than you found it*"). To the many people I have shared a coffee break, a train ride, a cry, a long walk home, a trip to Granada with wonderful vegetarian meals, a karaoke song, a tattoo, a shift at the beer stand, or a FIFA tournament with: you are family, and for that you have my deepest thanks.

To all model theorists, far and near: you have welcomed me, and offered me the irreplaceable gift of a community. Thanks to the groups in Münster, Paris, Dresden, and Vienna, for sharing a bit (or all) of this path with me. To so many of you I owe beautiful memories and hilarious jokes that will always be part of my personal mosaic.

I would like to particularly thank Matthias Aschenbrenner for his detailed report, comments and corrections to this manuscript.

Alla mia famiglia, per cui ho esaurito i modi di dire *grazie*. Al gruppo di Sanda (e dintorni), agli (ex?-)CVD, allə logichə italianə con cui ho condiviso conferenze e scuole estive in riva al lago, a tuttə coloro con cui ho speso giorni, mesi, o anche solo brevi momenti a parlare di matematica, accademia, o chissà cos'altro: a questo traguardo ci siamo arrivatə assieme.

Allə amichə che ho conosciuto da quando mi sono trasferito a Trieste: Baiardinə, masterinə, Sissinə. Quando mi chiedono se io sia contento di essere tornato in Italia, di

solito rispondo con voi. Grazie per questi mesi. A te, in particolare, che sai benissimo chi sei.

Ad maiora.

The author has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under the Excellence Strategy EXC 2044–390685587, *Mathematics Münster: Dynamics–Geometry–Structure*, via HI 2004/1-1 (part of the French-German ANR-DFG project *GeoMod*), and via the DFG project 495759320 – *Modelltheorie bewerteter Körper mit Endomorphismus*.

II. *First came the valuation:* INTRODUCTION

> She was also, by the standards of other people, lost. She would not see it like that. She knew where she was, it was just that everywhere else didn't.
>
> Terry Pratchett, *Equal Rites*

The first section of this introduction is meant to be a rather metaphorical and general introduction to the themes of this thesis, which can be read by anyone with some background in mathematics. The following sections dwelve more deeply into the results of the thesis, and require some knowledge of model-theoretic algebra.

2.1 | First came the valuation

It is a truth universally acknowledged that polynomials are complicated. If you disagree, I suggest the following experiment: choose a polynomial over $\mathbb{Z}$ that does not fit into the half a dozen criteria you might know to prove that it has roots in $\mathbb{Z}$. Can you prove it has roots? Can you exhibit them? You might be tempted to try and find out the solutions over $\mathbb{R}$, and then check if the solutions you found were integers. This strategy is not perfect and may not yield a complete answer, but at least you might be able to exhibit some of those solutions. Alternatively, you might check if the polynomial has a root when reduced modulo p: if the answer is no, then you know the original equation cannot be solved in $\mathbb{Z}$.[1]

What this experiment teaches us is that these hard arithmetical questions about polynomial equations become manifestly easier when the rings or fields we work with have additional structure. If we move over to the real numbers, then the tools of analysis come to the rescue, and we might argue about what the graph of the related polynomial function looks like. If you work with the p-adic valuation (in other words, start looking at things modulo p), you only need to check finitely many possible solutions for your equation to gain some information about your equation. Often, this additional structure – like the order on the reals – was actually hidden in the arithmetic all along; after all, if $\lambda \in \mathbb{R}$ then $\lambda \geqslant 0$ if and only if λ is a square. This theme of

[1] By the way, there is a good reason this problem is very hard. The question of whether a polynomial with coefficients in $\mathbb{Z}$ has a solution in $\mathbb{Z}$ cannot be answered algorithmically: in other words, Hilbert's Tenth Problem over $\mathbb{Z}$ has a negative answer. See [Dav73] or [Koe14] for more details.

structure hiding in the arithmetic will be a major point of CHAPTER 3; for now, it is enough to realize that if, beyond the slightly overwhelming arithmetical operations $+$ and $\times$, we are given a little bit more of spice – like an order or a valuation, then the resulting dish will have a more accessible taste profile. The focus of this thesis is on one of those possible spices: valuations.

At their core, valuations are nothing more than generalized absolute values. They then fulfil this role of enhancing the fields they are supported on to give us the ability to prove that certain equations have solutions (or do not). They give us a language to understand which elements are small, and which ones are not; they give us more tests and criteria to figure out whether polynomials have solutions or not; they give us a topology, and tools to talk geometrically about fields. It is thus perhaps not a mystery that valued fields are a major object of interest for model theorists: much of our job takes the form of understanding first-order information, which in the language of rings is well-represented by arithmetical questions. The development of the model theory of valued fields has run parallel to the development of model theory itself, taking a central stage from the very beginning; Abraham Robinson's work from the 1950s allows us to understand algebraically closed valued fields in detail, and we keep coming up with new ways to reinvent his insights to this very day. Many players have since then given their contributions to the model-theoretic understanding of these objects, so it would be impossible (and possibly endlessly frustrating) to try to recover a detailed historical account in these few pages. Instead, it might be fruitful to highlight a central philosophy that dominates the field (pun not intended), which takes two very different shapes in the two works that constitute this thesis, but remains a guiding principle for all those who dwell in the model theory of valued fields and associated topics.

As already emphasized, a field in itself might not be a particularly exciting object (in fact, it might be quite scary and mysterious). When endowed with a valuation, however, it acquires two fundamental gadgets, or invariants: an ordered (abelian) group, called the *value group*, and another field, called the *residue field*. These two invariants give us a way to shed light on the arithmetical structure of the field in many deep ways. In an ideal scenario, they contain all that there is to know about the valued field. One could then phrase the so-called *Ax-Kochen/Ershov principle* in the following, slightly optimistic, way:

Valued fields are nothing but a value group and a residue field in a trenchcoat[2].

This phrasing of the AK/E principle is as catchy as it is, of course, false. In most cases, the valuation is not good enough to completely capture the structure of the valued field in its invariants. Sometimes, this whole story of a valuation helping us solve polynomial equations is just, as most good stories, plainly untrue; we thus have to restrict to valuations that are rich and well-behaved enough. Finding out exactly

[2] https://knowyourmeme.com/memes/two-kids-in-a-trenchcoat.

what these are is a big question; words like *henselian*, *tame*, *σ-henselian*, ... will come up often in the later sections. For now, we can rephrase the AK/E principle to have at least some chance to be true:

If the valued field is good enough, any model-theoretic question about it can be answered using the model theory of its value group and the model theory of its residue field.

In a sense, CHAPTER 3 exploits the AK/E principle (and its many concrete iterations and consequences, like stable embeddedness of the residue field and value group) to say something about how much the arithmetic of a field already knows about its valuation; on the other hand, CHAPTER 4 sets out to establish yet another one of these iterations in the particular case of valued fields with an endomorphism.

A spoiler for the rest of this introduction: after an interlude on notation, each section will focus on one of the chapters, namely SECTION 2.3 will explain the main ideas and concepts of CHAPTER 3, whereas SECTION 2.4 will serve as an introduction to CHAPTER 4.

2.2 | A WARNING ABOUT NOTATION

The landscape of model theory of valued fields is extremely varied. In particular, different subcommunities use different notations for basic objects like residue field and value group. The two manuscripts that make up the mathematical content of this thesis speak to two of these subcommunities, which have rather different conventions. This leads to a moral, typographical, and existential conundrum: how should we uniformize the two chapters? We choose not to, and this section is meant to be a fair warning for what is about to come.

CHAPTER 3 focuses on definable valuations and uses heavily the machinery of tame fields. This requires considering, at times, several valuations on the same field; the notation then has to accomodate this, leading to the following choices: if (K, v) is a valued field, then we denote its value group by vK and its residue field by Kv.

CHAPTER 4, on the other hand, deals with valued difference fields and relative quantifier elimination. In this case, there will always only be one valuation on each field at a time, and it will always be denoted by v. It would thus be an obstacle to readability to keep track of v in the notation of residue field and value group. If (K, v) is a valued field, we thus denote by k_K the residue field and by Γ_K the value group.

We want to reiterate the ultimate goal of this choice: to allow members of both communities to access the part that more directly speaks to them, without impairing the readability of the other one (for example, it would be impossible to use the notation of CHAPTER 4 in CHAPTER 3).

Over the field of p-adic numbers $\mathbb{Q}_p$ the valuation ring is an $\mathcal{L}_{\text{ring}}$-definable subset, as first observed in the seminal work of Julia Robinson on Hilbert's Tenth Problem [Rob63]. Understanding this phenomenon is a classical topic in the model theory of valued fields, and it has striking applications in the investigations around dividing lines for fields, see e.g. Johnson's spectacular classification of dp-finite fields in [Joh20]. For a more thorough survey on definability of valuations we refer the reader to [FJ17].

In chapter 3, which is essentially a variant of [KRS24], in turn written together with Margarete Ketelsen and Piotr Szewczyk, we will focus on the problem of characterizing which fields admit a non-trivial definable henselian valuation. There are some clear obstructions: for example, if the field is separably closed, then the answer is always *no* (say, because separably closed fields are stable by [Woo79]); similarly, if the field admits *no* non-trivial henselian valuation, then we have no chance of finding a definable one. Once these cases are excluded, Jahnke and Koenigsmann isolate in [JK17] necessary and sufficient properties of the canonical henselian valuation v_K of a field K (see Subsection 3.2) for the existence of a definable non-trivial henselian valuation, under the assumption that $\operatorname{char} Kv_K = 0$.

This characterization (and its extension, which we will state in a moment) falls into the general philosophy of Ax-Kochen/Ershov principles. Typically, (model-theoretic) questions about a valued field (K, v) are answered using its residue field Kv and its value group vK. In this case, since the question is about a field K with no specified valuation, one ought to isolate a canonical one among all its henselian valuations; this is usually denoted by v_K. Answers to model-theoretic questions about K, then, should be given in terms of properties of Kv_K and $v_K K$.

The two directions of Jahnke and Koenigsmann's result are proved using different techniques: the existence of a definable non-trivial henselian valuation is proved directly in the various cases, combining a wealth of work from previous papers such as [FJ15] and [JK15]. On the other hand, the obstruction to the existence of such a valuation makes use of stable embeddedness in the case of residue characteristic zero henselian valued fields, a fact that essentially follows from Pas' relative quantifier elimination result in [Pas89].

We set out to the result of Jahnke and Koenigsmann by removing the assumption on the residue characteristic. Say that a field is *henselian* if it admits some non-trivial henselian valuation, and that it is *t-henselian* if it is elementarily equivalent (in the language of rings) to a henselian field. See section 3.2 for the definition of v_K. Then, we obtain the following:

Theorem A (theorem 3.5.1)**.** *Let K be perfect, not algebraically closed, and henselian. If $\operatorname{char} K = 0$ and $\operatorname{char} Kv_K = p > 0$, assume that $\mathcal{O}_{v_K}/p$ is semi-perfect. Then K admits a definable non-trivial henselian valuation if and only if at least one of the following conditions hold:*

1. *Kv_K is separably closed,*

2. *Kv_K is not t-henselian,*

3. *there is $L \succeq Kv_K$ such that $v_L L$ is not divisible,*

4. *$v_K K$ is not divisible,*

5. *(K, v_K) has defect,*

6. *there is $L \succeq Kv_K$ such that (L, v_L) has defect.*

Proving this result boils down to generalizing both of the directions of Jahnke and Koenigsmann's result, which once again requires different techniques. On the one hand, we exploit the technology of dependent and independent defect – as introduced and clarified in [KR23] – to produce a definable henselian valuation: this is essentially the content of SECTION 3.4. In positive characteristic, this works fine under the assumption of perfection; as for mixed characteristic, it is not enough to ask that the residue field be perfect, but rather one has to require that a certain quotient of the valuation ring is semi-perfect (this, in particular, implies that the residue field is perfect)[3]. Both of these conditions are also abundantly used for the other direction of the theorem, which heavily draws from the model theory of tame fields as developed in [Kuh16]; in particular, the various valuations around need to be tame in order to exploit the stable embeddedness results proved in [JS20].

2.4 | How I Learned to Stop Worrying and Love Endomorphisms

For the model-theoretically inclined, understanding valued fields has historically been a game of finding the correct invariants for their theories, reducing the possibly complicated valuational structure to something (at least a priori) easier to understand and classify. As already explained, given a valued field (K, v), one can naturally look at its value group Γ_K and residue field k_K as possible (model-theoretic) invariants, in the sense that one could hope to determine the theory of (K, v) using the theories of Γ_K and k_K. This concretized in the celebrated Ax-Kochen/Ershov results ([AK65], [Ers65]), which establish that – when (K, v) is henselian, and both K and k_K have characteristic zero – the theory of (K, v) is implied by the theories of Γ_K (as an ordered group) and k_K (as a field).

A great amount of work has since then blossomed, which seeks to generalize similar results in various directions. Recently, Rideau-Kikuchi and Vicaría in [RKV23] establish a similar transfer theorem for elimination of imaginaries in residue characteristic zero; in the positive characteristic world, the aforementioned work of Kuhlmann in [Kuh16]

[3] In the case of imperfect residue field, an argument of Scanlon's in [Jah, Proposition 3.6] seems to hint towards the fact that definable henselian valuations should (always?) exist.

provides similar answers to the question of axiomatizing a valued field using its residue field and value group.

In CHAPTER 4, which is essentially a variant of [Ram24], we tackle one such possible way of making the question harder, namely we enrich the structure of (K, v) with a valued field endomorphism σ. We call such a (K, v, σ) *inversive* if σ is surjective. Beyond the abstract interest for such structures, we note that they arise naturally in several interesting contexts: for example, if we let (K_p, v_p) be a separably closed valued field of characteristic p for every p, with no assumption on perfection, then we can expand one language of valued fields of our choice to be able to consider the structures (K_p, v_p, ϕ_p), where ϕ_p is a unary function that we interpret to be $\phi_p(x) = x^p$. If we take the ultraproduct along some non-principal ultrafilter $\mathcal{U}$ over the set of prime numbers, then, we end up with an algebraically closed valued field of residue characteristic zero equipped with some (possibly non-surjective) endomorphism that acts as an infinite non-standard Frobenius. These structures are now completely understood, first by work of Durhan[4] ([Azg10]), and then later in the harder case of positive characteristic[5] by work of Dor and Hrushovski ([DH22]) and Dor and Halevi ([DH23]), building on the twisted Lang-Weil estimates ([Hru], [SV21]). Such endomorphisms of valued fields are of a special kind, namely *ω-increasing*. At the other end of the spectrum sits the Witt-Frobenius, whose model theory we also understand thanks to [BMS07] and [AvdD11]. There, the endomorphism is surjective, so it is an automorphism; moreover, it is not just an automorphism of valued fields, but it is in fact valuation preserving, i.e. an *isometry*. In full generality, Durhan and Onay in [DO15] prove a relative quantifier elimination for a natural class of valued fields with an automorphism, showing that their theories can be understood in terms of their leading terms structures RV together with the induced automorphism.

The most recent work of Dor and Halevi ([DH23]) takes a step in the direction of removing the assumption of surjectivity of the endomorphism in question. They do so in the *absolute* case, namely identifying the model companion for valued fields enriched with an ω-increasing endomorphism. In turn, their work is based on the seminal results of Chatzidakis and Hrushovski ([CH04]), which isolate the model companion for fields with an endomorphism. Building on their work, and that of Durhan and Onay, we remove the surjectivity assumption in the *relative* case, proving a relative quantifier elimination down to the leading terms structure.

Proving relative quantifier elimination for (enriched) valued fields has essentially always been done along the same strategy. Namely, one first has to establish control of immediate extensions, what normally falls under the name of "Kaplansky theory": this means proving that, under reasonable assumptions, there is a unique immediate extension which is maximal in a specific category. This is the content of SECTION 4.3, and arguably the engine that makes the entire relative quantifier elimination work:

[4] Published under the surname Azgın.

[5] A characteristic $p > 0$ version of the example can be obtained by working with $\phi_q(x) = x^q$, where q is a power of p, and taking the ultraproduct over the set of powers of p.

Theorem B (THEOREM 4.3.21). *Suppose (K, v, σ) is a model of* FE *and $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Let (K_1, v, σ) and (K_2, v, σ) be two σ-separably σ-algebraically maximal, σ-separably σ-algebraic immediate extensions of K that are also models of* FE. *Then $K_1 \cong_K K_2$.*

In a back-and-forth scenario, it is then in a sense enough to reduce to the case where the extension we want to embed is immediate (so we can use the uniqueness theorem). This is the other crucial ingredient, which we adapt from [DO15] in SECTION 4.4.

It is important to note, at this point, that even when they are not inversive themselves, the valued difference fields in question have inversive residue difference field and value difference group. The realm of valued difference fields with non-inversive auxiliary endomorphisms seems rather hard to grasp: the uniqueness of a certain immediate extension fails, so the only other way would be the one followed by Kuhlmann in [Kuh16]. The two situations, however, are not so similar: crucially, being algebraically maximal is first-order, so one can consider an elementary class of algebraically maximal valued fields – on the other hand, being σ-algebraically maximal is not usually first-order (except for the ω-increasing case); also, the failure of Kaplansky theory in Kuhlmann's work is more a consequence of the residue field not satisfying Kaplansky's hypothesis (A) (what we call "linearly difference closed" in the difference world), since his residue fields are still perfect. It is not so clear what "σ-tame" should mean; a "σ-Abhyankar's inequality" does not hold in general, but it does in the ω-increasing world, so it seems like σ-tame should at least mean ω-increasing. Nevertheless, even in the general case parallels between non-surjective endomorphisms and imperfect positive characteristic valued fields can be drawn. Namely, models of our theory hVFE_0 seem to behave very closely to separably closed valued fields of finite degree of imperfection; for example, with a bit of saturation the inversive core is dense, and many techniques are σ-versions of the field theory of separable extensions. This is at least true when, following [DH23], we work with models of FE, namely difference fields (K, σ) where $\sigma(K)$ is algebraically closed in K. This is a technical assumption, which is clarified in SECTION 4.2.

Putting everything together, we are able to exhibit relative quantifier elimination for a wealth of languages, starting from the one using the leading terms structure (see DEFINITION 4.4.3 for the definition of the language $\mathcal{L}_{3,\mathrm{RV},\sigma}$). Note that all these languages must include (just like the languages considered in [DH23]) λ-functions parametrizing linear independence over $\sigma(K)$; see DEFINITION 4.4.6 and REMARK 4.4.7. See SECTION 4.3 for a discussion of σ-henselianity.

Theorem C (THEOREM 4.5.4). *Let* hVFE_0 *be the $\mathcal{L}_{3,\mathrm{RV},\sigma}$-theory of non-inversive weakly σ-henselian valued difference fields (K, v, σ) of residue characteristic zero, with $\sigma(K)$ algebraically closed in K. Then, modulo* hVFE_0*, every formula is equivalent to one where the quantifiers only range over variables from* **RV** *and* **VG**.

Once relative quantifier elimination using RV is established, one can immediately deduce a series of transfer theorems for models of hVFE_0 in a variety of languages, like ones with angular components, sections, or lifts of the residue field.

From this we can deduce a transfer theorem for NTP_2 along the strategy introduced in [CH14]. In particular, this gives a different proof of [DH23, Corollary 8.6], namely that the theory that [DH23] call $\widetilde{\mathrm{VFE}}_0$ is NTP_2.

Theorem D (THEOREM 4.5.26). *Let* $(K, v, \sigma) \models \mathrm{hVFE}_0$*. Then* (K, v, σ) *is* NTP_2 *if and only if* $(\mathrm{RV}_K, \sigma_{\mathrm{rv}})$ *is* NTP_2*.*

III. Definability of henselian valuations in positive residue characteristic

> So I go. I travel farther and faster and harder than most, and I read, and I write, and I love cities. To be alone in a crowd, apart and belonging, to have distance between what I see and what I am.
>
> Amal El-Mohtar, *This Is How You Lose the Time War*

This chapter almost completely overlaps with [KRS24]. It was written in joint work with Margarete Ketelsen and Piotr Szewczyk, to which we all contributed equally.

3.1 | The road ahead

The goal of this chapter is proving a generalization of the following result:

Theorem 3.1.1 ([JK17, Corollary 6.1])**.** *Let K be a henselian field that is not algebraically closed, with* $\operatorname{char} Kv_K = 0$*. Then K admits a definable non-trivial henselian valuation if and only if at least one of the following conditions holds:*

1. *Kv_K is separably closed,*
2. *Kv_K is not t-henselian,*
3. *there is $L \succeq Kv_K$ such that $v_L L$ is not divisible,*
4. *$v_K K$ is not divisible.*

The chapter is structured as follows:

- section 3.2 sets the stage, by fixing the notation and a few elementary facts.
- section 3.3 explores the two roles played by divisibility of the value group: on the one hand, in proposition 3.3.6, failure of divisibility in an elementary extension of Kv_K is exploited to ($\varnothing$-)define a valuation; on the other, in theorem 3.3.5, divisibility is the keystone to finding obstructions to the existence of definable valuations.

- SECTION 3.4 deploys the technology of independent defect to build definable henselian valuations from certain kinds of Galois defect extensions.

- SECTION 3.5 puts everything together, providing the full characterization as advertised (SECTION 2.3).

- SECTION 3.6 provides a few explicit examples of fields that fit into the main result, and discusses a few open questions.

3.2 | Preliminaries

Notation. Given a valued field (K, v), we denote by $\mathcal{O}_v$ or $\mathcal{O}_{(K,v)}$ its valuation ring with maximal ideal $\mathfrak{m}_v$, Kv its residue field and vK its value group. We will write vK_∞ for $vK \cup \{\infty\}$, and denote by $\mathrm{res}_v \colon \mathcal{O}_v \to Kv$ the residue map. Given a field K, we denote its separable closure by K^{sep}. A field will be called *henselian* if it admits a non-trivial henselian valuation. A ring of characteristic $p > 0$ is called *semi-perfect* if $x \mapsto x^p$ is surjective (but not necessarily injective). If $K \subseteq L$ is a Galois extension, we will write $\mathrm{Gal}(L|K)$ for the corresponding Galois group. We will mostly work with the two languages $\mathcal{L}_{\mathrm{ring}} = \{+, \cdot, -, 0, 1\}$, and $\mathcal{L}_{\mathrm{val}} = \mathcal{L}_{\mathrm{ring}} \cup \{\mathcal{O}\}$, where $\mathcal{O}$ is a unary predicate. If M is a first-order structure, I is an index set and $\mathcal{U}$ is a ultrafilter on I, then we denote by $M^{\mathcal{U}}$ the corresponding ultrapower. Unless otherwise stated, *definable* will mean definable *with parameters*. Otherwise, we say $\varnothing$-definable.

Coarsenings and compositions. We briefly recall some facts about coarsenings and refinements of valuations, which are explained in greater depth in [EP05, Section 2.3]. For two valuations v and w on K, we say that v is *finer* than w or is a *refinement* of w (that w is *coarser* than v or is a *coarsening* of v), if $\mathcal{O}_v \subseteq \mathcal{O}_w$. We identify v and w if they have the same valuation ring, and use v and $\mathcal{O}_v$ interchangeably. We say v and w are *comparable* if v is a coarsening of w, or w is a coarsening of v. Coarsenings of a fixed valuation ring $\mathcal{O}_v$ are linearly ordered by inclusion and in one-to-one correspondence with convex subgroups of vK. We denote by v_H the coarsening of v associated to a convex subgroup H of vK.

Given a valuation v and a coarsening w corresponding to the convex subgroup $H \subseteq vK$, we denote the *induced valuation* $\bar{v} \colon (Kw)^\times \to H$ on Kw by $\bar{v}$, with value group H and residue field Kv. Given a valuation w on K and a valuation u on its residue field Kw, we define the *composition* $v = u \circ w$ to be the valuation with valuation ring $\mathcal{O}_v = \mathrm{res}_w^{-1}(\mathcal{O}_u)$. One then has that $u(Kw)$ is a convex subgroup of vK with $vK/u(Kw) \cong wK$.

We can write the valuations as places and obtain the diagram

$$\begin{tikzcd} K \arrow[r, "w"] \arrow[rr, bend left, "v"] & Kw \arrow[r, "u"] & Kv = (Kw)u. \end{tikzcd}$$

The canonical henselian valuation. What follows is classical and can be found in greater detail in [EP05, Section 4.4]. For any field K, one can arrange henselian valuation rings on K according to whether their residue field is separably closed or not. Namely, one can define

$$H_1(K) = \{\mathcal{O}_v \mid (K, v) \text{ henselian and } (Kv)^{\text{sep}} \neq Kv\}$$

and

$$H_2(K) = \{\mathcal{O}_v \mid (K, v) \text{ henselian and } (Kv)^{\text{sep}} = Kv\}.$$

As above, we identify v with $\mathcal{O}_v$, so we will often write $v \in H_1(K)$ to mean $\mathcal{O}_v \in H_1(K)$.

The set $H_1(K)$ is linearly ordered by coarsening, with K as maximum. If $H_2(K) = \varnothing$, then there is a finest valuation in $H_1(K)$, which we denote by v_K. Otherwise, we let v_K be the coarsest valuation in $H_2(K)$. We call v_K the *canonical henselian valuation* on K.

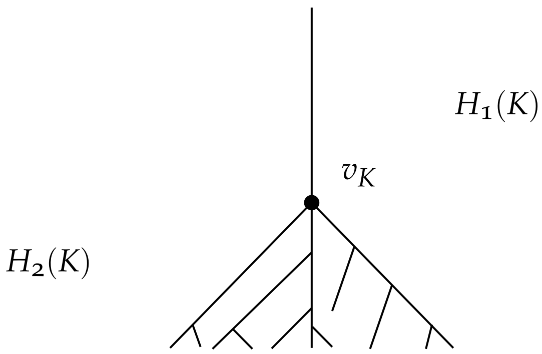

Every henselian valuation on K is comparable to v_K. Moreover, v_K is non-trivial if and only if K is henselian and not separably closed. By definition, $Kv_K \neq (Kv_K)^{\text{sep}}$ if and only if $v_K \in H_1(K)$ if and only if $H_2(K) = \varnothing$, and if these equivalent conditions hold, then all henselian valuation rings on K are linearly ordered and coarser than $\mathcal{O}_{v_K}$.

Remark 3.2.1. Let (K, v) be a henselian valued field with $v \in H_1(K)$. Then, $v_K = v_{Kv} \circ v$. Indeed, there is a correspondence between the set H of henselian refinements of v on K and henselian valuations on Kv. It is given by sending any refinement u of v to the induced valuation $\overline{u}$ on Kv, as in the following diagram (where valuations are written as places):

$$\begin{tikzcd} K \arrow[r, "v"] \arrow[rr, bend left, "u"] & Kv \arrow[r, "\overline{u}"] & (Kv)\overline{u} = Ku. \end{tikzcd}$$

Since the correspondence preserves residue fields, $H \cap H_1(K)$ is mapped to $H_1(Kv)$, and similarly for $H \cap H_2(K)$.

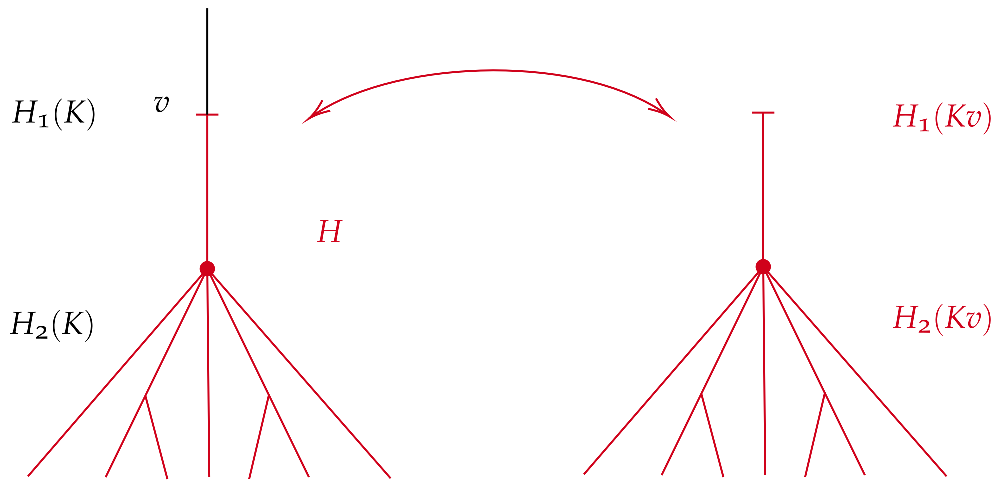

The correspondence preserves the coarsening relation. In particular, it follows that v_K is mapped to v_{Kv} and $v_K = v_{Kv} \circ v$.

This correspondence immediately allows us to prove that condition 6 of our Main Theorem trivializes in equicharacteristic zero.

Lemma 3.2.2. *If Kv_K is henselian, then it is separably closed.*

Proof. Suppose Kv_K admits a non-trivial henselian valuation v. Via REMARK 3.2.1 this gives rise to a proper refinement of v_K. Hence $H_2(K)$ is non-empty, so Kv_K is separably closed. □

Corollary 3.2.3. *Let $L \succeq Kv_K$. Then (L, v_L) cannot have mixed characteristic.*

Proof. Suppose there is $L \succeq Kv_K$ such that (L, v_L) has mixed characteristic. Restricting v_L down to Kv_K yields a henselian valuation (because Kv_K is algebraically closed in L) of mixed characteristic on Kv_K, which in particular must be non-trivial. By LEMMA 3.2.2, Kv_K must be separably closed. It follows that L is also separably closed, but this implies that v_L is trivial. □

Defect and tame fields. Let $(K, v) \subseteq (L, w)$ be a finite extension of a henselian valued field (note that w is uniquely determined). Let p be the *characteristic exponent* of Kv, namely $p = 1$ if $\mathrm{char}(Kv) = 0$, and $p = \mathrm{char}(Kv)$ otherwise. Then one has, by [EP05, Theorem 3.3.3],

$$[L : K] = d \cdot (wL : vK) \cdot [Lw : Kv], \tag{$\star$}$$

for some $d = p^\nu \in \mathbb{N}$. We say that the extension is *defectless* if $d = 1$ (or equivalently $\nu = 0$), otherwise we say that the extension has *defect* (or that it is a *defect extension*). The valued field (K, v) is called *defectless* if all of its finite extensions are defectless. We refer to [Kuh16] for the details on the various equivalent definitions of "tame".

Definition 3.2.4. Let (K, v) be a henselian valued field and let p be the characteristic exponent of the residue field. Then (K, v) is called *tame* if it is algebraically maximal[1], vK is p-divisible, and Kv is perfect. Equivalently, (K, v) is tame if it is defectless, vK is p-divisible, and Kv is perfect.

In the case of valued fields where $p = 1$ (i.e., valued fields of equicharacteristic zero), tame is equivalent to henselian.

Remark 3.2.5 (see [Kuh16, Section 7])**.** The class of tame valued fields is an elementary class in the language $\mathcal{L}_{\text{val}}$.

t-henselianity and saturation. A field K is said to be *t-henselian* if it is elementarily equivalent, in $\mathcal{L}_{\text{ring}}$, to a henselian field L. However, a field K can be t-henselian without being henselian. The next fact is well-known, and its proof follows from [PZ78, Lemma 3.3] and [Pre91, Page 203].

Fact 3.2.6. *Let K be t-henselian and $\aleph_1$-saturated. Then K is henselian. In particular, if K is t-henselian and $K \preceq L$ is an $\aleph_1$-saturated elementary extension, then L is henselian.*

Definability in Jahnke-Koenigsmann. For use in later sections, we also summarize two of the main theorems of [JK17], which will prove to be fundamental tools in our work. In particular, we highlight how the second part of the upcoming statement allows us to often assume that the value group we are working with is divisible.

Definition 3.2.7. Let p be a prime. An ordered abelian group Γ is *p-antiregular* if

1. no non-trivial quotient of Γ (by a convex subgroup) is p-divisible,

2. Γ has no rank 1 quotient.

Fact 3.2.8 ([JK17, Theorems A and B])**.** *Let K be a henselian field which is not separably closed. If $v_K K$ is not divisible, then K admits a definable (possibly with parameters) non-trivial henselian valuation. Assume further that K satisfies at least one of the following:*

1. *Kv_K is separably closed,*

2. *Kv_K is not t-henselian,*

3. *there is a prime p such that $v_K K$ is not p-divisible and not p-antiregular,*

then K admits an $\varnothing$-definable non-trivial henselian valuation.

[1] That is, if there is no proper algebraic extension $(K, v) \subseteq (L, w)$ where the canonical inclusions $vK \hookrightarrow wL$ and $Kv \hookrightarrow Lw$ are surjective.

Keisler-Shelah. In many occasions, it will be useful to turn elementary equivalence between two structure into an isomorphism between some very saturated elementary extensions. This kind of sorcery is achieved by the celebrated Keisler-Shelah isomorphism theorem:

Fact 3.2.9 ([Mar06, Theorem 2.5.36]). *Let $\mathcal{L}$ be a first-order language, and let M and N be $\mathcal{L}$-structures. Then, $M \equiv N$ if and only if there are an index set I and an ultrafilter $\mathcal{U}$ on I, both depending only on $|\mathcal{L}|$, such that $M^{\mathcal{U}} \cong N^{\mathcal{U}}$.*

3.3 | Divisibility

Divisibility as an obstruction. Using classical results on elimination of quantifiers for henselian valued fields of residue characteristic zero, we know that the induced structure of the value group is that of an ordered abelian group. The proof of theorem 3.1.1 crucially relies on this to obstruct the existence of definable valuations. We use the work of Jahnke and Simon in [JS20] to generalize this argument.

Definition 3.3.1. We let $\mathcal{L}_3$ be the three-sorted language with sorts:

1. **VF**, with language $\mathcal{L}_{\text{ring}}$ (see definition 4.4.5),
2. **RF**, with language $\mathcal{L}_{\text{ring}}$,
3. **VG**, with language $\mathcal{L}_{\text{oag}} \cup \{\infty\}$,

and connecting functions $\text{res}\colon \mathbf{VF} \to \mathbf{RF}$ and $v\colon \mathbf{VF} \to \mathbf{VG}$. Any valued field (K, v) can be interpreted as an $\mathcal{L}_3$-structure in the natural way.

Fact 3.3.2 ([JS20, Lemma 3.1]). *Let (K, v) be a tame field of positive residue characteristic. Then the value group vK is purely stably embedded as an ordered abelian group, i.e. for every $D \subseteq vK^n$ which is $\mathcal{L}_3$-definable (possibly with parameters), there is $\psi \in \mathcal{L}_{\text{oag}}(vK)$ which defines D.*

The following lemma constitutes the core of the arguments of Jahnke and Koenigsmann, and will be used repeatedly throughout our proofs as well. It relies on a classical fact about divisible ordered abelian groups, which follows from quantifier elimination.

Fact 3.3.3. *Let Γ be a divisible ordered abelian group. Then Γ has no proper, $\mathcal{L}_{\text{oag}}$-definable (even with parameters) non-trivial convex subgroup.*

Lemma 3.3.4. *Suppose (K, v) is a tame valued field of positive residue characteristic, with vK divisible. If w is a proper non-trivial coarsening of v, then w cannot be $\mathcal{L}_{\text{val}}$-definable. In particular, it cannot be $\mathcal{L}_{\text{ring}}$-definable.*

Proof. Suppose w is $\mathcal{L}_{\text{val}}$-definable (possibly with parameters); then the corresponding proper non-trivial convex subgroup

$$\Delta = \{x \in vK \mid \exists y(v(y) = x \wedge w(y) = 0)\},$$

is $\mathcal{L}_3$-definable (possibly with parameters) as a subset of vK. By FACT 3.3.2, Δ is $\mathcal{L}_{\text{oag}}(vK)$-definable. This is a contradiction, as divisible ordered abelian groups have no definable proper non-trivial convex subgroups. □

The proof of the following theorem is then adapted directly from [JK17, Corollary 6.1].

Theorem 3.3.5. *Let (K, v_K) be such that Kv_K is perfect of positive characteristic. Assume that K admits a definable non-trivial henselian valuation. Then, at least one of the following holds:*

1. *Kv_K is separably closed,*
2. *Kv_K is not t-henselian,*
3. *there is $L \succeq Kv_K$ such that $v_L L$ is not divisible,*
4. *$v_K K$ is not divisible,*
5. *(K, v_K) has defect,*
6. *there is $L \succeq Kv_K$ such that (L, v_L) has defect.*

Proof. We assume

$$\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge \neg 5 \wedge \neg 6,$$

and want to show that K then cannot admit a non-trivial definable henselian valuation. Note that $\neg 1$ means that $H_2(K) = \varnothing$. Thus, all henselian valuations on K are linearly ordered and coarser than v_K.

Now, let $(M, v) \succeq (K, v_K)$ be an $\aleph_1$-saturated extension in $\mathcal{L}_{\text{val}}$. Then, $Mv \succeq Kv_K$ is an $\aleph_1$-saturated elementary extension of a t-henselian field (by $\neg 2$), hence it is henselian by FACT 3.2.6. Since $L := Mv \equiv Kv_K$ is henselian and not separably closed (by $\neg 1$), v_L is non-trivial and in fact, by REMARK 3.2.1, we have the following diagram (where the arrows are intended as places),

$$\begin{tikzcd} M \arrow[r, "v"] \arrow[rr, bend left, "v_M"] & L \arrow[r, "v_L"] & Lv_L = Mv_M \end{tikzcd}$$

from which we can deduce that $v_M = v_L \circ v$ is a proper refinement of v. Now, since $v_L L$ is divisible (because of $\neg 3$), so is $v_M M$. Indeed,

$$vM = v_M M / v_L L,$$

and we know vM to be divisible, as $v_K K$ is (because of $\neg 4$).

Now, by assumption Kv_K is a perfect field of positive characteristic and thus so is $L = Mv \succeq Kv_K$. It follows that also the residue field $Lv_L = Mv_M$ of the valuation v_L on L is perfect. By assumptions $\neg 5$ and $\neg 6$, v and v_L are defectless. Thus v_M is defectless (as the composition of defectless valuations is defectless, see [AJ24, Lemma 2.9]), and (M, v_M) is a tame valued field.

Assume now that there is a non-trivial definable henselian valuation u on K, with valuation ring defined by the $\mathcal{L}_{\text{ring}}(K)$-formula $\psi(x)$. Because of $\neg 1$, u is a (not necessarily proper) coarsening of v_K. Now, via the elementary embedding, $\psi(x)$ defines a valuation u^* on M, which is a coarsening of v and hence a proper coarsening of v_M. This is a contradiction to LEMMA 3.3.4. □

Failure of divisibility as a source of definability. We now seek to extend the arguments given in [JK17, Proposition 5.5] of the case when $v_K K$ is itself divisible, but there is an elementary extension L of Kv_K which is henselian with non-divisible value group $v_L L$.

Proposition 3.3.6. *Let K be a henselian field, and let v_K be its canonical henselian valuation. If*

1. *there is $L \equiv Kv_K$ such that $v_L L$ is non-divisible, and*
2. *$v_K K$ is divisible,*

then K admits an $\emptyset$-definable non-trivial henselian valuation.

Proof. Let $L \equiv Kv_K$ be such that $v_L L$ is non-divisible. First, we notice that Kv_K is not separably closed, and hence by LEMMA 3.2.2 it does not admit any non-trivial henselian valuation. Indeed, if Kv_K were separably closed, L would also be separably closed, but separably closed valued fields always have divisible value group, and $v_L L$ is not divisible by assumption.

We now claim that $v_L \in H_1(L)$. Assume for a contradiction that it is not. Then by FACT 3.2.8, there is an $\emptyset$-definable non-trivial henselian valuation on L. Thus there is also an $\emptyset$-definable non-trivial henselian valuation u on Kv_K by elementary equivalence, a contradiction.

By the Keisler-Shelah isomorphism theorem (FACT 3.2.9), there are an index set I and an ultrafilter $\mathcal{U}$ on I such that $L^{\mathcal{U}} \cong (Kv_K)^{\mathcal{U}}$. Take the ultrapower $(L^{\mathcal{U}}, v_L^{\mathcal{U}}) := (L, v_L)^{\mathcal{U}}$ as an $\mathcal{L}_{\text{val}}$-structure. Then, $v_L^{\mathcal{U}} L^{\mathcal{U}} \equiv v_L L$ is not divisible. Since $v_L \in H_1(L)$, then $v_L^{\mathcal{U}} \in H_1(L^{\mathcal{U}})$ and thus $v_L^{\mathcal{U}}$ coarsens the canonical henselian valuation $v_{L^{\mathcal{U}}}$ on $L^{\mathcal{U}}$. In particular, also $v_{L^{\mathcal{U}}} L^{\mathcal{U}}$ cannot be divisible, say not q-divisible for some prime q. Moreover, if we let $(K^{\mathcal{U}}, v_K^{\mathcal{U}}) := (K, v_K)^{\mathcal{U}}$ as an $\mathcal{L}_{\text{val}}$-structure, then $K^{\mathcal{U}} v_K^{\mathcal{U}} = L^{\mathcal{U}}$ and $v_K^{\mathcal{U}} K^{\mathcal{U}} \equiv v_K K$ is divisible. Now,

- $K^{\mathcal{U}}$ is henselian and $v_{K^{\mathcal{U}}} = v_{L^{\mathcal{U}}} \circ v_K^{\mathcal{U}}$ (REMARK 3.2.1),
- $v_{L^{\mathcal{U}}} L^{\mathcal{U}} \subseteq v_{K^{\mathcal{U}}} K^{\mathcal{U}}$ is a convex subgroup and thus also $v_{K^{\mathcal{U}}} K^{\mathcal{U}}$ is not q-divisible, and

- $v_{K^{\mathcal{U}}} K^{\mathcal{U}} / v_{L^{\mathcal{U}}} L^{\mathcal{U}} \cong v_K^{\mathcal{U}} K^{\mathcal{U}}$ is divisible, hence $v_{K^{\mathcal{U}}} K^{\mathcal{U}}$ is not q-antiregular.

We can then apply FACT 3.2.8 to $(K^{\mathcal{U}}, v_{K^{\mathcal{U}}})$ (note the different valuation!) and find an $\varnothing$-definable non-trivial henselian valuation on $K^{\mathcal{U}}$. Since $K \preceq K^{\mathcal{U}}$, the same $\mathcal{L}_{\text{ring}}$-formula defines a non-trivial henselian valuation ring on K. $\square$

3.4 | Defect

Defect – i.e. the quantity d in equation ($\star$) from SECTION 3.2 – measures how much of the finite extension of a henselian valued field (K, v) is not induced by extensions of its value group vK and residue field Kv. Originally introduced by Ostrowski (see [Roq02] for a thorough summary of his work), it was then extensively studied by Franz-Viktor Kuhlmann in his thesis [Kuh90] and his subsequent work. While in many classical number theoretic situations, defect is trivial, it is a crucial invariant in the model theory of valued fields, in a certain sense *limiting* the Ax-Kochen/Ershov philosophy. In residue characteristic zero, every extension is defectless; this is why, for example, the classical Ax-Kochen/Ershov theorem can be proven for all henselian equicharacteric zero valued fields.

Valued fields of positive residue characteristic showcase wildly different behaviour. In particular, the fundamental equality can fail, i.e. there might be an extension with $d > 1$. This is usually a serious obstruction to most efforts in understanding complete theories of valued fields. Throughout this section, however, we will deploy such extensions to define henselian valuations.

A toy situation. Suppose that $v \in H_1(K)$ and there exists an $\mathcal{L}_{\text{ring}}$-definable subset D of K with

$$v(D) = \{\gamma \in vK \mid \gamma > H\},$$

where H is a proper convex subgroup of vK.

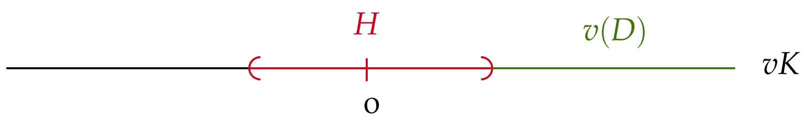

We will prove that whenever (K, v) is as above, then the coarsening of v corresponding to H is definable, see 3.2. This is a simplified version of the situation that will show up later when working with independent defect extensions. All the main ideas are already contained in this proof, with the advantage that we can avoid the technical subtleties of working modulo an interpretation of a finite field extension, which make the proof of THEOREM 3.4.13 cumbersome.

Remark 3.4.1. We will repeatedly use Beth's definability theorem ([Hod97, Theorem 5.5.4]). In our case, this means that in order to show that a certain subset $D \subseteq K^n$ is $\mathcal{L}_{\text{ring}}(\bar{c})$-definable on K, where $\bar{c}$ are some candidate parameters, it is enough to work in the augmented language $\mathcal{L} = \mathcal{L}_{\text{ring}}(\bar{c}) \cup \{P(X_1, \dots X_n)\}$, and show that whenever we take two structures $(L, \bar{c}', D_1), (L, \bar{c}', D_2) \equiv_{\mathcal{L}} (K, \bar{c}, D)$, we have $D_1 = D_2$. In particular, in the proof of LEMMA 3.4.2, we will consider valuation rings and thus unary predicates.

Lemma 3.4.2. *Let (K, v) be a valued field such that $v \in H_1(K)$. Let $D \subseteq K$ be an $\mathcal{L}_{\text{ring}}$-definable subset such that, for some proper convex subgroup $H \subseteq vK$ (possibly equal to the trivial subgroup),*

$$\{\alpha \in vK \mid \alpha > H\} = \{v(d) \mid d \in D\}.$$

Let v_H be the coarsening of v corresponding to H and denote by $\mathcal{O}_H$ its valuation ring. Then v_H is definable.

Proof. We will use Beth's definability theorem as explained in REMARK 3.4.1. We work in $\mathcal{L}_{\text{val}}$: let L and $\mathcal{O}_1, \mathcal{O}_2 \subseteq L$ be such that $(L, \mathcal{O}_1), (L, \mathcal{O}_2) \equiv (K, \mathcal{O}_H)$. In particular, both $\mathcal{O}_1$ and $\mathcal{O}_2$ are henselian valuation rings on L, corresponding to valuations v_1 and v_2. We want to show that $\mathcal{O}_1 = \mathcal{O}_2$. Note that, since $v_H \in H_1(K)$, the same holds for v_1 and v_2, and thus $\mathcal{O}_1$ and $\mathcal{O}_2$ are comparable. Without loss of generality, assume $\mathcal{O}_1 \subseteq \mathcal{O}_2$. This implies that $\mathcal{O}_1^\times \subseteq \mathcal{O}_2^\times$. We note that, if $D = \psi(K)$ for some $\mathcal{L}_{\text{ring}}$-formula ψ (possibly with parameters), then

$$(K, \mathcal{O}_H) \models \forall x \left(x \notin \mathcal{O}_H \iff \exists u \exists y \left(u \in \mathcal{O}_H^\times \wedge \psi(y) \wedge \left(x = u \frac{1}{y} \right) \right) \right).$$

The same $\mathcal{L}_{\text{val}}$-sentence is then true in both $(L, \mathcal{O}_1)$ and $(L, \mathcal{O}_2)$.

Now, towards proving $\mathcal{O}_1 = \mathcal{O}_2$, assume for a contradiction that there is $x \in \mathcal{O}_2$ but $x \notin \mathcal{O}_1$: then, there exist $u \in \mathcal{O}_1^\times \subseteq \mathcal{O}_2^\times$ and $d \in \psi(L)$ such that $x = u\frac{1}{d}$. In particular, then, by the $\mathcal{L}_{\text{val}}$-sentence above, $x \notin \mathcal{O}_2$, a contradiction. It follows that $\mathcal{O}_1 = \mathcal{O}_2$.

We can apply Beth's definability theorem and get that $\mathcal{O}_H = \varphi(K)$ for some $\mathcal{L}_{\text{ring}}$-formula $\varphi(x)$. □

Remark 3.4.3. Beth's definability theorem grants some control over parameters; namely, if D was $\varnothing$-definable to begin with, then so is $\mathcal{O}_H$. Moreover, the same argument works if we replace $\mathcal{L}_{\text{ring}}$ with some expansion $\mathcal{L}$; we then get that v_H is $\mathcal{L}$-definable.

Independent defect. Independent defect will be the central tool of our definability proof; the set Σ_L which will be defined in a moment, while not $\mathcal{L}_{\text{ring}}$-definable per se, will contain enough information to $\mathcal{L}_{\text{ring}}$-define a henselian valuation ring modulo the interpretation of a finite field extension.

Note that every equicharacteristic 0 valued field is defectless, so for the rest of this section we assume that our valued fields have residue characteristic $p > 0$.

Definition 3.4.4 ([KR23, Introduction]). Let $(K, v) \subseteq (L, v)$ be a Galois extension of degree p with defect. For any $\sigma \in \text{Gal}(L|K) \setminus \{\text{id}\}$, we let

$$\Sigma_\sigma := \left\{ v\left(\frac{\sigma f - f}{f} \right) \,\middle|\, f \in L^\times \right\}.$$

Fact 3.4.5 ([KR23, Theorems 3.4 and 3.5])**.** *Σ_σ does not depend on the choice of σ, and is a final segment of vK_∞. We denote it by Σ_L.*

Definition 3.4.6 ([KR23, Introduction])**.** Let $(K,v) \subseteq (L,v)$ be Galois of degree p with defect. The extension $(K,v) \subseteq (L,v)$ has *independent defect* if there is a (possibly trivial) proper convex subgroup $H \subseteq vK$ such that vK/H does not have a minimum positive element, and further

$$\Sigma_L = \{\alpha \in vK_\infty \mid \alpha > H\}.$$

We say that the extension has *dependent defect* otherwise.

Definition 3.4.7 ([KR23, Introduction])**.** Let (K,v) be a henselian valued field of positive residue characteristic p. If $\mathrm{char}(K) = 0$, then let $K' := K(\zeta_p)$ where ζ_p is a primitive pth root of unity. Otherwise let $K' := K$. We denote the unique extension of v to K' also by v. We say that (K,v) is an *independent defect field* if all Galois defect extensions of (K',v) of degree p have independent defect.

Remark 3.4.8. Any perfect field of positive characteristic is automatically an independent defect field: indeed, by [KR23, Proposition 3.9], the existence of a dependent defect extension requires the existence of an (immediate) purely inseparable extension (that in particular has to lie outside of the completion of the valued field).

We now prove a result on composition of independent defect valuations with defectless valuations, which will be quite useful later.

Lemma 3.4.9. *Let (K,v) be a henselian valued field with $v = \bar{v} \circ w$, such that (K,w) is defectless and $(Kw,\bar{v})$ is a perfect independent defect field. Then (K,v) is an independent defect field.*

Proof. Assume that $v = \bar{v} \circ w$, with w defectless and $(Kw,\bar{v})$ a perfect independent defect field. If (K,v) has mixed characteristic $(0,p)$ we assume without loss of generality that K contains a primitive p-th root of unity.

$$\begin{tikzcd}
K \arrow[r, "w", "\text{defectless}"'] \arrow[rr, bend left, "v"] & Kw \arrow[r, "\bar{v}", "\text{independent defect}"'] & (Kw)\bar{v} = Kv
\end{tikzcd}$$

Let (L,v) be a Galois defect extension of degree p of (K,v). In particular, this extension is immediate, i.e., $Lv = Kv$ and $vL = vK$. We want to show that $(L,v) \supseteq (K,v)$ has independent defect.

We first argue that $(Kw,\bar{v}) \subseteq (Lw,\bar{v})$ is a defect extension. We prove that it is immediate of degree p. As (K,w) is defectless, we have the fundamental equality

$$p = [L:K] = [Lw:Kw](wL:wK).$$

We argue that we must have $[Lw:Kw] = p$ and $wL = wK$. Suppose not, i.e. that $Lw = Kw$: then $\bar{v}(Lw) = \bar{v}(Kw)$, and thus

$$wL = {vL}\big/{\bar{v}(Lw)} = {vK}\big/{\bar{v}(Kw)} = wK,$$

contradicting $(wL : wK) = p$.

Using $wL = wK$ and $vL = vK$, we get

$$ {}^{vK}\!\big/_{\bar{v}(Kw)} = wK = wL = {}^{vL}\!\big/_{\bar{v}(Lw)} = {}^{vK}\!\big/_{\bar{v}(Lw)}, $$

and thus $\bar{v}(Lw) = \bar{v}(Kw)$. Moreover, $(Lw)\bar{v} = Lv = Kv = (Kw)\bar{v}$. This shows that $(Kw, \bar{v}) \subseteq (Lw, \bar{v})$ is an immediate extension of degree p, in particular a defect extension. Note that as Kw is perfect, the extension is separable; moreover, since $K \subseteq L$ is normal, the same holds for $Kw \subseteq Lw$ (for example, by [EP05, Proposition 3.2.16]), and thus the latter is Galois. Since $(Kw, \bar{v})$ is an independent defect field, the extension $(Kw, \bar{v}) \subseteq (Lw, \bar{v})$ has independent defect. By definition, there is a convex subgroup H of $\bar{v}(Kw)$ such that ${}^{\bar{v}(Kw)}\!\big/_{H}$ has no minimum positive element, and further

$$ \Sigma_{Lw} := \left\{ \bar{v}\left(\frac{f - \tau(f)}{f}\right) \,\middle|\, f \in (Lw)^{\times} \right\} = \{\gamma \in \bar{v}(Kw)_{\infty} \mid \gamma > H\}, $$

where $\tau \in \mathrm{Gal}(Lw|Kw)$ is a generator, i.e. $\langle \tau \rangle = \mathrm{Gal}(Lw|Kw)$. Since $\bar{v}(Kw) \subseteq vK$ is a convex subgroup, H is also a convex subgroup of the bigger group vK. Then, the embedding $\bar{v}(Kw) \subseteq vK$ gives rise to an embedding ${}^{\bar{v}(Kw)}\!\big/_{H} \subseteq {}^{vK}\!\big/_{H}$ as convex subgroup, thus ${}^{vK}\!\big/_{H}$ also has no minimum positive element. We will now show that

$$ \Sigma_{L} := \left\{ v\left(\frac{f - \sigma(f)}{f}\right) \,\middle|\, f \in L^{\times} \right\} = \{\gamma \in vK_{\infty} \mid \gamma > H\}, $$

where $\sigma \in \mathrm{Gal}(L|K)$ is such that $\langle \sigma \rangle = \mathrm{Gal}(L|K)$. This will prove that the defect in the extension $(L, v) \supseteq (K, v)$ is independent.

As Σ_L is a final segment in vK_{∞} (by FACT 3.4.5), it is enough to show that $\Sigma_L \cap \bar{v}(Kw)_{\infty} = \Sigma_{Lw}$.

Note that since (K, w) is henselian, $\mathcal{O}_{(L,w)}$ is fixed setwise by the action of $\mathrm{Gal}(L|K)$. Thus by [EP05, Proposition 3.2.16(3)], ϕ induces an automorphism $\bar{\phi} \in \mathrm{Gal}(Lw|Kw)$ given by

$$ \bar{\phi}(\mathrm{res}_w(f)) := \mathrm{res}_w(\phi(f)), \quad f \in \mathcal{O}_L. $$

We argue that the map

$$ \begin{cases} \mathrm{Gal}(L|K) & \to & \mathrm{Gal}(Lw|Kw) \\ \phi & \mapsto & \bar{\phi} \end{cases} $$

is a surjective group homomorphism (cf. the proof of [EP05, Lemma 5.2.6]): indeed, we write $Lw = Kw(\mathrm{res}_w(a))$, and we take any $\rho \in \mathrm{Gal}(Lw|Kw)$. We let $a_1, \dots a_p$ be the $\mathrm{Gal}(L|K)$-conjugates of a.

For some $i \leqslant p$, $\mathrm{res}_w(a_i) = \rho(\mathrm{res}_w(a))$, so there is $\sigma \in \mathrm{Gal}(L|K)$ such that

$$ \mathrm{res}_w(\sigma(a)) = \bar{\sigma}(\mathrm{res}_w(a)) = \rho(\mathrm{res}_w(a)). $$

Thus $\bar{\sigma}$ and ρ coincide on Lw.

Now, since $(\phi \mapsto \bar{\phi})$ is surjective between finite groups with the same order, it is an isomorphism. It follows that it maps generators of $\mathrm{Gal}(L|K)$ to generators of $\mathrm{Gal}(Lw|Kw)$.

We can now compare the sets Σ_L and Σ_{Lw}.

Step 1: $\Sigma_{Lw} \subseteq \Sigma_L \cap \bar{v}(Kw)_\infty$.

Fix $\sigma \in \mathrm{Gal}(L|K)$ such that $\langle\sigma\rangle = \mathrm{Gal}(L|K)$. Then, since the set Σ_{Lw} does not depend on the choice of the generator, we can write

$$\begin{aligned}
\Sigma_{Lw} &= \left\{ \bar{v}\left(\frac{f - \bar{\sigma}(f)}{f}\right) \,\middle|\, f \in (Lw)^\times \right\} \\
&= \left\{ \bar{v}\left(\frac{\mathrm{res}_w(f) - \bar{\sigma}(\mathrm{res}_w(f))}{\mathrm{res}_w(f)}\right) \,\middle|\, f \in \mathcal{O}^\times_{(L,w)} \right\} \\
&= \left\{ \bar{v}\left(\mathrm{res}_w\left(\frac{f - \sigma(f)}{f}\right)\right) \,\middle|\, f \in \mathcal{O}^\times_{(L,w)} \right\} \\
&= \left\{ v\left(\frac{f - \sigma(f)}{f}\right) \,\middle|\, f \in \mathcal{O}^\times_{(L,w)} \right\} \cap \bar{v}(Kw)_\infty \subseteq \Sigma_L \cap \bar{v}(Kw)_\infty,
\end{aligned}$$

since $\bar{v}(\mathrm{res}_w(x)) = v(x)$ for $x \in L$ such that $v(x) \in \bar{v}(Kw)$.

Step 2: $\Sigma_{Lw} \supseteq \Sigma_L \cap \bar{v}(Kw)_\infty$.

Let $v\left(\frac{f-\sigma(f)}{f}\right) \in \Sigma_L \cap \bar{v}(Kw)$, i.e. $f \in L^\times$, then because $wL = wK$, there is $g \in K^\times$ with $w(f) = w(g)$ and so for $h := \frac{f}{g} \in \mathcal{O}^\times_{(L,w)}$ we have

$$\frac{f - \sigma(f)}{f} = \frac{h - \sigma(h)}{h},$$

so

$$v\left(\frac{f - \sigma(f)}{f}\right) = v\left(\frac{h - \sigma(h)}{h}\right) \in \Sigma_{Lw}. \qquad \square$$

Question 3.4.10. Let (K, v) be a henselian valued field with $v = \bar{v} \circ w$, such that both (K, w) and $(Kw, \bar{v})$ are independent defect fields. Must then (K, v) be an independent defect field?

Defining valuations from independent defect. We now have all the tools needed to build a definable henselian valuation out of a (independent) defect extension. This subsection is structured in several steps:

- THEOREM 3.4.13 contains the core of the argument, showing how to deploy an independent defect extension (and its associated convex subgroup) to build a definable valuation.
- COROLLARY 3.4.15 shows how to go from our assumptions in the Main Theorem to the situation where the right defect extension actually appears.
- COROLLARY 3.4.17 and LEMMA 3.4.18 then specialize to the mixed characteristic situation, where some extra care is needed.

Finally, in SECTION 3.4, we will deal with a scenario that mirrors the one of PROPOSITION 3.3.6, namely with the situation where (K, v_K) is defectless, but some elementary extension L of Kv_K is henselian and v_L admits defect.

The following theorem should be thought of as a more elevated version of LEMMA 3.4.2, where the $\mathcal{L}_{\text{ring}}$-definable set D now appears as a subset of the cartesian product K^p, which we identify along an interpretation with a degree p Galois defect extension. For details on interpretations of finite field extensions, see for example [Cha09, Section 3.5].

We will often use [JK15, Theorem 3.10] to reduce to the case where our valuation of interest is in H_1. We briefly explain how this is done in the following remarks.

Remark 3.4.11. Note that if $F \subseteq E$ is a finite separable field extension of degree n, we can interpret E inside of F using the coefficients $\bar{c}$ of the minimal polynomial of a primitive element of the extension. Suppose we had a definable valuation v on E, say defined by $\varphi(X, \overline{d})$. Using the interpretation we can define $\mathcal{O}_v$ (as a subset of F^n) using some other formula $\widetilde{\varphi}(\overline{X}, \overline{d'})$, where now $\overline{d'}$ is given by the coordinates of the elements of $\overline{d}$ in a fixed F-basis of F^n, together with $\bar{c}$. The restriction $v|_F$ can then be defined using the formula $\widetilde{\varphi}(X, \underbrace{0, \ldots, 0}_{(n-1)\text{-times}}, \overline{d'})$. Assume now that the coefficients $\bar{c}$ of the minimal polynomial are integers: then, if v is $\emptyset$-definable, so is $v|_F$.

Remark 3.4.12. Suppose that K is not separably closed, and let $v \in H_2(K)$. Then, [JK15, Theorem 3.10] produces a non-trivial $\emptyset$-definable henselian valuation w on K. A closer look at the proof of [JK15, Theorem 3.10] allows us to argue that w is a coarsening of v. If $\text{char}(K) = p > 0$, or K contains a primitive p-th root of unity, then by construction w is a (possibly non-proper) coarsening of v_K, and thus in particular of v. Otherwise, one moves to $L = K(\zeta_p)$ for some primitive p-th root of unity ζ_p. On L, one finds a coarsening w' of v_L which is $\emptyset$-definable and non-trivial. Its restriction $w := w'|_K$ is then still a $\emptyset$-definable coarsening of v_K ([EP05, Theorem 4.4.3] – with the caveat that $H(K)$ means $H_1(K) \cup \{\mathcal{O}_{v_K}\}$ here).

Theorem 3.4.13. *Let (K, v) be a henselian valued field with* $\text{char}\, Kv = p > 0$ *such that (K, v) admits a Galois extension of degree p with independent defect. If* $\text{char}(K) = 0$, *we additionally assume that $\zeta_p \in K$, where ζ_p is a primitive p-th root of unity. Then K admits a non-trivial definable henselian valuation, coarsening v.*

Proof. Note that we can assume that $v \in H_1(K)$, by REMARK 3.4.12. Let $K \subseteq L$ be the degree p Galois extension with independent defect and take θ such that $L = K(\theta)$. Let σ be a generator of the (cyclic) Galois group $\text{Gal}(L|K)$, and let

$$D := \left\{ \frac{\sigma(f) - f}{f} \;\middle|\; f \in L^\times \right\}.$$

As the extension has independent defect, there is a (possibly trivial) convex subgroup $H \subsetneq vK$ such that

$$\Sigma_L = v(D) = \{\alpha \in vK_\infty \mid \alpha > H\}.$$

Recall that $L = K(\theta)$ can be interpreted in K via the K-linear isomorphism

$$f\colon \begin{cases} K^p & \to & L = K(\theta) \\ (a_0, \ldots, a_{p-1}) & \mapsto & \sum_{i=0}^{p-1} a_i \theta^i \end{cases}$$

Note that the coefficients of the minimal polynomial of the generator θ over K are needed as parameters to describe the multiplication. Now, the action of σ is definable in K (via the interpretation) using the coefficients of the change of basis matrix. If $\mathrm{char}(K) = p$ then the matrix has integer coefficients, so no additional parameter is needed. If $\mathrm{char}(K) = 0$, then we need to use ζ_p. We will denote the parameter tuple needed to define D in K by $\bar{c}$ (i.e. the coefficients of the minimal polynomial of the generator θ over K, and ζ_p if necessary).

Thus, $f^{-1}(D) \subseteq K^p$ is definable in K, using $\bar{c}$ as parameters; choose an $\mathcal{L}_{\mathrm{ring}}(\bar{c})$-formula $\psi(\overline{X}, \bar{c})$ so that $f^{-1}(D) = \psi(K^p, \bar{c})$. Denote by $\mathcal{O}_H \subseteq L$ the coarsening of (the unique extension of) v corresponding to H. We can see $\mathcal{O}_H$ as a subset $f^{-1}(\mathcal{O}_H)$ of K^p.

Claim: *$f^{-1}(\mathcal{O}_H) \subseteq K^p$ is $\mathcal{L}_{\mathrm{ring}}(\bar{c})$-definable.*

We will use Beth's definability theorem. We work in $\mathcal{L} := \mathcal{L}_{\mathrm{ring}}(\bar{c}) \cup \{P(\overline{X})\}$, where $P(\overline{X})$ is a p-ary predicate. Under the bijection f that establishes the interpretation, the predicate will represent the valuation subring $\mathcal{O}_H$ of $K(\theta)$. We now let $(M, \bar{c}', P_1), (M, \bar{c}', P_2) \equiv_{\mathcal{L}} (K, \bar{c}, f^{-1}(\mathcal{O}_H))$. Note that the interpretation given by f gives rise to an $\mathcal{L}_{\mathrm{ring}}$-structure N and a map f' which is establishing an interpretation of N in M.

Furthermore, because $(M, \bar{c}', P_1)$ and $(M, \bar{c}', P_2)$ are elementarily equivalent, we have that:

- $M \subseteq N$ are fields, and the polynomial with coefficients from $\bar{c}'$ is of degree p and irreducible over M; thus, the field extension $M \subseteq N$ is algebraic of degree p;
- $\mathcal{O}_1 = f'(P_1)$ and $\mathcal{O}_2 = f'(P_2)$ are non-trivial henselian valuation rings on N; we denote the corresponding valuations by v_1 and v_2, respectively;
- there is $\ell \geqslant 2$ such that there is a separable polynomial of degree ℓ over Lv_H with no root in Lv_H; thus, there are separable polynomials of degree ℓ over Nv_1 and Nv_2 with no root, and thus $v_1, v_2 \in H_1(N)$,
- for $i = 1, 2$, $x \notin \mathcal{O}_i$ if and only if there are $u \in \mathcal{O}_i^\times$ and $y \in \psi(M^p, \bar{c}')$ such that $x = u\frac{1}{f'(y)}$.

Since $\mathcal{O}_1, \mathcal{O}_2 \in H_1(N)$, they must be comparable. Assume, for example, that $\mathcal{O}_1 \subseteq \mathcal{O}_2$. Towards a contradiction, assume now that there is $x \in \mathcal{O}_2$ such that $x \notin \mathcal{O}_1$. In particular, this means that there are $u \in \mathcal{O}_1^\times$ and $y \in \psi(M^p, \bar{c}')$ such that $x = u\frac{1}{f'(y)}$. But since $\mathcal{O}_1^\times \subseteq \mathcal{O}_2^\times$, and thus $u \in \mathcal{O}_2^\times$, we get that $x \notin \mathcal{O}_2$, a contradiction. Then $\mathcal{O}_1 = \mathcal{O}_2$, so $P_1 = P_2$ and we can apply Beth's definability: there is an $\mathcal{L}_{\mathrm{ring}}$-formula $\theta(\overline{X}, \bar{c})$, such that $f^{-1}(\mathcal{O}_H) = \theta(K^p, \bar{c})$. $\blacksquare_{\text{Claim}}$

Now, we can define the (non-trivial) restriction $\mathcal{O}_H \cap K$ as follows: given $z \in K$, we have that

$$z \in \mathcal{O}_H \cap K \iff K \models \theta(z, \underbrace{0, \dots, 0}_{(p-1)\text{-times}}, \bar{c}).$$

This exhibits a non-trivial definable henselian valuation ring on K. □

We can now use this theorem to define a non-trivial henselian valuation on certain valued fields with defect, but first we need to produce the Galois defect extensions required to exhibit independent defect.

Lemma 3.4.14. *Let (K, v) be a perfect valued field of residue characteristic $p > 0$. If (K, v) has defect, then there is a finite extension $K \subseteq K'$ such that K' admits a defect Galois extension of degree p.*

Proof. Take a defect extension $K \subseteq L$. In particular, $p^n \mid [L : K]$ for some $n \geqslant 1$. We let N be the normal hull of L; note that $p^n \mid [N : K]$ and $K \subseteq N$ is Galois. Consider now $H \subseteq \mathrm{Gal}(N|K)$ to be a p-Sylow subgroup. Then $L' := N^H$ is such that $L' \subseteq N$ is Galois and it is a tower $L' = L_0 \subseteq L_1 \subseteq \cdots \subseteq L_n = N$ of normal degree p extensions. Moreover, since $K \subseteq L'$ has degree prime to p, it is defectless, hence $L' \subseteq N$ is defect. In particular, there is some $\ell \leqslant n$ such that $L_\ell \subseteq L_{\ell+1}$ is a defect Galois extension of degree p, so we take $K' = L_\ell$. □

Corollary 3.4.15. *Suppose (K, v) is a henselian valued field with* $\operatorname{char} Kv = p > 0$ *such that (K, v) has defect, and such that either*

1. *K is perfect, if* $\operatorname{char} K = p > 0$, *or*

2. *$\mathcal{O}_v/p$ is semi-perfect, if* $\operatorname{char} K = 0$.

Then K admits a non-trivial definable henselian valuation, coarsening v.

Proof. We may assume that $v \in H_1(K)$, by REMARK 3.4.12.

In the case where $\operatorname{char} K = 0$, we make the following two observations:

- since K always admits a definable valuation if v is finitely ramified, we may assume that it is infinitely ramified,

- the assumptions on (K, v) imply that it is a *roughly deeply ramified* valued field in the sense of [KR23, p. 2696] (note that by [KR23, Lemma 4.1], we only need to check that $\mathcal{O}_v/p$ is semi-perfect, without looking at the completion of K). Thus, by [KR23, Theorem 1.8] any finite extension of it is also roughly deeply ramified.

If $\operatorname{char} K = 0$, let $K_0 = K(\zeta_p)$, where ζ_p is some primitive pth root of unity. Otherwise, take $K_0 = K$. By LEMMA 3.4.14, we have a finite extension $K_0 \subseteq K_1$ which admits a Galois defect extension of degree p. Since $K \subseteq K_1$ is a finite separable field extension, we can find $b \in K_1$ such that $K_1 = K(b)$.

The field K_1 admits a Galois defect extension of degree p, i.e there is $c \in K_1$ which has no Artin-Schreier root in K if $\operatorname{char} K = p$ (resp. no pth root, if $\operatorname{char} K = 0$), and the extension $(K_1, v) \subseteq (K_1(\theta), v)$, where $\theta^p - \theta = c$ (resp. $\theta^p = c$), is an immediate defect extension. Denote by $L = K_1(\theta) = K(b, \theta)$. Then, we have to distinguish two cases:

- if $\operatorname{char} K = p$, then K_1 is a perfect field,
- if $\operatorname{char} K = 0$, then K_1 is a finite extension of a roughly deeply ramified field, and thus it is roughly deeply ramified itself.

In both cases, K_1 is an independent defect field (see [KR23, Theorem 1.10]), and thus the extension $(K_1, v) \subseteq (L, v)$ is a degree p defect extension with independent defect. By THEOREM 3.4.13, then, K_1 admits a non-trivial definable henselian valuation v_H, with valuation ring defined by $\varphi(X, \bar{c})$. As $K_1 = K(b)$ can be interpreted in K, by REMARK 3.4.11 the restriction of v_H to K is then $\mathcal{L}_{\text{ring}}$-definable with parameters. □

As we will later have to deal with definable valuations in some ultrapower of our field K, we immediately apply results of Anscombe and Jahnke from [AJ18] to get rid of parameters in our definition.

Remark 3.4.16. Note that if (K, v) is a henselian valued field of mixed characteristic, with $v \in H_1(K)$, then v_K is also of mixed characteristic.

Corollary 3.4.17. *Suppose (K, v) is a henselian valued field of mixed characteristic, such that (K, v) has defect and $\mathcal{O}_v/p$ is semi-perfect. Then K admits an $\varnothing$-definable non-trivial henselian valuation.*

Proof. We may assume that $v \in H_1(K)$, by REMARK 3.4.12. Now, v_K is also of mixed characteristic, thus COROLLARY 3.4.15 yields the existence of a definable non-trivial henselian valuation, and by [AJ18, Theorem 1.1.(B)] we can conclude that there is an $\varnothing$-definable non-trivial henselian one. □

Later, in the proof of PROPOSITION 3.4.19, we will find ourselves in the situation where we deal with a defect refinement v_1 of a valuation $v_K^{\mathcal{U}}$, where $K^{\mathcal{U}} v_K^{\mathcal{U}}$ is perfect.

Lemma 3.4.18. *Suppose (K, v) is a henselian valued field of mixed characteristic, such that (K, v) has defect, and there is a coarsening w of v such that (K, w) is defectless and Kw is perfect of characteristic $p > 0$. Then K admits an $\varnothing$-definable non-trivial henselian valuation.*

Proof. We may assume that $v \in H_1(K)$: otherwise, we can use [JK15, Theorem 3.10] to find an $\varnothing$-definable non-trivial henselian valuation. Using LEMMA 3.4.14, there is a finite extension $K(\zeta_p) \subseteq K_1$ such that (K_1, v) admits a defect extension of degree p, and we find $b \in K_1$ such that $K_1 = K(b)$.

We now show that (K_1, v) is an independent defect field. For this we consider the decomposition of (K_1, v) with respect to the coarsening (K_1, w) and we denote the induced valuation on $K_1 w$ by $\bar{v}$:

$$\begin{tikzcd} K_1 \arrow[r, "w"] \arrow[rr, bend left, "v"] & K_1 w \arrow[r, "\bar{v}"] & (K_1 w)\bar{v} = K_1 v. \end{tikzcd}$$

Note that (K_1, w) is defectless, because it is a finite extension of (K, w); moreover, $K_1 w$ is perfect, because it is a finite extension of Kw, and thus $(K_1 w, \bar{v})$ is an independent defect field. From LEMMA 3.4.9 it follows that also (K_1, v) is an independent defect field.

By THEOREM 3.4.13, then, K_1 admits a non-trivial definable henselian valuation v_H, with valuation ring defined by $\varphi(X, \bar{c})$. As $K_1 = K(b)$ can be interpreted in K, by REMARK 3.4.11 the restriction of v_H to K is then $\mathcal{L}_{\text{ring}}$-definable with parameters (namely, the coefficients of the minimal polynomial of b, and the coordinates of $\bar{c}$ in a chosen K-basis of K_1).

By our assumptions, v_K is also of mixed characteristic, thus [AJ18, Theorem 1.1.(B)] yields the existence of an $\varnothing$-definable non-trivial henselian valuation. □

Defect in elementary extensions of Kv_K. The following proposition is a direct result of LEMMA 3.4.18, while we will need some more work for positive characteristic.

Proposition 3.4.19. *Let K be a henselian field such that (K, v_K) has mixed characteristic and Kv_K is perfect. If*

1. *there is $L \equiv Kv_K$ such that (L, v_L) has defect*, and
2. *(K, v_K) is defectless,*

then K admits an $\varnothing$-definable non-trivial henselian valuation.

Proof. Choose L such that $L \equiv Kv_K$ and (L, v_L) has defect. By the Keisler-Shelah isomorphism theorem (FACT 3.2.9), take an index set I and an ultrafilter $\mathcal{U}$ on I such that $(Kv_K)^{\mathcal{U}} \cong L^{\mathcal{U}}$. Consider the ultrapowers $(K^{\mathcal{U}}, v_K^{\mathcal{U}}) := (K, v_K)^{\mathcal{U}}$ and $(L^{\mathcal{U}}, v_L^{\mathcal{U}}) := (L, v_L)^{\mathcal{U}}$ in $\mathcal{L}_{\text{val}}$. Since $v_L^{\mathcal{U}}$ is a valuation with defect, while $v_K^{\mathcal{U}}$ is defectless, the composition v_1 is a non-trivial henselian valuation on $K^{\mathcal{U}}$ with defect. The valuation v_1 has defect and of mixed characteristic, since $v_K^{\mathcal{U}}$ is, and its coarsening $v_K^{\mathcal{U}}$ is defectless and has perfect residue field $K^{\mathcal{U}} v_K^{\mathcal{U}} = L^{\mathcal{U}}$. Thus by LEMMA 3.4.18, there is an $\varnothing$-definable valuation on $K^{\mathcal{U}}$. By elementary equivalence, we can find an $\varnothing$-definable henselian valuation on K. □

In the case of positive characteristic, eliminating the parameters from the definition of the henselian valuation can be tricky; indeed, a theorem along the lines of [AJ18, Theorem 1.1.(B)] will not be true in general in positive characteristic. However, in our setting (particularly because we can assume $v_K K$ to be divisible) we manage to get an analogous result, where we take a valuation defined in an elementary extension $K \preceq K^*$ (with parameters possibly in K^*) and produce a valuation defined with parameters in K, thus allowing us to "push it down".

Fact 3.4.20 (cf. [Jah, Proposition 2.4])**.** *Let (K, v) be a henselian valued field with non-separably closed nor real closed residue field. Then v has an $\mathcal{L}_{\mathrm{ring}}$-definable refinement.*

Proposition 3.4.21. *Let K be a perfect henselian field of positive characteristic. If*

1. *$v_K K$ is divisible,* and
2. *there is $L \equiv K v_K$ such that (L, v_L) has defect,* and
3. *(K, v_K) is defectless,*

then K admits a definable non-trivial henselian valuation.

Proof. Note that we may assume that Kv_K is not separably closed; otherwise, FACT 3.2.8 gives an $\varnothing$-definable non-trivial henselian valuation.

By the Keisler-Shelah isomorphism theorem (FACT 3.2.9), take an index set I and an ultrafilter $\mathcal{U}$ on I such that $(Kv_K)^{\mathcal{U}} \cong L^{\mathcal{U}}$ in $\mathcal{L}_{\mathrm{ring}}$. Consider the ultrapowers $(K^{\mathcal{U}}, v_K^{\mathcal{U}}) := (K, v_K)^{\mathcal{U}}$ and $(L^{\mathcal{U}}, v_L^{\mathcal{U}}) := (L, v_L)^{\mathcal{U}}$. Since $v_L^{\mathcal{U}}$ is a valuation with defect, the composition $v_1 = v_L^{\mathcal{U}} \circ v_K^{\mathcal{U}}$ is a non-trivial henselian valuation on $K^{\mathcal{U}}$ with defect. Thus, by COROLLARY 3.4.15, there is an $\mathcal{L}_{\mathrm{ring}}(K^{\mathcal{U}})$-definable non-trivial henselian valuation v_1' coarsening v_1, say given by $\varphi(K^{\mathcal{U}}, \overline{\beta})$ for some parameters $\overline{\beta} \in (K^{\mathcal{U}})^{\ell}$, $\ell \geqslant 1$.

Since Kv_K is not separably closed, FACT 3.4.20 yields an $\mathcal{L}_{\mathrm{ring}}(K)$-definable valuation w on K (not necessarily henselian!) such that $\mathcal{O}_w \subsetneq \mathcal{O}_{v_K}$. Let $\psi(X, \overline{c})$ be the formula defining w. Denote by w^* the valuation corresponding to $\psi(K^{\mathcal{U}}, \overline{c})$. Note that since $(K, v_K) \preceq (K^{\mathcal{U}}, v_K^{\mathcal{U}})$, then we still have that w^* is a proper refinement of $v_K^{\mathcal{U}}$.

As in the proof of [EP05, Theorem 2.3.4], given two valuation rings $\mathcal{O}_1$ and $\mathcal{O}_2$, we let

$$\mathcal{O}_1 \cdot \mathcal{O}_2 := \left\{ \frac{x}{y} \,\middle|\, x \in \mathcal{O}_1,\ y \in \mathcal{O}_1 \backslash \mathfrak{m}_2 \right\}.$$

Note that $\mathcal{O}_1 \cdot \mathcal{O}_2$ contains both $\mathcal{O}_1$ and $\mathcal{O}_2$ and is always a valuation ring. Indeed, it is the finest valuation ring containing both: if $\mathcal{O}_1, \mathcal{O}_2 \subseteq \mathcal{O}_3$, then for any $x \in \mathcal{O}_1$ and $y \in \mathcal{O}_1 \backslash \mathfrak{m}_2$, one has that $y \notin \mathfrak{m}_3$, so $\frac{1}{y} \in \mathcal{O}_3$ and thus $\frac{x}{y} \in \mathcal{O}_3$. We now consider the $\mathcal{L}_{\mathrm{ring}}(\overline{c})$-definable set

$$X = \left\{ \overline{b} \in (K^{\mathcal{U}})^{\ell} \mid \varphi(K^{\mathcal{U}}, \overline{b}) \text{ is a valuation ring and } \mathcal{O}_{w^*} \cdot \varphi(K^{\mathcal{U}}, \overline{b}) \neq K^{\mathcal{U}} \right\}.$$

For each $\overline{b} \in X$, we consider the $\mathcal{L}_{\mathrm{ring}}(\overline{c}, \overline{b})$-definable valuation ring

$$\mathcal{O}_{\overline{b}} := \mathcal{O}_{w^*} \cdot \varphi(K^{\mathcal{U}}, \overline{b}),$$

with corresponding valuation $v_{\overline{b}}$. As all coarsenings of a given valuation are comparable, any $v_{\overline{b}}$ for $\overline{b} \in X$ is comparable with $\mathcal{O}_{v_K^{\mathcal{U}}}$.

Since $v_K^{\mathcal{U}}$ is a tame valuation with divisible value group, LEMMA 3.3.4 implies that $v_{\overline{b}}$ cannot be a proper coarsening. In particular, then, $\mathcal{O}_{v_{\overline{b}}} \subseteq \mathcal{O}_{v_K^{\mathcal{U}}}$. It follows that the union

$$\mathcal{O} := \bigcup_{\overline{b} \in X} \mathcal{O}_{\overline{b}}$$

is a non-trivial $\mathcal{L}_{\text{ring}}(\overline{c})$-definable (note the parameters!) valuation ring on $K^{\mathcal{U}}$.

To show that $\mathcal{O}$ is henselian, we will show that $\mathcal{O}$ contains $\varphi(K^{\mathcal{U}}, \overline{\beta})$, i.e. that $\overline{\beta} \in X$. It is enough to show that $\mathcal{O}_{\overline{\beta}} \neq K^{\mathcal{U}}$. Since $\varphi(K^{\mathcal{U}}, \overline{\beta})$ is comparable with $v_K^{\mathcal{U}}$, we have two possible cases. Either $\varphi(K^{\mathcal{U}}, \overline{\beta})$ is coarser than $v_K^{\mathcal{U}}$, in which case $\mathcal{O}_{\overline{\beta}} = \varphi(K^{\mathcal{U}}, \overline{\beta})$, or $\varphi(K^{\mathcal{U}}, \overline{\beta})$ is finer than $v_K^{\mathcal{U}}$, in which case $\mathcal{O}_{\overline{\beta}}$ is also a refinement of $v_K^{\mathcal{U}}$. Either way, $\mathcal{O}_{\overline{\beta}}$ is non-trivial.

We have just shown that $\mathcal{O}$ is a coarsening of $\varphi(K^{\mathcal{U}}, \overline{\beta})$, thus it is henselian. In particular, the same formula defines a non-trivial henselian valuation on K, as needed. □

3.5 | Proof of the main theorem

Theorem 3.5.1. *Let K be perfect, not separably closed, and henselian. If* $\operatorname{char} K = 0$ *and* $\operatorname{char} Kv_K = p > 0$, *assume that $\mathcal{O}_{v_K}/p$ is semi-perfect. Then K admits a definable non-trivial henselian valuation if and only if at least one of the following conditions hold:*

1. *Kv_K is separably closed,*
2. *Kv_K is not t-henselian,*
3. *there is $L \succeq Kv_K$ such that $v_L L$ is not divisible,*
4. *$v_K K$ is not divisible,*
5. *(K, v_K) has defect,*
6. *there is $L \succeq Kv_K$ such that (L, v_L) has defect.*

Remark 3.5.2. Note that in conditions 3 and 6, one automatically has that v_L is non-trivial (since the trivial valuation has divisible value group, and is defectless). Moreover, conditions 5 and 6 are trivial when $\operatorname{char} Kv_K = 0$ (for condition 6, see COROLLARY 3.2.3), and thus in that case the theorem reduces to THEOREM 3.1.1.

Proof of THEOREM 3.5.1. If $\operatorname{char} Kv_K = 0$, the statement is exactly THEOREM 3.1.1. This is because condition 5 is trivial (as equicharacteristic zero henselian valued fields are defectless), and COROLLARY 3.2.3 shows that condition 6 is also trivial in this case.

Assume $\operatorname{char} Kv_K = p > 0$ and note that Kv_K is perfect: in mixed characteristic, Kv_K is perfect since $\mathcal{O}_{v_K}/p$ is semi-perfect; in positive characteristic, Kv_K is perfect since K is. Thus, one direction is THEOREM 3.3.5.

In the other direction, we do a case distinction. First, assume $1 \vee 2 \vee 4$, then it follows from FACT 3.2.8 that K admits a definable non-trivial henselian valuation.

If $3 \wedge \neg 4$ then PROPOSITION 3.3.6 yields a existance of definable non-trivial henselian valuation.

Suppose now that 5 holds, i.e. (K, v_K) has defect. Then we get a non-trivial definable henselian valuation from COROLLARY 3.4.15.

Finally, assume that 6 holds, and since we have already established the case where $1 \vee 2 \vee 3 \vee 4 \vee 5$ holds, we can also assume $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge \neg 5$. Then by PROPOSITION 3.4.19 (in mixed characteristic) and PROPOSITION 3.4.21 (in positive characteristic), there is a definable non-trivial henselian valuation on K. $\square$

3.6 | Examples and questions

We isolate a few interesting examples where our result yields the existence of non-trivial definable henselian valuations. We then discuss the difficulties in building ones satisfying certain properties. We refer to the six conditions in the Main Theorem as conditions 1, 2, 3, 4, 5, and 6, and to their negations as conditions $\neg 1$, $\neg 2$, $\neg 3$, $\neg 4$, $\neg 5$, and $\neg 6$.

We start with the following observation.

Remark 3.6.1. Let (K, v) be a henselian valued field with Kv non-henselian and not separably closed. Then $v = v_K$ is the canonical henselian valuation on K. Indeed, there can be no proper henselian refinements of v, as they would correspond to non-trivial henselian valuations on Kv via REMARK 3.2.1. Thus $H_2(K)$ is empty and v is the finest valuation in $H_1(K)$.

The aim of this section is to give some examples for fields where our Main Theorem yields a definable non-trivial henselian valuation because of conditions 5 or 6 (while at the same time $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4$ holds). This boils down to finding a non-henselian t-henselian field that is not separably closed and such that we can control some properties of the canonical henselian valuation in some elementary extension (to access conditions 3 and 6).

The construction of non-henselian t-henselian fields goes back to Prestel and Ziegler, see [PZ78, p. 338]. Other instances can be found in [FJ15, Proposition 6.7] (note that the field constructed there is elementary equivalent to a field admitting a non-trivial henselian valuation with divisible value group) and in [JK17, Examples 3.8 and 5.4] (with an elementary extension L such that $v_L L$ is not divisible). All of these provide fields of characteristic 0. In [AJ18, Proposition 4.13], there is a construction of a non-henselian t-henselian field of positive characteristic that is elementary equivalent to a field admitting a non-trivial tame valuation with divisible value group. Such a field will serve as the residue field of the canonical henselian valuation to construct examples that satisfy condition 5, see EXAMPLE 3.6.8. We will also adapt their construction in LEMMA 3.6.13 and PROPOSITION 3.6.14 to serve as the residue field of the canonical henselian valuation of an example that satisfies condition 6.

The following is needed to check that the constructed examples satisfy condition ¬3.

Proposition 3.6.2. *Let K be non-henselian and t-henselian such that there is $K^* \equiv K$ admitting a non-trivial henselian valuation v^* with divisible value group. Then for all $L \equiv K$, $v_L L$ is divisible.*

Proof. We may assume that L is henselian and not separably closed. Otherwise v_L is trivial and $v_L L$ is the trivial group which is divisible. In particular, K is not separably closed.

By the Keisler-Shelah isomorphism theorem (FACT 3.2.9), there are an index set I and an ultrafilter $\mathcal{U}$ on I such that $F := (K^*)^{\mathcal{U}} \cong L^{\mathcal{U}}$. Let $(K^*, v^*)^{\mathcal{U}} =: (F, (v^*)^{\mathcal{U}})$ and $(L, v_L)^{\mathcal{U}} =: (F, (v_L)^{\mathcal{U}})$.

We claim that $(v^*)^{\mathcal{U}}$ and $(v_L)^{\mathcal{U}}$ are comparable. Suppose not: then, Lv_L is separably closed, and by FACT 3.2.8(1) L admits a non-trivial $\varnothing$-definable henselian valuation. The same formula then defines a $\varnothing$-definable non-trivial henselian valuation on K, contradicting the fact that K is non-henselian.

Now we have to treat two cases. First, if $(v_L)^{\mathcal{U}}$ is a coarsening of $(v^*)^{\mathcal{U}}$. Then $(v_L)^{\mathcal{U}}F$ is a quotient of $(v^*)^{\mathcal{U}}F$ modulo some convex subgroup. Since v^*K^* is divisible by assumption, so is $(v^*)^{\mathcal{U}}F = (v^*K^*)^{\mathcal{U}}$, which then implies that also $(v_L)^{\mathcal{U}}F = (v_L L)^{\mathcal{U}}$ and $v_L L$ are divisible.

Second, if $(v_L)^{\mathcal{U}}$ is a proper refinement of $(v^*)^{\mathcal{U}}$, then $(v^*)^{\mathcal{U}}F$ is a non-trivial quotient of $(v_L)^{\mathcal{U}}F$. We assume for a contradiction that $(v_L)^{\mathcal{U}}F$ is not p-divisible for some prime p. Now, $(v_L)^{\mathcal{U}}F$ is also not p-antiregular, since $(v^*)^{\mathcal{U}}F$ is a non-trivial quotient, which is divisible. Thus, by FACT 3.2.8(3), L and thus also K admit $\varnothing$-definable non-trivial henselian valuations, but K is non-henselian, a contradiction. Hence $(v_L)^{\mathcal{U}}F = (v_L L)^{\mathcal{U}}$ is divisible and so is $v_L L$. □

Remark 3.6.3. In the following, we will construct examples where K has positive characteristic. It would be interesting to also produce examples in mixed characteristic. This essentially would amount to finding a way to construct a valued field (K, v) of mixed characteristic, such that vK is divisible, $\mathcal{O}_v/p$ is semi-perfect, and Kv is some prescribed residue field. We are not aware of such a construction.

Puiseux series. We use the Puiseux series construction to build an example of a field satisfying conditions $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge 5$ of the Main Theorem.

Definition 3.6.4 ([AJ18, Definition 4.4])**.** A field K is of *divisible-tame type* if there exist $L \equiv K$ and a non-trivial valuation w on L such that (L, w) is tame with wL divisible.

Remark 3.6.5. In particular, every divisible-tame type field is perfect.

Fact 3.6.6 ([AJ18, Proposition 4.13])**.** *Let p be a prime or zero. There exists a non-henselian t-henselian field of characteristic p which is*

1. *not separably closed, and*
2. *of divisible-tame type.*

We will use the following folklore fact, whose proof we sketch out as we could not find a reference in the literature.

Lemma 3.6.7. *Let K_0 be a field of characteristic $p > 0$. Consider the Puiseux series over K_0,*

$$K := \bigcup_{n \geqslant 1} K_0((t^{\frac{1}{n}}))$$

together with the restriction v_t of the t-adic valuation from $K_0((t^{\mathbb{Q}}))$. Then, (K, v_t) is a henselian valued field of positive characteristic, v_t has defect, and if K_0 is perfect, then so is K.

Proof. As (K, v_t) is the increasing union of henselian fields, v_t is henselian. However, it is not algebraically maximal. Indeed, the equation $X^p - X - \frac{1}{t} = 0$ has no solution in K, however it admits a solution in $K_0((t^{\mathbb{Q}}))$, namely $a = \sum_{n \geqslant 0} t^{-\frac{1}{p^n}}$. Then, $K \subsetneq K(a) \subseteq K_0((t^{\mathbb{Q}}))$ is a tower of immediate extensions, thus $K \subsetneq K(a)$ is a proper algebraic immediate extension. Note that if K_0 is perfect, then $K_0((t^{\frac{1}{n}}))^{\frac{1}{p}} = K_0((t^{\frac{1}{np}}))$ for every n, and thus the union K is perfect. □

Example 3.6.8. Let K_0 be a non-henselian t-henselian field of characteristic $p > 0$ which is not separably closed, and such that there exist $L \equiv K_0$ and a non-trivial henselian valuation w on L such that wL is divisible (such a K_0 exists, for example, by FACT 3.6.6). Then,

$$K := \bigcup_{n \geqslant 1} K_0((t^{\frac{1}{n}}))$$

satisfies $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge 5$ from our Main Theorem. Indeed, we have $v_K = v_t$ by REMARK 3.6.1. Now conditions $\neg 1$, $\neg 2$ and $\neg 4$ are immediate from the construction, condition $\neg 3$ follows from PROPOSITION 3.6.2 and condition 5 holds because of LEMMA 3.6.7. It then follows that K admits a definable non-trivial henselian valuation.

Condition 6 is necessary. We adapt a construction from [AJ18] to produce non-henselian, t-henselian fields with defect in some elementary extension. We first introduce a series of weakenings of henselianity.

Definition 3.6.9 ([AJ18, Definition 3.3])**.** Let $n \geqslant 1$. Say that a valued field (K, v) is *$n_{\leqslant}$-henselian* if for every $f \in \mathcal{O}_v[X]$ of degree $\leqslant n$, and $a \in \mathcal{O}_v$, if $v(f(a)) > 0$ and $v(f'(a)) = 0$, then there is $b \in \mathcal{O}_v$ with $f(b) = 0$ and $v(b - a) > 0$.

Definition 3.6.10 ([EP05, Section 4.2])**.** Let q be a prime. Say that a valued field (K, v) is *q-henselian* if v extends uniquely to every Galois extension of K of q-power degree.

Remark 3.6.11. A valued field (K, v) is henselian if and only if it is $n_{\leqslant}$-henselian for all n. Being $n_{\leqslant}$-henselian is clearly a first-order property of (K, v); by [Koe95, Propositions 1.2 and 1.3], the same is true for q-henselianity.

Next, we isolate a new notion built along the lines of t-henselianity and divisible-tame type.

Definition 3.6.12. We call a field K *t-henselian of defect type* if there is an elementarily equivalent $L \equiv K$ which admits some henselian valuation v such that (L, v) has defect. We call a field K *t-henselian of divisible-defect type* if there is an elementarily equivalent $L \equiv K$ which admits some henselian valuation v such that (L, v) has defect with vL divisible.

We now replicate a construction that first appeared in [PZ78], and was later refined in [FJ15] and [AJ18]. The proofs are almost verbatim the same as in [FJ15] and [AJ18]; we thus give the appropriate references and explain the differences.

Lemma 3.6.13 ([AJ18, Lemma 4.8])**.** *Let $p > 0$ be a prime. Let K be a perfect field of characteristic p that contains all roots of unity. Let $n > p$ and let q be a prime with $q > n$. Then, there exists an equicharacteristic valued field (K', v) with*

- *$K'v = K$, $vK' = \mathbb{Q}$,*
- *K' is perfect,*
- *(K', v) is not q-henselian, but it is $n_{\leqslant}$-henselian,*
- *(K', v) admits a proper degree p immediate extension.*

Proof. We follow the proof of [AJ18, Lemma 4.8], but we substitute the generalized power series with the Puiseux series. Indeed, inside the Puiseux series $L := \bigcup_{n \geqslant 1} K(\!(t^{\frac{1}{n}})\!)$, endowed with the restriction v_t of the t-adic valuation on $K(\!(t^{\mathbb{Q}})\!)$, consider the subfield $F := K(t^{\nu} \mid \nu \in \mathbb{Q})$. Consider $\widetilde{F} := F^{\mathrm{alg}} \cap L$. Note that since L is henselian, so is $\widetilde{F}$. Then, arguing as in [AJ18, Proof of Lemma 4.8], there is a subgroup $G \leqslant \mathrm{Gal}(F^{\mathrm{sep}}|F)$ with $G \cong \mathbb{Z}_q$.

Let $E = \mathrm{Fix}(G) \subseteq F^{\mathrm{sep}}$, and consider $K' := E \cap \widetilde{F}$. Then, (K', v_t) has residue field K and value group $\mathbb{Q}$, and it is perfect. Moreover, the roots of $X^p - X - \frac{1}{t}$ give rise to immediate extensions of degree p over (K', v_t). Now, arguing as in [AJ18, Claim 4.8.1], (K', v_t) is not q-henselian. Similarly, arguing as in [AJ18, Claim 4.8.2], (K', v_t) is $n_{\leqslant}$-henselian (note that they argue that (K', v_t) is $(n!^2!)_{\leqslant}$-henselian, using that in their case $q > n!^2!$).

As for the last point, note that L admits an immediate extension of degree p, namely given by any root of the Artin-Schreier polynomial $X^p - X - \frac{1}{t}$. In particular, then, $\widetilde{F}$ also admits an immediate extension of degree p, and thus so does K'. □

The following construction follows very closely the proof of [AJ18, Proposition 4.13], using LEMMA 3.6.13 in place of [AJ18, Lemma 4.8] and diverging only in the last paragraph. We sketch out the construction for the convenience of the reader, but invite them to see [AJ18] and [FJ15] for the full details.

Proposition 3.6.14 ([AJ18, Proposition 4.13])**.** *Let $p > 0$ be a prime. There is a non-henselian, t-henselian of divisible-defect type perfect field of characteristic p which is not separably closed.*

Proof. For each $n \geqslant 0$, set $k_n = n + p + 1$ and choose a prime $q_n > k_n$. Let $K_0 = \mathbb{F}_p^{\mathrm{alg}}$. Then, using LEMMA 3.6.13, build a sequence $\{(K_{n+1}, \overline{v_n}) \mid n \geqslant 0\}$ of valued fields such that each $(K_{n+1}, \overline{v_n})$ has value group $\mathbb{Q}$, residue field K_n, is not q_n-henselian, but it is $(k_n)_{\leqslant}$-henselian, and it admits a proper degree p immediate extension.

For any $n > m \geqslant 0$, then, write $v_{n,m} := \overline{v_m} \circ \cdots \circ \overline{v_{n-1}}$ and denote by $\mathcal{O}_{n,m}$ the corresponding valuation ring on K_n. Denote by $\pi_{n,m} \colon \mathcal{O}_{n,0} \to \mathcal{O}_{m,0}$ the restriction of the residue map $\mathcal{O}_{n,m} \to K_m$. Then, the valuation rings $(\mathcal{O}_{n,0})_{n \geqslant 0}$ form a projective system together with the maps $\pi_{n,m}$. The limit $\mathcal{O}$ is again a valuation ring, together with natural projections $\pi_{\infty,n} \colon \mathcal{O} \to \mathcal{O}_{n,0}$. For each $n \geqslant 0$, consider the localization $\mathcal{O}_{v_n} := \mathcal{O}_{\ker(\pi_{\infty,n})}$. Each $\mathcal{O}_{v_n}$ is a valuation ring on $K = \mathrm{Frac}(\mathcal{O}) = \bigcup_{n \geqslant 0} \mathcal{O}_{v_n}$ with residue field K_n and non-trivial divisible value group. Indeed, for each $n \geqslant 0$, $\mathcal{O}_{v_n} \subseteq \mathcal{O}_{v_{n+1}}$, and v_n induces precisely $\overline{v_n}$ on $K_{n+1} = Kv_{n+1}$.

$$\begin{tikzcd} K \arrow[r, "v_{n+1}"'] \arrow[rr, bend left, "v_n"] & Kv_{n+1} = K_{n+1} \arrow[r, "\overline{v_n}"'] & K_n = Kv_n \end{tikzcd}$$

We now let (K^*, v^*) be an ultraproduct of the family $(K, v_n)_{n \geqslant 0}$. Then, K^* is perfect, and v^* is henselian with divisible value group. Moreover, K is not henselian and thus not separably closed.

Now, we diverge from the proof of [AJ18, Proposition 4.13] and argue that (K^*, v^*) admits a degree p immediate extension, in particular a defect extension. It is enough to show that each (K, v_n) admits one such generated by a root of an Artin-Schreier polynomial, since this can be expressed in a first-order sentence. We know that $(K_{n+1}, \overline{v_n})$ admits a degree p immediate extension generated by a root α of an Artin-Schreier polynomial $X^p - X - \xi$, for some $\xi \in K_{n+1}$. Denote by $\overline{u}$ the prolongation of $\overline{v_n}$ to $K_{n+1}(\alpha)$ that makes $(K_{n+1}(\alpha), \overline{u})$ immediate over $(K_{n+1}, \overline{v_n})$. For some $z \in K$ with $\mathrm{res}_{v_{n+1}}(z) = \xi$, we consider the Artin-Schreier polynomial $X^p - X - z$. Let w be a prolongation of v_{n+1} to K^{alg}: then there is $a \in K^{\mathrm{alg}}$ with $a^p - a - z = 0$ and $\mathrm{res}_w(a) = \alpha$, by henselianity of w. Then, if we denote by w again the restriction of w to $K(a)$, we have that $K(a)w = Kw(\alpha) = K_{n+1}(\alpha)$. We now let $u = \overline{u} \circ w$. Then, $K(a)u = K_{n+1}(\alpha)\overline{u} = Kv_n$. Moreover, $v_n K$ is divisible, so $uK(a) = v_n K$. □

Proposition 3.6.15. *Let K be a non-henselian, t-henselian of defect type, perfect field of characteristic p which is not separably closed. Let $L \equiv K$ be such that, for some henselian valuation v, (L, v) has defect. Then v_L has defect.*

Proof. We distinguish two cases. If v is a (possibly non-proper) coarsening of v_L, then we are done, since defect goes up in coarsenings. If v is a proper refinement of v_L, then Lv_L is separably closed, and since L is perfect, it is in particular algebraically closed. Thus, the valuation $\overline{v}$ induced by v on Lv_L is defectless. But then, since $v = \overline{v} \circ v_L$ has defect if and only at least one of v_L and $\overline{v}$ has defect ([AJ24, Lemma 2.9]), v_L must have defect. $\square$

Remark 3.6.16. The construction in PROPOSITION 3.6.14 shows that in equicharacteristic p, one cannot eliminate the parameters from COROLLARY 3.4.15. Indeed, assume that, on any perfect henselian valued field (F, v) of characteristic $p > 0$ such that v has defect, there is a non-trivial $\varnothing$-definable henselian coarsening w of v. Take a non-henselian, t-henselian of defect type, perfect field K (given by PROPOSITION 3.6.14), and let $L \equiv K$ be such that (L, v_L) has defect (which we can assume exists by PROPOSITION 3.6.15). Then, there is a coarsening w of v_L which is $\varnothing$-definable. In particular, using the same formula, we can find a non-trivial henselian valuation on K, a contradiction.

Example 3.6.17. Let $p > 0$ be a prime. Let K_0 be a non-henselian, t-henselian of divisible-defect type, perfect field of characteristic p which is not separably closed. Then,

$$K := K_0((t^{\mathbb{Q}}))$$

satisfies $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge \neg 5 \wedge 6$ from our Main Theorem. Indeed, we have $v_K = v_t$ by REMARK 3.6.1. Now conditions $\neg 1$, $\neg 2$, $\neg 4$, 6 are immediate from the construction, $\neg 5$ follows since $(K_0((t^{\mathbb{Q}})), v_t)$ is tame, and condition $\neg 3$ follows from PROPOSITION 3.6.2. It then follows that K admits a definable non-trivial henselian valuation.

Question 3.6.18. Let K be a non-henselian, t-henselian of divisible-tame type field of characteristic p which is not separably closed. Suppose that $L \equiv K$ admits a non-trivial henselian valuation. Must (L, v_L) be defectless? This is the defectless version of PROPOSITION 3.6.2, and it would yield that if K_0 is non-henselian, t-henselian of divisible-tame type, not separably closed, then $K_0((t^{\mathbb{Q}}))$ is an example for $\neg 1 \wedge \neg 2 \wedge \neg 3 \wedge \neg 4 \wedge \neg 5 \wedge \neg 6$ (and thus admits no non-trivial definable henselian valuation).

IV. Model theory of valued fields with an endomorphism

> Because the sunset, like survival, exists only on the verge of its own disappearing. To be gorgeous, you must first be seen, but to be seen allows you to be hunted.
>
> Ocean Vuong, *On Earth We're Briefly Gorgeous*

The content of this chapter is essentially contained in [Ram24].

4.1 | The road ahead

The chapter is structured as follows:

- section 4.2 introduces the necessary preliminaries in difference algebra and valued difference algebra, in particular focusing on the notion of σ-separability.
- section 4.3 develops the theory of weak σ-henselianity and immediate extensions, proving the crucial theorem 4.3.21 (theorem B).
- section 4.4 takes care of the other major ingredient for an embedding lemma, namely how to extend value (difference) group and residue (difference) field.
- section 4.5 proves the embedding lemma and deduces a wealth of variants and model-theoretic consequences, including relative quantifier elimination (theorem C), AKE-type theorems, and NTP_2 transfer (theorem D).
- section 4.6 shows why assuming that σ_{res} and σ_{val} are surjective seems necessary to have some control of the model-theoretic behaviour of valued difference fields.

4.2 | Preliminaries

Unless otherwise stated, all fields considered will be of characteristic zero (in particular, all residue fields will be of characteristic zero).

Notation. Let K be a field, and let $\sigma : K \to K$ be an endomorphism (in particular, σ is injective). We call the pair (K, σ) a *difference field*, and often omit σ whenever it is clear from context. We say that K (or (K, σ)) is *inversive* if σ is surjective. Write $\mathcal{L}_{\mathrm{ring},\sigma}$ for $\mathcal{L}_{\mathrm{ring}} \cup \{\sigma\}$.

Any difference field admits a minimal extension that is inversive:

Lemma 4.2.1. *Let (K, σ) be a difference field. Then, up to isomorphism over K, there is a unique inversive difference field (L, σ) extending (K, σ) and satisfying the following universal property: given any inversive difference field (N, σ) extending (K, σ), there is a (unique) K-embedding of (L, σ) into (N, σ).*

The proof of this result is straightforward. We denote the unique such (L, σ) by K^{inv} and call it the *inversive closure* of K. Explicitly, K^{inv} is the directed union of isomorphic copies of K, which we write as $K^{\mathrm{inv}} = \bigcup_{n \geqslant 0} \sigma^{-n}(K)$. Up to isomorphism over K, we can write that $a \in K^{\mathrm{inv}}$ if and only if there is $m \geqslant 0$ such that $\sigma^m(a) \in K$. We write K_{inv} for the unique difference field given by $K_{\mathrm{inv}} := \bigcap_{n \geqslant 0} \sigma^n(K)$.

If $(K, \sigma) \subseteq (L, \sigma)$ are difference fields, then for any $\alpha \in L$ we write

$$K(\alpha)_\sigma := \mathrm{Frac}(\langle K, \{\sigma^n(\alpha)\}_{n \geqslant 0} \rangle_{\mathcal{L}_{\mathrm{ring}}}) = \mathrm{Frac}(\langle K, \alpha \rangle_{\mathcal{L}_{\mathrm{ring},\sigma}}) \subseteq L,$$

and, if possible,

$$K(\alpha)_{\sigma, \sigma^{-1}} := \mathrm{Frac}(\langle K, \{\sigma^n(\alpha), \sigma^{-n}(\alpha)\}_{n \geqslant 0} \rangle_{\mathcal{L}_{\mathrm{ring}}}) \subseteq L.$$

If K is a field, we denote by K^{alg} its algebraic closure, and by K^{sep} its separable closure. We reserve $\overline{K}$ for something else (see DEFINITION 4.3.22). If Γ is an ordered group and $\sigma \in \mathrm{End}(\Gamma)$, we call (Γ, σ) an *ordered difference group*.

If (K, v) is a valued field, we write $\mathcal{O}_K = \{x \in K \mid v(x) \geqslant 0\}$ for its valuation ring, $\mathfrak{m}_K = \{x \in K \mid v(x) > 0\}$ for its maximal ideal, k_K for its residue field, Γ_K for its value group, and RV_K for its leading terms structure. We denote by $\widehat{K}$ its completion.

By a *valued difference field*, or valued field with endomorphism, we mean a valued field (K, v) together with an endomorphism of valued fields $\sigma\colon (K, v) \to (K, v)$. Equivalently, a valued field (K, v) together with a field endomorphism $\sigma\colon K \to K$ such that $a \in \mathcal{O}_K$ if and only if $\sigma(a) \in \mathcal{O}_K$. Some examples of valued difference fields are given in PROPOSITION 4.4.14.

The inversive closure can be endowed with a valued difference field structure in a unique way:

Lemma 4.2.2. *Let (K, v, σ) be a valued difference field. Then, up to isomorphism over K, there is a unique valuation v on $(K^{\mathrm{inv}}, \sigma)$ extending (K, v, σ) and satisfying the following universal property: given any inversive valued difference field (N, v, σ) extending (K, v, σ), there is a (unique) K-embedding of $(K^{\mathrm{inv}}, v, \sigma)$ into (N, v, σ).*

Once again, $(K^{\mathrm{inv}}, v, \sigma)$ is constructed as the directed union of isomorphic copies of (K, v, σ). Note that $\mathcal{O}_{K^{\mathrm{inv}}} = \bigcup_{n \geqslant 0} \sigma^{-n}(\mathcal{O}_K)$.

For any such $\sigma \in \mathrm{End}(K, v)$, we denote by σ_{val} the induced endomorphism of Γ_K, by σ_{res} the induced endomorphism of k_K, and by σ_{rv} the induced endomorphism of RV_K. We call $(k_K, \sigma_{\mathrm{res}})$ the *residue difference field* of (K, v, σ), $(\Gamma_K, \sigma_{\mathrm{val}})$ the *value difference group* of (K, v, σ), and $(\mathrm{RV}_K, \sigma_{\mathrm{rv}})$ the *leading terms difference structure* of (K, v, σ).

If $a \in K$ and $\gamma \in \Gamma_K$, we write

$$B_\gamma(a) = \{b \in K \mid v(a-b) > \gamma\}$$

for the *open ball around b of radius* γ, and

$$B_\gamma[a] = \{b \in K \mid v(a-b) \geqslant \gamma\}$$

for the *closed ball around b of radius* γ.

Given a valued difference field (K, v, σ), we can see its value difference group $(\Gamma, \sigma_{\mathrm{val}})$ as a $\mathbb{Z}[\sigma_{\mathrm{val}}]$-module in a natural way. Given any $I = i_0 + i_1\sigma_{\mathrm{val}} + \cdots + i_n\sigma_{\mathrm{val}}^n \in \mathbb{Z}[\sigma_{\mathrm{val}}]$, we write $|I| = i_0 + \cdots + i_n \in \mathbb{Z}$. Now, for any $\gamma \in \Gamma_K$, we let

$$I(\gamma) := i_0 \cdot \gamma + i_1 \cdot \sigma_{\mathrm{val}}(\gamma) + i_2 \cdot \sigma_{\mathrm{val}}^2(\gamma) + \cdots + i_n \cdot \sigma_{\mathrm{val}}^n(\gamma) \in \Gamma_K.$$

Assumption 4.2.3. Unless otherwise stated, whenever we write (K, v, σ) we assume that both σ_{val} and σ_{res} (or, equivalently, σ_{rv}) are surjective. Under this assumption, the extension $\sigma(K) \subseteq K$ is immediate. By induction, one then obtains that $K \subseteq K^{\mathrm{inv}}$ is also an immediate extension. Results that do not require this assumption are flagged by $\diamondsuit$.

Difference polynomials. Let (K, σ) be a difference field and fix countably many distinct formal indeterminates $(\sigma^n(X))_{n \geqslant 0}$. We say that $p(X)$ is a *difference polynomial over K* if, for some $n \geqslant 0$, there is a multivariable polynomial $P(X_0, \ldots X_n)$ over K such that $p(X) = P(X, \sigma(X), \ldots \sigma^n(X))$. We write

$$p(X) = \sum_I a_I X^I,$$

where for every $I = (i_0, \ldots i_n) \in \mathbb{N}^{n+1}$, $X^I = X^{i_0}\sigma(X)^{i_1} \cdots \sigma^n(X)^{i_n}$.

If (K, σ) is a difference field, then by $K[X]_\sigma$ we denote the ring of difference polynomials with coefficients from K. If $p(X) = \sum_I a_I X^I$ is a difference polynomial with coefficients in $\mathcal{O}_K$, we write $\mathrm{res}(p)(X)$ for the difference polynomial $\sum_I \mathrm{res}(a_I) X^I$ over k_K. Given any $p(X) = \sum_I \alpha_I X^I \in k_K[X]_\sigma$, an *exact lift* of $p(X)$ is a difference polynomial $q(X) = \sum_I a_I X^I \in \mathcal{O}_K[X]_\sigma$ such that $\mathrm{res}(a_I) = \alpha_I$, and if $\alpha_I = 0$ then $a_I = 0$.

Let $p(X)$ be a non-constant difference polynomial, say $p(X) = P(X, \sigma(X), \ldots \sigma^n(X))$, where X_n appears in P. Then, the *complexity* $I \in \mathbb{N}^3$ of $p(X)$ is given by

$$I = (n, \deg_{X_n}(P), \mathrm{totdeg}(P)),$$

where $\mathrm{totdeg}(P)$ is the total degree of P. If $p(X)$ is constant and non-zero, we say that it has complexity $(-\infty, 0, 0)$. The zero difference polynomial has complexity $(-\infty, -\infty, -\infty)$. We order complexities lexicographically, and declare that

$$(-\infty, -\infty, -\infty) < (-\infty, 0, 0) < \mathbb{N}^3.$$

Note that for any multivariable polynomial $P(X_0, \dots X_n)$ over K, there are unique polynomials $P_J(X_0, \dots X_n)$ such that, computing in $K[X_0, \dots X_n, Y_0, \dots Y_n]$,

$$P(\overline{X} + \overline{Y}) = \sum_{J \in \mathbb{N}^{n+1}} P_J(\overline{X}) Y_0^{j_0} \cdots Y_n^{j_n}.$$

If $I = (i_0, \dots i_n) \in \mathbb{N}^{n+1}$, we let $I! := i_0! \cdots i_n!$. Then, $I! P_I(X_0, \dots X_n) = \frac{\partial^I P}{\partial X^I}(X_0, \dots X_n)$. We define derivatives of difference polynomials by $p_J(X) := \frac{\partial^J P}{\partial X^J}(X, \sigma(X), \cdots \sigma^n(X))$.

For any $j = 0, \dots n$, if $E_j = (\delta_{i,j})_{i=0}^n$, we write $p_j(X) = p_{E_j}(X)$. We often write p' for p_0. Note that $p' = 0$ precisely when the variable X does not appear in the difference polynomial, a fact that will guide us throughout this chapter.

Shifting and changing variables. We devote a bit of space to the tedious operation of rigorously shifting difference polynomials back and forth using σ, so that we may keep track of how their derivatives change for later use.

First, we define the shift on coefficients. Computing in $K^{\mathrm{inv}}[X]_\sigma$, if $\ell \in \mathbb{Z}$, we define $\sigma^\ell(p)(X) := \sum_I \sigma^\ell(a_I) X^I$; then one has that, for any $\ell \in \mathbb{Z}$ and $I \in \mathbb{N}^{n+1}$, $(\sigma^\ell(p))_I(X) = \sigma^\ell(p_I)(X)$. Then, we define a change of variables that allows us to transform a difference polynomial p into another difference polynomial q such that $q' \neq 0$. Namely, if $\sigma^m(X)$ is the smallest iterate of σ appearing in p, we operate a change of variables $Y := \sigma^m(X)$.

Definition 4.2.4. For any non-constant $p(X)$ over K, let $m = m(p) \geqslant 0$ be the least such that there is $I = (i_0, \dots i_n)$ with $i_m \neq 0$ and $p_I \neq 0$. Equivalently, $m(p)$ is the smallest such that $p_m \neq 0$.[1] Then, we let $\mathfrak{S}(p)$ be the difference polynomial over K defined by $\mathfrak{S}(p)(Y) := p(\sigma^{-m(p)}(Y))$.

In a slight abuse of notation, if $I \in \mathbb{N}^{n+1}$, then by $\mathfrak{S}(p_I)(Y)$ we mean again $p_I(\sigma^{-m}(Y))$ (and not, as one might imagine, $p_I(\sigma^{-m(p_I)}(Y))$). We note that $\sigma^{-m}(\mathfrak{S}(p)) = \mathfrak{S}(\sigma^{-m}(p))$. We now compute the derivatives of $\mathfrak{S}(p)$ in terms of the derivatives of p. This is the key step in the computations necessary to prove LEMMA 4.3.6. We first define an operation on multi-indices.

Definition 4.2.5. Given any $I = (i_0, \cdots i_n) \in \mathbb{N}^{n+1}$, we let $I^{-m} = (i_m, \cdots i_n) \in \mathbb{N}^{n-m+1}$. Viceversa, for every $J = (j_0, \cdots j_{n-m}) \in \mathbb{N}^{n-m+1}$, let $J^{+m} = (0, \cdots 0, j_0, \cdots j_{n-m}) \in \mathbb{N}^{n+1}$.

For example, for every $m \leqslant j \leqslant n$, $E_j^{-m} = E_{j-m}$ and $E_j^{+m} = E_{j+m}$. Moreover, one has that $(I + J)^{\pm m} = I^{\pm m} + J^{\pm m}$, where $+$ is the pointwise sum on $\mathbb{N}^{n+1}$. By comparing Taylor expansions, we then obtain that, for $m = m(p)$ as above,

$$\mathfrak{S}(p)_J(Y) = \mathfrak{S}(p_{J^{+m}})(Y).$$

[1] In other words, m is the smallest number of iterations of $\sigma(X)$ that appears in p.

Given any $a \in K$, we can thus compute (possibly inside K^{inv}), for every $I \in \mathbb{N}^{n+1}$,

$$\sigma^{-m}(p_I(a)) = \sigma^{-m}(p_I)(\sigma^{-m}(a)) = \sigma^{-m}(\mathfrak{S}(p_I))(a) = \sigma^{-m}(\mathfrak{S}(p))_{I^{-m}}(a),$$

and viceversa $p_I(a) = \sigma^m(\sigma^{-m}(\mathfrak{S}(p))_{I^{-m}}(a))$. In particular, for $m \leqslant j \leqslant n$, $p_j(a) = \sigma^m(\sigma^{-m}(\mathfrak{S}(p))_{j-m}(a))$.

Remark 4.2.6. Note that for any $b \in K$ one has

$$p(b) = 0 \iff \sigma^{-m}(p(b)) = 0 \iff \sigma^{-m}(p)(\sigma^{-m}(b)) = 0 \iff \sigma^{-m}(\mathfrak{S}(p))(b) = 0,$$

in other words p and $\sigma^{-m}(\mathfrak{S}(p))$ have the same zeroes in K, with the advantage that now $\mathfrak{S}(p)$ and $\sigma^{-m}(\mathfrak{S}(p))$ satisfy that $(\mathfrak{S}(p))' \neq 0$ and $(\sigma^{-m}(\mathfrak{S}(p)))' \neq 0$.

σ-separability. Drawing heavily from [DH23, Section 4], which in turn builds on [CH04], we establish the corresponding notion to *separability* for difference field extensions. One should think of a non-inversive difference field as an imperfect positive characteristic field, and indeed most of the basic facts from separability transfer to our setting, as long as we assume that we work with a model of FE (see DEFINITION 4.2.12).

Definition 4.2.7. Let $(K, \sigma) \subseteq (L, \sigma)$ be a difference field extension. We say that the extension is *σ-separable* if K is linearly disjoint from $\sigma(L)$ over $\sigma(K)$. We say that the extension is *almost σ-separable* if K is algebraically free from $\sigma(L)$ over $\sigma(K)$. We say that the extension is *σ-separably σ-algebraic* if it is σ-separable and σ-algebraic[2].

Definition 4.2.8. We say that $(K, \sigma) \subseteq (L, \sigma)$ is *purely σ-inseparable* if, for every $a \in L$, there is $m \geqslant 0$ such that $\sigma^m(a) \in K$. In particular, any such extension can (up to isomorphism over K) be realized as a subextension of $K \subseteq K^{\mathrm{inv}}$.

Remark 4.2.9. With an induction argument, one sees that K is linearly disjoint from $\sigma(L)$ over $\sigma(K)$ if and only if K^{inv} is linearly disjoint from L over K.

Fact 4.2.10 ([DH23, Propositions 4.5, 4.29, 4.30, 4.31])**.** *Let $(K, \sigma) \subseteq (M, \sigma) \subseteq (L, \sigma)$ be a tower of difference fields.*

1. *If $(K, \sigma) \subseteq (M, \sigma)$ and $(M, \sigma) \subseteq (L, \sigma)$ are σ-separable (resp. almost σ-separable), then $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable (resp. almost σ-separable),*

2. *If $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable (resp. almost σ-separable), then $(K, \sigma) \subseteq (M, \sigma)$ is σ-separable (resp. almost σ-separable).*

3. *Let K_0 be a difference field extension of K, linearly disjoint from L over K. Then, $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable (resp. almost σ-separable) if and only if $(K_0, \sigma) \subseteq (L \otimes_K K_0, \sigma)$ is σ-separable (resp. almost σ-separable).*

Remark 4.2.11. Writing $L \otimes_K K_0$ is potentially ambiguous. Here, we follow [DH23, Remark 3.4] and use it to mean the difference field tensor product (and not, as usual, the ring-theoretic tensor product).

[2] Recall that $(K, \sigma) \subseteq (L, \sigma)$ is σ-algebraic if for every $a \in L$ there is $p(X) \in K[X]_\sigma$ such that $p(a) = 0$.

The theory FE**.** Here, we diverge from the intuition from the positive characteristic world: σ-separability does not behave well unless we assume that K is algebraically closed in K^{inv}. This is, at least partially, because we wish to be able to check σ-separability of a σ-algebraic extension by computing the first derivatives of the difference polynomials that witness algebraicity; however, ordinary minimal polynomials happen to also be difference polynomials, and since we work in characteristic zero, they are always separable, even if they are the minimal polynomial of an element of the inversive closure. For this reason, we always work in models of what [CH04] call T_σ. Let $\mathcal{L}_{\mathrm{ring},\sigma} := \{+, \cdot, 0, 1, -, \sigma\}$.

Definition 4.2.12. We let FE be the $\mathcal{L}_{\mathrm{ring},\sigma}$-theory of characteristic zero difference fields (K, σ) where $\sigma(K)$ is algebraically closed in K. Equivalently, of difference fields (K, σ) with K algebraically closed in K^{inv}.

We will often also say that (K, σ) *satisfies* FE, or that a valued difference field (K, v, σ) is *a model of* FE, to mean that $(K, \sigma) \models \mathrm{FE}$[3].

Remark 4.2.13. Any inversive difference field is automatically a model of FE. On the non-inversive side, any model of SCFE as defined in [CH04] is also a model of FE. If (K, v, σ) is a valued difference field which is henselian (as a valued field), then (K, σ) is automatically a model of FE (since $K \subseteq K^{\mathrm{inv}}$ is an immediate extension).

Fact 4.2.14 ([DH23, Propositions 4.5, 4.29, 4.30, 4.31])**.** *Let $(K, \sigma) \subseteq (M, \sigma) \subseteq (L, \sigma)$ be a tower of difference fields. Assume that (K, σ) is a model of* FE*. Then:*

1. *$(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only it is almost σ-separable,*
2. *if $L = K(\alpha)_\sigma$ is σ-algebraic over K, then $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only if there is $p(X) \in K[X]_\sigma$ with $p(\alpha) = 0$ and $p'(\alpha) \neq 0$*[4]*,*
3. *if $(K, \sigma) \subseteq (L, \sigma)$ is σ-algebraic, then it is σ-separable if and only if for every $\alpha \in L$ there exists $p(X) \in K[X]_\sigma$ with $p(\alpha) = 0$ and $p'(\alpha) \neq 0$,*
4. *if $(K, \sigma) \subseteq (M, \sigma)$ is σ-algebraic and (M, σ) also is a model of* FE*, then $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only if both $(K, \sigma) \subseteq (M, \sigma)$ and $(M, \sigma) \subseteq (L, \sigma)$ are.*

When a (valued) difference field does not satisfy FE, one can canonically move to a minimal extension that satisfies it, which we call the FE-closure.

Lemma 4.2.15. *Let (K, σ) be a difference field. Then (K, σ) admits an extension* $\mathrm{FE}(K)$ *that is a model of* FE *and with the following universal mapping property: for any extension (L, σ) of (K, σ) which is a model of* FE*, there is a unique embedding $f : \mathrm{FE}(K) \to L$ over K. If such an L exists that is σ-separable over K, then $K = \mathrm{FE}(K)$.*

[3]The acronym FE simply stands for *field with endomorphism*.
[4]We call such an α a *simple zero* of $p(X)$.

Proof. We let $\mathrm{FE}(K) := K^{\mathrm{alg}} \cap K^{\mathrm{inv}}$. Then, $\mathrm{FE}(K) \subseteq K^{\mathrm{inv}}$ gives a unique difference field structure on $\mathrm{FE}(K)$. Furthermore, since L is algebraically closed in L^{inv}, one gets a unique embedding of $\mathrm{FE}(K)$ into L over K. $\square$

Remark 4.2.16 (♢)**.** If (K, v, σ) is a valued difference field, then there is a unique valued field structure on $\mathrm{FE}(K)$, which then also satisfies the universal mapping property in the valued difference field category.

Remark 4.2.17. Note that $K \subseteq \mathrm{FE}(K)$ is an immediate extension.

Definition 4.2.18. We call $\mathrm{FE}(K)$ the FE-*closure* of K.

When building a back-and-forth system, it will be crucial that when we move to the FE-closure inside a σ-separable extension, this σ-separability carries over.

Lemma 4.2.19. *Suppose $(K, \sigma) \subseteq (L, \sigma)$ is a difference field extension, and (K, σ) is a model of* FE*. Let $(L, \sigma) \subseteq (L', \sigma)$ be a difference field extension, with $L \subseteq L'$ algebraic. Then $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only if $(K, \sigma) \subseteq (L', \sigma)$ is σ-separable. In particular, $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only if $(K, \sigma) \subseteq (\mathrm{FE}(L), \sigma)$ is σ-separable.*

Proof. As K is a model of FE, $(K, \sigma) \subseteq (L, \sigma)$ is σ-separable if and only if it is almost σ-separable, by FACT 4.2.14(1). The statement then follows from the fact that $L \subseteq L'$ is an algebraic extension. $\square$

Lemma 4.2.20. *Let $(K, \sigma) \subseteq (L, \sigma)$ be an extension of difference fields, both models of* FE*. Let $\tau \in L_{\mathrm{inv}}$ be σ-transcendental over K. Then $K(\tau)_{\sigma,\sigma^{-1}} \subseteq L$ is a σ-separable extension. In particular, $K(\tau)_{\sigma,\sigma^{-1}}$ is a model of* FE*.*

Proof. Let $K_0 := K(\tau)_{\sigma,\sigma^{-1}}$. Note that $K \subseteq K_0$ is a σ-separable extension. Then, by part (3) of FACT 4.2.10, $K_0 \subseteq L$ is σ-separable if and only if $K_0 \otimes_K K^{\mathrm{inv}} \subseteq L \otimes_K K^{\mathrm{inv}}$ is σ-separable. But this is trivially satisfied, since $K_0 \otimes_K K^{\mathrm{inv}}$ is inversive.
We then obtain that $K(\tau)_{\sigma,\sigma^{-1}}$ is a model of FE by LEMMA 4.2.15. $\square$

The σ-separably σ-algebraic closure. Let $(K, \sigma) \subseteq (L, \sigma)$. We say that (K, σ) is *σ-separably σ-algebraically closed* in (L, σ) if there is no proper intermediate $K \subset M \subseteq L$ with $(K, \sigma) \subset (M, \sigma)$ σ-separably σ-algebraic. If (K, σ) is a model of FE, this is equivalent to: for all $p(X) \in K[X]_\sigma$ with $p' \neq 0$, if $b \in L$ is such that $p(b) = 0$, then $b \in K$. Note that, in contrast with the case for fields, this is strictly stronger than simply requiring that whenever $p(X) \in K[X]_\sigma$ with $p' \neq 0$ has a solution in L, it also has a solution in K.

Example 4.2.21. (K, σ) is σ-separably σ-algebraically closed in $(K^{\mathrm{inv}}, \sigma)$.

Fact 4.2.22 ([DH23, Theorem 4.46(1)-(2)])**.** *Let (K, v, σ) be a model of* FE*. Let (L, σ) be an extension which is a model of* FE*. Let $\widetilde{K} := \{a \in L \mid \exists f(X) \in K[X]_\sigma (f(a) = 0 \wedge f'(a) \neq 0)\}$. Then,*

1. *$\widetilde{K}$ is σ-separably σ-algebraically closed in L,*

2. *$\widetilde{K}$ is σ-separably σ-algebraic over K and satisfies* FE,

3. *if $(K,\sigma) \subseteq (K',\sigma)$ is a σ-separably σ-algebraic extension, then every embedding of K' in L over K has image in $\widetilde{K}$.*

We call $\widetilde{K}$ as in FACT 4.2.22 the *σ-separably σ-algebraic closure* of K in L.

σ-separable generation. Drawing parallels with MacLane's structure theorem of finitely generated separable extensions, we recall the crucial results in [DH23] about finitely generated σ-separable extensions.

Definition 4.2.23. Let $(K,\sigma) \subseteq (L,\sigma)$ be an extension of difference fields. We say that $\mathcal{X} \subseteq L$ is a *σ-transcendence basis* of L over K if $\mathcal{X}$ is σ-algebraically independent over K and $K(\mathcal{X})_\sigma \subseteq L$ is σ-algebraic.

Definition 4.2.24 ([DH23, Definition 4.40])**.** Let $(K,\sigma) \subseteq (L,\sigma)$ be an extension of models of FE. We say that L is *σ-separably generated* over K if there is a σ-transcendence basis $\mathcal{X} \subseteq L$ of L over K such that $K(\mathcal{X})_\sigma \subseteq L$ is σ-separably σ-algebraic. If such $\mathcal{X}$ can be chosen finite, then we say that L is *finitely σ-separably generated* over K.

Remark 4.2.25. Suppose L is σ-separably generated over K (thus in particular the extension is σ-separable), and let $\mathcal{X}$ be a σ-transcendence basis witnessing that. Let $x \in \mathcal{X}$: then, $K(x)_\sigma \subseteq L$ is σ-separable. Since the composition of σ-separable extensions is σ-separable, it is enough to check that $K(x)_\sigma \subseteq K(\mathcal{X})_\sigma$ is σ-separable. Since x is transcendental over $K \otimes_{\sigma(K)} \sigma(K(\mathcal{X})_\sigma)$, this is true by [DH23, Proposition 4.50(2)].

Fact 4.2.26 ([DH23, Theorem 4.46(5)])**.** *Let $(K,\sigma) \subseteq (L,\sigma)$ be a σ-separable extension of models of* FE*. Suppose that L is finitely generated as a model of* FE *over K, meaning that there is a finite $\mathcal{X} \subseteq L$ such that $L =$ FE$(K(\mathcal{X})_\sigma)$. Then, L is finitely σ-separably generated over K.*

4.3 | σ-HENSELIANITY AND IMMEDIATE EXTENSIONS

This section constitutes the bulk of the tools that we will need for our relative quantifier elimination. The crucial ingredient for an embedding lemma is some notion of henselianity, i.e. a recipe for producing solutions to (difference) polynomial equations. We adapt the notion defined in [DO15]: since we deal with a possibly non-surjective σ, we have to rule out the possibility of solving the equation $\sigma(X) = a$, for $a \in K$. This is achieved by the definition of *weakly* σ-henselian.

σ-henselianity. In the pure valued field world, henselianity can be phrased in many ways. Crucially it can be seen as a first-order shadow of completeness, meaning that approximate roots of polynomials give rise to proper roots, granted that they are residually simple. In the valued difference world, since we allow all kinds of behaviour of σ_{val}, it is impossible to predict the behaviour of the values of $a, \sigma(a), \cdots \sigma^n(a)$, as

$v(a)$ grows larger; it is thus not enough to check that certain derivatives do not vanish residually, but one has to consider all of them at the same time. This leads to the definition introduced in [DO15] which is, at face value, more artificial and harder to parse than the one introduced in [BMS07] or [DH23] (where the behaviour of σ_{val} is fixed).

Definition 4.3.1. Let $p \in K[X]_\sigma$ be a non-constant difference polynomial, and $a \in K$. We say that p is in *σ-henselian configuration* at a if there are $\gamma \in \Gamma_K$ and $0 \leqslant i \leqslant n$ such that

1. $v(p(a)) = v(p_i(a)) + \sigma^i_{\mathrm{val}}(\gamma) \leqslant v(p_j(a)) + \sigma^j_{\mathrm{val}}(\gamma)$ for all $0 \leqslant j \leqslant n$,
2. $v(p_J(a)) + J(\gamma) < v(p_{J+L}(a)) + (J+L)(\gamma)$ whenever $J, L \neq 0$ and $p_J \neq 0$.

The γ as above is unique, and we denote it by $\gamma(p,a)$.

Remark 4.3.2. An instructive case to consider is the one where $p(X)$ is a non-constant difference polynomial, and a is such that $p(a) \neq 0$, $v(p(a)) > 0$, and for every I such that $p_I \neq 0$, $v(p_I(a)) = 0$. Then, p is in σ-henselian configuration at a.

[DO15] call the following definition *σ-Hensel scheme*.

Definition 4.3.3. We will say that (K, v, σ) is *strongly σ-henselian* if whenever p is in σ-henselian configuration at a, then there is $b \in K$ with $p(b) = 0$ and $v(b-a) = \gamma(p,a)$.

As the name suggests, strong σ-henselianity is too strong. In particular, one risks solving equations of the form $\sigma(X) = a$, implying that σ is surjective. We thus weaken it for our purposes.

Definition 4.3.4. We will say that (K, v, σ) is *weakly σ-henselian* if whenever p is in σ-henselian configuration at a, and $p' \neq 0$, then there is $b \in K$ with $p(b) = 0$ and $v(b-a) = \gamma(p,a)$.

Remark 4.3.5. Note that, since polynomials are in particular difference polynomials with non-vanishing first derivative (in the difference sense), then a weakly σ-henselian valued difference field is in particular henselian as a valued field.

We now explore for a moment the relationship between strong and weak σ-henselianity. By direct computation, one can prove the following lemma.

Lemma 4.3.6. *Let $n \in \mathbb{Z}$, and assume that p is in σ-henselian configuration at a. Then, $\sigma^n(p)$ is in σ-henselian configuration at $\sigma^n(a)$. Moreover, if m is the least such that $p_m \neq 0$ (see* SECTION 4.2*), then $\sigma^{-m}(\mathfrak{S}(p))$ is in σ-henselian configuration at a.*

Proof. We first claim that $q := \sigma^n(p)$ and $b := \sigma^n(a)$ are in σ-henselian configuration. Indeed, note that for every I, $v(q_I(b)) = \sigma^n_{\mathrm{val}}(v(p_I(a)))$, so if $\gamma = \gamma(p,a)$ and i are as in the definition of σ-henselian configuration, if we let $\tilde{\gamma} := \sigma^n_{\mathrm{val}}(\gamma)$, then

1. $v(q(b)) = \sigma^n_{\mathrm{val}}(v(p(a))) = \sigma^n_{\mathrm{val}}(v(p_i(a)) + \sigma^i_{\mathrm{val}}(\gamma)) = v(q_i(b)) + \sigma^i_{\mathrm{val}}(\tilde{\gamma}) \leqslant \sigma^n_{\mathrm{val}}(v(p_j(a)) + \sigma^j_{\mathrm{val}}(\gamma)) = v(q_j(b)) + \sigma^j_{\mathrm{val}}(\tilde{\gamma})$, for all $0 \leqslant j \leqslant n$,

2. $v(q_J(b)) + J(\tilde{\gamma}) = \sigma^n_{\mathrm{val}}(v(p_J(a))) + \sigma^n_{\mathrm{val}}(J(\gamma)) = \sigma^n_{\mathrm{val}}(v(p_J(a)) + J(\gamma)) < \sigma^n_{\mathrm{val}}(v(p_{J+L}(a)) + (J+L)(\gamma)) = v(q_{J+L}(b)) + (J+L)(\tilde{\gamma})$, whenever $J, L \neq 0$ and $p_J \neq 0$.

Next, we fix $n := -m$ and claim that $r := \mathfrak{S}(q)$ is in σ-henselian configuration at a (with $\gamma(p,a) = \gamma(r,a)$). Indeed, $r(a) = q(b)$, and more generally $r_I(a) = \sigma^{-m}(p_{I^{+m}}(a))$. Thus,[5]

1. $v(r(a)) = \sigma^{-m}_{\mathrm{val}}(v(p(a))) = \sigma^{-m}_{\mathrm{val}}(v(p_i(a))) + \sigma^{i-m}_{\mathrm{val}}(\gamma) = \sigma^{-m}_{\mathrm{val}}(v(\sigma^m(r_{i-m})(a))) + \sigma^{i-m}_{\mathrm{val}}(\gamma) = v(r_{i-m}(a)) + \sigma^{i-m}_{\mathrm{val}}(\gamma) \leqslant \sigma^{-m}_{\mathrm{val}}(v(p_j(a))) + \sigma^{j-m}_{\mathrm{val}}(\gamma) = v(r_{j-m}(a)) + \sigma^{j-m}_{\mathrm{val}}(\gamma)$, as j ranges between m and n,

2. for all $L \in \mathbb{N}^{n+1}$ of the form $L = I^{+m}$ for some $I \in \mathbb{N}^{n-m+1}$, we can compute $v(r_J(a)) + J(\gamma) = \sigma^{-m}_{\mathrm{val}}(p_{J^{+m}}(a) + J^{+m}(\gamma)) \leqslant \sigma^{-m}_{\mathrm{val}}(v(p_{J^{+m}+L}(a)) + (J^{+m} + L)(\gamma)) = v(r_{J+I}(a)) + (J+I)(\gamma)$.

This shows that $r = \sigma^{-m}(\mathfrak{S}(p))$ is in σ-henselian configuration at a. □

We can now use this to move between a valued difference field and its inversive closure.

Lemma 4.3.7. *Suppose (K, v, σ) is weakly σ-henselian and inversive. Then (K, v, σ) is strongly σ-henselian.*

Proof. Let p be in σ-henselian configuration at a. If $m = m(p)$, then let $q = \sigma^{-m}(\mathfrak{S}(p))$: by LEMMA 4.3.6, q is in σ-henselian configuration at a, and further $q' \neq 0$. As σ is surjective, q is still a difference polynomial over K, so by weak σ-henselianity we find $b \in K$ such that $q(b) = \sigma^{-m}(p)(\sigma^{-m}(b)) = \sigma^{-m}(p(b)) = 0$, and thus $p(b) = 0$. Moreover, $v(b-a) = \gamma(q,a) = \gamma(p,a)$. □

Lemma 4.3.8. *Let (K, v, σ) be a valued difference field that is a model of* FE. *Then, (K, v, σ) is weakly σ-henselian if and only if K^{inv} is strongly σ-henselian.*

Proof. Assume that K is weakly σ-henselian. Note that since being weakly σ-henselian forms an $\forall\exists$-theory, and K^{inv} is the union of copies of K (and hence models of this theory), K^{inv} is also weakly σ-henselian. Then, LEMMA 4.3.7 implies that K^{inv} is strongly σ-henselian.

Conversely, assume that K^{inv} is strongly σ-henselian. If $p \in K[X]_\sigma$ is in σ-henselian configuration at $a \in K$, and $p' \neq 0$, then we can find $b \in K^{\mathrm{inv}}$ such that $p(b) = 0$ and $v(b-a) = \gamma(p,a)$. However, $K \subseteq K^{\mathrm{inv}}$ is σ-separably σ-algebraically closed (EXAMPLE 4.2.21), thus $b \in K$ already. □

[5] By definition of m, whenever $i < m$ we have $p_i = 0$, thus we only need to check the first part of σ-henselian configuration for $i, j \geqslant m$. Analogously, we can check the second condition only for the Ls of the form I^{+m} for some I, since otherwise (by definition of m, again) $p_{J^{+m}+L} = 0$.

The authors of [DO15] consider valued difference fields which are (strongly) σ-henselian and furthermore have linearly difference closed residue difference fields. However, as observed in [Azg10], this already follows from (strong) σ-henselianity.

Lemma 4.3.9. *Let (K, v, σ) be weakly σ-henselian. Then $(k_K, \sigma_{\mathrm{res}})$ is* linearly difference closed, *i.e. if for every $\alpha_0, \ldots \alpha_n \in k_K$, at least one of them not zero, there is $z \in k_K$ such that*

$$1 + \alpha_0 z + \alpha_1 \sigma_{\mathrm{res}}(z) + \cdots + \alpha_n \sigma^n_{\mathrm{res}}(z) = 0.$$

Proof. Since $K \subseteq K^{\mathrm{inv}}$ is an immediate extension, it is enough to show that K^{inv} has linearly difference closed residue difference field. We let $\alpha_i = \mathrm{res}(a_i)$, for $i = 0, \ldots n$, and assume $S = \{j \leqslant n \mid \alpha_j \neq 0\} \neq \varnothing$. Consider the difference polynomial $p(X) := 1 + \sum_{i \in S} a_i X^i$. By construction, $p_i(X) = 0$ if and only if $i \notin S$, i.e. if and only if $\alpha_i = 0$, and $p_J = 0$ for any J with $|J| > 1$. We check that p is in σ-henselian configuration at 0: indeed, using $\gamma = 0$,

$$v(p(0)) = v(1) = 0 = v(p_j(0)) = v(a_j) = 0$$

for all $j \in S$. By LEMMA 4.3.7, K^{inv} is strongly σ-henselian, and thus we find $b \in K^{\mathrm{inv}}$ such that $p(b) = 0$ and $v(b) = \gamma(p, 0) = 0$. Then, $\beta = \mathrm{res}(b)$ is the required solution of the equation in k_K. □

σ-ramification. We establish that a weakly σ-henselian valued difference field is dense in its inversive closure, something similar to how a separably closed valued field is dense in its perfect hull. Introduced in [DH23], we call this phenomenon *deep σ-ramification*, and deduce that in $\aleph_0$-saturated models, the mirrored situation occurs where the inversive core is dense. Once again, this is reflected in $\aleph_0$-saturated separably closed valued fields, which contain a dense algebraically closed valued field, namely their perfect core.

Definition 4.3.10. We say that (K, v, σ) is *deeply σ-ramified* if $\sigma(K) \subseteq K$ is dense (equivalently, if $K \subseteq K^{\mathrm{inv}}$ is dense).

Lemma 4.3.11. *Suppose (K, v, σ) is weakly σ-henselian. Then for any $a \in \mathcal{O}_K$ and $\epsilon \in \mathfrak{m}_K \setminus \{0\}$, there is $b \in \mathcal{O}_K$ such that $\sigma(b) - \epsilon b - a = 0$.*

Proof. Let $p(X) := \sigma(X) - \epsilon X - a$. Note that $p' \neq 0$. Moreover, let $\beta \in k_K$ be such that $\sigma_{\mathrm{res}}(\beta) \neq \mathrm{res}(a)$, and let $b \in \mathcal{O}_K$ be a lift of β. We argue that p is in σ-henselian configuration at b: take $\gamma = 0$, then

1. $v(p(b)) = 0 = v(p_1(b)) = v(1) < v(p_0(b)) = v(\epsilon)$,
2. for all $J, L \neq 0$, $p_{J+L} = 0$, thus the second part is trivially satisfied.

By weak σ-henselianity, there is $b' \in K$ with $p(b') = 0$ and $v(b' - b) = 0$. In particular, $v(b') = v(b' + b - b) \geqslant \min\{0, v(b)\} \geqslant 0$. □

Corollary 4.3.12. *Suppose (K, v, σ) is weakly σ-henselian. Then, K is deeply σ-ramified.*

Proof. Given any $a \in K$ and $\gamma > 0$, we need to find $b \in K$ such that $\sigma(b) \in B_\gamma(a)$.

First, assume that $v(a) \geqslant 0$. We let $\epsilon \in K$ be such that $v(\epsilon) > \max(v(a), \gamma)$, and we let $b \in \mathcal{O}_K$ be such that $\sigma(b) - \epsilon b - a = 0$. Then, $v(\sigma(b) - a) = v(\epsilon b) \geqslant v(\epsilon) > \gamma$, so $\sigma(b) \in B_\gamma(a)$, as required.

If $v(a) < 0$, we rescale $a' = a\sigma(e)$ and $\gamma' = \gamma + \sigma_{\text{val}}(v(e))$, where $v(\sigma(e)) \geqslant -v(a) > 0$. We are then back to the first case: we obtain $b' \in \mathcal{O}_K$ with $\sigma(b') \in B_{\gamma'}(a')$, and then $b := b'/e$ satisfies $b \in B_\gamma(a)$. □

Lemma 4.3.13. *Let (K, v, σ) be weakly σ-henselian and $\aleph_0$-saturated. Then, K_{inv} is dense in K and strongly σ-henselian.*

Proof. By $\aleph_0$-saturation and the fact that $\sigma(K) \subseteq K$ is dense (by COROLLARY 4.3.12), one gets that $K_{\text{inv}} \subseteq K$ is dense (and thus, in particular satisfies ASSUMPTION 4.2.3). Furthermore, as σ is surjective on K_{inv}, it is enough by LEMMA 4.3.7 to show that K_{inv} is weakly σ-henselian. Note that σ^n gives an isomorphism $K \cong \sigma^n(K)$, thus each $\sigma^n(K)$ is weakly σ-henselian. Now, if $p \in K_{\text{inv}}[X]_\sigma$ is in σ-henselian configuration at $a \in K_{\text{inv}}$, then for every $n \geqslant 0$ we have $p \in \sigma^n(K)[X]_\sigma$ and $a \in \sigma^n(K)$, and thus we can find some $b_n \in \sigma^n(K)$ such that $p(b_n) = 0$ and $v(b_n - a) = \gamma(p, a)$. By saturation, then, we find $b \in K_{\text{inv}}$ such that $p(b) = 0$ and $v(b - a) = \gamma(p, a)$, as required. □

Immediate extensions. The next necessary tool to establish relative quantifier elimination is a reasonable theory of immediate extensions, which usually goes under the umbrella of *Kaplansky theory*.

Definition 4.3.14. We say that (K, v, σ) is *(σ-separably) σ-algebraically maximal* if it has no proper immediate (σ-separably) σ-algebraic extension. We say that (K, v, σ) is *(σ-separably) maximal* if it has no proper immediate (σ-separable) extension.

Example 4.3.15. If (K, v, σ) is a model of $\widetilde{\mathrm{VFE}}$ in the sense of [DH23, Definition 5.25], then [DH23, Proposition 8.14] shows that (K, v, σ) is σ-separably σ-algebraically maximal. In particular, any (non-principal) ultraproduct of separably closed, non-trivially valued, valued fields equipped with a Frobenius map is σ-separably σ-algebraically maximal, e.g.

$$(K, v, \sigma) := \prod_{p \in \mathcal{U}} ((\mathbb{F}_p^{\text{alg}}(t))^{\text{sep}}, v_t, x \mapsto x^p).$$

Lemma 4.3.16. *Suppose (K, v, σ) is a model of* FE. *Then, K is σ-separably σ-algebraically maximal if and only if K^{inv} is σ-algebraically maximal.*

Proof. Suppose K is σ-separably σ-algebraically maximal, and suppose K^{inv} is not σ-algebraically maximal. Let $K^{\text{inv}} \subseteq K^{\text{inv}}(\alpha)_{\sigma,\sigma^{-1}}$ be a proper immediate σ-algebraic extension. Let $f(X) \in K^{\mathrm{inv}}[X]_\sigma$ be a non-zero difference polynomial such that $f(\alpha) = 0$, and write

$$f(X) = \sum_I \sigma^{-n}(b_I) X^I,$$

for some $n \geqslant 0$, and $b_I \in K$. If $g = \mathfrak{S}(\sigma^n(f))$, then g is defined over K, $g' \neq 0$, and $g(\sigma^m(\alpha)) = 0$, where m is the least such that $\sigma^n(f)_m \neq 0$ (see SECTION 4.2). As $\alpha \notin K^{\mathrm{inv}}$, then $\sigma^m(\alpha) \notin K$, and thus $K \subseteq K(\sigma^m(\alpha))_\sigma$ is a proper extension. Moreover, it is a subextension of the immediate extension $K \subseteq K^{\mathrm{inv}}(\alpha)_{\sigma,\sigma^{-1}}$, so it is a proper, immediate and, by FACT 4.2.14(2), σ-separably σ-algebraic extension of K. This is a contradiction.

Conversely, suppose K^{inv} is σ-algebraically maximal, but K is not σ-separably σ-algebraically maximal. Let $K \subseteq K(\alpha)_\sigma$ be a proper σ-separably σ-algebraic immediate extension, thus in particular $\alpha \notin K^{\mathrm{inv}}$. Then, $K \subseteq K(\alpha)_\sigma \subseteq K(\alpha)_\sigma^{\mathrm{inv}}$ is a tower of immediate extensions which contains K^{inv}, thus in particular the subextension $K^{\mathrm{inv}} \subseteq K^{\mathrm{inv}}(\alpha)_\sigma$ is still immediate. As $\alpha \notin K^{\mathrm{inv}}$, this is a proper immediate σ-algebraic extension, a contradiction. $\square$

Lemma 4.3.17. *Suppose (K, v, σ) is a model of* FE. *Then, K is σ-separably maximal if and only if K^{inv} is maximal.*

Proof. Suppose K is σ-separably maximal, and suppose K^{inv} is not maximal. Let $K^{\mathrm{inv}} \subseteq K^{\mathrm{inv}}(\alpha)_{\sigma,\sigma^{-1}}$ be a proper immediate extension. If α is σ-algebraic over K^{inv}, we argue as in LEMMA 4.3.16. Otherwise, $K \subseteq K(\alpha)_\sigma$ is a proper σ-separable immediate extension, a contradiction.
Conversely, suppose K^{inv} is maximal and K is not σ-separably maximal. Then we can argue as in LEMMA 4.3.16 (note that the argument there does not use σ-algebraicity in any way). $\square$

Lemma 4.3.18. *Suppose (K, v, σ) is a model of* FE. *Then it admits a σ-separably σ-algebraically maximal, σ-separably σ-algebraic immediate extension K' that is a model of* FE.

Proof. Using Zorn's lemma, we let K' be maximal among the σ-separably σ-algebraic immediate extensions of K that be a model of FE. We argue that K' is σ-separably σ-algebraically maximal. Suppose not, i.e. there is some proper σ-separably σ-algebraic immediate extension $K' \subseteq L$. Then, by LEMMA 4.2.19, $K' \subseteq \mathrm{FE}(L)$ is still a proper σ-separably σ-algebraic immediate extension that is a model of FE, which is a contradiction. $\square$

We recall the fundamental result from [DO15] that sets the scene for our THEOREM 4.3.21.

Fact 4.3.19 ([DO15, Theorem 5.8])**.** *Suppose (K, v, σ) is inversive with $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Then all its σ-algebraically maximal immediate σ-algebraic extensions are isomorphic over K, and all its maximal immediate extensions are isomorphic over K.*

We are now ready to prove the main ingredient of LEMMA 4.5.1, namely the uniqueness of certain maximal immediate extensions. We first prove an asymmetric version of the uniqueness theorem.

Proposition 4.3.20. *Suppose (K, v, σ) is a model of* FE *and $(k_K, \sigma_{\text{res}})$ is linearly difference closed. Let (K_1, v, σ) be σ-separably σ-algebraically maximal, σ-separably σ-algebraic over K, and a model of* FE*. Let (K_2, v, σ) be σ-separably σ-algebraically maximal, σ-separable and immediate over K, and a model of* FE*. Then there is an embedding $\phi\colon K_1 \hookrightarrow K_2$ over K such that $\phi(K_1) \subseteq K_2$ is σ-separable.*

Proof. By LEMMA 4.3.16 and LEMMA 4.3.17, K_1^{inv} is σ-algebraically maximal and K_2^{inv} is maximal, and further both are immediate extensions of K, and thus of K^{inv}. Hence, by FACT 4.3.19, there is a K^{inv}-embedding $\phi\colon K_1^{\text{inv}} \to K_2^{\text{inv}}$. We use ϕ to embed K_1 into K_2 over K.

Given any $\alpha \in K_1$, we argue that $\phi(\alpha) \in K_2$. Suppose not, i.e. $\phi(\alpha) \in K_2^{\text{inv}} \backslash K_2$; equivalently, $K_2 \subseteq K_2(\phi(\alpha))_\sigma$ is a purely σ-inseparable extension. Since K_1 is σ-separably σ-algebraic over K, we let $f(X) \in K[X]_\sigma$ be a difference polynomial with $f(\alpha) = 0$ and $f'(\alpha) \neq 0$. Then, $\phi(\alpha)$ is also such that $f(\phi(\alpha)) = 0$ and $f'(\phi(\alpha)) \neq 0$. If we see $f(X)$ as a difference polynomial over K_2, we get that $K_2(\phi(\alpha))_\sigma$ is a σ-separably (σ-algebraic) immediate extension of K_2. This is a contradiction.

Thus, $\phi(K_1) \subseteq K_2$ and, since $K \subseteq \phi(K_1)$ is a σ-separably σ-algebraic extension, then $\phi(K_1) \subseteq K_2$ is σ-separable. □

Theorem 4.3.21. *Suppose (K, v, σ) is a model of* FE *and $(k_K, \sigma_{\text{res}})$ is linearly difference closed. Let (K_1, v, σ) and (K_2, v, σ) be two σ-separably σ-algebraically maximal, σ-separably σ-algebraic immediate extensions of K that are also models of* FE*. Then $K_1 \cong_K K_2$.*

Proof. From PROPOSITION 4.3.20, we get an embedding $\phi\colon K_1 \to K_2$ over K such that $\phi(K_1) \subseteq K_2$ is σ-separable. Then, since K_2 is σ-algebraic over K, $\phi(K_1) \subseteq K_2$ is σ-separably σ-algebraic and immediate. Thus ϕ must be surjective. □

Definition 4.3.22. Whenever K is as above, we denote the unique σ-separably σ-algebraically maximal, σ-separably σ-algebraic immediate extension that is a model of FE of K by $\overline{K}$.

Remark 4.3.23. Note that, by LEMMA 4.3.16, $\overline{K}^{\text{inv}}$ is σ-algebraically maximal.

Corollary 4.3.24. *Suppose (K, v, σ) is a model of* FE *and $(k_K, \sigma_{\text{res}})$ is linearly difference closed. Let (L, v, σ) be some σ-separably maximal σ-separable immediate extension of K. Then, there is a K-embedding $\phi\colon \overline{K} \to L$ such that L is σ-separable over $\phi(\overline{K})$.*

Proof. This is an application of PROPOSITION 4.3.20. □

Before continuing, we establish that (as one might expect) a σ-separably σ-algebraically maximal valued difference field is in fact weakly σ-henselian.

Fact 4.3.25 ([DO15, Corollary 5.6(2)]). *Let (K, v, σ) be σ-algebraically maximal, and $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Then (K, v, σ) is strongly σ-henselian.*

Lemma 4.3.26. *Let (K, v, σ) be a model of* FE*, and suppose $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Then, $\overline{K}$ is weakly σ-henselian.*

Proof. By FACT 4.3.25, $\overline{K}^{\mathrm{inv}}$ is strongly σ-henselian. As $\overline{K}$ is σ-separably σ-algebraically closed in $\overline{K}^{\mathrm{inv}}$ and a model of FE, $\overline{K}$ is weakly σ-henselian. □

Corollary 4.3.27. *Let (K, v, σ) be a model of* FE*, and suppose $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Then $\overline{K}$ is deeply σ-ramified.*

Proof. This is a consequence of LEMMA 4.3.26 and COROLLARY 4.3.12. □

Pseudo-Cauchy sequences and pseudolimits. So far, we have avoided talking about pseudo-Cauchy sequences, sweeping them under the rug of FACT 4.3.19. We now ought to make a stop and discuss them, for the sole purpose of establishing LEMMA 4.3.38.

Consider valued (difference) fields $K \subseteq L$. Given a limit ordinal λ, we call $(a_\rho)_{\rho<\lambda} \subseteq K$ a *pseudo-Cauchy sequence* if there is $\overline{\rho} < \lambda$ such that, for every $\rho_2 > \rho_1 > \rho_0 \geqslant \overline{\rho}$,

$$v(a_{\rho_2} - a_{\rho_1}) > v(a_{\rho_1} - a_{\rho_0}).$$

For any $\rho \geqslant \overline{\rho}$, we write $\gamma_\rho := v(a_{\rho+1} - a_\rho)$; note that $\gamma_\rho = v(a_\mu - a_\rho)$ for every $\mu > \rho$. We call the sequence $(\gamma_\rho)_{\rho<\lambda}$ the *radii* of the sequence. If $(\gamma_\rho)_{\rho<\lambda} \subseteq \Gamma_K$ is cofinal, we also say that $(a_\rho)_{\rho<\lambda}$ is a *Cauchy sequence*. We say that $a \in L$, for some extension L of K, is a *pseudolimit* of $(a_\rho)_\rho$ if $(v(a - a_\rho))_\rho$ is eventually strictly increasing (equivalently, $v(a - a_\rho) = \gamma_\rho$ for ρ big enough). We then write $a_\rho \rightsquigarrow a$. Given another $b \in L$, we have that $a_\rho \rightsquigarrow b$ if and only if $v(b - a) \geqslant \gamma_\rho$ eventually in ρ. We say that two sequences $(a_\rho)_\rho$ and $(b_\rho)_\rho$ are *equivalent*, and we write $(a_\rho)_\rho \sim (b_\rho)_\rho$, if for every extension L and $a \in L$, $a_\rho \rightsquigarrow a$ if and only if $b_\rho \rightsquigarrow a$.

In general, one does not have that, for any difference polynomial $p(X)$, if $a_\rho \rightsquigarrow a$, then $p(a_\rho) \rightsquigarrow p(a)$.[6] This can be easily fixed, however, by moving to an equivalent pseudo-Cauchy sequence.

Definition 4.3.28. We say that a difference field (K, σ) is *aperiodic* if for every $n > 0$ there is $\alpha \in K$ with $\sigma^n(\alpha) \neq \alpha$.

Remark 4.3.29. If (K, σ) is linearly difference closed, then it is in particular aperiodic.

Fact 4.3.30 ([DO15, Theorem 3.8]). *Suppose (K, v, σ) is a valued difference field and $(k_K, \sigma_{\mathrm{res}})$ is aperiodic. Suppose $(a_\rho)_\rho \subseteq K$ is a pseudo-Cauchy sequence. Take a in some extension of K with $a_\rho \rightsquigarrow a$. Let $\Sigma \subseteq K[X]_\sigma$ be finite. Then, there is a pseudo-Cauchy sequence $(b_\rho)_\rho \subseteq K$ with $(a_\rho)_\rho \sim (b_\rho)_\rho$ and such that for every non-constant $p \in \Sigma$, $p(b_\rho) \rightsquigarrow p(a)$.*

[6]This is true if σ_{val} is ω-increasing, i.e. for all $n > 0$ and $\gamma > 0$, $\sigma_{\mathrm{val}}(\gamma) > n\gamma$. See [Azg10, Lemma 4.10].

Given a difference polynomial $p(X)$, in particular, we have that $p(a_\rho) \rightsquigarrow 0$ if and only if $(v(p(a_\rho)))_\rho$ is eventually strictly increasing.

Definition 4.3.31. Let $(a_\rho)_\rho \subseteq K$ be a pseudo-Cauchy sequence. For $I \in \mathbb{N}^3$, denote by $K[X]_\sigma^I$ the set of difference polynomials of complexity I. Let

$$\mathrm{W}_I((a_\rho)_\rho) = \{p(X) \in K[X]_\sigma^I \mid \exists (b_\rho)_\rho \sim (a_\rho)_\rho (p(b_\rho) \rightsquigarrow 0)\},$$

and $\mathrm{W}((a_\rho)_\rho) = \bigcup_{I \in \mathbb{N}^3} \mathrm{W}_I((a_\rho)_\rho)$. We say that $(a_\rho)_\rho$ is:

1. of *σ-transcendental type* if $\mathrm{W}((a_\rho)_\rho) = \varnothing$,
2. of *σ-algebraic type* otherwise.

If $(a_\rho)_\rho$ is of σ-algebraic type, and J is the least such that $\mathrm{W}_J((a_\rho)_\rho)$ is non-empty, then we write $\mathrm{W}_{\min}((a_\rho)_\rho)$ for $\mathrm{W}_J((a_\rho)_\rho)$.

Definition 4.3.32. We say that $(a_\rho)_\rho$ is of *σ-separably σ-algebraic type* if it is of σ-algebraic type and there is $f \in \mathrm{W}_{\min}((a_\rho)_\rho)$ with $f' \neq 0$.

The following is a result which is well-known, and can be seen by direct computation; for example it is partially shown in [Azg10, Lemma 6.2] and [BMS07, Lemma 7.2].

Lemma 4.3.33 (♢)**.** *Let (K, v, σ) be a valued difference field, with $(k_K, \sigma_{\mathrm{res}})$ aperiodic, and let $(a_\rho)_\rho \subseteq K$ be a pseudo-Cauchy sequence without pseudolimit in K. Let $I \in \mathbb{N}^3$ be minimal with $\mathrm{W}_I((a_\rho)_\rho) \neq \varnothing$, if there is such an I; otherwise let $I := \infty > \mathbb{N}^3$. Let $f(X) \in K[X]_\sigma$ have complexity strictly smaller than I, and (L, v, σ) be some extension of K containing a pseudolimit a of $(a_\rho)_\rho$. Then,*

1. *$v(f(a)) \in \Gamma_K$, and*
2. *if $v(f(a)) \geqslant 0$, then $\mathrm{res}(f(a)) \in k_K$.*

We next show how to deal with adjoining pseudolimits.

Lemma 4.3.34 (♢)**.** *Suppose (K, v, σ) is a valued difference field, with $(k_K, \sigma_{\mathrm{res}})$ aperiodic, and let $(a_\rho)_\rho \subseteq K$ be a pseudo-Cauchy sequence without pseudolimit in K. Then,*

1. *if $(a_\rho)_\rho$ is of σ-transcendental type, there is an immediate extension $K(a)_\sigma$ of K such that $a_\rho \rightsquigarrow a$ and a is σ-transcendental over K, and for each pseudolimit b of $(a_\rho)_\rho$ in a valued difference field L extending K there is unique embedding $K(a)_\sigma \to L$ over K with $a \mapsto b$,*
2. *if $(a_\rho)_\rho$ is of σ-algebraic type and $p \in \mathrm{W}_{\min}((a_\rho)_\rho)$, then there is an immediate $K(a)_\sigma$ of K such that $a_\rho \rightsquigarrow a$ and $p(a) = 0$, and for each pseudolimit b of $(a_\rho)_\rho$ in a valued difference field L extending K with $p(b) = 0$, there is a unique embedding $K(a)_\sigma \to L$ over K with $a \mapsto b$.*

In 2., note that if $p' \neq 0$, then $p'(a) \neq 0$ since the complexity of p' is strictly smaller than that of p.

Proof. We briefly discuss how to adapt the proofs of [AvdD11, Lemmata 5.2 and 5.3] to our case. The proof of Lemma 5.2 from [AvdD11], which is actually given in [AvdD09, Lemma 6.2], translates verbatim.
The proof of Lemma 5.3, which is actually given in [AvdD09, Lemma 6.4], is non-trivial to begin with. Nevertheless, the necessary ingredients are still available in our setting, even when σ is non-inversive and σ_{val} is non-trivial. We sketch out the steps of the construction and highlight the tricky points.

1. First, we write $p(X) = P(X, \sigma(X), \dots \sigma^n(X))$ and note that P is irreducible over K. We take $K[\xi_0, \dots \xi_n] := K[X_0, \dots X_n]/(P)$, where $\xi_i := X_i + (P)$, and consider its field of fractions L.

2. We define v on $L^\times$ using a pseudolimit e of $(a_\rho)_\rho$ in some extension. Namely, for any $d = F(\xi_0, \dots \xi_n)/G(\xi_0, \dots \xi_{n-1}) \in L^\times$ where F has lower X_n-degree of P and $G \in K[X_0, \dots X_{n-1}]$ is non-zero, we note that $v(F(e, \sigma(e), \dots \sigma^n(e)))$ and $v(G(e, \sigma(e), \dots \sigma^{n-1}(e)))$ are both elements of Γ_K and
$$v(F(e, \sigma(e), \dots \sigma^n(e))) - v(G(e, \sigma(e), \dots \sigma^{n-1}(e)))$$
only depends on d. Note that this can be proven by switching to equivalent pseudo-Cauchy sequences, something that can also be done in our setting via FACT 4.3.30. We can then define v on $L^\times$ via
$$v(d) := v(F(e, \sigma(e), \dots \sigma^n(e))) - v(G(e, \sigma(e), \dots \sigma^{n-1}(e)))$$
and check that it is a valuation with the same residue field and value group as (K, v).

3. Contrary to the inversive case, it is not in general true that $K(\xi_1, \dots \xi_n)$ has transcendence degree n over K (namely, X_0 could very well not appear in P). In other words, we only have an embedding
$$\sigma \colon K(\xi_0, \dots \xi_{n-1}) \longrightarrow K(\xi_1, \dots \xi_n)$$
between subfields of L. We still note that is is an embedding of valued fields.

4. Unlike in the inversive case, L need not be an algebraic extension of $K(\xi_1, \dots \xi_n)$. However, L is still an algebraic immediate extension of $K(\xi_0, \dots \xi_{n-1})$, so we get that
$$L^h = K(\xi_0, \dots \xi_{n-1})^h.$$
By the universal property of the henselization, then, σ extends to an embedding
$$\sigma \colon L^h \longrightarrow \sigma(K)(\xi_1, \dots \xi_n)^h \subseteq L^h.$$

If we let $a := \xi_0$, then, we have that $\sigma(K(a)_\sigma) \subseteq K(a)_\sigma \subseteq L^h$, giving us the valued difference field we needed. □

Remark 4.3.35. If we had started with (K, v, σ) that is a model of FE, then the resulting immediate extension $K(a)_\sigma$ with $p(a) = 0$, $p'(a) \neq 0$ would be σ-separably σ-algebraic.

Corollary 4.3.36. *Let (K, v, σ) be σ-separably σ-algebraically maximal and that is a model of* FE*. Assume that $(k_K, \sigma_{\mathrm{res}})$ is aperiodic. Then, all pseudo-Cauchy sequences of σ-separably σ-algebraic type have a pseudolimit in K.*

The following result might seem underwhelming, but it is of fundamental importance in the final steps of the back-and-forth for LEMMA 4.5.1. In spirit, it comes from the theory of *dependent defect*, as developed in [KR23].

Remark 4.3.37. If K is deeply σ-ramified, i.e. $K \subseteq K^{\mathrm{inv}}$ is dense, then for any $a \in K^{\mathrm{inv}}$ there is a Cauchy sequence from K converging to a (see [AvdDvdH17, Lemma 2.2.26]). If $a \notin K$, then $v(a - K) = \{v(a - b) \mid b \in K\} \subseteq \Gamma_K$ is cofinal.

Lemma 4.3.38. *Suppose (K, v, σ) is σ-separably σ-algebraically maximal, is a model of* FE*, and $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed. Let $K \subseteq K(t)_\sigma$ be an immediate σ-transcendental extension, and let $(a_\rho)_\rho \subseteq K$ be a pseudo-Cauchy sequence with t as pseudolimit. Then $(a_\rho)_\rho$ is of σ-transcendental type.*

Proof. Suppose not. By LEMMA 4.3.34, we can find a proper immediate σ-algebraic extension $K \subseteq K(a)_\sigma$, with $a_\rho \rightsquigarrow a$. For n big enough, we can split the extension into the tower $K \subseteq K(\sigma^n(a))_\sigma \subseteq K(a)_\sigma$, where the bottom part is σ-separably σ-algebraic, and hence trivial. We are then left with a purely σ-inseparable extension $K \subseteq K(a)_\sigma$. Now, since K is deeply σ-ramified (COROLLARY 4.3.27), $v(a - K)$ is cofinal in Γ_K. On the other hand, since $K \subseteq K(t, a)_\sigma \subseteq K(t)_\sigma^{\mathrm{inv}}$ is a tower of immediate extensions, we can compute $v(a - t) \in \Gamma_K$. Since both are pseudolimits, $v(a - t) \geqslant v(a - K)$; as they are not isomorphic over K, because a is σ-algebraic and t is σ-transcendental, then $v(a - t) < \infty$. In particular, $v(a - K) \leqslant v(a - t)$ is not cofinal in Γ_K, a contradiction. □

4.4 | Setting up the embedding lemma

We now have almost all the tools to establish the embedding lemma and deduce relative quantifier elimination. We introduce the leading terms structure, the languages and theories, and take a brief detour to prove that one can embed $\overline{K}$ in saturated models. We then establish recipes for the auxiliary steps, namely increasing the residue difference field and value difference group.

The leading terms structure. Given a valued field (K, v), we start to define the leading terms structure of (K, v) by considering the set $\mathrm{RV}_K := (K^\times / (1 + \mathfrak{m}_K)) \cup \{0\}$. Denote by $\mathrm{RV}_K^\times := \mathrm{RV}_K \backslash \{0\}$. We write rv for the quotient map $K^\times \to \mathrm{RV}_K$, extended to K via $\mathrm{rv}(0) := 0$.

We have, for $a, b \in K^\times$,

$$\mathrm{rv}(a) = \mathrm{rv}(b) \iff v(a-b) > v(a) \iff v(a-b) > v(b).$$

In particular, $\mathrm{rv}(a) = \mathrm{rv}(b)$ implies that $v(a) = v(b)$, and so $v\colon K^\times \to \Gamma_K$ induces a map $v_{\mathrm{rv}}\colon \mathrm{RV}_K^\times \to \Gamma_K$. On the other hand, it gives rise to a short exact sequence of groups

$$1 \longrightarrow k_K^\times \xhookrightarrow{\iota} \mathrm{RV}_K^\times \xrightarrow{v_{\mathrm{rv}}} \Gamma_K \longrightarrow 0.$$

We endow $\mathrm{RV}_K^\times$ with the multiplicative structure inherited from $K^\times$, extended to RV_K by $a \cdot 0 = 0$ for all $a \in \mathrm{RV}_K$. Further, we define the ternary relation

$$\oplus(\alpha, \beta, \gamma) \iff \exists a, b \in K^\times (\mathrm{rv}(a) = \alpha \wedge \mathrm{rv}(b) = \beta \wedge \mathrm{rv}(a+b) = \gamma).$$

We call $(\mathrm{RV}_K, \oplus, \cdot, 0, 1)$ the *leading terms structure* of (K, v). We also write $\alpha \oplus \beta := \{\gamma \in \mathrm{RV}_K \mid \oplus(\alpha, \beta, \gamma)\}$. We say that $\alpha \oplus \beta$ is *well-defined* if $\alpha \oplus \beta = \{\gamma\}$ for some γ, in which case we write also $\alpha \oplus \beta = \gamma$.

For $\alpha_1, \dots \alpha_n \in \mathrm{RV}_K$, we write

$$\oplus(\alpha_1, \dots \alpha_n, \beta) \iff \exists a_1, \dots a_n \left(\bigwedge_{i=1}^{n} \mathrm{rv}(a_i) = \alpha_i \wedge \mathrm{rv}(a_1 + \cdots + a_n) = \beta \right).$$

We again say that $\alpha_1 \oplus \cdots \oplus \alpha_n$ is *well-defined* if there is exactly one β such that $\oplus(\alpha_1, \dots \alpha_n, \beta)$.

Lemma 4.4.1. *Let $\alpha_1, \dots \alpha_n \in \mathrm{RV}_K^\times$. Choose some representatives $a_1, \dots a_n \in K$ so that $\alpha_i = \mathrm{rv}(a_i)$, for $i = 1, \dots n$. Then $\alpha_1 \oplus \cdots \oplus \alpha_n$ is well-defined if and only if $v(a_1 + \cdots + a_n) = \min_{i=1,\dots n} v(a_i)$.*

Given $\sigma \in \mathrm{End}(K, v)$, then there is an induced $\sigma_{\mathrm{rv}} \in \mathrm{End}(\mathrm{RV}_K)$.

Remark 4.4.2. By applying the Short Five Lemma ([Wei94, Exercise 1.3.3]) to

$$\begin{tikzcd}
1 \arrow[r] & k_K^\times \arrow[r] \arrow[d, "\sigma_{\mathrm{res}}"] & \mathrm{RV}_K^\times \arrow[r] \arrow[d, "\sigma_{\mathrm{rv}}"] & \Gamma_K \arrow[r] \arrow[d, "\sigma_{\mathrm{val}}"] & 0 \\
1 \arrow[r] & k_K^\times \arrow[r] & \mathrm{RV}_K^\times \arrow[r] & \Gamma_K \arrow[r] & 0
\end{tikzcd}$$

we get that σ_{rv} is surjective if and only if both σ_{val} and σ_{res} are.

We call $(\mathrm{RV}_K, \oplus, \cdot, 0, 1, \sigma_{\mathrm{rv}})$ the *leading terms (difference) structure* of (K, v, σ).

Languages and theories. We now introduce the language that we will use to prove relative quantifier elimination.

Definition 4.4.3. We let $\mathcal{L}_{3,\mathrm{RV},\sigma}$ be the three-sorted language whose sorts are given as follows:

1. **VF** is the main sort, with language $\mathcal{L}_{\mathrm{ring},\sigma} = \{+,\cdot,-,0,1,\sigma\}$,
2. **RV** has language $\mathcal{L}_{\mathbf{RV}} := \{\oplus,\cdot,0,1,\sigma_{\mathrm{rv}}\}$,
3. **VG** has language $\mathcal{L}_{\mathrm{oag},\sigma} := \{+,\leqslant,0,\infty,\sigma_{\mathrm{val}}\}$,

with functions given by $\mathrm{rv}\colon \mathbf{VF} \to \mathbf{RV}$ and $v_{\mathrm{rv}}\colon \mathbf{RV} \to \mathbf{VG}$. We let $\mathcal{L}_{\mathbf{RV},\mathbf{VG}}$ be the reduct of $\mathcal{L}_{3,\mathrm{RV},\sigma}$ to the sorts **RV** and **VG**.

Remark 4.4.4. Note that one can define v in $\mathcal{L}_{3,\mathrm{RV},\sigma}$ as $v = v_{\mathrm{rv}} \circ \mathrm{rv}$. Moreover, in any valued difference field (K,v,σ), $(\Gamma_K,+,0,\leqslant,\infty,\sigma_{\mathrm{val}})$ and $(k_K,+,\cdot,-,0,1,\sigma_{\mathrm{res}})$ are interpretable in $(\mathrm{RV}_K,\oplus,\cdot,0,1,\sigma_{\mathrm{rv}})$.

Definition 4.4.5. We let $\mathcal{L}^{\lambda}_{\mathrm{ring},\sigma}$ be the expansion of $\mathcal{L}_{\mathrm{ring},\sigma}$ where we adjoin, for every $n \geqslant 1$, $i \in \{1,\dots n\}$, an $(n+1)$-ary function λ^i_n. We let $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$ be the expansion of $\mathcal{L}_{3,\mathrm{RV},\sigma}$ where we give **VF** the language $\mathcal{L}^{\lambda}_{\mathrm{ring},\sigma}$.

Definition 4.4.6. We let FE^{λ} be the $\mathcal{L}^{\lambda}_{\mathrm{ring},\sigma}$-theory of difference fields extending FE where we interpret λ^i_n as follows. Let $x_1,\dots x_n, y \in K$, and assume that $x_1,\dots x_n$ are $\sigma(K)$-linearly independent, with $y \in \mathrm{span}_{\sigma(K)}(x_1,\dots x_n)$. Then, $\lambda^1_n(\bar{x},y)\dots\lambda^n_n(\bar{x},y)$ are the unique elements of K such that

$$y = \sum_{i=1}^{n} \sigma(\lambda^i_n(x_1,\dots x_n,y))x_i.$$

Otherwise, set $\lambda^1_n(x_1,\dots x_n,y) = \dots = \lambda^n_n(x_1,\dots x_n,y) = 0$.

Remark 4.4.7. If $(L,\sigma) \models \mathrm{FE}^{\lambda}$ and $(K,\sigma) \subseteq (L,\sigma)$ is a difference subfield, then as $\mathcal{L}^{\lambda}_{\mathrm{ring},\sigma}$-structures we have that $(K,\sigma) \subseteq (L,\sigma)$ if and only if the extension is σ-separable. Indeed, (K,σ) is closed under the λ-functions precisely if and only if K is linearly disjoint from $\sigma(L)$ over $\sigma(K)$.

Definition 4.4.8. We let VFE be the $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-theory whose models $K = \langle K, \mathrm{RV}_K, \Gamma_K\rangle$ are non-inversive valued difference fields with leading terms difference structure RV_K, value difference group $(\Gamma_K, \sigma_{\mathrm{val}})$, and further (K,σ) is a model of FE^{λ}. We let VFE_0 extend VFE by further requiring that models have residue characteristic zero.

Definition 4.4.9. We let hVFE be the $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-theory extending VFE whose models (K,v,σ) are non-inversive weakly σ-henselian and such that $\sigma(K)$ is algebraically closed in K (i.e., they are also models of FE). We let hVFE_0 extend hVFE by further requiring that models have residue characteristic zero.

Remark 4.4.10. Note that an equicharacteristic zero inversive weakly σ-henselian valued difference field is precisely an equicharacteristic zero strongly σ-henselian valued difference field, as considered by [DO15].

We now prove a somewhat surprising lemma about the behaviour of λ-functions on subrings of our models; this mimics the analogous phenomenon in separably closed valued fields.

Lemma 4.4.11. *Let $K \models \mathrm{hVFE}$ and $A \leqslant K$. Then, there is a unique $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-structure on $\mathrm{Frac}(A)$ extending the one on A. In other words, given $A \leqslant K$, we may always assume that A is a field.*

Proof. As rv is multiplicative, it extends uniquely to $\mathrm{Frac}(A)$. The λ-functions also extend uniquely to $\mathrm{Frac}(A)$: indeed, if $A \subseteq K_{\mathrm{inv}}$, then $\mathrm{Frac}(A) \subseteq K_{\mathrm{inv}}$, and the λ-functions are trivial. Otherwise, assume that $A \nsubseteq K_{\mathrm{inv}}$: we show that $A = \mathrm{Frac}(A)$. Indeed, if $A \nsubseteq K_{\mathrm{inv}}$ then by induction we may assume that there is $a \in A \backslash \sigma(K)$, and so we can compute $\lambda^1_1(\sigma(a), 1) = \frac{1}{a} \in A$. For any $b \in \sigma(K) \cap A$, we have $ab \notin \sigma(K)$, and thus by our previous considerations $\frac{1}{ab} \in A$, so $\frac{1}{b} \in A$. □

Building models of hVFE_0. We exhibit a recipe for building models with prescribed residue difference field and value difference group. The standard Hahn constructions do not work, for reasons clarified in REMARK 4.4.15.

Definition 4.4.12. Let (K, σ) be a difference field. We say that it is *weakly linearly difference closed* if for every $b, a_0, \cdots a_n \in K$ with $a_0 \neq 0$, there is $x \in K$ with

$$b + a_0 x + a_1 \sigma(x) + \cdots + a_n \sigma^n(x) = 0.$$

Remark 4.4.13. If (K, σ) is weakly linearly difference closed, then K^{inv} is linearly difference closed.

Let $(k, \sigma_{\mathrm{res}})$ be a weakly linearly difference closed, non-inversive difference field of characteristic zero, and $(\Gamma, \sigma_{\mathrm{val}})$ a non-inversive ordered difference group. For example, one could take $(k, \sigma_{\mathrm{res}})$ to be a characteristic zero model of SCFE as defined in [CH04], and $(\Gamma, \sigma_{\mathrm{val}})$ to be given by $\mathbb{Z}$ together with the map $x \mapsto 2x$.

Consider the generalized power series field $K := k((t^{\Gamma}))$ with the t-adic valuation v and the lift of σ_{res} and σ_{val} given by

$$\sigma\left(\sum_{\gamma \in \Gamma} a_\gamma t^\gamma\right) = \sum_{\gamma \in \Gamma} \sigma_{\mathrm{res}}(a_\gamma) t^{\sigma_{\mathrm{val}}(\gamma)}.$$

Let $L := k^{\mathrm{inv}}((t^{\Gamma^{\mathrm{inv}}}))$. There is a natural embedding $K \subseteq L$[7], and L satisfies FE since it is inversive. Then,

[7]Note that L is not in general the inversive closure of K. Any element of K^{inv} lives into $\sigma^{-n}(K)$ for some n, in particular its coefficients will live in $\sigma^{-n}_{\mathrm{res}}(k)$. On the other hand, an element of L could have coefficients coming from $\sigma^{-m}_{\mathrm{res}}(k)$ for all $m \geqslant 0$.

Proposition 4.4.14. *The σ-separably σ-algebraic closure $\widetilde{\mathrm{FE}(K)}$ of* $\mathrm{FE}(K)$ *in L is a model of* hVFE_0*. It has residue difference field $(k^{\mathrm{inv}}, \sigma_{\mathrm{res}})$ and value difference group $(\Gamma^{\mathrm{inv}}, \sigma_{\mathrm{val}})$.*

Proof. Since k is weakly linearly difference closed, k^{inv} is linearly difference closed, and so (by FACT 4.3.25) L is strongly σ-henselian. Then, by construction, $\widetilde{\mathrm{FE}(K)}$ is weakly σ-henselian. In particular, it has inversive residue difference field and inversive value difference group by COROLLARY 4.3.12, thus they are equal to k^{inv} and Γ^{inv}. As $\mathrm{FE}(K)$ is not inversive, and $\widetilde{\mathrm{FE}(K)}$ is a σ-separable extension of $\mathrm{FE}(K)$, it is still not inversive. □

Remark 4.4.15. As $(k((t^\Gamma)), v)$, as defined above, is maximal as a valued field (i.e. it has no proper immediate extension), then whenever $\sigma \in \mathrm{End}(k((t^\Gamma)), v)$, we have that σ is surjective if and only if both σ_{res} and σ_{val} are. Indeed, if both σ_{res} and σ_{val} are surjective, $\sigma(k((t^\Gamma))) \subseteq k((t^\Gamma))$ is an immediate extension, and thus it must be trivial. In particular, this means that there is no non-inversive model of hVFE of the form $(k((t^\Gamma)), v, \sigma)$.

Embeddings in models. We explain how σ-separable immediate extensions can be embedded in saturated models.

Fact 4.4.16 ([DO15, Corollary 5.10])**.** *Let (K, v, σ) be an inversive valued difference field with linearly difference closed residue difference field. Let (K', v, σ) be a $|\Gamma_K|^+$-saturated strongly σ-henselian extension of K. Then, any maximal immediate extension of K embeds into K' over K.*

Lemma 4.4.17. *Suppose (K, v, σ) is a model of* FE *and $(k_K, \sigma_{\mathrm{res}})$ is linearly difference closed.*

Let $(K, v, \sigma) \subseteq (K^*, v, \sigma)$ be a σ-separable extension, where (K^*, v, σ) is a model of FE*, is $|\Gamma_K|^+$-saturated and weakly σ-henselian. Suppose $K \subseteq M$ is some σ-separable immediate extension that is a model of* FE*. Then there is a K-embedding $f : M \to K^*$ such that K^* is σ-separable over $f(M)$.*

Proof. We first argue that there is a K-embedding of $\overline{K}$ into K^* such that K^* is σ-separable over the image.
Consider a $|K|^+$-saturated (and thus in particular $|\Gamma_K|^+$-saturated) elementary extension L of $(K^*)^{\mathrm{inv}}$. Then, by LEMMA 4.3.8 and FACT 4.4.16, $(\overline{K})^{\mathrm{inv}}$ embeds into L over K^{inv}, in particular over K, along a map θ. Now, by saturation, it is enough[8] to exhibit a witness in K^* for every quantifier-free $\mathcal{L}_{3,\mathrm{RV},\sigma}$-formula φ over K such that, for some finite tuple $\alpha = (a_1, \dots a_r)$ in $\overline{K}$, $\overline{K} \vDash \varphi(\alpha)$. Upon strenghtening φ, we may assume all of its solutions are simple zeroes of difference polynomials over K.

The tuple $(\theta(a_1), \dots \theta(a_r)) \in L^r$ satisfies φ. As we have taken $(K^*)^{\mathrm{inv}} \preceq L$, then there exist elements $b_1, \dots b_r \in (K^*)^{\mathrm{inv}}$ such that $(K^*)^{\mathrm{inv}} \vDash \varphi(b_1, \dots b_r)$. But $b_1, \dots b_r$ are simple zeroes of difference polynomials over K, in particular over K^*, and thus they are in K^* already, as K^* is σ-separably σ-algebraically closed in its inversive closure.

[8] This clever trick is taken from [DH23, Proposition 5.8].

We now go back to M. We show that there is a K-embedding f of $\overline{M}$ into K^* such that K^* is σ-separable over $f(\overline{M})$. As $\overline{K} \subseteq \overline{M}$ (by PROPOSITION 4.3.20), it remains to find a way to extend the K-embedding $\overline{K} \to K^*$ to $\overline{M}$. By saturation, we only need to extend the embedding to subextensions $\overline{K} \subseteq N \subseteq \overline{M}$ which are finitely generated over $\overline{K}$ as models of FE. Then, by FACT 4.2.26, N is finitely σ-separably generated over $\overline{K}$, so there is a finite σ-transcendence basis $\mathcal{X} \subseteq N$ such that $\overline{K}(\mathcal{X})_\sigma \subseteq N$ is σ-separably σ-algebraic. Enumerate $\mathcal{X} = \{x_1, \cdots x_n\}$. Note that $\overline{K}(x_1)_\sigma \subseteq \overline{N}$ is σ-separable by REMARK 4.2.25. By LEMMA 4.3.38, if we take a pseudo-Cauchy sequence $(a_\rho)_\rho$ with no pseudolimit in $\overline{K}$ and such that $a_\rho \rightsquigarrow x_1$, then $(a_\rho)_\rho$ is of σ-transcendental type. We now argue as in [DH23, Theorem 7.1] to find an appropriate $y \in K^*$.

For the sake of the argument, denote by $\mathcal{B} \subseteq K^*$ the set of pseudolimits of $(a_\rho)_\rho$ in K^*. By saturation, $\mathcal{B}$ is non-empty, and the sequence $(\gamma_\rho)_\rho \subseteq \Gamma_{\overline{K}} \subseteq \Gamma_{K^*}$ of radii of the sequence is not cofinal. Pick $\delta \in \Gamma_{K^*}$ with $\delta > \gamma_\rho$ for every ρ. Moreover, pick some $b \in \mathcal{B}$: then $B_\delta(b) \subseteq \mathcal{B}$, thus $\mathcal{B}$ contains an open ball. On the other hand, consider the set $\Lambda \subseteq K^*$ of elements that are transcendental over $\overline{K} \otimes_{\sigma(\overline{K})} \sigma(K^*)$: by [DH23, Lemma 4.51], Λ is non-empty, and since $\sigma(K^*)\Lambda \subseteq \Lambda$, it must be dense in K^*. Thus there is $y \in \Lambda \cap \mathcal{B}$. Then, the isomorphism $\overline{K}(x_1)_\sigma \to \overline{K}(y)_\sigma$ over $\overline{K}$ gives the required embedding of $\overline{K}(x_1)_\sigma$ into K^*, with K^* σ-separable over $\overline{K}(y)_\sigma$ by [DH23, Proposition 4.50(2)].

We can now replace $\overline{K}$ by $\overline{\overline{K}(x_1)_\sigma} \subseteq \overline{N}$ and start again. □

Adding residues. We establish the two ways of increasing the residue difference field in a back-and-forth situation. Note that the proofs of the following two lemmas from [AvdD11] only rely on σ_{res} and σ_{val} being surjective, and thus apply also in our context. We let $(K, v, \sigma) \subseteq (L, v, \sigma)$ be a valued difference field extension.

We note that the proofs carry over verbatim from [AvdD11, Lemma 2.5] and [AvdD11, Lemma 2.6], changing $K(a)_{\sigma,\sigma^{-1}}$ with $K(a)_\sigma$ wherever necessary.

Fact 4.4.18 ([AvdD11, Lemma 2.5])**.** *Let $a \in \mathcal{O}_{L_{\text{inv}}}$ and assume that $\alpha = \text{res}(a)$ is σ-transcendental over k_K. Then,*

1. *$v(p(a)) = \min_I v(b_I)$, for each $p(X) = \sum_I b_I X^I$ over $\mathcal{O}_K$,*
2. *$K(a)_{\sigma,\sigma^{-1}}$ has residue difference field $k_K(\alpha)_{\sigma,\sigma^{-1}}$ and value difference group Γ_K,*
3. *if b is in $\mathcal{O}_{L'_{\text{inv}}}$, for some extension L' of K, such that $\text{res}(b)$ is σ-transcendental over k_K, then there is a valued difference field isomorphism $K(a)_{\sigma,\sigma^{-1}} \to K(b)_{\sigma,\sigma^{-1}}$ over K, sending a to b.*

Fact 4.4.19 ([AvdD11, Lemma 2.6])**.** *Let $a \in \mathcal{O}_L$ and assume that $\alpha = \text{res}(a)$ is σ-algebraic over k_K. Let $p(X) \in \mathcal{O}_K[X]_\sigma$ be such that $p(a) = 0$, and suppose $p(X)$ has the same complexity as $\text{res}(p)(X) \in k_K[X]_\sigma$ and $\text{res}(p)(X)$ is a nonzero difference polynomial in $k_K[X]_\sigma$ of minimal complexity such that $\text{res}(p)(\alpha) = 0$. Then,*

1. *$K(a)_\sigma$ has residue difference field $k_K(\alpha)_\sigma$ and value difference group Γ_K,*

2. *if b is in $\mathcal{O}_{L'}$, for some extension L' of K, such that $\operatorname{res} p$ is a difference polynomial of minimal complexity such that $\operatorname{res}(p)(\operatorname{res}(b)) = 0$, with $p(b) = 0$, then there is a valued difference field isomorphism $K(a)_\sigma \to K(b)_\sigma$ over K, sending a to b.*

Remark 4.4.20. Note that, since $(k_K, \sigma_{\mathrm{res}})$ is inversive, by exchanging $p(X)$ with the difference polynomial $\sigma^{-m}(\mathfrak{S}(p))(X)$ defined in SECTION 4.2 we can always choose $p(X)$ such that $p' \neq 0$. Thus, if (K, v, σ) is a model of FE, the resulting extension is σ-separably σ-algebraic over K. Upon passing to the FE-closure (thanks to REMARK 4.2.16), we may even assume that the extension is a model of FE, but it might not be generated by one element anymore.

Regular elements. There are several ways to increase the value difference group; here we take the same route as [DO15], which goes via adjoining a very generic element (as opposed to, for example, the route taken by [Rid17]).

Definition 4.4.21. Let $a \in K$ and $p(X) = \sum_I a_I X^I \in K[X]_\sigma$. We say that *$a$ is regular for p* if $v(p(a)) = \min_I\{v(a_I) + v(a^I)\}$. Let $(K, v, \sigma) \subseteq (L, v, \sigma)$ be an extension of valued difference fields. We say that $a \in L$ is *generic over K* if it is regular for every $p \in K[X]_\sigma$.

Proposition 4.4.22 ($\Diamond$)**.** *Let $K \models \mathrm{hVFE}_0$ and let $(E, v, \sigma) \subseteq (K, v, \sigma)$ be a valued difference subfield, possibly with non-inversive residue difference field and value difference group. Assume that K is $|k_E|^+$-saturated, and let $\gamma \in \Gamma_K$. Then, there is $a \in K_{\mathrm{inv}}$, generic over E, such that:*

1. *$v(a) = \gamma$,*

2. *$E(a)_{\sigma,\sigma^{-1}}$ has value group $\Gamma_E(\gamma)_{\sigma,\sigma^{-1}}$,*

3. *if b is another generic over E in some extension (L, v, σ), with $v(b) = \gamma$ and $b \in L_{\mathrm{inv}}$, then the map $a \mapsto b$ gives rise to a valued difference field isomorphism $E(a)_{\sigma,\sigma^{-1}} \to E(b)_{\sigma,\sigma^{-1}}$ over E.*

If, moreover, $(E, \sigma) \subseteq (K, \sigma)$ is σ-separable, then $E(a)_{\sigma,\sigma^{-1}} \subseteq K$ is still σ-separable. Analogously, if $(E, \sigma) \subseteq (L, \sigma)$ is σ-separable, then $E(b)_{\sigma,\sigma^{-1}} \subseteq L$ is still σ-separable.

Remark 4.4.23. Note that this is essentially a twisted version of FACT 4.4.18, although it is from a different point of view; FACT 4.4.18 shows how the isomorphism type is uniquely determined for a generic with $\gamma = 0$.

Proof of PROPOSITION 4.4.22. Note that $K_{\mathrm{inv}} \subseteq K$ is dense (LEMMA 4.3.13), and thus immediate. Choose $c \in K_{\mathrm{inv}}$ such that $v(c) = \gamma$. By saturation, there is $\beta \in k_K$ which is σ-transcendental over $k_{E(c)_{\sigma,\sigma^{-1}}}$. We let $b \in \mathcal{O}_{K_{\mathrm{inv}}}$ be such that $\operatorname{res}(b) = \beta$. Then, by construction, b is a generic of value zero, and $a := cb$ is still a generic. The rest of the properties then follow (arguing as in the proof of FACT 4.4.18) from the fact that if a and b are generics over E, then for every $p(X) = \sum_I a_I X^I \in E[X]_\sigma$ we can compute

$$v(p(a)) = \min_I\{v(a_I) + v(a^I)\} = \min_I\{v(a_I) + I(\gamma)\} = v(p(b)),$$

and thus the difference field isomorphism $E(a)_{\sigma,\sigma^{-1}} \to E(b)_{\sigma,\sigma^{-1}}$ can be upgraded to a valued field isomorphism. $\square$

Remark 4.4.24. PROPOSITION 4.4.22 is proved essentially in the same way as [DO15, Lemma 6.1]. Note, however, that the statement in that paper is incorrect: namely, it states that $E(a)_{\sigma,\sigma^{-1}}$ has the same residue difference field as E. This is false: if, for example, $\gamma = 0$, then we would be extending the residue difference field with a σ-transcendental element (the residue of the generic). This does not create issues in the surjective case, but in our setting it is precisely the reason why we need LEMMA 4.4.27.

We next prove that genericity of an element b is really a property of $\mathrm{rv}(b)$.

Lemma 4.4.25. *Let $K \models \mathrm{hVFE}_0$ and let $A \leqslant_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}} K$ be a valued difference subfield. Let $b \in K$ be generic over A, and let $b' \in K$ be such that $\mathrm{rv}(b') = \mathrm{rv}(b)$. Then b' is generic over A.*

Proof. Let $p(X) = \sum_I a_I X^I \in A[X]_\sigma$. Since b is generic over A,

$$v(p(b)) = \min_I(v(a_I) + v(b^I)),$$

and thus

$$\mathrm{rv}(p(b)) = \bigoplus_I \mathrm{rv}(a_I b^I) = \bigoplus_I \mathrm{rv}(a_I)\,\mathrm{rv}(b)^I,$$

so

$$\mathrm{rv}(p(b')) = \bigoplus_I \mathrm{rv}(a_I)\,\mathrm{rv}(b')^I = \bigoplus_I \mathrm{rv}(a_I (b')^I)$$

is well-defined, i.e. $v(p(b')) = \min_I(v(a_I) + v((b')^I))$. $\square$

Auxiliary surjectivity. In general, it is not clear if one can ensure that if E has inversive residue field, then $E(a)_{\sigma,\sigma^{-1}}$ also does. In particular, the residue field might turn non-inversive after adding a generic to E via PROPOSITION 4.4.22.

Example 4.4.26. Let $(E, v, \sigma) \subseteq (K, v, \sigma)$ be valued difference fields, with σ_{val} ω-increasing, i.e. for all $\gamma > 0$ and $n > 0$, $\sigma_{\mathrm{val}}(\gamma) > n\gamma$. Assume further that E is not inversive, but k_E is. Let $b \in K$ be generic over E, and assume that there is $c \in E \backslash \sigma(E)$ such that, for some $\ell \geqslant 0$, $v(c) = \sigma^{\ell}_{\mathrm{val}}(v(b))$. Without loss of generality, take $\ell = 0$, i.e. $v(c) = v(b)$. Then, $\frac{b}{c}$ is still a generic, thus $E(b)_\sigma = E(\frac{b}{c})_\sigma$ has residue difference field $k_E(\mathrm{res}\left(\frac{b}{c}\right))_\sigma$, generated by an element σ-transcendental over k_E; in particular, it is non-inversive. We now argue that $E(b)_{\sigma,\sigma^{-1}}$ also has non-inversive residue difference field. Indeed, consider $E(b)_\sigma(\sigma^{-1}(b))$: since σ_{val} is ω-increasing, this is an extension by a transcendental element with value outside of the divisible hull of $\Gamma_{E(b)_\sigma}$. In particular, the residue field does not change. We can see, by induction, that $E(b)_{\sigma,\sigma^{-1}}$ has residue field given by $k_E(\mathrm{res}\left(\frac{b}{c}\right))_\sigma$, which is non-inversive.

This issue is fixed by moving to a residually inversive hull, as we explain now.

Lemma 4.4.27 ($\Diamond$)**.** *Let (K, v, σ) be a weakly σ-henselian valued difference field that is also a model of* FE*. Let (F, v, σ) be a valued difference subfield such that $(F, \sigma) \subseteq (K, \sigma)$ is σ-separable*[9]*, and $(k_F, \sigma_{\mathrm{res}})$ is not necessarily inversive. Then, there is an extension $F \subseteq F' \subseteq K$ with the following properties:*

1. *$F \subseteq F'$ is σ-separably σ-algebraic,*
2. *F' is a model of* FE*, and thus $F' \subseteq K$ is σ-separable,*
3. *the residue difference field of F' is k_F^{inv},*
4. *if (L, v, σ) is another σ-separable extension of K which is a model of* FE *and is weakly σ-henselian, then F' embeds into L over K so that L is σ-separable over the image*[10]*.*

If, further, $(\Gamma_F, \sigma_{\mathrm{val}})$ is inversive, then $\Gamma_F = \Gamma_{F'}$.

Proof. We build F' recursively, alongside its F-embedding into L. Both (K, v, σ) and (L, v, σ) are in particular henselian, so we might replace F with $\mathrm{FE}(F^h)$ and assume that F is henselian as well. Note that $F \subseteq F^h$ is an algebraic extension, in particular because F is a model of FE, it is (σ-)separable, thus the same holds for $F \subseteq \mathrm{FE}(F^h)$ by LEMMA 4.2.19. Then, $\mathrm{FE}(F^h) \subseteq K$ is σ-separable. If $(\Gamma_F, \sigma_{\mathrm{val}})$ is inversive, then $F \subseteq \mathrm{FE}(F^h)$ is an unramified extension, since $\mathrm{FE}(F^h) \subseteq (F^h)^{\mathrm{inv}}$.

We may thus assume that (F, v) is henselian and is a model of FE. Then, by [DH23, Lemma 5.9], $(k_F, \sigma_{\mathrm{res}})$ is a model of FE as well. In particular, every $\alpha \in k_F \backslash \sigma_{\mathrm{res}}(k_F)$ is transcendental over $\sigma_{\mathrm{res}}(k_F)$.

Let $\alpha \in k_F$ be such that $\alpha \notin \sigma_{\mathrm{res}}(k_F)$. Given any $\epsilon \in \mathfrak{m}_F$ and any $a \in \mathcal{O}_F$ such that $\mathrm{res}(a) = \alpha$, consider $f(X) = \sigma(X) - \epsilon X - a$. By LEMMA 4.3.11, there is $b \in \mathcal{O}_K$ such that $\sigma(b) - \epsilon b - a = 0$ and so $\sigma_{\mathrm{res}}(\mathrm{res}(b)) = \alpha$. Note that $\sigma(X) - \alpha$ is of minimal complexity for $\mathrm{res}(b)$ over k_F, thus by repeating the same argument symmetrically in L we obtain, by FACT 4.4.19, an embedding $F(b)_{\sigma,\sigma^{-1}} \to L$ such that L is σ-separable over its image. Note that b is transcendental over F, and moreover $F(b)_{\sigma,\sigma^{-1}} = F(b)$, thus $F(b)_{\sigma,\sigma^{-1}}$ is a model of FE (as $\sigma(F(b)_{\sigma,\sigma^{-1}}) = \sigma(F)(\epsilon b - a) \subseteq F(\epsilon b - a) = F(b)$ is a regular extension). As $F \subseteq F(b)_{\sigma,\sigma^{-1}}$ is σ-separably σ-algebraic, it follows that $F(b)_{\sigma,\sigma^{-1}} \subseteq K$ is still σ-separable. Moreover, as $k_F \subseteq k_F(\mathrm{res}(b))$ and $F \subseteq F(b)$ are both purely transcendental extensions, the valuation is uniquely determined to be the Gauss valuation, and hence $\Gamma_{F(b)} = \Gamma_F$.

Upon replacing $F(b)_{\sigma,\sigma^{-1}}$ with $\mathrm{FE}(F(b)^h_{\sigma,\sigma^{-1}})$ again, we may restart the process, whose limit is the required F' together with an F-embedding into L. $\square$

Remark 4.4.28. Note that if F satisfies ASSUMPTION 4.2.3, then F^h is already FE.

Definition 4.4.29. We call F' as in LEMMA 4.4.27 a *residually inversive hull* of F.

[9] Note that then (F, σ) is a model of FE by LEMMA 4.2.15.

[10] Note that the embedding in question is not necessarily unique.

We now have all the tools necessary to prove the embedding lemma, and then deduce relative quantifier elimination and its consequences.

The embedding lemma. Unless otherwise stated, when we write $A \leqslant K$ we mean that A is an $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-substructure of K (i.e., if A is a valued difference subfield, that the extension $(A,\sigma) \subseteq (K,\sigma)$ is σ-separable). When we write an embedding $f\colon A \to L$, we mean that it is an $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-embedding (i.e., if A is a valued difference subfield, that L is σ-separable over $f(A)$), and we denote by f the embedding on $\mathbf{VF}(A)$, by $f_{\mathbf{RV}}$ the induced embedding on $\mathbf{RV}(A)$, and by $f_{\mathbf{VG}}$ the induced embedding on $\mathbf{VG}(A)$. If $\theta_{\mathrm{RV}}\colon \mathrm{RV}_A \to \mathrm{RV}_L$ and $\theta_{\Gamma}\colon \Gamma_A \to \Gamma_L$ are embeddings, we call them *compatible* if the diagram

$$\begin{tikzcd}
\mathbf{RV}(A) \arrow[r, "v_{\mathrm{rv}}"] \arrow[d, "\theta_{\mathrm{RV}}"'] & \mathbf{VG}(A) \arrow[d, "\theta_{\Gamma}"] \\
\mathbf{RV}(L) \arrow[r, "v_{\mathrm{rv}}"'] & \mathbf{VG}(L)
\end{tikzcd}$$

commutes.

Lemma 4.5.1 (Embedding lemma). *Let K and L be $(2^{\aleph_0})^+$-saturated models of* hVFE_0.
Let:

1. *$A \leqslant K$ be a countable substructure,*
2. *$f\colon A \hookrightarrow L$ be an embedding,*
3. *$\theta_{\mathrm{RV}}\colon \mathrm{RV}_K \hookrightarrow \mathrm{RV}_L$ and $\theta_{\Gamma}\colon \Gamma_K \hookrightarrow \Gamma_L$ be compatible embeddings over $\mathbf{RV}(A)$ and $\mathbf{VG}(A)$, extending $f_{\mathbf{RV}}$ and $f_{\mathbf{VG}}$.*

Then, for every $a \in K$, there is a substructure $A \leqslant A' \leqslant K$ together with an embedding $g\colon A' \hookrightarrow L$ extending f such that $a \in A'$, $g_{\mathbf{VG}} = \theta_{\Gamma}|_{\mathbf{VG}(A')}$, and $g_{\mathbf{RV}} = \theta_{\mathrm{RV}}|_{\mathbf{RV}(A')}$.

Remark 4.5.2. Note that this does not immediately yield that we can lift θ_{RV} and θ_{Γ} to an embedding of K itself in L; for this, K would need to be of size $(2^{\aleph_0})^+$, i.e. saturated in its own cardinality, which could be achieved for example by assuming the Continuum Hypothesis (this would anyway be a safe assumption, see [HK23]).

Remark 4.5.3. The embedding lemma is phrased in a bit of an odd way, to allow space for both relative quantifier elimination and proposition 4.5.16. As explained in the previous remark, this falls short of a purely algebraic embedding lemma due to saturation issues, essentially injected into the statement by the use of generics. One should think of applying this to an asymmetric back-and-forth scenario, where L is saturated in the cardinality of K, $f_{\mathbf{RV}}$ and $f_{\mathbf{VG}}$ are elementary, and thus they then

$$
\begin{tikzcd}
(\mathrm{RV}_K,\Gamma_K) \arrow[rr, "(\theta_{\mathrm{RV}},\theta_\Gamma)"] \arrow[d, dash] \arrow[ddr, dash] & & (\mathrm{RV}_L,\Gamma_L) \\
(\mathbf{RV}(A'),\mathbf{VG}(A')) \arrow[urr, hook, "(g_{\mathbf{RV}}|_{\mathbf{RV}(A')},g_{\mathbf{VG}}|_{\mathbf{VG}(A')})"] \arrow[dr, dash] & & \\
& (\mathrm{RV}_A,\Gamma_A) \arrow[uur, hook, "(f_{\mathbf{RV}},f_{\mathbf{VG}})"'] &
\end{tikzcd}
$$

$$
\begin{tikzcd}
K \arrow[d, dash] & & L \\
A' \arrow[urr, hook, "g"] \arrow[dr, dash] & & \\
& A(a)_\sigma \arrow[d, dash] & \\
& A \arrow[uuur, hook, "f"'] &
\end{tikzcd}
$$

Figure 4.1: The red elements of the diagram are the ones produced by LEMMA 4.5.1.

be extended by saturation and elementarity to θ_{RV} and θ_Γ in a compatible way. The embedding lemma then gives a recipe to extend f to any element of K which is not in A already, possibly extending the residue field along the way. In fact, this remark is a spoiler: this is precisely how the theorem will be used in THEOREM 4.5.4.

Proof of LEMMA 4.5.1. We proceed in steps, extending f recursively and checking at each stage that we still obtain an embedding as requested.

Warning. At each stage, we rename the newly obtained intermediate substructure as A again, and the extended embedding as f again, to save on notation. From the moment we can assume that we are working with subfields onwards, we check that $A \preccurlyeq K$ by checking that K is σ-separable over A; similarly, we check that a valued difference field embedding $f\colon A \to L$ is an embedding in $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$ by checking that L is σ-separable over $f(A)$. We also point out that, in this setting, the towers of extension $K_{\mathrm{inv}} \subseteq K \subseteq K^{\mathrm{inv}}$ and $L_{\mathrm{inv}} \subseteq L \subseteq L^{\mathrm{inv}}$ are all dense, and in particular immediate. We make use of this fact several times without further mention in the proof.

STEP 0. *We may assume that A is a field.*
By LEMMA 4.4.11, f extends uniquely to $\mathrm{Frac}(A)$. $\blacksquare_0$

STEP 1. *We may assume that A satisfies* FE.
By LEMMA 4.2.15, since K is σ-separable over A and satisfies FE, A must be a model of FE as well. $\blacksquare_1$

We are now working with a valued difference subfield (A, v, σ). Let A' be the smallest $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-substructure of A containing A and a. Note that, a priori, A' will not be equal to $A(a)_\sigma$, unless the extension $A(a)_\sigma \subseteq K$ happens to σ-separable.

STEP 2. *We may assume that for every $z \in A'$, $v(z) \in \Gamma_A$.*
Let $\gamma \in \Gamma_{A'} \backslash \Gamma_A$, and let $b \in K_{\mathrm{inv}}$ be generic over A with $v(b) = \gamma$ (such b exists by PROPOSITION 4.4.22). Let $\alpha = \mathrm{rv}(b)$ and $\alpha' = \theta_{\mathrm{RV}}(\alpha) \in \mathrm{RV}_L$. We now let $b' \in L_{\mathrm{inv}}$ be such that $\mathrm{rv}(b') = \alpha'$ (in particular, $v(b') = \theta_\Gamma(v(b))$). Then, by LEMMA 4.4.25, b' is still generic over $f(A)$, thus PROPOSITION 4.4.22 ensures an isomorphism $A(b)_{\sigma,\sigma^{-1}} \to f(A)(b')_{\sigma,\sigma^{-1}}$, extending f and $b \mapsto b'$, where $f(A)(b')_{\sigma,\sigma^{-1}} \subseteq L$ is still σ-separable, and thus this gives rise to an $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-embedding of $A(b)_{\sigma,\sigma^{-1}}$ into L over A. Moreover, since $A(b)_{\sigma,\sigma^{-1}} \subseteq K$ and $A(b')_{\sigma,\sigma^{-1}} \subseteq L$ are σ-separable (as $b \in K_{\mathrm{inv}}$ and $b' \in L_{\mathrm{inv}}$), using LEMMA 4.2.19 and since both (K, σ) and (L, σ) are models of FE, one gets automatically that $A(b)_{\sigma,\sigma^{-1}}$ and $A(b')_{\sigma,\sigma^{-1}}$ also do.
We can then repeat this procedure ω-many times to obtain a new countable substructure A_1 such that whenever $z \in A'$, we have that $v(z) \in \Gamma_{A_1}$. Note that, a priori, we do not have that whenever $z \in A_1(a)_\sigma$, $v(z) \in \Gamma_{A_1}$. Thus, we repeat this procedure countably many times to produce a chain of extensions $A = A_0 \leqslant A_1 \leqslant A_2 \leqslant \cdots$ such that their union, which is again countable and we call A_∞, now satisfies that whenever $z \in A_\infty(a)_\sigma$, $v(z) \in \Gamma_{A_\infty}$. The procedure might have made k_{A_∞} non-inversive, so we apply LEMMA 4.4.27 to replace A_∞ with a residually inversive hull. We rename A_∞ as A. $\blacksquare_2$

STEP 3. *We may assume that for every $z \in A'$, $\mathrm{res}(z) \in k_A$, and that $(k_A, \sigma_{\mathrm{res}})$ is linearly difference closed.*
Identify $k_A \subseteq \mathrm{RV}_A$, $k_L \subseteq \mathrm{RV}_L$, and $k_K \subseteq \mathrm{RV}_K$. Let $\alpha \in k_{A'} \backslash k_A$, and let $\alpha' = \theta_{\mathrm{RV}}(\alpha) \in \mathrm{RV}_L$. We distinguish two sub-cases.
Subcase 3.a: α is σ-transcendental over k_A.
Then, the same is true for α', so if we choose any $b \in K_{\mathrm{inv}}$ with $\mathrm{rv}(b) = \mathrm{res}(b) = \alpha$ and any $b' \in L_{\mathrm{inv}}$ with $\mathrm{rv}(b') = \mathrm{res}(b') = \alpha'$, by FACT 4.4.18 the map $b \mapsto b'$ extends f to an $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$-embedding (since both $A(b)_{\sigma,\sigma^{-1}} \subseteq K$ and $A(b')_{\sigma,\sigma^{-1}} \subseteq L$ are σ-separable, by LEMMA 4.2.20) $f \colon A(b)_{\sigma,\sigma^{-1}} \to L$ over A.
Subcase 3.b: α is σ-algebraic over k_A.
Then, we let $h(X) \in k_A[X]_\sigma$ be of minimal complexity such that $h(\alpha) = 0$. As $(k_A, \sigma_{\mathrm{res}})$ is inversive, we have that $h' \neq 0$ and, moreover, that whenever $h_J \neq 0$, $h_J(\alpha) \neq 0$. We now let $c \in \mathcal{O}_K$ be such that $\mathrm{res}(c) = \alpha$, and let $g(X) \in \mathcal{O}_A[X]_\sigma$ be an exact lift of h. Then, $g' \neq 0$, $v(g_J(c)) = 0$ for all J such that $g_J \neq 0$, and $v(g(c)) > 0$. As K is weakly σ-henselian, we find $b \in K$ with $\mathrm{res}(b) = \alpha$ and $g(b) = 0$. Analogously, on the other side, we find $b' \in L$ with $\mathrm{res}(b') = \alpha'$ and $g(b') = 0$.
By FACT 4.4.19, this gives rise to an isomorphism $A(b)_\sigma \to f(A)(b')_\sigma$, given by $b \mapsto b'$, which extends to an isomorphism $\mathrm{FE}(A(b)_\sigma) \to \mathrm{FE}(f(A)(b')_\sigma)$. As

$A \subseteq A(b)_\sigma$ and $f(A) \subseteq f(A)(b')_\sigma$ are σ-separably σ-algebraic extensions, the same holds for $A \subseteq \mathrm{FE}(A(b)_\sigma)$ and $f(A) \subseteq \mathrm{FE}(f(A)(b')_\sigma)$, and thus by the tower properties of σ-separability, $\mathrm{FE}(A(b)_\sigma) \subseteq K$ and $\mathrm{FE}(f(A)(b')_\sigma) \subseteq L$ are σ-separable extensions. In particular, we get an embedding $f\colon \mathrm{FE}(A(b)_\sigma) \to L$. We replace $\mathrm{FE}(A(b)_\sigma)$ with a residually inversive hull by LEMMA 4.4.27.

We can then repeat this procedure ω-many times to obtain a new countable substructure A_1 such that whenever $z \in \mathcal{O}_{A'}$, we have that $\mathrm{res}(z) \in k_{A_1}$. Since k_K is linearly difference closed, for any linear difference equation over k_A we find a solution in k_K. By repeating Subcase 3.b, we obtain a new countable substructure $\widetilde{A_1}$ such that whenever $z \in \mathcal{O}_{A'}$, $\mathrm{res}(z) \in k_{\widetilde{A_1}}$, and furthermore $(k_{\widetilde{A_1}}, \sigma_{\mathrm{res}})$ is linearly difference closed. Note that, as before, we do not necessarily have that whenever $z \in \mathcal{O}_{\widetilde{A_1}(a)_\sigma}$, $\mathrm{res}(z) \in k_{\widetilde{A_1}}$. Thus, we repeat this procedure countably many times to produce a chain of extensions $A = A_0 \leqslant \widetilde{A_1} \leqslant \widetilde{A_2} \leqslant \cdots$ such that their union, which is again countable and we call $\widetilde{A_\infty}$, now satisfies that whenever $z \in \mathcal{O}_{\widetilde{A_\infty}(a)_\sigma}$, $\mathrm{res}(z) \in k_{\widetilde{A_\infty}}$, and $(k_{\widetilde{A_\infty}}, \sigma_{\mathrm{res}})$ is linearly difference closed. Finally, we rename $\widetilde{A_\infty}$ as A. $\blacksquare_3$

What remains is a valued difference field extension $A \leqslant A'$ which is σ-separable and immediate.

STEP 4. *The immediate case.*
This is now LEMMA 4.4.17.

This extends f to the required embedding g. $\square$

We can now deduce the first and most crucial consequence of LEMMA 4.5.1, namely relative quantifier elimination.

Theorem 4.5.4. hVFE_0 *eliminates quantifiers resplendently relatively to* **RV** *and* **VG***.*

Proof. Work in the Morleyization of hVFE_0 with respect to the sorts **RV** and **VG**. Let K and L be two models, where K is $(2^{\aleph_0})^+$-saturated and L is $|K|^+$-saturated. Let $A \leqslant K$ be a countable substructure, and let $f\colon A \hookrightarrow L$ be an embedding. Since we are working in the Morleyization, $f_{\mathbf{RV}}$ and $f_{\mathbf{VG}}$ are elementary, and thus by saturation we can extend them to compatible $\theta_{\mathrm{RV}}\colon \mathrm{RV}_K \hookrightarrow \mathrm{RV}_L$ and $\theta_\Gamma\colon \Gamma_K \hookrightarrow \Gamma_L$. Now, for every $a \in K$, by LEMMA 4.5.1 we find an embedding $g\colon A' \hookrightarrow L$, for some $\mathcal{L}^\lambda_{3,\mathrm{RV},\sigma}$-substructure A' containing A and a, extending f such that $g_{\mathbf{VG}} = \theta_\Gamma|_{\mathbf{VG}(A(a)_\sigma)}$, and $g_{\mathbf{RV}} = \theta_{\mathrm{RV}}|_{\mathbf{RV}(A(a)_\sigma)}$. $\square$

Remark 4.5.5. In particular, one can work in the reduct of $\mathcal{L}^\lambda_{3,\mathrm{RV},\sigma}$ where we only consider the sorts **VF** and **RV**, see REMARK 4.4.4. Then, hVFE_0 (seen in this reduct) eliminates quantifiers resplendently down to **RV**.

Remark 4.5.6. Looking at the categorical arguments in [Rid17], it seems that a version of THEOREM 4.5.4 for the mixed characteristic unramified case should be a formal consequence; note, however, that there are not many interesting such examples. Indeed, in the standard decomposition of a $\aleph_1$-saturated model we would have a model of hVFE_0 as the residue characteristic zero part, whereas the rank 1 mixed characteristic part would be a Witt-Frobenius, as treated in [BMS07].

Angular components. Given a valued field (K, v), we say that a multiplicative group homomorphism $\mathrm{ac}\colon K^\times \to k_K^\times$ is an *angular component* if for every $u \in \mathcal{O}_K^\times$, $\mathrm{ac}(u) = \mathrm{res}(u)$. If $\sigma \in \mathrm{End}(K, v)$, we say that ac is *$\sigma$-equivariant* if $\sigma_{\mathrm{res}} \circ \mathrm{ac} = \mathrm{ac} \circ \sigma$. In the presence of a σ-equivariant angular component, we can construct a σ-equivariant section $s\colon \mathrm{RV}_K^\times \to k_K^\times$, thus the short exact sequence of $\mathbb{Z}[\sigma_{\mathrm{val}}]$-modules

$$1 \longrightarrow k_K^\times \longrightarrow \mathrm{RV}_K^\times \longrightarrow \Gamma_K \longrightarrow 0$$

splits and $\mathrm{RV}_K^\times \cong k_K^\times \times \Gamma_K$ along the map $\mathrm{rv}(a) \mapsto (\mathrm{ac}(a), v(a))$. This allows us to transfer relative quantifier elimination from $\mathcal{L}^\lambda_{3,\mathrm{RV},\sigma}$ to a more classical three-sorted language with k_K and Γ, at the cost of adding an angular component.

Definition 4.5.7. We let $\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}$ be the three-sorted language with sorts:

1. **VF**, with language $\mathcal{L}^\lambda_{\mathrm{ring},\sigma}$ (see DEFINITION 4.4.5),
2. **RF**, with language $\mathcal{L}_{\mathrm{ring},\sigma} = \{+, \cdot, -, 0, 1, \sigma_{\mathrm{res}}\}$,
3. **VG**, with language $\mathcal{L}_{\mathrm{oag},\sigma} := \{+, \leqslant, 0, \infty, \sigma_{\mathrm{val}}\}$,

and connecting functions $\mathrm{ac}\colon \mathbf{VF} \to \mathbf{RF}$ and $v\colon \mathbf{VF} \to \mathbf{VG}$. We let $\mathrm{hVFE}_0^{\mathrm{ac}}$ be the theory of non-inversive, equicharacteristic zero, weakly σ-henselian valued difference fields that are models of FE, endowed with a σ-equivariant angular component.

We now sketch how to deduce relative quantifier elimination for $\mathrm{hVFE}_0^{\mathrm{ac}}$ in $\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}$ from relative quantifier elimination for hVFE_0 in $\mathcal{L}^\lambda_{3,\mathrm{RV},\sigma}$.

1. Consider the expansion $\mathcal{L}^\lambda_{3,\mathrm{RV},\sigma} \subseteq \mathcal{L}^{\lambda,+}_{3,\mathrm{RV},\sigma}$ where the sort **RV** is enriched with a map $\mathrm{ac}_{\mathrm{rv}}\colon \mathbf{RV} \to \mathbf{RV}$, to be interpreted as the map induced by ac along the diagram

$$\begin{tikzcd} \mathbf{VF}^\times \arrow[rr, "\mathrm{ac}"] \arrow[dr, "\mathrm{rv}"'] & & \mathbf{RF}^\times \\ & \mathbf{RV} \arrow[ur, "\mathrm{ac}_{\mathrm{rv}}"'] & \end{tikzcd}$$

Then, resplendency of the relative quantifier elimination result means that the rephrasing hVFE_0^+ of hVFE_0 in $\mathcal{L}^{\lambda,+}_{3,\mathrm{RV},\sigma}$ still eliminates quantifiers relatively to **RV** and **VG**.

2. We now observe that, in models of hVFE_0^+, $\mathbf{RF}^\times$ and ac are definable in $\mathcal{L}^{\lambda,+}_{3,\mathrm{RV},\sigma}$. Namely, $y \in \mathbf{RF}^\times$ if and only if $v_{\mathrm{rv}}(y) = 0$, and $\mathrm{ac}(x) = z$ if and only if $\mathrm{ac}_{\mathrm{rv}}(\mathrm{rv}(x)) = z$. In the other direction, $\mathbf{RV}$ and $\mathrm{ac}_{\mathrm{rv}}$ are $\mathbf{VF}$-quantifier-free interpretable in $\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}$. Indeed, in the presence of ac one has the identification $\mathbf{RV}^\times = \mathbf{RF}^\times \times \mathbf{VG}$, where $\mathrm{ac}_{\mathrm{rv}}$ is the projection on the first coordinate. Thus, one can transfer relative quantifier elimination between these two languages.

Theorem 4.5.8. $\mathrm{hVFE}_0^{\mathrm{ac}}$ *eliminates quantifiers resplendently relatively to* $\mathbf{RF}$ *and* $\mathbf{VG}$.

Theorem 4.5.9. *Let K be a model of* hVFE_0. *Then $(\mathrm{RV}_K, \sigma_{\mathrm{rv}})$ and $(\Gamma_K, \sigma_{\mathrm{val}})$ are stably embedded, with induced structures given respectively by $\mathcal{L}_{\mathbf{RV}}$ and $\mathcal{L}_{\mathrm{oag},\sigma}$. Analogously, if K is a model of* $\mathrm{hVFE}_0^{\mathrm{ac}}$, *then $(k_K, \sigma_{\mathrm{res}})$ and $(\Gamma_K, \sigma_{\mathrm{val}})$ are stably embedded, with induced structures given respectively by $\mathcal{L}_{\mathrm{ring},\sigma}$ and $\mathcal{L}_{\mathrm{oag},\sigma}$.*

Lifting the residue field. We now turn to an enriched version of $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$. In many applications, it becomes useful to see the residue difference field $(k_K, \sigma_{\mathrm{res}})$ as a difference subfield of the valued difference field (K, v, σ). For example, this is successfully exploited in [HHYZ24]. First, we argue that such lifts exist.

Lemma 4.5.10. *Suppose (K, v, σ) is equicharacteristic zero, weakly σ-henselian, and $\aleph_0$-saturated. Let $F \subseteq k_K$ be an inversive difference subfield, and suppose there exists an embedding $\eta\colon F \hookrightarrow \mathcal{O}^\times_{K_{\mathrm{inv}}}$ such that $\mathrm{res} \circ \eta = \mathrm{id}$. Given any $\alpha \in k_K$, there is an embedding $\tilde{\eta}\colon F(\alpha)_{\sigma,\sigma^{-1}} \hookrightarrow \mathcal{O}^\times_{K_{\mathrm{inv}}}$, extending η and such that $\mathrm{res} \circ \tilde{\eta} = \mathrm{id}$.*

$$\begin{tikzcd} k_K \arrow[r, hook, "\tilde{\eta}"'] & K_{\mathrm{inv}} \arrow[l, bend right, "\mathrm{res}"'] \\ F \arrow[u, hook] \arrow[ur, hook, "\eta"] & \end{tikzcd}$$

Proof. Let $a \in \mathcal{O}^\times_{K_{\mathrm{inv}}}$ be such that $\mathrm{res}(a) = \alpha$. If for all $g(X) \in \eta(F)[X]_\sigma$ we have that $v(g(a)) = 0$, then a is σ-transcendental over $\eta(F)$ and α is σ-transcendental over F, so η extends to a difference field isomorphism $F(\alpha)_{\sigma,\sigma^{-1}} \to \eta(F)(a)_{\sigma,\sigma^{-1}} \subseteq \mathcal{O}^\times_{K_{\mathrm{inv}}}$. Suppose now that there is some $g(X) \in \eta(F)[X]_\sigma$ such that $v(g(a)) > 0$, and suppose that it is of minimal complexity such. Then, $g(X)$ is σ-henselian at a, so by strong σ-henselianity of K_{inv} (LEMMA 4.3.13) we find $b \in \mathcal{O}_{K_{\mathrm{inv}}}$ such that $g(b) = 0$ and $\mathrm{res}(b) = \alpha$. In particular, then, η extends to a difference field isomorphism $F(\alpha)_{\sigma,\sigma^{-1}} \to \eta(F)(b)_{\sigma,\sigma^{-1}} \subseteq \mathcal{O}^\times_{K_{\mathrm{inv}}}$ (see FACT 4.4.19). ☐

Corollary 4.5.11. *Suppose (K, v, σ) is equicharacteristic zero, weakly σ-henselian, and $\aleph_0$-saturated. Then there is a difference subfield $F \subseteq \mathcal{O}^\times_{K_{\mathrm{inv}}}$ such that $\mathrm{res}|_F\colon F \to k_K$ is a difference field isomorphism.*

We say that a map $s\colon \Gamma_K \to K$ is a *section* of the valuation if for all $x \in \Gamma_K$, $v(s(x)) = x$. We say that s is *σ-equivariant* if $s \circ \sigma = \sigma_{\mathrm{val}}^{-1} \circ s$.

Remark 4.5.12. Note that if s is a σ-equivariant section of the valuation, then

$$\mathrm{ac}(x) := \mathrm{res}\left(\frac{x}{s(v(x))}\right)$$

is a σ-equivariant angular component.

Definition 4.5.13. We let $\mathcal{L}_{3,\mathrm{RV},\sigma}^{\lambda,s,\iota}$ be the three-sorted language with sorts:

1. **VF**, with language $\mathcal{L}_{\mathrm{ring},\sigma}^{\lambda}$ (see DEFINITION 4.4.5),
2. **RF**, with language $\mathcal{L}_{\mathrm{ring},\sigma} = \{+,\cdot,-,0,1,\sigma_{\mathrm{res}}\}$,
3. **VG**, with language $\mathcal{L}_{\mathrm{oag},\sigma} := \{+,\leqslant,0,\infty,\sigma_{\mathrm{val}}\}$,

and connecting functions $s\colon \mathbf{VG} \to \mathbf{VF}$, $v\colon \mathbf{VF} \to \mathbf{VG}$, and $\iota\colon \mathbf{RF} \to \mathbf{VF}$. We let $\mathrm{hVFE}_0^{s,\iota}$ be the theory of non-inversive, equicharacteristic zero, weakly σ-henselian valued difference fields that are models of FE, and where s is interpreted as a σ-equivariant section of the valuation, and ι as an embedding of **RF** in **VF**, both with values in the inversive core of the model.

We now explain how to adapt the proof of LEMMA 4.5.1 to obtain relative quantifier elimination in this language; the argument is essentially as in [Kes24, Theorem 4.11], but we include a sketch of the proof for the convenience of the reader. Note that the proof is a direct adaptation of the proof for the case without σ, as for example proven in [vdD14, Theorem 5.11].

Theorem 4.5.14. $\mathrm{hVFE}_0^{s,\iota}$ *eliminates quantifiers resplendently relatively to* **RF** *and* **VG**.

Proof. We sketch out how to prove the following (cf. LEMMA 4.5.1):

> *Let K and L be $(2^{\aleph_0})^+$-saturated models of* $\mathrm{hVFE}_0^{s,\iota}$. *Let:*
>
> 1. *$A \leqslant K$ be a countable substructure,*
> 2. *$f\colon A \hookrightarrow L$ be an embedding,*
> 3. *$\theta\colon k_K \hookrightarrow k_L$ and $\eta\colon \Gamma_K \hookrightarrow \Gamma_L$ be compatible embeddings over* $\mathbf{RF}(A)$ *and* $\mathbf{VG}(A)$, *extending $f_{\mathbf{RF}}$ and $f_{\mathbf{VG}}$.*
>
> *Then, for every $a \in K$, there is a countable substructure $A \leqslant A' \leqslant K$ together with an embedding $g\colon A' \hookrightarrow L$ extending f such that $a \in A'$, $g_{\mathbf{VG}} = \eta|_{\mathbf{VG}(A')}$, and $g_{\mathbf{RF}} = \theta|_{\mathbf{RF}(A')}$.*

Arguing as in Step 0 and Step 1 of LEMMA 4.5.1, we may assume that A is a subfield that satisfies FE. We now show how to adapt Step 2 and Step 3.

Claim 1: (cf. [Kes24, Lemma 4.9]) *Suppose $\gamma \in \Gamma_K \backslash \Gamma_A$. Then, there is an isomorphism of valued difference fields $A(s(\gamma))_{\sigma,\sigma^{-1}} \to A(s(\eta(\gamma)))_{\sigma,\sigma^{-1}}$, and $\Gamma_{A(\iota(\alpha))_{\sigma,\sigma^{-1}}} = \Gamma_A$.*

Indeed, the extension $A \subseteq A(s(\gamma))$ of valued fields with lift and section is uniquely determined: if $\gamma \in \mathbb{Q} \otimes \Gamma_A$, then the fundamental inequality guarantees that the extension is unique up to isomorphism over A; if $\gamma \notin \mathbb{Q} \otimes \Gamma_A$, then $s(\gamma)$ is transcendental over A and so the extension is again unique. By induction the same is true for $A \subseteq A(s(\gamma))_{\sigma,\sigma^{-1}}$. $\blacksquare_1$

Claim 2: (cf. [Kes24, Lemma 4.10]) *Suppose $\alpha \in k_K \backslash k_A$. Then, there is an isomorphism of valued difference fields $A(\iota(\alpha))_{\sigma,\sigma^{-1}} \to A(\iota(\theta(\alpha)))_{\sigma,\sigma^{-1}}$, and $k_{A(\iota(\alpha))_{\sigma,\sigma^{-1}}} = k_A$.*

Indeed, the extension $A \subseteq A(\iota(\alpha))$ of valued fields with lift and section is uniquely determined: if α is algebraic over k_A, then the fundamental inequality guarantees that the extension is unique up to isomorphism over A; if α is transcendental over k_A, then $\iota(\alpha)$ is transcendental of valuation zero, and so the extension is again unique. By induction the same is true for $A \subseteq A(\iota(\alpha))_{\sigma,\sigma^{-1}}$. $\blacksquare_2$

One can then run the proof of LEMMA 4.5.1; note that the presence of the section avoids the use of generics, and that Step 4 translates verbatim. $\square$

Ax-Kochen/Ershov. We now establish transfer theorems for hVFE_0, $\mathrm{hVFE}_0^{\mathrm{ac}}$, and $\mathrm{hVFE}_0^{s,\iota}$, using THEOREM 4.5.4, THEOREM 4.5.8, and THEOREM 4.5.14.

Theorem 4.5.15. *Let K and L be models of hVFE_0. Then,*

$$K \equiv L \iff (\mathrm{RV}_K, \sigma_{\mathrm{rv}}) \equiv (\mathrm{RV}_L, \sigma_{\mathrm{rv}}).$$

If $K \leqslant L$, then

$$K \preceq L \iff (\mathrm{RV}_K, \sigma_{\mathrm{rv}}) \preceq (\mathrm{RV}_L, \sigma_{\mathrm{rv}}).$$

Analogously, let K and L be models of $\mathrm{hVFE}_0^{\mathrm{ac}}$ or $\mathrm{hVFE}_0^{s,\iota}$. Then,

$$K \equiv L \iff ((k_K, \sigma_{\mathrm{res}}) \equiv (k_L, \sigma_{\mathrm{res}}) \text{ and } (\Gamma_K, \sigma_{\mathrm{val}}) \equiv (\Gamma_L, \sigma_{\mathrm{val}})).$$

If $K \leqslant L$, then

$$K \preceq L \iff ((k_K, \sigma_{\mathrm{res}}) \preceq (k_L, \sigma_{\mathrm{res}}) \text{ and } (\Gamma_K, \sigma_{\mathrm{val}}) \preceq (\Gamma_L, \sigma_{\mathrm{val}})).$$

Proof. The proofs are agnostic to which one of the three languages we work in.

($\equiv$): this follows from THEOREM 4.5.4, since $(\mathbb{Q}, v_{\mathrm{triv}}, \mathrm{id}_{\mathbb{Q}})$ is a common substructure between any two models of hVFE_0.

($\preceq$): This is a direct consequence of relative quantifier elimination. $\square$

Proposition 4.5.16. *Let K and L be two models of hVFE_0 and $A \leqslant K, L$ be a common substructure. Then,*

$$\mathrm{Th}^{\exists}{}_{\mathcal{L}_{\mathbf{RV},\mathbf{VG}}(\mathbf{RV}(A) \cup \mathbf{VG}(A))}(\mathrm{RV}_K, \Gamma_K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}_{\mathbf{RV},\mathbf{VG}}(\mathbf{RV}(A) \cup \mathbf{VG}(A))}(\mathrm{RV}_L, \Gamma_L)$$

if and only if

$$\mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}(A)}(K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}(A)}(L).$$

Analogously, let K and L be models of $\mathrm{hVFE}_0^{\mathrm{ac}}$ *or* $\mathrm{hVFE}_0^{s,\iota}$*. Then,*

$$\mathrm{Th}^{\exists}{}_{\mathcal{L}_{\mathbf{RV},\mathbf{VG}}(\mathbf{RF}(A)\cup\mathbf{VG}(A))}(k_K,\Gamma_K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}_{\mathbf{RV},\mathbf{VG}}(\mathbf{RF}(A)\cup\mathbf{VG}(A))}(k_L,\Gamma_L)$$

if and only if

$$\mathrm{Th}^{\exists}{}_{\mathcal{L}(A)}(K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}(A)}(L)$$

for the appropriate $\mathcal{L} \in \{\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}, \mathcal{L}^{\lambda,s,\iota}_{3,\mathrm{RV},\sigma}\}$.

Proof. We give the argument for $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$. One direction is clear. For the other, without loss of generality we may assume that K and L are $(2^{\aleph_0})^+$-saturated and A is countable. Suppose that $\mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}(A)}(K) \nsubseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}(A)}(L)$. Let this be witnessed by some quantifier-free formula $\varphi(X)$ with parameters in A, and let $b \in K$ be such that $K \models \varphi(b)$, but $L \not\models \exists X\varphi(X)$. By assumption, we have compatible embeddings $\theta_{\mathrm{RV}}\colon \mathrm{RV}_K \to \mathrm{RV}_L$ and $\theta_\Gamma\colon \Gamma_K \to \Gamma_L$. Then, LEMMA 4.5.1 yields an embedding $g\colon C \to L$ of some countable substructure $A \leqslant C \leqslant L$ containing b, so $L \models \varphi(g(b))$. This is a contradiction, since then in particular $L \models \exists X\varphi(X)$. □

This implies two more recognizable results.

Theorem 4.5.17. *Let K and L be two models of* hVFE_0*, with $K \leqslant L$. Then, $K \preceq_\exists L$ if and only if $(\mathrm{RV}_K, \sigma_{\mathrm{rv}}) \preceq_\exists (\mathrm{RV}_L, \sigma_{\mathrm{rv}})$. Analogously, let K and L be two models of* $\mathrm{hVFE}_0^{\mathrm{ac}}$ *or* $\mathrm{hVFE}_0^{s,\iota}$*, with $K \leqslant L$. Then, $K \preceq_\exists L$ if and only if $(k_K, \sigma_{\mathrm{res}}) \preceq_\exists (k_L, \sigma_{\mathrm{res}})$ and $(\Gamma_K, \sigma_{\mathrm{val}}) \preceq_\exists (\Gamma_L, \sigma_{\mathrm{val}})$.*

Proof. One direction is clear. To prove the other, use PROPOSITION 4.5.16 with $A = K$. □

Theorem 4.5.18. *Let K and L be two models of* hVFE_0*. Then, $\mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}}(K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}}(L)$ if and only if $\mathrm{Th}^{\exists}(\mathrm{RV}_K, \sigma_{\mathrm{rv}}) \subseteq \mathrm{Th}^{\exists}(\mathrm{RV}_L, \sigma_{\mathrm{rv}})$. Analogously, let K and L be two models of* $\mathrm{hVFE}_0^{\mathrm{ac}}$ *or* $\mathrm{hVFE}_0^{s,\iota}$*. Then, $\mathrm{Th}^{\exists}{}_{\mathcal{L}}(K) \subseteq \mathrm{Th}^{\exists}{}_{\mathcal{L}}(L)$ if and only if $\mathrm{Th}^{\exists}(\Gamma_K, \sigma_{\mathrm{val}}) \subseteq \mathrm{Th}^{\exists}(\Gamma_L, \sigma_{\mathrm{val}})$ and $\mathrm{Th}^{\exists}(k_K, \sigma_{\mathrm{res}}) \subseteq \mathrm{Th}^{\exists}(k_L, \sigma_{\mathrm{res}})$, for the appropriate $\mathcal{L} \in \{\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}, \mathcal{L}^{\lambda,s,\iota}_{3,\mathrm{RV},\sigma}\}$.*

Proof. One direction is clear. To prove the other, use PROPOSITION 4.5.16 with $A = \mathbb{Q}$. □

NTP_2**.** We follow the strategy developed by Chernikov and Hils in [CH14]. We first show that dense pairs of algebraically closed valued fields are NIP, a folklore fact for which we could not find a proof in the literature; we include one for the convenience of the reader, using the same strategy deployed by Delon in [Del79] to show that separably closed, non-trivially valued, valued fields are NIP.

Definition 4.5.19. We let $\mathcal{L}_{3,\mathrm{ac}}$ be the three-sorted Denef-Pas language:

1. **VF**, with language $\{+,\cdot,-,0,1\}$,
2. **RF**, with language $\{+,\cdot,-,0,1\}$,
3. **VG**, with language $\{+,\leqslant,0,\infty\}$,

and connecting functions $\mathrm{ac}\colon \mathbf{VF}\to\mathbf{RF}$ and $v\colon \mathbf{VF}\to\mathbf{VG}$. We let $\mathcal{L}^*_{3,\mathrm{ac}}$ be an expansion of $\mathcal{L}_{3,\mathrm{ac}}$ where, on the sort **VF**, we add a unary predicate $P(x)$, and countably many functions $(\lambda^i_n)_{n\in\mathbb{N},1\leqslant i\leqslant n}$, where each λ^i_n is $(n+1)$-ary. Denote by $\mathcal{L}^\lambda_{\mathrm{ring,P}}$ the resulting language given by $\{+,\cdot,-,0,1\}\cup\{P(x)\}\cup\{\lambda^i_n \mid n\in\mathbb{N},1\leqslant i\leqslant n\}$.

Definition 4.5.20. We let PPairs be the $\mathcal{L}^*_{3,\mathrm{ac}}$-theory of proper pairs $(K,P(K))$ of valued fields, where $P(K)$ is algebraically closed and the λ-functions are interpreted as giving the coefficients for P-linear independence. We let $\mathrm{ACVF}_{\mathrm{dense}}$ be the $\mathcal{L}^*_{3,\mathrm{ac}}$-theory of proper dense pairs of algebraically closed valued fields.

Remark 4.5.21. Note that, when we interpret a valued difference field (K,v,σ) as an $\mathcal{L}^*_{3,\mathrm{ac}}$-structure with $P(K)=\sigma(K)$, the λ-functions from $\mathcal{L}^*_{3,\mathrm{ac}}$ do not coincide with the λ-functions introduced in DEFINITION 4.4.6. Namely, the λ-functions from $\mathcal{L}^*_{3,\mathrm{ac}}$ output the coefficient of a $P(K)$-linear combination, an element of $P(K)=\sigma(K)$, whereas the λ-functions in DEFINITION 4.4.6 output σ^{-1} of such a coefficient, so an element of K. The functions are however interdefinable (and in fact, both 0-definable in the respective reducts).

Proposition 4.5.22. $\mathrm{ACVF}_{\mathrm{dense}}$ *is* NIP.

Proof. Let $(M,P(M))\models\mathrm{ACVF}_{\mathrm{dense}}$ and let $(M,P(M))\preceq(\mathfrak{U},P(\mathfrak{U}))$ be a monster model. We show that for any $p\in S(M)$ there are at most $2^{|M|}$ global coheirs of p, i.e. $p\subseteq q\in S(\mathfrak{U})$ with q finitely satisfiable in M.

Indeed, take any $a\models p$ and let $N\succeq M$ be such that $a\in N$ and $|N|=|M|$. It is enough to show that $p'=\mathrm{tp}(N/M)$ has at most $2^{|M|}$ global coheirs.

Since pairs of algebraically closed fields are stable, $\mathrm{tp}_{\mathcal{L}^\lambda_{\mathrm{ring,P}}}(N/M)$ has a unique global coheir, call it r. If $N'\models r$, then N' is linearly disjoint from $\mathfrak{U}$ over M, and the λ-functions are uniquely determined on the compositum $N'\mathfrak{U}$. In other words, the $\mathcal{L}^\lambda_{\mathrm{ring,P}}$-structure on $N'\mathfrak{U}$ is uniquely determined, and the number of global coheirs of $\mathrm{tp}(N/M)$ is bounded by the number of global coheirs of the ACVF-type of N over M. Since ACVF is NIP by [Del79], there are at most $2^{|N|}=2^{|M|}$ many such. □

Lemma 4.5.23. *Any model of* PPairs *embeds in a model of* $\mathrm{ACVF}_{\mathrm{dense}}$.

Proof. Let $(K,P(K))$ be a model of PPairs. Without loss of generality, we may assume that K is algebraically closed, and thus so is $P(K)$. We can now make $P(K)$ dense in K by a chain construction: for every ball $B\subseteq K$ with $B\cap P(K)=\varnothing$, let t be a generic of B in the predicate. In particular, $K\subseteq K(t)$ and $P(K)\subseteq P(K)(t)$ are Gauss extensions. Then, $(K,P(K))\subseteq(K(t),P(K)(t))$ is an extension in $\mathcal{L}^*_{3,\mathrm{ac}}$: indeed, as $P(K)$ is algebraically closed, it is enough to check that $P(K)(t)$ is algebraically free from K over $P(K)$, which is clear. We then once again move to the algebraic closures. By iterating this construction, the limit is a model of $\mathrm{ACVF}_{\mathrm{dense}}$. □

Lemma 4.5.24. *Let $\varphi(x,y)$ be a quantifier-free $\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}$-formula. Then, modulo* VFE_0*, it is* NIP.

Proof. We proceed as in [CH14, Lemma 4.2]. We can write $\varphi(x,y)$ as

$$\psi(x,\sigma(x),\ldots\sigma^n(x),y,\sigma(y),\ldots\sigma^n(y)),$$

where $\psi(x_0,x_1,\ldots x_n,y_0,y_1,\ldots y_n)$ is a quantifier-free $\mathcal{L}^*_{3,\mathrm{ac}}$-formula.

Without loss of generality, we may assume that we are working with a non-surjective endomorphism; otherwise, the result is [CH14, Lemma 4.2]. Now, by LEMMA 4.5.23, to check that $\psi(\overline{x},\overline{y})$ is NIP in PPairs, it is enough to check if $\psi(\overline{x},\overline{y})$ is NIP in $\mathrm{ACVF}_{\mathrm{dense}}$. We conclude by PROPOSITION 4.5.22. □

Theorem 4.5.25. *Suppose K is a model of* $\mathrm{hVFE}_0^{\mathrm{ac}}$*. Then, K is* NTP_2 *in $\mathcal{L}^{\lambda,\mathrm{ac}}_{3,\mathrm{RV},\sigma}$ if and only if $(k_K,\sigma_{\mathrm{res}})$ is* NTP_2 *in $\mathcal{L}_{\mathrm{ring},\sigma}$ and $(\Gamma_K,\sigma_{\mathrm{val}})$ is* NTP_2 *in $\mathcal{L}_{\mathrm{oag},\sigma}$.*

Proof. The proof of [CH14, Theorem 4.1] can be followed verbatim, once one replaces their Lemma 4.2 with LEMMA 4.5.24. Note that while they work in the multiplicative case, this is really not necessary, as the same results hold in the general case by [DO15]. □

One can prove the same result in $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$ and $\mathcal{L}^{\lambda,s,\iota}_{3,\mathrm{RV},\sigma}$:

Theorem 4.5.26. *Suppose K is a model of* hVFE_0*. Then, K is* NTP_2 *in $\mathcal{L}^{\lambda}_{3,\mathrm{RV},\sigma}$ if and only if $(\mathrm{RV}_K,\sigma_{\mathrm{rv}})$ is* NTP_2 *in $\mathcal{L}_{\mathbf{RV}}$.*

Theorem 4.5.27. *Suppose K is a model of* $\mathrm{hVFE}_0^{s,\iota}$*. Then, K is* NTP_2 *in $\mathcal{L}^{\lambda,s,\iota}_{3,\mathrm{RV},\sigma}$ if and only if $(k_K,\sigma_{\mathrm{res}})$ is* NTP_2 *in $\mathcal{L}_{\mathrm{ring},\sigma}$ and $(\Gamma_K,\sigma_{\mathrm{val}})$ is* NTP_2 *in $\mathcal{L}_{\mathrm{oag},\sigma}$.*

Remark 4.5.28. Adapting appropriately the techniques of [DH23], one can also characterize forking and dividing, and give a different proof of NTP_2 transfer.

4.6 | A (DE)MOTIVATING (COUNTER)EXAMPLE

We briefly discuss the case where we drop the assumption of surjectivity on σ_{res}. For the sake of simplicity, we work with $\sigma_{\mathrm{val}} = \mathrm{id}_{\Gamma_K}$. As previously said, the key to an embedding theorem like LEMMA 4.5.1 is the uniqueness of certain maximal immediate extensions, i.e. THEOREM 4.3.21; in the case of valued difference fields where σ_{res} and σ_{val} are inversive (and σ_{res} is linearly difference closed) we obtain uniqueness for the σ-separably σ-algebraically maximal immediate ones. One could hope to obtain the same when the assumptions on σ_{res} and σ_{val} are dropped. We construct an explicit counterexample.

Remark 4.6.1. In [Azg10, Example 5.11], a counterexample is built that shows how uniqueness of the σ-algebraically maximal σ-algebraic immediate extension fails when σ_{res} is not linearly difference closed (which is in particularly true when σ_{res} is not surjective). One could then think that uniqueness of the σ-algebraically maximal σ-algebraic immediate extension might hold under the assumption that $(k_K, \sigma_{\mathrm{res}})$ is weakly linearly difference closed (DEFINITION 4.4.12). We thus build a counterexample that does not rely on a failure of being (weakly) linearly difference closed, but rather on a failure of σ_{res} being surjective.

Consider a non-inversive difference field $(k, \sigma_{\mathrm{res}})$ of characteristic zero such that $k = k_0(t)_\sigma$, for t σ-transcendental over k_0. We first build an isometric valued difference field (F_3, v, σ) with residue difference field $(k, \sigma_{\mathrm{res}})$ and with a σ-algebraic immediate extension generated by an element α that solves an equation of the form $\sigma(X) - \varepsilon X - c = 0$. We will then show that F_3 has no unique maximal immediate extension.

1. Start with $(k, \sigma_{\mathrm{res}})$ aperiodic, with $k = k_0(t)_\sigma$ and t σ-transcendental over k_0. Let $\Gamma := \mathbb{Z}^2$ and $F := k((\Gamma))$, endowed with the t-adic valuation v and with the lift σ of σ_{res} acting on coefficients (as defined in PROPOSITION 4.4.14). There is $F_0 \subseteq \mathcal{O}_F^\times$ such that, along the residue map, $(F_0, \sigma|_{F_0}) \cong (k_0, \sigma_{\mathrm{res}}|_{k_0})$. Choose $\varepsilon \in \mathrm{Fix}(K)$ with $v(\varepsilon) = (0, 1)$ and consider $F_1 := F_0(\varepsilon)$. This has residue difference field $(k_0, \sigma_{\mathrm{res}}|_{k_0})$ and value group given by the convex subgroup $\mathbb{Z} \equiv \{0\} \times \mathbb{Z} \leqslant \Gamma$.
2. Next, pick $\tau \in \mathcal{O}_F$ such that $(F_0(\tau)_\sigma, \sigma|_{F_0(\tau)_\sigma}) \cong (k, \sigma_{\mathrm{res}})$ along the residue map, and consider $\alpha := \frac{1}{\varepsilon}\tau$. Since τ is σ-transcendental over F_1, the same is true for α, and thus the same is also true for $c := \sigma(\alpha) - \varepsilon\alpha \in F$. Thus, we can consider $F_2 := F_1(c)_\sigma^h \subseteq F$, which has value group $\mathbb{Z}$ and residue difference field $(k, \sigma_{\mathrm{res}})$.
3. Since α is σ-algebraic over F_2, but not algebraic, and $F_2 \subseteq F_2(\alpha)_\sigma$ is an immediate extension, we can choose a pseudo-Cauchy sequence $(a_\rho)_{\rho<\lambda} \subseteq F_2$ witnessing this: namely $a_\rho \rightsquigarrow \alpha$, it has no pseudolimit in F_2, and it is of σ-algebraic type over F_2, say witnessed by $p_1(X) := \sigma(X) - \varepsilon X - c$. We can compute $v(p_1(a_\rho)) = v(p_1(\alpha) - p_1(a_\rho)) = v(\alpha - a_\rho) = \gamma_\rho$, for all $\rho < \lambda$. Moreover, $(a_\rho)_{\rho<\lambda}$ is not of algebraic type over $F_2 = F_2^h$. Now, pick $\delta \in \mathrm{Fix}(K)$ with $v(\delta) = (1, 1)$, and consider $F_3 = F_2(\delta)$, with value group $\mathbb{Z} \oplus_{\mathrm{lex}} \mathbb{Z}$. Note that α is still σ-algebraic over F_3 and not algebraic.
4. Now, $p_1(X) = \sigma(X) - \varepsilon X - c \in F_3[X]_\sigma$ is a minimal difference polynomial for $(a_\rho)_{\rho<\lambda}$ over F_3, the latter has no pseudolimit in F_3, $a_\rho \rightsquigarrow \alpha$, and $h(\alpha) = 0$. Thus, $F_3 \subseteq F_3(\alpha)_\sigma$ is an immediate extension by LEMMA 4.3.34.

By deforming the equation $\sigma(X) - \epsilon X - c = 0$, we now construct another immediate extension of F_3. Consider $z \in F$ such that $\mathrm{res}(z) \notin \sigma_{\mathrm{res}}(k)$, and let $p_2(X) = \sigma(X) - \varepsilon X - c - \delta z$. Let $\beta \in L$ in some extension with $p_2(\beta) = 0$. Now,

1. by additivity, $v(\beta - \alpha) = v(\delta) > \gamma_\rho$ for all $\rho < \lambda$, thus $a_\rho \rightsquigarrow \beta$,

2. $v(p_2(a_\rho)) = v(p_1(a_\rho) - \delta z) = v(p_1(a_\rho))$ eventually, so $p_2(a_\rho) \rightsquigarrow 0$,

3. $p_2(\beta) = 0$,

thus $F_3 \subseteq F_3(\beta)_\sigma$ is an immediate extension by LEMMA 4.3.34.

By construction, then, no immediate extension of F_3 may contain solutions of both $p_1(X) = 0$ and $p_2(X) = 0$. Indeed, if L is immediate over F_3 and contains solutions of both $p_1(X) = 0$ and $p_2(X) = 0$, denoting by θ a solution to $p_1(X) = 0$ and by ξ a solution to $p_2(X) = 0$, we would have

$$\operatorname{res}(z) = \sigma_{\mathrm{res}}\left(\operatorname{res}\left(\frac{\xi - \theta}{\delta}\right)\right) \in \sigma_{\mathrm{res}}(k),$$

a contradiction. In particular, if we take any maximal immediate extension of $F_3(\alpha)_\sigma$, it will not be F_3-isomorphic to any maximal immediate extension of $F_3(\beta)_\sigma$.

Bibliography

[AJ18] Sylvy Anscombe and Franziska Jahnke. Henselianity in the language of rings. *Ann. Pure Appl. Log.*, 169(9):872–895, 2018.

[AJ24] Sylvy Anscombe and Franziska Jahnke. Characterizing NIP henselian fields. *J. Lond. Math. Soc.*, 109(3):e12868, 2024.

[AK65] James Ax and Simon Kochen. Diophantine problems over local fields I. *Amer. J. Math.*, 87(3):605–630, 1965.

[AvdD09] Salih Azgin and Lou van den Dries. Equivalence of Valued Fields with Valuation Preserving Automorphism. *arXiv preprint arXiv:0902.0422*, 2009.

[AvdD11] Salih Azgın and Lou van den Dries. Elementary theory of valued fields with a valuation-preserving automorphism. *J. Inst. Math. Jussieu*, 10(1):1–35, 2011.

[AvdDvdH17] Matthias Aschenbrenner, Lou van den Dries, and Joris van der Hoeven. *Asymptotic Differential Algebra and Model Theory of Transseries*. Annals of Mathematics Studies. Princeton University Press, 2017.

[Azg10] Salih Azgın. Valued fields with contractive automorphism and Kaplansky fields. *J. Algebra*, 324(10):2757–2785, 2010.

[BMS07] Luc Bélair, Angus Macintyre, and Thomas Scanlon. Model theory of the Frobenius on the Witt vectors. *Amer. J. Math.*, 129(3):665–721, 2007.

[CH04] Zoé Chatzidakis and Ehud Hrushovski. Model theory of endomorphisms of separably closed fields. *J. Algebra*, 281(2):567–603, 2004.

[CH14] Artem Chernikov and Martin Hils. Valued difference fields and NTP_2. *Israel J. Math.*, 204(1):299–327, 2014.

[Cha09] Zoé Chatzidakis. Notes on the model theory of finite and pseudo-finite fields. https://www.math.ens.psl.eu/ zchatzid/papiers/Helsinki.pdf, 2009.

[Dav73] Martin Davis. Hilbert's tenth problem is unsolvable. *The American Mathematical Monthly*, 80(3):233–269, 1973.

[Del79] Françoise Delon. Types sur $\mathbb{C}((x))$. *Groupe d'étude de théories stables*, 2:1–29, 1978-1979.

[DH22] Yuval Dor and Ehud Hrushovski. Specialization of Difference Equations and High Frobenius Powers. *arXiv preprint arXiv:2212.05366*, 2022.

[DH23] Yuval Dor and Yatir Halevi. Contracting Endomorphisms of Valued Fields. *arXiv preprint arXiv:2305.18963*, 2023.

[DO15] Salih Durhan and Gönenç Onay. Quantifier elimination for valued fields equipped with an automorphism. *Sel. Math.*, 21:1177–1201, 2015.

[EP05] Antonio J Engler and Alexander Prestel. *Valued fields*. Springer Science & Business Media, 2005.

[Ers65] Yurii Ershov. On the elementary theory of maximal normed fields. In *Dokl. Akad.*, volume 165, pages 21–23. Russian Academy of Sciences, 1965.

[FJ15] Arno Fehm and Franziska Jahnke. On the quantifier complexity of definable canonical henselian valuations. *Math. Log. Q.*, 61(4-5):347–361, 2015.

[FJ17] Arno Fehm and Franziska Jahnke. Recent progress on definability of henselian valuations. *Cont. Math.*, 697:135–143, 2017.

[HHYZ24] Martin Hils, Ehud Hrushovski, Jinhe Ye, and Tingxiang Zou. Lang-Weil Type Estimates in Finite Difference Fields. *arXiv preprint arXiv:2406.00880*, 2024.

[HK23] Yatir Halevi and Itay Kaplan. Saturated models for the working model theorist. *Bulletin of Symbolic Logic*, 29(2):163–169, 2023.

[Hod97] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, 1997.

[Hru] Ehud Hrushovski. The Elementary Theory of the Frobenius Automorphism. *preprint, available at http://www.ma.huji.ac.il/ ehud/FROB.pdf*.

[Jah] Franziska Jahnke. Henselian expansions of NIP fields. *to appear in J. Math. Log.* https://doi.org/10.1142/S021906132350006X.

[JK15] Franziska Jahnke and Jochen Koenigsmann. Definable henselian valuations. *J. Symb. Log.*, 80(1):85–99, 2015.

[JK17] Franziska Jahnke and Jochen Koenigsmann. Defining coarsenings of valuations. *Proc. Edinb. Math. Soc.*, 60(3):665–687, 2017.

[Joh20] Will Johnson. Dp-finite fields VI: the dp-finite Shelah conjecture. *preprint arXiv:2005.13989*, 2020.

[JS20] Franziska Jahnke and Pierre Simon. NIP henselian valued fields. *Arch. Math. Logic*, 59(1-2):167–178, 2020.

[Kes24] Christoph Kesting. Tameness Properties in Multiplicative Valued Difference Fields with Lift and Section. *arXiv preprint arXiv:2409.10406*, 2024.

[Koe95] Jochen Koenigsmann. p-henselian fields. *manuscripta mathematica*, 87(1):89–99, 1995.

[Koe14] Jochen Koenigsmann. Undecidability in number theory. *Model theory in algebra, analysis and arithmetic*, 2111:159–195, 2014.

[KR23] Franz-Viktor Kuhlmann and Anna Rzepka. The valuation theory of deeply ramified fields and its connection with defect extensions. *Trans. Amer. Math. Soc.*, 376(04):2693–2738, 2023.

[KRS24] Margarete Ketelsen, Simone Ramello, and Piotr Szewczyk. Definable henselian valuations in positive residue characteristic. *arXiv preprint arXiv:2401.06884, to appear in J. Symb. Log.*, 2024.

[Kuh90] Franz-Viktor Kuhlmann. *Henselian function fields and tame fields*. Heidelberg, 1990. https://www.fvkuhlmann.de/Fvkpaper.html.

[Kuh16] Franz-Viktor Kuhlmann. The algebra and model theory of tame valued fields. *J. Reine Angew. Math.*, 2016(719):1–43, 2016.

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

[Pas89] Johan Pas. Uniform p-adic cell decomposition and local zeta functions. *J. Reine Angew. Math.*, 1989.

[Pre91] Alexander Prestel. Algebraic number fields elementarily determined by their absolute Galois group. *Isr. J. Math.*, 73(2):199–205, 1991.

[PZ78] Alexander Prestel and Martin Ziegler. Model theoretic methods in the theory of topological fields. *J. Reine Angew. Math.*, 1978(299-300):318–341, 1978.

[Ram24] Simone Ramello. Model theory of valued fields with an endomorphism. *arXiv preprint arXiv:2407.05043*, 2024.

[Rid17] Silvain Rideau. Some properties of analytic difference valued fields. *J. Inst. Math. Jussieu*, 16(3):447–499, 2017.

[RKV23] Silvain Rideau-Kikuchi and Mariana Vicaría. Imaginaries in equicharacteristic zero henselian fields. *arXiv preprint arXiv:2311.00657*, 2023.

[Rob63] Julia Robinson. The decision problem for fields. In J.W. Addison, Leon Henkin, and Alfred Tarski, editors, *The Theory of Models*, Studies in Logic and the Foundations of Mathematics, pages 299–311. North-Holland, 1963.

[Roq02] Peter Roquette. History of valuation theory, Part I. *Valuation theory and its applications*, 1:291–355, 2002.

[SV21] Kadattur V Shuddhodan and Yakov Varshavsky. The Hrushovski-Lang-Weil estimates. *arXiv preprint arXiv:2106.10682*, 2021.

[vdD14] Lou van den Dries. Lectures on the model theory of valued fields. In *Model theory in algebra, analysis and arithmetic*, volume 2111 of *Lecture Notes in Math.*, pages 55–157. Springer, Heidelberg, 2014.

[Wei94] Charles A Weibel. *An introduction to homological algebra*. Number 38. Cambridge University Press, 1994.

[Woo79] Carol Wood. Notes on the stability of separably closed fields. *J. Symb. Log.*, 44(3):412–416, 1979.