TATE'S ISOGENY THEOREM FOR ABELIAN VARIETIES OVER FINITE FIELDS

SAM LICHTENSTEIN

Contents

1. Introduction 1
2. Some motivation: applications of theorem 1.1 2
3. Initial reductions 4
4. Statement of the main theorem 6
5. Facts about semisimple algebras 7
6. The meat of the proof 8
6.1. Proof under the strong finiteness hypothesis 8
6.2. Idea of proof under the weak finiteness hypothesis 10
6.3. Recall: polarizations and their associated bilinear forms 11
6.4. Proof under the weak finiteness hypothesis 11
7. Loose ends 12
7.1. Existence of primes ℓ such that $F_\ell = \prod \mathbf{Q}_\ell$ 12
7.2. The finiteness hypothesis when $k = \mathbf{F}_q$ 13
7.3. Digression: Hilbert schemes, alternate proof of the finiteness hypothesis 13
References 15

1. Introduction

Let k be a field, and let A, B be abelian varieties over k. Let $G = \mathrm{Gal}(k_s/k)$. Write $\mathrm{Hom}(A, B)$ for the abelian group of k-homomorphisms $A \to B$, and for a prime $\ell \neq \mathrm{char}(k)$ let $\mathrm{Hom}_G(T_\ell(A), T_\ell(B))$ denote the $\mathbf{Z}_\ell$-module of G-equivariant maps between the ℓ-adic Tate modules of A and B.

Consider the map

$$(\star) \qquad \mathbf{Z}_\ell \otimes \mathrm{Hom}(A, B) \to \mathrm{Hom}_G(T_\ell(A), T_\ell(B))$$

associating to a homomorphism of abelian varieties the induced morphism of Tate modules.

A variant on $(\star)$ allowing $\ell = \mathrm{char}(k)$ is very important for some applications, but that requires p-divisible groups (to be discussed later in the fall) so we pass over it in silence here, apart from remark 2.5.

The map $(\star)$ is always injective; the proof is the same as for elliptic curves. (See [S].) What's the image?

Theorem 1.1 (Tate, 1966)**.** *If k is finite, then $(\star)$ is an isomorphism.*

One goal of this lecture is to prove this theorem. But another goal is to prove the implication Akshay mentioned in his talk: for abelian varieties over a field k, if k-isogeny classes are

known to consist of finitely many k-isomorphism classes, then the map $(\star)$ is an isomorphism. This is useful because once the appropriate propeties of the Faltings height are established, this finiteness property of isogeny classes (when k is a number field) will be clear. Moreover Tate's conjecture that $(\star)$ is an isomorphism whenever k is finitely generated over its prime field (e.g. k a number field) is helpful to our cause of proving Mordell's conjecture: it implies that two abelian varieties with the $\mathbf{Z}_\ell[G]$-isomorphic Tate module are isogenous. Together with the finiteness of isogeny classes and a finiteness result for Galois representations, this yields the Shafarevich conjecture, which gives Mordell via the Kodaira-Parshin trick.

In fact, for the purposes of proving the Tate conjecture, we can even relax the finiteness condition on k-isogeny classes assumption a bit. To be precise, we consider two finiteness hypotheses which might hold for an abelian variety A over a field k, given a fixed prime ℓ.

$\mathrm{Hyp}_s(A, k, \ell)$: Up to isomorphism, there are only finitely many abelian varieties B over k which are isogenous to A via a k-isogeny of ℓ-power degree.

$\mathrm{Hyp}_w(A, k, \ell)$: For any $d \geq 1$, there are only finitely many isomorphism classes of abelian varieties B over k such that

(i) B is isogenous to A via a k-isogeny of ℓ-power degree, and
(ii) there exists a polarization ψ of B (over k) of degree d^2.

(The subscripts s and w are for "strong" and "weak" respectively.)

We will see that Hyp_s is strong enough to prove Tate's conjecture, and for number fields it will be a consequence of the theory of Faltings height.

On the other hand, Hyp_w is not too difficult to show directly when k is finite. (Since it is somewhat separate from the main ideas of this talk, I'll do this only at the end.) However, it alone doesn't quite suffice to prove Tate's conjecture. But with a bit of trickery, it can still be leveraged to give the proof under an additional assumption which also holds when k is finite. (This is good, because Faltings height doesn't exist over finite fields, so without this alternate proof we would not know Tate's conjecture holds for such fields.)

Here's the structure of the talk. First we'll give some motivation, by mentioning some applications of Tate's theorem. Next we'll make a couple of reductions, and then state our main theorem precisely. Finally we'll prove the theorem.

2. Some motivation: applications of theorem 1.1

For our discussion of applications we assume throughout that $k = \mathbf{F}_q$ is finite. Thus the Galois group G is topologically generated by the q-Frobenius automorphism of $\overline{k}$. Let A and B be abelian varieties over k.

Tensoring $(\star)$ up with $\mathbf{Q}_\ell$ (over $\mathbf{Z}_\ell$) we get to the land of vector spaces:

$$(\star\star) \qquad \mathbf{Q}_\ell \otimes \mathrm{Hom}(A, B) \hookrightarrow \mathrm{Hom}_G(V_\ell(A), V_\ell(B)),$$

where by definition $V_\ell(A) = T_\ell(A) \otimes_{\mathbf{Z}_\ell} \mathbf{Q}_\ell$. Note that $(\star\star)$ is still injective, since localization is exact. As we will see soon, one can pretty much work with $(\star\star)$ rather than $(\star)$: if one is an isomorphism, so is the other.

A consequence of Tate's theorem is that

$$\mathrm{rank}(\mathrm{Hom}(A, B)) = \dim \mathrm{Hom}_G(V_\ell(A), V_\ell(B)).$$

A first set of applications uses this identity to relate *isogenies* and *point counting*.

Let's think about the right side above. Here's an important fact about abelian varieties, which Brian mentioned in his talk.

Proposition 2.1. Let ϕ be an endomorphism of an abelian variety A over a field k. For $\ell \neq \mathrm{char}(k)$, the characteristic polynomial f of $T_\ell(\phi)$ acting on $V_\ell(A)$ has integer coefficients, and is independent of ℓ.

Proof. This is shown for elliptic curves in [S, §V.2] and by essentially the same method for abelian varieties in [Mu, §19]. □

We can leverage this result using the following fact from linear algebra.

Proposition 2.2. Let α, β be semi-simple endomorphisms of two finite-dimensional vector spaces over a field K, with characeristic polynomials f_α and f_β having factorizations

$$f_\alpha = \prod_p p^{n_{\alpha,p}}, \qquad f_\beta = \prod_p p^{n_{\beta,p}}$$

as powers of distinct monic irreducible polynomials $p \in K[t]$. Then

$$\dim\{\gamma \in \mathrm{Hom}(V, W) : \gamma\alpha = \beta\gamma\} = r(f_\alpha, f_\beta) := \sum_p n_{\alpha,p} n_{\beta,p} \deg p.$$

Moreover the number $r(f_\alpha, f_\beta)$ is invariant under extension of the field K.

Proof. Exercise. (See [Mu, App. 1, Lemma 6, p. 251].) □

Returning to the situation of abelian varieties over a finite field $k = \mathbf{F}_q$, the q-Frobenius in G acts on $T_\ell(A)$ exactly as does the Frobenius $\pi_A \in \mathrm{End}(A)$. The induced action of Frobenius on $V_\ell(A)$ is *semisimple*; we'll check this later when we discuss semisimple algebras.

It follows (with a bit of thought) that

$$\dim_{\mathbf{Q}_\ell} \mathrm{Hom}_G(V_\ell(A), V_\ell(B)) = r(f_A, f_B)$$

where f_A (resp. f_B) is the characteristic polynomial of the q-Frobenius endomorphism π_A of A (resp. π_B of B). Since the righthand side can be computed over $\overline{\mathbf{Q}} \subset \overline{\mathbf{Q}}_\ell$, it follows that this dimension is independent of ℓ. (We'll use this fact – which of course also follows from Tate's theorem – later, in the course of the *proof* of the main theorem!)

Thus we have proved the first part of the following corollary.

Corollary 2.3. Let A, B be abelian varieties over a finite field k, f_A resp. f_B the characteristic polynomial of π_A resp. π_B on $V_\ell(A)$ resp. $V_\ell(B)$. Then

(a) $\mathrm{rank}\,\mathrm{Hom}(A, B) = r(f_A, f_B)$;
(b) The following are equivalent:
  (b1) B is k-isogenous to an abelian k-subvariety of A,
  (b2) $V_\ell(B)$ is G-isomorphic to a G-subrepresentation of $V_\ell(A)$ for some $\ell \neq \mathrm{char}(k)$,
  (b3) $f_B | f_A$ in $\mathbf{Q}[t]$;
(c) The following are equivalent:
  (c1) A and B are k-isogenous,
  (c2) $f_A = f_B$,
  (c3) A and B have the same number of points over any finite extension of k.

Proof. We've seen (a). Everything else is pretty easy. For (b2)⇒(b1), given a G-linear inclusion $u : V_\ell(B) \hookrightarrow V_\ell(A)$, we can by Tate's theorem find $\phi \in \mathbf{Q} \otimes \mathrm{Hom}(B, A)$ such that $V_\ell(\phi)$ is arbitrarily close to u in $\mathrm{Hom}(V_\ell(B), V_\ell(A))$. Since having rank at least a certain amount is an open condition on a linear map, we can ensure that $V_\ell(\phi)$ is injective. Clearing the denominator, we can find a multiple of ϕ which is actually a k-map $B \to A$ inducing an

injective map on V_ℓ. But the dimension of the kernel of $V_\ell(\phi)$ is twice the dimension of the abelian subvariety of B given as the connected component of the identity in the subgroup $\ker \phi$. Thus $\ker \phi$ must actually be finite, so ϕ is a k-isogeny. We leave the rest as an exercise. One must use Poincaré reducibility (recalled below) and the semisimplicity of the action of the q-Frobenius on $V_\ell(A)$. □

Why is this corollary cool? For one thing, it proves the hard direction of the fact that two elliptic curves (or abelian varieties) are isogenous, if and only if their zeta functions match. (This is precisely (c1)⇔(c3).) That's because the characteristic polynomial f_A of π_A determines the number of points of A over any finite k'/k.

The other important application of Tate's theorem is that is gives insight into the structure of the endomorphism algebra $\mathrm{End}^0(A) = \mathbf{Q} \otimes \mathrm{End}(A)$. Here is just a taste of what one can prove.

Corollary 2.4. Let A be an abelian variety of dimension $g > 0$ over a finite field k. Let $\pi = \pi_A$ be the Frobenius endomorphism of A relative to k, and f its characteristic polynomial (on one, or any, $V_\ell(A)$). Then $F = \mathbf{Q}[\pi]$ is the center $Z(E)$ of the endomorphism algebra $E = \mathrm{End}^0(A)$, and we have the inequality

$$2g \leq [E : \mathbf{Q}] = r(f, f) \leq (2g)^2.$$

Proof. It will come out in the proof of 1.1 that the subalgebra $F_\ell = \mathbf{Q}_\ell[T_\ell(\pi)]$ of $\mathrm{End}(V_\ell(A))$ generated by Frobenius (i.e. the Galois group of k) consists of all endomorphism of $V_\ell(A)$ which commute with $E_\ell = E \otimes \mathbf{Q}_\ell$. Thus $Z(E_\ell) = E_\ell \cap F_\ell = F_\ell$ since F_ℓ is obviously contained in E_ℓ. But this implies that $Z(E) = E \cap Z(E_\ell) = \mathrm{End}^0(A) \cap \mathbf{Q}_\ell[\pi] = \mathbf{Q}[\pi]$.

Now by part (a) of the previous corollary we have $[E : \mathbf{Q}] = r(f, f)$. If we factor f over an algebraically closed field as $f(t) = \prod (t - \pi_i)^{m_i}$ then $r(f, f) = \sum m_i^2$. But $\sum m_i = \deg f = \dim V_\ell(A) = 2g$, and we have

$$\sum m_i \leq \sum m_i^2 \leq \left(\sum m_i \right)^2.$$

□

Remark 2.5. One can actually say a lot more, especially in the cases when the upper or lower bound in the inequality above is achieved, or when the minimal polynomial of π over $\mathbf{Q}$ is irreducible, so F is a *field*. The latter situation turns out to be the case when A is isogenous to a power of a k-simple abelian variety. It's not hard to show that in this case E is a central simple F-algebra split at all finite places of F not dividing $p = \mathrm{char}(k)$, and ramified at all real places. To fully analyze the structure of E – which is relevant for **Honda-Tate theory** – we need to know its invariants at primes of F lying over p. This requires studying the p-divisible group of A, which plays the role of the Tate module when $\ell = p$.

3. Initial reductions

Consider again the map $(\star)$

$$(\star) \qquad \mathbf{Z}_\ell \otimes \mathrm{Hom}(A, B) \hookrightarrow \mathrm{Hom}_G(T_\ell(A), T_\ell(B)).$$

Tensoring up with $\mathbf{Q}_\ell$ (over $\mathbf{Z}_\ell$) as before, we obtain

$$(\star\star) \qquad \mathbf{Q}_\ell \otimes \mathrm{Hom}(A, B) \hookrightarrow \mathrm{Hom}_G(V_\ell(A), V_\ell(B)).$$

Since vector spaces over the field $\mathbf{Q}_\ell$ are (at least slightly) easier to deal with than modules over $\mathbf{Z}_\ell$, we begin with the following.

Lemma 3.1. For fixed ℓ, A, B, $(\star)$ is surjective if and only if $(\star\star)$ is so.

Proof. It's enough to show that $\operatorname{coker}(\star)$ is torsion-free, since if $(\star\star)$ is surjective then $\operatorname{coker}(\star)$ is torsion. So it's enough to show that the $\mathbf{Z}_\ell$-module

$$\operatorname{Hom}_{\mathbf{Z}_\ell}(T_\ell(A), T_\ell(B))/\operatorname{im}(\star)$$

is torsion-free, and it's enough to check this module has vanishing ℓ-torsion.

Suppose $f \in \operatorname{Hom}(T_\ell(A), T_\ell(B))$ is such that ℓf zero modulo $\operatorname{im}(\star)$. Then $\ell f = \sum \alpha_i T_\ell(\phi_i)$ for some $\alpha_i \in \mathbf{Z}_\ell, \phi_i \in \operatorname{Hom}(A, B)$. In particular, approximating each α_i by some $\alpha_{i,n} \in \mathbf{Z}$, we can find a sequence of maps $\phi^{(n)} : A \to B$ with $\phi^{(n)} = \sum_i \alpha_{i,n}\phi_i$ such that $\phi^{(n)} \to \sum \alpha_i T_\ell(\phi_i)$ in $\operatorname{Hom}(A, B) \otimes \mathbf{Z}_\ell$, and hence $T_\ell(\phi^{(n)}) \to \ell f$. In particular, by the nature of ℓ-adic convergence, we have $T_\ell(\phi^{(n)}) \in \ell \operatorname{Hom}(T_\ell(A), T_\ell(B))$ for large n.

Thus $T_\ell(\phi^{(n)}) = \ell f_n$ for some $f_n \in \operatorname{Hom}(T_\ell(A), T_\ell(B))$ and maps $T_\ell(A)$ to $\ell T_\ell(B)$. Moreover we must have $f_n \to f$ since $\ell f_n = T_\ell(\phi^{(n)}) \to \ell f$, and we can check equality after tensoring with $\mathbf{Q}_\ell$.

If $\lambda \in T_\ell(A)$ projects to some $a \in A[\ell](k^s)$ then by definition $T_\ell(\phi^{(n)})(\lambda)$ projects to $\phi^{(n)}(a) \in B[\ell](k^s)$. But $T_\ell(\phi^{(n)})(\lambda)$ is a multiple of ℓ in $T_\ell(B)$, so its projection onto $B[\ell](k^s)$ vanishes. Therefore the subgroup(scheme) $A[\ell]$ – i.e., the kernel of the multiplication-by-ℓ isogeny $A \to A$ – is contained in the (scheme-theoretic) kernel of the isogeny $\phi^{(n)}$. The theory of isogenies of abelian varieties is sufficiently well-behaved for this to ensure that $\phi^{(n)}$ factors through ℓ, and so $\phi^{(n)} = \psi^{(n)} \circ \ell$ for some isogeny $\psi^{(n)} : A \to B$; that is, $\phi^{(n)} = \ell\psi^{(n)} \in \ell \operatorname{Hom}(A, B)$ for certain homomorphisms $\psi^{(n)}$. Since the ball of multiples of ℓ in $\operatorname{Hom}(A, B) \otimes \mathbf{Z}_\ell$ is ℓ-adically closed and the $\phi^{(n)}$ converge, it follows that the limit $\sum \alpha_i\phi_i$ of the $\phi^{(n)}$s is a multiple of ℓ, and that $\psi^{(n)} \to \psi := \frac{1}{\ell}\sum \alpha_i\phi_i \in \operatorname{Hom}(A, B) \otimes \mathbf{Z}_\ell$.

Now we see that $T_\ell(\psi) = \lim T_\ell(\psi^{(n)}) = \frac{1}{\ell}\lim T_\ell(\phi^{(n)}) = f$. So $f \in \operatorname{im}(\star)$ as desired.

(Note that we've used here the fact one can check equality of elements of $\operatorname{Hom}(T_\ell(A), T_\ell(B))$ and $\operatorname{Hom}(A, B) \otimes \mathbf{Z}_\ell$, after tensoring up with $\mathbf{Q}_\ell$, which is true because these modules are torsion-free over $\mathbf{Z}_\ell$.) □

The next reduction takes us from vector spaces to the more structured realm of $\mathbf{Q}_\ell$-*algebras*, by restricting our attention to the case of endomorphisms of a single abelian variety, rather than maps between two of them. As we shall see, we will be able to exploit the algebra structure in the proof of the main result, theorem 4.1.

Lemma 3.2. Fix ℓ. Suppose that for all A over k the map

$$(\star\star\star) \qquad \mathbf{Q}_\ell \otimes \operatorname{End}(A) \to \operatorname{End}_G(V_\ell(A))$$

is an isomorphism. Then for all A, B, the map $(\star\star)$ – and hence $(\star)$ – is an isomorphism.

Proof. We leave this as an (easy) exercise in functoriality; use the fact that $\operatorname{End}(A \times B) = \operatorname{End}(A) \oplus \operatorname{End}(B) \oplus \operatorname{Hom}(A, B) \oplus \operatorname{Hom}(B, A)$ and a similar decomposition on the Tate module side. □

Remark 3.3. Lemma 3.2 reduces theorem 1.1 for A and B to the case of endomorphisms of $A \times B$. Hence we have *increased the dimension* of our abelian varieties at this step. In particular, to handle even the elliptic curve case of theorem 1.1, we consider abelian surfaces.

4. Statement of the main theorem

Now we can state the main result. Consider again the map

$$\mathbf{Q}_\ell \otimes \mathrm{End}(A) \to \mathrm{End}(V_\ell(A))$$

Note that I've dropped the $G = \mathrm{Gal}(k^s/k)$-linearity from the right side, so for a particular A, ℓ, k what we want to prove about this map is that the image is precisely the subspace of G-invariant elements.

Notation. Denote the image of this map by E_ℓ; that is,

$$E_\ell = \mathrm{End}(A) \otimes \mathbf{Q}_\ell \quad \text{regarded as a } \mathbf{Q}_\ell\text{-subalgebra of } \mathrm{End}(V_\ell(A)).$$

Denote by F_ℓ the $\mathbf{Q}_\ell$-subalgebra of $\mathrm{End}(V_\ell(A))$ generated by the Galois group G acting as automorphisms of $V_\ell(A)$.

Theorem 4.1. *Fix a prime $\ell \neq \mathrm{char}(k)$ and an abelian variety A over a field k.*

(i) *Assume $\mathrm{Hyp}_s(A \times A, k, \ell)$ holds. Then*

$$(\star\star\star) \qquad \mathbf{Q}_\ell \otimes \mathrm{End}(A) \to \mathrm{End}_G(V_\ell(A))$$

is an isomorphism.

Bonus result: if $\mathrm{Hyp}_s(A, k, \ell)$ holds then $V_\ell(A)$ is a semisimple representation of G, so the $\mathbf{Q}_\ell$-algebra F_ℓ is semisimple.

(ii) *Assume $\mathrm{Hyp}_w(A, k, \ell)$ holds, and assume further that F_ℓ is isomorphic as a $\mathbf{Q}_\ell$-algebra to a product of copies of $\mathbf{Q}_\ell$ (so in particular, it is semisimple). Then $(\star\star\star)$ is an isomorphism.*

Moreover,

(iii) *$\mathrm{Hyp}_w(A, k, \ell)$ holds for any A, ℓ when k is finite. Moreover when k is finite there exist (infinitely many) primes ℓ such that $F_\ell = \prod \mathbf{Q}_\ell$ as a $\mathbf{Q}_\ell$-algebra.*

To conclude from (ii) and (iii) of the theorem (plus our previous reductions 3.1 and 3.2) that Tate's conjecture is true over finite fields, we must prove that it is "independent of ℓ". Let's get this out of the way now via the following lemma.

Lemma 4.2. Fix A, B abelian varieties over k. To conclude $(\star\star)$ is surjective for any $\ell \neq \mathrm{char}(k)$, it's enough to prove it for one $\ell \neq \mathrm{char}(k)$ and to show that $\dim \mathrm{Hom}_G(V_\ell(A), V_\ell(B))$ is independent of ℓ.

Proof. Indeed, $(\star\star)$ is always injective, both sides are finite-dimensional $\mathbf{Q}_\ell$-vector spaces, and the dimension of the left side is independent of ℓ. □

Justifying the application of the previous lemma, we have:

Lemma 4.3. Suppose k is finite and A, B are abelian varieties over k. Then $\dim \mathrm{Hom}_G(V_\ell(A), V_\ell(B))$ is independent of $\ell \neq \mathrm{char}(k)$.

Proof. As discussed earlier this follows from 2.1 and 2.2 provided we know the Frobenius π_A acts semisimply on $V_\ell(A)$. We'll establish this in a moment, after reminding ourselves of the facts of life concerning semisimple algebras. □

5. Facts about semisimple algebras

Before starting in on the proof the main theorem, we digress to mention some facts in non-commutative algebra that will be useful.

Definition 5.1. Let K be a field. A K-algebra R is called **semisimple** if R is semisimple as a left module over itself; that is, if any left ideal of R admits a direct complement in R.

Below we will need the following basic facts about semisimple algebras.

Proposition 5.2. Let K be a field.

(0) A K-algebra R is semisimple if and only every left R-module is semisimple (i.e. is a direct sum of simple R-modules).
(i) A matrix algebra over a division K-algebra is semisimple.
(ii) A finite product of semisimple K-algebras is semisimple.
(iii) Let A be an abelian variety over a field k. Then $\mathrm{End}_k(A) \otimes \mathbf{Q}_\ell$ is a semisimple $\mathbf{Q}_\ell$-algebra.

Proof. (0), (i) and (ii) are standard facts, which can be found in [L] for example. For (iii), note that by Poincaré's complete reducibility theorem (shown in [Mu, §18] when $k = \overline{k}$), A is k-isogenous to a product $\prod A_i^{n_i}$ of powers of pairwise non-isogenous k-simple abelian varieties A_i. (Here *k-simple* means that $A_i \neq 0$ and A_i contains no abelian k-subvarieties other than itself and zero.) Thus $\mathrm{End}^0(A) := \mathrm{End}(A) \otimes \mathbf{Q}$ coincides with $\mathrm{End}^0(\prod A_i^{n_i})$. Now $\mathrm{Hom}(A_i^{n_i}, A_j^{n_j}) = 0$ for $i \neq j$ so $\mathrm{End}^0(\prod A_i^{n_i}) = \prod \mathrm{End}^0(A_i^{n_i})$. Finally it's easy to see that $\mathrm{End}^0(A_i^{n_i}) = \mathrm{Mat}_{n_i}(\mathrm{End}^0(A_i))$, so we are reduced to the case when A is k-simple, and we need to show the endomorphism algebra is a division ring. (The statement (iii) then follows by tensoring up to $\mathbf{Q}_\ell$ over $\mathbf{Q}$ and using (i) and (ii).) But if A is k-simple then any nonzero k-map $A \to A$ must be surjective (else the image would be a nontrivial abelian k-subvariety), hence an isogeny. Isogenies are all invertible in the ring $\mathrm{End}^0(A)$. Proof: if f is a self-isogeny of A then $\ker f$ is a finite commutative group scheme over a field. Any such is annihilated by its order n. The fact that $n(\ker f) = 0$ implies that mutiplication by n factors through f; in other words, f has an inverse up to multiplication by n, so we win. □

Example 5.3. Let's now show that π_A acts semisimply on $V_\ell(A)$.

This reduces to showing that $\mathbf{Q}[\pi_A] \subset \mathrm{End}^0(A)$ is semisimple, since then by injectivity of $(\star\star)$, we have that $\mathbf{Q}_\ell[\pi_A] = \mathbf{Q}_\ell \otimes \mathbf{Q}[\pi_A] \subset \mathrm{End}(V_\ell(A))$ is semisimple. (Note that semisimplicity is preserved under separable extensions of the base field; see [L, XVII§6].)

To see $\mathbf{Q}[\pi_A]$ is semisimple, note that it is contained in the center $Z(\mathrm{End}^0(A))$, since k-rational endomorphisms commute with Frobenius. But we already saw that $\mathrm{End}^0(A)$ is a product of matrix rings over division $\mathbf{Q}$-algebras, so its center is a product of fields, and hence reduced. A commutative $\mathbf{Q}$-algebra is semisimple if and only if it is reduced (easy exercise), so we are done.

Definition 5.4. Let $R' \subset R$ be an inclusion of associative K-algebras. The **commutant** of R' in R is

$$C_R(R') := \{r \in R : rr' = r'r, \forall r' \in R'\}.$$

Theorem 5.5 (bicommutation)**.** *Let R be a semisimple algebra over a field K, and let V be a faithful R-module, finite-dimensional as a K-vector space. Then $C_{\mathrm{End}_K(V)}(C_{\mathrm{End}_K(V)}(R)) = R$.*

Proof. This is a special case of [J, Thm. 4.10]. □

Actually "bicommutation" seems to be more commonly used to describe this result in the context of operator algebras. People also call this theorem the "Double Centralizer theorem".

6. The meat of the proof

6.1. **Proof under the strong finiteness hypothesis.** Let us begin by proving 4.1(i), following [Mi, IV.2.3ff.]. The idea is very easy. Let D denote the commutant of E_ℓ in $\mathrm{End}(V_\ell(A))$. By proposition 5.2(iii), E_ℓ is semisimple. So by the bicommutation theorem 5.5, E_ℓ is the commutant $B = C_{\mathrm{End}(V_\ell(A))}(D)$. We want to prove that $\mathrm{End}(V_\ell(A))^G \subset E_\ell$ (the reverse inclusion being obvious), so consider $\alpha \in \mathrm{End}(V_\ell(A))^G$; we will show that $\alpha \in B = E_\ell$ and be done. The graph of α

$$W = \{(x, \alpha x) : x \in V_\ell(A)\} \subset V_\ell(A) \times V_\ell(A) = V_\ell(A \times A)$$

is a G-stable $\mathbf{Q}_\ell$-linear subspace. Suppose we could find an element u of $\mathrm{End}(V_\ell(A \times A))$ which satisfies

(i) $uV_\ell(A \times A) = W$, and
(ii) u actually belongs to $\mathrm{End}(A \times A) \otimes \mathbf{Q}_\ell \subset \mathrm{End}(V_\ell(A \times A))$.

Since any $d \in D$ commutes with E_ℓ, the "diagonal" endomorphism $d \oplus d \in \mathrm{End}(V_\ell(A \times A))$ commutes with the image of $\mathrm{End}(A \times A) \otimes \mathbf{Q}_\ell = \mathrm{Mat}_2(\mathrm{End}(A)) \otimes \mathbf{Q}_\ell$ in $\mathrm{End}(V_\ell(A \times A))$. In particular, $d \oplus d$ commutes with u. So $(d \oplus d)W = (d \oplus d)uV_\ell(A \times A) = u(d \oplus d)V_\ell(A \times A) \subset W$. This means that for any $x \in V_\ell(A)$, the point $(x, \alpha x)$ in the graph W of α is mapped by $d \oplus d$ to another point $(dx, d\alpha x) \in W$. So $d\alpha x = \alpha dx$. This is true for all $x \in V_\ell(A)$, so it follows that as endomorphisms of $V_\ell(A)$, $d\alpha$ and αd coincide. But this says precisely that α commutes with d. Since d was arbitrary, α commutes with D, so it belongs to the commutant E_ℓ of D and we are done.

Thus it all comes down to producing the projection u satisfying (i) and (ii), given only the fairly arbitrary G-stable subspace W of $V_\ell(A \times A)$. That we can do this, leveraging the "strong" finiteness hypothesis, is a consequence of the following proposition.

Proposition 6.1.1. Fix A, k, ℓ and assume $\mathrm{Hyp}_s(A, k, \ell)$. Let W be a G-stable subspace of $V_\ell(A)$. Then there exists $u \in E_\ell = \mathrm{End}(A) \otimes \mathbf{Q}_\ell \subset \mathrm{End}(V_\ell(A))$ such that $uV_\ell(A) = W$.

The proof will require a lemma.

Lemma 6.1.2. Let $\Lambda \subset T_\ell(A)$ be a G-stable, finite-index $\mathbf{Z}_\ell$-submodule. There exists an abelian variety B over k and an isogeny $f : B \to A$ of ℓ-power degree, such that $\mathrm{im}\, T_\ell(f) = \Lambda$.

Proof. Since Λ has finite index, $\Lambda \supset \ell^n T_\ell(A)$ for some n. Let N be the image of Λ in $T_\ell(A)/\ell^n T_\ell(A) = A[\ell^n](k_s)$. Since Λ was G-stable, so is N.

Brian noted last time that we can form the quotient by such a subgroup $N \subset A(k_s)$ and get an abelian variety $B = A/N$ defined over the ground field k, with the quotient map $A \to B$ is a separable k-isogeny. Since $N \subset A[\ell^n]$, the multiplication-by-ℓ^n map $\ell^n : A \to A$ factors through the quotient B, so we get a map $f : B \to A$ over k so that $A \to B \xrightarrow{f} A$ is ℓ^n. In particular, the degree of f is a power of ℓ. It remains to show that $\mathrm{im}\, T_\ell(f) = \Lambda$.

Consider the resulting diagram on Tate modules:

$$\begin{array}{ccc} T_\ell(A) & & \\ \downarrow & \searrow^{\ell^n} & \\ T_\ell(B) & \xrightarrow[T_\ell(f)]{} & T_\ell(A) \end{array}$$

We want to prove $\operatorname{im} T_\ell(f) = \Lambda$. Certainly $\operatorname{im} T_\ell(f) \supset \ell^n T_\ell(A)$. So it's enough to prove $\operatorname{im} T_\ell(f)/\ell^n T_\ell(A) = N \subset A[\ell^n](k^s)$. But modulo ℓ^n, the image of $T_\ell(f)$ is just the image of $f : B[\ell^n](k^s) \to A[\ell^n](k^s)$. Since $B = A/N$, we have

$$B[\ell^n](k^s) = [\ell^n]_A^{-1}(N)/N$$

and the map f, being induced by $[\ell^n]_A$, sends the coset $a + N$ to $\ell^n a \in A[\ell^n](k^s)$. Since ℓ^n surjects $[\ell^n]_A^{-1}(N) \twoheadrightarrow N$, this shows that $T_\ell(f) \bmod \ell^n$ surjects $B[\ell^n](k^s)$ onto N. So the image of $T_\ell(f)$ is precisely Λ as desired. □

Proof of 6.1.1. Let $X_n = (T_\ell(A) \cap W) + \ell^n T_\ell(A) \subset T_\ell(A)$. This is a finite index, G-stable, $\mathbf{Z}_\ell$-submodule. So by 6.1.2 there is an isogeny $f_n : B_n \to A$ of ℓ-power degree with $\operatorname{im} T_\ell(f_n) = X_n$. By $\mathrm{Hyp}_s(A, k, \ell)$, the B_n's belong to only finitely many k-isomorphism classes, so by pigeonhole we can find a subsequence B_{n_i} of k-isomorphic abelian varieties. Choose k-isomorphisms $v_i : B_{n_1} \to B_{n_i}$. Form the element $u_i = f_{n_i} v_i f_{n_1}^{-1} \in \mathrm{End}^0(A)$, which is possible because we can divide by isogenies such as the f_n's in the isogeny category (i.e. after tensoring with $\mathbf{Q}$).

Now we compute $T_\ell(u_i)(X_{n_1})$. For notational simplicity, let's omit writing the functor T_ℓ when we apply it to morphisms. We have

$$u_i(X_{n_1}) = u_i f_{n_1}(T_\ell B_{n_1}) = f_{n_i} v_i T_\ell(B_{n_1}) = f_{n_i} T_\ell(B_{n_i}) = X_{n_i}.$$

Since by definition $X_{n_i} \subset X_{n_1}$, this says that the u_i all preserve the lattice X_{n_1} in $T_\ell(A)$. Thus they belong to a compact subspace $\mathrm{End}_{\mathbf{Z}_\ell}(X_{n_1}) \cap E_\ell \subset E_\ell = \mathrm{End}(A) \otimes \mathbf{Q}_\ell$ (with respect to the ℓ-adic topology).

By compactness, we can refine our subsequence X_{n_i} and assume that the sequence u_i converges to a limit $u \in \mathrm{End}(X_{n_1}) \cap E_\ell$. For any $x \in X_{n_1}$ the limit $u(x) = \lim u_i(x)$ is arbitrarily ℓ-adically close to $\bigcap X_n = T_\ell(A) \cap W$. Conversely for any $y \in T_\ell(A) \cap W$, we can find $x_i \in X_{n_i}$ with $u_i(x_i) = y$. Some subsequence of the x_is must converge, by the same compactness argument, to an $x \in X_{n_1}$, which must satisfy $u(x) = \lim u(x_i) = \lim u_i(x_i) = y$. So $u(X_{n_1}) = T_\ell(A) \cap W$, which implies $uV_\ell(A) = W$. □

Let's now prove the bonus result that F_ℓ is semisimple, assuming $\mathrm{Hyp}_s(A, k, \ell)$ holds, since it's a straightforward application of 6.1.1.

Let $W \subset V_\ell(A)$ be G-stable. We need to construct a G-stable complement W'. Now consider the right ideal of E_ℓ,

$$\mathfrak{a} = \{u \in E_\ell : uV_\ell(A) \subset W\}.$$

By 6.1.1 there is some $u \in \mathfrak{a}$ which satisfies $uV_\ell(A) = W$. Thus $\mathfrak{a}V_\ell(A) = W$.

We know that $\mathrm{End}^0(A)$ is a product of matrix algebras over division $\mathbf{Q}$-algebras, so E_ℓ is a product of matrix algebras over division $\mathbf{Q}_\ell$-algebras. By an explicit analysis (which we leave as an exercise) of the right ideals in $\mathrm{Mat}_n(D)$ for a division K-algebra D, one finds that each such ideal is generated by an idempotent. It follows easily that the same is true for a product of such algebras.

Thus we see that the right ideal $\mathfrak{a}$ of E_ℓ is generated by an idempotent e. Now $eV_\ell(A) = eE_\ell V_\ell(A) = \mathfrak{a}V_\ell(A) = W$, so e is actually a projection operator on $V_\ell(A)$, with image W. The orthogonal idempotent $1-e$ is therefore a projection onto a direct complement W' of W. So $V_\ell(A) = eV_\ell(A) \oplus (1-e)V_\ell(A) = W \oplus W'$, and since $W' = (1-e)V_\ell(A)$ is the image of an element of E_ℓ, which commutes with G, this complement W' is G-stable. □

6.2. **Idea of proof under the weak finiteness hypothesis.** Now we turn to the case where we have only the weak finiteness hypothesis to work with, plus the hypothesis that F_ℓ is a product of copies of $\mathbf{Q}_\ell$.

The natural thing to do is to try to prove 6.1.1 with our weaker hypotheses and repeat the proof above. Unfortunately this is too hard: we need to impose an additional assumption on the subspace W (besides G-stability) to produce the required projection u. Alas, the graph of a random G-invariant endomorphism of $V_\ell(A)$ will fail to satisfy this additional assumption. So we must modify our proof strategy.

First we state another reduction which will be helpful. Note that the hypothesis on F_ℓ ensures its semisimplicity. As before we let D denote the commutant

$$D := C_{\mathrm{End}(V_\ell(A))}(E_\ell).$$

Observe that $F_\ell \subset D$, since we know that the endomorphisms E_ℓ of $V_\ell(A)$ coming from (k-rational) isogenies of A commute with the Galois action.

Lemma 6.2.1. Fix A, ℓ, k, and assume F_ℓ is semisimple. Then

$$(\star\star\star) \qquad \mathbf{Q}_\ell \otimes \mathrm{End}(A) \to \mathrm{End}_G(V_\ell(A))$$

is bijective if and only if $F_\ell = D$.

Proof. The bijectivity of $(\star\star\star)$ says that $E_\ell = C_{\mathrm{End}(V_\ell(A))}(F_\ell)$. So by the bicommutation theorem 5.5 the lemma follows. □

We need to prove $D \subset F_\ell$, under the hypotheses of 4.1(ii).

Suppose we knew that every F_ℓ-eigenline (1-dimensional F_ℓ-stable subspace, which is to say, G-stable line) in $V_\ell(A)$, were also stable under D. Since $F_\ell = \prod \mathbf{Q}_\ell$ we can decompose $V_\ell(A) = \bigoplus V_i$ as a sum of nonzero subspaces corresponding to the factors of F_ℓ, and F_ℓ is precisely the subalgebra of $\mathrm{End}(V_\ell(A))$ consisting of those endomorphisms which act as a scalar on each subspace V_i. In particular, for any nonzero $v_i \in V_i$, v_i is an eigenvector of F_ℓ and hence of an arbitrary $d \in D$. That is, V_i is preserved by d, and in fact $V_i \setminus \{0\}$ consists entirely of eigenvectors of d.

Lemma 6.2.2. If T is an endomorphism of a vector space U for which every nonzero vector in U is an eigenvector, then T is a scalar.

Proof. Exercise. □

Thus d acts by a scalar on each V_i, so it belongs to F_ℓ.

This reduces our problem to showing that G-stable lines in $V_\ell(A)$ are D-stable. Now this is where Tate had a brilliant idea: G-stable *lines* have an additional structure beyond their G-stability, in virtue of their one-dimensionality: a line L is automatically **isotropic** for any alternating bilinear form β on $V_\ell(A)$. (Recall that this means $\beta|_{L\times L}$ vanishes.) It turns out that we can endow $V_\ell(A)$ with a suitable symplectic form, and that the condition of being (maximal) isotropic for this form is precisely what enables us to prove a version of proposition 6.1.1 under the weaker finiteness hypothesis Hyp_w.

6.3. **Recall: polarizations and their associated bilinear forms.** Last time Brian explained an abelian variety A over k admits a polarization $\theta : A \to \widehat{A}$ over k, i.e. a symmetric k-isogeny such that the pullback $\mathcal{L} := (1,\theta)^*\wp_A$ of the Poincaré bundle $\wp_A$ on $A \times \widehat{A}$, is ample. Moreover we saw that this data together with the natural maps $A[n] \times \widehat{A}[n] \to \mu_n$ yields, for any $\ell \neq \mathrm{char}(k)$, a non-degenerate, skew-symmetric, G-equivariant, $\mathbf{Z}_\ell$-bilinear pairing

$$\beta = \beta_\theta : T_\ell(A) \times T_\ell(A) \to \mathbf{Z}_\ell(1) := \varprojlim \mu_{\ell^n}.$$

Tensoring up, we obtain a G-equivariant symplectic form

$$\beta : V_\ell(A) \times V_\ell(A) \to \mathbf{Q}_\ell(1)$$

on $V_\ell(A)$.

6.4. **Proof under the weak finiteness hypothesis.**

Proposition 6.4.1. Suppose $\mathrm{Hyp}_w(A,k,\ell)$ holds and $W \subset V_\ell(A)$ is a G-stable, maximal isotropic subspace with respect to the symplectic form β_θ defined by a polarization θ of A (defined over k). Then there exists $u \in E_\ell$ such that $uV_\ell(A) = W$.

Proof. As in the proof of 6.1.1, form the submodules $X_n = (T_\ell(A) \cap W) + \ell^n T_\ell(A)$, and let $f_n : B_n \to A$ be the ℓ-power degree k-isogeny with $\mathrm{im}\, T_\ell(f_n) = X_n$ guaranteed to exist by 6.1.2. This time we need to keep track of degrees a bit more carefully. Note that since W is maximal isotropic for a symplectic form on $V_\ell(A)$, it has dimension $\frac{1}{2}\dim V_\ell(A) = g$, which implies that $X_n/\ell^n T_\ell(A)$ is a subgroup of $A[\ell^n](k^s) \approx (\mathbf{Z}/\ell^n\mathbf{Z})^{2g}$ of order ℓ^{ng}. (Proof: $(W \cap T_\ell(A) + \ell^n T_\ell(A))/\ell^n T_\ell(A) \approx (W \cap T_\ell(A))/(W \cap \ell^n T_\ell(A)) = (W \cap T_\ell(A))/\ell^n(W \cap T_\ell(A))$ since W is a $\mathbf{Q}_\ell$-vector space implies $W = \ell^n W$, and $W \cap T_\ell(A)$ is $\mathbf{Z}_\ell$-free of rank $\dim_{\mathbf{Q}_\ell} W = g$.) Thus the proof of 6.1.2 in fact shows that $\deg f_n = \ell^{ng}$.

In order to repeat the proof of 6.1.1 with only the hypothesis Hyp_w, we must equip each B_n with a polarization of some fixed degree d^2, which will turn out to be $\deg\theta$. The first step is to consider the isogeny

$$f_n^*\theta := \widehat{f}_n \theta f_n : B_n \to A \to \widehat{A} \to \widehat{B}_n.$$

It's easy to check this is a polarization (i.e. the symmetry and ampleness conditions), and we leave this as an exercise in chasing definitions.

Unfortunately $\deg f_n^*\theta = \deg(f_n)^2 \deg\theta = \ell^{2ng}\deg\theta$ depends on n. What we need to do is split off a factor of $[\ell^n]$ from $f_n^*\theta$, leaving a polarization of degree $\deg\theta$. For this we exploit the isotropicity of W.

Consider the non-degenerate alternating pairing

$$\beta_{f_n^*\theta} : T_\ell(B_n) \times T_\ell(B_n) \to \mathbf{Z}_\ell(1).$$

I claim that in fact $\beta_{f_n^*\theta}$ takes values in $\ell^n\mathbf{Z}_\ell(1)$. We have

$$\beta_{f_n^*\theta}(x,y) = e_{B_n,\ell^\infty}(x, \widehat{f}_n\theta f_n y) = e_{A,\ell^\infty}(f_n x, \theta f_n y) = \beta_\theta(f_n x, f_n y), \qquad x,y \in T_\ell(B_n),$$

(writing α for $T_\ell(\alpha)$ everywhere,) since the dual isogeny acts as the adjoint for the Weil e_n pairings. Thus $\beta_{f_n^*\theta}(x,y) = \beta_\theta(f_n x, f_n y) \in \beta_\theta(X_n, X_n)$ for each n. But since W is β_θ-isotropic, we have

$$\beta_\theta(X_n, X_n) = \beta_\theta(\ell^n T_\ell(A), X_n) \subset \ell^n \mathbf{Z}_\ell(1).$$

We omit the proof of the following fact, which uses the nondegeneracy of the Weil pairing.

Lemma 6.4.2. Let B be an abelian variety over a field k with a polarization ϕ defined over k, such that the ℓ-adic pairing β_ϕ (with $\ell \neq \mathrm{char}(k)$) takes values in $\ell^n \mathbf{Z}_\ell(1)$. Then $\phi = \psi \circ [\ell^n]$ for some polarization ψ over k. □

Thus the polarization $f_n^*\theta$ on B_n factors as $\psi_n \circ [\ell^n]$ where ψ_n is a polarization of degree $\deg(f_n^*\theta)/\deg[\ell^n] = \deg\theta =: d^2$.

The finiteness hypothesis $Hyp_w(A, k, \ell)$ now guarantees that the B_ns fall into finitely many k-isomorphism classes, and the rest of the proof of 6.4.1 is identical to that of 6.1.1. □

Corollary 6.4.3. Suppose that $\mathrm{Hyp}_w(A, k, \ell)$ holds, let θ be a polarization of A over k, and assume F_ℓ is a product of copies of $\mathbf{Q}_\ell$. If W is any G-stable, β_θ-isotropic subspace of $V_\ell(A)$, then W is also D-stable.

Proof. Descending induction on $\dim W$. When W is maximal isotropic, we apply proposition 6.4.1 to produce $u \in E_\ell$ with $uV_\ell(A) = W$. Then as in the proof of theorem 4.1(i) we compute that $DW = DuV_\ell(A) = uDV_\ell(A) \subset W$, since D is the commutant of E_ℓ and so commutes with u.

Now suppose $\dim W < g = \frac{1}{2}\dim_{\mathbf{Q}_\ell} V_\ell(A)$, so W is not maximal isotropic. Consider the orthocomplement $W^\perp$ of W with respect to β. Since the Weil pairing and the k-rational polarization θ are G-equivariant, the fact that W is G-stable implies $W^\perp$ is too.

Now we use the assumption on F_ℓ: it implies that any simple F_ℓ-module is 1-dimensional. Thus we can decompose $W^\perp = W \oplus W'$ into F_ℓ-stable subspaces, and $W' = \bigoplus L_i$ into F_ℓ-eigenlines. How many L_is must occur? Well $\dim W' = \dim W^\perp - \dim W$. Since β is non-degenerate, $\dim W + \dim W^\perp = \dim V_\ell(A) = 2g$. So $\dim W' = 2g - 2\dim W \geq 2$, as W is non-maximal. Thus there are at least two lines L_1, L_2. Now consider the F_ℓ-stable, β_θ-isotropic subspaces $W \oplus L_1, W \oplus L_2$. Each has dimension bigger than that of W, so by induction they are D-stable. Hence their intersection W is D-stable as well, completing the proof. □

Applying this corollary to F_ℓ-eigenlines as above, this completes the proof of theorem 4.1(ii).

7. Loose ends

We have proved parts (i) and (ii) of the main theorem 4.1. But we still need to show part (iii) in order to conclude Tate's theorem 1.1 using lemmas 4.2 and 4.3.

7.1. Existence of primes ℓ such that $F_\ell = \prod \mathbf{Q}_\ell$.

Assume k is finite.

Let $F \subset \mathrm{End}^0(A)$ be the subalgebra generated by π_A. This is central and commutative, so it's a product of fields. Moreover since π_A acts on $V_\ell(A)$ by the topological generator of G, we have $F_\ell = \mathbf{Q}_\ell \otimes F \subset \mathrm{End}(V_\ell(A))$.

Now if $F = \prod F_i$ is a product of number fields F_i, note that $F_i \otimes \mathbf{Q}_\ell = \prod_{v|\ell}(F_i)_v$ is a product of ℓ-adic fields. Each $(F_i)_v$ is isomorphic to $\mathbf{Q}_\ell$ provided that ℓ splits completely in F_i. Now it is an easy exercise in algebraic number theory to show that infinitely many ℓ satisfy this condition for all i. (Hint: by passing to the Galois closure of the compositum of the F_i's, reduce to the case of a single number field F; for this, use Chebotarev or argue directly.)

7.2. **The finiteness hypothesis when $k = \mathbf{F}_q$.** Here we verify that $Hyp_w(A, k, \ell)$ holds when k is finite.

The most direct argument rests upon the following fact. If X is a variety, a **g-cycle** on X is a formal linear combination of irreducible, g-dimensional subvarietis of X.

Theorem 7.2.1 (Existence of Chow variety, [R]). *Suppose k is perfect. The k-rational g-cycles of degree d in $\mathbf{P}^N_k$ are parametrized by the* rational points *of a projective k-varierty* $\mathrm{Chow}_{g,d}(\mathbf{P}^N)$. □

Here is the idea of the proof. Work over $\overline{k}$, and worry about rationality issues later. A generic codimension $g+1$ linear subspace of $\mathbf{P}^N$ will be disjoint from a given irreducible g-cycle V. It turns out that in $((\mathbf{P}^N)^\vee)^{r+1}$, the locus of $(r+1)$-tuples of hyperplanes which have a common intersection with V forms a hypersurface. The defining polynomial is called the **Chow form** of V, and it is multihomogeneous of degree $\deg_{\mathbf{P}^N}(V)$ in each factor. Its coefficients are known as the **Chow coordinates** of V, and they determine V. Now one has to do some work to show that the locus of possible Chow coordinates is actually a projective variety, but this is true. We emphasize, however, that this variety interacts well with rationality of cycles only over a perfect field.

Now how do we apply this to get Hyp_w? In fact we will show the stronger statement: *Over a finite field k, there are only finitely many isomorphism classes of g-dimensional abelian varieties A which admit a k-polarization of some fixed degree d^2.* This implies $\mathrm{Hyp}_w(A, k, \ell)$ since any B isogenous to a fixed A has dimension $g = \dim A$.

To get the finiteness result, we want to obtain from a polarization ψ of fixed degree d^2, a projective embedding of controlled degree. For then we can consider our abelian varieties as (irreducible) rational g-cycles in a projective space of controlled degree, which correspond to rational points of the corresponding Chow variety, and are thus finite in number.

Now $\mathcal{L} = (1, \psi)^* \wp_A$ is an ample line bundle on A. Over $\overline{k}$ any polarization comes from an ample line bundle, so $\psi_{\overline{k}} = \phi_{\mathcal{N}}$ for some ample $\mathcal{N}$ on $A_{\overline{k}}$. Moreover we have the relation that $\mathcal{L}_{\overline{k}} \cong \mathcal{N}^{\otimes 2} \bmod \mathrm{Pic}^0(A_{\overline{k}})$. Consequently the k-rational ample line bundle $\mathcal{L}$ satisfies $\deg \phi_{\mathcal{L}} = \deg \phi_{\mathcal{L}_{\overline{k}}} = \deg \phi_{\mathcal{N}^{\otimes 2}} = \deg([2]\phi_{\mathcal{N}}) = \deg([2]\psi) = (2^g d)^2$.

Now we must invoke another nontrivial fact:

Theorem 7.2.2 ([Mu]). *If $\mathcal{L}$ is an ample line bundle on an abelian variety over a field k with $\deg \phi_{\mathcal{L}} = d^2$, then $\mathcal{L}^{\otimes 3}$ is very ample and the corresponding projective embedding (via the complete linear system $|\mathcal{L}^{\otimes 3}|$) is a map*

$$A \hookrightarrow \mathbf{P}^{3^g d - 1}_k$$

of degree $g! 3^g d$. □

Thus any g-dimensional abelian variety over k with a polarization of degree d^2 is equipped with a projective embedding of degree $g! 6^g d$ into $\mathbf{P}^{6^g d - 1}_k$.

Hence our desired finiteness assertion follows from the fact that the Chow variety $\mathrm{Chow}_{g, g! 6^g d}(\mathbf{P}^{6^g d - 1}_k)$ exists and is of finite type.

7.3. **Digression: Hilbert schemes, alternate proof of the finiteness hypothesis.** Here we introduce the Hilbert scheme in order to give an alternate proof of Hyp_w for finite fields k. Hilbert schemes and related moduli spaces will be used later in the seminar, so it's worth getting a glimpse of them now. We refer to [FGAE] for details, proofs, and further examples.

Theorem 7.3.1 ([Mu], §16)**.** *Let X be an abelian variety over a field. Then for any line bundle $\mathcal{L} = \mathcal{O}_X(D)$ on X, $\chi(\mathcal{L}^{\otimes n}) = (D^{\cdot g})n^g/g!$. Here $(D^{\cdot g})$ denotes the g-fold self-intersection number of the divisor D. In particular, is $\mathcal{L}$ is very ample, the Hilbert polynomial of the embedding $|\mathcal{L}| : X \hookrightarrow \mathbf{P}^N$ is homogeneous, so is determined by the degree of the embedding.* □

Secretly we have already invoked this fact, since it is used for the proof of theorem 7.2.2. What this theorem tells us is that the Hilbert polynomial of the embedding

$$A \hookrightarrow \mathbf{P}_k^{6^g d - 1}$$

is determined by the degree $g!6^g d$. Thus rather than invoking the fact that this embedding makes A into a rational g-cycle on this projective space, and the existence of the Chow variety, we can use the Hilbert scheme instead:

Definition 7.3.2. The **Hilbert functor** of closed subschemes of $\mathbf{P}^N$ with Hilbert polynomial h is the functor $(Schemes) \to (Sets)^o$ defined by

$$Hilb^h_{\mathbf{P}^N}(T) = \{\text{finitely presented closed subschemes } Y \subset \mathbf{P}^N_T : Y \to T \text{ is flat, } h_{Y_t} = h, \forall t \in T\}.$$

Rather than asking for a Chow variety, whose "points" parametrize g-cycles on $\mathbf{P}^N$, we can ask for a Hilbert *scheme* which actually represents the functor above. This has the benefit that being a "fine moduli space" – i.e. actually representing a reasonable functor – makes such a scheme a lot easier to work with than something like a Chow variety. On the other hand, it has the cost that you must let yourself parametrize subschemes which are not actually varieties. (That is, $Y \subset \mathbf{P}^N$ might be non-integral, and it will still show up as a point of the corresponding moduli space.)

Here is a (special case of) the main existence theorem for Hilbert schemes.

Theorem 7.3.3 (Grothendieck)**.** *The functor $Hilb^h_{\mathbf{P}^N}$ is representable by a projective $\mathbf{Z}$-scheme* $\mathrm{Hilb}^h_{\mathbf{P}^N}$. □

(In fact one can replace the base $\operatorname{Spec}\mathbf{Z}$ by an arbitrary scheme S, and replace projective space with an arbitrary finitely presented projective morphism $\pi : Z \to S$ and a choice of relatively ample $\mathcal{L}$ on Z. One still has a representing scheme which is projective over S.)

Remark 7.3.4. Interestingly, Grothendieck's original proof of the above theorem used the notion of Chow coordinates introduced above. Mumford, however, discovered a proof which avoids this construction; see [FGAE].

We emphasize that in order to get a finite type representing scheme, you must pin down the Hilbert polynomial: if you ask merely for "all flat closed subschemes of $\mathbf{P}^N$" (or of some other projective scheme), the analogously defined functor is representable, but by a huge disjoint union of the Hilbert schemes corresponding to all possible Hilbert polynomials in $\mathbf{Q}[t]$.

However, as we saw above, an abelian variety A over a field k of fixed dimension g with a k-polarization of fixed degree d^2 automatically possesses a projective embedding $A \hookrightarrow \mathbf{P}^N_k$ over k, where $N = N_{d,g}$ and the Hilbert polynomial $h = h_{g,d}$ both depend only on d and g. Thus since the Hilbert scheme exists, each k-isomorphism class of such an A gives rise to a k-point of $\mathrm{Hilb}^{h_{g,d}}_{\mathbf{P}^{N_{g,d}}}$. Since the Hilbert scheme is finite type over $\mathbf{Z}$ and k is finite, there are only finitely many k-points. In particular, we have re-proved the finiteness hypothesis $\mathrm{Hyp}_w(A, k, \ell)$.

References

[FGAE] *FGA Explained*

[J] Jacobson, *Basic Algebra II.*

[L] Lang, *Algebra*

[Mi] Milne, Notes on abelian varieties.

[Mu] Mumford, *Abelian varieties*

[R] David Rydh, Master's thesis.

[S] Silverman, *Arithmetic of elliptic curves*

[T] Tate, Endomorphisms of abelian varieties over finite fields.