FINAL LECTURE: THE PROOF OF THE GALOIS TWISTING FORM AND THE ALGEBRAIC FORM OF THE MAIN THEOREM OF CM (05/24/12-05/31/12)

Recall the premise of the previous lecture. We considered (A, i, Φ) an abelian variety with CM structure, defined over $\overline{\mathbf{Q}}$, which (with all the relevant data) descended to A_0/K where K is a large enough finite Galois extension of E. For an element $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/E)$ the Galois twist $(A_0^\sigma, i^\sigma, \Phi)$ is another CM abelian variety of the same type, and Jeremy and Brandon constructed an explicit isogeny $\xi_{\sigma,\mathfrak{P}} \colon A_0 \to A_0^\sigma$, together with an L-isomorphism $\theta_{\sigma,\mathfrak{P}} \colon N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \simeq A_0^\sigma$ that made the following diagram commute:

$$\begin{tikzcd}
A_0 \arrow[r, "\xi_{\sigma,\mathfrak{P}}"] \arrow[dr] & A_0^\sigma \\
& N_\Phi(\mathfrak{p}) \otimes_{\mathcal{O}_L} A_0 \arrow[u, "\theta_{\sigma,\mathfrak{P}}"', "\simeq"]
\end{tikzcd}$$

The important thing is $\theta_{\sigma,\mathfrak{P}}$ - the identification of A_0^σ with a Serre tensor $N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0$.

Unfortunately, this is highly non-canonical. It depended on the choice of $\mathfrak{P}$ that satisfies $(\frac{K/E}{\mathfrak{P}}) = \sigma\mid_K$ - which as we recall produced $\xi_{\sigma,\mathfrak{P}}$ lifting the $\kappa(\mathfrak{P})$-Frobenius of the reduction mod $\mathfrak{P}$, and that i turn gave $Hom((A_0, i_0), (A_0^\sigma, i_0^\sigma))$ the structure of a fractional ideal of $\mathcal{O}_L$ identified with $N_\Phi(\mathfrak{p})^{-1}$. Even more, the construction depended on the choice of K, as well as the choice of descent (A_0, i_0).

It's not possible to make this canonical for our given setup - but it turns out that focusing, instead of $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/E)$, on ideles $s \in \mathbf{A}_{E,f}^\times$ realizing σ, allows us to define those identifications in a canonical way (depending on s rather than σ). As we see in the end of the lecture, the idelic interpretation is connected with the algebraic form of the main theorem.

Therefore our setup will be the following: given a finite idele $s \in \mathbf{A}_{E,f}^\times$ with $r_E(s) = \sigma\mid_{E^{\mathrm{ab}}}$, we will construct a canonical L-linear isomorphism $[N_\Phi(s)^{-1}]_L \otimes_{\mathcal{O}_L} A \xrightarrow{\sim} A^\sigma$.

This is strongly related to Jeremy's setup: our choice of $\mathfrak{P}$ implies that the adele $s_\mathfrak{p} = \mathfrak{i}_\mathfrak{p}(\pi_\mathfrak{p})$ - the idele whose components are all 1 except for the $\mathfrak{p}$-uniformizer at the $\mathfrak{p}$-th place, precisely realizes σ under the reciprocity map; and the ideal $[N_\Phi(s_\mathfrak{p})]_L$ is precisely $N_\Phi(\mathfrak{p})$. Therefore $\theta_{\sigma,\mathfrak{P}}$ is some sort of "K-level approximation" to the map we want, and we will use such $\theta_{\sigma,\mathfrak{P}}$ to construct it. However $\theta_{\sigma,\mathfrak{P}}$ will not necessarily be the descent of the canonical isomorphism to K - this is expected since $(\frac{K/E}{\mathfrak{P}})$ approximated σ only on K, not on E^{ab}. The difference between $(\theta_{\sigma,\mathfrak{P}})_\mathbf{Q}$ and the canonical isomorphism will be a factor in $N_\Phi(E^\times) \in T(\mathbf{Q}) \subset L$.

We will also use torsion, to assemble all torsions in the end to obtain a result about Tate modules that will give us the main theorem of CM in its algebraic form.

Here is the main theorem we will prove (the Galois twisting form of the main theorem of CM), in a preliminary form:

Theorem 0.1. *Let (A, i) be a CM abelian variety of type (L, Φ) over $\overline{\mathbf{Q}}$, and assume that it is principal i.e. $i^{-1}(\mathrm{End}(A)) = \mathcal{O}_L$. For any $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/E)$ and $s \in A_{E,f}^\times$ with $r_E(s) = \sigma\mid_{E^{\mathrm{ab}}}$ there is* *a unique L-linear isomorphism*

$$\theta_{\sigma,s}\colon [N_\Phi(s)]_L^{-1}\otimes_{\mathcal{O}_L} A \xrightarrow{\sim} A^\sigma$$

such that for all $M\geq 1$ the isomorphism $[\sigma]\colon A[M](\overline{\mathbf{Q}})\xrightarrow{\sim} A^\sigma[M](\overline{\mathbf{Q}})$ equals the composite

$$\begin{tikzcd}
A[M](\overline{\mathbf{Q}}) \arrow[r, "N_\Phi(s^{-1})", "\simeq"'] & ([N_\Phi(s^{-1})]_L/M[N_\Phi(s^{-1})]_L)\otimes_{\mathcal{O}_L/M\mathcal{O}_L} A[M](\mathbf{Q}) \arrow[d, "\simeq"] \\
& ([N_\Phi(s)]_L^{-1}\otimes_{\mathcal{O}_L} A)[M](\overline{\mathbf{Q}}) \arrow[d, "\theta_{\sigma,s}"', "\simeq"] \\
& A^\sigma[M](\overline{\mathbf{Q}})
\end{tikzcd}$$

There are additional specifications on the theorem, but we will define them later.

Remark 0.2. What do we mean by the map

$$N_\Phi(s^{-1})\colon A[M](\overline{\mathbf{Q}})\to ([N_\Phi(s^{-1})]_L/M[N_\Phi(s^{-1})]_L)\otimes_{\mathcal{O}_L/M\mathcal{O}_L} A[M](\overline{\mathbf{Q}})?$$

Note that for any idele $\alpha\in\mathbf{A}_{L,f}^\times$, and a set of places Σ, we can find by weak approximation an element $x\in L$ such that $([x]_L, M\mathcal{O}_L) = ([\alpha]_L, M\mathcal{O}_L)$ and $x\simeq\alpha \pmod{M\alpha\prod_{\mathfrak{p}\in\mathrm{Spec}(\mathcal{O}_L)})\mathcal{O}_\mathfrak{p}}$. This element x will project to a generator of $[\alpha]_L/M[\alpha]_L$ and moreover its image is independent of the choice of x. Then, multiplication by the idele x will just mean multiplication by the image of α modulo M. Alternatively, we can obtain this map as follows: any (integral) idele acts on the total Tate module, and we project this action to M-torsion; this makes sense for any idele that is "M-integral".

Here is the plan of attack:

Step 1. Go back to assembling the $\theta_{\sigma,\mathfrak{P}}$ maps at level K. We deal with the issue mentioned above that the maps $(\theta_{\sigma,\mathfrak{P}})$ are not quite descents of $\theta_{s,\mathfrak{P}}$, the error being a factor in $N_\Phi(E^\times)$. Here is how we find this factor:

Note that $N_\Phi(\mathfrak{p})^{-1}$ is not equal to $[N_\Phi(s^{-1})]_L$ - rather, because of Artin reciprocity, it is a multiple of it. This holds even before applying N_Φ, and so we can find a $c\in E^\times$ such that $N_\Phi(c\mathfrak{p}) = [N_\Phi(s)]_L$. Instead of $\theta_{s,\mathfrak{P}}$ we will let our approximation be the diagonal map in the following

$$\begin{tikzcd}
N_\Phi(\mathfrak{p})^{-1} \arrow[r, "\theta_{\sigma,\mathfrak{P}}", "\simeq"'] \arrow[d, "N_\Phi(c)^{-1}"', "\simeq"] & A_0^\sigma \\
N_\Phi(c\mathfrak{p})^{-1}\otimes_{\mathcal{O}_L} A_0 \arrow[ur, "\simeq"'] \arrow[d, equal] & \\
[N_\Phi(s)]_L^{-1}\otimes_{\mathcal{O}_L} A_0 &
\end{tikzcd}$$

There are also additional constraints on c in terms of the torsion parameter M, to de described later.

Now we base change back to $\overline{\mathbf{Q}}$. This diagonal map $\theta_{\sigma,\mathfrak{P},c,M}$ will depend on a lot of additional data - $\mathfrak{P}, K, A_0, i_0, c, M$ and more; we have to show that all dependences can be removed. For this, we must show two compatibility results: first, that they intertwine certain polarizations (step 2) and second that they agree on torsion (step 3). Together, these two compatibilities imply the maps are all the same.

Step 2. To eliminate the dependencies above, we will introduce a packet of polarizations which will be compatible with the maps $\theta_{\sigma,\mathfrak{P},c,M}$ and these compatibilities will force the $\theta_{\sigma,\mathfrak{P},c,M}$ to agree in the way we like.

Fix some polarization $\phi_0 \colon A_0 \to A_0^\vee$ (called ψ_0 in Brandon's notes).

We will then construct out of it $\mathbf{Q}$-polarizations $\phi_\mathfrak{a}$ of $\mathfrak{a} \otimes_{\mathcal{O}_L} A$ for certain fractional ideals $\mathfrak{a}$ of L (in particular ideals of the form $\mathfrak{a} = [N_\Phi(s)^{-1}]_L$ for $s \in \mathbf{A}_{E,f}^\times$). They will satisfy a simple compatibility property. This datum uniquely descends to K; the polarizations are called $\phi_{0,\mathfrak{a}}$ for each $\mathfrak{a}$.

Then we will show that $\theta_{\sigma,\mathfrak{P},c,s,M}$ intertwine $\phi_{0,[N_\Phi(s)^{-1}]_L}$ and ϕ_0^σ. It follows that the ratio of two such choices will be an automorphism of $(A^\sigma, i^\sigma, \phi^\sigma)$ and any automorphism of a $\mathbf{Q}$-polarized abelian variety (actual variety automorphism, not just in the isogeny category!) has finite order.

Step 3. We now use torsion to finalize the compatibility of the $\theta_{\sigma,\mathfrak{P},c,s,M}$. Recall the commutative diagram

$$\begin{tikzcd} A_0 \arrow[r, "\xi_{\sigma,\mathfrak{P}}"] \arrow[dr] & A_0^\sigma \\ & N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \arrow[u, "\theta_{\sigma,\mathfrak{P}}"', "\simeq"] \end{tikzcd}$$

When restricted to M-torsion, the map $\xi_{\sigma,\mathfrak{P}}$ becomes just the Galois twist $[\sigma] \colon A[M] \to A^\sigma[M]$. This is because it lifts the Frobenius map on $\kappa(\mathfrak{P})$-points and that is exactly $[\sigma]$ by the choice of $\mathfrak{P}$. (Here we implicitly assume good reduction and that the torsion is constant)

It follows that we get the diagram

$$\begin{tikzcd} A_0[M](K) \arrow[r, "{[\sigma]}"] \arrow[dr, "\simeq"'] & A_0^\sigma[M](K) \\ & (N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0)[M](K) \arrow[u, "\theta_{\sigma,\mathfrak{P}}"', "\simeq"] \end{tikzcd}$$

where the diagonal map is an isomorphism is M is coprime to p because upstairs it has p-power degree.

Now combine this with the defining diagram for $\theta_{\sigma,\mathfrak{P},c,s,M}$ to obtain

$$\begin{tikzcd}
A_0[M](\overline{\mathbf{Q}}) \arrow[rr, "{[\sigma]}"] \arrow[d, "\simeq"'] & & A_0^\sigma[M](\overline{\mathbf{Q}}) \\
N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[urr, "\theta_{\mathfrak{P},\sigma}"] \arrow[d, "\simeq"', "N_\Phi(c)^{-1}\otimes 1"] & & \\
{[N_\Phi(s)]_L^{-1}} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[uurr, "\theta_{\sigma,\mathfrak{P},c,s,M}"'] & &
\end{tikzcd}$$

The composite of the top horizontal maps is just multiplication by $N_\Phi(s)^{-1}$ in the sense of the remark above; in particular its inverse is multiplication by $N_\Phi(s)$ and so we flesh out the diagram

$$\begin{tikzcd}
A_0[M](\overline{\mathbf{Q}}) \arrow[r, "{[\sigma]}"] & A_0^\sigma[M](\overline{\mathbf{Q}}) \\
{[N_\Phi(s)]_L^{-1}} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[u, "N_\Phi(s)"] \arrow[ur, "\theta_{\sigma,\mathfrak{P},c,s,M}"'] &
\end{tikzcd}$$

This writes $\theta_{\sigma,\mathfrak{P},c,s,N}$ as a composition of maps that are independent of all the extra choices we need, and thus our maps coincide at least on M-torsion.

It remains to note that since they are of finite order, coinciding on M-torsion for $M \geq 3$ implies they coincide on the entire Tate module by an elementary procedure, and we are done.

Remark 0.3. In the torsion arguments above, we will actually have to assume certain "largeness" properties of K with respect to M, so it is important that we base change back to $\overline{\mathbf{Q}}$ in order to combine different values of M into the limit.

Now we begin the proof of the theorem, according to the plan sketched above.

Let's fix (A, i, ϕ) over $\overline{\mathbf{Q}}$, where $\phi\colon A \to A^\vee$ is an L-linear polarization, and an integer $M \geq 3$. We descend (A, i, ϕ) to (A_0, i_0, ϕ_0) over K, and we choose K large enough so that the following are satisfied: K is a Galois extension of the reflex field E containing all embeddings of L into $\overline{\mathbf{Q}}$, and the group $A_0[M]$ is constant (isomorphic to $((\mathbf{Z}/M\mathbf{Z})^{2g})_K$).

Choose the prime $\mathfrak{P}$ of K as in Jeremy's lecture, with the additional assumption that $\mathfrak{P} \nmid M$. Recall that the restriction to $\mathfrak{P}$ in E is called $\mathfrak{p}$, and in $\mathbf{Q}$ is just p.

On M-torsion, Jeremy's diagram becomes

$$\begin{tikzcd}
A_0[M](K) \arrow[r, "\xi_{\sigma,\mathfrak{P}}"] \arrow[dr, "\simeq"'] & A_0^\sigma[M](K) \\
& (N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0)[M](K) \arrow[u, "\simeq", "\theta_{\sigma,\mathfrak{P}}"']
\end{tikzcd}$$

The diagonal map becomes an isomorphism because the original map $A_0 \to N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0$ had p-power degree - since we can compose further to go down to $\frac{1}{p^k}\mathcal{O}_L \otimes_{\mathcal{O}_L} A_0$ and the map $A_0 \to \frac{1}{p_k}\mathcal{O}_L \otimes_{\mathcal{O}_L} A_0$ can be identified with $[p^k]$. It follows that the rank of the kernel of the map of M-torsion is also a power of p, but it has to divide M^{2g}, so the map is injective hence an isomorphism.

Now we claim that the top map $\xi_{\sigma,\mathfrak{P}}$ is actually equal to the twist $[\sigma]$.

Indeed, the varieties have good reduction at $\kappa(\mathfrak{P})$ and so the diagram (of constant group schemes) is equivalent to the corresponding diagrams over $\kappa(\mathfrak{P})$ but there $\xi_{\sigma,\mathfrak{P}}$ was defined by the Frobenius map which is exactly $[\sigma]$.

In a more rigorous form, this argument is as follows. Consider the Neron model $\mathscr{A}_0$ of A over the discrete valuation ring $(\mathcal{O}_K)_{\mathfrak{P}}$. Then the torsion $\mathscr{A}_0[M]$ is finite etale over $(\mathcal{O}_K)_{\mathfrak{P}}$, and by the Neron mapping property (plus functoriality of the Serre tensor construction) the diagram expands to

$$\begin{tikzcd}
\mathscr{A}_0[M] \arrow[r, "\xi_{\sigma,\mathfrak{P}}"] \arrow[dr, "\simeq"'] & \mathscr{A}_0^{\sigma}[M] \\
& (N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} \mathscr{A}_0)[M] \arrow[u, "\simeq", "\theta_{\sigma,\mathfrak{P}}"']
\end{tikzcd}$$

Over the generic fiber, this is our old diagram. Over the special fiber, $\xi_{\sigma,\mathfrak{P}}$ is by definition equal to $\mathrm{Frob}_{\overline{A_0}/\kappa(\mathfrak{P}),q}$ whose effect on the $\kappa(\mathfrak{P})$-points is exactly the action induced by the arithmetic Frobenius element of $\mathrm{Gal}(\kappa(\mathfrak{P})/\kappa(\mathfrak{p}))$ and that element equals σ. But the torsion groups $\mathscr{A}_0[M]$ are constant, so we can identify morphisms over the special fiber and morphisms over the generic fiber; which shows that $\xi_{\sigma,\mathfrak{P}}$ and $[\sigma]$ indeed coincide.

It only remains to prove that $\mathscr{A}_0[M]$ is indeed a constant group scheme, knowing that its generic fiber is:

Proposition 0.4. *Let R be a normal noetherian ring with total ring of fractions K. The functor $X \rightsquigarrow X_K$ from finite etale R-schemes to finite etale K-schemes is fully faithful. In particular, if X_K is a constant K-scheme then X is a constant R-scheme.*

Proof. Faithfulness is obvious since X is R-flat and affine, so the problem is to prove that for finite etale R-schemes X and X', any K-morphism $X'_K \to X_K$ extends to an R-morphism. The graph of an R-morphism $X' \to X$ is the same thing as a section to $\mathrm{pr}_1\colon X' \times X \to X'$ so it suffices to show that a section of this map over X'_K uniquely extends to a section over X'. But X' is normal noetherian affine since it is finite etale over R, and X'_K is its "scheme of generic points" (i.e., Spec of the total ring of fractions of its coordinate ring), and similarly for $X' \times X$ and $X'_K \times X_K$. Thus, we can replace $X \to \mathrm{Spec}(R)$ with pr_1 to reduce to showing that the natural map $X(R) \to X(K)$ is bijective. For this purpose it is harmless to treat the connected comppnents of X separately, so X is connected. But X is also normal noetherian and R-finite, so by connectedness $X = \mathrm{Spec}(R')$ for a domain R' finite flat over R. Hence, $\mathrm{Frac}(R') = R' \otimes_R K$, so if $R' \neq R$ then there is no K-point (hence nothing to do) whereas if $R' = R$ then the problem is trivial. $\square$

1. Correcting the approximation maps

Recall that for $\ell \nmid N$, the reflex norm on $N_\Phi \colon A_{E,f}^\times \to A_{L,f}^\times$ sends ℓ-integral ideles to ℓ-integral ideles (i.e. sends $\mathcal{O}_{E_\ell}^\times = \prod_{w|\ell} E_w^\times$-components into $\mathcal{O}_{L_\ell}$-components).

In particular, $U \subset \mathbf{A}_{E,f}^\times = \prod U_\ell$ is an open subgroup where

$$U_\ell = \mathcal{O}_{E_\ell}^\times \bigcap N_{\Phi,\ell}^{-1}(\{u \in \mathcal{O}_{L_\ell}^\times \mid u \equiv 1 \pmod{M}\})$$

for all ℓ.

So we require K to contain the class field for the open subgroup $E^\times E_\infty^\times U$.

This condition will ensure that the kernel of the Artin map $\mathbf{A}_{E,f}^\times \to \mathrm{Gal}(K/E)^{\mathrm{ab}}$ is contained in $E^\times U$. We need this assumption because we want to approximate s well enough that our construction will indeed be independent of the approximation - as seen before, $[\mathfrak{p}]$ is way too coarse an approximation so the maps $\theta_{\sigma,\mathfrak{P}}$ are not quite right.

More precisely, we have $\sigma \mid_K = (\frac{K/E}{\mathfrak{P}})$ and $r_E(s) = \sigma \mid_{E^{\mathrm{ab}}}$. Because $\mathfrak{p}$ is unramified, if we pick a uniformizer $\pi_\mathfrak{p}$ and naturally regard it as an E-idele, then $s \cdot \pi_\mathfrak{p}^{-1}$ is in the Kernel of the Arine map and so $s = \pi_\mathfrak{p} uc$ where $c \in E^\times$ and $u \in U$. We see that $N_\Phi(u) \in \mathbf{A}_{L,f}^\times$ is a local integral unit, congruent to 1 modulo M (as an integral idele in L).

The adjustment is that instead of $N_\Phi(\mathfrak{p})^{-1}$ we need to look at $N_\Phi(c\mathfrak{p})^{-1}$. More precisely we have the following commutative diagram

$$\begin{tikzcd}
N_\Phi(\mathfrak{p})^{-1} \arrow[r, "\theta_{\sigma,\mathfrak{P}}", "\simeq"'] \arrow[d, "N_\Phi(c)^{-1}"', "\simeq"] & A_0^\sigma \\
N_\Phi(c\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \arrow[ur, "\simeq"'] \arrow[d, equal] & \\
{[N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A_0} &
\end{tikzcd}$$

where $N_\Phi(c\mathfrak{p}) = [N_\Phi(c\pi_\mathfrak{p})]_L = [N_\Phi(s)]_L$ as u is a unit. The map $N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A_0 \to A_0^\sigma$ produced above will be our "right" approximation. We write it as $\theta_{\sigma,\mathfrak{P},\pi_p,c,s,M}$ because it depends on all these parameters (u is defined by them too).

The next step is to show that after base change to $\overline{\mathbf{Q}}$, these maps agree and only depend on σ, s.

2. Interlude on basic properties of the Serre tensor and Tate modules

Here we describe some easy basic things we will need later; but it is worth just pointing them out.

Recall that there is a natural $\mathbf{A}_{\mathbf{Q},f}^\times$-module structure on $V_f(A)$; integral adeles restrict their action to $T_f(A)$. Since we also have an $\mathcal{O}_L$-action on $T_f(A)$ and an L-action on $V_f(A)$, we see that $V_f(A)$ is canonically an $L \otimes_\mathbf{Q} \mathbf{A}_{\mathbf{Q},f}^\times = \mathbf{A}_{L,f}^\times$-module, and integral adeles restrict their action to $T_f(A)$. (In fact, we showed before that $V_f(A)$ is 1-dimensional free over $\mathbf{A}_{L,f}^\times$ but we do not need it.) We will implicitly use this below.

Some of the stuff we'll be proving is actually unnecessary: Weil pairing are not needed to define the polarizations, so Proposition 2.1 is actually not needed. The Weil pairing properties part is just for fun.

Proposition 2.1. (a) *The Serre tensor construction is functorial in the module component: if $\phi\colon M \to N$ then there is a map $M \otimes_{\mathcal{O}} A \to N \otimes_{\mathcal{O}} A$, and the dependence is functorial in the obvious way. Further, there is a canonical isomorphism $M \otimes_{\mathcal{O}} T_f(A) \simeq T_f(M \otimes_{\mathcal{O}} A)$ and under this identification, the map $M \otimes_{\mathcal{O}} T_f(A) \to N \otimes_{\mathcal{O}} T_f(A)$ corresponding to ϕ is just $\phi \otimes 1$.*

(b) *If $\mathfrak{a}, \mathfrak{b}$ are fractional ideals of L, and $\mathfrak{a} \subset \mathfrak{b}$ then there is a natural isogeny $j_{\mathfrak{a},\mathfrak{b}}\colon \mathfrak{a} \otimes_{\mathcal{O}_L} A \to \mathfrak{b} \otimes_{\mathcal{O}_L} A$, whose degree is the index $[\mathfrak{b} : \mathfrak{a}]$. Without the inclusion requirement, there is a map in the isogeny category $j_{\mathfrak{a},\mathfrak{b}} \in \mathrm{Hom}^0(\mathfrak{a} \otimes_{\mathcal{O}_L} A.\mathfrak{b} \otimes_{\mathcal{O}_L} A)$ of "degree" equal to the "index" (both defined by multiplicativity) - equivalently this number is the positive rational generator of the $\mathbf{Q}$-ideal $N_{L/\mathbf{Q}}(\mathfrak{a}\mathfrak{b}^{-1})$. Further, $j_{\mathfrak{b},\mathfrak{c}} \circ j_{\mathfrak{a},\mathfrak{b}} = j_{\mathfrak{a},\mathfrak{c}}$.*

(c) *Set $j_{\mathfrak{a}} = j_{\mathcal{O}_L,\mathfrak{a}}$. The map $j_{\mathfrak{a}}$ induces an isomorphism of the rational Tate modules $V_f(A) \xrightarrow{\sim} V_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A)$, in the sense of the commutative diagram below. Under this identification, $T_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A)$ inside $V_f(A)$ is identified with $\mathfrak{a} \cdot T_f(A)$. If $\mathcal{O}_L \subset \mathfrak{a}$ so $j_{\mathfrak{a}}$ is an actual isogeny then it becomes the natural inclusion $T_f(A) \to \mathfrak{a} \cdot T_f(A)$ obtained by tensoring the map $\mathcal{O}_L \hookrightarrow \mathfrak{a}$.*

Moreover, the following diagram of $L \otimes_{\mathbf{Q}} A^{\times}_{\mathbf{Q},f}$-module isomorphisms is commutative:

$$\begin{tikzcd}
\mathfrak{a} \otimes_{\mathcal{O}_L} V_f(A) \arrow[r] \arrow[d, "mult"'] & V_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A) \\
V_f(A) \arrow[ur, "V_f(j_{\mathfrak{a}})"'] &
\end{tikzcd}$$

where the horizontal map is natural from the Serre contruction; the vertical map is just multiplication on pure tensors, and the diagonal map is induced by the $\mathbf{Q}$-isogeny $j_{\mathfrak{a}}$ on the rational total Tate module.

(d) *Choose $x \in L$. The map $[x]\colon A \to A$ (which is an isogeny if $x \in \mathcal{O}_L$ and in the isogeny category in general), on the Tate modules equals the map $[x]\colon T_f(A) \to T_f(A)$ which is multiplication by the principal idele x.*

(e) *Suppose $x\mathfrak{a} \subset \mathfrak{b}$. Then the isogeny $j_{x\mathfrak{a},\mathfrak{b}} \circ [x]$ from $\mathfrak{a} \otimes_{\mathcal{O}_L} A \to \mathfrak{a} \otimes_{\mathcal{O}_L} A$, with respect to the identification in c), produces the map $\mathfrak{a} \cdot V_f(A) \to \mathfrak{b} \cdot V_f(A)$ equal to multiplication by x.*

(f) *$(\mathfrak{a} \otimes_{\mathcal{O}_L} A)^\vee \simeq (\mathfrak{a}^{-1})^* \otimes_{\mathcal{O}_L} A^\vee$ and $j_{\mathfrak{a}}^\vee = j^{-1}_{(\mathfrak{a}^{-1})^*}$.*

Proof. a) On the functor of points get the natural map $M \otimes_{\mathcal{O}} A(T) \to N \otimes_{\mathcal{O}} A(T)$.

For the result on Tate modules, let's work over $\mathbf{C}$ so assume A is a $\mathbf{C}$-abelian variety. Its complex points form a commutative complex lie group, and consider the exponential (uniformization) map $\exp\colon V \to A(\mathbf{C})$ where V is the tangent space to A at the identity, whose kernel is the lattice Λ. It is easy to show tha for any integer $n \geq 1$, the n-torsion of A is canonically identified with $\frac{1}{n}\Lambda/\Lambda \subset A(\mathbf{C})$.

For any $m \in M$, the scheme map $A \to M \otimes_{\mathcal{O}} A$ defined functorially by $x \mapsto m \otimes x$ induces on $\mathbf{C}$-points the map $A(\mathbf{C}) \to M \otimes_{\mathcal{O}} A(\mathbf{C}) = (M \otimes_{\mathcal{O}} A)(\mathbf{C})$ of complex tori given on points by $x \mapsto m \otimes x$.

This lifts uniquely to analytic uniformizations $T_m\colon V \to V_M$ where $\exp_M\colon V_M \to (M \otimes_{\mathcal{O}} A)(\mathbf{C}) = M \otimes_{\mathcal{O}} A(\mathbf{C})$ is the exponential uniformization of $(M \otimes_{\mathcal{O}} A)(\mathbf{C})$.

We make V into an $\mathcal{O}$-module (linear over its $\mathbf{C}$-structure) by unique lifting from $A(\mathbf{C})$, and similarly for V_M. By uniqueness, T_m is automatically $\mathcal{O}$-linear and varies $\mathcal{O}$-linearly in m, since this is true for the maps $A(\mathbf{C}) \to M \otimes_{\mathcal{O}} A(\mathbf{C})$ that T_m lift.

It follows that the maps T_m for varying map define an $\mathcal{O}$-linear map $\omega_M\colon M \otimes_{\mathcal{O}} V \to V_M$. We claim that this map is an isomorphism that carries $M \otimes_{\mathcal{O}} \Lambda$ isomorphically onto $\Lambda_M := ker(\exp_M)$. Moreover, by construction this is all functorial in the $\mathcal{O}$-module M.

To prove these maps we first note that T_M is compatible with direct sums in M, due to the compatibility of exponential uniformization with respect to direct products, and so $T_{M\oplus M'}$ is $T_M \oplus T_{M'}$. Since every projective module is a direct summand of a free module, we have thus reduced to the case $M = \oplus\mathcal{O}$. By compatibility with direct sums again, this reduces to the case $M = \mathcal{O}$, in which case it is quite easy to see that under the canonical identification $\mathcal{O} \otimes_{\mathcal{O}} V \simeq V$ the map T_M becomes the identity (because T_m is just multiplication by m, since this is true functorially using the definition of Serre tensor on the functor of points), and the conclusion trivially holds.

After this is done, we see that T_M induces isomorphisms between $M \otimes_{\mathcal{O}} (\frac{1}{n}\Lambda/\Lambda)$ and $\frac{1}{n}\Lambda_M/\Lambda_M$ for all $n \geq 1$ and passing to the direct limit we obtain our desired identification of total Tate modules $M \otimes_{\mathcal{O}} T_f(A) \xrightarrow{\sim} T_f(M \otimes_{\mathcal{O}} A)$.

Finally, if $\phi\colon M \to N$ is a map of $\mathcal{O}$-modules call $j_\phi\colon M \otimes_{\mathcal{O}} A \to N \otimes_{\mathcal{O}} A$ the map associated to ϕ. We need to show that $T_f(j_\phi)$ corresponds to $\phi\otimes 1$ under the identifications $M \otimes_{\mathcal{O}} T_f(A) \simeq T_f(M \otimes_{\mathcal{O}} A)$ and $N \otimes_{\mathcal{O}} T_f(A) \simeq T_f(M \otimes_{\mathcal{O}} A)$ just exhibited.

It is quite obvious by definition that the following diagram commutes

$$\begin{tikzcd} V \arrow[r, "T_m"] \arrow[dr, "T_{\phi(m)}"'] & V_M \arrow[d, "T(j_\phi)"] \\ & V_N \end{tikzcd}$$

in which the vertical map $T(j_\phi)$ being the map of tangent spaces corresponding to j_ϕ under the previous discussion (alternatively the lift of j_ϕ on $\mathbf{C}$-points).

By $\mathcal{O}$-linearity, the combine together into the diagram

$$\begin{tikzcd} M \otimes_{\mathcal{O}} V \arrow[r, "T_M"] \arrow[d, "\phi\otimes 1"'] & V_M \arrow[d, "T(j_\phi)"] \\ N \otimes_{\mathcal{O}} V \arrow[r, "T_N"] & V_N \end{tikzcd}$$

Reducing this diagram to the n-torsion and then taking the inverse limit in n identifies $\phi \otimes 1$ with $T_f(\phi)$ as desired.

b) Apply part a). For the degree, base change to $\mathbf{C}$, and identify A with its analytification (these operations preserve degrees of maps). We can identify A^{an} with V/Λ for a torsion-free finitely generated $\mathcal{O}_L$-module Λ of rank 1. In that case, $\mathfrak{a} \otimes_{\mathcal{O}_L} A$ is identified with $V/(\mathfrak{a}\cdot\Lambda)$ and $\mathfrak{b} \otimes_{\mathcal{O}_L} A$ with $V/(\mathfrak{b}\cdot\Lambda)$ - this fact is easy to prove using the functoriality of the Serre tensor construction. If $\mathfrak{a} \subset \mathfrak{b}$ the map is just induced by the inclusion $\mathfrak{a} \hookrightarrow \mathfrak{b}$ and the kernel is identified with $\mathfrak{b} \setminus \mathfrak{a}$ whose size is the index of $\mathfrak{a}$ in $\mathfrak{b}$. Also this is $N_{L/\mathbf{Q}}(\mathfrak{a}\mathfrak{b}^{-1})$ regarded as a positive integer. For the general case, pre-multiply by an integer to get an honest isogeny, and use the fact that multiplication has the degree we want.

c) The identification $\mathfrak{a} \otimes_{\mathcal{O}_L} T_f(\mathfrak{a}) \simeq T_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A)$ constructed in a), when tensored with L over $\mathcal{O}_L$ becomes $\mathfrak{a} \otimes_{\mathcal{O}_L} V_f(\mathfrak{a}) \simeq V_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A)$. Since $\mathfrak{a} \otimes_{\mathcal{O}_L} L \simeq L$ and the isomorphism is obtained by the multiplication, we obtain the following diagram of isomorphisms

$$\begin{tikzcd}
\mathfrak{a} \otimes_{\mathcal{O}_L} V_f(A) \arrow[r] \arrow[d, "\text{mult}"'] & V_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A) \\
V_f(A) \arrow[ur, "?"'] &
\end{tikzcd}$$

We need to figure out what the map labeled by "?" is.

Assume first that $\mathcal{O}_L \subset \mathfrak{a}$ so $j_{\mathfrak{a}}$ is an actual isogeny induced by $\iota\colon \mathcal{O}_L \in \mathfrak{a}$. Then the map *mult* has as inverse the map $V_f(A) = V_f(\mathcal{O}_L \otimes_{\mathcal{O}_L} A) = \mathcal{O}_L \otimes_{\mathcal{O}_L} V_f(A) \xrightarrow{\iota\otimes 1}$ (it is obvious that it is a left inverse, hence it is also a right inverse since we know *mult* is an isomorphism). It follows that under the identification of $\mathfrak{a} \otimes_{\mathcal{O}_L} V_f(A)$ with $V_f(\mathfrak{a} \otimes_{\mathcal{O}_L} A)$, "?" corresponds to $\iota \otimes 1$; but we have shown in a) that the map corresponding to $\iota \otimes 1$ is precisely $V_f(j_{\mathfrak{a}})$.

This finishes the proof in the case when $j_{\mathfrak{a}}$ is an actual isogeny. In the general case, just multiply with some $c \in \mathbf{Z}$ to make it an honest map, and apply the discussion above. The factor c will simplify since everything is $\mathbf{Q}$-linear.

d) By definition.

e) Follows by the previous two parts.

f) On points, the dual abelian variety represents line bundles etc. plus we have an action of $\mathcal{O}_L$. Alternatively, go to $\mathbf{C}$, perform analytification and use the fact that $(\mathfrak{a})^\vee$ is functorially identified with $(\mathfrak{a})^{-1}$. □

Now we discuss how polarizations enter the picture.

Definition 2.2. Assume that $\theta \in \mathrm{Hom}^0(A,B)$ and $\phi\colon B \to B^\vee$ is a $\mathbf{Q}$-polarization. Define $\theta^*\phi\colon A \to A^\vee$ be the $\mathbf{Q}$-polarization equal to $\theta^\vee \circ \phi \circ \theta$ i.e. which makes the following diagram commute

$$\begin{tikzcd}
A \arrow[r, "\theta^*\phi"] \arrow[d, "\theta"'] & A^\vee \\
B \arrow[r, "\phi"'] & B^\vee \arrow[u, "\theta^\vee"']
\end{tikzcd}$$

If $\theta \in \mathrm{Hom}(A,B)$ and ϕ is a polarization, then $\theta^*\phi$ is also a polarization.

This definition also commutes with composition.

Proposition 2.3. *If ϕ is an L-polarization then $[x]^*\phi = (xx^*)\phi$ for $x \in L$.*

Proof. The map $[x]^\vee$ equals just $[x^*]$ (i.e. the action of x^* on $A^\vee$) by construction of the dual action. Therefore $[x]^*\phi = [x^*] \circ \phi \circ [x]$ and we can swap ϕ and x because ϕ commutes with the L-action. □

Definition 2.4. For a number field K, let I_K denote the group of fractional ideals of K. Define the "Serre torus of ideals" $T(I_L)$ to consist of those fractional ideals $\mathfrak{a}$ of L, for which $\mathfrak{a}\mathfrak{a}^*$ is actually an ideal of $\mathbf{Q}$; the positive rational generator of this ideal is written $[\mathfrak{a}\mathfrak{a}^*]^+_{\mathbf{Q}}$. Alternatively, $T(I_L)$ can be defined as $[T(\mathbf{A}^\times_{\mathbf{Q},f})]_L$, the L-ideals corresponding to the adelic points of the Serre torus.

The Weil pairing stuff, in particular the next proposition is completely optional and can be skipped.

Recall that every $\mathbf{Q}$-polarization ϕ induces a Weil pairing $e_\phi\colon V_f(A) \times V_f(A) \to \mathbf{A}^\times_{\mathbf{Q},f}$ and if it is an actual polarization it restricts to $e_\phi\colon T_f(A) \times T_f(A) \to \hat{\mathbf{Z}}$. This is a priori a $\mathbf{Q}$-bilinear pairing, but if ϕ is an L-linear polarization we actually get the $\mathbf{A}^\times_{L,f}$-linear adelic Weil pairing $e_\phi\colon V_f(A) \times V_f(A) \to \mathbf{A}^\times_{L,f}$.

Proposition 2.5 (optional)**.** (a) *Assume that $\theta = j_{x\mathfrak{a},\mathfrak{b}} \circ [x]$ and $x \in T(\mathbf{A}^\times_{f,\infty})$. Then the pairing $e_{\theta^*\phi}\colon V_f(A) \times V_f(A)$ equals $xx^* e_\phi$ In particular if $\theta = j_{\mathfrak{a},\mathfrak{b}}$ then the Weil pairing remains the same (but the integral Tate module inside $V_f(A)$ may change)*

(b) *If $\theta = [x] \cdot \theta'$ then $e_\theta = x e_{\theta'}$ for $x \in \mathbf{Q}$*

(c) *If e is a Weil pairing, then the image $e(\mathfrak{a} \cdot T_f(A), \mathfrak{a} \cdot T_f(A))$ equals $\mathfrak{a}\mathfrak{a}^* e(T_f(A), T_f(A))$ for $\mathfrak{a} \in$ "$T(I_{\mathbf{Q}})$" (meaning ideals that when multiplied with their conjugate become $\mathbf{Q}$-ideals).*

Proof. a) First assume there is no j map i.e. $x\mathfrak{a} = \mathfrak{b}$. If x would be principal, then $\theta^*\phi$ would be $xx^* \circ \phi$ because ϕ commutes with the L-action, and $xx^* \in \mathbf{Q}$ so the equality holds. The same holds when $x \in \mathbf{A}^\times_{\mathbf{Q},f}$ and now apply $T(L)\mathbf{A}^\times_{\mathbf{Q},f} = T(\mathbf{A}^\times_{\mathbf{Q},f})$

Now for the j map, apply part c) of the previous proposition to see that j preserves $V_f(A)$ and so composition with j preserves the Weil pairing. Apply also part f) for composition with the dual.

b) Trivial by linearity.

c) Just write $\mathfrak{a}$ times the integral ideal as x times the integral ideles where x is as in a). Now we apply part a). ∎

3. The construction of the polarizations

Now let us construct the actual polarizations ϕ_s.

Theorem 3.1. *Let (A, i) be a CM abelian variety of type (L, Φ) over $\overline{\mathbf{Q}}$ with CM order $\mathcal{O}_L$, and let ϕ be an L-linear $\mathbf{Q}$-polarization on A. There is a unique way to assign to each $\mathfrak{a} \in T(I_L)$ an L-linear polarization $\phi_\mathfrak{a}$ of $\mathfrak{a} \otimes_{\mathcal{O}_L} A$ over $\overline{\mathbf{Q}}$ such that*

(1) $\phi_1 = \phi$,

(2) $\phi_\mathfrak{a} = \frac{1}{[\mathfrak{a}\mathfrak{a}^*]^+_{\mathbf{Q}}} (j_\mathfrak{a}^{-1})^*\phi$,

(3) (optional) *if $h\colon \mathfrak{a} \to \mathfrak{b}$ equals the map $[x] \circ j_{\mathfrak{a},x^{-1}\mathfrak{b}} = j_{x\mathfrak{a},\mathfrak{b}} \circ [x]$ for $x \in T(\mathbf{Q}) \subset L$ then $h^*\phi_\mathfrak{b} = \frac{[xx^*\mathfrak{a}\mathfrak{a}^*]^+_{\mathbf{Q}}}{[\mathfrak{b}\mathfrak{b}^*]^+_{\mathbf{Q}}} \phi_\mathfrak{a}$,*

(4) *if everything descends to K, and $\mathfrak{a} = [N_\Phi(s)]_L^{-1}$ for $s \in \mathbf{A}_{E,f}^\times$, then the maps $\theta_{\sigma,\mathfrak{P},\pi_\mathfrak{p},c,s,M}$ previously constructed intertwine ϕ^σ and $\phi_\mathfrak{a}$; i.e., the following diagram commutes:*

$$\begin{tikzcd} {[N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A_0} \arrow[r, "\phi_\mathfrak{a}"] \arrow[d, "\theta_{\sigma,\mathfrak{P},\pi_\mathfrak{p},c,s,M}"'] & {([N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A_0)^\vee} \\ A_0^\sigma \arrow[r, "\phi^\sigma"'] & (A_0^\sigma)^\vee \arrow[u, "\theta^\vee_{\sigma,\mathfrak{P},\pi_\mathfrak{p},c,s,M}"'] \end{tikzcd}$$

so (in terms of descent to K) $\phi_\mathfrak{a} = \theta^*_{\sigma,\mathfrak{P},\pi_\mathfrak{p},c,s,M}\phi^\sigma$.

(5) (optional) *Identifying $V_f([\mathfrak{a}] \otimes_{\mathcal{O}_L} A)$ with $V_f(A)$, the Weil pairing $e_{\phi_\mathfrak{a}}$ equals $\frac{1}{[\mathfrak{a}\mathfrak{a}^*]^+_\mathbf{Q}} e_\phi$*

Moreover, deg(ϕ_α) = deg(ϕ) and these $\mathbf{Q}$-polarizations descend uniquely to avery subfield K over which the variety and its CM-type descend.

Proof. Note that $j_\mathfrak{a}^{-1} = j_{\mathfrak{a},\mathcal{O}_L}$. The degree equality follows from (2) because $j^*_{\mathfrak{a},\mathcal{O}_L}(\phi) = j_{(\mathfrak{a}^*)^{-1}} \circ \phi \circ j_{\mathfrak{a},\mathcal{O}_L}$. As computed before, $j_{\mathfrak{a},\mathcal{O}_L}$ has degree $N_{L/\mathbf{Q}}(\mathfrak{a})$ and similarly $j_{(\mathfrak{a}^*)^{-1}}$ has degree $N_{L/\mathbf{Q}}(\mathfrak{a}^*)$ so together we get $N_{L/\mathbf{Q}}(\mathfrak{a}\mathfrak{a}^*) = ([\mathfrak{a}\mathfrak{a}^*]^+_\mathbf{Q})^{2g}$ since the ideal $\mathfrak{a}\mathfrak{a}^*$ is in $\mathbf{Q}$ and on those ideals norm is just raising to $[L : \mathbf{Q}]$-power. We offset this by dividing by $[\mathfrak{a}\mathfrak{a}^*]^+_\mathbf{Q}$ which has degree precisely $([\mathfrak{a}\mathfrak{a}^*]^+_\mathbf{Q})^{2g}$ yielding the cancellation.

For descent, we know that descent theory is effective for polarizations. This was done in Arnav's talk, where we could descend the line bundles and ampleness was also well-behaved with descent, etc. Now also descent theory is "effective" for the Serre tensor construction in the obvious way (if everything descends so does the Serre tensor). In particular, the descent of $\phi_\mathfrak{a}$ to the Serre tensor over K, satisfies the same properties.

Now (1) and (2) can be taken as the definition of $\phi_\mathfrak{a}$; they also cover uniqueness.

Condition (3) is just a slight generalization of (2), and follows from it. Note that the maps j and their duals all commute with the L-action. Also note that that the map h can be regarded as $j_{\mathfrak{a},\mathfrak{b}} \circ [x]$ if we regard $[x]$ as the $\mathbf{Q}$-map $\mathfrak{a} \otimes_{\mathcal{O}_L} A \to \mathfrak{a} \otimes_{\mathcal{O}_L} A$ induced from $[x] \in Hom_L(\mathfrak{a}, \mathfrak{a})$. Finally, write $j_{\mathfrak{a},\mathfrak{b}} = j_\mathfrak{b} \circ j_\mathfrak{a}^{-1}$ to obtain $h = j_\mathfrak{b} \circ j_\mathfrak{a}^{-1} \circ [x]$

Thus, $h^*(\phi_\mathfrak{b}) = [x]^*((j_\mathfrak{a}^{-1})^*(j_\mathfrak{b}^*)(\frac{1}{[\mathfrak{b}\mathfrak{b}^*]^+_\mathbf{Q}}(j_\mathfrak{b}^{-1})^*(\phi)))$

Move $\frac{1}{[\mathfrak{b}\mathfrak{b}^*]^+_\mathbf{Q}}$ in front by commutativity, $j_\mathfrak{b}^*$ and $j_\mathfrak{b}^{-1}$ cancel each other, and recall that $[x]^*$ is just multiplication by $[xx^*]^+_\mathbf{Q}$. The result is therefore $\frac{[xx^*]^*_\mathbf{Q}}{[\mathfrak{b}\mathfrak{b}^*]}(j_\mathfrak{a}^{-1})^*(\phi)$ and since $\phi_\mathfrak{a} = \frac{1}{[\mathfrak{a}\mathfrak{a}^*]}(j_\mathfrak{a}^{-1})^*(\phi)$ we deduce (3).

We move to (4). This is just a computation, once we write $\theta_{\sigma,\mathfrak{P},c,s,M}$ as a composite in the isogeny category:

Combine the two essential diagrams:

$$
\begin{tikzcd}
 & A_0 \arrow[dl, "j_{N_\Phi(\mathfrak{p})^{-1}}"'] \arrow[d, "\xi_{\sigma,\mathfrak{P}}"] \\
N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \arrow[r, "\theta_{\sigma,\mathfrak{P}}"] \arrow[d, "N_\Phi(c)^{-1}"'] & A_0^\sigma \\
N_\Phi(c\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \arrow[ur, "\theta_{\sigma,\mathfrak{P},c,s,M}"'] &
\end{tikzcd}
$$

Going across the boundary and inverting the arrows, we deduce that $\theta_{\sigma,b,c,s,M}$ equals $\theta_{\sigma,\mathfrak{P}} \circ j_{N_\Phi(cN_\Phi(\mathfrak{p})^{-1},N_\Phi(\mathfrak{p})^{-1})}[N_\Phi(c)]$ but $\theta_{\sigma,\mathfrak{P}} = \xi_{\sigma,\mathfrak{P}} \circ (j_{N_\Phi(\mathfrak{p})^{-1}})^{-1}$ and since $(j_{N_{\Phi(\mathfrak{p})}^{-1}})^{-1} \circ j_{N_\Phi(cN_\Phi(\mathfrak{p})^{-1},N_\Phi(\mathfrak{p})^{-1})}$ is just $j_{N_\Phi(cN_\Phi(\mathfrak{p}))^{-1},\mathcal{O}_L}$ we obtain $\theta_{\sigma,\mathfrak{P},c,s,M}$ as

$$[N_{\Phi(s)}]_L^{-1} \otimes_{\mathcal{O}_L} A_0 = N_\Phi(c\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \xrightarrow{N_\Phi(c)} N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0 \xrightarrow{j_{N_\Phi(\mathfrak{p})^{-1},\mathcal{O}_L}} A_0 \xrightarrow{\xi_{\sigma,\mathfrak{P}}} A_0^\sigma$$

Now let's compute the pullback of ϕ^σ with respect to this map. We already know that $\xi_{\sigma,\mathfrak{P}}^{-1}(\phi^\sigma) = q\phi$ for some $q \in \mathbf{Q}^+$. The constant q was given as degree of $\xi_{\sigma,\mathfrak{P}}$ to the power $1/g$, and since $\theta_{\sigma,\mathfrak{P}}$ is an isomorphism, the degree of $\xi_{\sigma,\mathfrak{P}}$ equals the degree of $j_{N_\Phi(\mathfrak{p})^{-1}}$ i.e. the positive rational generator of $Nm_{L/\mathbf{Q}}(N_\Phi(\mathfrak{p}))$. Now whatever this norm is, it is clearly also the norm of its conjugate so by taking the geometric mean we get

$$q = (Nm_{L/\mathbf{Q}}(N_\Phi(\mathfrak{p})(N_\Phi)^*(\mathfrak{p})))^{\frac{1}{2g}} = (Nm_{L/\mathbf{Q}}(Nm_E/\mathbf{Q}(\mathfrak{p})))^{\frac{1}{2g}} = Nm_{E/\mathbf{Q}}(\mathfrak{p})$$

[by abuse of notation, we identify fractional ideals of $\mathbf{Q}$ with their positive generator].

After this, we pull back $q\phi_1$ with respect to the map $j_{N_\Phi(\mathfrak{p})^{-1},\mathcal{O}_L} \circ [N_\Phi(c)]$ which, according to condition (3), equals q times $\frac{[cc^*(c\mathfrak{p})^{-1}((c\mathfrak{p})^*)^{-1}]^+_{\mathbf{Q}}}{[\mathcal{O}_L\mathcal{O}_L^*]^+_{\mathbf{Q}}}\phi_{N_\Phi(c\mathfrak{p})^{-1}} = [(N_\Phi(\mathfrak{p})N_\Phi(\mathfrak{p})^*)^{-1}]^+_{\mathbf{Q}}\phi_{[N_\Phi(c\mathfrak{p})^{-1}]} = Nm_{E/\mathbf{Q}}(\mathfrak{p})^{-1}\phi_{[N_\Phi(c\mathfrak{p})^{-1}]}$ and since $q = Nm_{E/\mathbf{Q}}(\mathfrak{p})$ the cancellation gives $\phi_{[N_\Phi(c\mathfrak{p})^{-1}]} = \phi_{[N_\Phi(s)]_L^{-1}}$

Condition (5) follows immediately from the definition of $\phi_{\mathfrak{a}}$ given by (2) and proposition 2.5.a). □

4. Torsion and polarization compatibility implies compatibility

Now assemble the results before to obtain the compatibility of the maps $\theta_{\sigma,\mathfrak{P},c,s,M}$.

We got the diagram

$$
\begin{tikzcd}
A_0[M](K) \arrow[r, "{[\sigma]}"] \arrow[dr, "\simeq"'] & A_0^\sigma[M](K) \\
 & (N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0)[M](K) \arrow[u, "\simeq"', "\theta_{\sigma,\mathfrak{P}}"]
\end{tikzcd}
$$

where the diagonal map is an isomorphism if M is coprime to p because upstairs it has p-power degree. In fact it's not just an isomorphism; it is the identity once we identify $A_0[M](K)$ with $T_f(A)[M]$

and $(N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0)[M](K)$ with $([N_\Phi(\mathfrak{p})^{-1}]_L \cdot T_f(A))[M] = ((([N_\Phi(\mathfrak{p})^{-1}]_L/M([N_\Phi(\mathfrak{p})^{-1}]_L) \otimes_{\mathcal{O}_L/M\mathcal{O}_L} T_f(A)$; the two are equal because the natural map $\mathcal{O}_L/M\mathcal{O}_L \to ([N_\Phi(\mathfrak{p})^{-1}]_L)$ is an isomorphism if $(M,p) = 1$.

Now combine this with the defining diagram for $\theta_{\sigma,\mathfrak{P},c,s,M}$ to obtain

$$\begin{tikzcd}
A_0[M](\overline{\mathbf{Q}}) \arrow[rr, "{[\sigma]}"] \arrow[d, "\simeq"] & & A_0^\sigma[M](\overline{\mathbf{Q}}) \\
N_\Phi(\mathfrak{p})^{-1} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[urr, "\theta_{\mathfrak{P},\sigma}"] \arrow[d, "N_\Phi(c)^{-1}\otimes 1"', "\simeq"] & & \\
{[N_\Phi(s)]_L^{-1}} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[uurr, "\theta_{\sigma,\mathfrak{P},c,s,M}"'] & &
\end{tikzcd}$$

The composite of the top horizontal maps is just multiplication by $N_\Phi(c)^{-1}$ in the sense discussed above; this equals multiplication by $N_\Phi(s)^{-1}$ because s differs from c by a factor of $\pi_\mathfrak{p}$, which acts trivially on M-torsion because $\pi_p - 1$ is M times an integral idele, and u is 1 modulo M so is also trivial on M-torsion. In particular its inverse is multiplication by $N_\Phi(s)$ and so we flesh out the diagram

$$\begin{tikzcd}
A_0[M](\overline{\mathbf{Q}}) \arrow[r, "{[\sigma]}"] & A_0^\sigma[M](\overline{\mathbf{Q}}) \\
{[N_\Phi(s)]_L^{-1}} \otimes_{\mathcal{O}_L} A_0[M](\overline{\mathbf{Q}}) \arrow[u, "N_\Phi(s)"] \arrow[ur, "\theta_{\sigma,\mathfrak{P},c,s,M}"'] &
\end{tikzcd}$$

This writes $\theta_{\sigma,\mathfrak{P},c,s,M}$ as a composition of maps that are independent of all the extra choices we need, and thus our maps coincide at least on M-torsion.

Also, the maps do the same to $\phi_{[N_\Phi(s)]_L^{-1}}$ (both take it to ϕ^σ). For for a different choice $\theta_{\sigma,\mathfrak{P}',c',s,M'}$ the error $\theta_{\sigma,\mathfrak{P}',c',s,M'} \circ \theta^{-1}_{\sigma,\mathfrak{P},c,s,M}$ agrees on $gcd(M,M')$-torsion using the diagram above and is an automorphism of the polarized CM abelian variety $(A_0^\sigma, i_0^\sigma, \phi_0^\sigma)$.

Lemma 4.1. *The automorphism group of a principally polarized abelian variety is finite.*

Proof. Recall from Arnav's talk, the Rosati involution $\star$ on $\Delta := \mathrm{End}^0(A)$ that is induced from our polarization. If ψ is the polarization and $x\colon A \to A$ is an isogeny, then $x^\star$ is defined to be $\psi^{-1} \circ x^\vee \circ \psi$. It is also proved in Mumford that the map $x \mapsto \mathrm{Tr}(xx^*)$ is positive-definite, where $\mathrm{Tr} = \mathrm{Tr}_{Z/\mathbf{Q}} \circ \mathrm{Trd}_{\Delta/Z}$ for the center Z of Δ.

Now if $x^*\psi = \psi$ (i.e. $x^\vee \circ \psi \circ x = \psi$) then $xx^\star = x\psi^{-1}x^\vee\psi$ equals 1. Indeed, $\psi^{-1} = x^{-1}\psi^{-1}(x^\vee)^{-1}$ and multiplying by x to the left and $x^\vee$ to the right we get $x\psi^{-1}x^\vee = \psi^{-1}$ and now multiply by ψ to the right to get $xx^\star x\psi^{-1}x^\vee\psi = 1$. Therefore $\mathrm{Tr}(xx^\star) = [Z : \mathbf{Q}]\sqrt{[\Delta : Z]}$, and since $\mathrm{Tr}(xx^\star)$ is positive definite it follows that the level set

$$\{x \in \Delta_{\mathbf{R}} \,|\, \mathrm{Tr}(xx^\star) = [Z:\mathbf{Q}]\sqrt{[\Delta:Z]}\}$$

is compact. However $\mathrm{End}(A)$ is a lattice in $\Delta_{\mathbf{R}}$ is a lattice, so it is discrete and hence the overlap

$$\mathrm{End}(A) \bigcap \{x \in \Delta_{\mathbf{R}} \,|\, \mathrm{Tr}(xx^\star) = [Z:\mathbf{Q}]\sqrt{[\Delta:Z]}\}$$

is compact and discrete, hence finite ◻

In particular, our discrepancy automorphism $\theta_{\sigma,\mathfrak{P}',c',s,M'} \circ \theta^{-1}_{\sigma,\mathfrak{P},c,s,M}$ preserves the $\gcd(M,M')$-torsion and has finite order. We conclude that $\theta_{\sigma,\mathfrak{P}',c',s,M'} = \theta^{-1}_{\sigma,\mathfrak{P},c,s,M}$ once $\gcd(M,M') \geq 3$ by the following:

Proposition 4.2. *Assume that σ is a finite-order automorphism of an abelian variety A over a field k. If $M \geq 3$ is not divisible by the characteristic of k and σ is trivial on $A[M](\overline{k})$, then it is trivial.*

Proof. Every integer greater than 2 has a prime power divisor greater than 2. Therefore it suffices to assume $M = p^n \geq 3$ is a prime power $p \neq \mathrm{char}(k)$. We also may and do assume k is algebraically closed.

Recall that $\mathrm{End}_{\overline{\mathbf{Q}}}(A)$ injects into $\mathrm{End}_{\mathbf{Z}_p}(T_p(A))$ so it suffices to show that σ is trivial on $T_p(A) = \varprojlim A[p^m]$. Since $A[p^n]$ is a constant group scheme $(\mathbf{Z}/p^n\mathbf{Z})^{2g}$, any automorphism of $T_p(A)$ is just an element of $\mathrm{GL}_{2g}(\mathbf{Z}_p)$. We are left to prove the following elementary assertion:

Lemma 4.3. *Let p be a prime and choose $n \geq 1$ such that $p^n \geq 3$. If $S \in \mathrm{GL}_m(\mathbf{Z}_p)$ has finite order and $S \equiv I_m \pmod{p^n}$ then $S = 1$.*

To prove the lemma, recall that for any p-adic field K and $q \in \mathcal{O}_K$ with $\mathrm{ord}_p(q) > \frac{1}{p-1}$ the log and exp series give inverse homomorphisms between the multiplicative group $1 + q\mathrm{Mat}_m(\mathcal{O}_K)$ and the additive group $q\mathrm{Mat}_m(\mathcal{O}_K)$. In particular since $q\mathrm{Mat}_m(\mathcal{O}_K)$ is torsion-free, so is $1+q\mathrm{Mat}_m(\mathcal{O}_K)$. Now plug in $q = p^n$ and $K = \mathbf{Q}_p$; the condition $p^n \geq 3$ implies that the condition $\mathrm{ord}_p(q) > \frac{1}{p-1}$ is satisfied.

◻

5. The Galois twisting form the Main theorem of CM.

Theorem 5.1. *Let (A,i) be a CM abelian variety of type (L,Φ) over $\overline{\mathbf{Q}}$, and assume that it is principal i.e. $i^{-1}(\mathrm{End}(A)) = \mathcal{O}_L$. For any $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/E)$ and $s \in A^\times_{E,f}$ with $r_E(s) = \sigma\,|_{E^{\mathrm{ab}}}$ there is a unique L-linear isomorphism*

$$\theta_{\sigma,s} \colon [N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A \xrightarrow{\sim} A^\sigma$$

such that for all $M \geq 1$ the isomorphism $[\sigma] \colon A[M](\overline{\mathbf{Q}}) \xrightarrow{\sim} A^\sigma[M](\overline{\mathbf{Q}})$ equals the composite

$$\begin{array}{ccc} A[M](\overline{\mathbf{Q}}) & \xrightarrow[\simeq]{N_\Phi(s^{-1})} & ([N_\Phi(s^{-1})]_L/M[N_\Phi(s^{-1})]_L) \otimes_{\mathcal{O}_L/M\mathcal{O}_L} A[M](\mathbf{Q}) \\ & & \downarrow \simeq \\ & & ([N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A)[M](\overline{\mathbf{Q}}) \\ & & \theta_{\sigma,s} \downarrow \simeq \\ & & A^\sigma[M](\overline{\mathbf{Q}}) \end{array}$$

For an L-linear polarization ϕ on A, the $\mathbf{Q}$-polarization $\phi_{[N_\Phi(s)]_L^{-1}}$ of $[N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A$ and the $\mathbf{Q}$-polarization ϕ^σ of A^σ are intertwined by the isomorphism $\theta_{\sigma,s}$. Also for any $r \in E^\times$ we have $\theta_{\sigma,rs} = N_\Phi(r)\theta_{\sigma,s}$ and the formation of $\theta_{\sigma,s,(A,i)}$ is natural in the principal CM abelian variety (A,i)

Proof. Our map is going to be $\theta_{\sigma,s} = \theta_{\sigma,\mathfrak{P},c,s,M}$ from the previous discussion. We have checked the torsion condition with mild conditions on M: M was required to be not divisible by $\mathfrak{p}$ (also recall that secretly the descent field K also depends on M via the class field theory condition). However $\mathfrak{P}$ (and hence $\mathfrak{p}$) was chosen according to the Chebotarev density theorem; there are infinitely such $\mathfrak{P}$ so we can pick one that does not divide M, in which case the map $\theta_{\sigma,\mathfrak{P},c,s,M}$ will satisfy the torsion condition, and we already know this map is canonical. Also we had $M \geq 3$, but this is not a problem since the condition on 2-torsion follows from the condition on 4-torsion.

Uniqueness follows from the explicit requirements on every torsion. The behavior with respect to $s \to rs$ comes from the uniqueness, as on torsion this is obvious. Alternatively, multiplying s by $r \in E$, we go back to the step of the construction where we write $s = c\pi_p u$ hence $rs = (rc)\pi_p u$ and from here we immediately see that $\theta_{\sigma,\mathfrak{P},c,s,M}$ differs from $\theta_{\sigma,\mathfrak{P},rc,rs,M}$ by the factor $N_\Phi(r)$, because the contribution of c to it was precisely multiplication by $N_\Phi(c)$. Naturality also follows from torsion level descriptions which are natural. $\square$

Here is a nice optional result: the maps $\theta_{\sigma,s}$ are compatible with composition in σ and multipication in s:

Proposition 5.2 (optional). *Assume $\sigma, \sigma' \in \mathrm{Gal}(\overline{\mathbf{Q}}/E)$ and $s, s' \in \mathbb{A}_{E,f}^\times$ satisfy $r_E(s) = \sigma\,|_{E^{\mathrm{ab}}}$, $r_E(s') = \sigma'\,|_{E^{\mathrm{ab}}}$. In particular $r_E(s's) = (\sigma'\sigma)\,|_{E^{\mathrm{ab}}}$. Then the map $\theta_{\sigma'\sigma,s's,(A,i)}$ equals the composite*

$$\begin{tikzcd}
{[N_\Phi(s's)_L^{-1}] \otimes_{\mathcal{O}_L} A} \arrow[r, "\simeq"] & {[N_\Phi(s')]_L^{-1} \otimes_{\mathcal{O}_L} ([N_\Phi(s)_L^{-1}] \otimes_{\mathcal{O}_L} A)} \arrow[d, "1\otimes\theta_{\sigma,s,(A,i)}"] \\
 & {[N_\Phi(s)_L^{-1}] \otimes_{\mathcal{O}_L} A^\sigma} \arrow[d, "\theta_{\sigma',s',(A^\sigma,i^\sigma)}"] \\
A^{\sigma'\sigma} & (A^\sigma)^{\sigma'} \arrow[l, "\simeq"']
\end{tikzcd}$$

Proof. It is enough to check that the above composite equals the same thing on torsion. That easily follows from the explicitly writing down the composite and using the fact that (via i and i^σ) the Galois twists commute with the action of L hence the action of L-ideles. We also use the fact that $[\sigma] \circ N_\Phi = N_\Phi$, as well as the fact that the canonical identifications $\mathfrak{a} \otimes_{\mathcal{O}_L} A[M] \simeq (\mathfrak{a} \otimes_{\mathcal{O}_L} A)[M]$ behave well with respect to composition - more exactly if this isomorphism is written as $\omega_\mathfrak{a}$ then we have $(\omega_\mathfrak{a} \otimes 1)\omega_\mathfrak{b} = \omega_{\mathfrak{ab}}$. This is simple verification on the functor of points. $\square$

6. The proof of the main theorem of CM

In the above theorem, we combine all the torsions and tensor over $\mathbf{Z}$ with $\mathbf{Q}$ (or over $\mathcal{O}_L$ with L) to obtain the diagram of isomorphisms

$$\begin{tikzcd}
V_f(A) \arrow[r, "N_\Phi(s)^{-1}"] \arrow[drr, "{[\sigma]}"'] & {[N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} V_f(A)} \arrow[r] & V_f([N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A) \arrow[d, "V_f(\theta_{\sigma,s})"] \\
 & & V_f(A^\sigma)
\end{tikzcd}$$

But recall that $V_f([N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A)$ is canonically isomorphic to $V_f(A)$ - the diagram from 2.1.c.) augments our diagram

$$\begin{tikzcd}
V_f(A) \arrow[r, "N_\Phi(s)^{-1}"] \arrow[dr, "N_\Phi(s)^{-1}"'] & {[N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} V_f(A)} \arrow[r] \arrow[d, "mult"'] & V_f([N_\Phi(s)]_L^{-1} \otimes_{\mathcal{O}_L} A) \arrow[d, "V_f(\theta_{\sigma,s})"] \\
 & V_f(A) \arrow[ur, "V_f(j_{[N_\Phi(s)]_L^{-1}})"] \arrow[r, "{[\sigma] \circ N_\Phi(s)}"'] & V_f(A^\sigma)
\end{tikzcd}$$

The important part is the rightmost triangle: we see that the L-linear $\mathbf{Q}$-isogeny $\lambda_{\sigma,s} = \theta_{\sigma,s} \circ j_{[N_\Phi(s)]_L^{-1}}$ satisfies $V_f(\lambda_{\sigma,s}) \cdot N_\Phi(s)^{-1} = [\sigma]$. Further, $\theta_{\sigma,s}$ pulls ϕ^σ back to $\phi_{[N_\Phi(s)]_L^{-1}}$, and by construction $j_{[N_\Phi(s)]_L^{-1}}$ pulls $\phi_{[N_\Phi(s)]_L^{-1}}$ back to ϕ times $[N_\Phi(s)^{-1} N_\Phi^{-1}(s)^*]_{\mathbf{Q}}^+ = [Nm_{E/\mathbf{Q}}(s)]_{\mathbf{Q}}^+$ i.e. to a positive rational multiple of ϕ. This means that $\lambda_{\sigma,s}^{-1} = \varphi_\sigma$ satisfies the requirements in the construction of the map

$$\mu \colon E^\times \setminus \mathbf{A}_{E,f}^\times \to \mathrm{Gal}(E^{\mathrm{ab}}/E) \to T(\mathbf{Q}) \setminus T(\mathbf{A}_{\mathbf{Q},f}^\times)$$

(see Brandon's lecture notes)

Recall that $\mu(s)$ was defined to be the coset of the element inducing the map of rational total Tate modules $V_f(A) \xrightarrow{[\sigma]} V_f(A^\sigma) \xrightarrow{V_f(\varphi_\sigma)} V_f(A)$. But $V_f(\varphi_\sigma)$ is $N_\Phi(s)^{-1} \circ [\sigma]^{-1}$ and hence the above composition is just multiplication by $N_\Phi(s)^{-1}$. This finishes the proof of the main theorem of CM.