J. Inst. Math. Jussieu (2025), **24**(5), 2047–2106
doi:10.1017/S1474748025000167

CONTRACTING ENDOMORPHISMS OF VALUED FIELDS

YUVAL DOR[1], YATIR HALEVI[2] AND WITH AN APPENDIX BY ITAY KAPLAN

[1] *Apple Inc., Israel*
(yuval.dor4@gmail.com)
[2] *Department of Mathematics, University of Haifa, 199 Abba Khoushy Avenue, Haifa, Israel*
(ybenarih@campus.haifa.ac.il)

(Received 5 August 2023; revised 2 April 2025; accepted 2 April 2025; first published online 8 May 2025)

In memory of Zoé Chatzidakis, 1955–2025.

Abstract We prove that the class of separably algebraically closed valued fields equipped with a distinguished Frobenius endomorphism $x \mapsto x^q$ is decidable, uniformly in q. The result is a simultaneous generalization of the work of Chatzidakis and Hrushovski (in the case of the trivial valuation) and the work of the first author and Hrushovski (in the case where the fields are algebraically closed).

The logical setting for the proof is a model completeness result for valued fields equipped with an endomorphism σ which is locally infinitely contracting and fails to be onto. Namely, we prove the existence of a model complete theory $\widetilde{\mathrm{VFE}}$ amalgamating the theories SCFE and $\widetilde{\mathrm{VFA}}$ introduced in [5] and [11], respectively. In characteristic zero, we also prove that $\widetilde{\mathrm{VFE}}$ is NTP_2 and classify the stationary types: they are precisely those orthogonal to the fixed field and the value group.

1. Introduction

1.1. Motivation and background

By a *transformal valued field* we mean a valued field K equipped with a distinguished endomorphism $\sigma\colon K \to K$ of fields such that $\sigma^{-1}(\mathcal{O}) = \mathcal{O}$ and $\sigma^{-1}(\mathcal{M}) = \mathcal{M}$; here, $\mathcal{O}$ and $\mathcal{M}$ denote the valuation ring and the maximal ideal, respectively.

Of particular interest is the case where K is of positive characteristic p and $\sigma = p^n$ coincides with a power of the Frobenius endomorphism of K; in this case, we say that K is a *Frobenius transformal valued field*.

The following is a logical corollary of our main results:

Theorem 1 (Theorems 7.16 and 7.17). *The elementary theory of the class of separably algebraically closed Frobenius transformal valued fields is decidable.*

The second author was partially supported by ISF grants No. 555/21 and 290/19 and by the Fields Institute for Research in Mathematical Sciences. The author of the appendix would like to thank the Israel Science Foundation for their support of this research (grant no. 1254/18 and 804/22)

Keywords: transformal; valued fields; frobenius; ultraproduct; model complete

2020 Mathematics subject classification: 12h10; 12l12; 03c60

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Theorem 1 follows from an abstract model completeness result; a key geometric input here is the the twisted Lang-Weil estimates of Hrushovski [16] and Shuddhodan-Varshavsky [22]. More precisely, we will define a theory $\widetilde{\mathrm{VFE}}$ of transformal valued fields which is closely related to the elementary theory of separably algebraically closed Frobenius transformal valued fields; Theorem 1 is then deduced from a quantifier elimination statement for $\widetilde{\mathrm{VFE}}$. Before we turn to the precise statement, it will be useful to give an overview of the model theory of the Frobenius endomorphism.

1.1.1. ACFA **and the limit theory of the Frobenius action.** The seminal work of Hrushovski [16] initiated the study of transformal geometry. Whereas algebraic geometry takes place over a field, transformal geometry takes place over a *difference field* – that is, a field K equipped with a distinguished endomorphism $\sigma\colon K \to K$.

The connection between transformal geometry and algebraic geometry is mediated by Hrushovski's Frobenius specialization functors. Model theoretically, the class of difference fields admits a model companion ACFA, which agrees precisely with the limit theory of the Frobenius action on an algebraically closed field. This justifies the useful heuristic that the generic endomorphism σ can be regarded as an infinitely large prime power.

The theory ACFA was studied in detail in the work of Chatzidakis and Hrushovski [4].

1.1.2. $\widetilde{\mathrm{VFA}}$**.** In Hrushovski's work [16], a key role is played by a moving lemma and a theory of specialization of difference varieties. This leads naturally to the use of valuations, and hence of transformal valued fields of the following form.

Let K be a transformal valued field. We say that K is ω-*increasing* if the induced action of σ on Γ is ω-increasing in the sense that $n\alpha < \sigma\alpha$ for all $0 < \alpha \in \Gamma$ and all $n \in \mathbf{N}$. The Frobenius transformal valued fields are not ω-increasing, but a nonprincipal ultraproduct of them is.

The class of ω-increasing transformal valued fields admits a model companion $\widetilde{\mathrm{VFA}}$. In characteristic zero, this is due to Durhan [2] and in all characteristics due to Hrushovski and the first author [11]. Furthermore, the theory $\widetilde{\mathrm{VFA}}$ is precisely the limit theory of the Frobenius action on an algebraically closed and nontrivially valued field.

1.1.3. Transformally separable extensions. Let K be a difference field. We say that K is *inversive* or *transformally perfect* if σ is on K. We say that an embedding L/K of difference fields is *transformally separable* if K and L^σ are linearly disjoint over K^σ; when $\sigma = p^n$ is the Frobenius endomorphism, this corresponds to the usual notion of a separable field extension.

1.1.4. One of the major goals of the present work is to initiate a systematic study of transformally separable extensions (with no valuation); this is carried out in Section 4. Here, we build and expand on the ideas of Chatzidakis-Hrushovski ([5]).

In order to obtain a well-behaved theory, we restrict attention to a universal theory of difference fields introduced in [5] which we call here FE; see Definition 4.11 for a precise statement. The class of models of FE enjoys excellent functorial properties, and includes

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

algebraically closed fields and Frobenius ultraproducts, as well as all those difference fields which are perfect and inversive.

For transformally separable extensions of FE, we obtain a structure theory which is in many ways analogous to the classical structure theory of separable field extensions; see Theorem 4.48. For example, we exhibit a functorial relative transformal separable algebraic closure, characterized set theoretically as the set of simple roots of 'twisted' difference polynomials; the theory of separable generation can likewise be generalized (see Theorem 4.48).

Using the results of Section 4, functorial constructions available in the inversive realm descend to the category of models of FE; we use this repeatedly in order to apply the machinery of [11]. We expect the results to be useful in other settings; see, for example, the work of Simone Ramello [21], which vastly generalizes our quantifier elimination results in characteristic zero.

1.1.5. Notation. For any theory of T difference fields, possibly enriched, one can define an expansion T_λ of T by definitions in which the models are the same, but embeddings are required to be transformally separable (see 4.11).

1.1.6. SCFE**.** In [5], Chatzidakis and Hrushovski introduce a theory SCFE of difference fields which are not inversive. The theory SCFE is simple, and the expansion SCFE_λ is model complete. The nonprincipal ultraproduct of separably algebraically closed and Frobenius difference fields is a model of SCFE.

1.2. Main results

We let VFE be the theory of ω-increasing transformal valued fields with underlying difference field a model of FE.

We prove the following:

Theorem 2 (See Theorem 7.9).

(1) *The class of models of* VFE *whose inversive hull is a model of* $\widetilde{\mathrm{VFA}}$ *is elementary; write* $\widetilde{\mathrm{VFE}}$ *for the common first order theory.*

(2) *The theory* $\widetilde{\mathrm{VFE}}_\lambda$ *is model complete. The models of* $\widetilde{\mathrm{VFE}}$ *which fail to be inversive are uniquely characterized up to elementary equivalence by the characteristic, the degree of imperfection and the isomorphism type of the action of* σ *on the algebraic closure of the prime field (and all possibilities occur).*

Using Theorem 2 and the main results of [16], it is not difficult to verify the following:

Theorem 3 (Theorem 7.17). *Let* $\widetilde{T}$ *be the asymptotic theory of separably algebraically closed and nontrivially valued Frobenius transformal valued fields – that is, the set of sentences true in all such transformal valued fields, outside a finite set of exceptional prime powers. Then* $\widetilde{T}$ *is decidable and* $\widetilde{T}_\lambda$ *is model complete; indeed, every model of* $\widetilde{T}$ *is a model of* $\widetilde{\mathrm{VFE}}$.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

We prove various other results:

- The theory $\widetilde{\mathrm{VFE}}$ admits an explicit axiomatization which parallels that of $\widetilde{\mathrm{VFA}}$ and is consequently decidable (Theorem 7.16).
- In $\widetilde{\mathrm{VFE}}$, the residue field field and the value group are stably embedded and model theoretically fully orthogonal; the induced structure is that of a pure model of ACFA and $\widetilde{\omega\mathrm{OGA}}$, respectively (Theorem 7.13).
- The model theoretic algebraic closure in $\widetilde{\mathrm{VFE}}$ is determined in Theorem 7.10; it is as small as can possibly be expected given the behavior in $\widetilde{\mathrm{VFA}}$ and SCFE.
- The theory $\widetilde{\mathrm{VFE}}$ admits elimination of quantifiers to the level of the field theoretic algebraic closure in a language where the transformal λ-functions are considered basic; see Corollary 7.7 and Corollary 7.8 for a precise statement. This level of quantifier reduction is likewise optimal given its behavior in $\widetilde{\mathrm{VFA}}$ and SCFE.
- In characteristic zero, we study stability theoretic properties of $\widetilde{\mathrm{VFE}}$. We show that $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ enjoys the NTP_2; over highly saturated models, forking is witnessed by formulas with the NIP (8.23). A basic intuition on NTP_2 theories is that they are NIP theories with 'random noise' (see, for example, [18]); our results support this.

In characteristic zero, we are able to pin down the stationary types:

Theorem 4 (Corollary 8.11). *In the theory $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$, a global invariant type is stationary if and only if it is orthogonal to the fixed field and to the value group.*

In $\mathrm{SCFE}_{\mathbf{Q}}$, this follows from the trichotomy for $\mathrm{ACFA}_{\mathbf{Q}}$ and the stationarity theorem of [5]; we deduce this here in the presence of a valuation using domination results.

1.2.1. The proof that $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ (and $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$) enjoys the NTP_2 uses a domination result by the residue field and the value group over spherically complete models; in order to lift it to the general case, we use an observation due to Chernikov [7]. Itay Kaplan relaxed Chernikov's argument to show that if forking of a global type is always witnessed by an NIP formula, then the theory is NTP_2. We thank him for his permission to include it here as an appendix.

2. Preliminaries

2.1. Conventions and notation

2.1.1. The characteristic exponent of a field. Recall that the *characteristic exponent* of a field is equal to p if the field is of positive characteristic $p > 0$ and is equal to 1 if the field if of characteristic zero. By convention, a natural number p which is either prime or equal to 1 is fixed through this text, and all valued fields under consideration are assumed of equal characteristic, of characteristic exponent equal to p.

2.1.2. Field theory. For basic reference on field theory, we refer the reader to [19, 12]. Fix a field k of characteristic exponent p. We write $\phi x = x^p$ for the Frobenius

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

endomorphism of k; thus, it is the identity map in characteristic exponent $p = 1$. We write $k^{p^{-\infty}}$, k^{sep} and k^{alg} for the perfect hull and for the choice of a separable and algebraic closure of k, respectively. The *Ershov invariant* or the *imperfection degree* of k is the natural number e with the property that $[k\colon k^p] = p^e$ if the former quantity is finite, and equal to ∞ otherwise.

Let K/k be a field extension. We say that K is *separable* over k if the fields K and $k^{p^{-\infty}}$ are linearly disjoint over k in $K^{p^{-\infty}}$. We say that K is *separably generated* over k if there is a transcendence basis x of K over k such that K is separably algebraic over $k(x)$. If K is separably generated over k, then it is separable over K; the converse holds when K is a finitely generated field extension of k. We say that K is *primary* over k if the field k is relatively separably algebraically closed in K; that is, every element of K separably algebraic over k already lies in k. It is evident that purely inseparably algebraic extensions are primary. Finally, the field K is *regular* over k if it is both separable and primary over k.

Let K/k be a field extension. We say that K is *relatively perfect* over k or that *the perfect hull of K splits over k* if the field K is separable over k and one has $K \otimes_k k^{p^{-\infty}} = K^{p^{-\infty}}$. Equivalently,[1] the field extension K/k is relatively perfect if it is separable and some (every) p-basis of k is a p-basis of K. Beware that in [19], relatively perfect extensions are *not* assumed separable. If k is of finite Ershov invariant, then a separable extension K of k is relatively perfect precisely in the event that the Ershov invariants coincide.

2.1.3. Difference algebra. We follow the conventions and notation of [16] and [11]. See Section 3 of [11] for a summary of the basic facts from the model theory of difference fields we will use.

By a *difference ring* we mean a (commutative and unital) ring R equipped with a distinguished (injective) endomorphism σ of rings; a *homomorphism* of difference rings is a ring homomorphism compatible with the action of σ in the evident manner. By a *difference field* we mean a difference ring whose underlying ring is a field.

Let K be a difference field. We use exponential notation for the action of σ; thus, for an element (or a tuple) $a \in K$, we write $a^\sigma = \sigma(a)$. We also write $K^\sigma = \sigma(K)$ for the image of K under the endomorphism σ and regard K^σ as a subfield of K.

Let K be a difference field. We say that K is *inversive* if σ is an automorphism of K. In the literature, this is sometimes taken as a part of the definition, but the whole point of the present work is to dispense with this assumption. In any event, if K is a difference field, then there is an inversive difference field $K^{\sigma^{-\infty}}$ over K which embeds uniquely over K in any other. In this situation, we say that $K^{\sigma^{-\infty}}$ is the *inversive hull* of K; it is unique up to a unique isomorphism of difference fields over K. The inversive hull of K is given explicitly as the directed union of isomorphic copies of K, so any property of difference fields stable in directed unions passes from K to its inversive hull.

[1] We are thankful to the referee for this remark.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

2.2. Notation

We write $\mathbf{N}[\sigma]$ for the commutative semiring of finite formal sums of the form $\lambda_n\sigma^n + \ldots + \lambda_1\sigma + \lambda_0$ where the λ_i are natural numbers. If K is a difference field, then we write $K\left[x^{\mathbf{N}[\sigma]}\right]$ for the ring of difference polynomials over K. For a an element (or a tuple) in an ambient difference field extension of K, we write $K\left(a^{\mathbf{N}[\sigma]}\right)$ for the smallest difference field extension of K containing a. We write $K^{\sigma \cdot p^{-\infty}}$ for the perfect hull of K^σ.

Let K be a difference field and a an element in an ambient extension. We say that a is *transformally algebraic* over K if it is a zero of a nonzero difference polynomial over K; otherwise, the element a is said to be *transformally transcendental* over K. The notion of a transformal transcendence basis and transformal independence of a tuple behaves as expected; see [11], Section 3 for more.

3. Descent, base extension, and the field theoretic tensor product

In modern algebraic geometry and commutative algebra, the method of descent along a morphism is ubiquitous. In the present work, in order to avoid excess generality, we restrict ourselves exclusively to the case of fields. Grothendieck's relative point of view is nevertheless valid: it is often possible to study an embedding L/K of difference fields after extending K, and many statements true over K inversive and algebraically closed admit relative versions. The purpose of this section is to set up the notation and terminology necessary in order to make a systematic use of descent. The material here is standard, and no claim to originality is made.

3.0.1. Linear disjointness

Let $B \hookleftarrow A \hookrightarrow C$ be difference fields, jointly embedded over A in an ambient extension Ω. We say that B and C are *linearly disjoint over* A in Ω if whenever the sequence $a_1, \ldots, a_n \in B$ is linearly independent over A, then it remains linearly independent over C. For more on this notion, see [12, Section 2.5].

3.0.2. Abstract linear disjointness

Linear disjointness is strictly speaking meaningless in the absence of an ambient extension Ω. Nevertheless, we will encounter the situation where the choice of Ω is irrelevant. Thus, the following definition will be useful:

Definition 3.1. Let $B \hookleftarrow A \hookrightarrow C$ be embeddings of difference fields.

(1) By a *compositum* of B and C over A we mean a difference field F over A, equipped with embeddings of B and C over A, such that F is generated as a difference field by the images of B and C.

(2) Let F be a compositum of B and C over A. We say that F is a *linearly free compositum* if B and C are linearly disjoint over A in F in the sense of 3.0.1.

(3) We say that B and C are *abstractly linearly disjoint* over A or simply *linearly disjoint* over A if they admit a linearly free compositum over A.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Remark 3.2. Let B and C be difference fields over A, jointly embedded over A in an ambient extension Ω. Then whether or not B and C are linearly disjoint over A in Ω depends only on the underlying field structure and not on the action of σ. By contrast, abstract linear disjointness depends on the difference field structure. Namely, the category of fields admits linearly disjoint amalgamation over an algebraically closed base, whereas the category of difference fields only admits linearly disjoint amalgamation over a base which is algebraically closed and inversive.

Let B and C be difference fields over A. Then a linearly free compositum F of B and C over A, if it exists, is unique up to a unique isomorphism. More precisely, the difference field F is characterized by the following universal mapping property. Let E be a difference field extension of A. Then the datum of an embedding of F in E over A is equivalent to the datum of embeddings of B and C in E over A, so as to render them linearly disjoint over A in the image.

3.0.3. Tensor product of difference fields

Definition 3.3. Let B and A_0 be difference fields, abstractly linearly disjoint over A. Then we write $B_0 = B \otimes_A A_0$ for the linearly free compositum of A; by the discussion of 3.0.1, it is uniquely determined up to a unique isomorphism. In this situation, we say that $B_0 = B \otimes_A A_0$ is the *(difference) field theoretic tensor product* of B and A_0 over A.

In the settings of Definition 3.3, we say that B_0/A_0 is obtained from B/A via *base extension*.

Remark 3.4. Let K be an abstract field. We may regard K as a difference field via the identity automorphism. In this way, the definitions apply to abstract fields as well.

Remark 3.5 (Warning). The notation of Definition 3.3 is potentially ambiguous as we now explain. Let us use temporarily the notation $\overline{\otimes}$ for the field theoretic tensor product. Let L and M be difference fields over K and write $R = L \otimes_K M$ for the usual tensor product of rings. The ring R comes equipped with a canonical endomorphism, compatible with both inclusions. The difference fields L and M are abstractly linearly disjoint over K if and only if the ring R is an integral domain and σ is injective on R. In this situation, the endomorphism σ of R lifts uniquely to the field of fractions F of R, and the field theoretic tensor product is just the field of fractions of the usual tensor product (i.e, we have $L \overline{\otimes}_K M = \mathrm{Frac}(R)$). However, in this work, we will deal exclusively with fields, as opposed to rings, so no confusion should arise.

Remark 3.6. The field theoretic tensor is associative and compatible with directed unions whenever defined. This follows immediately from the universal mapping property of the linearly free compositum of 3.0.2.

Remark 3.7. Let B and A_0 be difference fields, abstractly linearly disjoint over A, and write $B_0 = B \otimes_A A_0$ for the field theoretic tensor product. In model theory, it is common to regard linear disjointness as an abstract independence relation, obeying properties similar to forking independence. Moreover, we have given a description of the field theoretic

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

tensor product which is of category theoretic flavor. However, despite the symmetry of the construction, the role played by B and A_0 in our application is not symmetric: the idea is that for A_0 an arbitrary extension of A, properties of the extension B/A are inherited by the base extension B_0/A_0 and vice versa; see the discussion of 3.0.4.

3.0.4. Invariance under base extension

Here is the key definition.

Definition 3.8. Let $\mathcal{P}$ be a property of embeddings of difference fields. Let B/A be an extension of difference fields and assume that B_0/A_0 is obtained from B/A via base extension; thus, $B_0 = B \otimes_A A_0$ (in particular, B and A_0 are linearly disjoint over A). We say that the property $\mathcal{P}$ is *invariant under base extension* if the two extensions B/A and B_0/A_0 simultaneously enjoy the property $\mathcal{P}$.

We will only be interested in properties of embeddings of difference fields which are invariant under base extension; this is the case for many natural classes of interest; see Fact 3.9. Moreover, we will encounter the following situation. Suppose we wish to prove a statement about an extension B/A of difference fields. If we can show that the hypothesis and the conclusion are both invariant under base extension, then we can pick an extension A_0 of A abstractly linearly disjoint from B over A, and study the base extension B_0/A_0 instead, with no loss of generality. In practice, we will almost always take A_0 to be the perfect or inversive hull of A.

3.0.5. Basic properties

The following elementary properties of abstract linear disjointness will be used repeatedly, often without mention. See, for example, [12, Chapter 2].

Fix a tower $A \subseteq B \subseteq C$ of difference fields, and let A_0 be a difference field extension of A. Let us assume that A_0 and C are linearly disjoint over A. Then A_0 and B are linearly disjoint over A. Furthermore, the difference field $B_0 = A_0 \otimes_A B$ is linearly disjoint from C over A. The converse holds as well using associativity of the field theoretic tensor product.

Moreover, linear disjointness is of finite character in the following sense. Let $B \hookleftarrow A \hookrightarrow C$ be embeddings of difference fields and assume that B is the directed union of difference field extensions B_i of A; then B and C are linearly disjoint over A if and only if the B_i are all linearly disjoint from C over A.

3.0.6. Descent for field extensions

Recall that the field extension K/k is *relatively perfect* if it is separable and one has $K \otimes_k k^{p^{-\infty}} = K^{p^{-\infty}}$.

Fact 3.9. The following properties of field extensions are invariant under base extension: separable extensions, relatively perfect extensions, primary extensions and regular extensions.

Proof. The results are well known, so we only sketch the proofs. Let K/k be a field extension and M/m a base extension (i.e, we have $M = K \otimes_k m$).

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Separable extensions. We want to prove that K/k is separable if and only if M/m is so. If M is separable over m, then by definition, the fields M and $m^{p^{-\infty}}$ are linearly disjoint over m. By transitivity of linear disjointness, the fields K and $m^{p^{-\infty}}$ are linearly disjoint over k. In view of the inclusion $k^{p^{-\infty}} \subseteq m^{p^{-\infty}}$, it follows in particular that K and $k^{p^{-\infty}}$ are linearly disjoint over k, which is to say that K is separable over k.

Now assume conversely that K is separable over k; we want to prove that M is separable over m. By finite character of separable extensions and the discussion of 3.0.5, we may assume that K is finitely generated over k. So K is separably generated over k (i.e, there is a transcendence basis x of K over k such that K is separably algebraic over $k(x)$). It then follows immediately from the definition of linear disjointness that x is also a separating transcendence basis of M over m.

Primary extensions. First, assume that M is primary over m. Then M and m^{sep} are linearly disjoint over m. By transitivity of linear disjointness, the fields K and m^{sep} are linearly disjoint over k; hence, in particular, K and k^{sep} are linearly disjoint over k.

Now assume that K is primary over k; we show that M is primary over m. We will give a model theoretic argument. By finite character, we may assume that $K = k(\alpha)$ for a finite sequence α. Recall that the theory of algebraically closed fields eliminates imaginaries in the field language and that the definable closure of a subfield is the perfect hull. It is moreover a stable theory. So K is primary over k if and only if the type $\mathrm{tp}_{\mathrm{ACF}}(\alpha/k)$ admits a unique global nonforking extension. Now in the theory ACF, forking is characterized by a drop in the transcendence degree. Since K and m are linearly disjoint over k, a transcendence basis of K over k lifts to a transcendence basis of M over m. So as $\mathrm{tp}(\alpha/m)$ does not fork over k, we find that it admits a unique global nonforking extension too. So $\mathrm{tp}(\alpha/m)$ is stationary, whence M is primary over m.

Regular extensions. This follows from the fact that an extension is regular if and only if it is both primary and separable.

Relatively perfect extensions. We want to prove that K/k is relatively perfect if and only if M/m is so. Note that relatively perfect extensions are in particular separable; we have already shown invariance of separable extensions, so assume that one (and hence both) of these extensions is separable.

Let us assume first that k and m are both perfect. By definition, a relatively perfect extension of a perfect field is just a perfect extension. Thus, we must show that if one of K or M is perfect, then so is the other. If K is perfect, then M is perfect, since it is the tensor product of the perfect fields K and m. Conversely, assume that M is perfect. The field M is separable over K as it is obtained from the separable extension m/k via base extension. Since no separable extension of an imperfect field can possibly be perfect, we find that K is perfect.

Next, consider the case where $m = k^{p^{-\infty}}$ is the perfect hull of k. If K is relatively perfect over k, then by definition, M is the perfect hull of K, and hence perfect; so it is a relatively perfect extension of m. Conversely, assume M is perfect. The field M is purely inseparably algebraic over K, since it is obtained from the purely inseparably algebraic extension m/k via base extension. Since it is perfect, it is also the maximal purely inseparably algebraic extension of K (i.e, we have $M = K^{p^{-\infty}}$, which means that K/k is relatively perfect).

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

The general case follows from transitivity, combining the two special cases considered above together. □

The class of separable extensions is in general not transitive in towers. However, transitivity holds if the extension on the bottom is relatively perfect:

Fact 3.10. Let $K \subseteq M \subseteq L$ be a tower of fields and assume M is relatively perfect over K. Then L is separable over K if and only if it is separable over M.

Proof. Let us assume first that L is separable over M. We must prove that L is separable over K. Since M/K is relatively perfect, it is in particular separable. The field extension L/K is therefore separable as the composition of such.

Let us assume conversely that L is separable over K and that M is relatively perfect over K; we must prove that L is separable over M. Let $\overline{K} = K^{p^{-\infty}}$ be the perfect hull of K. Let us set $\overline{M} = M \otimes_K \overline{K}$ and likewise $\overline{L}$. Since separable and relatively perfect extensions are both invariant under base extension (Fact 3.9), replacing everything by the overlined data, we may assume that K is perfect. But then it follows immediately from the definition of relatively perfect extensions that M is perfect. In this situation, the field L is automatically separable over M, as wanted. □

4. Transformally separable extensions

Let K/k be a field extension. Recall that K is said to be *separable* over k if the fields K and $k^{p^{-\infty}}$ are linearly disjoint over k. The purpose of this section is to study this notion in the transformal setting.

4.1. Transformally separable extensions

The following definition will play a central role in this work.

Definition 4.1. Let L/K be an extension of difference fields.

(1) We say that L is *transformally separable* over K if the difference fields K and L^σ are linearly disjoint over K^σ.

(2) We say that L is *transformally separably algebraic* over K if it is transformally separable and transformally algebraic over K.

(3) We say that L is *purely transformally inseparably algebraic* over K if $K^{\sigma^{-\infty}} = L^{\sigma^{-\infty}}$; that is, for all $a \in L$, we have $a^{\sigma^n} \in K$ for $n \gg 0$ sufficiently large.

Remark 4.2. Let L/K be an extension of difference fields. One checks that the following conditions are equivalent:

(1) The difference field extension L/K is transformally separable.

(2) The difference fields $K^{\sigma^{-1}}$ and L are linearly disjoint over K.

(3) The difference fields $K^{\sigma^{-\infty}}$ and L are linearly disjoint over K.

Remark 4.3. Let L/K be a transformally separable extension. Then $K^{\sigma^{-\infty}}$ and L are linearly disjoint over K; hence, $K^{\sigma^{-\infty}} \cap L = K$, the intersection taken in $L^{\sigma^{-\infty}}$. This

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

condition is, however, not in general sufficient. See [12, Remark 2.7.6] for an example where σ coincides with the Frobenius endomorphism.

Remark 4.4. The class of transformally separable extensions is in general not transitive in towers. More precisely, let $K \subseteq M \subseteq L$ be a tower of difference fields; then it may be that L/K is transformally separable but L/M is not. For example, let K be an inversive difference field, let M be a difference field extension of K which is not inversive, and let L be the inversive hull of M. Since K is inversive, the extension L/K is transformally separable, but L/M is not. See, however, 4.7 for a partial converse.

The following summarizes some basic formal properties of transformally separable extensions; it will be used repeatedly throughout this work.

Proposition 4.5. *Let $K \subseteq M \subseteq L$ be a tower of difference fields.*

(1) *The composition of transformally separable extensions is again one. More precisely, if M/K and L/M are transformally separable, then L/K is transformally separable.*

(2) *The class of transformally separable extensions of a fixed base goes downwards. More precisely, if L/K is transformally separable, then M/K is transformally separable.*

(3) *The class of transformally separable extensions is of finite character. More precisely, the extension L/K is transformally separable if and only if it is the directed union of finitely generated transformally separable extensions.*

(4) *The class of transformally separable extensions is invariant under base extension. More precisely, let L/K be a difference field extension. Let K_0 be a difference field extension of K, linearly disjoint from L over K, and set $L_0 = L \otimes_K K_0$; then L/K is transformally separable if and only if L_0/K_0 is so.*

Proof. This follows from formal properties of linear disjointness. Let us prove, for example, that the class of transformally separable extensions descends along base extension; we use the notation of (4).

Let us assume first that L_0/K_0 is transformally separable. We will show that L/K is transformally separable. The assumption that L_0/K_0 is transformally separable means that L_0 and $K_0^{\sigma^{-\infty}}$ are linearly disjoint over K_0. By transitivity of linear disjointness, the fields L and $K_0^{\sigma^{-\infty}}$ are linearly disjoint over K. In view of the embedding $K^{\sigma^{-\infty}} \subseteq K_0^{\sigma^{-\infty}}$, it follows in particular that L and $K^{\sigma^{-\infty}}$ are linearly disjoint over K, which is to say that L/K is transformally separable, as wanted.

Let us assume conversely that L/K is transformally separable; we must show that L_0 is transformally separable over K_0. By transitivity of linear disjointness and Remark 4.2, it is enough to prove that L is linearly disjoint from $(K_0)^{\sigma^{-1}}$ over K.

So let $x_1, \ldots, x_n \in L$ be linearly dependent over $(K_0)^{\sigma^{-1}}$, that is, there are $m_i^{\sigma^{-1}} \in (K_0)^{\sigma^{-1}}$, not all zero, such that the equality $\sum m_i^{\sigma^{-1}} x_i = 0$ holds. We must show that the x_i are linearly dependent over K already. Applying σ, we get that $\sum m_i x_i^{\sigma} = 0$. By linear disjointness of L from K_0 over K, there are $k_i \in K$ with $\sum k_i x_i^{\sigma} = 0$, and hence, $\sum k_i^{\sigma^{-1}} x_i = 0$. So the x_i are linearly dependent over $K^{\sigma^{-1}}$; since L is transformally

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

separable over K, they are linearly dependent over K already, and hence in particular over K_0. □

4.2. Difference fields with twists

Let K be a difference field of characteristic exponent p. We write $\phi x = x^p$ for the Frobenius endomorphism of K; thus, ϕ is the identity map in characteristic exponent $p = 1$.

Definition 4.6. Let K be a difference field. We say that K is *closed under twists* if every power of the Frobenius endomorphism of K factors through σ. In other words, for every $n \in \mathbf{N}$, there is an endomorphism $\tau\colon K \to K$ of fields such that $\tau \circ \phi^n = \sigma$.

In the situation of Definition 4.6, the endomorphism τ is unique, and we write $\tau = \frac{\sigma}{p^n}$.

Remark 4.7. Let K be a difference field. Then the following are equivalent:

(1) The difference field K is closed under twists.
(2) For every $x \in K$ and every $n \in \mathbf{N}$, the polynomial $X^{p^n} - x^\sigma$ splits in K.
(3) The field K contains the perfect hull of K^σ.
(4) The field $K^{\sigma^{-1}}$ contains the perfect hull of K.

Remark 4.8. Let K be a difference field which is closed under twists. If K is inversive, then K is perfect. Indeed, by definition, we have the inclusions $K \subseteq K^{p^{-\infty}} \subseteq K^{\sigma^{-1}}$; since $K = K^{\sigma^{-1}}$, the inclusion $K \subseteq K^{p^{-\infty}}$ is an equality.

Remark 4.9. The compositum of difference fields closed under twists inside an ambient difference field is closed under twists, regardless of linear disjointness assumptions.

Proposition 4.10. *The following holds:*
(1) *Let K be a difference field. Then there is a difference field $\widetilde{K}$ over K closed under twists which embeds uniquely, over K in every other extension with this property. The field $\widetilde{K}$ is purely inseparably algebraic over K.*
(2) *Let K be a difference field which is closed under twists. Let L be an abstract purely inseparably algebraic field extension of K. Then σ lifts uniquely to L, and equipped with this lift, the difference field L is closed under twists.*
(3) *The class of difference fields closed under twists is stable in directed unions, relatively perfect extensions and algebraic extensions.*

Proof. (1) Let $\widetilde{K}$ be the splitting field of the polynomials of the form $X^{p^n} - x^\sigma$ over K, for $x \in K$ and $n \in \mathbf{N}$ varying. Then $\widetilde{K}$ is a normal purely inseparably algebraic extension of K. By construction, if the polynomial $f \in K[x]$ splits in $\widetilde{K}$, then the same is true of the polynomial obtained by applying σ to the coefficients. By uniqueness of splitting fields, the endomorphism σ of K lifts to $\widetilde{K}$. Since $\widetilde{K}$ is purely inseparably algebraic over K, the lifting is unique. We claim that $\widetilde{K}$ is closed under twists, and that it embeds uniquely over K in every other extension with this property.

First, we prove that $\widetilde{K}$ is closed under twists. Fix an element $a \in \widetilde{K}$; we must show that a^σ has all p-th power roots in $\widetilde{K}$. It is enough to prove this after replacing a by a^{p^n} for

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

$n \gg 0$. Since $\widetilde{K}$ is purely inseparably algebraic over K, we may therefore assume without loss of generality that the element a lies in K. But then it follows from the construction that the polynomials $X^{p^n} - a^\sigma$ all split in $\widetilde{K}$.

Now let L be an extension of K closed under twists. By Remark 4.7, for every $x \in K$ and every $n \in \mathbf{N}$, the polynomial $X^{p^n} - x^\sigma$ splits in L. By uniqueness of splitting fields, there is an abstract field embedding of $\widetilde{K}$ in L over K; since $\widetilde{K}$ is purely inseparably algebraic over K, this embedding is unique and compatible with the action of σ.

(2) The lift of σ is unique if it exists, as L is purely inseparably algebraic over K. Moreover, by assumption, we have the series of inclusions $K \subseteq L \subseteq K^{p^{-\infty}} \subseteq K^{\sigma^{-1}}$; if $a \in L$, then a^σ lies in K, and hence in L. So L is a difference field. Moreover, we have the inclusions $L \subseteq L^{p^{-\infty}} = K^{p^{-\infty}} \subseteq K^{\sigma^{-1}} \subseteq L^{\sigma^{-1}}$, so L is closed under twists.

(3) Stability under directed unions is clear. Next, assume that L is relatively perfect over K and that K is closed under twists; we want to prove that L is closed under twists as well. By assumption, we have $L^{\sigma \cdot p^{-\infty}} = L^\sigma \otimes_{K^\sigma} K^{\sigma \cdot p^{-\infty}}$; so the inclusion $K^{\sigma \cdot p^{-\infty}} \subseteq K$ gives an inclusion $L^{\sigma \cdot p^{-\infty}} \subseteq L$. Finally, assume that L is algebraic over K. There is a canonical factorization $K \subseteq M \subseteq L$ where M/K is separably algebraic and L/M purely inseparable; in view of (2), we may assume that L is separably algebraic over K. In this situation, the field L is relatively perfect over K; this case has already been dealt with, so we are done. □

4.3. Twisted difference polynomials

In this subsection, we introduce twisted difference polynomials.

4.3.1. Recall that $\mathbf{N}[\sigma]$ is the commutative semiring of finite formal sums of the form $\lambda_0 + \lambda_1 \cdot \sigma + \ldots + \lambda_N \cdot \sigma^N$ where the λ_n are natural numbers, with pointwise addition and multiplication given by the rule $\sigma^n \cdot \sigma^m = \sigma^{n+m}$ for $n,m \in \mathbf{N}$. We write $\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]$ for the union $\bigcup_{n=0}^\infty \mathbf{N}\left[\frac{\sigma}{p^n}\right]$, so every element of $\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]$ can be written in the form $\lambda_0 + \lambda_1 \cdot \frac{\sigma}{q_1} + \ldots + \lambda_N \cdot \frac{\sigma^N}{q_N}$, where the q_i are powers of p.

4.3.2. The ring of twisted difference polynomials. Let K be a difference field which is closed under twists. We now define the ring of *twisted difference polynomials* over K, denoted by $A = K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$. The elements of A are formal finite K-linear combinations of twisted difference monomials of the form x^ν where $\nu \in \mathbf{N}\left[\frac{\sigma}{p^\infty}\right]$. Multiplication is given by the rule $x^\nu \cdot x^\mu = x^{\nu+\mu}$, and we have an action of σ on $K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ given by $(x^\nu)^\sigma = x^{\sigma \cdot \nu}$. The difference ring A is transformally integral in the sense that the underlying ring is an integral domain and σ is injective. It is, moreover, closed under twists in the sense that for every $f \in A$ and $n \in \mathbf{N}$, there is a unique $g \in A$ with $g^{p^n} = f^\sigma$. Moreover, for an element $f \in A$, there is an evident evaluation map $K \to K$ denoted by $a \mapsto fa$.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

4.3.3. Derivatives of twisted difference polynomials. Let K be a difference field, closed under twists. Let $A = K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be the difference ring of twisted difference polynomials over K. The conditions[2] $dx = 1$ and $dx^{\sigma^n} = 0$ for $n > 0$ specify a unique K-derivation $d\colon A \to A$. For $f \in A$, we write $df = f'$ and refer to f' as the *transformal derivative* or simply the *derivative* of f.

4.4. The theory FE

We will be interested in the following class of difference fields, for which the notion of a transformally separable extension is particularly well behaved. In [5], this theory is denoted by T_σ.

Definition 4.11. Let FE denote the following first order theory of difference fields K:

(1) The difference field K is closed under twists in the sense of Definition 4.6; that is, the field K contains the perfect hull of K^σ.
(2) The field K is a primary extension of K^σ; that is, the field K^σ is relatively separably algebraically closed in K.

Remark 4.12. We shall regard models of FE as structures for the language of difference fields; in this language, an embedding of models of FE is an embedding of difference fields – that is, a ring homomorphism compatible with the action of σ.

It is clear that the EC models of FE are inversive; to rule this out, we introduce the following definition:

Definition 4.13. We let FE_λ be the theory whose models are difference fields as in 4.11 but embeddings are required to be transformally separable. This can be made first order either with relations expressing linear independence of tuples in F^n over F^σ or in a functional language – namely, with the transformal λ-functions (see 4.11). The exact choice of language will not matter for us; regardless of the choice, the theory FE_λ is a conservative expansion by definitions of FE (i.e, every model of FE admits a unique expansion to a model of FE_λ).

Remark 4.14. Let K be a difference field. Then the following are equivalent:

(1) The field K^σ is separably algebraically closed in K.
(2) The field K is separably algebraically closed in $K^{\sigma^{-1}}$.
(3) The field K is separably algebraically closed in its inversive hull $K^{\sigma^{-\infty}}$.

Lemma 4.15. *Let K be a difference field. Then there is a difference field extension E of K which is a model of* FE *and which embeds uniquely over K in any other. Moreover, if K is closed under twists, then E is separably algebraic over K, and at all events it is algebraic over K.*

[2] In exponential characteristic $p \neq 1$, the requirement that $dx^\sigma = 0$ is redundant; we have $x^\sigma = y^p$ where $y = x^{\frac{\sigma}{p}}$, so $dx^\sigma = p \cdot dy = 0$

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

In the situation of Lemma 4.15, we will say (for lack of better terminology) that E is the *smallest model of* FE *over* K.

Proof. We may assume that K is closed under twists, by Proposition 4.10.

Let $K^{\sigma^{-\infty}}$ denote the inversive hull of K, and let E be the separable algebraic closure of K inside. Then E is a difference field: if $a \in K^{\sigma^{-\infty}}$ is separably algebraic over K, then a^σ is separably algebraic over K^σ, and hence over K. We claim that E is advertised; for this, we must prove E is a model of FE which embeds uniquely over K in any other.

First, we check that E is a model of the theory FE. By construction, the field E is separably algebraically closed in $K^{\sigma^{-\infty}}$; however, we have $E^{\sigma^{-\infty}} = K^{\sigma^{-\infty}}$, so E is separably algebraically closed in its inversive hull. Since E is an algebraic extension of K, which is closed under twists, it is again closed under twists, using Proposition 4.10 again. It follows that E is a model of the theory FE.

To finish, we must prove that E embeds uniquely over K in every model L of FE over K. So fix such a model L. The embedding of K in L lifts uniquely to an embedding of $K^{\sigma^{-\infty}}$ in $L^{\sigma^{-\infty}}$. Since the field L is a model of FE, it is separably algebraically closed in its inversive hull; so the restriction of this embedding to E factors through L, by definition of the relative separable algebraic closure. So E embeds uniquely over K in L, and the proof is finished. □

4.4.1. Closure Properties

Lemma 4.16. *Let K be a model of* FE *and let L be an abstract purely inseparably algebraic extension of K. Then σ lifts uniquely to L, and this lift renders L a model of* FE.

Proof. Since K is closed under twists, it follows from Proposition 4.10 that σ lifts uniquely to L, and that L is again closed under twists when equipped with this lift. It remains to be shown that the algebraic closure of L in $L^{\sigma^{-\infty}}$ is purely inseparable over L. But L is purely inseparably algebraic over K; in view of the isomorphism $K^{\sigma^{-\infty}} = L^{\sigma^{-\infty}}$, the claim follows from the fact that L is purely inseparably algebraic over K. □

Lemma 4.17. *Let K be a model of* FE *and let x be an element which is transformally transcendental over K, in an ambient extension. Then the difference field $E = K\left(x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$ is a model of* FE.

Proof. Implicit in the choice of notation is the fact that E is closed under twists. To finish, we must prove that E is primary over E^σ, or equivalently over $E^{\sigma \cdot p^{-\infty}}$. The fields K and $E^{\sigma \cdot p^{-\infty}}$ are linearly disjoint over $K^{\sigma \cdot p^{-\infty}}$ so by descent for primary extensions, it is enough to prove that E is primary over $K \otimes_{K^{\sigma \cdot p^{-\infty}}} E^{\sigma \cdot p^{-\infty}}$. But this is evident from the construction: the field E is generated over the latter by the single element x which is algebraically transcendental. □

Example 4.18. Let L, M be models of FE linearly disjoint over a base model K of FE; then the tensor product $L \otimes_K M$ need not be a model of FE. For example, let $K = \mathbf{Q}\left(x^{\mathbf{N}[\sigma]}\right)$ where x is transformally transcendental, let L be the inversive hull of K, and

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

let $M = \mathbf{Q}\left(y^{\mathbf{N}[\sigma]}\right)$ where $y^2 = x$. Then $N = L \otimes_K M$ is not a model of FE; its inversive hull is an algebraic extension, obtained by adjoining square roots of the elements x^{σ^n} for all $n \in \mathbf{Z}$. See, however, Proposition 4.20 for a partial converse when K is separably algebraically closed and L,M are transformally separable over K.

We will make use of the following (see [4]):

Fact 4.19 (The Independence Theorem). Let E be an algebraically closed field and let A,B,C be algebraically closed field extensions of E, linearly disjoint over E in some ambient extension. Then the field $(AC)^{\mathrm{alg}}(BC)^{\mathrm{alg}}$ is a regular field extension of AB.

Proposition 4.20 (The amalgamation property). *Let A be a model of* FE *which is separably algebraically closed. Let B,C be models of* FE *transformally separable over A. Then B and C are linearly disjoint over A, and the difference field $B \otimes_A C$ is a model of* FE.

Remark 4.21. The following Remark will be used in the proof without further mention. Let $\mathcal{C}$ be the class of difference fields F with the property that F is relatively separably algebraically closed in its inversive hull $F^{\sigma^{-\infty}}$. Then $\mathcal{C}$ is downwards closed in transformally separable extensions and in primary extensions. More precisely, let E/F be an extension of difference fields which is transformally separable or primary. If E is a member of $\mathcal{C}$, then F is also a member of $\mathcal{C}$.

Proof. The category of difference fields admits linearly disjoint amalgamation over a base which is inversive and algebraically closed. Since A is closed under twists and separably algebraically closed, the inversive hull of A is algebraically closed. By transitivity of linear disjointness, using the fact that B,C are transformally separable over A, we find that they are linearly disjoint over A.

It remains to be shown that $B \otimes_A C$ is a model of FE. The tensor product of difference fields closed under twists is again closed under twists; so we must prove that $B \otimes_A C$ is relatively separably algebraically closed in its inversive hull. We prove this in several steps.

Step 1. We may assume that A is perfect. Let $A_0 = A^{p^{-\infty}}$ be the perfect hull of A. Since transformally separable extensions of models of FE are algebraically separable (see Lemma 4.30), the fields B,C are algebraically separable over A. Let $B_0 = B \otimes_A A_0$ and likewise C_0. Since purely inseparably algebraic extensions of models of FE are models of FE (Lemma 4.16), we have that B_0,C_0 are models of FE. Using Remark 4.21, we can replace everything by the subscripted data. By a similar argument, we can assume that B is perfect.

Step 2. We may assume that the field B is separably algebraically closed. Namely, let $B_0 = B^{\mathrm{sep}}$. Then B_0 is a model of FE, and it is transformally separable over B. Using Remark 4.21, it is enough to prove the claim after B has been replaced by B_0.

Combining these two steps and using symmetry, we may assume that the fields A,B,C are all algebraically closed.

Step 3. We may assume that A is inversive. Let $A_0 = A^{\sigma^{-\infty}}$ be the inversive hull of A. Let $B_0 = (B \otimes_A A_0)^{\mathrm{alg}}$ and likewise $C_0 = (C \otimes_A A_0)^{\mathrm{alg}}$. Then B_0,C_0 are algebraically

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

closed and hence models of FE. Moreover, the fields B_0, C_0 are transformally separable over A_0, as the latter is inversive. By Fact 4.19, the field $B_0 \otimes_{A_0} C_0$ is regular over $B \otimes_A C$, so using Remark 4.21, it is enough to prove the claim for the subscripted data.

Step 4. We may assume that B is inversive. Let $F = B \otimes_A C$, let $B_0 = B^{\sigma^{-\infty}}$ be the inversive hull of B and let $F_0 = B_0 \otimes_A C$. Since B is algebraically closed, the field B_0 is regular over B. Since regular extensions are invariant under base extension (Fact 3.9), the field F_0 is regular over F. Using Remark 4.21, we can replace B by B_0.

Combining all these steps and using symmetry, we may assume that A, B, C are all inversive (and hence perfect). So $B \otimes_A C$ is perfect and inversive as the tensor product of fields with this property; it is then a model of FE, which is what we wanted. □

Remark 4.22. Fix $e \in \mathbf{N}$ and write $\mathrm{FE}_{\leq e}$ for the class of models of FE of degree of imperfection at most e. Let A be a model of $\mathrm{FE}_{\leq e}$ which is separably algebraically closed. Let A_1, A_2 be models of FE_e transformally separable over A and set $B = A_1 \otimes_A A_2$. Then B is a model of FE transformally separable over the A_i, and hence over A. If (x_i) and (y_j) are relative p-bases of A_1 and A_2 over A, respectively, then (x_i, y_j) is a relative p-basis of B over A; so the Ershov invariant of B will exceed e in general (see 4.9.1). However, there is model $\widetilde{B}$ of FE purely inseparably algebraic over B and transformally separable over the A_i which is a model of $\mathrm{FE}_{\leq e}$. To prove this, assume for simplicity of notation that A is perfect and that the A_i are of Ershov invariant exactly e. Then the purely inseparably algebraic extension $\widetilde{B}$ of B obtained by adjoining all p-th power roots to the elements $z_1, \ldots, z_e$ is of Ershov invariant equal to e; here, $z_i = x_i y_i$ for $i = 1, \ldots, e$. The field $\widetilde{B}$ is then evidently as promised.

4.5. Almost transformally separable extensions

The notion of an *almost transformally separable* extension is a soft variant of the notion of a transformally separable extension. This allows breaking up many arguments into two parts, using the criterion of Proposition 4.32.

4.5.1. Algebraic disjointness. Let $B \hookleftarrow A \hookrightarrow C$ be fields, jointly embedded over A in an ambient extension Ω. We say that B and C are *algebraically disjoint* over A if whenever the sequence $a_1, \ldots, a_n \in B$ is algebraically independent over A, then it remains algebraically independent over C. If the transcendence degree of B over A is finite, then algebraic disjointness is equivalent to the requirement that $\mathrm{tdeg}_A B = \mathrm{tdeg}_C CB$.

Fact 4.23 [12, Lemma 2.6.7]. Let $L \hookleftarrow K \hookrightarrow M$ be field extensions, jointly embedded over K in an ambient extension Ω. If L and M are linearly disjoint over K, then they are algebraically disjoint over K; the converse holds provided that L is regular over K.

4.5.2. Almost transformally separable extensions.

Definition 4.24. Let L/K be an extension of difference fields. We say that L is *almost transformally separable* over K if the fields K and L^σ are algebraically disjoint over K^σ in L the sense of 4.5.1.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Remark 4.25. Let L/K be an extension of difference fields. If L is transformally separable over K, then it is almost transformally separable over K; this follows immediately from Fact 4.23.

Remark 4.26. Let L/K be an extension of difference fields. Then the following conditions are equivalent:

(1) The difference field extension L/K is almost transformally separable.
(2) The difference fields $K^{\sigma^{-1}}$ and L are algebraically disjoint over K in $L^{\sigma^{-1}}$.
(3) The difference fields $K^{\sigma^{-\infty}}$ and L are algebraically disjoint over K in $L^{\sigma^{-\infty}}$.

Remark 4.27. Let K be an algebraically closed difference field; then an extension of K is almost transformally separable if and only if it is transformally separable. This is an immediate consequence of Fact 4.23, as every field extension of an algebraically closed field is a regular extension. More generally, let L/K be an extension of difference fields. Then L/K is almost transformally separable if and only if L^{alg} is an almost transformally separable extension of K^{alg}, for some (any) choice of an extension of σ from L to L^{alg}.

Proposition 4.28. *Let $K \subseteq M \subseteq L$ be a tower of difference fields.*

(1) *The composition of almost transformally separable extensions is again one. More precisely, if M/K and L/M are almost transformally separable, then L is almost transformally separable over K.*
(2) *The class of almost transformally separable extensions is downwards closed. More precisely, if L/K is almost transformally separable, then M/K is almost transformally separable.*
(3) *The class of almost transformally separable extensions is of finite character. More precisely, the extension L/K is almost transformally separable if and only if it is the directed union of finitely generated almost transformally separable extensions.*
(4) *The class of almost transformally separable extensions is invariant under base extension. More precisely, let L/K be a difference field extension. Let $\overline{K}$ be a difference field extension of K, linearly disjoint from L over K, and set $\overline{L} = L \otimes_K \overline{K}$; then L/K is almost transformally separable if and only if $\overline{L}/\overline{K}$ is almost transformally separable.*

Proof. As in Proposition 4.5, this follows from formal properties of algebraic disjointness. □

4.6. A criterion for transformally separable extensions

In this subsection, we give a criterion for extensions to be transformally separable when the base is a model of FE.

We start with the following; see also [16, Lemma 5.1].

Lemma 4.29. *Let M/K be a transformally algebraic extension of difference field with K inversive. Then M is algebraic over M^σ.*

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. The following elegant proof is due to the anonymous referee. By a limit argument, we may assume that M is finitely generated over K; since M is transformally algebraic over K, the transcendence degree $\mathrm{tdeg}_K M$ is finite. Since K is inversive, we have the inclusions $K \subseteq M^\sigma \subseteq M$; the claim follows by comparing transcendence degrees in the displayed tower. □

Lemma 4.30. *Let K be a difference field which is closed under twists. Then a transformally separable extension of K is algebraically separable.*

Proof. Let L/K be transformally separable with K closed under twists. Then $K^{\sigma^{-1}}$ and L are linearly disjoint over K. Since K is closed under twists, we have the inclusions $K \subseteq K^{p^{-1}} \subseteq K^{\sigma^{-1}}$; by formal properties of linear disjointness, we find that $K^{p^{-1}}$ and L are linearly disjoint over K, whence L/K is algebraically separable. □

Lemma 4.31. *Let L/K be a transformally separably algebraic extension of difference fields, with K and L closed under twists. Then the perfect hull of L splits over K.*

Proof. Everything here is invariant under base extension (see Fact 3.9 and Proposition 4.5), so assume that K is inversive (Remark 4.8). By Lemma 4.29, the field L is algebraic over L^σ, and hence in particular algebraic over $L^{\sigma \cdot p^{-\infty}}$. But algebraic extensions of perfect fields are perfect, so L is perfect, as wanted. □

The following criterion will be used repeatedly:

Proposition 4.32. *Let L/K be an extension of difference fields with K a model of* FE. *Then L/K is transformally separable if and only if it is almost transformally separable and algebraically separable.*

Proof. If K is a model of FE, then it is closed under twists; a transformally separable extension is then algebraically separable by Lemma 4.30. Moreover, by Remark 4.25, an extension which is transformally separable is almost transformally separable, no matter the assumption on the ground field. Thus, an extension of a model of FE which is transformally separable is algebraically separable and almost transformally separable; it remains to show the converse.

We first reduce this to the case where K is perfect, using the discussion of 3.0.4. By elementary properties of transformally separable and almost transformally separable extensions (Proposition 4.28 and Proposition 4.5), the hypothesis and the conclusion are both invariant under base extension. Furthermore, by Lemma 4.16, a purely inseparably algebraic extension of a model of FE is a model of FE. We may therefore assume without loss of generality that K is perfect.

It remains to be shown that if K is a perfect model of FE, then an extension is transformally separable if and only if it is almost transformally separable. Since K is a model of FE, the field $K^{\sigma^{-1}}$ is a primary extension of K; our assumption that K is perfect implies that $K^{\sigma^{-1}}$ is in fact regular over K. By Fact 4.23, linear disjointness form $K^{\sigma^{-1}}$ over K is equivalent to algebraic disjointness, which is precisely what we needed. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

4.7. Transitivity in towers

The class of transformally separable extensions is not in general transitive in towers; see Remark 4.4. In this subsection, we investigate to what extent the converse holds.

Proposition 4.33. *The class of almost transformally separable extensions is transitive in towers, provided that the bottom extension is transformally algebraic. More precisely, let $K \subseteq M \subseteq L$ be a tower of difference fields and assume that M is transformally algebraic over K. Then L/K is almost transformally separable if and only if M/K and L/M are so.*

Proof. It is enough to prove that L/M is almost transformally separable; the other directions follow formally from Proposition 4.28.

Let us assume first that K is inversive. Since M is transformally algebraic over K, it follows from Lemma 4.29 that the field M is an algebraic extension of M^σ. So L is automatically almost transformally separable over M, as wanted.

It remains to reduce to the case where K is inversive. Let L^{alg} be an algebraic closure of L in the category of difference fields; replacing L by L^{alg} and M,K by the algebraic closures inside L^{alg} does not affect the statement and the conclusion descends, using Remark 4.27; thus, we may assume that K is algebraically closed, and hence that L/K is transformally separable. Let K_0 be the inversive hull of K. Since L/K is transformally separable, the difference field L is linearly disjoint from K_0 over K.

Let us set $M_0 = M \otimes_K K_0$ and $L_0 = L \otimes_K K_0$. By Proposition 4.28, the class of almost transformally separable extensions is invariant under base extensions; one checks, moreover, that the same is true of the class of transformally algebraic extensions. It is therefore harmless to replace K,M and L by K_0,M_0 and L_0, respectively, which brings us back to the situation already dealt with. □

Recall that an extension L/K of difference fields is said to be transformally separably algebraic if it is transformally separable and transformally algebraic.

Proposition 4.34. *Let L/K be a transformally separably algebraic extension of models of* FE.

(1) *The perfect hull of L splits over K.*

(2) *If the difference field K is inversive, then so is L.*

Remark 4.35. It follows in particular that the Ershov invariant does not change in a transformally separably algebraic extension of models of FE.

Note that for (2), any transformally algebraic extension of an inversive field is transformally separably algebraic.

Proof. (1) This is a special case of Lemma 4.31, since models of FE are in particular closed under twists.

(2) Assume that K is inversive; we show that L is inversive, or equivalently that the inclusion $L^\sigma \subseteq L$ is an equality. For this, it is enough to show that L is transformally separable over L^σ. By the criterion of Proposition 4.32, in order to prove that L is

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

transformally separable over L^σ, it will be sufficient to prove that L is algebraically separable and almost transformally separable over L^σ.

First, we show that L is algebraically separable over L^σ. By Remark 4.8, inversive models of FE are perfect; so using (1), we find that L^σ is perfect, as the isomorphic image of the perfect field L. So L is automatically algebraically separable over L^σ. Since K is inversive, the extension L of K is almost transformally separably algebraic; this class is transitive in towers, using Proposition 4.33, so L is almost transformally separable over L^σ. This concludes the proof. □

Proposition 4.36. *The class of transformally separable extensions of models of* FE *is transitive in towers, provided that the bottom extension is transformally algebraic. More precisely, let $K \subseteq M \subseteq L$ be a tower of models of* FE *with M/K transformally algebraic; then L/K is transformally separable if and only if L/M and M/K are so.*

Remark 4.37. This should be compared with [16, Lemma 5.2(2)].

Proof. By formal properties of transformally separable extensions as in Proposition 4.5, it will be sufficient to prove the following: if L/K is a transformally separable extension of models of FE, and $K \subseteq M \subseteq L$ is a model of FE transformally separably algebraic over K, then L/M is transformally separable. Using the criterion of Proposition 4.32, it is enough to prove that L/M is almost transformally separable and algebraically separable.

The fact that L/M is almost transformally separable follows from Proposition 4.33; so we must show that L/M is algebraically separable. By Proposition 4.34, the perfect hull of M splits over K, so the claim follows from Fact 3.10. □

4.8. Simple roots and transformally separably algebraic extensions

Let E/F be a field extension. Then E is separably algebraic over F if and only if every element of E is a simple root of a polynomial with coefficients in F. The purpose of this subsection is to generalize this to transformally separably algebraic extensions of models of FE. Recall the discussion of 4.3.2 on twisted difference polynomials and their derivatives.

Definition 4.38. Let L/K be an extension of difference fields closed under twists, let $fx \in K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted difference polynomial, and let $a \in L$ be an element; then a is said to be a *simple root* of f if $fa = 0$ and $f'a \neq 0$.

Remark 4.39. Let L/K be an extension of difference fields closed under twists and let $fx \in K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted difference polynomial. Let $a \in L$ be an element and let $gx \in K[x]$ be the polynomial obtained from fx via the substitutions $x^{\sigma^n \cdot p^m} \mapsto a^{\sigma^n \cdot p^m}$ for $n > 0$ and $m \in \mathbf{Z}$; then $f'a = g'a$. In particular, if we regard an ordinary polynomial over K as a twisted difference polynomial, then the derivative defined above has its usual meaning.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Lemma 4.40. *Let K be a difference field and let $L = K\left(a^{\mathbf{N}[\sigma]}\right)$, where a is a finite tuple. Then the following are equivalent:*

(1) *The difference field L is transformally algebraic and almost transformally separable over K.*

(2) *The tuple a is algebraic over $K\left(a^{\sigma\cdot\mathbf{N}[\sigma]}\right)$.*

Proof. (1) $\Longrightarrow$ (2) We may assume here that a is a single element. By assumption, the element a is transformally algebraic over K, so for some $n \in \mathbf{N}$, the sequence $a,\ldots,a^{\sigma^n}$ is algebraically dependent over K; let n be chosen minimal with respect to this property. Then by minimality of n the sequence, $a,\ldots,a^{\sigma^{n-1}}$ is algebraically independent over K. Since L is almost transformally separable over K, the sequence $a^{\sigma},\ldots,a^{\sigma^n}$ must then likewise be algebraically independent over K; by Steinitz exchange, the element a is algebraic over $K\left(a^{\sigma},\ldots,a^{\sigma^n}\right)$, and hence in particular algebraic over $K\left(a^{\sigma\cdot\mathbf{N}[\sigma]}\right)$.

(2) $\Longrightarrow$ (1) Since the tuple a is finite, there is some large natural number N with $a \in K\left(a^{\sigma},\ldots,a^{\sigma^N}\right)^{\mathrm{alg}}$. Since σ is an endomorphism of fields, it follows inductively that for all $n \in \mathbf{N}$, we have $a^{\sigma^n} \in K^{\sigma^n}\left(a^{\sigma^{n+1}},\ldots,a^{\sigma^{n+N}}\right)^{\mathrm{alg}}$, and so in particular that

$$a^{\sigma^n} \in K\left(a^{\sigma^{n+1}},\ldots,a^{\sigma^{n+N}}\right)^{\mathrm{alg}}.$$

Enlarging a by the tuple $a\ldots a^{\sigma^{N-1}}$ does not affect the statement; using the analysis of the previous paragraph, we then have that $a \in K\left(a^{\sigma}\right)^{\mathrm{alg}}$ and inductively that $a,\ldots,a^{\sigma^{n-1}} \in K\left(a^{\sigma^n}\right)^{\mathrm{alg}}$ for all $n \in \mathbf{N}$. So the transcendence degree of $K\left(a^{\mathbf{N}[\sigma]}\right)$ over K is finite; namely, it is bounded by the length of the tuple a. This means that a is transformally algebraic over K.

It remains to be shown that L/K is almost transformally separable: in other words, that the fields K and L^{σ} are algebraically disjoint over K^{σ}. By assumption, the field L is algebraic over KL^{σ}. Moreover, we have shown that L is transformally algebraic over K, so the transcendence degree $\mathrm{tdeg}_K L$ is finite. We then have

$$\mathrm{tdeg}_K L = \mathrm{tdeg}_K KL^{\sigma} \leq \mathrm{tdeg}_{K^{\sigma}} L^{\sigma}.$$

But σ is a field endomorphism, so $\mathrm{tdeg}_{K^{\sigma}} L^{\sigma} = \mathrm{tdeg}_K L$; the displayed inequalities must therefore be equalities, so L^{σ} and K are algebraically disjoint over K^{σ}. □

Here is the key statement:

Proposition 4.41. *Let L/K be a transformally algebraic extension of difference fields with K a model of* FE *and L closed under twists. Then the following are equivalent:*

(1) *The difference field L is transformally separably algebraic over K.*

(2) *Every element of L is a simple root of a twisted difference polynomials with coefficients in K.*

(3) *The difference field L is generated as a twisted difference field by simple roots of twisted difference polynomials with coefficients in K.*

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

(4) *The difference field L is a directed union of difference fields of the form $K\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$, where a is a finite tuple separably algebraic over $K\left(a^{\sigma\cdot\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$.*

Proof. (1) $\implies$ (2) Fix an element $a \in L$; we must show that it is a root of a simple twisted difference polynomial over K. Since L is transformally algebraic over K, there is a finite nonempty set $I \subseteq \mathbf{N}\left[\frac{\sigma}{p^\infty}\right]$ and nonzero elements $c_\nu \in K$ such that $\sum c_\nu a^\nu = 0$; indeed, this is precisely what it means for L to be transformally algebraic over K. If $fx = \sum c_\nu x^\nu$, then f' is nonzero precisely in the event that some $\nu \in I$ has constant term prime to p; let us show how to arrange this.

First, assume that every element $\nu \in I$ has constant term zero. Then $\sum c_\nu^{\frac{1}{\sigma}} a^{\frac{\nu}{\sigma}} = 0$, so the elements $a^{\frac{\nu}{\sigma}}$ are linearly dependent over $K^{\sigma^{-1}}$; since L is transformally separable over K, they must have been linearly dependent over K already. Thus, inductively, we can assume that some $\nu \in I$ has a nonzero constant term. Now assume that the constant term of every $\nu \in I$ is divisible by p; then the elements $a^{\frac{\nu}{p}}$ are linearly dependent over $K^{p^{-1}}$. The latter is contained in $K^{\sigma^{-1}}$ as K is a model of FE, and by assumption, L is transformally separable over K; thus we may assume that at least one $\nu \in I$ has constant term prime to p.

We have found a nonzero twisted difference polynomial f such that $f' \neq 0$ and $fa = 0$. If $f'a \neq 0$, then we are done. Otherwise, replace f by f' and argue inductively; this cannot go on forever.

(2) $\implies$ (3) This is clear.

(3) $\implies$ (4) Let us assume that L is generated as a twisted difference field by simple roots of twisted difference polynomials over K. Then L is the directed union of difference fields $K\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$, where $a = (a_1, \ldots, a_n)$ is a finite tuple and each a_i is a simple root of a twisted difference polynomials over K. Then each a_i is separably algebraic over $K\left(a_i^{\sigma\cdot\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$ by Remark 4.39; it follows in particular that a is separably algebraic over $K\left(a^{\sigma\cdot\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$, hence the claim.

(4) $\implies$ (1) Since transformally separably algebraic extensions are of finite character, we may assume that $L = K\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$, where a is a finite tuple separably algebraic over $K\left(a^{\sigma\cdot\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$. We have already seen in Lemma 4.40 that L is then almost transformally separable over K. By Proposition 4.32, it is enough to prove that L/K is algebraically separable. We will show a stronger statement – namely, that L is relatively perfect over K.

By Proposition 4.32 and Lemma 4.40, the extension $L^{\sigma\cdot p^{-\infty}}$ of $K^{\sigma\cdot p^{-\infty}}$ is transformally separable. Let

$$E = K \otimes_{K^{\sigma\cdot p^{-\infty}}} L^{\sigma\cdot p^{-\infty}}$$

be the displayed difference field. The field extension $L^{\sigma\cdot p^{-\infty}}$ of $K^{\sigma\cdot p^{-\infty}}$ is relatively perfect; by descent, the field E is relatively perfect over K. Since L is separably algebraic over E,

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

it follows from elementary properties of relatively perfect extensions that L is relatively perfect over K, which is what we wanted. □

4.9. Transformal p-independence, transformal p-bases, and transformal separable generation

The purpose of this subsection is to generalize the notion of separable generation to the transformal setting.

4.9.1. Relative and absolute p-bases. Let L/K be a separable field extension. Recall that the sequence (x_i) of elements of L is said to be *p-independent* over K in L if the polynomials $X^p - x_i^p$ are irreducible over $K \otimes_{K^p} L^p$, with linearly disjoint splitting fields. This is equivalent to the requirement that the sequence of monomials $\prod x_i^{m_i}$ for $0 \leq m_i \leq p-1$ is linearly independent over $K \otimes_{K^p} L^p$, as these monomials span the splitting field. Such a sequence is automatically algebraically independent over K, and p-independence enjoys the Steinitz exchange property.

By a *relative p-basis* of L over K we mean a maximal sequence (x_i) of elements of L which is p-independent over K. So if (x_i) is p-independent over K in L, then L is separable over $K(x_i)$, and if (x_i) is a relative p-basis, then the perfect hull of L splits over $K(x_i)$. There is also an absolute version – namely, by an (absolute) *p-basis* of L we mean a relative p-basis over the prime field, or any other perfect subfield. Every set of generators of L over a perfect subfield contains an absolute p-basis.

4.9.2. Transformal separable generation. Let L/K be a field extension. Recall that L is *separably generated* over K if there is a transcendence basis $x = (x_i)_{i \in I}$ of L over K with L separably algebraic over $K(x)$; we then say that x is a *separating transcendence basis* of L over K. If L is finitely generated over K, then it is separable if and only if it is separably generated, in which case a transcendence basis is separating if and only if it is a relative p-basis as in 4.9.1.

Definition 4.42. Let L/K be an extension of models of FE. Then L is *transformally separably generated* over K if there is a transformal transcendence basis $x = (x_i)_{i \in I}$ of L over K such that L is transformally separably algebraic over $K\left(x^{\mathbf{N}[\frac{\sigma}{p^\infty}]}\right)$. If x can be chosen finite, then L is *finitely transformally separably generated* over K. The transformal transcendence basis x is then said to be a *separating transformal transcendence basis* of L over K.

Remark 4.43. Let L/K be an extension of models of FE which is transformally separably generated. Then L is transformally separable over K. Indeed, purely transformally transcendental extensions are transformally separable, and the composition of transformally separable extensions is again one.

Remark 4.44. If L is transformally separable over K, then it is not in general transformally separably generated over K. For example, let $K = \mathbf{Q}$ and $L = K\left(x^{\mathbf{N}[\sigma^{\pm 1}]}\right)$ for x an element transformally transcendental over K.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Lemma 4.45. *Let L/K be a transformally separable extension with K a model of* FE *and L closed under twists. Let us set $E = K \otimes_{K^{\sigma \cdot p^{-\infty}}} L^{\sigma \cdot p^{-\infty}}$.*

(1) *The field L is an algebraically separable field extension of E.*

(2) *Let us assume that L is a model of* FE *that is finitely generated over K as a model of* FE *(i.e, there is a finite tuple a such that L is the smallest model of* FE *over $K\left(a^{\mathbf{N}[\sigma]}\right)$; then L is separably generated over E).*

Proof. (1) We must show that L is algebraically separable over E. The field extension $L^{\sigma \cdot p^{-\infty}}$ of $K^{\sigma \cdot p^{-\infty}}$ is relatively perfect by definition. By descent for relatively perfect extensions, the field E is relatively perfect over K. Since L is algebraically separable over K, it follows from Fact 3.10 that L is algebraically separable over E.

(2) Fix a finite tuple $a \in L$ which generates L as a model of FE over K. Let us set $L_0 = K\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$ and $E_0 = K \otimes_{K^{\sigma \cdot p^{-\infty}}} L_0^{\sigma \cdot p^{-\infty}}$. Then by (1), the field L_0 is algebraically separable over E_0. It is moreover finitely generated over E_0 as a field, namely by the tuple a, hence separably generated over E_0. Now the field L is separably algebraic over L_0; namely, it is the separable algebraic closure of L_0 in its inversive hull (Lemma 4.15). Moreover, the field E is separably algebraic over E_0. Thus, a separating transcendence basis of L_0 over E_0 is also a separating transcendence basis of L over E, as wanted. □

Proposition 4.46. *Let L/K be an extension of models of* FE*. Then L is transformally separable over K if and only if it is the directed union of models of* FE *which are finitely transformally separably generated over K.*

Remark 4.47. This should be compared with [4, Lemma 3.3].

Proof. The easy direction follows from Remark 4.43. For the converse, fix a finite tuple $a \in L$ and let L_0 be the smallest model of FE over K inside L containing a; we will show that L_0 is transformally separably generated over K. Relabeling, we may assume that $L = L_0$. Using the notation of Lemma 4.45, the field L is separably generated over E; after possibly enlarging a within L, we can fix a subtuple $b \subseteq a$ which is a separating transcendence basis for L over $E = K \otimes_{K^{\sigma \cdot p^{-\infty}}} L^{\sigma \cdot p^{-\infty}}$.

We claim that b is a separating transformal transcendence basis for L over K. For this, we must check that b is transformally independent over K and that L is transformally separably algebraic over $M = K\left(b^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$.

First, we verify that the tuple b is transformally independent over K. By construction, the tuple b is algebraically independent over E; it follows from the definition of transformal independence that it is transformally independent over K in L; thus, M is a model of FE, using Lemma 4.17.

Next, we verify that b is a separating transformal transcendence basis of L over K – namely, that L is transformally separably algebraic over M. Since b is an (algebraic) separating transcendence basis for L over E, we again see formally that every element of L is simple root of a twisted difference polynomial with coefficients in M; by Proposition 4.41, the field L is transformally separably algebraic over M. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

4.10. Summary

For ease of later reference, we collect here a summary of the main results of the previous section.

Theorem 4.48.

(1) *(The relative transformal separable algebraic closure) Let E/F be an embedding of models of* FE*. Then there is a model $\widetilde{F}$ of* FE *with $F \subseteq \widetilde{F} \subseteq E$ transformally separably algebraic over F characterized uniquely by the following property: whenever L is a model of* FE *transformally separably algebraic over F, then every embedding of L in E over F lands inside $\widetilde{F}$. We say that $\widetilde{F}$ is the* ***relative transformal separable algebraic closure*** *of F in E.*

(2) *(Transformally separably algebraic extensions and simple roots) Let E/F be an embedding of models of* FE*. Then the set of elements of E which are simple roots of twisted difference polynomials over F is a difference field, which is a model of* FE*. In fact, it is precisely the relative transformal separable algebraic closure of F in E.*

(3) *(Transitivity) Let $F \subseteq \widetilde{F} \subseteq E$ be a tower of models of* FE *where $\widetilde{F}$ is the relative transformal separable algebraic closure of F in E; then E is transformally separable over $\widetilde{F}$ if and only if it is transformally separable over F.*

(4) *(Descent over the inversive hull) Let F be a model of* FE *and let $F^{\sigma^{-\infty}}$ be its inversive hull. Then there is a canonical equivalence of categories between models of* FE *transformally separably algebraic over F and perfect, inversive difference fields transformally algebraic over $F^{\sigma^{-\infty}}$.*

(5) *(Transformal separable generation) Let E/F be a transformally separable extension of models of* FE*, and assume that E is finitely generated over F as a model of* FE*. Then E is transformally separably generated over F. To wit, there is a transformal transcendence basis $x \in E$ of E over F such that E is transformally separably algebraic over $F\left(x^{\mathbf{N}\left[\frac{\sigma}{p^{\infty}}\right]}\right)$.*

(6) *(The amalgamation property and closure under tensor products) Let A be a model of* FE *which is separably algebraically closed and let A_1, A_2 be models of* FE *transformally separable over A. Then A_1 and A_2 are linearly disjoint over A, and the difference field $A_1 \otimes_A A_2$ is a model of* FE*.*

Proof. The truth of (1) and (2) follows immediately from Proposition 4.41. Namely, let $\mathcal{X}$ be the set of elements of E which are simple roots of twisted difference polynomials over F. The smallest model $\widetilde{F}$ of FE inside E containing $\mathcal{X}$ is transformally separably algebraic over F; however, all elements of $\widetilde{F}$ are simple roots over F, so $\mathcal{X} = \widetilde{F}$.

The truth of (3) follows from Proposition 4.36. In (4), the equivalence is given by associating to a model E of FE transformally separably algebraic over F its inversive hull $E^{\sigma^{-\infty}}$. Indeed, the inversive hull of a model of FE is again one, and by Proposition 4.34, all models of FE transformally algebraic over $F^{\sigma^{-\infty}}$ are inversive. We can recover E inside $E^{\sigma^{-\infty}}$ as the relative transformal separable algebraic closure of F inside, since $E^{\sigma^{-\infty}}$ is by construction purely transformally inseparably algebraic over E. The proof of (5) is given in Proposition 4.46, and finally, (6) is Proposition 4.20. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Definition 4.49. Let M/K be an extension of models of FE. We say that M is *relatively inversive* over K if it is transformally separable over K, and whenever L is a model of FE with $K \subseteq M \subseteq L$ and L is transformally separable over K, then L is also transformally separable over M.

Remark 4.50. Let M/K be a relatively inversive extension of models of FE.

(1) If K is an inversive model of FE, then M must likewise be inversive (in the definition, consider the inversive hull of M).
(2) The field M is relatively perfect over K.
(3) Transformally separably algebraic extensions of models of FE are relatively inversive.

4.11. The transformal λ-functions

Let K be a difference field. We now define the *transformal λ-functions* of K. Let $n \in \mathbf{N}$ be fixed for the moment; we will define the function $\lambda = \lambda_n$. This function takes as an input an $n+1$-tuple of elements of K and returns an n-tuple of elements of K. To define it, fix a tuple $(a_1,\ldots,a_n,b) \in K^n \times K$. If the $a_1,\ldots,a_n$ are linearly dependent over K^σ, or else if b is not a K^σ-linear combination of the a_i, then $\lambda(a_1,\ldots,a_n,b) = 0$. Otherwise, the coordinates of $\lambda(a_1,\ldots,a_n,b)$ are the unique elements $c_1,\ldots,c_n \in K$ satisfying $\sum_{i=1}^n c_i^\sigma a_i = b$. Note that the transformal λ-functions of K are definable in the pure difference field language.

Let $F \subseteq K$ be an extension of difference fields. Then K is transformally separable over F if and only if F is closed under the transformal λ-functions of K. For suppose that F and K^σ are linearly disjoint over F^σ. Fix $(a_1,\ldots,a_n,b) \in F^n \times F$ with the $a_1,\ldots,a_n$ linearly independent over K^σ and b in the K^σ-linear span of the a_i. Then the tuple $(a_1,\ldots,a_n,b)$ is linearly dependent over K^σ; by linear disjointness, it must be linearly dependent over F^σ already, whence $\lambda(a_1,\ldots,a_n,b)$ has all coordinates in F. For the converse, assume F and K^σ are not linearly disjoint over F^σ. Let $a_1,\ldots,a_n,b$ be a tuple of elements of F linearly dependent over K^σ but linearly independent over F^σ. We may assume that n is chosen minimal with respect to this property. Then one of the coordinates of $\lambda(a_1,\ldots,a_n,b)$ is not in F; hence, F is not closed under the transformal λ-functions of K.

4.11.1. Let E be a model of FE. Recall that the *Frobenius twists* of E are the difference fields $(E,\sigma\circ\phi^m)$ for $m \in \mathbf{Z}$ where $\phi x = x^p$ is the Frobenius endomorphism of E.

Definition 4.51. Let E be a model of FE. We say that E is a model of SCFE if, in addition, it obeys the following two conditions:

(1) The field E is separably algebraically closed as an abstract field.
(2) The difference field E and all of its Frobenius twists obey the following property. Let V be an absolutely irreducible affine algebraic variety over E. Let $C \subseteq V \times V^\sigma$ be a locally closed absolutely irreducible subvariety and assume that the first projection is dominant and the second projection $C \to V^\sigma$ is finite étale; then there is a point $a \in V(E)$ such that $(a,a^\sigma) \in C(E)$.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Remark 4.52. The axiomatization of Definition 4.51 is slightly different from the axiomatization of [5], but as their proof shows, it is sufficient to handle the case where the second projection is finite. Note that by Proposition 4.41, if E and a are as in Definition 4.51, then the difference field $E(a^{\mathbf{N}[\sigma]})$ is transformally separably algebraic over E, so models of SCFE are precisely the models of FE which are existentially closed for transformally separably algebraic extensions.

Proposition 4.53. *Let E be a model of* FE*. Then*

(1) *The difference field E is a model of* SCFE *if and only if the inversive hull of E is a model of* ACFA*. Thus, if L/E is purely transformally inseparably algebraic, then L is a model of* SCFE *if and only if E is so.*

(2) *If E is relatively transformally separably algebraically closed in a model of of* SCFE*, then it is a model of* SCFE*.*

Proof. (1) The forward direction is [5, Proposition 5.2]; we leave the other direction to the reader.

(2) We may replace E and L by their inversive hulls, using (1); so E is transformally algebraically closed in a model of ACFA and hence itself a model of ACFA. □

Proposition 4.54. *Let K/F be an extension of models of* FE*.*

(1) *Let us assume that K is transformally separable over F and fix an element $x \in K$ not in $F \otimes_{F^p} K^p$; then x is transformally transcendental over F and K is transformally separable over $F\big(x^{\mathbf{N}[\frac{\sigma}{p^\infty}]}\big)$.*

(2) *Let us assume that K is almost transformally separable over F and fix an element $x \in K$ algebraically transcendental over $F \otimes_{F^{\sigma \cdot p^{-\infty}}} K^{\sigma \cdot p^{-\infty}}$; then x is transformally transcendental over F and K is almost transformally separable over $F\big(x^{\mathbf{N}[\frac{\sigma}{p^\infty}]}\big)$.*

Proof. Let us set $E = F\big(x^{\mathbf{N}[\frac{\sigma}{p^\infty}]}\big)$.

(1) In fact, it will be sufficient to assume that K is merely closed under twists. The requirement that x is not in $F \otimes_{F^p} K^p$ means that the sequence $1, x, \ldots, x^{p-1}$ is linearly independent over $F \otimes_{F^p} K^p$. This condition is invariant under base extension; using descent for the various properties considered here, we may thus assume that F is inversive.

We want to show that x is transformally transcendental over F. The element x is not a simple root of any polynomial with coefficients in K^p. Since F is inversive, the element x is not a simple root of any twisted difference polynomial over F; in other words, the element x is transformally transcendental over F.

Now we want to show that K is transformally separable over E. Since E is a model of FE, it is enough to show that K is algebraically separable and almost transformally separable over E (Proposition 4.32).

The fact that K is algebraically separable over E follows immediately from our choice of x; the extension E of E^p is normal of degree p, and since x is not in K^p, the minimal polynomial of x remains irreducible over K^p; thus, E and K^p are linearly disjoint

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

over E^p. Furthermore, since x is transformally transcendental over F, it is algebraically transcendental over K^σ; so E and K^σ are algebraically disjoint over E^σ. It follows that K is almost transformally separable over E, and we are done.

(2) We are free to replace F and K by their perfect hulls; now K is transformally separable over F, and the proof is similar to (1). □

Lemma 4.55. *Let E be a sufficiently saturated model of* FE *which fails to be inversive. Let $F \subseteq E$ be a small model of* FE *contained in E and with E almost transformally separable over F. Then we can find an element $a \in E$ transcendental over $F \otimes_{F^{\sigma \cdot p^{-\infty}}} E^{\sigma \cdot p^{-\infty}}$. Moreover, assume E is of infinite degree of imperfection and that F is transformally separable over E outright; then we can find $a \in E$ not in $F \otimes_{F^p} E^p$.*

Proof. If E is of infinite degree of imperfection, then by saturation, E is of large dimension as a vector space over E^p; as F is small, the inclusion $F \otimes_{F^p} E^p \subseteq E$ must be strict. Now we claim that in any case, we can find an element $a \in E$ not algebraic over $F \otimes_{F^{\sigma \cdot p^{-\infty}}} E^{\sigma \cdot p^{-\infty}}$. Since E fails to be inversive, there is an element $a \in E$ not in E^σ. If E is perfect, then as it is a model of FE, the field E^σ is algebraically closed in E, so a is transcendental over E. In this situation for m fixed, the sequence $a, a^q, a^{q^2}, \ldots, a^{q^m}$ approximates the type of a sequence of m elements of E algebraically independent over E^σ, and we conclude by compactness and saturation – similarly if E is imperfect by taking a to be an element of $E^{\sigma \cdot p^{-\infty}}$ of large degree over E^σ. □

5. Transformal valued fields

Definition 5.1. Let ωOGE be the the theory of ordered $\mathbf{Z}[\frac{\sigma}{p^\infty}]$-modules – that is, ordered abelian groups Γ equipped with an order preserving endomorphism $\sigma \colon \Gamma \to \Gamma$ of abelian groups such that $n\alpha < \sigma\alpha$ for all $0 < \alpha \in \Gamma$ and all $n \in \mathbf{N}$ with the further property that $\sigma\alpha$ has all p-th power roots in Γ.

Fact 5.2 [11, Section 2]. The theory ωOGE admits a model companion $\widetilde{\omega\text{OGA}}$ axiomatized by the requirement that Γ is nonzero and divisible as a module over $\mathbf{Z}[\frac{\sigma}{p^\infty}]$ (in case we say that Γ is *transformally divisible*). The theory $\widetilde{\omega\text{OGA}}$ is complete and o-minimal. If Γ is nonzero, then there is a model of $\widetilde{\omega\text{OGA}}$ over Γ which embeds uniquely over Γ in any other – namely, the transformal divisible hull $\Gamma \otimes \mathbf{Q}(\sigma)$.

Definition 5.3. Let Γ be a model of ωOGE. We say that Γ is *tamely transformally divisible* if we have $\nu\Gamma = \Gamma$ for all $\nu \in \mathbf{Z}[\frac{\sigma}{p^\infty}]$ with a nonzero constant term. Equivalently, the module Γ is tamely transformally divisible if and only if its inversive hull is transformally divisible.

Definition 5.4. By a *transformal valued field* we mean a valued field F equipped with an endomorphism $\sigma \colon F \to F$ of fields with the property that $\sigma^{-1}(\mathcal{O}) = \mathcal{O}$ and $\sigma^{-1}(\mathcal{M}) = \mathcal{M}$. By an *embedding* of transformal valued fields we mean an embedding of fields which is simultaneously an embedding of the underlying difference fields and of the underlying valued fields. The transformal valued field F is said to be ω-*increasing* if the induced

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

action of σ on the value group is ω-increasing; to wit, for all $0 < \alpha \in \Gamma$ and $n \in \mathbf{N}$, we have the inequality $n\alpha < \sigma\alpha$. We say that K is a *Frobenius transformal valued field* if it is of positive characteristic $p > 0$ and $\sigma = p^n$ is a power of the Frobenius endomorphism of K.

The class of ω-increasing transformal valued fields which are perfect and inversive was studied in [11]. Here, we seek to omit these assumptions.

Remark 5.5. In Definition 5.4, assume that F is inversive; then the compatibility condition on $\mathcal{M}$ is automatic. However, outside the inversive realm, the condition cannot be dispensed with.

More precisely, let $\mathcal{O}$ be a valuation ring and let σ be an *automorphism* of $\mathcal{O}$ in the category of rings. Then σ is automatically a local homomorphism (i.e, it carries non-units to non-units, or equivalently we have $\sigma^{-1}(\mathcal{M}) = \mathcal{M}$).

However, the condition that σ is an automorphism is necessary for this to hold, even when σ is injective.

For example, let $\mathcal{O}$ be a valuation ring which is abstractly isomorphic to its own localization $\widetilde{\mathcal{O}}$ at a prime ideal $\mathfrak{p}$. Thus, after relabeling $\widetilde{\mathcal{O}}$ as $\mathcal{O}$, we obtain an endomorphism σ of $\mathcal{O}$ which carries elements of positive but 'small' valuation (i.e., lying outside $\mathfrak{p}$) to elements of valuation zero; the field of fractions K when equipped with this endomorphism is *not* a transformal valued field according to Definition 5.4. For instance, the residue field does not inherit the action of σ.

5.1. The theory VFE

We will be interested in the following class of ω-increasing transformal valued fields:

Definition 5.6.

(1) Let VFE be the first order theory of ω-increasing transformal valued fields F with underlying difference field a model of FE. The language is an expansion of the language of the language of valued fields by a unary function symbol for the action of σ, so an embedding is an embedding of transformal valued fields as in 5.4 (i.e., an embedding need not be transformally separable).

(2) The definition of VFE_λ is as in 4.13; in this language, an embedding is an embedding of transformal valued fields which is required to be transformally separable.

(3) For $e \in \mathbf{N}_\infty$, we let $\mathrm{VFE}_{\leq e}$ be the theory of models of VFE of degree of imperfection at most e

(4) Finally, we write VFA for the theory models of VFE which are perfect and inversive.

In [11], the theory $\widetilde{\mathrm{VFA}}$ was studied in detail; we summarize below the main results that we will need.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Fact 5.7.

(1) The class of models of VFA admits a model companion $\widetilde{\text{VFA}}$.

(2) In $\widetilde{\text{VFA}}$, the residue field and the value group are stably embedded and fully orthogonal; the induced structure is that of a model of ACFA and a pure model of $\widetilde{\omega\text{OGA}}$, respectively.

(3) Let K be a model of VFA which is algebraically Henselian and assume that K has no nontrivial finite σ-invariant Galois extensions – that is, no difference field extensions which are finite and Galois as abstract field extensions. Then the theory of models of $\widetilde{\text{VFA}}$ over K is complete, and in particular, K admits a unique algebraic closure in the category of models of VFA, up to a (generally non unique) isomorphism. The converse holds if the various powers σ^n of σ for $0 < n \in \mathbf{N}$ are considered (that is the lack of such extensions is equivalent to the uniqueness of the lifts of the maps σ^n for all n).

(4) The theory $\widetilde{\text{VFA}}$ is precisely the asymptotic theory of algebraically closed and nontrivially valued Frobenius transformal valued fields – that is, the set of sentences true over all algebraically closed and nontrivially valued Frobenius transformal valued fields outside a finite set of exceptional prime powers.

(5) Let K be a model of VFA which is algebraically closed and transformally Henselian (see Definition 5.14). Then the category of models of VFA admits linearly disjoint amalgamation over K; that is, if L and M are models of VFA over K, then there is a model N of VFA over K in which L and M are jointly embedded over K and linearly disjoint over K. Moreover, the residue fields in the amalgamation can be taken linearly disjoint, and the transformal divisible hulls of the value groups rationally independent.

(6) Let K be an algebraically closed model of VFA and let $L = K\left(x^{\mathbf{N}[\sigma]}\right)$. Assume the residue class of x is transformally transcendental over k or else the valuation of x is not in $\Gamma \otimes \mathbf{Q}(\sigma)$. Then the inversive, algebraically Henselian hull of L has no nontrivial finite σ-invariant Galois extension. Moreover, the quantifier free type of x is determined in the first case by the fact that the residue class is transformally transcendental, and in the second case by the cut of vx over $\Gamma \otimes \mathbf{Q}(\sigma)$.

(7) Let K be a model of $\widetilde{\text{VFA}}$. Then K is transformally algebraically maximal; that is, if L is a model of VFA transformally algebraic over K with the same value group and residue field, then $K = L$. Furthermore, there is an elementary immediate extension of K which is spherically complete.

(8) The theory $\widetilde{\text{VFA}}$ admits the amalgamation property over an algebraically closed base. Thus, if K is a model of VFA, then completions of the theory $\widetilde{\text{VFA}}_K$ of models of $\widetilde{\text{VFA}}$ over K are determined by an isomorphism type of an extension of the transformal valued field structure to an algebraic closure of K and $\widetilde{\text{VFA}}$ eliminates quantifiers in the language of transformal valued fields provided that formulas ranging over the field theoretic algebraic closure are permitted.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. All references are from [11]. (1) See Theorem 9.20, (2) See Theorem 9.14, (3) See Proposition 4.29, (4) See Theorem 9.10, (5) See Theorem 9.1, (6) See Proposition 6.16, (7) See Proposition 9.9. □

Remark 5.8.

(1) Let K be a model of VFE. Then the residue field k of K need not be a model of FE. It follows from Lemma 5.9 below that VFE is not in general closed under Henselian hulls. Indeed, Example 9.1 of [13] shows that if F is a countable field, then every separably algebraic extension $\widetilde{F}$ of F occurs as the residue field of $F(x,y)$ where x,y are elements of $\widetilde{F}((\mathbf{Z}))$ algebraically independent over F. One can easily modify this example to show that if F is a countable (inversive) difference field of characteristic zero and $\widetilde{F}$ is a countable transformally algebraic extension of F, then $\widetilde{F}$ is the residue field of a difference subfield L of $\widetilde{F}((\mathbf{Z}[\sigma]))$, where L is generated by a pair of transformally independent elements over F. By 4.17, the field L is a model of VFE; since $\widetilde{F}$ is essentially arbitrary, it need not be a model of FE.

(2) Let K be a non-principal ultraproduct of Frobenius transformal valued fields. Then K is a model of VFE. For K of this form, the pathology of the previous item cannot occur since K is stationary over K^σ in the sense of ACVF; that is, for all separably algebraic valued field extensions L of K^σ, the valuation lifts uniquely to $K \otimes_{K^\sigma} L$.

We do have, however, the following:

Lemma 5.9. *Let K be a model of* VFE *which is algebraically Henselian. Then the residue field k of K is a model of* FE.

Proof. It is clear that k is closed under twists; for this, we do not need K to be algebraically Henselian. Furthermore, the embedding $K^\sigma \hookrightarrow K$ is a primary extension of Henselian valued fields; the induced extension $k^\sigma \hookrightarrow k$ of residue fields must therefore also be primary, by Hensel lifting. □

Definition 5.10. Let F be a model of VFE. We say that F is a *strict amalgamation basis* if the perfect, inversive, algebraically Henselian hull of F admits no nontrivial finite σ-invariant Galois extensions.

Remark 5.11. Let F be a model of VFE. Let us assume that F is a strict amalgamation basis. Then a separable algebraic closure of F in the category of models of VFE is uniquely determined up to isomorphism. Indeed, by [11, Proposition 4.29], an inversive, algebraically closed hull E of F is uniquely determined up to isomorphism; the relative separable algebraic closure of F in E is then a separably algebraically closed hull of F, and is likewise uniquely determined up to isomorphism. Conversely, assume that for all $0 < n \in \mathbf{N}$, a separable algebraic closure of (F,σ^n) is uniquely determined up to isomorphism; then F is a strict amalgamation basis. This is again a consequence of [11, Proposition 4.29].

Remark 5.12. Let F be a model of VFE. Let E be a model of VFE which is primary over F. Fix separable algebraic closures F^{sep} and E^{sep} of F and E, respectively, in the category of valued fields. Let us assume the following:

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

(1) The field E is a strict amalgamation basis; that is, for each n, the group $\mathrm{Aut}(E^{\mathrm{sep}}/E)$ acts transitively on the set of lifts of σ^n from E to E^{sep}.

(2) The valued field structure on the amalgamation $F^{\mathrm{sep}} \otimes_F E$ is uniquely determined.

Then F must likewise be a strict amalgamation basis. Indeed, under the assumptions, the restriction map $\mathrm{Aut}(E^{\mathrm{sep}}/E) \to \mathrm{Aut}(F^{\mathrm{sep}}/F)$ is well defined and surjective.

Condition (2) need not hold in general assuming only that E is primary over F (for example, the residue field of E might be a proper separably algebraic extension of the residue field of F, while E is still primary over F). However, it holds whenever $\mathrm{qftp}_{\widetilde{\mathrm{VFA}}}(E/F)$ admits a global quantifier free extension which does not split over F.

Remark 5.13. Let F be a model of VFE which is separably algebraically closed and let a be a tuple in an ambient extension of F. Let E be the smallest model of VFE over F containing a. Then the following conditions are equivalent:

(1) The field E is a strict amalgamation basis in the sense of Definition 5.10

(2) Fix $0 < n \in \mathbf{N}$. Then the type $\mathrm{qftp}_{\widetilde{\mathrm{VFA}}_n}\left(a, \ldots, a^{\sigma^{n-1}}/F\right)$ is complete.

Here, the notation $\widetilde{\mathrm{VFA}}_n$ refers to the theory of models of $\widetilde{\mathrm{VFA}}$ over (F, σ^n). This follows from Remark 5.11 and Fact 5.7(3).

5.2. The transformal Henselization

The following Definition will play an important role in this work.

Definition 5.14. Let K be a model of VFE. We say that K is *transformally Henselian* if it obeys the following transformal analogue of Hensel's lemma. Let $fx \in \mathcal{O}\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted difference polynomial and let $a \in \mathcal{O}$ be such that $vfa > 0$ and $vf'a = 0$; then an element $b \in \mathcal{O}$ is found with $fb = 0$ and whose residue class coincides with that of a.

In the situation of Definition 5.14, we say that b is the *canonical Hensel lift* of its residue class. Hensel lifts are unique if they exist.

Remark 5.15. Let K be a model of VFE. Then K is transformally Henselian if and only if the Frobenius twists of K all obey the transformal analogue of Hensel's lemma for ordinary difference polynomials (not necessarily twisted). In particular, for models of VFA, Definition 5.14 agrees with the Definition given in [11].

Proposition 5.16. *Let L/K be an extension of models of* VFE*. Let us assume that L is transformally Henselian and that K is transformally separably algebraically closed in L; then K is transformally Henselian. Moreover, the difference field k is transformally separably algebraically closed in l.*

Remark 5.17. Implicit in the statement of Proposition 5.16 is the assumption that the residue field of a transformally Henselian model of VFE is a model of FE. Since a transformally Henselian model is algebraically Henselian, Lemma 5.9 applies.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. Let $fx \in \mathcal{O}_K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted difference polynomial and let $a \in \mathcal{O}$ be such that $vfa > 0$ and $vf'a = 0$; we must find $b \in \mathcal{O}_K$ with $fb = 0$ and whose residue class coincides with that of a. Since L is transformally Henselian, such b can be found in $\mathcal{O}_L$. By assumption, however, the field K is transformally separably algebraically closed in L; since b is a simple root of a twisted difference polynomial with coefficients in K, the element b lies in K by Theorem 4.48(2). It follows that K is transformally Henselian.

We have seen that K is transformally Henselian; we now show that k is transformally separably algebraically closed in l. Fix an element $\alpha \in l$ which is transformally separably algebraic over k; then $g\alpha = 0$ and $g'\alpha \neq 0$ for some twisted difference polynomial $gx \in k\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$. Let $fx \in \mathcal{O}_K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted difference polynomial lifting gx. Since L is transformally Henselian, there is an element $a \in \mathcal{O}_L$ with residue class equal to α such that $fa = 0$ and $f'a \neq 0$. Since K is transformally separably algebraically closed in L, the element a lies in K; hence, its residue class lies in k, as wanted. □

Lemma 5.18. *Let L/K be a purely transformally inseparably algebraic extension of models of* VFE. *Then K and L are simultaneously transformally Henselian.*

Proof. It is enough to prove this when L is the inversive hull of K. First, assume that K is transformally Henselian; then L is the directed union of copies of K and hence transformally Henselian as the directed union of transformally Henselian models of VFE. The converse follows from Proposition 5.16. □

Definition 5.19. Let E/F be an extension of models of VFE. We say that E is a *separably algebraically closed, transformally Henselian hull* of F if the following conditions are satisfied:

(1) The field E is separably algebraically closed and transformally Henselian.
(2) Let $F \subseteq E_0 \subseteq E$ be a model of VFE which is separably algebraically closed and transformally Henselian; then $E_0 = E$.

Lemma 5.20. *Let F be a model of* VFE *which is a strict amalgamation basis. Then there is up to isomorphism a unique separably algebraically closed and transformally Henselian hull E of F. Moreover, the field E is transformally separable over F.*

Proof. Let us assume first that F is inversive. In this situation, an algebraically closed, transformally Henselian hull of F is unique up to isomorphism, using [11, Corollary 5.4]. Thus, F admits an inversive, algebraically closed, transformally Henselian hull E in the category of models of VFE which is unique up to isomorphism. Using Lemma 5.18, the relative transformal separable algebraic closure of F in E is a separably algebraically closed, transformally Henselian hull of F. □

Proposition 5.21 (Amalgamation for Models of VFE). *Let F be a model of* $\mathrm{VFE}_{\leq e}$ *which is separably algebraically closed and transformally Henselian. Let E, K be models of* $\mathrm{VFE}_{\leq e}$ *transformally separable over F. Then there is a model L of* $\mathrm{VFE}_{\leq e}$ *transformally separable over F in which K and E jointly embed over F. The fields E and K can be taken*

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

linearly disjoint over F in L, and L can be taken transformally separable over both E and K.

Proof. By Proposition 4.20 and Remark 4.22, the only problem is to equip the difference field $E \otimes_F K$ with the structure of a model of VFE subject to the evident compatibility constraints. If F is inversive, then it is algebraically closed and transformally Henselian; the result then follows from Fact 5.7. In general, by Lemma 5.18, the inversive hull of F is algebraically closed and transformally Henselian; since E and K are transformally separable over F, the truth of the statement with all fields replaced by their inversive hulls descends to the original settings, whence the proposition. □

Proposition 5.22. *Let K be a model of* VFE *whose residue field k is a model of* FE*. Let us assume that we are given a model $\widetilde{k}$ of* FE *transformally separably algebraic over k. Then there is a transformally Henselian model $\widetilde{K}$ of* VFE *transformally separably algebraic over K, reproducing the embedding $\widetilde{k}/k$ at the level of residue fields, and enjoying the following universal mapping property. Let L be a transformally Henselian model of* VFE *over K; then every embedding of $\widetilde{k}$ in l over k lifts to an embedding of $\widetilde{K}$ in L over K, and the lifting is unique. This extension does not enlarge the value group. Furthermore, if K is a strict amalgamation basis and $\widetilde{k}$ is separably algebraically closed, then $\widetilde{K}$ is a strict amalgamation basis.*

Proof. When K is inversive, this follows from [11, Theorem 5.2]. In general, the usual descent argument shows that the relative transformal separable algebraic closure of K inside the promised extension of the inversive hull will do. □

Definition 5.23. Let K be a model of VFE.

(1) We say that K is *strictly transformally Henselian* if it is transformally Henselian, and the residue field k of K is a model of SCFE.

(2) In the situation of Proposition 5.22, assume that $\widetilde{k}$ is a model of SCFE; then $\widetilde{K}$ is said to be a *strict transformal Henselization* of K (with respect to the embedding $\widetilde{k}/k$).

Remark 5.24. Let L/K be a purely transformally inseparably algebraic extensions of models of VFE. Then L is strictly transformally Henselian if and only if K is so; this follows immediately from Proposition 4.53 and Proposition 5.16.

5.3. The theory $\widetilde{\text{VFE}}$

Definition 5.25. Let K be a model of VFE. We say that K is a model of $\widetilde{\text{VFE}}$ if the following requirements are satisfied:

(1) The field K is strictly transformally Henselian in the sense of Definition 5.23.

(2) The valuation group Γ of K is nonzero and *tamely transformally divisible* (i.e., for all $\nu \in \mathbf{Z}\left[\frac{\sigma}{p^\infty}\right]$ with nonzero constant term, we have $\nu\Gamma = \Gamma$).

(3) Let $\tau x \in K\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$ be a twisted additive difference operator with $\tau' \neq 0$; then τ is onto on K-points.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

For $e \in \mathbf{N}_\infty$, we let $\widetilde{\mathrm{VFE}}_e$ be the theory of models of $\widetilde{\mathrm{VFE}}$ which fail to be inversive, of imperfection degree equal to e.

Remark 5.26. It will follow from Lemma 5.29 that if K is a model of $\widetilde{\mathrm{VFE}}$, then it lies dense in its inversive hull; so the residue field is in fact a model of ACFA, and the value group is a model of $\widetilde{\omega\mathrm{OGA}}$, using Remark 5.24.

Proposition 5.27.

(1) *Let K be a model of* VFE*. Then K is a model of $\widetilde{\mathrm{VFE}}$ if and only if the inversive hull of K is a model of $\widetilde{\mathrm{VFA}}$. Thus, if L is a model of* VFE *purely transformally inseparably algebraic over K, then L is a model of $\widetilde{\mathrm{VFE}}$ if and only if K is so.*

(2) *Let us assume that K is nontrivially valued and relatively transformally separably algebraically closed in a model of $\widetilde{\mathrm{VFE}}$; then K is itself a model of $\widetilde{\mathrm{VFE}}$.*

Proof. Clear using Remark 5.24 and Proposition 4.41. □

Definition 5.28. Let F be a model of VFE. We say that F is *deeply transformally ramified* if is nontrivially valued and lies dense in its inversive hull.

Lemma 5.29. *Let F be a model of* VFE *which is nontrivially valued. Let us assume that for all nonzero $t \in \mathcal{M}$, the operator $x^\sigma - tx$ is surjective on F-points. Then F is deeply transformally ramified. In particular, if F is a model of $\widetilde{\mathrm{VFE}}$, then F is deeply transformally ramified.*

Proof. By induction, it is enough to prove that F lies dense in $F^{\sigma^{-1}}$. Fix an element $a \in F$; we need to show that $a^{\sigma^{-1}}$ can be approximated to arbitrary precision by elements of F. Let $fx = x^\sigma - a$ be the displayed difference polynomial. Then fx can be regarded as the limit of the difference polynomials $f_t(x) = x^\sigma - tx - a$ for $t \in \mathcal{M}$ nonzero in the sense of the valuation topology. Namely, by assumption, we can find $b_t \in F$ with $f_t b_t = 0$; since $f_t \to f$ coefficient-wise and nearby polynomials have nearby roots, as $vt \to \infty$, we will have $vfb_t \to \infty$, whence the statement. □

Proposition 5.30. *Let K be a saturated model of $\widetilde{\mathrm{VFE}}$. Then K is existentially closed for transformally separably algebraic extensions in the following sense. Let F be a strictly amalgamative model of* VFE *contained in K and such that K is transformally separable over F. Let E be a small model of* VFE *transformally separably algebraic over F. Then there is an embedding of E in K over F. Moreover, such an embedding automatically renders K transformally separable over the image.*

Proof. If an embedding exists, then it automatically renders K transformally separable over the image, using Theorem 4.48(3). It is therefore enough to find an embedding. By compactness, this means that if $a \in E$ is a finite tuple and φ is a quantifier free formula in the language of transformal valued fields with parameters in F such that $\models \varphi(a)$, then φ has a solution in K.

For simplicity of notation, we assume that φ has a single free variable. Replacing K by the relative transformal separable algebraic closure of F inside, we may assume that K

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

is transformally separably algebraic over F. So let E be a model of VFE transformally separably algebraic over F and let $a \in E$ be an element such that $\models \varphi(a)$. Then a is a simple root of a twisted difference polynomial $fx \in F\left[x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right]$; strengthening φ we may assume that all solutions of φ are simple roots of fx.

Since F is strictly amalgamative, the theory of models of $\widetilde{\text{VFA}}$ over F is complete (Fact 5.7); the inversive hull of K is a model of $\widetilde{\text{VFA}}$, so by model completeness, the formula φ has a solution over the inversive hull of K. But the simple roots of f over K and over its inversive hull coincide, using Theorem 4.48; thus, there is a solution in K. □

6. Genericity in a ball

In this section, we fix a large enough saturated model $\mathbf{K}$ of $\widetilde{\text{VFA}}$; all transformal valued fields in this section live inside $\mathbf{K}$.

6.0.1. The model theoretic settings

We work in the language of transformal valued fields – that is, the expansion of the language $\mathcal{L}_{\text{VF}}$ by a unary function symbol for the action of σ. Let VF denote the valued field sort. By a *definable set* we mean definable in the language of transformal valued fields.

6.0.2. Split and definable balls

Let $B \subseteq \text{VF} = \text{VF}^1$ be a definable set, in one variable. We say that B is a *closed ball* if $B = a + \gamma\mathcal{O}$ for some $a \in \text{VF}$ and $\gamma \in \Gamma$. We say that B is an *open ball* if $B = \text{VF}$ or $B = a + \gamma\mathcal{M}$ for some $\gamma \in \Gamma$ and $a \in \text{VF}$. By a *possibly degenerate ball* we mean either a ball or a singleton. Finally, by an ∞-definable ball we mean the intersection of a (small) family of balls linearly ordered under reverse inclusion, regarded as a partial quantifier free type.

Let $F \subseteq \mathbf{K}$ be a small model of VFE contained in $\mathbf{K}$. Let $B \subseteq \text{VF}$ be a ball. We say that B is *definable* over F if it is definable with parameters in F and in the language of transformal valued fields or equivalently if it is setwise fixed by automorphisms of $\mathbf{K}$ over F. If B is an ∞-definable ball, then we say that it is ∞-definable over F if it is the intersection of balls definable over F.

Remark 6.1. Let F be a model of VFE which is separably algebraically closed and transformally Henselian. The theory of models of $\widetilde{\text{VFA}}$ over F is complete using Fact 5.7, so the notion of a ball definable over F makes sense intrinsically, independently of the choice of $\mathbf{K}$. In order to avoid ambiguity, we will only discuss definable balls over a model of VFE which is separably algebraically closed and transformally Henselian.

Let $B \subseteq \text{VF}$ be a ball. We say that B is *split* over F if it takes the form $a + \gamma\mathcal{O}$ for some $a \in F$ and $\gamma \in \Gamma_F$. If B is split over F, then it is evidently definable over F, but the converse is of course false. Nevertheless we have the following:

Proposition 6.2. *Let F be a model of $\widetilde{\text{VFE}}$. Then all F-definable balls are split over F.*

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. Let us assume first that F is inversive. Then F is a model of $\widetilde{\text{VFA}}$; by model completeness, every F-definable ball has a point over F, and in particular, all F-definable balls are split over F. For the general case, let E be the inversive hull of F. Then every element of E is definable over F, so there is no distinction between F-definable balls and E-definable balls. Moreover, by Lemma 5.29, the field F lies dense in E; so a ball is split over F if and only if it is split over E; this reduces us to the inversive case already dealt with, and the proof is finished. □

6.0.3. Genericity in a ball

Lemma 6.3. *Let F be a model of* VFE *contained in* **K**. *Let $B \subseteq \text{VF}$ be a closed ball definable over F. Then B enjoys the following compactness property. Let $fx \in F\left[x^{\mathbf{N}[\sigma]}\right]$ be a difference polynomial; then the function $x \mapsto vfx$ admits a minimum on B.*

Proof. If F is increased, then the conclusion descends; so assume that F is a model of $\widetilde{\text{VFA}}$. Then the ball B is split over F; hence, there is an affine map φ with coefficients in F which carries B to the unit ball $\mathcal{O}$. Replacing f by $f \circ \varphi$, we may therefore assume that $B = \mathcal{O}$. If $f = 0$, then the claim is trivial, so assume that $f \neq 0$. Replacing f by a scalar multiple shifts the function $x \mapsto vfx$ by an additive constant, so we may assume that f has coefficients in $\mathcal{O}_F$ and at least one of the coefficients is of valuation zero. Then $vfa \geq 0$ for all $a \in \mathcal{O}$, with equality in the case where the residue class of a is transformally transcendental over k_F. □

Definition 6.4. Let F be a model of VFE which is separably algebraically closed and transformally Henselian. Let B be an F-definable ball or a properly infinite intersection. Let a be an element of B.

(1) Let us assume that B is closed. Then a is *generic* in B over F if for all difference polynomials $fx \in F\left[x^{\mathbf{N}[\sigma]}\right]$, the quantity vfa is the generic value of vfx on the ball B (i.e., the minimum, which is attained as explained in in Lemma 6.3).

(2) Let us assume that B is open or a properly infinite intersection; then a is *generic* in B over F if it avoids U, whenever $U \subseteq B$ is a (possibly degenerate) closed subball definable over F.

Proposition 6.5. *Let F be a model of* VFE *which is separably algebraically closed and transformally Henselian. Let B be an F-definable ball or a properly infinite intersection. Then genericity in B over F determines a complete, consistent quantifier free type over F in the language of transformal valued fields, a realization of which is transformally transcendental over F. Furthermore, let a be generic in B over F and set $E = F\left(a^{\mathbf{N}\left[\frac{\sigma}{p^{\infty}}\right]}\right)$; then E is a model of* VFE *which is a strict amalgamation basis.*

Proof. By Remark 4.17, the field E is a model of VFE. Now [11, Corollary 6.19] essentially gives the result. However, it is not directly quotable in the present setting for the following reasons:

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

(1) The paper [11] deals exclusively with the case of models of VFA which are perfect and inversive, whereas here we deal with models of VFE. Thus, it does not apply to E or F directly, but rather to their perfect, inversive hulls.

(2) The definition of a 'generic type' in [11, Definition 6.13] is slightly more restrictive compared to the definition of generic considered here. More precisely, a generic type in the sense of [11] corresponds to the generic type of a *split* ball, while the statement here refers to definable (or ∞-definable) balls.

Now (1) is completely formal: the field E is a strict amalgamation basis if and only if the inversive, perfect hull of E is a strict amalgamation basis. For (2), using Remark 5.13 (and the notation there), the statement here refers to the completeness of a type in the theory $\widetilde{\mathrm{VFA}}$ (and the various reducts $\widetilde{\mathrm{VFA}}_n$ for $0 < n \in \mathbf{N}$ varying). The completeness of a type can be checked after base extension; replacing F by a sufficiently saturated model of $\widetilde{\mathrm{VFA}}$, we can assume the ball B splits over F. In this situation, [11, Corollary 6.19] applies. □

Proposition 6.6. *Let F be a model of $\widetilde{\mathrm{VFE}}$. Let E be a model of* VFE *transformally separable over F. Let us assume that k_F is transformally separably algebraically closed in k_E; then F is transformally separably algebraically closed in E. In fact, every element $a \in E$ not in F is generic over F in an F-definable ball or a properly infinite intersection.*

Remark 6.7. In the situation of Proposition 6.6. Recall that we work in an ambient $\widetilde{\mathrm{VFA}}$; so 'definable' means in the sense of this ambient model. Fix an element $a \in E$ not in F. Note that a is not definable over F (i.e., it is not in any degenerate F-definable subball): the only elements definable over F and not in F are the elements of the inversive hull (as the latter is a model of $\widetilde{\mathrm{VFA}}$), whereas by assumption, the field E is transformally separable over F. Also, since F lies dense in its inversive hull, the difference field k_F is inversive, so the assumption implies that no element of k_E is transformally algebraic over k_F, unless it already lies in k_F.

Proof. Let $\mathcal{B}$ be the family of all F-definable balls containing a, and B the intersection of the members of $\mathcal{B}$. If $\mathcal{B}$ has a maximal element under reverse inclusion, which is open, or else no minimal element under reverse inclusion, then a is generic over F in B, using Remark 6.7. We may therefore assume that B is closed. Since $F \models \widetilde{\mathrm{VFE}}$, the ball B is split over F, so without loss, we have $B = \mathcal{O}$. By Remark 6.7, the residue class of a must then be transformally transcendental over k_F, which precisely means that a is generic over F in $\mathcal{O}$. □

7. The model companion

The purpose of this section is to prove that the theory $\widetilde{\mathrm{VFE}_e}$ is the model companion of $\mathrm{VFE}_{\leq e}$ in an expansion of the language of transformal valued fields by the transformal λ-functions.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

7.1. Density of generic types and deep ramification

In this subsection, we prove a key statement towards the proof of model completeness of $\widetilde{\mathrm{VFE}_e}$. Before we turn to the statement and its proof, let us explain the motivation.

Fix a saturated model K of $\widetilde{\mathrm{VFE}_e}$. Let us consider definable subsets of K in one variable. There are at least two notions of largeness for such definable sets. The first notion of largeness comes from the valuation topology: if $B \subseteq K$ is an open ball, then a definable subset X of B is large if it is not contained in any proper definable subball. The second notion of largeness comes from the failure of σ to be surjective: the field K is an infinite dimensional vector space over K^σ, so a definable subset of K is large if it is not contained in any finite dimensional K^σ-subspace of K.

The main observation is that these two notions of genericity are not in conflict with each other, in the sense made precise below:

Proposition 7.1. *Let K be a saturated model of $\widetilde{\mathrm{VFE}_e}$. Let $F \subseteq K$ be a small model of* VFE *contained in K, separably algebraically closed and transformally Henselian. Let us assume that K is almost transformally separable over F. Let B be an F-definable ball or a properly infinite intersection of F-definable balls. Let $E = F\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$, where a is generic over F in B.*

(1) *Then there is an embedding of E in K over F so as to render K almost transformally separable over the image.*

(2) *Let us assume that K is transformally separable over F outright and that the perfect hull of K does not split over F; then there is an embedding of E in K over F so as to render K transformally separable over the image.*

Proof. We prove the first point; the second point is similar, using the fact that K^p lies dense in K. By compactness, we may assume that B is an honest ball, rather than an ∞-definable ball. Moreover, we deal with the case of a closed ball; the open case is similar.

Using Proposition 4.54, our task is to find an element $x \in K$ simultaneously obeying the following two properties. First, we want x to be algebraically transcendental over $F \otimes_{F^{\sigma \cdot p^{-\infty}}} K^{\sigma \cdot p^{-\infty}}$. Second, we want x to be generic over F in B.

Increasing F only makes these tasks more difficult, so assume F is an elementary substructure of K. Then F is a model of $\widetilde{\mathrm{VFE}}$, and K is transformally separable over F. By Proposition 6.2, the ball B splits over F. We may therefore assume without loss of generality that $B = \mathcal{O}$.

For the sake of the argument, say that $y \in K$ is *topologically generic* over F if $vy = 0$ and the residue class is transformally transcendental over k_F (i.e., it is generic over F in $\mathcal{O}$). Moreover, say that $z \in K$ is *λ-generic* over F if it is algebraically transcendental over $F \otimes_{F^{\sigma \cdot p^{-\infty}}} K^{\sigma \cdot p^{-\infty}}$. Finally, say that x is *generic* over K if it is simultaneously λ-generic and topologically generic over F. Our task, then, is to find an element of K which is generic over F.

Claim 1. The set X_v of elements of K which are topologically generic over F is nonempty. Indeed, it contains an open ball.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. The fact that X_v is nonempty follows from the fact that k is a saturated model of ACFA. Given that it is nonempty, it is by definition the union of residue classes, and hence open. □

Claim 2. The set X_λ of elements of K which are λ-generic over F is nonempty. Indeed, it lies dense in K for the valuation topology.

Proof. Since K fails to be inversive, it follows from saturation that X_λ is nonempty; see Lemma 4.55. It is clear that X_λ is stable under affine translation with coefficients in K^σ. By Lemma 5.29, the field K^σ lies dense in K for the valuation topology; the same must then be true of X_λ.

Using the notation of the above claims, the set X_v contains an open ball, and the set X_λ is dense for the valuation topology; the intersection $X_v \cap X_\lambda$ is therefore nonempty. By definition, every element of the intersection is generic over F, whence the proposition. □

□

7.2. The model companion

Proposition 7.2. *Let F be a model of $\mathrm{VFE}_{\leq e}$. Then there is a model of $\widetilde{\mathrm{VFE}_e}$ transformally separable over F.*

Proof. We work inside an ambient model $\mathbf{K}$ of $\widetilde{\mathrm{VFA}}$ over F. We gradually replace F by a model E of $\mathrm{VFE}_{\leq e}$ inside $\mathbf{K}$, transformally separable over F, so as to make sure that all the axioms are satisfied. Basic closure properties of the category of models of FE and elementary properties of transformally separable extensions will be used without mention.

Step 1. We may assume that F fails to be inversive. Indeed, in this case, the field F must also be perfect. Let $x \in \mathbf{K}$ be transformally transcendental over F and let $E = F\left(x^{\mathbf{N}[\sigma]}\right)^{\mathrm{alg}}$ be the displayed transformal valued field. Then E is a perfect model of VFE which by construction fails to be inversive.

Step 2. We may assume that F has the correct Ershov invariant (i.e., that $[F\colon F^p] = p^e$). Indeed, let $x = (x_1, \ldots, x_d) \in \mathbf{K}$ be a tuple transformally transcendental over over F and set $E = F\left(x^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$. Then x is a relative p-basis of E over F; hence, $[E\colon E^p] = [F\colon F^p] \cdot p^d$. In this way, we can increase the Ershov invariant which settles the case $e < \infty$; in the case $e = \infty$ take a directed union.

Step 3. We may assume that F is nontrivially valued. Indeed, let $x \in \mathbf{K}$ be transformally transcendental over F with $vx > 0$. Then $E = F\left(x^{\mathbf{N}\left[\sigma^{\pm 1}, p^{\pm 1}\right]}\right)$ is a model of VFE transformally separable over F; by construction, it is relatively perfect over F and hence a model of $\mathrm{VFE}_{\leq e}$ as well.

Now let E be the relative transformal separable algebraic closure of F in $\mathbf{K}$. Then E is nontrivially valued and transformally separably algebraically closed in $\mathbf{K}$; by Proposition 5.27, the field E is a model of $\widetilde{\mathrm{VFE}}$. By Proposition 4.34, the perfect hull of E splits over F, so the imperfection degrees are the same; and of course E fails to be inversive if F does. Thus, E is a model of $\widetilde{\mathrm{VFE}_e}$, and we are done. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Theorem 7.3. *Let K be a saturated model of $\widetilde{\mathrm{VFE}}_e$. Then K is existentially closed for transformally separable extensions of the same Ershov invariant in the following sense. Let $F \subseteq K$ be a small strictly amalgamative model of $\mathrm{VFE}_{\leq e}$ with K transformally separable over F. Let us assume that we are given a small model E of $\mathrm{VFE}_{\leq e}$ which is transformally separable over F. Then there is an embedding of E in K over F so as to render K transformally separable over the image.*

Before we turn to the proof of Theorem 7.3, we wish to make several elementary remarks.

Remark 7.4. In the situation of Theorem 7.3, assume that the perfect hull of E does not split over F. Then the perfect hull of K does not descend to F either. Indeed, a transformally separable extension of models of FE is also algebraically separable. In the case $e = \infty$, the vector space K is of large dimension over K^p; the inclusion $F \otimes_{F^p} K^p \subseteq K$ must therefore be proper as F is small. If $e < \infty$, then the inequality $\dim_{E^p} E \leq \dim_{K^p} K$ holds by assumption, and again we are in the clear.

Remark 7.5. Let us assume that L is a model of $\mathrm{VFE}_{\leq e}$ transformally separable over E. Suppose we are given given an embedding of L in K over F so as to render K transformally separable over the image. By abuse of notation, let us regard L – and hence E – as a transformal valued subfield of K over F. Then as K/L and L/E are transformally separable, the field extension K/E is transformally separable as the composition of such. It is therefore harmless to increase E by a transformally separable extension; using Proposition 7.2, we may therefore assume that E is a model of $\widetilde{\mathrm{VFE}}_e$, and in particular that it is separably algebraically closed and transformally Henselian.

Remark 7.6. We may assume, of course, that E properly extends F, for otherwise, the claim is trivial. Now assume further $F \subseteq E_0 \subseteq E$ is a model of VFE which is strictly amalgamative and that E is transformally separable over E_0. If E_0 properly extends F, then the truth of the statement for E/E_0 and E_0/F separately implies it for E/F. By transfinite induction and compactness, we can therefore assume that no such E_0 can be found, unless it coincides with E.

Proof of Theorem 7.3. Let F, K and E be as in the statement. By Remark 7.5, it is harmless to assume that E is a model of $\widetilde{\mathrm{VFE}}_e$.

We first reduce ourselves to the case where F is transformally separably algebraically closed in E. Indeed, let $F \subseteq E_0 \subseteq E$ be the relative transformal separable algebraic closure of F in E. The field E_0 is separably algebraically closed, as it is relatively separably algebraically closed inside the separably algebraically closed field E; in particular, the field E_0 is a strict amalgamation basis. By Proposition 5.30, the models of $\widetilde{\mathrm{VFE}}_e$ are existentially closed for transformally separably algebraic extensions, so there is an embedding of E_0 in K over F, and such an embedding automatically renders K transformally separable over the image. Thus, inductively using Remark 7.6, we may assume that F is transformally separably algebraically closed in E.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

By Proposition 5.27(2), F is a model of $\widetilde{\mathrm{VFE}}$, and so using Proposition 5.22, Proposition 6.6 implies that an element $a \in E$ lying outside F is generic over F in an F-definable ball or a properly infinite intersection of F-definable balls.

Next, we claim that E can be taken to be a relatively perfect extension of F. If E is not relatively perfect over F, fix an element $a \in E$ not in $F \otimes_{F^p} E^p$ and set $E_0 = F\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$. Then a is generic over F in a ball, so by Proposition 6.5, the field E_0 is a strict amalgamation basis. By Proposition 4.54, the field E is transformally separable over E_0. Using Remark 7.6, we can replace E by E_0; in this situation, however, Proposition 7.1 on the density of the generic types applies.

Now E is relatively perfect over F. In this situation, an embedding of E in K over F will *automatically* render K algebraically separable over the image by Fact 3.10; using the criterion of Proposition 4.32, it is therefore enough to exhibit an embedding of E in K over F with K almost transformally separable over the image. We now turn to this task. In order to achieve it, we reinstall our settings and weaken our assumption to allow ourselves to argue inductively. Thus, we forsake the assumption that K is transformally separable over F, and merely assume that it is almost transformally separable over F. Moreover, we do not make any assumptions on the degree of imperfection of E, and allow E to be an arbitrary model of VFE transformally separable over F.

By compactness and finite character of almost transformally separable extensions, we may assume that E is transformally separably generated over F, as Theorem 4.48. Fix an element $a \in E$ forming a part of a separating transformal transcendence basis of E over F; then E is transformally separable over E_0, where $E_0 = F\left(a^{\mathbf{N}\left[\frac{\sigma}{p^\infty}\right]}\right)$. The field E_0 is strictly amalgamative; arguing as in Remark 7.6 and using the fact that almost transformally separable extensions are closed under composition, we may assume that $E = E_0$. But Proposition 7.1(2) applies, and the proof is finished. □

Corollary 7.7. *Let F be a model of $\mathrm{VFE}_{\leq e}$ which is a strict amalgamation basis. Then there is up to elementary equivalence a unique model of $\widetilde{\mathrm{VFE}_e}$ which is transformally separable over F.*

Proof. This follows immediately from Theorem 7.3 and Proposition 7.2. □

Let $\mathcal{L}_\lambda$ be the language of models of VFE with the transformal λ-functions. Thus, for $K \models \widetilde{\mathrm{VFE}_\lambda}$, an $\mathcal{L}_\lambda$-substructure of K is a model F of VFE contained in K such that K is transformally separable over F.

Corollary 7.8. *In the theory $\widetilde{\mathrm{VFE}_\lambda}$, every formula is equivalent to a boolean combination of formulas of the form $\exists y \varphi(x,y)$ where $\varphi \in \mathcal{L}_\lambda$ is quantifier free and each y_i obeys a monic (algebraic) polynomial with coefficients in $\mathbf{Z}[x]$.*

Proof. Let K be a saturated model of $\widetilde{\mathrm{VFE}_\lambda}$. By definition, an $\mathcal{L}_\lambda$-substructure of K is a model F of VFE contained in K such that K is transformally separable over F. Let $\mathcal{S}$ be the set of formulas in the statement. We are given an isomorphism $f\colon F \cong F'$ between models of VFE contained in K, with K transformally separable over both, and that the isomorphism preserves the truth value of formulas in $\mathcal{S}$; we must show that f lifts to a

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

global automorphism. Now using formulas in $\mathcal{S}$, we can determine the isomorphism type of the action of σ on the separable algebraic closure of both F and F' inside K (see, for example, the proof of [11, Proposition 9.17]). Thus, we may extend f to the relative separable algebraic closures inside K. By elementary properties of transformally separable extensions, this does not disturb the property that K is transformally over F, and likewise F'; and now Corollary 7.7 applies. □

Theorem 7.9. *Let $T = \mathrm{VFE}_{\leq e}$ and $\widetilde{T} = \widetilde{\mathrm{VFE}}_e$. Then $\widetilde{T}_\lambda$ is the model companion of T_λ.*

Proof. This follows immediately from Corollary 7.7 and Proposition 7.2. □

Theorem 7.10. *Let K be a model of $\widetilde{\mathrm{VFE}}_e$ and let $F \subseteq K$ be a model of* VFE *which is contained in K. Then F is model theoretically algebraically closed in K if and only if it is separably algebraically closed, transformally Henselian, and closed under the transformal λ-functions of K.*

Proof. We may assume that K is sufficiently saturated and that F is small.

Recall that K is transformally separable over F if and only if F is closed under the transformal λ-functions of K, and that the transformal λ-functions of K are definable in the language of difference fields. With this description, it is clear that the conditions are necessary. For the converse, assume the conditions of the theorem, and fix an element $a \in K$ not in F; we will show that it admits infinitely many conjugates over F.

We may assume that a enumerates a small separably algebraically closed and transformally Henselian model E of VFE over F, closed under the transformal λ-functions of K; so K is transformally separable over E. By model completeness, Proposition 5.21 and saturation, there is a sequence $(E_n)_{n=0}^{\infty}$ of linearly disjoint copies of $E = E_0$ over F inside K all closed under the transformal λ-functions of K. By linear disjointness, we have $E_n \cap E_m = F$ for $n \neq m$. By Corollary 7.7, the isomorphism $E_n \cong E_m$ for $n \neq m$ lifts to a global automorphism over F. This concludes the proof. □

Lemma 7.11. *Let K be a model of $\widetilde{\mathrm{VFE}}_e$ which is ω-saturated and let $F = K^{\sigma^\infty}$ be the maximal inversive difference subfield of K. Then F is a model of $\widetilde{\mathrm{VFA}}$ which lies dense in K.*

Proof. Since K lies dense in its inversive hull, it follows by compactness and saturation that F lies dense in K; in particular, it is nontrivially valued. The transformal algebraic closure of F in K is inversive by Proposition 4.34, so by maximality, it is relatively transformally algebraically closed in K. Thus, by Proposition 5.27, the field F is a model of $\widetilde{\mathrm{VFA}}$. □

Remark 7.12. The field F is not stably embedded in K. Indeed,

Claim 1. Every element of K is the unique realization of its type over F in K; thus, in the language of [17], every imaginary of K is F-comprehended.

Proof. The field F lies dense in K for the valuation topology, so if $a \in K$ lies outside F, it is the unique element of K inside the intersection of all F-definable balls containing a. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Claim 2. No element of K is definable over F unless it already lies in F.

Proof. This follows from the characterization of the model theoretic algebraic closure in $\widetilde{\mathrm{VFE}}$ above (Theorem 7.10). By [17], a sort D in a theory T is stably embedded if and only if every imaginary which is D-comprehended is definable over D; the embedding of F in K provides an extreme example of the opposite scenario. Thus, F is not stably embedded in K. □

In the following, we wish to understand the induced structure on the residue field and on the value group. For the latter, recall the theory $\widetilde{\omega\mathrm{OGA}}$ of ω-increasing transformally divisible ordered abelian groups from [11, Section 2] It is complete and admits quantifier elimination.

Theorem 7.13. *In the theory $\widetilde{\mathrm{VFE}_e}$, the residue field and the value group are stably embedded and fully orthogonal. The induced structure is the pure difference field structure and the pure ordered transformal module structure, respectively.*

Proof. Let K be a sufficiently saturated model of $\widetilde{\mathrm{VFE}_e}$. We first verify that the residue field and value group are stably embedded and fully orthogonal as a difference field structure and an ordered transformal module structure, respectively, but possibly with some added constants. After doing that, we show that indeed the $\emptyset$-induced structure is that of pure structures.

It will be sufficient to prove the following: given a strict amalgamation basis $F \models \mathrm{VFE}_{\leq e}$, with K/F transformally separable, a tuple of residue elements α and a tuple of value group elements γ, we have

$$\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_F) \cup \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_F) \vdash \mathrm{tp}(\alpha,\gamma/F).$$

We show this by proving that any pair of automorphisms of the residue field and value group (as a difference field and an ordered transformal module) lifts to an automorphism of K by the aid of Theorem 7.3. To this end, it will be sufficient to show that these automorphisms arise as an automorphism (over F) of some model E of VFE_e extending F, which is transformally separable over F.

We build this extension E step by step by repeatedly taking care of the following kind of elements: (1) elements constituting γ, (2) elements of α which are transformally transcendental over k_F and (3) elements of α which are transformally algebraic over k_F. Starting with $E = F$, at each step, we will extend E and the automorphism on it to induce the desired automorphism on the residue field and value group; moreover, when we take care of residue elements, we will not extend the value group, and vice-versa.

(1) Let $\gamma_1,\gamma_2 \in \Gamma_K$ (singletons) be value group elements with $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma_1/\Gamma_F) = \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma_2/\Gamma_F)$ and γ_i not in the transformal divisible hull of Γ_F. For $i = 1,2$, let p_{γ_i} be the quantifier-free type of a generic in the closed ball of radius γ_i around 0, as supplied by Proposition 6.5. By density of K^σ in K, Lemma 5.29, the type p_{γ_i} is consistent with the maximal inversive subfield K^{σ^∞} of K. Let $a_i \in K^{\sigma^\infty}$ be a realization of p_{γ_i} inside the maximal inversive subfield of K. It follows from the proof of Proposition 6.5 that the

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

field $E_i = F\big(a_i^{\mathbf{N}[\sigma^{\pm 1}, p^{\pm 1}]}\big)$ is strictly amalgamative. It is obviously a model of VFE_e and as $(E_i)^{\sigma^{-\infty}} = F^{\sigma^{-\infty}}\big(a_i^{\mathbf{N}[\sigma^{\pm 1}, p^{\pm 1}]}\big)$, K is transformally separable over E_i. Finally, since $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma_1/\Gamma_F) = \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma_2/\Gamma_F)$, the fields E_1 and E_2 are isomorphic as difference valued fields over F. Note that by construction, Γ_{E_i} lies in the transformal divisible hull of $\Gamma_F \cup \{\gamma_i\}$ and that $k_{E_i} = k_F$.

If the γ_i are in the transformal divisible hull of Γ_F, then they are in its definable closure, and so are automatically preserved.

(2) Let $\alpha_1, \alpha_2 \in k_K$ (singletons) with $\mathrm{tp}_{\mathrm{ACFA}}(\alpha_1/k_F) = \mathrm{tp}_{\mathrm{ACFA}}(\alpha_2/k_F)$ and α_i transformally transcendental over k_F. We now proceed similarly as in (1), realizing the generic of the closed ball of radius 0 around 0. Note that, in the notation of (1), we get $k_{E_i} = k_F(\alpha_i^{\mathbf{N}[\sigma^{\pm 1}, p^{\pm 1}]})$ and $\Gamma_{E_i} = \Gamma_F$.

(3) Let $\alpha_1, \alpha_2 \in k_K$ be singletons with $\mathrm{tp}_{\mathrm{ACFA}}(\alpha_1/k_F) = \mathrm{tp}_{\mathrm{ACFA}}(\alpha_2/k_F)$ and α_i are in the relative transformal algebraic closure of k_F in k_K. Let $\widetilde{k_F}$ be the relative transformal separable algebraic closure of k_F in k_K. By [11, Lemma 3.6], the α_i are in the inversive hull of $\widetilde{k_F}$. Hence, there exists a twisted difference polynomial over k_F and some natural number n such that $\alpha_1^{\sigma^n}$ and $\alpha_2^{\sigma^n}$ are both simple roots of this twisted difference polynomial. Since $\alpha_i^{\sigma^n}$ and α_i are interdefinable, there is not harm in the replacing the latter with the former. So assume that the α_i are simple roots of twisted difference polynomials over k_F.

By [11, Proposition 9.13], the difference field $K^{\sigma^{-\infty}}$ is saturated in κ^+ if K is. Indeed, their residue fields and value groups are equal, respectively, and a nested collection of balls of cofinality κ in $K^{\sigma^{-\infty}}$ have centers in K by Lemma 5.29, so since they have nonempty intersection in K, it follows for $K^{\sigma^{-\infty}}$ as well. As a result, $K^{\sigma^{-\infty}}$ is saturated as well. By stable embeddedness of k in $\widetilde{\mathrm{VFA}}$ [11, Theorem 9.14], there exists an automorphism ρ of $K^{\sigma^{-\infty}}$ over F which maps α_1 to α_2. By Theorem 4.48, ρ restricts to $\widetilde{F}$, the relative transformal separable algebraic closure of F inside $K^{\sigma^{-\infty}}$. It is a model of $\widetilde{\mathrm{VFE}}$ by Proposition 5.27. In particular, it is strictly amalgamative and K is transformally separable over F. Finally, by model completeness, $\widetilde{F} \prec K$, and thus since $\widetilde{F}$ is relatively transformally separably algebraically closed in K, the same is true for $k_{\widetilde{F}}$ in k; so $\alpha_i \in k_{\widetilde{F}}$.

We now verify the induced structure is precisely the pure difference field and ordered module structure; that is, we prove that $\emptyset$-definable subsets of k^d definable in the language of VFE are already $\varnothing$-definable in the language of difference fields and likewise for Γ. For this, we show that given small models $k_0 \prec k_K$ and $\Gamma_0 \prec \Gamma_K$, any pair of automorphism of k_0 and Γ_0 lifts to an automorphisms of K.

Given such models k_0 and Γ_0, the Hahn series $K_0 = k_0((t^{\Gamma_0}))$ is a model of $\widetilde{\mathrm{VFA}}$ (see for example the proof of Proposition 9.14 in [11]). We claim that K_0 embeds in K.

Let us assume for a moment that this was shown and finish the proof. The map $\mathrm{Aut}(K_0) \to \mathrm{Aut}(k_0) \times \mathrm{Aut}(\Gamma_0)$ is evidently surjective (indeed, it even admits a section). The field K_0 is algebraically closed and inversive, and hence a strict amalgamation basis; using Corollary 7.7, every automorphism of K_0 lifts to an automorphism of K.

It remains to be shown that K_0 can be embedded in K (over the prime field – that is, over $\varnothing$).

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Let $\mathbf{F}$ be the algebraic closure of the prime field of k_0. By Hensel's lemma, it identifies with the algebraic closure of the prime field in K_0. Since $\mathbf{F}$ is a strict amalgamation basis, by Theorem 7.3, we may embed K_0 in K over $\mathbf{F}$, and we are done. □

Corollary 7.14. *Let K be a model of $\widetilde{\mathrm{VFE}}_e$, and let $\tau = \sigma^n p^m$ with $n,m \in \mathbf{N}$, with $n > 0$ and n,m relatively prime. Let F be the fixed field of τ. Then every definable subset of F^d which is definable with parameters in K is already definable with parameters in F and in the language of fields.*

Proof. This follows from the corresponding statement in ACFA (see [6]), stable embeddedness of the residue field and transformal Henselianity. □

Corollary 7.15. *The theory $\widetilde{\mathrm{VFE}}_e$ is complete relative to the theory of the residue field. Equivalently, the completions of $\widetilde{\mathrm{VFE}}_e$ are given by specifying the isomorphism type of the action of σ on the algebraic closure of the prime field.*

Proof. By Hensel's lemma, the algebraic closure of the prime field identifies via the residue map with the algebraic closure of the prime field of the residue field. So this follows immediately from Corollary 7.7 by taking F to be the relative algebraic closure of the prime field and the corresponding statement for ACFA. □

Theorem 7.16. *The theory $\widetilde{\mathrm{VFE}}$ is decidable.*

Proof. Let p be a fixed characteristic exponent and let $e \in \mathbf{N}_\infty$ be the degree of imperfection, so $e = 0$ when $p = 1$. We will prove that the theory $T_{p,e}$ of models of $\widetilde{\mathrm{VFE}}$ of characteristic exponent p and degree of imperfection e is decidable; the general case follows from compactness. By Corollary 7.15, the theory $T_{p,e}$ of models of $\widetilde{\mathrm{VFE}}$ is complete relative to the residue field. By Theorem 7.13, the induced structure on the residue field is that of a pure model of ACFA; since ACFA is decidable (see [4]), we find that $T_{p,e}$ is also decidable.

Now fix a prime p. Let T be the theory of models of $\widetilde{\mathrm{VFE}}$ of characteristic p and T_e the theory of models of $\widetilde{\mathrm{VFE}}$ of characteristic p and Ershov invariant e; we will prove that T is decidable. Given a sentence θ in the language of transformal valued fields, we must determine whether or not θ holds in every model of T. Using the result of the previous paragraph, we can effectively determine whether or not θ holds in every model of T_∞. If θ fails in some model of T_∞, then we are done; so assume that $T_\infty \vdash \theta$. By compactness, there is an effective $N = N(\theta)$ such that $T_e \vdash \theta$ for all $e > N$; this reduces our problem to verifying whether or not $T_e \vdash \theta$ for finitely many choices of e.

This shows that the theory of models of $\widetilde{\mathrm{VFE}}$ of any fixed characteristic is decidable. Using the reasoning of the previous paragraph, we conclude that $\widetilde{\mathrm{VFE}}$ is decidable in general. □

$\widetilde{\mathrm{VFE}}$ and the asymptotic theory of the Frobenius action. By a *Frobenius transformal valued field* we mean a transformal valued field of positive characteristic p whose distinguished endomorphism coincides with a power of the Frobenius endomorphism $x \mapsto x^p$. We regard Frobenius transformal valued fields as structures for the language of

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

transformal valued fields; this expands the language of valued fields by a unary function symbol for the action of $\sigma = p^n$.

Let $\mathcal{C}$ be the class of separably algebraically closed Frobenius transformal valued fields. By the *asymptotic theory* of $\mathcal{C}$ we mean the set of sentences true in all members of $\mathcal{C}$, outside a finite set of exceptional prime powers.

Theorem 7.17. *Let T be the asymptotic theory of separably algebraically closed, nontrivially valued, Frobenius transformal valued fields. Then T is logically equivalent to $\widetilde{\mathrm{VFE}}$, together with a scheme of axioms asserting that if the characteristic is p, then σ and the Frobenius are simultaneously onto.*

Proof. Let K be a model of $\widetilde{\mathrm{VFE}}$ obeying the scheme of axioms of the theorem; we must prove that K is elementarily equivalent to an ultraproduct of Frobenius transformal valued fields. If K is inversive, then the result follows from Fact 5.7, so assume that K fails to be inversive. Then either K is of positive characteristic $p > 0$ and Ershov invariant $e > 0$ or else it is of characteristic zero. We deal with the positive characteristic case; the case $p = 0$ is similar. By [16], the residue field k of K is elementarily equivalent to an ultraproduct of algebraically closed Frobenius difference fields $(k_{q_i}, x \mapsto x^{q_i})$ for powers q_i of p with $q_i \to \infty$. For each i, let K_{q_i} be a separably algebraically closed and nontrivially valued field of degree of imperfection e and with residue field k_{q_i}. It follows from Corollary 7.15 that the ultraproduct of the $(K_{q_i}, x \mapsto x^{q_i})$ is elementarily equivalent to K, and we are done. □

8. Stability theoretic phenomena

In the final section, we study stability theoretic phenomena in the theory $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ (we restrict our selves to the characteristic zero case for simplicity). The characterization of forking independence is the same in finite characteristic, but the classification of the stationary types is not. Before we turn to the details, it will be useful to give an overview of stability theoretic phenomena in the various reducts of $\widetilde{\mathrm{VFE}}$.

Stability theory in ACVF. The theory ACVF of algebraically closed and nontrivially valued fields is unstable but enjoys NIP. Thus, forking independence in ACVF is of fundamentally geometric character; namely, a global type over an algebraically closed set does not fork precisely in the event that it is invariant.

In [14], a fundamental role is played by stably dominated types; for such types, the entire forking calculus is available; namely, we have uniqueness and symmetry of forking. In ACVF, they can be equivalently characterized as the generically stable types (equivalently the stationary types, since NIP), and the types orthogonal to the value group Γ. Thus, in some loose sense, all instability in ACVF can be traced back to the value group.

The orthogonality to the value group, however, is only directly apparent over spherically complete models (which are also of fundamental importance in [14]). Over a spherically complete model, all types are dominated by the value group and the residue field. We will see below that the principle of spherically complete domination by the value group and the residue field remains valid in $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ as well.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Stability theory in $\mathrm{ACFA}_{\mathbf{Q}}$. The theory ACFA is supersimple but unstable. Forking independence is governed by linear disjointness – that is, by forking in the reduct to the field language. For an extension B/A of algebraically closed inversive fields to be stationary in ACFA, it is sufficient that $B \otimes_A C$ has no nontrivial finite σ-invariant Galois extensions, for all algebraically closed inversive extensions C of A; types of this form are called *superficially stable* in [4]. In characteristic zero, it turns out that superficially stable types are precisely those orthogonal to the fixed field (see Fact 8.10 below). In finite characteristic, there exist non-superficially stable types which are modular; they are related to additive equations involving σ and the Frobenius (see [4, Example 6.5]).

Remark 8.1. [3]We work in ACFA. Let B/A be an extension of inversive algebraically closed difference fields. Say that B/A is *generically stable* if for all algebraically closed and inversive C linearly disjoint from B over A, the difference field $B \otimes_A C$ admits a unique lift of σ to the algebraic closure, up to isomorphism. If B/A is superficially stable, then it is generically stable; conversely in $\mathrm{ACFA}_{\mathbf{Q}}$, it can be shown that there is no distinction between superficial stability and generic stability. It is not known if this holds for modular types in finite characteristic; the question is related to the the classification of unstable types in $\mathrm{ACFA}_{\mathbf{F}_p}$ orthogonal to all fixed fields.

Situation. In what follows, we will study the following situation. Let A be a model of $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ which is spherically complete. Let B and C be models of VFA over A, jointly embedded over A in an ambient extension. Let us assume that the intersection $\Gamma_B \cap \Gamma_C = \Gamma_A$ and that the residue fields k_B and k_C are linearly disjoint over k_A. Let b be a tuple enumerating the elements of B. Moreover, let α and γ be tuples enumerating the residue field elements and the value group elements of B, respectively.

Lemma 8.2. *Notation as above. Then* $\mathrm{qftp}_{\widetilde{\mathrm{VFA}}}(b/A) \cup \mathrm{qftp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C) \cup \mathrm{qftp}_{\mathrm{ACFA}}(\alpha/k_C) \vdash \mathrm{qftp}_{\widetilde{\mathrm{VFA}}}(b/C)$. *Moreover,*

(1) *The fields B and C are linearly disjoint over A.*

(2) *Let $N = B \otimes_A C$; then N has residue field $k_B \otimes_{k_A} k_C$, and its value group is the cofibered sum $\Gamma_B \oplus_{\Gamma_A} \Gamma_C$.*

Proof. This follows formally from spherically complete domination in ACVF; see [14, Proposition 12.11]. Namely, the quantifier free type in the language of abstract valued fields is determined; hence, by functoriality, the lift of σ is unique also (or just argue directly using the linear disjointness of B and C over A). The 'moreover' clauses are part of the proof of [14, Proposition 12.11]. □

Lemma 8.3. *Let K be a model of* $\mathrm{VFA}_{\mathbf{Q}}$. *Assume that the value group of K is algebraically divisible. Then completions of the theory* $\widetilde{\mathrm{VFA}}_K$ *of models of* $\widetilde{\mathrm{VFA}}$ *over K correspond bijectively to completions of* ACFA_k, *the theory of models of* ACFA *over k. In particular, if k has no nontrivial finite σ-invariant Galois extensions, then the theory of models of* $\widetilde{\mathrm{VFA}}$ *over K is complete.*

[3]We are grateful to Zoé Chatzidakis and Ehud Hrushovski for useful discussions around this issue.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. Using [11, Proposition 4.7], we may assume that K is Henselian; this does not change the value group or the residue field. By Henselianity, the valuation lifts uniquely to finite extensions of K. Using divisibility of Γ and elementary ramification theory, the category of algebraic extensions of K is in canonical bijection with the category of algebraic extensions of k via the residue map; by functoriality, the category of transformal valued fields algebraic over K is equivalent to the category of difference fields algebraic over k. So we finish using Fact 5.7(8). □

Lemma 8.4. *Notation as in Lemma 8.2 (and recall that the characteristic is zero). Then we have*

$$\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/A) \cup \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C) \cup \mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C) \vdash \mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C).$$

Proof. Increasing B and C, we can assume both are algebraically closed and the conclusion descends (the linear disjointness assumption on the residue fields is not violated by Fact 4.23, for example). Let $N = B \otimes_A C$ be the displayed model of VFA. By Lemma 8.2, the transformal valued field structure on N is determined; hence, using Fact 5.7 (8), the type $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C)$ is determined by the additional data of the choice of an isomorphism type of an extension of the transformal valued field structure of N to an algebraic closure. So we finish using Lemma 8.3, together with Fact 5.7(8). □

For the sake of the next corollary, we keep the notation A, B, C, α, γ and b. Elementary properties of forking independence and the characterization of forking in ACFA and $\widetilde{\omega\mathrm{OGA}}$ will be used freely.

Corollary 8.5. *Let A be a spherically complete model of $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ and $B, C \supseteq A$ two models of* VFA *over A. Then the type $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C)$ does not fork over A if and only if the following conditions are satisfied:*

(1) *The type $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C)$ does not fork over Γ_A.*
(2) *The type $\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C)$ does not fork over k_A.*

Moreover, in these two (equivalent) situations, we have

$$\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/A) \cup \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C) \cup \mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C) \vdash \mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C).$$

Proof. Assume first that $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C)$ does not fork over A. Then $\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C)$ does not fork over k_A by elementary properties of forking independence, and likewise for Γ. This shows the easy direction.

Conversely, assume that (1) and (2) hold. We show that $\mathrm{tp}(b/C)$ does not fork over A. Towards this end, there is no harm in increasing C. We may therefore assume that C is sufficiently saturated. In this situation, forking equals dividing for types over C, so it will be enough to show that this type does not divide over A.

Using the characterization of forking in ACFA and full embeddedness of the residue field in $\widetilde{\mathrm{VFA}}$, we find that k_B and k_C are linearly disjoint over k_A. Similarly, using full embeddedness of Γ, we find that $\Gamma_B \cap \Gamma_C = \Gamma_A$. Using Lemma 8.2, we have that

$$\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/A) \cup \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C) \cup \mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C) \vdash \mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C).$$

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

The goal is to prove that the right-hand side does not divide over A. Since the left-hand side implies the right, it is sufficient to show that any conjunction of formulas from the left-hand side does not divide over A.

By stable embeddedness and orthogonality of the residue field and value group in $\widetilde{\mathrm{VFA}}$, $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(\gamma,\alpha/C)$ is equivalent $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C)\cup\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C)$ and so does not divide over A.

Consider a conjunction $\phi(x)\wedge\psi(y,z,c)$, where $\phi(x)$ is an A-formula from $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/A)$ (so x is a tuple of field sort variables) and $\psi(y,z,c)$ is a formula from $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(\gamma,\alpha/C)$. The tuples x,y and z correspond to finite subtuples b_0,γ_0 and α_0, respectively. Since α is an enumeration of the residue elements of b and γ is an enumeration of the value group of b, increasing the tuples b_0,γ_0 and α_0, we may assume that

$$b_0 \models \psi'(x,y,z,c) := \phi(x)\wedge v(x)=y\wedge \mathrm{res}(x)=z\wedge\psi(y,z,c).$$

The formula $\exists x\psi'(x,y,z,c)$ lies in $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(\gamma,\alpha/C)$ and so does not divide over A. Since ϕ is over A, the definition of dividing implies that $\phi(x)\wedge\psi(y,z,c)$ does not divide over A as well. □

The following was obtained by Chernikov and Hils in [8] in a relative setting.

Corollary 8.6. *We work in the theory* $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$. *Let A be a small model and p a global type which forks over A; then there is an NIP formula φ dividing over A such that $p\vdash\varphi$.*

Proof. Fix a small but sufficiently saturated model C over A, so a type over C forks over A if and only if it divides over A. Let us assume that $p=\mathrm{tp}(b/C)$ divides over A; we will prove that $p\vdash\varphi$ where φ is an NIP formula which divides over A. Without loss of generality, replacing b by $\mathrm{dcl}(Ab)$, we can take b to enumerate a small model B of VFA over A.

First, assume that A is spherically complete. In this situation, by Corollary 8.5, either $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C)$ divides over Γ_A or $\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C)$ divides over k_A. In any of these situations, the formula witnessing this is an NIP formula inside $\mathrm{tp}(b/C)$ dividing over A.

The reduction to the spherically complete case is formal. Let $A\subseteq A_0\subseteq C$ be a spherically complete model of $\widetilde{\mathrm{VFA}}$. Replacing A_0 by an isomorphic copy over A, we can assume that p continues to divide over A_0. By the spherically complete case, there is an NIP formula $p\vdash\varphi$ dividing over A_0; this formula divides over A as well, and we are done. □

Corollary 8.7. *The theory* $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ *has NTP$_2$.*

Proof. By Corollary 8.6, $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ has NTP_2-forking in the sense of Proposition A.3. Thus, by Theorem A.6, $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ has NTP_2. □

Using the characterization of forking, we can pin down the stationary types in $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$.

Definition 8.8. Let T be any theory and let p be a global type. We say that p is *stationary* if there is a small set C_0 with the following property: for every small set C containing C_0, the type p is the unique global nonforking extension of $p|_C$.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Remark 8.9. Let p be a global type which is stationary and let C_0 be a small set witnessing this. Then p does not fork over C_0. In fact, since the forking ideal is invariant under automorphisms of the universal domain, the type p does not split over C_0 (i.e., it is invariant under global automorphisms fixing C_0 pointwise).

The following is a Corollary of [4] and [5] worth making explicit. It also follows from [3, Lemma 3.6], we sketch here another proof.

Fact 8.10. In the theory $\mathrm{ACFA}_{\mathbf{Q}}$, a global invariant type is stationary if and only if it is orthogonal to the fixed field.

Proof. There are no non-algebraic stationary types in the theory of pseudofinite fields, so the condition is necessary. For the converse, assume C is an algebraically closed extension of an algebraically closed field A orthogonal to the fixed field; we will show that C is superficially stable over A in the sense of [4]; such types are stationary, using the characterization of forking in ACFA. First, assume that C is purely transformally transcendental over A. In this situation, it is stationary over A using [5, Theorem 5.3]. Next, assume that C is transformally algebraic over A. In this situation, it is superficially stable over A using [4, Theorem 4.11, Theorem 4.2.5]. We have shown the claim in the case C transformally algebraic over A and purely transformally transcendental over A. But if $A \subseteq B \subseteq C$ is a tower with B/A and C/B stationary, then it follows formally that C/A is likewise stationary. Applying this with B the relative transformal algebraic closure of A in C, we obtain the general case. □

Corollary 8.11. *In the theory $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$, a global invariant type is stationary if and only if it is orthogonal to the fixed field and the value group.*

Proof. The conditions are necessary: the value group is linearly ordered, so the only generically stable types are algebraic. Orthogonality to the fixed field is likewise necessary using (the easy direction of) Fact 8.10.

The converse follows immediately from Fact 8.10 and Corollary 8.5. □

Remark 8.12. In finite characteristic, a classification of this form cannot hold; see [4, Example 6.5].

Let K be a valued field. We say that K is *separably maximal* if the perfect hull of K is spherically complete; this is equivalent to the requirement that K admits no immediate separable extensions. In [15], valued fields of this form were used to establish domination over the residue field and the value group in the theory SCVF_e of separably algebraically closed valued fields of *finite* degree of imperfection. It is not difficult to see that a domination result of this form cannot hold in infinite Ershov invariant. Similarly, in $\widetilde{\mathrm{VFE}}$, we cannot hope for a domination theorem relative to the value group and the residue field. In order to transfer the domination statement, we will make use of the following definition:

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Definition 8.13. Let K be a model of VFE.

(1) We say that K is *transformally separably algebraically maximal* if whenever $\widetilde{K}$ is a model of VFE transformally separably algebraic over K with $k = \widetilde{k}$ and $\widetilde{\Gamma} = \Gamma$, then $\widetilde{K} = K$; in other words, the field K has no nontrivial immediate transformally separably algebraic extensions.

(2) We say that K is *transformally separably maximal* if whenever $\widetilde{K}$ is a model of VFE transformally separable over K with $k = \widetilde{k}$ and $\widetilde{\Gamma} = \Gamma$, then $\widetilde{K} = K$; in other words, the field K has no nontrivial immediate transformally separable extensions.

Remark 8.14. Let K be a model of VFE. It is clear that if K is transformally separably maximal, then K is transformally separably algebraically maximal.

Remark 8.15. Let k be an inversive algebraically closed difference field which is trivially valued and x transformally transcendental over k with $vx > 0$; by [16, Proposition 6.23], the completed algebraic closure of $k\left(x^{\mathbf{N}[\sigma]}\right)$ is transformally algebraically maximal.

Proposition 8.16. *Let K be a model of $\widetilde{\text{VFE}}$. Then K is transformally separably algebraically maximal.*

Proof. Let L be a model of VFE transformally separably algebraic and immediate over K. Let $\widetilde{K}$ and $\widetilde{L}$ denote the inversive hulls of K and L, respectively. Since the value group and the residue field of K are inversive, the same is true of L; hence, $\widetilde{L}$ is an immediate extension of $\widetilde{K}$. But $\widetilde{K}$ is a model of $\widetilde{\text{VFA}}$, and hence transformally algebraically maximal by Fact 5.7. It follows that $\widetilde{K} = \widetilde{L}$; that is, the inversive hulls of K and L coincide. So L is simultaneously purely transformally inseparably algebraic and transformally separably algebraic over K; it follows that $K = L$, which is what we wanted. □

Remark 8.17. Let K be a model of VFA which is algebraically closed; then K is maximal in the category of models of VFA if and only if it is maximal in the category of valued fields; see [11, Proposition 4.18].

Lemma 8.18. *Let K be a model of $\widetilde{\text{VFE}}$. Then K is transformally separably maximal if and only if the inversive hull of K is spherically complete.*

Proof. If the inversive hull of K is spherically complete, then the proof of Proposition 8.16 shows that K is transformally separably maximal. Conversely, if the inversive hull M of K is not spherically complete, let L be an extension witnessing this. Since M is transformally algebraically maximal, we can find some x in L transformally transcendental over M; then $K\left(x^{\mathbf{N}[\frac{\sigma}{p^\infty}]}\right)$ is an immediate proper transformally separable extension of K, whence K is not transformally separably maximal. □

Proposition 8.19. *Let K be a model of $\widetilde{\text{VFE}}_e$. Then there is a model M of $\widetilde{\text{VFE}}_e$, immediate and transformally separable over K, which is transformally separably maximal.*

Proof. Let E be a spherically complete immediate extension of K (Fact 5.7(7)); then E is a model of $\widetilde{\text{VFA}}$ and in particular inversive. We will find the desired M inside E; in fact, the field E will be the inversive hull of M.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Let (x_i) be a transformal transcendence basis of E over K. Let M be the relative transformal separable algebraic closure of $K\left(x_i^{\mathbf{N}[\sigma,p^{-1}]}\right)$ inside E. We claim that M is as desired. For this, we must check that M is transformally separably maximal, that M is transformally separable over K, and that M is a model of $\widetilde{\mathrm{VFE}}_e$.

By construction, the perfect hull of M splits over K, so it is a model of $\mathrm{VFE}_{\leq e}$ immediate over K; it remains to be shown that M is transformally separably maximal, and that it is a model of $\widetilde{\mathrm{VFE}}_e$. We claim that E is the inversive hull of M. Indeed, since the x_i form a transformal transcendence basis of E over K, the field E is transformally algebraic over M. Furthermore, by our choice of M, the field M is transformally separably algebraically closed in E. So E is inversive and purely transformally inseparably algebraic over M; hence, it is the inversive hull of M. So the inversive hull of M is a spherically complete model of $\widetilde{\mathrm{VFA}}$; it follows from Lemma 8.18 and Proposition 5.27 that M is as desired. □

Let $\mathrm{ACF}^2 = \mathrm{ACF}^2_{\mathbf{Q}}$ be the theory of pairs (B,A) of algebraically closed fields with B a proper extension of A. This theory is model complete and stable if we expand the language by relation symbols expressing linear independence of tuples of elements of B viewed as a vector space over A [10]. The theory ACF^2 is a reduct of $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ by setting $B = K$ and $A = K^\sigma$. The transformal λ-functions are then definable in the reduct to ACF^2.

Lemma 8.20. *We work in* $\mathrm{SCFE}_{\mathbf{Q}}$. *Let A be an algebraically closed difference field and B,C be algebraically closed extensions of A inside a saturated model* $\mathbf{K}$. *Let us assume that* $\mathbf{K}$ *is transformally separable over B and C. Then B and C are independent over A (in the sense of simplicity theory) if and only if they are linearly disjoint over A and* $\mathbf{K}$ *is transformally separable over $B \otimes_A C$. This independence relation is moreover given by stable formulas – namely, by independence in the sense of* ACF^2.

Proof. See [5, Definition 4.9 and Proposition 4.10]; the fields B and C are independent over A in the sense of SCFE if and only if they are linearly disjoint over A and a transcendence basis of B over $A \otimes_{A^\sigma} B^\sigma$ lifts to a transcendence basis of $B \otimes_A C$ over $B^\sigma \otimes_{A^\sigma} A \otimes_{A^\sigma} C^\sigma$. □

Lemma 8.21. *Let A be a model of* $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ *and* $\mathbf{K}$ *a saturated model of* $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ *transformally separable over A. Let B,C be models of* VFE *over A inside* $\mathbf{K}$, *closed under the transformal λ-functions of* $\mathbf{K}$, *and assume that B and C are independent over A in the sense of* $\mathrm{SCFE} = \mathrm{SCFE}_{\mathbf{Q}}$. *Let b be a tuple enumerating the elements of B. Then* $\mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/A) \cup \mathrm{tp}_{\mathrm{ACF}^2}(b/C) \cup \mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C) \vdash \mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/C)$.

Proof. This is a formal consequence of Lemma 8.20 and quantifier elimination in $\widetilde{\mathrm{VFE}}$ (Corollary 7.7). Namely, by the independence assumption (encoded in the type $\mathrm{tp}_{\mathrm{ACF}^2}(b/C)$), the fields B and C are linearly disjoint over A, and the field $N = B \otimes_A C$ is closed under the transformal λ-functions of $\mathbf{K}$; given this, the type $\mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/C)$ is determined by its quantifier free part and the choice of an extension of the transformal valued field structure to the algebraic closure. This data is encoded in $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C)$. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Let A be a model of $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ which is transformally separably maximal. Let B,C be models of VFE over A inside an ambient saturated transformally separable extension $\mathbf{K} \models \widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ of A, and assume B and C are closed under the transformal λ-functions of $\mathbf{K}$. Let b,α,γ be tuples enumerating the elements of B, the residue field elements of B, and the value group elements of B, respectively.

Proposition 8.22. *Notation as above (and recall the characteristic is zero). Then the type* $\mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/C)$ *does not fork over* A *if and only if the following conditions are satisfied:*

(1) *The type* $\mathrm{tp}_{\mathrm{ACF}^2}(b/C)$ *does not fork over* A*; that is,* B *and* C *are independent over* A *in the sense of* SCFE.

(2) *The type* $\mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C)$ *does not fork over* k_A*; that is, the fields* k_B *and* k_C *are linearly disjoint over* k_A.

(3) *The type* $\mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C)$ *does not fork over* Γ_A*; that is, it admits a global extension which does not split over* Γ_A.

Moreover, in these two (equivalent) cases, we have

$$\mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/A) \cup \mathrm{tp}_{\mathrm{ACF}^2}(b/C) \cup \mathrm{tp}_{\widetilde{\omega\mathrm{OGA}}}(\gamma/\Gamma_C) \cup \mathrm{tp}_{\mathrm{ACFA}}(\alpha/k_C) \vdash \mathrm{tp}_{\widetilde{\mathrm{VFE}}}(b/C).$$

Proof. The three conditions are necessary in order to guarantee nonforking. Namely, (1) is expressed using stable formulas, and independence is the generic condition. Moreover, (2) and (3) follow from stable embeddedness of the value group and the residue field. To show the conditions are sufficient it will be enough to establish the 'moreover' part (as in the proof of Corollary 8.5). Using Lemma 8.18, the inversive hull of A is a spherically complete model of $\widetilde{\mathrm{VFA}}$. So working over the inversive hull of A and using Corollary 8.5, the left-hand side axiomatizes $\mathrm{tp}_{\widetilde{\mathrm{VFA}}}(b/C)$; using Lemma 8.21, this data is sufficient to axiomatize the right-hand side. □

Corollary 8.23. *In the theory* $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$*, forking over a model is always witnessed by an* NIP *formula; hence,* $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$ *is* NTP_2.

Proof. The proof for the case of $\widetilde{\mathrm{VFA}}_{\mathbf{Q}}$ goes through, using Proposition 8.22 to replace maximal models with transformally separably maximal models and noting that the theory ACF^2 is stable (so forking is always witnessed by a stable formula). □

Remark 8.24. Simone Ramello proved a vast generalization of Corollary 8.23 using the proof strategy from [8]; see [21, Section 5.5].

Corollary 8.25. *In the theory* $\widetilde{\mathrm{VFE}}_{\mathbf{Q}}$*, a global invariant type is stationary if and only if it is orthogonal to the fixed field and the value group.*

Proof. This follows immediately from Proposition 8.22, Fact 8.10 and the uniqueness of nonforking extensions in ACF^2. □

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Appendix: NTP$_2$-forking

by Itay Kaplan

Let T be a first-order theory and let $\mathcal{U} \models T$ be a large saturated model. Chernikov proved in [7, Proposition 4.14(1)] that assuming that whenever $p(x) \in S(N)$ divides over $M \prec N$, then this is witnessed by an NIP (dependent) formula,[4] then T is NTP$_2$. We generalize this slightly by relaxing the assumption on T from all types over models to global types (or types over saturated enough models; see Proposition A.3 below).

Definition A.1. A formula $\phi(x,y)$ (perhaps with parameters) is *TP$_2$* if there is an array $(a_{i,j})_{i,j<\omega}$ and $k<\omega$ such that

- every row is k-inconsistent (i.e., for every $i<\omega$, $\{\phi(x,a_{i,j}) \mid j<\omega\}$ is k-inconsistent) and
- every vertical path is consistent (i.e., for every $\eta:\omega\to\omega$, $\{\phi(x,a_{i,\eta(i)}) \mid i<\omega\}$ is consistent).

If ϕ does not have TP$_2$, we say that ϕ is *NTP$_2$*.

Definition A.2. Say that a (complete first-order) theory T has *NTP$_2$-forking over a (small) model M* if whenever $\phi(x,a)$ forks over M, then there are NTP$_2$-formulas $\phi_i(x,z_i)$ over M and b_i for $i<n$ such that $\phi(x,a)\vdash \bigvee_{i<n}\phi_i(x,b_i)$ and $\phi_i(x,b_i)$ divides over M. In other words, forking over M is witnessed by dividing NTP$_2$-formulas.

Say that T has *NTP$_2$-forking* if it has NTP$_2$ forking over all models M.

Proposition A.3. *The following are equivalent for a theory T and $M\models T$.*

(1) *T has NTP$_2$-forking over M.*

(2) *If $p(x)$ is a global type which divides/forks over M, then there is some NTP$_2$-formula $\phi(x,y)$ over M and b such that $\phi(x,b)\in p$ and $\phi(x,b)$ divides over M.*

(3) *If $N\succ M$ is $|M|^+$-saturated and $p(x)\in S(N)$ divides/forks over M, then there is some NTP$_2$-formula $\phi(x,y)$ over M and b such that $\phi(x,b)\in p$ and $\phi(x,b)$ divides over M.*

Proof. (1) implies (3) follows from saturation, and (3) implies (2) is clear. In order to prove that (2) $\implies$ (1), we assume that (1) fails and show that (2) must fail also. So suppose that $\phi(x,a)$ forks over M but does not imply a disjunction of dividing NTP$_2$-formulas over M as in the definition; then $\{\neg\psi(x,b) \mid \psi(x,y)\in L(M) \text{ is NTP}_2, b\in\mathcal{U} \text{ and } \psi(x,b) \text{ divides over } M\}$ is consistent with $\phi(x,a)$, and thus, any completion gives a contradiction to (2). □

We need the notions of strict invariance and strict Morley sequences.

Definition A.4. We say that a global type $q(x)\in S(\mathcal{U})$ is *strictly invariant over a model M*, if q is M-invariant and if $c\models q|_C$ where $C\supseteq M$, then $C \mathop{\downarrow}_M c$ (i.e., $\text{tp}(C/Mc)$ does not fork over M).

[4] In fact, the proof only uses that the witness is NTP$_2$.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

We write $a \mathop{\smile\!\!\!\!\!\!|}^{ist}_{M} C$ for the following: there exists a global strictly invariant over M extension of $\mathrm{tp}(a/MC)$.

Theorem A.5. *Suppose that T has NTP_2-forking over M. Then forking equals dividing over M: if $\phi(x,b)$ forks over M, then $\phi(x,b)$ divides over M.*

It follows that if T has NTP_2-forking, then forking equals dividing over models.

Proof. The proof is essentially the same as the proof of [9, Theorem 1.1]. See also [23, Theorem 5.49] and [1]. The proof follows the following stages:

(1) (Generic witness) Suppose that $\phi(x,y) \in L(M)$ is NTP_2 and that $\phi(x,a)$ divides over a model M. Then there is a global coheir $q(y)$ extending $\mathrm{tp}(a/M)$ such that some/any Morley sequence in q over M witnesses the dividing of $\phi(x,a)$. In other words, if $(a_i)_{i<\omega} \vDash q^{(\omega)}|_M$, then $\{\phi(x,a_i) \mid i<\omega\}$ is inconsistent. This follows from the same proof as in [9, Lemma 3.12] (see also [23, Lemma 5.35]): note that although there NTP_2 is assumed globally, it is only used for the specific formula in question.

(2) (Forking implies quasi-dividing) If $\phi(x,a)$ forks over M, then there are finitely many conjugates of a over M, $a = a_0, \ldots, a_{m-1}$ such that $\{\phi(x,a_i) \mid i<n\}$ is inconsistent. Indeed, by assumption, there are NTP_2-formulas $\phi_i(x,z_i) \in L(M)$ and $b_i \in \mathcal{U}$ for $i<n$ such that $\phi_i(x,b_i)$ divides over M and $\phi(x,a) \vdash \bigvee_{i<n} \phi_i(x,b_i)$. By (1), there are global coheirs $q_i(z_i) \in S(\mathcal{U})$ such that any Morley sequence in q_i over M witnesses that $\phi(x,b_i)$ divides. This allows us to apply the 'broom lemma', [9, Lemma 3.1] (see also [23, Lemma 5.36]) which in [9] does not require NTP_2 and allows us to conclude exactly as in the proof of [9, Corollary 3.13].

(3) (Existence of strictly invariant extensions) if $p(x) \in S(M)$, then there is a strictly invariant over M type (in fact, even a coheir, but we will not use this) $q(x) \in S(\mathcal{U})$ extending p. This follows from [9, Proposition 3.7(1)] which does not require NTP_2 but does assume (2).

(4) (Kim's lemma for NTP_2-formulas) If $\phi(x,y) \in L$ is NTP_2 and $\phi(x,a)$ divides over a model M, and $q(y) \supseteq \mathrm{tp}(a/M)$ is a global strictly invariant type over M, then any Morley sequence in q over M witnesses dividing. In other words, if $(a_i)_{i<\omega}$ is a Morley sequence of q over M, then $\{\phi(x,a_i) \mid i<\omega\}$ is inconsistent. This follows from (the proof of) [9, Lemma 3.14].

(5) (Conclusion) If $\phi(x,a)$ forks over M, then there are NTP_2-formulas $\phi_i(x,z_i) \in L(M)$ and $b_i \in \mathcal{U}$ for $i<n$ such that $\phi_i(x,b_i)$ divides over M and $\phi(x,a) \vdash \bigvee_{i<n} \phi_i(x,b_i)$. By concatenating and adding dummy variables (these operations clearly preserves being NTP_2), we may assume that $b_i = a$ and $z_i = z$ for all $i<n$. By (3), there exists a global M-strictly invariant type $q(z) \supseteq \mathrm{tp}(a/M)$. Let $(a_j)_{j<\omega}$ be a Morley sequence in q. By (4), $\{\phi_i(x,a_j) \mid j<\omega\}$ is inconsistent for all $i<n$. It follows by pigeonhole that $(a_j)_{j<\omega}$ witnesses that $\phi(x,a)$ divides over M (see also the proof of [9, Corollary 3.16]). □

Now we can conclude:

Theorem A.6. *Let T be any theory. Then T has NTP_2-forking if and only if T is NTP_2.*

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

Proof. Right to left is obvious, so we prove left to right.

Suppose not: T has NTP_2 forking and TP_2. By [7, Theorem 4.9], there is a model M and sequence $(b_i)_{i<\kappa}$ for $\kappa = (|M|+|T|)^+$, such that $b_i \mathop{\smile\!\!\!\!\!\!|}^{ist}_{M} b_{<i}$ and some a such that $\mathrm{tp}(a/Mb_i)$ divides for all $i<\kappa$. For each $i<\kappa$, let $\phi_i(x,b_i) \in \mathrm{tp}(a/Mb_i)$ divide over M. By assumption, there are NTP_2-formulas $(\psi_{i,j}(x,y_{i,j}))_{j<n_i}$ over M and $(c_{i,j})_{j<n_i}$ such that $\phi_i(x,b_i) \vdash \bigvee_{j<n_i} \psi_{i,j}(x,c_j)$. Let y_i be the concatenation of all the variables $y_{i,j}$, and let c_i be the appropriate concatenation of the $c_{i,j}$'s (we cannot incorporate b_i as well since this may cause the loss of the strict independence over M). Note that $\psi_i(x,y_i) := \bigvee_{j<n_i} \psi_{i,j}(x,y_{i,j})$ is NTP_2 by [7, Lemma 7.1].

By Theorem A.5, $\psi_i(x,c_i)$ divides over M for all $i<\kappa$. Thus, for each $i<\kappa$, there is an M-indiscernible sequence $(d_{i,j})_{j<\omega}$ starting with $d_{i,0}=c_i$ and some number $k_i<\omega$ such that $\{\psi_i(x,d_{i,j}) \mid j<\omega\}$ is k_i-inconsistent. Trimming, we may assume that $k_i =: k$, $\phi_i =: \phi$ and $\psi_i(x,y_i) =: \psi(x,y)$ for all $i<\omega$. For each $d_{i,j}$, find $b_{i,j}$ such that $b_{i,j}d_{i,j} \equiv_M b_ic_i$ and $b_{i,0}=b_i$. In particular, $\phi_i(x,b_{i,j}) \vdash \psi_i(x,d_{i,j})$. We may assume, by Ramsey and applying an automorphism, that $(b_{i,j})_{j<\omega}$ is M-indiscernible. By [7, Lemma 4.4], there are M-mutually indiscernible sequences $((e_{i,j})_{j<\omega})_{i<\omega}$ such that $e_{i,0}=b_{i,0}=b_i$ and $(e_{i,j})_{j<\omega} \equiv_M (b_{i,j})_{j<\omega}$ for all $i<\omega$.

For each $i<\omega$, find $(f_{i,j})_{j<\omega}$ such that $(e_{i,j}f_{i,j})_{j<\omega} \equiv_M (b_{i,j}d_{i,j})_{j<\omega}$. We claim that the array $(f_{i,j})_{i,j<\omega}$ and k witnesses that $\psi(x,y)$ has TP_2.

(1) Every row is k-inconsistent (i.e., $\{\psi(x,f_{i,j}) \mid j<\omega\}$ is k-inconsistent) because for each $i<\omega$, $(f_{i,j})_{j<\omega} \equiv_M (d_{i,j})_{j<\omega}$ and choice of the latter sequence.

(2) Every vertical path is consistent: as $e_{i,0}=b_i$ and $\phi(a,b_i)$ holds for all $i<\omega$, it follows from mutual indiscernibility of the array $(e_{i,j})_{i,j<\omega}$ that for every $\eta:\omega\to\omega$, $\{\phi(x,e_{i,\eta(i)}) \mid i<\omega\}$ is consistent, and let a_η be a realization. As $\phi(x,b_{i,j}) \vdash \psi(x,d_{i,j})$, $\phi(x,e_{i,j}) \vdash \psi(x,f_{i,j})$ and hence $\psi(a_\eta,f_{i,\eta(i)})$ holds, and in particular, the set $\{\psi(x,f_{i,\eta(i)}) \mid i<\omega\}$ is consistent, as required.

Together, we get a contradiction. □

Remark A.7.

(1) From the proof of (2) in the proof of Theorem A.5, one can deduce the following statement, which may be interesting on its own. Let T be any theory, and let $\phi(x,y)$ be an NIP formula. Suppose that M is a model and that $\pi(x)$ is a partial type over $\mathcal{U}$ which is M-invariant. Then there is some $\pi^* \supseteq \pi$ which is M-invariant and such that $\pi^* \restriction \phi$ is a complete global ϕ-type. Indeed, it is enough to show that $\pi(x) \cup \{\phi(x,c) \leftrightarrow \phi(x,b) \mid b \equiv_M c\}$ is consistent. If not, then $\pi \vDash \bigvee_{i<n} \neg(\phi(x,c_i) \leftrightarrow \phi(x,b_i))$. For each $i<n$, let q_i be a global coheir extending $\mathrm{tp}(b_i/M) = \mathrm{tp}(c_i/M)$, and let $a_i \vDash q_i|_{Mb_ic_i}$. Then we get that $\pi \vDash \bigvee_{i<n} \neg(\phi(x,a_i) \leftrightarrow \phi(x,b_i)) \vee \neg(\phi(x,a_i) \leftrightarrow \phi(x,c_i))$. Note that as ϕ is NIP, $\neg(\phi(x,a_i) \leftrightarrow \phi(x,b_i))$ divides over M as witnessed by any Morley sequence in $q_i^{(2)}$ over M (and the same is true replacing b_i by c_i), and that $\psi(x,yz) := \neg(\phi(x,y) \leftrightarrow \phi(x,z))$ is NIP thus NTP_2. So we have that $\pi(x)$ implies a disjunction of dividing over M NTP_2 formulas. By the 'broom lemma', [9, Lemma 3.1], we get that π is inconsistent, a contradiction.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

(2) Call a formula $\phi(x,y)$ *simple* if it does not have the tree property (see, for example, [20, Definition 2.3.5]). Say that T has *simple forking* if whenever a formula forks over a model M, this is witnessed by simple formulas over M, as in Definition A.2. Every simple formula is NTP_2, so if T has simple forking, it has NTP_2 forking and thus forking equals dividing over models. A disjunction of simple formulas is still simple (see [20, Lemmas 2.3.2 and 2.3.4(3)]), so similar ideas as in the proof of Theorem A.5 imply that in this case, T is simple (if not, then by lack of local character, there is some a and an increasing continuous sequence $(M_i)_{i<|T|^+}$ of models of size $\leq |T|$ such that $\mathrm{tp}(a/M_{i+1})$ forks over M_i, as witnessed by $\psi_i(x,d_i)$ where ψ is over M_i. By the above, there is a simple formula $\phi_i(x,y_i)$ over M_i and some c_i such that $\psi_i(x,d_i) \vdash \phi_i(x,c_i)$ and $\phi_i(x,c_i)$ divides over M_i. By Fodor and pigeonhole, we may assume that for all $i<\omega$, $\phi_i = \phi$ and $\psi_i = \psi$ which are over M_0 and that $\phi(x,c_i)$ k-divides over M_i for some constant $k<\omega$. Now, exactly as in the proof of Theorem A.6, we conclude by constructing a tree of instances of ϕ, contradiction). This is analogous to Chernikov's [7, Proposition 4.14(2)] in the same way that Theorem A.6 is analogous to [7, Proposition 4.14(1)].

Acknowledgements. We would like to thank Zoé Chatzidakis, Ehud Hrushovski and Simone Ramello for useful discussions. Lastly, we would like to thank the anonymous referee for his\her thorough reading and numerous helpful comments and suggestions, and for catching several inaccuracies in some of the proofs.

Competing interest. The authors have no competing interests to declare.

Data availability statement. Data sharing is not applicable to this article as no new data were created or analyzed in this study.

References

[1] Adler H (2014) Kim's lemma for NTP_2 theories: A simpler proof of a result by Chernikov and Kaplan. *Rend. Semin. Mat. Univ. Politec. Torino* **72**(34), 121–126.
[2] Azgin S (2010) Valued fields with contractive automorphism and Kaplansky fields. *J. Algebra* **324**(10), 2757–2785.
[3] Bustamante-Medina RF (2019) Definable groups in DCFA. *Rev. Mat. Teor. Apl.* **26**(2), 179–195.
[4] Chatzidakis Z and Hrushovski E (1999) Model theory of difference fields. *Trans. Amer. Math. Soc.* **351**(8), 2997–3071.
[5] Chatzidakis Z and Hrushovski E (2004) Model theory of endomorphisms of separably closed fields. *J. Algebra* **281**(2), 567–603.
[6] Chatzidakis Z, Hrushovski E and Peterzil Y (2002) Model theory of difference fields, ii: Periodic ideals and the trichotomy in all characteristics. *Proceedings of the London Mathematical Society* **85**(2), 257–311.
[7] Chernikov A (2014) Theories without the tree property of the second kind. *Ann. Pure Appl. Logic* **165**(2), 695–723.
[8] Chernikov A and Hils M (2014) Valued difference fields and ntp2. *Israel J. Math.* **204**(1), 299–327.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press

[9] Chernikov A and Kaplan I (2012) Forking and dividing in NTP_2 theories. *J. Symbolic Logic* **77**(1), 1–20.
[10] Delon F (2012) Élimination des quantificateurs dans les paires de corps algébriquement clos. *Confluentes Math.* **4**(2), 1250003, 11.
[11] Dor Y and Hrushovski E (2022) Specialization of difference equations and high frobenius powers. arXiv preprint 2212.05366.
[12] Fried MD and Jarden M (2008) *Field Arithmetic*, vol. 11, 3rd edn. Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]. Berlin: Springer-Verlag. Revised by Jarden.
[13] Haskell D, Hrushovski E and Macpherson D (2008) *Stable Domination and Independence in Algebraically Closed Valued Fields*, no. 30. Cambridge: Cambridge University Press.
[14] Haskell D, Hrushovski E and Macpherson D (2008) *Stable Domination and Independence in Algebraically Closed Valued Fields*, vol 30. Lecture Notes in Logic. Association for Symbolic Logic, Chicago, IL. Cambridge: Cambridge University Press.
[15] Hils M, Mand Rideau S (2018) Imaginaries in separably closed valued fields. *Proc. Lond. Math. Soc.* **116**(6),1457–1488.
[16] Hrushovski E (2022) The elementary theory of the frobenius automorphisms. arXiv preprint 0406514.
[17] Hrushovski E and Tatarsky A (2006) Stable embeddedness in algebraically closed valued fields. *J. Symb. Log.* **71**(3), 831–862.
[18] Johnson W (2022) Forking and dividing in fields with several orderings and valuations. *J. Math. Log.* **22**(1), Paper No. 2150025, 43.
[19] Karpilovsky G (1989) *Topics in Field Theory*, vol. 155. North-Holland Mathematics Studies. Amsterdam: North-Holland Publishing Co. Notas de Matemática [Mathematical Notes], 124.
[20] Kim B (2014) *Simplicity Theory*, vol. 53. *Oxford Logic Guides*. Oxford: Oxford University Press.
[21] Ramello S (2024) Model theory of valued fields with an endomorphism. arXic preprint.05043.
[22] Shuddhodan KV and Varshavsky Y (2022) The Hrushovski-Lang-Weil estimates. *Algebr. Geom.* **9**(6), 651–687.
[23] Simon P (2015) *A Guide to NIP Theories*, vol. 44. *Lecture Notes in Logic*. Association for Symbolic Logic, Chicago, IL. Cambridge: Cambridge Scientific Publishers.

https://doi.org/10.1017/S1474748025000167 Published online by Cambridge University Press