Étale cohomology and Galois Representations

Tom Lovering

June 22, 2012

Abstract

In this essay we briefly introduce the main ideas behind the theory of studying algebraic varieties over a number field by constructing associated Galois representations, and see how this can be understood naturally in the context of an extension of monodromy theory from geometry. We then, following Deligne's original method, use some of these ideas to prove the Riemann Hypothesis for varieties over a finite field, and make some remarks about the implications this result has for the general theory of Galois representations.

Contents

1 Introduction

In classical algebraic topology, a good general philosophy is that if one understands the combinatorics of covers of a space, then one can extract some information about the space itself. In particular, given a sufficiently nice base space, there exists a fundamental cover, which somehow also encodes all the intermediate connected covers and so renders the subject conceptually reasonably 'simple': the existence of a discrete fundamental group whose subgroups correspond to connected covers (all of which are subcovers of the universal cover).

However, in practice these groups can be difficult to compute or work with directly. One solution is to work with their *abelianisations*, shifting the emphasis from the fundamental group towards cohomology with constant coefficients, where in some loose sense only the covers with abelian automorphism group are really under consideration. This viewpoint was later seen to be overly restrictive, and with the theory of sheaves came the possibility of studying wider families of covers: those whose fibres can be viewed as possessing any nice algebraic structure. For example, the theory of local systems (locally constant sheaves of finite dimensional vector spaces), which arose naturally from studying the solutions to linear differential equations, allows one to consider all covers whose automorphism group is *linear* group of matrices acting on a finite dimensional vector space.

In this essay, we shall see how these ideas from classical topology give rise to powerful ideas in arithmetic algebraic geometry. The theory here is initially made difficult by the fact that there is no longer any kind of universal cover. Indeed, consider the classical covers of $\mathbb{C}^*$. The universal cover is given by the exponential map $\mathbb{C} \to \mathbb{C}^*$, which has automorphism group $\mathbb{Z}$: so all the other covers are precisely those given by nth power maps $\mathbb{C}^* \to \mathbb{C}^*$. While the nth power maps can be given by polynomial equations $t \mapsto t^n$, the exponential map certainly cannot. In such a way, we see that while the scheme $(\mathbb{G}_m)_{/\mathbb{C}} :=$ Spec $\mathbb{C}[x, x^{-1}]$ possesses a collection of finite algebraic covers which suggest that it ought to have a fundamental cover with automorphism group $\mathbb{Z}$, this cover itself is absent from the category of schemes. To fix this problem, we will often find ourselves taking various inverse limits over families of finite covers, passing to single *topological* (profinite) groups which fairly naively glue together all the information from the individual finite covers.

However, in spite of this caveat, the theory is a very powerful context for studying schemes over number fields because it is sensitive to the Galois action. For example, in the Kummer exact sequence (which replaces the classical exponential exact sequence for the reasons explained above)

$$0 \to (\mu_m)_{/\mathbb{Q}} \to (\mathbb{G}_m)_{/\mathbb{Q}} \stackrel{t \mapsto t^m}{\to} (\mathbb{G}_m)_{/\mathbb{Q}} \to 0$$

where $(\mu_m)_{/\mathbb{Q}} :=$ Spec $\mathbb{Q}[x]/(x^m - 1)$ is the scheme of mth roots of unity, any element of the absolute Galois group $G_\mathbb{Q}$ of $\mathbb{Q}$ gives an automorphism of the exact sequence on geometric points (for some fixed $\mathbb{Q} \to \bar{\mathbb{Q}}$)

$$0 \to \mu_m \to \bar{\mathbb{Q}}^* \stackrel{t \mapsto t^m}{\to} \bar{\mathbb{Q}}^* \to 0.$$

In particular, we can read off the action on the fibre over 1 as $G_{\mathbb{Q}} \to Aut(\mu_m)$. Choosing $m = l^k$ for l some fixed prime, and k varying over all positive integers, and making (noncanonical, but a compatible system of) identifications $\mathbb{Z}/m\mathbb{Z} \cong \mu_m$, these combine (taking the inverse limit) to give a map $\chi^{cyc} : G_{\mathbb{Q}} \to \mathbb{Z}_l^*$. Tensoring up to $\mathbb{Q}_l$, this gives a 1-dimensional Galois representation, the so-called *l-adic cyclotomic character*. Thus, studying a family of covers of $(\mathbb{G}_m)_{/\mathbb{Q}}$ gave us a natural representation of $G_{\mathbb{Q}}$, which is in fact well-defined (one can check that making a different choice of identifications $\mathbb{Z}/m\mathbb{Z} \cong \mu_m$ has no effect on the Galois action).

A very similar example arises in the study of elliptic curves E over a field K. What are the finite covers of an elliptic curve? By a Riemann-Hurwitz argument, any finite cover of an elliptic curve is itself an elliptic curve, and after choosing an appropriate basepoint can be viewed as an isogeny $E' \to E$. But there is a *dual isogeny* $E \to E'$ easily obtained by pullback of line bundles via the canonical identification of an elliptic curve with its own Picard variety, and this has the property that $E \to E' \to E$ is a multiplication-by-m map for some m. Thus in some sense elliptic curves are 'finitely uniformised' by themselves.

As with the Kummer exact sequence, we get for each m, an exact sequence on geometric points:

$$0 \to E[m] \to E(\bar{K}) \to E(\bar{K}) \to 0,$$

where the *torsion points* $E[m]$ are known to be a finite abelian group, non-canonically isomorphic to $\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$. But again, there is an action of G_K on $E[m]$, and now taking an appropriate profinite limit gives the *Tate module* $T_l(E)$, a rank 2 free $\mathbb{Z}_l$-module with a continuous linear G_K-action. Tensoring up to $\mathbb{Q}_l$ recovers a two-dimensional Galois representation (well-defined up to conjugacy)

$$G_K \to GL_2(\mathbb{Q}_l).$$

This Galois representation turns out to have remarkable properties. Indeed, in [18, III.7] it is noted that provided $l \neq charK$, then (where the RHS is homomorphisms of $\mathbb{Z}_l$ modules which are compatible with the G_K action),

$$Hom_K(E_1, E_2) \otimes_{\mathbb{Z}} \mathbb{Z}_l \to Hom_{G_K}(T_l(E_1), T_l(E_2))$$

is injective. In other words, a map between elliptic curves over K is uniquely determined by the map it induces on Tate modules!

Furthermore, if K is a finite field, Tate showed this is an isomorphism, and some deep work of Faltings shows that this is actually an isomorphism of K is a number field. In fact, in both cases one can deduce that if two elliptic curves give isomorphic Galois representations, there is an isogeny between them. Compare this with classical algebraic topology, where there is absolutely no guarantee that spaces with the same homotopy or homology groups will be in any way 'the same'. By working in the algebraic category and keeping track of the Galois action (for finite fields or number fields), we have an exact dictionary between a certain class of geometrical objects (elliptic curves modulo isogeny) and a certain class of linear-algebraic objects (2-dimensional l-adic Galois representations obtained by constructing Tate modules). Proving results via such correspondences seems to be a theme of modern number theory: if you want to show that every elliptic curve over $\mathbb{Q}$ is modular, the easiest way seems to be to show its Tate module must be a Galois representation cut out by a modular form.

If $K = \mathbb{F}_q$ is a finite field with q elements (and l does not divide q), these representations yield up geometrical data surprisingly explicitly. Recall that G_K will in this case be procyclic with canonical generator given by the *geometric Frobenius* $\phi = (x \mapsto x^q)^{-1}$, so a continuous Galois representation is just a vector space together with a Frobenius automorphism. It is not that difficult to prove (see [18, V]) that if α_1, α_2 are the eigenvalues of the Frobenius automorphism on the Tate module of E,

$$|E(\mathbb{F}_{q^r})| = 1 + q^r - \alpha_1^r - \alpha_2^r.$$

and in fact that $|\alpha_1| = |\alpha_2| = \sqrt{q}$, giving a reasonably tight bound on how far from $1 + q^r$ the number of $\mathbb{F}_{q^r}$-rational points can get.

The most natural way to prove statements like the above is to package all the data $|E(\mathbb{F}_{q^r})|$ into a kind of generating function called a *Zeta function*. We define

$$t\frac{d}{dt}log\ \zeta(E,t) = \sum_{r\geq 1} |E(\mathbb{F}_{q^r})|t^r.$$

It is possible to prove that this formal power series is in fact a rational function, given by

$$\zeta(E,t) = \frac{(1-\alpha_1 t)(1-\alpha_2 t)}{(1-t)(1-qt)}.$$

Observe that the numerator is simply the (inverse) characteristic polynomial of the Frobenius morphism of the Tate module (and the denominators can be interpreted as representing the one dimensional cohomology groups H^0 and H^2, as we shall see shortly). Much of this theory admits easy generalisations to arbitrary abelian varieties. However, it is unclear how to generate Galois representations in cases where there is no underlying abelian group structure.

In the first half of the essay we shall explore this question andsee how similar precise statements about zeta functions of general algebraic varieties over finite fields can be made. In particular, we shall realise the Galois representations obtained above as special cases of those naturally arising in the context of l-adic local systems, which are formed by manufacturing inverse systems of covering spaces whose fibres are artificially equipped with the structure of finitely generated $\mathbb{Z}/l^t\mathbb{Z}$-modules. By some foundational results from SGA, we will be able to start with the trivial covers $X \times \mathbb{Z}/l^t\mathbb{Z}$ of a projective variety X over k and pass to a finite family of nontrivial but still finitely generated covers of Spec k, which under the inverse limit will give an analagous theory to the classical theory of cohomology with $\mathbb{Z}$ coefficients, and give a conceptual explanation for the rational form of the Zeta function above (generalised to any variety).

In the second half of the essay we follow Deligne's original argument [4] to prove a vast generalisation of the above statement that $|\alpha_1| = |\alpha_2| = \sqrt{q}$, the so-called Riemann Hypothesis for varieties over finite fields. This is important for many reasons, but in particular one consequence is that the characteristic polynomials of Frobenius acting on representations like the Tate module have rational coefficients (a priori they need only have been in $\mathbb{Q}_l$). The proof of this theorem makes full use of the machinery developed, as well as containing some definitely arithmetical ingredients. We shall need to first establish an analogue of the Cebotarev density theorem, in order to establish an auxiliary rationality result, and the Riemann Hypothesis then follows by a curious mixture of geometrical and analytic (or almost combinatorial) arguments.

2 The Étale Fundamental Group

In this section we state precisely what we mean by 'covers' in this context, and construct the étale fundamental group, a profinite group whose finite quotients are automorphism groups of certain ('Galois') covers. This construction generalises the absolute Galois group from field theory. We then develop some basic properties of these groups, in particular, some functorialities and the 'homotopy exact sequence' which allows one to decompose an étale fundamental group into a 'geometric' part and a field-theoretic part (the absolute Galois group of the base field). We mainly follow Szamuely [19, 5.3-5.6], with some ideas imported from Conrad's notes [2, 1.2].

Firstly, what is the appropriate generalisation of covering map? In the theory of Riemann surfaces, recall that arbitrary holomorphic maps fail to be covering spaces in the sense of topology because they may be ramified at a discrete set of points. We need our covering maps to have analogous properties to the topological case, so must insist that they are *unramified*: the correct algebraic definition being that a map $f : X \to Y$ locally of finite type is unramified if the maps on stalks $\mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ give rise to finite separable extensions of residue fields $k(f(x)) \to \mathcal{O}_{X,x}/\mathfrak{m}_{f(x)}\mathcal{O}_{X,x}$.

Additionally, we need to insist that our covering maps are flat morphisms, which is a technical commutative algebraic condition to ensure that the stalks vary continuously. In the cases which we shall study, the flatness will be obvious, and unramifiedness will agree with our geometric intuition. For further discussion of these points I refer the reader to [16, Ch 2]. In any case, we define an *étale morphism* to be one that is locally of finite type, flat and unramified. We will usually be dealing with *finite* étale morphisms, so the first condition is also never really an issue for us.

These have some important 'stability' properties. Firstly given $X \xrightarrow{f} Y \xrightarrow{g} Z$, if f and g are étale, so is $g \circ f$, and also if $g \circ f$ and g are étale then so is f (i.e. any map between two étale Z-schemes is étale). Also, given $X \to S$ étale and

$S' \to S$ arbitrary, the base change $X \times_S S' \to S'$ is étale. These facts will be used implicitly throughout the essay and are proved, for example, in [15, 3].

One very important property (closely related to what in [8, 1.5] Grothendieck calls "the fundamental property of étale morphisms") of unramified maps $X \to Y$ is that the diagonal map $X \to X \times_Y X$ is an open immersion (for a proof, see, for example, [15, 3.5]). This property allows us to prove the following crucial lemma, which takes the place of uniqueness of path liftings in classical topology. Note that it also generalises the familiar (but trivial) fact from field theory that for L/K finite separable, $Aut_K(L) \hookrightarrow Hom_K(L, \bar{K})$.

Lemma 2.1 (Rigidity Lemma)**.** *Let X and Y be finite étale S-schemes, with X connected and Y separated. Given two S-morphisms $f, g : X \to Y$, and a geometric point $\bar{x} \to X$, if $f(\bar{x}) = g(\bar{x})$ then $f = g$.*

Proof. Note that since $Y \to S$ is étale and separated, $\Delta : Y \to Y \times_S Y$ is both open and closed, so we may write $Y \times_S Y = \Delta(Y) \coprod Z$ as a disjoint union. But under the hypotheses, $f \times g : X \to Y \times_S Y$ intersects the diagonal, and since X is connected, this map has scheme theoretical image in $\Delta(Y)$, which implies that $f = g$. □

Equipped with this lemma, it is clear that for $X \to S$ connected finite étale, for any geometric point $\bar{x} \in S$ and a choice of point $\bar{x'}$ in the fibre, the map $f \mapsto f(\bar{x'})$ gives an injection $Aut(X/S) \hookrightarrow X_{\bar{x}}$. Generalising the definition from field theory, when this map is actually an isomorphism we say that $X \to S$ is *Galois*. In fact, a large amount of classical field theory generalises. We refer the interested reader to Szamuely [19, 5.3] for proofs of the following. For now and the rest of the next two sections unless otherwise specified, we shall assume that S is connected.

Proposition 2.2 (Generalised Galois Theory)**.**

1. *(Fundamental Theorem: [19, 5.3.6-7]) If $X \to S$ is (finite) Galois with automorphism group $G = Aut(X/S)$, there is an order-preserving correspondence*

$$\{\textit{connected étale subcovers } X \to Y \to S\} \leftrightarrow \{\textit{subgroups of } G\}.$$

This is given by the functors $Y \mapsto Aut_Y(X)$, and $H \mapsto X^H$, the fixed scheme of H, whose construction is a nontrivial step. This can be rephrased as an equivalence of categories

$$\{\textit{connected étale subcovers } X \to Y \to S\} \leftrightarrow \{\textit{transitive } G\textit{-sets}\}.$$

2. *(Existence of Galois closures: [19, 5.3.9]) Given $X \to S$ finite étale, there is a unique $X' \to X$ such that $X' \to S$ is Galois and such that any other Galois cover $X'' \to S$ factors through $X' \to S$.*

One of Grothendieck's key insights was to describe the above situation in terms of the (geometric) *fibre functor*, defined for a geometric point $\bar{x} \to S$ by

$$Fib_{\bar{x}} : (X \to S) \mapsto \mathrm{Hom}_S(\bar{x}, X) =: X(\bar{x}),$$

between the category of finite étale covers of S and the category of (finite) sets. This encodes simultaneously all the finite étale covers and their geometric fibres over a fixed basepoint, which as we remarked above are, by rigidity, sufficient avatars for measuring the symmetries of these spaces. In classical topology, the existence of a universal cover says that the fibre functor is representable. Here, it may not be representable, but it is nevertheless an object that tells us about all the covers of S. It is therefore natural to consider the automorphism group of the *functor* itself (i.e. the group of all compatible systems of permutations of the fibres of *every* finite étale cover of S). Indeed, we define the *étale fundamental group* $\pi_1(X, \bar{x})$ to be precisely this automorphism group.

Example 1 ($S =$ Spec k, k algebraically closed.)**.** *By unwrapping the commutative algebraic consequences of the definition of an étale map, the finite étale covers of a field are precisely Spec A for A a finite direct product of finite separable extensions of k. In particular, for k algebraically closed, they are just algebras of the form $A = k \oplus k \oplus ... \oplus k$. Considering k-homomorphisms of such algebras, it is not hard to see that the category of finite étale covers of Spec k is actually equivalent to the category of finite sets. In particular, there are a lot of morphisms on the domain category, so an automorphism of the fibre functor will be very heavily constrained.*

Indeed, let $(\alpha_n : Hom_k(Spec\ k, Spec\ k^n) \to Hom_k(Spec\ k, Spec\ k^n))_n$ be such an automorphism. For each $j \in \{1, ..., n\}$ let f_j be the map of schemes corresponding to projection from the jth co-ordinate. Since it is an automorphism of functors, we must have that $f_j = f_j \circ \alpha_1(id_k) = \alpha_n(f_j \circ id_k) = \alpha_n(f_j)$, hence α_n acts trivially for all n. So we have shown that $\pi_1(Spec\ k, id_k) = 1$. This may seem like a trivial example, but it is very important, and helps to highlight that the automorphism group of a fibre functor grows as the number of arrows in the category of finite étale covers diminishes. One may also generalise this argument to show that π_1 acts on the geometric points of each individual connected component independently, a fact we will need to simplify arguments later.

Though I find this abstract definition very beautiful, it is fortunate that one can also identify it somewhat more concretely as follows.

Proposition 2.3 (Pro-representability of the fibre functor)**.** *Let $\{(X_\alpha, \bar{x}_\alpha)\}$ be the inverse system of all finite Galois covers of S together with a compatible set of chosen basepoints lying over $\bar{x}$. Then*

$$Fib_{\bar{x}} \cong \varinjlim Hom_S(X_\alpha, -).$$

Furthermore,

$$\pi_1(S, \bar{x}) \cong \varprojlim Aut_S(X_\alpha)^{op}.$$

Proof. The first part is essentially a series of routine checks using the rigidity lemma. There are obvious compatible maps (given compatible basepoints have been fixed) $Hom_S(X_\alpha, Y) \to Fib_{\bar{x}}(Y)$, which combine to give a map in one direction. One then constructs a functorial inverse by reducing to the connected case and taking an appropriate map from a Galois closure of Y (using the fact that a map $\bar{y} \to Y$ can always be factored through a Galois closure $\bar{y} \to X_\alpha \to Y$, and the point this factorisation passes through corresponds (relative to the chosen basepoint $\bar{x}_\alpha$) to a unique S-automorphism of X_α, which in turn gives an element of $Hom_S(X_\alpha, Y)$. For further details, see [19, 5.4].

The second part then follows formally via the full-faithfulness of the Yoneda embedding $X \mapsto Hom_S(X, -)$. □

Example 2 (The absolute Galois group)**.** *Taking k an arbitrary field, recall that its absolute Galois group is given by $Gal(k^s/k) = \varprojlim Gal(L/k)$, where the limit is taken over all finite Galois extensions of k. However, 2.3 (recalling that by definition a Galois cover is connected) identifies this group with the group $\pi_1(Spec\ k, \bar{x})$ for any choice $\bar{x}$ of algebraic closure. Given we shall be interested in Galois representations, this is one of the main motivations for using étale fundamental groups. Also notice how the change in direction in passing from fields to schemes 'cancels' with the change of direction coming from the contravariance occuring in viewing $Fib_{\bar{x}}$ as pro-representable in the left term.*

Now, just as there is an infinite version of the fundamental theorem of Galois theory (and just as the fundamental group classifies all connected covers), we get the following result, which is probably the most important fact about the étale fundamental group.

Theorem 2.4. *The group $\pi_1(S, \bar{x})$ is profinite, acting continuously on $Fib_{\bar{x}}(X)$ for all $X \to S$ finite étale, with orbits corresponding to the fibres of each connected component. The fibre functor induces an equivalence of categories*

$$\{\textit{finite étale covers of } S\} \xrightarrow{\sim} \{\textit{finite continuous left } \pi_1(S, \bar{x})\textit{-sets}\}.$$

Proof. The first part of the theorem is obvious from what we have already done, and the statement about orbits is a straightforward consequence of the existence of Galois closures, the rigidity lemma and some previous remarks, so we just check the statement about $Fib_{\bar{x}}$ inducing an equivalence of categories, the main content of which is the descent-theoretic finite theorem 2.2.

Firstly, for essential surjectivity, note that given a finite $\pi_1(S, \bar{x})$-set Σ, wlog transitive (we can then get to arbitrary sets by taking a disjoint union), by continuity any point stabiliser is open, so must contain some element of a basis of open sets, in particular a normal subgroup giving rise to a quotient of the form $G = Aut_S(P)$ for some $P \to S$ Galois. Thus $\pi_1(S, \bar{x})$ acts through such a finite quotient, and by (2.2) the image H of the stabiliser under $\pi_1(S, \bar{x}) \to G$ gives rise to some finite étale $X = P^H \to S$, which by the rigidity lemma has geometric fibre isomorphic to the coset space $(G : H)$. But by transitivity and the orbit-stabiliser theorem, this is in turn isomorphic to Σ, and hence $\Sigma \cong X(\bar{x})$.

Now we check full faithfulness, so must show that the induced $Hom_S(X, Y) \to Hom_{\pi_1(S,\bar{x})}(X(\bar{x}), Y(\bar{x}))$ is an isomorphism. Injectivity follows immediately from the rigidity lemma. For surjectivity, we may suppose X wlog connected (by restricting attention to a single orbit of $X(\bar{x})$), and so also Y connected (by then taking a disjoint union in the second argument). Suppose we are given a map of $\pi_1(S,\bar{x})$-sets $f : X(\bar{x}) \to Y(\bar{x})$. By the $\pi_1(S,\bar{x})$-equivariance of f, for any $\bar{x}' \in X(\bar{x})$, $Stab(\bar{x}') \subseteq Stab(f(\bar{x}'))$. Taking a Galois closure $P \to X$, and noting that if $\sigma \in \pi_1(S,\bar{x})$ acts trivially on P, then it certainly stabilises both $\bar{x}'$ and $f(\bar{x}')$ (so in particular by rigidity one can use 2.2 to prove there is also a map $P \to Y$), we can pass to a finite quotient $Aut(P/X)$ and use the equivalence of categories at the finite level (2.2) to get a map $X \to Y$ with the required property. □

So far we have defined the étale fundamental group, quoted some generalisations of finite Galois theory, and used them to prove that the étale fundamental group can be computed as an inverse limit of finite automorphism groups, and as such is profinite. Using all the structure we were then able to prove a general Galois correspondence in the 'representation theoretic' form above. One notable feature was that even a single point Spec k, when over a non-algebraically closed field, picked up a nontrivial étale fundamental group, namely the absolute Galois group of the k. In the remainder of this chapter we develop some functorial properties of these groups and sketch the proof of an important theorem which allows the arithmetic and geometric content of these groups to be studied separately.

Firstly, it will be useful for us to understand the role of basepoints. Our group was defined by fixing a geometric point $\bar{x}$ and considering automorphisms of its fibre functor. If S is connected, it is easy to use pro-representability to noncanonically construct an isomorphism $\lambda : Fib_{\bar{x}} \to Fib_{\bar{x}'}$ between different fibre functors, and such a thing is called a *path* between $\bar{x}'$ and $\bar{x}$, since of course it defines an isomorphism

$$\pi_1(S,\bar{x}) \ni \sigma \mapsto \lambda^{-1}\sigma\lambda \in \pi_1(S,\bar{x}'),$$

well-defined up to conjugacy (corresponding to different choices of isomorphism λ).

This procedure will be very important for us in the following context. Given a connected scheme W of finite type over $\mathbb{Z}$ (e.g. any open subset of a ring of integers in a number field, or a variety over a finite field), all the closed points have finite residue fields, so the maps Spec $k(x) \to W$ embedding points induce, fixing an algebraic closure $\bar{x}$ of $k(x)$, a map $\pi_1(k(x),\bar{x}) \to \pi_1(W,\bar{x})$ and $\pi_1(k(x),\bar{x}) = G_{k(x)} \cong \hat{\mathbb{Z}}$ has a canonical topological generator, namely the *geometric Frobenius element* $(x \mapsto x^{|k(x)|})^{-1}$, so this in turn gives rise to a canonical element in $\pi_1(W,\bar{x})$.

Suppose we are (as we often shall be) studying $G = \pi_1(W,\bar{w})$ for some fixed basepoint $\bar{w}$. Then since W is connected we can obtain as above an isomorphism $\pi_1(W,\bar{x}) \to \pi_1(W,\bar{w})$, and hence get a canonical element in G, almost. Recall that in fact we could have chosen different isomorphisms, so we actually get a canonical conjugacy class, called the *Frobenius class at x*, a representative of which is denoted ϕ_x. In other words, we have constructed a map

$$\{\text{closed points of } W\} \to \{\text{conjugacy classes in } \pi_1(W, \bar{w})\},$$

which will play a crucial role in what follows. Note that if we are considering *representations* of G, elements in the same conjugacy class behave essentially identically: in particular they have the same characteristic polynomials.

Next, in classical algebraic topology one of the cool things about fundamental groups is that maps between spaces induce maps between fundamental groups. Here something similar happens. Given any map $f : S' \to S$, there is a *base change functor* $BC_{S' \to S} : X \mapsto X \times_S S'$ from finite étale S-schemes to finite étale S'-schemes. If $\bar{x}'$ maps to $\bar{x}$, the universal property of pullbacks gives a natural isomorphism

$$Hom_{S'}(\bar{x}', X \times_S S') \cong Hom_S(\bar{x}, X).$$

In other words, $Fib_{\bar{x}} = Fib_{\bar{x}'} \circ BC_{S' \to S}$, so any automorphism of $Fib_{\bar{x}'}$ induces an automorphism of $Fib_{\bar{x}}$, so we have an induced map

$$f_* : \pi_1(S', \bar{x}') \to \pi_1(S, \bar{x}).$$

By its very definition, it has the following useful property.

Proposition 2.5 (Compatibility of base change with restriction)**.** *Given $(S', \bar{x}') \to (S, \bar{x})$ a map of pointed schemes, the functors $BC_{S' \to S}$ from finite étale S-schemes to finite étale S'-schemes and Res_{f_*} from finite continuous $\pi_1(S, \bar{x})$-sets to finite continuous $\pi_1(S', \bar{x}')$-sets (given by acting via the morphism $f_* : \pi_1(S', \bar{x}') \to \pi_1(S, \bar{x})$) commute with the fibre functors.*

We now conclude the chapter with a proof of the following theorem (with some of the more tedious steps omitted, again [19] is an excellent reference). The main groundwork for the proof will be to ask questions like "When is the induced map on fundamental groups injective?", and show that such questions can be answered interestingly in terms of topological or algebraic-geometrical properties of the situation. We can then apply our answers to what will soon be an obvious sequence of two maps to consider.

Theorem 2.6 (Homotopy Exact Sequence)**.** *Let X_0 be a quasicompact geometrically connected scheme over a field k, and let $X = X_0 \otimes_k k^s$ denote its base change to k^s, the separable closure of k, and $\bar{x} : Spec\ k^s \to X$ any geometric point. Then there is a short exact sequence*

$$1 \to \pi_1(X, \bar{x}) \to \pi_1(X_0, \bar{x}) \to Gal(k^s/k) \to 1.$$

We start with a lemma characterising general exactness properties of induced maps in terms of geometrical criteria.

Lemma 2.7. *Given $\phi : (S', \bar{x}') \to (S, \bar{x})$ a map of pointed schemes, we consider the induced map $\phi_* : \pi_1(S', \bar{x}') \to \pi_1(S, \bar{x})$.*

1. *The map ϕ_* is trivial iff the base change of any finite étale cover is trivial.*

2. *The map ϕ_* is surjective iff whenever $X \to S$ is connected finite étale, so is $X \times_S S'$.*

3. *The map ϕ_* is injective iff for all $X' \to S'$ finite étale, there exists $X \to S$ finite étale such that for some connected component Z of $X \times_S S'$, there is a map $Z \to X'$. In particular, if every connected cover is obtained by base change, then ϕ_* is injective.*

4. *Given additionally $\psi : (S, \bar{x}') \to (S'', \bar{x})$, the sequence*

$$\pi_1(S', \bar{x}') \to \pi_1(S, \bar{x}) \to \pi_1(S'', \bar{x}'')$$

is exact iff the following two conditions hold.

   (a) *('Ker $\supseteq$ Im') For each $X'' \to S''$ finite étale, $X'' \times_{S''} S' \to S'$ is trivial.*

   (b) *('Ker $\subseteq$ Im') If $X \to S$ is such that $X \times_S S' \to S'$ has a section, then there is a finite étale map $X'' \to S''$ with some component of $X'' \times_{S''} S$ mapping to X.*

Proof. We shall only prove the first part, in order to illustrate roughly the flavour of the kind of arguments used. For the remainder of the proofs (which I admit are increasingly difficult and involved) see [19, 5.5].

To prove the first part, let us note that a finite étale cover is trivial iff it corresponds to a finite set with trivial action of the fundamental group, so one half is clear. It remains to prove that ϕ_* is trivial assuming that every pullback of a cover is trivial. Suppose it is not, so there must exist an open subgroup U of $\pi_1(S, \bar{x})$ which does not contain the image of ϕ_* (since the intersection of all the open subgroups is easily checked to be trivial). This corresponds under 2.4 to a finite étale cover $X \to S$, which will give a nontrivial base change $X \times_S S' \to S'$, contradicting the hypothesis. □

This lemma gives us a strategy for proving the theorem. We shall consider the sequence of maps $X_0 \times_{\text{Spec } k} \text{Spec } k^s \to X_0 \to \text{Spec } k$, and check it has all the properties required by the lemma for 2.6 to hold. If one tries to do this straight away, one does encounter a problem, and the following 'compactness' result is needed to fix it.

Lemma 2.8. *Let X_0 be quasicompact over k, and X its base change to k^s. Then any finite étale map $Y \to X$ is obtained by base change from a finite étale map $Y_L \to X_L$ of L-schemes, where L is a finite separable extension of k. Also, any element of $Aut(Y/X)$ comes from an element of $Aut(Y_L/X_L)$ for L sufficiently large.*

Proof. This is a fairly simple-minded "there are only finitely many equations with finitely many variables, so just adjoin them all to your base field" type proof. We leave the easy details to the reader, emphasising that quasicompactness is obviously an important hypothesis for such a strategy to work. Alternatively, see [19, 5.6.2]. □

We now have all the tools assembled to prove the theorem.

Proof of theorem 2.6. Lemma 2.7 gives us four things to check, which we can do in turn.

First let us deal with injectivity of the map $\pi_1(X, \bar{x}) \to \pi_1(X_0, \bar{x})$. We need to check that whenever $Y \to X$ is finite étale, there is some $Y_0 \to X_0$ also finite étale with a connected component of its base change mapping into Y. Given lemma 2.8 this is not too hard: we know that $Y = Y_L \otimes_L k^s$ for some finite separable L/k, which may be taken wlog to be Galois. Since Spec $L \to$ Spec k is also finite étale, in fact $Y_L \to$ Spec k is finite étale, and

$$Y_L \otimes_k k^s \cong Y_L \otimes_k L \otimes_L k^s,$$

and since L/k is Galois (so $L \otimes_k L \cong L \times ... \times L$), this decomposes into connected components each of which is k^s-isomorphic to $Y_L \otimes_L k^s \cong Y$.

Next, let us deal with surjectivity of $\pi_1(X_0, \bar{x}) \to Gal(k^s/k)$. It suffices to prove that $X_0 \otimes_k F$ is connected for all F/k finite separable. This is immediate, since we insisted that X_0 was geometrically connected.

Finally, we deal with the two exactness conditions. The first one is clear, because there is a commuting square via the morphism $X \to$ Spec $k^s \to$ Spec k, the middle term of which has trivial fundamental group. Verifying the second condition is slightly more fiddly, and it is here we will use integrality to argue over the generic fibre. Indeed, suppose $Y_0 \to X_0$ is finite *Galois* and such that its base change $Y \to X$ has a section $p : X \to Y$. Since X_0 (hence also Y_0, being irreducible finite étale) is integral, the generic fibre of $Y_0 \to X_0$ is a finite Galois extension K of the function field $k(X_0)$ which splits into a direct product of copies of $k^s(X_0)$ when tensored up to k^s (by the section and that $Y \to X$ is Galois). Thus, we must have $K = k(X_0) \otimes_k L$ for some finite Galois extension L/k. Hence $X_L = X_0 \otimes_k L$ has the same function field as Y_0, whence they are birational. So for some open dense $U_0 \subset X_0$, we have an isomorphism $Y_0 \times_{X_0} U_0 \overset{\cong}{\to} X_L \times_{X_0} U_0$ of finite locally-free U_0-schemes, which extends uniquely over X_0 to $Y_0 \overset{\cong}{\to} X_L$, as required. □

By now we should have a good idea of what the étale fundamental group can do for us. By keeping track of the automorphisms of all covers, and allowing us to replace geometrical objects (finite étale covers) with apparently algebraic or combinatorial objects (sets with a group action), they give us the possibility of transforming otherwise very fiddly arguments involving sections and connected components into more apparently algebraic arguments. Of course, the caveat is that the groups themselves are usually very large and complicated. In the next section we will approach finite étale covers via *sheaves*, which approach the task of studying covers algebraically from a different angle, and will eventually provide us with a slick tool for studying representations of étale fundamental groups, via the correspondence we established in this section.

3 Local Systems and Monodromy Representations

In this section we will study what are arguably the simplest nontrivial examples of sheaves on a space. In very classical cohomology theories, the coefficients are fixed (some coefficient ring like $\mathbb{Z}$ or $\mathbb{C}$). It was realised, with the invention of sheaves, that it is possible to actually let the coefficients vary from point to point on the space, for example you may want to use 'modules of functions near a point' as your coefficients (e.g. coherent sheaves). However, there is a step between these extremes: insisting that each point has a neighbourhood on which constant coefficients are used (so in some sense the coefficients *can* vary over the space, but must do so very 'slowly'). For setting up l-adic cohomology, this (together with a couple more tricks) is the level of sophistication required. We will eventually end up with something that is an interesting theory of cohomology for (almost) locally constant sheaves of $\mathbb{Q}_l$-vector spaces. Most significantly, passing via some of the work in the last section, we will show that such sheaves correspond precisely to finite dimensional π_1-representations, and in particular obtain a new perspective on Galois representations. Our main references are [2] and [6].

It is most convenient to use the language of sheaves, though in our context (thanks to another great idea of Grothendieck) we will not be talking about sheaves on a topological space, but rather sheaves on the 'étale site'. The general idea is to view étale morphisms $U \to X$ as taking the place of open sets, and to forget all about points. Given a scheme S, let Et_S be the category of all étale maps $U \to S$. Then an '*étale presheaf* of $\mathcal{C}$' on S is a contravariant functor $\mathcal{F} : Et_S \to \mathcal{C}$. For example, you could take $\mathcal{C} = (Sets)$ and map everything to the empty set. However, you could not map $S \neq U \to S$ to the empty set, but $S \to S$ to a singleton set, because there is an arrow $U \to S$, which would have to map to an arrow $\{*\} \to \emptyset$, and no such arrows exist.

A presheaf is called a *sheaf* if "whenever an element is defined étale-locally in a sensible, compatible way, it defines a unique element étale-globally." More precisely, suppose $U \in Et_S$, and there is some cover of U by étale maps $\{U_i \to U\}$. Then in all such situations, the diagram (which exists because $\mathcal{F}$ is a presheaf)

$$\mathcal{F}(U) \to \prod_i \mathcal{F}(U_i) \rightrightarrows \prod_{i,j} \mathcal{F}(U_i \times_U U_j)$$

must be an equaliser.

An element of $\mathcal{F}(U)$ tends to be called a *section*, and note also that for any $j : U \to X$ étale, we can define the restriction, written $j^*\mathcal{F}$ or $\mathcal{F}|_U$, just by restricting the category Et_X to the full subcategory Et_U (and it remains a sheaf).

There are many important examples of étale sheaves, but actually proving that almost anything is an étale sheaf is pretty technical. This is one of the main purposes of descent theory, with which we do not want to get bogged down. For a brief review see [6, Appendix III], for a more in depth survey [21], and for the canonical reference [8]. However, the two main results from descent theory which we shall need are the following.

Proposition 3.1 (Results from descent theory). 1. *The presheaves*

$$\bar{V} : U \mapsto Hom_S(U, V)$$

are all in fact étale sheaves. In other words, every representable presheaf is a sheaf (or every étale scheme is an étale sheaf).

2. *A sheaf is representable if and only if it is étale-locally representable (every geometric point has an étale neighbourhood over which the restriction of the sheaf is representable).*

So we now have a large stock of examples: any étale S-scheme defines (in fact can be identified with) a sheaf. There is also no reason for these sheaves to all just be sets (though of course this is how they naturally first appear): there could be some additional structure. For example, we can define the sheaf of nth roots of unity as that represented by $(\mu_n)_S := Spec(\mathbb{Z}[X]/(X^n - 1)) \times_{\text{Spec }(\mathbb{Z})} S$, and this has a natural abelian group structure. Indeed, evaluating on a geometric point $\bar{x}$ over the closed point x, we get $Hom_{k(x)}(k(x)[X]/(X^n - 1), k(x)^s) \cong \mathbb{Z}/n\mathbb{Z}$, a group under multiplication in a natural way.

A key example is the *constant sheaf* $\bar{\Sigma}$, which is represented by the scheme $\Sigma \times S$ (a disjoint union of S's indexed by elements of the set Σ). Note that if you try to define a constant sheaf in the naive way, by setting $\mathcal{F}(U) = \Sigma$ for all U, you only get a presheaf (in the classical case this is because some U may have multiple connected components - in the étale case it is because of this and even more reasons!). However, given a presheaf, there is a *sheafification functor* $\mathcal{F} \mapsto Sheaf(\mathcal{F})$ which is a left adjoint to the forgetful functor from sheaves to presheaves (so in some sense produces the closest possible sheaf to the given presheaf), and the sheaf thus obtained from the naive constant presheaf turns out to be the same as the one we just defined.

One very important construction on étale sheaves is the *inverse image functor*. Given a map of schemes $f : S \to T$ and an étale sheaf $\mathcal{F}$ on T, we can construct the sheaf

$$f^*\mathcal{F} := Sheaf(U \mapsto \varinjlim_{V} \mathcal{F}(V)),$$

where the limit is over the inverse system of all étale T-schemes V with a map $U \to V$ commuting with the map $f : S \to T$ on base schemes.

A prototypical and highly significant example of this construction is when $f = \bar{x}$ is a geometric point $Spec(k^s) \to T$. In this case, we write $\mathcal{F}_{\bar{x}}$ to denote $\bar{x}^*\mathcal{F}$ (or, only very slightly abusing notation, this sheaf evaluated on a point), and we call it the *stalk* at the geometric point $\bar{x}$. In fact, notice that the sheaf axioms on $Spec(k^s)$ are trivially always satisfied by any presheaf that respects disjoint unions, so no sheafification step takes place, and this construction can be performed for any presheaf. Furthermore, one may show that there is a natural map of presheaves $\mathcal{F} \to Sheaf(\mathcal{F})$ (the unit of adjunction), and that it induces an isomorphism $\mathcal{F}_{\bar{x}} \cong Sheaf(\mathcal{F})_{\bar{x}}$. The stalks are good at detecting local behaviour, and sheaves are designed to behave with respect to local behaviour, so in fact it turns out that a map of sheaves $\mathcal{F} \to \mathcal{G}$ is an isomorphism if and only if the induced maps on all stalks are isomorphisms (clearly unlike for presheaves, as our example shows). When we consider abelian sheaves, this will be even more important (as concepts like exactness become easiest to work with at the level of stalks).

We note in passing that the inverse image functor also generalises the restriction functor. If $j : U \to X$ is étale, then j^* defined just above is the same as j^* $(= -|_U)$ defined earlier in this section. Note that there are several other important sheaf operations we shall see later, which take a sheaf on S and 'push it forward' to a sheaf on T, but the inverse image functor is probably the most fundamental. One reason for this is noted in $[6, p28]$: the pullback of the sheaf represented by a space $X \to T$ along a morphism $f : S \to T$ is the sheaf represented by $S \times_T X \to S$, exactly as you might hope[1]

It is now time to make contact with the work we did in the previous section, where we classified all the finite étale covers of a base scheme S in terms of continuous $\pi_1(S, \bar{x})$-sets. The discussion above illustrates that every finite étale cover in fact gives rise to a sheaf. Now we explicitly identify the precise family of sheaves to which they correspond. Say that an étale sheaf $\mathcal{F}$ is *lcc* ('locally constant constructible') if every geometric point has an étale neighbourhood U such that $\mathcal{F}|_U$ is constant with values in a finite set (represented by $\Sigma_U \times U$ for Σ_U finite).

Theorem 3.2 (The correspondence between lcc sheaves and finite étale maps). *The map $X \mapsto Hom_S(-, X)$ associating a scheme to a sheaf induces an equivalence of categories between the category of finite étale covers of S and the category of lcc sheaves on S.*

Proof. As with a lot of the results in the previous section, much of the content here is again really in the descent theory. We already know by 3.1 that indeed this functor gives a sheaf, and it is fully faithful by the classical Yoneda lemma, so it remains to check that the sheaves it gives are indeed lcc, and that any lcc sheaf is representable by a finite étale scheme.

For the first part, it will suffice to show that for any $X \to S$ finite étale we can find an étale cover $\{S_i\}$ of S with each $S_i \times_S X$ a finite trivial cover of S_i. By separating into disjoint Zariski open subsets of S, assume $X \to S$ has constant degree n. The strategy will be induction on the degree, at each time splitting off one component by finding a section, in a way that generalises the classical construction of splitting fields.

Pass from $X \to S$ to $X \times_S X \to X$. Since $X \to S$ is separated and étale, this admits a section onto a connected component, so we get $X \times_S X = X \coprod Z \to X$,

[1] I should thank Teruyoshi Yoshida for pointing out that this follows formally as follows. The pullback $f^*\mathcal{X}$ of the T-sheaf $\mathcal{X}$ represented by X satisifes the canonical isomorphism $Hom_{Sh(S)}(f^*\mathcal{X}, \mathcal{F}) \cong Hom_{Sh(T)}(\mathcal{X}, f_*\mathcal{F}) \cong f_*\mathcal{F}(X) = \mathcal{F}(X \times_T S) \cong Hom_{Sh(S)}(\mathcal{X}', \mathcal{F})$, where $\mathcal{X}'$ is the sheaf represented by $X \times_T S$. Since $\mathcal{F}$ is arbitrary, it follows that $f^*\mathcal{X} \cong \mathcal{X}'$.

and $Z \to X$ has constant degree $n-1$. Thus by induction there is some $Y \to X$ finite étale with $Y \times_X Z \to Y$ a trivial cover. Then obviously

$$Y \times_X (X \times_S X) = Y \times_X (X \coprod Z) = Y \coprod Y \times_X Z \to Y \cong \coprod_n Y.$$

So $Y \to S$ is a trivialising étale cover, and thus X defines a locally constant étale sheaf.

For the converse, we again need simply apply descent theory. An lcc sheaf is by definition étale-locally representable, hence representable by 3.1. To check that the scheme X representing it is finite, note that by more descent theory, finiteness of a morphism is étale local on the base (see for example [12, Descent, 19.21]), and X is étale-locally finite over S by construction and full faithfulness of the Yoneda embedding. □

Note that previously (2.4) we showed a correspondence between the finite étale covers of S and the category of finite $\pi_1(S, \bar{x})$-sets. So passing via finite étale covers, it is easy to check we obtain the following result.

Corollary 3.3 (The finite monodromy correspondence)**.** *For any connected scheme S, there is an equivalence of categories*

$$\{\textit{lcc sheaves on } S\} \leftrightarrow \{\textit{finite continuous left } \pi_1(S, \bar{x})\textit{-sets}\}.$$

It is given by the stalk functor $\mathcal{F} \mapsto \mathcal{F}_{\bar{x}}$.

Furthermore, there is a functorial compatibility between pullback of sheaves and restriction of group action.

Recall that in the introduction we noted that objects like the Tate module of an elliptic curve are great news because they are nice modules with an action of the Galois group, so give rise to representations, but that we were lucky because the zero fibre of a map of elliptic curves is a kernel of a group homomorphism, so has a group structure. One of our major goals is therefore to come up with a more systematic way of generating such situations. And now we have one staring us in the face. Instead of looking at arbitrary lcc sheaves of sets, let us restrict our attention to lcc sheaves of abelian groups or (eventually) modules. These will precisely correspond to finite étale covers with a group structure on geometric fibres. We therefore spend the rest of the chapter, and much of the rest of the essay, studying abelian lcc sheaves in their own right.

Given a scheme S, it is well-known that the category of abelian étale sheaves on S is an abelian category (the maps between them have abelian group structures bilinear under composition, binary products and coproducts exist and coincide, and the concepts of exactness make sense: in particular there is a 'first isomorphism theorem'). In fact this category has the property that a sequence $\mathcal{F}' \to \mathcal{F} \to \mathcal{F}''$ is exact iff for every geometric point $\bar{x}$, the sequence of abelian groups $\mathcal{F}'_{\bar{x}} \to \mathcal{F}_{\bar{x}} \to \mathcal{F}''_{\bar{x}}$ is exact. It was also proved by Grothendieck ([7]) that this category has enough injectives, so there is a good theory of right derived functors.

Given a map $f : X \to Y$ there are a few crucial functors we can define between their categories $Ab(X)$ and $Ab(Y)$ of abelian étale sheaves. We already mentioned $f^* : Ab(Y) \to Ab(X)$, and this is an exact functor (since it obviously preserves all stalks functorially). We can also define the *direct image functor* (or 'pushforward') $f_* : Ab(X) \to Ab(Y)$ in a fairly straightforward way: it is given by $f_*\mathcal{F}(V) := \mathcal{F}(X \times_Y V)$. For a finite morphism, a stalk of finite size will tend to have its size multiplied by the degree of the map, so it's not entirely unresaonable to picture this as some kind of integral or sum.

Similarly to f^*, perhaps the most important special case is the structure map $f : X \to \text{Spec } k$ for X a scheme over an algebraically closed field. Because sheaves on the latter are really just abelian groups, this functor gives an abelian group $\mathcal{F}(X)$, usually called the *global sections* of $\mathcal{F}$ and written $\Gamma(X, \mathcal{F})$ or $H^0(X, \mathcal{F})$.

Unlike f^* note that the functor f_* is generally only left exact, so admits nontrivial right-derived functors, called the *higher direct images* and written $R^i f_* : Ab(X) \to Ab(Y)$. In the special case where Y is the spectrum of an algebraically closed field, it is traditional to denote these (evaluated at a single point) by $H^i(X, \mathcal{F})$, and call this the *cohomology* of X with coefficients in $\mathcal{F}$, in line with the definition of classical sheaf cohomology.

There is another important operation $f_! : Ab(X) \to Ab(Y)$ we will have to consider, but it is somewhat more complicated than the previous two. Let us start with considering an open immersion $j : U \to X$. Given a sheaf $\mathcal{F}$ on U, we can define a presheaf

$$\mathcal{F}_!(g : V \to X) = \begin{cases} \mathcal{F}(V \to U) & \text{if } g(V) \subset U \\ 0 & \text{otherwise} \end{cases}.$$

This need not be a sheaf, so we define the *extension-by-zero* sheaf to be $j_!\mathcal{F} = Sheaf(\mathcal{F}_!)$. Now let us return to the (more) general case.

Suppose $j : X \hookrightarrow \bar{X}$ is an open immersion and that there is a proper morphism $\bar{f} : \bar{X} \to Y$ with $\bar{f} \circ j = f$. For $\mathcal{F}$ a sheaf on X, we *define*

$$R^i f_! \mathcal{F} := R^i \bar{f}_*(j_! \mathcal{F}).$$

This is called the *higher direct image with proper support*, and for discussion of why it is well-defined and called this, see [6, 1.8] or [15, VI.3]. We should warn the reader that this is not a derived functor of $f_! = \bar{f}_* \circ j_!$. However, in many cases we care about it will behave sufficiently like it is.

For example, in the case of the structure morphism of a variety over an algebraically closed field, $X \to \text{Spec } k$, the Nagata compactification theorem gives such a proper $\bar{f} : \bar{X} \to \text{Spec } k$ (for a modern proof of this see [3]), so the above procedure works, and the groups obtained by evaluating $R^i f_! \mathcal{F}$ on a single point are denoted by $H^i_c(X, \mathcal{F}) := H^i(\bar{X}, j_!\mathcal{F})$, and called the *cohomology with proper support.* Of course, it is clear that we also get a long exact sequence of cohomology groups with proper support (commencing with the *sections with proper support* $H^0(\bar{X}, j_!\mathcal{F})$). We should also immediately remark that when X

itself is proper, $f_* = f_!$ and $H^i_c = H^i$, but when X is not proper it will be important to have both functors around.

There are a few key facts it will be worth knowing. Firstly, it is the case that f^* has f_* as its right adjoint (so a map $f^*\mathcal{F} \to \mathcal{G}$ is the same as a map $\mathcal{F} \to f_*\mathcal{G}$. Next, all these functors 'are themselves functorial', so two composable maps of schemes will give rise to composable functors that behave as nicely as you might hope. Also, one should note that $f_!$ usually gives a subsheaf of f_*, and in fact there is a very useful exact sequence. If $j : U \hookrightarrow X \leftarrow Z : i$ are an open and a closed immersion respectively of a Zariski-open set and its complement, there is an exact sequence of abelian sheaves

$$0 \to j_! j^* \mathcal{F} \to \mathcal{F} \to i_* i^* \mathcal{F} \to 0.$$

This is easy to prove by just checking on stalks.

In our case, we note that for a sheaf $\mathcal{G}$ on U that $j^* j_* \mathcal{G} = \mathcal{G}$, so get

$$0 \to j_! \mathcal{G} \to j_* \mathcal{G} \to i_* i^* j_* \mathcal{G} \to 0$$

and since $\bar{f}_*$ is left exact (and $\bar{f}_* \circ j_* = f_*$), we get $f_! \mathcal{G} \hookrightarrow f_* \mathcal{G}$.

It will be important for us to have some compatibility between the pullback functors and the higher direct image functors, at least under certain nice conditions. Consider an arbitrary commutative diagram of schemes as below, and start with a sheaf $\mathcal{F}$ on X.

$$\begin{array}{ccc} X' & \xrightarrow{g'} & X \\ \downarrow{\scriptstyle f'} & & \downarrow{\scriptstyle f} \\ S' & \xrightarrow{g} & S. \end{array}$$

It is possible (and not very hard) in general to construct a *base change morphism*

$$g^* R^i f_* \mathcal{F} \to R^i f'_* (g'^* \mathcal{F}).$$

That said, there is no reason it should generally have any interesting properties. However, if we insist that the square is a pullback square, and impose certain conditions on f, it becomes an isomorphism. The following two results are crucial to the foundations of the theory of étale cohomology, and are very difficult to prove (the original source is [1], but abridgements of parts of the proofs are given in [6] and [5]).

Theorem 3.4 (Proper base change and smooth base change)**.** *Let*

$$\begin{array}{ccc} X' & \xrightarrow{g'} & X \\ \downarrow{\scriptstyle f'} & & \downarrow{\scriptstyle f} \\ S' & \xrightarrow{g} & S \end{array}$$

be a pullback square of schemes, and let $\mathcal{F}$ be an lcc abelian sheaf on X. Suppose we have one of the following:

- *The map f is proper.*
- *The map g is smooth, f is quasicompact, and $\mathcal{F}$ has torsion prime to the residue characteristics of all points on the base S.*

Then the base change morphism is an isomorphism.

A key example we will be making use of a lot is

$$\begin{array}{ccc} X_{\bar{x}} & \longrightarrow & X \\ \big\downarrow{\scriptstyle f'} & & \big\downarrow{\scriptstyle f} \\ \text{Spec } k^s & \xrightarrow{\bar{x}} & S \end{array}$$

where the base change morphism is an isomorphism identifying fibres of a higher direct image with cohomology of a geometric fibre

$$(R^i f_* \mathcal{F})_{\bar{x}} \stackrel{\cong}{\to} H^i(X_{\bar{x}}, \mathcal{F}').$$

A common special case of this is of course when $S = \text{Spec } k$, in which case we learn that we can extract the cohomology of the base change of a variety over k to its algebraic closure by taking a fibre of the higher direct images to Spec k. This is not unimportant: sheaves on Spec k are Galois representations, and this process is one explanation for the Galois action on cohomology.

There are some less trivial but important applications of the proper base change theorem, sketched in Deligne [5, Arcata, IV]. Firstly, note that as a result of our definition of cohomology with compact support, whenever there is a Nagata compactification, we get a base change isomorphism!

Corollary 3.5 (Proper base change for higher direct images with proper support)**.** *Let*

$$\begin{array}{ccc} X' & \xrightarrow{g'} & X \\ \big\downarrow{\scriptstyle f'} & & \big\downarrow{\scriptstyle f} \\ S' & \xrightarrow{g} & S \end{array}$$

be a pullback square of schemes, and let $\mathcal{F}$ be an lcc abelian sheaf on X. Suppose further more that $f : X \to S$ is a separated morphism of finite type, of noetherian schemes. Then there is a base change map, and it is an isomorphism:

$$g^* R^i f_! \mathcal{F} \to R^i f'_!(g'^* \mathcal{F}).$$

This has the following far from trivial but not terribly hard consequences.

Theorem 3.6 (Vanishing theorem)**.** *Let $f : X \to S$ be a separated morphism of finite type of noetherian schemes, with fibres of dimension at most n, and $\mathcal{F}$ be lcc. Then $R^i f_! \mathcal{F} = 0$ for all $i > 2n$.*

Proof. See Deligne [5, Arcata, IV-6] for a sketch. The basic strategy seems to be some kind of induction, fibring over curves and at each stage using proper base change. □

Theorem 3.7 (Finiteness theorem and Ehresmann's Fibration Theorem)**.** *Let $f : X \to S$ be smooth and proper, and $\mathcal{F}$ be an lcc sheaf with torsion prime to the residue characteristics at points of S. Then the higher direct images $R^i f_* \mathcal{F}$ are also lcc.*

Proof. The finiteness theorem (which does not prove it is locally constant) is again sketched in Deligne [5, Arcata, IV-6]. Conrad [2, 1.3.7] goes on to prove the appropriate version of the fibration theorem, which guarantees our sheaves remain locally constant. □

This last result in particular is suddenly unbelievably powerful. Recall that I said f_* could be considered as some kind of sum or integral, and for large finite morphisms might produce large finite covering spaces. What 3.7 tells us is that actually as the dimensions get higher, everything doesn't blow up horribly. Instead, what seems to be happening is that the extra dimensions are filling up more of the higher direct image sheaves, but this effect too seems to be bounded, by 3.6. We therefore have, for a proper smooth variety, a very helpful intuition for bounding how its étale cohomology (with coefficients in lcc sheaves) can behave. In particular, we know that we can start with a finite étale cover on X and taking higher direct images will generate a nice family of finite étale covers of S.

Now let us conclude the chapter by taking our final steps towards generalised Tate modules. So far all our covering spaces have been locally constant with fibres being finite groups. In order to do linear algebra, we would ideally like our sheaves to actually be sheaves of finite dimensional vector spaces over some field of characteristic zero. Inspired by how we constructed the Tate module, we proceed as follows.

Firstly, let us fix a prime l for all time, and make sure if possible that it is coprime to all characteristics of residue fields of all schemes we are interested in (this will make life nicer, because then the hypothesis of smooth base change will hold). We set $\Lambda_n := \mathbb{Z}/(l^{n+1})\mathbb{Z}$, and restrict our attention to sheaves of finitely generated Λ_n-modules on a scheme S. We now define a *smooth $\mathbb{Z}_l$-sheaf* $\mathcal{F}$ on S to be a projective system of lcc abelian sheaves $(\mathcal{F}_n)_n$ on S, where each $\mathcal{F}_n$ is an lcc Λ_n-module, and where the maps $\mathcal{F}_n \to \mathcal{F}_{n-1}$ each induce an isomorphism

$$\mathcal{F}_n \otimes_{\Lambda_n} \Lambda_{n-1} \stackrel{\cong}{\to} \mathcal{F}_{n-1}.$$

Recall that by (3.3), if we fix a geometric point $\bar{x}$ of S and take the fibres of each of the sheaves at every level, we get an inverse system of Λ_n-modules with a continuous left $\pi_1(S, \bar{x})$-action. Therefore, we can take the inverse limit of this system, and get a $\mathbb{Z}_l$-module with a continuous action of $\pi_1(S, \bar{x})$, and this is suddenly looking very much like the Tate module. Now, we will wish to manipulate these 'sheaves' (we should remember that they are actually *inverse systems* of sheaves) as if they were honest sheaves, so it is probably worth noting that category of $\mathbb{Z}_l$-sheaves is an abelian category, and it respects the functors f_* and f^*. Moreover, from [6, 1.12.15] we have as a consequence of 3.7 the following key result.

Proposition 3.8 (Higher direct images with proper support of l-adic sheaves are l-adic)**.** *Let $f : X \to S$ be a separated morphism of finite type of noetherian schemes, and $\mathcal{F} = (\mathcal{F}_n)$ a smooth $\mathbb{Z}_l$-sheaf. Then for each i, $R^i f_! \mathcal{F} := (R^i f_! \mathcal{F}_n)$ is a smooth $\mathbb{Z}_l$-sheaf.*

So we now in particular have found a very good way to generate finite dimensional $G_k = \pi_1(\text{Spec } k, \bar{x})$-representations (with $\mathbb{Z}_l$-coefficients): write down any variety over k, and take the higher direct images of the constant $\mathbb{Z}_l$ sheaf to get some potentially highly nonconstant $\mathbb{Z}_l$-sheaves on Spec k, which correspond via 3.3 to nontrivial G_k-representations. These are precisely the l-adic cohomology groups, which we will study in more detail in the next section.

Of course, we will then often want to tensor up $\otimes_{\mathbb{Z}_l} \mathbb{Q}_l$ and obtain finite dimensional representations over the vector space $\mathbb{Q}_l$. Such a sheaf (a $\mathbb{Z}_l$-sheaf whose stalks are evaluated by taking an inverse limit and tensoring by $\mathbb{Q}_l$) is called a $\mathbb{Q}_l$-*sheaf.* A *smooth* $\mathbb{Q}_l$-sheaf is one obtained in such a way from a smooth $\mathbb{Z}_l$-sheaf (an inverse system of *lcc* sheaves). Since we will often suppress their finite construction and study $\mathbb{Z}_l$ and $\mathbb{Q}_l$ sheaves as objects in their own right, it is worth noting that we can extend the finite monodromy correspondence 3.3 to the following. For R a topological ring and G a topological group, we let $(Rep_R(G))$ denote the category of finite rank free R-modules with a continuous action of G.

Corollary 3.9 (The $\mathbb{Z}_l$-monodromy correspondence)**.** *Let S be connected. Taking the stalk (as an inverse limit of stalks) gives an equivalence of categories*

$$\{\textit{smooth } \mathbb{Z}_l\textit{-sheaves on } S\} \leftrightarrow (Rep_{\mathbb{Z}_l}(\pi_1(S, \bar{x}))).$$

Proof. This is immediate, since the category on the right may easily shown to be isomorphic to the category of inverse systems of Λ_n-modules with continuous $\pi_1(S, \bar{x})$-action, via the mutually inverse functors $(M_n) \mapsto \varprojlim M_n$ and $M \mapsto (M/l^{n+1})$. □

Corollary 3.10 (The $\mathbb{Q}_l$-monodromy correspondence)**.** *Let S be connected. Taking the stalk (by taking an inverse limit and tensoring) gives an equivalence of categories*

$$\{\textit{smooth } \mathbb{Q}_l\textit{-sheaves on } S\} \leftrightarrow (Rep_{\mathbb{Q}_l}(\pi_1(S, \bar{x}))).$$

Proof. Let V be a finite dimensional $\mathbb{Q}_l$-representation of $\pi_1(S, \bar{x})$, and pick M an arbitrary $\mathbb{Z}_l$-submodule which spans V over $\mathbb{Q}_l$. Since M is open in V, it is stable under the action of some open (in particular finite index) subgroup H of $\pi_1(S, \bar{x})$. Picking coset representatives $g_1, ..., g_n$ for H, $\pi_1(S, \bar{x})$ therefore acts continuously on the $\mathbb{Z}_l$-module $N = (g_1(M), g_2(M), ..., g_n(M))$. Since all $\mathbb{Z}_l$-submodules of V are free, any other such $\mathbb{Z}_l$-module is $\mathbb{Q}_l$-isomorphic to this one as a $\pi_1(S, \bar{x})$-representation, which establishes an equivalence of categories, making this corollary equivalent to the previous one. □

It is also possible to generalise the entire story to any algebraic extension of $\mathbb{Q}_l$ (more generally, by considering inverse systems of $\mathcal{O}/\mathfrak{p}^r$-modules), and taking all such constructions together obtain the following. We shall not do the details, but they can be found in [2] or [13].

Corollary 3.11 (The $\bar{\mathbb{Q}}_l$-monodromy correspondence)**.** *Let S be connected. Taking the stalk (by taking an inverse limit and tensoring) gives an equivalence of categories*

$$\{\textit{smooth } \bar{\mathbb{Q}}_l\textit{-sheaves on } S\} \leftrightarrow (Rep_{\mathbb{Z}_l}(\pi_1(S, \bar{x}))).$$

This final corollary is an analogue of the classical monodromy correspondence, which it would be nice to briefly mention. Given an nth order linear ODE (say we are looking at functions with a single complex variable), we can solve it in any small ball to get an n-dimensional complex vector space of solutions, and these glue together to form a *local system*: a locally constant sheaf of n-dimensional vector spaces. However, if you take a solution and analytically continue it around a loop, you might change it when you get back, and these changes correspond to linear automorphisms of the stalk over your starting point. One can show that this process only depends on the homotopy class of the loop, so this really is an action of the classical fundamental group on the stalk. The classical monodromy theorem tells us that this process extends to a natural correspondence between local systems and representations of the fundamental group.

In our case, we did no such analytic continuation, and our fundamental group does not involve walking around paths, but we have managed to extract an equally fantastic and perhaps somewhat surprising theorem. These locally constant sheaves, even though they can vary slowly across our space, can be completely classified by their stalk at a single point and how the étale fundamental group acts, and furthermore for any such action there is a unique such sheaf. In the next section, we will restrict our attention to étale cohomology groups, which will end up being l-adic sheaves on Spec k, for k an algebraically closed field (though often viewed as stalks of a sheaf over a nonalgebraically closed field - i.e. as Galois representations). However, when we come to prove Deligne's theorem, having understanding of the fuller machinery of l-adic 'local systems' available to us will be vital.

4 A Hitchhiker's Guide to Étale Cohomology

In this short section we shall give a quick tour, with little proof, of the main features of (for want of a better name) 'classical' étale cohomology.

Recall that an elliptic curve E like $y^2 = x^3 - x$ is (in projective space) a genus one plane curve, so over $\mathbb{C}$ it is topologically a doughnut. One could apply one's favourite classical theory of homology or cohomology, and the groups one expects to compute are $H^0(E, \mathbb{Q}) = H^2(E, \mathbb{Q}) = \mathbb{Q}$, and $H^1(E, \mathbb{Q}) = \mathbb{Q}^2$ (with all other groups trivial), and these are supposed to give you topological information about the space. The dimensions of these $\mathbb{Q}$-vector spaces are called the *Betti numbers* b_i (so here $b_0 = b_2 = 1, b_1 = 2, b_i = 0\ \forall i \neq 0, 1, 2$). For a compact d-dimensional complex manifold, they were proved generally to be finite positive integers and vanish outside $0 \leq i \leq 2d$. Another basic feature of the theory is that maps $f : X \to Y$ induce maps $H^i(Y, \mathbb{Q}) \to H^i(X, \mathbb{Q})$, and these might have information which the raw Betti numbers fail to give you.

In the first half of the twentieth century, this theory was developed into a fearsome arsenal of tools for studying spaces via their cohomology. The groups are related by a *cup product* $\cup : H^i(X, \mathbb{Q}) \times H^j(X, \mathbb{Q}) \to H^{i+j}(X, \mathbb{Q})$, which makes $H^*(X, \mathbb{Q}) = \oplus_i H^i(X, \mathbb{Q})$ into a graded algebra. One could write a *Kunneth formula* $H^*(X \times Y, \mathbb{Q}) \cong H^*(X, \mathbb{Q}) \otimes H^*(Y, \mathbb{Q})$ to give the cohomology of a product of spaces. A closer study of the cup product gave the *Poincaré duality* theorem that for an oriented compact d-fold, $\cup : H^i(X, \mathbb{Q}) \times H^{2d-i}(X, \mathbb{Q}) \to H^{2d}(X, \mathbb{Q}) = \mathbb{Q}$ is a perfect pairing, so in fact $b_i = b_{2d-i}$. Another key tool for computing Betti numbers of algebraic varieties is the *weak Lefschetz theorem*, which allows one to compute $b_0, ..., b_{d-2}$ and get a bound on b_{d-1} by computing the cohomology of a hyperplane section (in particular a variety of smaller dimension). These tools combine to focus all the attention on the middle cohomology group $H^d(X, \mathbb{Q})$.

In order to study the middle cohomology of a projective variety, the *Lefschetz pencil* was devised, where you choose a rotating system of hyperplanes and study all of their sections varying simultaneously as an algebraic variety fibred over $\mathbb{P}^1$, and study a *local system* on $\mathbb{P}^1$ whose stalks are the cohomology of the hyperplane sections. This local system, which behaves well except at the finitely many singular hyperplane sections, can then be studied in detail at singular points using the theory of *vanishing cycles*. These essentially boil down to studying the cohomology classes corresponding to loops which get smaller and smaller as the hyperplane sections approach some singular point, at which they vanish. Being functionals of loops, these interact in an interesting way with the monodromy group. All this information can then (hopefully) be used to reconstruct the middle cohomology group.

There were also tools developed to study *maps* more thoroughly, most notably the *Lefschetz Fixed Point theorem*, which gives a relation between the number of fixed points of an appropriate endomorphism of a space and the traces of the induced endomorphisms on the cohomology groups. It was the hope of applying this formula to the Frobenius morphism which really inspired the birth of étale cohomology as a means of solving the Weil Conjectures (which will be discussed in detail in the next chapter and are our ultimate goal in this essay).

The problem with classical cohomology is that much of it is intrinsically topological. You get to it by considering formal linear sums of paths, or using other techniques from differential geometry or analysis. Even where this is not an issue, there are big problems with making it work in positive characteristic. In particular, if you try things like replacing the coefficient ring $\mathbb{Q}$ by $\mathbb{F}_q$, the Lefschetz fixed point formula becomes meaningless (and the Betti numbers might not agree with classical results). The reason étale cohomology was invented and why it is useful is that by proceeding via a finitary 'take finite étale covers type' construction and then taking a limit, it somehow sidesteps all of these difficulties.

Let $\pi : X \to \mathrm{Spec}\ k$ be a variety (reduced integral scheme of finite type) of dimension d over an algebraically closed field k, possibly of positive characteristic p. Let $l \neq p$ be prime. We define the étale cohomology groups as in the previous section $H^i(X, \mathbb{Q}_l) := R^i\pi_*\mathbb{Q}_l$, where $\mathbb{Q}_l$ is the *constant* $\mathbb{Q}_l$*-sheaf* attached to the *constant inverse system* ... $\to \Lambda_2 \to \Lambda_1 \to \Lambda_0$. Similarly, we can define *cohomology with compact support* $H^i_c(X, \mathbb{Q}_l) := R^i\pi_!\mathbb{Q}_l$. We now state without proof the main theorems of étale cohomology in this simple case. They are proved in [1] for torsion coefficients, and the passage to $\mathbb{Z}_l$ and $\mathbb{Q}_l$ coefficients is checked in [9].

Theorem 4.1 (Basic properties of étale cohomology)**.** *We have:*

1. *(Finiteness) For $\mathcal{F}$ a smooth $\mathbb{Q}_l$-sheaf, the groups $H^i(X, \mathcal{F})$ and $H^i_c(X, \mathcal{F})$ vanish for $i > 2d$, and the groups $H^i_c(X, \mathcal{F})$ are finite dimensional $\mathbb{Q}_l$-vector spaces.*
2. *(Functoriality) The maps $X \mapsto H^i(X, \mathbb{Q}_l)$, $X \mapsto H^i_c(X, \mathbb{Q}_l)$ are contravariant functors in X.*
3. *(Cup product) There is a cup product structure $H^i(X, \mathbb{Q}_l) \times H^j(X, \mathbb{Q}_l) \to H^{i+j}(X, \mathbb{Q}_l)$ defined for all $i, j \geq 0$.*
4. *(Proper smooth base change) If $X \to Y$ is smooth and proper, and Y is connected, then the dimension of $H^i(X_y, \mathbb{Q}_l)$ is constant as y varies on Y.*
5. *(Kunneth formula) For X and Y proper varieties over k,*

$$H^n(X \times Y, \mathbb{Q}_l) \cong \oplus_{p+q=n} H^p(X, \mathbb{Q}_l) \otimes H^q(Y, \mathbb{Q}_l).$$

Note some special cases of the difficult theorems 3.7 and 3.6.

As classically, we have a Poincaré duality theorem. Before we can state it, we need to briefly discuss the correct analogue of orientability. In the theory of complex manifolds, any compact manifold is orientable, a choice of orientation corresponds to a choice of i, and changing the orientation could change the answer given by Poincaré duality by a sign.

In our theory, the correct analogue is choosing an isomorphism between the roots of unity μ_∞ on the base field and the abstract group $\mathbb{Q}/\mathbb{Z}$. Therefore, to avoid having to make such a non-canonical identification at the outset, we introduce the sheaf $\mathbb{Z}_l(1) := \varprojlim \mu_{l^n}(k)$, viewing it as a $\mathbb{Z}_l$-sheaf in an obvious way. We write $\mathbb{Z}_l(n)$ for its nth tensor power, allow negative numbers to correspond to taking duals, and write $\mathcal{F}(n) := \mathcal{F} \otimes_{\mathbb{Z}_l} \mathbb{Z}_l(n)$. As groups, the sheaves $\mathbb{Z}_l(n)$ are all isomorphic to $\mathbb{Z}_l$, but the subtle difference of their being roots of unity is very important once we start viewing them as Galois representations, since they will start to have nontrivial Galois action. In [4, 2] these issues are discussed in more depth.

We can now state our theorem, both in its most elementary case, and then in some more general cases which we need.

Theorem 4.2 (Poincaré Duality)**.** *If X is smooth and proper of dimension d, then there is a canonical isomorphism $H^{2d}(X, \mathbb{Q}_l) \cong \mathbb{Q}_l(-d)$, and for each i, the cup product induces a perfect pairing*

$$H^i(X, \mathbb{Q}_l) \times H^{2d-i}(X, \mathbb{Q}_l) \to H^{2d}(X, \mathbb{Q}_l) \cong \mathbb{Q}_l(-d).$$

More generally, if X is smooth (not necessarily proper) of dimension d, and $\mathcal{F}$ a smooth $\mathbb{Q}_l$ sheaf, $\check{\mathcal{F}}$ its dual, there is a perfect pairing

$$H^i(X, \mathcal{F}) \times H_c^{2d-i}(X, \check{\mathcal{F}}(d)) \to H_c^{2d}(X, \mathcal{F} \otimes \check{\mathcal{F}}(d)) \to H_c^{2d}(X, \mathbb{Q}_l(d)) \to \mathbb{Q}_l.$$

Note that in the second form, where X is no longer necessarily proper, the duality is between ordinary cohomology and cohomology with compact support. We will also need a slight modification which deals with how sheaves behave under an open immersion of a curve into its compactification. All these results are quoted directly from Deligne's paper [4], in which he references SGA4 [1] and SGA5 [9].

Theorem 4.3 (Poincaré Duality after an open immersion (for curves))**.** *Let X be a smooth connected projective curve over k algebraically closed, and $j : U \hookrightarrow X$ an open immersion. Then for $\mathcal{F}$ a smooth $\mathbb{Q}_l$-sheaf on U, we have a perfect pairing*

$$H^i(X, j_*\mathcal{F}) \times H^{2-i}(X, j_*\check{\mathcal{F}}(1)) \to H^2(X, j_*\mathbb{Q}_l(1)) = \mathbb{Q}_l.$$

The tools for computing Betti numbers of X by reducing to computing Betti numbers of hyperplane sections also generalise.

Theorem 4.4 (Weak Lefschetz Theorem)**.** *Let $X \subset \mathbb{P}^n$ be a projective scheme of dimension d over an algebraically closed field, and $\mathcal{F}$ a smooth $\mathbb{Q}_l$-sheaf. Let H be a hyperplane, intersecting X at a closed subvariety $i : A \hookrightarrow X$. Provided the complement $U = X \backslash A$ is smooth, the maps*

$$H^i(X, \mathcal{F}) \to H^i(A, i^*\mathcal{F})$$

are bijective for $i < d - 1$ and injective for $i = d - 1$.

This theorem is proved by considering the long exact sequence coming from taking cohomology of the short exact sequence of sheaves (writing $j : U \hookrightarrow X$)

$$0 \to j_! j^* \mathcal{F} \to \mathcal{F} \to i_* i^* \mathcal{F} \to 0,$$

and using the fact that cohomology actually vanishes above $n > d$ for U affine, so by Poincaré duality the cohomology with proper supports vanishes below d.

It is also worth extracting the Poincaré dual of the case $n = d-1$, as a useful fact for later.

Corollary 4.5 (Gysin morphism of a hyperplane section is surjective)**.** *Let $X \subset \mathbb{P}^n$ be a projective scheme of dimension d over an algebraically closed field, and $\mathcal{F}$ a smooth $\mathbb{Q}_l$-sheaf. Let H be a hyperplane, intersecting X at a closed subvariety $i : A \hookrightarrow X$. Suppose further that the complement $U = X \backslash A$ is smooth. Then the map*

$$H^{d-1}(A, i^*\mathcal{F})(-1) \to H^{d+1}(X, \mathcal{F})$$

is surjective.

At least in theory, this machinery should allow one to focus on the middle cohomology group $H^d(X, \mathbb{Q}_l)$, and just as people did classically, we shall then pick this apart using *Lefschetz pencils* and the theory of vanishing cycles, which also generalises well to étale cohomology.

Firstly, a Lefschetz pencil is a construction of the following kind. The reader is advised for now to visualise this over the real numbers in three dimensions. Take a smooth projective variety $X \hookrightarrow \mathbb{P}^N$, and fix a codimension 2 linear subspace $L \subset \mathbb{P}^N$. The space of all hyperplanes $\{H_d\}$ containing L is parameterised by $D := \mathbb{P}^1$ (if you like, if you parameterise an arbitrary hyperplane and solve for those which contain L, only two homogeneous parameters survive - or just think of L as an axis of a single degree of rotation). We now construct the variety $\tilde{X} := \{(x, d) \in X \times D | x \in H_d\}$. This comes with projection maps:

$$X \leftarrow \tilde{X} \xrightarrow{f} D.$$

The map $\tilde{X} \to X$ is the identity at all points of X except those which lie on L, at which point the fibres are $\mathbb{P}^1$. This is what algebraic geometers call a 'blow up of X by $\mathbb{P}^1$ at L', and for our purposes it suffices to know that provided L is transverse to X, $\tilde{X}$ is still smooth, and that the induced map $H^d(X, \mathbb{Q}_l) \to H^d(\tilde{X}, \mathbb{Q}_l)$ is injective.

The interesting map is the map $f : X \to \mathbb{P}^1$. Its fibres are the hyperplane sections $f^{-1}(d) = X \cap H_d$. We would expect most of these sections to be smooth (given X is smooth), but it also seems likely that some could have singularities. A *Lefschetz pencil* is such a fibration where all but finitely many of the hyperplane sections are smooth, and those which are not smooth just have one ordinary double point. We let $U \subset D$ be the smooth locus, and S the hyperplanes whose section contains a singular point. It is a fact that for any smooth projective variety, a Lefschetz pencil can be constructed (although in positive characteristic not necessarily for every embedding of the variety into projective space).

Let us now get back to studying the cohomology. We were interested in $H^d(X, \mathbb{Q}_l)$, but admit that it is hard to study, and have now embedded it into $H^d(\tilde{X}, \mathbb{Q}_l)$, which admits a map $\tilde{X} \to D$. Since we are having difficulty studying the derived functors of the pushforward along $X \to \operatorname{Spec} k$, the idea is to break the task up by factoring this map $(X \to)\tilde{X} \to D \to \operatorname{Spec} k$ and studying the derived functors from each pushforward individually and then combining them to give the pushforward of the composite. Given that 'derived functors' are a fairly complicated bit of homological algebra, and we are now in some sense studying 'derived functors squared', we get an even more complicated bit of homological algebra for working the details of this out, but it *is* pure homological algebra. To be more specific, there is a *Leray spectral sequence*

$$H^p(D, R^q f_* \mathcal{F}) \Rightarrow H^{p+q}(\tilde{X}, \mathcal{F}).$$

We shall explain this in more detail when we need it, but for now be convinced that knowing $H^p(D, R^q f_* \mathcal{F})$ for all p, q suffices to determine $H^r(\bar{X}, \mathcal{F})$ for all r.

Now, for studying cohomology with constant coefficients, the only really interesting term will be $H^1(D, R^n f_* \mathbb{Q}_l)$ where $n = d - 1$ is the dimension of the fibres, so we will be led to study the $\mathbb{Q}_l$-sheaf $R^n f_* \mathbb{Q}_l$ on D. By proper base change, $(R^n f_* \mathbb{Q}_l)_u \cong H^n(\tilde{X}_u, \mathbb{Q}_l)$, so studying these local systems is to study the cohomology of all fibres simultaneously. But by proper smooth base change, these groups will be isomorphic for all $u \in U$, so the only thing we don't know about this local system is what happens away from the smooth locus. Since all our singularities are of a nice form, this has a nice answer in terms of *vanishing cycles* (whose theory for l-adic cohomology is developed in detail in [6, 3]).

Theorem 4.6 (Vanishing cycles in a Lefschetz pencil)**.** *Let everything be as above, and fix $u \in U$. Also for technical reasons assume $\operatorname{char} k \neq 2$ and n is odd. Then there is a subspace $E \subset H^n(\tilde{X}_u, \mathbb{Q}_l)$ of 'vanishing cycles', which has the following properties.*

If $E = 0$, then the sheaves $R^i f_* \mathbb{Q}_l$ are constant for all $i \neq n+1$, and the $i = n+1$ sheaf fits into an exact sequence

$$0 \to \oplus_{s \in S} \mathbb{Q}_l(-\frac{n+1}{2})_s \to R^{n+1} f_* \mathbb{Q}_l \to \mathcal{G} \to 0,$$

where $\mathcal{G}$ is constant, and $(\mathbb{Q}_l)_s := s_*(\mathbb{Q}_l)$ is the direct image of a sheaf on a point (known as a skyscraper sheaf).

If $E \neq 0$ (as typically is the case),

1. *For $i \neq n$, $R^i f_* \mathbb{Q}_l$ is constant on D.*
2. *Letting $j : U \hookrightarrow D$, we have $R^n f_* \mathbb{Q}_l = j_* j^* R^n f_* \mathbb{Q}_l$.*
3. *The space E is stable under the action of $\pi_1(U, u)$, and the action on $H^n(\tilde{X}_u, \mathbb{Q}_l)/E$ is trivial.*
4. *Letting $E^\perp := H^n(\tilde{X}_u, \mathbb{Q}_l)^{\pi_1(U,u)}$, we have that the representation of $\pi_1(U, u)$ on $E/(E \cap E^\perp)$ is absolutely irreducible.*
5. *There are 'vanishing cycles exact sequences' involving the associated subsheaves $\mathcal{E}, \mathcal{E}^\perp$, which are given by either (if $E \subset E^\perp$)*

$$0 \to j_* \mathcal{E} \to R^n f_*(\mathbb{Q}_l) \to j_*(j^* R^n f_*(\mathbb{Q}_l)/\mathcal{E}^\perp) \to \oplus_{s \in S} \mathbb{Q}_l(-\frac{n+1}{2})_s \to 0,$$

or (if $E \not\subset E^{\perp}$) two short exact sequences

$$0 \to j_*\mathcal{E} \to R^n f_*(\mathbb{Q}_l) \to j_*(j^* R^n f_*(\mathbb{Q}_l)/\mathcal{E}) \to 0,$$

$$0 \to j_*(\mathcal{E} \cap \mathcal{E}^{\perp}) \to j_*\mathcal{E} \to j_*(\mathcal{E}/(\mathcal{E} \cap \mathcal{E}^{\perp})) \to 0.$$

6. *The cup product induces an alternating pairing $E \times E \to \mathbb{Q}_l(-n)$ with kernel $E \cap E^{\perp}$, so there is an induced perfect alternating $\pi_1(U_0, u)$-equivariant pairing*
$$\psi : \frac{E}{E \cap E^{\perp}} \times \frac{E}{E \cap E^{\perp}} \to \mathbb{Q}_l(-n)$$
and the map $\pi_1(U, u) \to Sp(E/(E \cap E^{\perp}), \psi)$ has open dense image.

That is quite a lot to take on board, and a much fuller explanation is provided in [4], together with references to the proofs. The significant facts for us are (assuming $E \neq 0$) that to study $R^n f_* \mathbb{Q}_l$ it suffices to consider its restriction to U (by 2), which by the monodromy correspondence 3.10 reduces us to studying $H^n(\tilde{X}_u, \mathbb{Q}_l)$ as a $\pi_1(U, u)$-representation, and by 3 and 4 the filtration

$$0 \subset E \cap E^{\perp} \subset E \subset H^n(\tilde{X}_u, \mathbb{Q}_l)$$

splits this representation into two trivial parts and one absolutely irreducible part. The exact sequences allow us to relate these representations to the local systems on $\mathbb{P}^1$, and the alternating pairing gives us valuable information for studying the nontrivial part.

So we now have a large kit of tools for computing cohomology groups, including the more difficult middle one. Let us now turn to the main theorem on endomorphisms, the Lefschetz Fixed Point formula, which in some sense is where our story really starts. It will be useful for us to have the notation that if $\phi : X \to X$ is an endomorphism of a variety, we write $Tr(\phi, H^i(X, \mathbb{Q}_l))$, $Det(\phi, H^i(X, \mathbb{Q}_l))$ to be the trace and determinant of the linear map induced by ϕ on the ith cohomology group (and similarly for H^i_c).

Theorem 4.7 (Lefschetz Fixed Point Formula)**.** *Let X be a smooth proper variety over an algebraically closed field k, and consider a morphism $\phi : X \to X$. Then*
$$(\Gamma_\phi . \Delta) = \sum_{i \geq 0} (-1)^i Tr(\phi, H^i(X, \mathbb{Q}_l)),$$
where $(\Gamma_\phi . \Delta)$ denotes the intersection number of the graph $\Gamma_\phi = \{(x, \phi(x)) | x \in X\}$ of ϕ with the image of the diagonal $\Delta : X \to X \times_k X$.

This is exactly the analogue one would expect from the classical theory, and for these traces to make sense it was crucial that $\mathbb{Q}_l$ had characteristic zero. In this section we have stated some of the main results of étale cohomology, which (I think) miraculously are true and analogous to results from classical topology. In fact, there is also a comparison theorem which tells us for example, that the Betti numbers we compute using étale cohomology and those from classical cohomology are the same, and in general that this theory really is very compatible with classical theories (in spite of the proofs of all these theorems looking quite different from their classical counterparts).

Of course, the strengths of the new theory are its application to positive characteristic, and its sensitivity to the Galois action. For the remainder of the essay, we will therefore focus on the case with simplest Galois action: varieties defined over finite fields, and follow Deligne in [4] to use the tools developed or stated so far to prove a subtle but important and very beautiful theorem about such varieties.

5 The Weil Conjecture

At the midpoint of the twentieth century, André Weil [22] observed fascinating patterns emerging amongst the number of points on varieties over a finite field. It is clear, by a direct counting argument, that projective n-space over $\mathbb{F}_q$ has $q^n + q^{n-1} + ... + 1$ $\mathbb{F}_q$-rational points, and by the time of Weil's paper it had been proved that for any smooth projective curve C of genus g over $\mathbb{F}_q$, one has the bound (on how far the number of points differs from that on $\mathbb{P}^1$)

$$||C(\mathbb{F}_{q^n})| - (1 + q^n)| \leq 2g\sqrt{q^n}.$$

For any variety X_0 over $\mathbb{F}_q$, Hasse and Weil had defined the *Zeta function* as the formal power series

$$\zeta(X_0, t) := \exp(\sum_{r\geq 1} |X_0(\mathbb{F}_{q^r})| \frac{t^r}{r}),$$

as a convenient way to store all the data of how many points X_0 has on each algebraic extension of $\mathbb{F}_q$.

To demonstrate how convenient this is (in spite of appearances), let us take our explicit calculations for $\mathbb{P}^n_{/\mathbb{F}_q}$ and substitute them in.

$$\begin{aligned} \log \zeta(\mathbb{P}^n_{/\mathbb{F}_q}, t) &= \sum_{r\geq 1} (1 + q^r + .. + q^{rn}) \frac{t^r}{r} \\ &= \sum_{r\geq 1} \frac{t^r}{r} + \sum_{r\geq 1} \frac{(qt)^r}{r} + ... + \sum_{r\geq 1} \frac{(q^n t)^r}{r} \\ &= -log(1 - t) - log(1 - qt) - ... - log(1 - q^n t). \end{aligned}$$

So we conclude

$$\zeta(\mathbb{P}^n_{/\mathbb{F}_q}, t) = \frac{1}{(1 - t)(1 - qt)...(1 - q^n t)}.$$

In particular, notice that $\zeta(\mathbb{P}^n_{/\mathbb{F}_q}, t)$ is a *rational function*. In particular, it defines an analytic function everywhere, and we know exactly where its zeroes and poles are. By using the theory of Tate modules of Jacobians, they also managed to prove a similar result for smooth projective curves C_0 of genus g: they proved that the Zeta function takes the form

$$\zeta(C_0, t) = \frac{P(t)}{(1-t)(1-qt)},$$

where $P(t)$ is a polynomial of degree $2g$ with rational coefficients and roots of absolute value $q^{-1/2}$. With these theorems, and partial results for higher dimensional varieties starting to be proved, Weil was led to the following amazing conjecture.

Theorem 5.1 (Weil Conjecture)**.** *Let X_0 be a smooth projective geometrically irreducible variety over $k = \mathbb{F}_q$ of dimension d. Then we have the following.*

- *The Zeta function is a rational function of the form*

$$\zeta(X_0, t) = \frac{P_1(t)...P_{2d-1}(t)}{P_0(t)...P_{2d}(t)}.$$

- *It satisfies a functional equation*

$$\zeta(X_0, \frac{1}{q^d t}) = \pm q^{-d\chi/2} t^{-\chi} \zeta(X_0, t),$$

  where $\chi = \sum_{i=0}^{2d} (-1)^i deg P_i$.

- *('Riemann Hypothesis') Each individual $P_i(t)$ is a polynomial with rational coefficients, and with roots of absolute value $q^{-i/2}$.*

- *If X_0 is the good reduction of a variety in characteristic zero, $deg P_i$ is equal to the ith Betti number of the original variety in classical cohomology.*

This conjecture is very strongly suggesting the existence of a cohomology theory for dealing with such varieties, and was the starting point for work on étale cohomology. In particular, it was noted that $X_0(\mathbb{F}_q)$ is precisely the set of fixed points of the *Frobenius endomorphism* $F : X_0 \to X_0$ induced by $(x_0 : ... : x_n) \mapsto (x_0^q : ... : x_n^q)$, so we should be able to take $P_i(t)$ to be the characteristic polynomial $Det(1 - tF, H^i(X, \mathbb{Q}_l))$ (where $X = X_0 \otimes_k \bar{k}$), and deduce the first part from the Lefschetz fixed point formula. With this in place, the second part should then also follow neatly from Poincaré duality. For the rest of this section, we shall show that this is the case, and remark on the generalisation of Zeta functions to L functions, where similar ideas will also be useful. We also note that the final part follows easily from the third part and the comparison theorems, so after this section the rest of this essay will be devoted to proving the third, most difficult, part.

The following was originally established by Dwork using p-adic methods (which can be very useful for getting more explicit results via computations with de Rham cohomology). However, we give Grothendieck's approach via étale cohomology, which seems to be the easiest tool for completing the rest of the proof. We shall need a slight generalisation of the trace formula in the case of the Frobenius morphism to cohomology with compact supports.

Theorem 5.2 (Grothendieck's Trace Formula). *Let X_0 be a variety over $\mathbb{F}_q$, and X its scalar extension to $\bar{\mathbb{F}}_q$. For any r, letting F^r denote the rth iterate of the Frobenius morphism, we have the trace formula:*

$$|X_0(\mathbb{F}_{q^r})| = \sum_i (-1)^i Tr(F^r, H^i_c(X, \mathbb{Q}_l)).$$

We shall use this to deduce the first part of the Weil Conjecture. More precisely:

Proposition 5.3. *Let X_0 be a variety over $\mathbb{F}_q$, and X its scalar extension to $\bar{\mathbb{F}}_q$. Then for $l \not| q$,*

$$\zeta(X_0, t) = \prod_i Det(1 - tF, H^i_c(X, \mathbb{Q}_l))^{(-1)^{i+1}}.$$

In particular it is a rational function (at least with $\mathbb{Q}_l$-coefficients).

Proof. This is a direct calculation, not unlike the calculation above for projective spaces.

$$\begin{aligned} \log \zeta(X_0, t) &= \sum_{r \geq 1} |X_0(\mathbb{F}_{q^r})| \frac{t^r}{r} \\ &= \sum_{r \geq 1} (\sum_i (-1)^i Tr(F^r, H^i_c(X, \mathbb{Q}_l))) \frac{t^r}{r} \\ &= \sum_i (-1)^i (\sum_{r \geq 1} Tr(F^r, H^i_c(X, \mathbb{Q}_l)) \frac{t^r}{r}) \end{aligned}$$

But for any endomorphism $\phi : V \to V$ of a finite dimensional vector space it is clear by putting in upper triangular form (allowing coefficients in an algebraic closure will not alter the trace or determinant) that the familiar one dimensional identity $-\log(1 - at) = at + a^2t^2/2 + a^3t^3/3 + ...$ generalises (taking a sum of such identities) to give

$$-\log \det(1 - t\phi) = \sum_{r \geq 1} \mathrm{tr}(\phi^r) \frac{t^r}{r},$$

whence the conclusion. □

Note that so far we have been talking about the Frobenius morphism, and although it looks superficially the same as the Frobenius map from Galois theory, we have not made explicit the relationship between them. For the rest of the essay it will be very important for us to have clear in our minds how the different kinds of Frobenius action are related. We just state what is true, and for a detailed discussion of these matters Deligne [4] refers the reader to [9, XV, 1,2].

Proposition 5.4 (Frobenius compatibilities)**.** *Let X_0 be an algebraic variety over $\mathbb{F}_q$ (with structure morphism $\pi : X_0 \to Spec\, \mathbb{F}_q$), and $\mathcal{F}_0$ be a smooth sheaf on X_0. Take their base changes $X, \mathcal{F}$ along a fixed algebraic closure $a : \mathbb{F}_q \to \bar{\mathbb{F}}_q$. Let $\phi \in Gal(\bar{\mathbb{F}}_q/\mathbb{F}_q)$ be the geometric Frobenius. The following three endomorphisms of $H^i_c(X, \mathcal{F})$ are in fact the same:*

1. *The map induced by the Frobenius morphism $F : X \to X$.*
2. *The map induced by the Galois endomorphism $X = X_0 \times_{Spec\, \mathbb{F}_q} Spec\, \bar{\mathbb{F}}_q \stackrel{id\times\phi}{\to} X_0 \times_{Spec\, \mathbb{F}_q} Spec\, \bar{\mathbb{F}}_q = X$.*
3. *(whenever applicable) The map induced by the monodromy action of ϕ under the base change isomorphism $(R^i\pi_!\mathcal{F}_0)_a \to H^i_c(X, \mathcal{F})$.*

Let x_0 be a closed point of X_0, corresponding to a closed (and geometric) point of X, and let $deg(x_0) = [k(x_0) : \mathbb{F}_q]$ be the number of $G_{\mathbb{F}_q}$-conjugate points corresponding to x_0. The stalk $\mathcal{F}_x = (\mathcal{F}_0)_x$ also carries an action by Frobenius which can be interpreted in the following (equivalent) ways:

1. *Fixing an embedding $x \to X$, the map induced by the Frobenius morphism $F^{deg(x_0)} : X \to X$.*
2. *Similarly, the map induced by the Galois endomorphism $X = X_0 \times_{Spec\, \mathbb{F}_q} Spec\, \bar{\mathbb{F}}_q \stackrel{id\times\phi^{deg(x_0)}}{\to} X_0 \times_{Spec\, \mathbb{F}_q} Spec\, \bar{\mathbb{F}}_q = X$.*
3. *If $Spec\, k(x_0) \to X_0$ is the embedding of our closed point, the monodromy action of the image of the geometric Frobenius under the induced map $G_{k(x_0)} \to \pi_1(X_0, x)$ on fundamental groups. In particular, note that the conjugacy class ϕ_x as we defined in section 2 will have the same characteristic polynomial as this actual Frobenius action.*

With this clear, it is fairly easy to deduce the second part of the Weil Conjecture from Poincaré duality (recalling that X_0 is projective, so $H_c = H$). Indeed, if α is an eigenvalue of Frobenius acting on $H^i(X, \mathbb{Q}_l)$, by Poincaré duality it is a Frobenius eigenvalue of $Hom_{\mathbb{Q}_l}(H^{2d-i}(X, \mathbb{Q}_l), \mathbb{Q}_l(-d))$. But recalling the definition of $\mathbb{Q}_l(1)$ (tensoring up $\mathbb{Z}_l(1) := \varprojlim \mu_{l^r}$), and noting that geometric Frobenius acts on lth roots of unity by $(\zeta \mapsto \zeta^q)^{-1}$, it acts on $\mathbb{Q}_l(-d)$ by multiplication by q^d. We therefore deduce that q^d/α is an eigenvalue of the Frobenius action on $H^{2d-i}(X, \mathbb{Q}_l)$. In other words, in the statement of the Weil conjecture, the roots of $P_i(t)$ are related to those of $P_{2d-i}(t)$ by $\alpha \mapsto q^d/\alpha$. The functional equation now follows by simple algebraic manipulations.

This is all well and good, and in some sense miraculous, but there is a significant unsatisfying detail in what we have proved so far. Everything in the above formulae has coefficients in $\mathbb{Q}_l$, and it is not clear that a different choice of l will not give some totally different-looking expression. Of course, in all the examples the coefficients are actually rational and independent of the choice of l. There are two obvious conjectures, and one is significantly stronger than the other. We could ask whether $\zeta(X_0, t)$ has rational coefficients, or more generally whether each individual $P_i(t)$ has rational coefficients and is independent of the choice of l. We are now in a position where proof of the former is relatively straightforward, and the latter statement can be reduced to a statement about absolute values of eigenvalues.

Proposition 5.5. *For any variety X_0 over a finite field, the zeta function $\zeta(X_0, t)$ is a rational function with $\mathbb{Q}$-coefficients.*

Proof. By definition $\zeta(X_0, t) = \sum_{i \geq 0} a_i t^i \in \mathbb{Z}[[t]]$, and by the rationality result established above, $\zeta(X_0, t) \in \mathbb{Q}_l(t)$. One can show purely algebraically that this implies that $det((a_{i+j+k})_{0 \leq i,j \leq M})(k > N)$ vanishes for all k and fixed M, N sufficiently large. However, one can also show that if such conditions hold, then actually $\sum_{i \geq 0} a_i t^i$ comes from an element of $\mathbb{Q}(t)$ (this algebraic trick is known as the study of 'Hankel determinants'). The result follows, and more generally we observe the useful fact that $\mathbb{Z}[[t]] \cap \mathbb{Q}_l(t) = \mathbb{Q}(t)$. □

So we now have in place most of the Weil conjectures. The key remaining ingredient is the following theorem.

Theorem 5.6 (Riemann Hypothesis for varieties over finite fields)**.** *Let X_0 be a smooth projective variety over a finite field, X its base change to an algebraic closure. Then all the eigenvalues of Frobenius acting on $H^i(X, \mathbb{Q}_l)$ have absolute value $q^{i/2}$ for every embedding $\bar{\mathbb{Q}}_l \hookrightarrow \mathbb{C}$.*

Indeed, it turns out this will suffice to prove the strong rationality result.

Proposition 5.7. *Under the Riemann Hypothesis, it follows that the characteristic polynomials $Det(1 - Ft, H^i(X, \mathbb{Q}_l))$ are coprime and each individually have rational coefficients.*

Proof. The roots of distinct such polynomials have distinct absolute values in every embedding $i : \bar{\mathbb{Q}}_l \to \mathbb{C}$, whence they are obviously coprime. We know that $\zeta(X_0, t)$ has rational coefficients, so is stable under $Gal(\bar{\mathbb{Q}}_l/\mathbb{Q})$. However, suppose α is a root of $Det(1 - Ft, H^i(X, \mathbb{Q}_l))$. Then for any $\sigma \in Gal(\bar{\mathbb{Q}}_l/\mathbb{Q})$, and $i : \bar{\mathbb{Q}}_l \to \mathbb{C}$, we have (by the Riemann Hypothesis) that

$$q^{-i/2} = |(i \circ \sigma)(\alpha)| = |i(\sigma\alpha)|,$$

and therefore $\sigma\alpha$ cannot lie amongst the roots of any of the other polynomials in the factorisation of ζ, so again must be a root of $Det(1 - Ft, H^i(X, \mathbb{Q}_l))$. We have thus shown that $Det(1 - Ft, H^i(X, \mathbb{Q}_l))$ is stable under $Gal(\bar{\mathbb{Q}}_l/\mathbb{Q})$, so it has rational coefficients. □

In the final part of this section, we discuss why the Riemann Hypothesis for finite fields has its name, define the L-functions attached to arbitrary smooth $\bar{\mathbb{Q}}_l$-sheaves, and state an important generalisation of (5.2), also due to Grothendieck. These will turn out to be important tools for proving (5.6).

The Zeta function we defined earlier can be re-written as a 'formal Euler product' as follows (writing $|X_0|$ for the set of closed points of X_0, and noting that each closed point x_0 gives $deg(x_0)$ $\mathbb{F}_q^{deg(x_0)}$-rational points).

$$\begin{aligned}\log \zeta(X_0,t) &= \sum_{r\geq 1} |X_0(\mathbb{F}_{q^r})| \frac{t^r}{r} \\ &= \sum_{x_0 \in |X_0|} (\sum_{r \geq 1} \frac{t^{rdeg(x_0)}}{r}) \\ &= \sum_{x_0 \in |X_0|} -\log(1-t^{deg(x_0)}).\end{aligned}$$

So in fact we get the Euler product

$$\zeta(X_0,t) = \prod_{x_0 \in X_0} (1-t^{deg(x_0)})^{-1}.$$

Making the substitution $t = q^{-s}$, we recover the general formula for the Zeta function of any scheme W of finite type over $\mathbb{Z}$ (where $N(x)$ is the size of the (finite) residue field at x):

$$\zeta(W,s) = \prod_{x \in |W|} (1-N(x)^{-s})^{-1}.$$

In particular, for $W = \text{Spec } \mathbb{Z}$ this is the classical Riemann Zeta function.

So (5.6) is (more or less) equivalent to the statement that when W is a projective smooth d-dimensional variety over a finite field, all the poles of $\zeta(W,s)$ lie on the lines $Res = 0, 1, ..., d$ and all the zeroes on the lines $Res = \frac{1}{2}, \frac{3}{2}, ..., \frac{2d-1}{2}$, so the analogy with the classical Riemann hypothesis is now clear.

With things in this form, we also now have clear generalisations of Artin L-functions. Recall that an Artin L-function is given by some finite dimensional Galois representation of a number field. In this context, we know quite a lot about analogous representations, namely that they correspond to smooth $\bar{\mathbb{Q}}_l$-sheaves, so it is natural to define (where x is always some geometric point at x_0, and this Frobenius action on stalks is that discussed in (5.4))

$$L(X_0,\mathcal{F}_0,t) := \prod_{x_0 \in X_0} det(1-t^{deg(x_0)}F^{deg(x_0)}, \mathcal{F}_x)^{-1}.$$

Just as in classical number theory, these obviously have the property that $\zeta(X_0,t) = L(X_0,\mathbb{Q}_l,t)$ (the zeta function is just the L-function of the trivial representation). They also have the 'direct sum of representations' property $L(X_0,\mathcal{F}\oplus\mathcal{G},t) = L(X_0,\mathcal{F},t)L(X_0,\mathcal{G},t)$, and given a G-Galois étale covering $Y_0 \to X_0$ we have that $\zeta(Y_0,t)$ is the L-function of the regular representation of G on X_0, and so get the corresponding decomposition into a product over all irreducible representations of G:

$$\zeta(Y_0,t) = \prod_{\rho} L(X_0,\mathcal{F}_{0,\rho},t)^{\dim \rho},$$

where $\mathcal{F}_\rho$ is the smooth $\bar{\mathbb{Q}}_l$ sheaf corresponding to the $\pi_1(X_0, x)$-representation obtained by factoring $\pi_1(X_0, x) \twoheadrightarrow G$ through ρ. Studying these functions will be an important first step in our proof of the Riemann Hypothesis (we will need estimates on the locations of the poles and zeroes of these functions in order to obtain the relevent Cebotarev density theorem).

Their study is made much easier by the following generalisation of the trace formula, also due to Grothendieck. Essentially, the Lefschetz fixed point formula can be generalised, at least in the case of the Frobenius morphism, and then formally rearranges as above to the following result.

Proposition 5.8 (Cohomological interpretation of L-functions)**.** *Let X_0 be a variety over $\mathbb{F}_q$, $\mathcal{F}_0$ a smooth $\bar{\mathbb{Q}}_l$ sheaf on it, with $l \not| q$, and $X, \mathcal{F}$ their scalar extensions to $\bar{\mathbb{F}}_q$. Then*

$$L(X_0, \mathcal{F}_0, t) = \prod_i Det(1 - tF, H^i_c(X, \mathcal{F}))^{(-1)^{i+1}}.$$

In particular it is a rational function with $\bar{\mathbb{Q}}_l$-coefficients.

In fact, the entire Weil conjecture admits a generalisation to L-functions of large classes of sheaves (including a 'generalised Riemann hypothesis'), also stated and proved by Deligne in a later paper, but this important result is now usually proved using the generalised Fourier transform, as detailed in [13, 1]. However, this lies beyond the scope of this essay.

At last we have set up all our machinery and stated our problem. We can now finally set out to prove the Riemann Hypothesis for varieties over a finite field. The proof will proceed in four steps. First, we study L-functions in much more depth, in particular establishing vanishing results analogous to those used in Dirichlet's theorem on primes lying in arithmetic progressions, and so deduce the Cebotarev density theorem, which tells us that the Frobenius conjugacy classes are actually dense in $\pi_1(U_0, x)$ for U_0 a curve, an important result known to Deligne which I view as in some sense the 'base case' of his proof. Secondly, we start to attack the problem properly, using an inductive argument and many of the tools from chapter 4 to reduce this large and scary question to an apparently much simpler question about the cohomology of a certain smooth $\mathbb{Q}_l$ sheaf on an open subset of the projective line. Thirdly, we use the Cebotarev density theorem and an intricate arithmetical argument to establish that the characteristic polynomials of Frobenius acting on the stalks of this sheaf actually have rational coefficients. Then finally we prove Deligne's 'main lemma' which shows that such a sheaf must have Frobenius acting on the stalks with eigenvalues of an appropriate absolute value and allows us to conclude the result.

6 The Cebotarev Density Theorem for Curves

In this section we use some elementary analytic arguments to establish a density theorem for the Frobenius conjugacy classes in $\pi_1(U_0, u)$, where U_0 is any geometrically irreducible smooth curve over $\mathbb{F}_q$. In spite of this theorem seeming to pay a key role in Deligne's proof, it seems to get little mention (in the form we need) in the literature, and Deligne gives no reference. The two references [11] and [17] which the author managed to find give respectively 'elementary' and 'analytic' proofs, both of which depend on (or at least are not obviously independent of) the Riemann Hypothesis for curves! Since it would be nice for our proof to also contain the curves case, following a useful discussion with Tony Scholl, the author will try to give a reasonably short analytic proof that uses some of the machinery of the previous chapters but is independent of the Riemann hypothesis for curves. Since this section is mostly 'original work' (albeit presumably well known to the experts), readers are warned that there is a higher likelihood of error here than elsewhere in the essay.

For this entire section, we let U_0 be a geometrically irreducible smooth curve over $\mathbb{F}_q$, equipped with u_0 an $\mathbb{F}_q$-rational point. Let U be its base change to $\bar{\mathbb{F}}_q$, and u the geometric point over u_0 (henceforth in the essay we shall continue to freely use this notation of a 0 subscript for objects over $\mathbb{F}_q$, which is dropped when passing to the algebraic closure). Recall that for any other closed point $x_0 \in |U_0|$ (with x a geometric point over it), the map on fundamental groups induced by inclusion of the residue field $G_{k(x_0)} \to \pi_1(U_0, x)$ composed with a path from x to u gives rise to a Frobenius conjugacy class $[\phi_{x_0}]$, which is independent of the choice of path.

We shall need a notion of the *density* of a set of closed points, and it will be useful if it relates nicely to the definitions of zeta and L-functions. The following classical definition used by Dirichlet in establishing the existence of infinitely many primes in arithmetic progressions is convenient. Let $A \subset |U_0|$. The *Dirichlet density* $\delta(A)$ is defined by

$$\delta(A) = \lim_{t \to q^{-1}} \frac{\sum_{x_0 \in A} t^{deg(x_0)}}{\sum_{x_0 \in |U_0|} t^{deg(x_0)}}.$$

Note that for $|t| < q^{-1}$ the infinite sums converge absolutely and uniformly because by Noether normalisation $|U_0(\mathbb{F}_{q^r})| = O(q^r)$. The limit is always taken from the disc $|t| < q^{-1}$ to the point $t = q^{-1}$, and we shall not worry about whether the limit always exists or is well-defined for general A: in all cases we consider we shall see that it does.

Our main result is the following.

Theorem 6.1 (Cebotarev Density theorem, finite form)**.** *Let $V_0 \to U_0$ be a Galois (finite) étale cover, with automorphism group G (so there is a map $p : \pi_1(U_0, u) \twoheadrightarrow G$). Fix a conjugacy class $C \subset G$, and let $A = \{x_0 \in |U_0| : p([\phi_{x_0}]) \subseteq C\}$. Then*

$$\delta(A) = \frac{|C|}{|G|}.$$

We shall prove this fact by studying the family of L-functions mentioned briefly in the last section. Let ρ be an irreducible $\bar{\mathbb{Q}}_l$-representation of G. This gives rise to a finite image irreducible representation of $\pi_1(U_0, u)$, which corresponds to a smooth $\bar{\mathbb{Q}}_l$-sheaf $\mathcal{F}_{0,\rho}$ on U_0, under the $\bar{\mathbb{Q}}_l$-monodromy correspondence (3.11). We therefore obtain an L-function $L(U_0, \mathcal{F}_{0,\rho}, t)$, which is the analogue of the classical Artin L-function, and so is the key to the analogue of the classical density theorem.

Indeed, by the same trace-determinant switching trick as in the proof of (5.3),

$$\begin{aligned}
\log L(U_0, \mathcal{F}_{0,\rho}, t) &= \sum_{x_0 \in |U_0|} -\log\det(1 - t^{deg(x_0)} F^{deg(x_0)}, \mathcal{F}_{\rho,x}) \\
&= \sum_{x_0 \in |U_0|} \sum_{r \geq 1} Tr(F^{deg(x_0)r}, \mathcal{F}_{\rho,x}) \frac{t^{deg(x_0)r}}{r} \\
&= \sum_{x_0 \in |U_0|} \sum_{r \geq 1} \chi(p(\phi_{x_0})) \frac{t^{deg(x_0)r}}{r}
\end{aligned}$$

where χ is the character corresponding to the representation ρ.

By basic finite representation theory, we know that the characters of the irreducible $\bar{\mathbb{Q}}_l$ representations are a basis for the space of class functions $ccl(G) \to \bar{\mathbb{Q}}_l$. In particular, the indicator function of C can be written as a sum of characters of irreducible representations:

$$1_C = \lambda_1\chi_1 + \lambda_2\chi_2 + ... + \lambda_k\chi_k.$$

Hence,

$$\begin{aligned}
\sum_{x_0 \in A} t^{deg(x_0)} &= \sum_{x_0 \in |U_0|} 1_C(p(\phi_{x_0})) t^{deg(x_0)} \\
&= \sum_{x_0 \in |U_0|} (\lambda_1\chi_1(p(\phi_{x_0})) + ... + \lambda_k\chi_k(p(\phi_{x_0}))) t^{deg(x_0)}.
\end{aligned}$$

The key observation is that this gives us an expression for the Dirichlet density of A in terms of a sum of leading order terms of the L-functions corresponding to each character. Observe further that if χ_1 is the trivial representation,

$$\lambda_1 = \langle \chi_1, 1_C \rangle = \frac{1}{|G|} \sum_{g \in G} 1_C(g) = \frac{|C|}{|G|},$$

so actually

$$\sum_{x_0 \in A} t^{deg(x_0)} - \frac{|C|}{|G|} \sum_{x_0 \in |U_0|} t^{deg(x_0)} = \sum_{x_0 \in |U_0|} (\lambda_2\chi_2(p(\phi_{x_0})) + ... + \lambda_k\chi_k(p(\phi^r_{x_0}))) t^{deg(x_0)}.$$

It will therefore suffice to prove that the RHS is small compared with $\sum_{x_0 \in |U_0|} t^{deg(x_0)}$ as $t \to q^{-1}$, and to do this we will relate it to the value of the L-functions at $t = q^{-1}$ (by bounding the higher order terms) and then use Grothendieck's trace formula, some simple cohomological arguments, and a further analytic trick to prove that the L-functions do not have a pole or zero at $t = q^{-1}$ (so their logarithm is bounded), in contrast to $\sum_{x_0 \in |U_0|} t^{deg(x_0)}$, which we will show tends to infinity as $t \to q^{-1}$.

Lemma 6.2. *The expressions*

$$|\log L(U_0, \mathcal{F}_{0,\rho}, t) - \sum_{x_0 \in U_0} \chi(p(\phi_{x_0}))t^{deg(x_0)}| = |\sum_{x_0 \in U_0} \sum_{r \geq 2} \chi(p(\phi_{x_0}^r)) \frac{t^{deg(x_0)r}}{r}|$$

are bounded as $t \to q^{-1}$ (in fact are absolutely convergent with a fixed bound).

Proof. Again, using Noether normalisation we get for our variety U_0 a constant C such that $|U_0(\mathbb{F}_{q^s})| \leq Cq^s$. Also since G is finite, the irreducible characters satisfy the trivial bound $|\chi(p(\phi_{x_0}^r))| \leq |G|$.

Hence for any fixed $r \geq 2$, and for $|t| \leq q^{-1}$ we have the crude estimate

$$\sum_{x_0 \in U_0} |\chi(p(\phi_{x_0}^r)) \frac{t^{deg(x_0)r}}{r}| \leq C|G| \sum_{s \geq 1} (q^s t^{rs}) = C|G| \frac{qt^r}{1 - qt^r} \leq \frac{C|G|q}{1 - q^{-1}} t^r.$$

And now summing over such expressions, the required boundedness is clear. □

It therefore remains to prove the following.

Theorem 6.3 ('$L(1, \chi) \neq 0$'). *The functions $L(U_0, \mathcal{F}_{0,\rho}, t)$ (as rational functions with $\bar{\mathbb{Q}}_l$-coefficients), for ρ an irreducible representation of G, have neither a zero nor a pole at $t = q^{-1}$, unless ρ is the trivial representation in which case there is a simple pole.*

Grothendieck's formula (5.8) tells us that these are rational functions with $\bar{\mathbb{Q}}_l$ coefficients and that to gain more information about them it suffices to study the groups $H_c^i(U, \mathcal{F}_\rho)$. Since U is a curve, it is easy to explicitly compute these groups for $i \neq 1$, so we can prove the following.

Lemma 6.4. *Let ρ be an irreducible representation of G. Then*

$$L(U_0, \mathcal{F}_{0,\rho}, t) = \frac{P_\rho(t)}{Q_\rho(t)},$$

where $P_\rho(t)$ is some $\bar{\mathbb{Q}}_l$-polynomial, and we can explicitly identify

$$Q_\rho(t) = \begin{cases} 1 & \text{if } \rho \text{ nontrivial,} \\ 1 - qt & \text{if } U \text{ is affine, and } \rho \text{ is trivial,} \\ (1-t)(1-qt) & \text{if } U \text{ is complete, and } \rho \text{ is trivial.} \end{cases}$$

Proof. We simply compute the (at most one-dimensional) outer cohomology groups. Firstly, since $\mathcal{F}_\rho$ is locally constant and U geometrically irreducible, it is clear that $H^0_c(U, \mathcal{F})$ is trivial unless U is complete (for any compactification j, a global section of $j_!\mathcal{F}$ will be locally zero, and so globally zero).

Now, note further that by the monodromy correspondence (and that clearly the trivial representation corresponds to the constant $\mathbb{Q}_l$-sheaf), we have that $H^0(U, \check{\mathcal{F}}_\rho) = (\check{\mathcal{F}}_\rho)_u^{\pi_1(U,u)}$, which (since $\mathcal{F}_\rho$ is irreducible) vanishes unless ρ is trivial, so by Poincaré duality $H^2_c(U, \mathcal{F})$ is trivial in this case. In the case where ρ is trivial, of course $H^0(U, \mathbb{Q}_l) = \mathbb{Q}_l$, and Poincaré duality gives $H^2_c(U, \mathcal{F}) = \mathbb{Q}_l(-1)$.

Putting all these observations together (and by one more application of Poincaré duality or a direct computation for the U complete case), we deduce the lemma from Grothendieck's rationality formula (5.8). $\square$

Corollary 6.5. *If ρ is a nontrivial irreducible representation, then $L(U_0, \mathcal{F}_{0,\rho}, t)$ has no poles, in particular no pole at $t = q^{-1}$.*

Now, recall the product formula we mentioned at the end of the previous chapter, which in the context of our Cebotarev density theorem for the Galois covering $V_0 \to U_0$ is the following (the product is over all irreducible representations of G):

$$\zeta(V_0, t) = \prod_\rho L(U_0, \mathcal{F}_{0,\rho}, t)^{\dim \rho},$$

which follows formally from the (easily checkable: see [17]) fact that, for such a Galois cover, $\zeta(V_0, t)$ is the L-function over U_0 of the regular representation of G. Combining this with our previous lemma (6.4), the result has been further reduced to proving the following proposition.

Proposition 6.6. *For any smooth geometrically irreducible curve U_0 over $\mathbb{F}_q$, $\zeta(U_0, t)$ has a simple pole at $t = q^{-1}$.*

Indeed, if this proposition is true, then no L-function in the above factorisation can possibly have a zero at $t = q^{-1}$, since by our previous lemma the product contains at most one factor of $(1 - qt)$ in the denominator, namely that which comes from $\zeta(U_0, t)$.

Now we have reduced the problem to ζ-functions it really feels much more tractable, since we can avoid thinking about the middle cohomology group explicitly and instead investigate it via the entire zeta function and its interpretation as counting $\mathbb{F}_{q^r}$-rational points. As we shall see, this is also an important strategy in several steps of Deligne's proof. It will be convenient to have the following notation: let $\epsilon = 1$ if U_0 is complete, and $\epsilon = 0$ otherwise.

By (6.4), let us write

$$(1 - t)^\epsilon (1 - qt)\zeta(U_0, t) = P_1(t) = (1 - \alpha_1 t)...(1 - \alpha_k t),$$

obviously a polynomial with rational coefficients (since the LHS is). We are required to prove that it is not possible for some α_i to be equal to q. In fact, we can prove something more general.

Proposition 6.7 ('Prime number theorem'). *For any smooth geometrically irreducible curve U_0 over $\mathbb{F}_q$, the rational polynomial $P_1(t)$ has no zeroes inside the closed disc $|t| \leq q^{-1}$. In other words, $\zeta(t)$ is holomorphic and nowhere vanishing on the closed disc $|t| \leq q^{-1}$ save for a simple pole at $t = q^{-1}$.*

This will be a relatively easy consequence of the following analytic/combinatorial lemma. Note that since everything in sight is algebraic, it may be possible to replace this lemma by thinking about how roots of rational polynomials must be algebraic numbers and thus not give rise to complex numbers with an irrational argument, but since the lemma is quite easy and pretty, we proceed this way.

Lemma 6.8. *Let $\beta_1, ..., \beta_l$ be complex numbers of unit modulus that are roots of a real polynomial. For any $r_0 \in \mathbb{N}$, there exists $r \geq r_0$ with the property that for each conjugate pair β_i, β_j (including the trivial case $i = j$)*

$$\beta_i^r + \beta_j^r \geq \frac{3}{2}.$$

Proof. Let $\theta = (\theta_1, ..., \theta_l)$ be their arguments, considered as an element of the space $(\mathbb{R}/2\pi\mathbb{Z})^l$. For some large integer k, partition this space into k^l small hypercubes $C_u := [\frac{2\pi u_1}{k}, \frac{2\pi(u_1+1)}{k}) \times [\frac{2\pi u_2}{k}, \frac{2\pi(u_2+1)}{k}) \times ... \times [\frac{2\pi u_l}{k}, \frac{2\pi(u_l+1)}{k})$.

Consider the vectors $r_0\theta, 2r_0\theta, ..., (k^l + 1)r_0\theta$. By the pigeonhole principle, two of them must land in the same hypercube, and taking their difference we get some $r\theta \in [\frac{-2\pi}{k}, \frac{2\pi}{k}] \times [\frac{-2\pi}{k}, \frac{2\pi}{k}] \times ... \times [\frac{-2\pi}{k}, \frac{2\pi}{k}]$, where r is a positive multiple of r_0.

This implies that

$$\beta_i^r + \beta_j^r \geq 2 \cos \frac{2\pi}{k}.$$

For k sufficiently (not very) large, this proves the result. □

Proof of (6.7). Note that by the definition of $\zeta(U_0, t)$, it is easy to see that

$$|U_0(\mathbb{F}_{q^r})| = \epsilon + q^r - (\alpha_1^r + ... + \alpha_k^r). \qquad (*)$$

Suppose for contradiction that $P_1(t)$ does have a zero in the closed disc $|t| \leq q^{-1}$. Wlog, let $|\alpha_1| \geq q$. By the lemma, we can choose arbitrarily large r such that the sum of each conjugate pair of terms satisfies $\alpha_i^r + \alpha_j^r \geq \frac{3}{2}|\alpha_i|^r$. If α_1 is complex, we have $\alpha_1^r + \bar{\alpha}_1^r \geq \frac{3}{2}q^r$, so $|U_0(\mathbb{F}_{q^r})| < 0$, which is absurd. But even if α_1 is real, we still deduce that

$$|U_0(\mathbb{F}_{q^r})| = \epsilon - (\alpha_2^r + ... + \alpha_k^r) \leq 1,$$

for arbitrarily large r, which is also absurd. This contradiction proves the result. □

With this proved, we have at last established all the theorems in this section, including the finite Cebotarev density theorem. It is worth noting that we can combine the finite statements for *all* Galois covers into the following statement about $\pi_1(U_0, u)$.

Theorem 6.9 (Cebotarev Density theorem, profinite version). *Let $C \subseteq \pi_1(U_0, u)$ be an open subset closed under conjugation, and $A = \{x_0 \in |U_0| : [\phi_{x_0}] \subseteq C\}$. Then $\delta(A)$ is equal to the normalised Haar measure of C. In particular, the elements of Frobenius conjugacy classes are dense in $\pi_1(U_0, u)$.*

Proof. Firstly, let us remark on the behaviour of the Haar measure. Since there are at most countably many finite Galois covers of U_0 (they are each given by finitely many algebraic equations with a finite coefficient ring in some finite number of variables in at least one way), the basis of open sets given by cosets of fundamental groups of finite Galois covers for $\pi_1(U_0, u)$ is countable, so we can write C as a countable union $C = \bigcup_{n=1}^{\infty} F_n$, where each open $F_n = C_n + G_n$ is a conjugacy class of cosets of a fundamental group G_n of some finite Galois cover of U_0. Since the Haar measure μ is countably additive,

$$\mu(C) = \lim_{N\to\infty} \mu(\bigcup_{n=1}^{N} F_n).$$

But we note that $\bigcup_{n=1}^{N} F_n$ is a finite union of some number k_N cosets of the intersection $H_N = \bigcap_{i=1}^{N} G_n$, which is also a Galois group of some $V_{N,0} \to U_0$, hence a normal open subgroup of $\pi_1(U_0, u)$. Considering the projection $p_N : \pi_1(U_0, u) \twoheadrightarrow Aut(V_{N,0}/U_0)^{op}$, and noting the image of $\bigcup_{n=1}^{N} F_n$ must be closed under conjugation by normality of H_N, we can apply the finite Cebotarev density theorem to deduce that

$$\delta(\{x_0 \in |U_0| : [\phi_{x_0}] \subseteq \bigcup_{n=1}^{N} F_n\}) = \frac{k_N}{|Aut(V_{N,0}/U_0)|} = \mu(\bigcup_{n=1}^{N} F_n).$$

Letting $N \to \infty$ and using the previous remarks, we recover the result. □

7 Geometrical Reductions

We now begin our attack on the Riemann Hypothesis for a smooth projective variety X_0 over a finite field (5.6), following Deligne [4]. Note that we shall make constant use of the easily checked fact that the category of $G_{\mathbb{F}_q}$-representations whose Frobenius eigenvalues are confined to lie within a certain set is abelian (so kernels, cokernels, direct sums, taking cohomology of complexes, etc. all preserve the set in which the Frobenius eigenvalues can lie). In this section we will essentially be following the classical strategy outlined in section 4 for reducing the computation of cohomology groups of complex varieties to the study of a local system on $\mathbb{P}^1$ and cohomology groups of varieties of smaller dimension. However, we have only to keep track of the Frobenius eigenvalues, so much of this will be computationally easy.

It is clear that the statement of the Riemann hypothesis is unchanged under finite extensions of the base field (the Frobenius element and the number q are simultaneously taken to some higher power). The cohomology also splits into a sum over irreducible components which can be checked individually. Putting these together we may assume X_0 is absolutely irreducible, in particular of some pure dimension d, and we may proceed by induction on d.

By a further finite extension of the base field if necessary, we may take a smooth hyperplane section Y_0 defined over $\mathbb{F}_q$, and by the Weak Lefschetz Theorem, $H^i(X, \mathbb{Q}_l) \hookrightarrow H^i(Y, \mathbb{Q}_l)$ for $i = 0, 1, ..., d-1$, so the Riemann hypothesis for these groups follows by induction on dimension (noting that the base case of finite sets of points is trivial). Also, Poincaré duality now allows us to deduce the Riemann hypothesis for the groups $H^{2d-i}(X, \mathbb{Q}_l) : i = 0, 1, ..., d-1$. It therefore remains to study the group $H^d(X, \mathbb{Q}_l)$.

Everything so far has been fairly routine, but the next argument is more subtle. It will be convenient much later in the proof for us to have some 'breathing space', so rather than proving the eigenvalues have absolute value $q^{d/2}$ on the nose we will only be able to show their absolute values lie in a small interval $[q^{d/2-C}, q^{d/2+C}]$ where C is an absolute constant. We also adopt a strategy that only works for even dimensional varieties. The next argument shows that actually this is enough to deduce the full result.

Proposition 7.1. *It suffices to prove that for all even-dimensional geometrically irreducible smooth projective varieties Y_0 there is an absolute constant C such that the eigenvalues α of Frobenius acting on $H^d(Y, \mathbb{Q}_l)$ satisfy*

$$q^{d/2-C} \leq |\alpha| \leq q^{d/2+C}.$$

Proof. The idea is to use the Kunneth formula with $Y_0 = X_0 \times X_0 \times ... \times X_0$ where the product is taken $2k$ times. By the Kunneth formula, we have the embedding (as a direct summand)

$$H^d(X, \mathbb{Q}_l)^{\otimes 2k} \hookrightarrow H^{2kd}(Y, \mathbb{Q}_l).$$

So if α is a Frobenius eigenvalue for $H^d(X, \mathbb{Q}_l)$, by hypothesis we get

$$q^{kd-C} \leq |\alpha|^{2k} \leq q^{kd+C}.$$

Thus

$$q^{d/2-C/2k} \leq |\alpha| \leq q^{d/2+C/2k},$$

and letting $k \to \infty$, we deduce $|\alpha| = q^{d/2}$.

The additional step that we may assume Y_0 is geometrically irreducible can be reached as above for X_0. □

This idea of taking a large product to get more manouvring space will be even more significant in the next section (where it will be vital to take large tensor powers of representations). Armed with the proposition, we can continue the classical algorithm for computing cohomology (but now using our even-dimensional Y_0). Again, we shall prove the claim in the proposition by induction on (even) dimension d with the fixed value $C = \frac{1}{2}$.

We imitate the classical trick, and form a Lefschetz pencil $Y \leftarrow \tilde{Y} \xrightarrow{f} \mathbb{P}^1$. Recall from section 4 that this involves fixing an embedding of Y in projective space and a codimension 2 linear subspace A transverse to Y, and the map $\tilde{Y} \to \mathbb{P}^1$ has fibres equal to hyperplane sections of hyperplanes passing through A. These are smooth over some open subset $U \subset \mathbb{P}^1$, and have a single ordinary double point over the closed $S = \mathbb{P}^1 - U$. Since $H^d(Y, \mathbb{Q}_l) \hookrightarrow H^d(\tilde{Y}, \mathbb{Q}_l)$, it will suffice to consider the latter.

Recall that these Lefschetz pencils always exist (for a suitable embedding in projective space) and taking another extension of the base field we may assume that the spaces A, S, the sheaves $\mathcal{E}, \mathcal{E}^{\perp}$ and the singular points of fibres over S are in fact defined over $\mathbb{F}_q$, and that there is an $\mathbb{F}_q$-rational basepoint $u_0 \in U_0$, whose fibre $f_0^{-1}(u_0) \subset \tilde{Y}_0$ admits a smooth hyperplane section Z_0 defined over $\mathbb{F}_q$. Note that $\tilde{Y}_0$ is d-dimensional, $f_0^{-1}(u_0)$ $d-1$-dimensional so Z_0 is $d-2$-dimensional (so in particular we will be able to apply the induction hypothesis to its cohomology).

We now call in the Leray spectral sequence

$$H^p(\mathbb{P}^1, R^q f_* \mathbb{Q}_l) \Rightarrow H^{p+q}(\tilde{Y}, \mathbb{Q}_l).$$

A spectral sequence is an algorithm which takes the data (E^{pq}) on the left hand side, and at each stage of the algorithm considers E^{pq} as part of a chain complex $\ldots \to E' \to E^{pq} \to E'' \to \ldots$ and replaces E^{pq} with its cohomology in this complex: so takes a subspace and then quotients it out by a smaller subspace. As we already remarked, such an operation certainly preserves the property of having eigenvalues of Frobenius in the range required. In this case, the sequence will halt: after some finite number of such operations, one can recover

$$H^r(\tilde{Y}, \mathbb{Q}_l) = \oplus_{p+q=r} E^{pq}.$$

Therefore, passing through the spectral sequence, it will suffice for us to prove that the terms $H^p(\mathbb{P}^1, R^{d-p} f_* \mathbb{Q}_l)$ have Frobenius eigenvalues in the desired range. By an appropriate vanishing theorem[2], all terms for $p > 2$ vanish, so we are left to consider three individual terms. As always, the middle one is where the real problem lies. As you might expect, the rest of this chapter will rely firmly on properties of vanishing cycles (4.6), and we set $n = d - 1$ to be the dimension of the hyperplane sections, and note that it is odd (which was important to get our statement of the properties of vanishing cycles to work).

Proposition 7.2. *The groups $H^2(\mathbb{P}^1, R^{n-1} f_* \mathbb{Q}_l)$ and $H^0(\mathbb{P}^1, R^{n+1} f_* \mathbb{Q}_l)$ both have Frobenius eigenvalues with absolute values in the required range.*

Proof. First we consider $H^2(\mathbb{P}^1, \mathbb{Q}_l)$. By (4.6), $R^{n-1} f_* \mathbb{Q}_l$ is constant, and by evaluating it on the fibre over u (via proper base change) and using that $H^2(\mathbb{P}^1, \mathbb{Q}_l) = \mathbb{Q}_l(-1)$ (if you like, by Poincaré duality), we deduce that

$$H^2(\mathbb{P}^1, R^{n-1} f_* \mathbb{Q}_l) = H^{n-1}(Y_u, \mathbb{Q}_l)(-1).$$

[2]Strictly speaking we need more than (3.6). We need a vanishing theorem for 'constructible sheaves', which are more general than lcc sheaves, but in fact the sources we gave for the proof of (3.6) in fact actually prove this more general case.

And by Weak Lefschetz this embeds into $H^{n-1}(Z, \mathbb{Q}_l)(-1)$, so we are done by induction.

Now let us deal with $H^0(\mathbb{P}^1, R^{n+1}f_*\mathbb{Q}_l)$. Here it is important to split into the cases $E \neq 0$ and $E = 0$. In the former case, we know again that $R^{n+1}f_*\mathbb{Q}_l$ is constant, so the result follows similarly to that above by induction and the surjectivity of the Gysin morphism (4.5) $H^{n-1}(Z, \mathbb{Q}_l)(-1) \to H^{n+1}(Y_u, \mathbb{Q}_l)$.

If $E = 0$, we must be more careful, since now $R^{n+1}f_*\mathbb{Q}_l$ is not constant, but taking the stalk at u of the exact sequence it fits into in (4.6), one gets (from the long exact sequence of cohomology) an exact piece

$$\oplus_{s\in S}\mathbb{Q}_l(-d/2) \to H^0(\mathbb{P}^1, R^{n+1}f_*\mathbb{Q}_l) \to H^{n+1}(Y_u, \mathbb{Q}_l).$$

The left term obviously has Frobenius eigenvalues $q^{d/2}$, and the right term can be dealt with as in the previous case. □

It now remains to consider the middle term $H^1(\mathbb{P}^1, R^n f_*\mathbb{Q}_l)$. In the case where $E = 0$ this is easy because again $R^n f_*\mathbb{Q}_l$ is constant, and we know (or it is very easy to check, perhaps most quickly via Grothendieck's formula) that $H^1(\mathbb{P}^1, \mathbb{Q}_l) = 0$, so in fact this term vanishes! The crux of the problem is when there are vanishing cycles, and when there is some corresponding nontrivial geometrical monodromy on the local system $R^n f_*\mathbb{Q}_l$. It might be that the latter fails to happen, in which case $E \subset E^\perp$, but we shall soon see that this case is also easy to deal with.

It will now help for us to replace these $\pi_1(U, u)$-subrepresentations $E, E^\perp$ of $H^n(Y_u, \mathbb{Q}_l)$ by the corresponding smooth $\mathbb{Q}_l$-subsheaves $\mathcal{E}, \mathcal{E}^\perp$ of $j^*R^n f_*\mathbb{Q}_l$ (which itself is a smooth $\mathbb{Q}_l$-sheaf by (3.7)). Since $R^n f_*\mathbb{Q}_l = j_*j^*R^n f_*\mathbb{Q}_l$, these allow us to get a filtration (whose usefulness was remarked upon during the discussion of vanishing cycles)

$$0 \subseteq j_*(\mathcal{E} \cap \mathcal{E}^\perp) \subseteq j_*\mathcal{E} \subseteq R^n f_*\mathbb{Q}_l.$$

In the degenerate case $E \subset E^\perp$, the vanishing cycles exact sequence gives us two short exact sequences, setting

$$\mathcal{F} = Coker(j_*\mathcal{E}^\perp \hookrightarrow R^n f_*\mathbb{Q}_l) = Ker(j_*(j^*R^n f_*(\mathbb{Q}_l)/\mathcal{E}^\perp) \twoheadrightarrow \oplus_{s\in S}\mathbb{Q}_l(-\frac{d}{2})_s).$$

Since $j_*\mathcal{E}^\perp$ and $j_*(j^*R^n f_*(\mathbb{Q}_l)/\mathcal{E}^\perp)$ are both constant, we deduce (from long exact sequences on cohomology) that $H^1(\mathbb{P}^1, \mathcal{F})$ is the image of a representation which trivially has Frobenius eigenvalues $q^{d/2}$, and that $H^1(\mathbb{P}^1, R^n f_*\mathbb{Q}_l) \hookrightarrow H^1(\mathbb{P}^1, \mathcal{F})$, whence the required result.

When this does not happen, we can use the two vanishing cycles short exact sequences, together with the facts that $j_*(\mathcal{E}\cap\mathcal{E}^\perp)$ and $j_*(j^*R^n f_*\mathbb{Q}_l/\mathcal{E})$ are both constant to show that

$$H^1(\mathbb{P}^1, j_*\mathbb{E}) \twoheadrightarrow H^1(\mathbb{P}^1, R^n f_*\mathbb{Q}_l) \text{ ,and } H^1(\mathbb{P}^1, j_*\mathbb{E}) \hookrightarrow H^1(\mathbb{P}^1, j_*(\mathcal{E}/(\mathcal{E} \cap \mathcal{E}^\perp))).$$

So the entire problem now rests on the group $H^1(\mathbb{P}^1, j_*(\mathcal{E}/(\mathcal{E} \cap \mathcal{E}^\perp)))$. The sheaf $j_*(\mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp))$ still need not be smooth, so we do one final step to reduce the problem to that of a smooth sheaf. Firstly note that by Poincaré duality, if α is a Frobenius eigenvalue for $H^1(\mathbb{P}^1, j_*(\mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp)))$, then so is q^{n+1}/α (using the bilinear pairing ψ from (4.6) to identify the dual sheaf and in particular the factor q^n), so we need only bound the eigenvalues in one direction.

We then note that the embedding $j_!(\mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp)) \hookrightarrow j_*\mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp)$ has a cokernel which is supported on a zero dimensional variety, so by a vanishing theorem (and perhaps a look at the first couple of terms of a Leray spectral sequence) this does not admit any H^1, and thus the long exact sequence of cohomology gives us

$$H^1_c(U, \mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp)) \twoheadrightarrow H^1(\mathbb{P}^1, j_*(\mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp))).$$

So at last we have reduced our problem to studying a smooth $\mathbb{Q}_l$ sheaf defined on a curve. Let us summarise what we have shown.

Proposition 7.3. *To prove the Riemann Hypothesis, it will suffice to show that for any embedding $\bar{\mathbb{Q}}_l \hookrightarrow \mathbb{C}$, the eigenvalues α of Frobenius acting on $H^1_c(U, \mathcal{E}/(\mathcal{E}\cap\mathcal{E}^\perp))$ satisfy*

$$|\alpha| \le q^{\frac{d+1}{2}}.$$

We now write $\mathcal{F}_0 = \mathcal{E}_0/(\mathcal{E}_0\cap\mathcal{E}_0^\perp)$ for the sheaf in question, a smooth $\mathbb{Q}_l$-sheaf on U_0. Recall that studying smooth $\mathbb{Q}_l$ sheaves is equivalent to studying $\pi_1(U_0, u)$-representations, and that is now the direction in which our attention turns. What do we know about the corresponding representation $V = \frac{E}{E\cap E^\perp}$?

By the vanishing cycles theorem (4.6), we know that there is an alternating perfect pairing

$$\psi : V \times V \to \mathbb{Q}_l(-n),$$

which is preserved by the geometric component of the monodromy $\pi_1(U, u)$, and conversely that almost every linear map preserving this pairing comes from geometric monodromy. More precisely, $\pi_1(U, u)$ has dense image in $Sp(V, \psi)$. These are both quite strong properties, and tell us a lot about the representation, but there is still obviously one key ingredient missing. We have remarked previously that the Riemann Hypothesis implies that the characteristic polynomial of Frobenius acting on $H^d(X, \mathbb{Q}_l)$ has *rational* coefficients, in particular its eigenvalues are algebraic integers. It is to the 'algebraic' nature of the representation V which we now turn.

8 The Rationality Theorem

In this section we shall go back and examine features of our geometrical argument to establish the following 'arithmetical' theorem about how the Frobenius elements act on V.

Theorem 8.1 (Rationality theorem)**.** *Let V be the $\pi_1(U_0, u)$-representation defined at the end of the previous section. The Frobenius conjugacy classes $[\phi_x] \subset \pi_1(U_0, u)$ all have characteristic polynomials with rational coefficients.*

Equivalently, the polynomials $Det(1 - F^{\deg(x_0)}t, \mathcal{F}_x)$ *all have rational coefficients.*

To prove this, we will go back to viewing $\mathcal{F}_x$ as part of a filtration of $H^n(Y_x, \mathbb{Q}_l)$ and the strategy will be to deduce the rationality of the relevent characteristic polynomial from the rationality of the zeta function of each (smooth projective) fibre Y_x, and by establishing the rationality of every other factor in the zeta function except that which corresponds to $\mathcal{F}_x$.

The key is the following lemma, which follows easily from the monodromy theory we developed in the opening chapters.

Lemma 8.2. *Let* $\mathcal{G}_0$ *be a sheaf on* U_0 *such that its base change* $\mathcal{G}$ *to* U *is constant. Then the action* $\pi_1(U_0, u) \to GL(\mathcal{G}_u)$ *factors*

$$\pi_1(U_0, u) \to G_{\mathbb{F}_q} \to GL(\mathcal{G}_u).$$

In particular, there are constant l-adic units $\alpha_1, ..., \alpha_k$ *such that for all* $x_0 \in |U_0|$,

$$\det(1 - t\phi_{x_0}, \mathcal{G}_u) = \prod_{i=1}^{k}(1 - \alpha_i^{\deg(x_0)}t).$$

Proof. Note that $\mathcal{G}$ is the pullback of $\mathcal{G}_0$ along the extension of scalars map, so by the compatibility noted in (3.3), the representation $\mathcal{G}_u$ restricted to $\pi_1(U, u)$ is that associated to a constant $\mathbb{Q}_l$-sheaf on U, which is a trivial representation. Therefore $\pi_1(U, u)$ acts trivially, so by the homotopy exact sequence

$$0 \to \pi_1(U, u) \to \pi_1(U_0, u) \xrightarrow{v} G_{\mathbb{F}_q} \to 0$$

we get the claimed factorisation. The second part follows by noting that $v(\phi_{x_0})$ is always the $deg(x_0)$-th power of the geometric Frobenius, since the embedding Spec $k(x_0) \to U_0$ is an $\mathbb{F}_q$-map. □

Equipped with this, we now immediately note that almost every sheaf in the game is geometrically constant. Indeed, recall that by Grothendieck's trace formula and proper base change (taking $Y_{x,0}$ to be the variety defined over $k(x_0)$ in the obvious way)

$$\begin{aligned}\zeta(Y_{x,0}, t) &= \prod_i Det(1 - tF^{deg(x_0)}, H^i(Y_x, \mathbb{Q}_l))^{(-1)^{i+1}} \\ &= \prod_i Det(1 - tF^{deg(x_0)}, (R^i f_* \mathbb{Q}_l)_x)^{(-1)^{i+1}}.\end{aligned}$$

By properties of vanishing cycles, each $R^i f_* \mathbb{Q}_l$ is constant except for $i = n$, and in that case we can split the characteristic polynomial up into pieces corresponding to the sheaves $R^n f_* \mathbb{Q}_l / \mathcal{E}$, $\mathcal{E}/\mathcal{E} \cap \mathcal{E}^\perp$ and $\mathcal{E} \cap \mathcal{E}^\perp$, and again the outer two of these are constant. By the lemma (and that n is odd), we have proved the following.

Proposition 8.3. *There are fixed l-adic units $\alpha_1, ..., \alpha_r$ and $\beta_1, ..., \beta_s$ (wlog with no $\alpha_i = \beta_j$), such that for all $x_0 \in |U_0|$,*

$$\zeta(Y_{x,0}, t) = \frac{(1 - \alpha_1^{deg(x_0)}t)\ ...\ (1 - \alpha_r^{deg(x_0)}t)}{(1 - \beta_1^{deg(x_0)}t)\ ....\ (1 - \beta_s^{deg(x_0)}t)} det(1 - \phi_{x_0}t, V).$$

Note that we know they are l-adic units because they come from a representation on $\mathbb{Z}_l$ and the geometric Frobenius has an inverse. This is worth keeping track of for later, because for any element in $\hat{\mathbb{Z}}$ (not just ordinary $\mathbb{Z}$) we can take powers of l-adic units.

So as when we were proving Cebotarev, we can now avoid having to look at the hardest object V by aiming to prove that everything *except* the factor $det(1-\phi_{x_0}t, V)$ forms an element of $\mathbb{Q}(t)$, and the result will follow. Since we know this for $\zeta(Y_{x,0}, t)$, it will suffice to consider the rational function involving the α_is and β_js. Our approach will be to use the Cebotarev density theorem to find points with certain convenient constraints on their degree, and whose Frobenii do avoid certain eigenvalues (by showing the set of corresponding elements in $\pi_1(U_0, u)$ is open with positive measure). We can then deduce rationality by a purely algebraic argument as follows (here our method more closely follows that taken by Freitag-Kiehl [6] rather than Deligne [4]).

Lemma 8.4 (Rationality Criterion)**.** *Fix $\tau \in Gal(\bar{\mathbb{Q}}_l/\mathbb{Q})$, and suppose there is a closed point $x_0 \in |U_0|$ with the following properties:*

1. *The integer $k = deg(x_0)$ does not have the property that there exist α_i and β_j with α_i/β_j a kth root of unity.*
2. *Furthermore, $\alpha_i/\tau(\alpha_j)$ and $\beta_i/\tau(\beta_j)$ are not kth roots of unity except when they are 1.*
3. *The numbers $\alpha_1^k, ..., \alpha_r^k, \beta_1^k, ..., \beta_s^k$ and their images under τ^{-1} are not amongst the eigenvalues of ϕ_{x_0} acting on V.*

Then $\frac{(1-\alpha_1^{deg(y_0)}t)\ ...\ (1-\alpha_r^{deg(y_0)}t)}{(1-\beta_1^{deg(y_0)}t)\ ....\ (1-\beta_s^{deg(y_0)}t)}$ is invariant under τ for all $y_0 \in |U_0|$.

Proof. Let $\gamma_1, ..., \gamma_u$ be the eigenvalues of ϕ_{x_0} acting on V. Then since the zeta function is rational, the following expression is invariant under τ:

$$\frac{(1 - \alpha_1^k t)\ ...\ (1 - \alpha_r^k t)}{(1 - \beta_1^k t)\ ....\ (1 - \beta_s^k t)}(1 - \gamma_1 t)...(1 - \gamma_u t).$$

From property 3 it is clear that τ cannot possibly send γ_i to an α_j^k or β_j^k, so it acts separately on the product $(1 - \gamma_1 t)...(1 - \gamma_u t)$. Also, by property 1, note that the remaining fraction cannot cancel, so τ permutes the α_i^ks amongst themselves and the β_j^ks amongst themselves. Finally, by property 2 the same action can be observed amongst the α_is and β_js, whence the result. □

We now need to establish the existence of such points. Fix the isomorphism $G_{\mathbb{F}_q} \cong \hat{\mathbb{Z}}$ identifying the geometric Frobenius with 1, and let $v : \pi_1(U_0, u) \to \hat{\mathbb{Z}}$ be the corresponding projection map. The key result we need (to set us up for applying Cebotarev) is the following.

Proposition 8.5. *For any l-adic unit $\lambda \in \bar{\mathbb{Q}}_l$, the set $Z = \{\sigma \in \pi_1(U_0, u) : \lambda^{v(\sigma)}$ is an eigenvalue of σ acting on $V\}$ is a closed subset of $\pi_1(U_0, u)$ with zero Haar measure.*

Proof. Recall that $\rho : \pi_1(U, u) \to Sp(V, \psi)$ has open image. Let $GSp(V, \psi)$ be the group of *symplectic similitudes*

$$GSp(V, \psi) = \{\alpha \in GL(V) : \psi(\alpha(x), \alpha(y)) = \mu(\alpha)\psi(x, y)\ \forall x, y \in V\},$$

where the scaling factor $\mu(\alpha)$ can be any constant depending only on α.

Note that by the homotopy exact sequence and $\pi_1(U_0, u)$-equivariance of ψ, any $\sigma \in \pi_1(U_0, u)$ acts as a symplectic similitude with scaling factor $q^{v(\sigma)}$. We therefore (noting that q is an l-adic unit) let $H = \{(\alpha, n) \in GSp(V, \psi) \times \hat{\mathbb{Z}} : \mu(\alpha) = q^n\}$, a closed subgroup of $GSp(V, \psi) \times \hat{\mathbb{Z}}$ and there is an induced map

$$\rho_1 : \pi_1(U_0, u) \ni \sigma \mapsto (\rho(\sigma), v(\sigma)) \in H.$$

To make it clear what is going on, we shall now restrict to the case where q is a quadratic residue modulo $l \neq 2$, so that we can form an honest isomorphism $H \ni (\alpha, n) \mapsto (q^{-n/2}\alpha, n) \in Sp(V) \times \hat{\mathbb{Z}} =: H'$ (the general argument in [4, 6] and [6, IV.3] is similar but slightly more technical). Now it is obvious that ρ_1 has open image $\rho(\pi_1(U, u)) \times \hat{\mathbb{Z}}$, so we can compare the topology and Haar measure of $\pi_1(U_0, u)$ with those of H', which are much easier to understand. In particular it will now suffice to show that $Z' = \{(\alpha, n) \in H' : (q^{-1/2}\lambda)^n$ is an eigenvalue of $\alpha\}$ is a closed subset of H' with zero Haar measure. For convenience, let $\lambda' = q^{-1/2}\lambda$.

Firstly, it is obviously closed (it is determined by the equation $det(\lambda'^n - \alpha) = 0$ which varies continuously in (α, n)). Also, if we fix n, view $Sp(V, \psi)$ as an algebraic group over $\mathbb{Q}_l$, and let $P_n(t)$ be the minimal polynomial of λ'^n over $\mathbb{Q}_l$, the set $Z'_n = \{\alpha \in Sp(V, \psi) : \lambda'^n$ is an eigenvalue of $\alpha\}$ is a subset of the proper algebraic subspace of $Sp(V, \psi)$ defined by $\{\det P_n(\alpha) = 0\}$. Since such algebraic subspaces (of algebraic groups over an infinite base field) have Haar measure zero, it follows (by Fubini's theorem) that

$$\mu(Z') = \int_{n \in \hat{\mathbb{Z}}} \mu(Z'_n) = 0.$$

□

Equipped with this proposition, we can now use the Cebotarev density theorem to find points satisfying the hypotheses of the rationality criterion (8.4), and thus deduce $\frac{(1-\alpha_1^{deg(y_0)}t)\ \ldots\ (1-\alpha_r^{deg(y_0)}t)}{(1-\beta_1^{deg(y_0)}t)\ \ldots.\ (1-\beta_s^{deg(y_0)}t)}$ is rational and complete our proof of the main theorem (8.1).

Proposition 8.6. *For any $\tau \in Gal(\bar{\mathbb{Q}}_l/\mathbb{Q})$, there is a point $x_0 \in |U_0|$ satisfying the hypotheses of (8.4).*

Proof. By the Cebotarev density theorem, it suffices to prove that the set of $\sigma \in \pi_1(U_0, u)$ with the following three properties contains an open set of positive Haar measure.

1. If $v(\sigma)$ is an actual integer k, it does not have the property that there exist α_i and β_j with α_i/β_j a kth root of unity.

2. Furthermore, $\alpha_i/\tau(\alpha_j)$ and $\beta_i/\tau(\beta_j)$ are not kth roots of unity except when they are 1.

3. The numbers $\alpha_1^k, ..., \alpha_r^k, \beta_1^k, ..., \beta_s^k$ and their images under τ^{-1} are not amongst the eigenvalues of σ acting on V.

But this is easy by what we have already proved. The previous proposition implies that the set of elements satisfying 3 is actually open with *full* Haar measure. Also, the set of elements eliminated by the first two properties is certainly contained in the inverse image under $\pi_1(U_0, u) \twoheadrightarrow \hat{\mathbb{Z}}$ of $\bigcup_{i=1}^m n_i\hat{\mathbb{Z}}$ for $n_1, ..., n_m$ some finite collection of integers greater than 1. This is a proper closed subset of $\hat{\mathbb{Z}}$, so its complement is obviously open with positive Haar measure, and this property lifts to the inverse image.

□

So we have at last established the rationality result desired. It now feels like we know a lot about the arithmetic properties of the sheaf $\mathcal{F}_0$. Surely the Riemann hypothesis cannot be far away?

9 Deligne's Main Lemma

In this section, we present Deligne's ingenious 'main lemma' (following his original argument [4, 3], which shows that the properties of $\mathcal{F}_0$ we have now established are indeed enough to deduce that in fact Frobenius acts on the *stalks* with eigenvalues of absolute value precisely $q^{n/2\deg(x)}$ (for any embedding of $\bar{\mathbb{Q}}_l$ into $\mathbb{C}$.), and this will be enough to deduce the conditions of (7.3) and so complete the proof of the Riemann hypothesis and the full Weil Conjecture.

A smooth $\mathbb{Q}_l$-sheaf is called *pure of weight* $\beta \in \mathbb{Z}$ if for any embedding $\tau : \bar{\mathbb{Q}}_l \hookrightarrow \mathbb{C}$, the eigenvalues α of Frobenius acting on the stalks at every point x_0 of the base space satisfy $|\tau(\alpha)| = q^{(\beta/2)\deg(x)}$. For example, $\mathbb{Q}_l(m)$ has weight $-2m$. We shall henceforth suppress the τ from the notation, writing $|\alpha|$ to mean $|\tau(\alpha)|$, and all our estimates will be independent of the choice of τ.

Theorem 9.1 (Main lemma). *Let U_0 be an open subset of $\mathbb{P}_0^1$, the projective line on $\mathbb{F}_q$, u an $\mathbb{F}_q$-rational geometric point, and $\mathcal{F}_0$ any smooth $\mathbb{Q}_l$-sheaf on U_0. Suppose for some $\beta \in \mathbb{Z}$ we have the following:*

1. *There is a $\pi_1(U_0, u)$-equivariant nondegenerate alternating bilinear form*

$$\psi : \mathcal{F}_0 \times \mathcal{F}_0 \to \mathbb{Q}_l(-\beta).$$

2. *The induced map $\pi_1(U, u) \to Sp(\mathcal{F}_x, \psi)$ has open dense image.*

3. *For every geometric point x of U, the characteristic polynomial $Det(1 - F^{deg(x_0)}t, \mathcal{F}_x)$ has rational coefficients.*

Then $\mathcal{F}_0$ is pure of weight β.

Note from the outset that we assume $\mathcal{F}$ is nonzero (else the problem is trivial) and that U is affine (if not, apply the result twice with different points removed). The idea of the proof is to replace the sheaf $\mathcal{F}$ by $\mathcal{F}^{\otimes 2k}$, so that we can exploit the theory of the symplectic group to determine the outer cohomology groups, and we then conclude by some convergence properties of the L-function which exploit the rationality assumption. Again, notice that we manage to avoid ever having to directly work with $H^1_c(U, -)$.

We start with two useful lemmas that are easy consequences of the rationality assumption (we remain throughout under the hypotheses of the main lemma).

Lemma 9.2 (Positivity of Euler factors)**.** *The Euler factors*

$$Det(1 - F^{deg(x_0)}t^{\deg(x_0)}, \mathcal{F}_x^{\otimes 2k})^{-1}$$

can be written as formal power series with positive rational coefficients.

Proof. Let $T = t^{\deg(x_0)}$. Observe that (recalling the determinant-trace formula used in section 5)

$$\begin{aligned} T\frac{d}{dT}\log Det(1 - F^{deg(x_0)}T, \mathcal{F}_x^{\otimes 2k})^{-1} &= \sum_{r\geq 1} Tr(F^{deg(x_0)r}, \mathcal{F}_x^{\otimes 2k})T^r \\ &= \sum_{r\geq 1} Tr(F^{deg(x_0)r}, \mathcal{F}_x)^{2k}T^r. \end{aligned}$$

By condition (3) the coefficients are squares of rational numbers, so this is a formal power series with positive rational coefficients. Further noting that dividing by T, formally integrating and then substituting into the exponential series do not alter this property, we conclude the result. □

Lemma 9.3 (Radius of convergence)**.** *If α is a pole of $L(U_0, \mathcal{F}^{\otimes 2k}, t)$, and α_x a zero of $Det(1 - F^{deg(x_0)}t^{deg(x_0)}, \mathcal{F}_x^{\otimes 2k})$, it follows that $|\alpha| \leq |\alpha_x|$.*

Proof. Note that as a formal power series $L(U_0, \mathcal{F}^{\otimes 2k}, t) = 1 + \sum_{i\geq 1} a_i t^i$, and that it is a product of formal power series

$$f_{x_0}(t) := 1 + \sum_{i\geq 1} a_{i\deg(x_0),x_0} t^{\deg(x_0)i} := Det(1 - F^{\deg(x_0)}t^{\deg(x_0)}, \mathcal{F}_x^{\otimes 2k})^{-1}$$

with positive rational coefficients and constant coefficient 1 (with the convention that $a_{i,x_0} = 0$ if $deg(x_0) \not| i$).

But then it is clear that $a_i \geq a_{i,x_0}$ for all i, x_0, so the radius of convergence of $\sum_{i\geq 1} a_i t^i$ is less than that of any $f_{x_0}(t) = \sum_{i\geq 1} a_{i,x_0} t^i$. The result now follows. □

It will also be useful for us to have the following result from Weyl's 'The Classical Groups' [23, VI,1] about the representation theory of symplectic groups. Let V be a $\mathbb{Q}_l$-vector space with a symplectic form $\psi : V \times V \to \mathbb{Q}_l$. Then the symplectic group acts naturally on $V^{\otimes 2k}$. For each set $P = \{\{i_1, j_1\}, ..., \{i_k, j_k\}\}$ of k pairs of indices partitioning $\{1, ..., 2k\}$, let $\psi_P(x) : V^{\otimes 2k} \to \mathbb{Q}_l$ be the linear extension of $\psi_P(x_1 \otimes ... \otimes x_{2k}) = \prod_{r=1}^{k} \psi(x_{i_r}, x_{j_r})$. These clearly give surjections onto 1-dimensional spaces which are invariant under the symplectic group. Weyl's result tells us that in some sense all such spaces come from this construction.

Lemma 9.4 (Coinvariants of the symplectic group on tensor powers)**.** *With everything as above, there exists a family $\mathcal{P}$ of sets P of pairs such that the induced map through the group of $Sp(V, \psi)$-coinvariants*

$$\psi_{\mathcal{P}} = \prod_{P\in\mathcal{P}} \psi_P : (V^{\otimes 2k})_{Sp(V,\psi)} \to \mathbb{Q}_l^{|\mathcal{P}|}$$

is an isomorphism.

With these in place, we can start the full proof. Again, we shall heavily exploit Grothendieck's cohomological interpretation of L-functions.

Proof of Theorem 9.1. We shall attempt to pick apart the L-function $L(U_0, \mathcal{F}_0^{\otimes 2k}, t)$ by analysing the outer cohomology groups. Since we assumed U affine, as we showed in the proof of (6.4), $H_c^0(U, \mathcal{F}^{\otimes 2k}) = 0$, and $H_c^2(U, \mathcal{F}^{\otimes 2k})(1)$ is the dual of $H^0(U, \check{\mathcal{F}}^{\otimes 2k}) = (\check{\mathcal{F}}^{\otimes 2k})^{\pi_1(U,u)}$.

But for V any k-vector space with a G-action, and W another k-vector space with trivial G-action noting the canonical isomorphisms $Hom_k(V, W) \cong Hom_k(\check{W}, \check{V})$ and that $(-)^G, (-)_G$ are right and left adjoints respectively for the functor $(Vec_k) \to (Rep_k(G))$ equipping a vector space with the trivial G-action, it is generally true (by full faithfulness of the Yoneda embedding) that $\check{V}^G$ is the dual of V_G. In particular, we conclude that

$$H_c^2(U, \mathcal{F}^{\otimes 2k}) \cong (\mathcal{F}_u^{\otimes 2k})_{\pi(U,u)}(-1).$$

But we assumed that $\pi(U, u) \to Sp(\mathcal{F}_u, \psi)$ has dense image, so (since the action is continuous), we can use the above lemma 9.4 to deduce that for some N

$$H_c^2(U, \mathcal{F}^{\otimes 2k}) \cong (\mathcal{F}_u^{\otimes 2k})_{Sp(\mathcal{F}_u,\psi)}(-1) \cong \mathbb{Q}_l(-k\beta - 1)^N$$

So by Grothendieck's formula, we know that the L-function has denominator $(1 - q^{k\beta+1}t)^N$, so in particular all its poles have absolute value $q^{-k\beta-1}$.

Thus by 9.3, any eigenvalue α of Frobenius on the stalk at x_0 must satisfy $q^{(-k\beta-1)\deg(x_0)} \leq |\alpha^{2k}|^{-1}$, which rearranges to $|\alpha| \leq q^{(\beta/2+1/2k)\deg(x)}$. But by the nondegenerate pairing ψ on such a stalk, whenever α is such an eigenvalue, so is q^{β}/α, so in fact we obtain bounds on both sides

$$q^{(\beta/2-1/2k)\deg(x)} \leq |\alpha| \leq q^{(\beta/2+1/2k)\deg(x)}.$$

Letting $k \to \infty$, we deduce the theorem. □

And at last we can complete our proof of the Riemann Hypothesis, with a beautiful final step that closely resembles the classical proof that $\zeta(s)$ has no zeroes in $Re(s) > 1$.

Corollary 9.5. *Any eigenvalue α of $H^1_c(U, \mathcal{F})$ is algebraic and all its conjugates satisfy the bound*

$$|\alpha| \leq q^{\beta/2+1}.$$

Proof. Since $\mathcal{F}$ has no constant subsheaf ($Sp(V, \psi)$ will never fix a 1-dimensional subspace of V), $H^2_c(U, \mathcal{F}) = 0$, and the L-function becomes

$$L(U_0, \mathcal{F}_0, t) = Det(1 - Ft, H^1_c(U, \mathcal{F})).$$

But the left hand side is, as a formal power series, a product of polynomials with rational coefficients, so by the fact that $\mathbb{Q}[[t]] \cap \mathbb{Q}_l[t] = \mathbb{Q}[t]$ which again follows from Hankel determinants (and being a polynomial rather than a rational function by the non-existence of poles over an algebraic closure), we deduce that $Det(1 - Ft, H^1_c(U, \mathcal{F})) \in \mathbb{Q}[t]$, so α is certainly algebraic.

For the second part (where we shall use the main lemma) note that α is a zero of

$$\begin{aligned} L(U_0, \mathcal{F}_0, t) &= \prod_{x_0 \in |U_0|} Det(1 - F^{deg(x_0)} t^{\deg(x_0)}, \mathcal{F}_x)^{-1}. \\ &=: \prod_{x_0 \in |U_0|} \prod_{i=1}^{r} (1 - t^{\deg(x_0)} \alpha_{x_0,i})^{-1} \end{aligned}$$

factorising each Euler factor in terms of its roots, which (by the main lemma) satisfy $|\alpha_{x_0,i}| = q^{\deg(x_0)\beta/2}$. Recall from analysis that this product is absolutely convergent iff the following sum is.

$$\sum_{x_0 \in |U_0|} r q^{\deg(x_0)\beta/2} t^{\deg(x_0)} \leq \sum_{s \geq 1} r q^s q^{s\beta/2} t^s.$$

And this is clearly so for $|t| < q^{-\beta/2-1}$. Thus there are no zeroes and poles in this region, so $|\alpha|^{-1} \geq q^{-\beta/2-1}$, which implies the result. □

Noting that the sheaf $\mathcal{F}_0 = \mathcal{E}_0/(\mathcal{E}_0 \cap \mathcal{E}_0^{\perp})$ was proved to satisfy all the hypotheses of this section, we have thus verified the condition of (7.3) so have at last established the full Weil conjecture.

10 Conclusion: The Wider Story

Now we have the Weil conjecture as a theorem, we can reflect on possible generalisations, stand back and examine the story it is telling, and discuss the interesting way it fits into the theory of Galois representations of a *number field*.

Just as the classical Riemann hypothesis for the Riemann zeta function is a special case of a 'generalised Riemann hypothesis' for L-functions, so too here do we have a generalisation of the Weil conjectures, first proved also by Deligne six years later using a similar method (but with the 'Cebotarev density' step needing to be heavily altered). It has the benefit that it can be proved by induction in a more obvious way, and this is fully exploited in the modern proof given in the first chapter of [13]. The statement is as follows (we proved this in the case $Y_0 = \text{Spec } \mathbb{F}_q$, $\mathcal{F}_0 = \mathbb{Q}_l$).

Theorem 10.1 (Generalised Riemann Hypothesis)**.** *Let $f_0 : X_0 \to Y_0$ be a smooth proper morphism of algebraic varieties over $\mathbb{F}_q$, and $\mathcal{F}_0$ a smooth $\bar{\mathbb{Q}}_l$ sheaf pure of weight β on X_0. Then for any $i \geq 0$, $R^i f_{0!}(\mathcal{F}_0)$ is pure of weight $\beta + i$.*

This more general statement is (at least according to a famous review of [6] by Katz) the one that is now most often used, especially in applications to getting tight bounds on certain exponential sums. However, the classical statement remains very beautiful and probably easier to prove.

We could ask whether there are any significantly different proofs out there. The Riemann hypothesis for curves C has, as well as a proof via Jacobians, a different proof using the intersection pairing on $C \times C$ and the Hodge index theorem. In the 1960s, before Deligne found his proof, Grothendieck worked hard to find an intersetion-theoretic proof of the general case. In a contemporary paper [14], Kleiman explains how the Weil conjectures follow from *Grothendieck's standard conjectures* on algebraic cycles, and it is not unreasonable to assume that Grothendieck considered the Weil Conjecture to be dependent on these. However, while the Weil conjecture has been a theorem for 40 years, the standard conjectures remain wide open.

One of the key purposes of the standard conjectures is to give weight to the *theory of motives*, proposed by Grothendieck as a kind of universal cohomology theory to explain why all the different l-adic cohomology theories seem to behave very similarly. It is therefore natural to ask whether the Weil conjecture gives us anything in this direction. After a bit of thought, one sees a glimmer of such ideas. As we emphasised early on, one important aspect of the Weil conjecture is that the individual characteristic polynomials of Frobenius acting on cohomology groups have honest *rational* coefficients and are intrinsically defined as factors of the zeta function by looking at the absolute values of their roots. Since the zeta function is also defined independently of l, we therefore see that the choice of l does not affect the dimensions of the cohomology groups, or the eigenvalues of the Frobenius action. Perhaps the present author is mistaken, but without the Weil conjecture, neither of these statements seems entirely obvious[3].

The Weil conjecture also has something fascinating to say about Galois representations of number fields, which are at the centre of modern number theory. For simplicity we work with $\mathbb{Q}$, but everything we say works for an arbitrary number field. Let us fix embeddings $\bar{\mathbb{Q}} \hookrightarrow \bar{\mathbb{Q}}_p$, inducing injections of '*decomposition groups*' $G_{\mathbb{Q}_p} \hookrightarrow G_{\mathbb{Q}}$. Also recall that $G_{\mathbb{Q}_p} \twoheadrightarrow G^{ur}_{\mathbb{Q}_p} \cong G_{\mathbb{F}_p}$, and the kernel of this map is called the *inertia subgroup* I_p, which has an interesting further internal structure. If this group acts trivially on a representation V, we say V is *unramified at* p and in this case the local action is just an action of $G_{\mathbb{F}_p}$.

The study of Galois representations $\rho : G_{\mathbb{Q}} \to GL(V)$ includes describing the 'local behaviour': how the representation behaves when restricted to $G_{\mathbb{Q}_p}$. In early parts of this essay we discussed in detail what is still the most important source of Galois representations in modern number theory. A representation is called *geometric* if it arises as a subquotient of some $H^i(X, \bar{\mathbb{Q}}_l(j))$ for X a smooth projective variety over $\mathbb{Q}$. We can study the local behaviour of such a representation at all but finitely many primes using the proper smooth base change theorem as follows.

Proposition 10.2 (Local properties of geometric Galois representations at primes of good reduction)**.** *Let X be a proper smooth variety over $\mathbb{Q}$, and suppose it has a proper smooth model $\mathfrak{X}$ over $R = \mathbb{Z}[\frac{1}{N}]$ (i.e. the generic fibre of $\mathfrak{X}$ is X). Then for any $p \not| Nl$, the $G_{\mathbb{Q}}$-representation $H^i(X_{\bar{\mathbb{Q}}}, \bar{\mathbb{Q}}_l)$ is unramified at p, and we have an isomorphism of $G_{\mathbb{F}_p}$-representations:*

$$H^i(X_{\bar{\mathbb{Q}}}, \bar{\mathbb{Q}}_l)|_{G_{\mathbb{Q}_p}} \cong H^i(\mathfrak{X}_{\bar{\mathbb{F}}_p}, \bar{\mathbb{Q}}_l).$$

Proof. Firstly, consider the following square.

$$\begin{array}{ccc} X_{\mathbb{Q}_p} & \longrightarrow & X \\ \downarrow & & \downarrow \\ \text{Spec } \mathbb{Q}_p & \longrightarrow & \text{Spec } \mathbb{Q} \end{array}$$

By proper base change and because pullbacks of sheaves correspond to restrictions of representations, we get an isomorphism of $G_{\mathbb{Q}_p}$-representations:

$$H^i(X_{\bar{\mathbb{Q}}}, \bar{\mathbb{Q}}_l)|_{G_{\mathbb{Q}_p}} \cong H^i(X_{\bar{\mathbb{Q}_p}}, \bar{\mathbb{Q}}_l).$$

We now need to pass through an integral model, and they do not tend to have enough étale covers for the elegant general yoga of étale local systems to work, so we need to instead explicitly work with base changes to algebraic closures and check that everything is functorial to realise the Galois action in terms of maps $X \times_k \bar{k} \overset{id\times\sigma}{\to} X \times_k \bar{k}$.

[3]Though I would guess maybe the dimension property could follow by taking a smooth lift to characteristic zero and using some kind of comparison theorem.

Let $\mathfrak{p}$ be a prime of $\bar{\mathbb{Q}}$ lying above p. Then we can form its valuation ring T, with generic point $\bar{\mathbb{Q}}_p$ and special point $\bar{\mathbb{F}}_p$. Furthermore, there is an obvious map $R \to T$ and we can base-change to get $\mathfrak{X}_T$ proper smooth over T. Now, by proper smooth base change, there is an isomorphism, functorial in $\mathfrak{X}_T$, between the cohomology of the two fibres:

$$H^i(\mathfrak{X}_{\bar{\mathbb{F}}_p}, \bar{\mathbb{Q}}_l) \cong H^i(X_{\bar{\mathbb{Q}}_p}, \bar{\mathbb{Q}}_l).$$

Note that both varieties are obtained as base changes from $\mathfrak{X} \times_S T$ and that $G_{\mathbb{Q}_p}$ acts naturally on T in a way that extends to the action on $\bar{\mathbb{Q}}_p$. Hence, the functoriality of the above isomorphism gives us that the $G_{\mathbb{Q}_p}$-action induced by such a map on the right hand side is the same as that on the left. In particular, $G_{\mathbb{Q}_p}$ acts through the quotient $G_{\mathbb{F}_p}$, so I_p acts trivially on *both* sides. By our opening remarks, the RHS is $H^i(X_{\bar{\mathbb{Q}}}, \bar{\mathbb{Q}}_l)|_{G_{\mathbb{Q}_p}}$, so we have proved that it is unramified p, and by what we have shown the rest of the proposition is also clear.

□

Combining this with the Weil conjecture, we obtain the following remarkable result.

Theorem 10.3 (Purity of geometric Galois representations)**.** *Let $\rho : G_{\mathbb{Q}} \to GL(V)$ be an l-adic geometric Galois representation. Then there is some integer $i - 2j$ such that for all but finitely many primes p, ρ is unramified at p, and all the eigenvalues of $\rho(Frob_p)$ have absolute value q^{i-2j}.*

As noted by Taylor [20], this is particularly surprising in the context of the following important conjecture about Galois representations.

Conjecture 1 (Fontaine-Mazur)**.** *Let $\rho : G_{\mathbb{Q}} \to GL(V)$ be an l-adic Galois representation with the following properties:*

- *It is unramified outside a finite set of primes p.*
- *At the prime $p = l$ it is de Rham (a technical condition from p-adic Hodge theory).*

Then ρ is geometric.

If this conjecture were true, then the purity of geometric Galois representations really is a strikingly regular and beautiful pattern of behaviour for such a large class of Galois representations to possess.

In this essay, we have developed very powerful general machinery for generating and studying Galois representations, and with this machinery in place have gone on to prove what is in many ways a very surprising theorem. To me, perhaps the most striking feature of what we did is how indirect the proof was, constantly switching between different techniques, combining fearsome machinery with an array of elementary tricks. At one moment we would be doing representation theory, the next algebraic geometry, and later still thinking about analytic properties of L-functions, but all the time getting slowly but perceptibly closer to our eventual proof of the Weil conjecture, which both as a conjecture and as a theorem has to be one of the highlights of the mathematical landscape of the twentieth century.

Acknowledgements

I am very grateful to Teruyoshi Yoshida, for setting this essay, inspiring me to start learning étale cohomology several months ago, and useful general conversations and advice. The chapter on the Cebotarev density theorem was heavily influenced by a discussion with Tony Scholl. Several important details were clarified by illuminating discussions with Chris Blake, Brian Conrad (Stanford) and Freddie Manners. I also had more indirectly helpful conversations with Lukas Brantner, Julio Brau, James Cranch (Sheffield), Chris Elliott (Northwestern), Saul Glasman (MIT), Wansu Kim, Jonathan Lee, Daniel Litt (Stanford), James Newton and Jack Shotton (Imperial College). I have also given several seminars on related topics this year, and should thank all attendees of those. Finally, for motivating my choice of topic I should thank Matt Emerton for his post[4] on Terence Tao's blog, and Andrew Wiles for telling me about the connection between this theory and classical Iwasawa theory.

References

[1] Artin, Michael; Alexandre Grothendieck, Jean-Louis Verdier, eds. Théorie des topos et cohomologie étale des schémas - (SGA 4).

[2] Conrad, Brian. Étale cohomology, online notes.

[3] Conrad, Brian. Nagata's compactification theorem (via schemes), online.

[4] Deligne, Pierre (1974). "La conjecture de Weil: I". Publications Mathématiques de l'IHÉS 43: 273307.

[5] Deligne, Pierre. SGA 4 1/2, Cohomologie étale, Springer Lecture Notes 1977.

[6] E. Freitag, R. Kiehl. Étale cohomology and the Weil conjecture, Springer, 1988.

[7] Grothendieck, Alexander (1957). "Sur quelques points d'algèbre homologique", The Tohoku Mathematical Journal. Second Series 9: 119221.

[8] Grothendieck, A. Revêtements étales et groupe fondamental (SGA 1).

[9] Grothendieck, Alexandre 1965-66. Cohomologie l-adique et Fonctions L - (SGA 5)

[4] http://terrytao.wordpress.com/career-advice/learn-and-relearn-your-field/

[10] R. Hartshorne. Algebraic Geometry, Springer, 1977.

[11] M. Jarden. The Cebotarev Density Theorem for Function Fields: An Elementary Approach, Math. Ann. 261, 467-475 (1982).

[12] J. de Jong. The Stacks Project.

[13] R. Kiehl, R.Weissauer. Weil conjectures, Perverse Sheaves and the l'adic Fourier Transform, Springer, 2000.

[14] S.L. Kleiman. Algebraic Cycles and the Weil Conjectures (1967), Dix exposés sur la cohomologie des schémas.

[15] Milne, J.S. Étale cohomology, Princeton University Press, 1980.

[16] Milne, J.S. Lectures on Étale cohomology, online notes.

[17] Serre, J.P. "Zeta and L-functions", Arithmetical Algebraic Geometry ed. Schilling, 1963.

[18] J.H. Silverman. The arithmetic of elliptic curves, Springer, 1985.

[19] Szamuely, T. Galois Groups and Fundamental Groups, Cambridge University Press, 2009.

[20] Taylor, Richard. Galois Representations, Proceedings of ICM 2002, volume I, 449-474.

[21] Vistoli, Angelo. Notes on Grothendieck topologies, fibered categories and descent theory.

[22] Weil, André (1949). "Numbers of solutions of equations in finite fields", Bulletin of the American Mathematical Society 55 (5): 497508.

[23] Weyl, H. The Classical Groups, Princeton Univ. Press.