SIX LECTURES ON MODEL THEORY AND DIFFERENTIAL-ALGEBRAIC GEOMETRY

RAHIM MOOSA

Contents

What follows is a write-up of some lectures I gave in the Fall of 2021 at the Fields Institute in Toronto, as part of the Thematic Programme on Trends in Pure and Applied Model Theory. The actual lectures were given in four 90 minute instalments, but I present them here as six lectures because the material organised itself better that way. I have taken other liberties as well: corrected errors, filled gaps, and improved exposition.

The goal of this module was to give a quick introduction to the model theory of differential fields that puts differential-algebraic geometry at the center. As such, fundamental algebraic and model theoretic aspects of the subject, that would normally form the core of such a course, are entirely omitted. Instead, I have tried to keep my focus on the birational geometry of algebraic vector fields, and more generally D-varieties in the sense of Buium [1]. Applications of model theory to differential-algebraic geometry is an active area of research, and the approach I have taken here is meant to both whet the student's appetite and prepare them for work in the subject.

The style of the lectures are rather informal, lacking in both rigour and detail. While I make every attempt to explain the central ideas, I leave many proofs to the student and I give essentially no references. Moreover, I am content to take as black boxes many important theorems, especially when I have nothing to add to their exposition as it already exists in the literature.

There are a number of resources out there on the model theory of differential fields that the reader can consult for a more thorough, and more traditional, introduction. Among them let me only mention Dave Marker's very influential treatment of the subject in [5], and Anand Pillay's chapter in [2]. (It is maybe interesting to note that the latter is based on another, much longer, course on differential

---

fields held during another thematic programme at the Fields Institute some 25 years earlier.) These articles also contain the references that I have generally omitted, coming, for example, from the differential algebra literature. There is very little overlap between the lecture notes I am presenting here and either of these precedents.

I am grateful to the Fields Institute for hosting what was, for me personally, a very productive and stimulating programme. It was especially successful as an early effort to return to normal in-person academic life and collaborative work.

1. From geometry to algebra to model theory

My goal in this first lecture is to describe how the classical geometry-algebra correspondence gives rise to the model theory of differentially closed (and indeed difference closed) fields.

But first, as a review and to set the stage, let us consider the model theory of pure algebraically closed fields from this approach. At its core, and speaking very loosely, the geometry-algebra correspondence is between a geometric space and the algebra of functions on the space. In particular, to affine n-space $\mathbb{A}^n$ over a field k of characteristic zero we associate the polynomial ring $k[x]$ in the n variables $x = (x_1, \ldots, x_n)$, with coefficients in k. These we naturally view as functions on $\mathbb{A}^n$ in the sense that we can evaluate a polynomial at any n-tuple from any field extension of k. To an algebraic subvariety $V \subseteq \mathbb{A}^n$ over k we associate the co-ordinate ring $k[V] = k[x]/I(V)$, where $I(V)$ is the ideal of polynomials vanishing on V. That is, we restrict the polynomials to functions on V, identifying polynomials if they agree on V – or rather on $V(K)$ for all field extensions $K \supseteq k$. So the algebraic counterpart of the variety V is the finitely generated k-algebra $k[V]$, which is an integral domain if V is irreducible. In fact, every finitely generated integral k-algebra, R, arises in this way. Indeed, fixing generators $a_1, \ldots, a_n \in R$ we obtain a surjective k-algebra homomorphism $\phi : k[x] \to R$ taking each x_i to a_i, so that R is isomorphic to $k[x]/\ker \phi$, and letting $V \subseteq \mathbb{A}^n$ be the subvariety determined by setting the polynomials in $\ker \phi$ to zero, we have exhibited R as the co-ordinate ring of V. So, if you are interested in the geometry of (embedded) irreducible affine algebraic varieties over k then you should study finitely generated integral k-algebras. A model-theorist will, of course, do so in the natural language $\mathcal{L}_k = \{0, 1, +, -, \times, (\lambda_r)_{r \in k}\}$ of k-algebras. Here λ_r denotes the unary function symbol to be interpreted as sclar multiplication by $r \in k$. Being finitely generated is not axiomatisable, but we can at least consider the (universal) theory T of integral k-algebras in this language. To better understand T we should look to the *existentially closed models*: those $M \models T$ with the property that whenever a system of polynomial equations and inequations (i.e., a conjunction of atomic and negated atomic formulas) with coefficients from M has a realisation in some model of T extending M then it already has a realisation in M. The class of existentially closed models of T is itself axiomatisable; given by the theory ACF_k of algebraically closed fields extending k. It is thus that the study of affine algebraic varieties over k leads to the first order theory ACF_k. Moreover, we recover the co-ordinate rings we were originally interested in as precisely the finitely generated substructures of the models of ACF_k.

The path we have just described, from geometry to algebra to model theory, and then back again, is the template for the expansion of algebraic geometry that we now study. We will add additional structure to the algebraic varieties that we consider, right at the beginning of the above process, and then trace where this leads to algebraically and model-theoretically. That additional structure is a vector field. Namely, we are interested in algebraic varieties, $V \subseteq \mathbb{A}^n$, equipped with a polynomial function that picks out, for each point $v \in V$, a tangent vector to V at v in $\mathbb{A}^n$. Let's make this precise:

Definition 1.1. Suppose $V \subseteq \mathbb{A}^n$ is an affine variety over k. By the *tangent bundle* to V we mean the subvariety $TV \subseteq \mathbb{A}^{2n}$ over k defined, in co-ordinates $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$, by

$$
\begin{aligned}
f(x) &= 0 \\
\sum_{i=1}^{n} \frac{\partial f}{\partial x_i} y_i &= 0
\end{aligned}
$$

for all $f \in I(V)$. The projection onto the x co-ordinates gives us a surjective morphism $\pi : TV \to V$.

An *algebraic vector field on V* is then a morphism $s : V \to TV$ over k which is a section to π, that is, $\pi \circ s = \mathrm{id}_V$. We will also sometimes refer to the pair (V, s) as an *algebraic vector field* over k.

For any field extension $K \supseteq k$, and any point $v \in V(K)$, the fibre of $\pi : TV \to V$ over v is a linear subspace of $\mathbb{A}^n$ defined over the field $k(v)$. Staring at the equations, we see that it is in fact the familiar *tangent space to V at v* which we denote by $T_v V$. The algebraic vector field $s : V \to TV$ is given by $s = (\mathrm{id}_V, s_1, \ldots, s_n)$ where $s_1, \ldots s_n \in k[V]$ and we have that $(s_1(v), \ldots, s_n(v)) \in T_v V(K)$.

I hope it is clear that algebraic vector fields are inherently of interest, and I will make no effort to justify this claim. In any case, we will take them as our basic geometric objects of study. Our first question is: what is the algebraic counterpart to (V, s)? In other words, what algebraic structure does s induce on the co-ordinate ring $k[V]$? The answer (you will have guessed) is a derivation.

Definition 1.2. A *derivation* on a (commutative) ring R is a function $\delta : R \to R$ that is additive, namely satisfying $\delta(a + b) = \delta(a) + \delta(b)$, and satisfies the Leibniz rule $\delta(ab) = \delta(a)b + a\delta(b)$, for all $a, b \in R$. By the *constants* of (R, δ) we mean $R^\delta := \{a \in R : \delta(a) = 0\}$.

Note that the constants R^δ form a subring. Assuming R is nontrivial we must have that R^δ contains the integers: $\delta(1) = \delta(1 \cdot 1) = \delta(1)1 + 1\delta(1) = 2\delta(1)$ forces $\delta(1) = 0$. Note also that if R is a k-algebra, then a derivation δ on R will be k-linear if and only if $k \subseteq R^\delta$. Indeed, k-linearity forces $\delta(\lambda) = \delta(\lambda 1) = \lambda\delta(1) = 0$ for all $\lambda \in k$, and conversely, $k \subseteq R^\delta$ implies $\delta(\lambda a) = \delta(\lambda)a + \lambda\delta(a) = \lambda\delta(a)$ for all $a \in R$.

Proposition 1.3. *Suppose $V \subseteq \mathbb{A}^n$ is a subvariety over k. If $s = (\mathrm{id}_V, s_1, \ldots, s_n)$ is an algebraic vector field on V over k then there is a unique k-linear derivation δ_s on $k[V]$ such that $\delta_s(x_i + I(V)) = s_i$ for all $i = 1, \ldots, n$.*

Moreover, every k-linear derivation on $k[V]$ is of the form δ_s for some algebraic vector field s on V over k.

Proof. First we define δ_s on the polynomial ring $k[x]$. Write each $s_i = g_i + I(V)$ for some $g_i \in k[x]$. Then there is a unique k-linear derivation on $k[x]$ satisfying

$\delta_s(x_i) = g_i$. Indeed, uniqueness is clear because k-linearity and the Leibniz rule ensure that a k-linear derivation on a k-algebra is determined by its action on generators. For existence, I give you the formula and leave it to you to check that it works:

$$\delta_s f := \sum_{i=1}^{n} \frac{\partial f}{\partial x_i} g_i \tag{1}$$

for each $f \in k[x]$.

The next thing to observe is that $I(V)$ is a δ_s-*ideal*: it is closed under the action of δ_s. Here we use that $s : V \to TV$. Indeed, for every field extension $K \supseteq k$, and every point $v \in V(K)$, we have that $s(v) \in TV(K)$, and hence

$$\sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(v) s_i(v) = 0$$

for all $f \in I(V)$. Since the above identity holds for all K-points of V as we range over all field extensions K, and since $s_i = g_i + I(V)$, we get that

$$\sum_{i=1}^{n} \frac{\partial f}{\partial x_i} g_i \in I(V).$$

By (1), $\delta_s f \in I(V)$ for all $f \in I(V)$, as desired.

So $\delta_s : k[x] \to k[x]$ induces a k-linear derivation on $k[x]/I(V) = k[V]$, which we also denote by δ_s. This has the desired property of $\delta_s(x_i + I(V)) = s_i$ for all $i = 1, \ldots, n$, by construction. Uniqueness follows as before.

Finally, for the "moreover" clause, suppose we begin with a k-linear derivation δ on $k[V]$. Write $\delta(x_i + I(V)) =: s_i$ for all $i = 1, \ldots, n$. Using k-linearity and the Leibniz rule one verifies that for any $f \in k[x]$, any $K \supseteq k$ a field extension, and any $v \in V(K)$,

$$\delta(f + I(V))(v) = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(v) s_i(v).$$

Applying this to $f \in I(V)$ we get that $(s_1(v), \ldots, s_n(v)) \in T_v V(K)$. That is, $s := (\mathrm{id}_v\, s_1, \ldots, s_n)$ is an algebraic vector field on V. That $\delta = \delta_s$ is clear from construction. $\qquad\square$

We see therefore, using the geometry-algebra correspondence, that to study irreducible algebraic vector fields we should consider the class of finitely generated integral k-algebras equipped with a k-linear derivation. The natural language for this is $\mathcal{L}_{k,\delta} = \{0, 1, +, -, \times, (\lambda_r)_{r \in k}, \delta\}$. We still cannot express the property of being finitely generated as a k-algebra, but we do have a universal theory, $T_{k,\delta}$, of integral k-algebras equipped with a k-linear derivation. It turns out that the class of existentially closed models of $T_{k,\delta}$ is also elementary; its theory, that of *differentially closed fields* which contain k and where the derivation vanishes on k, is denoted by DCF_k. Just as the study of algebraic varieties lead us to ACF, the study of algebraic vector fields motivates the model-theoretic consideration of DCF.

But what about going back again? That is, how do we recover the differential co-ordinate rings $(k[V], \delta_s)$ of algebraic vector fields that we were originally interested in? This time, studying the finitely generated substructures of models of DCF_k won't do the trick as these co-ordinate rings are outright finitely generated as k-algebras, and not just as differential k-algebras. Nevertheless, the differential rings

$(k[V], \delta_s)$ can be detected from DCF_k, as the *finite dimensional* substructures of models. More on this later.

In these lectures I will be focusing on DCF. However, while we have this template set-up, let us consider one other variant breifly. Instead of expanding $V \subseteq \mathbb{A}^n$ by an algebraic vector field, consider instead an *algebraic dynamical system* on V, i.e., a dominant morphism $\phi : V \to V$ over k. Recall that being dominant means $\phi(V(K)) \subseteq V(K)$ is Zariski dense, for any (equivalently some) algebraically closed field extension $K \supseteq k$. Again, I take for granted the intrinsic interest in algebraic dynamical systems (V, ϕ). To understand the algebraic counterpart of this geometric object we need to think about what ϕ induces on the co-ordinate $k[V]$. Writing $\phi = (f_1, \ldots, f_n)$ with $f_1, \ldots, f_n \in k[x]$, we obtain a k-linear endomorphism $\phi^* : k[V] \to k[V]$ given by $\phi^*(g) = g(f_1, \ldots, f_n)$. That is, viewing g as a function on V, $\phi^*(g)$ is obtained by pre-composing with ϕ. In diagrams:

$$
\begin{array}{ccc}
V & \xrightarrow{\phi} & V \\
& {\scriptstyle \phi^*(g)} \searrow & \downarrow {\scriptstyle g} \\
& & \mathbb{A}^1
\end{array}
$$

Indeed, the geometry-algebra correspondence is a functor and this is how it acts on endomorphism. The fact that ϕ is dominant implies that ϕ^* is injective: if $\phi^*(g) = 0$ then g must vanish on the image of ϕ, which being Zariski dense forces $g = 0$. The algebraic counterpart to (V, ϕ) is the *difference k-algebra* $(k[V], \phi^*)$. And again, all such arise in this way. That is, every injective k-linear endomorphism of $k[V]$ is of the form ϕ^* for some algebraic dynamical system ϕ on V over k. To study irreducible affine algebraic dynamical systems over k is, therefore, to study finitely generated integral k-algebras equipped with an injective k-linear endomorphism. The latter are studied model theoretically by the universal theory $T_{k,\sigma}$, in the language $\mathcal{L}_{k,\sigma} = \{0, 1, +, -, \times, (\lambda)_{r \in k}, \sigma\}$, of integral k-algebras equipped with an injective k-linear endomorphism. The existentially closed models of $T_{k,\sigma}$ are axiomatisable, by the theory ACFA_k of *difference closed fields* extending (k, id). So, the same path leading from varieties to ACF, and from algebraic vector fields to DCF, takes us from algebraic dynamical systems to ACFA.

2. D-VARIETIES

In the first lecture we worked exclusively over a fixed field k on which our derivations were assumed to act trivially. This is the so-called *autonomous* situation. But the constraint is somewhat artificial, at least form the model-theoretic point of view. In any case, even if we are interested primarily in the autonomous case, we will sometimes have to take base extensions to nontrivial differential fields in order to see the full geometric picture. So, let us now fix a differential field (k, δ) of characteristic zero.

Definition 2.1. Suppose $V \subseteq \mathbb{A}^n$ is an affine variety over k. By the *prolongation* of V over (k, δ) we mean the subvariety $\tau V \subseteq \mathbb{A}^{2n}$ over k defined, in co-ordinates $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$, by

$$
\begin{aligned}
f(x) &= 0 \\
f^\delta(x) + \sum_{i=1}^{n} \frac{\partial f}{\partial x_i} y_i &= 0
\end{aligned}
$$

for all $f \in I(V)$. Here f^δ denotes the polynomial obtained form f by applying δ to its coefficients. The projection onto the x co-ordinates gives us a surjective morphism $\pi : \tau V \to V$.

An affine D-*variety over* (k, δ) is a pair (V, s) where $V \subseteq \mathbb{A}^n$ is an affine variety over k and $s : V \to \tau V$ is a morphism over k which is a section to π.

Prolongations are the appropriate modification of tangent bundles in the presence of a derivation on the base field. In particular, if $\delta = 0$ on k then $\tau V = TV$. In general, for each $v \in V$, the fibre of the prolongation, $\tau_v V$, is a coset in $\mathbb{A}^n$ of the tangent space $T_v V$. The point here is that when working in the possibly nonautonomous case, the basic geometric objects of interest are D-varieties rather than vector fields.

Where do the equations for the prolongation space come from? The key differential-algebraic fact is the following elementary computation, the autonomous case of which was implicit in the proof of Proposition 1.3.

Lemma 2.2. *Suppose* (R, δ) *is a differential ring,* $f \in R[x]$ *is a polynomial in* $x = (x_1, \ldots, x_n)$, *and* $a = (a_1, \ldots, a_n) \in R^n$. *Then*

$$\delta(f(a)) = f^\delta(a) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(a)\delta(a_i).$$

Proof. We give a sketch, leaving the computations to the reader. First prove the result, by induction on total degree, in the case when f is a monomial (and so, in particular $f^\delta = 0$). Next, consider a polynomial with only one term, say $f = bg$ where g is a monomial and $b \in R$. The result follows for such f using the Leibniz rule and the case of monomials. Finally, for the general case, note that all the operators involved in the desired identity – namely $f \mapsto \delta(f(a))$, $f \mapsto \frac{\partial f}{\partial x_j}(a)$ and $f \mapsto f^\delta(a)$ – are additive in f. So, as every polynomial is a sum of polynomials of the form already dealt with, the Lemma is proven. $\square$

Corollary 2.3. *Suppose* (V, s) *is a D-variety over* (k, δ) *and* $(K, \delta) \supseteq (k, \delta)$ *is a differential field extension. Then* $a \mapsto (a, \delta(a))$ *defines a map* $\nabla : V(K) \to \tau V(K)$.

Proof. It suffices to show that $f^\delta(a) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(a)\delta(a_i) = 0$ for each $f \in I(V) \subseteq k[x]$ and $a \in V(K)$. But, as $f(a) = 0$, and hence $\delta(f(a)) = 0$, this is just Lemma 2.2 applied to $(R, \delta) = (K, \delta)$. $\square$

Now, the geometry-algebra correspondence yields a bijective correspondence between D-varieties over (k, δ) and finitely generated reduced k-algebras equipped with a derivation extending δ. That is, we have the following generalisation of Proposition 1.3.

Proposition 2.4. *Suppose* $V \subseteq \mathbb{A}^n$ *a subvariety over k. If* $s = (\mathrm{id}_V, s_1, \ldots, s_n)$ *is a D-variety structure on V over k then there is a unique derivation δ_s on $k[V]$ extending δ on k such that $\delta_s(x_i + I(V)) = s_i$ for all $i = 1, \ldots, n$.*

Moreover, every derivation on $k[V]$ extending δ on k is of the form δ_s for some D-variety s on V over k.

Proof. The proof of Proposition 1.3 readily generalises to this nonautonomous setting, and I leave the details to you. $\square$

Fix a D-variety (V, s) over (k, δ).

Definition 2.5. A *D-subvariety* of (V, s) is an algebraic subvariety $W \subseteq V$ over k such that s restricts to a D-variety structure on W. That is, $s \restriction_W \colon W \to \tau W$.

It is worth thinking about what this means algebraically. We have a differential structure δ_s on $k[V]$ induced by s, given by Proposition 2.4, and we have an ideal $I(W) \subseteq k[V]$ that defines the subvariety W.

Lemma 2.6. *A subvariety $W \subseteq V$ is a D-subvariety of (V, s) if and only if $I(W)$ is a δ_s-ideal of $k[V]$.*

Proof. We are being a bit imprecise here. If $V \subseteq \mathbb{A}^n$ and $x = (x_1, \ldots, x_n)$ are co-ordinates for affine n-space then $I(W)$ is an ideal of $k[x]$ containing $I(V)$. But these are in bijective correspondence with ideals of $k[V] = k[x]/I(V)$, and it is in this sense that we view $I(W)$ in $k[V]$.

Write $s = (\mathrm{id}_V(s_1, \ldots, s_n)$ and each $s_i = g_i + I(V)$ for some $g_i \in k[x]$. Suppose $f \in I(W)$. Then, by the constructuion of δ_s and Lemma 2.2, we have

$$\delta_s(f + I(V)) = f^\delta + \sum_{i=1}^n \frac{\partial f}{\partial x_i} f g_i + I(V).$$

So $I(W)$ is a δ_s-ideal if and only if the right-hand-side is in $I(W)$ for all $f \in I(W)$. That is, we need to show that the following are equivalent:

(i) W is a D-subvariety,
(ii) for all $f \in I(W)$, $K \supseteq k$ a field extension, and all $a \in W(K)$,

$$f^\delta(a) + \sum_{i=1}^n \frac{\partial f}{\partial x_i} f(a) g_i(a) = 0.$$

But, by definition τW is defined by $f^\delta + \sum_{i=1}^n \frac{\partial f}{\partial x_i} f y_i$ as f varies in $I(W)$, and the g_i define s. So (ii) is precisely expressing that s maps W to τW, which is what it means to be a D-subvariety. $\square$

Let us a say a word about base extension. Model-theoretically we are used to passing to a larger parameter set without comment. But in algebraic geometry it is more standard (and prudent) to distinguish notationally between the given variety V over k and its base extension V_K to a field extension $K \supseteq k$. The co-ordinate ring of V_K is the tensor product $k[V] \otimes_k K$, and the ideal of V_K is the extension ideal $I(V)K[V_K]$. So V_K is just the variety V viewed as being over K rather than k. Suppose, now, that $(K, \delta) \supseteq (k, \delta)$ is a differential field extension. Then prolongations commute with base extension: $\tau(V_K) = (\tau V)_K$. This follows from the fact (which we have not stated nor proved) that in the defining equations of the prolongation of V, given in Definition 2.1, we could have restricted to any fixed set of generators for $I(V)$ instead of ranging over all $f \in I(V)$. I leave this as an exercise for you. It follows that (V_K, s_K) is naturally a D-variety over (K, δ). We will usually continue the model-theoretic habit of dropping these subscripts, and simply view (V, s) as a D-subvariety over (K, δ) as well. This allows us to use terminology like "a D-subvariety of (V, s) over K" when what is really meant is a D-subvariety of (V_K, s_K).

3. Finite dimensional types

Now let us pass to a differentially closed field. These were discussed (motivated) in the first lecture, but I will deviate slightly from the conventions established there. We will use the simpler language $\mathcal{L}_\delta := \{0, 1, +, -, \times, \delta\}$ of differential rings, rather than, say, the language $\mathcal{L}_{\mathbb{Q},\delta}$ of differential $\mathbb{Q}$-algebras. And instead of working with the theory $T_{\mathbb{Q},\delta}$ of integral $\mathbb{Q}$-algebras equipped with a (necessarilly $\mathbb{Q}$-linear) derivation, we consider the (still universal) $\mathcal{L}_\delta$-theory T_δ of differential integral domains of characteristic zero. Note that if $(R, \delta) \models T_\delta$, and $R_{\mathbb{Q}}$ denotes the localisation at $\mathbb{Z} \setminus \{0\}$, then δ extends uniquely to $R_{\mathbb{Q}}$ and $(R_{\mathbb{Q}}, \delta) \models T_{\mathbb{Q},\delta}$. So this change is really very harmless. In particular, T_δ and $T_{\mathbb{Q},\delta}$ have the "same" existentially closed models. We denote the theory of these existentially closed models by DCF_0, and call it the the theory of *differentially closed fields of characteristic zero*. Note that I never proved that the class of existentially closed models of T_δ (or $T_{\mathbb{Q},\delta}$ for that matter) is elementary; it is, but I forego all discussion of the axiomatisation here. Instead we will use existential closedness directly whenever we want to establish any properties of the models of DCF_0. For example, let me point out that differentially closed fields are algebraically closed:

Lemma 3.1. *If $(K, \delta) \models \mathrm{DCF}_0$ then K is an algebraically closed field.*

Proof. Despite the terminology, we have not yet observed that K is a field. So let's do that first. To see that every nonzero $a \in K$ is invertible, just apply existential closedness to the formula $\phi(x)$ given by $xa = 1$. Indeed, the quotient rule – itself an immediate consequence of the Leibniz rule – gives us a unique extension of δ to the localisation of K at a, which thus yields a model of T_δ extending (K, δ) in which $\phi(x)$ has a realisation. Hence $\phi(x)$ is already realised in (K, δ), as desired.

As similar approach gives that K is algebraically closed. Let $K(a)$ be a simple algebraic extension of K and f the minimal polynomial of a over K. The there is a unique extension of δ to $K(a)$. Indeed, uniqueness – which we don't actually need in this proof – is by Lemma 2.2 which dictates that $\delta(a)$ must equal $-\frac{f^\delta(a)}{f'(a)}$. In any case, existence takes a bit more work. First, view the derivation δ on K as being $K(a)$-valued. Then use the freeness of the polynomial ring to extend δ to a derivation $\delta : K[x] \to K(a)$ by setting $\delta(x) := -\frac{f^\delta(a)}{f'(a)}$. Then observe, using Lemma 2.2, that $\delta(f) = 0$. We thus obtain an induced derivation $\delta : K[x]/(f) \to K(a)$. The natural identification of $K[x]/(f)$ with $K(a)$ gives us our desired extension of (K, δ) to $(K(a), \delta)$. The formula $f(x) = 0$ has a realisation in $(K(a), \delta) \models T_\delta$, and hence by existential closedness, in (K, δ). This forces $K(a) = K$, and we have shown that K is algebraically closed. $\square$

My goal for the rest of this lecture is to show how D-varieties capture precisely the finite dimensional fragment of DCF_0. Let us fix from now on a sufficiently saturated model $(\mathbb{K}, \delta) \models \mathrm{DCF}_0$. By convention all differential fields we consider are differential subfields of $\mathbb{K}$ of cardinality strictly less than the level of saturation, unless explicitly stated otherwise. (In fact, DCF_0 is ω-stable and hence admits saturated models of arbitrary large cardinality, so we can assume $\mathbb{K}$ is saturated in its own cardinality.) Similarly all parameters sets are assume to be small unless explicitly stated otherwise.

Fix a D-variety (V, s) over a differential field (k, δ).

Definition 3.2. By a *D-point of* (V, s) we mean a point $a \in V(\mathbb{K})$ such that the singleton $\{a\}$ is a *D*-subvariety of (V, s) over $(\mathbb{K}, \delta)$. We denote the set of all *D*-points by $(V, s)^\sharp$.

Note that we are implicitly taking a base extension as discussed earlier; the singleton $\{a\}$ is really a *D*-subvariety of the base extension of (V, s) to $(\mathbb{K}, \delta)$. In order to preserve our convention that all parameters sets be small, we could instead take the base extension of (V, s) to the differential field generated by a over k, namely $k\langle a \rangle := k(a, \delta a, \delta^2 a, \dots)$. In any case, we have the following characterisation of $(V, s)^\sharp$.

Lemma 3.3. $(V, s)^\sharp = \{a \in V(\mathbb{K}) : s(a) = \nabla(a)\}$.

Proof. The lemma follows immediately from the definitions once we observe that, for any $a \in V(\mathbb{K})$, the prolongation of $\{a\}$, viewed as a subvariety of V over $F := k\langle a \rangle$, is precisely $\{\nabla(a)\}$. And this can be checked directly: Suppose $V \subseteq \mathbb{A}^n$ and $a = (a_1, \dots, a_n)$. Then the ideal of $\{a\}$ is generated by the linear polynomials $x_1 - a_1, x_2 - a_2, \dots, x_n - a_n$ in $F[x_1, \dots, x_n]$. So $\tau\{a\} \subseteq \mathbb{A}^{2n}$ is defined in the variables (x, y) by the equations $x_j - a_j = 0$, for each $j = 1, \dots, n$, along with the equation $(x_j - a_j)^\delta + \sum_{i=1}^{n} \frac{\partial(x_j - a_j)}{\partial x_i} y_i = 0$. But this latter is just $y_j - \delta(a_j) = 0$. So the prolongation of $\{a\}$ is the singleton $\{(a_1, \dots, a_n, \delta a_1, \dots, \delta a_n) = \nabla(a)\}$. $\square$

It follows that $(V, s)^\sharp$ is a *k*-definable set in $(\mathbb{K}, \delta)$. Indeed, writing the section as $s = (\mathrm{id}_V, s_1, \dots, s_n)$, and each $s_i = g_i + I(V)$ for some $g_1, \dots, g_n \in k[x]$, we have that $(V, s)^\sharp$ is defined by the formula $(x \in V) \wedge \bigwedge_{i=1}^{n} \left(\delta x_i = g_i(x) \right)$. Note that this formula is of a particularly simple form. First of all, it is quantifier-freee. But this is not surprising as in fact DCF$_0$ admits quantifier elimination. Moreover, it is a conjunction of δ-*polynomial* equations; namely polynomial equations over k in the variables $x, \delta x, \delta^2 x, \dots$. Such definable sets are called *Kolchin closed* in analogy with the Zariski closed sets of algebraic geometry. In addition, $(V, s)^\sharp$ is of *order* 1 in the sense that only the first derivative of the variables actually appear.

Next, assuming that V is irreducible, we associate to (V, s) a certain complete type in $(\mathbb{K}, \delta)$. Let $p(x)$ be the collection of formulae over k asserting that x is a *D*-point of (V, s) not contained in any proper subvariety of V over k.

I claim that $p(x)$ is consistent. That is, given a proper subvariety $W \subsetneq V$ over k, there is a *D*-point of (V, s) in $U := V \setminus W$. Our assumption that V is irreducible ensures that U is a dense Zariski open subset of V. As you might expect, we need to use existential closedness of $(\mathbb{K}, \delta)$ to show that there is a *D*-point in U, or indeed that there are any *D*-points of (V, s) at all. This is done as follows: Using Proposition 2.4, let δ_s be the derivation on $k[V]$ extending δ on k that is induced by s. Let $a := x + I(V)$ be the generators of $k[V] = k[x]/I(V)$ coming form the variables of the ambient polynomial ring. By definition of δ_s we have that $\delta_s(a_i) = s_i$ for all $i = 1, \dots, n$, in the differential ring $(k[V], \delta_s)$. That is, a is a realisation in $(k[V], \delta_s)$ of the formula defining $(V, s)^\sharp$. Moroever, a also realises $x \notin W$ as $I(W) \supsetneq I(V)$ since W is a proper subvariety of V. That is, the formula defining $(V, s)^\sharp \cap U$ has a realisation in $(k[V], \delta_s)$. Now, as V is irreducible, $k[V]$ is an integral domain, and hence $(k[V], \delta_s)$ is a model of T_δ extending (k, δ). It

follows by existential closedness that the formula defining $(V, s)^\sharp \cap U$ must have a realisation in $(\mathbb{K}, \delta)$, as desired.

Next, I claim that $p(x)$ determines a complete type. Using the fact that DCF_0 admits quantifier elimination (whose proof I also forego!), it suffices to prove that if $a \models p(x)$ then $\mathrm{tp}_{\mathcal{L}}(a, \delta a, \delta^2 a, \ldots / k)$ is determined, where recall that $\mathcal{L}$ is just the language of rings (as opposed to the language $\mathcal{L}_\delta$ of differential rings). But as a is a D-point of (V, s), the derivative δa is given by polynomials in a over k, and hence so is $\delta^\ell a$ for all $\ell \geq 1$ by Lemma 2.2. So $\mathrm{tp}_{\mathcal{L}}(a, \delta a, \delta^2 a, \ldots / k)$ is determined by $\mathrm{tp}_{\mathcal{L}}(a/k)$. By quantifer elimination in ACF_0, this is in turn determined by $\mathrm{Z\text{-}loc}(a/k)$, the *Zariski locus* of a over k; namely, the smallest Zariski closed subset of V over k that contains a. But as $p(x)$ ensures that a is not contained in any proper sunbvariety of V over k, we must have that $\mathrm{Z\text{-}loc}(a/k) = V$. So $\mathrm{tp}_{\mathcal{L}}(a/k)$ is determined.

Definition 3.4. We call this $p(x)$ the *generic type of* (V, s) *over* k. It is the type asserting that x is a D-point of (V, s) and that the Zariski locus of x over k is V. We call a realisation of $p(x)$ a *generic D-point of* (V, s) *over* k.

We have thus associated a complete type to every irreducible D-variety. Not all complete types arise in this way. To see this, let us introduce dimension for types as follows:

Definition 3.5. Suppose (k, δ) is a differential field and $p \in S(k)$ is a complete type over k. By the *dimension of* p, denoted by $\dim(p)$, we mean the non-decreasing sequence of non-negative integers

$$(\mathrm{tr.\,deg.}(a/k), \mathrm{tr.\,deg.}(a, \delta a/k), \mathrm{tr.\,deg.}(a, \delta a, \delta^2 a/k), \ldots)$$

where $a \models p$. We also write this as $\dim(a/k)$. If this sequence eventually stabilises then we say that p is *finite dimensional* and we write $\dim(p) = d$ where d is that eventual value of the sequence.

Note that p is finite dimensional if and only if for some (equivalently any) $a \models p$ we have that the differential field $k\langle a \rangle = k(a, \delta a, \delta^2 a, \ldots)$ generated by a over k has finite transcendence degree, and in that case $\dim(p) = \mathrm{tr.\,deg.}(k\langle a \rangle / k)$.

Lemma 3.6. *If (V, s) is an irreducible D-variety over (k, δ) then its generic type over k is finite dimensional and of dimension $\dim V$.*

Proof. If a is a generic D-point of (V, s) over k then $\mathrm{tr.\,deg.}(a/k) = \dim V$ since $V = \mathrm{Z\text{-}loc}(a/k)$. On the other hand, $k(a, \delta a, \ldots, \delta^\ell a) = k(a)$ for all $\ell \geq 1$ as a is a D-point and hence δa is a polynomial in a over k. So $\dim(a/k) = (\dim V, \dim V, \ldots)$. $\square$

In fact, every finite dimensional type arises in this way. Well, at least up to inter-definability. Here, we say that $p, q \in S(k)$ are *interdefinable* if for all (equivalently some) $a \models p$ there is $b \models q$ such that $\mathrm{dcl}(ka) = \mathrm{dcl}(kb)$. Using quantifier-elimination for DCF_0 one can show that $\mathrm{dcl}(ka) = k\langle a \rangle$, and hence finite dimensionality is an interdefinability invariant of complete types. But note that dimension itself is *not* an interdefinability invariant: if $a \in \mathbb{K}$ is differentially-transcendental over k in the sense that $(a, \delta a, \delta^2 a, \ldots)$ is an algebraically independent sequence then $\dim(a/k) = (1, 2, 3, \ldots)$ while $\dim(a, \delta a/k) = (2, 3, 4, \ldots)$, though $\mathrm{tp}(a/k)$ and $\mathrm{tp}(a, \delta a/k)$ are interdefinable.

Theorem 3.7. *Suppose (k, δ) is a differential field. Every finite dimensional complete type over k is interdefinable with the generic type of an irreducible D-variety over (k, δ).*

Proof. I will give only a sketch, in the special case of 1-types, from which you will see how to proceed in general and in detail. We are given a type $p = \mathrm{tp}(a/k)$ where $a \in \mathbb{K}$ and $k\langle a \rangle$ is of finite transcendence degree over k. It follows that, for some $\ell \geq 0$, $\delta^\ell a \in k(a, \delta a, \ldots, \delta^{\ell-1}a)^{\mathrm{alg}}$. Let P be the minimal polynomial of $\delta^\ell a$ over $k(a, \delta a, \ldots, \delta^{\ell-1}a)$. Since $P(\delta^\ell a) = 0$, differentiating both sides we get from Lemma 2.2 that
$$0 = \delta(P(\delta^\ell a)) = P^\delta(\delta^\ell a) + P'(\delta^\ell a)\delta^{\ell+1}a.$$
By minimality, $P'(\delta^\ell a) \neq 0$, and hence we get that
$$\delta^{\ell+1}a = -\frac{P^\delta(\delta^\ell a)}{P'(\delta^\ell a)} = f(a, \delta a, \ldots, \delta^\ell a)$$
for some rational function $f \in k(x^{(0)}, x^{(1)}, \ldots, x^{(\ell)})$. Now let $b := (a, \delta a, \ldots, \delta^\ell a)$. Then
$$\nabla(b) = (a, \delta a, \ldots, \delta^\ell a, \delta a, \delta^2 a, \ldots, \delta^{\ell+1}a) = (b, \pi_1(b), \pi_2(b), \ldots, \pi_\ell(b), f(b))$$
where the π_i are the ith co-ordinate projections. That is, $\nabla(b) = s(b)$ where $s = (\mathrm{id}, \pi_1, \ldots, \pi_\ell, f)$. If we set $V := \mathrm{Z\text{-}loc}(b/k) \setminus \mathrm{pole}(f) \subseteq \mathbb{A}^{\ell+1}$, then $s : V \to \tau V$ is a regular section, b is a generic D-point of (V, s) over k, and $\mathrm{tp}(a/k)$ and $\mathrm{tp}(b/k)$ are interdefinable. I am cheating, of course, because V is not a closed subvariety of $\mathbb{A}^{\ell+1}$, but by working with one more variable this can be remedied. $\square$

The study of the birational geometry of D-varieties thus coincides with the model theory of the finite dimensional fragment of DCF_0.

4. STABILITY AND INDEPENDENCE

I have omitted proofs of some of the most fundamental properties of DCF_0; in particular, that it exists (i.e., the fact that the existentially closed models of T_δ form an elementary class) and that it admits quantifier elimination. This is largely because you can find the proofs elsewhere, and I had nothing to add. In this lecture I want to discuss some further model theoretic properties, around stability, and this time I will give at least some proofs.

Before talking about stability, let me say a few words about the elimination of imaginaries, another important property that DCF_0 enjoys. This is the statement that every definable set D has a *code*; that is, a finite tuple e such that for all $\sigma \in \mathrm{Aut}(\mathbb{K}, \delta)$, $\sigma(D) = D$ if and only if $\sigma(e) = e$. Equivalently, for some formula $\phi(x, y)$, $D = \phi(\mathbb{K}, e)$ but $D \neq \phi(\mathbb{K}, e')$ for any $e' \neq e$. That is, a code for a definable set is a kind of minimal and canonical parameter. Using quantifier elimination and the noetherianity of the Kolchin topology (both facts not proved here), it is not hard to reduce the verification that all definable sets have codes to showing that all Kolchin closed sets have codes. Now, a Kolchin closed set is of the form
$$D = \{a \in \mathbb{K}^n : \nabla_\ell(a) := (a, \delta a, \delta^2 a, \ldots, \delta^\ell a) \in V(\mathbb{K})\}$$
for some $\ell \geq 0$ and some algebraic subvariety $V \subseteq \mathbb{A}^{(\ell+1)n}$. Moreover, replacing V by the Zariski closure of $\nabla_\ell(D)$, we may assume that $\nabla_\ell(D)$ is Zariski dense in V. From algebraic geometry we know that V has a minimal field of definition, say L,

that L is finitely generated, say $L = \mathbb{Q}(e)$, and that e is a code for $V(\mathbb{K})$ in the pure field structure $\mathbb{K}$. I claim that e will in fact be a code for D in $(\mathbb{K}, \delta)$. It is clear that D is defined over e, and so it suffices to show that if $\sigma \in \mathrm{Aut}(\mathbb{K}, \delta)$ preserves D then $\sigma(e) = e$. But as σ commutes with δ, if $\sigma(D) = D$ then $\sigma(\nabla_\ell(D)) = \nabla_\ell(D)$, and hence σ preserves the Zariski closure V, which in turn forces $\sigma(e) = e$.

Now let us pass to stability.

Theorem 4.1. DCF_0 *is ω-stable: there are only countably many types over count-ably many parameters.*

Proof. It suffices to count 1-types. Actually, the proof sketch I gave of Theorem 3.7 already suggests how to count the 1-types over a differential field (k, δ), and we follow that suggestion now. First of all, there is a unique 1-type of a *differentially transcendental* element over k; that is, of an element $a \in \mathbb{K}$ such that $(a, \delta a, \delta^2 a, \dots)$ is an algebraically independent sequence over k. So it remains to count *differentially algebraic* 1-types (namely those that are not differentially transcendental).

Suppose therefore that $a \in \mathbb{K}$ is differentially algebraic over k. Let $\ell \geq 0$ be least such that $\delta^\ell a \in k(a, \delta a, \dots, \delta^{\ell-1} a)^{\mathrm{alg}}$ and let $P(t)$ be the minimal polynomial of $\delta^\ell a$ over $k(a, \delta a, \dots, \delta^{\ell-1} a)$. After clearing denominators we can write $P = g(a, \delta a, \dots, \delta^{\ell-1} a, t)$ where $g \in k[x^{(0)}, x^{(1)}, \dots, x^{(\ell)}, t]$. I claim that $\mathrm{tp}(a/k)$ is determined by the pair (ℓ, g). But before proving this let us observe that $k\langle a \rangle = k(a, \delta a, \dots, \delta^\ell a)$. Indeed, we saw in the proof of Theorem 3.7 that $\delta^\ell a$ being algebraic over $k(a, \delta a, \dots, \delta^{\ell-1} a)$ implies that $\delta^{\ell+1} a$ is contained in $k(a, \delta a, \dots, \delta^\ell a)$. In particular, $\delta^{\ell+1} a$ is algebraic over $k(a, \delta a, \dots, \delta^\ell a)$ and hence

$$\delta^{\ell+2} a \in k(a, \delta a, \dots, \delta^{\ell+1} a) = k(a, \delta a, \dots, \delta^\ell a).$$

Iterating gives us the desired fact that $k\langle a \rangle = k(a, \delta a, \dots, \delta^\ell a)$.

Suppose now that $b \in \mathbb{K}$ is differentially algebraic over k and gives rise to the same data (ℓ, g). I want to show that $\mathrm{tp}(a/k) = \mathrm{tp}(b/k)$. To do so, I will exhibit a differential-field-isomorphism from $k\langle a \rangle$ to $k\langle b \rangle$ over k, that takes a to b, and this will suffice by quantifier elimination. But we know that

$$k\langle a \rangle = k(a, \delta a, \dots, \delta^\ell a)$$

and

$$k\langle b \rangle = k(b, \delta a, \dots, \delta^\ell b).$$

So it suffices to exhibit a field-isomorphism $\alpha : k(a, \delta a, \dots, \delta^\ell a) \to k(b, \delta b, \dots, \delta^\ell b)$, over k, which satisfies $\alpha(\delta^i a) = \delta^i b$ for all $i = 0, \dots, \ell+1$. (It is not a typo here that we have to check all the way up to $i = \ell + 1$.) Indeed, I am using the fact, which I leave to you to check, that if you have a field isomorphism between differential fields which commutes with the derivation on the field-generators of the domain, then it must be a differential-field-isomorphism. (Hint: This too rests on the infinitely useful Lemma 2.2.)

First of all, $\mathrm{tp}_{\mathcal{L}}(a, \delta a, \dots, \delta^{\ell-1} a/k) = \mathrm{tp}_{\mathcal{L}}(b, \delta b, \dots, \delta^{\ell-1} b/k)$ as it is the field-type of an algebraically independent ℓ-tuple. So we have a field isomorphism $\alpha : k(a, \delta a, \dots, \delta^{\ell-1} a) \to k(b, \delta b, \dots, \delta^{\ell-1} b)$, over k, taking $\delta^i a$ to $\delta^i b$ for all $i = 0, \dots, \ell-1$. We want to extend α to $\delta^\ell a$. To do so, note that, by construction, $P = g(a, \delta a, \dots, \delta^{\ell-1} a, t)$ is an irreducible polynomial over $k(a, \delta a, \dots, \delta^{\ell-1} a)$ of which $\delta^\ell a$ is a root. On the other hand, $\delta^\ell b$ is a root of $Q := g(b, \delta b, \dots, \delta^{\ell-1} b, t)$, and Q is the transform of P by α. So we can extend α to

$$\alpha : k(a, \delta a, \dots, \delta^\ell a) \to k(b, \delta b, \dots, \delta^\ell b)$$

by sending $\delta^\ell a$ to $\delta^\ell b$.

It remains only to check that $\alpha(\delta^{\ell+1}a) = \delta^{\ell+1}b$. But this also follows from the proof of Theorem 3.7, where we saw that $\delta^{\ell+1}a = -\frac{P^\delta(\delta^\ell a)}{P'(\delta^\ell a)}$ and $\delta^{\ell+1}b = -\frac{Q^\delta(\delta^\ell b)}{Q'(\delta^\ell b)}$. I leave it to you to check that α takes P^δ to Q^δ and P' to Q'. Hence α takes $\delta^{\ell+1}a$ to $\delta^{\ell+1}b$, as desired.

We have proved that the differentially algebraic 1-type $\mathrm{tp}(a/k)$ is determined by the pair (ℓ, g). If k is countable then there are only countably many possible such pairs, and hence only countably many differentially algebraic 1-types over k (and only one differentially transcendental 1-type). As every countable set of parameters is contained in a countable differential field (namely the differential field it generates), it follows that over countably many parameters we have only countably many complete 1-types. That is, DCF_0 is ω-stable. $\qquad\square$

What I like about the above proof is that it uses very little differential algebra. In particular, somewhat unexpectedly, no use is made of the Ritt-Raudenbush basis theorem which says that every radical differential ideal in a differential polynomial ring over a differential field is finitely generated (as a radical differential ideal). Differential algebra is a useful and beautiful subject, but it is interesting to note how little of it one really needs to do model theory in DCF_0.

Once we have ω-stability, the full machinery of geometric stability theory becomes available. In particular we have the good behaviour of Shelah's nonforking independence. I will not give the abstract definition of nonforking, but rather specialise to what it means in DCF_0.

Definition 4.2. Given a tuple a and subsets $B \subseteq A$ of $\mathbb{K}$, we say that a *is independent from A over B*, denoted by $a \underset{B}{\downarrow} A$, to mean that $\dim(a/\mathrm{dcl}(A)) = \dim(a/\mathrm{dcl}(B))$. In this case we also say that $\mathrm{tp}(a/A)$ a *does not fork over B* or that $\mathrm{tp}(a/A)$ is a *nonforking extension* of $\mathrm{tp}(a/B)$.

Recall that the dimension of a type over a differential field was defined in 3.5 as a certain infinite sequence of nondecreasing integers, and it is as such that equality is to be understood above. Namely, if $F := \mathrm{dcl}(A)$ and $k := \mathrm{dcl}(B)$, then $a \underset{B}{\downarrow} A$ if and only if $\mathrm{tr.\,deg.}(a, \delta a, \ldots, \delta^\ell a/F) = \mathrm{tr.\,deg.}(a, \delta a, \ldots, \delta^\ell a/k)$, for all $\ell \geq 0$. It is also worth pointing out that this is equivalent to saying that $k\langle a \rangle$ is *algebraically disjoint*, in the sense of pure field theory, from F over k.

Maybe the best way to see that this agrees with Shelah's nonforking is to prove that it satisfies the usual desired properties (nontriviality, finite character, automorphism invariance, symmetry, transitivity, existence of nonforking extensions, uniqueness of nonforking extensions over algebraically closed sets), using the analogous properties for algebraic independence in pure fields, and then use the fact that Shelah's nonforking in stable theories is characterised among all abstract independence relations by these properties. But we do not go into that here.

How does this specialise to generic types of D-varieties (and hence, by Theorem 3.7, to finite dimensional types)? Here is the simple answer:

Proposition 4.3. *Suppose (V, s) is a D-variety over (k, δ) with V absolutely irreducible[1]. Let a be a generic D-point of (V, s) over k. Then, for any differential field extension $F \supseteq k$, $a \underset{k}{\downarrow} F$ if and only if a is a generic D-point of (V, s) over F.*

---

[1]A variety is *absolutely irreducible* if its base extension to any field extension is irreducible.

Proof. By Lemma 3.6, the dimension of the generic type of (V, s) over k is $\dim V$, and so is the dimension of the generic type of (V, s) over F. (Note that we are using absolute irreducibility here so that we can apply Lemma 3.6 to the base extension of (V, s) to F, which remains irreducible.) This proves the right-to-left direction. For the converse, assume $a \underset{k}{\downarrow} F$. As we already know that a is a D-point of (V, s), it remains to verify that the Zariski locus of a over F, say W, is equal to V. But W is a Zariski closed subset of V over F and it is of dimension $\operatorname{tr.deg.}(a/F) = \operatorname{tr.deg.}(a/k) = \dim V$. Hence, by the irreducibility of V over F, we must have $W = V$. □

The assumption in the above proposition that V be absolutely irreducible may seem a little unnatural. At the very least we should give a model-theoretic explanation of absolute irreducibility:

Proposition 4.4. *Suppose (V, s) is an irreducible D-variety over (k, δ). Then V is absolutely irreducible if and only if the generic type of (V, s) over k is* stationary: *it has a unique nonforking extension to every set containing k.*

Proof. If V is absolutely irreducible, then by Proposition 4.3 the nonforking extension of its generic type, p, over k, to a set $A \supseteq k$, is precisely the generic type of (V, s) over $F := \operatorname{dcl}(A)$. Conversely, suppose the generic type p of (V, s) over k is stationary and let $F := k^{\operatorname{alg}}$. For absolute irreducibility it suffices to verify that V is irreducible over F. Note that F is also a differential subfield of $(\mathbb{K}, \delta)$, since, as observed in the proof of Lemma 3.1, if $a \in k^{\operatorname{alg}}$ with minimal polynomial f then $\delta(a) = -\frac{f^{\delta}(a)}{f'(a)}$. Now, each irreducible component W of V over F is of dimension $\dim V$. I claim that it suffices to check that W is a D-subvariety of V over F. Indeed, in that case the generic type of $(W, s \restriction_W)$ over F would be the nonforking extension of p to F, by Proposition 4.3, and hence all irreducible components of V over F would share a generic type, forcing V itself to be irreducible over F.

So we have reduced to showing that the irreducible components of V over F are D-subvarieties. This is of independent interest. Recall, from algebraic geometry, that the irreducible components of a variety correspond to the minimal prime ideals containing the ideal of that variety. By Lemma 2.6, we need to show that those minimal prime ideals are δ_s-ideals. This follows from the following fundamental (but easy) fact of differential algebra: *If (F, δ) is a differential field and R is a finitely generated F-algebra equipped with a δ-ring structure extending (F, δ), and I is a radical δ-ideal of R, then every minimal prime ideal containing I is also a δ-ideal.* To see this, let $I = P_1 \cap \cdots \cap P_\ell$ be the prime decomposition of I. We need to observe that each of the P_j is a δ-ideal. That is, fixing $j = 1, \ldots, \ell$, and $a \in P_j$, we need to show that $\delta(a) \in P_j$. For each $r \neq j$ choose $b_r \in P_r \setminus P_j$ and let b be the product of all the b_r's. Then $ab \in I$, and hence $\delta(a)b + a\delta(b) = \delta(ab) \in I \subseteq P_j$. Since $a \in P_j$ this implies $\delta(a)b \in P_j$. But by construction $b \notin P_j$, and hence $\delta(a) \in P_j$ as desired. □

5. Around the constants

By the *field of constants* we mean $\mathcal{C} := \{a \in \mathbb{K} : \delta(a) = 0\}$. Note that $\mathcal{C}$ is the set of D-points of the trivial vector field on the affine line; that is, $\mathcal{C} = (\mathbb{A}^1, 0)^{\sharp}$. Here 0 here denotes the zero vector field on the affine line given by $a \mapsto (a, 0)$. It is

also easy to verify that $\mathcal{C}$ is a subfield of $\mathbb{K}$. Let us observe that it is algebraically closed: If $f(t)$ is the minimal polynomial over $\mathcal{C}$ of an element $a \in \mathbb{K}$ that is algebraic over $\mathcal{C}$, then by differentiating $0 = f(a)$ we get $0 = f'(a)\delta(a)$ as $f^\delta = 0$. But $f'(a) \neq 0$ by minimality, so that we must have $\delta(a) = 0$. This shows that $\mathcal{C}$ is relatively algebraically closed in $\mathbb{K}$. But we have seen that $\mathbb{K}$ is algebraically closed (Lemma 3.1), and hence so is $\mathcal{C}$.

Proposition 5.1. *The constants form a* stably embedded pure *algebraically closed field. That is, if $D \subseteq \mathcal{C}^n$ is definable in $(\mathbb{K}, \Delta)$ then it is definable in $(\mathcal{C}, +, \times)$.*

Proof. Note that for us "definable" means "definable with parameters". So the first step is to show that D is definable in $(\mathbb{K}, \Delta)$ with parameters from $\mathcal{C}$. We will use the fact (not proven in this course) that all types of a stable theory are definable. Suppose D is defined by $\phi(a, y)$ where $\phi(x, y)$ is a $\mathcal{L}_\delta$-formula and $a \in \mathbb{K}^m$, and let $\mathbf{p}(x) := \mathrm{tp}(a/\mathcal{C})$. I use boldface here because $\mathbf{p}$ is not a proper type according to our conventions where we are only to allow parameter sets that are of cardinality less then the level of saturation. But it is OK, such *global types* are also definable. In particular, we have a formula $\mathrm{d}_{\mathbf{p}}\phi(y)$ over $\mathcal{C}$ such that $\phi(x, b) \in \mathbf{p}$ if and only if $\models \mathrm{d}_{\mathbf{p}}\phi(b)$ for all $b \in \mathcal{C}^n$. But this implies that, for all $b \in \mathbb{K}^n$,

$$b \in D \iff \models \mathrm{d}_{\mathbf{p}}\phi(b) \wedge (\delta b = 0)$$

which shows that D is definable with parameters in $\mathcal{C}$.

Next, easy manipulations using quantifier elimination reduces us further to the case when D is defined by a conjunction $(\delta y = 0) \wedge (P(y, \delta y, \ldots, \delta^\ell y) = 0)$ where P is a polynomial with coefficients in $\mathcal{C}$. But that formula is equivalent to

$$(\delta y = 0) \wedge (P(y, 0, \ldots, 0) = 0),$$

and the latter defines a set which is clearly definable in $(\mathcal{C}, +, \times)$. $\square$

In this way, pure algebraic geometry lives definably in the finite dimensional fragment of DCF_0, precisely as the induced structure on the constants. (Notice that this is saying something different, more meaningful, than that $(\mathbb{K}, \Delta)$ is an expansion of the pure algebraically closed field $\mathbb{K}$.) The model theorist's approach can now be described as follows: Understand the fine structure of a finite dimensional type by studying its relationship to the field of constants. In other words, study differential-algebraic geometry in relation to the algebraic geometry living therein.

Such an approach will, of necessity, say nothing about those types that have no definable relationship to the constants. Here we have to be careful about what we might mean by "having no definable relation to the constants", in particular with respect to parameters. Consider, for example, the set $D \subseteq \mathbb{K}$ defined by the equation $\delta(x) = 1$. Note that $D = (\mathbb{A}^1, 1)^\sharp$ where 1 here denotes the constant vector field on the affine line given by $a \mapsto (a, 1)$. The elements of D are independent of the constants over the empty set. Indeed, if $d \in D$ and $F \subseteq \mathcal{C}$ is any subfield then, as $d \notin F^{\mathrm{alg}} \subseteq \mathcal{C}$, $\dim(d/F) = 1 = \dim(d/\mathbb{Q})$. It follows that $d \underset{\perp}{\,} F$, which expresses the fact that there are no nontrivial 0-definable relations between D and $\mathcal{C}$. On the other hand, if we fix a solution $d_0 \in D$, and let $k := \mathbb{Q}\langle d_0 \rangle$, then D is definably isomorphic to $\mathcal{C}$ over k as $D = d_0 + \mathcal{C}$. In fact, we have an isomorphism of D-varieties $(\mathbb{A}^1, 1) \to (\mathbb{A}^1, 0)$ over k, given by translation by $-d_0$. So parameters matter very much, leading to the following two natural implementations of "having no definable relation to the constants".

Definition 5.2. Suppose (k, δ) is a differential field and $p \in S(k)$ is a complete stationary type. We say that p is *weakly orthogonal to* $\mathcal{C}$, denoted by $p \perp^w \mathcal{C}$, to mean that whenever $a \models p$ and c is a tuple from $\mathcal{C}$ then $a \mathop{\smile}\limits_k c$. We say that p is *orthogonal to* $\mathcal{C}$, denoted $p \perp \mathcal{C}$, if every nonforking extension of p is weakly orthogonal to $\mathcal{C}$.

That is, $p \perp \mathcal{C}$ means that for any $B \supseteq k$, any $a \models p$ with $a \mathop{\smile}\limits_k B$, and any tuple c from $\mathcal{C}$, $a \mathop{\smile}\limits_B c$. For example, what follows from the above discussion is that the type over $\mathbb{Q}$ of any solution to $\delta(x) = 1$ is weakly $\mathcal{C}$-orthogonal, but not $\mathcal{C}$-orthogonal. In fact, $\delta(x) = 1$ fails orthogonality to the constants in a particularly strong way:

Definition 5.3. Suppose (k, δ) is a differential field and $p \in S(k)$ is a complete stationary type. We say that p is (*almost*) $\mathcal{C}$-*internal* if there is $B \supseteq k$, $a \models p$ with $a \mathop{\smile}\limits_k B$, and c a tuple from $\mathcal{C}$, such that $a \in \mathrm{dcl}(Bc)$ (respectively, $a \in \mathrm{acl}(Bc)$).

The type of any solution to $\delta(x) = 1$, while being weakly $\mathcal{C}$-orthogonal, is at the same time $\mathcal{C}$-internal.

Examples of $\mathcal{C}$-orthogonality (or of non-almost-$\mathcal{C}$-internality, for that matter) are harder to verify as they require considering all possible extensions of parameters. But examples do exist, even in order one: the type of any nonzero solution to $\delta(x) = \frac{x}{x+1}$ is orthogonal to the constants. Such equations were studied by Kolchin, Rosenlicht, and Shelah; see Dave Marker's treatment of the above equation in [5, §II.6].

Orthogonality to $\mathcal{C}$ implements the idea of having essentially no definable relation to the constants (even after passing to additional parameters), while almost $\mathcal{C}$-internality captures the opposite extreme of having a very significant definable relation with the constants (after possibly adding parameters). There is, of course, a lot of room in-between.

Specialising to finite dimensional types, let us give the geometric meaning of $\mathcal{C}$-orthogonality and almost $\mathcal{C}$-internality for the generic types of D-varieties. To do so we need a little more differential-algebraic geometry. The following notions are very natural and I could have, probably should have, discussed them in the second lecture.

Definition 5.4. Suppose (V_1, s_1) and (V_2, s_2) are D-varieties over a differential field (k, δ). Then a D-*rational map* $f : (V_1, s_1) \to (V_2, s_2)$ is a rational map $f : V_1 \to V_2$ over k such that $\tau f \circ s_1 = s_2 \circ f$. That is, the diagram

$$
\begin{array}{ccc}
\tau V_1 & \xrightarrow{\ \tau f\ } & \tau V_2 \\
{\scriptstyle s_1}\big\uparrow & & \big\uparrow{\scriptstyle s_2} \\
V_1 & \xrightarrow{\ \ f\ \ } & V_2
\end{array}
$$

of rational maps, commutes.

In the above definition I am using implicitly the fact that prolongation is a functor, so that $\tau f : \tau V_1 \to \tau V_2$ makes sense. When f is a morphism it is quite clear how to define τf, namely: letting $\Gamma \subseteq V_1 \times V_2$ be the subvariety given by the graph of f, verify that $\tau \Gamma \subseteq \tau(V_1 \times V_2) = \tau V_1 \times \tau V_2$ is in fact the graph of a morphism from τV_1 to τV_2, and then define τf to be that morphism. It is not too difficult to extend this construction to rational maps, but I leave the details to you.

The following is a useful way to check when a rational map is D-rational:

Lemma 5.5. *Suppose (V_i, s_i) are D-varieties over (k, δ) and $f : V_1 \to V_2$ is a rational map. Then the following are equivalent:*

(i) *f is D-rational,*
(ii) *$f\big((V_1, s_1)^\sharp \cap \mathrm{dom}(f)\big) \subseteq (V_2, s_2)^\sharp$, and*
(iii) *$f(a) \in (V_2, s_2)^\sharp$ for some (equivalently any) generic D-point a of (V_1, s_1).*

Proof. I leave it to you to check that the following diagram

$$
\begin{array}{ccc}
\tau V_1(\mathbb{K}) & \xrightarrow{\ \tau f\ } & \tau V_2(\mathbb{K}) \\
\nabla \uparrow & & \uparrow \nabla \\
V_1(\mathbb{K}) & \xrightarrow{\ \ f\ \ } & V_2(\mathbb{K})
\end{array}
$$

always commutes. From this, and the fact that ∇ agrees with s_i on $(V_i, s_i)^\sharp$, the implication (i) $\Longrightarrow$ (ii) follows easily.

Note that if a is a generic D-point of (V_1, s_1) over k then it is Zariski-generic in V_1 over k, and hence the rational map f is defined at a. Hence, that (ii) implies the "for any" version of (iii) is clear.

Now, assume (iii) holds of some generic D-point a. This implies that $\tau f \circ s_1$ and $s_2 \circ f$ agree on a. But a is Zariski-generic in $\mathrm{dom}(f)$ over k, and such agreement is a Zariski closed condition over k. It follows that they agree on all of $\mathrm{dom}(f)$, as required for (i). $\qquad\square$

OK, now back to orthogonality and internality for finite dimensional types. Recall that the constants can be viewed as the set of D-points on the trivial vector field $(\mathbb{A}^1, 0)$. So, from the differential-algebraic geometric point of view, a natural notion of "interaction" between a D-variety (V, s) and $(\mathbb{A}^1, 0)$ would be the existence of a nonconstant D-rational map $f : (V, s) \to (\mathbb{A}^1, 0)$. By "nonconstant" here we mean that f is not a constant function on V, that as an element of the rational function field, $k(V)$, it is not in k. However, the fact that f is a D-rational map to $(\mathbb{A}^1, 0)$ does mean, exactly, that f is a constant of the derivation δ_s on $k(V)$. (You should check this.) The constants of $(k(V), \delta_s)$ are often called *rational first integrals* for (V, s). In any case, here are the promised geometric characterisations:

Theorem 5.6. *Suppose (V, s) is an absolutely irreducible D-variety over (k, δ) with generic type $p \in S(k)$. Then*

(a) *p is nonorthogonal to the constants if and only if there is a differential field extension $F \supseteq k$ and a nonconstant D-rational map $f : (V, s) \to (\mathbb{A}^1, 0)$ over F,*
(b) *p is almost internal to the constants if and only if there is a differential field extension $F \supseteq k$ and a D-rational map $f : (V, s) \to (\mathbb{A}^{\dim V}, 0)$ over F which is dominant and generically finite-to-one.*

Proof. Let us consider the right-to-left direction of part (a). Fix $a \in (V, s)^\sharp$ a generic D-point over F. Then $a \models p$, and $a \underset{k}{\downarrow} F$ by Proposition 4.3. We have $f(a) \in (\mathbb{A}^1, 0)^\sharp = \mathcal{C}$. Since f is not constant on V, the Zariski locus of a over $F\langle f(a) \rangle$ is a proper subset of V, and hence $a \underset{F}{\not\downarrow} f(a)$, again by 4.3. This witnesses $p \not\perp \mathcal{C}$.

Note that the same proof also gives the right-to-left direction of part (b); indeed, in that case $f(a) \in \mathcal{C}^{\dim V}$ and f being dominant and generically finite-to-one witnesses that $a \in \mathrm{acl}(Ff(a))$, which yeilds almost $\mathcal{C}$-internality.

For the converse of part (a), let $L \supseteq k$ be a differential field extension and $c = (c_1, \ldots, c_m)$ a tuple from $\mathcal{C}$ such that $a \mathop{\smile\hskip-0.8em\vert\,}_k L$ and $a \mathop{\not\smile\hskip-0.8em\vert\,}_L c$. That is, L and c witness that $p \mathop{\not\perp} \mathcal{C}$. (Note that by taking definable closures we can always assume that the parameter extension witnessing nonorthogonality is a differential field extension.) Moreover, suppose m is minimal such. So, if we set $F = L(c_1, \ldots, c_{m-1})$, then $a \mathop{\smile\hskip-0.8em\vert\,}_k F$ and $a \mathop{\not\smile\hskip-0.8em\vert\,}_F c_m$. As c_m is a constant, its dimension over any differential field is at most 1, and hence the fact that $c_m \mathop{\not\smile\hskip-0.8em\vert\,}_F a$ implies that $c_m \in \mathrm{acl}(Fa) \setminus \mathrm{acl}(F)$. Let E be the finite orbit of c_m under the action of the automorphisms of $(\mathbb{K}, \delta)$ that fix F and a pointwise. And let e be a code for E. (In a theory of fields, codes for finite sets always exist.) So $e \in \mathrm{dcl}(Fa) \setminus \mathrm{acl}(F)$. Moreover, as $E \subseteq \mathcal{C}$, we have that e is itself a tuple of constants, say $e = (e_1, \ldots, e_\ell)$. Re-indexing, we may assume that $e_1 \in \mathrm{dcl}(Fa) \setminus \mathrm{acl}(F)$. Note that $\mathrm{dcl}(Fa) = F(a)$ and a is Zariski-generic on V over F, so that $e_1 = f(a)$ for some rational function f on V over F. That f is nonconstant follows from the fact that $e_1 \notin \mathrm{acl}(F)$. Since $f(a) \in \mathcal{C} = (\mathbb{A}^1, 0)^\sharp$, Lemma 5.5 tells us that f is a D-rational map, as desired.

The left-to-right direction of part (b) is proved similarly. Let $L \supseteq k$ and $c = (c_1, \ldots, c_m)$ witness the almost $\mathcal{C}$-internality of p so that $a \mathop{\smile\hskip-0.8em\vert\,}_k L$ and $a \in \mathrm{acl}(Lc)$. Re-indexing we can find $0 \le \ell < m$ such that $(c_1, \ldots, c_\ell)$ is an acl-basis for c over La. Letting $F = L(c_1, \ldots, c_\ell)$, we have $a \mathop{\smile\hskip-0.8em\vert\,}_k F$ and a interalgebraic with $(c_{\ell+1}, \ldots, c_m)$ over F. Let e be the code of the orbit of $(c_{\ell+1}, \ldots, c_m)$ over Fa so that $e \in \mathrm{dcl}(Fa)$ and $a \in \mathrm{acl}(Fe)$. So $e = f(a)$ where f is a generically finite-to-one dominant rational map from V to $W := \mathrm{Z\text{-}loc}(e/F)$. Exactly as before we see that $f : (V, s) \to (W, 0)$ is D-rational. As $\dim W = \dim(e/F) = \dim(a/F) = \dim(a/k) = \dim V$, we can compose f with a finite-to-one co-ordinate projection dominantly onto $(\mathbb{A}^{\dim V}, 0)$. $\square$

6. The dichotomy

In this final lecture it is my intention to articulate the Zilber dichotomy as it is manifest in DCF_0, in terms of the birational geometry of D-varieties. I will not prove anything here; neither the dichotomy itself, nor even that the geometric formulation I give is equivalent to the usual model-theoretic formulation. The former (a proof of the dichotomy) is certainly beyond the scope of these lectures, while the latter I leave to the reader as it can be derived, with some work, using the various translations between model theory and geometry that we have discussed throughout these lectures.

The dichotomy will be a statement about D-varieties that are not "covered by a family of proper infinite D-subvarieties". More precisely:

Definition 6.1. Suppose (V, s) is an absolutely irreducible D-variety over (k, δ). We say that (V, s) is *simple*[2] if the following holds: for all irreducible D-varieties (W, t) over (k, δ), and all proper irreducible D-subvarieties Z of $(V \times W, s \times t)$ over k, if Z projects dominantly onto both V and W then $\dim Z = \dim W$.

---

[2]My terminology is inspired by the bimeromorphic geometry of compact complex manifolds, where "simplicity" is used to describe the analogous property.

Let me explain a little how to think about this in terms of covering families of D-subvarieties. Given an irreducible D-variety (W, t) over k and an irreducible D-subvarieties Z of $(V \times W, s \times t)$ over k projecting dominantly onto W, we view Z as a *family of D-subvarieties of (V, s) parametrised by (W, t)* in the following way: to each $b \in (W, t)^\sharp$ we can associate the fibre

$$Z_b := \{x \in V : (x, b) \in Z\}.$$

It is a D-subvariety of (V, s) over the differential field $k(b)$; this uses that b is a D-point, Z is a D-subvariety, and the co-ordinate projection is a D-morphism. It is this family of D-subvarieties of (V, s), namely $\{Z_b : b \in (W, t)^\sharp\}$, that we have in mind. By a *generic member* of this family, or a *generic fibre*, we mean a fibre of the form Z_b where b is a generic D-point of (W, t). Note that the irreducibility of Z ensures that this generic fibre Z_b is irreducible over $k(b)$. If Z is a proper subvariety of $V \times W$ then the generic fibre Z_b is a proper subvariety of V. We say that (V, s) is *generically covered* by the family if Z projects dominantly onto V as well. The reason for this terminology is that it implies (indeed, precisely says) that if a is a generic D-point of (V, s) then $a \in Z_b$ for some (generic) D-point b of (W, t). Note also that $\dim Z$ is the sum of $\dim W$ and the dimension of the generic fibre. In particular, $\dim Z = \dim W$ if and only if the generic fibre is finite. So Definition 6.1 is saying that (V, s) admits no generically covering family of D-subvarieties whose generic members are infinite and proper.

Model-theoretically, (V, s) being simple is equivalent to the generic type of (V, s) being *minimal*, that is, all its forking extensions are algebraic. While we do not verify this here, let me say a few words that may be of use. First of all, one has to convince oneself that the extra parameters needed to witness nonminimality can always be taken to be themselves finite dimensional. So the generic D-points of the (W, t) appearing in the definition of simple are the putative parameters for the forking extensions. Secondly, if $\mathrm{tp}(a/k)$ is the generic type of (V, s) and $\mathrm{tp}(a/kb)$ is a forking extension, with $\mathrm{tp}(b/k)$ the generic type of (W, t), then the Zariski locus of (a, b) over k is a family of proper D-subvarieties of (V, s) parametrised by (W, t) that generically covers (V, s). Moreover, the converse holds as well; every such family gives rise to a forking extension. Beyond that, one only has to verify that $\mathrm{tp}(a/kb)$ is algebraic if and only if the projection of $\mathrm{Z\text{-}loc}(a, b/k)$ onto W is generically finite-to-one.

Every D-*curve*, by which I mean a D-variety (V, s) where $\dim V = 1$, is simple. This is because *every* proper subvariety of a curve is finite, let alone those that come from generically covering families. But there are simple D-varieties of higher dimension. Here is an example without proof: Consider the surface $V \subseteq \mathbb{A}^3$ defined by the equation $xz = 1$ in co-ordinate variables (x, y, z), and let s be the section to the tangent bundle given by $s(x, y, z) = (x, y, z, y, yz, -yz^2)$. Then (V, s) is simple. Indeed, this D-variety is the order 2 differential equation $xx'' - x' = 0$ in disguise, studied and shown to have minimal generic type by Poizat (see, for example, the treatment in [5, §II.5]).

In fact, simple D-varieties are ubiquitous in all dimensions. Moreover, they are, in some real but subtle sense that I will not go into here, the building blocks for all D-varieties. Suffice it to say that they are very much worth understanding. The Zilber dichotomy will tell us that they either come from pure algebraic geometry or that they are geometrically very tame in a sense that we will now discuss.

Definition 6.2. Suppose (V, s) is an absolutely irreducible D-variety over (k, δ). A family Z of D-subvarieties of (V, s) parameterised by (W, t) is said to be *rich* if the following conditions hold:

(i) for generic $b \in (W, t)^\sharp$, Z_b is absolutely irreducible, and
(ii) for generic $a \in (V, s)^\sharp$ there are infinitely many distinct generic fibres Z_b that pass through a.

Here by "distinct generic fibres" I mean simply generic D-points b_1, b_2 of (W, t) such that $Z_{b_1} \neq Z_{b_2}$ as subvarieties of V.

Note that in the case of a trivial vector field, every subvariety over constant parameters is a D-subvariety. This is because the zero section takes every subvariety to its tangent bundle, which, when we are over constant parameters, agrees with the prolongation. So trivial vector fields are a natural place to look for rich families of D-subvarieties, at least if the dimension is greater than 1. For example, in $(\mathbb{A}^2, 0)$, we have the family of lines $y = mx + b$, which defines a D-subvariety of $(\mathbb{A}^2, 0) \times (\mathbb{A}^2, 0)$ in the variables (x, y, m, b), and hence a family of D-subvarieties of $(\mathbb{A}^2, 0)$ parameterises by $(\mathbb{A}^2, 0)$. This family is rich because there are infinitely many lines through each point in the plane.

What the above example also illustrates, is that while simple D-varieties (more or less vaciously) admit no rich families of D-subvarieties, cartesian powers of a simple D-variety may: $(\mathbb{A}^1, 0)$ is simple but $(\mathbb{A}^2, 0)$ admits rich families of D-curves. However, there are simple D-varieties all of whose cartesian powers admit no rich families: the Kolchin equation $x' = \frac{x}{x+1}$ and the Poizat equation $xx'' = x'$, both of which we have already mentioned, are examples in dimension 1 and dimension 2 respectively. The absence of rich families in all cartesian powers is a strong structural constraint on the differential-algebraic geometry of the D-variety.

The Zilber dichotomy states that $(\mathbb{A}^1, 0)$ is essentially the only simple D-variety that does not satisfy this strong structural constraint.

Theorem 6.3 (Zilber dichotomy in DCF_0, geometric formulation). *Suppose (V, s) is a simple D-variety over (k, δ). Then either*

(1) *there is a generically finite to one D-rational map $f : (V, s) \to (\mathbb{A}^1, 0)$ over some differential field extension of k, or*
(2) *whenever $n \geq 1$ and X is an absolutely irreducible D-subvariety of (V^n, s^n) over k that projects dominantly onto V in each co-ordinate, then X admits no rich families.*

For the model theorist who is not used to seeing the dichotomy described this way, recall from Theorem 5.6(b) that case (1) corresponds to the generic type of (V, s) being almost $\mathcal{C}$-internal. (Because of simplicity, in this case we actually get that $\dim V = 1$.) Case (2) is equivalent to the generic type of (V, s) being 1-*based*, though I have not said anything about that. In any case, this is the usual dichotomy in DCF_0, but presented as a theorem about the birational geometry of D-varieties.

More is known about case (2). It splits into two more cases (thus forming the Zilber trichotomy); one coming from certain D-group structures on universal vectorial extensions of simple abelian varieties that do not descend to the constants (the *Manin kernel* case), and the other being when the cartesian powers of (V, s) are truly devoid of any structure in that the absence of rich families is replaced by the absence of any infinite families at all (the *relationally trivial* case). But that is a story for another course.

The Zilber dichotomy was first proved for DCF_0 by Hrushovski and Sokolovic [3], relying on the theory of Zariski geometries developed by Hrushovski and Zilber [4]. Later, a new simpler proof was found by Pillay and Ziegler [6] using differential jet spaces, and having to do with the canonical base property, itself inspired by the model theory of compact complex manifolds. But that too is a story for another course.

REFERENCES

[1] Alexandru Buium. *Differential function fields and moduli of algebraic varieties*. Springer-Verlag, 1986.

[2] Bradd Hart and Matthew Valeriote, editors. *Lectures on algebraic model theory*, volume 15 of *Fields Institute Monographs*. American Mathematical Society, Providence, RI, 2002.

[3] Ehud Hrushovski and Željko Sokolović. Strongly minimal sets in differentially closed fields. *unpublished manuscript*, 1993.

[4] Ehud Hrushovski and Boris Zilber. Zariski geometries. *Bull. Amer. Math. Soc. (N.S.)*, 28(2):315–323, 1993.

[5] David Marker, Margit Messmer, and Anand Pillay. *Model theory of fields*, volume 5 of *Lecture Notes in Logic*. Association for Symbolic Logic, La Jolla, CA; A K Peters, Ltd., Wellesley, MA, second edition, 2006.

[6] Anand Pillay and Martin Ziegler. Jet spaces of varieties over differential and difference fields. *Selecta Math. (N. S.)*, 9 (4):579–599, 2003.

RAHIM MOOSA, UNIVERSITY OF WATERLOO, DEPARTMENT OF PURE MATHEMATICS, 200 UNIVERSITY AVENUE WEST, WATERLOO, ONTARIO N2L 3G1, CANADA

Email address: `rmoosa@uwaterloo.ca`