Weil Conjectures (Deligne's Purity Theorem)

David Sherman, Ka Yu Tam

June 7, 2017

Let $\kappa = \mathbf{F}_q$ be a finite field of characteristic $p > 0$, and k be a fixed algebraic closure of κ. We fix a prime $\ell \neq p$, and an isomorphism $\tau : \overline{\mathbf{Q}}_\ell \to \mathbf{C}$. Whenever we want to denote something (e.g. scheme, sheaf, morphism, etc.) defined over κ, we will put a subscript $_0$ (e.g. X_0 is a scheme over κ, $\mathcal{F}_0$ is a Weil sheaf defined over X_0, etc.), and when the subscript is dropped, this means the corresponding base change to k (e.g. $X = X_0 \times_\kappa k$, $\mathcal{F}$ is the pullback of $\mathcal{F}_0$ via $X \to X_0$, etc.). For a closed point $a \in |X_0|$, we'll use the subscript $\mathcal{F}_{\overline{a}}$ to denote the stalk of a sheaf $\mathcal{F}$ at a geometric (k)-point $\overline{a}$ over a. This will avoid any confusion, for example, about the stalks at 0 for a sheaf on $\mathbf{A}^1_0$.

1 Goal

Let us recall part of Deligne's purity theorem (Theorem 12 in Lecture 18), which is our goal today:

Theorem 1.1. *Let $f_0 : X_0 \to Y_0$ be a morphism of κ-varieties, and $\mathcal{F}_0$ be a τ-mixed sheaf on X_0 whose largest weight is β. Then, for all i, $R^i(f_0)_!\mathcal{F}_0$ is τ-mixed, and its weights are at most $\beta + i$.*

By Poincaré duality one gets an immediate corollary, which leads to Weil's Riemann hypothesis (Theorem 10 and Corollary 14 in Lecture 18):

Corollary 1.2. *Let $f_0 : X_0 \to Y_0$ be smooth and proper, and $\mathcal{F}_0$ be a τ-pure lisse sheaf on X_0 of weight β. Then, for all i, $R^i(f_0)_*\mathcal{F}_0$ is τ-pure of weight $\beta + i$.*

2 First Reduction

In this section we will reduce Theorem 1.1 to the following special case:

Theorem 2.1. *Let $\mathcal{F}_0$ be a τ-pure lisse sheaf of weight β on a smooth geometrically irreducible curve X_0 over κ. Then, for all i, the weights of $H^i_{c,\text{ét}}(X, \mathcal{F})$ are at most $\beta + i$.*

Now we prove Theorem 1.1 using Theorem 2.1. We remark that we can perform the following reductions:

1. The assertion is trivial if f_0 is quasi-finite: in this case $R^i(f_0)_!\mathcal{F}_0 = 0$ for $i > 0$ (Theorem 3.1 in Lecture 10), and the case $i = 0$ is trivial.

2. We may replace X_0 by any non-empty open $U_0 \subseteq X_0$: if S_0 is the complement of U_0, then the excision sequence (Theorem 2.5(3) in Lecture 10) says it suffices to prove the result for $U_0 \to Y_0$ and $S_0 \to Y_0$, and by Noetherian induction on X_0 the result holds for S_0.

3. If the result holds for g_0 and h_0, then it also holds for $f_0 = g_0 \circ h_0$: this follows from the Leray spectral sequence (Theorem 2.5(2) in Lecture 10).

4. We can replace f_0 by $f_0|_{f_0^{-1}(U_0)} : f_0^{-1}(U_0) \to U_0$ for a non-empty open U_0 in Y_0: if the image of f_0 is not dense in Y_0 then we may replace Y_0 by a closed subscheme (the scheme-theoretic image of f_0) and use Noetherian induction on Y_0. If f_0 has dense image, then $f_0^{-1}(U_0)$ is non-empty, so reduction 2 says it suffices to prove this for $f_0^{-1}(U_0) \to Y_0$. This map factors as $f_0^{-1}(U_0) \to U_0 \to Y_0$, and the theorem holds for $U_0 \to Y_0$ (reduction 1), so it suffices to deal with the case $f_0^{-1}(U_0) \to U_0$ (reduction 3).

We claim that, with these reductions, we may assume that $f_0 : X_0 \to Y_0$ is a surjective affine smooth morphism whose fibers are geometrically irreducible curves, and that $\mathcal{F}_0$ is lisse and τ-pure of weight β. Indeed, by reductions 2 and 4 we may assume that X_0 and Y_0 are affine, so in particular f_0 is affine. If η is a generic point of Y_0, then $(X_0)_\eta \to \operatorname{Spec} \kappa(\eta)$ is affine, so by Noether normalization lemma there is a finite map $(X_0)_\eta \to \mathbf{A}^n_{\kappa(\eta)}$, and we can spread this out to get a finite map $f_0^{-1}(U_0) \to \mathbf{A}^n_{U_0}$ for some open affine $U_0 \subseteq Y_0$. Hence, by replacing Y_0 with U_0 (reduction 4), we may assume that there is a finite map $X_0 \to \mathbf{A}^n_{Y_0}$. Using reductions 1 and 3 we reduce to the case $X_0 = \mathbf{A}^1_{Y_0} \to Y_0$ where Y_0 is affine. Now, there exists a stratification of X_0 such that $\mathcal{F}_0$ is lisse when restricted to this stratification (proved in Lecture 16). In particular, there is a non-empty affine open $U_0 \subseteq \mathbf{A}^1_{Y_0}$ for which $\mathcal{F}_0|_{U_0}$ is lisse, so by replacing X_0 with U_0 (reduction 2) we may assume that $\mathcal{F}_0$ is lisse, X_0 is open in $\mathbf{A}^1_{Y_0}$, and f_0 is smooth affine. In particular, the image $f_0(X_0)$ is open in Y_0. Replacing Y_0 with this image (reduction 4) we can assume that f_0 is surjective smooth affine. Note that X_0 is still an open subset of $\mathbf{A}^1_{Y_0}$, so the fibers are non-empty open in $\mathbf{A}^1$ and hence are geometrically irreducible curves. By passing to the irreducible subquotients of $\mathcal{F}_0$ using the long exact sequence for the cohomological δ-functor $R^\bullet(f_0)_!$, we may assume $\mathcal{F}_0$ is irreducible and hence τ-pure.

Let $\overline{y}$ be a geometric point of Y_0, with underlying point y, and let $C_0 = f_0^{-1}(y)$. Note that the above reductions guarantee that C_0 is an affine smooth geometrically irreducible curve. The estimates for the weights clearly follow from Theorem 2.1, which says that the weights of $\left(R^i(f_0)_!\mathcal{F}_0\right)_{\overline{y}} \simeq H^i_{c,\text{ét}}(C, \mathcal{F}|_C)$ are at most $\beta + i$.

For τ-mixedness, it suffices to show that $R^i(f_0)_!\mathcal{G}_0$ is τ-mixed where $\mathcal{G}_0 = \mathcal{F}_0 \oplus \overline{\mathcal{F}}_0$. Since we know that τ-real sheaves are τ-mixed (Lecture 22), it suffices to show that $R^i(f_0)_!\mathcal{G}_0$ is τ-real, so we want to show that $\tau \det\left(1 - tF | H^i_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$ has real coefficients. Note that $\mathcal{G}_0$ is τ-real and τ-pure of weight β. Note also that, since C is affine, $H^0_{c,\text{ét}}(C, \mathcal{G}|_C) = 0$, so the Grothendieck-Lefschetz trace formula says

$$\frac{\tau \det\left(1 - tF | H^1_{c,\text{ét}}(C, \mathcal{G}|_C)\right)}{\tau \det\left(1 - tF | H^2_{c,\text{ét}}(C, \mathcal{G}|_C)\right)} = \tau L(C_0, \mathcal{G}_0, t) \tag{2.1}$$

has real coefficients because $\mathcal{G}_0$ is by assumption τ-real. Using Poincaré duality together with the assumption that $\mathcal{G}_0$ is τ-pure of weight β, one sees that every root α of $\tau \det\left(1 - tF | H^2_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$ satisfies

$$|\alpha| = N(y)^{-(\beta+2)/2}.$$

On the other hand, the weight estimate above implies that every root α' of the polynomial $\tau \det\left(1 - tF | H^1_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$ satisfies

$$|\alpha'| \geq N(y)^{-(\beta+1)/2},$$

so for every root α of $\tau \det\left(1 - tF | H^2_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$, neither α nor its complex conjugate $\overline{\alpha}$ is a root of $\tau \det\left(1 - tF | H^1_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$. Essentially, the left-hand side of Equation (2.1) is in "lowest terms." Equation (2.1) now shows that both of $\tau \det\left(1 - tF | H^1_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$ and $\tau \det\left(1 - tF | H^2_{c,\text{ét}}(C, \mathcal{G}|_C)\right)$ have real coefficients, since these two polynomials each have leading coefficient 1.

3 Second Reduction

In this section we will further reduce Theorem 2.1 to

Theorem 3.1. *Let X_0 be a smooth, geometrically irreducible projective curve over κ, and let $j_0 : U_0 \to X_0$ be a non-empty open subscheme. Let $\mathcal{F}_0$ be a τ-pure lisse sheaf of weight β on U_0. Then, for all i, $H^i_{\text{ét}}(X, j_*\mathcal{F})$ is τ-pure of weight $\beta + i$.*

To perform the reduction, let $\overline{X}_0$ be the regular completion of X_0 (which is smooth since κ is perfect), and $j_0 : X_0 \to \overline{X}_0$ be the inclusion. The reduction is then obvious from the following lemma (where we take U_0 to be X_0 and X_0 to be $\overline{X}_0$):

Lemma 3.2. *Let X_0 be a smooth, geometrically irreducible projective curve over κ, and let $j_0 : U_0 \to X_0$ be a non-empty open subscheme. Let $\mathcal{F}_0$ be a τ-pure lisse sheaf of weight β on U_0. Then, for all i, the following are equivalent:*

(a) *the weights of* $H^i_{\text{ét}}(X, j_*\mathcal{F})$ *are at most* $\beta + i$.

(b) *the weights of* $H^i_{c,\text{ét}}(U, \mathcal{F})$ *are at most* $\beta + i$.

Proof of Lemma 3.2. Let $\mathcal{H}_0 = (j_0)_*\mathcal{F}_0/(j_0)_!\mathcal{F}_0$, so it is concentrated on the finite complement S_0 of U_0 in X_0. Applying the cohomological δ-functor $H^\bullet_{\text{ét}}(X, -)$ to the exact sequence

$$0 \longrightarrow (j_0)_!\mathcal{F}_0 \longrightarrow (j_0)_*\mathcal{F}_0 \longrightarrow \mathcal{H}_0 \longrightarrow 0$$

gives us an exact sequence

$$H^{i-1}_{\text{ét}}(X, \mathcal{H}) \longrightarrow H^i_{c,\text{ét}}(U, \mathcal{F}) \longrightarrow H^i_{\text{ét}}(X, j_*\mathcal{F}) \longrightarrow H^i_{\text{ét}}(X, \mathcal{H}) \tag{3.1}$$

Since $\mathcal{H}_0$ is concentrated on the finite set S_0, $H^i_{\text{ét}}(X, \mathcal{H}) = 0$ for $i > 0$. Note that

$$H^0_{\text{ét}}(X, \mathcal{H}) \simeq \bigoplus_{s \in S_0} ((j_0)_*\mathcal{F}_0)_{\overline{s}}.$$

Since $(j_0)_*\mathcal{F}_0$ is lisse, by semi-continuity of weights (Corollary 2.3.0.1 in Lecture 19) we see that the weights of $\mathcal{H}_0$ are at most β. The equivalence is now immediate from the exact sequence (3.1). □

4 Third Reduction

Analogously to the deduction of Corollary 1.2 from Theorem 1.1 (see Lecture 18), we will now use Poincaré duality to reduce Theorem 3.1 to an upper bound statement on the weights in first cohomology:

Theorem 4.1. *In the situation of Theorem 3.1, let* α *be an eigenvalue of geometric Frobenius*

$$F : \mathrm{H}^1(X, j_*\mathcal{F}) \to \mathrm{H}^1(X, j_*\mathcal{F}).$$

Then

$$|\tau(\alpha)|^2 \le q^{\beta+1}.$$

Before proceeding further, we may assume $\mathcal{F}_0$ is actually a $\overline{\mathbf{Q}}_\ell$-sheaf on U_0 (rather than just a Weil sheaf). Indeed, by considering the determinant weights, one can twist away by a quasicharacter χ_b of $W(k/\kappa)$ to make $\mathcal{F}_0$ a $\overline{\mathbf{Q}}_\ell$-sheaf (see Corollary 1.1.14 in Lecture 19). Since such quasicharacters are geometrically trivial, they pass straight through the whole computation.

Now, because $\mathcal{F}_0$ begins life as a lisse sheaf on U_0, not on all of X_0, we need a local duality result of Deligne before we can apply Poincaré duality. Essentially, we want to identify the dual of the pushforward with the pushforward of the dual. In this section, we'll be careful to deduce what we need from statements on finite level without needlessly invoking the derived category of ℓ-adic sheaves (although this will be unavoidable when we use the ℓ-adic Fourier transform in the next section).

Theorem 4.2 (Deligne, [1], Exposé [Dualité], Théorème 1.3)**.** *Let X be a regular $\mathbf{Z}[1/n]$-scheme of pure dimension 1, $j : U \to X$ be an open dense subset, and $\mathcal{G}$ be an locally constant constructible sheaf of $\mathbf{Z}/n$-modules on U. Then*

$$\mathrm{R}\mathcal{H}om(j_*\mathcal{G}, \mathbf{Z}/n) \simeq j_*\mathcal{H}om(\mathcal{G}, \mathbf{Z}/n) =: j_*\mathcal{G}^\vee$$

inside the (bounded, constructible) derived category of $\mathbf{Z}/n$-sheaves on X.

Deligne gives a brief proof of this theorem in the above reference, so we won't give one here. The main ideas are: (1) one can first reduce to the case of $X = \operatorname{Spec} R$ for R a DVR with residue characteristic p prime to n, (2) in this case, one can consider instead of our representation the invariants of wild inertia. Then the desired vanishing of higher $\mathcal{E}xt$ groups follows from computing the (co)-invariants of tame inertia. We also remark that the theorem remains valid if we replace $\mathbf{Z}/n$ by $\Lambda_k = \Lambda/\mathfrak{m}^{k+1}$, where Λ is a DVR of residue characteristic ℓ (and now X is a $\mathbf{Z}[1/\ell]$-scheme, of course).

Let's put Theorem 4.2 together with Poincaré duality (see Lectures 12 and 13) in our smooth projective situation.

Lemma 4.3. *Let X be a smooth projective curve over an algebraically closed field k of characteristic p (prime to ℓ). Let $j : U \to X$ be an open dense subset, and $\mathcal{G}$ be an lcc sheaf of flat Λ_k-modules on U. Then the cup product induces a perfect pairing*

$$\mathrm{H}^i_{\acute{e}t}(X, j_*\mathcal{G}^\vee(1)) \otimes \mathrm{H}^{2-i}_{\acute{e}t}(X, j_*\mathcal{G}) \to \mathrm{H}^2_{\acute{e}t}(X, \Lambda_k(1)) \simeq \Lambda_k.$$

Proof. (sketch) As in the proof of Poincaré duality itself, one can obtain an abstract isomorphism by using the Ext version of Poincaré duality and plugging in Theorem 4.2 to the local-to-global Ext spectral sequence. But for functoriality purposes (e.g. taking a projective limit, as we're about to do), one needs to do a computation and be sure that the cup product induces this isomorphism.

Luckily for us, the key computation already appears in the lcc case. That is, similarly to the proof of Poincaré duality itself, one can reduce to the case of the constant sheaf $\mathcal{G} = (\underline{\Lambda_k})_U$. But then $j_*\mathcal{G} = (\underline{\Lambda_k})_X$, which is lcc on X. Thus, the cup product computation for usual Poincaré duality for the constant sheaf Λ_k on X gives us the result we want.[1] □

Given Lemma 4.3, we can repeat the proof of the Λ-adic version of Poincaré duality (see Lecture 16) to obtain the following

Corollary 4.4. *Let X_0 be a smooth projective curve over κ, $j_0 : U_0 \to X_0$ be an open dense subset, and $\mathcal{F}_0$ a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U_0. Then the cup product induces a perfect pairing*

$$\mathrm{H}^i_{\text{ét}}(X, j_*\mathcal{F}^\vee(1)) \otimes \mathrm{H}^{2-i}_{\text{ét}}(X, j_*\mathcal{F}) \to \mathrm{H}^2_{\text{ét}}(X, \overline{\mathbf{Q}}_\ell(1)) \simeq \overline{\mathbf{Q}}_\ell.$$

[1] Note that this is the key case, especially in degree $i = 1$, for the proof of Poincaré duality, see Lectures 12 and 13.

We're now ready to prove our reduction.

Proof. (of Theorem 3.1, given Theorem 4.1) As explained above, we may assume that $\mathcal{F}_0$ is a lisse $\overline{\mathbf{Q}}_\ell$-sheaf on U_0. For $i > 2$, the cohomology vanishes, so the claim is true vacuously. For $i = 0$, we have

$$\mathrm{H}^0_{\text{ét}}(X, j_*\mathcal{F}) = \Gamma(X, j_*\mathcal{F}) = \Gamma(U, \mathcal{F}).$$

But certainly an eigenvalue of the latter must be an eigenvalue of each stalk $\mathcal{F}_{\overline{x}} = \mathcal{F}_{0,\overline{x}}$. Then $\mathrm{H}^0_{\text{ét}}(X, j_*\mathcal{F})$ inherits τ-purity with weight β from $\mathcal{F}_0$.

By Corollary 4.4 in the case $i = 0$, we have

$$\mathrm{H}^2_{\text{ét}}(X, j_*\mathcal{F}) = \left(\mathrm{H}^0_{\text{ét}}(X, j_*\mathcal{F}^\vee(1))\right)^\vee.$$

Since $\mathcal{F}_0$ is τ-pure of weight β, its dual is τ-pure of weight $-\beta$. Thus, applying the $i = 0$ case of the theorem, which we just obtained, $\mathrm{H}^2_{\text{ét}}(X, j_*\mathcal{F})$ is τ-pure of weight $-(-\beta - 2) = \beta + 2$.

Similarly, Corollary 4.4 in the case $i = 1$ says that

$$\mathrm{H}^1_{\text{ét}}(X, j_*\mathcal{F}) = \left(\mathrm{H}^1_{\text{ét}}(X, j_*\mathcal{F}^\vee(1))\right)^\vee.$$

Apply Theorem 4.1 to $\mathcal{F}_0^\vee(1)$. Then the eigenvalues α of geometric Frobenius on the left-hand side satisfy

$$|\tau(\alpha)|^2 \leq q^{-(-\beta-2+1)} = q^{\beta+1}.$$

On the other hand, we get the opposite inequality by applying Theorem 4.1 to $\mathcal{F}_0$. Hence $\mathrm{H}^1_{\text{ét}}(X, j_*\mathcal{F})$ is τ-pure of weight $\beta + 1$. □

5 ℓ-adic Fourier Transform and Conclusion

In this section, we'll prove Theorem 4.1 by reducing to the case $X_0 = \mathbf{P}^1_0$ and using the ℓ-adic Fourier transform. First, let's recall some facts about the ℓ-adic Fourier transform. The usual caveats about our cavalier treatment of the "derived category" of ℓ-adic sheaves apply (see, for example, Section II.5 of [2] or Section §4.5-4.8 of Exposé [Rapport] of [1] for details).

Let $\psi : \kappa \to \overline{\mathbf{Q}}_\ell^*$ be a character, and define $\psi_x : y \mapsto \psi(xy)$. Then we have the corresponding rank 1 sheaf $\mathcal{L}_0(\psi)$ constructed using the Artin-Schreier covering of $\mathbf{A}^1_0$ (see Lecture 23 for details). If $\mathcal{F}_0$ is an $\overline{\mathbf{Q}}_\ell$-sheaf (or even an object in the derived category) on $\mathbf{A}^1_0$, then we have the ℓ-adic Fourier transform

$$\mathrm{FT}_\psi(\mathcal{F}_0) := \mathrm{R}\pi^1_!\left(\pi^{2*}_0(\mathcal{F}_0) \otimes m_0^*\mathcal{L}_0(\psi)\right)[1],$$

where π^1_0, π^2_0 are the two projections $\mathbf{A}^1_0 \times \mathbf{A}^1_0 \to \mathbf{A}^1_0$ and m_0 is the multiplication map. Also recall that $\mathrm{FT}_\psi(\mathcal{F}_0)$ has the following properties:

1. (Geometric stalks) If $a \in \mathbf{A}^1(k)$, then

$$(\mathrm{FT}_\psi(\mathcal{F}_0))_{\overline{a}} = \mathrm{R}\Gamma_c(\mathcal{F} \otimes \mathcal{L}(\psi_a))[1].$$

   In particular, if $\mathcal{F}_0$ is a sheaf, then the geometric stalk above $0 \in \mathbf{A}^1$ of its Fourier transform is $\mathrm{H}^1_c(\mathbf{A}^1, \mathcal{F})$. Or, under the sheaf-function correspondence, if κ_n is the degree n extension of κ, then[2]

$$f_{\mathrm{FT}_\psi(\mathcal{F}_0)}(0) = -\sum_{x \in \kappa_n} f_{\mathcal{F}_0}(x).$$

2. (Fourier inversion)

$$(\mathrm{FT}_{\psi^{-1}} \circ \mathrm{FT}_\psi)(\mathcal{F}_0) = \mathcal{F}_0(-1).$$

3. (Plancherel formula)

$$\|f_{\mathrm{FT}_\psi(\mathcal{F}_0)}\|_n = q^{n/2} \|f_{\mathcal{F}_0}\|_n,$$

   where we define the L^2 norm as in Lecture 21 (or Section I.2.12 of [2]).

Since we didn't give full proofs of these facts in, we mention something about their proofs now (see Section I.5 of [2] for more). The Plancherel formula here is exactly the Plancherel formula for functions on finite groups (with proof as usual using the orthogonality relations and switching the order of summation). To prove 1 and 2, we mimic the proofs for finite groups, although we again have to make sheaf-theoretic constructions instead (as in the construction of the Fourier transform itself). In particular, we prove a sheaf-theoretic version of the orthogonality relations as the key lemma. One last remark: the (inverse) Tate twist in Fourier inversion and the factor $q^{n/2}$ in the Plancherel formula are consequences of our normalization; we have no good way to "divide by $\sqrt{q}$" on the sheaf side of the correspondence.

Proof. (of Theorem 4.1)

More reductions: The first part of the proof consists of getting into a situation where we can use the ℓ-adic Fourier transform. For example, we'd like to reduce to the case $X_0 = \mathbf{P}^1_0$. If $j_0' : U_0' \to U_0$ is a smaller open set, then we have the projection formula $(j_0 \circ j_0')_*(j_0')^*(\mathcal{F}_0) = j_{0*}(\mathcal{F}_0)$. This means we can replace U_0 and $\mathcal{F}_0$ by U_0' and $(j_0')^*(\mathcal{F}_0)$, respectively (here we've used "Permanence Property (1)" of weights under pullback in Section I.1.2 of [2]). In particular, we may assume U_0 is affine. Then, by Noether normalization, U_0 is finite over the affine space $\mathbf{A}^1_0$. Since X_0 is a curve, we can extend this finite map $U_0 \to \mathbf{P}^1_0$ over the finitely many missing points, obtaining a finite morphism $X_0 \to \mathbf{P}^1_0$.

By "Permanence Property (2)" in Section I.1.2 of [2], the pushforward of $\mathcal{F}_0$ to $\mathrm{im}(U_0) \subset \mathbf{A}^1_0$ is lisse and τ-pure of weight β. Especially, the theorem is insensitive to finite morphisms. Thus, we may now assume $X_0 = \mathbf{P}^1_0$.

[2] The minus sign here won't matter for us, but it's an artifact of our choice of Artin-Schreier cover $\{y^q - y - x = 0\} \subset \mathbf{A}^2_0$ (versus, say $\{y^q - y + x = 0\}$).

Next, since the complement $\mathbf{P}^1_0 \smallsetminus U_0$ is finite, we have the surjection

$$\mathrm{H}^1_c(U, \mathcal{F}) = \mathrm{H}^1(\mathbf{P}^1, j_! \mathcal{F}) \twoheadrightarrow \mathrm{H}^1(\mathbf{P}^1, j_* \mathcal{F})$$

coming from the long exact sequence in cohomology. Thus, it's enough to prove an upper bound of Frobenius eigenvalues on $\mathrm{H}^1_c(U, \mathcal{F})$.

Our next goal is to reduce to the case where the lisse sheaf $\mathcal{F}_0$ is (i) geometrically irreducible, (ii) geometrically non-constant, and (iii) unramified at ∞. If ρ is the representation of $\pi_1(U_0, \overline{a})$ (for $\overline{a}$ a geometric point of U_0) corresponding to $\mathcal{F}_0$, then (i,ii) mean that $\rho|_{\pi_1(U,\overline{a})}$ is irreducible and non-trivial, respectively. Also, (iii) means that $\rho|_{I_\infty}$ is trivial. Equivalently, $\mathcal{F}_0$ extends to a lisse sheaf on $U \cup \{\infty\}$.

Note that by "Permanence Property (3)" in Section I.1.2 of [2], τ-purity is preserved (along with the weight β) under base change to a finite extension of κ. Thus, for the purposes of this theorem we're free to make such a base change at any time.

Since $\mathrm{H}^1_c(U, -)$ is half-exact, it suffices to prove our theorem on the irreducible constituents of $\mathcal{F}_0$. There is some finite extension κ' of κ such that all constituents of $(\mathcal{F}_0)'$ all are even geometrically irreducible, allowing us to assume (i) from now on.

If we make another extension of κ, then U_0 necessarily contains some rational point s. Then $\mathcal{F}|_{U_0 \smallsetminus \{s\}}$ is certainly unramified at s. Thus, delete s from U_0 and make a coordinate change to move s to ∞. We're now in case (iii).

To get into the geometrically non-constant case (ii), we just prove the theorem directly in the case that $\mathcal{F}_0$ is geometrically constant. Recall that $\mathbf{P}^1$ has no non-trivial connected finite étale covers. In particular, if $\mathcal{F}$ is constant, then $\mathrm{H}^1(\mathbf{P}^1, j_*\mathcal{F}) = 0$. Thus, by looking at the long exact sequence relating the cohomology of $j_!\mathcal{F}$ and $j_*\mathcal{F}$, we have

$$\prod_{s \in \mathbf{P}^1 \smallsetminus U} (j_{0*}\mathcal{F}_0)_{\overline{s}} \twoheadrightarrow \mathrm{H}^1_c(U, \mathcal{F}).$$

Hence, the weights β of $\mathcal{F}_0$ bound the weights of $\mathrm{H}^1_c(U, \mathcal{F})$ in this case.

Fourier transform: We've reduced to the case where $X_0 = \mathbf{P}^1_0$, and $\mathcal{F}_0$ satisfies properties (i,ii,iii). So we're finally ready to use the ℓ-adic Fourier transform. As described in Akshay's lecture (Lecture 23), one should think of $\mathrm{FT}_\psi(j_!\mathcal{F}_0)$ like a function whose value at 0 is the value we care about (the eigenvalues on $\mathrm{H}^1_c(\mathbf{P}^1, \mathcal{F})$). Then we bound our value by controlling the L^2 norm of our function. Accordingly, the key work for us left to do is showing that the (upper) weights of $\mathrm{FT}_\psi(j_!\mathcal{F}_0)$ are indeed determined by its L^2 norm. We'll accomplish in part by showing that $\mathrm{FT}_\psi(j_!\mathcal{F}_0)$ is in fact a τ-mixed sheaf.

Pursuant to this strategy, we make three claims about $\widehat{\mathcal{F}}_0 := \mathrm{FT}_\psi(j_!\mathcal{F}_0)$: (a) $\widehat{\mathcal{F}}_0$ is an actual sheaf (i.e. it's supported only in degree zero), (b) $\widehat{\mathcal{F}}_0$ has no compactly-supported sections, and (c) $\widehat{\mathcal{F}}_0$ is τ-mixed.

Proof of (a): We may prove this stalk-wise. Using our formula for the geometric stalks of $\widehat{\mathcal{F}}_0$, we want to show that the twists $\mathcal{F} \otimes \mathcal{L}(\psi_x)$ have vanishing compactly supported cohomology except in degree 1. But these are lisse sheaves on the affine U, so $\mathrm{H}^0_c(U, \mathcal{F} \otimes \mathcal{L}(\psi_x))$ necessarily vanishes for each x. On the other hand, by Poincaré duality, we have

$$\mathrm{H}^2_c(U, \mathcal{F} \otimes \mathcal{L}(\psi_x)) = (\rho \otimes \psi_x)_{\pi_1(U,\overline{a})}(-1),$$

where the subscript denotes co-invariants. But ρ (hence $\rho \otimes \psi_x$) is geometrically irreducible, so the right-hand side must either be zero (as we're hoping) or all of $\rho \otimes \psi_x$. In the latter case, $\rho \otimes \psi_x$ is geometrically trivial. For $x = 0$ this contradicts property (ii) of ρ. For $x \neq 0$, property (iii) then implies that ψ_x is unramified at ∞ (this is a condition on the inertia I_∞, which is contained in the geometric $\pi_1(U, \overline{a})$). But ψ_x came from the Artin-Schreier cover of $\mathbf{A}^1_0$; since this cover is geometrically irreducible, it can't be extended to a connected étale cover of $\mathbf{P}^1_0$ (which has no geometric covers at all!).

Proof of (c): Here we reuse the trick that τ-real sheaves are τ-mixed (Lecture 22). When we wanted to use this trick above for the pushforward of a τ-mixed sheaf, we essentially used the following idea. If f is a complex-valued function, then $f + \overline{f}$ is real-valued. As a construction of sheaves, we wrote that $\mathcal{F}_0$ was a direct summand of the τ-real sheaf $\mathcal{F}_0 \oplus \mathcal{F}_0^\vee \otimes \mathcal{L}_b$ (for an appropriate choice of $b \in \overline{\mathbf{Q}}_\ell^*$). Here, we must modify the idea slightly. If f' denotes the inverse Fourier transform, then $\widehat{f} + \overline{f}'$ is real-valued.

Then using property (a) together with the Grothendieck-Lefschetz trace formula (to prove that $\mathrm{R}\pi^1_!$ of our τ-real sheaf remains τ-real), we see that $\mathrm{FT}_\psi(j_!\mathcal{F}_0)$ is a direct summand of the τ-real sheaf

$$\mathrm{FT}_\psi(j_!\mathcal{F}_0) \oplus \mathrm{FT}_{\psi^{-1}}(j_!\mathcal{F}_0^\vee) \otimes \mathcal{L}_b,$$

hence it's τ-mixed.

Proof of (b): This is a straightforward application of the Fourier inversion formula. First, using our formula for geometric stalks,

$$\mathrm{H}^0_c(\mathbf{A}^1, \widehat{\mathcal{F}}) = \mathcal{H}^{-1}(\mathrm{FT}_{\psi^{-1}}(\widehat{\mathcal{F}}))_{\overline{0}},$$

where the $\mathcal{H}^{-1}$ denotes the degree -1 cohomology of the complex. But then

$$\begin{aligned}
\mathrm{H}^0_c(\mathbf{A}^1, \widehat{\mathcal{F}}) &= \mathcal{H}^{-1}(\mathrm{FT}_{\psi^{-1}} \circ \mathrm{FT}_\psi(j_!\mathcal{F}))_{\overline{0}} \\
&= \mathcal{H}^{-1}(j_!\mathcal{F}(-1))_{\overline{0}} \\
&= 0,
\end{aligned}$$

because $j_!\mathcal{F}(-1)$ is supported in degree zero.

Now that we have (a,b,c), we're ready to conclude. We know that $(\widehat{\mathcal{F}}_0)_{\overline{0}} = \mathrm{H}^1_c(U, \mathcal{F})$. Thus, the upper weights of $\widehat{\mathcal{F}}_0$ bound the eigenvalues of geometric Frobenius on $\mathrm{H}^1_c(U, \mathcal{F})$. But $\widehat{\mathcal{F}}_0$ is τ-mixed and has no compactly-supported sections, so its upper weights are the same as the L^2-norm of the associated function (as we saw using the computations of the radius of convergence of the L-function in Lecture 21). And the Plancherel formula implies the following equality of upper weights: $\mathrm{wt}(\widehat{\mathcal{F}}_0) = \mathrm{wt}(\mathcal{F}_0) + 1$. □

References

[1] P. Deligne, et al. *Cohomologie étale: Séminaire de Géométrie Algébrique du Bois-Marie SGA $4\frac{1}{2}$*. Springer-Verlag, Berlin, 1977.

[2] R. Kiehl & R. Weissauer. *Weil Conjectures, Perverse Sheaves and ℓ-adic Fourier Transform.* Springer-Verlag, Berlin, 2001.

[3] Multiple Authors. *Weil II Notes, Stanford Number Theory Learning Seminar.* `http://math.stanford.edu/~conrad/Weil2seminar/`