THE ELEMENTARY THEORY OF THE FROBENIUS AUTOMORPHISMS: EQUIVALENCE OF THE INFINITE COMPONENTS

by

Thomas Scanlon

Abstract. — We explain Section 13 of [**1**] in which Hrushovski completes the proof of the main theorem of this paper. The crucial lemma is a comparison between virtual intersection numbers and the true cardinalities of the solutions to difference equations in variable Frobenii.

Résumé **(La théorie élémentaire des automorphismes de Frobenius: Équivalence des composantes infinies)**

Nous expliquons la section 12 de [**1**]. Dans ce chapitre, Hrushovski montre la théorème principal de l'article. La lemme crucial est une comparaison entre le nombre d'intersection virtuel et les vraies cardinalités des ensembles de solutions d'équations aux différences quand les Frobenius varient.

With these notes we explain some of the details from Section 12 of [**1**]: Equivalence of the infinite components: asymptotic estimates. The main theorem of this section is a comparison theorem, Proposition 12.1, between the virtual intersection numbers and the actual number of points in the intersection of correspondences and graphs of Frobenius. These notes should be read as an annotated account of Hrushovski's argument. We follow the structure (and in most cases the details) of his argument filling in some of the steps which appear merely as references in [**1**], but we have made no attempt at simplification or improvement.

We recall the main theorem (Theorem 1B) to be proven.

Theorem 0.1. — *For any algebraically closed field K of characteristic $p > 0$, positive power q of p, irreducible affine variety $V \subseteq \mathbb{A}^n_K$, and irreducible correspondence $\Gamma \subseteq V \times V^{\phi_q}$ (by which we mean that both projections are generically finite and at least one is quasi-finite), the cardinality of the set $\{a \in V(K) \ : \ (a, \phi_q(a)) \in \Gamma(K)\}$ is $\delta q^{\dim(V)} + e$ where $\delta = ([K(\Gamma) : K(V)]/[K(\Gamma) : K(V^{\phi_q})]_{insep})$ and $|e| \leq C q^{\dim(V) - \frac{1}{2}}$ for some constant C depending only on n and the degrees of the (projective closures) of V and Γ (but not on K, p or q).*

Partially supported by NSF grants FRG DMS-0854998 and DMS-1001550.

There is an easy reformulation of Theorem 0.1 where we allow V to be projective, but it is not so easy to reduce Theorem 0.1 to the corresponding case for projective varieties for which one of the projections from Γ is quasi-finite. The point of this section is to check that the errors introduced by passing to the projective closure are swamped by the error term in Theorem 0.1 itself.

We begin by recalling the basic data presenting a difference equation with variable Frobenius. As we require only the statement of Theorem 1B, which refers to directly presented difference schemes, we restrict to this case rather than the more general presentation of Notation 10.18 to which Hrushovski refers on pages 119 and 120.

Hypotheses 0.2. —
- We are given an (affine) integral scheme B of finite type and a second closed integral subscheme $B' \subseteq B \times B$.
- We have a projective integral scheme $X \subseteq \mathbb{P}^n_B$.
- We have a closed integral subscheme $S \subseteq X \times X$.
- We presume that S lies over B' via the map $X \times X \to B \times B$.
- We denote the two projection maps $S \to X$ by $\pi_0 : S \to X$ and $\pi_1 : S \to X$.
- We denote the two projection maps $B' \to B$ by $\nu_0 : B' \to B$ and $\nu_1 : B' \to B$.

With these data, if (k, σ) is a difference ring and $b \in B'(k)$ satisfies $\sigma(\nu_0(b)) = \nu_1(b)$,

$$X_{\nu_0(b)}(S_b, \sigma)(k) := \{a \in X_{\nu_0(b)}(k) \ : \ (a, \sigma(a)) \in S_b(k)\}$$

Remark 0.3. — In some other works [**2**], this set is denoted as $(X_{\nu_0(b)}, S_b)^\sharp(k, \sigma)$. Hrushovski uses two other notations. He calls the difference scheme associated to S_b, $(S_b \star \Sigma)$, and on occasion (see page 121) $X(S_b)$. One should note that he uses "V" for the ambient space where we use "X" so that the symbol "X" in his notation does not refer to that space. We will not use this latter notation as it risks confusion with the standard notation for sets of points. That is, $X(S_b)$ ought to mean $\mathrm{Hom}(S_b, X)$. The notation $(S_b \star \Sigma)$ for the difference scheme directly presented by S_b does omit the important data of the maps π_0 and π_1, but as these are clearly present in our context, this should not cause any problems.

We impose some further requirements on these data.

Hypotheses 0.4. — There are natural numbers δ_0, δ_1, and d, so that for any algebraically closed field k and k-valued point $b \in B(k)$ we require:
- the fibres $X_{\nu_0(b)}$, $X_{\nu_1(b)}$ and S_b are irreducible varieties of dimension d,
- the projections $\pi_0 : S_b \to X_{\nu_0(b)}$ and $\pi_1 : S_b \to X_{\nu_1(b)}$ are dominant and generically finite of degrees δ_0 and δ_1, respectively, (we write δ for δ_0),
- $\pi_1 : S_b \to X_{\nu_1(b)}$ is generically smooth,
- $X_{\nu_0(b)}$ is smooth.

In approaching Theorem 1B, passing to projective closures, one reduces to a problem about correspondences on projective varieties. One reduces easily to the case where the fibres are irreducible varieties using the observation that irreducibility is a constructible condition. Likewise, decomposing the base constructibly, we may arrange that the degrees of the projection maps are constant. Taking π_1 to be generically smooth is somewhat subtler. We decompose the base into finitely many constructible pieces, one on which δ_1 is invertible and finitely many others having characteristic p with p dividing δ_1. When δ_1 is invertible, the extension $k(S_b)/k(X_{\nu_1(b)})$ is automatically separable so that $\pi_1 : S_b \to X_{\nu_1(b)}$ is generically smooth. When working over $\operatorname{Spec} \mathbb{F}_p$, Lemma 10.19 is used to show that via appropriate precomposition with the Frobenius it suffices to consider the case that $\pi_1 : \Gamma_b \to X_{\nu_1(b)}$ is generically smooth. Smoothing the closure of $X_{\nu_1(b)}$ in projective space passes through either Hironaka's resolution of singularities in characteristic zero or de Jong's theorem on smooth alterations in general (see Lemma 10.30). The key point is that for the purpose of counting the number of points in $(X_{\nu_0(b)}, \Gamma_b)^\sharp(K_q, \phi_q)$ qualitatively, it suffices to count the points on finite covers. Again, one needs to argue that the smoothing process, which we have applied at only one fibre, when applied to the generic fibre extends to a neighborhood in B so that one may constructibly achieve families satisfying these stronger hypotheses from which the truth of Theorem 1B restricted to these families one may deduce the full result.

Remark 0.5. — In [**1**], the notation $X_b(\Gamma, q)$ is used for $(X_{\nu_0(b)}, S_b)^\sharp(K_q, \phi_q)$ where q a prime power and K_q is an algebraically closed field extending $\mathbb{F}_q$ with $\phi_q : K_q \to K_q$ the q-power Frobenius.

Finally, when $b \in B'(k)$, q is a power of the (positive) characteristic of k, and $\phi_q(\nu_0(b)) = \nu_1(b)$, then we denote by Φ_q the graph of the Frobenius $X_{\nu_0(b)} \to X_{\nu_1(b)}$. From the hypothesis that $S_b \to X_{\nu_1(b)}$ is generically smooth and generically finite, the largest open subvariety $\widetilde{S_b} \subseteq S_b$ on which the projection map $\widetilde{S_b} \to X_{\nu_1(b)}$ is étale is dense in S_b. Using the openness of the quasi-finite and unramified loci, we may take $\widetilde{S_b}$ to be a fibre of an open subvariety $\widetilde{S} \subseteq S$, again possibly shrinking B.

The main theorem of this section is Proposition 12.1 of [**1**].

Theorem 0.6. — *With the data as above, there is a constant C so that for any field k of characteristic $p > 0$ and any point $b \in B(k)$, if q is a power of p and $\nu_1(b) = \phi_q(\nu_0(b))$, then $|(\Phi_q \cdot \Gamma_b) - \#(\Phi_q \cap \widetilde{S}_b)| < Cq^{d-1}$.*

Remark 0.7. — In light of the fact that the intersection $\Phi_q \cap \widetilde{S}_b$ is transverse, counting $\#(\Phi_q \cap \widetilde{S}_b)$ with or without multiplicities gives the same result.

Remark 0.8. — The set $\Phi_q \cap \widetilde{S}_b$ is in bijection with $(X_{\nu_0(b)}, \widetilde{S}_b)^\sharp(K_q, \phi_q)$. Hence, counting the number of intersection points is the same as counting the number of solutions to our difference equation.

Remark 0.9. — The main result of Section 11 is that the virtual intersection number $(\Phi_q \cdot S_b)$ differs from δq^d by an error term of size $O(q^{d-\frac{1}{2}})$. Combined with Theorem 0.6, we conclude that $\#(X_{\nu_0(b)}, \widetilde{S}_b)^\sharp(K_q, \phi_q)$ is also on the order of δq^d with an error term of size $O(q^{d-\frac{1}{2}})$.

In a nutshell, Theorem 0.6 is proven by replacing S_b with a (strongly reduced-rationally) equivalent correspondence for which all of the numbers in question differ by $O(q^{d-1})$ and for which we may remove lower dimensional subvarieties altering these numbers again by only $O(q^{d-1})$. We then check that all of the numbers match once an appropriate lower dimensional variety is removed and thereby finish the proof. As we will see, the theory of transformal valued fields plays a crucial rôle in the proof of the estimate for the number of points on a lower dimensional correspondence and is a tool in the proof of a moving lemma.

Notation 0.10. — Throughout the rest of these notes, we implicitly work with some parameter $b \in B'(k)$ where (k, σ) is a difference field satisfying $\sigma(\nu_0(b)) = \nu_1(b)$. We write Γ for S_b, $\widetilde{\Gamma}$ for $\widetilde{S}_b$, V for $X_{\nu_0(b)}$ and V^σ for $X_{\nu_1(b)}$. On the rare occasion where we need to remember the total family, we shall say so explicitly.

Notation 0.11. — Let us recall the notion of a correspondence and then specialize it to separable correspondences.

–
$$\begin{aligned} \mathcal{C}(V) \quad := \quad & \{T \subseteq V \times V^\sigma \ : \ T \text{ an irreducible subvariety such that} \\ & \dim(T) = d, \pi_0 : T \to V \text{ and } \pi_1 : T \to V^\sigma \text{ are dominant } \} \end{aligned}$$

–
$$\mathcal{C}(V)_{\text{sep}} \quad := \quad \{T \in \mathcal{C}(V) \ : \ \pi_1 : T \to V^\sigma \text{ is generically smooth } \}$$

– $\mathbb{Z}\mathcal{C}(V)$ is the free abelian group generated by $\mathcal{C}(V)$ and $\mathbb{Z}\mathcal{C}(V)_{\text{sep}}$ is the free abelian group generated by $\mathcal{C}(V)_{\text{sep}}$.
– For $T \in \mathcal{C}(V)_{\text{sep}}$ the map $\pi_1 : T \to V^\sigma$ is generically smooth and generically finite. We denote by $(V,T)_{\text{ét}}$ the largest open subvariety $U \subseteq V$ for which $\pi_1 \restriction (T \cap \pi^{-1}V^\sigma)$ is étale over V^σ.
– For $T \in \mathcal{C}(V)$ we let $(V,T)_{\text{inf}}$ be the largest closed subvariety $W \subseteq V$ for which $\dim(\pi_1 \restriction T)^{-1}(\sigma(a)) > 0$ for all points $a \in W^\sigma$. We set $(V,T)_{\text{fin}} := V \smallsetminus (V,T)_{\text{fin}}$.
– We define
$$\mathcal{C}(V)_f \quad := \quad \{T \in \mathcal{C}(V) \ : \ [\sigma]_k(V,T)_{\text{inf}} \cap (T \star \Sigma) = \varnothing\}$$
As above, $\mathbb{Z}\mathcal{C}(V)_f$ is the free abelian group generated by $\mathcal{C}(V)_f$.

Remark 0.12. — With the notation introduced before Theorem 0.6, $\widetilde{\Gamma} = \pi_1^{-1}(V,\Gamma)^\sigma_{\text{ét}}$. Shortly, we shall adjust the meaning of $\widetilde{\Gamma}$. As we noted earlier, $\widetilde{S_b}$ does vary constructible over B. However, $(X_b, S_b)_{\text{ét}}$ and $(X_b, S_b)_{\text{fin}}$ do not. Rather, they vary in a difference constructible family.

Remark 0.13. — Note that necessarily $(V,T)_{\text{ét}}$ is dense in V. While the notation $V_{\text{ét}}$ is often used for the étale site over V, we have no need to discuss this site here. Hrushovski writes $X_{\text{ét}}(T)$ for what we call $(V,T)_{\text{ét}}$.

Remark 0.14. — Hrushovski writes $X_{fin}(T)$ for $(V,T)_{\text{fin}}$. We shall adopt Hrushovski's alternate notation: $\widetilde{V} := (V,T)_{\text{fin}}$ when T is understood. We shall write $\widetilde{T} := T \cap (\widetilde{V} \times \widetilde{V}^\sigma)$. Note that the map $\pi_1 : \widetilde{T} \to \widetilde{T}^\sigma$ is quasifinite. Note also that this definition of $\widetilde{T}$ differs somewhat from what was given before Theorem 0.6 but these two varieties agree up to a subvariety of lower dimension so that the number of points in $(\widetilde{T} \star \Sigma)(K_q, \phi_q)$ will agree up to $O(q^{d-1})$.

Remark 0.15. — Hrushovski's definition of $\mathcal{C}(V)_f$ is slightly different. First of all, (adjusting for the names of the varieties) he writes $\mathrm{pr}_0[1](T) \cap (T \star \Sigma) = \varnothing$ for his defining condition and then says that this is an equivalent condition to $X(T) = T \star \Sigma$. Recall that $X(T) = (\widetilde{T} \star \Sigma)$. This first formulation is clearly missing some symbols as $\mathrm{pr}_0[1](T)$ is a variety while $(T \star \Sigma)$ is a difference scheme. This is rectified by regarding $\mathrm{pr}_0[1](T)$ as a difference scheme via the $[\sigma]$ functor from the category of schemes to the category of difference schemes. One should recall that if $f : U \to W$ is a morphism of schemes and n is a natural number, then $f[n](U)$ is the subscheme of W defined by $a \in f[n](U)$ just in case $\dim f^{-1}(a) \geq n$. From the equivalent definition, one surmises that $(V, T)_{\mathrm{fin}} = (\mathrm{pr}_1[1](T))^{\sigma^{-1}}$ was meant.

We recall the notation of the Frobenius fields and of the Frobenius specializations.

For $p > 0$ a prime and q a (positive) power of p, K_q is an algebraically closed field of characteristic p considered as a difference field with the map $\phi_q : K_q \to K_q$ given by $x \mapsto x^q$. Since (K_q, ϕ_q) is a definitional expansion of K_q considered as a field and the theory of algebraically closed fields of characteristic p is complete, it does not matter which field we choose to represent K_q, at least for the sake of the estimates proven here.

Recall that we have a functor M_q from the category of difference algebras to the category of commutative $\mathbb{F}_p$-algebras given by sending a difference algebra (D, σ) to the quotient difference algebra $D/\langle p, \{\sigma(a) - a^q \ : \ a \in D\}\rangle$ and then forgetting the distinguished endomorphism. Recall that the functor M_q induces a corresponding functor from the category of affine difference schemes to the category of affine schemes over $\mathbb{F}_p$ which localizes appropriately so that it extends to a functor from the category of difference schemes to the category of schemes. For a difference scheme Y over (D, σ) and a morphism of difference rings $h : (D, \sigma) \to (K_q, \phi_q)$, we write $Y_{q,h}$ for $\mathrm{M}_q(Y) \otimes_h K_q$.

Definition 0.16. — For $T \in \mathcal{C}(V)$ irreducible with T and V defined over some difference domain $D \subseteq k$ and $h : (D, \sigma) \to (K_q, \phi_q)$ a morphism of difference rings, we define $\eta_0(X, T, q, h) := \# \, \mathrm{M}_q((\widetilde{V}, \widetilde{V})^\sharp) \otimes_h K_q$. Note that here we do not count the points with multiplicity. Let δ' be the inseparable degree of the field extension from $K_q(X^\sigma \otimes_h K_q)$ to $K_q(T \otimes_h K_q)$. We then define $\eta(V, T, q, h) := \delta' \eta_0(V, T, q, h)$.

Definition 0.17. — For T and U in $\mathcal{C}(V)$ we say that T and U are *asymptotically equinumerous*, written $T \sim_a U$, if there is a finitely generated difference domain $D \subseteq k$ over which V, T and U are defined and a constant C so that for almost all q (meaning for all sufficiently high powers of p if $\mathrm{char}(\mathcal{O}_B) = p$ and for all powers of p for p a sufficiently large prime otherwise) and all maps of difference rings $h : (D, \sigma) \to (K_q, \phi_q)$, we have

$$|\eta(V, T, q, h) - \eta(X, U, q, h)| \leq C q^{d-1}$$

We say that $x, y \in \mathbb{Z}\mathcal{C}(X)$ are asymptotically equinumerous if $x - y$ belongs to the subgroup of $\mathbb{Z}\mathcal{C}(X)$ generated by the differences $[T] - [U]$ where $T \sim_a U$.

Remark 0.18. — It is easy to see that if asymptotic equinumerosity is witnessed by some finitely generated difference domain $D \subseteq k$ and $D \subseteq D' \subseteq k$ is a larger finitely generated difference domain, then D' also witnesses asymptotic equinumerosity.

Remark 0.19. — With Remark 12.2 Hrushovski proposes an extension of the definition of asymptotic equinumerosity for non-reduced schemes including geometric multiplicities in the counting function. Including such multiplicities in our case of varieties would not alter the definition of asymptotic equinumerosity by the Second Bounding Lemma, Corollary 9.24 of [**1**]. Hence, the difference between the counting function taking multiplicities into account and our more naïve function is $O(q^{d-1})$.

We shall show that Γ may be replaced by an asymptotically equinumerous cycle for which it is easier to compare the virtual intersection numbers with the counting functions by invoking the theory of reduced-rational equivalence as introduced in Section 10.2. To be completely honest, reduced-rational equivalence as defined there may be too weak a notion. For the sake of this section we revise the definition.

Definition 0.20. — Consider the subgroup $\mathbb{Z}\mathcal{C}(V)_{srr}$ of $\mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$ generated by cycles of the form $[T_\infty] - [T_0]$ where $T \subseteq (V \times V^\sigma) \times \mathbb{P}^1$ is a subvariety of dimension $d+1$ for which $T_\infty \in \mathcal{C}(V)_{\mathrm{sep}}$ and $T_0 \in \mathcal{C}(T)_{\mathrm{sep}}$. We say that x and y from $\mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$ are strongly reduced-rationally equivalent if $x - y \in \mathbb{Z}\mathcal{C}(V)_{srr}$.

The definition of reduced-rational equivalence in Section 10.2 allows for T_∞ and T_0 to be arbitrary d-cycles but the proof that reduced-rationally equivalent correspondences are asymptotically equinumerous requires strong reduced-rational equivalence. Fortunately, while Lemma 12.3 asserts merely that for $U \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$ it is possible to find $U' \in \mathbb{Z}\mathcal{C}(V)_f$ reduced-rationally equivalent to U, in fact, the "moreover" clause of the cited Lemma 10.7 gives strong reduced-rational equivalence. Let us state this fact as a proposition.

Proposition 0.1. — *For any $U \in \mathbb{Z}\mathcal{C}(V)_{sep}$ there is some strongly reduced-rationally equivalent $U' \in \mathbb{Z}\mathcal{C}(V)_f$.*

Since the intersection product is invariant under rational equivalence and strong reduced-rational equivalence is a stronger notion that rational equivalence, if $\Gamma' \in \mathbb{Z}\mathcal{C}(X)_f$ is obtained from Γ via Lemma 0.1, then for every difference domain D over which V, Γ, Γ' and the varieties $T \subseteq (V \times V^\sigma) \times \mathbb{P}^1$ witnessing strong reduced-rational equivalence are defined and almost every specialization $h : (D, \sigma) \to (K_q, \phi_q)$, we have $(\Gamma_{q,h} \cdot \Phi_q) = (\Gamma'_{q,h} \cdot \Phi_q)$. Here, the word "almost" is included as one may need to localize D to ensure that the specializations of the algebraic families of cycles witnessing strong reduced-rational equivalence still witness rational equivalence. Thus, if we can further establish that $\Gamma \sim_a \Gamma'$, then the main theorem would follow from an estimate of the expected form for Γ'.

Lemma 12.5 establishes that the generators of $\mathbb{Z}\mathcal{C}(V)_{srr}$ are asymptotically equivalent to 0. The general result then follows.

Lemma 0.21. — *If $T \subseteq (V \times V^\sigma) \times \mathbb{P}^1$ is an irreducible variety of dimension $d+1$, $t \in \mathbb{P}^1$ is the generic point, and T_0 is a correspondence which is generically smooth over X^σ, then $T_t \sim_a T_0$.*

Remark 0.22. — Of course, since t does not correspond to a k-valued point, our definition of asymptotic equivalence does not produce formally meaningful notion for the relation between T_t and T_0. Rather, we base change to $k(t)_\sigma$, the difference field of difference rational functions of $[\sigma]_k\mathbb{P}^1$ and consider asymptotic equivalence there.

Proof. — As is by now standard, we consider $k(t)_\sigma$ as a transformal valued field extending (k,σ) on which the valuation is trivial on k and has $0 < v(t) \ll v(\sigma(t)) \ll v(\sigma^2(t)) \ll \cdots$. Let $k[\check{t}]_\sigma$ be the valuation ring of $k(t)_\sigma$. Let $\mathcal{V} := V_{k[\check{t}]_\sigma}$ and let V_t and V_0 be the generic and special fibres, respectively. We let Y_t be the closure of [smallest closed difference subscheme containing] $(\widetilde{T}_t \star \Sigma)$ in V_t and we let $\mathcal{Y}$ be the closure of Y_t in $\mathcal{V}$.

Consider $E := V_0 \smallsetminus (V_0, \Gamma_0)_{\text{ét}}$ and set

$$^*E := \{a \in \mathcal{V}(k[\check{t}]_\sigma) \ : \ \operatorname{red}(a) \in E(k)\}$$

The set $^*E$ is definable in the transformal valued field $k(t)_\sigma$.

Claim 1: The inertial dimension of $^*E$ is less than d.
Proof of Claim: From Proposition 4.44(1), one knows that if Z is any absolutely irreducible k variety and $U \subseteq Z \times Z^\sigma$ is a correspondence, then any difference subscheme W of $(U \star \Sigma)$ for which $\sigma[W]$ is disjoint from $[\sigma]_k\pi_1[1](U)$ has total dimension at most $\dim(Z)$. By the very definition of $\widetilde{T}_t$, this condition holds (with $W = (\widetilde{T}_t \star \Sigma)$, $Z = V$ and $U = T_t$) so that the total dimension of $(\widetilde{T}_t \star \Sigma)$ is at most d. As taking closures preserves total dimension, the total dimension of Y_t is at most d.

Since $T_0 \in \mathcal{C}(V)$, Proposition 4.44(2), which asserts that if $Z \subseteq (T_0 \star \Sigma)$ is an algebraically integral sub difference scheme which is weakly Zariski dense in V, then Z is a component of $(T_0 \star \Sigma)$ implies that such a Z must be transformally reduced.

With these two conditions established, we may apply Proposition 8.9 letting V play the rôle of V itself, $\mathcal{Y}$ that of X and E that of W. We have just checked the first two conditions of Proposition 8.9, namely that the reduced total dimension of Y_t is at most d and that every sub difference scheme of Y_0 which is weakly dense in V is transformally reduced. Thus, we conclude that since E is a proper subvariety of V, $^*E$ has inertial dimension strictly less than d. ✠

We now fix some finitely generated difference subring $D \subseteq k$ over which T and V are defined. For each specialization $h : (D,\sigma) \to (K_q,\sigma)$ we shall extend h to a specialization $D(t)_\sigma \to K_q(t) \hookrightarrow L_q := K_q(t)^a$. We regard $K_q(t)$ as a discretely valued field with $v(t) > 0$ and trivial valuation on K_q and then extend v to a valuation on L_q.

Claim 2: There is a constant C (independent of h and q) so that if q is large enough (so that $(\tilde{T}_0 \star \Sigma)_{q,h}$ is necessarily a finite scheme), then $|\eta(V,T_t,q,h) - \eta(V,T_0,q,h)| \leq Cq^{d-1}$.

Proof of Claim: Let us note first that from our hypothesis that $T_0 \in \mathcal{C}(V)_{\text{sep}}$, $\eta(V, T_0, q, h) = \eta_0(V, T_0, q, h)$ and from our hypothesis that q is large enough, we may count the points in K_q or in L_q and will obtain the same finite number.

Secondly, we note that because $\mathcal{V}$ is projective, the sets $V_t(L_q)$ and $\mathcal{V}(\mathcal{O}_{L_q,v})$ may be identified so that a reduction map

$$\text{red} : V_t(L_q) \to V_0(L_q)$$

is defined. Restricting to $(Y_t)_{q,h}$ we have a map

$$r : (Y_t)_{q,h}(L_q) \to (Y_0)_{q,h}(L_q)$$

Let $Z := (Y_t)_{q,h} \smallsetminus {}^*E$ and let r' be the restriction of r to Z.

As ${}^*E$ is the set of points which reduce to points in E, the complement of $(V, T_0)_{\text{ét}}$, we see that r' takes its values in $((V, T_0)_{\text{ét}})_{g,h}$.

We check now that, in fact, r' is a bijection between these two sets. Indeed, injectivity is precisely an instance of Lemma 7.3: Z is a quasiprojective finitely presented difference scheme over a strictly increasing transformal valued field for which $Z \to B_1(Z)$ is unramified (which is the hypothesis of Lemma 7.3), so that the reduction map is injective (which is the conclusion of Lemma 7.3).

Surjectivity follows from Hensel's Lemma. Indeed, at this point we are dealing with ordinary schemes. If $a \in ((V, T_0)_{\text{ét}})_{g,h}(K_q)$, then we may regard a as a point in $\mathcal{V}(\mathcal{O}_{L_q,v})$ (via the inclusion of $K_q \hookrightarrow L_q$ which while not necessarily a point in Z is close to being a solution of the equations defining $(Y_t)_{q,h}$. That is, working locally, if $f(x_1, \ldots, x_n; y_1, \ldots, y_n) \in D[t][x_1, \ldots, x_n; y_1, \ldots, x_n]$ is a polynomial for which $f(\boldsymbol{x}; \boldsymbol{\sigma}(\boldsymbol{x})) \in I(Y_t)$, then $v(f^h(a; \phi(a))) > 0$ (where here we write $(\cdot)^h$ for the natural extension of the specialization to the polynomial ring). Choosing a system of defining equations $f_1, \ldots, f_\ell$, we see from the hypothesis that $a \in ((V, T_0)_{\text{ét}})_{g,h}(K_q)$ that matrix $(\frac{\partial f_i}{\partial x_j}(a, \phi_q(a)))$ has rank d. From the chain rule and the fact that the differential of the Frobenius is zero, we see that if we define $F_i(\boldsymbol{x}) := f_i(\boldsymbol{x}, \boldsymbol{x}^q)$, then $(\frac{\partial F_i}{\partial x_j}(a))$ has rank d as well. Hensel's Lemma now applies to give the lifting.

Thus, $|(V, T_0)_{q,h}(K_q) \smallsetminus E| = |(V, T_t)_{q,h}(L_q) \smallsetminus {}^*E|$.

The Second Bounding Lemma, Corollary 9.24, implies that $|(X, T_0)_{q,h}(K_q) \cap E| = O(q^{d-1})$ (since E is a proper subvariety of V and $T_0 \in \mathcal{C}(V)$ so that every algebraic component of $(\widetilde{T_0} \star \Sigma)$ has transformal multiplicity 0). Since the inertial dimension of ${}^*E$ is less than d, we see by Remark 7.18(2) and Example 7.19 that $|{}^*E| = O(q^{d-1})$. Thus, concluding the proof of the Claim and of the Lemma. ✠ □

Pivoting around the generic point we conclude that strong reduced-rational equivalence implies asymptotic equinumerosity.

Lemma 0.23. — *Cycles which are strongly reduced-rationally equivalent are asymptotically equinumerous.*

Proof. — It suffices to show that the generators of $\mathbb{Z}\mathcal{C}(V)_{srr}$ are asymptotically equinumerous to 0. Indeed, if $T \subseteq (V \times V^\sigma) \times \mathbb{P}^1$ is an irreducible variety of dimension $d+1$ for which $T_0 \in \mathcal{C}(V)_{\text{sep}}$ and $T_\infty \in \mathcal{C}(V)_{\text{sep}}$, then taking $t \in \mathbb{P}^1$ to be its generic point, by Lemma 0.21, we have $T_0 \sim_a T_t \sim_a T_\infty$ so that $[T_\infty] - [T_0] \sim_a 0$. □

We have all the pieces in place to finish the proof of Theorem 0.6 from which Theorem 0.1 follows.

Proof. — As we discussed above (using localization or the Frobenius precomposition trick from Lemma 10.19), it suffices to consider $\Gamma \in \mathcal{C}(V)_{\mathrm{sep}}$. By Proposition 0.1 we may find $\Gamma' \in \mathbb{Z}\mathcal{C}(V)_f$ strongly reduced-rationally equivalent to Γ. By Lemma 0.23, $\Gamma \sim_a \Gamma'$ so that it suffices to prove the theorem for Γ'. That is, we may assume that Γ. Now that both projections $\Gamma \to V$ and $\Gamma \to V^\sigma$ are quasifinite, we see that (again for q large enough), the intersection $\Gamma \cap \Phi_q$ is proper so that $(\Gamma_{q,h} \cdot \Phi_q)$ is equal to $\#(\Gamma_{q,h} \cap \Phi_q)$ counted with intersection multiplicity. Over $((X,\Gamma)_{\text{ét}})_{q,h}$, the intersection multiplicity is one. Thus, it suffices to show that even counted with geometric multiplicity (which may be greater than the intersection multiplicity) the remaining points could contribute at most $O(q^{d-1})$ to the virtual intersection number. As we argued in the proof of Lemma 0.21, there are only $O(q^{d-1})$ points in lying outside of this locus. Indeed, $Z := V \smallsetminus (V,\Gamma)_{\text{ét}}$ is a proper subvariety of V so that by Corollary 9.24, even counting multiplicities, $|(V,\Gamma)_{q,h} \cap Z| = O(q^{d-1})$. □

Remark 0.24. — We have been working with a single fibre of the family presented with Hypotheses 0.2 but in Theorem 0.1 we ask for bounds which are independent of the choice of parameters. As far as I can see, such uniformity is not a formal consequence of the non-uniform version, but tracing the sources of these constants, it is not hard to see that they are uniform. The bounds for $(\Gamma \cdot \Phi_q)_{q,h}$ in Theorem 11.2 are explicit and are clearly independent of the parameters. It follows by compactness that the move from a general $\Gamma \in \mathcal{C}(V)$ to $\Gamma' \in \mathbb{Z}\mathcal{C}(V)_f$ is uniform in families. The implicit constants in the proof of Lemma 0.21 come from the bounds on the size of ${}^*E$ and from Corollary 9.24 to bound $(V,T_0)_{q,h}(K_q) \cap E$. The set ${}^*E$ is uniformly definable in the theory of strictly increasing transformal valued fields. As such, it admits a uniformly definable analysis in the residue field from which the bounds may be computed. Likewise, the proof of Corollary 9.24 is constructive and though explicit bounds depending only on the geometric data are not given, in principle, they may be extracted from the proof.

References

[1] Ehud Hrushovski, The elementary theory of the Frobenius automorphism, 24 July 2012, preprint, available at `http://math.huji.ac.il/~ehud/FROB.pdf`

[2] Alice Medvedev and Thomas Scanlon, Invariant varieties for polynomial dynamical systems, *Ann. of Math.*, to appear

THOMAS SCANLON • *E-mail :* `scanlon@math.berkeley.edu`, Department of Mathematics, University of California, Berkeley, Evans Hall, Berkeley, CA 94720-3840, USA
Url : `http://www.math.berkeley.edu/~scanlon`