JOURNAL OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 29, Number 1, January 2016, Pages 81–94
S 0894-0347(2014)00820-1
Article electronically published on October 31, 2014

BERTINI IRREDUCIBILITY THEOREMS OVER FINITE FIELDS

FRANÇOIS CHARLES AND BJORN POONEN

1. Introduction

The classical Bertini theorems over an infinite field k state that if a subscheme $X \subseteq \mathbb{P}^n_k$ has a certain property (smooth, geometrically reduced, geometrically irreducible), then a sufficiently general hyperplane section over k has the property too. In [Poo04], an analogue of the Bertini smoothness theorem for a finite field $\mathbb{F}_q$ was proved, in which hyperplanes were replaced by hypersurfaces of degree tending to infinity.

The goal of the present article is to prove Bertini irreducibility theorems over finite fields. The proof of the Bertini irreducibility theorem over infinite fields k [Jou83, Théorème 6.3(4)] relies on the fact that a dense open subscheme of $\mathbb{P}^n_k$ has a k-point, so the proof fails over a finite field: see also the end of the introduction in [Ben11], where this problem is mentioned. The proof of the Bertini smoothness theorem over finite fields in [Poo04] depends crucially on the fact that smoothness can be checked analytically locally, one closed point at a time; in contrast, irreducibility and geometric irreducibility are not local properties. Therefore our proof must use ideas beyond those used in proving the earlier results. Indeed, our proof requires ingredients that are perhaps unexpected: resolution of singularities for surfaces, cones of curves in a surface, and the function field Chebotarev density theorem.

1.1. **Results for subschemes of projective space.**

Let $\mathbb{F}_q$ be a finite field of size q. Let $\mathbb{F}$ be an algebraic closure of $\mathbb{F}_q$. Let $S = \mathbb{F}_q[x_0, \ldots, x_n]$ be the homogeneous coordinate ring of $\mathbb{P}^n_{\mathbb{F}_q}$, let $S_d \subset S$ be the $\mathbb{F}_q$-subspace of homogeneous polynomials of degree d, and let $S_{\mathrm{homog}} = \bigcup_{d=0}^{\infty} S_d$. For each $f \in S_{\mathrm{homog}}$, let H_f be the subscheme $\operatorname{Proj}(S/(f)) \subseteq \mathbb{P}^n$, so H_f is a hypersurface (if f is not constant). Define the *density* of a subset $\mathcal{P} \subseteq S_{\mathrm{homog}}$ by

$$\mu(\mathcal{P}) := \lim_{d \to \infty} \frac{\#(\mathcal{P} \cap S_d)}{\# S_d},$$

if the limit exists. Define *upper* and *lower density* similarly, using lim sup or lim inf in place of lim.

Received by the editors November 27, 2013 and, in revised form, July 04, 2014, July 07, 2014, and September 15, 2014.

2010 *Mathematics Subject Classification.* Primary 14J70; Secondary 14N05.

Key words and phrases. Bertini irreducibility theorem, finite field.

This material is based upon work supported by the National Science Foundation under grant number DMS-1069236. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Theorem 1.1. *Let X be a geometrically irreducible subscheme of $\mathbb{P}^n_{\mathbb{F}_q}$. If $\dim X \geq 2$, then the density of*

$$\{f \in S_{\mathrm{homog}} : H_f \cap X \text{ is geometrically irreducible}\}$$

is 1.

We can generalize by requiring only that X be defined over $\mathbb{F}$; we still intersect with hypersurfaces over $\mathbb{F}_q$, however. Also, we can relax the condition of geometric irreducibility in both the hypothesis and the conclusion: the result is Theorem 1.2 below. To formulate it, we introduce a few definitions. Given a noetherian scheme X, let $\operatorname{Irr} X$ be its set of irreducible components. If $f \in S_{\mathrm{homog}}$ and X is a subscheme of $\mathbb{P}^n_L$ for some field extension $L \supseteq \mathbb{F}_q$, let X_f be the L-scheme $(H_f)_L \cap X$.

Theorem 1.2. *Let X be a subscheme of $\mathbb{P}^n_{\mathbb{F}}$ whose irreducible components are of dimension at least* 2*. For f in a set of density* 1*, there is a bijection $\operatorname{Irr} X \to \operatorname{Irr} X_f$ sending C to $C \cap X_f$.*

Remarks 1.3.

(a) For some f, the specification $C \mapsto C \cap X_f$ does not even define a map $\operatorname{Irr} X \to \operatorname{Irr} X_f$; i.e., $C \in \operatorname{Irr} X$ does not imply $C \cap X_f \in \operatorname{Irr} X_f$.
(b) If the specification does define a map, then the map is surjective.
(c) Given a geometrically irreducible subscheme $X \subseteq \mathbb{P}^n_{\mathbb{F}_q}$, Theorem 1.1 for X is equivalent to Theorem 1.2 for $X_{\mathbb{F}}$. Thus Theorem 1.2 is more general than Theorem 1.1.
(d) It may seem strange to intersect a scheme X over $\mathbb{F}$ with hypersurfaces defined over $\mathbb{F}_q$, but one advantage of Theorem 1.2 is that it implies the analogue of Theorem 1.1 in which "geometrically irreducible" is replaced by "irreducible" in both places. More generally, Theorem 1.2 implies an $\mathbb{F}_q$-analogue of itself, namely Corollary 1.4 below.

Corollary 1.4. *Let X be a subscheme of $\mathbb{P}^n_{\mathbb{F}_q}$ whose irreducible components are of dimension at least* 2*. For f in a set of density* 1*, there is a bijection $\operatorname{Irr} X \to \operatorname{Irr} X_f$ sending C to $C \cap X_f$.*

Proof. Apply Theorem 1.2 to $X_{\mathbb{F}}$, and identify $\operatorname{Irr} X$ with the set of $\operatorname{Gal}(\mathbb{F}/\mathbb{F}_q)$-orbits in $\operatorname{Irr} X_{\mathbb{F}}$, and likewise for X_f. ☐

1.2. **Results for morphisms to projective space.** In [Jou83] one finds a generalization (for an infinite field k) in which the subscheme $X \subseteq \mathbb{P}^n$ is replaced by a k-morphism $\phi \colon X \to \mathbb{P}^n$. Specifically, Théorème 6.3(4) of [Jou83] states that for a morphism of finite-type k-schemes $\phi \colon X \to \mathbb{P}^n$ with X geometrically irreducible and $\dim \overline{\phi(X)} \geq 2$, almost all hyperplanes $H \subseteq \mathbb{P}^n$ are such that $\phi^{-1} H$ is geometrically irreducible; here $\overline{\phi(X)}$ is the Zariski closure of $\phi(X)$ in $\mathbb{P}^n$, and "almost all" refers to a dense open subset of the moduli space of hyperplanes. The following example shows that we cannot expect such a generalization to hold for a density 1 set of hypersurfaces over a finite field.

Example 1.5. Let $n \geq 2$, and let $\phi \colon X \to \mathbb{P}^n_{\mathbb{F}_q}$ be the blowing up at a point $P \in \mathbb{P}^n(\mathbb{F}_q)$. The density of the set of f that vanish at P is $1/q$, and for any such nonzero f, the scheme $\phi^{-1} H_f$ is the union of the exceptional divisor $\phi^{-1} P$ and the strict transform of H_f, so it is not irreducible, and hence not geometrically irreducible.

We can salvage the result by disregarding irreducible components of $X_{\mathbb{F}}$ that are contracted to a point. To state the result, Theorem 1.6, we introduce the following terminology: given a morphism $\phi\colon X \to \mathbb{P}^n$, a subscheme Y of X (or $X_{\mathbb{F}}$) is *vertical* if $\dim \overline{\phi(Y)} = 0$, and *horizontal* otherwise. Let $\operatorname{Irr}_{\mathrm{horiz}} Y$ be the set of horizontal irreducible components of Y, and let Y_{horiz} be their union. Define $X_f := \phi^{-1} H_f$, viewed as a scheme over the same extension of $\mathbb{F}_q$ as X; this definition extends the earlier one.

Theorem 1.6. *Let X be a finite-type $\mathbb{F}$-scheme. Let $\phi\colon X \to \mathbb{P}^n_{\mathbb{F}}$ be an $\mathbb{F}$-morphism such that $\dim \overline{\phi(C)} \geq 2$ for each $C \in \operatorname{Irr} X$. For f in a set of density 1, there is a bijection $\operatorname{Irr} X \to \operatorname{Irr}_{\mathrm{horiz}} X_f$ sending C to $(C \cap X_f)_{\mathrm{horiz}}$.*

Alternatively, we can obtain a result for all irreducible components, but with only positive density instead of density 1.

Corollary 1.7. *Under the hypotheses of Theorem 1.6, if V is the set of closed points $P \in \mathbb{P}^n_{\mathbb{F}_q}$ such that $\phi^{-1}P$ is of codimension 1 in X, then the density of the set of f such that there is a bijection $\operatorname{Irr} X \to \operatorname{Irr} X_f$ sending C to $C \cap X_f$ is $\prod_{P \in V} \left(1 - q^{-\deg P}\right)$.*

Proof. By Lemma 3.1, we may assume that f does not vanish on any irreducible component of X. Then $\operatorname{Irr}_{\mathrm{horiz}} X_f = \operatorname{Irr} X_f$ if and only if f does not vanish at any element of V. □

1.3. **Strategy of proofs.** To prove a statement such as Theorem 1.1 for a variety X, we bound the number of geometrically reducible divisors of X cut out by hypersurfaces of large degree d. It is easier to count decompositions into effective Cartier divisors (as opposed to arbitrary subvarieties) since we can count sections of line bundles. So it would be convenient if X were smooth. If $\dim X = 2$, then a resolution of singularities $\widetilde{X} \to X$ is known to exist, and we can relate the problem for X to a counting problem on $\widetilde{X}$: see Proposition 4.1. If $\dim X > 2$, then resolution of singularities for X might not be known, and using an alteration seems to destroy the needed bounds, so instead we use induction on the dimension. The idea is to apply the inductive hypothesis to $J \cap X$ for some hypersurface J, but this requires $J \cap X$ itself to be geometrically irreducible, and finding such a J is dangerously close to the original problem we are trying to solve. Fortunately, finding *one* such J of unspecified degree is enough, and this turns out to be easier: see Lemma 5.3(a).

1.4. **Applications.** Two applications for our theorems existed even before the theorems were proved:

- Alexander Duncan and Zinovy Reichstein observed that Theorem 1.6 can be used to extend Theorem 1.4 of their article [DR14] to the case of an arbitrary ground field k; originally they proved it only over an infinite k. Their theorem compares variants of the notion of essential dimension for finite subgroups of $\mathrm{GL}_n(k)$. They need to use the Bertini theorems to construct hypersurface sections passing through a finite set of points [DR14, Theorem 8.1].
- Ivan Panin observed that Theorem 1.1 can be used to extend a result concerning the Grothendieck-Serre conjecture on principal G-bundles. The conjecture states that if R is a regular local ring, K is its fraction field, and G is a reductive group scheme over R, then $H^1_{\mathrm{et}}(R, G) \to H^1_{\mathrm{et}}(K, G)$

has trivial kernel. In [Pan14a, Pan14b, Pan14c], Panin proves such a statement for regular semilocal domains containing a field k; his proof relies on Theorem 1.1 when k is finite.

Here is another application, in the same spirit as [DR14, Theorem 8.1]. By *variety* we mean a separated finite-type scheme over a field.

Theorem 1.8. *Let X be a geometrically irreducible variety of dimension $m \geq 2$ over a field k. Let $F \subset X$ be a finite set of closed points. Then there exists a geometrically irreducible variety $Y \subseteq X$ of dimension $m-1$ containing F.*

Proof. By Nagata's embedding theorem (see, e.g., [Con07]), X embeds as a dense open subscheme of a proper k-scheme $\overline{X}$. If we find a suitable Y for $(\overline{X}, F)$, then $Y \cap X$ solves the problem for (X, F) (if necessary, we enlarge F to be nonempty to ensure that $Y \cap X$ is nonempty). So we may assume that X is proper.

Chow's lemma provides a projective variety X' and a birational morphism $\pi\colon X' \to X$. Enlarge F, if necessary, so that F contains a point in an open subscheme $U \subseteq X$ above which π is an isomorphism. Choose a finite set of closed points F' of X' such that $\pi(F') = F$. If Y' solves the problem for (X', F'), then $\pi(Y')$ solves the problem for (X, F) (we have $\dim \pi(Y') = \dim Y'$ since Y' meets $\pi^{-1}U$). Thus we may reduce to the case in which X is embedded in a projective space.

When k is infinite, one can use the classical Bertini irreducibility theorem as in [Mum70, p. 56] to complete the proof. So assume that k is finite.

We will let $Y := H_f \cap X$ for some f of high degree. For f in a set of positive density, H_f contains F, and Theorem 1.1 shows that for f outside a density 0 set, $H_f \cap X$ is geometrically irreducible (and of dimension $m-1$). As a consequence, we can find f such that H_f satisfies both conditions. $\square$

Corollary 1.9. *For X, m, and F as in Theorem 1.8, and for any integer y with $1 \leq y \leq m$, there exists a y-dimensional geometrically irreducible variety $Y \subseteq X$ containing F.*

Proof. Use Theorem 1.8 iteratively. $\square$

1.5. **An anti-Bertini theorem.** The following theorem uses a variant of the Bertini smoothness theorem to provide counterexamples to the Bertini irreducibility theorem over finite fields if we insist on *hyperplane* sections instead of higher-degree hypersurface sections. This discussion parallels [Poo04, Theorem 3.1].

Theorem 1.10 (Anti-Bertini theorem)**.** *Fix a finite field $\mathbb{F}_q$. For every sufficiently large positive integer d, there exists a geometrically irreducible degree d surface $X \subseteq \mathbb{P}^3_{\mathbb{F}_q}$ such that $H \cap X$ is reducible for every $\mathbb{F}_q$-plane $H \subseteq \mathbb{P}^3_{\mathbb{F}_q}$.*

Proof. Let $P = \mathbb{P}^3_{\mathbb{F}_q} - \mathbb{P}^3(\mathbb{F}_q)$. Let Z be the union of all the $\mathbb{F}_q$-lines in $\mathbb{P}^3$. By [Poo08, Theorem 1.1(i)] applied to P and Z, for any sufficiently large d, there exists a degree d surface $X \subseteq \mathbb{P}^3_{\mathbb{F}_q}$ containing Z such that $X \cap P$ is smooth of dimension 2. If $X_{\overline{\mathbb{F}}}$ were reducible, then $X_{\overline{\mathbb{F}}}$ would be singular along the intersection of two of its irreducible components; such an intersection would be of dimension at least 1, so this contradicts the smoothness of $X \cap P$. Thus X is geometrically irreducible.

On the other hand, for any $\mathbb{F}_q$-plane $H \subseteq \mathbb{P}^3_{\mathbb{F}_q}$, the 1-dimensional intersection $H \cap X$ contains all the $\mathbb{F}_q$-lines in H, so $H \cap X$ is reducible. $\square$

Remark 1.11. One could similarly construct, for any d_0, examples in which no hypersurface section of degree less than or equal to d_0 is irreducible: just include higher degree curves in Z. Also, one could give higher-dimensional examples, hypersurfaces X in $\mathbb{P}^n_{\mathbb{F}_q}$ for $n > 3$ containing all $(n-2)$-dimensional $\mathbb{F}_q$-subspaces L: a straightforward generalization of [Poo08, Theorem 1.1(i)] proves the existence of such X whose singular locus has codimension 2 in X.

2. Notation

If x is a point of a scheme X, let $\kappa(x)$ be its residue field. If X is an irreducible variety, let $\kappa(X)$ be its function field. If $L \supseteq k$ is an extension of fields, and X is a k-scheme, let X_L be the L-scheme $X \times_{\operatorname{Spec} k} \operatorname{Spec} L$. If X is a reduced subscheme of a projective space, let $\overline{X}$ be its Zariski closure. Given a finite-type scheme X over a field, let X_{red} be the associated reduced subscheme, let X^{smooth} be the smooth locus of X, and let X^{sing} be the closed subset $X \setminus X^{\mathrm{smooth}}$.

3. Lemmas

Lemma 3.1. *Let X be a positive-dimensional subscheme of $\mathbb{P}^n$ over $\mathbb{F}_q$ or $\mathbb{F}$. The density of $\{f \in S_{\mathrm{homog}} : f \text{ vanishes on } X\}$ is 0.*

Proof. We may assume that X is over $\mathbb{F}_q$; if instead X is over $\mathbb{F}$, replace X by its image under $\mathbb{P}^n_{\mathbb{F}} \to \mathbb{P}^n_{\mathbb{F}_q}$. If x is a closed point of X, the upper density of the set of f vanishing on X is bounded by the density of the set of f vanishing at x, which is $1/\#\kappa(x)$. Since X is positive-dimensional, we may choose x of arbitrarily large degree. □

Lemma 3.2. *Let X be a positive-dimensional subscheme of $\mathbb{P}^n$ over $\mathbb{F}_q$ or $\mathbb{F}$. The density of $\{f \in S_{\mathrm{homog}} : H_f \cap X \neq \emptyset\}$ is 1.*

Proof. Again we may assume that X is over $\mathbb{F}_q$. Given $r \in \mathbb{R}$, let $X_{<r}$ be the set of closed points of X of degree $< r$. The density of $\{f \in S_{\mathrm{homog}} : H_f \cap X_{<r} = \emptyset\}$ equals the finite product

$$\prod_{P \in X_{<r}} \left(1 - q^{-\deg P}\right) = \zeta_{X_{<r}}(1)^{-1},$$

which diverges to 0 as $r \to \infty$, since $\zeta_X(s)$ converges only for $\operatorname{Re}(s) > \dim X$. □

Lemma 3.3. *Let X and ϕ be as in Theorem 1.6. Let $U \subseteq X$ be a dense open subscheme. For f in a set of density 1, there is a bijection $\operatorname{Irr}_{\mathrm{horiz}} X_f \to \operatorname{Irr}_{\mathrm{horiz}} U_f$ sending D to $D \cap U$.*

Proof. Lemma 3.1 shows that the set of f vanishing on $\overline{\phi(D)}$ for some $D \in \operatorname{Irr}_{\mathrm{horiz}}(X \setminus U)$ has density 0. After excluding such f, every $D \in \operatorname{Irr}_{\mathrm{horiz}} X_f$ meets U (if not, then $D \in \operatorname{Irr}_{\mathrm{horiz}}(X \setminus U)$ and $f(\overline{\phi(D)}) = 0$). Then the sets $\operatorname{Irr}_{\mathrm{horiz}} X_f$ and $\operatorname{Irr}_{\mathrm{horiz}} U_f$ are in bijection: the forward map sends D to $D \cap U$, and the backward map sends D to its closure in X_f. □

Lemma 3.4. *Let X be a smooth finite-type $\mathbb{F}$-scheme with a morphism $\phi \colon X \to \mathbb{P}^n_{\mathbb{F}}$ such that $\dim \overline{\phi(C)} \geq 2$ for every $C \in \operatorname{Irr} X$. Let $f \in S_{\mathrm{homog}} \setminus \{0\}$. The following are equivalent:*

(a) *There is a bijection $\operatorname{Irr} X \to \operatorname{Irr}_{\mathrm{horiz}} X_f$ sending C to $(C_f)_{\mathrm{horiz}} = (C \cap X_f)_{\mathrm{horiz}}$.*
(b) *For every $C \in \operatorname{Irr} X$, the scheme $(C_f)_{\mathrm{horiz}}$ is irreducible.*

Proof.

(a)⇒(b): If the map is defined, then $(X'_f)_{\mathrm{horiz}}$ is irreducible by definition.

(b)⇒(a): The assumption (b) implies that the map in (a) is defined. It is surjective since any irreducible component of X_f is contained in an irreducible component of X. Smoothness implies that the irreducible components of X are disjoint, so the map is injective. □

Lemma 3.5. *Let X be a subscheme of $\mathbb{P}^n_{\mathbb{F}}$ that is smooth of dimension m. For f in a set of density* 1*, the singular locus $(X_f)^{\mathrm{sing}}$ is finite.*

Proof. We follow the proof of [Poo04, Lemma 2.6 and Theorem 3.2]. The formation of $(X_f)^{\mathrm{sing}}$ is local, so we may replace $\mathbb{P}^n_{\mathbb{F}}$ by $\mathbb{A}^n_{\mathbb{F}}$, and replace each f by the corresponding dehomogenization in the $\mathbb{F}_q$-algebra $A := \mathbb{F}_q[x_1, \ldots, x_n]$. For $d \geq 1$, let $A_{\leq d}$ be the set of $f \in A$ of total degree at most d. Let $\mathscr{T}_X$ be the tangent sheaf of X; identify its sections with derivations.

Let $X_1, \ldots, X_n$ be the distinct $\mathrm{Gal}(\mathbb{F}/\mathbb{F}_q)$-conjugates of X in $\mathbb{A}^n_{\mathbb{F}}$. For each nonnegative integer k, let B_k be a set of points in $\mathbb{A}^n(\mathbb{F})$ contained in exactly k of the X_i. For $x \in B_k$, let V_x be the span of the tangent spaces $T_x X_i$ in $T_x \mathbb{A}^n_{\mathbb{F}}$ for all X_i containing x. For $r \geq m$, define $C_{k,r} := \{x \in B_k : \dim V_x = r\}$ (it would be empty for $r < m$). Each B_k and each $C_{k,r}$ is the set of $\mathbb{F}$-points of a locally closed subscheme of $\mathbb{A}^n_{\mathbb{F}_q}$; from now on, B_k and $C_{k,r}$ refer to these subschemes. There is a rank r subbundle $\mathscr{V}$ of $\mathscr{T}_{\mathbb{A}^n_{\mathbb{F}_q}}|_{C_{k,r}}$ whose fiber at $x \in C_{k,r}(\mathbb{F})$ is V_x. If $x \in (X_f)^{\mathrm{sing}}(\mathbb{F}) \cap C_{k,r}(\mathbb{F})$, and y is its image in $\mathbb{A}^n_{\mathbb{F}_q}$, then $T_y H_f \supseteq \mathscr{V} \otimes \kappa(y)$ as subspaces of $T_y \mathbb{A}^n_{\mathbb{F}_q}$.

Suppose that y is a closed point of the $\mathbb{F}_q$-scheme $C_{k,r}$. Choose global derivations $D_1, \ldots, D_r \colon A \to A$ whose images in $\mathscr{T}_{\mathbb{A}^n_{\mathbb{F}_q}}|_{C_{k,r}}$ form a basis for $\mathscr{V}$ on some neighborhood U of y in $C_{k,r}$. From now on, we use only $D_1, \ldots, D_m$. Choose $t_1, \ldots, t_m \in A$ such that $D_i(t_j)$ is nonzero at x if and only if $i = j$. After shrinking U, we may assume that the values $D_i(t_i)$ are invertible on U. Let $U_X := U_{\mathbb{F}} \cap X$. If $P \in U_X \cap (X_f)^{\mathrm{sing}}$, then $f(P) = (D_1 f)(P) = \cdots = (D_m f)(P) = 0$.

Let $\tau = \max_i(\deg t_i)$ and $\gamma = \lfloor (d - \tau)/p \rfloor$. Select $f_0 \in A_{\leq d}$, $g_1 \in A_{\leq \gamma}$, $\ldots$, $g_m \in A_{\leq \gamma}$ uniformly and independently at random, and define

$$f := f_0 + g_1^p t_1 + \cdots + g_m^p t_m.$$

Then the distribution of f is uniform over $A_{\leq d}$, and $D_i f = D_i f_0 + g_i^p (D_i t_i)$ for each i. For $0 \leq i \leq m$, define the $\mathbb{F}$-scheme

$$W_i := U_X \cap \{D_1 f = \cdots = D_i f = 0\}.$$

Thus $U_X \cap (X_f)^{\mathrm{sing}} \subseteq W_m$.

In the remainder of this proof, the big-O and little-o notations indicate the behavior as $d \to \infty$, and the implied constants may depend on n, X, U, and the D_i, but not on f_0 or the g_i (and of course not on d). We claim that for $0 \leq i \leq m-1$, conditioned on a choice of f_0, g_1, $\ldots$, g_i for which $\dim W_i \leq m - i$, the probability that $\dim W_{i+1} \leq m-i-1$ is $1 - o(1)$ as $d \to \infty$. First, let $Z_1, \ldots, Z_\ell$ be the $(m-i)$-dimensional irreducible components of $(W_i)_{\mathrm{red}}$. As in the proof of [Poo04, Lemma 2.6], we have $\ell = O(d^i)$ by Bézout's theorem, and the probability that $D_{i+1} f$ vanishes on a given Z_j is at most $q^{-\gamma-1}$, so the probability that the inequality $\dim W_{i+1} \leq m - i - 1$ *fails* is at most $\ell q^{-\gamma-1} = o(1)$, as claimed.

By induction on i, the previous paragraph proves that $\dim W_i \le m - i$ with probability $1-o(1)$ as $d \to \infty$, for each i. In particular, W_m is finite with probability $1-o(1)$. Thus $U_X \cap (X_f)^{\mathrm{sing}}$ is finite with probability $1-o(1)$. Finally, X is covered by U_X for finitely many U (contained in different $C_{k,r}$), so $(X_f)^{\mathrm{sing}}$ is finite with probability $1-o(1)$. □

4. Surfaces over a finite field

Proposition 4.1. *Let X be a 2-dimensional closed integral subscheme of $\mathbb{P}^n_{\mathbb{F}_q}$. For f in a set of density 1, there is a bijection $\operatorname{Irr} X_{\mathbb{F}} \to \operatorname{Irr}(X_f)_{\mathbb{F}}$ sending C to $C \cap X_f$.*

Before beginning the proof of Proposition 4.1, we prove a lemma.

Lemma 4.2. *Let Y be a smooth projective irreducible surface over a field k. Let B be a big and nef line bundle on Y (see, e.g.,* [Laz04, Definitions 2.2.1 and 1.4.1]*). Then for every line bundle L on Y,*

$$h^0(Y, L) \le \frac{(L.B)^2}{2B.B} + O(L.B) + O(1),$$

where the implied constants depend on Y and B, but not L.

Proof. If C is an effective curve on Y, then $C.B \ge 0$ since B is nef. In particular, if L has a nonzero section, then $L.B \ge 0$. Thus if $L.B < 0$, then $h^0(Y, L) = 0$.

Now suppose that $L.B \ge 0$. Since B is big, we may replace B by a power to assume that $B = \mathscr{O}(D)$ for some effective divisor D. Taking global sections in

$$0 \to L(-D) \to L \to L|_D \to 0$$

yields

$$h^0(Y, L) \le h^0(Y, L(-D)) + h^0(D, L|_D) = h^0(Y, L(-D)) + L.B + O(1), \tag{1}$$

while $L(-D).B = L.B - B.B$. Combine (1) for L, $L(-D)$, $L(-2D)$, $\ldots$, until reaching $L(-mD)$ such that $L(-mD).B < 0$; then $h^0(Y, L(-mD)) = 0$ by the first paragraph. We have $m \le \lfloor L.B/B.B \rfloor + 1$, and the result follows by summing an arithmetic series. □

Proof of Proposition 4.1. Let $\pi \colon \widetilde{X} \to X$ be a resolution of singularities of X. Let B be a divisor with $\mathscr{O}(B) \simeq \pi^* \mathscr{O}_X(1)$. Then B is big and nef, so there exists $b \in \mathbb{Z}_{>0}$ such that bB is linearly equivalent to $A + E$ with A ample and E effective. Consider $f \in S_d$. By Lemma 3.1, we may discard the f vanishing on any one curve in X. In particular, given a positive constant d_0, we may assume that H_f does not contain any 1-dimensional irreducible component of $\pi(E)$ or any of the finitely many curves C on X with $C.\mathscr{O}_X(1) < d_0$. (In fact, we could have chosen A and E so that $\pi(E)$ is finite.)

The linear map

$$H^0(\mathbb{P}^n_{\mathbb{F}_q}, \mathscr{O}(d)) \to H^0(X, \mathscr{O}_X(d))$$

is surjective for large d, and $h^0(X, \mathscr{O}_X(d)) \to \infty$. Thus densities can be computed by counting divisors X_f (corresponding to elements of $\mathbb{P}H^0(X, \mathscr{O}_X(d))$) instead of polynomials f.

If X_f is reducible, then the Cartier divisor $\pi^* X_f$ may be written as $D + D'$ where D and D' are effective divisors such that $D.B, D'.B \ge d_0$ and D is irreducible (and hence horizontal relative to π since $D.B > 0$). Let L be the line bundle $\mathscr{O}(D)$

on $\widetilde{X}$. Since $\mathscr{O}(\pi^* X_f) \simeq \mathscr{O}(dB)$, we have $\mathscr{O}(D') \simeq \mathscr{O}(dB) \otimes L^{-1}$. Let $n := L.B$. Then $d_0 \le n \le dB.B - d_0$.

We will bound the number of reducible X_f by bounding the number of possible L for each n, and then bounding the number of pairs (D, D') for a fixed L. The assumption on H_f implies that D and E have no common components, so $L.E = D.E \ge 0$. Thus $L.A \le L.(A+E) = L.(bB) = nb$. The numerical classes of effective L are lattice points in the closed cone $\overline{\mathrm{NE}}(\widetilde{X}) \subseteq (N_1 \widetilde{X})_{\mathbb{R}}$, on which A is positive (except at 0), so the number of such L up to numerical equivalence satisfying $L.A \le nb$ is $O(n^\rho)$ as $n \to \infty$ for some ρ. The number of L within each numerical equivalence is at most $\# \mathbf{Pic}^\tau_{\widetilde{X}}(\mathbb{F}_q)$, where $\mathbf{Pic}^\tau_{\widetilde{X}}$ is the finite-type subgroup scheme of the Picard scheme parametrizing line bundles with torsion Néron-Severi class. Thus the number of possible L is $O(n^\rho)$.

Now fix L. Recall that $n = L.B$. By Lemma 4.2 with $Y := \widetilde{X}$,

$$
\begin{aligned}
&h^0(\widetilde{X}, L) + h^0(\widetilde{X}, \mathscr{O}(dB) \otimes L^{-1}) \\
&\le \frac{(L.B)^2}{2B.B} + O(L.B) + O(1) + \frac{(dB.B - L.B)^2}{2B.B} + O(dB.B - L.B) + O(1) \\
&\le \frac{B.B}{2} d^2 - \frac{n(dB.B - n)}{B.B} + O(d).
\end{aligned}
$$

Summing over all $n \in [d_0, dB.B - d_0]$ and all L shows that the number of pairs (D, D') is at most

$$
\begin{aligned}
\sum_{n=d_0}^{dB.B-d_0} O(n^\rho) q^{h^0(\widetilde{X},L)+h^0(\widetilde{X},\mathscr{O}(dB)\otimes L^{-1})} &\le O(d^\rho) \sum_{n=d_0}^{dB.B-d_0} q^{\frac{B.B}{2}d^2 - \frac{n(dB.B-n)}{B.B} + O(d)} \\
&\le d^\rho q^{\frac{B.B}{2}d^2 + O(d)}\, 2 \sum_{n=d_0}^{dB.B/2} q^{-\frac{n(dB.B-n)}{B.B}} \\
&\le q^{\frac{B.B}{2}d^2 + O(d)} \sum_{n=d_0}^{dB.B/2} q^{-nd/2} \\
&\le q^{\frac{B.B}{2}d^2 - \frac{d_0 d}{2} + O(d)}.
\end{aligned}
$$

The number of reducible X_f is at most this. On the other hand, the total number of X_f is

$$
\#\mathbb{P}H^0(X, \mathscr{O}_X(d)) = q^{\frac{B.B}{2}d^2 + O(d)} \tag{2}
$$

since $\deg X = B.B$. Dividing yields a proportion that tends to 0 as $d \to \infty$, provided that d_0 was chosen large enough.

Finally we must bound the number of irreducible X_f such that the conclusion of Proposition 4.1 fails; i.e., there is not a bijection $\operatorname{Irr} X_{\mathbb{F}} \to \operatorname{Irr}(X_f)_{\mathbb{F}}$ sending C to $C \cap X_f$. Consider such an f. Let $Y \in \operatorname{Irr} \widetilde{X}_{\mathbb{F}}$. Let $\mathbb{F}_r$ be the field of definition of Y. Redefine Y as an element of $\operatorname{Irr} \widetilde{X}_{\mathbb{F}_r}$. If we view the $\mathbb{F}_r$-scheme Y as an $\mathbb{F}_q$-scheme (by composing with $\operatorname{Spec} \mathbb{F}_r \to \operatorname{Spec} \mathbb{F}_q$), then the morphisms $Y \to \widetilde{X} \to X$ are birational $\mathbb{F}_q$-morphisms. Thus $Y \times_{\mathbb{F}_q} \mathbb{F}$ and $X_{\mathbb{F}}$ share a common smooth dense open subscheme. Applying Lemma 3.3 twice lets us deduce that, excluding f in a set of density 0, there still is not a bijection $\operatorname{Irr}(Y \times_{\mathbb{F}_q} \mathbb{F}) \to \operatorname{Irr}_{\mathrm{horiz}}(Y_f \times_{\mathbb{F}_q} \mathbb{F})$ sending C to $(C_f)_{\mathrm{horiz}} = (C \cap Y_f)_{\mathrm{horiz}}$. Then, by Lemma 3.4 applied to $Y \times_{\mathbb{F}_q} \mathbb{F}$, there exists $C \in \operatorname{Irr}(Y \times_{\mathbb{F}_q} \mathbb{F})$ such that $(C_f)_{\mathrm{horiz}}$ is reducible. But C is the base change of Y by *some* $\mathbb{F}_q$-homomorphism $\mathbb{F}_r \to \mathbb{F}$, so the $\mathbb{F}_r$-scheme $(Y_f)_{\text{horiz}}$ is not geometrically irreducible. On the other hand, by Lemma 3.2, for f in a set of density 1, the scheme X_f meets X^{smooth}, in which case $(Y_f)_{\text{horiz}}$ viewed as $\mathbb{F}_q$-scheme is birational to X_f, so $(Y_f)_{\text{horiz}}$ is irreducible. Thus we have a map from the set $\mathscr{X}$ of irreducible X_f such that the conclusion of Proposition 4.1 fails (excluding the density 0 sets above) to the set $\mathscr{Y}$ of irreducible and geometrically reducible schemes of the form Z_{horiz} for $Z \in \mathbb{P}H^0(Y, \mathscr{O}(d))$, namely the map sending X_f to $(Y_f)_{\text{horiz}}$. This map is injective since X_f is determined as a Cartier divisor by any of its dense open subschemes, and X_f and $(Y_f)_{\text{horiz}}$ share a subscheme that is dense and open in both. Therefore it suffices to bound $\#\mathscr{Y}$.

For $e \geq 2$, let $\mathscr{Y}_e$ be the set of $Z_{\text{horiz}} \in \mathscr{Y}$ such that there exists an effective Cartier divisor D on $Y_{\mathbb{F}_{r^e}}$ such that

$$(Z_{\mathbb{F}_{r^e}})_{\text{horiz}} = \sum_{\sigma \in \operatorname{Gal}(\mathbb{F}_{r^e}/\mathbb{F}_r)} {}^{\sigma}D.$$

Then $\mathscr{Y} = \bigcup_{e \geq 2} \mathscr{Y}_e$.

Let L be the line bundle $\mathscr{O}(D)$ on $Y_{\mathbb{F}_{r^e}}$. Let $B_Y = B|_Y$. Then

$$B.B = [\mathbb{F}_r : \mathbb{F}_q] B_Y.B_Y$$

since the local self-intersection numbers of B have the same sum on each conjugate of Y. Let $n := L.B_Y$. Then $ne = Z_{\text{horiz}}.B_Y = Z.B_Y = dB_Y.B_Y$, so $n = (d/e)B_Y.B_Y$. The number of numerical classes of effective L with $L.B_Y = n$ is $O(n^\rho)$ as before, uniformly in e (because they are so bounded even over $\mathbb{F}$). We have $\#\,\mathbf{Pic}^{\tau}_Y(\mathbb{F}_{r^e}) = (r^e)^{O(1)}$. Thus the number of possible L is $O(n^\rho)(r^e)^{O(1)}$. Applying Lemma 4.2 to $Y_{\mathbb{F}}$ shows that

$$\begin{aligned}
h^0(Y_{\mathbb{F}_{r^e}}, L) &= h^0(Y_{\mathbb{F}}, L) \\
&\leq \frac{(L.B_Y)^2}{2B_Y.B_Y} + O(L.B_Y) + O(1) \\
&\leq \frac{(d/e)^2(B_Y.B_Y)^2}{2B_Y.B_Y} + O(d/e) \\
&\leq \frac{(d/e)^2 B_Y.B_Y}{2} + O(d/e) \\
&\leq \frac{(d/e)^2 B.B}{2[\mathbb{F}_r : \mathbb{F}_q]} + O(d/e),
\end{aligned}$$

so

$$\#H^0(Y_{\mathbb{F}_{r^e}}, L) = (r^e)^{h^0(Y_{\mathbb{F}_{r^e}}, L)} \leq q^{d^2 B.B/2e + O(d)}$$

since $r = O(1)$, and

$$\#\mathscr{Y}_e \leq O(n^\rho)(r^e)^{O(1)} q^{d^2 B.B/2e + O(d)} \leq q^{d^2 B.B/2e + O(d)}$$

since n and e are $O(d)$, so

$$\#\mathscr{Y} = \sum_{e=2}^{O(d)} \#\mathscr{Y}_e \leq O(d) q^{d^2 B.B/4 + O(d)}.$$

This divided by the quantity (2) tends to 0 as $d \to \infty$. $\square$

5. Reductions

Lemma 5.1. *Let X and Y be irreducible finite-type $\mathbb{F}$-schemes, with morphisms $X \xrightarrow{\pi} Y \xrightarrow{\psi} \mathbb{P}^n_{\mathbb{F}}$ such that π is finite étale, ψ has relative dimension s at each point of Y, and $\dim \overline{\psi(Y)} \geq 2$. For f in a set of density 1, the implication*

$$Y_f \text{ irreducible} \implies X_f \text{ irreducible}$$

holds.

(The reverse implication holds for *all* f. Later we will prove that both sides hold for f in a set of density 1.)

Proof of Lemma 5.1. Irreducibility of X_f becomes harder to achieve only if X is replaced by a higher finite étale cover of Y. In particular, we may replace X by a cover corresponding to a Galois closure of $\kappa(X)/\kappa(Y)$. So assume from now on that $X \to Y$ is Galois étale, say with Galois group G.

Choose a finite extension $\mathbb{F}_r$ of $\mathbb{F}_q$ with a morphism $\psi' \colon Y' \to \mathbb{P}^n_{\mathbb{F}_r}$, and a Galois étale cover $\pi' \colon X' \to Y'$ whose base extensions to $\mathbb{F}$ yield ψ and π. Let $m := \dim \overline{\psi(Y)}$. Then $\dim Y' = \dim Y = s + m$.

Given a closed point $y \in Y'$, let Frob_y be the associated Frobenius conjugacy class in G. We will prove that the following claims hold for f in a set of density 1:

Claim 1. The Frob_y for $y \in Y'_f$ cover all conjugacy classes of G.

Claim 2. The scheme X'_f contains two closed points whose degrees over $\mathbb{F}_r$ are coprime.

Let C be a conjugacy class in G. Let $c := \#C/\#G$. In the arguments below, for fixed X', Y', ψ', π', G, and C, the expression $o(1)$ denotes a function of e that tends to 0 as $e \to \infty$. By a function field analogue of the Chebotarev density theorem [Lan56, last display on p. 393] (which, in this setting, follows from applying the Lang-Weil estimates to all twists of the cover $Y' \to X'$), the number of closed points $y \in Y'$ with residue field $\mathbb{F}_{r^e}$ satisfying $\operatorname{Frob}_y = C$ is $(c + o(1))\, r^{(s+m)e}/e$. Since ψ has fibers of dimension s, there exists $c' > 0$ such that the images of these points in $\mathbb{P}^n_{\mathbb{F}_q}$ are at least $(c' + o(1))\, r^{me}/e$ closed points of $\overline{\psi(Y)}$ with residue field of size at most r^e, so the density of the set of f such that H_f misses all these points is at most $(1 - r^{-e})^{(c'+o(1))r^{me}/e}$, which tends to 0 as $e \to \infty$ since $m \geq 2$.

Proof of Claim 1: The previous sentence shows that for f in a set of density 1, there exists a closed point $y' \in Y'_f$ with $\operatorname{Frob}_y = C$; apply this to every C.

Proof of Claim 2: The same sentence shows that the lower density of the set of f such that there exists $y \in Y'_f$ with $\kappa(y) = \mathbb{F}_{r^e}$ and $\operatorname{Frob}_y = \{1\}$ tends to 1 as $e \to \infty$. For any such y, any preimage $x \in X'_f$ satisfies $\kappa(x) = \mathbb{F}_{r^e}$. Apply this to two coprime integers e and e', and then let $\min(e, e')$ tend to ∞.

To complete the proof of the lemma, we show that if Y_f is irreducible and Claims 1 and 2 hold, then X_f is irreducible. Assume that Y_f is irreducible, so Y'_f is geometrically irreducible. The only subgroup of G that meets all conjugacy classes is G itself, so Claim 1 implies that $X'_f \to Y'_f$ is a finite Galois *irreducible* cover (with Galois group G). If such a cover X'_f is not geometrically irreducible, then there is an integer $a > 1$ dividing the degrees (over $\mathbb{F}_r$) of all its closed points. Thus Claim 2 implies that X'_f is geometrically irreducible, so X_f is irreducible. $\square$

We now strengthen Lemma 5.1 to replace "finite étale" by "dominant" and to remove the relative dimension hypothesis.

Lemma 5.2. *Let X and Y be irreducible finite-type $\mathbb{F}$-schemes, with morphisms $X \xrightarrow{\pi} Y \xrightarrow{\psi} \mathbb{P}^n_{\mathbb{F}}$ such that π is dominant and $\dim \overline{\psi(Y)} \geq 2$. For f in a set of density 1, the implication*

$$(Y_f)_{\mathrm{horiz}} \text{ irreducible } \implies (X_f)_{\mathrm{horiz}} \text{ irreducible}$$

holds.

Proof. By Lemma 3.3, we may replace X and Y by dense open subschemes; thus we may assume that ψ has constant relative dimension and π factors as

$$X \xrightarrow{\text{surjective radicial}} V \xrightarrow{\text{finite étale}} W \xrightarrow{\text{dense open immersion}} \mathbb{A}^r_Y \longrightarrow Y$$

for some $r \in \mathbb{Z}_{\geq 0}$, because $\kappa(X)$ is a finite inseparable extension of a finite separable extension of a purely transcendental extension of $\kappa(Y)$, and we may spread this out. Then $(Y_f)_{\mathrm{horiz}} = Y_f$ and similarly for $\mathbb{A}^r_Y$, W, V, and X. Irreducibility of Y_f is equivalent to irreducibility of $(\mathbb{A}^r_Y)_f$; for f in a set of density 1, this is equivalent to irreducibility of W_f (Lemma 3.3), which implies irreducibility of V_f (Lemma 5.1 applied to $V \to W$), which is equivalent to irreducibility of X_f (homeomorphism). □

Part (a) of the following lemma and its proof are closely related to results of Lior Bary-Soroker [BS13]; see also [Neu98, Lemma 2.1], attributed to Wulf–Dieter Geyer.

Lemma 5.3. *Let X be a smooth irreducible subscheme of $\mathbb{P}^n_{\mathbb{F}}$ of dimension $m \geq 3$.*

(a) *There exists a hypersurface $J \subseteq \mathbb{P}^n_{\mathbb{F}_q}$ such that $\dim J \cap (\overline{X} \setminus X) \leq m-2$ and $J \cap X$ is irreducible of dimension $m-1$.*

(b) *For any such J, there exists a density 1 set of f for which the implication*

$$(J \cap X)_f \text{ irreducible } \implies X_f \text{ irreducible}$$

holds.

Proof.

(a) Inductively choose $h_0, \ldots, h_m \in S_{\mathrm{homog}}$ so that the common zero locus of $h_0, \ldots, h_r$ on $\overline{X}$ is of the expected dimension $m-r-1$ for $r \in \{0, 1, \ldots, m-1\}$ and empty for $r = m$. Replace every h_i by a power to assume that they have the same degree. Then $(h_0 : \cdots : h_m) \colon \mathbb{P}^n_{\mathbb{F}} \dashrightarrow \mathbb{P}^m_{\mathbb{F}}$ restricts to a morphism $\pi \colon \overline{X} \to \mathbb{P}^m_{\mathbb{F}}$ whose fiber above $(0 : \cdots : 0 : 1)$ is 0-dimensional, so π is generically finite. Since $\dim \overline{X} = m$, the morphism π is dominant.

Let Z be the set of points of $\mathbb{P}^m_{\mathbb{F}_q}$ above which the fiber has codimension 1 in $\overline{X}$. Since $m > 1$ and π is dominant, Z is finite. Let $B_1, \ldots, B_s$ be the images in $\mathbb{P}^m_{\mathbb{F}_q}$ of the irreducible components of $\overline{X} \setminus X$. Let Z' be the union of Z with all the 0-dimensional B_i. The density of homogeneous polynomials g on $\mathbb{P}^m_{\mathbb{F}_q}$ such that $H_g \cap Z' = \emptyset$ is $\prod_{z \in Z'}(1 - 1/\#\kappa(z)) > 0$. The density of such g such that also H_g is geometrically integral and does not contain any positive-dimensional B_i is the same by [Poo04, Proposition 2.7] and Lemma 3.1. For such g, the subscheme $X_g := \pi^{-1} H_g$ is horizontal and contains no irreducible component of $\overline{X} \setminus X$. Lemma 5.2 applied to $X \xrightarrow{\pi} \mathbb{P}^m_{\mathbb{F}} \xrightarrow{\mathrm{Id}} \mathbb{P}^m_{\mathbb{F}}$ shows that, after excluding a further density 0 set, X_g is irreducible of dimension $m-1$. Let J be the hypersurface in $\mathbb{P}^n_{\mathbb{F}_q}$ defined by $g(h_0,\dots,h_m)$. Then J contains no irreducible component of $\overline{X}\setminus X$, so $\dim J\cap(\overline{X}\setminus X)\le \dim(\overline{X}\setminus X)-1\le m-2$. Also, $J\cap X = X_g$, which is irreducible of dimension $m-1$.

(b) Consider f such that f does not vanish on any positive-dimensional irreducible component of $J\cap(\overline{X}\setminus X)$ and $(X_f)^{\mathrm{sing}}$ is finite. By Lemmas 3.1 and 3.5, this set has density 1.

Suppose that $(J\cap X)_f$ is irreducible but X_f is reducible, say $X_f = V_1\cup V_2$, where V_i are closed subsets of X_f, neither containing the other. Then $\dim V_i \ge m-1$ for each i. Since J is a hypersurface, $J\cap\overline{V}_i$ is nonempty and of dimension at least $m-2$. On the other hand, $J\cap(\overline{V}_i\setminus V_i)\subseteq J\cap(\overline{X}\setminus X)_f$, which is of dimension at most $m-3$. Thus $J\cap V_i$ is nonempty of dimension at least $m-2\ge 1$. Since $(J\cap X)_f = (J\cap V_1)\cup(J\cap V_2)$, one of the $J\cap V_i$ must contain the other, say $(J\cap V_2)\subseteq (J\cap V_1)$. Then $J\cap V_2\subseteq V_1\cap V_2\subseteq (X_f)^{\mathrm{sing}}$. This is a contradiction since $\dim J\cap V_2\ge 1$ and $(X_f)^{\mathrm{sing}}$ is finite. □

Proposition 5.4. *Let X be an irreducible subscheme of $\mathbb{P}^n_{\mathbb{F}_q}$ of dimension at least 2. For f in a set of density 1, there is a bijection $\operatorname{Irr} X_{\mathbb{F}}\to\operatorname{Irr}(X_f)_{\mathbb{F}}$ sending C to C_f.*

Proof. We use induction on $\dim X$. Replace X by X_{red}; then X is integral. The case $\dim X = 2$ follows from Proposition 4.1 for $\overline{X}$, by using Lemma 3.3 to pass from X to $\overline{X}$.

Now suppose $\dim X > 2$. Because of Lemma 3.3, we may shrink X to assume that X is smooth. Choose one $C\in\operatorname{Irr} X_{\mathbb{F}}$. Choose J as in Lemma 5.3 applied to C, so $J\cap C$ is irreducible. Then so is its image $J\cap X$ under $C\to X$. Also, $\dim(J\cap X)=\dim X-1\ge 2$.

We have $J\cap C\in\operatorname{Irr}((J\cap X)_{\mathbb{F}})$. The inductive hypothesis applied to $J\cap X$ shows that for f in a set of density 1, the scheme $(J\cap C)_f$ is irreducible. By Lemma 5.3(b), for f in a smaller density 1 set, this implies that C_f is irreducible. Each $C'\in\operatorname{Irr} X_{\mathbb{F}}$ is conjugate to C, so C'_f is irreducible for the same f. For these f, Lemma 3.4 implies the conclusion. □

Proposition 5.5. *Let X be a finite-type $\mathbb{F}_q$-scheme. Let $\phi\colon X\to\mathbb{P}^n_{\mathbb{F}_q}$ be a morphism such that $\dim\overline{\phi(C)}\ge 2$ for each $C\in\operatorname{Irr} X$. For f in a set of density 1, there is a bijection $\operatorname{Irr} X_{\mathbb{F}}\to\operatorname{Irr}_{\mathrm{horiz}}(X_f)_{\mathbb{F}}$ sending C to $(C_f)_{\mathrm{horiz}}$.*

Proof. Replace X by X_{red} to assume that X is reduced. Because of Lemma 3.3, we may shrink X to assume that X is smooth; then its irreducible components are disjoint. If the conclusion of Proposition 5.5 holds for each $C\in\operatorname{Irr} X$, then it holds for X. So assume that X is irreducible. Let $X' := \overline{\phi(X)}$. Consider $C\in\operatorname{Irr} X_{\mathbb{F}}$. Let $C' := \overline{\phi(C)}\in\operatorname{Irr} X'_{\mathbb{F}}$. For f in a set of density 1, Proposition 5.4 for X' implies irreducibility of C'_f, which, by Lemma 5.2 for $C\to C'\hookrightarrow\mathbb{P}^n_{\mathbb{F}}$, implies irreducibility of $(C_f)_{\mathrm{horiz}}$. By Lemma 3.4, there is a bijection $\operatorname{Irr} X_{\mathbb{F}}\to\operatorname{Irr}_{\mathrm{horiz}}(X_f)_{\mathbb{F}}$ sending C to $(C_f)_{\mathrm{horiz}}$. □

Proof of Theorem 1.6. We are given a finite-type $\mathbb{F}$-scheme X and a morphism $\phi\colon X\to\mathbb{P}^n_{\mathbb{F}}$ such that $\dim\overline{\phi(C)}\ge 2$ for each $C\in\operatorname{Irr} X$. Let $\mathbb{F}_r$ be a finite extension of $\mathbb{F}_q$ such that X, ϕ, and all irreducible components of X are defined over $\mathbb{F}_r$.

From now on, consider X and ϕ as objects over $\mathbb{F}_r$. We need to prove that there is a bijection $\operatorname{Irr} X_{\mathbb{F}} \to \operatorname{Irr}_{\text{horiz}}(X_f)_{\mathbb{F}}$ sending C to $(C \cap X_f)_{\text{horiz}}$.

The composition $X \to \operatorname{Spec} \mathbb{F}_r \to \operatorname{Spec} \mathbb{F}_q$ lets us reinterpret X as a finite-type $\mathbb{F}_q$-scheme $\mathcal{X}$ with a morphism ψ to $\mathbb{P}^n_{\mathbb{F}_q}$ fitting in a commutative diagram

$$\begin{tikzcd} X \arrow[r, equal] \arrow[d, "\phi"'] & \mathcal{X} \arrow[d, "\psi"] \\ \mathbb{P}^n_{\mathbb{F}_r} \arrow[r] \arrow[d] & \mathbb{P}^n_{\mathbb{F}_q} \arrow[d] \\ \operatorname{Spec} \mathbb{F}_r \arrow[r] & \operatorname{Spec} \mathbb{F}_q. \end{tikzcd}$$

Since X and $\mathcal{X}$ are equal as schemes (forgetting the base field), each irreducible component C of X is an irreducible component $\mathcal{C}$ of $\mathcal{X}$. The morphism $\mathbb{P}^n_{\mathbb{F}_r} \to \mathbb{P}^n_{\mathbb{F}_q}$ is finite, so $\dim \overline{\psi(\mathcal{C})} = \dim \overline{\phi(C)} \geq 2$. Proposition 5.5 applied to ψ yields a bijection $\operatorname{Irr} \mathcal{X}_{\mathbb{F}} \to \operatorname{Irr}_{\text{horiz}}(\mathcal{X}_f)_{\mathbb{F}}$. We now rephrase this in terms of X. The identification of X with $\mathcal{X}$ equates $X_f := \phi^{-1}(H_f)_{\mathbb{F}_r}$ with $\mathcal{X}_f := \psi^{-1} H_f$. Then we have a diagram of $\mathbb{F}$-schemes

$$\begin{tikzcd} (\mathcal{X}_f)_{\mathbb{F}} \arrow[r, equal] \arrow[d] & \coprod_\sigma {}^\sigma X_f \arrow[d] \\ \mathcal{X}_{\mathbb{F}} \arrow[r, equal] & \coprod_\sigma {}^\sigma X \end{tikzcd}$$

where σ ranges over $\mathbb{F}_q$-homomorphisms $\mathbb{F}_r \to \mathbb{F}$, and ${}^\sigma X$ denotes the corresponding base extension (and ${}^\sigma X_f$ is similar), and the vertical map on the right is induced by the inclusion $X_f \hookrightarrow X$. Thus for each σ, there is a bijection $\operatorname{Irr} {}^\sigma X \to \operatorname{Irr}_{\text{horiz}} {}^\sigma X_f$ sending each C to $(C \cap {}^\sigma X_f)_{\text{horiz}}$. Taking σ to be the inclusion $\mathbb{F}_r \hookrightarrow \mathbb{F}$ yields the conclusion of Theorem 1.6. $\square$

Acknowledgements

The authors thank Ivan Panin and Zinovy Reichstein for independently asking in June 2013 whether Theorem 1.1 is true. The authors thank also Lior Bary-Soroker for sending his letter [BS13], in which he developed a new explicit Hilbert irreducibility theorem to prove that for a geometrically irreducible quasi-projective variety X over $\mathbb{F}_q$ of dimension $r + 1 \geq 2$, there is *some* dominant rational map $\phi\colon X \dashrightarrow \mathbb{P}^{r+1}_{\mathbb{F}_q}$ such that for a density 1 set of hypersurfaces of the form $y = h(x_1, \ldots, x_r)$ in $\mathbb{A}^{r+1}_{\mathbb{F}_q}$, the closure of its preimage in X is geometrically irreducible. The authors thank Kęstutis Česnavičius for asking a question that led to Theorem 1.8. Finally, the authors thank the referees for many helpful suggestions.

References

[BS13] Lior Bary-Soroker, September 16, 2013. Letter to Bjorn Poonen, available at `http://www.math.tau.ac.il/~barylior/files/let_poonen.pdf`.

[Ben11] Olivier Benoist, *Le théorème de Bertini en famille* (French, with English and French summaries), Bull. Soc. Math. France **139** (2011), no. 4, 555–569. MR2869305 (2012m:14100)

[Con07] Brian Conrad, *Deligne's notes on Nagata compactifications*, J. Ramanujan Math. Soc. **22** (2007), no. 3, 205–257. MR2356346 (2009d:14002)

[DR14] Alexander Duncan and Zinovy Reichstein, *Pseudo-reflection groups and essential dimension*, February 28, 2014. Preprint, `arXiv:1307.5724v2`.

[Jou83] Jean-Pierre Jouanolou, *Théorèmes de Bertini et applications* (French), Progress in Mathematics, vol. 42, Birkhäuser Boston, Inc., Boston, MA, 1983. MR725671 (86b:13007)

[Lan56] Serge Lang, *Sur les séries L d'une variété algébrique* (French), Bull. Soc. Math. France **84** (1956), 385–407. MR0088777 (19,578c)

[Laz04] Robert Lazarsfeld, *Positivity in algebraic geometry. I*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics], vol. 48, Springer-Verlag, Berlin, 2004. Classical setting: line bundles and linear series. MR2095471 (2005k:14001a)

[Mum70] David Mumford, *Abelian varieties*, Tata Institute of Fundamental Research Studies in Mathematics, No. 5, Published for the Tata Institute of Fundamental Research, Bombay; Oxford University Press, London, 1970. MR0282985 (44 #219)

[Neu98] Konrad Neumann, *Every finitely generated regular field extension has a stable transcendence base*, Israel J. Math. **104** (1998), 221–260, DOI 10.1007/BF02897065. MR1622303 (99d:12002)

[Pan14a] Ivan Panin, *On Grothendieck-Serre conjecture concerning principal G-bundles over regular semi-local domains containing a finite field: I*, June 2, 2014. Preprint, `arXiv:1406.0241v1`.

[Pan14b] Ivan Panin, *On Grothendieck-Serre conjecture concerning principal G-bundles over regular semi-local domains containing a finite field: II*, June 2, 2014. Preprint, `arXiv:1406.1129v1`.

[Pan14c] Ivan Panin, *Proof of Grothendieck-Serre conjecture on principal G-bundles over regular local rings containing a finite field*, June 2, 2014. Preprint, `arXiv:1406.0247v1`.

[Poo04] Bjorn Poonen, *Bertini theorems over finite fields*, Ann. of Math. (2) **160** (2004), no. 3, 1099–1127, DOI 10.4007/annals.2004.160.1099. MR2144974 (2006a:14035)

[Poo08] Bjorn Poonen, *Smooth hypersurface sections containing a given subscheme over a finite field*, Math. Res. Lett. **15** (2008), no. 2, 265–271, DOI 10.4310/MRL.2008.v15.n2.a5. MR2385639 (2009c:14037)

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307; and CNRS, Laboratoire de Mathématiques d'Orsay, Université Paris-Sud, 91405 Orsay CEDEX, France

E-mail address: `francois.charles@math.u-psud.fr`

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307

E-mail address: `poonen@math.mit.edu`

URL: `http://math.mit.edu/~poonen/`