COUNTING POINTS ON IMAGES

BHARGAV BHATT

ABSTRACT. Let $f : X \to Y$ be a morphism of varieties over a finite field $\mathbf{F}_q$. The goal of this note is to prove rationality of the generating function counting the image of the rational points of X under f. More precisely, we show that this function lies in the $\mathbf{Q}$-subspace of $\mathbf{Q}[\![t]\!]$ spanned by rational functions having "geometric type." In particular, the zeroes and poles of the associated zeta function are q-Weil numbers. These results are well-known to the experts.

1. INTRODUCTION

Fix a prime number p, and a finite extension $\mathbf{F}_q$ of the prime field $\mathbf{F}_p$. Our goal in this note is to prove the theorem mentioned in the abstract. Hence, we make the following definition:

Definition 1.1. Let $f : X \to Y$ be a morphism of finite type $\mathbf{F}_q$-schemes. We define the sets $f(\mathbf{F}_{q^n}) = \mathrm{im}(X(\mathbf{F}_{q^n}) \to Y(\mathbf{F}_{q^n}))$ for all $n \geq 1$. Organising this data into a formal power series, we define

$$\mathcal{Z}(f,t) = \sum_{n\geq 1} \#f(\mathbf{F}_{q^n}) \cdot t^n \in \mathbf{Q}[\![t]\!].$$

We also set $\mathcal{Z}(Y,t) = \mathcal{Z}(\mathrm{id},t)$ where $\mathrm{id} : Y \to Y$ is the identity morphism.

The series $\mathcal{Z}(Y,t)$ is the logarithmic derivative of the zeta function of Y. We prefer working with $\mathcal{Z}(Y,t)$ (and $\mathcal{Z}(f,t)$) instead of the associated zeta function because of the direct connection with point counting. Our main theorem is that $\mathcal{Z}(f,t)$ is a rational function; we use the rationality of $\mathcal{Z}(Y,t)$ for a variety Y. In fact, we prove a more precise statement locating the zeroes and poles of the associated zeta function:

Theorem 1.2. *Let $f : X \to Y$ be a morphism of finite type $\mathbf{F}_q$-schemes. Then $\mathcal{Z}(f,t)$ is a rational function. Moreover, there exist finite type $\mathbf{F}_q$-schemes $Z_1, \dots, Z_k$ over $\mathbf{F}_q$, rational numbers $a_1, \dots, a_k \in \mathbf{Q}$, and a rational polynomial $g(t) \in \mathbf{Q}[t]$ such that*

$$\mathcal{Z}(f,t) = g(t) + \sum_{i=1}^{k} a_i \mathcal{Z}(Z_i, t).$$

The outline for the paper is as follows. We first study some elementary properties enjoyed by the functions $\mathcal{Z}(f,t)$ in §2. Next, we prove Theorem 1.2 in the case that f is finite in §3, the crucial case being that of finite étale covers; Johan de Jong pointed out that this case can also be deduced from Grothendieck's work on the rationality of L-series associated to certain non-constant ℓ-adic sheaves. Lastly, we prove Theorem 1.2 in §4 in the general case by using the Weil conjectures on geometrically irreducible varieties to reduce to the case studied in §3.

Remark 1.3. It is tempting to try to prove Theorem 1.2 by a sheaf theoretic approach. One way to implement this strategy is to construct a complex $\mathcal{F}$ of constructible ℓ-adic sheaves of geometric origin on X whose stalk at $x \in X(\mathbf{F}_{q^n})$ has Frobenius trace given by the size of the fibre:

$$\mathrm{Tr}(\mathrm{Frob}_x \mid \mathcal{F}_x) = \frac{1}{\#f^{-1}(f(x))(\mathbf{F}_{q^n})}.$$

We do not how to do this[1]. In fact, as stated, this idea is not feasible because the existence of such a sheaf would give rationality for the zeta function associated to $\mathcal{Z}(f,t)$, which we know is true only up to a factor of an exponential of a rational polynomial (see Example 2.8).

Acknowledgements. We thank Mircea Mustaţă for asking the question addressed in this note.

2. Preliminaries on $\mathcal{Z}$-functions

The purpose of this section is to collect some elementary properties that the functions $\mathcal{Z}(f,t)$ satisfy. First, we record the relation between $\mathcal{Z}$ and the zeta function.

Proposition 2.1. *Let Y be an $\mathbf{F}_q$-scheme of finite type. Then we have an equality of formal power series:*

$$t \cdot \frac{d}{dt} \log(\zeta_{Y/\mathbf{F}_q}(t)) = \mathcal{Z}(Y,t)$$

Proof. Recall that the ζ function is defined via

$$\zeta_{Y/\mathbf{F}_q}(t) = \exp(\sum_{n \geq 1} \frac{\#Y(\mathbf{F}_{q^n})}{n} \cdot t^n)$$

Taking logarithmic derivatives and multiplying by t gives the desired identity. □

Next, we explain the topological invariance of $\mathcal{Z}(f,t)$.

Proposition 2.2. *Let $f : X \to Y$ and $f' : X' \to Y'$ be two morphisms of finite type $\mathbf{F}_q$-schemes. If $g_X : X \to X'$ and $g_Y : Y \to Y'$ are homeomorphisms commuting with f and f' in the obvious sense, then*

$$\mathcal{Z}(f,t) = \mathcal{Z}(f',t).$$

Proof. The equality follows from the fact that the set of $\mathbf{F}_{q^n}$-rational points of any $\mathbf{F}_q$-scheme is invariant under homeomorphisms. □

In a slightly orthogonal direction, we point out that $\mathcal{Z}(f,t)$ is invariant under dominating f by any map that induces surjections on rational points; we produce such maps in Proposition 4.9.

Proposition 2.3. *Let $f : X \to Y$ and $g : X' \to X$ be two morphisms of finite type $\mathbf{F}_q$-schemes. Assume that $g(\mathbf{F}_{q^n}) = X(\mathbf{F}_{q^n})$. Then*

$$\mathcal{Z}(f,t) = \mathcal{Z}(f \circ g, t).$$

Proof. The assumption implies that $f(\mathbf{F}_{q^n}) = (f \circ g)(\mathbf{F}_{q^n})$ for all n, whence the claim. □

Next, we explain two "constructibility" properties enjoyed by $\mathcal{Z}(f,t)$. The first of these allows us to cut the target of f into pieces, and is extremely useful in the sequel.

Proposition 2.4. *Let $f : X \to Y$ be a morphism of finite type $\mathbf{F}_q$-schemes. If Y_1 and Y_2 are locally closed subschemes of Y with $Y = Y_1 \cup Y_2$ as sets, then*

$$\mathcal{Z}(f,t) = \mathcal{Z}(f|_{f^{-1}(Y_1)}, t) + \mathcal{Z}(f|_{f^{-1}(Y_2)}, t) - \mathcal{Z}(f|_{f^{-1}(Y_1 \cap Y_2)}, t).$$

Proof. This follows from intersecting the inclusion-exclusion identity

$$\#Y(\mathbf{F}_{q^n}) = \#Y_1(\mathbf{F}_{q^n}) + \#Y_2(\mathbf{F}_{q^n}) - \#(Y_1 \cap Y_2)(\mathbf{F}_{q^n})$$

with $f(\mathbf{F}_{q^n})$. □

The second constructibility property allows us to break the source of f in a predictable manner. We do not actually use this identity.

[1]Note, however, that the complex $f^* \mathrm{R} f_! \mathbf{Q}_\ell$ on X does have a stalk whose Frobenius traces are inverses of the desired ones.

Proposition 2.5. *Let $f_1 : X_1 \to Y$ and $f_2 : X_2 \to Y$ be morphisms of finite type $\mathbf{F}_q$-schemes. Let $f_\sqcup : X_1 \sqcup X_2 \to Y$ and $f_\cap : X_1 \times_Y X_2 \to Y$ denote the morphisms obtained from the f_i by taking disjoint unions and fibre products, respectively. Then we have an equality*

$$\mathcal{Z}(f_\sqcup, t) = \mathcal{Z}(f_1, t) + \mathcal{Z}(f_2, t) - \mathcal{Z}(f_\cap, t).$$

Proof. The natural diagram of finite sets

$$\begin{array}{ccc} f_\cap(\mathbf{F}_{q^n}) & \xrightarrow{a} & f_1(\mathbf{F}_{q^n}) \\ \downarrow{\scriptstyle b} & & \downarrow \\ f_2(\mathbf{F}_{q^n}) & \longrightarrow & f_\sqcup(\mathbf{F}_{q^n}) \end{array}$$

is a pushout diagram with a and b injective. The claim follows by counting cardinalities. □

Question 2.6. The functions $\mathcal{Z}(f, t)$ enjoy many "motivic" properties as discussed above. Hence, it seems reasonable seek a motivic explanation for Theorem 1.2. In particular, can one, a priori, predict the Weil numbers occuring in the associated zeta function?

Next, we introduce a definition to facilitate notational brevity: we single out the subspace of $\mathbf{Q}[\![t]\!]$ spanned by functions of type occuring in the conclusion of Theorem 1.2.

Definition 2.7. A power series $g \in \mathbf{Q}[\![t]\!]$ has *q-geometric type* (or simply *geometric type* if q is fixed) if it lies in the $\mathbf{Q}$-linear subspace of $\mathbf{Q}[\![t]\!]$ generated by the polynomial ring $\mathbf{Q}[t]$ and the series $\mathcal{Z}(W, t)$ as W runs through finite type $\mathbf{F}_q$-schemes.

By the rationality of the zeta function and Proposition 2.1, any power series having geometric type is a rational function. Hence, Theorem 1.2 can be reformulated as stating that $\mathcal{Z}(f, t)$ has geometric type for any morphism f of finite type $\mathbf{F}_q$-schemes. We close this section with an example.

Example 2.8. We give an explicit example of a map $f : X \to Y$ with the property that $\mathcal{Z}(f, t) - \mathcal{Z}(Y, t)$ is a non-trivial rational polynomial $g(t)$. It follows then that the zeta function associated to f differs multiplicatively from the zeta function of Y by $e^{\int g(t)}$. In particular, it is not a rational function. For the example, let $Y = \mathrm{Spec}(\mathbf{F}_q)$, and let $X = \mathbf{A}^1_{\mathbf{F}_q} - \mathbf{A}^1(\mathbf{F}_q)$. Then we have $\mathcal{Z}(f, t) - \mathcal{Z}(Y, t) = t$. One can also produce a smooth projective geometrically irreducible example by simply picking such a variety without a rational point; this can be done with curves of high genus over small finite fields (by a Bertini argument).

3. Finite morphisms

Our goal in this section is to explain the proof of Theorem 1.2 in the special case that the map f is *finite* and surjective. The bulk of the work lies in verifying the claims when f is finite étale; the rest is handled by devissage. To deal with the case of finite étale covers, we use some formal properties about classifying stacks of finite groups over finite fields. Hence, we review the relevant facts about these objects in §3.1. Having done that, we prove Theorem 1.2 in the case of finite étale maps in §3.2, and deal with general finite maps in §3.3.

3.1. **Classifying stacks of finite groups over finite fields.**

In this section we introduce an equivalence relation called m-equivalence (m for monodromy) on the points of the stacks in question, and study its effect on finite field valued points. The relevant definitions are:

Definition 3.1. Let Y be a noetherian connected scheme, let G be a finite group, and let $f_1 : X_1 \to Y$ and $f_2 : X_2 \to Y$ be two G-torsors over Y. We say that f_1 and f_2 are *equivalent* if the monodromy representations $\pi_1(X) \to G$ associated to f_1 and f_2 are conjugate to one another. We say that f_1 and f_2 are *m-equivalent* if the images of the monodromy representations $\pi_1(X) \to G$ associated to f_1 and f_2 are conjugate to one another.

The notion of equivalence introduced in Definition 3.1 coincides with isomorphisms for torsors and is standard. On the other hand, the m-equivalence relation introduced here is a little stranger. We use this notion to define and count certain loci in a family of G-torsors (see Theorem 3.8). When $X = \mathrm{Spec}(\mathbf{F}_q)$ (which is the only relevant case for applications), there is an explicit descriptions of m-equivalence classes:

Proposition 3.2. *Let $X = \mathrm{Spec}(\mathbf{F}_q)$ be the spectrum of a finite field, let G be a finite group, and let $B(G)(X)$ denote the groupoid of G-torsors on X.*

(1) *The set of equivalence classes $\pi_0(B(G)(X))$ of G-torsors is in bijective correspondence with the set G_{conj} of conjugacy classes in G. The set $\pi_1(B(G)(X), g)$ of automorphisms of a G-torsor corresponding to an element g is given by the centralizer $\mathrm{Cent}_G(g)$.*
(2) *Given two elements $g_1, g_2 \in G$, the G-torsors corresponding to the conjugacy classes of g_1 and g_2 are m-equivalent if and only if the subgroups they generate are conjugate, i.e., one has*

$$g_1 = h \cdot g_2^{\ell} \cdot h^{-1}$$

where $h \in G$ is some element, and ℓ is some integer coprime to the order of g_2 (and hence g_1). In particular, the m-equivalence relation respects the automorphism groups.

Proof. The fundamental group of X is identified with $\widehat{\mathbf{Z}}$ using the geometric Frobenius element $x \mapsto x^{-q}$. This description allows us to identify the groupoid $B(G)(X)$ with the groupoid $\mathrm{Map}(B(\widehat{\mathbf{Z}}), B(G))$ of morphisms $B(\widehat{\mathbf{Z}}) \to B(G)$. A model for this groupoid is provided by a category with object set G_{conj}, and automorphism groups $\mathrm{Cent}_G(g)$ at an object corresponding to g; this description implies the first half of the proposition.

An element $g \in G$ determines a homomorphism $\widehat{\mathbf{Z}} \to G$ whose image is the cyclic subgroup generated by g. Hence, two elements g_1 and g_2 determine m-equivalent torsors precisely when the subgroups they generate are conjugate. This last condition is easily seen to be equivalent to the condition in statement of Proposition. Lastly, the claim about automorphism groups being respected is simply the observation that $\mathrm{Cent}_G(g)$ depends only on conjugacy class of the subgroup generated by g. □

Warning 3.3. Proposition 3.2 describes m-equivalence classes of G-torsors over finite fields in terms of cyclic subgroups of G up to conjugation. In the sequel we will often abuse notation by identifying an m-equivalence class of a G-torsor over a finite field in terms of an element of G representing a cyclic subgroup in the m-equivalence class.

Remark 3.4. Fix a finite group G. By Proposition 3.2, the m-equivalence relation $H^1(\mathrm{Spec}(\mathbf{F}_q), G)$ corresponds to "conjugacy of group generated" under the identification of the cohomology set with G_{conj}. The gap between m-equivalence and conjugation-equivalence depends on the flavour of finite groups involved. For G a cyclic group, conjugating an element does nothing, while m-equivalance remembers only the order of the element. On the other hand, for symmetric groups, m-equivalence and conjugation-equivalence coincide. For applications to Theorem 1.2, it suffices to work with symmetric groups.

We close this section with two examples illustrating the utility of the formalism surrounding $B(G)$.

Example 3.5. One can use classifying stacks of finite groups over finite fields to prove group theoretic identities sometimes. For example, for a fixed finite group G, the groupoid $B(G)(\mathbf{F}_q)$ is identified with the groupoid $\mathrm{Map}(B(\mathbf{Z}), B(G))$ by Proposition 3.2. One model for this groupoid has object set G_{conj} with the set of automorphisms of the conjugacy class associated to an element g being the centralizer $\mathrm{Cent}_G(g)$. Counting points in the groupoid sense, we see that

$$\#B(G)(\mathbf{F}_q) = \sum_{g \in G_{\mathrm{conj}}} \frac{1}{\#\mathrm{Cent}_G(g)}.$$

On the other hand, by the Lefschetz trace formula and the $\mathbf{Q}$-acyclicity of $B(G)$, we see that $\#B(G)(\mathbf{F}_q) = 1$. Hence, we obtain the well-known identity

$$1 = \sum_{g \in G_{\mathrm{conj}}} \frac{1}{\#\mathrm{Cent}_G(g)}.$$

Of course, using the Lefschetz trace formula and algebraic stacks seems overkill for a simple group theory formula, but we believe it illustrates a useful organisational principle.

Example 3.6. We explain how to prove that the Brauer group of a finite field vanishes using the formalism of classifying spaces of group-stacks, i.e., with 2-stacks. Let μ_n be the group scheme of n-roots of unity over a finite field $\mathbf{F}_q$ with n prime to p. Then the 2-groupoid $BB(\mu_n)(\mathbf{F}_q)$ is the simplicial abelian group associated to the complex $\big(\tau_{\leq 2}\mathrm{R}\Gamma(\mathrm{Spec}(\mathbf{F}_q), \mu_n)\big)[2]$ via the Dold-Kan correspondence. Hence, its groupoid-cardinality is given by the formula

$$\#BB(\mu_n)(\mathbf{F}_q) = \sum_{\psi \in H^2(\mathrm{Spec}(\mathbf{F}_q), \mu_n)} \frac{\#H^1(\mathrm{Spec}(\mathbf{F}_q), \mu_n)}{\#H^0(\mathrm{Spec}(\mathbf{F}_q), \mu_n)}. \tag{1}$$

Using the Kummer sequence and Hilbert's theorem 90 shows that $H^0(\mathrm{Spec}(\mathbf{F}_q), \mu_n)$ and $H^1(\mathrm{Spec}(\mathbf{F}_q), \mu_n)$ are the kernel and cokernel of the n-th power map $\mathbf{F}_q^* \to \mathbf{F}_q^*$ of the finite abelian group $\mathbf{F}_q^*$. Hence, each of summands on the right hand side of formula (1) has size 1 which gives

$$\#BB(\mu_n)(\mathbf{F}_q) = \sum_{\psi \in H^2(\mathrm{Spec}(\mathbf{F}_q), \mu_n)} 1.$$

On the other hand, the 2-stack $BB(\mu_n)_{\overline{\mathbf{F}_q}}$ is $\mathbf{Q}$-acyclic. Hence, by the Lefschetz trace formula, it has cardinality 1. Thus, we see that

$$H^2(\mathrm{Spec}(\mathbf{F}_q), \mu_n) = 1$$

which recovers the well-known fact that the Brauer group of a finite field vanishes, up to p-torsion.

3.2. **The étale case.** We now prove Theorem 1.2 for finite étale morphisms. The idea is to first deal with the Galois case where we formulate and prove a more general statement than what is required. Using this extra generality, we reduce the general case to the Galois case. To formulate the more general statement in the Galois case, we define certain loci on the base of a family of G-torsors as follows.

Definition 3.7. Let $f : X \to Y$ be a finite étale cover that is a G-torsor for a finite group G. For any element $g \in G$, we define

$$Y(\mathbf{F}_{q^n})_g = \{y \in Y(\mathbf{F}_{q^n}) \mid \text{The G-torsor } f^{-1}(y) \text{ has m-equivalence class } g\}.$$

We assemble these sets into a generating function

$$\mathcal{Z}(f_g, t) = \sum_{n \geq 1} \#Y(\mathbf{F}_{q^n})_g \cdot t^n \in \mathbf{Q}[\![t]\!].$$

Note that $\mathcal{Z}(f_e, t) = \mathcal{Z}(f, t)$ where $e \in G$ is the identity element. Moreover, one can easily see that

$$\sum \mathcal{Z}(f_g, t) = \mathcal{Z}(Y, t)$$

where the sum takes place over all m-equivalence classes. In particular, the sum on the left has geometric type. The content of Theorem 1.2 for finite Galois covers is the statement that $\mathcal{Z}(f_e, t)$ has geometric type. We will show more generally that *each* of the summands of the sum on the left has geometric type.

Theorem 3.8. *Let G be a finite group, and $g \in G$ be an element. For an $\mathbf{F}_q$-scheme Y and a G-torsor $f : X \to Y$, the formal power series $\mathcal{Z}(f_g, t)$ has geometric type.*

Proof. We will show the rationality of $\mathcal{Z}(f_g, t)$ by induction on the order d of g (as noted in Proposition 3.2, the order of an element depends only on its m-equivalence class). To start the induction, let $g = e$ be the identity element. In this case, the sets $Y(\mathbf{F}_{q^n})_e$ are precisely the sets $f(\mathbf{F}_{q^n})$. Since f is G-equivariant, there is a natural surjection $X(\mathbf{F}_{q^n})/G \to f(\mathbf{F}_{q^n})$. This map is also injective since f is a G-cover. Hence, we find that

$$\#Y(\mathbf{F}_{q^n})_e = \#f(\mathbf{F}_{q^n}) = \#(X(\mathbf{F}_{q^n})/G).$$

Since f is étale, the G-action on $X(\mathbf{F}_{q^n})$ has no fixed points, which shows that

$$\#Y(\mathbf{F}_{q^n})_e = \#f(\mathbf{F}_{q^n}) = \#(X(\mathbf{F}_{q^n})/G) = \frac{\#X(\mathbf{F}_{q^n})}{\#G}.$$

Thus, we obtain

$$\mathcal{Z}(f_e, t) = \frac{1}{\#G} \cdot \mathcal{Z}(X, t)$$

showing that $\mathcal{Z}(f_e, t)$ has geometric type.

Proceeding inductively, fix an element $g \in G$ of order d, and assume that $\mathcal{Z}(f_h, t)$ is known to have geometric type for all m-equivalence classes represented by elements $h \in G$ with smaller order. Let $i_g : B(\mathbf{Z}/d) \to B(G)$ be the 1-morphism of algebraic stacks (over $\mathbf{Z}$) defined by g. This morphism is representable by algebraic spaces because d is the order of g. The G-torsor f can be viewed as a map $\phi : Y \to B(G)$. Taking fibre products, we obtain a cartesian diagram

$$\begin{array}{ccc} \mathrm{Fib}(\phi, i_g) := B(\mathbf{Z}/d) \times_{B(G)} Y & \longrightarrow & B(\mathbf{Z}/d) \\ \downarrow & & \downarrow{\scriptstyle i_g} \\ Y & \xrightarrow{\quad\phi\quad} & B(G). \end{array}$$

Since i_g was representable, the stack $\mathrm{Fib}(\phi, i_g)$ is actually an algebraic space. The definition of fibre products allows us to identify the set $\mathrm{Fib}(\phi, i_g)(\mathbf{F}_{q^n})$ with the discrete category of 3-tuples (a, b, c) where

- $a \in B(\mathbf{Z}/d)(\mathbf{F}_{q^n})$ is a rational point of $B(\mathbf{Z}/d)$.
- $b \in Y(\mathbf{F}_{q^n})$ is a rational point of Y.
- $c : i_g(a) \xrightarrow{\sim} \phi(b)$ is an isomorphism between the images of a and b in $B(G)(\mathbf{F}_{q^n})$.

In particular, there is a natural composite map

$$\pi : \mathrm{Fib}(\phi, i_g)(\mathbf{F}_{q^n}) \to \pi_0(B(\mathbf{Z}/d)(\mathbf{F}_{q^n})) = H^1(\mathrm{Spec}(\mathbf{F}_{q^n}), \mathbf{Z}/d) \to H^1(\mathrm{Spec}(\mathbf{F}_{q^n}), \mathbf{Z}/d)/\sim$$

where "$\sim$" denotes the m-equivalence relation on $H^1(\mathrm{Spec}(\mathbf{F}_{q^n}), \mathbf{Z}/d)$. The group $H^1(\mathrm{Spec}(\mathbf{F}_{q^n}), \mathbf{Z}/d)$ can be identified with $\mathbf{Z}/d$ with the element 1 in the latter corresponding to the non-trivial degree d extension of $\mathbf{F}_{q^n}$. Under this identification, the set of m-equivalence classes of $H^1(\mathrm{Spec}(\mathbf{F}_{q^n}), \mathbf{Z}/d)$ is identified with the set of subgroups of $\mathbf{Z}/d$ which, in turn, can be viewed as the set of divisors i of d (with i corresponding to the unique subgroup of order i). Hence, we can use the preceding map to write

$$\#\mathrm{Fib}(\phi, i_g)(\mathbf{F}_{q^n}) = \sum_{i|d} \#\pi^{-1}(i).$$

A little bit of thought reveals now that $\pi^{-1}(i)$ can be identified with the set

$$\Big(Y(\mathbf{F}_{q^n})_{g^{\frac{d}{i}}} \times \mathrm{Cent}_G(g^{\frac{d}{i}})\Big)/\mathrm{Cent}_{\mathbf{Z}/d}(\frac{d}{i}).$$

The centralizer $\mathrm{Cent}_{\mathbf{Z}/d}(\frac{d}{i})$ is $\mathbf{Z}/d$ as $\mathbf{Z}/d$ is abelian. Hence, after reindexing with $j = \frac{d}{i}$, we find

$$\#\mathrm{Fib}(\phi, i_g)(\mathbf{F}_{q^n}) = \sum_{j|d} \#Y(\mathbf{F}_{q^n})_{g^j} \cdot \frac{\#\mathrm{Cent}_G(g^j)}{d}.$$

Since the preceding equality holds for all n, we obtain an equality of formal power series

$$\mathcal{Z}(\mathrm{Fib}(\phi, i_g), t) = \sum_{j|d} \frac{\#\mathrm{Cent}_G(g^j)}{d} \cdot \mathcal{Z}(f_{g^j}, t).$$

By induction, we may assume that $\mathcal{Z}(f_{g^j}, t)$ has geometric type for $j \mid d$ with $j \neq 1$. Stratifying the algebraic space $\mathrm{Fib}(\phi, i_g)$ by finite type $\mathbf{F}_q$-schemes and using Proposition 2.4 shows that $\mathcal{Z}(\mathrm{Fib}(\phi, i_g), t)$ also has gemetric type. As the coefficient $\frac{\#\mathrm{Cent}_G(g)}{d}$ of $\mathcal{Z}(f_g, t)$ occuring on the right hand side of formula above is non-zero, it follows that $\mathcal{Z}(f_g, t)$ has geometric type. □

Question 3.9. Can Theorem 3.8 be refined to count equivalence classes rather than m-equivalence classes? We do not know if the corresponding generating function has geometric type, or is even rational.

We now show how to prove Theorem 1.2 in the case of finite étale morphisms by reducing to the Galois case and using Proposition 3.8.

Corollary 3.10. *Let $f : X \to Y$ be a finite étale cover. Then the formal power series $\mathcal{Z}(f, t)$ has geometric type.*

Proof. By Proposition 2.4, we may assume that f is surjective and that Y is connected. Let d denote the degree of f, and let Σ_d denote the symmetric group on d elements. The map f can be viewed as a twisted form of the of map $\sqcup_{i=1}^{d} Y \to Y$, and thus defines an Σ_d-torsor over Y. More explicitly, the Σ_d-torsor $\phi : X' \to Y$ corresponding to f is the map taking the Galois closure of f in the fibres (see Construction 3.11). For any map $y : \mathrm{Spec}(\mathbf{F}_{q^n}) \to Y$, the pullback of X_y of X along y has a $\mathbf{F}_{q^n}$-rational point if and only if conjugacy class in Σ_d determined by X_y can be represented by a partition of d containing 1 as one of its entries. Since conjugacy classes in Σ_d coincide with m-equivalence classes, we see that $X_y \to \mathrm{Spec}(\mathbf{F}_{q^n})$ has a section if and only if the m-equivalence class of $X'_y \to \mathrm{Spec}(\mathbf{F}_{q^n})$ is represented by a partition which contains 1 as one of its entries. Thus, we obtain an identity

$$\mathcal{Z}(f, t) = \sum_{\substack{p \in (\Sigma_d)_{\mathrm{conj}}, \\ p \text{ fixes a point.}}} \mathcal{Z}(\phi_p, t)$$

where the sum takes place over all conjugacy classes p in the symmetric group corresponding to a partition containing a 1. By Theorem 3.8, the right side has geometric type, and thus so does the left side. □

We explain a construction used in the proof of Corollary 3.10: the Galois closure of a finite étale cover.

Construction 3.11. Let $f : X \to Y$ be a finite étale degree d cover. We will explain an inductive construction of the Galois closure of f[2]. We define inductively a sequence of Y-schemes $X(1)$, $X(2)$, …, $X(d)$, actions of Σ_j on $X(j) \to Y$, and maps $\Delta_j^i : X(i) \to X(i) \times_Y X$ indexed by $1 \leq j \leq i$ as follows: set $X(1) = X$ with Δ_1^1 the diagonal map and the unique Σ_1-action, and define

$$X(i) = X(i-1) \times_Y X - \sqcup_{j=1}^{i-1} \Delta_j^{i-1}(X(i-1)).$$

The scheme $X(i)$ can be viewed as the subscheme of i-fold fibre product of X over Y defined by the complement of all the diagonals. With this intrepretation, the action of Σ_i on $X(i) \to Y$ is clear, while the maps Δ_j^i are projection maps onto the j-th co-ordinate. Since f is étale, all the maps Δ_j^i occuring above are also étale. Hence, we compute easily that $X(j) \to Y$ is a finite étale cover of degree $d \cdot (d-1) \cdot \cdots \cdot (d-j+1)$. In particular, $X(d) \to Y$ is a finite étale cover of degree $d!$ admitting an action of Σ_d. All these constructions are compatible with base changes on Y, and an easy calculation then shows that $X(d) \to Y$ is actually an Σ_d-torsor for the natural Σ_d action on $X(d)$. We call $X(d) \to Y$ the Galois closure of $X \to Y$.

[2] One could simply define $X(d) \to Y$ as the space over Y parametrising d-distinct points in the fibres of f, but we prefer the inductive approach due to its computational utility.

Remark 3.12. We briefly indicate how one can use Theorem 3.8 to prove a version of the Chebotarev density theorem for function fields. Let L/K be a finite Galois extension with group G of function fields in one variable over the finite field $\mathbf{F}_q$, and let $f : C_L \to C_K$ be the corresponding map of smooth projective curves over $\mathbf{F}_q$. We let $U_K \subset C_K$ be the complement of the discriminant of f, and let $U_L = f^{-1}(U_K)$; the resulting map $f : U_L \to U_K$ is finite étale with group G. For a closed point x of U_K, we let $\mathrm{Art}_{L/K}(x)$ be the Artin symbol of x in the group G. If y is a closed point of U_L above x, then $\mathrm{Art}_{L/L}(x)$ may be viewed as the conjugacy class in G associated to a geometric Frobenius element generating $\mathrm{Gal}(\kappa(y)/\kappa(x))$. The Chebotarev density theorem predicts that

$$\lim_{n\to\infty} \frac{\#\{x \in U_K(\mathbf{F}_{q^n}) \mid \mathrm{Art}_{L/K}(x) = \{e\}\}}{\#U_K(\mathbf{F}_{q^n})} = \frac{1}{\#G}.$$

Geometrically, the condition that $\mathrm{Art}_{L/K}(x) = \{e\}$ translates to the condition that $f^{-1}(x)$ is a trivial G-torsor. Hence, the limit occuring above may be rewritten as

$$\begin{aligned}
\lim_{n\to\infty} \frac{\#\{x \in U_K(\mathbf{F}_{q^n}) \mid \mathrm{Art}_{L/K}(x) = \{e\}\}}{\#U_K(\mathbf{F}_{q^n})} &= \lim_{n\to\infty} \frac{\#\{x \in U_K(\mathbf{F}_{q^n}) \mid f^{-1}(x)(\mathbf{F}_{q^n}) \neq \emptyset\}}{\#U_K(\mathbf{F}_{q^n})} \\
&= \lim_{n\to\infty} \frac{\#f(\mathbf{F}_{q^n})}{\#U_K(\mathbf{F}_{q^n})} \\
&= \lim_{n\to\infty} \frac{\#U_L(\mathbf{F}_{q^n})}{\#U_K(\mathbf{F}_{q^n})} \cdot \frac{1}{\#G}
\end{aligned}$$

where the last equality used first step of Theorem 3.8 (the base case of induction). The geometric content of Chebotarev is therefore the statement that the generating functions associated to U_L and U_K grow similarly. This follows from standard theorems in étale cohomology: the Lefschetz trace formula and the Weil bounds reduce us to showing that generating functions associated to the Galois representations $H^2_c(U_K \times_{\mathbf{F}_q} \overline{\mathbf{F}_q}, \mathbf{Q}_\ell)$ and $H^2_c(U_L \times_{\mathbf{F}_q} \overline{\mathbf{F}_q}, \mathbf{Q}_\ell)$ grow similarly. Poincare duality identifies both these representations with $\mathbf{Q}_\ell(-1)$, and the claim follows.

3.3. **The general finite case.** We use the results of §3.2 to handle arbitrary finite morphisms in the present section. First, we deal with the flat ones.

Proposition 3.13. *Let $f : X \to Y$ be a finite flat morphism of finite type $\mathbf{F}_q$-schemes. Then the formal power series $\mathcal{Z}(f, t)$ has geometric type.*

Proof. We prove the claim by induction on $\dim(Y)$. By Proposition 2.2, we may assume that Y is reduced. Let $U \subset Y$ denote the largest normal open subscheme of Y, and let Z be its complement. Since Y is reduced, the open subscheme U is dense in Y, and $\dim(Z) < \dim(Y)$. By Proposition 2.4 and induction on dimension, it suffices to show that $\mathcal{Z}(f|_{f^{-1}(U)}, t)$ is rational, i.e., we may assume that Y is normal.

The map f can then be factored as $X \xrightarrow{g} X' \xrightarrow{h} Y$ with g purely inseparable, and h generically separable. Since $X \to X'$ is a universal homeomorphism, we see that $X(\mathbf{F}_{q^n}) = X'(\mathbf{F}_{q^n})$ by Proposition 2.2. Thus, after replacing X with X' and f with h, we may assume that the map f being considered is generically separable. In this case, there is an open subscheme $U \subset Y$ with complement Z such that $f|_{f^{-1}(U)}$ is finite étale, and $\dim(Z) < \dim(Y)$. The desired claim now follows by induction using Proposition 2.4 and Corollary 3.10. $\square$

We can now handle arbitrary finite maps by devissage.

Theorem 3.14. *Let $f : X \to Y$ be a finite morphism of finite type $\mathbf{F}_q$-schemes. Then the formal power series $\mathcal{Z}(f, t)$ has geometric type.*

Proof. We proceed by induction on $\dim(Y)$. As in the first half of the proof of Proposition 3.13, we reduce to the case that Y is normal. By generic flatness, there exists an open subscheme $U \subset Y$ with complement Z

such that $f|_{f^{-1}(U)}$ is finite flat, and $\dim(Z) < \dim(Y)$. The claim now follows by induction on dimension using Propositions 2.4 and 3.13. □

4. The main theorem

Our goal in the present section is to prove Theorem 1.2. We do so by first examining the case of geometrically irreducible varieties in §4.1, extending the preceding to families in §4.2, and then proving the main theorem in §4.3.

4.1. Finding points on geometrically irreducible varieties.

We show that geometrically irreducible varieties over finite fields always have points after a sufficiently large field extension that depends only on the topology of the variety. This result is well-known, but we include a proof for completeness.

Proposition 4.1. *Let X be a geometrically irreducible $\mathbf{F}_q$-scheme of finite type of dimension d, and fix an auxilliary prime ℓ different from p. Then there exists a positive integer n_0 depending only on the compactly supported ℓ-adic Betti numbers of X such that for all $n \geq n_0$, the set $X(\mathbf{F}_{q^n})$ is non-empty.*

Proof. We fix an embedding $\iota : \overline{\mathbf{Q}_\ell} \hookrightarrow \mathbf{C}$ that allows us to talk about archimedean absolute values of ℓ-adic numbers. By the Lefschetz trace formula, we have

$$X(\mathbf{F}_{q^n}) = \sum_{i \geq 0} (-1)^i \mathrm{Tr}(\mathrm{Frob}^n \mid H^i_c(X_{\overline{\mathbf{F}_q}}, \mathbf{Q}_\ell)).$$

Artin-Grothendieck theory tells us the following:

- The groups $H^i_c(X_{\overline{\mathbf{F}_q}}, \mathbf{Q}_\ell)$ vanish for $i > 2d$.
- There is a natural isomorphism of representations $H^{2d}_c(X_{\overline{\mathbf{F}_q}}, \mathbf{Q}_\ell) \simeq \mathbf{Q}_\ell(-d)$.

These facts allow us to rewrite the trace formula as

$$X(\mathbf{F}_{q^n}) = q^{nd} + \sum_{i=0}^{2d-1} a_i$$

with a_i being the signed trace of Frobenius on $H^i_c(X_{\overline{\mathbf{F}_q}}, \mathbf{Q}_\ell)$. Hence, $X(\mathbf{F}_{q^n}) = \emptyset$ if and only if

$$q^{nd} = -\sum_{i=0}^{2d-1} a_i.$$

Deligne's results give the inequality

$$|a_i| \leq b_i(X) \cdot q^{\frac{ni}{2}}$$

where $b_i(X) = \dim H^i_c(X_{\overline{\mathbf{F}_q}}, \mathbf{Q}_\ell)$ is the i-th compactly supported Betti number of X. Hence, if $X(\mathbf{F}_{q^n}) = \emptyset$, then

$$q^{nd} = -\sum_{i=0}^{2d-1} a_i \leq |\sum_{i=0}^{2d-1} a_i| \leq \sum_{i=0}^{2d-1} |a_i| \leq \sum_{i=0}^{2d-1} b_i(X) \cdot q^{\frac{ni}{2}}$$

As the left hand side grows much faster than the right side as a function of n, the preceding inequality exists only for finitely many values for n once the Betti numbers $b_i(X)$ have been fixed. Picking an n_0 bigger than any integer occuring in this finite set therefore does the job. □

Remark 4.2. The proof of Proposition 4.1 given above only uses the knowledge of the asymptotic number of points on a d-dimensional geometrically irreducible variety over a finite field (and a crude estimate for the error term). Hence, it can be deduced from the Lang-Weil estimates which are considerably more elementary than the full Weil conjectures used above.

Remark 4.3. The integer n_0 provided by Proposition 4.1 can be chosen to be independent of the auxilliary prime ℓ. In fact, one can show that there are only finitely many possibilities for the compactly supported ℓ-adic Betti numbers of X as the prime ℓ varies as follows: if X is smooth and proper, then the Weil conjectures show that its Betti numbers are independent of ℓ and the claim follows. If X is proper but not necessarily smooth, we can find a proper hypercovering of X by smooth projective varieties thanks to de Jong's theorems. Such a hypercover allows us to bound the ℓ-adic Betti numbers of X in terms of the Betti numbers of a finite number of auxilliary smooth projective varieties. As Betti numbers are always non-negative, the claim follows. In general, fixing a closed immersion of X into a proper variety $\overline{X}$ allows us to bound the Betti numbers of X in terms of those of $\overline{X}$ and $\overline{X} - X$, whence the claim follows.

Remark 4.4. The bounds provided by Proposition 4.1 can be made quite explicit in specific examples. For example, if C is a smooth projective curve of genus g over $\mathbf{F}_q$, then we have $a_0 = 1$, $a_2 = q$, and $|a_1| \leq 2g\sqrt{q}$ (in the notation of Proposition 4.1). Thus, we find that $C(\mathbf{F}_q) \neq \emptyset$ as long as $2g \leq \sqrt{q}$.

4.2. **Bounding homotopy types in a family.** To extend the results of §4.1 to families with geometrically irreducible fibres, we need to ensure that only finitely homotopy types can occur as the fibres of such a family. We show a truncated version that is sufficient for applications. Note that the required boundedness also follows from the Lang-Weil estimates for points on varieties in an algebraic family, but we prefer giving a self-contained and slightly more general argument.

Proposition 4.5. *Let $f : X \to S$ be a proper morphism of noetherian schemes. Fix a prime number ℓ invertible on S (assumed to exist), and an auxilliary positive integer k (the "level"). Then there exists a flag of closed subschemes $\emptyset = S_0 \subset S_1 \subset \cdots \subset S_m \subset S_{m+1} = S$ in S such that for any two geometric points a, b lying in $S_i - S_{i-1}$, the k-truncated pro-ℓ homotopy types of X_a and X_b are canonically isomorphic.*

Proof. The idea is to prove the general case by using alterations and hypercoverings to reduce to the case of proper smooth morphisms, which follows from work of Artin and Mazur.

In more detail, since the statement to be proven is invariant under passage to finite covers of S, we may assume that S is integral. By noetherian induction, it suffices to show that there exists a single open subscheme $U \subset S$ such that the k-truncated pro-ℓ homotopy type is constant in the geometric fibres of $X_U \to U$. Let η denote the generic point of U. After replacing S by a finite cover, we may use de Jong's theorems to construct a k-truncated proper hypercover $X_{\cdot,\eta} \to X_\eta$ with each $X_{i,\eta}$ smooth and proper over η. By spreading out, we can find an open subscheme $U \subset S$ and a k-truncated proper hypercover $X_{\cdot,U} \to X_\eta$ with each $X_{i,U} \to U$ smooth and proper. By the Artin-Mazur theorem (see [AM69, Corollary 12.13]), the pro-ℓ homotopy types of the geometric fibres of $X_{i,U} \to U$ are canonically isomorphic[3]. The canonicity implies that face and degeneracy maps of $X_{\cdot,U}$ also induce the same maps at the level of pro-ℓ homotopy types of the geometric fibres. It follows then that the k-truncated pro-ℓ homotopy types of the geometric fibres of $X_U \to U$ are also canonically identified. □

Using Proposition 4.5, we bound the Betti numbers occuring in the family of algebraic varieties.

Corollary 4.6. *Let $f : X \to S$ be a proper morphism of noetherian schemes. Fix a prime number ℓ invertible on S (assumed to exist). Then there exists a flag of closed subschemes $\emptyset = S_0 \subset S_1 \subset \cdots \subset S_m \subset S_{m+1} = S$ in S such that for any two geometric points a, b lying in $S_i - S_{i-1}$, the complexes $\mathrm{R}\Gamma(X_a, \mathbf{Z}_\ell)$ and $\mathrm{R}\Gamma(X_b, \mathbf{Z}_\ell)$ are quasi-isomorphic.*

Proof. The idea is to use the Artin-Grothendieck vanishing theorems in étale cohomology to recover $\mathrm{R}\Gamma(X, \mathbf{Z}_\ell)$ from the k-truncated pro-ℓ homotopy type of an algebraic variety X for an appropriate choice of k, and then use Proposition 4.5.

[3] Artin and Mazur only state their theorem over discrete valuation rings, but one can reduce to that case easily. The canonicity is not explicitly spelled out in the statement, but follows from the proof of the proper and smooth base change theorems in étale cohomology.

In more detail, fix an integer n. By stratifying S and working with the strata, we may assume that all fibres of f have dimension the same dimension d. Now let s be a geometric point of S. Then Artin-Grothendieck vanishing tells us that $\mathrm{R\Gamma}(X_s, \mathbf{Z}/\ell^n) \in D^{[0,2d]}(\mathbf{Z}/\ell^n)$. Hence, we have an equivalence

$$\tau_{\geq -2d}\big(\mathrm{R\Gamma}(X_s, \mathbf{Z}/\ell^n)[2d]\big) \simeq \mathrm{R\Gamma}(X_s, \mathbf{Z}/\ell^n)[2d]. \tag{2}$$

On the other hand, for any integer k, we have an identification

$$\mathrm{DK}\Big(\tau_{\geq -k}\big(\mathrm{R\Gamma}(X_s, \mathbf{Z}/\ell^n)[k]\big)\Big) \simeq \mathrm{Map}(\widehat{X_s}, K(\mathbf{Z}/\ell^n, k)) \simeq \mathrm{Map}(\tau_{\leq k}\widehat{X_s}, K(\mathbf{Z}/\ell^n, k)). \tag{3}$$

Here $\mathrm{DK}(\cdot)$ denotes the Dold-Kan functor from non-positively graded cohomological chain complexes to simplicial abelian groups, $K(\mathbf{Z}/\ell^n, k)$ denotes the Eilenberg-Maclane space of degree k on the abelian group $\mathbf{Z}/\ell^n$, the space $\widehat{X_s}$ is the pro-ℓ homotopy type of X_s, the object $\mathrm{Map}(\cdot,\cdot)$ denotes the mapping space in the profinite homotopy category, and τ denotes the Postnikov truncation functor (on either complexes or spaces, with opposite conventions).

The upshot of formula (3) is that the $\mathbf{Z}/\ell^n$-cohomology of X_s, up to degree k, can be recovered from its k-truncated homotopy type $\tau_{\leq k}\widehat{X_s}$. Using formula (2), we see that the entire $\mathbf{Z}/\ell^n$-cohomology of X_s can be recovered from the $2d$-truncated homotopy type $\tau_{\leq 2d}\widehat{X_s}$ where $d = \dim(X_s)$. The desired claim now follows from Proposition 4.5 by taking an inverse limit over n. □

Remark 4.7. Proposition 4.5 remains valid for $k = \infty$, i.e., one can finitely partition S in terms of the *entire* homotopy types of the fibres of f. This follows more-or-less formally from Corollary 4.6, the fact that the quasi-isomorphisms in question are induced by actual specialisation maps of homotopy types, the fact that a non-abelian version of Corollary 4.6 is also true, and the Hurewicz theorem in topology. We omit the details here as we don't need them.

We can now show the promised extension of Proposition 4.1 to families.

Proposition 4.8. *Let $f : X \to Y$ be a proper surjective morphism of finite type $\mathbf{F}_q$-schemes with geometrically irreducible fibres. Then $\mathcal{Z}(f,t) - \mathcal{Z}(Y,t)$ is a polynomial with integer coefficients.*

Proof. By Corollary 4.6, the Betti numbers of the fibres of f are bounded as we vary the base point. By Proposition 4.1, there exists an integer n_0 depending on the finite list of Betti numbers occuring in the fibres of f such that for all $n \geq n_0$ and any rational point $s \in S(\mathbf{F}_{q^n})$, the fibre X_s has an $\mathbf{F}_{q^n}$-rational point. It follows that $\mathcal{Z}(S,t) - \mathcal{Z}(f,t)$ is an integral polynomial of degree at most n_0. □

4.3. **The proof of Theorem 1.2.** In this section, we complete the proof of Theorem 1.2. In order to do so, we first explain how one can always reduce to the case of normal schemes.

Proposition 4.9. *Let X be a noetherian excellent scheme. Then there exists a normal scheme X' and a finite surjective morphism $\pi : X' \to X$ such that $X'(k) \to X(k)$ is surjective for any field k.*

Proof. We may assume X is reduced. Let $\emptyset = X_n \subset X_{n-1} \subset \cdots \subset X_1 \subset X_0 = X$ be the flag of reduced closed subschemes defined inductively by the following recipe: $X_0 = X$ and X_i is the non-normal locus of X_{i-1} given its reduced structure. We note that n is finite by the noetherianness of X. We set Y_i to be the normalisation in X_i of the open normal subscheme $U_i = X_i - X_{i-1} \hookrightarrow X_i$. The induced map $\pi_i : Y_i \to X_i$ is surjective and, by excellence, finite as well. Moreover, by construction, the map on rational points $Y_i(k) \to X_i(k)$ restricts to a surjective map on $U_i(k)$ for any field k. Setting $X' = \sqcup_i Y_i$ and $\pi = \sqcup \pi_i$ we obtain a finite surjective map $\pi : X' \to X$ which, at the level of k-rational points for a field k, induces a surjective map onto $U_i(k)$ for all i. Since $X = \cup_i U_i$, the claim follows. □

We now finish the proof.

Proof of Theorem 1.2. Let $f : X \to Y$ be a morphism of finite type $\mathbf{F}_q$-schemes. Assume first that f is proper. In this case, we will prove the theorem by induction on the dimension of Y. The case of dimension 0 that follows already includes all the essential steps of the argument.

Let's assume that $\dim(Y) = 0$. By Propositions 4.9 and 2.3, we may assume that X is normal. Let

$$X \xrightarrow{f'} Y' \xrightarrow{g} Y$$

be the Stein factorisation of f. The scheme Y' being reduced is the spectrum of a product of finite fields by finiteness of g. By the normality of X, the definition of the Stein factorisation, and the perfectness of finite fields, we see that all fibres of f' are geometrically irreducible. By Proposition 4.8, we conclude that $f'(\mathbf{F}_{q^n}) = Y'(\mathbf{F}_{q^n})$ for all n sufficiently large. Hence, $\mathcal{Z}(f, t) - \mathcal{Z}(g, t)$ is an integral polynomial. By Theorem 3.14, we know that $\mathcal{Z}(g, t)$ has geometric type, whence the claim follows.

For an arbitrary proper morphism f, we reduce as above to the case that X is normal. By induction and Propositions 2.2 and 2.4, we may shrink Y to assume that Y is normal and connected. Let

$$X \xrightarrow{f'} Y' \xrightarrow{g} Y$$

be the Stein factorisation of f. The scheme Y' is normal, and the map f' satisfies $f'_*\mathcal{O}_X \simeq \mathcal{O}_{Y'}$.

For each generic point $\eta' \in Y'$, the fibre $X_{\eta'}$ is easily seen to be normal and geometrically connected. Passing to the separable closure of the generic points preserves both normality and geometric connectedness of the fibres of f', while passing to inseparable extensions does not change the topology. Hence, each generic fibre of f' is geometrically irreducible. By [Gro66, Theorem 9.7.7], there exists a dense open $U' \subset Y'$ such that the fibres of f' over points of U' are geometrically irreducible. Shrinking U' a little if necessary and using that g is finite surjective, we may assume that $U' = g^{-1}(U)$ where $U \subset Y$ is an open dense subscheme. By induction on dimension and Proposition 2.4, we may replace Y with U and Y' with U' to assume that all fibres of f' are geometrically irreducible. Proposition 4.8 then gives $f'(\mathbf{F}_{q^n}) = Y'(\mathbf{F}_{q^n})$ for all n sufficiently large. We then conclude as in the case of dimension 0 above.

Now we move beyond the proper case. If f is separated but not proper, we can still reduce to the case that X is normal as above. Using Nagata compactification and excellence to find a normal compactification of f, and Stein factorising the resulting map to S allows us to find a commutative diagram

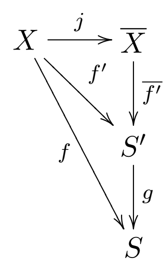

with $\overline{X}$ normal, j an open dense immersion, and the vertical maps on the right being the Stein factorisation of the proper map $g \circ \overline{f'}$. By the argument from the proper case above, we see that $\overline{f'}$ has geometrically irreducible generic fibres. By density and [Gro66, Theorem 9.7.7] as above, the same is true for f' over a dense open in S'. Now we can repeat the argument given above in the proper case.

In general, we first reduce to the case that S is affine (by Proposition 2.4), and then to the case that X is affine (by covering X by open affines and using Proposition 2.3), and then use use the separated case.

□

References

[AM69] M. Artin and B. Mazur. *Etale homotopy*. Lecture Notes in Mathematics, No. 100. Springer-Verlag, Berlin, 1969.

[Gro66] A. Grothendieck. Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. III. *Inst. Hautes Études Sci. Publ. Math.*, (28):255, 1966.