A contribution to the model theory of fields with free operators

A thesis submitted to the University of Manchester
for the degree of
Doctor of Philosophy
in the Faculty of Science and Engineering

2024

Shezad Mohamed

Department of Mathematics
School of Natural Sciences

Blank page

Contents

Word count: 28524.

The University of Manchester

Shezad Mohamed
Doctor of Philosophy
A contribution to the model theory of fields with free operators
October 2, 2024

This thesis makes a contribution to the model theory of fields with free operators, as introduced by Moosa and Scanlon. The classical Weil restriction, a result of algebraic geometry, establishes the existence of a left adjoint to base extension of algebras. Generalising the corresponding differential result of León Sánchez and Tressl, we extend this to the case of algebras equipped with free operators – given an extension of rings with free operators whose underlying extension of rings is free and of finite rank, and subject to a mild algebraic condition on the endomorphisms definable in the free operator structure, we show that there is a unique sequence of free operators on the classical Weil restriction that ensures the unit and counit of the classical adjunction preserve the free operator structure. Thus base change in the category of algebras with free operators has a left adjoint, which we call the $\mathcal{D}$-Weil restriction. Properties of the free operator structure preserved under the $\mathcal{D}$-Weil restriction are investigated, including triviality of the associated endomorphisms and commutativity of the operators, and a partial converse to the main adjunction result is shown: the existence of a left adjoint to base change over a field implies the associated endomorphisms must have the aforementioned algebraic condition.

The theory $\mathrm{UC}_{\mathcal{D}}$ in the language of rings with free operators is introduced as a suitable weakening of the geometric axiom of Moosa and Scanlon's theory of $\mathcal{D}$-closed fields $\mathcal{D}$-CF_0, the model companion of the theory of fields of characteristic zero with free operators. We show that whenever T is a model complete theory of difference large fields of characteristic zero – a notion of Cousins – $T \cup \mathrm{UC}_{\mathcal{D}}$ is the model companion of the theory $T \cup$"free operators", establishing the existence of the uniform companion for theories of difference large fields of characteristic zero with free operators, following Tressl's result in the differential context. We show that quantifier elimination transfers from T to $T \cup \mathrm{UC}_{\mathcal{D}}$ – from which it immediately follows that stability and NIP do as well – and we use the $\mathcal{D}$-Weil restriction to show that the algebraic closure of a model of $\mathrm{UC}_{\mathcal{D}}$ is a model of $\mathcal{D}$-CF_0.

We provide an axiomatic framework for proving the transfer of various neostability properties from theories of fields to theories of fields with operators, show that this unifies many proofs of stability and simplicity of theories of fields with operators existing already in the literature, and use it to characterise forking in the theory of separably differentially closed fields of infinite differential degree of imperfection, as defined by Ino and León Sánchez.

Finally, we introduce the class of bounded pseudo $\mathcal{D}$-closed fields in analogy to the class of bounded pseudo-differentially closed fields as a case study for some of the general results just described.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright Statement

Acknowledgements

First of all I would like to thank my supervisor Omar León Sánchez. It goes without saying that this project would not have been possible without his guidance; his mathematical influence can be seen throughout this thesis. Derek Goldrei, who sadly passed away in 2022. I would never have considered model theory without first seeing his love for it. All the friends I've made on the Atrium Bridge and all those from before Manchester. It's been quite a transitional time for me; there's no better group of people to have experienced it with. And lastly, I would like to thank Mum, Dad, Lena, Hassaan, Omar, Alex, and Tasha.

Introduction

This thesis is primarily concerned with $\mathcal{D}$-fields. These structures were introduced by Rahim Moosa and Thomas Scanlon in their trilogy [51, 52, 53] in order to provide a common framework for fields equipped with operators. As noted by Gogolok and Kowalski in [20], there have been several other attempts at formulating such a framework: Białynicki-Birula's *fields with operators* [5]; Buium's *jet operators* [8]; Hardouin's *iterative q-difference operators* [23]; Takeuchi's *C-ferential operators* [63]; and of course Gogolok and Kowalski's *$\mathcal{B}$-operators* [20]. All of these frameworks have strengths suited to the tasks their authors' introduced them for – differential Galois theory seems to be a common theme. Moosa and Scanlon introduced theirs to unify many model-theoretic properties common to the theories DCF_0 and ACFA_0; their series of papers culminated in proving the Zilber dichotomy for finite-dimensional minimal types in $\mathcal{D}$-CF_0, adapting the jet space methods of Pillay and Ziegler [56] (itself an adaptation of jet space arguments of Campana [9] and Fujiki [19] in the setting of complex manifolds) to the setting of $\mathcal{D}$-fields using their earlier construction of $\mathcal{D}$-jet spaces from [52]. This thesis will attempt to fill in some of the remaining model theory surrounding $\mathcal{D}$-fields of characteristic zero.

The uniform companion

Fix a base field k, a finite-dimensional k-algebra $\mathcal{D}$, and a k-algebra homomorphism $\pi\colon \mathcal{D} \to k$. A $\mathcal{D}$-field is then a field K extending k equipped with a k-algebra homomorphism $\partial\colon K \to K \otimes_k \mathcal{D}$ which is a section to $\mathrm{id}_K \otimes \pi$. Thus ordinary differential fields are an instance of this framework – a map $\delta\colon K \to K$ is a k-linear derivation if and only if the map

$$\begin{aligned} K &\to K[\varepsilon]/(\varepsilon^2) \\ a &\mapsto a + \delta(a)\varepsilon \end{aligned}$$

is a k-algebra homomorphism, as are ordinary difference rings – a map $\sigma \colon K \to K$ is a k-linear endomorphism if and only if the map

$$\begin{aligned} K &\to K \times K \\ a &\mapsto (a, \sigma(a)) \end{aligned}$$

is a k-algebra homomorphism.

In [53], Moosa and Scanlon proved that, subject to the condition that every maximal ideal of $\mathcal{D}$ has residue field k, the theory of $\mathcal{D}$-fields of characteristic zero, axiomatisable in the language of rings together with unary function symbols $\partial_1, \ldots, \partial_l$ (here $l+1$ is the dimension of $\mathcal{D}$ as a k-vector space), admits a model companion $\mathcal{D}\text{-CF}_0$, unifying the proofs establishing model companions for the theories of differential fields and difference fields of characteristic zero. A natural question to ask is whether other theories of $\mathcal{D}$-fields admit model companions.

In [65], Tressl constructs a theory of differential fields with m commuting derivations, UC_m, such that $T \cup \mathrm{UC}_m$ is the model companion of $T \cup$ "differential fields with m commuting derivations" whenever T is a model complete theory of large fields of characteristic zero. This will be our blueprint.

There is a complication however. Every $\mathcal{D}$-field has a sequence of definable endomorphisms, $\sigma_1, \ldots, \sigma_t$, called the associated endomorphisms.[1] Thus if some theory of $\mathcal{D}$-fields has a model companion, so must the reduct to the language of difference fields. In [33], Kikyo and Shelah prove that if T is model complete and has the strict order property, then the theory $T \cup$ "σ is an automorphism" has no model companion. This immediately implies that the theory of real closed fields equipped with an endomorphism has no model companion, and hence the theory of real closed fields equipped with a $\mathcal{D}$-field structure has no model companion whenever there is at least one nontrivial associated endomorphism. We resolve this dialectical tension by proving the following: *if T is a model complete theory of difference large fields of*

[1] This is the case subject to the condition that every maximal ideal of $\mathcal{D}$ has residue field k, one associated endomorphism for each maximal ideal except the one corresponding to π.

characteristic zero, then $T \cup \mathrm{UC}_{\mathcal{D}}$ *is the model companion of* $T \cup$ *"$\mathcal{D}$-fields".*

Following Cousins [15], a difference field $(K, \sigma_1, \ldots, \sigma_t)$ is called difference large if *for any pair of* K*-irreducible varieties* V *and* W *such that*

(i) $W \subseteq V \times V^{\sigma_1} \times \cdots \times V^{\sigma_t}$,

(ii) *the projections* $W \to V^{\sigma_i}$ *are dominant for all* $i = 0, \ldots, t$, *and*

(iii) W *has a smooth* K*-rational point,*

W *has a Zariski-dense set of* K*-rational points of the form* $(a, \sigma_1(a), \ldots, \sigma_t(a))$ *for* $a \in V(K)$. Difference largeness is just a suitable weakening of the geometric axiom of $\mathrm{ACFA}_{0,t}$; the only examples of difference large fields currently known to the author are models of $\mathrm{ACFA}_{0,t}$. However, in the case of fields (that is, when $t = 0$), K is difference large if and only if it is large: *every* K*-irreducible variety with a smooth* K*-rational point has a Zariski-dense set of* K*-rational points.* There are plenty of model complete large fields; Section 1.4 details some of them.

Similarly, the axiom scheme $\mathrm{UC}_{\mathcal{D}}$ is a suitable weakening of the geometric axiom scheme of $\mathcal{D}$-CF_0: *for every pair of* K*-irreducible varieties* V *and* W *such that*

(i) $W \subseteq \tau V$,

(ii) *the projections* $W \to V^{\sigma_i}$ *are dominant for each* $i = 0, \ldots, t$, *and*

(iii) W *has a smooth* K*-rational point,*

W *has a Zariski-dense set of* K*-rational points of the form* $\nabla(a)$ *for* $a \in V(K)$.

We establish the following two facts about the theory $\mathrm{UC}_{\mathcal{D}}$, analogously to Tressl's result for UC_m in [65]:

Theorem A.
1. *Suppose* $M, N \models \mathrm{UC}_{\mathcal{D}}$ *contain a common* $\mathcal{D}$*-subfield* A. *If the associated difference fields of* M *and* N *have the same existential theory over* A *as difference fields, then* M *and* N *have the same existential theory over* A *as* $\mathcal{D}$*-fields.*
2. *Every* $\mathcal{D}$*-field whose associated difference field is difference large can be extended to a model of* $\mathrm{UC}_{\mathcal{D}}$, *and the extension of their associated difference fields is elementary.*

The theorem establishing that $\mathrm{UC}_{\mathcal{D}}$ is indeed the uniform companion follows immediately.

Theorem B. *Let T be a model complete theory of difference large fields, and suppose it is the model companion of some T_0. Then*

(i) $T \cup \mathrm{UC}_{\mathcal{D}}$ *is the model companion of* $T_0 \cup$ *"$\mathcal{D}$-fields";*

(ii) *if T is the model completion of T_0, then $T \cup \mathrm{UC}_{\mathcal{D}}$ is the model completion of $T_0 \cup$ "$\mathcal{D}$-fields"; and*

(iii) *if T has quantifier elimination in some expansion by definitions, then $T \cup \mathrm{UC}_{\mathcal{D}}$ has quantifier elimination in the same expansion.*

In the case $\mathcal{D}$ is local, there are no nontrivial associated endomorphisms: the associated difference field is just the underlying field and difference largeness is just largeness; our result in this case is precisely the $\mathcal{D}$-field analogue of Tressl's. This yields the uniform companion in the following cases:

- several (not necessarily commuting) derivations;
- truncated, non-iterative higher derivations; and
- operators combining these two.

In particular, $\mathrm{RCF} \cup \mathrm{UC}_{\mathcal{D}}$ is the model companion of $\mathrm{RCF} \cup$ "$\mathcal{D}$-fields", and $\mathrm{Th}(\mathbb{Q}_p) \cup \mathrm{UC}_{\mathcal{D}}$ is the model companion of $\mathrm{Th}(\mathbb{Q}_p) \cup$ "$\mathcal{D}$-fields".

Remark. Restricted to the setting of derivations, the above result coincides with Tressl's only for the case of a single derivation. For that of several derivations, Tressl's deals with the commuting case, ours with the noncommuting case. However, the case of noncommuting derivations does appear in a recent paper of Fornasiero and Terzo [17] where they consider generic derivations on algebraically bounded structures – a wider context than the large and model complete fields considered here.

We then explore some equivalent characterisations in the local case: one in terms of D-varieties, and one analagous to the notion of *differential largeness* from [38]. Like Theorem 5.11 of that paper, we prove that algebraic extensions of models of $\mathrm{UC}_{\mathcal{D}}$ whose underlying field is large are again models of $\mathrm{UC}_{\mathcal{D}}$. This requires establishing the appropriate Weil restriction functor in the category of $\mathcal{D}$-algebras, and hence an interlude into non-model-theoretic geometry.

The $\mathcal{D}$-Weil descent

Suppose $T \to S$ is a morphism of schemes. Given a scheme Y over T, we define $W_{T/S}(Y)$ to be the scheme over S representing the functor

$$\begin{aligned}\mathsf{Sch}_S &\to \mathsf{Set}\\ U &\mapsto \mathrm{Hom}_T(U \times_S T, Y),\end{aligned}$$

if it exists. $W_{T/S}(Y)$ is uniquely determined, and we call it the Weil restriction of Y with respect to $T \to S$; see Section 1.3 of [67] for the original statement by Weil and [21] for Grothendieck's generalisation.

We will only be interested in the case when $S = \mathrm{Spec}(A)$ for some commutative ring A and $T = \mathrm{Spec}(B)$ where B is a finite and free A-algebra. Then $W_{T/S}$ is actually a functor on affine schemes

$$W_{B/A} \colon \mathsf{Aff}_B \to \mathsf{Aff}_A$$

which is right adjoint to base change. Hence we will often work with its algebraic dual, $W \colon \mathsf{Alg}_B \to \mathsf{Alg}_A$, which is left adjoint to base change

$$\begin{aligned}F \colon \mathsf{Alg}_A &\to \mathsf{Alg}_B\\ R &\mapsto R \otimes_A B.\end{aligned}$$

If $T \to S$ comes from a finite separable field extension, then $W_{T/S}(Y)$ is an abelian variety if Y is. In [48], Milne uses this fact to show that the Birch–Swinnerton-Dyer conjecture holds for Y if and only if it holds for $W_{T/S}(Y)$. Thus one can reduce the full conjecture from its statement over number fields to one over $\mathbb{Q}$.

This classical Weil restriction is also fundamental to the construction of prolongation spaces in the sense of Moosa and Scanlon [51], which we make extensive use of in establishing the uniform companion. For example, if K is a field, the tangent bundle of a K-variety V (a special case of their prolongation) can be seen as the Weil restriction of $V \times_K K[\varepsilon]/(\varepsilon^2)$ over $K \to K[\varepsilon]/(\varepsilon^2)$.

Furthermore, since $W_{T/S}(U)$ represents the functor $U \mapsto \mathrm{Hom}_T(U \times_S T, Y)$, we obtain a bijection between the T-points of U and the S-points of $W_{T/S}(U)$. This fact is used by Pop in [58] to show that algebraic extensions of large fields are large.

In [39] the case of differential algebras is considered. The authors show that the differential base change functor, F^δ, has a left adjoint, which they call the differential Weil descent functor, W^δ. More precisely, they show that if (A,∂) is a differential ring and (B,d) an (A,∂)-algebra, where B is finite and free as an A-module, then for any (B,d)-algebra (D,δ), there exists a unique derivation δ^W on $W(D)$ making the unit of the classical adjunction into a differential ring homomorphism. The authors then use this result in a similar way to Pop to show that algebraic extensions of differentially large fields are again differentially large (see [38]).

It is natural, then, to explore whether the difference base change functor, F^σ – here difference rings are rings equipped with a not necessarily injective endomorphism – also has a left adjoint. In general, it does not. Let A be a commutative ring with identity and consider the case when $B = A[\varepsilon]/(\varepsilon^2)$ for an indeterminate ε. Let $\tau\colon B \to B$ be given by $\tau(a + b\varepsilon) = a$ so that $(A, \mathrm{id}_A) \leq (B,\tau)$. Let $R = B[x]$ and let $\rho\colon R \to R$ be the unique endomorphism extending τ and sending $x \mapsto \varepsilon$. If $F^\sigma = F^\sigma_{B/A}$ had a left adjoint W^σ, then the unit of this adjunction at R

$$\eta^\sigma_R\colon R \to F^\sigma W^\sigma(R)$$

would be a difference ring homomorphism. In particular

$$\eta^\sigma_R(\rho(x)) = (\theta \otimes \tau)(\eta^\sigma_R(x)) \tag{$\star$}$$

where θ is the endomorphism of $W^\sigma(R)$. Let λ_1 and λ_2 be the coordinate projections with respect to the A-basis $\{1,\varepsilon\}$ of B. Then equation $(\star)$ translates to

$$\begin{bmatrix} \lambda_1(\rho(x)) \\ \lambda_2(\rho(x)) \end{bmatrix} = \begin{bmatrix} \lambda_1(\tau(1)) & \lambda_1(\tau(\varepsilon)) \\ \lambda_2(\tau(1)) & \lambda_2(\tau(\varepsilon)) \end{bmatrix} \begin{bmatrix} \theta(\lambda_1(\eta^\sigma_R(x))) \\ \theta(\lambda_2(\eta^\sigma_R(x))) \end{bmatrix}.$$

See Lemma 2.6.1 for details on this. Using the facts $\rho(x) = \varepsilon$, $\tau(1) = 1$, and $\tau(\varepsilon) = 0$, the above yields

$$\begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \theta(\lambda_1(W^\sigma_R(x))) \\ \theta(\lambda_2(W^\sigma_R(x))) \end{bmatrix},$$

which is clearly inconsistent. Hence, equation $(\star)$ cannot hold, and the left adjoint W^σ cannot exist. The issue here is that the 2×2 matrix on the right-hand side that we associate to (B,τ) is not invertible. In this case we say that τ does not have invertible matrix. We will see in the course of Section 2.4 that τ having invertible matrix is sufficient for a left adjoint to exist, and, in Section 2.5, that in the case when A is a field, it is also necessary.

Theorem C. *Let (A, σ) be a difference ring and (B, τ) a difference (A, σ)-algebra where B is finitely generated and free as an A-module. Assume that τ has invertible matrix. If (C, ρ) is a difference (B, τ)-algebra, then there is a unique endomorphism ρ^W on the classical Weil restriction, $W(C)$, making $(W(C), \rho^W)$ into a difference (A, σ)-algebra and the unit of the classical adjunction $\eta_C \colon C \to W(C) \otimes_A B$ into a difference ring homomorphism $(C, \rho) \to (W(C) \otimes_A B, \rho^W \otimes \tau)$. The assignment $(C, \rho) \mapsto (W(C), \rho^W)$ is the left adjoint to the difference base change functor.*

One might initially think to define $\rho^W = W(\rho)$. However, while ρ is a ring endomorphism, it is not in general a B-algebra homomorphism, and thus the functor W cannot be applied to it. There is a natural way to make ρ into a B-algebra homomorphism though: let C^τ be the B-algebra which, as a ring, is just C, but whose B-algebra structure is given by $b \mapsto \tau(b) \in C$; then ρ is a B-algebra homomorphism considered as a map $C \to C^\tau$. Applying W gives an A-algebra homomorphism $W(\rho) \colon W(C) \to W(C^\tau)$. However, this does not correspond to an endomorphism of $W(C)$. If we had an A-algebra homomorphism $W(C^\tau) \to W(C)^\sigma$, then composing with $W(\rho)$ gives an A-algebra homomorphism $W(C) \to W(C)^\sigma$, which corresponds to an endomorphism of $W(C)$ extending σ. In Section 2.3 we will see that such a map $W(C^\tau) \to W(C)^\sigma$ exists if τ has invertible matrix and in Section 2.4 that it yields the left adjoint.

As this is a thesis on $\mathcal{D}$-rings, we prove the above theorem in this more general setting. Let (A, e) be a $\mathcal{D}$-ring and (B, f) an (A, e)-algebra where B is finite and free as an A-module. Subject to the condition that every maximal ideal of $\mathcal{D}$ has residue field k, the $\mathcal{D}$-structure on B has associated endomorphisms and, as in the difference case, if the associated endomorphisms of (B, f) do not have invertible matrix, then the left adjoint to the $\mathcal{D}$-base change functor (see Definition 2.2.2) does not generally exist. Nonetheless, our main result states that this is indeed the main obstacle: if the associated endomorphisms of (B, f) have invertible matrix, then the $\mathcal{D}$-base change functor has a left adjoint, and if A is a field, this condition is necessary. See Theorem 2.4.5 and Corollary 2.5.21. Section 2.5 contains results on properties preserved under the $\mathcal{D}$-Weil descent. For instance, we show that if the $\mathcal{D}$-operators of some (B, f)-algebra pairwise commute, then the same is true of its $\mathcal{D}$-Weil restriction, and that if some associated endomorphism is an automorphism, then the same is true of its $\mathcal{D}$-Weil restriction. Thus our main result truly generalises the differential Weil descent of [39] and establishes the existence of Weil descent functors in the categories of difference rings with several (commuting) endomorphisms and difference rings with several (commuting) automorphisms. In addition we also prove the partial converse to the main theorem alluded to above: that if A is a field and the left adjoint exists, the associated endomorphisms must necessarily have invertible matrix.

Derivation-like theories and neostability

In Chapter 4 we return to a model-theoretic analysis and examine what neostability properties of models of $\mathrm{UC}_{\mathcal{D}}$ are determined by its underlying field. For instance, it is immediate from the transfer of quantifier elimination in Theorem B that if T is stable or NIP, then so is $T \cup \mathrm{UC}_{\mathcal{D}}$. For simplicity, we must do more work. To apply the Kim–Pillay theorem, we need to understand what nonforking independence looks like in the underlying field – we use the notion of slimness from [32]. The authors show that model complete, large fields are very slim in the language of rings, and hence that, in such a field, algebraic independence is an independence relation (in the sense of Adler [2]). From the proof of the Kim–Pillay theorem, we then get that if two $\mathcal{D}$-fields are independent in the sense of nonforking, they are algebraically independent as fields. This fact will allow us to amalgamate independent (in the sense of nonforking) $\mathcal{D}$-fields, and thus prove that if the independence theorem holds in T, it must also hold in $T \cup \mathrm{UC}_{\mathcal{D}}$.

In fact, the methods used are not reliant on the particular behaviour of $\mathcal{D}$-fields: they work for any suitable theory of fields with operators. Hence we formulate these results using an axiomatic approach. Given a complete and model complete $\mathfrak{L}$-theory T and a monster model $\mathbb{U}$ that has some relation $\mathop{\smile\!\!\!\!\!\!|}^0$ on triples of small subsets, we say that a $\mathfrak{D}$-theory Δ (for $\mathfrak{D} \supseteq \mathfrak{L}$) is *derivation-like* (with respect to T and $\mathop{\smile\!\!\!\!\!\!|}^0$) if the following four conditions hold:

(a) if $M \models \Delta$ and $M \leq_{\mathfrak{L}} N \models T$, then there is a $\mathfrak{D}$-structure on N extending the one on M such that $N \models \Delta$;

(b) if $M \models T \cup \Delta$ and $A \leq_{\mathfrak{D}} M \leq_{\mathfrak{L}} \mathbb{U}$, then $\mathrm{acl}_T(A) \leq_{\mathfrak{D}} M$ and $\mathrm{acl}_T(A) \models \Delta$; moreover, this is the only $\mathfrak{D}$-structure on $\mathrm{acl}_T(A)$ extending the one on A that makes $\mathrm{acl}_T(A)$ into a model of Δ;

(c) if $M \models T \cup \Delta$ with $M \leq_{\mathfrak{L}} \mathbb{U}$ and A and B are two models of Δ which are $\mathfrak{D}$-substructures of M with a common acl_T-closed $\mathfrak{D}$-substructure C such that $A \mathop{\smile\!\!\!\!\!\!\mid}^0_C B$, then $\langle AB\rangle_{\mathfrak{L}} \leq_{\mathfrak{D}} M$ and $\langle AB\rangle_{\mathfrak{L}} \models \Delta$; moreover, this is the only $\mathfrak{D}$-structure on $\langle AB\rangle_{\mathfrak{L}}$ extending the ones on A and B and making it into a model of Δ; and

(d) if A and B are two models of Δ which are $\mathfrak{L}$-substructures of $\mathbb{U}$ with a common acl_T-closed $\mathfrak{D}$-substructure C such that $A \mathop{\smile\!\!\!\!\!\!\mid}^0_C B$, then there is a $\mathfrak{D}$-structure on $\langle AB\rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} \mathbb{U}$ extending the ones on A and B that makes $\langle AB\rangle_{\mathfrak{L}}$ into a model of Δ.

If T is the theory of a very $\mathcal{L}_{\mathrm{ring}}(C)$-slim field of characteristic zero (where C is some set of constant symbols) and $\mathop{\smile\!\!\!\!\!\!\mid}^0$ is algebraic independence, then the following are examples of derivation-like theories:

1. differential fields with m noncommuting derivations;
2. differential fields with m commuting derivations;
3. $\mathcal{D}$-fields (where $\mathcal{D}$ is a local k-algebra); and
4. $\mathcal{D}$-fields (where $\mathcal{D}$ is a local k-algebra) with pairwise commuting operators.

Endomorphisms are not examples of derivation-like operators with respect to this choice of T (or even with $T = \mathrm{ACF}_0$): they fail axiom (a). In the case of characteristic $p > 0$, we can take $T = \mathrm{SCF}^{\lambda}_{p,\infty}$, the theory of separably closed fields of infinite degree of imperfection in the language of rings expanded by the λ-functions (see Proposition 27 of [16]), with nonforking independence $\mathop{\smile\!\!\!\!\!\!\mid}^0$. Then the theory of differential fields is also derivation-like with respect to this choice of T. If $T = \mathrm{ACFA}_{0,t}$ and $\mathop{\smile\!\!\!\!\!\!\mid}^0$ is nonforking independence, then $\mathcal{D}$-fields (where each maximal ideal of $\mathcal{D}$ has residue field k) is derivation-like with respect to T.

The main result of Chapter 4 is the following.

Theorem D. *Suppose Δ is derivation-like with respect to T and that $T \cup \Delta$ has a model companion T^+. Let $\mathfrak{C}$ be a monster model of (some completion of) T^+, and define the following relation on triples of subsets of $\mathfrak{C}$:*

$$A \mathop{\smile\!\!\!\!\!\!\mid}^*_C B \iff \operatorname{acl}(AC) \mathop{\smile\!\!\!\!\!\!\mid}^0_{\operatorname{acl}(C)} \operatorname{acl}(BC).$$

Then

(i) *if* $\mathop{\smile\!\!\!\!\!\!|}^0$ *is an abstract independence relation, so is* $\mathop{\smile\!\!\!\!\!\!|}^+$*;*

(ii) *if* $\mathop{\smile\!\!\!\!\!\!|}^0$ *is a strict independence relation, so is* $\mathop{\smile\!\!\!\!\!\!|}^+$*;*

(iii) *for some parameter set* M*, if* $\mathop{\smile\!\!\!\!\!\!|}^0$ *is an independence relation that satisfies the independence theorem over* M*, so is* $\mathop{\smile\!\!\!\!\!\!|}^+$*; and*

(iv) *for some parameter set* M*, if* $\mathop{\smile\!\!\!\!\!\!|}^0$ *is an independence relation that satisfies stationarity over some* M*, so is* $\mathop{\smile\!\!\!\!\!\!|}^+$*.*

Thus simplicity and stability of T are transferred to T^+. This unifies many of the proofs of stability and simplicity of theories of fields with operators occurring in the literature (the simplicity of $\mathcal{D}$-CF_0 as proved in Theorem 5.9 of [53] for instance). One novel result stemming from this axiomatic work is that we may characterise nonforking independence in $\mathrm{SDCF}_{p,\infty}$, the theory of separably differentially closed fields of characteristic $p > 0$ and infinite differential degree of imperfection, defined and analysed by Ino and León Sánchez in [28], as p-disjointness plus algebraic independence, analogously to the field-theoretic case of $\mathrm{SCF}_{p,\infty}$, the theory of separably closed fields of characteristic $p > 0$ and infinite degree of imperfection, characterised by Srour in [61].

Pseudo $\mathcal{D}$-closed fields

Finally, Chapter 5 functions as a case study for many of the general results stated throughout the thesis. We study the PAC substructures in $\mathcal{D}$-CF_0 using the definition from [25] of being existentially closed in every $\mathcal{L}_{\mathrm{ring}}(\partial)$-regular extension (that is, an extension of $\mathcal{D}$-fields $A \leq B$ where $\mathrm{acl}(A) \cap \mathrm{dcl}(B) = \mathrm{dcl}(A)$), and we show that they are characterised as those $\mathcal{D}$-fields which are models of $\mathrm{UC}_{\mathcal{D}}$ and PAC as fields. We use this in conjunction with Chapter 4 to prove simplicity and elimination of imaginaries for the theory of a bounded $\mathcal{D}$-field which is a PAC substructure in $\mathcal{D}$-CF_0, extending the corresponding differential results from Section 5 of [26] to the case of $\mathcal{D}$-fields.

Conventions. All rings are commutative with identity. Ring homomorphisms preserve the identity.

Chapter 1

Preliminaries

1.1 Model theory

While model theory appears throughout this thesis, it is not so thoroughly ingrained that a non-model-theorist is incapable of reading it. I will give a brief exposition of the model theory that makes an appearance; all of it can be found in [24], [43], and [64].

Predicate logic

First we fix a language $\mathcal{L}$. This language is a set of predicate symbols, function symbols, and constant symbols. Each predicate symbol and function symbol comes with a particular arity – some $n \in \mathbb{N}$. An $\mathcal{L}$-structure $\mathcal{M}$ is defined by the following data:

- a nonempty set M – the universe of the $\mathcal{L}$-structure;
- for each predicate symbol P of arity n, a subset $P^{\mathcal{M}} \subseteq M^n$;
- for each function symbol f of arity n, a function $f^{\mathcal{M}} \colon M^n \to M$; and
- for each constant symbol c, an element $c^{\mathcal{M}} \in M$.

These are the *interpretations* of the symbols of $\mathcal{L}$ in $\mathcal{M}$. We will also assume that there is always a 2-ary predicate symbol $=$ which is always interpreted as equality; since every structure has such a predicate, we do not include it in the language $\mathcal{L}$. The language $\mathcal{L}$ also contains an infinite set of variables x_i.

The $\mathcal{L}$-terms are defined as follows:

- the variables x_i are terms;
- the constant symbols are terms; and
- if $t_1, \ldots, t_n$ are terms and f is an n-ary function symbol, then $f(t_1, \ldots, t_n)$ is a term.

The atomic $\mathcal{L}$-formulas are defined as follows:

- if s, t are terms, then $s = t$ is an atomic formula; and
- if $t_1, \ldots, t_n$ are terms and P is an n-ary predicate symbol, then $P(t_1, \ldots, t_n)$ is an atomic formula.

The $\mathcal{L}$-formulas are defined as follows:

- the atomic formulas are formulas;
- if ϕ and ψ are formulas, then $\phi \wedge \psi$, $\phi \vee \psi$, $\neg\phi$, $\phi \to \psi$, $\phi \leftrightarrow \psi$ are formulas; and
- if ϕ is a formula and x_i is a variable, then $\exists x_i \phi$, and $\forall x_i \phi$ are formulas.

We write $\phi \in \mathcal{L}$ to mean that ϕ is an $\mathcal{L}$-formula.

If x_i is a variable appearing in some formula ϕ, then it is a bound variable if it only ever occurs within the scope of some quantifier $\exists x_i$ or $\forall x_i$. It is a free variable otherwise. A formula with no free variables is called a sentence. We write $\phi(x_1, \ldots, x_n)$ to stress that the free variables of ϕ are contained in the tuple $(x_1, \ldots, x_n)$.

If $\phi(x_1, \ldots, x_n)$ is an $\mathcal{L}$-formula and $(a_1, \ldots, a_n)$ is a tuple from $\mathcal{M}$, then we write $\mathcal{M} \models \phi(\bar{a})$ if $\phi(\bar{a})$ is true in $\mathcal{M}$ in the natural sense.[1] The formula $\phi(\bar{x})$ defines a subset of M^n in a natural way:

$$\phi(\mathcal{M}) = \{\bar{a} \in M^n \colon \mathcal{M} \models \phi(\bar{a})\}.$$

If we partition the variables of ϕ as $\phi(\bar{x}, \bar{y})$, where $\bar{x}$ has length n and $\bar{y}$ length m, then for any $\bar{b} \in M^m$ we can form the $\bar{b}$-definable set

$$\phi(\mathcal{M}, \bar{b}) = \{\bar{a} \in M^n \colon \mathcal{M} \models \phi(\bar{a}, \bar{b})\}.$$

[1] Tarski's definition of truth: the symbols $\neg$, $\wedge$, $\vee$, $\to$, $\leftrightarrow$ mean "not", "and", "or", "implies", "if and only if" respectively.

If T is a set of $\mathcal{L}$-sentences, then we say that T is consistent if there is some $\mathcal{L}$-structure $\mathcal{M}$ such that $\mathcal{M} \models \phi$ for each $\phi \in T$. In this case, we write $\mathcal{M} \models T$ and say that $\mathcal{M}$ is a model of T.

The Compactness Theorem. *A set of $\mathcal{L}$-sentences T is consistent if and only if every finite subset of it is consistent.*

A consistent set of $\mathcal{L}$-sentences is also called an $\mathcal{L}$-theory. If T is a set of $\mathcal{L}$-sentences and ϕ is an $\mathcal{L}$-sentence, then we say that T logically implies ϕ if every model of T is also a model of ϕ, and we write $T \models \phi$.

The $\mathcal{L}$-theory of an $\mathcal{L}$-structure $\mathcal{M}$ is the set of all $\mathcal{L}$-sentences true in $\mathcal{M}$.

$$\mathrm{Th}(\mathcal{M}) \coloneqq \{\phi \in \mathrm{Sent}(\mathcal{L})\colon \mathcal{M} \models \phi\}.$$

If two $\mathcal{L}$-structures $\mathcal{M}$ and $\mathcal{N}$ have the same $\mathcal{L}$-theory, then they are elementarily equivalent and we write $\mathcal{M} \equiv \mathcal{N}$. A maximally consistent set of $\mathcal{L}$-sentences is called a complete $\mathcal{L}$-theory. A consistent theory is complete if and only if for any two models $\mathcal{M}$ and $\mathcal{N}$, we have $\mathcal{M} \equiv \mathcal{N}$.

Many-sorted predicate logic

Everything defined in the previous section has been for one-sorted structures. However, we could have started with the notion of a *many-sorted structure*. A many-sorted language also has a decomposition into predicate, function, and constant symbols, but also has a set of sorts S. The arity of a predicate symbol is no longer some $n \in \mathbb{N}$, but a tuple $(s_1, \ldots, s_n) \in S^n$. The arity of a function symbol is $(s_1, \ldots, s_n, s) \in S^{n+1}$, and the arity of a constant symbol is $s \in S$.

If $\mathcal{L}$ is a many-sorted language, then an $\mathcal{L}$-structure $\mathcal{M}$ consists of the following data:

- for each sort $s \in S$, a nonempty set M_s;
- for each predicate symbol P of arity $(s_1, \ldots, s_n)$, a subset $P^{\mathcal{M}} \subseteq M_{s_1} \times \cdots \times M_{s_n}$;
- for each function symbol f of arity $(s_1, \ldots, s_n, s)$, a function $f^{\mathcal{M}}\colon M_{s_1} \times \cdots \times M_{s_n} \to M_s$;
- for each constant symbol c of arity s, an element $c^{\mathcal{M}} \in M_s$.

For each sort we have an infinite set of variables (whose arity is that sort), and terms and formulas are defined in the usual way but now with the extra condition that all the sorts are compatible. Unless discussing elimination of imaginaries, all our languages will be one-sorted.

Maps between structures

Suppose $\mathcal{M}$ and $\mathcal{N}$ are two $\mathcal{L}$-structures and $g\colon M \to N$ is a map between their underlying universes. We say that g is an $\mathcal{L}$-embedding $\mathcal{M} \to \mathcal{N}$ when the following conditions hold:

- $(a_1, \ldots, a_n) \in P^{\mathcal{M}} \iff (g(a_1), \ldots, g(a_n)) \in P^{\mathcal{N}}$ for every predicate symbol P (including equality);
- $g(f^{\mathcal{M}}(a_1, \ldots, a_n)) = f^{\mathcal{N}}(g(a_1), \ldots, g(a_n))$ for every function symbol f; and
- $g(c^{\mathcal{M}}) = c^{\mathcal{N}}$ for every constant symbol c.

If $M \subseteq N$, then $\mathcal{M}$ is an $\mathcal{L}$-substructure of $\mathcal{N}$ if inclusion is an $\mathcal{L}$-embedding. In this case we write $\mathcal{M} \leq \mathcal{N}$.

Remark 1.1.1. Suppose $\mathcal{M}$ is an $\mathcal{L}$-structure and A is a subset of M. Then for A to be an $\mathcal{L}$-substructure of $\mathcal{M}$ we need each constant $c^{\mathcal{M}}$ to lie in A and each function $f^{\mathcal{M}}$ to restrict to a function on A^n.

An $\mathcal{L}$-isomorphism is a bijective $\mathcal{L}$-embedding. An $\mathcal{L}$-embedding $g\colon \mathcal{M} \to \mathcal{N}$ is called $\mathcal{L}$-elementary if for any $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ and any $(a_1, \ldots, a_n) \in M^n$,

$$\mathcal{M} \models \phi(\bar{a}) \iff \mathcal{N} \models \phi(g(\bar{a})).$$

If the inclusion map $\mathcal{M} \leq \mathcal{N}$ is $\mathcal{L}$-elementary, say that $\mathcal{M}$ is an elementary substructure of $\mathcal{N}$ and write $\mathcal{M} \preceq \mathcal{N}$. If $A \subseteq M$ and $B \subseteq N$ and $g\colon A \to B$ is a map of sets, we say g is a partial $\mathcal{L}$-elementary map if for any $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ and any $(a_1, \ldots, a_n) \in A^n$,

$$\mathcal{M} \models \phi(\bar{a}) \iff \mathcal{N} \models \phi(g(\bar{a})).$$

Note that if $\mathcal{M} \leq \mathcal{N}$ and ϕ is an $\mathcal{L}$-sentence, then

1. if ϕ is existential, we have

$$\mathcal{M} \models \phi \implies \mathcal{N} \models \phi;$$

2. if ϕ is universal, we have

$$\mathcal{M} \models \phi \iff \mathcal{N} \models \phi.$$

Diagrams

Suppose $\mathcal{M}$ is an $\mathcal{L}$-structure and $A \leq \mathcal{M}$ is an $\mathcal{L}$-substructure. We can form the language $\mathcal{L}(A) = \mathcal{L} \cup \{c_a \colon a \in A\}$ by adding new constant symbols, and consider $\mathcal{M}$ as an $\mathcal{L}(A)$-structure $\mathcal{M}_A$ by interpreting each new constant symbol c_a as a. Then the (quantifier-free) diagram of A is the set of quantifier-free $\mathcal{L}(A)$-sentences true in $\mathcal{M}_A$. The diagram of A does not depend on which $\mathcal{M}$ it is computed in. We denote it by $\operatorname{diag}_{\mathcal{L}}(A)$ or $\operatorname{diag}(A)$ if no confusion arises. The key fact is that $\mathcal{N} \models \operatorname{diag}(A)$ if and only if there is some embedding $A \to \mathcal{N}$. One also defines the complete or elementary diagram of A as the set of $\mathcal{L}(A)$-sentences true in $\mathcal{M}_A$. Then $\mathcal{N} \models \operatorname{eldiag}(A)$ if and only if there is an elementary embedding $A \to \mathcal{N}$.

We will say that an $\mathcal{L}(A)$-formula is an $\mathcal{L}$-formula with parameters from A.

The Löwenheim–Skolem theorems

One aspect of first-order model theory is that it cannot distinguish between different sizes of infinity. The following two theorems formalise this idea.

The Downward Löwenheim–Skolem Theorem. *Suppose $\mathcal{M}$ is an $\mathcal{L}$-structure, $A \subseteq M$, and κ an infinite cardinal with $|\mathcal{L}| + |A| \leq \kappa \leq |M|$. Then $\mathcal{M}$ has an elementary substructure of cardinality κ containing A.*

The Upward Löwenheim–Skolem Theorem. *Let $\mathcal{M}$ be an infinite $\mathcal{L}$-structure and κ an infinite cardinal with $\kappa \geq |\mathcal{L}| + |M|$. Then $\mathcal{M}$ has an elementary extension of cardinality κ.*

Existentially closed models, model companions, and quantifier elimination

If $\mathcal{M} \leq \mathcal{N}$ is an extension of $\mathcal{L}$-structures, then $\mathcal{M}$ is existentially closed in $\mathcal{N}$ if for every existential $\mathcal{L}$-formula ϕ with parameters from M we have

$$\mathcal{N} \models \phi \implies \mathcal{M} \models \phi.$$

That is, for every quantifier-free $\mathcal{L}$-formula $\phi(\bar{x}, \bar{y})$ and every $\bar{a} \in M$,

$$\mathcal{N} \models \exists\bar{x}\phi(\bar{x}, \bar{a}) \implies \mathcal{M} \models \exists\bar{x}\phi(\bar{x}, \bar{a}).$$

If T_0 is some $\mathcal{L}$-theory, then $\mathcal{M} \models T_0$ is an existentially closed model of T_0 if for every $\mathcal{N} \models T_0$ with $\mathcal{N} \geq \mathcal{M}$, we have that $\mathcal{M}$ is existentially closed in $\mathcal{N}$.

An $\mathcal{L}$-theory T is called model complete if for any two models $\mathcal{M} \leq \mathcal{N}$, we have $\mathcal{M} \preceq \mathcal{N}$. So T is model complete if and only if for any model $\mathcal{M}$, we have that the $\mathcal{L}(\mathcal{M})$-theory $T \cup \mathrm{diag}(\mathcal{M})$ is complete.[2] If T is model complete, then every $\mathcal{L}$-formula $\phi(\bar{x})$ is T-equivalent to an existential $\mathcal{L}$-formula $\psi(\bar{x})$ and a universal $\mathcal{L}$-formula $\theta(\bar{x})$:

$$\begin{aligned} T &\models \forall\bar{x}(\phi(\bar{x}) \leftrightarrow \psi(\bar{x})); \\ &\models \forall\bar{x}(\phi(\bar{x}) \leftrightarrow \theta(\bar{x})). \end{aligned}$$

Suppose T_0 and T are two $\mathcal{L}$-theories. Then we say that T is a model companion of T_0 if the following conditions hold:

1. every model of T_0 embeds in a model of T;
2. every model of T embeds in a model of T_0;
3. T is model complete.

Conditions (1) and (2) together are equivalent to the fact that T_0 and T have the same universal theory.

If T_0 is inductive (that is, axiomatised by $\forall\exists$-sentences), then T_0 has a model companion if and only if the class of existentially closed models of T_0 is axiomatisable by some $\mathcal{L}$-theory. In this case, its model companion is given by the $\mathcal{L}$-theory axiomatising its existentially closed models.

Remark 1.1.2. Suppose T_0 is inductive and T is the model companion of T_0. Then every model of T is a model of T_0.

T is a model completion of T_0 if for any $\mathcal{M} \models T_0$, we have that $T \cup \mathrm{diag}(\mathcal{M})$ is complete. T has quantifier elimination if for any $\mathcal{M} \models T$ and any $\mathcal{L}$-substructure $\mathcal{A} \leq \mathcal{M}$, we have that $T \cup \mathrm{diag}(\mathcal{A})$ is complete. Fortunately, quantifier elimination does correspond to being able to eliminate quantifiers: T has quantifier elimination

[2]This is how Robinson first defined model completeness [60] and is where the name comes from.

if and only if for every $\mathcal{L}$-formula $\phi(\bar{x})$, there is some quantifier-free $\mathcal{L}$-formula $\psi(\bar{x})$ such that $T \models \forall \bar{x}(\phi(\bar{x}) \leftrightarrow \psi(\bar{x}))$.[3]

For use in Chapters 3 and 4, we summarise below.

Fact 1.1.3. *Suppose T_0 and T are two $\mathcal{L}$-theories with the same universal $\mathcal{L}$-theory. Then*

(i) *T is the model companion of T_0 if for every $\mathcal{M} \models T$, $T \cup \mathrm{diag}(\mathcal{M})$ is complete;*

(ii) *T is the model completion of T_0 if for every $\mathcal{M} \models T_0$, $T \cup \mathrm{diag}(\mathcal{M})$ is complete; and*

(iii) *T has quantifier elimination if for every $\mathcal{A} \leq \mathcal{M} \models T$, $T \cup \mathrm{diag}(\mathcal{A})$ is complete.*

Historical Remark. The notion of a model companion was first introduced by Robinson. It has seemed to be most useful in studying the model theory of fields – conveniently the topic of this thesis.

Expansions by definitions and Morleyisation

Suppose $\mathcal{L} \subseteq \mathcal{L}^*$ are two languages, T is an $\mathcal{L}$-theory, and $T^* \supseteq T$ is an $\mathcal{L}^*$-theory.

Definition 1.1.4. T^* is an expansion by definitions of T if

- for every new predicate symbol P of arity n in $\mathcal{L}^*$, there is some $\mathcal{L}$-formula $\phi_P(x_1, \ldots, x_n)$,
- for every new function symbol f of arity n in $\mathcal{L}^*$, there is some $\mathcal{L}$-formula $\phi_f(x_1, \ldots, x_n, y)$ such that $T \models \forall x_1 \ldots x_n \exists! y \; \phi_f(x_1, \ldots, x_n, y)$,[4]
- for every new constant symbol c in $\mathcal{L}^*$, there is some $\mathcal{L}$-formula $\phi_c(y)$ such that $T \models \exists! y \; \phi_c(y)$,

and T^* is logically equivalent to

$$\begin{aligned} T &\cup \{\forall x_1 \ldots x_n (P(x_1, \ldots, x_n) \leftrightarrow \phi_P(x_1, \ldots, x_n)) \colon P \text{ is a new predicate}\} \\ &\cup \{\forall x_1 \ldots x_n y (f(x_1, \ldots, x_n) = y \leftrightarrow \phi_f(x_1, \ldots, x_n, y)) \colon f \text{ is a new function}\} \\ &\cup \{\phi_c(c) \colon c \text{ is a new constant}\} \end{aligned}$$

[3] As long as $\bar{x}$ is not the empty tuple (that is, ϕ is not a sentence), ψ can always be taken so that its free variables also appear among $\bar{x}$. If ϕ is a sentence, then ψ may need to contain a free variable – if $\mathcal{L}$ has no constant symbols, then there are no quantifier-free $\mathcal{L}$-sentences.

[4] Here $\exists! y \phi(\bar{x}, y)$ means "there exists a unique y such that $\phi(\bar{x}, y)$ holds" and is just an abbreviation of $\exists y(\phi(\bar{x}, y) \wedge \forall z(\phi(\bar{x}, z) \to z = y))$.

Remark 1.1.5. 1. Note that the defining $\mathcal{L}$-formulas, ϕ_P, ϕ_f, and ϕ_c, are not allowed to contain any parameters.

2. If T^* is an expansion by definitions of T and $M \models T$, then M can be uniquely expanded to an $\mathcal{L}^*$-structure which is a model of T^*.

Example 1.1.6. There is a natural example of an expansion by definitions for any $\mathcal{L}$-theory T. For every $n \in \omega$ and every $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$, let R_ϕ be a new n-ary predicate symbol. Let $\mathcal{L}^* = \mathcal{L} \cup \{R_\phi \colon \phi(x_1, \ldots, x_n)$ is an $\mathcal{L}$-formula$\}$, and let T^* be the $\mathcal{L}^*$-theory

$$T^* := T \cup \{\forall x_1 \ldots x_n (R_\phi(x_1 \ldots x_n) \leftrightarrow \phi(x_1 \ldots x_n))\}.$$

T^* is called the Morleyisation of T, and it has quantifier elimination.

We can always Morleyise a theory, but it is often overkill; we may be able to find a more reasonable language in which our theory eliminates quantifiers. This also gives us a better understanding of the definable sets in our theory.

Example 1.1.7. Consider the $\mathcal{L}_{\text{ring}}$-theory of real closed fields, the model companion of the theory of formally real fields. This theory is model complete, but it does not have quantifier elimination: the set of non-negative elements, defined by $\exists y\ x = y^2$, is not quantifier-free $\mathcal{L}_{\text{ring}}$-definable – any quantifier free $\mathcal{L}_{\text{ring}}$-formula in a single variable defines a finite or cofinite set in a field.

Let $\mathcal{L}^* := \mathcal{L}_{\text{ring}} \cup \{\leq\}$, and let RCF* be the $\mathcal{L}^*$-theory RCF $\cup \{\forall x \forall y (x \leq y \leftrightarrow \exists z\ y - x = z^2)$. Then RCF* has quantifier elimination by the Tarski–Seidenberg theorem; see Section 3.3 of [43] for instance.

So every set definable in a real closed field is a Boolean combination of solution sets of polynomial equations and inequalities: a semialgebraic set.

Example 1.1.8. Consider the $\mathcal{L}_{\text{ring}}$-theory pCF of p-adically closed fields (of p-rank 1), the model companion of the theory of formally p-adic fields (of p-rank 1). Expanding by definitions to the language $\mathcal{L}_{\text{ring}}(\mathcal{O}, (P_n)_{n \in \mathbb{N}})$ by defining $\mathcal{O}$ as the valuation ring and each unary predicate P_n as the set of nth powers, this theory has quantifier elimination. See Theorem 5.6 of [59].

Types

Let $\mathcal{M}$ be an $\mathcal{L}$-structure. For some parameter set $A \subseteq \mathcal{M}$ and tuple $\bar{b} \in \mathcal{M}$, the (complete) type of $\bar{b}$ over A is

$$\mathrm{tp}(\bar{b}/A) = \{\phi(\bar{x}) \in \mathcal{L}(A) \colon \mathcal{M} \models \phi(\bar{b})\}.$$

Here $\bar{x}$ is a fixed tuple of variables of the same length (possibly infinite) as $\bar{b}$.

Remark 1.1.9. The definition above *prima facie* depends on $\mathcal{M}$. But note that if $\mathcal{N} \succeq \mathcal{M}$, then the type of $\bar{b} \in \mathcal{M}$ over $A \subseteq \mathcal{M}$ is the same computed in either $\mathcal{M}$ or $\mathcal{N}$.

Now let $p(\bar{x})$ be a set of $\mathcal{L}$-formulas with parameters from $A \subseteq \mathcal{M}$. We say that p is a type over A if there is some $\mathcal{N} \succeq \mathcal{M}$ and $\bar{b} \in \mathcal{N}$ such that $p = \mathrm{tp}(\bar{b}/A)$. If $\bar{b}$ can be taken in $\mathcal{M}$, then $\mathcal{M}$ realises p; otherwise it omits p. A set $\pi(\bar{x})$ of $\mathcal{L}$-formulas with parameters from $A \subseteq \mathcal{M}$ is called a partial type over A if it a subset of a type over A. Note then that partial types over A are precisely those sets of $\mathcal{L}$-formulas with parameters from A which are finitely satisfiable in $\mathcal{M}$ and that types are the partial types which are maximal (with respect to inclusion) among sets of $\mathcal{L}$-formulas which are finitely satisfiable in $\mathcal{M}$.

Remark 1.1.10. If $\bar{b}$ and $\bar{c}$ have the same type over A, we write $\bar{b} \equiv_A \bar{c}$. The map that fixes A and sends $\bar{b}$ to $\bar{c}$ is a partial elementary map.

We write $S(A)$ for the set of all types over A.

The monster

Algebraic geometry used to be conducted inside some *universal domain*: a large algebraically closed field containing all the points of all the varieties geometers were interested in. It has since progressed past this to the more sophisticated machinery of schemes. Our version of the universal domain is the *monster model*. The power of the monster model $\mathfrak{C}$ lies in the fact that it contains all the objects we would ever want: $\mathfrak{C}$ will contain realisations of all types over small subsets, any small model elementarily equivalent to $\mathfrak{C}$ will embed elementarily inside $\mathfrak{C}$, and any two tuples with the same type over some small parameter set A will be conjugate by some $\sigma \in \mathrm{Aut}(\mathfrak{C}/A)$.

This subsection follows the construction of Hodges in Section 10.4 of [24] and is also one of the two approaches taken by Tent and Ziegler in [64]. It requires some minor set theory, for which [29] is a good reference.

Definition 1.1.11. For a cardinal κ, an $\mathcal{L}$-structure $\mathcal{M}$ is called κ-saturated if whenever $A \subseteq M$ with $|A| < \kappa$ and p is a type over A, then p is realised in $\mathcal{M}$.

If $\mathcal{M}$ is of infinite cardinality κ, then it is called special if it is the union of an elementary chain $\bigcup_{\mu<\kappa} \mathcal{M}_\mu$ where each $\mathcal{M}_\mu$ is μ^+-saturated (here the elementary chain is indexed over *cardinals*, not ordinals).

Definition 1.1.12. A cardinal κ is a strong limit cardinal if $\mu < \kappa$ implies $2^\mu < \kappa$.

Example 1.1.13. Recall that the beth numbers are defined as follows.

- $\beth_0 = \aleph_0$;
- $\beth_{\alpha+1} = 2^{\beth_\alpha}$; and
- $\beth_\lambda = \bigcup_{\alpha<\lambda} \beth_\alpha$ for limit ordinals λ.

Let λ be any limit ordinal. Then $\beth_\lambda$ is a strong limit: if $\mu < \beth_\lambda = \bigcup_{\alpha<\lambda} \beth_\alpha$, then $\mu < \beth_\alpha$ for some $\alpha < \lambda$, and so $2^\mu \leq 2^{\beth_\alpha} = \beth_{\alpha+1} < \beth_\lambda$.

Definition 1.1.14. For an ordinal α, its cofinality $\mathrm{cf}(\alpha)$ is the least ordinal β such that there is an unbounded function $\beta \to \alpha$. A regular cardinal is one whose cofinality is itself.

Theorem 10.4.2 of [24]. *Let $\mathcal{M}$ be an infinite $\mathcal{L}$-structure and κ a strong limit cardinal greater than $|\mathcal{M}| + |\mathcal{L}|$. Then $\mathcal{M}$ has a special elementary extension of cardinality κ.*

We now examine the properties of special structures.

Definition 1.1.15. $\mathcal{M}$ is called strongly κ-homogeneous if every partial elementary map $A \to B$ with $|A|, |B| < \kappa$ extends to an automorphism of $\mathcal{M}$.

$\mathcal{M}$ is called κ-universal if every structure of cardinality strictly less than κ which is elementarily equivalent to $\mathcal{M}$ can be elementarily embedded in $\mathcal{M}$.

Theorem 10.4.5 of [24]. *If $\mathcal{M}$ is special of cardinality κ and A is a set of elements of size less than $\mathrm{cf}(\kappa)$, then $\mathcal{M}_A$ is special as an $\mathcal{L}(A)$-structure.*

The following is clear from the definition. It will be important in Chapter 4.

Theorem 1.1.16. *Every reduct of a special structure is special.*

Fact 1.1.17. *Special structures of cardinality κ are κ^+-universal, strongly* $\mathrm{cf}(\kappa)$*-homogeneous, and* $\mathrm{cf}(\kappa)$*-saturated.*

Proof. κ^+-universality is by Corollary 10.4.10 of [24]; strong $\mathrm{cf}(\kappa)$-homogeneity is by Corollary 10.4.6; $\mathrm{cf}(\kappa)$-saturation is by Corollary 10.4.12(a). ■

Remark 1.1.18. Suppose F is a normal function (that is, a strictly increasing continuous class function on ordinals). Then $\mathrm{cf}(F(\lambda)) = \mathrm{cf}(\lambda)$ for any limit ordinal λ. The map $\alpha \mapsto \beth_\alpha$ is such a function.

Now given a complete theory T, fix a regular cardinal γ larger than all the models of T and parameter sets we wish to consider. Use Theorem 10.4.2 of [24] to construct a special model $\mathfrak{C} \models T$ of cardinality $\beth_\gamma$. Then by Fact 1.1.17 and Remark 1.1.18, $\mathfrak{C}$ is γ-universal, strongly γ-homogeneous, and γ-saturated. This $\mathfrak{C}$ is our monster model for T.

Algebraic closure and definable closure

Let $\mathfrak{C}$ be the monster model. For $a \in \mathfrak{C}$ and $A \subseteq \mathfrak{C}$, a is said to be algebraic over A if it is contained in some finite A-definable set. The algebraic closure of A is

$$\mathrm{acl}(A) := \{a \in \mathfrak{C} \colon a \text{ is algebraic over } A\}.$$

An element $a \in \mathfrak{C}$ is definable over A if $\{a\}$ is A-definable. The definable closure of A is

$$\mathrm{dcl}(A) := \{a \in \mathfrak{C} \colon a \text{ is definable over } A\}.$$

A type $p = \mathrm{tp}(a/A)$ is algebraic if a is algebraic over A; equivalently if p has only finitely many realisations in $\mathfrak{C}$.

Elimination of imaginaries

Suppose X is some definable set in $\mathfrak{C}$. We say that a finite tuple d is a canonical parameter (or a code) for X if for every $\sigma \in \mathrm{Aut}(\mathfrak{C})$, σ fixes X setwise if and only if it fixes d pointwise.

$\mathrm{Aut}(\mathfrak{C})$ acts on $S(\mathfrak{C})$ in a natural way:

$$\sigma \cdot p(\bar{x}) = p^{\sigma}(\bar{x}) \coloneqq \{\phi(\bar{x}, \sigma(\bar{b}))\colon \phi(\bar{x}, \bar{b}) \in p\}.$$

A set B is the canonical base of $p \in S(\mathfrak{C})$ if it is fixed pointwise by every automorphism that fixes p.

Definition 1.1.19. A theory T eliminates imaginaries if for any 0-definable equivalence relation E, every equivalence class e/E has a canonical parameter.

It is often convenient to work in a theory that eliminates imaginaries, but not all theories do. However, one can always pass to T^{eq} which will eliminate imaginaries.

Enumerate all 0-definable equivalence relations on $\mathfrak{C}$ as $(E_i)_{i \in I}$ where E_i has arity n_i. We form the many-sorted structure $\mathfrak{C}^{\mathrm{eq}} = (\mathfrak{C}, \mathfrak{C}^{n_i}/E_i : i \in I)$ in the language of $\mathfrak{C}$ with extra function symbols $\pi_i \colon \mathfrak{C}^{n_i} \to \mathfrak{C}^{n_i}/E_i$ interpreted as the natural projections. The elements of $\mathfrak{C}$ are called real elements, and the elements of $\mathfrak{C}^{n_i}/E_i$ are called imaginary elements. T^{eq} is the theory of $\mathfrak{C}^{\mathrm{eq}}$; $\mathfrak{C}^{\mathrm{eq}}$ is its monster model.

Proposition 8.4.5 of [64]. *T^{eq} eliminates imaginaries.*

Definition 1.1.20. *T eliminates finite imaginaries* if every finite set of tuples has a canonical parameter.

T has *weak elimination of imaginaries* if for every imaginary e there is a real c such that $e \in \mathrm{dcl}^{\mathrm{eq}}(c)$ and $c \in \mathrm{acl}(e)$.

Corollary 8.4.6 of [64]. *T eliminates imaginaries if and only if in T^{eq} every imaginary is interdefinable with a real.*

Lemma 8.4.10 of [64]. *T eliminates imaginaries if and only if it has weak and finite elimination of imaginaries.*

The following is very useful for us.

Lemma 1.1.21. *If T is a theory of fields, then T eliminates finite imaginaries.*

Proof. This is shown in the proof of Corollary 8.4.12 of [64]. ■

Simplicity

Let T be a complete theory with infinite models and monster model $\mathfrak{C}$.

Definition 1.1.22. We say that a formula $\phi(x, b)$ k-divides over $A \subseteq \mathfrak{C}$ if there is a sequence $(b_i)_{i\in\omega}$ of realisations of $\operatorname{tp}(b/A)$ such that $\{\phi(x, b_i) : i \in \omega\}$ is k-inconsistent. It divides over A if it k-divides for some k. A partial type $\pi(x)$ divides over A if it implies some formula which divides over A.

The partial type $\pi(x)$ forks over A if it implies a finite disjunction of formulas each of which divides over A.

We write $A \mathop{\smile\!\!\!\!\!\!\mid}_C B$ and say A is (nonforking) independent from B over C if $\operatorname{tp}(A/BC)$ does not fork over C.

T is simple if nonforking independence is symmetric:

$$A \mathop{\smile\!\!\!\!\!\!\mid}_{C} B \iff B \mathop{\smile\!\!\!\!\!\!\mid}_{C} A.$$

In this thesis, we will follow Adler [2] in his treatment of *abstract independence relations.*

Definition 1.1.23. A relation $\mathop{\smile\!\!\!\!\!\!\mid}^*$ on triples of small subsets of $\mathfrak{C}$ is called an abstract independence relation if it is invariant under automorphisms and satisfies the following conditions.

1. *normality:* $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B \implies X \mathop{\smile\!\!\!\!\!\!\mid}^*_A AB$;
2. *monotonicity:* $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B \implies X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B'$ for $B' \subseteq B$;
3. *base monotonicity:* $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A D \implies X \mathop{\smile\!\!\!\!\!\!\mid}^*_B D$ for $A \subseteq B \subseteq D$;
4. *transitivity:* $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B$ and $X \mathop{\smile\!\!\!\!\!\!\mid}^*_B D \implies X \mathop{\smile\!\!\!\!\!\!\mid}^*_A D$ for $A \subseteq B \subseteq D$;
5. *symmetry:* $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B \iff B \mathop{\smile\!\!\!\!\!\!\mid}^*_A X$;
6. *full existence:* for any X, A, B there is $X' \equiv_A X$ such that $X' \mathop{\smile\!\!\!\!\!\!\mid}^*_A B$ (recall that $X' \equiv_A X$ means that X' and X have the same type over A);
7. *finite character:* if $X_0 \mathop{\smile\!\!\!\!\!\!\mid}^*_A B$ for all finite $X_0 \subseteq X$, then $X \mathop{\smile\!\!\!\!\!\!\mid}^*_A B$;
8. *local character:* there is a cardinal κ such that for all X and A, there is $A_0 \subseteq A$ with $|A_0| < \kappa$ such that $X \mathop{\smile\!\!\!\!\!\!\mid}^*_{A_0} A$.

There are three extra properties that an abstract independence relation $\mathop{\smile\!\!\!\!\!\!\mid}^*$ might satisfy that we are interested in:

9. *strictness:* if $b \mathop{\smile\!\!\!\!\!\!\mid}^*_A b$, then $b \in \operatorname{acl}(A)$;
10. *independence theorem over M:* if $A_1 \mathop{\smile\!\!\!\!\!\!\mid}^*_M A_2$, $a_1 \mathop{\smile\!\!\!\!\!\!\mid}^*_M A_1$, $a_2 \mathop{\smile\!\!\!\!\!\!\mid}^*_M A_2$, and $a_1 \equiv_M a_2$, then there is $a \models \operatorname{tp}(a_1/MA_1) \cup \operatorname{tp}(a_2/MA_2)$ with $a \mathop{\smile\!\!\!\!\!\!\mid}^*_M A_1A_2$;
11. *stationarity over M:* whenever $A \supseteq M$, $a, b \in \mathfrak{C}$ with $a \equiv_M b$, $a \mathop{\smile\!\!\!\!\!\!\mid}^*_M A$ and $b \mathop{\smile\!\!\!\!\!\!\mid}^*_M A$, then $a \equiv_A b$.

For 10 and 11, M is usually an acl-closed set or a model.

These axioms also appear in various forms throughout the literature – often with *normality*, *monotonicity*, *base monotonicity*, and *transitivity* combined into a single axiom, and sometimes with *extension* instead of *full existence* – see Theorem 7.3.13 of [64] and Definition 4.1 of [35].

One of the results that sparked widespread interest in the study of simple theories was the Kim–Pillay theorem [35]. This gave a semantic way of proving a theory was simple and characterising the behaviour of nonforking independence.

The Kim–Pillay Theorem. *Suppose $\mathop{\smile\!\!\!\!\!\!\mid}^*$ is an abstract independence relation on T which satisfies the independence theorem over models. Then T is simple and $\mathop{\smile\!\!\!\!\!\!\mid}^*$ coincides with nonforking independence $\mathop{\smile\!\!\!\!\!\!\mid}$.*

Remark 1.1.24. In any theory T, nonforking independence satisfies normality, monotonicity, base monotonicity, finite character, and strictness. If nonforking independence satisfies transitivity, symmetry, or local character, then T is simple.

Stability

Let T be a complete theory with infinite models, and $\mathfrak{C}$ a monster model. Let κ be some infinite cardinal. We say that T is κ-stable if for any parameter set $A \subseteq \mathfrak{C}$ with $|A| \leq \kappa$, we have $|S(A)| \leq \kappa$. We say ω-stable instead of $\aleph_0$-stable.

T is stable if it is κ-stable for some κ. An equivalent characterisation is if no formula has the order property.

Definition 1.1.25. Let $\phi(x, y)$ be a formula whose free variables are partitioned into two tuples x and y. We say that $\phi(x, y)$ has the order property if there are sequences of tuples $(a_i)_{i\in\omega}$ and $(b_i)_{i\in\omega}$ from $\mathfrak{C}$ such that

$$\mathfrak{C} \models \phi(a_i, b_j) \iff i < j.$$

T is stable if no formula has the order property.

We will be more interested in a third characterisation, one in terms of nonforking independence, similar to the Kim–Pillay theorem. Indeed, simple theories initially received interested because it was observed that much of the desirable behaviour of nonforking independence in stable theories also held in simple theories.

Fact 2.1.4 of [34]. *Suppose $\mathop{\smile\!\!\!\!\!\!\mid}^*$ is an abstract independence relation on $\mathfrak{C}$ which satisfies stationarity over models. Then T is stable and $\mathop{\smile\!\!\!\!\!\!\mid}^*$ coincides with nonforking independence.*

The independence property

As before, let T be a complete theory and $\mathfrak{C}$ a monster model. We say that a formula $\phi(x, y)$ has the *independence property* (IP) if there are $(a_i)_{i\in\omega}$ and $(b_I)_{I\subseteq\omega}$ in $\mathfrak{C}$ such that

$$\mathfrak{C} \models \phi(a_i, b_I) \iff i \in I.$$

T is NIP if no formula has the independence property.

Remark 1.1.26. As is usually the case with combinatorial neostability properties, it does not matter whether we allow the formula $\phi(x, y)$ to contain parameters from $\mathfrak{C}$ or not.

1.2 Algebraic geometry

In this section we will briefly explain the geometric perspective we will take throughout this thesis. Analysing the model theory of fields lends itself to the classical viewpoint of algebraic geometry: that of affine varieties being solution sets of polynomials. However, the construction of the prolongation of a variety is more naturally done in the language of schemes. Hence we will lay out the classical viewpoint first, and then briefly explain how to translate concepts into the scheme-theoretic viewpoint. This final aspect is taken from Section 10.8 of [18].

The classical viewpoint

Let $\mathbb{U}$ be an algebraically closed field which is κ-saturated, strongly κ-homogeneous, and κ-universal for some large enough cardinal κ. This is our universal domain.

Affine n-space, $\mathbb{A}^n$, is the set of all n-tuples of elements of $\mathbb{U}$. Let K be a small subfield of $\mathbb{U}$ and $X = (X_1, \ldots, X_n)$ a tuple of variables. For each subset $\mathfrak{a} \subseteq K[X]$, we define the K-algebraic (or K-closed) set

$$V(\mathfrak{a}) := \{x \in \mathbb{A}^n \colon f(x) = 0 \text{ for all } f \in \mathfrak{a}\}.$$

Then the K-algebraic sets form the closed sets of a topology on $\mathbb{A}^n$. This is the K-Zariski topology on $\mathbb{A}^n$.

Now for any $A \subseteq \mathbb{A}^n$, define the following ideal of $K[X]$

$$I_K(A) := \{f \in K[X] \colon f(x) = 0 \text{ for all } x \in A\}.$$

Then V and I_K define an inclusion-reversing correspondence between K-closed subsets of $\mathbb{A}^n$ and radical ideals of $K[X]$.

A K-closed set V is called K-reducible if it can be written as the union of two proper K-closed sets. It is K-irreducible otherwise. A K-closed set V is K-irreducible if and only if $I_K(V)$ is a prime ideal in $K[X]$. In this case, the coordinate ring of V is $K[V] = K[X]/I_K(V)$, and the function field of V is the quotient field of its coordinate ring, denoted by $K(V)$.

Suppose $f_1, \ldots f_m \in K[X]$ so that $V = V(f_1, \ldots, f_m)$ is a K-closed set. Then for any $L \geq K$, $f_i \in L[X]$, and hence V is also an L-closed set. However, if V is K-irreducible, it might not be L-irreducible. We say that V is absolutely irreducible if it is L-irreducible for every $L \geq K$, or equivalently, if it is $\tilde{K}$-irreducible (here $\tilde{K}$ is the algebraic closure of K).

If V is K-irreducible, then $K[V] = K[X]/I_K(V)$ embeds in $\mathbb{U}$. Let x be the image of $X + I_K(V)$ under this embedding. The point $x \in V$ is called a K-generic point of V. Now given any point $x \in \mathbb{A}^n$, there is a K-irreducible K-variety V such that x is a K-generic point of V; take $V = V(I_K(x))$. We write $V = \operatorname{loc}(x/K)$.

For a field extension $L \geq K$, a point $x \in V$ is L-rational if its entries are in L. The L-rational points of V correspond to K-algebra homomorphisms $K[V] \to L$. We write the set of L-rational points of V as $V(L)$. Thus $V(\mathbb{U}) = V$.

Given a K-variety $V = V(f_1, \ldots, f_m)$ and a homomorphism of fields $\sigma \colon K \to L$, we can form the *conjugation* of V by σ:

$$V^\sigma := V(f_1^\sigma, \ldots, f_m^\sigma),$$

where $f_i^\sigma \in L[X]$ is the polynomial formed from f_i by applying σ to its coefficients. Since $\mathbb{U}$ is strongly κ-homogeneous, σ extends to an automorphism of $\mathbb{U}$ and induces a map

$$\begin{aligned} V(\mathbb{U}) &\to V^\sigma(\mathbb{U}) \\ (x_1, \dots, x_n) &\mapsto (\sigma(x_1), \dots, \sigma(x_n)). \end{aligned}$$

Translating to the language of schemes

Now for each K-closed set V, there is an associated reduced affine scheme of finite type over K. Since $V = V(\mathfrak{a})$ for some radical ideal $\mathfrak{a} \subseteq K[X]$, let $R = K[X]/\mathfrak{a}$; this ring is reduced and finitely generated over K. The associated scheme over K is $V' = \operatorname{Spec} R \to \operatorname{Spec} K$. Now each $x \in V$ corresponds to a K-algebra homomorphism $R \to \mathbb{U}$. The kernel of this homomorphism defines a prime ideal of R, and hence a point of V'. So K-isomorphism classes of points of V correspond to points of V'.

For $L \geq K$, V is also an L-closed set, but V' is not a scheme over L. Hence we consider the base change of V' to L, given by $V' \times_K L$ (formally $V' \times_{\operatorname{Spec} K} \operatorname{Spec} L$). An L-rational point of V' is a K-morphism $\operatorname{Spec} L \to V'$. This corresponds to a K-algebra homomorphism $R \to L$, and hence a point in $V(L)$. Note that a K-morphism $\operatorname{Spec} L \to V'$ corresponds to an L-morphism $\operatorname{Spec} L \to V' \times_K L$. Hence $V'(L) = (V' \times_K L)(L)$.

1.3 The Weil restriction

In this section we briefly go over the details of the construction of the classical Weil descent. We will not give proofs, but the reader can consult Section 7.6 of [6] and Section 2 of [51] for further details. Our approach is modelled after [39], so the reader can also consult there for a more in-depth explanation.

Let A be a ring, and B an A-algebra. For any A-algebra R we can form the base change[5] of R to B, namely $R \otimes_A B$, where the B-algebra structure is given by $b \mapsto 1 \otimes b$. This base change naturally extends to a functor $F\colon \mathsf{Alg}_A \to \mathsf{Alg}_B$ where we set $F(\phi) = \phi \otimes \mathrm{id}_B$. If we let $G\colon \mathsf{Alg}_B \to \mathsf{Alg}_A$ be the scalar restriction functor, where $G(C)$ is the A-algebra given by composing $A \to B \to C$, then G is right

[5]This is just the dual of the geometric base change given in the previous section.

adjoint to F. More importantly, if B is free and of finite rank as an A-module, then F has a left adjoint: Weil restriction $W\colon \mathsf{Alg}_B \to \mathsf{Alg}_A$.

We state the following useful fact about adjunctions from Theorem 2 and Corollaries 1 and 2 of [41].

Theorem 1.3.1. *Let $F\colon \mathcal{X} \to \mathcal{Y}$ be a functor, and suppose that for each $C \in \mathcal{Y}$, there is some $W(C) \in \mathcal{X}$ and $\eta_C\colon C \to F(W(C))$ in $\mathcal{Y}$ such that the assignment $g \mapsto F(g) \circ \eta_C$ is a bijection $\mathrm{Hom}_{\mathcal{X}}(W(C), R) \to \mathrm{Hom}_{\mathcal{Y}}(C, F(R))$. Then W extends to a functor $\mathcal{Y} \to \mathcal{X}$ which is left adjoint to F. The unit of this adjunction is given by η_C.*

In particular, for a morphism $C \xrightarrow{f} C'$ in $\mathcal{Y}$, $W(f)$ is defined to be the unique morphism $W(C) \xrightarrow{g} W(C')$ such that $F(g) \circ \eta_C = \eta_{C'} \circ f$.

This fact will allow us to construct the left adjoint using only the data of its object map and unit. This fact is also the method of proof for the differential Weil descent in [39].

We now explain the situation in the classical setup. Let $b_1, \ldots, b_r$ be an A-basis of B. For each $i = 1, \ldots, r$, let $\lambda_i\colon B \to A$ be the A-module homomorphism with $\lambda_i\left(\sum_{j=1}^r a_j b_j\right) = a_i$. If R is an A-algebra, we consider the base change of λ_i to R – the R-module homomorphism $\mathrm{id}_R \otimes \lambda_i\colon R \otimes_A B \to R$. Note that $\mathrm{id}_R \otimes \lambda_i$ simply picks out the coefficient of the basis element $1 \otimes b_i$. We will write λ_i for $\mathrm{id}_R \otimes \lambda_i$ throughout, but it will be clear from context which we mean.

Now let T be a set of indeterminates, and define

$$W(B[T]) = A[T]^{\otimes r} = A[T] \otimes_A \cdots \otimes_A A[T]$$

For each i and $t \in T$, let $t(i) = 1 \otimes \cdots \otimes 1 \otimes t \otimes 1 \otimes \cdots \otimes 1$, where the t occurs in the ith position. We also let $\eta_{B[T]}$ be the B-algebra homomorphism

$$\begin{aligned} \eta_{B[T]}\colon B[T] &\to F(W(B[T])) \\ t &\mapsto \sum_{i=1}^r t(i) \otimes b_i \end{aligned}$$

These choices make the following map $\tau(B[T], R)$ a bijection for each A-algebra

R:

$$\begin{aligned} \mathrm{Hom}_{\mathsf{Alg}_A}(A[T]^{\otimes r}, R) &\to \mathrm{Hom}_{\mathsf{Alg}_B}(B[T], R \otimes_A B) \\ \phi &\mapsto F(\phi) \circ \eta_{B[T]} \end{aligned}$$

where the compositional inverse is defined as follows. For a B-algebra homomorphism $\psi\colon B[T] \to R \otimes_A B$, let ϕ be the unique A-algebra homomorphism with $\phi(t(i)) = \lambda_i(\psi(t))$.

Now let C be a B-algebra, and take a surjective B-algebra homomorphism $\pi_C\colon B[T] \to C$ for some set of indeterminates T. Let I_C be the ideal of $W(B[T])$ generated by all the $\lambda_i(\eta_{B[T]}(f))$ where f ranges over $\ker \pi_C$. Now define $W(C) = W(B[T])/I_C$ and $W(\pi_C)\colon W(B[T]) \to W(C)$ as the residue map.

Then we induce a map $\tau(C,R)\colon \mathrm{Hom}_{\mathsf{Alg}_A}(W(C), R) \to \mathrm{Hom}_{\mathsf{Alg}_B}(C, F(R))$ which makes the following diagram commute:

$$\begin{array}{ccc} \mathrm{Hom}_{\mathsf{Alg}_A}(W(C), R) & \xrightarrow{\tau(C,R)} & \mathrm{Hom}_{\mathsf{Alg}_B}(C, F(R)) \\ \Big\downarrow{\scriptstyle - \circ W(\pi_C)} & & \Big\downarrow{\scriptstyle - \circ \pi_C} \\ \mathrm{Hom}_{\mathsf{Alg}_A}(W(B[T]), R) & \xrightarrow{\tau(B[T],R)} & \mathrm{Hom}_{\mathsf{Alg}_B}(B[T], F(R)) \end{array}$$

Let $\eta_C = \tau(C, W(C))(\mathrm{id}_{W(C)})$, and note that

$$\eta_C(\pi_C(t)) = \sum_{i=1}^{r} W(\pi_C)(t(i)) \otimes b_i$$

From this we see that $\tau(C,R)(\phi) = F(\phi) \circ \eta_C$ and that $\tau(C,R)$ is a bijection, satisfying the conditions of Theorem 1.3.1. Then W is a functor which is left adjoint to F with unit η_C. This W is the classical Weil descent functor.

1.4 Some field theory

We will not need much field theory, but we record some notions and facts that will be important later.

Linear and algebraic disjointness

Let $\mathbb{U}$ be a special algebraically closed field containing small subfields $F \leq L$, and $F \leq K$. The compositum of L and K inside $\mathbb{U}$ is the smallest subfield of $\mathbb{U}$ containing both L and K.

We say that L and K are linearly disjoint over F if every finite subset of L which is linearly independent over F remains linearly independent over K. This is equivalent to saying that the multiplication map

$$\begin{aligned} L \otimes_F K &\to LK \\ a \otimes b &\mapsto ab \end{aligned}$$

is an injection. Note then that linear disjointness is symmetric in L and K.

We say that L and K are algebraically independent, or algebraically disjoint, or free, over F if every finite subset of L that is algebraically independent over F remains algebraically independent over K. If this is the case, we write $L \mathop{\smile\!\!\!\!\!\!|}\nolimits_F^{\text{alg}} K$.

Fact 1.4.1. *If L and K are linearly disjoint over F, then they are algebraically independent. The converse holds if at least one of the extensions $F \leq L$ or $F \leq K$ is a regular extension, that is, relatively algebraically closed and separable.*

Large fields

For Chapter 3 we will need the notion of a large field. These were first introduced by Pop [58] as fields over which regular inverse Galois problems could be solved. Model-theoretically, most “tame” fields are large. We recall the definition, equivalent characterisations, and some examples.

Definition 1.4.2. A field K is called large if every K-irreducible variety with a smooth K-rational point has a Zariski-dense set of K-rational points.

Proposition 1.1 of [58]. *The following are equivalent.*

1. *K is large;*
2. *if a curve over K has a smooth K-rational point, then it has infinitely many K-rational points; and*
3. *K is existentially closed in the Laurent series field $K((t))$.*

Example 1.4.3. 1. Algebraically closed fields, real closed fields, and fields that admit a nontrivial Henselian valuation are all large. Pseudo-algebraically closed fields, pseudo-real closed fields, and pseudo-p-adically closed fields are large.

2. By Falting's theorem, number fields are not large.

Large fields play an important role in Tressl's uniform companion [65]. They will play a similar role in the results of Chapter 3 once we generalise to the notion to difference largeness.

Fields of positive characteristic

We now explain some of the algebra and model theory of separably closed fields of characteristic $p > 0$. This part is heavily based on [14], [16], and [46].

Definition 1.4.4. Let K be a field of characteristic $p > 0$. For $n \in \mathbb{N}$, let p^n be the set of n-tuples with entries from $\{0, \ldots, p-1\}$.

a) For a finite set $\bar{a} = (a_1, \ldots, a_n) \subseteq K$, the p-monomials over $\bar{a}$ are the elements $m_i(\bar{a}) := a_1^{i(1)} \cdots a_n^{i(n)} \in K$ for $i \in p^n$. The finite set $\bar{a}$ is p-independent in K if the set of p-monomials over $\bar{a}$ is linearly independent over K^p. An infinite set is p-independent if each finite subset is.

b) The set $\bar{a}$ is p-independent in K over $F \subseteq K$ if it is linearly independent over FK^p.

c) A p-basis of K is a maximal p-independent subset of K. The cardinality of a p-basis is called the degree of imperfection.

d) A field extension $F \subseteq K$ is separable if each p-independent set in F remains p-independent in K, equivalently if there is some p-basis of F that is p-independent in K, equivalently if F and K^p are linearly disjoint over F^p.

e) K is separably closed if it has no proper separable algebraic extension.

The theory of separably closed fields of characteristic p and degree of imperfection $e \in \mathbb{N} \cup \{\infty\}$ is denoted $\mathrm{SCF}_{p,e}$. In Proposition 27 of [16], Delon finds a language in which $\mathrm{SCF}_{p,e}$ has quantifier elimination.

If e is finite, define the language

$$\mathcal{L}^\lambda := \{+, -, \cdot, 0, 1\} \cup \{\lambda_i \colon i \in p^e\}$$

where each λ_i is an $(e+1)$-ary function symbol. We let $\mathrm{SCF}^{\lambda}_{p,e}$ be the expansion by definitions of $\mathrm{SCF}_{p,e}$ to $\mathcal{L}^{\lambda}$ given by defining the functions λ_i as follows. If $\bar{y}$ is p-dependent, then $\lambda_i(x,\bar{y}) = 0$ for each $i \in p^e$. Otherwise,

$$x = \sum_{i \in p^e} \lambda_i(x,\bar{y})^p \, m_i(\bar{y}).$$

If e is infinite, we let $\mathcal{L}^{\lambda}$ be the language

$$\mathcal{L}^{\lambda} \coloneqq \{+,-,\cdot,0,1\} \cup \{\lambda_{n,i} \colon n \in \omega, i \in p^n\}$$

where $\lambda_{n,i}$ is an $(n+1)$-ary function symbol. We let $\mathrm{SCF}^{\lambda}_{p,\infty}$ be the expansion by definitions of $\mathrm{SCF}_{p,\infty}$ to $\mathcal{L}^{\lambda}$ given by defining the functions $\lambda_{n,i}$ as follows. If $\bar{y}$ is p-dependent or $(x,\bar{y})$ is p-independent, then $\lambda_{n,i}(x,\bar{y}) = 0$ for each $n \in \omega$ and $i \in p^n$. Otherwise,

$$x = \sum_{i \in p^n} \lambda_{n,i}(x,\bar{y})^p \, m_i(\bar{y}).$$

The key fact about the λ functions is that $\mathcal{L}^{\lambda}$-extensions are precisely the separable extensions. See Lemma 1.9 of [11].

Fact 1.4.5. *Let K be a field of characteristic p and degree of imperfection e in the language $\mathcal{L}$. Let F be a subfield of K. Then F is an $\mathcal{L}$-substructure if and only if K/F is a separable field extension.*

Definition 1.4.6. Suppose that $F \leq K, L \leq \mathbb{U}$ are all separable extensions. We say that K and L are p-disjoint over F (inside $\mathbb{U}$) if every subset of K which is p-independent over F in $\mathbb{U}$ remains p-independent over L in $\mathbb{U}$.

Equivalently, there are p-bases B_F, B_K, and B_L of F, K and L, respectively, with $B_F \subseteq B_K, B_L$ such that $B_K \cup B_L$ is p-independent in $\mathbb{U}$.

This notion guarantees that composita of separable subfields remain separable.

Fact 1.4.7. *Suppose K and L are p-disjoint over F inside $\mathbb{U}$. Then $KL \subseteq \mathbb{U}$ is separable.*

Proof. $B_K \cup B_L$ is a p-basis of KL. By p-disjointness, $B_K \cup B_L$ is p-independent in $\mathbb{U}$. ■

The following fact characterises nonforking independence in $\mathrm{SCF}_{p,\infty}$. See Theorem 13 of [61] and the paragraph after its proof.

Fact 1.4.8. *Suppose $\mathbb{U} \models \mathrm{SCF}_{p,\infty}$ is a monster model and that $F \leq K, L \leq \mathbb{U}$ are all separable extensions. Then K and L are nonforking independent over F if and only if they are algebraically independent and p-disjoint over F.*

1.5 Differential fields and difference fields

Differential fields and difference fields are the main motivating examples for this thesis's work on fields with operators.

Differential fields

Definition 1.5.1. $(K, \delta_1, \ldots, \delta_m)$ is a differential field if K is a field and each δ_i is a *derivation*:

$$\begin{aligned} \delta_i(a+b) &= \delta_i(a) + \delta_i(b) \\ \delta_i(ab) &= a\delta_i(b) + \delta_i(a)b \end{aligned}$$

It is an ordinary differential field if $m = 1$. We can axiomatise the theory of differential fields with m derivations in the language

$$\mathcal{L}_{\mathrm{ring}}(\delta) = \{+, -, \cdot, 0, 1, \delta_1, \ldots, \delta_m\}$$

where each δ_i is a unary function symbol.

The theory of differential fields of characteristic zero with m commuting derivations has a model completion, $\mathrm{DCF}_{0,m}$. This theory has quantifier elimination, elimination of imaginaries and is ω-stable [45].

The theory $\mathrm{DCF}_{0,1}$ can be axiomatised in a geometric fashion; see [54]. Our axioms for the uniform companion will also be geometric, so it is useful to see the simplest case of ordinary differential fields. But first, we need a preliminary notion.

Definition 1.5.2. Let (K, δ) be an ordinary differential field. Let $V \subseteq \mathbb{A}^n$ be an affine K-variety. The δ-prolongation of V is the affine K-variety defined as

$$\tau_\delta V = \left\{ (x, y) \in \mathbb{A}^{2n} \colon x \in V \text{ and } f^\delta(x) + \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(x) y_i = 0 \text{ for each } f \in I_K(V) \right\}.$$

Here $f^\delta(x)$ is the polynomial obtained by applying δ to the coefficients of f. There is a natural algebraic morphism $\tau_\delta V \to V$ given by projecting onto the first n coordinates.

A differential field $(K, \delta) \models \mathrm{DCF}_{0,1}$ if and only if:

1. K is algebraically closed; and
2. for every affine K-irreducible variety V, and every affine K-irreducible $W \subseteq \tau_\delta V$ such that $W \to V$ is dominant, W has a Zariski dense set of K-rational points of the form $(a, \delta(a)) \in W(K)$.

Remark 1.5.3. In [37], León Sánchez gives a geometric characterisation of $\mathrm{DCF}_{0,m}$. The commutativity of the derivations means the axiomatisation is more complex than ours will be in Chapter 3.

Difference fields

Definition 1.5.4. $(K, \sigma_1, \ldots, \sigma_m)$ is a difference field if K is a field and each σ_i is a field endomorphism. The endomorphisms do not need to commute.

The theory of difference fields has a model companion, $\mathrm{ACFA}_{0,m}$; see [27]. It has elimination of imaginaries and is simple.

$(K, \sigma_1, \ldots, \sigma_m) \models \mathrm{ACFA}_{0,m}$ if and only if

1. K is algebraically closed; and
2. for every affine K-irreducible variety V, and every affine K-irreducible $W \subseteq V \times V^{\sigma_1} \times \cdots \times V^{\sigma_m}$ such that each $W \to V^{\sigma_i}$ is dominant, W has a Zariski dense set of K-rational points of the form $(a, \sigma_1(a), \ldots, \sigma_m(a)) \in W(K)$.

The two geometric axioms described above are very similar; this is part of the reason Moosa and Scanlon developed their theory of prolongations in [51] and their theory of fields with free operators in [53].

1.6 Fields with free operators

Fix a base field k. Let $\mathcal{D}$ be a finite-dimensional k-algebra, and let $\varepsilon_0, \ldots, \varepsilon_l$ be a k-basis of $\mathcal{D}$. We require that there exists a k-algebra homomorphism $\pi\colon \mathcal{D} \to k$ that sends $\varepsilon_0 \mapsto 1$ and $\varepsilon_i \mapsto 0$ for $i = 1, \ldots, l$. If R is a k-algebra, $1 \otimes \varepsilon_0, \ldots, 1 \otimes \varepsilon_l$ is an

R-basis of $R \otimes_k \mathcal{D}$. Write $\pi^R \colon R \otimes_k \mathcal{D} \to R$ for the k-algebra homomorphism $\mathrm{id}_R \otimes \pi$. Recall that rings and algebras are commutative and unital and that homomorphisms preserve the unit.

Definition 1.6.1. Let R be a k-algebra and $\partial_i \colon R \to R$ a sequence of unary functions on R for $i = 1, \ldots, l$. We say that $(R, \partial_1, \ldots, \partial_l)$ is a $\mathcal{D}$-ring if the map $\partial \colon R \to R \otimes_k \mathcal{D}$ given by

$$r \mapsto r \otimes \varepsilon_0 + \partial_1(r) \otimes \varepsilon_1 + \cdots + \partial_l(r) \otimes \varepsilon_l$$

is a k-algebra homomorphism. Equivalently, we will say that (R, ∂) is a $\mathcal{D}$-ring if $\partial \colon R \to R \otimes_k \mathcal{D}$ is a k-algebra homomorphism such that $\pi^R \circ \partial = \mathrm{id}_R$.

If R is a k-algebra and S is an R-algebra given by $a \colon R \to S$, we say that $\partial \colon R \to S \otimes_k \mathcal{D}$ is a $\mathcal{D}$-operator along $a \colon R \to S$ if it is a k-algebra homomorphism and $\pi^S \circ \partial = a$. Then (R, ∂) is a $\mathcal{D}$-ring if and only if ∂ is a $\mathcal{D}$-operator along id_R.

The ring structure of $\mathcal{D}$ determines the additive and multiplicative rules of the functions ∂_i. Indeed, let $a_{ijk}, b_i \in k$ be the elements defined by $\varepsilon_i \varepsilon_j = \sum_{k=0}^{l} a_{ijk} \varepsilon_k$ and $1_{\mathcal{D}} = \sum_{i=0}^{l} b_i \varepsilon_i$. Then k-linearity of ∂ corresponds to k-linearity of each ∂_i. Multiplicativity of ∂ corresponds to the following "product rule" being satisfied for each k: $\partial_k(rs) = \sum_{i,j=0}^{l} a_{ijk} \partial_i(r) \partial_j(s)$. That ∂ preserves the unit corresponds to the equation $\partial_i(1_R) = b_i$.

Note that being a $\mathcal{D}$-ring imposes no additional relations between the functions ∂_i. For example, commutativity of the operators is not imposed by being a $\mathcal{D}$-ring (though a particular $\mathcal{D}$-ring may indeed have $\partial_i \partial_j = \partial_j \partial_i$).

We can axiomatise the theory of $\mathcal{D}$-rings in the language

$$\mathcal{L}_{\text{ring}}(\partial) = \{+, -, \cdot, 0, 1, (c_a)_{a \in k}, \partial_1, \ldots, \partial_l\},$$

where c_a is a constant symbol for the element $a \in k$.

Example 1.6.2.

1. Take $\mathcal{D}$ to be the algebra of dual numbers, $k[\varepsilon]/(\varepsilon^2)$, with the standard k-algebra structure, basis $\{1, \varepsilon\}$, and $\pi \colon \mathcal{D} \to k$ the map that quotients by ε. Then (R, ∂_1) is a $\mathcal{D}$-ring precisely when R is a k-algebra and ∂_1 is a k-linear derivation of R.
2. Let $\mathcal{D} = k[\varepsilon_1, \ldots, \varepsilon_l]/(\varepsilon_1, \ldots, \varepsilon_l)^2$ with basis $\{1, \varepsilon_1, \ldots, \varepsilon_l\}$ and π the map that quotients by $(\varepsilon_1, \ldots, \varepsilon_l)$. Then $(R, \partial_1, \ldots, \partial_l)$ is a $\mathcal{D}$-ring if R is a k-algebra and each ∂_i is a k-linear derivation of R. As explained before, these derivations will in general be noncommuting.

3. Take $\mathcal{D} = k^{l+1}$ with the product k-algebra structure, the standard basis, and π the projection to the first coordinate. Then $(R, \partial_1, \ldots, \partial_l)$ is a $\mathcal{D}$-ring if and only if R is a k-algebra and each ∂_i is a k-linear endomorphism of R. These endomorphisms will in general be noncommuting.

4. We can combine the above examples. Let $\mathcal{D} = k[\varepsilon]/(\varepsilon^2) \times k$ with basis $\{(1,0), (\varepsilon, 0), (0,1)\}$ and π the map which projects to the first coordinate and then quotients by ε. Then a $\mathcal{D}$-ring $(R, \partial_1, \partial_2)$ is a k-algebra with a derivation ∂_1 and an endomorphism ∂_2.

5. Let $\mathcal{D} = k[\varepsilon]/(\varepsilon^{l+1})$ with basis $\{1, \varepsilon, \ldots, \varepsilon^l\}$ and π the map that quotients by ε. Then $\mathcal{D}$-rings are k-algebras with non-iterative, truncated higher derivations $(\partial_1, \ldots, \partial_l)$. That is, they satisfy the following higher-order Leibniz rule: $\partial_i(xy) = \sum_{r+s=i} \partial_r(x)\partial_s(y)$.

The reader is referred to [53] for more examples.

Since $\mathcal{D}$ is a finite-dimensional k-algebra, it is artinian and can be written as a finite product of local finite-dimensional k-algebras $\mathcal{D} = \prod_{i=0}^{t} B_i$. For each i let $\mathfrak{m}_i$ be the unique maximal ideal of B_i. Then the residue field is a finite field extension of k: $B_i/\mathfrak{m}_i = k[x]/(P_i)$ for some k-irreducible polynomial P_i. We define the k-algebra homomorphism $\pi_i \colon \mathcal{D} \to k[x]/(P_i)$ by the composition $\mathcal{D} \to B_i \to k[x]/(P_i)$, and we let $\pi_i^R = \mathrm{id}_R \otimes \pi_i$ be the k-algebra homomorphism $R \otimes_k \mathcal{D} \to R[x]/(P_i)$ for any k-algebra R. Note that the k-algebra homomorphism $\pi \colon \mathcal{D} \to k$ gives a maximal ideal of $\mathcal{D}$ with residue field k. So π must correspond to one of the π_i. By renaming if necessary, say π corresponds to π_0, and hence B_0 has residue field k.

Definition 1.6.3. Suppose $\partial \colon R \to S \otimes_k \mathcal{D}$ is a $\mathcal{D}$-operator along $a \colon R \to S$. Composing ∂ and the map π_i^S gives the following k-algebra homomorphism:

$$R \xrightarrow{\partial} S \otimes_k \mathcal{D} \xrightarrow{\pi_i^S} S[x]/(P_i).$$

This is called the ith associated homomorphism, σ_i, of ∂.

Now, $\sigma_0 = \pi_0^S \circ \partial = \pi^S \circ \partial = a$ and the associated homomorphism corresponding to B_0 is always a.

Suppose now that (R, ∂) is a $\mathcal{D}$-ring. If $\alpha \in R$ is a root of P_i, we have a map $R[x]/(P_i) \to R$. The composition of σ_i with this map gives an endomorphism of R, $\sigma_{i,\alpha} \colon R \to R$. This endomorphism is uniformly quantifier-free α-definable in $\mathcal{L}_{\text{ring}}(\partial)$.

In [53], the authors impose the following condition on the k-algebra $\mathcal{D}$.

Assumption A. For each $i = 0, 1, \ldots, t$, the field $B_i/\mathfrak{m}_i$, which is necessarily a finite extension of k, is k itself.

As a consequence of this assumption, all the associated homomorphisms of a $\mathcal{D}$-ring (R, ∂) are now endomorphisms $R \to R$. In this thesis, we will often impose this assumption, or the stronger assumption that $\mathcal{D}$ is a local ring. In the latter case, there are no nontrivial associated endomorphisms.

For the construction of the uniform companion, we will need to understand how to extend $\mathcal{D}$-structures. We will follow the proof of Lemma 2.7 of [4], which is based on the notions of being 0-smooth and 0-étale; see Section 25 of [44].

Definition 1.6.4. Let $a\colon R \to S$ be an R-algebra. We say that S is 0-smooth over R if it has the following property: for any R-algebra C, any nilpotent ideal N of C, and any R-algebra homomorphism $u\colon S \to C/N$, there is some lifting of u to an R-algebra homomorphism $v\colon S \to C$. That is, given a diagram of ring homomorphisms

$$\begin{array}{ccc} S & \xrightarrow{u} & C/N \\ {\scriptstyle a}\uparrow & & \uparrow \\ R & \longrightarrow & C \end{array}$$

there is some v such that

$$\begin{array}{ccc} S & \xrightarrow{u} & C/N \\ {\scriptstyle a}\uparrow & \overset{v}{\searrow} & \uparrow \\ R & \longrightarrow & C \end{array}$$

S is 0-unramified over R if there is at most one such v, and it is 0-étale if there is exactly one such v.

Lemma 1.6.5. *Suppose R is a k-algebra, S is an R-algebra given by $a\colon R \to S$, T is an S-algebra given by $b\colon S \to T$ and $\partial\colon R \to T \otimes_k \mathcal{D}$ is a $\mathcal{D}$-operator along ba (that is, $\pi^T \circ \partial = ba$). Let $\sigma_i\colon R \to T[x]/(P_i)$ be the associated homomorphisms of ∂, $\pi_i^T \circ \partial$. Let $\tau_i\colon S \to T[x]/(P_i)$ be k-algebra homomorphisms extending σ_i. If S is 0-smooth over R, there is an extension of ∂ to a $\mathcal{D}$-operator $\partial'\colon S \to T \otimes_k \mathcal{D}$ along b with associated homomorphisms τ_i. If S is 0-étale over R, there is a unique such extension.*

Proof. We have a diagram

$$\begin{array}{ccc} S & \xrightarrow{\tau_i} & T[x]/(P_i) \\ {\scriptstyle a}\big\uparrow & & \big\uparrow \\ R & \longrightarrow & T \otimes_k B_i \end{array}$$

where the composition anticlockwise is σ_i. Note that the vertical map on the right is surjective with nilpotent kernel $T \otimes_k \mathfrak{m}_i$. Since $R \to S$ is 0-smooth (0-étale), there is a (unique) homomorphism $S \to T \otimes_k B_i$ fitting into this diagram. Let ∂' be the product of these. This gives a (unique) map $S \to T \otimes_k \mathcal{D}$ whose associated homomorphisms are τ_i. The commutativity of the lower triangle implies that ∂' extends ∂. ■

Remark 1.6.6.
1. Separable extensions are 0-smooth (Theorem 26.9). Separable algebraic extensions are 0-étale (Theorem 25.3). If a field extension shares a p-basis, then it is 0-étale (Theorem 26.7). Localisations are 0-étale. All references are to [44].
2. When $\mathcal{D}$ is local, this lemma appears as Lemma 2.7 in [4].
3. In Chapter 3, we will only need this result in the case $T = S$ and $b = \mathrm{id}_S$. The extra generality will be necessary in Chapter 4.

1.7 The prolongation of an affine variety

In Section 1.5, we saw that the axiomatisations of $\mathrm{DCF}_{0,1}$ and $\mathrm{ACFA}_{0,t}$ required the geometric objects $\tau_\delta V$ and $V \times V^{\sigma_1} \times \cdots \times V^{\sigma_t}$. These objects are sometimes called prolongations. In [51], the authors develop a generalisation of these objects to the case of $\mathcal{D}$-fields; they then apply them in their model-theoretic analysis in [53]. These prolongations will also play a key role in the uniform companion developed in this thesis.

Let (K, ∂) be a $\mathcal{D}$-field, and V an affine K-variety. Since $\partial\colon K \to K \otimes_k \mathcal{D}$ is a k-algebra homomorphism, it is also a K-algebra homomorphism considered as a map $\partial\colon K \to \mathcal{D}^\partial(K)$, where $\mathcal{D}^\partial(K)$ is the K-algebra whose underlying ring is $K \otimes_k \mathcal{D}$ and whose K-algebra structure map is ∂. Thus we may consider the base change of V from K to $\mathcal{D}^\partial(K)$; denote this $V \times_K \mathcal{D}^\partial(K)$. This object is a scheme over $\mathcal{D}(K) = K \otimes_k \mathcal{D}$ given by projecting onto the second coordinate. Since $\mathcal{D}(K)$ is a finite-dimensional K-algebra, we can take the Weil restriction of this scheme from

$\mathcal{D}(K)$ to K.

$$\tau V := W_{\mathcal{D}(K)/K}\left(V \times_K \mathcal{D}^{\partial}(K)\right).$$

The functor $V \mapsto \tau V$ can also be seen as the right adjoint to the functor which is defined (algebraically) as:

$$\begin{aligned} \mathsf{Alg}_K &\to \mathsf{Alg}_K \\ R &\mapsto \mathcal{D}^{\partial}(R). \end{aligned}$$

Here $\mathcal{D}^{\partial}(R)$ is the ring $R \otimes_k \mathcal{D}$ with K-algebra structure given by $K \xrightarrow{\partial} K \otimes_k \mathcal{D} \to R \otimes_k \mathcal{D}$. See the discussion after Remark 2.10 of [4].

We can construct τV more explicitly. For a polynomial $g \in K[X]$, where $X = (X_1, \dots, X_n)$, g^{∂} means the polynomial in $\mathcal{D}(K)[X]$ given by applying ∂ to the coefficients of g. Now compute the polynomials $g^{(0)}, \dots, g^{(l)} \in K[X^{(0)}, \dots, X^{(l)}]$ which make the following true in $\mathcal{D}(K)[X^{(0)}, \dots, X^{(l)}]$.

$$g^{\partial}\left(\sum_{i=0}^{l} X^{(i)}\varepsilon_i\right) = \sum_{i=0}^{l} g^{(i)}(X^{(0)}, \dots, X^{(l)})\varepsilon_i.$$

If $V = \operatorname{Spec}(K[y]/I)$, then $\tau V = \operatorname{Spec}(K[y^{(0)}, \dots, y^{(l)}]/I')$ where I' is the ideal generated by the $g^{(0)}, \dots, g^{(l)}$ as g ranges over I.

Suppose Assumption **A** holds. Then each associated endomorphism $\sigma_i \colon K \to K$ for $i = 0, \dots, t$ induces an algebraic morphism $\hat{\pi}_i \colon \tau V \to V^{\sigma_i}$; see Section 4.1 of [51].

We summarise the crucial facts in the following.

Fact 1.7.1. *Suppose (K, ∂) is a $\mathcal{D}$-field, V is a scheme over K, and τV its prolongation. Then*

1. *if V is affine, so is τV;*
2. *if V is of finite type, so is τV;*
3. *if $L \geq K$ is a field extension, there is an identification $\tau V(L) \leftrightarrow V(\mathcal{D}^{\partial}(L))$; and*
4. *if $(L, \delta) \geq (K, \partial)$ is a $\mathcal{D}$-field extension, there is a (nonalgebraic) map*

$$\nabla \colon V(L) \to \tau V(L);$$

with respect to the above coordinates, it is given as

$$\nabla(a) = (a, \delta_1(a), \ldots, \delta_l(a)).$$

Proof. 1 and 2 are clear by the construction above. 3 is by Lemma 4.5 of [51]. See the discussion after Proposition 4.6 of [51] for 4. ■

1.8 The theory $\mathcal{D}$-CF$_0$

Suppose Assumption **A** holds. By axiomatising the existentially closed models, Moosa and Scanlon prove that the theory of $\mathcal{D}$-fields of characteristic zero has a model companion, $\mathcal{D}$-CF$_0$.

Theorem 4.6 of [53]. $(K, \partial) \models \mathcal{D}\text{-CF}_0$ *if and only if*

- $K \models \mathrm{ACF}_0$*;*
- *the associated endomorphisms* $\sigma_1, \ldots, \sigma_t \colon K \to K$ *are all automorphisms; and*
- *for any* K*-irreducible varieties* V *and* W *with* $W \subseteq \tau V$ *such that each projection* $\hat{\pi}_i \colon W \to V^{\sigma_i}$ *is dominant, there is some* $a \in V(K)$ *such that* $\nabla a \in W(K)$*.*

They then prove that $\mathcal{D}$-CF$_0$ eliminates imaginaries and that every completion of $\mathcal{D}$-CF$_0$ is simple, where A and B are nonforking independent over C exactly when $\mathrm{acl}(AC)$ is linearly disjoint from $\mathrm{acl}(BC)$ over $\mathrm{acl}(C)$.

Remark 1.8.1. The appendix to [53] gives an axiomatisation of the existentially closed $\mathcal{D}$-fields in the absence of Assumption **A**.

Theorem 4.6 of [53] can be seen as a proof that $\mathrm{ACFA}_{0,t} \cup$ "$\mathcal{D}$-fields" has a model companion under Assumption **A**. In Chapter 3, we will see that a similar result holds when replacing $\mathrm{ACFA}_{0,t}$ by any model complete theory of difference large fields.

Chapter 2

The Weil descent functor in the category of algebras with free operators

In this chapter we explore the existence of the Weil descent functor in the appropriate categories of $\mathcal{D}$-rings. In Sections 2.1 and 2.2, we establish the objects and morphisms in the category of $\mathcal{D}$-algebras, as well as define the appropriate notion of $\mathcal{D}$-base change. Following on from the example in the Introduction that shows the difference base change functor cannot always have a left adjoint, in Section 2.3 we associate to every finite and free extension of $\mathcal{D}$-rings a matrix whose invertibility corresponds to the invertibility of a certain natural transformation. Section 2.4 contains the main theorem of this chapter: we construct the $\mathcal{D}$-Weil restriction using the classical Weil restriction together with the functorial nature of $\mathcal{D}$-ring structures. In Section 2.5, we show that several properties of a $\mathcal{D}$-ring such as commutativity of its individual operators and triviality of its associated endomorphisms are preserved under the $\mathcal{D}$-Weil descent as well as a partial converse to the main theorem. Section 2.6 details an explicit construction of the $\mathcal{D}$-Weil descent and explicates any algebraic notions necessary for it. The content of this chapter appears in the author's published paper [50].

2.1 Some more $\mathcal{D}$-algebra

Let k be a field of arbitrary characteristic, and let $\mathcal{D}$ be a finite-dimensional k-algebra. Recall Assumption **A**: Since $\mathcal{D}$ is a finite-dimensional k-algebra, we may decompose it as a finite product of local finite-dimensional k-algebras, say $\mathcal{D} = B_1 \times \cdots \times B_t$. We assume that the residue field of each B_i is actually k.

For any k-algebra R, we define $\mathcal{D}(R) = R \otimes_k \mathcal{D}$ to be the base change of $\mathcal{D}$ to R. Note that $\mathcal{D}(R)$ remains free and finite as an R-module. We will often identify a k-basis of $\mathcal{D}$ with the corresponding R-basis of $\mathcal{D}(R)$. By a slight abuse of notation, we think of $\mathcal{D}$ also as a functor $\mathsf{Alg}_k \to \mathsf{Alg}_k$, where for a k-algebra homomorphism $\phi\colon R \to S$, $\mathcal{D}(\phi) = \phi \otimes \mathrm{id}_{\mathcal{D}}$.

For this chapter only, a $\mathcal{D}$-ring is a k-algebra R equipped with a k-algebra homomorphism $e\colon R \to \mathcal{D}(R)$. That is, we **do not** require that e is a section to the k-algebra homomorphism $\pi^R\colon \mathcal{D}(R) \to R$. Relaxing this definition changes the behaviour of some of our examples. In Section 2.5, we will see that our results also work with the original definition of a $\mathcal{D}$-ring.

Example 2.1.1. 1. Take $\mathcal{D}$ to be the algebra of dual numbers, $k[\varepsilon]/(\varepsilon^2)$, with the standard k-algebra structure. If (R, e) is a $\mathcal{D}$-ring, let σ and δ be such that $e(r) = \sigma(r) + \delta(r)\varepsilon$. Then σ is a k-linear endomorphism of R, and δ is a k-linear derivation on R which is twisted by σ. Indeed, the k-linearity of e implies k-linearity of σ and δ, and multiplicativity implies that

$$\sigma(rs) + \delta(rs)\varepsilon = \sigma(r)\sigma(s) + (\sigma(r)\delta(s) + \delta(r)\sigma(s))\varepsilon$$

Note that if a $\mathcal{D}$-ring has $\sigma = \mathrm{id}_R$ (that is, it satisfies our original definition of a $\mathcal{D}$-ring), then it is a differential k-algebra.

2. Take $\mathcal{D} = k^l$ with the product k-algebra structure. If (R, e) is a $\mathcal{D}$-ring, let $e(r) = \sum_i \sigma_i(r)\varepsilon_i$ where ε_i is the standard basis of $\mathcal{D}$. Then $\mathcal{D}$-rings are precisely rings with l (not necessarily commuting) k-linear endomorphisms $\sigma_1, \ldots, \sigma_l$.

3. We can combine the above two examples. Let $\mathcal{D} = k[\varepsilon]/(\varepsilon^2) \times k$. Then $\mathcal{D}$-rings can be viewed as rings with two endomorphisms σ_1 and σ_2, and a derivation δ twisted by the first endomorphism σ_1. A $\mathcal{D}$-ring with $\sigma_1 = \mathrm{id}$ is then a ring with an endomorphism and a derivation which do not necessarily commute.

4. Let $\mathcal{D} = k[\varepsilon]/(\varepsilon^l)$. Coordinatising the $\mathcal{D}$-structure of a $\mathcal{D}$-ring (R, e) as

$e(r) = \sigma(r) + \delta_1(r)\varepsilon + \ldots + \delta_{l-1}(r)\varepsilon^{l-1}$, we see that $\mathcal{D}$-rings are rings with non-iterative, truncated Hasse-Schmidt derivations $(\delta_1, \ldots, \delta_{l-1})$ twisted by the endomorphism σ:

$$\delta_k(xy) = \delta_k(x)\sigma(y) + \sigma(x)\delta_k(y) + \sum_{\substack{i,j>0 \\ i+j=k}} \delta_i(x)\delta_j(y).$$

We now specify the morphisms of the categories we are working in. These were defined in Section 3.1 of [52].

Definition 2.1.2. If (R, e) and (S, f) are two $\mathcal{D}$-rings, then $\phi\colon (R, e) \to (S, f)$ is a $\mathcal{D}$-homomorphism if it is a k-algebra homomorphism and the following diagram commutes:

$$\begin{array}{ccc} \mathcal{D}(R) & \xrightarrow{\mathcal{D}(\phi)} & \mathcal{D}(S) \\ {\scriptstyle e}\uparrow & & \uparrow{\scriptstyle s} \\ R & \xrightarrow{\quad\phi\quad} & S \end{array}$$

If S is an R-algebra, then we will call (S, f) an (R, e)-algebra if the structure map $R \to S$ is a $\mathcal{D}$-homomorphism. If (S, f) and (T, g) are both (R, e)-algebras and $\phi\colon S \to T$ is a map between them, then we say that ϕ is a (R, e)-algebra homomorphism if it is an R-algebra homomorphism and a $\mathcal{D}$-homomorphism.

Remark 2.1.3. Note that in the context of Example 2.1.1(1) above, under our original definition of a $\mathcal{D}$-ring where σ is the identity map, a map being a $\mathcal{D}$-homomorphism is equivalent to it being a differential ring homomorphism. In the context of Example 2.1.1(2), being a $\mathcal{D}$-homomorphism is equivalent to being a difference ring homomorphism for each endomorphism.

From now on the category of (R, e)-algebras with (R, e)-algebra homomorphisms is denoted by $\mathsf{Alg}_{(R,e)}$.

2.2 The tensor product of $\mathcal{D}$-structures

We now need the correct notion of base change in the context of $\mathcal{D}$-algebras. That is, given a $\mathcal{D}$-ring (R, e) and an (R, e)-algebra (T, g), for any (R, e)-algebra (S, f) we need a $\mathcal{D}$-ring structure on $S \otimes_R T$ that makes $S \otimes_R T$ into a (T, g)-algebra. In [4], it is proved that there exists a unique $\mathcal{D}$-structure, $f \otimes g$ (called $(\tilde{f}, \tilde{g})$ in [4]), on

$S \otimes_R T$ which makes the natural maps $\phi_S \colon S \to S \otimes_R T$ and $\phi_T \colon T \to S \otimes_R T$ into $\mathcal{D}$-homomorphisms. We recall the definition of this structure:

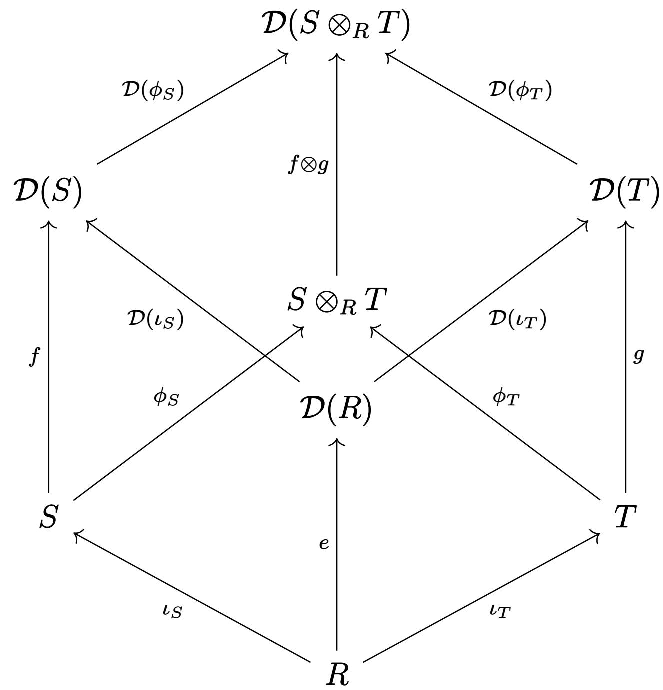

Explicitly,

$$(f \otimes g)(s \otimes t) = (\mathcal{D}(\phi_S) \circ f(s)) \cdot (\mathcal{D}(\phi_T) \circ g(t))$$

where $\cdot$ is the product in $\mathcal{D}(S \otimes_R T)$. The existence and uniqueness of this map $f \otimes g$ follow from the fact that $S \otimes_R T$ is the pushout in the category of k-algebras.

Remark 2.2.1. A short computation shows that this agrees with the correct notions of derivations on tensor products: $(\delta \otimes d)(s \otimes t) = \delta(s) \otimes t + s \otimes d(t)$ (see page 21 of [7]), and endomorphisms on tensor products: $(\sigma \otimes \tau)(s \otimes t) = \sigma(s) \otimes \tau(t)$.

Definition 2.2.2. Let (R, e) be a $\mathcal{D}$-ring and let (T, g) be an (R, e)-algebra. The $\mathcal{D}$-base change functor from (R, e) to (T, g) is defined as follows.

$$\begin{aligned} F^{\mathcal{D}} \colon \mathsf{Alg}_{(R,e)} &\to \mathsf{Alg}_{(T,g)} \\ (S, f) &\mapsto (S \otimes_R T, f \otimes g) \\ (S \xrightarrow{\theta} U) &\mapsto (S \otimes_R T \xrightarrow{\theta \otimes \mathrm{id}_T} U \otimes_R T) \end{aligned}$$

Proposition 2.2.3. *The map $F^{\mathcal{D}}$ is a functor.*

Proof. If $\theta \colon (S, f) \to (U, h)$ is an (R, e)-algebra homomorphism, then $\theta \otimes \mathrm{id}_T$ is a

T-algebra homomorphism. It remains to show it is also a $\mathcal{D}$-homomorphism, that is, that the following square commutes.

$$\begin{array}{ccc} \mathcal{D}(S \otimes_R T) & \xrightarrow{\mathcal{D}(\theta \otimes \mathrm{id}_T)} & \mathcal{D}(U \otimes_R T) \\ {\scriptstyle f\otimes g}\uparrow & & \uparrow{\scriptstyle h\otimes g} \\ S \otimes_R T & \xrightarrow{\theta \otimes \mathrm{id}_T} & U \otimes_R T \end{array}$$

Now consider the following diagram of k-algebra homomorphisms.

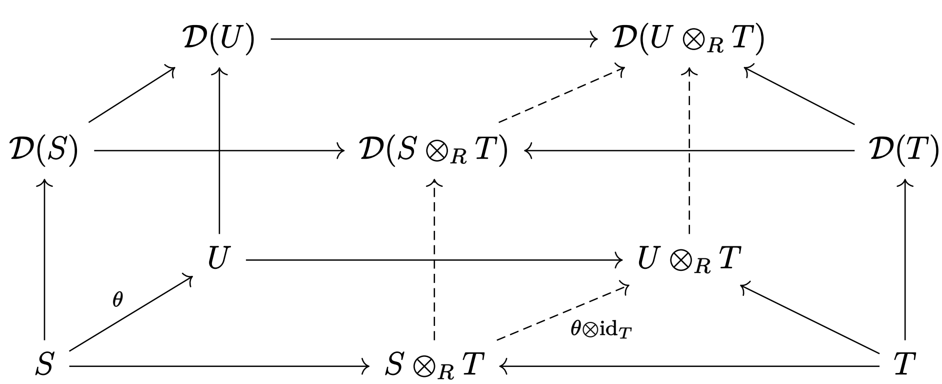

We should also include the R-algebra structure maps and squares expressing that they are $\mathcal{D}$-homomorphisms, but they have been omitted to declutter.

Every square except the dashed one commutes since θ is an (R, e)-algebra homomorphism or by the result above. Now consider a path $S \to \mathcal{D}(U \otimes_R T)$ and a path $T \to \mathcal{D}(U \otimes_R T)$ both avoiding the dashed square. These two paths agree on R, and hence there is a unique map $S \otimes_R T \to \mathcal{D}(U \otimes_R T)$ through which they factor. But these paths also factor through both directions along the dashed square. By uniqueness, both directions must be equal, and $\theta \otimes \mathrm{id}_T$ is a $\mathcal{D}$-homomorphism. ∎

We finish this section with the following lemma which will be used in Section 2.5. It describes the associated endomorphisms of the $\mathcal{D}$-structure on a tensor product. Recall that since Assumption **A** is in force, we have k-algebra homomorphisms $\pi_i \colon \mathcal{D} \to B_i \to k$ for each $i = 1, \ldots, t$ given by quotienting $\mathcal{D}$ by each of its finitely many maximal ideals. We can lift these to k-algebra homomorphisms $\pi_i^R := \mathrm{id}_R \otimes_k \pi_i \colon \mathcal{D}(R) \to R$ for any k-algebra R. Then if (R, e) is a $\mathcal{D}$-ring, each composition $\pi_i^R \circ e$ is a k-linear endomorphism of R. These are the associated endomorphisms $\sigma_1, \ldots, \sigma_t$ of the $\mathcal{D}$-ring (R, e).

Lemma 2.2.4. *Let (R, e) be a $\mathcal{D}$-ring and $(S, f), (T, g) \in \mathsf{Alg}_{(R,e)}$. If the ith associated endomorphism of (S, f) is σ_i and that of (T, g) is τ_i, then the ith associated endomorphism of $(S \otimes_R T, f \otimes g)$ is $\sigma_i \otimes \tau_i$.*

Proof. Using the notation above, we have

$$\begin{aligned}\pi_i^{S\otimes_R T} \circ (f \otimes g)(s \otimes t) &= \pi_i^{S\otimes_R T}(\mathcal{D}(\phi_S) \circ f(s)) \cdot \pi_i^{S\otimes_R T}(\mathcal{D}(\phi_T) \circ g(t)) \\ &= (\pi_i^S \circ f(s) \otimes 1) \cdot (1 \otimes \pi_i^T \circ g(t)) \\ &= \sigma_i(s) \otimes \tau_i(t)\end{aligned}$$

■

2.3 The matrix associated to a free and finite $\mathcal{D}$-ring

In this section we establish some technical results that will be needed to construct a left adjoint to $F^{\mathcal{D}}$ in Section 2.4. We carry forward the notation from the previous section. In particular, k is a field, $\mathcal{D}$ is a finite-dimensional k-algebra, and Assumption **A** still holds.

Recall from the example in the introduction that, in general, the difference base change functor had no left adjoint. There, the nonexistence of the left adjoint is due to the fact that the matrix associated to the endomorphism,

$$\begin{bmatrix} \lambda_1(f(1)) & \lambda_1(f(\varepsilon)) \\ \lambda_2(f(1)) & \lambda_2(f(\varepsilon)) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

is not invertible.

We will show in Section 2.4 that if the associated matrix is invertible, then we can construct a left adjoint to $F^{\mathcal{D}}$. The next subsection investigates conditions under which the associated matrix is invertible.

The matrix associated to an endomorphism

As before, let A be a ring and B an A-algebra which is finite and free as an A-module. We fix a ring endomorphism $\sigma\colon B \to B$ with $\sigma(A) \subseteq A$.

Definition 2.3.1. For an A-basis $b = (b_1, \dots, b_r)$ of B, let M_b^σ be the following matrix associated to σ:

$$M_b^\sigma = \begin{bmatrix} \lambda_1(\sigma(b_1)) & \lambda_1(\sigma(b_2)) & \cdots & \lambda_1(\sigma(b_r)) \\ \lambda_2(\sigma(b_1)) & \lambda_2(\sigma(b_2)) & \cdots & \lambda_2(\sigma(b_r)) \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_r(\sigma(b_1)) & \lambda_r(\sigma(b_2)) & \cdots & \lambda_r(\sigma(b_r)) \end{bmatrix}$$

where λ_i is the ith coordinate projection $B \to A$ with respect to the basis b. Note that the maps λ_i are dependent on the basis $b_1, \dots, b_r$ and hence will change if the basis changes.

We will say that σ has invertible matrix with respect to the basis $b = (b_1, \dots, b_r)$ if M_b^σ is invertible in $\mathrm{Mat}_{r\times r}(A)$.

Proposition 2.3.2. *The following are equivalent:*

(i) *σ has invertible matrix with respect to some A-basis of B;*

(ii) *σ has invertible matrix with respect to every A-basis of B*

(iii) *if $b_1, \dots, b_r$ is an A-basis of B, then $\sigma(b_1), \dots, \sigma(b_r)$ is also an A-basis of B; and*

(iv) $\mathrm{span}_A(\sigma(B)) = B$.

Proof. (ii) $\Rightarrow$ (i) and (iii) $\Rightarrow$ (iv) are clear.

For (ii) $\Leftrightarrow$ (iii), note that M_b^σ is just the change of basis matrix between the two tuples $b_1, \dots, b_r$ and $\sigma(b_1), \dots, \sigma(b_r)$.

For (i) $\Rightarrow$ (ii), say σ has invertible matrix with respect to $b_1, \dots, b_r$, and let $\beta_1, \dots, \beta_r$ be some other basis. Let X be the change of basis matrix from the b to the β, that is, $\beta_i = \sum_j x_{ji} b_j$, let $Y = X^{-1}$, and let μ_i be the A-module homomorphism with $\mu_i(\sum_j a_j \beta_j) = a_i$. Then

$$\begin{aligned} \sigma(\beta_i) &= \sum_j \sigma(x_{ji})\sigma(b_j) \\ &= \sum_j \sum_k \sigma(x_{ji})\lambda_k(\sigma(b_j))b_k \\ &= \sum_j \sum_k \sum_n \sigma(x_{ji})\lambda_k(\sigma(b_j))y_{nk}\beta_n \end{aligned}$$

and so $\mu_n(\sigma(\beta_i)) = \sum_j \sum_k \sigma(x_{ji})\lambda_k(\sigma(b_j))y_{nk}$, that is, $M^\sigma_\beta = YM^\sigma_b\sigma(X)$. Now since X is invertible, $\sigma(X)$ is invertible in $\mathrm{Mat}_{r\times r}(A)$. So M^σ_β is invertible.

For (iv) $\Rightarrow$ (iii), assume $b_1,\dots,b_r$ is an A-basis of B. Any $b \in B$ has $b = \sum_i a_i\sigma(\beta_i)$ for some $\beta_i \in B$. Also, $\beta_i = \sum_j \alpha_{ij}b_j$ since the b_i are a basis, and so $b = \sum_i\sum_j a_i\sigma(\alpha_{ij})\sigma(b_j)$. Then $\sigma(b_1),\dots,\sigma(b_r)$ spans B over A. Now write X for the matrix where $\sigma(b_i) = \sum_j x_{ji}b_j$, and Y for the matrix where $b_i = \sum_j y_{ji}\sigma(b_j)$. Then, since $b_1,\dots,b_r$ is a basis, we have that $XY = I$, and so by taking determinants, we see that X and Y are invertible in $\mathrm{Mat}_{r\times r}(A)$. Then $\sigma(b_1),\dots,\sigma(b_r)$ is an A-basis of B. ∎

Definition 2.3.3. As a result of this proposition, having invertible matrix is independent of the choice of A-basis of B. We will say that σ has invertible matrix if any of the above conditions hold.

The following results explain the connection between the endomorphism σ having invertible matrix and being an automorphism.

Lemma 2.3.4. *If $\sigma|_A\colon A \to A$ is an automorphism, then σ is an automorphism on B if and only if σ has invertible matrix.*

Proof. Define B^σ to be the A-algebra with underlying ring B, but A-algebra structure map $a \mapsto \sigma(a)$. Since $\sigma|_A$ is an automorphism, B^σ is a finite and free A-algebra of the same rank as B; in fact, if $b_1,\dots,b_r$ is a basis of B, then it is also a basis of B^σ. Now the map $f\colon B \to B^\sigma$ given by $f(b) = \sigma(b)$ is actually A-linear, with

$$\begin{aligned} f(b_i) &= \sum_j \lambda_j(\sigma(b_i))b_j \\ &= \sum_j \sigma(\sigma|_A^{-1}\lambda_j(\sigma(b_i)))b_j \end{aligned}$$

and so the matrix of the A-linear map f is $\sigma|_A^{-1}(M^\sigma_b)$. Then f is an isomorphism if and only if $\sigma|_A^{-1}(M^\sigma_b)$ is invertible, if and only if σ has invertible matrix. ∎

Lemma 2.3.5. *If σ is an automorphism on B, then $\sigma|_A$ is an automorphism on A.*

Proof. It is enough to show that $\sigma|_A$ is surjective onto A. Note that since σ is surjective onto B, the A-linear span of $\{\sigma(b_1),\dots,\sigma(b_r)\}$ is B, and by a similar argument to the proof of (iv) $\Rightarrow$ (iii) in Proposition 2.3.2, it must be an A-basis. Now, let $a \in A$. Then there is a $b \in B$ such that $a\sigma(b_1) = \sigma(b)$. Writing $b = \sum_{i=1}^r a_ib_i$

for some $a_i \in A$, we get $a\sigma(b_1) = \sum_{i=1}^{r} \sigma(a_i)\sigma(b_i)$. Since $\{\sigma(b_1), \ldots, \sigma(b_r)\}$ is an A-basis, we get that $a = \sigma(a_1)$, and hence $\sigma|_A$ is surjective onto A. ■

As a result, we see that if σ is an automorphism of B, then it has invertible matrix. It turns out the converse is not true, as we point out in the following example.

Example 2.3.6. Let $A = \mathbb{R}(x_1, x_2, \ldots)$, $B = \mathbb{C}(x_1, x_2, \ldots)$, with basis $b_1 = 1, b_2 = i$, $\sigma|_{\mathbb{C}} = \mathrm{id}_{\mathbb{C}}$, and $\sigma(x_i) = x_{i+1}$. Note that A and B are fields and that σ and $\sigma|_A$ are not surjective. However, the associated matrix is

$$M_b^\sigma = \begin{bmatrix} \lambda_1(\sigma(b_1)) & \lambda_1(\sigma(b_2)) \\ \lambda_2(\sigma(b_1)) & \lambda_2(\sigma(b_2)) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

which is invertible.

On the other hand, one can have an injective endomorphism σ that does not have invertible matrix.

Example 2.3.7. Let K be a field, $A = K[x]$ and $B = A[\varepsilon]/(\varepsilon^2)$ with $\sigma(p(x) + q(x)\varepsilon) = p(x) + xq(x)\varepsilon$. Then with respect to the basis $b = \{1, \varepsilon\}$, we have

$$M_b^\sigma = \begin{bmatrix} \lambda_1(\sigma(b_1)) & \lambda_1(\sigma(b_2)) \\ \lambda_2(\sigma(b_1)) & \lambda_2(\sigma(b_2)) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & x \end{bmatrix}$$

which is not invertible in $\mathrm{Mat}_{2\times 2}(K[x])$.

The matrix associated to a $\mathcal{D}$-ring

We now extend the ideas of the previous subsection to the more general case of $\mathcal{D}$-rings. Just as we can associate a matrix to an endormophism of B, we can associate a matrix to a $\mathcal{D}$-ring structure on B which, when invertible, will allow us to construct a left adjoint to $F^{\mathcal{D}}$ in Section 2.4. Here, we analyse this matrix and the conditions on its invertibility.

Let (A, e) be a $\mathcal{D}$-ring and let (B, f) be an (A, e)-algebra, where B is a finite and free A-algebra. For any k-basis $\varepsilon_1, \ldots, \varepsilon_l$ of $\mathcal{D}$ and any A-basis $b_1, \ldots, b_r$ of B, consider the following $rl \times rl$ matrix with entries in A:

$$M = \begin{bmatrix} M_{11} & M_{12} & \cdots & M_{1,l} \\ M_{21} & M_{22} & \cdots & M_{2,l} \\ \vdots & \vdots & \ddots & \vdots \\ M_{l,1} & M_{l,2} & \cdots & M_{l,l} \end{bmatrix},$$

where M_{mj} is the $r \times r$ matrix given by

$$(M_{mj})_{ni} = \sum_{k=1}^{l} a_{jkm} \lambda_n(f_k(b_i)).$$

Recall that $\lambda_n \colon B \to A$ is the coordinate b_n-projection. The elements $a_{jkm} \in k$ are defined by $\varepsilon_j \varepsilon_k = \sum_{m=1}^{l} a_{jkm}\varepsilon_m$, and $f_k \colon B \to B$ is the coordinate of f with respect to ε_k given by $f(r) = \sum_{k=1}^{l} f_k(r)\varepsilon_k$. We call M the *matrix associated to* (B, f).

We will now briefly explain where this matrix comes from and why we need to consider its invertibility. Define the functor $\mathcal{D}^e \colon \mathsf{Alg}_A \to \mathsf{Alg}_A$ by setting $\mathcal{D}^e(R)$ to be the ring $\mathcal{D}(R)$ but with A-algebra structure given by the composition of $e \colon A \to \mathcal{D}(A)$ with the natural map $\mathcal{D}(A) \to \mathcal{D}(R)$; we defined this A-algebra structure in Section 1.7. On morphisms, $\mathcal{D}^e(\alpha) = \mathcal{D}(\alpha)$. We define $\mathcal{D}^f \colon \mathsf{Alg}_B \to \mathsf{Alg}_B$ similarly. Suppose $u \colon R \to \mathcal{D}(R)$ is a $\mathcal{D}$-ring structure on the A-algebra R. Then (R, u) is an (A, e)-algebra if and only if u is an A-algebra homomorphism $R \to \mathcal{D}^e(R)$.

We now define a natural transformation $\mu \colon F\mathcal{D}^e \to \mathcal{D}^f F$ in the following way: for any A-algebra R, we have a natural A-algebra homomorphism $\mathcal{D}^e(R) \to \mathcal{D}^f(R \otimes_A B)$ and an A-algebra homomorphism $B \to \mathcal{D}^f(R \otimes_A B)$ coming from the composition of f with the natural map. Since $\mathcal{D}^e(R) \otimes_A B$ is the coproduct of A-algebras, we get an A-algebra homomorphism $\mu_R \colon \mathcal{D}^e(R) \otimes_A B \to \mathcal{D}^f(R \otimes_A B)$, which is also a B-algebra homomorphism. It is clear from its construction that μ is natural in R.

Lemma 2.3.8. *The component of μ at R, $\mu_R \colon \mathcal{D}^e(R) \otimes_A B \to \mathcal{D}^f(R \otimes_A B)$, is an R-linear map of free R-modules with the natural R-module structure. With respect to the R-bases $\{\varepsilon_n \otimes b_m\}$ of $\mathcal{D}^e(R) \otimes_A B$ and $\{1 \otimes b_n \varepsilon_m\}$ of $\mathcal{D}^f(R \otimes_A B)$, the matrix representation of μ_R is M, the matrix associated to (B, f). In particular, μ is a natural isomorphism if and only if M is invertible.*

Proof. That μ_R is R-linear is clear from construction. The explicit formula for μ_R is given by

$$\begin{aligned}\sum_{i=1}^{r}\left(\sum_{j=1}^{l} r_{ij}\varepsilon_j\right)\otimes b_i \mapsto \sum_{i=1}^{r}\left(\sum_{j=1}^{l} r_{ij}\otimes 1\ \varepsilon_j\right)\cdot\left(\sum_{k=1}^{l} 1\otimes f_k(b_i)\varepsilon_k\right) \\ = \sum_{i=1}^{r}\sum_{j=1}^{l}\sum_{k=1}^{l}\sum_{m=1}^{l} a_{jkm}r_{ij}\otimes f_k(b_i)\ \varepsilon_m \\ = \sum_{i=1}^{r}\sum_{j=1}^{l}\sum_{k=1}^{l}\sum_{m=1}^{l}\sum_{n=1}^{r} a_{jkm}\lambda_n(f_k(b_i))r_{ij}\otimes b_n\ \varepsilon_m\end{aligned}$$

which immediately shows that M is the matrix of μ_R with respect to the aforementioned bases. ■

From the lemma, we see that if M is invertible, we have a natural transformation $W\mathcal{D}^f \to W\mathcal{D}^f FW \to WF\mathcal{D}^e W \to \mathcal{D}^e W$ coming from the composition of μ^{-1} with the unit and counit of the classical Weil restriction adjunction $W \dashv F$. If $g\colon C \to \mathcal{D}^f(C)$ is a B-algebra homomorphism, then composing the above natural transformation with the morphism $W(g)\colon W(C) \to W\mathcal{D}^f(C)$ gives an A-algebra homomorphism $g^W : W(C) \to \mathcal{D}^e W(C)$. In the next section, we will see that this $\mathcal{D}$-structure on $W(C)$ yields the left adjoint of $F^{\mathcal{D}}$. For now, we study the invertibility of M.

Note that M depends on the choice of the k-basis of $\mathcal{D}$ and the A-basis of B. The following result shows us that invertibility of M is actually independent of the k-basis of $\mathcal{D}$. After the proof of Theorem 2.3.10, we will see that invertibility of M is also independent of the A-basis of B.

Proposition 2.3.9. *Suppose we have two bases $\varepsilon = \{\varepsilon_1, \dots, \varepsilon_l\}$ and $\omega = \{\omega_1, \dots, \omega_l\}$ of $\mathcal{D}$, with X the change of basis matrix from the ε to the ω; that is, $\omega_i = \sum_{j=1}^{l} x_{ji}\varepsilon_j$. Let $\tilde{X}$ be the $rl\times rl$ matrix obtained from X by replacing each entry x by the $r\times r$ block xI, where I is the $r \times r$ identity matrix. Write M^ε for the matrix M corresponding to the basis ε and similarly for M^ω. Then*

$$M^\omega = \tilde{X}^{-1}M^\varepsilon\tilde{X}.$$

Proof. Let a_{ijk} be the product coefficients of the ε and α_{ijk} for the ω. Also, write f_i^ε for the ith operator with respect to the basis ε and similarly for f_i^ω. We can obtain a relation between these by noting that the homomorphism $f\colon B \to \mathcal{D}(B)$ they induce must be the same; that is

$$\sum_{i=1}^{l} f_i^{\varepsilon}(b)\ \varepsilon_i = \sum_{i=1}^{l} f_i^{\omega}(b)\ \omega_i \text{ for all } b \in B.$$

To ease notation, let $\tilde{Y} = \tilde{X}^{-1}$. Let $N = \tilde{Y} M^{\varepsilon} \tilde{X}$. Then the mj block of N is

$$\begin{aligned} N_{mj} &= \sum_p \sum_q \tilde{Y}_{mp} M^{\varepsilon}_{pq} \tilde{X}_{qj} \\ &= \sum_p \sum_q y_{mp} M^{\varepsilon}_{pq} x_{qj}. \end{aligned}$$

Then the ni element of N_{mj} is

$$\begin{aligned} (N_{mj})_{ni} &= \sum_p \sum_q y_{mp} x_{qj} (M^{\varepsilon}_{pq})_{ni} \\ &= \sum_p \sum_q \sum_k y_{mp} x_{qj} a_{qkp} \lambda_n(f_k^{\varepsilon}(b_i)) \\ &= \sum_p \sum_q \sum_k y_{mp} x_{qj} a_{qkp} \lambda_n \left(\sum_r x_{kr} f_r^{\omega}(b_i) \right) \\ &= \sum_p \sum_q \sum_k \sum_r x_{kr} y_{mp} x_{qj} a_{qkp} \lambda_n(f_r^{\omega}(b_i)) \\ &= \sum_r \left(\sum_p \sum_q \sum_k x_{kr} y_{mp} x_{qj} a_{qkp} \right) \lambda_n(f_r^{\omega}(b_i)). \end{aligned}$$

We now claim that $\alpha_{jrm} = \sum_{p,q,k} x_{kr} y_{mp} x_{qj} a_{qkp}$. Indeed, we have

$$\begin{aligned} \omega_j \omega_r &= \left(\sum_q x_{qj} \varepsilon_q \right) \left(\sum_k x_{kr} \varepsilon_k \right) \\ &= \sum_{q,k,p} x_{qj} x_{kr} a_{qkp} \varepsilon_p \\ &= \sum_{q,k,p,u} x_{qj} x_{kr} a_{qkp} y_{up} \omega_u. \end{aligned}$$

Then the claim follows.

Now

$$\begin{aligned}(N_{mj})_{ni} &= \sum_r \alpha_{jrm}\lambda_n(f_r^\omega(b_i)) \\ &= \sum_k \alpha_{jkm}\lambda_n(f_k^\omega(b_i)) \\ &= (M_{mj}^\omega)_{ni},\end{aligned}$$

and hence $M^\omega = \tilde{Y}M^\varepsilon\tilde{X}$. ■

This proposition tells us that invertibility of M is independent of which k-basis of $\mathcal{D}$ we choose. We now construct a particular basis of $\mathcal{D}$ that allows us to characterise invertibility of M in Theorem 2.3.10 below. This basis is constructed as follows. Write $\mathcal{D} = B_1 \times \cdots \times B_t$ where each B_i is a local finite-dimensional k-algebra with residue field k (see Assumption **A**). Let $\mathfrak{m}_i$ be the unique maximal ideal of B_i. Nakayama's Lemma tells us that $\mathfrak{m}_i$ is nilpotent: say d_i is minimal such that $\mathfrak{m}_i^{d_i+1} = 0$. It then follows that for each B_i we can find a k-basis $\mathcal{B}_i = \bigcup_{j=0}^{d_i} \mathcal{B}_i^j$ where $\mathcal{B}_i^j/\mathfrak{m}_i^{j+1}$ is a k-basis of $\mathfrak{m}_i^j/\mathfrak{m}_i^{j+1}$. Note that since the residue field of B_i is k, we may choose $\mathcal{B}_i^0 = \{1\}$. Embed these bases inside $\mathcal{D}$ in the usual way, that is, if $x \in \mathcal{B}_i$, send x to the element of $\mathcal{D}$ with x in the ith position and zeros elsewhere. Then the union of these forms a basis $\mathcal{B}$ of $\mathcal{D}$. Order $\mathcal{B} = \bigcup_{i=1}^t \bigcup_{j=0}^{d_i} \mathcal{B}_i^j$ lexicographically on i and j. The ordering of each $\mathcal{B}_i^j$ does not matter. We will write the elements of $\mathcal{B}$ as $\varepsilon_1, \ldots, \varepsilon_l$ according to this order. Let a_{jkm} be the product coefficients of $\mathcal{B}$; that is, $\varepsilon_j\varepsilon_k = \sum_{m=1}^l a_{jkm}\varepsilon_m$.

By the construction of the basis, we know that $\varepsilon_j\varepsilon_k = 0$ whenever ε_j and ε_k come from different $\mathcal{B}_i$. If they come from the same $\mathcal{B}_i$, then $\varepsilon_j\varepsilon_k$ can be expressed as a linear combination of $\mathcal{B}_i$, and so if ε_m does not come from $\mathcal{B}_i$, it will not appear in this linear combination. So we see that $a_{jkm} = 0$ unless ε_j, ε_k, and ε_m all come from the same $\mathcal{B}_i$.

Furthermore, if $\varepsilon_j \in \mathcal{B}_i^n$ and $\varepsilon_k \in \mathcal{B}_i^p$, then $\varepsilon_j\varepsilon_k \in \operatorname{span}(\bigcup_{q=n+p}^{d_i} \mathcal{B}_i^q)$. Hence, if $\varepsilon_m \in \mathcal{B}_i^q$ for $q < n+p$, $a_{jkm} = 0$. From these facts we can deduce the values of a_{jkm} in specific cases:

1. $m < j$ and $p > 0$: $a_{jkm} = 0$.

   Since $m < j$, $q \leq n$, and hence $q < n + p$. By the above, $a_{jkm} = 0$.

2. $m < j$ and $p = 0$: $a_{jkm} = 0$.

Since $p = 0$, ε_k is the 1 in B_i. Then $\varepsilon_j \varepsilon_k = \varepsilon_j \neq \varepsilon_m$.

3. $m = j$ and $p > 0$: $a_{jkm} = 0$.

   Again, as $m = j$, $q = n$ and so $q < n + p$.

4. $m = j$ and $p = 0$: $a_{jkm} = 1$.

   $\varepsilon_j \varepsilon_k = \varepsilon_j = \varepsilon_m$. So $a_{jkm} = 1$.

Now, recall the definition of the matrix M:

$$M = \begin{bmatrix} M_{11} & M_{12} & \cdots & M_{1,l} \\ M_{21} & M_{22} & \cdots & M_{2,l} \\ \vdots & & \ddots & \\ M_{l,1} & M_{l,2} & \cdots & M_{l,l} \end{bmatrix},$$

where

$$(M_{mj})_{ni} = \sum_{k=1}^{l} a_{jkm} \lambda_n(f_k(b_i)).$$

With respect to the chosen basis, $\mathcal{B} = \{\varepsilon_1, \ldots, \varepsilon_l\}$, we now investigate the shape of each block M_{mj} for $m \leq j$. Consider first the case when $m < j$. As pointed out above, if ε_j and ε_m belong to different $\mathcal{B}_i$, then $a_{jkm} = 0$ for all k. Otherwise, we are in cases (1) or (2) above, and hence $a_{jkm} = 0$ for all k. Hence, the block M_{mj} is 0.

Now consider the case $m = j$, that is, the block M_{jj}. Again, if ε_j and ε_m belong to different B_i, then $a_{jkj} = 0$ for all k. If they belong to the same B_i, then case (3) tells us that $a_{jkj} = 0$ when $p > 0$, and (4) tells us that $a_{jkj} = 1$ when $p = 0$. In conclusion, $(M_{jj})_{ni} = \lambda_n(f_k(b_i))$ where k is such that $\varepsilon_k \in \mathcal{B}_r^0$ and $\varepsilon_j \in \mathcal{B}_r$.

From Definition 1.6.3 we see that the ith projection map π_i is just the map that projects onto the coefficient of ε_k where $\varepsilon_k \in \mathcal{B}_i^0$. Hence, the ith associated endomorphism of (B, f), denoted σ_i, is just f_k. Note that σ_i has this form because of the chosen basis of $\mathcal{D}$.

So in all, M is a block lower triangular matrix whose diagonal $r \times r$ blocks M_{jj}

are of the form

$$M_b^{\sigma_i} = \begin{bmatrix} \lambda_1(\sigma_i(b_1)) & \lambda_1(\sigma_i(b_2)) & \cdots & \lambda_1(\sigma_i(b_r)) \\ \lambda_2(\sigma_i(b_1)) & \lambda_2(\sigma_i(b_2)) & \cdots & \lambda_2(\sigma_i(b_r)) \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_r(\sigma_i(b_1)) & \lambda_r(\sigma_i(b_2)) & \cdots & \lambda_r(\sigma_i(b_r)) \end{bmatrix}$$

where i is determined by $\varepsilon_j \in \mathcal{B}_i$.

Note that $M_b^{\sigma_i}$ is the matrix associated to the endomorphism σ_i as in the previous subsection. Hence, we have proved the following important result:

Theorem 2.3.10. *M is invertible if and only if each associated endomorphism of (B, f) has invertible matrix (in the sense of the previous subsection).*

Remark 2.3.11. Combining this theorem with Propositions 2.3.2 and 2.3.9, we see that invertibility of M is independent of the choice of bases of $\mathcal{D}$ and B.

2.4 Weil descent for $\mathcal{D}$-algebras

In this section we prove the main theorem: Theorem 2.4.5 below. As before, we let (A, e) be a $\mathcal{D}$-ring, (B, f) an (A, e)-algebra where B is a finite and free A-algebra.

The proofs in this section make use of the natural transformation $\mu \colon F\mathcal{D}^e \to \mathcal{D}^f F$ defined in the previous section whose invertibility is equivalent to the invertibility of the matrix M – the matrix associated to (B, f) – by Lemma 2.3.8. Furthermore, recall that in Theorem 2.3.10 we proved that M is invertible if and only if the associated endomorphisms of (B, f) have invertible matrix. For the remainder of this section, in addition to Assumption **A**, we make the following assumption:

Assumption 2.4.1. The associated endomorphisms of (B, f) all have invertible matrix. Equivalently, μ is a natural isomorphism.

The following is part of the content of our main theorem.

Theorem 2.4.2. *The $\mathcal{D}$-base change functor, $F^{\mathcal{D}}$, has a left adjoint $W^{\mathcal{D}}$. More precisely, for a (B, f)-algebra (C, g), there exists a unique $\mathcal{D}$-structure g^W on $W(C)$ that makes the unit of the classical adjunction, η_C, into a $\mathcal{D}$-homomorphism. $W^{\mathcal{D}}$ has the form $W^{\mathcal{D}}(C, g) = (W(C), g^W)$.*

Before proving this result, we fix some notation. Since $W \dashv F$, we have the natural transformations given by the unit, $\eta\colon \mathrm{id}_{\mathsf{Alg}_B} \to FW$, and the counit, $\varepsilon\colon WF \to \mathrm{id}_{\mathsf{Alg}_A}$. We do not need to refer to the k-basis of $\mathcal{D}$ in this section, so we will use ε to denote the counit. We will often identify a functor with the identity natural transformation on that functor. Recall the functors $\mathcal{D}^e\colon \mathsf{Alg}_A \to \mathsf{Alg}_A$ and $\mathcal{D}^f\colon \mathsf{Alg}_B \to \mathsf{Alg}_B$ defined in the previous section where $\mathcal{D}^e(R)$ is the ring $\mathcal{D}(R)$ but with A-algebra structure given by the composition of $e\colon A \to \mathcal{D}(A)$ with the natural map $\mathcal{D}(A) \to \mathcal{D}(R)$, and on morphisms, $\mathcal{D}^e(\alpha) = \mathcal{D}(\alpha)$. Recall also that a $\mathcal{D}$-ring structure on R making it into an (A,e)-algebra is equivalent to an A-algebra homomorphism $R \to \mathcal{D}^e(R)$.

Remark 2.4.3. Suppose (R,u) is an (A,e)-algebra so that $u\colon R \to \mathcal{D}^e(R)$ is an A-algebra homomorphism. Then $\mu_R \circ F(u)\colon F(R) \to \mathcal{D}^f F(R)$ is the $\mathcal{D}$-ring structure on $F(R)$ corresponding to $u \otimes f$ from the $\mathcal{D}$-base change functor in Section 2.2.

We now use the natural isomorphism μ to define a suitable $\mathcal{D}$-ring structure on $W(C)$. For ease of notation we first define the natural transformation

$$\zeta\colon W\mathcal{D}^f \to \mathcal{D}^e W$$

by the composition

$$W\mathcal{D}^f \xrightarrow{W\mathcal{D}^f\eta} W\mathcal{D}^f FW \xrightarrow{W\mu^{-1}W} WF\mathcal{D}^e W \xrightarrow{\varepsilon\mathcal{D}^e W} \mathcal{D}^e W \tag{$\ominus$}$$

Now suppose (C,g) is a (B,f)-algebra, so that g corresponds to the B-algebra homomorphism $g\colon C \to \mathcal{D}^f(C)$. Let $g^W := \zeta_C \circ W(g)\colon W(C) \to \mathcal{D}^e W(C)$. Then $(W(C), g^W)$ is an (A,e)-algebra. We define the functor $W^{\mathcal{D}}$ as

$$W^{\mathcal{D}}(C,g) := (W(C), g^W);$$
$$W^{\mathcal{D}}(\alpha) := W(\alpha).$$

Since both W and ζ are natural, it is clear that if α is a $\mathcal{D}$-homomorphism, then $W(\alpha)$ is a $\mathcal{D}$-homomorphism, so that $W^{\mathcal{D}}$ is actually a functor. We now need to show that $W^{\mathcal{D}}$ is left adjoint to $F^{\mathcal{D}}$ by showing that the natural bijection coming from the classical adjunction

$$\begin{aligned} \mathrm{Hom}_{\mathsf{Alg}_A}(W(C), R) &\to \mathrm{Hom}_{\mathsf{Alg}_B}(C, F(R)) \\ \phi &\mapsto F(\phi) \circ \eta_C \\ \varepsilon_R \circ W(\psi) &\leftarrow\!\shortmid \psi \end{aligned}$$

restricts to a natural bijection

$$\begin{aligned} \mathrm{Hom}_{\mathsf{Alg}_{(A,e)}}(W^{\mathcal{D}}(C,g), (R,u)) &\to \mathrm{Hom}_{\mathsf{Alg}_{(B,f)}}((C,g), F^{\mathcal{D}}(R,u)) \\ \phi &\mapsto F^{\mathcal{D}}(\phi) \circ \eta_C \\ \varepsilon_R \circ W^{\mathcal{D}}(\psi) &\leftarrow\!\shortmid \psi \end{aligned}$$

We will do this by showing that both η_C and ε_R are $\mathcal{D}$-homomorphisms with the appropriate $\mathcal{D}$-structures defined above. Consider the following diagram of natural transformations.

$$\begin{array}{ccccccccc} FW\mathcal{D}^f & \xrightarrow{FW\mathcal{D}^f\eta} & FW\mathcal{D}^f FW & \xrightarrow{FW\mu^{-1}W} & FWF\mathcal{D}^e W & \xrightarrow{F\varepsilon\mathcal{D}^e W} & F\mathcal{D}^e W & \xrightarrow{\mu W} & \mathcal{D}^f FW \\ {\scriptstyle \eta\mathcal{D}^f}\uparrow & & {\scriptstyle \eta\mathcal{D}^f FW}\uparrow & & {\scriptstyle \eta F\mathcal{D}^e W}\uparrow & \diagup\!\!\diagup & & & \\ \mathcal{D}^f & \xrightarrow{\mathcal{D}^f\eta} & \mathcal{D}^f FW & \xrightarrow{\mu^{-1}W} & F\mathcal{D}^e W & & & & \end{array}$$

The squares commute due to naturality of η, and the equality is due to the adjunction axiom: $F\varepsilon \circ \eta F = F$. The composition along the top row is $\mu W \circ F\zeta$. By naturality of η, we get

$$\begin{array}{ccc} FW(C) & \xrightarrow{FW(g)} & FW\mathcal{D}^f(C) \\ {\scriptstyle \eta_C}\uparrow & & \uparrow{\scriptstyle \eta_{\mathcal{D}^f(C)}} \\ C & \xrightarrow{g} & \mathcal{D}^f(C) \end{array}$$

and putting these together we get

$$\begin{array}{ccc} \mathcal{D}^f(C) & \xrightarrow{\mathcal{D}^f(\eta_C)} & \mathcal{D}^f FW(C) \\ {\scriptstyle g}\uparrow & & \uparrow{\scriptstyle \mu_{W(C)} \circ F(g^W)} \\ C & \xrightarrow{\eta_C} & FW(C) \end{array}$$

so that η_C is a $\mathcal{D}$-homomorphism by Remark 2.4.3.

Suppose now that g' is a $\mathcal{D}$-ring structure $W(C) \to \mathcal{D}^e W(C)$ making η_C into a $\mathcal{D}$-homomorphism, so that the following diagram of B-algebras commutes.

$$\begin{array}{ccc} \mathcal{D}^f(C) & \xrightarrow{\mathcal{D}^f(\eta_C)} & \mathcal{D}^f FW(C) \\ g\uparrow & & \uparrow \mu_{W(C)} \circ F(g') \\ C & \xrightarrow{\eta_C} & FW(C) \end{array}$$

Since μ is an isomorphism, this is equivalent to the following diagram of B-algebras commuting:

$$\begin{array}{ccc} \mathcal{D}^f(C) & \xrightarrow{\mu^{-1}_{W(C)} \circ \mathcal{D}^f(\eta_C)} & F\mathcal{D}^e W(C) \\ g\uparrow & & \uparrow F(g') \\ C & \xrightarrow{\eta_C} & FW(C) \end{array} \qquad (\diamond)$$

Consider now the diagram of A-algebras

$$\begin{array}{ccccc} W\mathcal{D}^f(C) & \xrightarrow{W(\mu^{-1}_{W(C)} \circ \mathcal{D}^f(\eta_C))} & WF\mathcal{D}^e W(C) & \xrightarrow{\varepsilon_{\mathcal{D}^e W(C)}} & \mathcal{D}^e W(C) \\ W(g)\uparrow & & \uparrow WF(g') & & g'\uparrow \\ W(C) & \xrightarrow{W(\eta_C)} & WFW(C) & \xrightarrow{\varepsilon_{W(C)}} & W(C) \end{array}$$

Note that the left square commutes by applying W to square $(\diamond)$, and the right square commutes by naturality of ε. By the adjunction axiom $\varepsilon W \circ W\eta = W$, the composition along the bottom is $\mathrm{id}_{W(C)}$, and the composition along the top is ζ_C by definition. So $g^W = g'$, and we have proved the following.

Lemma 2.4.4. *g^W is the unique $\mathcal{D}$-structure on $W(C)$ making $(W(C), g^W)$ into an (A, e)-algebra and the unit, η_C, into a $\mathcal{D}$-homomorphism.*

The adjunction axioms tell us that $F\varepsilon \circ \eta F = F$, so that $\mathcal{D}^f F\varepsilon \circ \mathcal{D}^f \eta F = \mathcal{D}^f F$.

Since μ is natural, the following diagram commutes :

$$\begin{array}{ccc}
F\mathcal{D}^e W F & \xrightarrow{F\mathcal{D}^e\varepsilon} & F\mathcal{D}^e \\
\downarrow{\scriptstyle \mu W F} & & \downarrow{\scriptstyle \mu} \\
\mathcal{D}^f F W F & \xrightarrow{\mathcal{D}^f F\varepsilon} & \mathcal{D}^f F \\
\uparrow{\scriptstyle \mathcal{D}^f\eta F} & \diagup\!\!\diagup & \\
\mathcal{D}^f F & &
\end{array}$$

Now apply W and use naturality of the counit to get

$$\begin{array}{ccc}
\mathcal{D}^e W F & \xrightarrow{\mathcal{D}^e\varepsilon} & \mathcal{D}^e \\
\uparrow{\scriptstyle \varepsilon\mathcal{D}^e W F} & & \uparrow{\scriptstyle \varepsilon\mathcal{D}^e} \\
W F\mathcal{D}^e W F & \xrightarrow{W F\mathcal{D}^e\varepsilon} & W F\mathcal{D}^e \\
\downarrow{\scriptstyle W\mu W F} & & \downarrow{\scriptstyle W\mu} \\
W\mathcal{D}^f F W F & \xrightarrow{W\mathcal{D}^f F\varepsilon} & W\mathcal{D}^f F \\
\uparrow{\scriptstyle W\mathcal{D}^f\eta F} & \diagup\!\!\diagup & \\
W\mathcal{D}^f F & &
\end{array}$$

Note that the composition up the left is precisely ζF. So $\varepsilon\mathcal{D}^e = \mathcal{D}^e\varepsilon \circ \zeta F \circ W\mu$.

Naturality of ε gives

$$\begin{array}{ccccc}
W F(R) & \xrightarrow{W F(u)} & W F\mathcal{D}^e(R) & \xrightarrow{\zeta_{F(R)}\circ W(\mu_R)} & \mathcal{D}^e W F(R) \\
\downarrow{\scriptstyle \varepsilon_R} & & \downarrow{\scriptstyle \varepsilon_{\mathcal{D}^e(R)}} & \swarrow{\scriptstyle \mathcal{D}^e(\varepsilon_R)} & \\
R & \xrightarrow{u} & \mathcal{D}^e(R) & &
\end{array}$$

and since the composition along the top row is $(\mu_R \circ F(u))^W$, the counit ε_R is a $\mathcal{D}$-homomorphism.

We have thus proved the following.

Theorem 2.4.5 (The $\mathcal{D}$-Weil Descent). *Suppose (A, e) is a $\mathcal{D}$-ring and (B, f) is an (A, e)-algebra, where B is a finite and free A-algebra. Suppose also that the associated endomorphisms of (B, f) all have invertible matrix. Then the $\mathcal{D}$-base change functor, $F^{\mathcal{D}}\colon \mathsf{Alg}_{(A,e)} \to \mathsf{Alg}_{(B,f)}$ has a left adjoint denoted $W^{\mathcal{D}}$ called the $\mathcal{D}$-Weil descent. More precisely, $W^{\mathcal{D}}(C, g) = (W(C), g^W)$ where g^W is the $\mathcal{D}$-ring structure defined by $\zeta_C \circ W(g)$ and $\zeta\colon W\mathcal{D}^f \to \mathcal{D}^e W$ is the natural transformation* *defined in equation (⊖).*

In fact, the natural bijection $\tau(C,R)$ from the classical adjunction restricts to a natural bijection:

$$\tau^{\mathcal{D}}((C,g),(R,u))\colon \mathrm{Hom}_{\mathsf{Alg}_{(A,e)}}(W^{\mathcal{D}}(C,g),(R,u)) \to \mathrm{Hom}_{\mathsf{Alg}_{(B,f)}}((C,g),F^{\mathcal{D}}(R,u))$$

Remark 2.4.6. If we apply this theorem to the case when $\mathcal{D}=k$, we get what we call the difference Weil descent and denote it W^{σ}. In this case, $\mathcal{D}$-rings are rings with a single (not necessarily injective) endomorphism.

2.5 Further remarks

In this section we investigate three further aspects. Firstly, we make some observations about properties of the associated endomorphisms that are transferred by the $\mathcal{D}$-Weil descent. In particular, we prove that if the ith associated endomorphism of (C,g) is trivial, then the same is true of the $\mathcal{D}$-Weil descent, $(W(C),g^W)$. Secondly, we prove results about the composition of a $\mathcal{D}_1$-structure and a $\mathcal{D}_2$-structure and their Weil descents. In particular, we will show that commutativity of these structures is preserved after taking the Weil descent. These two subsections imply that the result of this paper is an actual generalisation of the case of several commuting derivations from [39]. Thirdly, we explore the necessity of the condition that the associated endomorphisms of (B,f) have invertible matrix for the existence of the $\mathcal{D}$-Weil descent.

Throughout this section, unless stated otherwise, (A,e) is a $\mathcal{D}$-ring, (B,f) is an (A,e)-algebra, where B is finite and free over A, and (C,g) is a (B,f)-algebra. Assumption **A** is still in force.

Transferred properties of the associated endomorphisms

Recall from Definition 1.6.3 the projection maps for $\mathcal{D}$. If $\mathcal{D}=\prod_{i=1}^{t} B_i$ where each B_i is a local k-algebra with residue field k, then $\pi_i\colon \mathcal{D}\to B_i\to k$ is the composition of the projection onto the ith component of $\mathcal{D}$ with the residue map onto k. These π_i lift to R-algebra homomorphisms $\pi_i^R\colon \mathcal{D}(R)\to R$. Then the associated endomorphisms of a $\mathcal{D}$-ring (R,e) are defined by $\pi_i^R\circ e$ for each $i=1,\dots,t$.

Lemma 2.5.1. *Let (C, g) be a (B, f)-algebra, and suppose that the associated endomorphisms of (B, f) have invertible matrix. Then the associated endomorphisms of the $\mathcal{D}$-Weil descent of (C, g) are the difference Weil descents of the associated endomorphisms of (C, g). In particular, if an associated endomorphism of (C, g) is trivial, then so is the corresponding one of $W^{\mathcal{D}}(C, g)$.*

Proof. Let (σ_i), (τ_i), (υ_i), and (ρ_i) be the associated endomorphisms of (A, e), (B, f), (C, g), and $(W(C), g^W)$, respectively. We need to show that $\rho_i = \upsilon_i^{W^\sigma}$. Consider the following diagrams for each $i = 1, \ldots, t$:

$$\begin{array}{ccc} C & \xrightarrow{\eta_C} & F(W(C)) \\ {\scriptstyle \pi_i^C}\uparrow & & \uparrow{\scriptstyle \pi_i^{F(W(C))}} \\ \mathcal{D}(C) & \xrightarrow{\mathcal{D}(\eta_C)} & \mathcal{D}(F(W(C))) \\ {\scriptstyle g}\uparrow & & \uparrow{\scriptstyle g^W \otimes f} \\ C & \xrightarrow{\eta_C} & F(W(C)) \end{array}$$

The compositions of the vertical maps on the left are υ_i by definition. On the right they are $\rho_i \otimes \tau_i$ by Lemma 2.2.4. Hence ρ_i is a difference structure on $W(C)$ that makes $(W(C), \rho_i)$ into an (A, σ_i)-algebra and η_C into a (B, τ_i)-algebra homomorphism. Since τ_i has invertible matrix, Lemma 2.4.4 tells us that such a difference structure is unique, and so we must have $\rho_i = \upsilon_i^{W^\sigma}$.

For the in particular clause, since the following square commutes

$$\begin{array}{ccc} C & \xrightarrow{\eta_C} & F(W(C)) \\ {\scriptstyle \mathrm{id}_C}\uparrow & & \uparrow{\scriptstyle \mathrm{id}_{W(C)} \otimes \mathrm{id}_B} \\ C & \xrightarrow{\eta_C} & F(W(C)) \end{array}$$

we must have $(\mathrm{id}_C)^W = \mathrm{id}_{W(C)}$ by the uniqueness of the difference structure on $W(C)$ making it an (A, id_A)-algebra and η_C a (B, id_B)-algebra homomorphism. Note here that id_B has invertible matrix. ■

Remark 2.5.2. This lemma tells us that we may apply our $\mathcal{D}$-Weil descent result (Theorem 2.4.5) in the context of [53], this thesis's original definition of a $\mathcal{D}$-ring. Recall that, there and in Section 1.6, a $\mathcal{D}$-ring (R, e) must have e be a section to π_1^R, and hence the first associated endomorphism must be the identity. Thus, if (A, e) and (B, f) have trivial first associated endomorphism, we may consider the category of (A, e)-algebras (R, u) where u has trivial first associated endomorphism, and similarly for (B, f)-algebras. Denote these subcategories $\mathsf{Alg}^*_{(A,e)}$ and $\mathsf{Alg}^*_{(B,f)}$. One checks that $F^{\mathcal{D}}$ can now be considered as a functor $\mathsf{Alg}^*_{(A,e)} \to \mathsf{Alg}^*_{(B,f)}$, and the previous lemma tells us that $W^{\mathcal{D}}$ restricts to a functor $\mathsf{Alg}^*_{(B,f)} \to \mathsf{Alg}^*_{(A,e)}$ which is still left adjoint to $F^{\mathcal{D}}$. In particular, our result is an actual generalisation of the single derivation case from [39], since the category of differential A-algebras is equal to $\mathsf{Alg}^*_{(A,e)}$ when we take $\mathcal{D} = k[\varepsilon]/(\varepsilon^2)$.

We point out here that $W^{\mathcal{D}}$ does not in general preserve injectivity of the associated endomorphisms. That is, if the ith associated endomorphism of g is injective, the ith associated endomorphism of g^W may no longer be injective.

Example 2.5.3. Let $\mathcal{D} = k$ so that the associated endomorphism of a $\mathcal{D}$-ring structure is just the $\mathcal{D}$-ring structure itself. Let $A = \mathbb{F}_2$ be the field with two elements, and let $B = \mathbb{F}_2[\varepsilon]/(\varepsilon^2)$. Let id_A and id_B be the $\mathcal{D}$-ring structures on A and B respectively. Note then that if (C, ρ) is a (B, id_B)-algebra, ρ is a B-algebra endomorphism of C making the following diagram commute:

$$\begin{array}{ccc} C & \xrightarrow{\eta_C} & F(W(C)) \\ \uparrow{\scriptstyle \rho} & & \uparrow{\scriptstyle \rho^W \otimes \mathrm{id}_B} \\ C & \xrightarrow{\eta_C} & F(W(C)) \end{array}$$

Note also that since ρ is a B-algebra endomorphism, it is a morphism in Alg_B, and so we may apply the classical Weil descent to it. Theorem 1.3.1 tells us that $W(\rho) = \rho^W$.

Let $C = B[t]$ and let ρ be the unique map extending id_B on B and sending $t \mapsto t^2$. Then ρ is injective. Recall from Section 1.3 that $W(B[t]) = A[t] \otimes_A A[t]$ and that $\eta_C(t) = t(1) \otimes 1 + t(2) \otimes \varepsilon$. Then

$$\begin{aligned} \eta_C(\rho(t)) &= \eta_C(t^2) \\ &= \eta_C(t)^2 \\ &= t(1)^2 \otimes 1 \end{aligned}$$

where the last equality holds because $\varepsilon^2 = 0$ and we are in characteristic 2.

Also

$$(W(\rho) \otimes \mathrm{id}_B)(\eta_C(t)) = W(\rho)(t(1)) \otimes 1 + W(\rho)(t(2)) \otimes \varepsilon.$$

By the commutativity of the diagram, we have $W(\rho)(t(1)) = t(1)^2$ and $W(\rho)(t(2)) = 0$. Hence $W(\rho)$ is not injective.

Remark 2.5.4.

1. This example tells us that in general the difference Weil descent functor does not restrict to the categories of algebras equipped with an injective endomorphism. However, Corollary 2.5.9 will tell us that the difference Weil descent does preserve automorphisms, and hence will restrict to a functor in the categories of inversive difference algebras (see [40]).
2. The example above uses in an essential way the fact that the characteristic is positive. We are not currently aware of such an example in characteristic zero.

The composition of a $\mathcal{D}_1$-structure and a $\mathcal{D}_2$-structure

Suppose we now have two finite-dimensional k-algebras $\mathcal{D}_1 = \prod_{i=1}^{t_1} B_i$ and $\mathcal{D}_2 = \prod_{j=1}^{t_2} C_j$ where each B_i and each C_j is local with residue field k. Then $\mathcal{D}_2 \otimes_k \mathcal{D}_1 = \prod_{i=1}^{t_1} \prod_{j=1}^{t_2} C_j \otimes_k B_i$. From [62] we know that $C_j \otimes_k B_i$ is local with residue field k, and hence $\mathcal{D}_2 \otimes_k \mathcal{D}_1$ satisfies Assumption **A**. We may then consider the category of $\mathcal{D}_2 \otimes_k \mathcal{D}_1$-rings. We will write these as $\mathcal{D}_1\mathcal{D}_2$-rings since

$$(\mathcal{D}_2 \otimes_k \mathcal{D}_1)(R) = R \otimes_k (\mathcal{D}_2 \otimes_k \mathcal{D}_1) \cong (R \otimes_k \mathcal{D}_2) \otimes_k \mathcal{D}_1 = \mathcal{D}_1(\mathcal{D}_2(R))$$

for a k-algebra R.

If some k-algebra R has a $\mathcal{D}_1$-structure e_1 and a $\mathcal{D}_2$-structure e_2, we can form a $\mathcal{D}_1\mathcal{D}_2$-structure on R by the k-algebra homomorphism

$$\mathcal{D}_1(e_2) \circ e_1 \colon R \to \mathcal{D}_1\mathcal{D}_2(R).$$

We now investigate the Weil descent of this composition of $\mathcal{D}_1$-structures and $\mathcal{D}_2$-structures. Suppose R, S, and T all have a $\mathcal{D}_1$-structure e_1, f_1, g_1 and a $\mathcal{D}_2$-structure e_2, f_2, g_2 that make (S, f_1) and (T, g_1) into (R, e_1)-algebras and (S, f_2) and (T, g_2) into (R, e_2)-algebras. We can then define $\mathcal{D}_1\mathcal{D}_2$-structures on each of them as above.

Lemma 2.5.5. *Assuming the notation of the paragraph above, we have*

$$\mathcal{D}_1(f_2 \otimes g_2) \circ (f_1 \otimes g_1) = (\mathcal{D}_1(f_2) \circ f_1) \otimes (\mathcal{D}_1(g_2) \circ g_1)$$

as $\mathcal{D}_1\mathcal{D}_2$-structures on $S \otimes_R T$.

Proof. Consider the following diagram.

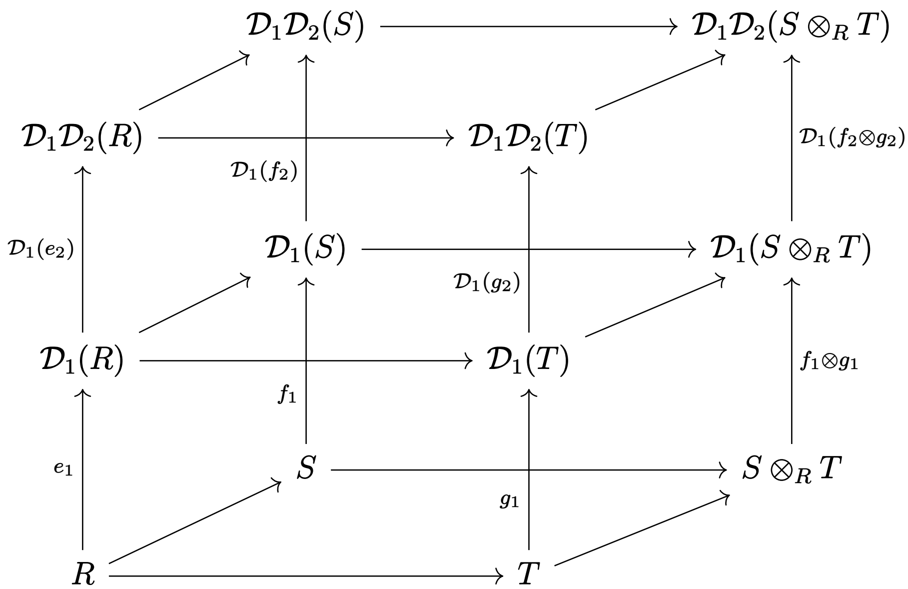

The horizontal maps are just the natural maps. The lower cube commutes due to the definition of the tensor product of $\mathcal{D}_1$-algebras, and the upper cube commutes by applying $\mathcal{D}_1$ to the cube that commutes due to the definition of the tensor product of $\mathcal{D}_2$-algebras. This means that $\mathcal{D}_1(f_2 \otimes g_2) \circ (f_1 \otimes g_1)$ is a $\mathcal{D}_1\mathcal{D}_2$-structure on $S \otimes_R T$ that extends the ones on S and T, and hence by uniqueness of the tensor product of $\mathcal{D}_1\mathcal{D}_2$-structures, we must have that

$$\mathcal{D}_1(f_2 \otimes g_2) \circ (f_1 \otimes g_1) = (\mathcal{D}_1(f_2) \circ f_1) \otimes (\mathcal{D}_1(g_2) \circ g_1).$$

∎

We now return to the case when B is a finite and free A-algebra.

Definition 2.5.6. For any B-algebra C, let $\mathcal{D}\text{-Str}_B(C)$ be the collection of triples (e, f, g) where e is a $\mathcal{D}$-structure on A, f one on B, and g one on C such that (B, f) is an (A, e)-algebra and (C, g) is a (B, f)-algebra, and the associated endomorphisms of

(B, f) have invertible matrix. For any A-algebra R, let $\mathcal{D}\text{-Str}_A(R)$ be the collection of pairs (e, u) where e is a $\mathcal{D}$-structure on A and u one on R such that (R, u) is an (A, e)-algebra. The $\mathcal{D}$-Weil descent then tells us that we have a map

$$\begin{aligned}(\,\cdot\,)^{W^{\mathcal{D}}}\colon \mathcal{D}\text{-Str}_B(C) &\to \mathcal{D}\text{-Str}_A(W(C)) \\ (e, f, g) &\mapsto (e, g^{W^{\mathcal{D}}}).\end{aligned}$$

Unless we need to be precise, we will drop the tuple notation and just use g for (e, f, g) and u for (e, u). We will also suppress the $\mathcal{D}$ notation in the map $(\,\cdot\,)^{W^{\mathcal{D}}}$ and just write $(\,\cdot\,)^{W}$. In what follows, we will make use of these maps for $\mathcal{D}_1$, $\mathcal{D}_2$, and $\mathcal{D}_1\mathcal{D}_2$, but it will be clear from context which we mean: $(\cdot)^{W^{\mathcal{D}_1}}$ will be applied only to $\mathcal{D}_1$-structures, $(\,\cdot\,)^{W^{\mathcal{D}_2}}$ only to $\mathcal{D}_2$-structures, and $(\,\cdot\,)^{W^{\mathcal{D}_1\mathcal{D}_2}}$ only to $\mathcal{D}_1\mathcal{D}_2$-structures.

Lemma 2.5.7. *The following map is well-defined.*

$$\begin{aligned}\Theta_B\colon \mathcal{D}_1\text{-Str}_B(C) \times \mathcal{D}_2\text{-Str}_B(C) &\to \mathcal{D}_1\mathcal{D}_2\text{-Str}_B(C) \\ ((e_1, f_1, g_1), (e_2, f_2, g_2)) &\mapsto (\mathcal{D}_1(e_2) \circ e_1, \mathcal{D}_1(f_2) \circ f_1, \mathcal{D}_1(g_2) \circ g_1).\end{aligned}$$

Proof. Since $(e_1, f_1, g_1) \in \mathcal{D}_1\text{-Str}_B(C)$, the following diagram commutes:

$$\begin{array}{ccccc} \mathcal{D}_1(A) & \longrightarrow & \mathcal{D}_1(B) & \longrightarrow & \mathcal{D}_1(C) \\ {\scriptstyle e_1}\uparrow & & {\scriptstyle f_1}\uparrow & & {\scriptstyle g_1}\uparrow \\ A & \longrightarrow & B & \longrightarrow & C \end{array}$$

Since $(e_2, f_2, g_2) \in \mathcal{D}_2\text{-Str}_B(C)$, we get a similar diagram. Apply $\mathcal{D}_1$ to this second diagram and compose the vertical maps to get the following commuting diagram:

$$\begin{array}{ccccc} \mathcal{D}_1\mathcal{D}_2(A) & \longrightarrow & \mathcal{D}_1\mathcal{D}_2(B) & \longrightarrow & \mathcal{D}_1\mathcal{D}_2(C) \\ {\scriptstyle \mathcal{D}_1(e_2)\circ e_1}\uparrow & & {\scriptstyle \mathcal{D}_1(f_2)\circ f_1}\uparrow & & {\scriptstyle \mathcal{D}_1(g_2)\circ g_1}\uparrow \\ A & \longrightarrow & B & \longrightarrow & C \end{array}$$

So these $\mathcal{D}_1\mathcal{D}_2$-structures make the algebra structure maps $A \to B$ and $B \to C$ into $\mathcal{D}_1\mathcal{D}_2$-homomorphisms.

Finally, we need to check that the associated endomorphisms of $(B, \mathcal{D}_1(f_2) \circ f_1)$ have invertible matrix. Recall that the associated endomorphisms are defined using the projection maps $\mathcal{D}_2 \otimes_k \mathcal{D}_1 \to k$. For $1 \le i \le t_1$ and $1 \le j \le t_2$, we will say that the (i, j)th projection map for $\mathcal{D}_2 \otimes_k \mathcal{D}_1$ is the composition $\mathcal{D}_2 \otimes_k \mathcal{D}_1 \to C_j \otimes_k B_i \to k$. Then we claim that the (i, j)th associated endomorphism of $(B, \mathcal{D}_1(f_2) \circ f_1)$ is $\sigma_j \tau_i$ where τ_i is the ith associated endomorphism of (B, f_1) and σ_j is that of (B, f_2). To see this, consider the following commuting diagram:

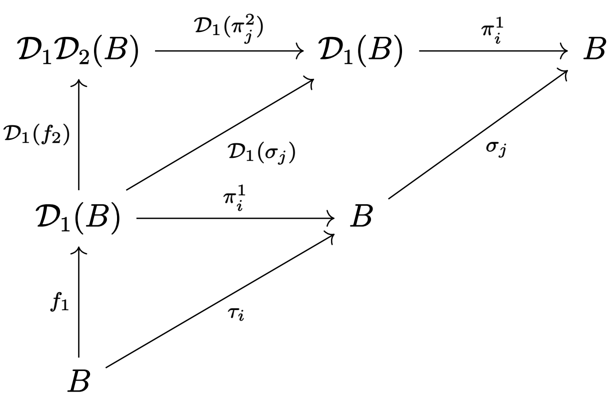

where π_i^1 is the ith projection map for $\mathcal{D}_1$ and π_j^2 is the jth projection map for $\mathcal{D}_2$.

The lower triangle commutes due to the definition of τ_i. The triangle in the upper left commutes by applying $\mathcal{D}_1$ to the definition of σ_j. It remains to show that the composition along the top row is the (i, j)th projection map for $\mathcal{D}_2 \otimes_k \mathcal{D}_1$. But this follows from the commutativity of the following diagram

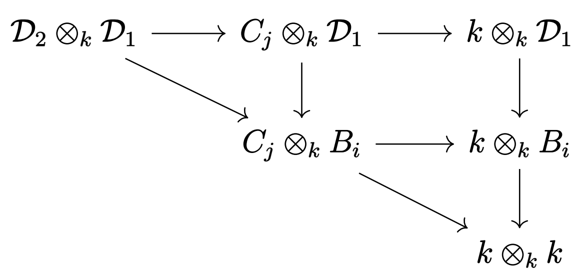

where the composition along the top row is $\mathcal{D}_1(\pi_j^2)$, the composition along the right column is π_i^1 and the diagonal composition is the (i, j)th projection map for $\mathcal{D}_2 \otimes_k \mathcal{D}_1$.

Recall that Proposition 2.3.2 says that an endomorphism has invertible matrix if and only if it sends any A-basis of B to another A-basis. Then, since τ_i and σ_j both have invertible matrix, $\sigma_j \tau_i$ must as well. ■

A similar proof also shows that we have a well-defined map

$$\begin{aligned}\Theta_A\colon\ \mathcal{D}_1\text{-}\mathrm{Str}_A(R)\times\mathcal{D}_2\text{-}\mathrm{Str}_A(R)&\to\mathcal{D}_1\mathcal{D}_2\text{-}\mathrm{Str}_A(R)\\ ((e_1,u_1),(e_2,u_2))&\mapsto(\mathcal{D}_1(e_2)\circ e_1,\mathcal{D}_1(u_2)\circ u_1).\end{aligned}$$

We also get maps

$$\mathcal{D}_2\text{-}\mathrm{Str}_B(C)\times\mathcal{D}_1\text{-}\mathrm{Str}_B(C)\to\mathcal{D}_2\mathcal{D}_1\text{-}\mathrm{Str}_B(C)$$

and

$$\mathcal{D}_2\text{-}\mathrm{Str}_A(R)\times\mathcal{D}_1\text{-}\mathrm{Str}_A(R)\to\mathcal{D}_2\mathcal{D}_1\text{-}\mathrm{Str}_A(R)$$

by exchanging the roles of $\mathcal{D}_1$ and $\mathcal{D}_2$. We will also denote these maps by Θ_B and Θ_A, but it will be clear from context which one we mean.

Theorem 2.5.8. *For $g_1\in\mathcal{D}_1\text{-}\mathrm{Str}_B(C)$ and $g_2\in\mathcal{D}_2\text{-}\mathrm{Str}_B(C)$,*

$$\Theta_B(g_1,g_2)^W=\Theta_A(g_1^W,g_2^W).$$

Proof. Using the $\mathcal{D}_1$-Weil descent and the $\mathcal{D}_2$-Weil descent, the following squares commute:

$$\begin{array}{ccc}\mathcal{D}_1(C)&\xrightarrow{\mathcal{D}_1(\eta_C)}&\mathcal{D}_1(FW(C))\\ {\scriptstyle g_1}\big\uparrow&&\big\uparrow{\scriptstyle g_1^W\otimes f_1}\\ C&\xrightarrow{\ \eta_C\ }&FW(C)\end{array}$$

$$\begin{array}{ccc}\mathcal{D}_2(C)&\xrightarrow{\mathcal{D}_2(\eta_C)}&\mathcal{D}_2(FW(C))\\ {\scriptstyle g_2}\big\uparrow&&\big\uparrow{\scriptstyle g_2^W\otimes f_2}\\ C&\xrightarrow{\ \eta_C\ }&FW(C)\end{array}$$

Apply $\mathcal{D}_1$ to the second square and compose the vertical maps so that the following square commutes:

$$\begin{array}{ccc}\mathcal{D}_1\mathcal{D}_2(C)&\xrightarrow{\mathcal{D}_1\mathcal{D}_2(\eta_C)}&\mathcal{D}_1\mathcal{D}_2(FW(C))\\ {\scriptstyle \mathcal{D}_1(g_2)\circ g_1}\big\uparrow&&\big\uparrow{\scriptstyle \mathcal{D}_1(g_2^W\otimes f_2)\circ(g_1^W\otimes f_1)}\\ C&\xrightarrow{\ \eta_C\ }&FW(C)\end{array}$$

By Lemma 2.5.5, the right vertical map is equal to $(\mathcal{D}_1(g_2^W)\circ g_1^W)\otimes(\mathcal{D}_1(f_2)\circ f_1)$.

And hence, by the uniqueness of the $\mathcal{D}_1\mathcal{D}_2$-structure on $W(C)$ that makes it into an $(A, \mathcal{D}_1(e_2) \circ e_1)$-algebra and η_C into a $(B, \mathcal{D}_1(f_2) \circ f_1)$-algebra homomorphism, we must have

$$(\mathcal{D}_1(g_2) \circ g_1)^W = \mathcal{D}_1(g_2^W) \circ g_1^W.$$

■

We now specialise this theorem to the difference case. Let $\mathcal{D} = \mathcal{D}_1 = \mathcal{D}_2 = k$. Then $\mathcal{D}_2 \otimes_k \mathcal{D}_1 = k$ and Θ_A and Θ_B are just composition of endomorphisms. $\mathcal{D}\text{-Str}_B(C)$ is a monoid with composition Θ_B and identity $(\mathrm{id}_A, \mathrm{id}_B, \mathrm{id}_C)$. Similary, $\mathcal{D}\text{-Str}_A(R)$ is a monoid under Θ_A and $(\mathrm{id}_A, \mathrm{id}_R)$.

Corollary 2.5.9. *In the notation of the above paragraph,*

$$(\,\cdot\,)^W \colon \mathcal{D}\text{-Str}_B(C) \to \mathcal{D}\text{-Str}_A(W(C))$$

is a monoid homomorphism.

Proof. We have that $\Theta_B(g_1, g_2)^W = (g_1 \circ g_2)^W$ and $\Theta_A(g_1^W, g_2^W) = g_1^W \circ g_2^W$. Then Theorem 2.5.8 tells us that $(g_1 \circ g_2)^W = g_1^W \circ g_2^W$. Lemma 2.5.1 then tells us that $(\mathrm{id}_C)^W = \mathrm{id}_{W(C)}$. ■

Remark 2.5.10. Corollary 2.5.9 tells us that the difference Weil descent restricts to the categories of inversive difference algebras, that is, algebras equipped with an automorphism. Indeed, if (A, e), (B, f) and (C, g) are all inversive difference algebras, applying $(\,\cdot\,)^W$ to the equations $g \circ g^{-1} = \mathrm{id}_C = g^{-1} \circ g$ tells us that g^W is also an automorphism on $W(C)$.

We now further develop these results to study the commutativity of a $\mathcal{D}_1$-structure and a $\mathcal{D}_2$-structure. Let Γ be the canonical isomorphism

$$\Gamma \colon \mathcal{D}_2 \otimes_k D_1 \to \mathcal{D}_1 \otimes_k \mathcal{D}_2$$
$$\alpha_2 \otimes \alpha_1 \mapsto \alpha_1 \otimes \alpha_2.$$

For any k-algebra S, Γ lifts to $\Gamma^S \colon S \otimes_k \mathcal{D}_2 \otimes_k \mathcal{D}_1 \to S \otimes_k \mathcal{D}_1 \otimes_k \mathcal{D}_2$ in the usual way. Therefore, Γ induces maps $\mathcal{D}_1\mathcal{D}_2\text{-Str}_B(C) \to \mathcal{D}_2\mathcal{D}_1\text{-Str}_B(C)$ and $\mathcal{D}_1\mathcal{D}_2\text{-Str}_A(R) \to \mathcal{D}_2\mathcal{D}_1\text{-Str}_A(R)$ by applying the appropriate Γ coordinate-wise. We will also denote these maps Γ. It should be clear from context which we mean.

Definition 2.5.11. Let S be a k-algebra, equipped with a $\mathcal{D}_1$-structure e_1 and a $\mathcal{D}_2$-structure e_2. We will say that e_1 commutes with e_2 if

$$\Gamma^S \circ \mathcal{D}_1(e_2) \circ e_1 = \mathcal{D}_2(e_1) \circ e_2.$$

For $g_1 \in \mathcal{D}_1\text{-Str}_B(C)$ and $g_2 \in \mathcal{D}_2\text{-Str}_B(C)$, we will say that g_1 commutes with g_2 if

$$\Gamma \circ \Theta_B(g_1, g_2) = \Theta_B(g_2, g_1).$$

Similarly, for $u_1 \in \mathcal{D}_1\text{-Str}_A(R)$ and $u_2 \in \mathcal{D}_2\text{-Str}_A(R)$, we will say that u_1 commutes with u_2 if $\Gamma \circ \Theta_A(u_1, u_2) = \Theta_A(u_2, u_1)$.

Remark 2.5.12. If we choose bases of $\mathcal{D}_1$ and $\mathcal{D}_2$ and think of e_1 and e_2 as their corresponding sequence of free operators, the condition

$$\Gamma^S \circ \mathcal{D}_1(e_2) \circ e_1 = \mathcal{D}_2(e_1) \circ e_2$$

says that every operator of e_1 commutes with every operator of e_2.

We now prove a modification of Theorem 2.5.8 that includes Γ.

Lemma 2.5.13. *For* $g_1 \in \mathcal{D}_1\text{-Str}_B(C)$ *and* $g_2 \in \mathcal{D}_2\text{-Str}_B(C)$,

$$(\Gamma \circ \Theta_B(g_1, g_2))^W = \Gamma \circ \Theta_A(g_1^W, g_2^W).$$

Proof. Firstly, suppose e_1, f_1, g_1 and e_2, f_2, g_2 are $\mathcal{D}_1$- and $\mathcal{D}_2$-structures on R, S, and T making (S, f_1) and (T, g_1) into (R, e_1)-algebras and (S, f_2) and (T, g_2) into

(R, e_2)-algebras. Consider the following diagram:

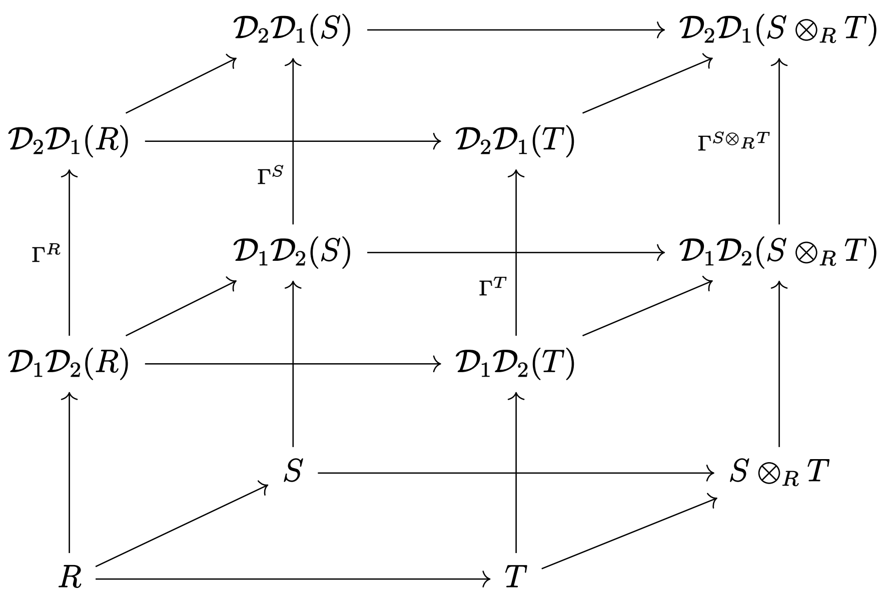

where the horizontal maps are the usual ones and the vertical ones in the lower cube are the compositions of the $\mathcal{D}_1$-structure and $\mathcal{D}_2$-structure. By the uniqueness of the $\mathcal{D}_2\mathcal{D}_1$-structure on $S \otimes_R T$, we have that

$$\Gamma^{S\otimes_R T} \circ \mathcal{D}_1(f_2 \otimes g_2) \circ (f_1 \otimes g_1) = \left(\Gamma^S \circ \mathcal{D}_1(f_2) \circ f_1\right) \otimes \left(\Gamma^T \circ \mathcal{D}_1(g_2) \circ g_1\right).$$

Now, returning to the original context, the following diagram also commutes:

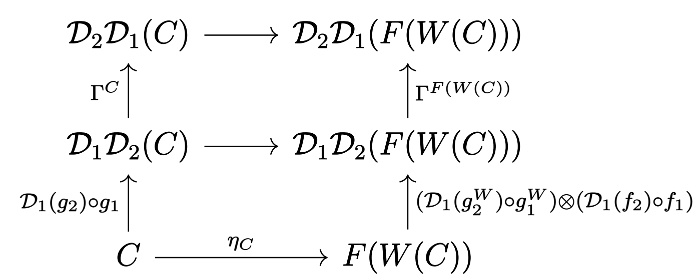

Hence, $\Gamma^{W(C)} \circ \mathcal{D}_1(g_2^W) \circ g_1^W$ is a $\mathcal{D}_2\mathcal{D}_1$-structure on $W(C)$ making it into an $(A, \Gamma^A \circ \mathcal{D}_1(e_2) \circ e_1)$-algebra and η_C into a $(B, \Gamma^B \circ \mathcal{D}_1(f_2) \circ f_1)$-algebra homomorphism. If the associated endomorphisms of $\Gamma^B \circ \mathcal{D}_1(f_2) \circ f_1$ all had invertible matrix, then by the uniqueness of such a $\mathcal{D}_2\mathcal{D}_1$-structure, we must have that $(\Gamma^C \circ \mathcal{D}_1(g_2) \circ g_1)^W = \Gamma^{W(C)} \circ \mathcal{D}_1(g_2^W) \circ g_1^W$, from which the result follows.

Now, note that the (i, j)th projection map for $\mathcal{D}_1 \otimes_k \mathcal{D}_2 = \prod_{i=1}^{t_2} \prod_{j=1}^{t_1} B_j \otimes C_i$

is $\mathcal{D}_1 \otimes_k \mathcal{D}_2 \to B_j \otimes_k C_i \to k$. Let $\pi^{\mathcal{D}_1\otimes\mathcal{D}_2}_{(i,j)}$ denote the (i,j)th projection map for $\mathcal{D}_1 \otimes_k \mathcal{D}_2$, and let $\pi^{\mathcal{D}_2\otimes\mathcal{D}_1}_{(i,j)}$ denote the (i,j)th projection map for $\mathcal{D}_2 \otimes_k \mathcal{D}_1$: $\mathcal{D}_2 \otimes_k \mathcal{D}_1 \to C_j \otimes_k B_i \to k$. Then $\pi^{\mathcal{D}_1\otimes\mathcal{D}_2}_{(i,j)} \circ \Gamma = \pi^{\mathcal{D}_2\otimes\mathcal{D}_1}_{(j,i)}$. Thus, the (i,j)th associated endomorphism of $\Gamma^B \circ \mathcal{D}_1(f_2) \circ f_1$ is the (j,i)th associated endomorphism of $\mathcal{D}_1(f_2) \circ f_1$, $\sigma_i \tau_j$, which has invertible matrix. ■

Corollary 2.5.14. *Let $g_1 \in \mathcal{D}_1\text{-Str}_B(C)$ and $g_2 \in \mathcal{D}_2\text{-Str}_B(C)$. If g_1 commutes with g_2, then g_1^W commutes with g_2^W.*

Proof. If g_1 commutes with g_2, then $\Gamma \circ \Theta_B(g_1, g_2) = \Theta_B(g_2, g_1)$. Applying $(\,\cdot\,)^W$ to this equation and using Theorem 2.5.8 and Lemma 2.5.13, we get $\Gamma \circ \Theta_A(g_1^W, g_2^W) = \Theta_A(g_2^W, g_1^W)$. Hence, g_1^W commutes with g_2^W. ■

For a k-algebra S, we will say that a $\mathcal{D}$-structure e on S commutes if $\Gamma^S \circ \mathcal{D}(e) \circ e = \mathcal{D}(e) \circ e$. Note that this is equivalent to saying that, with respect to a fixed basis of $\mathcal{D}$, the free operators corresponding to e pairwise commute. For $g \in \mathcal{D}\text{-Str}_B(C)$, we will say g commutes if $\Gamma \circ \Theta_B(g, g) = \Theta_B(g, g)$, and similarly for $u \in \mathcal{D}\text{-Str}_A(R)$, u commutes if $\Gamma \circ \Theta_A(u, u) = \Theta_A(u, u)$. An immediate consequence of Corollary 2.5.14 is the following.

Corollary 2.5.15. *Let $g \in \mathcal{D}\text{-Str}_B(C)$. If g commutes, then g^W commutes.*

These results allow us to deduce that commutativity is preserved by the $\mathcal{D}$-Weil descent in several cases. We give details for the case of m endomorphisms and n derivations.

Example 2.5.16. Suppose $\mathcal{D} = k[x_1, \ldots, x_n]/(x_1, \ldots, x_n)^2 \times k^m$ and that for every $\mathcal{D}$-structure, the first associated endomorphism is trivial (unless $n = 0$, in which case we do not impose that any associated endomorphism is trivial). Then a $\mathcal{D}$-structure is a collection of n derivations and m endomorphisms. Suppose also that for a given A, B, C, all of the derivations and endomorphisms pairwise commute. Then, by Corollary 2.5.15, we have that the Weil descents of all the derivations and endomorphisms pairwise commute.

Remark 2.5.17. The $n = 0$ case also follows from Corollary 2.5.9. The $m = 0$ case appears in [39].

Remark 2.5.18. It seems possible that Corollary 2.5.15 could be extended to a more general context where commutativity is replaced by an iterativity condition as in

Section 2.2 of [52]. We leave this for future work as it goes beyond the scope of this thesis.

On the necessity of having invertible matrix

It is a natural question to ask whether a converse to our main theorem holds.

Question 2.5.19. If $F^{\mathcal{D}}$ has a left adjoint, must every associated endomorphism of (B, f) have invertible matrix?

We do not yet know the answer in general, but we do have the following partial converse which imposes some mild conditions on such a left adjoint. We use the following notation. For each $z \in \mathcal{D}(B)$, let $g^z\colon B[t] \to \mathcal{D}(B[t])$ be the $\mathcal{D}$-structure on $B[t]$ that extends f on B and sends $t \mapsto z$.

Theorem 2.5.20. *Suppose $F^{\mathcal{D}}$ has a left adjoint, $W^{\mathcal{D}}$, and that for each $z \in \mathcal{D}(B)$ the underlying A-algebra of $W^{\mathcal{D}}(B[t], g^z)$ is a faithfully flat A-module. Then the associated endomorphisms of (B, f) all have invertible matrix.*

Proof. Note that by Section 1.9 of [66], for any R-algebra S, S is a faithfully flat R-module if and only if S is a flat R-module and every linear system of equations defined over R which has a solution in S already has a solution in R.

For $z \in \mathcal{D}(B)$, consider the unit of the adjunction $\eta\colon (B[t], g^z) \to F^{\mathcal{D}}W^{\mathcal{D}}(B[t], g^z)$. That this is a $\mathcal{D}$-homomorphism at t means that

$$\sum_{m=1}^{l}\sum_{n=1}^{r} \lambda_n(\eta(g^z_m(t))) \otimes b_n\, \varepsilon_m = \sum_{i=1}^{r}\sum_{j=1}^{l}\sum_{k=1}^{l}\sum_{n=1}^{r}\sum_{m=1}^{l} a_{jkm}\lambda_n(f_k(b_i))h^z_j(\lambda_i(\eta(t))) \otimes b_n\, \varepsilon_m \qquad (*)$$

where h^z is the $\mathcal{D}$-structure on $W^{\mathcal{D}}(B[t], g^z)$. Write $z = \sum_m \beta_m \varepsilon_m$ and $\beta_m = \sum_n a_{nm} b_n$. Then $\lambda_n(\eta(g^z_m(t))) = \lambda_n(\eta(\beta_m)) = \lambda_n(\beta_m) = a_{nm}$ since η is a B-algebra homomorphism.

Let $\bar{a}$ be the vector in A^{rl} of the elements a_{nm}. Then equation $(*)$ tells us that we have a solution in $W^{\mathcal{D}}(B[t], g^z)$ to the linear system $\bar{a} = M\bar{x}$. Since $W^{\mathcal{D}}(B[t], g^z)$ is faithfully flat, we have a solution in A, and hence M is onto as a linear map $A^{rl} \to A^{rl}$. Then M is invertible. ∎

If A is a field, then $W^{\mathcal{D}}(B[t], g^z)$ is a free A-module – hence faithfully flat – and so Theorems 2.4.5 and 2.5.20 yield the following:

Corollary 2.5.21. *Suppose A is a field. Then $F^{\mathcal{D}}$ has a left adjoint if and only if the associated endomorphisms of (B, f) all have invertible matrix.*

This result specialises to the difference case:

Corollary 2.5.22. *Suppose $(K, \sigma) \leq (L, \tau)$ is an extension of difference fields where L/K is finite and σ is an automorphism. Then the difference base change functor has a left adjoint (the difference Weil descent).*

Proof. Note that by Lemma 2.3.4, τ is an automorphism if and only if it has invertible matrix. Since σ is an automorphism, L/K is a finite-dimensional K-vector space, and τ is injective, τ must also be an automorphism. ■

2.6 An explicit construction of the $\mathcal{D}$-Weil descent

While the construction of the $\mathcal{D}$-ring structure g^W given in Section 2.4 is very natural, it does not yield an explicit or computational construction. In this appendix we will sketch a construction that parallels the classical one. Let $\varepsilon_1, \dots, \varepsilon_l$ be a k-basis of $\mathcal{D}$ and $b_1, \dots, b_r$ an A-basis of B. We continue to impose Assumption **A**.

Recall the following notation. If (R, e) is a $\mathcal{D}$-ring, $e_i \colon R \to R$ denotes the i-th coordinate map of e with respect to the basis ε. That is, the maps e_i are the additive operators of R such that $e(r) = \sum_{i=1}^{l} e_i(r)\varepsilon_i$ for all $r \in R$. We also have that $\lambda_n \colon B \to A$ is the A-module homomorphism given by $b = \sum_{i=1}^{r} \lambda_i(b) b_i$.

The matrix M is defined as follows.

$$M = \begin{bmatrix} M_{11} & M_{12} & \cdots & M_{1,l} \\ M_{21} & M_{22} & \cdots & M_{2,l} \\ \vdots & \vdots & \ddots & \vdots \\ M_{l,1} & M_{l,2} & \cdots & M_{l,l} \end{bmatrix}$$

where

$$(M_{mj})_{ni} = \sum_{k=1}^{l} a_{jkm} \lambda_n(f_k(b_i)).$$

We need some notation and a technical result relating the matrix M to whether an algebra homomorphism is also a $\mathcal{D}$-homomorphism.

Notation. For a collection of elements $\{x_{ij} \colon 1 \leq i \leq r, 1 \leq j \leq l\}$ in some A-algebra, we write (x_{ij}) for the rl-vector ordered reverse lexicographically on the indices i and j. We write $M \cdot (x_{ij})$ to denote the standard matrix multiplication of an $rl \times rl$ matrix with an rl-vector resulting in an rl-vector.

Lemma 2.6.1. *Let (C, g) be a (B, f)-algebra, (R, u) an (A, e)-algebra, and $\phi \colon C \to F(R) = R \otimes_A B$ a B-algebra homomorphism. Then ϕ is a (B, f)-algebra homomorphism if and only if the following equation holds for every $c \in C$:*

$$(\lambda_i \phi g_j(c)) = M \cdot (u_j \lambda_i \phi(c)) \tag{$*$}$$

As a result, when M is invertible, the values $u_j \lambda_i \phi(c)$ are uniquely determined.

Proof. ϕ is a (B, f)-algebra homomorphism if and only if it is a $\mathcal{D}$-homomorphism, if and only if the following diagram commutes:

$$\begin{array}{ccc}
\mathcal{D}(C) & \xrightarrow{\mathcal{D}(\phi)} & \mathcal{D}(F(R)) \\
{\scriptstyle g}\uparrow & & \uparrow{\scriptstyle u \otimes f} \\
C & \xrightarrow{\quad \phi \quad} & F(R)
\end{array}$$

Now expand both compositions and equate coefficients of the $b_n \varepsilon_m$. ■

Remark 2.6.2. If some $S \subseteq C$ generates C as a B-algebra, then it is enough to ask for equality $(*)$ to hold for every $s \in S$.

Our explicit construction of the $\mathcal{D}$-Weil descent parallels the classical construction. So we need the algebraic notions of $\mathcal{D}$-ideals, $\mathcal{D}$-quotients, and $\mathcal{D}$-polynomial rings.

Definition 2.6.3. Let (R, e) be a $\mathcal{D}$-ring, and let I be an ideal of R. We define $\mathcal{D}(I)$ to be the k-submodule of $\mathcal{D}(R)$ given by $\mathcal{D}(I) := I \otimes_k \mathcal{D}$. We say that I is a $\mathcal{D}$-ideal if $e(I) \subseteq \mathcal{D}(I)$. Note that $\mathcal{D}(I)$ is an ideal of $\mathcal{D}(R)$: if $IR \subseteq I$, then

$$\mathcal{D}(I) \cdot \mathcal{D}(R) = (I \otimes_k \mathcal{D}) \cdot (R \otimes_k \mathcal{D}) \subseteq I \otimes_k \mathcal{D}.$$

Remark 2.6.4. In the context of Example 2.1.1(1) where σ is trivial, I is a $\mathcal{D}$-ideal if and only if it is a differential ideal, that is, if $\delta(I) \subseteq I$. For Example 2.1.1(2), $\mathcal{D}$-ideals are ideals with $\sigma_i(I) \subseteq I$ for each i.

Lemma 2.6.5. *Let (R, e) and (S, f) be two $\mathcal{D}$-rings and suppose $\phi\colon R \to S$ is a $\mathcal{D}$-homomorphism. Then* $\ker\phi$ *is a $\mathcal{D}$-ideal.*

Proof. Since ϕ is a $\mathcal{D}$-homomorphism, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{D}(R) & \xrightarrow{\mathcal{D}(\phi)} & \mathcal{D}(S) \\ {\scriptstyle e}\uparrow & & \uparrow{\scriptstyle f} \\ R & \xrightarrow{\phi} & S \end{array}$$

For $g \in \ker\phi$, $f \circ \phi(g) = 0$, and so $\mathcal{D}(\phi) \circ e(g) = 0$. Then $e(g) \in \ker\mathcal{D}(\phi)$. Consider the standard kernel-cokernel exact sequence for ϕ:

$$0 \longrightarrow \ker\phi \longrightarrow R \xrightarrow{\phi} S \longrightarrow \operatorname{coker}\phi \longrightarrow 0$$

$\mathcal{D}$ is a free (and hence flat) k-module, so tensoring is exact:

$$0 \longrightarrow \mathcal{D}(\ker\phi) \longrightarrow \mathcal{D}(R) \xrightarrow{\mathcal{D}(\phi)} \mathcal{D}(S) \longrightarrow \mathcal{D}(\operatorname{coker}\phi) \longrightarrow 0$$

We also have the kernel-cokernel exact sequence for $\mathcal{D}(\phi)$:

$$0 \longrightarrow \ker\mathcal{D}(\phi) \longrightarrow \mathcal{D}(R) \xrightarrow{\mathcal{D}(\phi)} \mathcal{D}(S) \longrightarrow \operatorname{coker}\mathcal{D}(\phi) \longrightarrow 0$$

Now $\mathcal{D}(\ker\phi) \subseteq \ker\mathcal{D}(\phi)$ and hence we get the following commuting diagram with exact rows:

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{D}(\ker\phi) & \longrightarrow & \mathcal{D}(R) & \xrightarrow{\mathcal{D}(\phi)} & \mathcal{D}(S) \\ \downarrow & & \downarrow\scriptstyle{\subseteq} & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \ker\mathcal{D}(\phi) & \longrightarrow & \mathcal{D}(R) & \xrightarrow{\mathcal{D}(\phi)} & \mathcal{D}(S) \end{array}$$

Then the four lemma tells us that the inclusion is onto, so $\ker\mathcal{D}(\phi) = \mathcal{D}(\ker\phi)$ and

$e(g) \in \mathcal{D}(\ker \phi)$. ■

Lemma 2.6.6. *Let I be an ideal of R. Then $\mathcal{D}(R)/\mathcal{D}(I) \cong \mathcal{D}(R/I)$.*

Proof. We have the following exact sequence:

$$0 \longrightarrow I \longrightarrow R \longrightarrow R/I \longrightarrow 0$$

Since $\mathcal{D}$ is a flat k-module, tensoring is exact:

$$0 \longrightarrow \mathcal{D}(I) \longrightarrow \mathcal{D}(R) \longrightarrow \mathcal{D}(R/I) \longrightarrow 0$$

So $\mathcal{D}(R/I) \cong \mathcal{D}(R)/\mathcal{D}(I)$. ■

Lemma 2.6.7. *Let (R, e) be a $\mathcal{D}$-ring, (S, f) an (R, e)-algebra, and I a $\mathcal{D}$-ideal of S. Then there exists a unique $\mathcal{D}$-ring structure on the quotient S/I given by*

$$\begin{aligned} \bar{f} \colon S/I &\to \mathcal{D}(S/I) \\ s + I &\to f(s) + \mathcal{D}(I) \end{aligned}$$

which makes the quotient map $q \colon S \to S/I$ into an (R, e)-algebra homomorphism.

Proof. First, note that, by Lemma 2.6.6, $\bar{f}$ is indeed a well-defined k-algebra homomorphism since I is a $\mathcal{D}$-ideal of S. Consider the following diagram:

$$\begin{array}{ccccc} \mathcal{D}(R) & \longrightarrow & \mathcal{D}(S) & \xrightarrow{\mathcal{D}(q)} & \mathcal{D}(S/I) \\ {\scriptstyle e}\big\uparrow & & {\scriptstyle f}\big\uparrow & & {\scriptstyle \bar{f}}\big\uparrow \\ R & \longrightarrow & S & \xrightarrow{q} & S/I \end{array}$$

The left square commutes since (S, f) is an (R, e)-algebra, and the right square commutes because of Lemma 2.6.6. Then q is a $\mathcal{D}$-homomorphism. Note that the composition of the lower horizontal maps is the R-algebra structure on S/I, and hence S/I is an (R, e)-algebra. This shows that q is an (R, e)-algebra homomorphism. ■

We now need the natural notion of a $\mathcal{D}$-polynomial ring. These have been defined in Section 3.1 of [52] (implicitly) and in Remark 3.8 of [53]. We expand on the details here.

Definition 2.6.8. We denote by Θ the set of all finite words on the alphabet $\{1, \ldots, l\}$. For a $\mathcal{D}$-ring (R, e) and $\theta \in \Theta$, we will write e_θ for the corresponding composition of coordinatised $\mathcal{D}$-operators. For example, if $\theta = 123$, then $e_\theta = e_3 \circ e_2 \circ e_1$. Note then that $e_{\theta_1\theta_2} = e_{\theta_2} \circ e_{\theta_1}$.

Definition 2.6.9. Let (R, e) be a $\mathcal{D}$-ring and $T = (t)_{t\in T}$ a collection of indeterminates. The $\mathcal{D}$-polynomial algebra in indeterminates T over (R, e) with respect to ε is the ring

$$R\{T\}^{\varepsilon}_{\mathcal{D}} = R[t^\theta : t \in T \text{ and } \theta \in \Theta]$$

where $(t^\theta)_{t\in T,\theta\in\Theta}$ is a new family of indeterminates, equipped with $\mathcal{D}$-ring structure

$$\begin{aligned} e' \colon R\{T\}^{\varepsilon}_{\mathcal{D}} &\to \mathcal{D}(R\{T\}^{\varepsilon}_{\mathcal{D}}) \\ t^\theta &\mapsto t^{\theta 1}\varepsilon_1 + t^{\theta 2}\varepsilon_2 + \cdots t^{\theta l}\varepsilon_l \\ r &\mapsto e(r) \end{aligned}$$

This makes $(R\{T\}^{\varepsilon}_{\mathcal{D}}, e')$ an (R, e)-algebra.

Suppose (S, f) is an (R, e)-algebra and $X \subseteq S$. We denote by $R\{X\}_{\mathcal{D}}$ the $\mathcal{D}$-ring generated in S by X over (R, e). This is a well-defined notion since the intersection of a collection of $\mathcal{D}$-subrings is a $\mathcal{D}$-subring.

Lemma 2.6.10. *Suppose that (S, f) is an (R, e)-algebra which is generated as a $\mathcal{D}$-ring by the (possibly infinite) tuple $\bar{a} = (a_i)_{i\in I}$ over (R, e), so that $S = R\{\bar{a}\}_{\mathcal{D}}$. Let $\bar{t} = (t_i)_{i\in I}$ be a tuple of indeterminates. Then there exists a unique surjective (R, e)-algebra homomorphism $\mathrm{ev}_{\bar{a}} \colon R\{\bar{t}\}^{\varepsilon}_{\mathcal{D}} \to S$ which maps $t_i \mapsto a_i$ for each $i \in I$.*

Proof. Define $\mathrm{ev}_{\bar{a}}(t_i^\theta) = f_\theta(a_i)$ (see Definition 2.6.8). Then $\mathrm{ev}_{\bar{a}}$ is clearly a surjective R-algebra homomorphism. To show it is a $\mathcal{D}$-homomorphism, we need to show that the following diagram commutes:

$$\begin{array}{ccc} \mathcal{D}(R\{\bar{t}\}^{\varepsilon}_{\mathcal{D}}) & \xrightarrow{\mathcal{D}(\mathrm{ev}_{\bar{a}})} & \mathcal{D}(S) \\ {\scriptstyle e}\uparrow & & \uparrow{\scriptstyle f} \\ R\{\bar{t}\}^{\varepsilon}_{\mathcal{D}} & \xrightarrow{\mathrm{ev}_{\bar{a}}} & S \end{array}$$

As both $R\{\bar{t}\}^{\varepsilon}_{\mathcal{D}}$ and S are (R, e)-algebras, it is sufficient to check commutativity of this diagram for each t_i^θ. We have

$$\mathcal{D}(\mathrm{ev}_{\bar{a}}) \circ e(t_i^\theta) = \mathcal{D}(\mathrm{ev}_{\bar{a}})(\sum_{j=1}^{l} t_i^{\theta j} \varepsilon_j) = \sum_{j=1}^{l} f_{\theta j}(a_i)\varepsilon_j$$

and

$$f \circ \mathrm{ev}_{\bar{a}}(t_i^\theta) = f(f_\theta(a_i)) = \sum_{j=1}^{l} f_j(f_\theta(a_i))\varepsilon_j = \sum_{j=1}^{l} f_{\theta j}(a_i)\varepsilon_j.$$

For uniqueness, note that if ϕ is also an (R, e)-algebra homomorphism with $\phi(t_i) = a_i$, then, since it is a $\mathcal{D}$-homomorphism, $\phi(t_i^\theta) = f_\theta(a_i)$ for all $\theta \in \Theta$. Then ϕ and $\mathrm{ev}_{\bar{a}}$ agree on the generators of the polynomial algebra and are both R-algebra homomorphisms, so must be equal. ■

This lemma yields the following.

Corollary 2.6.11. *Suppose that ε and ω are two bases of $\mathcal{D}$. Then $R\{T\}_{\mathcal{D}}^{\varepsilon}$ and $R\{T\}_{\mathcal{D}}^{\omega}$ are isomorphic as (R, e)-algebras.*

As a result of this corollary, we omit the superscript and just write $R\{T\}_{\mathcal{D}}$.

Remark 2.6.12. Combining the above results, we see that any (R, e)-algebra is a quotient of some $\mathcal{D}$-polynomial algebra over (R, e) by a $\mathcal{D}$-ideal.

We now return to the construction of the $\mathcal{D}$-Weil restriction. As usual, (A, e) is a $\mathcal{D}$-ring, and (B, f) is an (A, e)-algebra where B is finite and free as an A-module with basis $b_1, \dots, b_r$. Recall that the component of the unit of the classical adjunction at the polynomial algebra $B\{T\}_{\mathcal{D}}$ is

$$\eta_{B\{T\}_{\mathcal{D}}}\left(t^\theta\right) = \sum_{i=1}^{r} t^\theta(i) \otimes b_i$$

We first construct the $\mathcal{D}$-Weil descent for a $\mathcal{D}$-polynomial algebra over (B, f).

Lemma 2.6.13. *Let T be a set of indeterminates. Then there exists a $\mathcal{D}$-structure s on $W(B\{T\}_{\mathcal{D}})$ making $(W(B\{T\}_{\mathcal{D}}), s)$ into an (A, e)-algebra and $\eta_{B\{T\}_{\mathcal{D}}}$ into a $\mathcal{D}$-homomorphism.*

Proof. As A-algebras, we have $W(B\{T\}_{\mathcal{D}}) = A\{T\}_{\mathcal{D}}^{\otimes r}$, and applying Lemma 2.6.1 with $\eta_{B\{T\}_{\mathcal{D}}}$ tells us that $\eta_{B\{T\}_{\mathcal{D}}}$ is a $\mathcal{D}$-homomorphism if and only if

$$(\lambda_i \eta_{B\{T\}_{\mathcal{D}}} h_j(t^\theta)) = M \cdot (s_j \lambda_i \eta_{B\{T\}_{\mathcal{D}}}(t^\theta))$$

for every t^θ. Here h is the $\mathcal{D}$-structure on the $\mathcal{D}$-polynomial ring $B\{T\}_\mathcal{D}$ described in Definition 2.6.9. Now, $(s_j\lambda_i\eta_{B\{T\}_\mathcal{D}}(t^\theta)) = s_j(t^\theta(i))$, and $(\lambda_i\eta_{B\{T\}_\mathcal{D}}h_j(t^\theta)) = t^{\theta j}(i)$. Since M is invertible we have $s_j(t^\theta(i)) = M^{-1}\cdot(t^{\theta j}(i))$. This gives an explicit expression for s_j on each generator of $A\{T\}_\mathcal{D}^{\otimes r}$ and hence an explicit expression for s. Since $A\{T\}_\mathcal{D}^{\otimes r}$ is a polynomial algebra, this gives a $\mathcal{D}$-ring structure on $A\{T\}_\mathcal{D}^{\otimes r}$ making $\eta_{B\{T\}_\mathcal{D}}$ into a $\mathcal{D}$-homomorphism. ■

Remark 2.6.14. Note that $W(B\{T\}_\mathcal{D})$ is a polynomial algebra, but that, in general, s is not the $\mathcal{D}$-ring structure that makes $(W(B\{T\}_\mathcal{D}), s)$ a $\mathcal{D}$-polynomial algebra as in Definition 2.6.9 – it is twisted by M^{-1}. The same occurs in the differential case; see the proof of Theorem 3.2 of [39].

Now let (C, g) be a (B, f)-algebra. By Lemma 2.6.10, there is a set of indeterminates T and a surjective (B, f)-algebra homomorphism $\pi_C\colon B\{T\}_\mathcal{D} \to C$, where $B\{T\}_\mathcal{D}$ has the standard $\mathcal{D}$-structure h extending f with $h(t^\theta) = t^{\theta 1}\varepsilon_1 + \cdots + t^{\theta l}\varepsilon_l$. The component of the unit of the classical adjunction at C, η_C, is given by

$$\eta_C\left(\pi_C(t^\theta)\right) = \sum_{i=1}^{r} W(\pi_C)(t^\theta(i)) \otimes b_i.$$

Recall from Section 1.3 the definition of the ideal I_C. This ideal is generated by the elements $\lambda_i(\eta_{B\{T\}_\mathcal{D}}(\gamma))$ as γ ranges over $\ker \pi_C$, and $W(\pi_C)$ is the residue map of this ideal.

Lemma 2.6.15. *The ideal I_C of $W(B\{T\}_\mathcal{D})$ is a $\mathcal{D}$-ideal for the $\mathcal{D}$-structure s given in Lemma 2.6.13.*

Proof. Let $\gamma \in \ker \pi_C$. By definition of I_C, we need to show $s(\lambda_i(\eta_{B\{T\}_\mathcal{D}}(\gamma))) \in \mathcal{D}(I_C)$ for each i, that is, that the vector $(s_j\lambda_i\eta_{B\{T\}_\mathcal{D}}(\gamma)) \in I_C$.

Since $\eta_{B\{T\}_\mathcal{D}}$ is a $\mathcal{D}$-homomorphism, we have

$$(\lambda_i\eta_{B\{T\}_\mathcal{D}}h_j(\gamma)) = M\cdot(s_j\lambda_i\eta_{B\{T\}_\mathcal{D}}(\gamma))$$

Now $\ker \pi_C$ is a $\mathcal{D}$-ideal for $h\colon B\{T\}_D \to \mathcal{D}(B\{T\}_D)$ – the standard $\mathcal{D}$-polynomial structure – and so $h_j(\gamma) \in \ker \pi_C$. Then, by construction of I_C, $(\lambda_i\eta_{B\{T\}_\mathcal{D}}h_j(\gamma))$ is in I_C. Since M is invertible, $(s_j\lambda_i\eta_{B\{T\}_\mathcal{D}}(\gamma)) \in I_C$. ■

Lemma 2.6.7 and Lemma 2.6.15 together imply that the s from Lemma 2.6.13 induces a $\mathcal{D}$-structure g^W on $W(C) = W(B\{T\}_\mathcal{D})/I_C$ which makes it an (A, e)-

algebra and $W(\pi_C)$ an (A,e)-algebra homomorphism. We now check it makes η_C into a $\mathcal{D}$-homomorphism by an argument similar to Theorem 3.2 of [39].

Lemma 2.6.16. *The $\mathcal{D}$-structure g^W on $W(C)$ makes η_C into a $\mathcal{D}$-homomorphism.*

Proof. Consider the following cube:

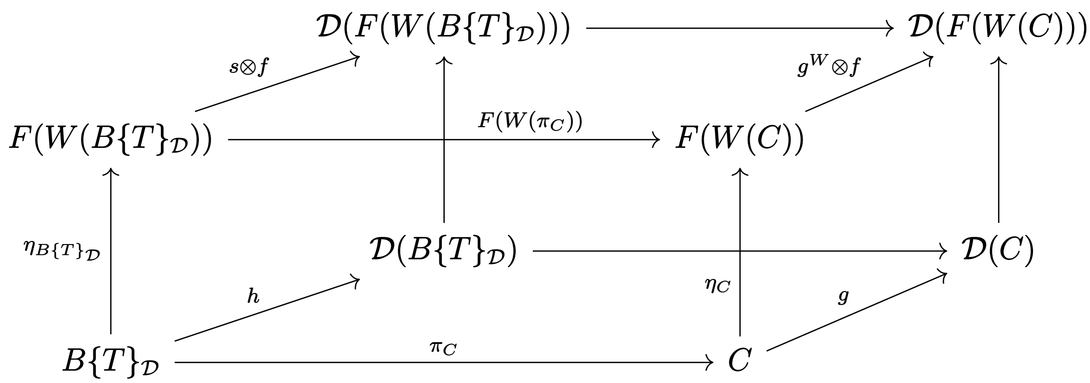

The maps on the back face are just $\mathcal{D}(\phi)$ for ϕ the corresponding map on the front face.

We want to show that the right-hand face of the cube commutes.

1. The front face commutes due to the classical Weil descent.
2. The back face commutes since the front one does: it is just applying the functor $\mathcal{D}$ to the front face.
3. The left face commutes due to choice of s.
4. The bottom face commutes because π_C is a $\mathcal{D}$-homomorphism.
5. Since $W(\pi_C)$ is an (A,e)-algebra homomorphism, $F^\mathcal{D}(W(\pi_C))$ is a (B,f)-algebra homomorphism, and hence the top face commutes.

Since π_C is surjective, the right face of the cube also commutes. ■

Remark 2.6.17. Note that g^W is necessarily the unique $\mathcal{D}$-structure on $W(C)$ making η_C into a $\mathcal{D}$-homomorphism. This is a consequence of Lemma 2.4.4, but it can also be seen from the statement at the end of Lemma 2.6.1.

Therefore, we have provided an explicit way to construct the $\mathcal{D}$-Weil descent $W^\mathcal{D}(C,g) = (W(C), g^W)$.

Chapter 3

The uniform companion for theories of $\mathcal{D}$-fields in characteristic zero

In this chapter, we return to a model-theoretic analysis of $\mathcal{D}$-fields. Recall from the introduction that we aim to construct a $\mathcal{L}_{\text{ring}}(\partial)$-theory, $\text{UC}_{\mathcal{D}}$, such that whenever T is a model complete theory of difference large fields of characteristic zero, then $T \cup \text{UC}_{\mathcal{D}}$ is the model companion of $T \cup$ "$\mathcal{D}$-fields". Here "$\mathcal{D}$-fields" is the $\mathcal{L}_{\text{ring}}(\partial)$-theory of $\mathcal{D}$-fields defined in Section 1.6 together with additional axioms saying that the associated endomorphisms of such a $\mathcal{D}$-field coincide with the endomorphisms of T.[1]

We first take Cousins's definition of difference largeness from [15] and extract its geometric and model-theoretic content. In Section 3.2, we define the axiom scheme $\text{UC}_{\mathcal{D}}$ and show that it has the properties from Theorem A, from which Theorem B immediately follows. In particular, we prove that various theories of $\mathcal{D}$-fields have model companions in Corollary 3.2.8. In Section 3.3, we give alternative characterisations of $\text{UC}_{\mathcal{D}}$ in the case $\mathcal{D}$ is local, and we use the $\mathcal{D}$-Weil descent of the previous chapter to show that the algebraic closure of a model of $\text{UC}_{\mathcal{D}}$ must be a model of $\mathcal{D}$-CF_0. Finally, in Section 3.4, we examine what can be said about the uniform companion in the absence of Assumption **A**: theories of large $\mathcal{D}$-fields can

[1]Recall from Definition 1.6.3 that the associated endomorphisms are uniformly quantifier-free $\mathcal{L}_{\text{ring}}(\partial)$-definable. This means that there is a set of elements $(\gamma_{i,j} \in k\colon i = 1, \dots, t, j = 0, \dots, l)$ such that in every $\mathcal{D}$-field (K, ∂), the ith associated endomorphism is given by $\sigma_i = \sum_{j=0}^{l} \gamma_{i,j}\partial_j$. Then these additional axioms are just $\forall x\ \sigma_i(x) = \sum_{j=0}^{l} \gamma_{i,j}\partial_j(x)$ for every $i = 1, \dots, t$.

only be uniformly companionised when $\mathcal{D}$ is local.

This chapter forms part of the content of the author's [49], currently submitted for publication.

Conventions. For this chapter, all rings are of characteristic zero.

3.1 Difference largeness

Recall our setup: k is a field of characteristic zero, $\mathcal{D}$ is a finite-dimensional k-algebra, there is some k-algebra homomorphism $\pi\colon \mathcal{D} \to k$, the algebra $\mathcal{D}$ has a decomposition as $\prod_{i=0}^{t} B_i$ where each B_i is a local finite-dimensional k-algebra, and, in contrast to the previous chapter, (K, ∂) is a $\mathcal{D}$-ring only when $\partial\colon K \to K \otimes_k \mathcal{D}$ is a section to $\mathrm{id}_K \otimes \pi$. We also impose Assumption **A**: the residue field of each B_i, which is necessarily a finite field extension of k, is k itself. Thus all the associated homomorphisms are actually endomorphisms $\sigma_i\colon K \to K$ for $i = 0, \dots, t$. Recall also from Section 1.6 that, since ∂ must now be a section to $\mathrm{id}_K \otimes \pi$, the associated endomorphism corresponding to π is $\sigma_0 = \mathrm{id}_K \otimes \pi \circ \partial = \mathrm{id}_K$.

Since our uniform companion will be given "relative" to the associated difference field, to simplify notation we will also work with $\mathcal{E}$-operators where

$$\mathcal{E} = k^{t+1}$$

so that $\mathcal{E}$-fields are precisely fields with t endomorphisms that do not necessarily commute.

Recall also that we have k-algebra homomorphisms $\pi_i\colon \mathcal{D} \to k$ given by the composition of the projection to B_i and then the residue map to k. Let $\alpha\colon \mathcal{D} \to \mathcal{E}$ be the product of the maps π_i. Then if (K, ∂) is a $\mathcal{D}$-ring and (K, σ) is its associated difference ring thought of as an $\mathcal{E}$-ring – so that $\sigma\colon K \to K^{t+1}$ is given by $r \mapsto (r, \sigma_1(r), \dots, \sigma_t(r))$ – then $\alpha \circ \partial = \sigma$. By Section 4.1 of [51], α induces a morphism of varieties $\hat{\alpha}\colon \tau_{\mathcal{D}} X \to \tau_{\mathcal{E}} X = X \times X^{\sigma_1} \times \cdots \times X^{\sigma_t}$ such that the following diagram (of nonalgebraic maps) commutes:

$$\begin{array}{ccc} \tau_{\mathcal{D}} X(K) & \xrightarrow{\hat{\alpha}} & \tau_{\mathcal{E}} X(K) \\ & {\scriptstyle \nabla_{\mathcal{D}}} \nwarrow \quad \nearrow {\scriptstyle \nabla_{\mathcal{E}}} & \\ & X(K) & \end{array}$$

Note also that $\hat{\alpha}$ is the product of the morphisms $\hat{\pi}_i$. As mentioned in Fact 1.7.1, if (L,σ) is some $\mathcal{E}$-field extension of (K,σ), then $\nabla_{\mathcal{E}}$ is also a map $X(L) \to \tau_{\mathcal{E}} X(L)$. Since such maps are compatible, we will not distinguish them.

Lemma 3.1.1. *Let (K,∂) be a $\mathcal{D}$-field, and (K,σ) its associated difference field thought of as an $\mathcal{E}$-field. Suppose X and $Y \subseteq \tau_{\mathcal{D}} X$ are irreducible varieties over K. Let b be a K-generic point of Y. Then the following are equivalent:*

(1) *Y has a Zariski-dense set of K-rational points whose projections to $\tau_{\mathcal{E}} X(K)$ are in $\nabla_{\mathcal{E}}(X(K))$;*

(2) *there is some difference field (L,σ) containing the function field $K(b)$ in which $\hat{\alpha}(b)$ is in $\nabla_{\mathcal{E}}(X(L))$ and which is a difference field elementary extension of (K,σ).*

Proof. (1) $\Longrightarrow$ (2). Working with respect to the coordinates in Section 1.7, saying that $\hat{\alpha}(b)$ is in the image of $\nabla_{\mathcal{E}}$ is equivalent to saying that $\sigma_i(\hat{\pi}_0(b)) = \hat{\pi}_i(b)$ for each $i = 1, \dots, t$.

Consider the following set of formulas with parameters from K in the language of difference rings:

$$p(x) = \mathrm{qftp}(b/K) \cup \{\sigma_i(\hat{\pi}_0(x)) = \hat{\pi}_i(x) \colon i = 1, \dots, t\}.$$

Since b is K-generic in Y and Y has a Zariski-dense set of K-rational points c with $\sigma_i(\hat{\pi}_0(c)) = \hat{\pi}_i(c)$, $p(x)$ is finitely satisiable in (K,σ), and hence is a partial type. So there is some difference field $(L,\sigma) \succeq (K,\sigma)$ with a realisation of $p(x)$. This is precisely (2).

(2) $\Longrightarrow$ (1). Suppose (L,σ) is the difference field given by (2). Let $U \subseteq Y$ be any nonempty Zariski-open subset of Y. It is enough to check for U basic and K-Zariski-open. Now b is K-generic in Y and hence is in U. So U contains an L-rational point whose projection under $\hat{\alpha}$ is in the image of $\nabla_{\mathcal{E}}$. That is

$$(L,\sigma) \models \exists x \left(x \in U \wedge \bigwedge_{i=1}^{t} \sigma_i(\hat{\pi}_0(x)) = \hat{\pi}_i(x) \right).$$

Since $(K,\sigma) \preceq (L,\sigma)$, this sentence is true in (K,σ), and hence U contains a K-rational point whose projection to $\tau_{\mathcal{E}} X$ is in the image of $\nabla_{\mathcal{E}}$. ■

Recall from the introduction that we cannot hope to uniformly find the model companion for $\mathcal{D}$-fields whose underlying field is large; we need to take into account the associated difference field. The following definition facilitates this.

Definition 3.1.2. A difference field $(K, \sigma_1, \dots, \sigma_t)$ is difference large if for any pair of K-irreducible varieties V and W such that

(i) $W \subseteq V \times V^{\sigma_1} \times \cdots \times V^{\sigma_t}$,

(ii) the projections $W \to V^{\sigma_i}$ are dominant for all $i = 0, \dots, t$, and

(iii) W has a smooth K-rational point,

then W has a Zariski-dense set of K-rational points of the form $(a, \sigma_1(a), \dots, \sigma_t(a))$ for $a \in V(K)$.

Remark 3.1.3.

1. If (K, σ) is difference large, then K is large. If V is K-irreducible and has a smooth K-rational point, then V is absolutely irreducible, $V \times V^{\sigma_1} \times \cdots \times V^{\sigma_t}$ is absolutely irreducible and has a smooth K-rational point.
2. If $t = 0$, that is, $\mathcal{D}$ is local, then difference largeness is precisely largeness. If $t > 0$, then the only examples of difference large fields known to the author are models of $\mathrm{ACFA}_{0,t}$; hence we will focus on the local case in Section 3.3 and Chapter 5.
3. This notion first appeared (for $t = 1$) in Cousins's thesis [15].

3.2 The uniform companion

In this section we define the axiom scheme $\mathrm{UC}_{\mathcal{D}}$ and show it has the desired properties from Theorem A in the Introduction. We continue to impose Assumption **A**.

Definition 3.2.1. Let (K, ∂) be a $\mathcal{D}$-field. We say that (K, ∂) is a model of $\mathrm{UC}_{\mathcal{D}}$ if for every pair of K-irreducible varieties V and W such that

(i) $W \subseteq \tau V$,

(ii) the projections $\hat{\pi}_i \colon W \to V^{\sigma_i}$ are dominant for each $i = 0, \dots, t$, and

(iii) W has a smooth K-rational point,

then W has a Zariski-dense set of K-rational points of the form $\nabla(a)$ for $a \in V(K)$.

Remark 3.2.2. If $(K,\partial) \models \mathrm{UC}_{\mathcal{D}}$, then the associated difference field (K,σ) is difference large. If X and $Y \subseteq X \times X^{\sigma_1} \times \cdots \times X^{\sigma_t}$ are K-irreducible and Y has a smooth K-rational point, then let $Y' \subseteq \tau X$ be a K-irreducible component of $\hat{\alpha}^{-1}(Y)$ that projects dominantly to Y under $\hat{\alpha}$. Since Y has a smooth K-rational point and $\hat{\alpha}\colon Y' \to Y$ is dominant, so does Y'. Now apply the $\mathrm{UC}_{\mathcal{D}}$ axiom and project the Zariski-dense set of K-rational points of Y' to Y. This idea is similar to Proposition 4.12 of [53].

This axiom scheme can be expressed in a first-order fashion in the language $\mathcal{L}_{\mathrm{ring}}(\partial)$. This is nowadays a standard argument, but we provide some details following the argument used in [65]. We will make use of Theorem 4.2 from there, which collects results about ideals of polynomials from [66].

Theorem 4.2 of [65]. *Let $n, d \in \mathbb{N}$, $x = (x_1, \ldots, x_n)$. Then there are bounds $B = B(n,d)$, $C = C(n,d)$, and $E = E(n,d)$ in $\mathbb{N}$ such that for each field K, each ideal I of $K[x]$ generated by polynomials of degree at most d, and all $f_1, \ldots, f_p \in K[x]$ of degree at most d, the following are true.*

(i) *If I is generated by $f_1, \ldots, f_p$, then for every $g \in I$ of degree at most d, there are $c_1, \ldots, c_p \in K[x]$ of degree at most E such that $g = c_1 f_1 + \cdots + c_p f_p$.*

(ii) *I is prime if and only if $1 \notin I$ and for all $f, g \in K[x]$ of degree at most B, if $fg \in I$, then $f \in I$ or $g \in I$.*

(iii) *For all $m \in \{1, \ldots, n\}$, the ideal $I \cap K[x_1, \ldots, x_m]$ is generated by at most C polynomials of degree at most C.*

Let $n, d, m \in \mathbb{N}$, $x = (x_1, \ldots, x_n)$, $f_1, \ldots, f_m \in K[x]$, $g_1, \ldots, g_m, h \in K[x^0, \ldots, x^l]$, all of degree at most d. The polynomials f_j generate the ideal that defines X, I_X, the polynomials g_j generate the ideal that defines Y, I_Y, and the nonvanishing of the polynomial h will define the open subset U. Then also the polynomials $f_1^{\sigma_i}, \ldots, f_m^{\sigma_i}$ generate the ideal that defines X^{σ_i}. The fact that K-irreducibility of X and Y can be expressed as first-order $\mathcal{L}_{\mathrm{ring}}$-axioms comes from (i) and (ii). That $Y \subseteq \tau X$ can be checked by verifying that the polynomials that define τX are elements of I_Y; that this is a first-order condition follows from (i). Indeed the polynomials that define τX can be computed in terms of the coefficients of the f_i (see the discussion at the end of Section 3 of [53]). That $\hat{\pi}_0\colon Y \to X$ is dominant says that $I_Y \cap K[x] = I_X$. By (iii) there is a bound, depending on n and d, on the number and degree of the polynomials needed to generate $I_Y \cap K[x]$. That this equality is first-order comes from (i). A similar argument shows that dominance of $\hat{\pi}_i \colon Y \to X^{\sigma_i}$ is a first-order condition. The existence of a smooth K-point for Y can be verified by the Jacobian condition on the g_i. Note also that the dimension of Y is the Krull dimension of I_Y, which is definable in terms of the coefficients of the g_i. Finally, that there is a point $\nabla(a)$ in the nonempty Zariski open set U is equivalent to the fact that that the polynomial h either is an element of I_Y, which is first-order by (i), or does not vanish at some $\nabla(a)$.

The fact that the theory $\mathrm{UC}_{\mathcal{D}}$ is the desired uniform companion will follow from the next two theorems (Theorem 3.2.3 and Theorem 3.2.4). These form the $\mathcal{D}$-field analogue of Theorem 6.2 from [65], where the differential counterpart is stated.

Theorem 3.2.3. *Let $M, N \models \mathrm{UC}_{\mathcal{D}}$ be $\mathcal{D}$-fields and A a common $\mathcal{D}$-subring. If M and N have the same existential theory over A as difference fields, then they do as $\mathcal{D}$-fields.*

Proof. Let F_1 and F_2 be the quotient fields of A inside M and N, respectively. Since the associated endomorphisms of A are injective, they extend uniquely to F_1 and F_2. By Lemma 1.6.5, the $\mathcal{D}$-structure on A extends uniquely to F_1 and F_2, and so they are isomorphic as $\mathcal{D}$-fields. Then we may assume $F = F_1 = F_2$ is contained in both M and N. Let L_1 be the relative algebraic closure of F in M and similarly for L_2 in N. Then the associated endomorphisms of M and N restrict to endomorphisms of L_1 and L_2 respectively. Since M and N have the same existential theory over A as difference fields, L_1 and L_2 must be isomorphic as difference fields. The $\mathcal{D}$-field structures on M and N restrict to ones on L_1 and L_2, and by Remark 1.6.6, they must be isomorphic as $\mathcal{D}$-fields; we may then assume $L = L_1 = L_2$ is contained in both M and N.

Suppose that $M \models \exists \bar{x} \phi(\bar{x})$ where $\phi(\bar{x})$ is a quantifier-free $\mathcal{L}_{\mathrm{ring}}(\partial)$-formula with parameters in A and $\bar{x} = (x_1, \dots, x_m)$. As usual, we can assume that ϕ is of the form $\bigwedge_{i=1}^{n} f_i(\bar{x}) = 0$, where the f_i are $\mathcal{L}_{\mathrm{ring}}(\partial)$-terms with coefficients in F. Let c_0 be such that $M \models \phi(c_0)$. Let Ξ be the set of all finite words on $\{\partial_1, \dots, \partial_l\}$. For each r, let Ξ_r be an enumeration of the words of length at most r such that Ξ_r is an initial segment of Ξ_{r+1}, let $n_r = |\Xi_r|$, and let $\nabla_r \colon M \to M^{n_r}$ be given by $b \mapsto (\xi(b) \colon \xi \in \Xi_r)$. Let r be minimal such that $\phi(M) = \{b \in M^m \colon \nabla_r(b) \in Z\}$ where $Z \subseteq M^{mn_r}$ is a Zariski closed set over F. If $r = 0$, then ϕ is in fact an $\mathcal{L}_{\mathrm{ring}}$-formula and since M and N have the same existential theory over A as fields, we have a solution in N. So assume $r > 0$ and let $c = \nabla_{r-1}(c_0)$.

Let $X = \mathrm{loc}(c/L)$, $Y = \mathrm{loc}(\nabla c/L)$. Note that $Y \subseteq \tau X$ and that $\hat{\pi}_i(Y)$ is Zariski dense in X^{σ_i}. Let $g_1, \ldots, g_s$ be polynomials that generate the vanishing ideal of Y over L. By the primitive element theorem, let $\alpha \in L$ be such that $F(\alpha)$ is a field of definition for X and Y. After clearing denominators, we can rewrite the polynomials $g_i(\bar{u})$ instead as $G_i(v, \bar{u}) \in F[v, \bar{u}]$, where $G_i(\alpha, \bar{u}) = g_i(\bar{u})$. Let $\mu(v) \in F[v]$ be the minimal polynomial of α.

We claim that $(\alpha, \nabla c) \in M$ is a smooth solution to the system

$$\mu(v) = 0, G_1(v, \bar{u}) = 0, \ldots, G_s(v, \bar{u}) = 0. \tag{$\blacklozenge$}$$

As ∇c is L-generic in Y, it must be a smooth solution to the system $g_1(\bar{u}) = 0, \ldots, g_s(\bar{u}) = 0$. Let $J(\nabla c)$ be the Jacobian for $g_1, \ldots, g_s$ at ∇c, and let d be its rank. Since $\mu(v)$ contains none of the $\bar{u}$ variables, the Jacobian of the system ($\blacklozenge$) at $(\alpha, \nabla c)$ is of the form

$$\begin{pmatrix} \frac{d\mu}{dv}(\alpha) & 0 \\ \star & J(\nabla c) \end{pmatrix}.$$

Since μ is the minimal polynomial of α, $\frac{d\mu}{dv}(\alpha) \neq 0$, and hence this matrix has rank $d + 1$. Note also that the variety defined by ($\blacklozenge$) is a finite union of conjugates of Y, and hence has the same dimension as Y. So $(\alpha, \nabla c)$ is a smooth point of ($\blacklozenge$).

Consider the quantifier-free $\mathcal{L}_{\mathrm{ring}}(\sigma_1, \ldots, \sigma_t)$-type of $(\alpha, \nabla c) \in M$ over F. Since M and N have the same existential theory over F as difference fields, this partial type is finitely satisfiable in N. We may also assume that N is sufficiently saturated. Then there is a realisation $(\beta, b) \in N$ of this partial type. This induces a difference field F-isomorphism $\theta \colon F\langle\alpha\rangle_\sigma \to F\langle\beta\rangle_\sigma$ where $F\langle\alpha\rangle_\sigma$ means the difference field generated by F and α and likewise for β. We also have that b is a smooth point of Y^θ.

Both $F\langle\alpha\rangle_\sigma$ and $F\langle\beta\rangle_\sigma$ are algebraic extensions of F, and hence by Remark 1.6.6, the $\mathcal{D}$-field structure on F extends uniquely to $\mathcal{D}$-field structures on $F\langle\alpha\rangle_\sigma$ and $F\langle\beta\rangle_\sigma$. So θ is a $\mathcal{D}$-field isomorphism between $\mathcal{D}$-subfields of M and N.

Since Y is L-irreducible, Y^θ is L-irreducible, and since L is relatively algebraically closed in N, Y^θ is N-irreducible. Likewise, X^θ is N-irreducible. Proposition 4.8 of [51] tells us that $\tau X^\theta \simeq \tau X$, and that this isomorphism respects points of the form ∇z. We also have that $Y^\theta \subseteq \tau X^\theta$ and that $\hat{\pi}_i(Y^\theta)$ is Zariski dense in $(X^\theta)^{\sigma_i}$. Since θ fixes F it also fixes Z.

Since $N \models \mathrm{UC}_{\mathcal{D}}$, there is a point $a \in X^\theta(N)$ with $\nabla a \in Y^\theta(N)$. Let a_0 be the first m coordinates of a. We claim that a_0 is a realisation of ϕ, which will conclude the proof. As in the proof of Theorem 4.5 of [53], we prove first that $\nabla_{r-1}(a_0) = a$. Write $a = (a_\xi : \xi \in \Xi_{r-1})$. We prove by induction on the length of ξ that $\xi(a_0) = a_\xi$. For $\xi = \mathrm{id}$ this is clear. Suppose now that $\xi = \partial_i \xi'$. Since $\nabla_{r-1}(c_0) = c$, we have that $\partial_i c_{\xi'} = c_\xi$. This is an algebraic fact about ∇c over F. Since ∇a satisfies all the algebraic relations ∇c does over F (since θ fixes F), we also have $\partial_i a_{\xi'} = a_\xi$. By the inductive hypothesis, $\partial_i a_{\xi'} = \partial_i \xi'(a_0) = \xi(a_0)$.

Since c_0 realises ϕ, $\nabla_r(c_0) \in Z$. This is an algebraic fact about $\nabla c = \nabla\nabla_{r-1}(c_0)$ over F. Since ∇a satisfies all the algebraic relations ∇c does over F, we have $\nabla_r(a_0) \in Z$. So $N \models \phi(a_0)$. ∎

Theorem 3.2.4. *Every $\mathcal{D}$-field whose associated difference field is difference large has a $\mathcal{D}$-field extension which is a model of* $\mathrm{UC}_{\mathcal{D}}$ *and an elementary extension at the level of difference fields.*

Proof. Let (F, ∂) be a $\mathcal{D}$-field that is difference large as a difference field, and let X and Y be F-irreducible varieties where $Y \subseteq \tau X$, $\hat{\pi}_i(Y)$ is Zariski dense in X^{σ_i}, and Y has a smooth F-rational point. Let U be a nonempty Zariski-open subset of Y defined over F.

Let $b \in Y(L)$ be an F-generic point in some field extension L of F. Since $\hat{\pi}_i(Y)$ is Zariski dense in X^{σ_i}, we get that $\hat{\pi}_i(b)$ is F-generic in X^{σ_i}. Let $\hat{\alpha}\colon \tau X \to X \times X^{\sigma_1} \times \cdots \times X^{\sigma_t}$ be the product of the morphisms $\hat{\pi}_i$. Let Z be the Zariski-closure of $\hat{\alpha}(Y)$ in $X \times X^{\sigma_1} \times \cdots \times X^{\sigma_t}$. Then $\hat{\alpha}\colon Y \to Z$ is dominant. Let V be the F-open subset of smooth points of Z. By dominance, V has a point in the image of $\hat{\alpha}$. Then $\hat{\alpha}^{-1}(V)$ is a nonempty F-open set. Since F is large and Y has a smooth F-rational point, $\hat{\alpha}^{-1}(V)$ has an F-rational point. So V has an F-rational point – that is, Z has an F-rational smooth point.

Let $W \subseteq Y$ be some open subset of Y. Since Y is irreducible, W is dense in Y. Then $\hat{\alpha}(W)$ is dense in Z. As Z has a smooth F-rational point and F is difference large, Z has a Zariski dense set of F-rational points of the form $(a, \sigma_i(a), \dots, \sigma_t(a))$. So Y has a Zariski-dense set of F-rational points whose projections under $\hat{\alpha}$ have the form $(a, \sigma_i(a), \dots, \sigma_t(a))$.

So by Lemma 3.1.1 there is some difference field (K, σ) containing $F(b)$ which is an elementary extension of (F, σ) and in which $\sigma_i(\hat{\pi}_0(b)) = \hat{\pi}_i(b)$. We will define a $\mathcal{D}$-field structure on K whose associated difference field is (K, σ).

As mentioned in Fact 1.7.1, there is an identification $\tau X(K) \leftrightarrow X(K \otimes_k \mathcal{D})$. Let

b' be the tuple from $K \otimes_k \mathcal{D}$ that corresponds to $b \in \tau X(K)$ under this identification. Note that $\pi_i(b') = \sigma_i(\pi_0(b'))$ in K because $\hat{\pi}_i(b) = \sigma_i(\hat{\pi}_0(b))$. Write $a = \hat{\pi}_0(b)$ for the F-generic point of X.

As in the proof of Theorem 4.5 of [53], we can extend $\partial \colon F \to F \otimes_k \mathcal{D} \subseteq K \otimes_k \mathcal{D}$ to a k-algebra homomorphism $\partial \colon F[a] \to K \otimes_k \mathcal{D}$ with $\partial(a) = b'$.

Indeed, since $b \in \tau X(K)$, we have $p^{\partial}(b') = 0$ for all $p \in I(X/F)$. Since a is F-generic in X, $I(X/F) = I(a/F)$. As $p^{\partial}(b') = 0$ for all $p \in I(a/F)$, ∂ extends to $F[a] = F[x]/I(a/F)$ by setting $\partial(a) = b'$. Since $\hat{\pi}_0(b) = a$, we have that $\pi_0^K(b') = a$ so $\pi_0^K \circ \partial \colon F[a] \to K$ is inclusion. We also have that $\pi_i \circ \partial(a) = \pi_i(b') = \sigma_i(\pi_0(b')) = \sigma_i(a)$. So $\pi_i \circ \partial = \sigma_i$. Now we can use Lemma 1.6.5 to extend ∂ to a $\mathcal{D}$-field structure on K extending the one on F whose associated endomorphisms are precisely the σ_i. In (K, ∂) we will also have $\nabla(a) = b$. Since b is F-generic in Y, we must have that $b = \nabla a \in U(K)$.

Then we can iterate this construction transfinitely to get an extension of F that is elementary as an extension of difference fields, which is also a model of $\mathrm{UC}_{\mathcal{D}}$. ■

Having established Theorems 3.2.3 and 3.2.4, the following result is proved in precisely the same way as in Proposition 6.3 of [65].

Proposition 3.2.5. *The $\mathcal{L}_{\mathrm{ring}}(\partial)$-theory $\mathrm{UC}_{\mathcal{D}}$ is inductive. If U is an $\mathcal{L}_{\mathrm{ring}}(\partial)$-theory of difference large $\mathcal{D}$-fields satisfying the properties in the previous two theorems (Theorems 3.2.3 and 3.2.4), then U contains $\mathrm{UC}_{\mathcal{D}}$. If U is in addition inductive, $U = \mathrm{UC}_{\mathcal{D}} \cup$ "difference large fields", where containment and equality here are as deductively closed sets of sentences.*

Proof. It is clear that the union of an increasing chain of models of $\mathrm{UC}_{\mathcal{D}}$ is also a model of $\mathrm{UC}_{\mathcal{D}}$, and likewise for difference large fields. Hence the theories $\mathrm{UC}_{\mathcal{D}}$ and "difference large fields" are inductive.

Let U be another theory of difference large $\mathcal{D}$-fields satisfying Theorems 3.2.3 and 3.2.4. Let $M \models U$. Since the associated difference field of M is difference large, by Theorem 3.2.4, it embeds as a $\mathcal{D}$-field in some $N \models \mathrm{UC}_{\mathcal{D}}$ such that the extension of difference fields is elementary. By this last fact, N is also difference large, and hence embeds in some $M' \models U$, since U satisfies the property in Theorem 3.2.4. In addition, the associated difference field of M' is an elementary extension of that of N. Thus $M \leq M'$ are two models of U whose associated difference fields have the same existential theory over M. By the property in Theorem 3.2.3 for U, M and M'

have the same existential theory over M as $\mathcal{D}$-fields; that is, M is existentially closed in M', and hence in N. Since $N \models \mathrm{UC}_{\mathcal{D}}$ and $\mathrm{UC}_{\mathcal{D}}$ is inductive, we have $M \models \mathrm{UC}_{\mathcal{D}}$. So $\mathrm{UC}_{\mathcal{D}} \subseteq U$.

If U is also inductive, then repeat the above proof with U and $\mathrm{UC}_{\mathcal{D}}\cup$"difference large fields" exchanged to get $U = \mathrm{UC}_{\mathcal{D}} \cup$ "difference large fields". ■

The following theorem is also proved in like its differential counterpart: Theorem 7.1 of [65].

Theorem 3.2.6. *Let C be a set of new constant symbols and T a model complete theory of difference large fields in the language $\mathcal{L}_{\mathrm{ring}}(C)(\sigma)$. Let T^* be an expansion by definitions of T to a language $\mathcal{L}^* \supseteq \mathcal{L}_{\mathrm{ring}}(C)(\sigma)$. In addition, let A be an $\mathcal{L}^*(\partial)$-structure such that when A is viewed as an $\mathcal{L}_{\mathrm{ring}}(\partial)$-structure it is a $\mathcal{D}$-field.*

If $T^* \cup \mathrm{diag}(A|_{\mathcal{L}^*})$ is complete, then $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$ is complete.

Proof. We first prove a series of claims.

Claim 1. $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$ is consistent.

Proof of claim. $T^* \cup \mathrm{diag}(A|^*_{\mathcal{L}})$ is consistent so $A|^*_{\mathcal{L}}$ is a substructure of a model of T^*. Since the model of T^* has a difference structure, use Lemma 1.6.5 to extend the $\mathcal{D}$-field structure on A to one on the model of T^*. Then by Theorem 3.2.4 we get the claim.

Claim 2. $T^* \cup \mathrm{UC}_{\mathcal{D}}$ is model complete.

Proof of claim. By the previous claim, $T \cup \mathrm{UC}_{\mathcal{D}}$ is consistent. Consider any extension $M \leq N$ of models of $T \cup \mathrm{UC}_{\mathcal{D}}$. By model completeness of T, M is existentially closed in N as difference fields. By Theorem 3.2.3, M is existentially closed in N as $\mathcal{D}$-fields. So $T \cup \mathrm{UC}_{\mathcal{D}}$ is model complete. As T^* is an expansion by definitions of T, $T^* \cup \mathrm{UC}_{\mathcal{D}}$ is model complete.

Claim 3. If T^* has quantifier elimination, then so does $T^* \cup \mathrm{UC}_{\mathcal{D}}$.

Proof of claim. Let $M \models T^* \cup \mathrm{UC}_{\mathcal{D}}$ and let A be an $\mathcal{L}^*(\partial)$-substructure. We will show that $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$ is complete. So let $N \models T^* \cup \mathrm{UC}_{\mathcal{D}}$ containing A as an $\mathcal{L}^*(\partial)$-substructure. As T^* has quantifier elimination, we know that $M \equiv_A N$ as $\mathcal{L}^*$-structures, and so they have the same existential theory over A as difference fields. As they are both models of $\mathrm{UC}_{\mathcal{D}}$, Theorem 3.2.3 tells us that they have the same existential theory over A as $\mathcal{D}$-fields. Now any $\mathcal{L}^*(\partial)$-formula is equivalent modulo $T^* \cup \mathrm{UC}_{\mathcal{D}}$ to an $\mathcal{L}_{\mathrm{ring}}(C)(\partial)$-formula, and to an existential $\mathcal{L}_{\mathrm{ring}}(C)(\partial)$-formula since $T \cup \mathrm{UC}_{\mathcal{D}}$ is model complete. Then M and N have the same $\mathcal{L}^*(\partial)$-theory over A, and $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$ is complete.

Now we are able to prove the theorem. Let M and N be two models of $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$. We need to show that M and N are elementarily equivalent as $\mathcal{L}^*(\partial)(A)$-structures, that is, that $M \equiv_A N$.

Let $\mathcal{L}^+$ be the language extending $\mathcal{L}^*$ with new k-ary predicate symbols R_ϕ for each existential $\mathcal{L}^*$-formula ϕ in k free variables. Let T^+ be T^* with the sentences $\forall u_1, \dots, u_k (R_\phi(\bar{u}) \leftrightarrow \phi(\bar{u}))$. T^+ is an expansion by definitions of T^* and hence of T. As T is model complete, so is T^+. T^+ proves that every formula is equivalent to an existential one, and that every existential one is equivalent to a quantifier-free one. So it has quantifier elimination. By the previous claim, $T^+ \cup \mathrm{UC}_{\mathcal{D}}$ has quantifier elimination.

Let $\tilde{M}$ and $\tilde{N}$ be the unique expansions of M and N to models of $T^+ \cup \mathrm{UC}_{\mathcal{D}}$. Since $T^* \cup \mathrm{diag}(A|_{\mathcal{L}^*})$ is complete, $M \restriction_{\mathcal{L}^*} \equiv_A N \restriction_{\mathcal{L}^*}$. By definition of T^+, $\tilde{M}$ and $\tilde{N}$ induce the same structure on A. By quantifier elimination of $T^+ \cup \mathrm{UC}_{\mathcal{D}}$, $\tilde{M} \equiv_A \tilde{N}$, and so also $M \equiv_A N$. ■

In the next theorem we argue why Theorem 3.2.6 justifies calling $\mathrm{UC}_{\mathcal{D}}$ the uniform companion for theories of difference large $\mathcal{D}$-fields. This is the precise formulation of Theorem B from the introduction. Recall that "$\mathcal{D}$-fields" is the $\mathcal{L}_{\mathrm{ring}}(\partial)$-theory consisting of the axioms for a $\mathcal{D}$-field together with additional axioms saying that the associated endomorphisms of such a $\mathcal{D}$-field coincide with the endomorphisms of the $\mathcal{L}_{\mathrm{ring}}(\sigma_1, \dots, \sigma_t)$-theory T.

Theorem 3.2.7. *Let C be a set of new constant symbols and T a model complete theory of difference large fields in the language $\mathcal{L}_{\mathrm{ring}}(C)(\sigma)$. Let T^* be an $\mathcal{L}^*$-theory which is an expansion by definitions of T.*

Assume T^* is a model companion of an $\mathcal{L}^*$-theory T_0^* which extends the theory of difference fields. Then:

(i) *$T^* \cup \mathrm{UC}_{\mathcal{D}}$ is a model companion of the $\mathcal{L}^*(\partial)$-theory $T_0^* \cup$ "$\mathcal{D}$-fields";*

(ii) *if T^* is a model completion of T_0^*, then $T^* \cup \mathrm{UC}_{\mathcal{D}}$ is a model completion of $T_0^* \cup$ "$\mathcal{D}$-fields";*

(iii) *if T^* has quantifier elimination, then $T^* \cup \mathrm{UC}_{\mathcal{D}}$ has quantifier elimination;*

(iv) *if T is complete and M is a $\mathcal{D}$-field which is a model of T, then $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(\mathcal{C})$ is complete, where $\mathcal{C}$ is the $\mathcal{L}_{\mathrm{ring}}(C)(\partial)$-substructure generated by $\emptyset$ in M, that is, the $\mathcal{D}$-subring of M generated by the elements $(c^M)_{c \in C}$.*

Proof. First note that $T^* \cup \mathrm{UC}_{\mathcal{D}}$ and $T_0^* \cup$ "$\mathcal{D}$-fields" have the same universal theory (equivalently, that a model of one can be embedded in a model of the other). Let $M \models T^* \cup \mathrm{UC}_{\mathcal{D}}$. Since T^* and T_0^* have the same universal theory, there is an $\mathcal{L}^*$-structure N such that $M \restriction_{\mathcal{L}^*} \leq N$. By Lemma 1.6.5, we can extend the $\mathcal{D}$-structure on M to one on N, so that $M \leq N$ as $\mathcal{L}^*(\partial)$-structures. For the other direction, let $M \models T_0^* \cup$ "$\mathcal{D}$-fields". Then there is some $\mathcal{L}^*$-structure $N \models T^*$ such that $M \restriction_{\mathcal{L}^*} \leq N$. Use Lemma 1.6.5 to extend the $\mathcal{D}$-structure on M to one on N, so that $N \models T^* \cup$ "$\mathcal{D}$-fields" and then use Theorem 3.2.4 to embed this in a model of $T^* \cup \mathrm{UC}_{\mathcal{D}}$.

Now to show (i), it is enough to show that $T^* \cup \mathrm{UC}_{\mathcal{D}}$ is model complete, or equivalently, that if $M \models T^* \cup \mathrm{UC}_{\mathcal{D}}$, then $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(M)$ is complete. Since T^* is model complete, $T^* \cup \mathrm{diag}(M|_{\mathcal{L}^*})$ is complete. Then $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(M)$ is complete by Theorem 3.2.6.

For (ii), let M be a model of $T_0^* \cup$ "$\mathcal{D}$-fields". We need to show that $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(M)$ is complete. But $M \models T_0^*$, and so $T^* \cup \mathrm{diag}(M|_{\mathcal{L}^*})$ is complete since T^* is a model completion of T_0^*. Then apply Theorem 3.2.6.

For (iii), let $M \models T^* \cup \mathrm{UC}_{\mathcal{D}}$, and let $A \leq M$ be an $\mathcal{L}^*(\partial)$-substructure. We need to show that $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(A)$ is complete. By quantifier elimination for T^*, we have that $T^* \cup \mathrm{diag}(A|_{\mathcal{L}^*})$ is complete; the result follows by Theorem 3.2.6.

For (iv), since T is complete, $T^* \cup \mathrm{diag}(\mathcal{C}|_{\mathcal{L}^*}) \subseteq T^*$ is complete. By Theorem 3.2.6, $T^* \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(\mathcal{C})$ is complete. ■

We can now collect the consequences of these theorems, similarly to the differential set-up in Section 8 of [65].

Corollary 3.2.8. (1) $\mathrm{ACFA}_{0,t} \cup \mathrm{UC}_{\mathcal{D}}$ *is* $\mathcal{D}$-CF_0 *from* [53].

(2) *If $\mathcal{D}$ is local, then* $\mathrm{RCF} \cup \mathrm{UC}_{\mathcal{D}}$ *is complete and is the model companion of the theory of real closed $\mathcal{D}$-fields. It admits quantifier elimination in* $\mathcal{L}_{\mathrm{ring}}(\leq)(\partial)$.

(3) *Suppose $\mathcal{D}$ is local. The theory* $p\mathrm{CF}_d$ *of p-adically closed fields of fixed p-rank d has quantifier elimination in the language* $\mathcal{L}_{\mathrm{ring}}(\mathcal{O}, c_1, \ldots, c_d, (P_n)_{n\in\mathbb{N}})$, *where $\mathcal{O}$ is a predicate for the valuation ring, $c_1, \ldots, c_n$ are constants that form a $\mathbb{Z}/p$-basis for $\mathcal{O}/p$, and P_n is a predicate for the nth powers – see* [42] *for $d = 1$ and* [59] *for any finite d. Then* $p\mathrm{CF}_d \cup \mathrm{UC}_{\mathcal{D}}$ *is complete and is the model companion of p-adically closed $\mathcal{D}$-fields of fixed rank d. It has quantifier elimination in* $\mathcal{L}_{\mathrm{ring}}(\mathcal{O}, c_1, \ldots, c_d, (P_n)_{n\in\mathbb{N}})(\partial)$.

(4) *Suppose $\mathcal{D}$ is local. Let* $\mathrm{Psf}_0(C)$ *be the* $\mathcal{L}_{\mathrm{ring}}(C)$*-theory of pseudofinite fields of characteristic zero with sentences saying that the polynomial* $x^n + c_{n,1}x^{n-1} + \cdots + c_{n,n}$ *is irreducible for each* $n > 1$. *Then* $\mathrm{Psf}_0(C)$ *is model complete – this is Proposition 2.7 of* [12]. *We then get that* $\mathrm{Psf}_0(C) \cup \mathrm{UC}_{\mathcal{D}}$ *is the model companion of* $\mathrm{Psf}_0(C) \cup$ *"$\mathcal{D}$-fields".*

3.3 Alternative characterisations of the uniform companion

In this section we will describe some additional characterisations of $\mathrm{UC}_{\mathcal{D}}$ in the case $\mathcal{D}$ is local. One in particular will use the notion of a D-variety, and will allow us to show that an algebraic extension of a large field which is a model of $\mathrm{UC}_{\mathcal{D}}$ is also a model of $\mathrm{UC}_{\mathcal{D}}$. In particular, the algebraic closure of such a $\mathcal{D}$-field will be a model of $\mathcal{D}$-CF_0 from [53]. For this section, we impose the following.

Assumption B. The k-algebra $\mathcal{D}$ is local.

Example 3.3.1. The algebras in (1), (2), and (5) from Example 1.6.2 are local. We can combine these algebras using fibred products and tensor products to form more local examples. See Examples 3.4 and 3.5 of [53].

Since $\mathcal{D}$ is local, any $\mathcal{D}$-ring R has only one associated homomorphism: the identity id_R; the associated difference ring is then just the underlying ring. Hence for any affine K-variety X, there is only one induced morphism $\tau X \to X$, which we call $\hat{\pi}$. With respect to the coordinates described in Section 1.7, this is just the morphism induced by the inclusion $K[x]/I \to K[x^0, \dots, x^l]/I'$ where $x \mapsto x^0$.

Definition 3.3.2. Let (K, ∂) be a $\mathcal{D}$-ring. A D-variety over (K, ∂) is a pair (V, s) where V is a variety over K and $s\colon V \to \tau V$ is an algebraic morphism over K which is a section to the canonical projection $\hat{\pi}\colon \tau V \to V$. We say that (V, s) is K-irreducible if V is K-irreducible, affine if V is affine, etc.

Given a $\mathcal{D}$-field extension (L, δ) of (K, ∂), the (L, δ)-rational sharp points of (V, s) are defined as $(V, s)^\sharp(L, \delta) = \{a \in V(L)\colon \nabla a = s(a)\}$.

As before, we will mainly be interested in affine D-varieties. If V is an affine variety, a D-variety structure on V is equivalent to a $\mathcal{D}$-ring structure on its coordinate ring, $K[V]$. A K-rational sharp point is equivalent to a $\mathcal{D}$-ring homomorphism

$K[V] \to K$. This is the natural $\mathcal{D}$-field analogue of D-varieties as defined for differential rings; see for example [36].

We now establish some basic results about D-varieties. Recall that, if (R, ∂) is a $\mathcal{D}$-ring and $\mathfrak{a}$ is an ideal of R, $\mathfrak{a}$ is called a $\mathcal{D}$-ideal if $\partial(\mathfrak{a}) \subseteq \mathfrak{a} \otimes_k \mathcal{D}$, or equivalently, if $\partial_i(\mathfrak{a}) \subseteq \mathfrak{a}$ for each $i = 1, \ldots, l$.

Lemma 3.3.3. *Let (R, ∂) be a $\mathcal{D}$-ring, and $\mathfrak{a} \subseteq R$ a radical $\mathcal{D}$-ideal. Then the minimal prime ideals above $\mathfrak{a}$ are $\mathcal{D}$-ideals.*

Proof. Let $\mathfrak{p}$ be a minimal prime ideal above $\mathfrak{a}$, and consider the localisation $R_\mathfrak{p}$. Since $\mathfrak{a}$ is radical, so is $\mathfrak{a}R_\mathfrak{p}$ (see Proposition 3.11 of [3]). Suppose $\mathfrak{q} \subseteq \mathfrak{p}R_\mathfrak{p}$ is a prime ideal of $R_\mathfrak{p}$ that also lies above $\mathfrak{a}R_\mathfrak{p}$. Then by part iv) of the same proposition, we must have $\mathfrak{q} = \mathfrak{p}R_\mathfrak{p}$, and hence $\mathfrak{p}R_\mathfrak{p}$ is a minimal prime above $\mathfrak{a}R_\mathfrak{p}$. It is also the unique maximal ideal of $R_\mathfrak{p}$, and hence is the only prime ideal lying above $\mathfrak{a}R_\mathfrak{p}$. Then $\mathfrak{a}R_\mathfrak{p} = \sqrt{\mathfrak{a}R_\mathfrak{p}} = \mathfrak{p}R_\mathfrak{p}$ (the radical of an ideal is the intersection of the prime ideals lying above it).

By Remark 1.6.6 we know that ∂ extends uniquely to a $\mathcal{D}$-structure on $R_\mathfrak{p}$ with $\partial(\frac{a}{b}) = \partial(a)\partial(b)^{-1}$. Since $\mathfrak{a}$ is a $\mathcal{D}$-ideal it is clear that $\mathfrak{a}R_\mathfrak{p}$ is also a $\mathcal{D}$-ideal.

Then $\mathfrak{p}R_\mathfrak{p}$ is a $\mathcal{D}$-ideal, and hence its contraction to R, $\mathfrak{p}$, is also a $\mathcal{D}$-ideal. ■

Lemma 3.3.4. *Let (V, s) be an affine D-variety over (K, ∂). Then*

a) *any nonempty Zariski-open $U \subseteq V$ defined over K is a D-subvariety of (V, s);*

b) *any K-irreducible component of V is a D-subvariety of (V, s).*

Proof. a) Let $K[V]$ be the coordinate ring of V. Then s corresponds to $\partial_s \colon K[V] \to K[V] \otimes_k \mathcal{D}$. Let U be a basic open subset of V given by the nonvanishing of some f. By Remark 1.6.6 we then get that ∂_s extends uniquely to $K[V]_f \to K[V]_f \otimes_k \mathcal{D}$. That is, s restricts to $U \to \tau U$. Now if $U = \bigcup_{i \in I} U_i$ is a union of basic open subsets, s restricts to $U_i \to \tau U_i \subseteq \tau U$, and these restrictions agree on $U_i \cap U_j$ since this is also a basic open. Glueing the morphisms $U_i \to \tau U$ gives a morphism $U \to \tau U$ which is a restriction of s.

b) by Lemma 3.3.3. ■

Theorem 3.3.5. *Suppose (K, ∂) is a $\mathcal{D}$-field and K is large. Then the following are equivalent:*

(1) $(K, \partial) \models \mathrm{UC}_\mathcal{D}$;

(2) *whenever (V, s) is an affine, K-irreducible D-variety, if V has a smooth K-rational point, then the set of K-rational sharp points of (V, s) is Zariski dense in V;*

(3) *whenever (V, s) is an affine, K-irreducible D-variety, if V has a smooth K-rational point, then (V, s) has a K-rational sharp point;*

(4) *whenever (V, s) is a smooth, affine, K-irreducible D-variety, if V has a K-rational point, then (V, s) has a K-rational sharp point; and*

(5) *whenever (L, δ) is a $\mathcal{D}$-field extension of (K, ∂) such that K is existentially closed in L as a field, then (K, ∂) is existentially closed in (L, δ) as a $\mathcal{D}$-field.*

Proof. (1) $\implies$ (2): Suppose $(K, \partial) \models \mathrm{UC}_{\mathcal{D}}$ and let (V, s) be a K-irreducible D-variety with a smooth K-rational point. Let $X = V$ and $Y = s(V)$. Note that X and Y are isomorphic. Then Y has a smooth K-rational point, $Y \subseteq \tau X$, and $\hat{\pi}\colon Y \to X$ is an isomorphism. So, since $K \models \mathrm{UC}_{\mathcal{D}}$, Y has a Zariski dense set of K-rational points of the form $\nabla(a)$ for $a \in X(K)$, and hence for each such $a \in X(K)$, $\nabla(a) = s(a)$.

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

(3) $\implies$ (1): Let X, Y, U be as in the statement of $\mathrm{UC}_{\mathcal{D}}$. Let $b \in L \geq K$ be a K-generic point of Y, so that $a = \hat{\pi}(b) \in L$ is K-generic in X by dominance. Since $b \in \tau X(K(b))$, let $b' \in K(b) \otimes_k \mathcal{D}$ be the point corresponding to b under the identification $\tau X(K(b)) \leftrightarrow X(K(b) \otimes_k \mathcal{D})$. Then $P^{\partial}(b') = 0$ for all $P \in I(X/K)$, and so ∂ extends to a homomorphism $\partial\colon K[a] \to K(b) \otimes_k \mathcal{D}$ with $\partial(a) = b'$. Extend this to a $\mathcal{D}$-ring structure $\partial\colon K(b) \to K(b) \otimes_k \mathcal{D}$ using Lemma 1.6.5. In this $\mathcal{D}$-ring structure, $\nabla(a) = b$. Now each $\partial_i(b_j) \in K(b)$ so $\partial_i(b_j) = \frac{P_{ij}(b)}{Q_{ij}(b)}$ for some polynomials $P_{ij}, Q_{ij} \in K[x]$. Let $Q \in K[x]$ be the product of all Q_{ij}. Note that ∂ restricts to $K[b] \to K[b]_{Q(b)} \otimes_k \mathcal{D}$. Again by Lemma 1.6.5, we must have that ∂ extends to $K[b]_{Q(b)} \to K[b]_{Q(b)} \otimes_k \mathcal{D}$. Let U' be the open subset of Y corresponding to $Q(x)$. This extension of ∂ gives a D-variety structure $s\colon U' \to \tau U'$.

Since K is large and V has a smooth K-point, $U \cap U'$ has a smooth K-point. By Lemma 3.3.4, $(U \cap U', s|_{U \cap U'})$ is a K-irreducible D-variety with a smooth K-rational point. By (3) there is $(c, d_1, \ldots, d_l) \in (U \cap U')(K)$ with $\nabla(c, d_1, \ldots, d_l) = s(c, d_1, \ldots, d_l)$. Then $c \in X(K)$ with $\nabla(c) = (c, d_1, \ldots, d_l) \in U(K)$.

(3) $\implies$ (4) is clear.

(4) $\implies$ (3): Let (V, s) be a D-variety over K with V K-irreducible and $a \in V(K)$ a smooth K-rational point. Let $W \subseteq V$ be the smooth locus of V. Then W is a

smooth, K-irreducible D-subvariety of V. The point a is a K-rational point of W and so by (4), W has a K-rational sharp point. Then V has a K-rational sharp point.

(1) $\implies$ (5): Let (L,δ) be a $\mathcal{D}$-field extension of $(K,\partial) \models \mathrm{UC}_{\mathcal{D}}$ such that K is existentially closed in L as a field. Then there is a field extension $L \leq M$ such that M is an elementary extension of K as a field; note that M is then also a large field. Extend the $\mathcal{D}$-field structure on L to one on M, and use Theorem 3.2.4 to find a $\mathcal{D}$-field extension $(N,d) \models \mathrm{UC}_{\mathcal{D}}$ such that $K \prec M \prec N$ as fields. This last fact implies that K and N have the same existential theory as fields over K. So by Theorem 3.2.3, they have the same existential theory as $\mathcal{D}$-fields over (K,∂) – recall that since $\mathcal{D}$ is local, the associated difference field is just the underlying field. Then (K,∂) is existentially closed in (N,d), and hence in (L,δ).

(5) $\implies$ (1): Assume (K,∂) has property (5). By Theorem 3.2.4, there is $(L,\delta) \models \mathrm{UC}_{\mathcal{D}}$ extending it such that $K \prec L$ as fields. Then K is existentially closed in L as fields, and so (K,∂) is existentially closed in (L,δ) as $\mathcal{D}$-fields by (5). Since $\mathrm{UC}_{\mathcal{D}}$ is inductive, we must also have $(K,\partial) \models \mathrm{UC}_{\mathcal{D}}$. ■

We will now show that algebraic extensions of large models of $\mathrm{UC}_{\mathcal{D}}$ are again large and models of $\mathrm{UC}_{\mathcal{D}}$. Similar to the differential case (Theorem 5.11 of [38]), this will rely on the $\mathcal{D}$-Weil descent, established in Chapter 2.

We recall only the necessary properties of the $\mathcal{D}$-Weil descent in their geometric form. Let $(L,\delta)/(K,\partial)$ be an extension of $\mathcal{D}$-fields where L/K is a finite field extension. Let (V,s) be an affine D-variety over (L,δ); as mentioned above, this is equivalent to a $\mathcal{D}$-ring structure, δ^s, on the coordinate ring, $L[V]$, extending δ on L. The classical Weil descent of V, V^W, is a K-variety such that there is a natural bijection

$$V(L) \leftrightarrow V^W(K).$$

Stated algebraically, this is equivalent to the natural bijection

$$\mathrm{Hom}_L(L[V],L) \leftrightarrow \mathrm{Hom}_K(K[V^W],K).$$

In the previous chapter we showed that there is a unique $\mathcal{D}$-ring structure, ∂^s, on $K[V^W]$ extending ∂ on K such that the above natural bijection restricts to a natural bijection

$$\mathrm{Hom}_{(L,\delta)}((L[V],\delta^s),(L,\delta)) \leftrightarrow \mathrm{Hom}_{(K,\partial)}((K[V^W],\partial^s),(K,\partial)).$$

The $\mathcal{D}$-ring structure ∂^s corresponds to $s^W : V^W \to \tau(V^W)$ and makes (V^W, s^W) into a D-variety over (K,∂). As mentioned above, a $\mathcal{D}$-ring homomorphism $L[V] \to L$ corresponds to an L-rational sharp point of (V,s). Geometrically then, we have that the first natural bijection restricts to the natural bijection

$$(V,s)^\sharp(L,\delta) \leftrightarrow (V^W,s^W)^\sharp(K,\partial).$$

Theorem 3.3.6. *Let $(L,\delta)/(K,\partial)$ be an algebraic extension of $\mathcal{D}$-fields where (K,∂) is a model of* $\mathrm{UC}_{\mathcal{D}}$ *and K is a large field. Then (L,δ) is a model of* $\mathrm{UC}_{\mathcal{D}}$ *and L is large.*

Proof. Consider first the case when L/K is a finite extension. We verify condition (4) of Theorem 3.3.5. Let (V,s) be a smooth, L-irreducible D-variety defined over (L,δ) with an L-rational point. Now apply the $\mathcal{D}$-Weil descent to get a D-variety (V^W,s^W) over (K,∂). Since V is affine and smooth, V^W is affine and smooth (see Proposition 5 of Section 7.6 of [6]). By the bijection $V(L) \leftrightarrow V^W(K)$, V^W has a K-rational point. Let (U,t) be the irreducible component of (V^W,s^W) containing the K-rational point. Since (K,∂) satisfies condition (4), (U,t) has a K-rational sharp point, and hence (V^W,s^W) has a K-rational sharp point. By the bijection $(V,s)^\sharp(L,\delta) \leftrightarrow (V^W,s^W)^\sharp(K,\partial)$, (V,s) has an L-rational sharp point.

If L/K is algebraic, let F be an intermediate extension such that V, s, and the L-rational point are all defined over F and F/K is finite. Then by the above, $(V,s)^\sharp(F,\delta) \neq \emptyset$, and hence $(V,s)^\sharp(L,\delta) \neq \emptyset$. ■

3.4 The non-local case

Recall that throughout this chapter we assumed that either the k-algebra $\mathcal{D}$ was a local ring or each component in its local decomposition had residue field k. In this section we make some observations about the existence of model companions of $\mathcal{D}$-fields in the case when neither assumption holds. Without Assumption **A** the associated homomorphisms of a $\mathcal{D}$-field are not necessarily *endo*morphisms, and hence it does not make sense to ask whether $T \cup$ "$\mathcal{D}$-fields" has a model companion when T is a theory of difference fields. However, it does make sense to ask the question as T varies over theories of fields. The main result of this section says that when the base field k is finitely generated over $\mathbb{Q}$, we get a full characterisation of when the uniform companion for large $\mathcal{D}$-fields exists: it exists if and only if $\mathcal{D}$ is local.

We start with the general case: k is a field of characteristic zero, $\mathcal{D}$ is a finite-dimensional k-algebra, and $\mathcal{D} = \prod_{i=0}^{t} B_i$ where each B_i is a local finite-dimensional k-algebra. We no longer impose Assumption **B** – that $\mathcal{D}$ is local – or even Assumption **A** – that the residue field of each B_i is k. For $i > 0$, the residue field of B_i is $k[x]/(P_i)$ for some k-irreducible polynomial P_i, and that of B_0 is k. For an $\mathcal{L}$-theory T, the $\mathcal{L}(\partial)$-theory $T \cup$ "$\mathcal{D}$-fields" is denoted by $T_{\mathcal{D}}$, and the $\mathcal{L}(\sigma)$-theory $T \cup$ "σ is an endomorphism" is denoted by T_σ.

A result of Kikyo and Shelah [33] states that if a model complete theory has the strict order property, then the theory obtained by adding an automorphism has no model companion. In particular, if $\mathcal{D} = k \times k$, $\mathcal{D}$-fields correspond to fields with an endomorphism, and so $\mathrm{RCF}_{\mathcal{D}} = \mathrm{RCF}_\sigma$ and $\mathrm{Th}(\mathbb{Q}_p)_{\mathcal{D}} = \mathrm{Th}(\mathbb{Q}_p)_\sigma$ have no model companion. In fact, the Kikyo–Shelah theorem implies that $T_{\mathcal{D}}$ has no model companion when $\mathcal{D}$ is not local and T has a model in which one of the polynomials P_i has a root. We first prove this for the case when $\mathcal{D}$ has at least one local component with residue field k, and then reduce the more general statement to this case.

Proposition 3.4.1. *Assume $\mathcal{D}$ is such that one of the local components B_i has residue field k for $i > 0$. If $T_{\mathcal{D}}$ has a model companion, then T_{σ_i} has a model companion.*

Proof. Note that by a particular choice of the basis $\varepsilon_0, \ldots, \varepsilon_l$, we may assume that the associated endomorphism σ_i corresponding to B_i appears as one of the operators ∂_j. So $\mathcal{L}(\sigma_i) \subseteq \mathcal{L}(\partial)$.

Write T^+ for the model companion of $T_{\mathcal{D}}$ and T^- for its reduct to $\mathcal{L}(\sigma_i)$. We will show that T^- is the model companion of T_{σ_i}; clearly their universal parts coincide, so it suffices to prove T^- is model complete.

Let $(K, \sigma) \models T^-$. We will show that $T^- \cup \mathrm{diag}_{\mathcal{L}(\sigma_i)}(K)$ is complete. Use Lemma 1.6.5 to equip K with a $\mathcal{D}$-ring structure whose ith associated homomorphism is σ_i and whose jth associated homomorphism is inclusion $K \to K[x]/(P_j)$ for $j \neq i$. Then $K \models T_{\mathcal{D}}$, and it embeds in some $L \models T^+$. Since T^+ is model complete, $T^+ \cup \mathrm{diag}_{\mathcal{L}(\partial)}(L)$ is complete, and hence its reduct to $\mathcal{L}(\sigma_i)(K)$, $T^- \cup \mathrm{diag}_{\mathcal{L}(\sigma_i)}(K)$, is complete. ■

We now weaken the assumption that the residue field of some B_i is k to the assumption that T has a model K in which one of the polynomials P_i has a root. If $T_{\mathcal{D}}$ has a model companion, then $T_{\mathcal{D}} \cup \text{diag}(K)$ has a model companion. Let $\mathcal{E}$ be the K-algebra $\mathcal{D} \otimes_k K$. As mentioned in the proof of Theorem 3.2 of [4], if L is an $\mathcal{E}$-field, then $\mathcal{E}$-field extensions of L coincide with $\mathcal{D}$-field extensions of L. Hence if $T_{\mathcal{D}} \cup \text{diag}(K)$ has a model companion, so does $T_{\mathcal{E}}$. But $\mathcal{E}$ now satisfies the assumption in Proposition 3.4.1. So we have proved the following.

Corollary 3.4.2. *If T is model complete and has a model with the strict order property in which one of the P_i has a root, then $T_{\mathcal{D}}$ has no model companion.*

Real closed fields and $\mathbb{Q}_p$ have the strict order property, and so this result means if any P_i has a root in some real closed field or some p-adically closed field, there is no uniform companion. In particular, if the base field k is a finitely generated field extension of $\mathbb{Q}$, we get a full converse to the main theorem.

Corollary 3.4.3. *Suppose k is a finitely generated field extension of $\mathbb{Q}$. Then there is a uniform companion for theories of large $\mathcal{D}$-fields if and only if $\mathcal{D}$ is a local ring.*

Proof. If $\mathcal{D}$ is local, $\mathcal{D}$-fields whose associated difference field is difference large correspond precisely to $\mathcal{D}$-fields whose underlying field is large. The uniform companion then exists by Section 3.2.

Suppose $\mathcal{D}$ is not local. Then the splitting field of the polynomial $P_1 \in k[x]$ is a finitely generated extension of $\mathbb{Q}$, and so by Theorem 1 of [10], embeds in some $\mathbb{Q}_p$. Hence P_1 has a root in $\mathbb{Q}_p$. Then by Corollary 3.4.2, $\text{Th}(\mathbb{Q}_p)_{\mathcal{D}}$ has no model companion. ■

Remark 3.4.4. The base field k does have an impact on when the uniform companion exists. If k is algebraically closed, then the only model complete theory of fields containing k is ACF_0, and hence the existence of a uniform companion for $\mathcal{D}$-fields is equivalent to the existence of the model companion of $\text{ACF}_0 \cup$ "$\mathcal{D}$-fields"; this exists for all $\mathcal{D}$ by Theorem 7.6 of [53].

However, for other fields k the question is still open. For instance, suppose $k = \mathbb{R}$. No model of $\text{Th}(\mathbb{Q}_p)$ can be an $\mathbb{R}$-algebra, and so $\text{Th}(\mathbb{Q}_p) \cup$ "$\mathcal{D}$-fields" is inconsistent. Hence the above method does not show that there is no uniform companion in the case $\mathcal{D} = \mathbb{R} \times \mathbb{C}$, for instance.

Chapter 4

Derivation-like theories and neostability

4.1 Introduction

We saw in the previous chapter that if T is a model complete theory of difference large fields, then $T \cup$ "$\mathcal{D}$-fields" has a model companion: $T \cup \mathrm{UC}_{\mathcal{D}}$. In this chapter we will investigate model-theoretic properties transferred from the theory of difference fields to the theory of $\mathcal{D}$-fields. For instance it is immediate from the transfer of quantifier elimination from an expanion by definitions T^* to $T^* \cup \mathrm{UC}_{\mathcal{D}}$ in Theorem 3.2.7 that NIP also transfers.

Corollary 4.1.1. *Let C be a set of new constant symbols, and suppose that T is the complete, model complete $\mathcal{L}_{\mathrm{ring}}(C)(\sigma)$-theory of a difference large field of characteristic zero. If T is NIP, so is (any completion of) $T \cup \mathrm{UC}_{\mathcal{D}}$.*

Proof. Let T^* be an expansion by definitions of T with quantifier elimination where the $\mathcal{L}^*$-terms are the same as the $\mathcal{L}_{\mathrm{ring}}(C)(\sigma)$-terms (for instance, if T^* is the Morleyisation of T). Let $\mathfrak{C}$ be a monster model of $T^* \cup \mathrm{UC}_{\mathcal{D}}$ whose reduct to $\mathcal{L}^*$ is a monster model of T^*.

Suppose $\phi(x, y)$ is an $\mathcal{L}^*(\partial)$-formula with IP: so there are $(a_i)_{i\in\omega}$, $(b_I)_{I\subseteq\omega}$ in $\mathfrak{C}$ with $\mathfrak{C} \models \phi(a_i, b_I) \iff i \in I$. By Theorem 3.2.7, $T^* \cup \mathrm{UC}_{\mathcal{D}}$ has quantifier elimination, and we may assume $\phi(x, y)$ is quantifier-free. Now, since the $\mathcal{L}^*$-terms are the same as the $\mathcal{L}_{\mathrm{ring}}(C)$-terms, the $\mathcal{L}^*(\partial)$-terms in the variables x, y are then just polynomials in x, y, C, and any application of the operators to these. So there are $r \in \mathbb{N}$ and a quantifier-free $\mathcal{L}^*$-formula ϕ^* such that $\mathfrak{C} \models \phi(x, y)$ if and only if $\mathfrak{C} \models \phi^*(\nabla_r(x), \nabla_r(y))$. Then

$$\mathfrak{C} \models \phi^*(\nabla_r(a_i), \nabla_r(b_I)) \iff \mathfrak{C} \models \phi(a_i, b_I) \iff i \in I.$$

Therefore the sequences $(\nabla_r(a_i))_{i \in \omega}$ and $(\nabla_r(b_I))_{I \subseteq \omega}$ witness that ϕ^* has IP. ■

Remark 4.1.2.

1. This result generalises the fact of Michaux and Rivière that CODF is NIP from Theorem 2.2 of [47].
2. In Corollary 4.3 of [22], Guzy and Point show that NIP is transferred from a topological field (possibly with extra structure) to the model companion of the field with a derivation. The imposition of a topological structure allows them to consider fields with genuine extra structure, as opposed to the definitional expansions considered here.

A similar argument to Corollary 4.1.1 shows that stability transfers via its characterisation of no formula having the order property. But in the case when $\mathcal{D}$ is local, stability yields something stronger.

Lemma 4.1.3. *Suppose $\mathcal{D}$ is local and that T is the complete, model complete $\mathcal{L}_{\text{ring}}(C)$-theory of a large field of characteristic zero. If T is stable, then $T \cup \mathrm{UC}_{\mathcal{D}} = \mathcal{D}\text{-}\mathrm{CF}_0$.*

Proof. A stable, large field of characteristic zero is algebraically closed by Theorem D of [30]. The result then follows as $\mathrm{ACF}_0 \cup \mathrm{UC}_{\mathcal{D}} = \mathcal{D}\text{-}\mathrm{CF}_0$. ■

We now turn our sights to simplicity. Like NIP and stability, simplicity has a combinatorial characterisation.

Definition 4.1.4. A formula $\phi(x, y)$ has the tree property if there is a tree of parameters $(a_s : s \in \omega^{<\omega})$ and some $k \geq 2$ such that

1. for every $\eta \in \omega^{<\omega}$ the set $\{\phi(x, a_{\eta i}) : i \in \omega\}$ is k-inconsistent; and
2. for every $\sigma \in \omega^\omega$, the set $\{\phi(x, a_{\sigma|_n}) : n \in \omega\}$ is consistent.

A complete theory T is simple if no formula has the tree property.

However, it is not clear to the author whether the previous proof will adapt to the tree property. Instead, we tackle the problem using the more semantic characterisation of simplicity given in the preliminaries. Recall from Section 1.1 the following definition.

Definition 1.1.23. Let T be a complete theory and $\mathfrak{C}$ a monster model. A relation $\mathop{\smile\!\!\!\!\!\mid}^*$ on triples of small subsets of $\mathfrak{C}$ is called an abstract independence relation if it is invariant under automorphisms and satisfies the following conditions.

1. *normality:* $X \mathop{\smile\!\!\!\!\!\mid}^*_A B \implies X \mathop{\smile\!\!\!\!\!\mid}^*_A AB$;
2. *monotonicity:* $X \mathop{\smile\!\!\!\!\!\mid}^*_A B \implies X \mathop{\smile\!\!\!\!\!\mid}^*_A B'$ for $B' \subseteq B$;
3. *base monotonicity:* $X \mathop{\smile\!\!\!\!\!\mid}^*_A D \implies X \mathop{\smile\!\!\!\!\!\mid}^*_B D$ for $A \subseteq B \subseteq D$;
4. *transitivity:* $X \mathop{\smile\!\!\!\!\!\mid}^*_A B$ and $X \mathop{\smile\!\!\!\!\!\mid}^*_B D \implies X \mathop{\smile\!\!\!\!\!\mid}^*_A D$ for $A \subseteq B \subseteq D$;
5. *symmetry:* $X \mathop{\smile\!\!\!\!\!\mid}^*_A B \iff B \mathop{\smile\!\!\!\!\!\mid}^*_A X$;
6. *full existence:* for any X, A, B there is $X' \equiv_A X$ such that $X' \mathop{\smile\!\!\!\!\!\mid}^*_A B$ (recall $X' \equiv_A X$ means that X' and X have the same type over A);
7. *finite character:* if $X_0 \mathop{\smile\!\!\!\!\!\mid}^*_A B$ for all finite $X_0 \subseteq X$, then $X \mathop{\smile\!\!\!\!\!\mid}^*_A B$;
8. *local character:* there is a cardinal κ such that for all X and A, there is $A_0 \subseteq A$ with $|A_0| < \kappa$ such that $X \mathop{\smile\!\!\!\!\!\mid}^*_{A_0} A$.

There are three extra properties that an abstract independence relation $\mathop{\smile\!\!\!\!\!\mid}^*$ might satisfy that we are interested in:

9. *strictness:* if $b \mathop{\smile\!\!\!\!\!\mid}^*_A b$, then $b \in \operatorname{acl}(A)$;
10. *independence theorem over M:* if $A_1 \mathop{\smile\!\!\!\!\!\mid}^*_M A_2$, $a_1 \mathop{\smile\!\!\!\!\!\mid}^*_M A_1$, $a_2 \mathop{\smile\!\!\!\!\!\mid}^*_M A_2$, and $a_1 \equiv_M a_2$, then there is $a \models \operatorname{tp}(a_1/MA_1) \cup \operatorname{tp}(a_2/MA_2)$ with $a \mathop{\smile\!\!\!\!\!\mid}^*_M A_1A_2$.
11. *stationarity over M:* whenever $A \supseteq M$, $a, b \in \mathfrak{C}$ with $a \equiv_M b$, $a \mathop{\smile\!\!\!\!\!\mid}^*_M A$ and $b \mathop{\smile\!\!\!\!\!\mid}^*_M A$, then $a \equiv_A b$.

T is simple if and only if there is an abstract independence relation on $\mathfrak{C}$ satisfying 1–8 and the independence theorem over models. Thus, we will use this notion of independence for T to define one for $T \cup \mathrm{UC}_{\mathcal{D}}$ and show it has all the required properties.

It will turn out that the transfer of these properties is not specific to the setting of $\mathcal{D}$-fields. Hence, we will work in more generality.

4.2 Derivation-like theories

Let $\mathfrak{L} \subseteq \mathfrak{D}$ be two languages, T a complete and model complete $\mathfrak{L}$-theory, and Δ an inductive $\mathfrak{D}$-theory. Let $\mathbb{U}$ be a monster model for T, and let $\mathop{\smile\!\!\!\!\!\mid}^0$ be some relation on triples of small subsets of $\mathbb{U}$.

Example. The reader should have in mind the following typical example.

- the language of fields as $\mathfrak{L}$;
- the language of differential fields as $\mathfrak{D}$;
- the theory of some large, model complete field as T;
- the theory of differential fields as Δ; and
- algebraic independence or linear disjointness as $\mathop{\smile\!\!\!\!\!\!\mid}^0$.

Remark 4.2.1. Suppose T^* is an expansion by definitions of T. If $\mathop{\smile\!\!\!\!\!\!\mid}^0$ satisfies any of 1–11 in T, then it also does in T^*. This is because acl is the same taken in T and T^* and equality of T-types implies equality of T^*-types.

We write $M \leq_{\mathfrak{L}} N$ if M is an $\mathfrak{L}$-substructure of N and $M \leq_{\mathfrak{D}} N$ is M is a $\mathfrak{D}$-substructure of N. For $A \subseteq M$ where M is an $\mathfrak{L}$-structure, we write $\langle A \rangle_{\mathfrak{L}}$ for the $\mathfrak{L}$-substructure generated by A inside M. By acl_T we mean the model-theoretic algebraic closure in the sense of T.

We say that Δ is derivation-like (with respect to T and $\mathop{\smile\!\!\!\!\!\!\mid}^0$) if the following conditions hold:

(a) if $M \models \Delta$ and $M \leq_{\mathfrak{L}} N \models T$, then there is a $\mathfrak{D}$-structure on N extending the one on M such that $N \models \Delta$;

(b) if $M \models T \cup \Delta$ and $A \leq_{\mathfrak{D}} M \leq_{\mathfrak{L}} \mathbb{U}$, then $\mathrm{acl}_T(A) \leq_{\mathfrak{D}} M$ and $\mathrm{acl}_T(A) \models \Delta$; moreover, this is the only $\mathfrak{D}$-structure on $\mathrm{acl}_T(A)$ extending the one on A that makes $\mathrm{acl}_T(A)$ into a model of Δ;

(c) if $M \models T \cup \Delta$ with $M \leq_{\mathfrak{L}} \mathbb{U}$ and A and B are two models of Δ which are $\mathfrak{D}$-substructures of M with a common acl_T-closed $\mathfrak{D}$-substructure C such that $A \mathop{\smile\!\!\!\!\!\!\mid}^0_C B$, then $\langle AB \rangle_{\mathfrak{L}} \leq_{\mathfrak{D}} M$ and $\langle AB \rangle_{\mathfrak{L}} \models \Delta$; moreover, this is the only $\mathfrak{D}$-structure on $\langle AB \rangle_{\mathfrak{L}}$ extending the ones on A and B and making it into a model of Δ; and

(d) if A and B are two models of Δ which are $\mathfrak{L}$-substructures of $\mathbb{U}$ with a common acl_T-closed $\mathfrak{D}$-substructure C such that $A \mathop{\smile\!\!\!\!\!\!\mid}^0_C B$, then there is a $\mathfrak{D}$-structure on $\langle AB \rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} \mathbb{U}$ extending the ones on A and B that makes $\langle AB \rangle_{\mathfrak{L}}$ into a model of Δ.

Remark 4.2.2. Suppose Δ is derivation-like with respect to T. If T^* is the Morleyisation of T, then Δ is derivation-like with respect to T^*.

4.3 Transferring neostability properties

For the rest this section, we assume that Δ is derivation-like with respect to T and $\mathop{\smile\!\!\!\!\!\!|}^0$ and that $T \cup \Delta$ has a model companion, T^+. Using Fact 1.1.16, we may find a $\mathfrak{D}$-structure $\mathfrak{C}$ such that $\mathfrak{C}$ is a monster model for T^+ and $\mathfrak{C}|_{\mathfrak{L}}$ is a monster model for T. In this section we will prove that many neostability properties transfer from T to T^+. Recall that acl_T refers to model-theoretic algebraic closure in the sense of T; acl refers to algebraic closure in the sense of T^+ and tp to types in the sense of T^+.

Theorem 4.3.1. *Suppose $\mathop{\smile\!\!\!\!\!\!|}^0$ satisfies full existence. Let T^* be an expansion by definitions of T to a language $\mathfrak{L}^*$. Let $K \models T^* \cup T^+$ and $A \leq_{\mathfrak{L}^*(\mathfrak{D})} K$. If $T^* \cup \mathrm{diag}_{\mathfrak{L}^*}(A)$ is complete, then $T^* \cup T^+ \cup \mathrm{diag}_{\mathfrak{L}^*(\mathfrak{D})}(A)$ is complete.*

Proof. We first note that $T^* \cup T^+$ is model complete. Given any extension of models, it must be a $\mathfrak{D}$-elementary extension since both are models of T^+, and hence a $\mathfrak{L}^*(\mathfrak{D})$-elementary extension since any symbol from $\mathfrak{L}^*$ can be defined with an $\mathfrak{L}$-formula.

Let $K, L \models T^* \cup T^+$ and let A be a common $\mathfrak{L}^*(\mathfrak{D})$-substructure. Both K and L are models of T^* and A is a common $\mathfrak{L}^*$-substructure. Since $T^* \cup \mathrm{diag}_{\mathfrak{L}^*}(A)$ is complete, the bijection $A \to A$ is a partial $\mathfrak{L}^*$-elementary map from K to L. This map then extends to a partial $\mathfrak{L}^*$-elementary bijection $\mathrm{acl}^K_{T^*}(A) \to \mathrm{acl}^L_{T^*}(A)$; see Lemma 5.6.4 of [64]. This is an $\mathfrak{L}^*$-isomorphism. By property (b), $\mathrm{acl}^K_{T^*}(A)$ and $\mathrm{acl}^L_{T^*}(A)$ are $\mathfrak{D}$-substructures of K and L, respectively, and are both models of Δ. Pushing the first $\mathfrak{D}$-structure through the $\mathfrak{L}^*$-isomorphism, the moreover clause of (b) tells us that it must also be a $\mathfrak{D}$-isomorphism. So we may assume that A is relatively $\mathfrak{L}^*$-algebraically closed in K and L.

By completeness of $T^* \cup \mathrm{diag}_{\mathfrak{L}^*}(A)$, we may think of K and L both as $\mathfrak{L}^*$-substructures of $\mathfrak{C}|_{\mathfrak{L}}$. Now use full existence of $\mathop{\smile\!\!\!\!\!\!|}^0$ to replace L by a copy with $K \mathop{\smile\!\!\!\!\!\!|}^0_A L$ inside $\mathfrak{C}|_{\mathfrak{L}}$. Let $M \preceq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$ be some $\mathfrak{L}$-elementary substructure containing both K and L. Since A is relatively algebraically closed in K and L, we can use (d) to amalgamate the $\mathfrak{D}$-structures on K and L to one on $\langle KL \rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} M$ making it a model of Δ. Then use (a) to extend this to a $\mathfrak{D}$-structure on M so that $M \models \Delta$. Since T^+ is the model companion of $T \cup \Delta$, extend $M \models T \cup \Delta$ to some $N \models T^+$ which then uniquely expands to a model of $T^* \cup T^+$. Since $T^* \cup T^+$ is model complete, $K \preceq N \succeq L$, and so $K \equiv_A L$. ■

As in Theorem 3.2.7, we collect some immediate consequences.

Corollary 4.3.2. *Suppose $\mathop{\smile\!\!\!\!\!\!\!\mid}^0$ satisfies full existence. Let T^* be an expansion by definitions of T which is the model companion of some inductive T_0^* in the language $\mathfrak{L}^*$. Then*

1. *$T^* \cup T^+$ is the model companion of $T_0^* \cup \Delta$;*
2. *if T^* is the model completion of T_0^*, then $T^* \cup T^+$ is the model completion of $T_0^* \cup \Delta$; and*
3. *if T^* has quantifier elimination, then $T^* \cup T^+$ has quantifier elimination.*

Proof. First we show that $T^* \cup T^+$ and $T_0^* \cup \Delta$ have the same universal theory. Let $M \models T_0^* \cup \Delta$. Since T^* is the model companion of T_0^*, there is some $N \models T^*$ containing M as an $\mathfrak{L}^*$-substructure. By (a), N expands to a model of Δ, and since T^+ is the model companion, there is some $N' \models T^+$ extending N. Now $N' \models T$ as well and hence uniquely expands to a model of T^*. Since $N \leq_{\mathfrak{L}} N'$ are models of the model complete theory T, this extension is $\mathfrak{L}$-elementary, and hence must be an $\mathfrak{L}^*$-extension. So every model of $T_0^* \cup \Delta$ embeds in a model of $T^* \cup T^+$. The converse is similar.

Now the statements follow immediately from Theorem 4.3.1. ■

Proposition 4.3.3. *Suppose $\mathop{\smile\!\!\!\!\!\!\!\mid}^0$ is invariant and satisfies full existence, monotonicity, and strictness. For any $A \subseteq \mathfrak{C}$, we have $\operatorname{acl}(A) = \operatorname{acl}_T(\langle A \rangle_{\mathfrak{D}})$.*

Proof. Let $F = \operatorname{acl}_T(\langle A \rangle_{\mathfrak{D}})$. By (b), F is a $\mathfrak{D}$-substructure of $\mathfrak{C}$. Clearly $F \subseteq \operatorname{acl}(A)$. For the converse, suppose $x \notin F$. We will show that $\operatorname{tp}(x/F)$ is not algebraic.

Suppose for a contradiction that $\operatorname{tp}(x/F)$ has only finitely many realisations. Let K be some small $\mathfrak{D}$-elementary substructure of $\mathfrak{C}$ containing F and all the realisations of $\operatorname{tp}(x/F)$. Now use full existence of $\mathop{\smile\!\!\!\!\!\!\!\mid}^0$ to find $L \equiv_F K$ as $\mathfrak{L}$-structures with $L \mathop{\smile\!\!\!\!\!\!\!\mid}^0_F K$ inside $\mathfrak{C}|_{\mathfrak{L}}$. Let M be some $\mathfrak{L}$-elementary substructure of $\mathfrak{C}|_{\mathfrak{L}}$ containing K and L. The partial $\mathfrak{L}$-elementary map $\alpha \colon K \to L$ fixing F is an $\mathfrak{L}$-isomorphism. Use this $\mathfrak{L}$-isomorphism to define an isomorphic $\mathfrak{D}$-structure on L. Now F is acl_T-closed, and so use (d) to amalgamate the $\mathfrak{D}$-structures on K and L to one on $\langle KL \rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} M$ making it a model of Δ. By (a), M expands to a model of Δ. Since T^+ is the model companion of $T \cup \Delta$, there is some $N \models T^+$ extending M. Now both N and $\mathfrak{C}|_{\mathfrak{L}}$ are models of the complete $\mathfrak{L}$-theory $T \cup \operatorname{diag}_{\mathfrak{L}}(M)$ – completeness is by model completeness of T. So we may embed N inside $\mathfrak{C}|_{\mathfrak{L}}$ over M, and thus, without loss of generality, $M \leq_{\mathfrak{D}} N \leq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$.

By model completeness of T^+, there is a $\mathfrak{D}$-elementary embedding $j\colon N \to \mathfrak{C}$ that fixes K. Now

$$\mathrm{tp}^{\mathfrak{C}}(x/F) = \mathrm{tp}^{K}(x/F) = \mathrm{tp}^{L}(\alpha(x)/F) = \mathrm{tp}^{N}(\alpha(x)/F) = \mathrm{tp}^{\mathfrak{C}}(j\alpha(x)/F)$$

and hence $j\alpha(x)$ is a realisation of $\mathrm{tp}(x/F)$. Since K contained all such realisations, we have $j\alpha(x) \in K$. We also have $j(L) \mathop{\downarrow}^0_F K$ by invariance. Then monotonicity gives $j\alpha(x) \mathop{\downarrow}^0_F j\alpha(x)$, and strictness gives $j\alpha(x) \in \mathrm{acl}_T(F) = F$. But now j and α both fixed F, so we must have had $x \in F$, a contradiction. ■

Theorem 4.3.4. *Define the following relation on triples of small subsets of* $\mathfrak{C}$*:*

$$A \mathop{\downarrow}^+_C B \iff \mathrm{acl}(AC) \mathop{\downarrow}^0_{\mathrm{acl}(C)} \mathrm{acl}(BC).$$

Then

(i) *if* $\mathop{\downarrow}^0$ *is an abstract independence relation, so is* $\mathop{\downarrow}^+$*;*

(ii) *if* $\mathop{\downarrow}^0$ *is a strict independence relation, so is* $\mathop{\downarrow}^+$*;*

(iii) *for some parameter set* M*, if* $\mathop{\downarrow}^0$ *is an independence relation that satisfies the independence theorem over* M*, so is* $\mathop{\downarrow}^+$*; and*

(iv) *for some parameter set* M*, if* $\mathop{\downarrow}^0$ *is an independence relation that satisfies stationarity over some* M*, so is* $\mathop{\downarrow}^+$*.*

Proof. For (i) and (ii), *invariance, normality, monotonicity, transitivity, symmetry, finite character, local character, and strictness* are either by definition or follow from the corresponding property of $\mathop{\downarrow}^0$. For *base monotonicity*, suppose $A \mathop{\downarrow}^+_C B$ and $C \subseteq D \subseteq B$. We may also assume that $A \supseteq C$ by normality. Then $\mathrm{acl}(A) \mathop{\downarrow}^0_{\mathrm{acl}(C)} \mathrm{acl}(B)$. By monotonicity, we have $\mathrm{acl}(A) \mathop{\downarrow}^0_{\mathrm{acl}(C)} \mathrm{acl}(D)$. By (c), $\langle \mathrm{acl}(A)\,\mathrm{acl}(D)\rangle_{\mathfrak{L}}$ is a $\mathfrak{D}$-substructure. So $\langle AD\rangle_{\mathfrak{D}} \subseteq \langle \mathrm{acl}(A)\,\mathrm{acl}(D)\rangle_{\mathfrak{L}}$, and so $\mathrm{acl}(AD) \subseteq \mathrm{acl}_T(\mathrm{acl}(A)\,\mathrm{acl}(D))$. By base monotonicity and normality for $\mathop{\downarrow}^0$, we get $\mathrm{acl}(A)\,\mathrm{acl}(D) \mathop{\downarrow}^0_{\mathrm{acl}(D)} \mathrm{acl}(B)$. By full existence, we get $\mathrm{acl}_T(\mathrm{acl}(A)\,\mathrm{acl}(D)) \mathop{\downarrow}^0_{\mathrm{acl}(A)\,\mathrm{acl}(D)} \mathrm{acl}(B)$, and by transitivity and monotonicity, $\mathrm{acl}(AD) \mathop{\downarrow}^0_{\mathrm{acl}(D)} \mathrm{acl}(B)$. That is, $A \mathop{\downarrow}^+_D B$.

Full existence. Suppose a, A, B are given inside $\mathfrak{C}$. Let M be a small $\mathfrak{D}$-elementary substructure of $\mathfrak{C}$ containing these, and let $C = \mathrm{acl}(A)$. Use full existence for $\mathop{\downarrow}^0$ to find $M' \mathop{\downarrow}^0_C M$ with $M' \equiv^{\mathfrak{L}}_C M$ inside $\mathfrak{C}|_{\mathfrak{L}}$. Let N be some $\mathfrak{L}$-elementary substructure of $\mathfrak{C}|_{\mathfrak{L}}$ containing M and M'. Use the $\mathfrak{L}$-isomorphism $\alpha\colon M \to M'$ that fixes C to put a $\mathfrak{D}$-structure on M'. Using (d), the $\mathfrak{D}$-structures on M and M' then amalgamate to one on $\langle MM'\rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} N$ which makes it into a model of Δ. Then by (a), N expands to a model of Δ, which then embeds inside some $N' \models T^+$. As T is model complete, $T \cup \mathrm{diag}_{\mathfrak{L}}(N)$ is complete, and both N' and $\mathfrak{C}|_{\mathfrak{L}}$ are models of it. So we may embed N' inside $\mathfrak{C}|_{\mathfrak{L}}$ over N and thus assume that $N \leq_{\mathfrak{D}} N' \preceq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$. And now since $N' \models T^+$, we may $\mathfrak{D}$-elementarily embed N' inside $\mathfrak{C}$ over M via $j\colon N' \to \mathfrak{C}$. By invariance, we get $j(M') \mathop{\smile\!\!\!\!|}^0_C M$ and by monotonicity $\mathrm{acl}(Aj\alpha(a)) \mathop{\smile\!\!\!\!|}^0_C \mathrm{acl}(AB)$, that is, $j\alpha(a) \mathop{\smile\!\!\!\!|}^+_A B$. In addition,

$$\mathrm{tp}^{\mathfrak{C}}(a/C) = \mathrm{tp}^{M}(a/C) = \mathrm{tp}^{M'}(\alpha(a)/C) = \mathrm{tp}^{N'}(\alpha(a)/C) = \mathrm{tp}^{\mathfrak{C}}(j\alpha(a)/C).$$

Independence theorem. Let $M \models T^+$, $A_1 \mathop{\smile\!\!\!\!|}^+_M A_2$, $a_1 \mathop{\smile\!\!\!\!|}^+_M A_1$, $a_2 \mathop{\smile\!\!\!\!|}^+_M A_2$, and $\mathrm{tp}(a_1/M) = \mathrm{tp}(a_2/M)$. We will show that there is $a \mathop{\smile\!\!\!\!|}^+_M A_1A_2$ realising $\mathrm{tp}(a_1/A_1) \cup \mathrm{tp}(a_2/A_2)$. Let $N \models T^+$ be some elementary substructure of $\mathfrak{C}$ containing all of the above subsets.

Claim 1. We may assume that A_1, A_2, a_1, and a_2 are all models of T^+ containing M.

Proof of claim. By Löwenheim–Skolem, find an elementary substructure $\bar{A}_1' \prec N$ containing A_1 and M. By full existence, find $\bar{A}_1 \equiv_{MA_1} \bar{A}_1'$ with $\bar{A}_1 \mathop{\smile\!\!\!\!|}^+_{MA_1} A_2$. Then find an elementary $\bar{A}_2' \prec N$ containing A_2 and M, and by full existence $\bar{A}_2 \equiv_{MA_2} \bar{A}_2'$ with $\bar{A}_2 \mathop{\smile\!\!\!\!|}^+_{MA_2} \bar{A}_1$. Then

$$\bar{A}_1 \mathop{\smile\!\!\!\!|}^+_{MA_1} A_2 \text{ and } A_1 \mathop{\smile\!\!\!\!|}^+_{M} A_2 \implies \bar{A}_1 \mathop{\smile\!\!\!\!|}^+_{M} A_2 \text{ by transitivity}$$
$$\bar{A}_1 \mathop{\smile\!\!\!\!|}^+_{M} A_2 \text{ and } \bar{A}_2 \mathop{\smile\!\!\!\!|}^+_{MA_2} \bar{A}_1 \implies \bar{A}_1 \mathop{\smile\!\!\!\!|}^+_{M} \bar{A}_2$$

Do the same with a_1 and a_2. Löwenheim–Skolem constructs the models $\bar{a}_i'$ by closing Ma_i under Skolem functions. The elementary map $Ma_1 \mapsto Ma_2$ extends to the closures of Ma_i under Skolem functions, and so we will have $\bar{a}_1' \equiv_M \bar{a}_2'$ and hence $\bar{a}_1 \equiv_M \bar{a}_2$. So we may assume that A_1, A_2, a_1, a_2 are all models of T^+ containing M.

Claim 2. There is some $a \in \mathfrak{C}|_{\mathfrak{L}}$ with $a \mathop{\smile\!\!\!\!|}^0_M N$ with $a \models \mathrm{tp}_{\mathfrak{L}}(a_1/A_1) \cup \mathrm{tp}_{\mathfrak{L}}(a_2/A_2)$.
Proof of claim. By the independence theorem for $\mathop{\smile\!\!\!\!|}^0$, there is $a \in \mathfrak{C}|_{\mathfrak{L}}$ with $a \mathop{\smile\!\!\!\!|}^0_M A_1A_2$ and $a \models \mathrm{tp}_{\mathfrak{L}}(a_1/A_1) \cup \mathrm{tp}_{\mathfrak{L}}(a_2/A_2)$. Now by full existence for $\mathop{\smile\!\!\!\!|}^0$, we can find $a' \equiv^{\mathfrak{L}}_{A_1A_2} a$ such that $a' \mathop{\smile\!\!\!\!|}^0_{A_1A_2} N$. Let α be the $\mathfrak{L}$-automorphism of N' sending $a \mapsto a'$ and fixing A_1A_2. By invariance, $a' \mathop{\smile\!\!\!\!|}^0_M A_1A_2$, and by transitivity, $a' \mathop{\smile\!\!\!\!|}^0_M N$.

Renaming a' to a, we have that $a \mathop{\smile\!\!\!\!\!\mid}^0_M N$.

Claim 3. Inside $\mathfrak{C}|_{\mathfrak{L}}$, there are $\mathfrak{L}$-isomorphic copies of N, N_1 and N_2, both containing a, with $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_a N_2$ and $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1A_2} N_1N_2$.

Proof of claim. For $i = 1, 2$, let N_i' be the copy of N coming from the $\mathfrak{L}$-automorphism $A_ia_i \mapsto A_ia$. By full existence for $\mathop{\smile\!\!\!\!\!\mid}^0$, let $N_i \equiv_{A_ia} N_i'$ with $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_{A_1a} N$ and $N_2 \mathop{\smile\!\!\!\!\!\mid}^0_{A_2a} NN_1$. Then $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1} N_1$ and $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_2} N_2$ by transitivity. From $a \mathop{\smile\!\!\!\!\!\mid}^0_M N$ we get $a \mathop{\smile\!\!\!\!\!\mid}^0_{A_1} A_2$, and so $A_1a \mathop{\smile\!\!\!\!\!\mid}^0_{A_1} A_2$. Along with $A_1 \mathop{\smile\!\!\!\!\!\mid}^0_M A_2$, transitivity gives $A_1a \mathop{\smile\!\!\!\!\!\mid}^0_M A_2$, so that $A_1a \mathop{\smile\!\!\!\!\!\mid}^0_a A_2$ by base monotonicity. This implies $A_1 \mathop{\smile\!\!\!\!\!\mid}^0_a A_2$ and $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_a A_2$. This last part implies $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_a A_2a$ and along with $N_2 \mathop{\smile\!\!\!\!\!\mid}^0_{A_2a} NN_1$ implies $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_a N_2$. Also, $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1A_2} N_1$ by base monotonicity since $A_1A_2 \subseteq N$. From $NN_1 \mathop{\smile\!\!\!\!\!\mid}^0_{A_2a} N_2$, we get $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_2aN_1} N_2$, and hence $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_2N_1} N_2$ since $a \in N_1$. Combining this with $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1A_2} N_1$ gives $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1A_2} N_1N_2$.

Claim 4. There is some model of $T \cup \Delta$ which is a $\mathfrak{D}$-extension of N, N_1, and N_2.

Proof of claim. Define $\mathfrak{D}$-structures on N_1 and N_2 such that (N_i, A_i, a) is $\mathfrak{D}$-isomorphic to (N, A_i, a_i) under $\alpha_i \colon N \to N_i$. So $N_i \models T^+$ for $i = 1, 2$. Let P be some $\mathfrak{L}$-elementary substructure of $\mathfrak{C}|_{\mathfrak{L}}$ containing N, N_1, and N_2. Note that since a_i is a $\mathfrak{D}$-substructure of N, a is also a $\mathfrak{D}$-substructure of N_i. Now $N_1 \mathop{\smile\!\!\!\!\!\mid}^0_a N_2$, and a is acl_T-closed – it is a model of the model complete theory T – so their $\mathfrak{D}$-structures can be amalgamated to one on $\langle N_1N_2 \rangle \leq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$ making it into a model of Δ using (d). By (c) and the fact that $A_1 \mathop{\smile\!\!\!\!\!\mid}^0_M A_2$, we have that $\langle A_1A_2 \rangle_{\mathfrak{L}}$ is a $\mathfrak{D}$-substructure of $\mathfrak{C}$. And hence by (b), $\mathrm{acl}_T(A_1A_2)$ is a $\mathfrak{D}$-substructure of $\mathfrak{C}$. Now $N \mathop{\smile\!\!\!\!\!\mid}^0_{A_1A_2} N_1N_2$, and so $N \mathop{\smile\!\!\!\!\!\mid}^0_{\mathrm{acl}_T(A_1A_2)} \langle N_1N_2 \rangle_{\mathfrak{L}}$ by base monotonicity and full existence. Now amalgamate the $\mathfrak{D}$-structures on N and $\langle N_1N_2 \rangle_{\mathfrak{L}}$ using (d) to one on $\langle NN_1N_2 \rangle_{\mathfrak{L}} \leq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$ making it into a model of Δ. By (a), P expands to a model of Δ extending the $\mathfrak{D}$-structures on N, N_1, and N_2.

Now P extends to some $S \models T^+$. Again since $T \cup \mathrm{diag}_{\mathfrak{L}}(P)$ is complete, we may assume that $P \leq_{\mathfrak{D}} S \preceq_{\mathfrak{L}} \mathfrak{C}|_{\mathfrak{L}}$. Now let $j \colon S \to \mathfrak{C}$ be the $\mathfrak{D}$-elementary embedding of S in $\mathfrak{C}$ that fixes N. Then

$$\mathrm{tp}^{\mathfrak{C}}(a_1/A_1) = \mathrm{tp}^N(a_1/A_1) = \mathrm{tp}^{N_1}(a/A_1) = \mathrm{tp}^S(a/A_1) = \mathrm{tp}^{\mathfrak{C}}(j(a)/A_1)$$

and similarly we have $j(a) \equiv_{A_2} a_2$. By construction of a, we had $a \mathop{\smile\!\!\!\!\!\mid}^0_M N$, and by monotonicity and invariance, we get $j(a) \mathop{\smile\!\!\!\!\!\mid}^0_M \mathrm{acl}(A_1A_2)$, and so $j(a) \mathop{\smile\!\!\!\!\!\mid}^+_M A_1A_2$.

Stationarity. By Morleyising T and using Corollary 4.3.2, we may assume that

T^+ has quantifier elimination. Let $M \prec N \prec \mathfrak{C}$, $a, b \in \mathfrak{C}$ such that $\mathrm{tp}(a/M) = \mathrm{tp}(b/M)$, and $a \mathop{\smile\!\!\!\!\!\mid}^{+}_{M} N$ and $b \mathop{\smile\!\!\!\!\!\mid}^{+}_{M} N$. We need to show that $\mathrm{tp}(a/N) = \mathrm{tp}(b/N)$. Write $K_a = \mathrm{acl}(Ma)$ and $K_b = \mathrm{acl}(Mb)$. By definition of $\mathop{\smile\!\!\!\!\!\mid}^{+}$, $K_a \mathop{\smile\!\!\!\!\!\mid}^{0}_{M} N$ and $K_b \mathop{\smile\!\!\!\!\!\mid}^{0}_{M} N$. By stationarity for $\mathop{\smile\!\!\!\!\!\mid}^{0}$, $\mathrm{tp}_{\mathfrak{L}}(K_a/N) = \mathrm{tp}_{\mathfrak{L}}(K_b/N)$, and hence there is an $\mathfrak{L}$-isomorphism $\langle K_a N\rangle_{\mathfrak{L}} \to \langle K_b N\rangle_{\mathfrak{L}}$ that fixes N. Now by (c), both $\langle K_a N\rangle_{\mathfrak{L}}$ and $\langle K_b N\rangle_{\mathfrak{L}}$ are $\mathfrak{D}$-substructures of $\mathfrak{C}$ which are models of Δ, and by its moreover clause, this $\mathfrak{L}$-isomorphism must be a $\mathfrak{D}$-isomorphism. By quantifier elimination for T^+, we must have $\mathrm{tp}(a/N) = \mathrm{tp}(b/N)$. ■

Remark 4.3.5. The moreover clauses of axioms (b) and (c) are only necessary to prove stationarity transfers from $\mathop{\smile\!\!\!\!\!\mid}^{0}$ to $\mathop{\smile\!\!\!\!\!\mid}^{+}$.

4.4 Examples

In this section we will see some applications of the above framework, both to existing proofs of simplicity and stability in the literature, and to new ones.

$\mathcal{D}$-fields are derivation-like over $\mathrm{ACFA}_{0,t}$

Let $\mathcal{D}$ be a finite-dimensional k-algebra satisfying Assumption **A**: that each maximal ideal of $\mathcal{D}$ has residue field k. Recall then that every $\mathcal{D}$-field (K, ∂) has a sequence $\sigma_1, \ldots, \sigma_t$ of associated endomorphisms which are uniformly $\mathcal{L}_{\mathrm{ring}}(\partial)$-definable in every $\mathcal{D}$-field.

Example 4.4.1. Let T be the simple $\mathcal{L}_{\mathrm{ring}}(\sigma_1, \ldots, \sigma_t)$-theory $\mathrm{ACFA}_{0,t}$, and let $\mathop{\smile\!\!\!\!\!\mid}^{0}$ be nonforking independence. Let $\mathfrak{D}$ be the language $\mathcal{L}_{\mathrm{ring}}(\partial_1, \ldots, \partial_l)$ and Δ the theory of $\mathcal{D}$-fields whose associated endomorphisms are $\sigma_1, \ldots, \sigma_t$. Then Δ is derivation-like with respect to T.

Axiom (a) is by Lemma 1.6.5. For axiom (b), let M be a $\mathcal{D}$-field and A some $\mathcal{D}$-structure. Now the $\mathcal{D}$-structure on A is a $\mathcal{D}$-operator $A \to M$, which extends by Lemma 1.6.5 to a $\mathcal{D}$-operator $\mathrm{acl}_T(A) \to M$. We also have a $\mathcal{D}$-operator $\mathrm{acl}_T(A) \to M$ given by restricting the $\mathcal{D}$-field structure on M. Since $\mathrm{acl}_T(A) = A^{\mathrm{alg}}$ is 0-étale over A, these must agree by Lemma 1.6.5. Axiom (c) comes from the multiplicative rules for $\mathcal{D}$-fields, and its moreover clause and (d) come from Lemma 5.1 of [53]. Note also that forking independence in $\mathrm{ACFA}_{0,t}$ is precisely linear disjointness after closing under acl.

Then Theorem 4.3.4 gives a different, less algebraic proof of Theorem 5.9 of [53].

Very $\mathcal{L}_{\text{ring}}(C)$-slim fields

In order to show that axiom (d) held in the previous subsection, we needed a full algebraic characterisation of nonforking independence. We can relax this requirement using a modified notion of Junker and Koenigsmann – that of very slim fields [32]. In these fields, algebraic independence is an abstract independence relation, and thus nonforking independence always implies algebraic independence.

Let $\mathfrak{L}$ be some language expanding the language of rings and T some $\mathfrak{L}$-theory of fields satisfying the following assumption.

Assumption 4.4.2. Suppose A and B are $\mathfrak{L}$-substructures of some $K \models T$. Then the field compositum of A and B is also an $\mathfrak{L}$-substructure of K.

Example 4.4.3. The above is satisfied if $\mathfrak{L} = \mathcal{L}_{\text{ring}}(C)$.

Definition 4.4.4. Let K be a field in the language $\mathfrak{L}$. We say that K is $\mathfrak{L}$-slim if for every $\mathfrak{L}$-substructure F, we have $\text{acl}^K_{\mathfrak{L}}(F) = F^{\text{alg}}$. Equivalently, if for every subset A, we have $\text{acl}^K_{\mathfrak{L}}(A) = \langle A \rangle^{\text{alg}}_{\mathfrak{L}}$. By F^{alg} we always mean the relative, field-theoretic algebraic closure of F in K.

We say that K is very $\mathfrak{L}$-slim if every $\mathfrak{L}$-structure elementarily equivalent to K is $\mathfrak{L}$-slim.

Remark 4.4.5.

1. We recover the definition of (very) slim from [32] by setting $\mathfrak{L} = \mathcal{L}_{\text{ring}}$ and only considering fields in the language of rings. In other words, a field with no extra structure is (very) slim exactly when it is (very) $\mathcal{L}_{\text{ring}}$-slim.
2. As mentioned in [32] for slim fields, to check whether K is very $\mathfrak{L}$-slim it is enough to check whether a sufficiently saturated model of its theory is $\mathfrak{L}$-slim.

The authors of [32] then show that algebraic independence in very slim fields is an abstract independence relation – in general, algebraic independence does not satisfy full existence (called existence in [32]). Here we need to modify algebraic independence slightly to account for the extra structure.

Definition 4.4.6. For a monster model of $\text{Th}(K)$, define the following relation on triples of subsets of K:

$$A \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{L}}_{D} B \iff \langle AD \rangle_{\mathfrak{L}} \mathop{\smile\!\!\!\!\!\!|}^{\text{alg}}_{\langle D \rangle_{\mathfrak{L}}} \langle BD \rangle_{\mathfrak{L}}.$$

We say that A and B are $\mathfrak{L}$-algebraically independent over D.

Theorem 2.1 of [32] says that a field is very slim if and only if algebraic independence is an independence relation. The following result is the very $\mathfrak{L}$-slim analogue.

Lemma 4.4.7. *K is very $\mathfrak{L}$-slim if and only if $\mathop{\smile\!\!\!\!\!\mid}^{\mathfrak{L}}$ is a strict independence relation.*

Proof. Invariance, *normality*, *monotonicity*, *symmetry*, *transitivity*, *finite character*, *local character*, and *strictness* are either by definition or follow from the corresponding property of $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}$. *Base monotonicity* follows from Assumption 4.4.2.

We now follow the proof of Theorem 2.1 of [32] to show that K is very $\mathfrak{L}$-slim if and only if $\mathop{\smile\!\!\!\!\!\mid}^{\mathfrak{L}}$ satisfies existence. Let a, B, C be given, and consider $\mathrm{tp}(a/\langle C\rangle)$. If it is algebraic, $a \in \mathrm{acl}(\langle C\rangle)$, and by very $\mathfrak{L}$-slimness, $a \in \langle C\rangle^{\mathrm{alg}}$. Then $a \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle} \langle B\rangle$. If it is not algebraic, then among the infinitely many realisations, there is one which is transcendental over $\langle B\rangle$ by compactness. Thus we have shown that if a is a tuple of length one, then there is some $a' \models \mathrm{tp}(a/\langle C\rangle)$ with $a' \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle} \langle B\rangle$. If $a = (a_1, \dots, a_n)$, then find $(a'_1, \dots, a'_n) \equiv_{\langle C\rangle} (a_1, \dots, a_n)$ with $a'_1 \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle} \langle B\rangle$. By induction, find $(a''_2, \dots, a''_n) \equiv_{\langle C\rangle a'_1} (a'_2, \dots, a'_n)$ with $a''_2, \dots, a''_n \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle a'_1} \langle B\rangle$. By transitivity, $(a'_1, a''_2 \dots, a''_n) \equiv_{\langle C\rangle} (a_1, \dots, a_n)$ and $a'_1, a''_2 \dots, a''_n \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle} \langle B\rangle$. So for all finite tuples a, there is some $a' \equiv_{\langle C\rangle} a$ with $a' \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{\langle C\rangle} \langle B\rangle$. By compactness this is true for infinite a. Now apply it with $\langle Ca\rangle$ to get $a' \mathop{\smile\!\!\!\!\!\mid}^{\mathfrak{L}}_{C} B$.

If K is not $\mathfrak{L}$-slim, then there is some $\mathfrak{L}$-substructure k and $a \in \mathrm{acl}(k)$ with a transcendental over k. Let C consist of the finitely many conjugates of a over k. Then there is no realisation of $\mathrm{tp}(a/k)$ which is algebraically independent from C over k. But if we did have $b \models \mathrm{tp}(a/k)$ with $b \mathop{\smile\!\!\!\!\!\mid}^{\mathfrak{L}}_{k} C$, then we would also have $b \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{alg}}_{k} C$ by monotonicity, a contradiction. ■

Remark 4.4.8. Remark 1.20 of [1] says that nonforking independence (which is not in general an independence relation) implies any independence relation. Indeed this fact is implicit in the proof of the Kim–Pillay theorem, see Theorem 4.2 Claim I of [35].

We now restrict to the case when $\mathfrak{L} = \mathcal{L}_{\mathrm{ring}}(C)$ for some set of constants C.

Remark 4.4.9. Let $\mathcal{C}$ be the field generated by the constants C inside K. Then by Lemma 2.12 of [31], K is very $\mathcal{L}_{\mathrm{ring}}(C)$-slim if and only if it is algebraically bounded over $\mathcal{C}$.

Lemma 4.4.10. *Let K be an $\mathcal{L}_{\mathrm{ring}}(C)$-structure which is a perfect field. Suppose K is large and its $\mathcal{L}_{\mathrm{ring}}(C)$-theory is model complete. Then K is very $\mathcal{L}_{\mathrm{ring}}(C)$-slim.*

Proof. The same proof as in Theorem 5.4 of [32] works here. In part 2 of that proof, when they take a subfield k, we instead take an $\mathfrak{L}$-substructure k, that is, a subfield containing the constants C. Model completeness then implies that ϕ is an existential $\mathcal{L}_{\text{ring}}(C)$-formula with parameters from k. But this is the same as an existential $\mathcal{L}_{\text{ring}}$-formula with parameters from k since k contains C. The rest of the proof is the same. ■

Example 4.4.11. Suppose T is the simple $\mathcal{L}_{\text{ring}}(C)$-theory of a very $\mathcal{L}_{\text{ring}}(C)$-slim field of characteristic zero and $\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^{0}$ is nonforking independence. Let $\mathcal{D}$ be some local finite-dimensional k-algebra with residue field k, let $\mathfrak{D} := \mathcal{L}_{\text{ring}}(C) \cup \{\partial_1, \dots, \partial_l\}$ and let Δ be one of the two following $\mathfrak{D}$-theories:

- the theory of $\mathcal{D}$-fields; or
- the theory of $\mathcal{D}$-fields where the operators pairwise commute.

Then Δ is derivation-like with respect to T and $\mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^{0}$.

Proof. If Δ is just the theory of $\mathcal{D}$-fields, then axioms (a), (b), and (c) hold for the same reason they do in Example 4.4.1. For axiom (d), since T is the theory of a very $\mathcal{L}_{\text{ring}}(C)$-slim field, Remark 4.4.8 tells us that if $A \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^{0}_{C} B$ where C is acl_T-closed, then A and B are algebraically independent over C, and hence they are linearly disjoint over C. Now we amalgamate using Lemma 5.1 of [53].

If Δ is the theory of $\mathcal{D}$-fields where the operators pairwise commute, the only things left to check are that in (a), the $\mathcal{D}$-field extension N may be taken to have commuting operators if M does, and that in (d), the $\mathcal{D}$-field structure on AB may also be taken to have commuting operators if A and B do.

For the first, let (M, ∂) be a $\mathcal{D}$-field with commuting operators, and let N be some field extension. Let T be a transcendence basis for N/M. Define δ on $M(T)$ by setting $\delta(t) = t \otimes 1 \in N \otimes_k \mathcal{D}$. So δ is a $\mathcal{D}$-operator along $M(T) \subseteq N$ whose operators pairwise commute. This is equivalent to the following diagram commuting.

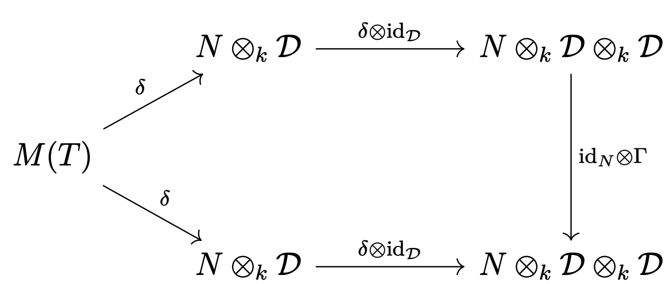

where $\Gamma\colon \mathcal{D} \otimes_k \mathcal{D} \to \mathcal{D} \otimes_k \mathcal{D}$ is the map that swaps the two coordinates. Now δ extends to a $\mathcal{D}$-field structure on N by Lemma 1.6.5 since $N/M(T)$ is 0-smooth. Then $(\delta \otimes \mathrm{id}_{\mathcal{D}}) \circ \delta$ and $(\mathrm{id}_N \otimes \Gamma) \circ (\delta \otimes \mathrm{id}_{\mathcal{D}}) \circ \delta$ are two $\mathcal{D} \otimes_k \mathcal{D}$-structures on N extending $(\delta \otimes \mathrm{id}_{\mathcal{D}}) \circ \delta = (\mathrm{id}_M \otimes \Gamma) \circ (\delta \otimes \mathrm{id}_{\mathcal{D}}) \circ \delta$ on $M(T)$. Since $N/M(T)$ is 0-étale, they must agree. So δ is a $\mathcal{D}$-field structure on N with pairwise commuting operators. Axiom (d) is a similar argument using the uniqueness guaranteed in Lemma 5.1 of [53]. ■

Corollary 4.4.12. 1. *If the model companion of* PAC *fields with a (pairwise commuting) $\mathcal{D}$-field structure exists, then it is simple.*

2. *If the model companion of algebraically closed fields with a pairwise commuting $\mathcal{D}$-field structure exists, then it is stable.*

Remark 4.4.13. In the next chapter, we will prove that the theory of PAC fields with a $\mathcal{D}$-field structure has a model companion and apply the results of this section.

Separably differentially closed fields of infinite differential degree of imperfection

In [28], Ino and León Sánchez study the class of ordinary differential fields which are existentially closed in every differential field extension which is separable as an extension of fields. They show that this class is elementary, and they denote its $\mathcal{L}_{\mathrm{ring}}(\delta)$-theory by SDCF: separably differentially closed fields. The theory of separably differentially closed fields of some fixed characteristic p is denoted by SDCF_p. Note that SDCF_0 is precisely DCF_0.

Recall that, in the case of fields, SCF_p is not a complete theory: one needs to specify the degree of imperfection $e \in \mathbb{N} \cup \{\infty\}$. Likewise, the authors define the *differential degree of imperfection* of a differential field (K, δ) of characteristic p to be $\epsilon \in \mathbb{N} \cup \{\infty\}$ such that $[C_K \colon K^p] = p^\epsilon$.

Definition 4.4.14. A tuple $\bar{a} \subseteq C_K$ is called differentially p-independent if the p-monomials over $\bar{a}$ are linearly independent over K^p. The tuple $\bar{a}$ is called a differential p-basis if the p-monomials form a basis for C_K over K^p.

As the authors mention in Remark 5.3, (K, δ) has differential degree of imperfection ϵ if and only if it has a differential p-basis of size ϵ.

Definition 4.4.15. 1. The $\mathcal{L}_{\text{ring}}(\delta)$-theory of differential fields of characteristic p and differential degree of imperfection ϵ is denoted $\mathrm{DF}_{p,\epsilon}$.

2. The $\mathcal{L}_{\text{ring}}(\delta)$-theory of separably differentially closed fields of characteristic p and differential degree of imperfection ϵ is denoted $\mathrm{SDCF}_{p,\epsilon}$.

Recall now the $\mathcal{L}_{\text{ring}}$-definable functions $(\lambda_{n,i} \colon n \in \omega, i \in p^n)$ from Section 1.4. Field extensions which preserve these functions are precisely the separable ones. In Section 6.2 of [28], the authors define the *differential* λ-functions, $(\ell_{n,i} \colon n \in \omega, i \in p^n)$, and show that $\mathrm{SDCF}^{\ell}_{p,\epsilon}$ is the model companion of $\mathrm{SCF}^{\ell}_{p,\infty} \cup \mathrm{DF}^{\ell}_{p,\epsilon}$. We require the analogous result using the algebraic λ-functions. The argument is essentially the same, but we will provide details anyway.

Fact 4.4.16. $\mathrm{SDCF}^{\lambda}_{p,\infty}$ *is the model companion of* $\mathrm{SCF}^{\lambda}_{p,\infty} \cup \{\delta$ *is a derivation*$\}$.

Proof. Firstly, every model of $\mathrm{SDCF}^{\lambda}_{p,\infty}$ is also a model of $\mathrm{SCF}^{\lambda}_{p,\infty} \cup \{\delta$ is a derivation$\}$ by Lemma 4.5 of [28]. By Lemma 5.6 of [28], any model of $\mathrm{SCF}^{\lambda}_{p,\infty} \cup \{\delta$ is a derivation$\}$ has a separable extension with infinite differential degree of imperfection, and by Proposition 5.10, this has a separable extension which is a model of $\mathrm{SDCF}_{p,\infty}$. Since all extensions are separable, they must preserve the λ-functions.

Now suppose $(K, \delta) \leq (L, d)$ is an extension of models of $\mathrm{SDCF}^{\lambda}_{p,\epsilon}$. Since this field extension preserves the λ-functions, L/K is a separable field extension. Now expand both K and L by the differential λ-functions. Since L/K is separable, the extension $K \leq L$ will also preserve these differential λ-functions; see Lemma 6.5. Then by model completeness for $\mathrm{SDCF}^{\ell}_{p,\epsilon}$ (Theorem 6.6), $K \preceq L$ as differential fields. Since the λ-functions are $\mathcal{L}_{\text{ring}}$-definable, $K \preceq L$ as models of $\mathrm{SDCF}^{\lambda}_{p,\epsilon}$. ■

In Theorem 6.8 of [28], Ino and León Sánchez prove that the $\mathcal{L}_{\text{ring}}(\delta)$-theory $\mathrm{SDCF}_{p,\infty}$ is stable by counting types. They do not characterise forking. The remainder of this subsection is devoted to showing that the theory of differential fields is derivation-like with respect to $\mathrm{SCF}^{\lambda}_{p,\infty}$ and hence that we may use the results of this chapter to characterise forking in $\mathrm{SDCF}_{p,\infty}$.

Let $\mathfrak{L}$ be the language of rings expanded by the λ-functions $\lambda_{n,i}$. Let $T = \mathrm{SCF}^{\lambda}_{p,\infty}$. This theory has quantifier elimination and is stable. Let $\mathop{\smile\!\!\!\!\!\!|}^0$ be nonforking independence. Srour characterises this in [61]:

$$A \mathop{\smile\!\!\!\!\!\!|}^0_C B \iff A \text{ and } B \text{ are algebraically independent and } p\text{-disjoint over } C.$$

Proposition 4.4.17. *The theory of differential fields is derivation-like with respect to* $\mathrm{SCF}^{\lambda}_{p,\infty}$.

Proof. Axiom (a) holds since derivations can always be extended to separable extensions; (b) holds since they extend uniquely to separably algebraic extensions. For axiom (c), if $A \mathop{\smile\!\!\!\!\!\!|}^{0}_{C} B$, then by p-disjointness, AB is a separable subfield of M; that is $\langle AB \rangle_{\mathfrak{L}} = AB$. Now by the Leibniz rule for derivations, AB is closed under δ if both A and B are. For axiom (d) and the moreover clause of (c), the same argument shows that $\langle MN \rangle_{\mathfrak{L}} = MN$. Also by Srour's characterisation of $\mathop{\smile\!\!\!\!\!\!|}^{0}$, M and N are linearly disjoint over A. Then the field compositum MN is the quotient field of $M \otimes_A N$, and thus by Lemma 5.1 of [53] there is a unique derivation on MN extending the ones on M and N. ■

Corollary 4.4.18. $\mathrm{SDCF}_{p,\infty}$ *is stable and in the language* $\mathcal{L}^{\lambda}(\delta)$ *nonforking independence is given by*

$$A \mathop{\smile\!\!\!\!\!\!|}^{+}_{C} B \iff \mathrm{acl}(A) \textit{ is linearly disjoint and } p\textit{-disjoint from } \mathrm{acl}(B) \textit{ over } \mathrm{acl}(C).$$

Chapter 5

Pseudo $\mathcal{D}$-closed fields

We now apply some of the results obtained in Chapters 3 and 4 to the study of PAC substructures in the theory $\mathcal{D}$-CF_0. We again impose Assumption **B** – that $\mathcal{D}$ is local. Recall that a field K is called pseudo algebraically closed (PAC) if every absolutely irreducible variety over K has a K-rational point. PAC fields are large – a K-irreducible variety with a smooth K-rational point is absolutely irreducible – and a field is PAC if and only if it is existentially closed in every regular extension.

In [13], Chatzidakis and Pillay show that if T_C is the $\mathcal{L}_{\mathrm{ring}}(\lambda)(C)$-theory of a bounded PAC field with λ interpreted by the λ-functions and the constants C naming coefficients of irreducible polynomials that encode all the finitely many Galois extensions of a fixed degree, then T_C is simple and that if it, in addition, has finite degree of imperfection, then it eliminates imaginaries after naming constants for a p-basis. Hoffman and León Sánchez in [26] then prove the analogous results for bounded pseudo differentially closed fields of characteristic zero. Their result gives an example of a differential field whose theory is simple and unstable. In this chapter we will prove analogous results in the case of $\mathcal{D}$-fields.

5.1 PAC substructures in $\mathcal{D}$-CF_0

PAC substructures of a given theory have been defined as generalisations of PAC fields in various ways. We use the definition presented in [25].

Definition 5.1.1. Let T be an arbitrary complete L-theory, and $\mathfrak{C}$ a monster model. An extension of L-substructures $A \leq B$ of $\mathfrak{C}$ is called L-regular if $\mathrm{dcl}^{\mathrm{eq}}(B) \cap$

$\mathrm{acl}^{\mathrm{eq}}(A) = \mathrm{dcl}^{\mathrm{eq}}(A)$. An L-substructure A of $\mathfrak{C}$ is called a PAC substructure if A is existentially closed in every L-regular extension.

Consider now the $\mathcal{L}_{\mathrm{ring}}(\partial)$-theory $\mathcal{D}$-CF$_0$. This theory eliminates imaginaries (see Theorem 5.12 of [53]), $\mathrm{acl}(A)$ is the (full) field-theoretic algebraic closure of the $\mathcal{D}$-field generated by A (Proposition 5.5 of [53]), and $\mathrm{dcl}(A)$ is the $\mathcal{D}$-field generated by A (this is not in [53] since there $\mathcal{D}$-fields may have associated endomorphisms; in the case when $\mathcal{D}$ is local, this fact is proved in the same way as for differential fields). Then an extension of $\mathcal{D}$-fields is $\mathcal{L}_{\mathrm{ring}}(\partial)$-regular exactly when the field extension is field-theoretically, relatively algebraically closed (and so regular in the field sense since we are in characteristic zero).

We now prove three conditions equivalent to being a PAC substructure in $\mathcal{D}$-CF$_0$.

Theorem 5.1.2. *Let (K, ∂) be a $\mathcal{D}$-field. The following are equivalent:*

(1) *(K, ∂) is a* PAC *substructure in the theory* $\mathcal{D}$-CF$_0$*;*

(2) *K is a* PAC *field and $(K, \partial) \models \mathrm{UC}_{\mathcal{D}}$;*

(3) *if (V, s) is a D-variety over K and V is absolutely irreducible, then (V, s) has a K-rational sharp point; and*

(4) *(K, ∂) is existentially closed in every $\mathcal{D}$-field extension (L, δ) which is R-regular, that is, where $\mathrm{tp}^{\mathcal{D}\text{-}\mathrm{CF}_0}(a/K)$ is stationary for every finite tuple $a \in L$.*

Proof. (1) $\Longrightarrow$ (2). Let L be any regular field extension of K, and let δ be any $\mathcal{D}$-structure on L extending ∂. Then (K, ∂) is existentially closed in (L, δ) as $\mathcal{D}$-fields, and hence K is existentially closed in L as fields. So K is PAC. Now since K is large, there is a $\mathcal{D}$-field extension $(F, \gamma) \models \mathrm{UC}_{\mathcal{D}}$ of (K, ∂) such that K is elementary in F as fields. In particular, $K \subseteq F$ is regular. By (1), (K, ∂) is existentially closed in (F, γ). Since $\mathrm{UC}_{\mathcal{D}}$ is inductive, $(K, \partial) \models \mathrm{UC}_{\mathcal{D}}$.

(2) $\Longrightarrow$ (1). Let (L, δ) be an $\mathcal{L}_{\mathrm{ring}}(\partial)$-regular $\mathcal{D}$-field extension of (K, ∂) so that L/K is a regular field extension. Since K is PAC, K is existentially closed in L as fields. By characterisation (5) of Theorem 3.3.5, (K, ∂) is existentially closed in (L, δ) as $\mathcal{D}$-fields.

(2) $\Longrightarrow$ (3). If V is absolutely irreducible, then $K(V)/K$ is a regular extension. V has a smooth $K(V)$-rational point, and since K is existentially closed in $K(V)$, V has a smooth K-rational point. By characterisation (3) of $\mathrm{UC}_{\mathcal{D}}$ in Theorem 3.3.5, (V, s) has a K-rational sharp point.

(3) $\Longrightarrow$ (2). Let V be an absolutely irreducible variety defined over K. Extend the $\mathcal{D}$-field structure on K to one on $K(V)$ using Lemma 1.6.5. As in the proof of Theorem 3.3.5 (3) $\Longrightarrow$ (1), there is an open affine subset $U \subseteq V$ defined over K such that this $\mathcal{D}$-ring structure restricts to one on $K[U]$. This gives a D-variety structure s on U, making (U, s) an absolutely irreducible D-variety defined over (K, ∂). By (3), (U, s) has a K-rational sharp point, and hence V has a K-rational point. So K is a PAC field. We again use characterisation (3) of Theorem 3.3.5 and the fact that a K-irreducible variety with a smooth K-rational point is absolutely irreducible to get that $(K, \partial) \models \mathrm{UC}_{\mathcal{D}}$.

(1) $\Longleftrightarrow$ (4) is the content of Lemma 3.36 in [25]; R-regular extensions are the same as $\mathcal{L}_{\mathrm{ring}}(\partial)$-regular extensions since $\mathcal{D}$-CF_0 is stable and eliminates imaginaries. ■

We say that a $\mathcal{D}$-field is pseudo $\mathcal{D}$-closed if any of the equivalent conditions of Theorem 5.1.2 hold.

Remark 5.1.3. Apart from condition (3), this is just the $\mathcal{D}$-field analogue of Theorem 5.16 from [38]. There the authors need to consider differential varieties as they work with several commuting derivations. In a single derivation, it is enough to consider D-varieties; see Proposition 5.6 of [55] for instance.

5.2 The model theory of bounded pseudo $\mathcal{D}$-closed fields

Theorem 5.2 of [26] states that the theory of a bounded pseudo differentially closed field (that is, a PAC substructure of $\mathrm{DCF}_{0,m}$) is simple and eliminates imaginaries. We will now prove the $\mathcal{D}$-field analogue. Let (K, ∂) be a bounded pseudo $\mathcal{D}$-closed field. For each $n > 1$, let $N(n)$ be the degree over K of the Galois extension composite of all Galois extensions of K of degree n. Let $C = (c_{n,i})_{n>1, 0 \leq i < N(n)}$ be the set of constant symbols in our language $\mathcal{L} = \mathcal{L}_{\mathrm{ring}}(C)$, and consider the set of $\mathcal{L}$-sentences $\Sigma_C = \{\sigma_n \colon n > 1\}$ where σ_n says that the polynomial $x^{N(n)} + c_{n,N(n)-1}x^{N(n)-1} + \cdots + c_{n,0}$ is irreducible and the extension this polynomial defines is Galois and contains all Galois extensions of K of degree n. This is the same set-up used by Chatzidakis and Pillay in Section 4 of [13] in their treatment of bounded PAC fields. Let $T^+ = \mathrm{Th}(K, \partial) \cup \Sigma_C$. Note then that $T^+ \supseteq \mathrm{Th}_{\mathcal{L}_{\mathrm{ring}}}(K) \cup \Sigma_C \cup \mathrm{UC}_{\mathcal{D}}$.

For the next two proofs, we will at times need to refer to notions both in the sense of T^+ and in the sense of $\mathcal{D}$-CF_0. In the second case, we will always include this as a superscript; if no superscript is given, the notion should be understood in the sense of whatever model of T^+ we are working in. The full, field-theoretic algebraic closure of A is denoted by $\tilde{A}$, and A^{alg} denotes the relative, field-theoretic algebraic closure of A. If $A \supseteq C$, then A^{alg} is equal to $\mathrm{acl}_{\mathcal{L}}(A)$ since models of T^+ are very $\mathcal{L}$-slim.

Lemma 5.2.1. *Let* $(F, \partial, C) \models T^+$, *and* $A \leq B \leq F$ *with* A acl-*closed in the sense of* T^+. *Then*

$$\mathrm{acl}^{\mathcal{D}\text{-}\mathrm{CF}_0}(B) = \mathrm{acl}(B) \cdot \mathrm{acl}^{\mathcal{D}\text{-}\mathrm{CF}_0}(A).$$

Proof. Since B is a $\mathcal{D}$-field containing C and F is very $\mathcal{L}$-slim, $\mathrm{acl}(B) = B^{\mathrm{alg}}$, and $\mathrm{acl}^{\mathcal{D}\text{-}\mathrm{CF}_0}(B) = \tilde{B}$. So we need to show $\tilde{B} = B^{\mathrm{alg}} \cdot \tilde{A}$. The proof of Proposition 4.6(2) of [13] shows that the restriction maps $\mathrm{Gal}(F) \to \mathrm{Gal}(A)$ and $\mathrm{Gal}(F) \to \mathrm{Gal}(\mathrm{acl}(B))$ are isomorphisms, and hence the restriction map $\mathrm{Gal}(\mathrm{acl}(B)) \to \mathrm{Gal}(A)$ is an isomorphism. Therefore, any automorphism of $\tilde{B}$ that fixes $B^{\mathrm{alg}} \cdot \tilde{A}$ must also fix $\tilde{B}$. Since we are in characteristic zero, $\tilde{B}/B^{\mathrm{alg}} \cdot \tilde{A}$ is a Galois extension, and so $\tilde{B} = B^{\mathrm{alg}} \cdot \tilde{A}$. ■

Remark 5.2.2. A similar result occurs in Lemma 3.8 of [57]. In [25], the author requires this fact as an assumption to prove his analogue of the following theorem.

Theorem 5.2.3. *Let* $(F, \partial, C) \models T^+$ *and* $(E, \partial, C) \subseteq (F, \partial, C)$. *Then*

(1) $\mathrm{acl}(E) = E^{\mathrm{alg}}$;

(2) *if* $E = \mathrm{acl}(E)$, *then* $T^+ \cup \mathrm{diag}(E)$ *is complete;*

(3) T^+ *is model complete;*

(4) *the independence theorem holds for* T^+ *over algebraically closed sets;*

(5) T^+ *is simple and forking is given by forking independence in* $\mathcal{D}$-CF_0;

(6) T^+ *has elimination of imaginaries.*

Proof. (1). By Proposition 4.3.3 since (F, C) is very $\mathcal{L}_{\mathrm{ring}}(C)$-slim (it is model complete, large, and characteristic zero).

(2). By Proposition 4.6(2) of [13], $\mathrm{Th}_{\mathcal{L}_{\mathrm{ring}}}(K) \cup \Sigma_C \cup \mathrm{diag}(E \restriction_{\mathcal{L}})$ is complete. Then Theorem 3.2.6 tells us that $\mathrm{Th}_{\mathcal{L}_{\mathrm{ring}}}(K) \cup \Sigma_C \cup \mathrm{UC}_{\mathcal{D}} \cup \mathrm{diag}(E)$ is complete. So $T^+ \cup \mathrm{diag}(E)$ is complete.

(3). By Theorem 3.2.7(i) since $\mathrm{Th}_{\mathcal{L}_{\mathrm{ring}}}(K) \cup \Sigma_C$ is model complete (Proposition 4.6(1) of [13]).

(4). This follows from Theorem 4.3.4, Example 4.4.11, and the fact that the independence theorem over algebraically closed sets holds for bounded PAC fields (Theorem 4.7 of [13]).

(5). By Theorem 4.3.4, Example 4.4.11, and the corresponding result for bounded PAC fields (Corollary 4.8 of [13]) we know that T^+ is simple and forking independence is given by linear disjointness after closing under acl – the relative algebraic closure of the $\mathcal{D}$-field it generates. We can then use general properties of linear disjointness of regular extensions, along with Lemma 5.2.1, to show that, if A, B, and D are all acl-closed, then A and B are linearly disjoint over D if and only if $\tilde{A}$ is linearly disjoint from $\tilde{B}$ over $\tilde{D}$. This is precisely forking independence in $\mathcal{D}$-CF_0 (see Theorem 5.9 of [53]).

(6). This proof is essentially a combination of Theorem 5.12 of [53], Theorem 4.36 of [25], and Theorem 5.6 of [26]. Nonetheless, some details will be provided. We will assume that (F, ∂, C) is a monster model of (some completion of) T^+, and that $(\mathfrak{D}, \partial)$ is a monster model of $\mathcal{D}$-CF_0 extending it. We write $\mathop{\smile\!\!\!\!\!\!|}$ for nonforking independence in (F, ∂, C). If we omit a superscript from an operator, we mean in the sense of (F, ∂, C).

We need the notion of dimension from Definition 5.10 of [53]. If K is a $\mathcal{D}$-field, then $\dim_{\mathcal{D}}(a/K) = (\mathrm{trdeg}(\nabla_r(a)/K) \colon r < \omega) \in \omega^\omega$, where $\nabla_r(a)$ is the tuple applying words of length at most r in the language $\{\partial_1, \ldots, \partial_l\}$ to a. We order dimensions with the lexicographic order on ω^ω. Note that $\dim_{\mathcal{D}}(a/K) = \dim_{\mathcal{D}}(a/\tilde{K})$. Using Lemma 5.11 of [53], we then get that if L/k is a regular extension, $\dim_{\mathcal{D}}(a/k) = \dim_{\mathcal{D}}(a/L)$ if and only if $\mathrm{acl}^{\mathcal{D}\text{-}\mathrm{CF}_0}(ka)$ is linearly disjoint from $\tilde{L}$ over $\tilde{k}$ if and only if $a \mathop{\smile\!\!\!\!\!\!|}_k L$.

Let $e \in (F, \partial, C)^{\mathrm{eq}}$ given by a 0-definable function f and a finite real tuple $a \in F$, that is, $f(a) = e$. Let $E = \mathrm{acl}^{\mathrm{eq}}(e) \cap F$ and let Q be the set of realisations of $\mathrm{tp}(a/E)$. We now follow the proof in Theorem 5.12 of [53] to find some $u \in Q$ such that $f(u) = e$ and $u \mathop{\smile\!\!\!\!\!\!|}_E a$.

As in Theorem 5.12 of [53], Neumann's lemma tells us that there is some $b_0 \models \mathrm{tp}(a/Ee)$ such that $\mathrm{acl}^{\mathrm{eq}}(Ea) \cap \mathrm{acl}^{\mathrm{eq}}(Eb_0) \cap F = E$. Then $f(b_0) = e$ and $b_0 \in Q$. We want $b \in F$ such that

- $b \models \mathrm{tp}(a/Ee)$;

- $\operatorname{acl}^{\mathrm{eq}}(Ea) \cap \operatorname{acl}^{\mathrm{eq}}(Eb) \cap F = E$;
- b has maximal $\dim_{\mathcal{D}}$ over $\operatorname{acl}(Ea)$.

By the same argument as in Theorem 5.12 of [53], we can in fact choose such a b.

Now let $u \models \operatorname{tp}(b/\operatorname{acl}(Ea))$ with $u \mathop{\smile\!\!\!\!\!\!|}_{Ea} b$. Then $f(u) = e$. It remains to show $u \mathop{\smile\!\!\!\!\!\!|}_{E} a$. Since $u \mathop{\smile\!\!\!\!\!\!|}_{Ea} b$, we have that $\operatorname{acl}^{\mathrm{eq}}(Eau) \cap \operatorname{acl}^{\mathrm{eq}}(Eab) \subseteq \operatorname{acl}^{\mathrm{eq}}(Ea)$. Then $\operatorname{acl}^{\mathrm{eq}}(Eu) \cap \operatorname{acl}^{\mathrm{eq}}(Eb) \cap F \subseteq \operatorname{acl}^{\mathrm{eq}}(Ea) \cap \operatorname{acl}^{\mathrm{eq}}(Eb) \cap F = E$. Let d be such that $\operatorname{tp}(bu/Ee) = \operatorname{tp}(ad/Ee)$. Then $d \models \operatorname{tp}(a/Ee)$ and $\operatorname{acl}^{\mathrm{eq}}(Ea) \cap \operatorname{acl}^{\mathrm{eq}}(Ed) \cap F = E$. By maximality, $\dim_{\mathcal{D}}(d/\operatorname{acl}(Ea)) \leq \dim_{\mathcal{D}}(b/\operatorname{acl}(Ea))$, and since $\dim_{\mathcal{D}}$ is automorphism invariant, $\dim_{\mathcal{D}}(u/\operatorname{acl}(Eb)) \leq \dim_{\mathcal{D}}(b/\operatorname{acl}(Ea))$. We also have that

$$\dim_{\mathcal{D}}(u/\operatorname{acl}(Eb)) \geq \dim_{\mathcal{D}}(u/\operatorname{acl}(Eab)) = \dim_{\mathcal{D}}(u/\operatorname{acl}(Ea)) = \dim_{\mathcal{D}}(b/\operatorname{acl}(Ea)).$$

The first equality is true since $u \mathop{\smile\!\!\!\!\!\!|}_{Ea} b$ and the second since u and b have the same type over $\operatorname{acl}(Ea)$. Hence all these dimensions are equal, and $\dim_{\mathcal{D}}(u/\operatorname{acl}(Eb)) = \dim_{\mathcal{D}}(u/\operatorname{acl}(Eab))$, that is, $u \mathop{\smile\!\!\!\!\!\!|}_{Eb} a$.

Let $p = \operatorname{tp}^{\mathcal{D}\text{-CF}_0}(u/\operatorname{acl}(Eab))$. Since $\operatorname{acl}(Eab)$ is regular in F and $u \in F$, p is stationary (recall that R-regular and $\mathcal{L}_{\mathrm{ring}}(\partial)$-regular extensions are the same). Then $\operatorname{Cb}(p) \subseteq \operatorname{dcl}^{\mathcal{D}\text{-CF}_0}(\operatorname{acl}(Eab)) = \operatorname{acl}(Eab) \subseteq F$. From $u \mathop{\smile\!\!\!\!\!\!|}_{Ea} b$, we get $\operatorname{Cb}(p) \subseteq \operatorname{acl}^{\mathcal{D}\text{-CF}_0}(Ea)$, and from $u \mathop{\smile\!\!\!\!\!\!|}_{Eb} a$, we get $\operatorname{Cb}(p) \subseteq \operatorname{acl}^{\mathcal{D}\text{-CF}_0}(Eb)$. So

$$\begin{aligned}\operatorname{Cb}(p) &\subseteq \operatorname{acl}^{\mathcal{D}\text{-CF}_0}(Ea) \cap \operatorname{acl}^{\mathcal{D}\text{-CF}_0}(Eb) \cap F \\ &= \operatorname{acl}(Ea) \cap \operatorname{acl}(Eb) \\ &= E.\end{aligned}$$

So p does not fork over $\tilde{E}$ and $u \mathop{\smile\!\!\!\!\!\!|}_{E} ab$. Then $u \mathop{\smile\!\!\!\!\!\!|}_{E} a$. This completes the claim. We now follow the rest of the argument in Theorem 5.6 of [26]. Let $D = \{d \in Q \colon f(d) = e\}$. If $D = Q$, then $e \in \operatorname{dcl}^{\mathrm{eq}}(E)$ and we get weak elimination of imaginaries.

If $D \subsetneq Q$, let $d_0 \in Q \setminus D$ and $d \equiv_E d_0$ with $d \mathop{\smile\!\!\!\!\!\!|}_{E} D$. If $f(d) = e$, then $d \in D$ and hence $d \in \operatorname{acl}(E) = E$. So $d \in \operatorname{acl}^{\mathrm{eq}}(e)$. Since $f(d) = e$, $e \in \operatorname{dcl}^{\mathrm{eq}}(d)$, and we get weak elimination of imaginaries.

So assume $f(d) \neq e$. Now $u \equiv_E d$, $u \mathop{\smile\!\!\!\!\!\!|}_{E} a$, and $u \mathop{\smile\!\!\!\!\!\!|}_{E} d$. By the independence theorem over algebraically closed sets, we get $m \models \operatorname{tp}(u/Ea) \cup \operatorname{tp}(d/Eu)$ with $m \mathop{\smile\!\!\!\!\!\!|}_{E} au$. But this contradicts $f(d) \neq e$. Finally, since we are in a theory of fields and we have weak elimination of imaginaries, we have elimination of imaginaries. ■

Bibliography

[1] Hans Adler. *Explanation of Independence.* PhD thesis, University of Freiburg, 2005.

[2] Hans Adler. A geometric introduction to forking and thorn-forking. *Journal of Mathematical Logic*, 9(1):1–20, 2009.

[3] Michael F. Atiyah and Ian G. Macdonald. *Introduction to Commutative Algebra.* Addison-Wesley, 1969.

[4] Özlem Beyarslan, Daniel Max Hoffmann, Moshe Kamensky, and Piotr Kowalski. Model theory of fields with free operators in positive characteristic. *Transactions of the American Mathematical Society*, 372(8):5991–6016, 2019.

[5] A. Białynicki-Birula. On Galois Theory of Fields with Operators. *American Journal of Mathematics*, 84(1):89–109, 1962.

[6] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. *Néron Models*, volume 21 of *Ergebnisse der Mathematik und ihrer Grenzgebiete.* Springer Berlin, Heidelberg, 1990.

[7] Alexandru Buium. *Differential algebra and Diophantine geometry.* Actualités mathématiques. Hermann, Paris, 1994.

[8] Alexandru Buium. Arithmetic analogues of derivations. *Journal of Algebra*, 198(1):290–299, 1997.

[9] F. Campana. Algébricité et compacité dans l'espace des cycles d'un espace analytique complexe. *Mathematische Annalen*, 251:7–18, 1980.

[10] J.W.S. Cassels. An embedding theorem for fields. *Bulletin of the Australian Mathematical Society*, 14(2):193–198, 1976.

[11] Zoé Chatzidakis. Generic automorphisms of separably closed fields. *Illinois Journal of Mathematics*, 45(3):693–733, 2001.

[12] Zoé Chatzidakis, Lou van den Dries, and Angus MacIntyre. Definable sets over finite fields. *Journal für die reine und angewandte Mathematik*, 427:107–136, 1992.

[13] Zoé Chatzidakis and Anand Pillay. Generic structures and simple theories. *Annals of Pure and Applied Logic*, 95(1):71–92, 1998.

[14] Zoé Chatzidakis. Simplicity and independence for pseudo-algebraically closed fields. In S. Barry Cooper and John K Truss, editors, *Models and Computability*, volume 259 of *London Mathematical Society Lecture Note Series*, pages 41 —- 62. Cambridge University Press, 1999.

[15] Gregory Cousins. *Some model theory of fields and differential fields*. PhD thesis, University of Notre Dame, 2019.

[16] Françoise Delon. Idéaux et types sur les corps séparablement clos. *Mémoires de la Société Mathématique de France*, 33:1–76, 1988.

[17] Antongiulio Fornasiero and Giuseppina Terzo. Generic derivations on algebraically bounded structures. *Preprint*, arXiv:2310.20511, 2023.

[18] Michael D. Fried and Moshe Jarden. *Field Arithmetic*, volume 11 of *A Series of Modern Surveys in Mathematics*. Springer Berlin, Heidelberg, 1986.

[19] Akira Fujiki. On the Douady space of a compact complex space in the category. *Nagoya Mathematical Journal*, 85:189–211, 1982.

[20] Jakub Gogolok and Piotr Kowalski. Operators coming from ring schemes. *Journal of the London Mathematical Society*, 106(3):1725–1758, 2022.

[21] Alexander Grothendieck. Technique de descente et théorèmes d'existence en géométrie algébrique. I. Généralités. Descente par morphismes fidèlement plats. In *Séminaire Bourbaki*, volume 5, Exp. 190, pages 299–327. Société Mathématique de France, Paris, 1995.

[22] Nicolas Guzy and Françoise Point. Topological differential fields. *Annals of Pure and Applied Logic*, 161(4):570–598, 2010.

[23] Charlotte Hardouin. Iterative q-difference galois theory. *Journal für die reine und angewandte Mathematik*, 2010(644):101–144, 2010.

[24] Wilfrid Hodges. *Model Theory.* Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.

[25] Daniel Max Hoffman. Model theoretic dynamics in a Galois fashion. *Annals of Pure and Applied Logic*, 170(7):755–804, 2019.

[26] Daniel Max Hoffman and Omar León Sánchez. Model theory of differential fields with finite group actions. *Journal of Mathematical Logic*, 22(1):2250002, 2022.

[27] Ehud Hrushovski. The Manin–Mumford conjecture and the model theory of difference fields. *Annals of Pure and Applied Logic*, 112(1):43–115, 2001.

[28] Kai Ino and Omar León Sánchez. Separably differentially closed fields. *arXiv:2302.11319, Preprint*, 2023.

[29] Thomas Jech. *Set Theory.* Pure and Applied Mathematics: A Series of Monographs and Textbooks. Academic Press, 1978.

[30] Will Johnson, Chieu-Minh Tran, Erik Walsberg, and Jinhe Ye. The étale-open topology and the stable fields conjecture. *Journal of the European Mathematical Society*, 2023.

[31] Will Johnson and Jinhe Ye. A note on geometric theories of fields. *Model Theory*, 2(1):121–132, 2023.

[32] Markus Junker and Jochen Koenigsmann. Schlanke Körper (slim fields). *The Journal of Symbolic Logic*, 75(2):481–500, 2010.

[33] Hirotaka Kikyo and Saharon Shelah. The strict order property and generic automorphisms. *Journal of Symbolic Logic*, 67:214–216, 2002.

[34] Byunghan Kim. *Simplicity Theory*, volume 53 of *Oxford Logic Guides.* Oxford University Press, 2013.

[35] Byunghan Kim and Anand Pillay. Simple theories. *Annals of Pure and Applied Logic*, 88(2):149–164, 1997.

[36] Piotr Kowalski and Anand Pillay. Quantifier elimination for algebraic D-groups. *Transactions of the American Mathematical Society*, 358(1):167–181, 2006.

[37] Omar León Sánchez. Algebro-geometric axioms for $\mathrm{DCF}_{0,m}$. *Fundamenta Mathematicae*, 01 2016.

[38] Omar León Sánchez and Marcus Tressl. Differentially large fields. *to appear in Algebra and Number Theory*, arXiv:2005.00888, 2020.

[39] Omar León Sánchez and Marcus Tressl. Differential Weil descent. *Communications in Algebra*, 50(1):104–114, 2022.

[40] Alexander Levin. *Difference Algebra*, volume 8 of *Algebra and Applications*. Springer Dordrecht, 2008.

[41] Saunders Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer New York, 1978.

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

[43] David Marker. *Model Theory: An Introduction*, volume 217 of *Graduate texts in mathematics*. Springer, 2002.

[44] Hideyuki Matsumura. *Commutative Ring Theory*, volume 8 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1989.

[45] Tracey McGrail. The model theory of differential fields with finitely many commuting derivations. *Journal of Symbolic Logic*, 65(2):885–913, 2000.

[46] Margit Messmer. Some model theory of separably closed fields. In *Model Theory of Fields*, Lecture Notes in Logic, page 135–152. Cambridge University Press, 2017.

[47] Christian Michaux and Cédric Rivière. Quelques remarques concernant la théorie des corps ordonnés différentiellement clos. *Bulletin of the Belgian Mathematical Society*, 12(3):341–348, 2005.

[48] James S. Milne. On the arithmetic of abelian varieties. *Inventiones Mathematicae*, 17(3):177–190, 1972.

[49] Shezad Mohamed. The uniform companion for fields with free operators in characteristic zero. *Preprint*, arXiv:2311.01856, 2024.

[50] Shezad Mohamed. The Weil descent functor in the category of algebras with free operators. *Journal of Algebra*, 640:216–252, 2024.

[51] Rahim Moosa and Thomas Scanlon. Jet and prolongation spaces. *Journal of the Institute of Mathematics of Jussieu*, 9(2):391–430, 2010.

[52] Rahim Moosa and Thomas Scanlon. Generalized Hasse-Schmidt varieties and their jet spaces. *Proceedings of the London Mathematical Society*, 103(2):197–234, 2011.

[53] Rahim Moosa and Thomas Scanlon. Model theory of fields with free operators in characteristic zero. *Journal of Mathematical Logic*, 14(2):1450009, 2014.

[54] David Pierce and Anand Pillay. A note on the axioms for differentially closed fields of characteristic zero. *Journal of Algebra*, 204(1):108–115, 1998.

[55] Anand Pillay and Dominika Polkowska. On PAC and bounded substructures of a stable structure. *The Journal of Symbolic Logic*, 71(2):460–472, 2006.

[56] Anand Pillay and Martin Ziegler. Jet spaces of varieties over differential and difference fields. *Selecta Mathematica, New Series*, 9:579–599, 2003.

[57] O. P. Nicholas Marie Polkowska. On simplicity of bounded pseudoalgebraically closed structures. *Journal of Mathematical Logic*, 07(02):173–193, 2007.

[58] Florian Pop. Embedding Problems Over Large Fields. *Annals of Mathematics*, 144(1):1–34, 1996.

[59] Alexander Prestel and Peter Roquette. *Formally p-adic fields.* Number 1050) in Lecture notes in mathematics. Springer, 1984.

[60] Abraham Robinson. *Complete Theories.* Studies in logic and the foundations of mathematics. North-Holland Publishing Company, 1956.

[61] G. Srour. The independence relation in separably closed fields. *The Journal of Symbolic Logic*, 51(3):715–725, 1986.

[62] Moss Eisenberg Sweedler. When is the tensor product of algebras local? *Proceedings of the American Mathematical Society*, 48(1):8–10, 1975.

[63] Mitsuhiro Takeuchi. A Hopf algebraic approach to the Picard–Vessiot theory. *Journal of Algebra*, 122:481–509, 1989.

[64] Katrin Tent and Martin Ziegler. *A Course in Model Theory*. Lecture Notes in Logic. Cambridge University Press, 2012.

[65] Marcus Tressl. The uniform companion for large differential fields of characteristic 0. *Transactions of the American Mathematical Society*, 357(10):3933–3951, 2005.

[66] L. van den Dries and K. Schmidt. Bounds in the theory of polynomial rings over fields. a nonstandard approach. *Inventiones Mathematicae*, 76(1):77–91, 1984.

[67] André Weil. *Adeles and Algebraic Groups*, volume 23 of *Progress in Mathematics*. Birkhäuser Boston, 1982.