A short proof of Deligne's bound for Frobenius eigenvalues, after Alpöge

Version 2, 24 November 2026 (first version: 21 November 2026) CC0 1.0

Abstract

We write out in full the proof, sketched by L. Alpöge in a footnote of his note of 20 November 2026, of Deligne's bound for the eigenvalues of Frobenius. Let X be a separated scheme of finite type over a finite field with q elements, and let $\mathcal{F}$ be a constructible ℓ-adic sheaf on X whose Frobenius eigenvalues at each closed point x have absolute value at most $|\kappa(x)|^{w/2}$. Then the eigenvalues of Frobenius on $H^i_c(X,\mathcal{F})$ have absolute value at most $q^{(w+i)/2}$, and they have absolute value exactly $q^{(w+i)/2}$ if X is smooth and proper and $\mathcal{F}$ is lisse and pure of weight w. The proof is short. Granted the formalism of étale cohomology (the trace formula, the Künneth formula, Poincaré duality and Artin's vanishing theorem) and a weak Lefschetz theorem for generic hyperplane sections of affine varieties, it takes six pages, and its only estimate is the Cauchy–Schwarz inequality for a character sum over a finite field. Lefschetz pencils, monodromy and Rankin's method are not used. The proof was produced by GPT-6 Astra (ChatGPT Pro mode) from that sketch. It was checked by Claude Opus 5.5 at effort high and, independently, by Claude Fable 5.2 at effort max, which also wrote this text. In this second version the presentation follows, where it is cleaner, the exposition of the case of the affine line by D. Zhang (arXiv:2611.30749), and a section of remarks is added: his formulation with a single cohomology group, and three statements of what the method does not give. Apart from that exposition, we know of no check of the proof by a human.

hexagon:2611.00789v2

1 The theorem

Let $\mathbf{F}_q$ be a finite field, ℓ a prime not dividing q, and $\iota : \overline{\mathbf{Q}}_\ell \hookrightarrow \mathbf{C}$ an embedding of fields, fixed throughout. The absolute value of an element of $\overline{\mathbf{Q}}_\ell$ always means the absolute value of its image under ι. Let X be a separated scheme of finite type over $\mathbf{F}_q$ and $\mathcal{F}$ a constructible $\overline{\mathbf{Q}}_\ell$-sheaf on X. The groups $H^i_c(X,\mathcal{F})$ are those of $X \otimes \overline{\mathbf{F}}_q$; they are finite-dimensional, and the geometric Frobenius F acts on them. At a closed point x of X, with residue field $\kappa(x)$, the geometric Frobenius of $\kappa(x)$ acts on the stalk $\mathcal{F}_{\bar{x}}$. For a real number w, we say that $\mathcal{F}$ *has weights at most* w if at every closed point x all eigenvalues of this local Frobenius have absolute value at most $|\kappa(x)|^{w/2}$, and that $\mathcal{F}$ is *pure of weight* w if they all have absolute value exactly $|\kappa(x)|^{w/2}$.

Theorem 1. *Let X be a separated scheme of finite type over $\mathbf{F}_q$ and $\mathcal{F}$ a constructible $\overline{\mathbf{Q}}_\ell$-sheaf on X.*

(a) If $\mathcal{F}$ has weights at most w, then for every i every eigenvalue of F on $H^i_c(X,\mathcal{F})$ has absolute value at most $q^{(w+i)/2}$.

(b) If X is smooth and proper and $\mathcal{F}$ is lisse and pure of weight w, then for every i every eigenvalue of F on $H^i(X,\mathcal{F})$ has absolute value $q^{(w+i)/2}$.

Part (a) says that the passage from the stalks of $\mathcal{F}$ to its cohomology in degree i raises the weights by at most i. Part (b) says that on a smooth proper scheme nothing is lost: the cohomology in degree i of a pure sheaf is pure, of weight $w+i$. For the constant sheaf $\overline{\mathbf{Q}}_\ell$ on a smooth projective variety X the hypothesis of (b) holds with $w=0$ for every embedding ι, so each eigenvalue of F on $H^i(X, \overline{\mathbf{Q}}_\ell)$ has absolute value $q^{i/2}$ under every embedding. This is the Riemann hypothesis for X.

Both statements are theorems of Deligne. Part (b) for the constant sheaf is [3, Théorème 1.6], apart from the assertions made there about integrality and independence of ℓ, which we do not treat. Part (a) is the statement about absolute values contained in [4, Théorème 3.3.1] for a scheme over a finite field. That theorem concerns a morphism $f : X \to S$ and a mixed sheaf, and says that $R^i f_! \mathcal{F}$ is mixed of weights at most $w+i$. For absolute values the relative statement follows from (a), applied to the fibres of f, because the stalks of $R^i f_! \mathcal{F}$ are the groups H^i_c of the fibres. Part (a) asks only for a bound at each closed point, for the one embedding ι, and does not assume that $\mathcal{F}$ is mixed. We do not claim this as a gain in generality. By results of L. Lafforgue, an irreducible lisse sheaf on a smooth curve over $\mathbf{F}_q$ whose determinant has finite order is pure of weight zero for every embedding [9, Théorème VII.6 (ii)], and on a normal scheme of finite type over $\mathbf{F}_q$ every lisse sheaf is, in a unique way, a direct sum $\bigoplus_\chi \sigma_\chi \otimes \chi$ of twists by characters χ [9, Corollaire VII.8]. Nothing is proved here about the other parts of Deligne's theory: that the sheaves $R^i f_! \mathcal{F}$ are mixed, the weight filtration, the local monodromy theorem and the hard Lefschetz theorem of [4], the rationality and the functional equation of zeta functions, or the semisimplicity of Frobenius.

2 The proof, its source and its length

The source. The note [1] of L. Alpöge mentions a proof of the theorems of Weil and Deligne that uses only a sign, based on an identity $qN = M^k + R$, and sketches it in a footnote [1, p. 2 and footnote 3]. In the sketch N is a weighted count of the zeros of the form $x_1x_2 + \cdots + x_{k-1}x_k$, written as a sum over additive characters, and M^k is the term of the trivial character. An argument with a hyperplane section gives the contribution of middle cohomology a definite sign, one degree away from that of M^k. The remainder R has square-root size, by Parseval's identity applied k times. Because of the sign, the inverse roots cannot cancel one another, so the bound for R is a bound for each of them. In higher dimension vectors replace the scalars x_i.

The present note writes this argument out in full. Its ideas are those of the sketch. What is supplied here is the detail: the reduction to the lowest cohomology group of a smooth affine variety, a finite map to affine space that turns the count into the character sum, the weak Lefschetz theorem that provides the sign, and the passage to all degrees.

An exposition. D. Zhang has written an exposition of this proof, dated 22 November 2026, for a lisse sheaf on the affine line [16]. In that case the auxiliary choices of the general argument disappear, and Zhang organizes the cohomological half of the proof around a single cohomology group (Remark 10). His note lists no error in the first version of the present text; it describes that version as not well polished. The present version follows his presentation in three places. The lemma on power sums is stated for one space, and the passage to k-th roots is made where it is used, in the proof of Theorem 9 [16, Lemma 2.1, and the deduction of Theorem 1.1 from Proposition 1.5]. The weak Lefschetz property has a name and one definition (Section 5; compare [16, Definition 2.3]). And the sketch below begins with the affine line. Two of his remarks are recalled at their places: on the form of the equation, after Lemma 3, and on the single cohomology group, in Remark 10. For an open subscheme of the affine line, Zhang reduces to the affine line itself by Abhyankar's finite étale map $t \mapsto t + h(t)^{-q}$ [16, Remark 1.2].

In the present note the finite map π of Section 7 plays this role; it need not be étale, because only trace functions are pushed forward along it. Zhang describes the present note as a direct proof of Deligne's general theorem [16, Remark 1.3]. Its exact scope is that of Section 1: the bound for absolute values is proved, and the other assertion of that theorem, that the sheaves $R^i f_! \mathcal{F}$ are mixed, is not.

The length. The proof is short, and we say in what sense. We take for granted the formalism of ℓ-adic cohomology as it stood before [3]: the trace formula, the Künneth formula, Poincaré duality, the vanishing of the cohomology of an affine scheme above its dimension, and the exact sequence attached to an open subscheme and its complement. We also take for granted one theorem of a less elementary kind: for a generic hyperplane section Y of a smooth affine variety Z, and a lisse sheaf on Z, restriction from Z to Y is bijective on cohomology below the dimension of Y and injective in that degree. The proof of it that we use is Deligne's, printed in the appendix of [6]. It is a few pages long and rests on his generic base change theorem and on Artin's vanishing theorem; Section 6 recalls it. With these inputs, the proof of Theorem 1 occupies Sections 3 to 8, six pages, and its only estimate is the Cauchy–Schwarz inequality, applied to a character sum over a finite field (Lemma 3).

For comparison, [3] uses the same formalism together with the theory of Lefschetz pencils, a theorem of Kazhdan and Margulis on their monodromy, and Rankin's method. In [4] the general statement is reduced to the case of curves, which is proved by a study of monodromy and by the method of Hadamard and de la Vallée-Poussin. None of these enters below. Laumon [10] gave another proof of the main theorem of [4], by the ℓ-adic Fourier transform, and Kedlaya [8] carried that proof over to p-adic cohomology. Katz gave a variant of Laumon's proof for curves that uses Larsen's alternative [7]; we take this description from [16, Remark 1.4], which also observes that all these proofs pass through tensor powers, as the present one does.

For curves, the estimate for the number of points has an elementary proof, by the method of Stepanov [2]. Scholl [11] reduced the Riemann hypothesis for smooth projective varieties to the case of hypersurfaces, and remarked that an elementary proof of the point-count estimate for hypersurfaces, in the manner of Stepanov, would be of great interest. The proof below is not of that kind. Its estimate is elementary, but it concerns one explicit hypersurface in a power of the variety, and the passage from that estimate to the eigenvalues uses cohomology throughout.

The idea, on the affine line. The case of a lisse sheaf $\mathcal{F}$ on the affine line $\mathbf{A}^1$ over $\mathbf{F}_q$, with weights at most w, already contains the whole idea; it is the case treated in [16]. The group to be bounded is $V = H^1_c(\mathbf{A}^1, \mathcal{F})$. For a finite extension of $\mathbf{F}_q$ with Q elements, write $t(x)$ for the trace of Frobenius on the stalk of $\mathcal{F}$ at a point x of $\mathbf{A}^1$ with values in that extension, so that $|t(x)|$ is at most a constant times $Q^{w/2}$. The trace formula expresses $\sum_x t(x)$ through the traces of Frobenius on V and on $H^2_c(\mathbf{A}^1, \mathcal{F})$, and the trivial estimate of this sum gives for V only the weights at most $w+2$. To gain one, a sum has to be shown to cancel, and the sum $\sum_x t(x)$ itself offers no hold.

The way out is to compare two counts on the k-fold product $\mathbf{A}^k$, for large k. Let $Y \subset \mathbf{A}^k$ be the hypersurface

$$x_1x_2 + x_2x_3 + \cdots + x_{k-1}x_k + a_1x_1 + \cdots + a_kx_k + c = 0,$$

and put $M = \sum_x t(x)$ and $N = \sum_{x \in Y} t(x_1) \cdots t(x_k)$, the sums running over the points with values in the field with Q elements. A hypersurface in general position should carry one Q-th of the total M^k, so we put $R = QN - M^k$ and call it the remainder. The proof has two halves.

The sign. For generic coefficients $a_1, \ldots, a_k$ and c the hypersurface Y is smooth and has the weak Lefschetz property: every compactly supported cohomology group of $\mathbf{A}^k$ above the

lowest degree k is the image, bijectively or surjectively, of a group of Y (Lemma 8). Written with the trace formula, R is an alternating sum over the groups of $\mathbf{A}^k$ and of Y, and along these maps the terms cancel in pairs. Two terms survive. One is the lowest group of $\mathbf{A}^k$, which by the Künneth formula is $V^{\otimes k}$. The other is the kernel K of the map out of the lowest group of Y. The first sits in degree k and is subtracted; the second sits in degree $k-1$ and is added; so they enter R with the same sign (Proposition 4). Up to that one sign, R is therefore the trace of Frobenius on an actual vector space, the sum of $V^{\otimes k}$ and K, and not a difference of two traces. This matters because of an elementary fact about power sums (Lemma 2): if the traces of the powers of an automorphism are at most $C\,B^n$ in absolute value, then every one of its eigenvalues has absolute value at most B. For a difference of two traces this is false, since eigenvalues can cancel (Remark 11).

The estimate. Detecting the equation of Y by additive characters writes QN as a sum over the characters $\psi(h\,\cdot)$ of the field. The term with $h = 0$ is exactly M^k, so R is the sum of the other terms. For $h \neq 0$ the term is

$$\psi(hc) \sum_x t_1(x_1) \cdots t_k(x_k)\, \psi\big(h(x_1x_2 + \cdots + x_{k-1}x_k)\big), \qquad t_i(x) = t(x)\, \psi(ha_i x).$$

It is a product of $k-1$ Fourier transforms on the field, interlaced with the multiplications by the functions t_i. A Fourier transform is $Q^{1/2}$ times a unitary operator, which is Parseval's identity, and the multiplications are bounded by a constant times $Q^{w/2}$. So each term is at most a constant to the power k, times $Q^{kw/2}\,Q^{(k+1)/2}$ (Lemma 3), and $|R|$ is at most a constant to the power k, times $Q^{k(w+1)/2}$, times $Q^{3/2}$. This is square-root size, up to the factor $Q^{3/2}$, which does not depend on k.

Putting the two halves together, every eigenvalue of Frobenius on $V^{\otimes k}$, in particular α^k for an eigenvalue α on V, has absolute value at most $q^{k(w+1)/2+3/2}$. Taking k-th roots gives $|\alpha| \leq q^{(w+1)/2} \cdot q^{3/(2k)}$, and letting k grow removes the loss. This last step is known as the tensor power trick.

In general. Three things are added to the case of the affine line. First, for a smooth affine scheme U of pure dimension d and a lisse sheaf $\mathcal{F}$ on U, the same argument bounds the lowest group $H^d_c(U, \mathcal{F})$ (Theorem 9). One chooses a finite map π from U to the affine space $\mathbf{A}^d$ and a nondegenerate bilinear form b on $\mathbf{A}^d$, and takes for $Y \subset U^k$ the hypersurface

$$\sum_{i=1}^{k-1} b\big(\pi(u_i), \pi(u_{i+1})\big) + \sum_{i=1}^{k} g_i(u_i) + c = 0,$$

where each g_i is a linear combination of coordinates of U. Consecutive factors are coupled through b, and there is no other coupling. Collecting the points of U according to their images under π turns the terms with $h \neq 0$ into sums of the same kind on $\mathbf{A}^d$, and the loss becomes $Q^{d/2+1}$. That such a Y can be chosen with the weak Lefschetz property is not obvious, because the bilinear part of its equation is fixed in advance. But Y is a hyperplane section of U^k for the embedding given by the coordinates together with the fixed bilinear expression, and a generic hyperplane whose coefficient on that expression is 1 will do (Lemma 8). Secondly, the groups $H^i_c(U, \mathcal{F})$ with $i > d$ are images of groups of a hyperplane section of U, of smaller dimension, again by the weak Lefschetz theorem; with the first point this gives part (a) for U and $\mathcal{F}$, by induction on d. Thirdly, for general X and $\mathcal{F}$ one cuts X into pieces and uses induction on the dimension (Section 8). Part (b) follows from part (a) by Poincaré duality.

3 Power sums of an actual module

Let W be a finite-dimensional vector space over $\overline{\mathbf{Q}}_\ell$ with an automorphism Φ. Then $\mathrm{Tr}(\Phi^n \mid W) = \sum_j m_j \lambda_j^n$, where the λ_j are the distinct eigenvalues of Φ and the $m_j \geq 1$ their algebraic

multiplicities. The next lemma is the form in which we use the positivity of the m_j.

Lemma 2. *Let W be a finite-dimensional $\overline{\mathbf{Q}}_\ell$-vector space with an automorphism Φ, and let $C > 0$ and $B > 0$ be real numbers such that*

$$|\iota \operatorname{Tr}(\Phi^n \mid W)| \le C\, B^n \qquad \text{for all } n \ge 1.$$

Then every eigenvalue μ of Φ satisfies $|\iota\mu| \le B$.

Proof. Let $\mu_1, \dots, \mu_s$ be the images under ι of the distinct eigenvalues of Φ, and $m_1, \dots, m_s \ge 1$ their multiplicities, so that $\iota \operatorname{Tr}(\Phi^n \mid W) = \sum_j m_j \mu_j^n$ for all $n \ge 1$. The μ_j are distinct, because ι is injective, and nonzero. By the hypothesis, the power series $\sum_{n\ge1} \iota \operatorname{Tr}(\Phi^n \mid W)\, z^n$ converges in the disc $|z| < 1/B$. Near $z = 0$ it is equal to the rational function $\sum_j m_j \mu_j z/(1 - \mu_j z)$. Since the μ_j are distinct and no m_j is zero, this rational function has a pole at each of the points $1/\mu_j$. It agrees on the disc with a holomorphic function, so none of these points lies in the disc. Hence $|\mu_j| \le B$ for every j. □

The lemma will be applied, for each k, to a space of the form $W = V^{\otimes k} \oplus K$, with Φ a power of Frobenius. If $v \ne 0$ and $Fv = \alpha v$ on V, then $v \otimes \cdots \otimes v$ is an eigenvector on $V^{\otimes k}$ with eigenvalue α^k, so a bound for the eigenvalues on W is a bound for α^k. The lemma is false for a difference of two traces: the coefficients m_j can then be negative or zero, and eigenvalues can cancel (Remark 11).

4 A character sum along a path

The estimate of this note is the following. The sum runs over k points of a finite vector space; each point is weighted by an arbitrary bounded function, and the phase couples only consecutive points.

Lemma 3. *Let E be a finite field with Q elements, ψ a nontrivial additive character of E, $d \ge 1$, $A = E^d$ and $k \ge 2$. Let $b_1, \dots, b_{k-1}$ be nondegenerate bilinear forms on A with values in E, and let $t_1, \dots, t_k : A \to \mathbf{C}$ be functions with $|t_i(x)| \le T$ for all i and x. Then for every $h \in E^\times$ and $c \in E$*

$$\Big| \sum_{x \in A^k} t_1(x_1) \cdots t_k(x_k)\, \psi\Big(h\Big(\sum_{i=1}^{k-1} b_i(x_i, x_{i+1}) + c\Big)\Big)\Big| \le T^k Q^{d(k+1)/2}.$$

Proof. For $i < k$ let K_i be the matrix with rows and columns indexed by A and entries $K_i(x, y) = \psi(h\, b_i(x, y))$. The entry of $K_i K_i^*$ at (x, x') is $\sum_{y \in A} \psi(h\, b_i(x - x', y))$. If $x \ne x'$, the map $y \mapsto h\, b_i(x - x', y)$ is a nonzero linear form on A, because b_i is nondegenerate and $h \ne 0$, and the sum of a nontrivial character over A is zero. If $x = x'$ the sum is Q^d. So $K_i K_i^* = Q^d$, and K_i is $Q^{d/2}$ times a unitary matrix. Let D_i be the diagonal matrix with entries $t_i(x)$; its norm is at most T. Regard t_1 and t_k as vectors; their norms are at most $TQ^{d/2}$. The sum in the lemma equals $\psi(hc)\, t_1^{\mathsf{T}} K_1 D_2 K_2 \cdots D_{k-1} K_{k-1} t_k$, and by the Cauchy–Schwarz inequality its absolute value is at most

$$\|t_1\|\, \|t_k\| \prod_{i=1}^{k-1} \|K_i\| \prod_{i=2}^{k-1} \|D_i\| \le (TQ^{d/2})^2\, Q^{d(k-1)/2}\, T^{k-2} = T^k Q^{d(k+1)/2}.$$ □

The trivial bound for the sum is $T^k Q^{dk}$. The lemma gains a factor $Q^{d/2}$ for each of the $k - 1$ couplings. For $k = 2$ and $d = 1$ it is the form over a finite field of Vinogradov's bilinear estimate, as Zhang remarks [16, Remark 2.7]. He also observes there that the path matters. If the k points were coupled in $k/2$ disjoint pairs, by the phase $b(x_1, x_2) + b(x_3, x_4) + \cdots$, the sum would be a product of $k/2$ sums of the case $k = 2$, and the gain would be only $Q^{dk/4}$. Along the path the gain is $Q^{d(k-1)/2}$, which is what Section 7 needs, up to a factor that does not depend on k.

5 The remainder of a point count

We recall what duality gives. Let Z be a smooth scheme of pure dimension m over a finite field $\mathbf{F}_r$ and $\mathcal{G}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on Z. Poincaré duality [12, XVIII, §3] is a perfect pairing, compatible with Frobenius, between $H^i_c(Z,\mathcal{G})$ and $H^{2m-i}(Z,\mathcal{G}^\vee)(m)$. Here (m) is the Tate twist; on a space twisted by (-1) the eigenvalues of Frobenius are multiplied by r. If Z is affine, then $H^i(Z,\mathcal{G}^\vee)=0$ for $i>m$ by Artin's theorem [12, XIV, 3.1], and therefore $H^i_c(Z,\mathcal{G})=0$ for $i<m$. If $Y\subset Z$ is a smooth closed subscheme of pure dimension $m-1$, we use duality on Z and on Y: the transpose of the restriction map $H^a(Z,\mathcal{G}^\vee)\to H^a(Y,\mathcal{G}^\vee)$ is a map

$$H^j_c(Y,\mathcal{G})(-1)\longrightarrow H^{j+2}_c(Z,\mathcal{G}),\qquad j=2m-2-a,$$

compatible with Frobenius. We call it the *Gysin map*. It is bijective when the restriction map is bijective, and surjective when the restriction map is injective.

We say that Y has the *weak Lefschetz property* for $\mathcal{G}$ if, over $\overline{\mathbf{F}}_r$, the restriction map $H^a(Z,\mathcal{G}^\vee)\to H^a(Y,\mathcal{G}^\vee)$ is bijective for $a<m-1$ and injective for $a=m-1$. Equivalently, the Gysin map $H^j_c(Y,\mathcal{G})(-1)\to H^{j+2}_c(Z,\mathcal{G})$ is bijective for $j>m-1$ and surjective for $j=m-1$. The empty subscheme is allowed, and the same definition is used over an algebraically closed field. The property is named for the sheaf whose compactly supported cohomology is studied, as in [16, Definition 2.3]; it is a condition on the cohomology of the dual sheaf.

Proposition 4. *Let U be a smooth affine scheme of pure dimension $d\geq 1$ over a finite field $\mathbf{F}_r$, and $\mathcal{F}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U. Let $k\geq 1$, $X=U^k$, $D=kd$ and $\mathcal{F}_k=\mathcal{F}\boxtimes\cdots\boxtimes\mathcal{F}$ on X. Let $Y\subset X$ be a smooth closed subscheme of pure dimension $D-1$, possibly empty, defined over $\mathbf{F}_r$, with the weak Lefschetz property for $\mathcal{F}_k$: over $\overline{\mathbf{F}}_r$ the restriction map*

$$H^a(X,\mathcal{F}_k^\vee)\longrightarrow H^a(Y,\mathcal{F}_k^\vee)$$

is an isomorphism for $a<D-1$ and injective for $a=D-1$. For $n\geq 1$ let M_n be the sum of the traces $\mathrm{Tr}(F_{u,n}\mid\mathcal{F}_{\bar u})$ over $u\in U(\mathbf{F}_{r^n})$, let N_n be the corresponding sum for $\mathcal{F}_k$ over $Y(\mathbf{F}_{r^n})$, and put $R_n=r^nN_n-M_n^k$. Let F be the geometric Frobenius of $\mathbf{F}_r$, and let K be the kernel of the Gysin map $H^{D-1}_c(Y,\mathcal{F}_k)(-1)\to H^{D+1}_c(X,\mathcal{F}_k)$. Then for all $n\geq 1$

$$(-1)^{D-1}R_n=\mathrm{Tr}\big(F^n\mid H^d_c(U,\mathcal{F})^{\otimes k}\big)+\mathrm{Tr}\big(F^n\mid K\big).$$

Proof. The schemes U, X and Y are smooth and affine, of pure dimensions d, D and $D-1$. By what was recalled, the groups $H^j_c(U,\mathcal{F})$ vanish for $j<d$, the groups $H^j_c(X,\mathcal{F}_k)$ for $j<D$, and the groups $H^j_c(Y,\mathcal{F}_k)$ for $j<D-1$. The hypothesis on restriction says that the Gysin map $H^j_c(Y,\mathcal{F}_k)(-1)\to H^{j+2}_c(X,\mathcal{F}_k)$ is an isomorphism for $j>D-1$ and surjective for $j=D-1$.

The stalk of $\mathcal{F}_k$ at a point $(u_1,\dots,u_k)$ is the tensor product of the stalks of $\mathcal{F}$ at the u_i, so the sum of the traces of $\mathcal{F}_k$ over $X(\mathbf{F}_{r^n})=U(\mathbf{F}_{r^n})^k$ is M_n^k. By the trace formula [13, Rapport, Théorème 3.2], applied on X and on Y,

$$M_n^k=\sum_{b\geq 0}(-1)^{D+b}\,\mathrm{Tr}\big(F^n\mid H^{D+b}_c(X,\mathcal{F}_k)\big),$$
$$r^nN_n=\sum_{a\geq 0}(-1)^{D-1+a}\,\mathrm{Tr}\big(F^n\mid H^{D-1+a}_c(Y,\mathcal{F}_k)(-1)\big).$$

Therefore, in $(-1)^{D-1}R_n$, the trace on $H^{D-1+a}_c(Y,\mathcal{F}_k)(-1)$ enters with the sign $(-1)^a$ and the trace on $H^{D+b}_c(X,\mathcal{F}_k)$ with the sign $(-1)^b$. For $a\geq 1$ the Gysin map identifies $H^{D-1+a}_c(Y,\mathcal{F}_k)(-1)$ with $H^{D+1+a}_c(X,\mathcal{F}_k)$, which enters with the opposite sign, and the two traces cancel. Three terms remain: the trace on $H^{D-1}_c(Y,\mathcal{F}_k)(-1)$, minus the trace on

$H_c^{D+1}(X, \mathcal{F}_k)$, plus the trace on $H_c^D(X, \mathcal{F}_k)$. The second space is the quotient of the first by K, so the first two terms together are the trace on K. Finally, by the Künneth formula [12, XVII, 5.4.3] the group $H_c^D(X, \mathcal{F}_k)$ is the sum of the tensor products $H_c^{j_1}(U, \mathcal{F}) \otimes \cdots \otimes H_c^{j_k}(U, \mathcal{F})$ with $j_1 + \cdots + j_k = D$. Since the factors vanish for $j_i < d$, only $j_1 = \cdots = j_k = d$ contributes, and $H_c^D(X, \mathcal{F}_k) = H_c^d(U, \mathcal{F})^{\otimes k}$. □

Nothing is assumed about the higher cohomology groups of U or of Y: they cancel through actual isomorphisms, and the surjection at the border leaves an actual kernel. The two spaces of the proposition are the two steps of a filtration of one cohomology group (Remark 10).

Example 5. Let $U = \mathbf{A}^1$, $\mathcal{F} = \overline{\mathbf{Q}}_\ell$ and $k = 2$, and let $Y \subset \mathbf{A}^2$ be the conic $xy + ax + by + c = 0$ with $ab \neq c$. It is isomorphic to the multiplicative group, so $N_n = r^n - 1$, while $M_n = r^n$; hence $R_n = -r^n$ and $(-1)^{D-1}R_n = r^n$. On the other side $H_c^1(\mathbf{A}^1, \overline{\mathbf{Q}}_\ell) = 0$, and $K = H_c^1(Y, \overline{\mathbf{Q}}_\ell)(-1)$ has dimension one, with Frobenius acting by r. The group $H_c^2(Y, \overline{\mathbf{Q}}_\ell)(-1)$ is identified with $H_c^4(\mathbf{A}^2, \overline{\mathbf{Q}}_\ell)$ and has cancelled.

Remark 6. The space K is the analogue, for compact supports on an affine variety, of the part of the cohomology of a hyperplane section that the Gysin map kills. In Grothendieck's notes on the standard conjectures in the Montpellier archive this part is his *partie évanescente* E^{n-1} (cote 16, scan 40, in the transcription [15]). The comparison plays no role in the proof.

6 Hypersurfaces with the weak Lefschetz property

Theorem 7 (Deligne; [6, Appendix, A.5, with A.3 and A.4])**.** *Let Z be a separated scheme of finite type over an algebraically closed field of characteristic different from ℓ, let $\varphi : Z \to \mathbf{P}^N$ be a quasi-finite morphism, and let $\mathcal{K}$ be a perverse sheaf on Z. Then there is a dense open set of hyperplanes $H \subset \mathbf{P}^N$ for which the restriction map*

$$H^a(Z, \mathcal{K}) \longrightarrow H^a(\varphi^{-1}H, \mathcal{K})$$

is an isomorphism for $a < -1$ and injective for $a = -1$.

Katz attributes this proof to Deligne [6, p. 152]. It runs as follows. Let $\alpha : Z \smallsetminus \varphi^{-1}H \to Z$ be the inclusion. The assertion is equivalent to the vanishing of $H^a(Z, \alpha_!\alpha^*\mathcal{K})$ for $a < 0$. For generic H, the generic base change theorem [13, Théorèmes de finitude, 1.9], applied over the space of hyperplanes, identifies these groups with the compactly supported cohomology of the affine scheme $\mathbf{P}^N \smallsetminus H$ with coefficients in a direct image of $\alpha^*\mathcal{K}$; duality and Artin's theorem then give the vanishing. No weights occur. If Z is smooth of pure dimension D and $\mathcal{G}$ is a lisse sheaf on Z, then $\mathcal{G}^\vee[D]$ is perverse, and the conclusion reads: the restriction map $H^a(Z, \mathcal{G}^\vee) \to H^a(\varphi^{-1}H, \mathcal{G}^\vee)$ is an isomorphism for $a < D - 1$ and injective for $a = D - 1$. When $\varphi^{-1}H$ is smooth, this is the weak Lefschetz property of $\varphi^{-1}H$ for $\mathcal{G}$. For affine space, the dual form of this statement is [16, Theorem 2.4].

We need this for hypersurfaces whose equation has a fixed part. Corollary 3.4.1 of [6] states exactly that form, but it is deduced there from the main theorem of that article, whose proof is by the Fourier transform. The next lemma deduces it from Theorem 7 instead, so that only the inputs named above are used. It also gives the smoothness of the hypersurface, which the trace formula on Y and duality on Y require and which does not follow from the weak Lefschetz theorem.

Lemma 8. *Let V be a smooth affine scheme of pure dimension $D \geq 1$ over $\overline{\mathbf{F}}_q$, let $y_1, \ldots, y_M$ be functions on V such that $\Gamma = (y_1, \ldots, y_M) : V \to \mathbf{A}^M$ is a closed immersion, and let $\mathcal{G}$ be a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on V. Consider the hypersurfaces*

$$Y_{\lambda,\mu} = V \cap \Big(\sum_{m=1}^{M} \lambda_m y_m + \mu = 0 \Big), \qquad \lambda \neq 0.$$

For (λ, μ) in a dense open subset of the space of all parameters, and also for (λ, μ) in a dense open subset of the space of parameters with $\lambda_M = 1$, the scheme $Y_{\lambda,\mu}$ is smooth, of pure dimension $D-1$ or empty on each connected component of V, and it has the weak Lefschetz property for $\mathcal{G}$: the restriction map $H^a(V, \mathcal{G}^\vee) \to H^a(Y_{\lambda,\mu}, \mathcal{G}^\vee)$ is an isomorphism for $a < D-1$ and injective for $a = D-1$.

Proof. Regard $\mathbf{A}^M$ as the complement of the hyperplane at infinity in $\mathbf{P}^M$. The hyperplanes of $\mathbf{P}^M$ other than the one at infinity are the zero sets of the equations $\sum_m \lambda_m y_m + \mu = 0$ with $\lambda \neq 0$, and the inverse image of such a hyperplane under Γ is $Y_{\lambda,\mu}$.

The restriction property holds for a dense open set of hyperplanes by Theorem 7, applied to $\varphi = \Gamma$, which is quasi-finite, and to the perverse sheaf $\mathcal{G}^\vee[D]$.

Smoothness holds for a dense open set of hyperplanes by a count of dimensions. Consider the pairs (v, H) of a point v of V and a hyperplane H that contains the tangent space of $\Gamma(V)$ at $\Gamma(v)$. Since Γ is an immersion, this tangent space has dimension D, and the hyperplanes containing it form a projective space of dimension $M-D-1$. So the set of such pairs has dimension at most $D + (M-D-1) = M-1$, and its image in the space of hyperplanes, which has dimension M, is not dense. If H lies outside the closure of this image, the equation of H restricts to a nonzero linear form on the tangent space of V at each point of $\Gamma^{-1}H$, so $\Gamma^{-1}H$ is smooth, and of pure dimension $D-1$ or empty on each component of V. This argument is valid in every characteristic, because Γ is an immersion; compare [5, I.6.10].

The hyperplanes with both properties form a dense open subset of the space of all hyperplanes. It meets the open subset $\lambda_M \neq 0$ in a dense open subset, and on that subset we may scale the equation so that $\lambda_M = 1$. □

7 The lowest degree of a smooth affine variety

Theorem 9. *Let U be a smooth affine scheme of pure dimension $d \geq 1$ over $\mathbf{F}_q$ and $\mathcal{F}$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U. Suppose that there are a constant $C > 0$ and a real number w with*

$$|\iota \operatorname{Tr}(F_{u,n} \mid \mathcal{F}_{\bar{u}})| \leq C\, q^{nw/2} \qquad \text{for all } n \geq 1 \text{ and } u \in U(\mathbf{F}_{q^n}).$$

Then every eigenvalue of F on $H^d_c(U, \mathcal{F})$ has absolute value at most $q^{(w+d)/2}$.

Proof. Choices. By Noether normalization [14, Tag 00OY] there is a finite morphism $\pi : U \to \mathbf{A}^d$ over $\mathbf{F}_q$. Let B be the number of elements in a set of generators of the coordinate ring $\mathcal{O}(U)$ as a module over that of $\mathbf{A}^d$. Over any field, each fibre of π is the spectrum of an algebra of dimension at most B, so it has at most B rational points. Choose generators $z_1, \ldots, z_s$ of $\mathcal{O}(U)$ as an $\mathbf{F}_q$-algebra, so that $z = (z_1, \ldots, z_s) : U \to \mathbf{A}^s$ is a closed immersion, and a nondegenerate bilinear form b on $\mathbf{A}^d$ defined over $\mathbf{F}_q$. Fix $k \geq 2$ and put $D = kd$.

The hypersurface. For parameters $a = (a_{ij})$, with $1 \leq i \leq k$ and $1 \leq j \leq s$, and c, let $P_{a,c}$ be the function on U^k

$$P_{a,c}(u_1, \ldots, u_k) = \sum_{i=1}^{k-1} b\big(\pi(u_i), \pi(u_{i+1})\big) + \sum_{i=1}^{k} \sum_{j=1}^{s} a_{ij}\, z_j(u_i) + c,$$

and let $Y_{a,c} \subset U^k$ be its zero set. Over $\overline{\mathbf{F}}_q$ the map

$$\Gamma : \; u \longmapsto \Big(z(u_1), \ldots, z(u_k),\; \sum_{i=1}^{k-1} b\big(\pi(u_i), \pi(u_{i+1})\big)\Big)$$

is a closed immersion of U^k into $\mathbf{A}^{sk+1}$, because its first sk coordinates already are one. The $Y_{a,c}$ are exactly the hypersurfaces of Lemma 8 for Γ whose coefficient on the last coordinate is

1. By that lemma, applied to the sheaf $\mathcal{F}_k$, there is a dense open set of parameters (a, c) for which $Y_{a,c}$ is smooth of pure dimension $D-1$, or empty, and has the weak Lefschetz property for $\mathcal{F}_k$, as required in Proposition 4. A dense open subset of an affine space over $\overline{\mathbf{F}}_q$ has a point with coordinates in a finite field. Fix such a point (a, c), with coordinates in $\mathbf{F}_r$, $r = q^e$, and write $P = P_{a,c}$ and $Y = Y_{a,c}$. The exponent e and the point (a, c) depend on k.

The remainder as a character sum. Let $n \geq 1$, $Q = r^n$ and $E = \mathbf{F}_Q$, and let M_n, N_n and $R_n = QN_n - M_n^k$ be formed over $\mathbf{F}_r$ as in Proposition 4. Let ψ be a nontrivial additive character of E, and for $u \in U(E)$ put $\tau(u) = \iota \operatorname{Tr}(F_{u,en} \mid \mathcal{F}_{\bar{u}})$; by hypothesis $|\tau(u)| \leq CQ^{w/2}$. For a point $u = (u_1, \ldots, u_k)$ of U^k with values in E, the sum $\sum_{h \in E} \psi(hP(u))$ is Q if $P(u) = 0$ and 0 otherwise. Therefore

$$Q \cdot \iota N_n = \sum_{h \in E} \sum_{u \in U(E)^k} \tau(u_1) \cdots \tau(u_k)\, \psi\big(hP(u)\big).$$

In the inner sum we collect the u_i according to their images $x_i = \pi(u_i)$ in E^d. With

$$t_{i,h}(x) = \sum_{u \in U(E),\ \pi(u) = x} \tau(u)\, \psi\Big(h \sum_{j=1}^{s} a_{ij} z_j(u)\Big),$$

the inner sum becomes

$$\psi(hc) \sum_{x \in (E^d)^k} t_{1,h}(x_1) \cdots t_{k,h}(x_k)\, \psi\Big(h \sum_{i=1}^{k-1} b(x_i, x_{i+1})\Big).$$

For $h = 0$ all phases are 1, and the inner sum is $\big(\sum_{u \in U(E)} \tau(u)\big)^k = (\iota M_n)^k$. For $h \neq 0$, each $t_{i,h}(x)$ is a sum of at most B terms of absolute value at most $CQ^{w/2}$, so Lemma 3 applies with $T = BCQ^{w/2}$ and all b_i equal to b. There are $Q - 1$ such h, and we obtain

$$|\iota R_n| = \big|Q \cdot \iota N_n - (\iota M_n)^k\big| \leq (Q-1)\,(BC)^k Q^{kw/2}\, Q^{d(k+1)/2} \leq (BC)^k\, q^{en\{k(w+d)/2 + d/2 + 1\}}.$$

The eigenvalues of a tensor power. The geometric Frobenius of $\mathbf{F}_r$ is F^e. Put $V = H_c^d(U, \mathcal{F})$ and $W = V^{\otimes k} \oplus K$, with K as in Proposition 4, and let Φ be the automorphism of W induced by F^e on both summands. By that proposition, applied over $\mathbf{F}_r$,

$$(-1)^{D-1}\, \iota R_n = \iota \operatorname{Tr}\big(F^{en} \mid V^{\otimes k}\big) + \iota \operatorname{Tr}\big((F^e)^n \mid K\big) = \iota \operatorname{Tr}\big(\Phi^n \mid W\big)$$

for all $n \geq 1$. With the estimate just proved, $|\iota \operatorname{Tr}(\Phi^n \mid W)| \leq (BC)^k\, B_k^n$ for all $n \geq 1$, where $B_k = q^{e\{k(w+d)/2 + d/2 + 1\}}$. By Lemma 2 every eigenvalue of Φ on W has absolute value at most B_k.

Conclusion. Let α be an eigenvalue of F on V and v an eigenvector. Then $v \otimes \cdots \otimes v$ is an eigenvector of Φ in $V^{\otimes k} \subset W$, with eigenvalue α^{ke}. Hence $|\iota\alpha|^{ke} \leq q^{e\{k(w+d)/2 + d/2 + 1\}}$, that is,

$$|\iota\alpha| \leq q^{(w+d)/2} \cdot q^{(d/2+1)/k}.$$

This holds for every $k \geq 2$: the field $\mathbf{F}_r$, the hypersurface Y and the space K depend on k, and α does not. Letting k tend to infinity gives $|\iota\alpha| \leq q^{(w+d)/2}$. □

8 Proof of Theorem 1

A bound for the eigenvalues may be proved after replacing $\mathbf{F}_q$ by a finite extension $\mathbf{F}_{q^e}$: the eigenvalues of F^e are the e-th powers of those of F, so a bound $q^{e(w+i)/2}$ for the former gives the bound $q^{(w+i)/2}$ for the latter.

Part (a), for X smooth and affine of pure dimension d and $\mathcal{F}$ lisse. We prove the statement by induction on d, for all finite fields at once. Let $d = 0$. Then $H^0_c(X, \mathcal{F})$ is the sum of the stalks of $\mathcal{F}$ at the geometric points of X. If x is a closed point of degree m, the Frobenius F permutes the m geometric points above x cyclically, and F^m acts on each of their stalks as the local Frobenius at x. So for an eigenvalue α of F on the part of H^0_c above x, α^m is an eigenvalue of the local Frobenius, and $|\alpha|^m \leq q^{mw/2}$.

Let $d \geq 1$. For $i < d$ the group $H^i_c(X, \mathcal{F})$ vanishes (Section 5). For $i = d$ we apply Theorem 9, with C the rank of $\mathcal{F}$: if u is a point with values in $\mathbf{F}_{q^n}$ over a closed point x of degree m, then $F_{u,n}$ is the (n/m)-th power of the local Frobenius at x, so its eigenvalues have absolute value at most $q^{nw/2}$. For $i > d$ we cut by a hyperplane. Choose a closed immersion of X into an affine space. By Lemma 8, applied to this immersion, to all hyperplanes and to the sheaf $\mathcal{F}$, there is, over some finite extension $\mathbf{F}_{q^e}$, a hyperplane section Y of X that is smooth and affine of pure dimension $d-1$, or empty, and has the weak Lefschetz property for $\mathcal{F}$. So the Gysin map

$$H^{i-2}_c(Y, \mathcal{F})(-1) \longrightarrow H^i_c(X, \mathcal{F})$$

is surjective for $i = d+1$ and bijective for $i > d+1$ (Section 5). By the induction hypothesis for Y over $\mathbf{F}_{q^e}$, the eigenvalues of F^e on $H^{i-2}_c(Y, \mathcal{F})$ have absolute value at most $q^{e(w+i-2)/2}$. The twist multiplies them by q^e. So the eigenvalues of F^e on $H^i_c(X, \mathcal{F})$ have absolute value at most $q^{e(w+i)/2}$, which is the bound.

Part (a) in general. The cohomology of X is that of the reduced scheme, so we may assume X reduced. We use induction on the dimension of X; in dimension 0 the scheme X is smooth and affine and $\mathcal{F}$ is lisse, which is the case already treated. Since $\mathbf{F}_q$ is perfect, X has a dense open subscheme U that is smooth and affine, of pure dimension on each of its connected components, and on which $\mathcal{F}$ is lisse. Its complement Z has smaller dimension, and the restriction of $\mathcal{F}$ to Z again has weights at most w. The sequence

$$H^i_c(U, \mathcal{F}) \longrightarrow H^i_c(X, \mathcal{F}) \longrightarrow H^i_c(Z, \mathcal{F})$$

is exact and compatible with Frobenius; it is part of the long exact sequence of $0 \to j_! j^* \mathcal{F} \to \mathcal{F} \to i_* i^* \mathcal{F} \to 0$, where j and i are the inclusions of U and Z. So every eigenvalue of F on $H^i_c(X, \mathcal{F})$ is an eigenvalue on $H^i_c(U, \mathcal{F})$ or on $H^i_c(Z, \mathcal{F})$. The first are bounded by the case already treated, applied to the components of U, and the second by the induction hypothesis.

Part (b). Treating the connected components of X separately, we may assume that X has pure dimension d. Since X is proper, $H^i(X, \mathcal{F}) = H^i_c(X, \mathcal{F})$, and part (a) gives the upper bound. The dual sheaf $\mathcal{F}^\vee$ is lisse and pure of weight $-w$, because the eigenvalues of a local Frobenius on its stalk are the inverses of those on the stalk of $\mathcal{F}$; here we use that $\mathcal{F}$ is pure at every point, and not only bounded. Poincaré duality is a perfect pairing of $H^i(X, \mathcal{F})$ with $H^{2d-i}(X, \mathcal{F}^\vee)$ into $\overline{\mathbf{Q}}_\ell(-d)$, compatible with Frobenius. So if α is an eigenvalue of F on $H^i(X, \mathcal{F})$, then q^d/α is an eigenvalue of F on $H^{2d-i}(X, \mathcal{F}^\vee)$. Part (a), applied to $\mathcal{F}^\vee$, gives $|q^d/\alpha| \leq q^{(-w+2d-i)/2}$, that is, $|\alpha| \geq q^{(w+i)/2}$. □

9 Remarks on the method

The four remarks of this section are not used in the proof. The first is a reformulation, due to Zhang. The other three say what the method does not give; they come from other runs on the same subject, and their status is given under "Provenance and checks".

Remark 10 (One cohomology group). In the situation of Proposition 4, the spaces $H^d_c(U, \mathcal{F})^{\otimes k}$ and K are the two steps of a filtration of one cohomology group. Let $i : Y \to X$ be the inclusion and $j : X \smallsetminus Y \to X$ that of the complement. The distinguished triangle

$i_*i^!\mathcal{F}_k \to \mathcal{F}_k \to Rj_*j^*\mathcal{F}_k$ and the purity isomorphism $i^!\mathcal{F}_k \simeq i^*\mathcal{F}_k(-1)[-2]$ give a long exact sequence, compatible with Frobenius,

$$\cdots \to H_c^e(X, \mathcal{F}_k) \to H_c^e(X, Rj_*j^*\mathcal{F}_k) \to H_c^{e-1}(Y, \mathcal{F}_k)(-1) \to H_c^{e+1}(X, \mathcal{F}_k) \to \cdots,$$

in which the third arrow is, up to sign, the Gysin map. Since $H_c^e(X, \mathcal{F}_k)$ vanishes for $e < D$ and $H_c^e(Y, \mathcal{F}_k)$ for $e < D-1$, the subscheme Y has the weak Lefschetz property for $\mathcal{F}_k$ if and only if $H_c^e(X, Rj_*j^*\mathcal{F}_k) = 0$ for all $e \neq D$, and then the sequence

$$0 \to H_c^d(U, \mathcal{F})^{\otimes k} \to H_c^D(X, Rj_*j^*\mathcal{F}_k) \to K \to 0$$

is exact. So Proposition 4 says that $(-1)^{D-1}R_n$ is the trace of F^n on the single space $H_c^D(X, Rj_*j^*\mathcal{F}_k)$. This is how Zhang organizes the argument, for $X = \mathbf{A}^k$ [16, the sequence (†), Lemma 2.2 and Definition 2.3]; he quotes the purity isomorphism from [12, XVI, 3.7, 3.8 and 3.10]. The proof above keeps the formulation with restriction maps, which needs only the trace formula on X and on Y and duality.

Remark 11 (The sign). The sign enters the proof at one point: by Proposition 4 the number $(-1)^{D-1}R_n$ is the trace of F^n on one space, so that Lemma 2 applies. Can a phase, or an operator of finite order, take the place of the scalar $(-1)^{D-1}$? Let M be a finite-dimensional $\overline{\mathbf{Q}}_\ell$-vector space, Φ an automorphism of M, and s an automorphism of M of finite order m.

(a) If s commutes with Φ, then $(s\Phi)^{mn} = \Phi^{mn}$. So the traces of the powers $(s\Phi)^{mn}$ are power sums with positive multiplicities, and a bound $|\iota \operatorname{Tr}((s\Phi)^{mn})| \leq C\, B^{mn}$ for all $n \geq 1$ gives $|\iota\alpha| \leq B$ for every eigenvalue α of Φ, by Lemma 2 applied to Φ^m.

(b) A single insertion of s is not enough, even when s commutes with Φ. In that case $\operatorname{Tr}(s\Phi^n) = \sum_\mu \mu^n \operatorname{Tr}(s \mid M(\mu))$, where μ runs over the eigenvalues of Φ and $M(\mu)$ is the generalized eigenspace; the coefficient $\operatorname{Tr}(s \mid M(\mu))$ can vanish. For $M = M' \oplus M'$ with Φ acting in the same way on both summands and $s = (1, -1)$, all these traces are zero. Here is a less artificial case. Let a be a further automorphism of M, of order 3, commuting with Φ, with $sas^{-1} = a^{-1}$. Then s exchanges the eigenspaces of a for the two primitive cube roots of unity, and Φ preserves them, so $\operatorname{Tr}(s\Phi^n \mid M) = \operatorname{Tr}(s\Phi^n \mid M^a)$, where M^a is the space of invariants of a: the traces do not see the part of M on which a acts nontrivially. On the two-dimensional irreducible representation of the symmetric group S_3, twisted so that Frobenius acts by q^N, one has $\operatorname{Tr}(s\Phi^n) = 0$ for all n, and the same holds on all its tensor powers. Nor can s be made central: in a quotient of the group generated by a and s in which the image of s is central, the image of a is equal to its inverse and has order dividing 3, so it is trivial.

(c) Without commutation, (a) fails as well. Let Φ be the diagonal matrix with entries L and L^{-1}, with $L > 1$, and s the exchange of the two coordinates. Then $s\Phi s^{-1} = \Phi^{-1}$ and $(s\Phi)^2 = 1$, so $\operatorname{Tr}((s\Phi)^{2n}) = 2$ for all n, although Φ has the eigenvalue L.

So an operator can take the place of the sign only if it commutes with Frobenius, and then through its powers, as in (a), or when it acts as one scalar on the whole space, as $(-1)^{D-1}$ does in Proposition 4.

Remark 12 (Twists, and the Gauss sum). Theorem 1 turns a bound on the stalks into a bound on the cohomology. Twisting shows that the bound is sharp, and that the hypothesis on the stalks carries all the information. Let q be odd, let ψ be a nontrivial additive character of $\mathbf{F}_q$ with values in $\overline{\mathbf{Q}}_\ell$, and let $\mathcal{L}$ be the pullback by $x \mapsto x^2$ of the Artin–Schreier sheaf of ψ on $\mathbf{A}^1$. It is lisse of rank one, and its local trace at a point x with values in $\mathbf{F}_{q^n}$ is the root of unity $\psi(\operatorname{Tr}_{\mathbf{F}_{q^n}/\mathbf{F}_q}(x^2))$, so $\mathcal{L}$ is pure of weight 0 for every embedding. Its compactly supported cohomology is a line G_ψ, in degree 1 [13, Sommes trig., §4], and by the trace formula F acts on G_ψ by $\gamma_\psi = -g$, where $g = \sum_{x \in \mathbf{F}_q} \psi(x^2)$ is the quadratic Gauss sum. Theorem 9, with $d = 1$, $C = 1$ and $w = 0$, gives $|\iota g| \leq q^{1/2}$. Since $|\iota g|^2 = q$ by the classical direct computation, the bound of Theorem 1 is attained here. The example is also the simplest test of the proof:

for this sheaf the sums of Lemma 3 have $t_i(x) = \psi_E(x^2 + ha_i x)$, of absolute value 1, where ψ_E is ψ composed with the trace from E to $\mathbf{F}_q$; no weight enters anywhere.

Now let C be a finite-dimensional $\overline{\mathbf{Q}}_\ell$-vector space with an automorphism φ, regarded as a lisse sheaf on $\operatorname{Spec}\mathbf{F}_q$, and let $\mathcal{F}$ be the tensor product of $\mathcal{L}$ with the pullback of C to $\mathbf{A}^1$. The sheaf $\mathcal{F}$ has weights at most w if and only if $|\iota\alpha| \le q^{w/2}$ for every eigenvalue α of φ. Its only cohomology group is $H^1_c(\mathbf{A}^1, \mathcal{F}) = G_\psi \otimes C$, on which the eigenvalues of F are the products $\gamma_\psi \alpha$, of absolute value $q^{1/2}|\iota\alpha|$. So for these sheaves the hypothesis and the conclusion of Theorem 1(a) are the same statement: the theorem gives exactly as much about the cohomology as is assumed on the stalks. Every C occurs in this way, since C is the cohomology of the twist of $\mathcal{L}$ by $C \otimes G_\psi^\vee$. So a vector space with an automorphism cannot be bounded by realizing it as a cohomology group and appealing to the sign: the bound on the stalks has to come from elsewhere.

For these sheaves the half of the argument that concerns the sign holds in an explicit form. For the hypersurface $x_1x_2 + \cdots + x_{k-1}x_k = 0$ itself, without linear terms, the remainder is, up to sign and without any error term, the trace of Frobenius on a space that can be written down, for every even k such that the determinant of the $k \times k$ matrix with entries 1 on the diagonal, $h/2$ next to the diagonal and 0 elsewhere is a squarefree polynomial in h over $\mathbf{F}_q$. This is the case when the characteristic does not divide $k+1$, hence for all $k = 2pr$, where p is the characteristic.

Remark 13 (Traces and semisimplicity). Every quantity that enters the proof is the trace of a power of Frobenius: the local traces in M_n and N_n, and the two traces of Proposition 4. Such traces depend only on the semisimplification of the space on which Frobenius acts. So the method says nothing about the semisimplicity of Frobenius, nor about any statement that would imply it. An example makes this concrete. Let N be a nonzero nilpotent matrix with rational coefficients, and let $\mathcal{U}$ be the lisse sheaf on $\operatorname{Spec}\mathbf{F}_q$ on which the Galois group $\widehat{\mathbf{Z}}$ acts by $z \mapsto \exp(z_\ell N)$, where z_ℓ is the component of z in $\mathbf{Z}_\ell$; the action is continuous because $\exp(z_\ell N)$ is a polynomial in z_ℓ. The sheaf $\mathcal{U}$ is pure of weight 0, and the characteristic polynomial of Frobenius on it is that of the constant sheaf of the same rank, but Frobenius is not semisimple on it. If $\mathcal{F}$ satisfies the hypothesis of Theorem 1(a) or (b), so does its tensor product with the pullback of $\mathcal{U}$ to X. The cohomology of that sheaf is $H^i_c(X, \mathcal{F}) \otimes \mathcal{U}$, and Frobenius is not semisimple on it as soon as it is nonzero.

The same example bears on Grothendieck's notes on the standard conjectures. Take for N the nilpotent Jordan block of size 3. The algebra $\mathbf{Q}[N]$ has the involution $N \mapsto -N$, which is the adjoint with respect to a nondegenerate symmetric form, and for it the trace form $u \mapsto \operatorname{Tr}(u^\dagger u)$ vanishes at $u = N$. Yet for $\Phi = c\exp(N)$ the traces $\operatorname{Tr}(u\Phi^n)$, $u \in \mathbf{Q}[N]$, and all identities that follow from them by tensor products, exterior powers and duals, are those of the semisimple case $\Phi = c$. So trace identities of the kind used here cannot give the positivity $\operatorname{Tr}(u^\dagger u) > 0$ for nonzero algebraic correspondences u. That positivity implies that the algebra of correspondences is semisimple, and it is the condition at stake in those notes at cote 16, scan 49 [15].

Provenance and checks

The proof. The proof is Part I, Section 1 of the output of a research run of GPT-6 Astra (ChatGPT Pro mode) on 21 November 2026, whose task was to work out footnote 3 of [1]. The same day Claude Opus 5.5 at effort high checked every step of it. Claude Fable 5.2 at effort max, without seeing that check, re-proved every step and found no gap; it read the pages of [6] used here (pp. 152–158, 179–181 and 219–222) and re-proved Theorem 7 independently. One correction came from that second check and is adopted here: the weak Lefschetz input is taken from Deligne's proof in Katz's appendix, through Lemma 8, and not from the printed

proof of Katz's Corollary 3.4.1.

Version 2. The proof is that of the first version, of 21 November 2026. Its presentation was revised on 24 November 2026 by Claude Fable 5.2 at effort max, after reading the exposition [16] in full, in its source. That exposition lists no error in the first version. The changes are these. Lemma 2 is stated for one space; the eigenvector $v \otimes \cdots \otimes v$ and the passage to the limit in k, which were part of that lemma, are now the last two steps of the proof of Theorem 9. The weak Lefschetz property has one definition, stated for the sheaf itself, and Lemma 8 is applied to $\mathcal{F}_k$ and to $\mathcal{F}$ where the first version applied it to their duals. Section 2 has a paragraph on the exposition and begins its sketch with the affine line. The remarks after Lemma 3 and all of Section 9 are new. No step of the proof was removed, and the numbers of the statements are those of the first version. The same model read the revised proof again in full and compared it step by step with the first version. Claude Opus 5.5 at effort max read every change and the four remarks step by step before deposit. No human has checked the revision.

The remarks. Remark 10 is Zhang's formulation; Claude Fable 5.2 at effort max verified it in the generality of Proposition 4. Remark 11 states findings AM-S1 to AM-S3 of a run of GPT-6 Astra (ChatGPT Pro mode) on Alpöge's argument (run ALP_MACHINE, 21 and 22 November 2026), where the question arose for an operator of order two that inverts one of order three. Claude Opus 5.5 at effort high checked their proofs on 22 November 2026, and Claude Fable 5.2 at effort max re-derived them for this text. The first two paragraphs of Remark 12 were written and verified for this text by Claude Fable 5.2 at effort max; they put in the terms of this note the conclusion of finding F50 of a run of GPT-6 Astra (ChatGPT Pro mode) that examined Mulmuley's nonstandard Riemann hypothesis in the light of the sign argument (run GCT_NSRH, 21 November 2026). The third paragraph states the explicit part of F50. Claude Opus 5.5 at effort high checked the proof of F50 on 21 November 2026, and on 24 November recomputed the squarefreeness of the determinants for the odd primes up to 13 and all $k \leq 40$. Claude Fable 5.2 at effort max verified the argument for squarefreeness and the case $k = 2$ of the explicit formula, and did not re-derive the formula for larger k. Remark 13 states finding SN-14 of a run of GPT-6 Astra (ChatGPT Pro mode) on Grothendieck's notes (run SIGN_NEXT, 21 November 2026); the run re-proved it with a second instance and checked the matrix identities by exact computation, Claude Opus 5.5 at effort high checked its proof the same day, and Claude Fable 5.2 at effort max re-derived the statements quoted here. Scan 49 of cote 16 was not reopened for this text; what is said about it is the reading of that run.

The references. The other references were not reopened for this text. The foundational ones are cited with the locators of the run. What the first two sections say about [2], [3], [4], [8] and [10] is from memory. The two statements of [9] and the introduction of [11] were read at the source by Claude Opus 5.5 at effort max. What is said about [7], about Vinogradov's estimate and about Abhyankar's map is taken from [16] and was not checked at the sources; that Abhyankar's map is finite and étale was verified directly. The locator in [13] for Gauss sums is that of the run on F50.

No human has checked this text. The case of the affine line has been written out by a human reader [16].

The prompt, the output of the run and the checks of the proof are public at

`https://github.com/KokunoYumeto/grothendieck-montpellier-archive/tree/main/research/grh_20261120/deligne_sign_only`.

This text was written by Claude (Claude Fable 5.2 at effort max), using an earlier text by Claude Opus 5.5 at effort max. It is dedicated to the public domain (CC0 1.0).

References

[1] L. Alpöge, note on the Riemann hypothesis, 20 November 2026, `https://alpo.ge/remark.pdf`.

[2] E. Bombieri, Counting points on curves over finite fields (d'après S. A. Stepanov), Séminaire Bourbaki 1972/73, Exposé 430, Lecture Notes in Math. 383, Springer, 1974, 234–241.

[3] P. Deligne, La conjecture de Weil. I, Publ. Math. IHÉS 43 (1974), 273–307.

[4] P. Deligne, La conjecture de Weil. II, Publ. Math. IHÉS 52 (1980), 137–252.

[5] J.-P. Jouanolou, Théorèmes de Bertini et applications, Birkhäuser, 1983.

[6] N. M. Katz, Affine cohomological transforms, perversity, and monodromy, J. Amer. Math. Soc. 6 (1993), 149–222.

[7] N. M. Katz, L-functions and monodromy: four lectures on Weil II, Adv. Math. 160 (2001), 81–132.

[8] K. S. Kedlaya, Fourier transforms and p-adic "Weil II", Compos. Math. 142 (2006), 1426–1450.

[9] L. Lafforgue, Chtoucas de Drinfeld et correspondance de Langlands, Invent. Math. 147 (2002), 1–241.

[10] G. Laumon, Transformation de Fourier, constantes d'équations fonctionnelles et conjecture de Weil, Publ. Math. IHÉS 65 (1987), 131–210.

[11] A. J. Scholl, Hypersurfaces and the Weil conjectures, arXiv:0912.2073.

[12] M. Artin, A. Grothendieck, J.-L. Verdier, Théorie des topos et cohomologie étale des schémas (SGA 4), Lecture Notes in Math. 269, 270, 305, Springer, 1972–1973.

[13] P. Deligne, Cohomologie étale (SGA $4\frac{1}{2}$), Lecture Notes in Math. 569, Springer, 1977.

[14] The Stacks Project, `https://stacks.math.columbia.edu/tag/00OY`.

[15] Transcription of A. Grothendieck's papers in the Montpellier archive, `https://github.com/KokunoYumeto/grothendieck-montpellier-archive`.

[16] D. Zhang, On an elementary proof of Weil II, 22 November 2026, arXiv:2611.30749.