LECTURE 2: CM TYPES, REFLEX FIELDS, AND MAIN THEOREM (2/22/12)

NOTES BY I. BOREICO

1. The CM type in characteristic 0.

Recall the setting at the end of the previous lecture: A is an abelian variety over a field F of characteristic zero, with dimension $g > 0$. Let L be a number field of degree $2g$ over $\mathbf{Q}$ "acting" on A; that is, we have an embedding $L \hookrightarrow \mathrm{End}^0(A)$. The elements of L that actually act on A – i.e., $\mathcal{O} := L \bigcap \mathrm{End}(A)$ – form an order in L. Since $\mathrm{char}(F) = 0$, the action of $\mathcal{O}$ on $T_0(A)$ (a g-dimensional F-vector space) induces an action of $L = \mathcal{O} \otimes_{\mathbf{Z}} \mathbf{Q}$ on $T_0(A)$. This makes $T_0(A)$ into an $F \otimes_{\mathbf{Q}} L$-module (also $F \otimes_{\mathbf{Q}} L = F \otimes_{\mathbf{Z}} \mathcal{O}$, so we can get this action by another way).

Keep in mind that there are two actions of $\mathbf{Q}$ (or $\mathbf{Z}$) on $T_e(A)$: one coming from the embedding of $\mathbf{Q}$ in $\mathrm{End}^0(A)$ and another coming from the F-vector space structure (as A is an F-scheme). These actions coincide since both respect the underlying abelian group structure and the $\mathbf{Z}$-module structure on an abelian group is unique (so likewise for a $\mathbf{Q}$-vector space structure over this).

If $F = \mathbf{C}$, then A is a quotient V/Λ where V is a g-dimensional $\mathbf{C}$-vector space and Λ is a $\mathcal{O}$-module that is isomorphic to $\mathbf{Z}^{2g}$. In general when $g > 1$ this is *not* an invertible $\mathcal{O}$-module when $\mathcal{O} \neq \mathcal{O}_L$. (The case $g = 1$ is very misleading in this regard since invertibility does hold in such cases for any $\mathcal{O}$, due in part to the fact that quadratic orders are monogenic; higher-degree orders usually are not monogenic.) However, $H_1(A(\mathbf{C}), \mathbf{Q}) = \Lambda_{\mathbf{Q}} = \Lambda \otimes_{\mathbf{Z}} \mathbf{Q}$ is a 1-dimensional L-vector space. We try as much as possible to work in the isogeny category (or to pass to the case when the CM order is $\mathcal{O}_L$) so that we can sidestep the fact that the integral homology is usually not an invertible $\mathcal{O}$-module when $\mathcal{O} \neq \mathcal{O}_L$.

Observe that $T_0(A) = V = \mathbf{R} \otimes_{\mathbf{Z}} \Lambda \simeq \mathbf{R} \otimes_{\mathbf{Q}} \Lambda_{\mathbf{Q}}$ as $\mathbf{R} \otimes_{\mathbf{Q}} L$-modules, so it is a free $\mathbf{R} \otimes_{\mathbf{Q}} L$-module of rank 1. This description (involving only $\mathbf{R}$-linear structures) does not take into account the complex structure from $F = \mathbf{C}$; we would like to know the structure of $V = T_0(A)$ as a module over $\mathbf{C} \otimes_{\mathbf{Q}} L \simeq \prod\limits_{\phi\colon L \to \mathbf{C}} \mathbf{C}e_\phi$ (i.e., the $\mathbf{C}$-dimension of each isotypic part V_ϕ). This will allow us to prove the following assertion which was mentioned in the previous lecture:

Theorem 1.1. *Let A be an abelian variety over $\mathbf{C}$ of dimension $g > 0$, equipped with a number field $L \hookrightarrow \mathrm{End}^0(A)$ of $\mathbf{Q}$-degree $2g$. As a $\mathbf{C} \otimes_{\mathbf{Q}} L$-module, $T_0(A)$ is a direct sum $\prod\limits_{\phi \in \Phi} \mathbf{C}e_\phi$ where Φ is a set of g embeddings $\phi\colon L \to \mathbf{C}$ containing no conjugate pairs and $\mathbf{C}e_\phi$ is a 1-dimensional $\mathbf{C}$-space on which L acts via the embedding ϕ (thereby forcing L to have no real embeddings for $\mathbf{C}$-dimension reasons, so Φ is a set of representatives among conjugate pairs in* $\mathrm{Hom}(L, \mathbf{C})$*).*

Proof. We have $H_1(A(\mathbf{C}), \mathbf{C}) = \mathbf{C} \otimes_{\mathbf{Z}} \Lambda \simeq \mathbf{C} \otimes_{\mathbf{Q}} L$; i.e., the homology is free of rank 1 over $\mathbf{C} \otimes_{\mathbf{Q}} L$ and appears in the middle of the short exact sequence

$$\begin{tikzcd}
0 \arrow[r] & H^0(A(\mathbf{C}), \Omega^1) \arrow[r] \arrow[d, equal] & H^1(A(\mathbf{C}), \mathbf{C}) \arrow[r] & H^1(A(\mathbf{C}), \mathscr{O}_A) \arrow[r] & 0 \\
& T_0(A)^* & & &
\end{tikzcd}$$

This diagram is $\mathbf{C} \otimes_{\mathbf{Q}} L$-linear, where the action of $\mathbf{C}$ on $H^1(A(\mathbf{C}), \mathbf{C})$ comes from "coefficients" and the action of L comes from the "action" on $A(\mathbf{C})$. According to Hodge theory, inside $H^1(A(\mathbf{C}), \mathbf{C})$ (equipped with its natural complex conjugation via the $\mathbf{R}$-structure $H^1(A(\mathbf{C}), \mathbf{R})$) we have $H^0(A(\mathbf{C}), \Omega^1) \bigcap \overline{H^0(A(\mathbf{C}), \Omega^1)} = 0$ and $H^0(A(\mathbf{C}), \Omega^1)$ has half the dimension; i.e., g.

Because $H^1(A(\mathbf{C}), \mathbf{C}) \simeq \prod_{\phi} \mathbf{C} e_\phi$ as a $\mathbf{C} \otimes_{\mathbf{Q}} L$-module, $H^0(A(\mathbf{C}), \Omega^1)$ as a $\mathbf{C} \otimes_{\mathbf{Q}} L$-submodule of $\mathbf{C}$-dimension g must be $\prod_{\phi \in \Phi} \mathbf{C} e_\phi$ where Φ is a subset of g embeddings of L into $\mathbf{C}$. Disjointness of this $\mathbf{C}$-subspace from its conjugate now precisely states that Φ contains at most one element from each pair of conjugate pairs, and equality must hold (since the decomposition according to idempotents of $\mathbf{C} \otimes_{\mathbf{Q}} L$ is based on L-actions through embeddings into $\mathbf{C}$ and the left term in the exact sequence is the $\mathbf{C}$-linear dual of $T_0(A)$). □

Remark 1.2. For L a CM field, on $\mathrm{Hom}(L, \mathbf{C})$ there are two actions of complex conjugation: one coming from the action on the target $\mathbf{C}$ and another from the action on the source L (recall that complex conjugation on L is well-defined). These actions coincide, and we prefer to describe the action of complex conjugation only in terms of its action on L, as this will make sense when $\mathbf{C}$ is replaced with other algebraically closed fields of characteristic 0 (such as $\overline{\mathbf{Q}}$ or $\overline{\mathbf{Q}}_p$) on which there is no natural involution.

Corollary 1.3. *Let A be a CM abelian variety of dimension g over a field F of characteristic 0, and let $L \hookrightarrow \mathrm{End}^0(A)$ be a CM number field realizing the CM structure (which always exists). The action of L on $T_0(A)_{\overline{F}}$ is given by a set of g embeddings $\Phi \in \mathrm{Hom}(L, \overline{F})$ that contains exactly one embedding from each pair of complex-conjugate embeddings (relative to pre-composition against complex conjugation on L).*

Proof. We use the Lefschetz principle, in the following precise sense. The abelian variety A with its CM structure descends to an abelian variety A' over a subfield F' of F that is finitely generated over $\mathbf{Q}$ and contains a Galois closure of L over $\mathbf{Q}$. We can now embed such a subfield F' into $\mathbf{C}$. Since all embeddings of L into $\overline{F}$ or $\mathbf{C}$ land in F', we can read off the set of embeddings of interest by working with $A'_{\mathbf{C}}$. Now we apply the previous considerations in the complex-analytic setting (using Hodge theory). □

Example 1.4. Let p be an odd prime and consider the curve $y^p = x(1-x)$. It is defined over $\mathbf{Q}$, and has an action of μ_p via the action $\zeta.(x, y) = (x, \zeta y)$ defined over $\mathbf{Q}(\zeta_p)$. It follows that $J = \mathrm{Jac}(C)$ over $\mathbf{Q}(\zeta_p)$ has an action of $\mathbf{Z}[\zeta_p]$. Since the genus of C is $g = \frac{p-1}{2}$, and $\mathbf{Q}(\zeta_p)$ has dimension $2g = p - 1$, we see that the induced embedding $\mathbf{Q}(\zeta_p) \hookrightarrow \mathrm{End}^0(J)$ is a CM structure on J.

Let's compute its CM type; i.e., the set Φ of embeddings $\mathbf{Q}(\zeta_p) \hookrightarrow \mathbf{C}$ that occur on its tangent space at the identity. It is equivalent to work with the $\mathbf{C}$-linear dual $T_0(J)^* = H^0(J, \Omega^1_{J/\mathbf{C}}) = H^0(C, \Omega^1_{C/\mathbf{C}})$, and this has an explicit basis $\{\mathrm{d}x/y^j\}$ where j ranges from $\frac{p+1}{2}$ to $p-1$. (Note that y is a uniformizer at the origin hence $\mathrm{d}x$ has order $p-1$ there, so our list of forms has orders $0, 1, \ldots, g-1$ at the origin and so these must be linearly independent, hence a basis.) This decomposition respects the $\mathbf{Z}[\zeta_p]$-module structure: in fact, ζ acts by $\zeta^{-j} = \zeta^{p-j}$ there. Thus we have found our decomposition with the set of embeddings Φ being given by $\zeta \to \zeta^j$ for j ranging from 1 to $\frac{p-1}{2}$.

2. Examples.

Keep the notations of the previous section, without the assumption that F has characteristic 0. Let $p = \mathrm{char}(F) \neq 0$. Then $T_0(A)$ is acted upon by $\mathcal{O}/p\mathcal{O}$.

Example 2.1. Using Honda-Tate theory (which classifies simple abelian varieties over finite fields $\mathbf{F}_q$ up to $\mathbf{F}_p$-isogeny), we can construct a simple abelian surface A over $\mathbf{F}_{p^2}$, with $\mathbf{Z}[\zeta_5] \hookrightarrow \mathrm{End}(A)$ such that the action of the the p^2-Frobenius coincides with the action of $p\zeta_5$. (Keep in mind also that q-Frobenius endomorphisms are always zero on the tangent space over $\mathbf{F}_q$.) It can be shown that $\mathrm{End}^0(A) = \mathbf{Q}(\zeta_5)$, but we do not use this.

Suppose $p \equiv \pm 2 \pmod 5$, so p is inert in $\mathbf{Q}(\zeta_5)$. Since $\mathbf{Z}[\zeta_5]$ acts on $T_0(A)$ and (p) is a maximal ideal of $\mathbf{Z}[\zeta_5]$, we get an $\mathbf{F}_{p^2}$-linear action on $T_0(A)$ by the finite field $\kappa = \mathbf{Z}[\zeta_5]/(p)$ of size p^4. But $T_0(A)$ has dimension 2 over $\mathbf{F}_{p^2}$, so $T_0(A)$ is a 1-dimensional κ-vector space with a κ-linear action of $\mathbf{F}_{p^2}$ (i.e., the $\mathbf{F}_{p^2}$-action commutes with the κ-action. Since $\mathrm{End}_\kappa(T_0(A)) = \kappa$ canonically, as for a 1-dimensional vector space over any field, we thereby get a *canonical* embedding

$$\mathbf{F}_{p^2} \hookrightarrow \mathrm{End}_\kappa(T_0(A)) = \kappa.$$

Now consider the geometric tangent space $T_0(A) \otimes_{\mathbf{F}_{p^2}} \overline{\mathbf{F}}_{p^2}$. This is free of rank 1 as a module over

$$\kappa \otimes_{\mathbf{F}_{p^2}} \overline{\mathbf{F}}_{p^2} \simeq \prod_{j:\kappa \to \overline{\mathbf{F}}_{p^2}} \overline{\mathbf{F}}_{p^2},$$

where j varies through the two $\mathbf{F}_{p^2}$-emebddings of κ into $\overline{\mathbf{F}}_{p^2}$. These two embeddings are swapped by the $\mathbf{F}_{p^2}$-involution of κ, which is induced by the complex conjugation of $\mathbf{Z}[\zeta_5]$.

In particular, unlike the characteristic zero case, the geometric tangent space produces two embeddings of κ which are related through the complex conjugation of $\mathbf{Z}[\zeta_5]$. This shows that A together with the $\mathbf{Z}[\zeta_5]$ action *cannot* lift to characteristic zero. (This argument does not rule out the possibility of passing to an isogenous abelian variety with CM order $\mathbf{Z} + p\mathbf{Z}[\zeta_5]$, in which case the above obstruction would not arise. In fact this does happen, and there is a CM lift after applying a suitable $\mathbf{F}_{p^2}$-isogeny.

Before we take up the next example, we need to record a remarkable general fact that highlights the importance of CM fields in the general theory of complex multiplication (even though the actual definition of "complex multiplication" does not involve CM fields):

Theorem 2.2. *Let K be an arbitrary field, and A a K-simple abelian variety of dimension $g > 0$ over K. Assume A has complex multiplication. Let $\Delta = \mathrm{End}_K^0(A)$, and let the number field Z be the center of Δ.*

1. *If $\mathrm{char}(K) = 0$ then Δ is a CM field.*
2. *If $\mathrm{char}(K) > 0$ then either Z is totally real of degree g and Δ is a central quaternion division algebra over Z non-split at all (real) archimedean places, or Z is a CM field and Δ is a central division algebra over Z.*

The proof of this result will be discussed in Tripathy's lecture(s). An example of case (2) with Z a totally real field is a supersingular elliptic curve over an algebraically closed field of characteristic p, in which case $Z = \mathbf{Q}$. Honda–Tate theory provides examples over $\mathbf{F}_{p^{2n+1}}$ with $Z = \mathbf{Q}(\sqrt{p})$. In case (2) with Z a CM field, finer information is known about the splitting behavior at finite places (and more specifically, relations between the local invariants of Δ at places of Z related by complex conjugation).

Example 2.3. Over $\mathbf{C}$, from a CM type on A we can read off whether or not A is simple. The criterion is that A is simple if and only if (L, Φ) is "primitive" in the sense that it is not induced from a CM type on a proper CM subfield $L' \subset L$. (From (L', Φ') we can define an associated Φ on L as the set of embeddings $L \to \mathbf{C}$ whose restriction to L' lies in Φ'.)

Proof. It is harmless to apply an L-linear isogeny, so we may assume $A^{\mathrm{an}} = V/\Lambda$ with $V = (\mathbf{R} \otimes_{\mathbf{Q}} L)_\Phi$ and $\Lambda = \mathcal{O}_L$. Now assume that (L, Φ) is induced from (L', Φ') with L' a proper CM subfield of L. Then we have a CM abelian variety A' with $A'^{\mathrm{an}} = (\mathbf{R} \otimes_{\mathbf{Q}} L')_{\Phi'}/\mathcal{O}_{L'}$ with smaller dimension than A and CM order $\mathcal{O}_{L'}$ inside L'. Since Φ is defined by extending Φ, the natural inclusion $\mathbf{R} \otimes_{\mathbf{Q}} L' \to \mathbf{R} \otimes_{\mathbf{Q}} L$ is $\mathbf{C}$-linear when using the respective complex structures from Φ' and Φ. Thus, we have a $\mathbf{C}$-linear injection

$$(\mathbf{R} \otimes_{\mathbf{Q}} L')_{\Phi'} \hookrightarrow (\mathbf{R} \otimes_{\mathbf{Q}} L)_\Phi$$

that visibly carries $\mathcal{O}_{L'}$ into $\mathcal{O}_L$. Passing to the quotient by these lattices realizes A'^{an} as an isogeny factor of A^{an} of strictly smaller dimension, so A is not simple.

Conversely, assume that A is not simple. Since the maximal commutative semisimple $\mathbf{Q}$-subalgebra $L \subset \mathrm{End}^0(A)$ is a *field* (rather than a nontrivial product of fields), the semisimple $\mathbf{Q}$-algebra $\mathrm{End}^0(A)$ cannot be a nontrivial product of semisimple $\mathbf{Q}$-algebras; i.e., it must be simple. That is, A must be isotypic. Letting A' be a simple isogeny factor, we have $\mathrm{End}^0(A) = \mathrm{Mat}_e(\mathrm{End}^0(A'))$ where $A \sim A'^e$ and $\mathrm{End}^0(A')$ is a central division algebra over a number field. Recall from the first lecture that A' must be CM since A is. By Theorem 2.2(1), $L' := \mathrm{End}^0(A')$ is a CM field, so $\mathrm{End}^0(A) \simeq \mathrm{Mat}_e(L')$. In particular, L' *is* the center of $\mathrm{End}^0(A)$. More specifically, for *every* isogeny $A \sim A'^e$, the resulting "diagonal" inclusion $L' \hookrightarrow \mathrm{End}^0(A)$ is always the same (since any two such isogenies are related through conjugation by a unit in $\mathrm{End}^0(A'^e) = \mathrm{Mat}_e(L')$ due to the Skolem–Noether theorem).

This canonical identification of L' with the center of $\mathrm{End}^0(A)$ thereby defines a canonical inclusion of L' into the maximal commutative subfield $L \subset \mathrm{End}^0(A)$ (as L is its own centralizer, so it must contain the center). In this way, L' is a CM subfield of L, necessarily a proper subfield for

$\mathbf{Q}$-degree reasons. (Explicitly, $[L : L'] = e > 1$.) Applying an L'-linear isogeny to A', we may arrange that there is an L'-linear isomorphism

$$A' \simeq (\mathbf{R} \otimes_{\mathbf{Q}} L')_{\Phi'}/\mathscr{O}_{L'}$$

where Φ' is the CM type for L' acting $\mathbf{C}$-linearly on $T_0(A')$. We have seen above that a choice of inclusion $j : A' \hookrightarrow A$ is automatically linear over the inclusion $L' \hookrightarrow L$ and so also linear over the inclusion $\mathscr{O}_{L'} \hookrightarrow \mathscr{O}_L$ between the CM orders.

Although $\mathscr{O}_L$ is generally not a free module over $\mathscr{O}_{L'}$, it is locally free of finite rank. Hence, the "Serre tensor construction" (to be discussed in Tripathy's lecture(s)) can be applied: this defines an abelian variety $\mathscr{O}_L \otimes_{\mathscr{O}_{L'}} A'$ over $\mathbf{C}$ which has the expected universal property for $\mathscr{O}_{L'}$-linear maps from A' to $\mathscr{O}_L$-module objects (via a definition as representing a suitable "tensor product sheaf"), and moreover it passes through the exponential analytic uniformization in the expected manner: we have an $\mathscr{O}_L$-linear analytic isomorphism

$$(\mathscr{O}_L \otimes_{\mathscr{O}_{L'}} A')^{\mathrm{an}} \simeq (\mathscr{O}_L \otimes_{\mathscr{O}_{L'}} V')/(\mathscr{O}_L \otimes_{\mathscr{O}_{L'}} \Lambda') = (L \otimes_{L'} (\mathbf{R} \otimes_{\mathbf{Q}} L')_{\Phi'})/\mathscr{O}_L.$$

The tangent space is easily seen to be exactly $(\mathbf{R} \otimes_{\mathbf{Q}} L)_{\Psi}$ where Ψ is the CM type on L defined by lifting Φ' from L'. Thus, to show that $\Phi = \Psi$ (thereby showing that (L, Φ) is not primitive) it remains to check that the $\mathscr{O}_L$-linear map

$$\mathscr{O}_L \otimes_{\mathscr{O}_{L'}} A' \to A$$

induced by $j : A' \to A$ is an isogeny.

The induced map on rational degree-1 homology is the map $L \otimes_{L'} \mathrm{H}_1(A'(\mathbf{C}), \mathbf{Q}) \to \mathrm{H}_1(A(\mathbf{C}), \mathbf{Q})$ induced by linearization of the injective map $\mathrm{H}_1(A'(\mathbf{C}), \mathbf{Q}) \to \mathrm{H}_1(A(\mathbf{C}), \mathbf{Q})$ linear over $L' \hookrightarrow L$ between 1-dimensional vector spaces over these respective fields. Thus, the linearization is a nonzero L-linear map between 1-dimensional L-vector spaces, so it is an isomorphism. This isomorphism property on degree-1 rational homology forces the map of abelian varieties to be an isogeny. □

Corollary 2.4. *Over any algebraically closed field F of characteristic 0, a CM pair (A, L) is determined up to L-linear isogeny by the F-valued CM type.*

Proof. As in the proof of 1.3, $A = A'_F$ for an abelian variety A' with CM by L over a subfield F' of F finitely generated over $\overline{\mathbf{Q}}$. In particular, A' has the same CM type as A. We can do the same for another CM abelian variety with CM by L over F, at the cost of increasing F' a bit. It suffices to solve the problem over $\overline{F}'$ instead of over F. Thus, we may assume F has finite transcendence degree over $\mathbf{Q}$, so it admits an embedding into $\mathbf{C}$. Fix such an embedding.

If two abelian varieties A and B with action by L over F admit an L-linear isogeny between them over $\mathbf{C}$, then by descent to an L-linear isogeny $A_R \to B_R$ over a finitely generated F-subalgebra $R \subset \mathbf{C}$ and passage to the quotient by a maximal ideal of R (to return to working over F) we can "specialize" to an L-linear isogeny over F between A and B. This process retains the information of the CM type since everything is done with $\overline{\mathbf{Q}}$-algebras (and the CM type involves embeddings of L into $\overline{\mathbf{Q}} \subset F \subset \mathbf{C}$). Hence, in this way it suffices to solve the problem over $\mathbf{C}$. But this case has already been settled by analytic methods: we can reconstruct A^{an} up to L-linear isogeny from its CM type Φ as $(\mathbf{R} \otimes_{\mathbf{Q}} L)_{\Phi}/\mathscr{O}_L$. □

Proposition 2.5 (Tate). *Every abelian variety over a finite field* $\mathbf{F}_q$ *has CM over* $\mathbf{F}_q$.

Example 2.6. Let's try to show directly that if $\dim(A) = 1$ then $\mathrm{End}^0(A) \neq \mathbf{Q}$ for A over $\mathbf{F}_q$. If the q-Frobenius endomorphism $\mathrm{Frob}_{A/\mathbf{F}_q} : A \to A$ is not in $\mathbf{Z}$ then there is nothing to do, so we may assume it lies in $\mathbf{Z}$. By the Riemann Hypothesis, this integer must be $q^{1/2}$, so $q = p^{2n}$ with $\mathrm{Frob}_{A/\mathbf{F}_q} = \pm p^n$. According to Honda-Tate theory, for such q these Weil numbers *do* arise from an elliptic curve over $\mathbf{F}_q$, but these are exactly the supersingular elliptic curves over $\mathbf{F}_q$ whose geometric endomorphism algebra, a quaternion division algebra over $\mathbf{Q}$, is defined over $\mathbf{F}_q$. This quaternion division algebra contains many imaginary quadratic fields, so even in these exceptional cases a CM structure can be seen.

3. Arithmetic aspects

Example 3.1. Let $L = \mathbf{Q}(\zeta_7)$ with Galois group $G \simeq (\mathbf{Z}/7\mathbf{Z})^\times$. A CM type Φ is then a choice of representatives of $G/\langle \pm 1\rangle$. Note that Φ will be induced from a subfield $L' = L^H$ (H not containing -1 if L' is CM) if and only if the choice of representatives of $G/\langle \pm 1\rangle$ lifts a choice of representatives of $(G/H)/\langle \pm 1\rangle$ - that, as easily seen, is equivalent to Φ being invariant under the action of H in G. Now since G contains only one subgroup that does not contain -1, namely $\{1, 2, 4\}$ we see that, up to conjugation, there are choices for Φ that yield simple abelian 3-folds with CM - namely $\Phi = \{1, 2, 3\}$ and that yield non-simple abelian 3-folds with CM (e.g. $\Phi = \{1, 2, 4\}$), isogenous to the cube of a CM elliptic curve.

Both examples can be realized over $\overline{\mathbf{Q}}$ by Theorem 3.3 below, and hence over a number field K (since $\overline{\mathbf{Q}}$ is a direct limit of number fields). Since $K = \varinjlim \mathcal{O}_K[1/N]$, such examples arise as abelian schemes over $\mathcal{O}_K[1/N]$ for sufficiently divisible N, which is to say they have "good reduction" at $p \not| N$. (This latter spreading out to a smooth proper $\mathcal{O}_K[1/N]$-scheme is a general formalism unrelated to abelian varieties.)

Remark 3.2. For a simple abelian variety A over a finite field $\mathbf{F}_q$, Tate showed that the center Z of the division algebra $\mathrm{End}^0(A)$ is $\mathbf{Q}(\pi)$ where π is the q-Frobenius endomorphism of A. For any embedding $j : Z \to \mathbf{C}$ the Riemann Hypothesis gives that $j(\pi)\overline{j(\pi)} = q$. Hence, $\overline{j(\pi)} = q/j(\pi) = j(q/\pi)$, so Z is stable under the complex conjugation via j with its restriction to Z described by the formula $\pi \mapsto q/\pi$ that does not involve j. Thus, either Z is totally real (in which case $\pi = \pm p^n$ with $q = p^{2n}$) or Z is a CM field.

Theorem 3.3 (Shimura-Taniyama). *Let A be a CM abelian variety over F of characteristic* 0.

1. *If F is algebraically closed then A descends* (*with the CM structure*) *to* $\overline{\mathbf{Q}} \subset F$, *and hence to a number field.*
2. *If F is a number field then there is a finite extension F'/F such that $A_{F'}$ has good reduction at all finite places of F'.*

Phenomena related to the group scheme α_p in characteristic p yield counterexamples to (1) in positive characteristic using abelian surfaces that are isogenous to the square of a supersingular elliptic curve. Even in characteristic 0, (1) fails if we drop the condition that F is algebraically closed, even when $F \subset \mathbf{C}$. For example, if A is an abelian variety over $\overline{\mathbf{Q}}$ then the quadratic twist A' of $A_{\overline{\mathbf{Q}}(t)}$

by the generator of $\mathrm{Gal}(\overline{\mathbf{Q}}(\sqrt{t})/\overline{\mathbf{Q}}(t))$ makes the Galois action on any Tate module $T_\ell(A')$ for $\overline{\mathbf{Q}}(t)$ non-trivial, so this twist A' cannot descend to $\overline{\mathbf{Q}}$ (as otherwise its torsion would have trivial Galois action over $\overline{\mathbf{Q}}(t)$). By taking A to be a CM abelian variety, A' is also CM (since -1 commutes with everything), so it provides counterexamples to part (1) of the Shimura–Taniyama theorem when F is not assumed to be algebraically closed.

Sketch. Here is a sketch of the idea of the proof of Theorem 3.3. For part (1), we write F as a direct limit of its finitely generated $\overline{\mathbf{Q}}$-subalgebras R, so we descend A with its CM structure to an abelian scheme over such an R such that the tangent space as module over the ring

$$L \otimes_{\mathbf{Q}} R = \prod_{\phi : L \to \overline{\mathbf{Q}}} R_\phi$$

(making $c \in L$ act on R_ϕ via $\phi(c) \in \overline{\mathbf{Q}} \subset R$) has ϕ-component free of rank 1 over R when $\phi \in \Phi$ and 0 when $\phi \notin \Phi$.

Now reducing module a maximal ideal $\mathfrak{m}$ of R gives a CM abelian variety (A_0, L) over $\overline{\mathbf{Q}}$ with CM type Φ. Consider $(A_0)_F$. This has CM by L with CM type Φ since $\mathrm{Hom}(L, F) = \mathrm{Hom}(L, R) = \mathrm{Hom}(L, \overline{\mathbf{Q}})$ and the quotient map $R \to R/\mathfrak{m} = \overline{\mathbf{Q}}$ is of $\overline{\mathbf{Q}}$-algebras. Thus, by 2.4 it follows that $(A_0)_F$ is L-linearly isogenous to A. The kernel of such an isogeny $f : (A_0)_F \to A$ is a finite F-subgroup of $(A_0)_F$, so $\ker(f) \subset (A_0)_F[n] = A_0[n]_F$ for some $n > 0$. But $A_0[n]$ is constant over $\overline{\mathbf{Q}}$ since $\overline{\mathbf{Q}}$ is algebraically closed of characteristic 0, so $\ker(f) = H_F$ for some $H \subset A_0[n]$. Hence, $(A_0/H, L)$ is an L-linear $\overline{\mathbf{Q}}$-descent of A (though its CM order may be smaller than that of A_0).

The proof of part (2) lies much deeper, relying on the semistable reduction theorem for abelian varieties (whose proof involves either the full force of the theory of Néron models or a lot of work on the theory of integral models for curves and the link between such models for a curve and the Néron model of the Jacobian). □

Now we pose the following question: given a CM structure (A, L, Φ) over $\overline{\mathbf{Q}}$, what can we say about number fields $K \subset \overline{\mathbf{Q}}$ to which it descends? Can we “bound below” such fields? For elliptic curves the j-invariant provides some such information. In our case the action of $\mathrm{Gal}(\overline{\mathbf{Q}}/\mathbf{Q})$ on Φ gives a preliminary lower bound (much cruder than the j-invariant in the 1-dimensional case), as follows.

Let L be a CM field of degree $2g$, and $\Phi \subset \mathrm{Hom}(L, \overline{\mathbf{Q}})$ a CM type. Recall that L admits an intrinsic complex conjugation, and we use this to speak of “complex conjugation” on $\mathrm{Hom}(L, \overline{\mathbf{Q}})$ so as to define the concept of a $\overline{\mathbf{Q}}$-valued CM type on L. We do *not* fix embeddings of L or $\overline{\mathbf{Q}}$ into $\mathbf{C}$; these are best viewed as *abstract* fields for our present purposes.

Suppose $F \subset \overline{\mathbf{Q}}$ is a number field so that we can descend A to A_0 over F with CM type (L, Φ). As $L \otimes_{\mathbf{Q}} \overline{\mathbf{Q}}$-modules we have an identification

$$T_0(A_0) \otimes_F \overline{\mathbf{Q}} = \prod_{\phi \in \Phi} \overline{\mathbf{Q}} e_\phi.$$

For $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/F)$, it is easy to check that the twist $\sigma^*(A_0)$ (using scalar extension through $\sigma : \overline{\mathbf{Q}} \simeq \overline{\mathbf{Q}}$) with its induced L-action has CM type $\sigma(\Phi)$. Thus, provided that $\sigma(\Phi) = \Phi$, $\sigma^*(A_0)$ would be L-linearly isogenous to A_0 by Corollary 2.4.

Definition 3.4. For a $\overline{\mathbf{Q}}$-valued CM type (L, Φ), the *reflex field* $E \subset \overline{\mathbf{Q}}$ is the fixed field of the subgroup $\{\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/\mathbf{Q}) \mid \sigma\Phi = \Phi\}$

Next time we will give a "better" definition of the reflex field, in terms of operations with algebraic groups. In view of the discussion above, if A descends to F then $E \subset F$. Intrinsically, if F is any (abstract) field of characteristic 0 over which we are given a triple (A, L, Φ) with Φ consisting of $\overline{F}$-embeddings, then using $\overline{\mathbf{Q}} \subset \overline{F}$ we get $E \subset \overline{\mathbf{Q}}$. That is, although L is an abstract field, the reflex field lies inside the algebraic closure of $\mathbf{Q}$ in $\overline{F}$ by definition.

Considering $V_\Phi := \prod_{\phi\in\Phi} \overline{\mathbf{Q}} e_\phi$ as an $L \otimes_{\mathbf{Q}} \overline{\mathbf{Q}}$-module where L acts on e_ϕ through ϕ, it is an elementary exercise to check that $V_\Phi \simeq \overline{\mathbf{Q}} \otimes_K V_0$ for a subfield $K \subset \overline{\mathbf{Q}}$ and an $L \otimes_{\mathbf{Q}} K$-module V_0 if and only if K contains E, in which case V_0 is unique up to $L \otimes_{\mathbf{Q}} K$-linear isomorphism. Hence, the subfield $E \subset \overline{\mathbf{Q}}$ is the "minimal" field of definition over $\mathbf{Q}$ for the $L \otimes_{\mathbf{Q}} \overline{\mathbf{Q}}$-module V_Φ, and up to non-canonical isomorphism there is a unique descent to an $L \otimes_{\mathbf{Q}} E$-module t_Φ. Loosely speaking, of (L, Φ) arises from a CM abelian variety A over $\overline{\mathbf{Q}}$ then t_Φ is what the tangent space would be for an L-linear descent of A to E if such a descent were to exist (which usually is *not* the case).

Remark 3.5. By definition, E belongs to the Galois closure of L in $\overline{\mathbf{Q}}$ (since the group fixing the Galois closure fixes all embeddings of L). Note that this Galois closure makes intrinsic sense even though L is not given as a subfield of $\overline{\mathbf{Q}}$.

It follows that E is a CM field. Indeed, E is contained in the Galois closure which is CM (as composites of CM fields are CM), so it either CM or totally real. But it cannot be totally real since no complex conjugation preserves Φ (and hence cannot fix E pointwise inside $\mathbf{C}$ for a choice of embedding).

Example 3.6. If L is an "abstract" imaginary quadratic field (i.e., given just as an extension of $\mathbf{Q}$) acting on an elliptic curve A over a field F of characteristic 0 then the F-linear action of L on the 1-dimensional tangent space $T_0(A)$ over F defines an embedding $\phi\colon L \hookrightarrow \mathrm{End}_F(T_0(A)) = F$. Thus, for an algebraic closure $\overline{F}$ of F and its subfield $\overline{Q}$, ϕ is a map $\phi\colon L \hookrightarrow \overline{\mathbf{Q}}$. That is, E is $\phi(L) \subset \overline{\mathbf{Q}} \subset \overline{F}$.

Example 3.7. Choose p and odd prime, and let $L = \mathbf{Q}(\zeta_p)$ and Φ be a set of representatives of $(\mathbf{Z}/p\mathbf{Z})^\times/\langle\pm1\rangle$. Then $E \subset \mathbf{Q}(\zeta_p)$ is the fixed field of all $n \in (\mathbf{Z}/p\mathbf{Z})^\times$ that satisfy $n\Phi = \Phi$.

For example, if $\Phi = \{1, 2, \ldots, \frac{p-1}{2}\}$ then since $n\Phi$ is a progression of difference n, the only way $n\Phi$ can avoid $\{\frac{p+1}{2}, \ldots, p-1\}$ is if $n = 1$, so $E = L = \mathbf{Q}(\zeta_p) \subset \overline{\mathbf{Q}}$ in this case.

If we return to the example $p = 7$ and $\Phi = \{1, 2, 4\}$ which is a group, then Φ is its own stabilizer and the fixed field E associated to Φ is $\mathbf{Q}(\sqrt{-7})$.

It is usually impossible to descend even A, let alone the CM structure, to E. For example, in the case of elliptic curves this happens when the j invariant is not contained in E. There is a better bound, to be discussed next time, called the *field of moduli* $M(L, \Phi)$, that is contained in the Hilbert class of E. In a sense to be made precise later, this lower bound is optimal (generalizing the classical description of class field theory of imaginary quadratic fields in terms of various j-invariants).

4. The main theorem of complex multiplication

There are several related but different results that are called the "Main Theorem". Their proofs will be intertwined, but we state them separately here in a preliminary form. The precise form will be given next time, once we introduce some necessary concepts.

(1) Let (A, L, Φ) be a CM abelian variety over a number field F, and let $\Phi \subset Hom(L, \overline{F})$ be the CM type. Clearly $V_\ell(A) = T_\ell(A) \otimes_{\mathbf{Z}_\ell} \mathbf{Q}_\ell$ is a module over $L \otimes_{\mathbf{Q}} \mathbf{Q}_\ell$, and is invertible (by faithfulness) since both have dimension $2g$ over $\mathbf{Q}_\ell$.

The continuous $\mathbf{Q}_\ell$-linear action of the absolute Galois group G_F is linear over $L \otimes_{\mathbf{Q}} \mathbf{Q}_\ell$ (since the CM structure is "defined over F"), so the action is expressed as a continuous homomorphism $G_F \to (L \otimes_{\mathbf{Q}} \mathbf{Q}_\ell)^\times$ into the open subset of units in $L \otimes_{\mathbf{Q}} \mathbf{Q}_\ell$. This factors through the topological abelianization G_F^{ab}.

But later we will see that 2.4 is valid without the algebraically closed condition, so (A, L) is determined by Φ up to L-linear isogeny over F. Hence, the above character only depends on Φ! Composing with the global Artin reciprocity map (say with normalization that a uniformizer π_v is carried to an arithmetic Frobenius element in G_F^{ab}, we get a continuous homomorphism map $\mathbf{A}_F^\times \to (L \otimes_{\mathbf{Q}} \mathbf{Q}_p)^\times$ depending only on Φ. We want to describe this directly in terms of (L, Φ). This will involve an "algebraic Hecke character" $\mathbf{A}_F^\times \to L^\times$ that is independent of ℓ. (We will define this concept next time.)

(2) Consider (A, L) over $\overline{\mathbf{Q}}$ with $\overline{\mathbf{Q}}$-valued CM type Φ. For the reflex field $E \subset \overline{\mathbf{Q}}$, consider $\sigma \in \mathrm{Gal}(\overline{Q}/E)$. The twist $\sigma^*(A)$ with its natural L-action has CM type $\sigma\Phi = \Phi$ also. We seek an "algebro-geometric" recipe to build $\sigma^*(A)$ with its L-action from (A, L). The proof of this turns out to be the key to the proofs of items (1) and (3) here.

(3) For applications to Shimura varieties, we need an analytic variant on (2): give a direct analytic description of $(\sigma^*(A))^{\mathrm{an}}_{\mathbf{C}}$ in terms of $A^{\mathrm{an}}_{\mathbf{C}}$ (relative to a fixed embedding $\overline{\mathbf{Q}} \hookrightarrow \mathbf{C}$).

In the 1-dimensional case, this is related to the following classical question for CM elliptic curves with maximal CM order $\mathcal{O}_K$ (for an imaginary quadratic field K). Pick a nonzero ideal $\mathfrak{a}$ in $\mathcal{O}_K$ such that $A^{\mathrm{an}}_{\mathbf{C}} \simeq \mathbf{C}/\mathfrak{a}$, and consider a nonzero $\mathcal{O}_K$-ideal $\mathfrak{a}_\sigma$ such that $(\sigma^*(A))^{\mathrm{an}}_{\mathbf{C}} \simeq \mathbf{C}/a_\sigma$. These ideals are well-defined in the class group of K, so it makes sense to ask for a direct description of the operation $\mathfrak{a} \mapsto \mathfrak{a}_\sigma$ on the class group of K in terms of the adelic class field theory description of $\sigma|_{K^{\mathrm{ab}}}$.