ABELIAN VARIETIES: GEOMETRY, PARAMETER SPACES, AND ARITHMETIC

BRIAN CONRAD

In this article I wish to convey to beginners how analytic, algebro-geometric, and arithmetic techniques come together in the study of abelian varieties and their variation in families. For many modern arithmetic applications it is crucial that the entire theory admits an algebraic foundation (requiring the full force of the theory of schemes and beyond), but much of our geometric intuition and expectation is guided by the complex-analytic theory. It is for this reason that the first three sections §1-§3 are devoted to an overview of the analytic aspects of the theory, with an emphasis on those structures and examples that will be seen to admit natural analogues in the algebraic theory. With this experience behind us, §4-5 explain how to make the appropriate definitions in the algebraic setting and what sorts of results and interesting examples one obtains. The final two sections (§6 and §7) are devoted to the study of analytic and algebraic families of abelian varieties parameterized by some "modular" varieties, with the aim of explaining (via the Main Theorem of Complex Multiplication) how these varieties fit within the framework of Shimura varieties.

It is assumed that the reader has some prior exposure to the basics of algebraic geometry over an algebraically closed field, and notions such as complex manifold and vector bundle over a manifold (as well as bundle operations, such as dual, tensor product, and pullback). We will certainly have to allow ourselves to work with algebro-geometric objects over a field that is not algebraically closed (so the reader unfamiliar with such things will have to take a lot on faith), and the proofs of many of the algebraic theorems that we state without proof require a solid command of the theory of schemes. Hence, these notes should be understood to be merely a survey of important notions, examples, and results for a reader who is taking their first steps into this vast and beautiful subject.

Notation. Throughout, $\mathbf{C}$ denotes an algebraic closure of $\mathbf{R}$ that is fixed for all time. We write $\mathbf{Z}(1)$ to denote the kernel of the exponential map $\exp : \mathbf{C} \twoheadrightarrow \mathbf{C}^\times$, so this is a free $\mathbf{Z}$-module of rank 1 generated by $\pm 2\pi\sqrt{-1}$; of course, there is no canonical choice of $\sqrt{-1}$ in $\mathbf{C}$. For any $\mathbf{Z}$-module M we write $M(1)$ to denote $M \otimes_{\mathbf{Z}} \mathbf{Z}(1)$. For example, $\mathbf{R}(1)$ is the "imaginary axis" in $\mathbf{C}$ (*i.e.*, the -1-eigenspace for complex conjugation on the $\mathbf{R}$-vector space $\mathbf{C}$), $\mathbf{C}(1)$ is canonically isomorphic to $\mathbf{C}$ via the multiplication map $\mathbf{C} \otimes_{\mathbf{Z}} \mathbf{Z}(1) \to \mathbf{C}$, and for any positive integer N the function $e^{(\cdot)/N}$ identifies $(\mathbf{Z}/N\mathbf{Z})(1) = \mathbf{Z}(1)/N \cdot \mathbf{Z}(1)$ with the group $\mu_N(\mathbf{C})$ of Nth roots of unity in $\mathbf{C}$. (Thus, for any prime ℓ we have that $\mathbf{Z}_\ell(1) \simeq \varprojlim \mathbf{Z}(1)/\ell^n \cdot \mathbf{Z}(1)$ is isomorphic to $\varprojlim \mu_{\ell^n}(\mathbf{C})$ via the transition maps $\mu_{\ell^{n+1}}(\mathbf{C}) \twoheadrightarrow \mu_{\ell^n}(\mathbf{C})$ defined by $z \mapsto z^\ell$.)

By working systematically with $\mathbf{Z}(1)$ we can avoid using a choice of $\sqrt{-1}$ throughout the analytic theory (as we must be able to do if we are to give definitions that admit algebraic analogues, since the algebraic theory does not know the concept of orientation for the complex plane that is equivalent to making a choice of $\sqrt{-1}$). In this spirit, for any $z \in \mathbf{C}$ we write $\mathrm{Re}(z)$ to denote the "real part" $(z + \overline{z})/2 = (1/2)\mathrm{Tr}_{\mathbf{C}/\mathbf{R}}(z)$ and z_{im} to denote the "imaginary component" $(z - \overline{z})/2 = z - \mathrm{Re}(z)$; these are canonical, whereas the classical "imaginary part" $\mathrm{Im}(z)$ is not.

If M is a finite and free module over a commutative ring R (especially $\mathbf{Z}$ or a field) then we write $M^\vee$ to denote the dual module $\mathrm{Hom}_R(M, R)$. For example, if M is a finite free $\mathbf{Z}$-module then $M^\vee(1)$ is naturally identified with the module $\mathrm{Hom}(M, \mathbf{Z}(1))$ of $\mathbf{Z}(1)$-valued $\mathbf{Z}$-linear forms on M.

Date: July 8, 2005.

1991 *Mathematics Subject Classification.* Primary 14G22; Secondary 14H52.

This work was partially supported by grants from the NSF and the Alfred P. Sloan Foundation. I would like to thank the organizers of the summer school for the invitation to present an introductory overview of this subject.

Contents

1. Basic analytic definitions and examples

1.1. **Uniformizations.** The beginning of the theory is:

Definition 1.1.1. A *complex torus* is a compact connected complex Lie group. In the 1-dimensional case it is called an *elliptic curve*.

Example 1.1.2. Let V be a finite-dimensional $\mathbf{C}$-vector space with dimension $g \geq 1$ and let Λ be a lattice in V (*i.e.*, a co-compact discrete subgroup of V). Equivalently, Λ is a discrete subgroup of V that is finite and free of rank $2g$ as a $\mathbf{Z}$-module, or in yet another formulation it is a finite free $\mathbf{Z}$-submodule of V such that the induced $\mathbf{R}$-linear map $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda \to V$ is an isomorphism. The quotient V/Λ has a natural structure of compact Hausdorff commutative topological group and it admits a unique structure of complex manifold with respect to which the projection map $V \to V/\Lambda$ is a local analytic isomorphism. As such, V/Λ is a complex torus.

Remark 1.1.3. Beware that although the natural $\mathbf{R}$-linear map $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda \to V$ in Example 1.1.2 is an isomorphism, so a $\mathbf{Z}$-basis of Λ provides an $\mathbf{R}$-basis of V giving an $\mathbf{R}$-linear identification $V \simeq \mathbf{R}^{2g}$ carrying

Λ over to $\mathbf{Z}^{2g}$, there is generally no $\mathbf{C}$-basis of V with respect to which the description of Λ in V is equally straightforward. In other words, the $\mathbf{C}$-structure put on the $\mathbf{R}$-vector space $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ via its identification with V is a subtle structure. However, it can be described in somewhat concrete terms if we choose suitable bases. For example, in the classical case $g = 1$ we may pick a basis $\{\lambda, \lambda'\}$ of Λ and use λ' as a basis of V, so the resulting identification of V with $\mathbf{C}$ carries Λ to the lattice $\mathbf{Z}\tau \oplus \mathbf{Z}$ where $\tau \in \mathbf{C} - \mathbf{R}$ satisfies $\lambda = \tau\lambda'$ in V. Conversely, for any $\tau \in \mathbf{C} - \mathbf{R}$ certainly $\mathbf{Z}\tau \oplus \mathbf{Z}$ is a lattice in $\mathbf{C}$. Hence, in the case $g = 1$ we may describe all lattice embeddings $\Lambda \hookrightarrow V$ in the concrete form $\mathbf{Z}\tau \oplus \mathbf{Z} \subseteq \mathbf{C}$; the choice of τ is not intrinsic to the lattice (and, following Deligne, it determines an orientation of the $\mathbf{R}$-vector space $\mathbf{C}$ via the nonzero element $\tau \wedge 1 \in \wedge^2_{\mathbf{R}}(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda) = \wedge^2_{\mathbf{R}}(\mathbf{C})$; this is opposite to the "classical" orientation).

Example 1.1.4. Let $\Lambda \subseteq \mathbf{C}$ be a lattice. In the classical theory of elliptic functions it is proved that the Weierstrass $\wp$-function

$$\wp_\Lambda(z) = \frac{1}{z^2} + \sum_{\lambda \in \Lambda - \{0\}} \left(\frac{1}{(z-\lambda)^2} - \frac{1}{\lambda^2} \right)$$

is uniformly convergent on compacts in $\mathbf{C} - \Lambda$ and is meromorphic in $\mathbf{C}$ with Λ-periodicity and double poles along Λ. For the constants

$$g_2(\Lambda) = 60 \sum_{\lambda \in \Lambda - \{0\}} \frac{1}{\lambda^4}, \quad g_3(\Lambda) = 140 \sum_{\lambda \in \Lambda - \{0\}} \frac{1}{\lambda^6}$$

one has the non-linear differential equation ${\wp'_\Lambda}^2 = 4\wp_\Lambda^3 - g_2(\Lambda)\wp_\Lambda - g_3(\Lambda)$ for which the cubic has discriminant $\Delta(\Lambda) = g_2(\Lambda)^3 - 27g_3(\Lambda)^2$ that is *non-zero*. Hence, the projective curve E in $\mathbf{CP}^2$ with affine model $y^2 = 4x^3 - g_2(\Lambda)x - g_3(\Lambda)$ is smooth with $[0,1,0]$ as its unique point on the line at infinity and $z \mapsto (\wp_\Lambda(z), \wp'_\Lambda(z))$ uniquely extends to an analytic isomorphism $\mathbf{C}/\Lambda \simeq E$ carrying 0 to $[0,1,0]$. Moreover, the ratio

$$j(\Lambda) = \frac{1728 g_2(\Lambda)^3}{\Delta(\Lambda)}$$

is called the *j-invariant* of $\mathbf{C}/\Lambda$ and it determines the isomorphism class of the elliptic curve $\mathbf{C}/\Lambda$. Conversely, *every* smooth plane cubic $y^3 = 4x^3 - ax^2 - b$ with $a^3 - 27b^2 \neq 0$ arises from some Λ as above.

Rather generally, if E is an elliptic curve then a *Weierstrass model* for E is a *smooth* planar cubic curve $C \subseteq \mathbf{CP}^2$ of the form

$$y^2w + a_1xyw + a_3yw^2 = x^3 + a_2x^2w + a_4xw^2 + a_6w^3$$

equipped with an isomorphism $E \simeq C$ carrying the origin on E to the unique point $[0,1,0]$ where C meets the line at infinity $w = 0$. Many abstract concepts for complex tori can be made very explicit in the 1-dimensional case via Weierstrass models, though using such models in proofs tends to make things messy and there is no equally explicit analogue of them in the higher-dimensional case.

It is a fundamental fact that Example 1.1.2 gives all examples. However, not every complex torus "in nature" is given to us in such a form. (See Example 1.1.12.) Before we explore more examples, let us make precise the sense in which Example 1.1.2 is universal.

Theorem 1.1.5. *Every complex torus X is commutative, and for the tangent space $V = \mathrm{T}_e(X)$ at the identity e the complex-analytic exponential map $\exp_X : V \to X$ is a surjective group homomorphism whose kernel Λ is a lattice in V. In particular, it is a covering map and the induced map $V/\Lambda \to X$ is an isomorphism of complex Lie groups.*

Proof. See [Mum, pp. 1-2]. ■

We shall not need to use the Lie-theoretic exponential map beyond its functoriality and the properties given in this theorem, so the reader may take it as a black box. It coincides with the exponential map on the underlying real Lie group $X_{\mathbf{R}}$ (via the canonical identification of the tangent space to $X_{\mathbf{R}}$ at e with the underlying $\mathbf{R}$-vector space of $\mathrm{T}_e(X)$), and it is characterized by the condition $\exp_X(v) = \phi_v(1)$ for all $v \in V$, where $\phi_v : \mathbf{C} \to X$ is the unique Lie group homomorphism ("1-parameter subgroup") for which $\mathrm{d}(\phi_v)_0(\partial_t|_0) = v$.

The lattice Λ in Theorem 1.1.5 has a concrete interpretation, as follows. Since V is simply connected, the exponential map for X identifies $(V, 0)$ with a pointed universal covering space of (X, e) (endowed with its unique compatible complex structure), and as such the lattice Λ is identified with $\mathrm{H}_1(X, \mathbf{Z}) = \pi_1(X, e)$. (For each $\lambda \in \Lambda$, the $\mathbf{R}$-line segment in V from 0 to λ projects to an oriented loop in X based at e that represents the corresponding homology class $[\lambda] \in \pi_1(X, e)^{\mathrm{ab}} = \mathrm{H}_1(X, \mathbf{Z})$.) Since X is a compact connected complex manifold, its only global holomorphic functions are the constant functions. As with any complex Lie group, every global holomorphic 1-form on X is a holomorphic linear combination of invariant holomorphic 1-forms (with respect to left/right translation), and so by constancy of such holomorphic multiplier functions on X we conclude that all holomorphic 1-forms on X are invariant. Hence, the cotangent space $V^\vee$ for X at e may be identified with the $\mathbf{C}$-vector space of holomorphic 1-forms on X, and so V may be identified with the space of linear functionals on this space of 1-forms.

In particular, each $\lambda \in \Lambda \subseteq V$ gives rise to a functional on the space of holomorphic 1-forms on X, and this functional is precisely the operation $\int_{[\lambda]}$ of integration along the associated homology class. It is rather important for explicit descriptions of various abstract procedures in the theory that if W is a finite-dimensional complex vector space and $L \subseteq W$ is a lattice then for the complex torus $X = W/L$ we have canonically $\mathrm{T}_e(X) \simeq W$ and this isomorphism carries the exponential uniformization $\mathrm{T}_e(X) \twoheadrightarrow X$ to the *canonical* projection $W \twoheadrightarrow W/L$.

Remark 1.1.6. We have just seen that every nonzero holomorphic 1-form on a complex torus has no zeros. By the Riemann-Roch theorem, on any compact connected Riemann surface with genus $g > 1$ every nonzero holomorphic 1-form has $2g - 2$ zeros (with multiplicity). Since the case of genus 0 has no nonzero global holomorphic 1-forms, an elliptic curve must have genus 1. Conversely, in the classical theory of elliptic curves it is proved that any compact connected Riemann surface C with genus 1 and a chosen point $e \in C$ admits a unique structure of elliptic curve (*i.e.*, a unique analytic Lie group structure) with identity e. (These classical facts follow from Theorem 1.2.1 and Example 1.1.11.) Thus, an elliptic curve "is" precisely a genus-1 compact connected Riemann surface endowed with a marked point.

To generalize Remark 1.1.3 to the higher-dimensional case, we proceed as follows. Let $\overline{V} = \mathbf{C} \otimes_{\mathbf{C},\sigma} V$ denote the "conjugate space" (where $\sigma : \mathbf{C} \simeq \mathbf{C}$ is complex conjugation), and let $\overline{v} = 1 \otimes v \in \overline{V}$ for $v \in V$. The natural isomorphism $\mathbf{C} \otimes_{\mathbf{R}} \mathbf{C} \simeq \mathbf{C} \times \mathbf{C}$ of left $\mathbf{C}$-algebras defined by $a \otimes b \mapsto (ab, a\overline{b})$ induces a natural $\mathbf{C}$-linear isomorphism

$$\mathbf{C} \otimes_{\mathbf{Z}} \Lambda = \mathbf{C} \otimes_{\mathbf{R}} (\mathbf{R} \otimes_{\mathbf{Z}} \Lambda) \simeq \mathbf{C} \otimes_{\mathbf{R}} V = (\mathbf{C} \otimes_{\mathbf{R}} \mathbf{C}) \otimes_{\mathbf{C}} V \simeq V \oplus \overline{V} \tag{1.1.1}$$

(given by $a \otimes \lambda \mapsto (a\lambda, \overline{a\lambda})$). In fact, a moment's reflection with this calculation shows that any $\mathbf{Z}$-linear map $\iota : L \to W$ from a finite free $\mathbf{Z}$-module to a finite-dimensional $\mathbf{C}$-vector space likewise induces a $\mathbf{C}$-linear map $\iota_{\mathbf{C}} : \mathbf{C} \otimes_{\mathbf{Z}} L \to W \oplus \overline{W}$, and that ι is an injection onto a lattice if and only if $\iota_{\mathbf{C}}$ is an isomorphism. Hence, if we pick a $\mathbf{C}$-basis $\{e_1, \dots, e_g\}$ of V and a $\mathbf{Z}$-basis $\{\lambda_1, \dots, \lambda_{2g}\}$ of Λ then upon writing $\lambda_j = \sum c_{ij} e_i$ for unique $c_{ij} \in \mathbf{C}$ the condition that Λ be a lattice in V is precisely the condition that the $2g \times 2g$ matrix M with first g rows given by (c_{ij}) and last g rows given by $(\overline{c}_{ij})$ be invertible.

This matrix M is a bit unwieldy, but we can make it easier to work with by improving our choice of basis of Λ. Namely, since a $\mathbf{Z}$-basis of Λ spans the $\mathbf{C}$-vector space V we can certainly rearrange the λ_j's so that the ordered subset $\{\lambda_{g+1}, \dots, \lambda_{2g}\}$ in our ordered $\mathbf{Z}$-basis of Λ is a $\mathbf{C}$-basis of V. Taking $e_j = \lambda_{j+g}$, the matrix M acquires the block form

$$M = \begin{pmatrix} Z & 1_g \\ \overline{Z} & 1_g \end{pmatrix} \tag{1.1.2}$$

for a $g \times g$ matrix Z that describes $\lambda_1, \dots, \lambda_g$ as $\mathbf{C}$-linear combinations of $\lambda_{g+1}, \dots, \lambda_{2g}$ in V (and 1_g denotes the $g \times g$ identity matrix). Our lattice Λ in $V = \mathbf{C}^g$ is $Z(\mathbf{Z}^g) + \mathbf{Z}^g$ (necessarily a direct sum inside of V). By subtracting row $j + g$ from row j in M for $1 \leq j \leq g$, we see that $\det(M) = \det(Z - \overline{Z})$. Thus, the invertibility of M is equivalent to that of the matrix $Z_{\mathrm{im}} = (Z - \overline{Z})/2$ of "imaginary components" of Z. Conversely, given any $Z \in \mathrm{Mat}_{g \times g}(\mathbf{C})$ for which Z_{im} is invertible, we may reverse the procedure to build a lattice embedding of $\mathbf{Z}^{2g}$ onto $Z(\mathbf{Z}^g) \oplus \mathbf{Z}^g \subseteq \mathbf{C}^g$ giving a complex torus as the quotient.

Example 1.1.7. Let us use the preceding calculations to reformulate Example 1.1.2 from another point of view. What does it mean to put a $\mathbf{C}$-structure on the $\mathbf{R}$-vector space $\mathbf{R}\otimes_{\mathbf{Z}}\Lambda$, or equivalently to give an $\mathbf{R}$-algebra map $\mathbf{C}\to \mathrm{End}_{\mathbf{R}}(\mathbf{R}\otimes_{\mathbf{Z}}\Lambda)$? If V denotes the $\mathbf{C}$-vector space arising from such a choice then we get a natural surjective $\mathbf{C}$-linear map $\pi : \mathbf{C}\otimes_{\mathbf{Z}}\Lambda \twoheadrightarrow V$ by scalar multiplication with respect to this structure, and so the kernel F of π is a g-dimensional $\mathbf{C}$-subspace of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$. If we pass to the conjugate $\mathbf{C}$-structure $\overline{V}$, the resulting surjection $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda\twoheadrightarrow\overline{V}$ is the map $\overline{\pi}$ obtained through extension of scalars on π by complex conjugation, so its kernel is the conjugate subspace $\overline{F}\subseteq \mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ that is the image of F under the conjugation-linear involution of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ given by $a\otimes\lambda\mapsto \overline{a}\otimes\lambda$. The isomorphism condition (1.1.1) implies that the natural $\mathbf{C}$-linear map

$$F\oplus\overline{F}\to\mathbf{C}\otimes_{\mathbf{Z}}\Lambda \tag{1.1.3}$$

is an isomorphism, with $\overline{F}$ projecting isomorphically onto the quotient V of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ (resp. F projecting isomorphically onto the quotient $\overline{V}$ of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$) via the composition of (1.1.3) with π (resp. with $\overline{\pi}$).

This procedure can be reversed: if $F\subseteq\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ is a g-dimensional $\mathbf{C}$-linear subspace such that the natural $\mathbf{C}$-linear map (1.1.3) is an isomorphism then Λ maps onto a lattice in the $\mathbf{C}$-vector space $V=F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$ because the $\mathbf{R}$-linear map

$$(\mathbf{R}\otimes_{\mathbf{Z}}\Lambda)\oplus F\to\mathbf{C}\otimes_{\mathbf{Z}}\Lambda \tag{1.1.4}$$

is injective (as $(\mathbf{R}\otimes_{\mathbf{Z}}\Lambda)\cap F=\overline{F}\cap F=0$) and consequently is an isomorphism for $\mathbf{R}$-dimension reasons. We may therefore identify the quotient $X=V/\Lambda$ with the *double-coset* space

$$X=F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)/\Lambda. \tag{1.1.5}$$

To summarize, the condition that (1.1.3) is an isomorphism is *equivalent* to the condition that Λ maps to a lattice in $F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$. The data consisting of Λ equipped with such a $\mathbf{C}$-subspace $F\subseteq\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ is an example of a (pure) *Hodge structure*. Observe that in this description, $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ is precisely $\mathbf{C}\otimes_{\mathbf{Z}}\mathrm{H}_1(X,\mathbf{Z}) = \mathrm{H}_1(X,\mathbf{C})=\mathrm{H}^1(X,\mathbf{C})^\vee$.

Remark 1.1.8. In the special case that $X=E=\mathbf{C}/(\mathbf{Z}\tau\oplus\mathbf{Z})$ is an elliptic curve in classical form, what is the kernel F of the projection from $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda=\mathrm{H}_1(E,\mathbf{C})$ onto the tangent space $V=\mathbf{C}$ of E at the origin? Writing $[\lambda]$ to denote the homology class of $\lambda\in\Lambda=\mathbf{Z}\tau\oplus\mathbf{Z}$, a nonzero element of the kernel line F is $\tau[1]-[\tau]\in\mathrm{H}_1(E,\mathbf{C})$. In less coordinate-dependent terms, if we write $E=V/\Lambda$ and pick a $\mathbf{Z}$-basis $\{\lambda,\lambda'\}$ of Λ then we have $\lambda=\tau\lambda'$ in V for a unique $\tau\in\mathbf{C}-\mathbf{R}$ and $\tau\otimes[\lambda']-[\lambda]\in\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ is a basis for the $\mathbf{C}$-line F.

The advantage of the uniformization (1.1.5) over the more popular "V/Λ" exponential uniformization is that it is better-suited to the study of analytic families of complex tori. More specifically, the two descriptions (1.1.2) and (1.1.5) express rather different points of view when one tries to consider the variation of complex tori in analytic families. When using (1.1.2) one is led to consider varying Z, which comes down to moving the lattice Λ in the fixed vector space $\mathbf{C}^g$ (a well-known procedure in the classical case $g=1$, where it becomes the process of moving τ in $\mathbf{C}-\mathbf{R}$). However, when using (1.1.5) the natural tendency is to instead consider the lattice Λ as fixed and to move the subspace $F\subseteq\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$. This amounts to varying the *complex structure* on the fixed $\mathbf{R}$-vector space $\mathbf{R}\otimes_{\mathbf{Z}}\Lambda\simeq F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$. In other words, an analytic family of complex tori from the double-coset point of view becomes a variation of complex structure on a fixed vector space. These a rather different ways to looking at the same situation, and it is the latter viewpoint that leads to deeper generalization within the framework of Hodge theory. Curiously, the 19th-century and early 20th-century mathematicians who studied higher-dimensional complex tori prior to the advent of modern topology (with covering spaces, *etc.*) tended to use (1.1.5) rather than the exponential uniformization, though the emphasis on the interpretation as a variation of complex structure seems to be due to Deligne.

Example 1.1.9. Let C be a compact connected Riemann surface with genus g. Let $\Lambda=\mathrm{H}_1(X,\mathbf{Z})$ and let $W=\Omega^1(X)$ be the $\mathbf{C}$-vector space of global holomorphic 1-forms on C. By the classical theory of Riemann surfaces, W is a g-dimensional complex vector space and the natural map

$$W\oplus\overline{W}\to\mathrm{H}^1(C,\mathbf{C})\simeq\mathrm{H}_1(C,\mathbf{C})^\vee=\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^\vee \tag{1.1.6}$$

is an isomorphism. That is, every degree-1 deRham cohomology class for C with $\mathbf{C}$-coefficients can be uniquely represented by $\omega + \overline{\eta}$ where ω is a holomorphic 1-form and $\overline{\eta}$ is an anti-holomorphic 1-form. (In a local holomorphic coordinate z we have $\omega = f\mathrm{d}z$ and $\overline{\eta} = \overline{g}\mathrm{d}\overline{z}$ with holomorphic f and g.)

Since the identification of deRham cohomology with dual to $\mathbf{C}$-homology is defined via integration of forms along cycles, we conclude that the $\mathbf{C}$-linear isomorphism

$$(\iota, \overline{\iota}) : \mathbf{C} \otimes_{\mathbf{Z}} \Lambda \simeq W^\vee \oplus \overline{W}^\vee$$

dual to (1.1.6) is induced by the map $\iota : \Lambda \to W^\vee$ that assigns to each homology class σ the integration functional $\int_\sigma$. Thus, by Example 1.1.7, this integration mapping identifies Λ with a *lattice* in $W^\vee$. The associated g-dimensional complex torus $W^\vee/\Lambda$ is called the *Jacobian* of C and is denoted $\mathrm{Jac}(C)$. If we pick a $\mathbf{Z}$-basis $\{\sigma_1, \dots, \sigma_{2g}\}$ of homology cycles and a $\mathbf{C}$-basis $\{\omega_1, \dots, \omega_g\}$ of global holomorphic 1-forms then the $g \times 2g$ matrix describing the embedding of Λ into V is the *period matrix* $(\int_{\sigma_j} \omega_i)$. If we rearrange the σ_j's and choose the ω_i's so that $\int_{\sigma_{j+g}} \omega_i = \delta_{ij}$ (which we may certainly always do), then the period matrix takes the form $\begin{pmatrix} Z & 1_g \end{pmatrix}$ for a matrix Z with invertible imaginary component. This Z satisfies classical conditions called *Riemann's relations* (encoding, for example, a positive-definiteness property of the intersection form on $\mathrm{H}_1(C, \mathbf{Z})$).

Remark 1.1.10. For the reader who knows some basic Hodge theory and sheaf cohomology, we can refine the preceding considerations as follows. It is a pleasant exercise to check that the subspace F in (1.1.5) is precisely the linear dual of the quotient $\mathrm{H}^1(X, \mathscr{O}_X)$ of $\mathrm{H}^1(X, \mathbf{C})$, and so we obtain the cohomological double-coset description

$$X = \mathrm{H}^1(X, \mathscr{O}_X)^\vee \backslash (\mathbf{C} \otimes_{\mathbf{Z}} \mathrm{H}_1(X, \mathbf{Z}))/\mathrm{H}_1(X, \mathbf{Z}). \tag{1.1.7}$$

Likewise, in the double-coset language from Example 1.1.7, the kernel F of the integration map $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda = \mathrm{H}_1(C, \mathbf{C}) \twoheadrightarrow W^\vee$ is the linear dual of the quotient $\mathrm{H}^1(C, \mathscr{O}_C)$ of $\mathrm{H}^1(C, \mathbf{C})$. Thus, in cohomological language we may say

$$\mathrm{Jac}(C) = \mathrm{H}^1(C, \mathscr{O}_C)^\vee \backslash \mathrm{H}_1(C, \mathbf{C})/\mathrm{H}_1(C, \mathbf{Z}).$$

More generally, for *any* projective complex manifold M it follows from Hodge theory that integration along cycles identifies $\mathrm{H}_1(M, \mathbf{Z})$ with a lattice in the linear dual of the finite-dimensional $\mathbf{C}$-vector space $W_M = \Omega^1(M)$ of global holomorphic 1-forms on M; the resulting complex torus $W_M^\vee/\mathrm{H}_1(M, \mathbf{Z})$ is called the *Albanese variety* $\mathrm{Alb}(M)$ of M, and the kernel F of the integration mapping from $\mathbf{C} \otimes_{\mathbf{Z}} \mathrm{H}_1(M, \mathbf{Z})$ onto $W_M^\vee$ is the dual of the quotient $\mathrm{H}^1(M, \mathscr{O}_M)$ of $\mathrm{H}^1(M, \mathbf{C})$. The Albanese variety has very special properties when $\dim M = 1$, as we shall see later.

Example 1.1.11. Let C be a compact connected Riemann surface with genus g, so its Jacobian $\mathrm{Jac}(C)$ is a g-dimensional complex torus. Suppose $g > 0$, so $\mathrm{Jac}(C) \neq 0$, and pick $x_0 \in C$. For any $x \in C$ if we pick a path $\sigma : [0, 1] \to C$ with $\sigma(0) = x_0$ and $\sigma(1) = x$ then we get a well-defined integration functional $\int_\sigma$ on the vector space W of global holomorphic 1-forms on C. If σ' is another such path then $\int_\sigma - \int_{\sigma'}$ is integration along the loop ${\sigma'}^{-1} \cdot \sigma$, so since all holomorphic 1-forms are closed it follows from Stokes' theorem that this loop integration only depends on the homology class of the loop. Hence, $\int_\sigma$ and $\int_{\sigma'}$ coincide when considered as functionals on W modulo functionals arising from the lattice $\mathrm{H}_1(C, \mathbf{Z})$ embedded into W via integration. We may write $\int_{x_0}^x \in W^\vee/\mathrm{H}_1(C, \mathbf{Z}) = \mathrm{Jac}(C)$ to denote this well-defined element.

A local calculation near x shows that the map $i_{x_0} : C \to \mathrm{Jac}(C)$ defined by $x \mapsto \int_{x_0}^x$ is *analytic*, and it certainly carries x_0 to 0. Since changing x_0 to another point x_0' changes i_{x_0} to the additive translate $i_{x_0'} = i_{x_0'}(x_0) + i_{x_0}$, a further calculation using the Riemann–Roch theorem and the positivity of g shows that i_{x_0} is a closed embedding of complex manifolds. (In particular, when $g = 1$ we get an isomorphism $C \simeq \mathrm{Jac}(C)$ that thereby puts a Lie group structure on C with x_0 as the identity.) The isomorphism

$$\mathrm{H}_1(i_{x_0}) : \mathrm{H}_1(C, \mathbf{Z}) \simeq \mathrm{H}_1(\mathrm{Jac}(C), \mathbf{Z})$$

is *the same* as the one given by the analytic definition of $\mathrm{Jac}(C)$ as a coset space. (This agreement depends crucially on the directions of integration used in the definition of i_{x_0} and in the definition of the embedding of $\mathrm{H}_1(C, \mathbf{Z})$ into $W^\vee$.)

Example 1.1.12. In the discussion preceding Example 1.1.7, we saw that any uniformization V/Λ can be put in an especially simple form by identifying V with $\mathbf{C}^g$ using g vectors from a $\mathbf{Z}$-basis of Λ. In terms of this identification we identify $\Lambda \subseteq V$ with $Z(\mathbf{Z}^g) \oplus \mathbf{Z}^g$ for a $g \times g$ matrix $Z = (c_{ij})$ over $\mathbf{C}$ such that $Z_{\rm im}$ is invertible. If we fix a choice of basis $2\pi\sqrt{-1}$ for $\mathbf{Z}(1)$ then the exponential map $e^{2\pi\sqrt{-1}(\cdot)}$ identifies $\mathbf{C}/\mathbf{Z}$ with $\mathbf{C}^\times$ and so identifies $\mathbf{C}^g/\mathbf{Z}^g$ with $T = (\mathbf{C}^\times)^g$. Letting $q_j = (e^{2\pi\sqrt{-1}c_{ij}})_{1\le i\le g} \in T$, the q_j's must multiplicatively generate a discrete co-compact subgroup L in T such that $T/L \simeq V/\Lambda$. Hence, we see that any complex torus may be presented as a quotient of $(\mathbf{C}^\times)^g$ by a discrete co-compact subgroup that is finite free of rank g as a $\mathbf{Z}$-module.

The non-compact groups $(\mathbf{C}^\times)^g$ are the analytic versions of the "algebraic tori" from the theory of algebraic groups, and they are not to be confused with (compact!) complex tori in the sense we are considering. The presentation of a complex torus as an analytic quotient of an algebraic torus over $\mathbf{C}$ (modulo a discrete torsion-free subgroup) is fundamental in the theory of degenerations of complex tori. The importance of such a uniformization was first recognized by Tate in the 1-dimensional case, where it gives the so-called Tate uniformization $\mathbf{C}^\times/q^{\mathbf{Z}}$ for an elliptic curve (with $0 < |q| < 1$).

1.2. **Endomorphisms and isogenies.** Let X and X' be complex tori with respective dimensions g and g' and respective identity elements e and e'. Consider the commutative group $\mathrm{Hom}(X, X')$ of Lie group maps $X \to X'$. This Hom-group turns out to have strong finiteness properties.

Theorem 1.2.1. *Any analytic map $f : X \to X'$ that satisfies $f(e) = e'$ is necessarily a map of Lie groups, and the homology representation $f \mapsto \mathrm{H}_1(f)$ defined by*

$$\mathrm{H}_1 : \mathrm{Hom}(X, X') \to \mathrm{Hom}(\mathrm{H}_1(X, \mathbf{Z}), \mathrm{H}_1(X', \mathbf{Z})) \simeq \mathrm{Mat}_{2g'\times 2g}(\mathbf{Z})$$

is faithful (i.e., injective). In particular, the group law on X is uniquely determined by e and $\mathrm{Hom}(X, X')$ is a finite free $\mathbf{Z}$-module with rank $\le 4gg'$.

Proof. Consider the analytic exponential uniformizations V/Λ and V'/Λ' for X and X' respectively. We have seen that via the exponential uniformization $(V, 0)$ is (uniquely) identified with a pointed universal cover of (X, e), and likewise for $(V', 0)$ over (X', e'), so the analytic map $f : X \to X'$ carrying e to e' uniquely lifts to an analytic map $\widetilde{f} : V \to V'$ carrying 0 to 0. For any $\lambda \in \Lambda$, the map $v \mapsto \widetilde{f}(v + \lambda) - \widetilde{f}(v)$ projects to e' in X' and so its connected image in V' is contained in the lattice Λ'; this must therefore be the constant map to the image $\widetilde{f}(\lambda)$ of $v = 0$. Hence, $\widetilde{f}$ restricts to an additive map between Λ and Λ'. It also follows that for any $n \in \mathbf{Z}$ the map $\widetilde{f}_n : v \mapsto \widetilde{f}(nv) - n\widetilde{f}(v)$ is Λ-invariant and thus factors through an analytic map $V/\Lambda \to V'$ carrying e to 0. This latter analytic map from a connected *compact* complex manifold to a Euclidean space must be constant, and so it is the constant map to the origin. Thus, $\widetilde{f}_n = 0$.

This shows that $\widetilde{f}$ restricts to a $\mathbf{Q}$-linear map between the dense $\mathbf{Q}$-subspaces $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda$ and $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda'$, so by continuity $\widetilde{f}$ must be the $\mathbf{R}$-linear extension of its restriction to the lattices Λ and Λ'. Since this $\mathbf{R}$-linear $\widetilde{f}$ is complex-analytic, it must be $\mathbf{C}$-linear. As a $\mathbf{C}$-linear map from V to V' carrying Λ into Λ', the induced map $f : X \to X'$ is visibly a group homomorphism. By naturality, the restriction of $\widetilde{f}$ to the lattices is identified with $\mathrm{H}_1(f)$. Since this lattice map uniquely determines $\widetilde{f}$ (as its $\mathbf{R}$-linear extension) and hence uniquely determines f, we get the desired faithfulness of the homology representation. ■

From now on, a map of complex tori is understood to respect the identities and hence to be a group map.

Definition 1.2.2. A map of complex tori $f : X \to X'$ is an *isogeny* if it is surjective and has finite kernel. The common size of the fibers of f is called the *degree* of the isogeny.

Example 1.2.3. Let X be a complex torus and let $G \subseteq X$ be a finite subgroup. The quotient X/G in the sense of Lie groups is a complex torus with the same dimension and $X \to X/G$ is an isogeny (with degree equal to the size of G). In terms of the exponential uniformization V/Λ, the subgroup G corresponds to a lattice Λ' containing Λ inside of the rational vector space $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda = \mathrm{H}_1(X, \mathbf{Q})$ and the map $X \to X/G$ is identified with the natural quotient map $V/\Lambda \to V/\Lambda'$ whose kernel G is Λ'/Λ. For any map $f : X \to X'$ of complex tori such that $G \subseteq \ker(f)$, it is clear via consideration of the exponential uniformization (or general properties of quotients in the theory of Lie groups) that f uniquely factors through the isogeny $X \to X/G$.

Example 1.2.4. Let $n \in \mathbf{Z}$ be a nonzero integer. The mapping $[n]_X : X \to X$ is an isogeny with degree $n^{2\dim X}$. Indeed, in terms of the exponential uniformization V/Λ we see that $[n]_X$ lifts to multiplication by n on V, so it is surjective with kernel $(1/n)\Lambda/\Lambda$ that is isomorphic to $(\mathbf{Z}/n\mathbf{Z})^{2\dim(X)}$ since Λ is a finite free $\mathbf{Z}$-module with rank $2\dim(X)$.

If $f : X \to X'$ is a map between complex tori with the same dimension then f is surjective if and only if it has finite kernel. Indeed, first assume that f is surjective. The $\mathbf{C}$-linear map $\widetilde{f} : V \to V'$ has image that, together with Λ', additively spans V'. That is, the countable Λ' surjects onto the $\mathbf{C}$-vector space $\operatorname{coker}(\widetilde{f})$, so this cokernel must be 0. Hence, for $\mathbf{C}$-dimension reasons $\widetilde{f}$ is an isomorphism. The map $\widetilde{f}$ is the $\mathbf{R}$-linear extension of the map induced by f on rational homology, so this $\mathbf{Q}$-linear homology map must be an isomorphism. In other words, the induced map on integral homology is a finite-index lattice inclusion. The finite cokernel for $\mathrm{H}_1(f)$ is thereby identified with $\ker(f)$, so $\ker(f)$ is indeed finite. Conversely, if $\ker(f)$ is finite then for $\lambda \in \ker(\mathrm{H}_1(f))$ the subspace $\mathbf{C}\lambda \subseteq V$ maps to 0 in V' and thus maps to e' in X', so its image in V/Λ is contained in the finite $\ker f$. That is, $\mathbf{C}\lambda + \Lambda$ contains Λ with finite index, which is only possible if $\lambda = 0$. This says that $\mathrm{H}_1(f)$ is injective, so for $\mathbf{Z}$-rank reasons it is a finite-index lattice inclusion. The associated $\mathbf{R}$-linear map $\widetilde{f} : V \to V'$ is therefore an injective and thus an isomorphism (for dimension reasons), so f is indeed surjective.

Example 1.2.5. Let $f : X \to X'$ be a map between complex tori with the same dimension. This map is an isogeny if and only if there exists a map of complex tori $f' : X' \to X$ such that $f' \circ f = [n]_X$ for some nonzero $n \in \mathbf{Z}$, in which case f' is unique and is an isogeny. To prove this, first note that by working on homology lattices we see that such an f' is indeed unique, and moreover that the condition $f' \circ f = [n]_X$ is *equivalent* to the condition $f \circ f' = [n]_{X'}$. Next, observe that if f is an isogeny with finite kernel G then the map $X/G \to X'$ is an isomorphism of complex tori and so if G has order n then $[n]_X : X \to X$ factors through the projection f from X to $X/G \simeq X'$. The induced factorization $f' \circ f$ of $[n]_X$ thereby provides the desired f': the map f' is surjective (as $[n]_X$ is) and so it must be an isogeny. Conversely, if there exists f' such that $f' \circ f = [n]_X$ then clearly $\ker(f) \subseteq \ker([n]_X)$, so $\ker(f)$ is finite.

We say that two complex tori X and X' with the same dimension are *isogenous* if there is an isogeny between them (in either direction, in which case there is one in each direction, by the preceding example). The property of being isogenous is an equivalence relation on the set of isomorphism classes of complex tori, and the *isogeny class* of a complex torus X is the set of (isomorphism classes of) all complex tori that are isogenous to X. In view of Example 1.2.5, we deduce:

Theorem 1.2.6. *A map $f : X \to X'$ of complex tori is an isogeny if and only if the map*

$$\mathrm{H}_1(f) : \mathrm{H}_1(X, \mathbf{Z}) \to \mathrm{H}_1(X', \mathbf{Z})$$

is a finite-index inclusion, or equivalently if and only if the induced map $\mathrm{H}_1(f, \mathbf{Q})$ on $\mathbf{Q}$-homology is an isomorphism.

Observe that if L and L' are finitely generated $\mathbf{Z}$-modules (perhaps with nonzero torsion) and their associated $\mathbf{Q}$-vector spaces are $W = \mathbf{Q} \otimes_{\mathbf{Z}} L$ and $W' = \mathbf{Q} \otimes_{\mathbf{Z}} L'$ then the natural map

$$\mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{Hom}_{\mathbf{Z}}(L, L') \to \mathrm{Hom}_{\mathbf{Q}}(W, W') \tag{1.2.1}$$

is an isomorphism. (In the torsion-free case, this just says that any matrix over $\mathbf{Q}$ can be scaled by a sufficiently divisible nonzero integer so that it has entries in $\mathbf{Z}$.) Thus, whereas the category of complex tori may be viewed as the category of pairs (Λ, V) consisting of a finite-dimensional $\mathbf{C}$-vector space equipped with a lattice (using an evident notion of morphism for such pairs), we may replace lattices by rational vector spaces by using:

Definition 1.2.7. The *isogeny category* of complex tori is the category whose objects are complex tori and whose morphism-groups are $\mathrm{Hom}^0(X, X') = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{Hom}(X, X')$.

Remark 1.2.8. If we say that a map between finite free $\mathbf{Z}$-modules is an *isogeny* when it is a finite-index inclusion, then by (1.2.1) a concrete analogue of Definition 1.2.7 is to say that the category of finite-dimensional $\mathbf{Q}$-vector spaces "is" the isogeny category for the category of finite free $\mathbf{Z}$-modules. This analogy is an excellent source of intuition when working with the isogeny of category of complex tori. For example, just as tensoring with $\mathbf{Q}$ kills torsion in $\mathbf{Z}$-modules, finite-level torsion phenomena associated to complex tori become invisible when we work in the isogeny category of complex tori. A more fundamental example is the fact that in the category of finitely generated $\mathbf{Z}$-modules there are many non-split short exact sequences whereas tensoring with $\mathbf{Q}$ makes all sequences become split. Likewise, an important theorem of Poincaré to be discussed later assures us that (under a polarization condition) passing to the isogeny category of complex tori tends to "split" exact sequences of complex tori.

By Theorem 1.2.1, $\mathrm{Hom}^0(X, X')$ is a finite-dimensional $\mathbf{Q}$-vector space containing $\mathrm{Hom}(X, X')$ as a lattice of full rank. In particular, the functor associating to each complex torus X the same object considered in the isogeny category is a *faithful* functor. (That is, this functor induces an injection on Hom-sets.) By passing to the isogeny category we gain precisely the ability to "invert" isogenies. Indeed, by definition any map in the isogeny category may be multiplied by a sufficiently divisible nonzero integer such that it arises from a genuine map of complex tori, so a map of complex tori $f : X \to X'$ becomes an isomorphism in the isogeny category if and only if there exists a map of complex tori $f' : X' \to X$ such that $f' \circ f = [n]_X$ for some nonzero $n \in \mathbf{Z}$; by Example 1.2.5, this is precisely the condition that f is an isogeny (and then $(1/n) \otimes f'$ is its "inverse" in the isogeny category). Many interesting properties of complex tori are invariant under isogeny, and so working with the isogeny category is often technically convenient much as working with finite-dimensional $\mathbf{Q}$-vector spaces is easier than working with finitely generated $\mathbf{Z}$-modules.

1.3. CM tori.

Let X be a complex torus with dimension $g > 0$. The ring $\mathrm{End}(X) = \mathrm{Hom}(X, X)$ is an associative finite torsion-free $\mathbf{Z}$-algebra that admits an embedding into a $2g \times 2g$ matrix algebra (via the homology representation). It follows that if R is a *commutative domain* in $\mathrm{End}(X)$ then R is a finite $\mathbf{Z}$-module and its fraction field $K = \mathbf{Q} \otimes_{\mathbf{Z}} R$ admits an embedding into $\mathrm{End}^0(X) \subseteq \mathrm{End}_{\mathbf{Q}}(\mathrm{H}_1(X, \mathbf{Q}))$, so the $2g$-dimensional $\mathbf{Q}$-vector space $\mathrm{H}_1(X, \mathbf{Q})$ admits a structure of K-vector space. This forces $[K : \mathbf{Q}] \leq 2g$. Observe that $\mathrm{End}(X) \cap K$ is a $\mathbf{Z}$-finite subring of K that contains the lattice R whose $\mathbf{Z}$-rank is $[K : \mathbf{Q}]$, so $\mathrm{End}(X) \cap K$ is an order in the ring of integers $\mathscr{O}_K$. It is a very subtle problem to work with this order (which is usually not $\mathscr{O}_K$, and so not Dedekind), and in general it is a lot easier to work with K. In the isogeny category we may speak of K "acting" on X, and this is a very convenient point of view in practice.

The extreme case $[K : \mathbf{Q}] = 2g$ will be of much interest. Let us see that in this case there are non-trivial restrictions on K. Since $\mathrm{H}_1(X, \mathbf{Q})$ has a structure of K-vector space, for $\mathbf{Q}$-dimension reasons we see that it must be 1-dimensional over K. We have an $\mathbf{R}$-algebra embedding

$$\mathbf{R} \otimes_{\mathbf{Q}} K \hookrightarrow \mathbf{R} \otimes_{\mathbf{Q}} \mathrm{End}_{\mathbf{Q}}(\mathrm{H}_1(X, \mathbf{Q})) = \mathrm{End}_{\mathbf{R}}(\mathbf{R} \otimes_{\mathbf{Q}} \mathrm{H}_1(X, \mathbf{Q})) = \mathrm{End}_{\mathbf{R}}(\mathrm{T}_e(X))$$

that lands inside of $\mathrm{End}_{\mathbf{C}}(\mathrm{T}_e(X))$ since the action of the order $\mathrm{End}(X) \cap K$ on $\mathrm{T}_e(X)$ is $\mathbf{C}$-linear. But $\mathrm{H}_1(X, \mathbf{Q})$ is a free K-module of rank 1 so $\mathrm{T}_e(X) = \mathbf{R} \otimes_{\mathbf{Q}} \mathrm{H}_1(X, \mathbf{Q})$ is a free $\mathbf{R} \otimes_{\mathbf{Q}} K$-module of rank 1. If we consider the $\mathbf{R}$-algebra decomposition $\mathbf{R} \otimes_{\mathbf{Q}} K \simeq \prod_{v|\infty} K_v$ into the product of completions of K at its infinite places (where a complex place is a pair of conjugate non-real embeddings into $\mathbf{C}$), then $V = \mathrm{T}_e(X)$ admits a corresponding $\mathbf{R} \otimes_{\mathbf{Q}} K$-module decomposition $\prod_v V_v$ with each V_v of dimension 1 over K_v. Since V_v is the image of V under the action of the primitive idempotent in $\mathbf{R} \otimes_{\mathbf{Q}} K$ associated to the factor ring K_v, the commuting nature of the $\mathbf{C}$-module and $\mathbf{R} \otimes_{\mathbf{Q}} K$-module structures on V implies that each V_v is a $\mathbf{C}$-subspace of V. But the $\mathbf{R}$-vector space V_v is 1-dimensional over the field extension $K_v/\mathbf{R}$, so the existence of a $\mathbf{C}$-vector space structure on V_v rules out the possibility $K_v = \mathbf{R}$. This proves the first part of:

Theorem 1.3.1. *Let X be a g-dimensional complex torus, and $K \subseteq \mathrm{End}^0(X)$ a subfield of degree $2g$ over $\mathbf{Q}$. This number field must be totally complex. Moreover, it is its own centralizer in $\mathrm{End}^0(X)$ and the $\mathbf{Q}$-linear tangential action of $K \subseteq \mathrm{End}^0(X) = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{End}(X)$ on the $\mathbf{C}$-vector space $\mathrm{T}_e(X)$ has exactly g lines that are K-stable. The K-action on these lines is given by g pairwise distinct embeddings $\varphi : K \hookrightarrow \mathbf{C}$ that are pairwise non-conjugate.*

Proof. To see that K is its own centralizer in $\mathrm{End}^0(X)$, we embed $\mathrm{End}^0(X)$ into $\mathrm{End}_{\mathbf{Q}}(\mathrm{H}_1(X,\mathbf{Q}))$ to reduce to proving that for an n-dimensional vector space W over a field k (such as $\mathrm{H}_1(X,\mathbf{Q})$ with dimension $2g$ over $\mathbf{Q}$), any commutative subfield $F \subseteq \mathrm{End}_k(W)$ of degree n over k is its own centralizer in $\mathrm{End}_k(W)$. This follows from the fact that W must be 1-dimensional as a F-vector space, and so any F-linear endomorphism of W has to be a scalar multiplication by some $c \in F$.

As for the eigenline decomposition, we have seen that $V = \mathrm{T}_e(X)$ has a unique $\mathbf{R} \otimes_{\mathbf{Q}} K$-module decompsotion $\prod_v V_v$ with each V_v of dimension 1 over K_v (and so of dimension 2 over $\mathbf{R}$), and the $\mathbf{C}$-action on V respects this decomposition. Hence, each $\mathbf{R}$-vector space V_v has commuting $\mathbf{R}$-linear actions by K_v and $\mathbf{C}$, so this gives an $\mathbf{R}$-embedding of K_v into $\mathrm{End}_{\mathbf{C}}(V_v)$. But V_v must be 1-dimensional over $\mathbf{C}$ (for $\mathbf{R}$-dimension reasons), so $\mathrm{End}_{\mathbf{C}}(V_v) = \mathbf{C}$. Hence, we get a canonical $\mathbf{R}$-algebra isomorphism $K_v \simeq \mathbf{C}$ determined by the tangential representation for K on X, and these isomorphisms $K_v \simeq \mathbf{C}$ are precisely the data of g embeddings $K \hookrightarrow \mathbf{C}$ that define the full set of g complex places of the totally complex field K; that is, these are pairwise non-conjugate. Since any line in the $\mathbf{C}$-vector space V that is K-stable is also an $\mathbf{R} \otimes_{\mathbf{Q}} K$-submodule, it must be a direct sum of some of the V_v's and so for $\mathbf{R}$-dimension reasons this line is a single V_v. ■

Let K be a totally complex field of even degree $2g$ over $\mathbf{Q}$ and let $\Phi \subseteq \mathrm{Hom}(K,\mathbf{C})$ be a set of g pairwise non-conjugate embeddings. That is, for each archimedean place v of K we pick one of the two isomorphisms $K_v \simeq \mathbf{C}$ over $\mathbf{R}$. For any such Φ, the $\mathbf{R}$-algebra $\mathbf{R} \otimes_{\mathbf{Q}} K = \prod_v K_v$ acquires a structure of $\mathbf{C}$-vector space via Φ by identifying each K_v with $\mathbf{C}$ through the unique $\varphi \in \Phi$ inducing the place v. We write $(\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi$ to denote this K-vector space equipped with its K-linear action by $\mathbf{C}$ through Φ.

If X is any complex torus of dimension g equipped with an embedding $i : K \hookrightarrow \mathrm{End}^0(X)$, we have seen in the proof of Theorem 1.3.1 that $\mathrm{T}_e(X)$ equipped with its commuting actions of K and $\mathbf{C}$ is non-canonically identified with $(\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi$ for a unique Φ (that is determined by the K-action on $\mathrm{T}_e(X)$). The homology $\mathrm{H}_1(X,\mathbf{Q})$ is thereby identified with a 1-dimensional K-subspace in here, and so for a basis vector λ_0 of this K-vector space we see that $\mathbf{R} \cdot \lambda_0$ fills up $\mathbf{R} \otimes_{\mathbf{Q}} K$. Under the decomposition $\prod_v K_v$ it follows that λ_0 has *nonzero* components along each factor K_v, which is to say $\lambda_0 \in (\mathbf{R} \otimes_{\mathbf{Q}} K)^\times$. Scaling the choice of identifiction $\mathrm{T}_e(X) \simeq (\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi$ by the $\mathbf{C}$-linear multiplication by the inverse of this unit brings us to the case $\lambda_0 = 1$. That is, $\mathrm{H}_1(X,\mathbf{Q})$ is identified with $K \subseteq (\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi$ and so $\mathrm{H}_1(X,\mathbf{Z})$ is identified with a fractional ideal $\mathfrak{a} \subseteq K$ for the order $\mathscr{O} = \mathrm{End}(X) \cap K$ in $\mathscr{O}_K$. It is automatic that $\mathscr{O} = \mathrm{End}_{\mathbf{Z}}(\mathfrak{a}) \cap K$ inside of the matrix algebra $\mathrm{End}_{\mathbf{Q}}(K)$.

This procedure can be reversed, as follows. Let $\mathfrak{a} \subseteq K$ be any $\mathbf{Z}$-lattice whose endomorphism ring

$$\mathrm{End}_{\mathbf{Z}}(\mathfrak{a}) \subseteq \mathrm{End}_{\mathbf{Q}}(\mathbf{Q} \otimes_{\mathbf{Z}} \mathfrak{a}) = \mathrm{End}_{\mathbf{Q}}(K)$$

meets K in an order $\mathscr{O}$ of $\mathscr{O}_K$. (For example, $\mathfrak{a}$ can be a fractional $\mathscr{O}_K$-ideal and $\mathscr{O} = \mathscr{O}_K$.) Let $\Phi \subseteq \mathrm{Hom}(K,\mathbf{C})$ be any set of g pairwise non-conjugate embeddings. The quotient

$$X = (\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi/\mathfrak{a} \tag{1.3.1}$$

is a g-dimensional complex torus equipped with an embedding $K \hookrightarrow \mathrm{End}^0(X)$ such that $K \cap \mathrm{End}(X) = \mathscr{O}$.

Remark 1.3.2. The double-coset language provides a way to state (1.3.1) entirely "holomorphically" (that is, without the intervention of $\mathbf{R}$): it is $F\backslash(\mathbf{C} \otimes_{\mathbf{Q}} K)/\mathfrak{a}$ where $F \subseteq \mathbf{C} \otimes_{\mathbf{Q}} K$ is the $\mathbf{C}$-subspace spanned by eigencharacters for K that lie in the *complement* of Φ.

Definition 1.3.3. A *CM field* is a totally complex number field K that is a quadratic extension of a totally real field $K_0 \subseteq K$. A *CM type* for K is a subset $\Phi \subseteq \mathrm{Hom}(K,\mathbf{C})$ of representatives modulo complex conjugation. (CM is short for "complex multiplication".)

Imaginary quadratic fields are CM fields, as are non-real abelian extensions of $\mathbf{Q}$ (such as cyclotomic fields $\mathbf{Q}(\zeta_n)$ for $n > 2$), though many totally complex fields (*e.g.*, those that are Galois over $\mathbf{Q}$ with no central involution in the Galois group) are not CM. Note that the subfield K_0 in Definition 1.3.3 is unique: it is the maximal totally real subfield of K.

Since a CM field K has degree 2 over its maximal totally real subfield $K_0 \subseteq K$ and the non-trivial element σ of $\mathrm{Gal}(K/K_0)$ intertwines with complex conjugation with respect to any embedding $K \hookrightarrow \mathbf{C}$, it follows from the fact that σ is intrinsic to K that the notion of CM type for K is also intrinsic to K in the sense that it can be defined without reference to the topological structure on $\mathbf{C}$ (*i.e.*, complex conjugation on $\mathbf{C}$): for such K, a CM type for K is a set of representatives for the σ-orbits in $\mathrm{Hom}(K, \mathbf{C})$. In this sense, we can replace $\mathbf{C}$ with *any* field C of characteristic 0 that splits K (*e.g.*, C may be an algebraic closure of $\mathbf{Q}_p$) and we may define the concept of a CM type in the set $\mathrm{Hom}(K, C)$ of size $[K : \mathbf{Q}]$ by working with σ-orbits; nothing of the sort can be done if K is a non-CM totally complex field.

Although the preceding construction (1.3.1) of tori X with a "large" commutative subfield in the endomorphism algebra $\mathrm{End}^0(X)$ makes sense for arbitrary totally complex fields K without any further restriction on K, we shall see later that it is the case when K is a CM field that is especially interesting.

Definition 1.3.4. A *CM torus* is a pair (X, i) where X is a complex torus and $i : K \hookrightarrow \mathrm{End}^0(X)$ is an embedding from a CM field with degree $2\dim(X)$ over $\mathbf{Q}$. The subset $\Phi \subseteq \mathrm{Hom}(K, \mathbf{C})$ encoding the $\mathbf{C}$-linear K-action on $\mathrm{T}_e(X)$ is the *CM-type* of (X, i), and the field K is the *underlying field* of the CM torus (or of the CM type). The order $i^{-1}(\mathrm{End}(X)) \subseteq \mathscr{O}_K$ is the *CM order* of the CM torus.

Remark 1.3.5. Beware that it is crucial to keep track of both the CM field K and the data of i that encodes how it "acts" on X (in the isogeny category). In some cases it turns out that $\mathrm{End}^0(X)$ is a commutative field, and then i is forced to be an isomorphism so one can say that $\mathrm{End}^0(X)$ with its tautological action (i is the identity) "is" the CM type. However, often one is interested in studying CM tori with a *fixed* abstract underlying CM field K, and so the data of the embedding $i : K \hookrightarrow \mathrm{End}^0(X)$ is part of the structure even when such an embedding is forced to be an isomorphism. For example, if σ is a non-trivial automorphism of K then $(X, i \circ \sigma)$ is a *different* CM torus, and it usually does not admit an isogeny to (X, i) respecting the K-actions (in the isogeny category).

Remark 1.3.6. It cannot be stressed too much that when working with a CM torus we should always consider the underlying CM field K to be an *abstract* field; it is not endowed with a preferred embedding into $\mathbf{C}$, for example. In is some older works the avoidance of this abstract point of view occasionally makes some proofs unnecessarily complicated. Such *a priori* preferred embeddings should be avoided at all times, and it clarifies the situation to be vigilant about this.

Example 1.3.7. Let (X, i) and (X', i') be two CM tori with the same underlying CM field K (and hence the same dimension). Suppose that X and X' are K-linearly isogenous; that is, there is an isogeny $f : X \to X'$ satisfying $i'(c) \circ f = f \circ i(c)$ in $\mathrm{Hom}^0(X, X')$ for all $c \in K$. In this case $\mathrm{T}_e(f) : \mathrm{T}_e(X) \to \mathrm{T}_{e'}(X')$ is a $\mathbf{C}$-linear isomorphism (as is the tangent map for any isogeny) and it is also $\mathbf{R} \otimes_{\mathbf{Q}} K$-linear. The CM type is determined by the underlying $\mathbf{R} \otimes_{\mathbf{Q}} K$-module and $\mathbf{C}$-module structures on the tangent space, so it follows that (X, i) and (X', i') have the *same* CM type Φ. If σ is the non-trivial involution of K over its maximal totally real subfield K_0 then $(X, i \circ \sigma)$ has CM type $\overline{\Phi} = \{\varphi \circ \sigma = \overline{\varphi} \,|\, \varphi \in \Phi\}$ that is conjugate to the one for (X, i), and so (X, i) and $(X, i \circ \sigma)$ are never K-linearly isogenous since $\overline{\Phi} \neq \Phi$ (in fact, $\overline{\Phi}$ is the complement of Φ in $\mathrm{Hom}(K, \mathbf{C})$).

The preceding example can be strengthened:

Theorem 1.3.8. *A pair of CM complex tori (X, i) and (X', i') with the same CM field K are K-linearly isogenous if and only if they have the same CM type. That is, K-linear isogeny classes of CM tori with CM field K are in bijective correspondence with CM types for K.*

Both this theorem and the preceding example (as well as Definition 1.3.4) work with any totally complex field, not merely CM fields. It is only in the further study of the theory that we shall see the significance of the case of CM fields.

Proof. The "only if" direction was proved above, so for the converse suppose the CM types coincide with some $\Phi \subseteq \mathrm{Hom}(K, \mathbf{C})$. We get isomorphisms of complex tori

$$X \simeq (\mathbf{R} \otimes_{\mathbf{Q}} K)_{\Phi}/\mathfrak{a}, \;\; X' \simeq (\mathbf{R} \otimes_{\mathbf{Q}} K)_{\Phi}/\mathfrak{a}'$$

that are K-linear in the isogeny category, with $\mathfrak{a}, \mathfrak{a}' \subseteq K$ two lattices that are stable by (possibly distinct) orders in $\mathscr{O}_K$. In particular, these lattices are commensurable with each other, so $N\mathfrak{a} \subseteq \mathfrak{a}'$ for a suitably divisible nonzero integer N. We may replace $\mathfrak{a}$ with $N\mathfrak{a}$, so $\mathfrak{a} \subseteq \mathfrak{a}'$. Hence, we get a unique isogeny of tori $X \to X'$ that respects the tangent space identifications with $(\mathbf{R}\otimes_{\mathbf{Q}} K)_\Phi$, and so by working with the tangent spaces we see that this is an K-linear isogeny. ■

Example 1.3.9. If (X, i) is a CM torus with underlying CM field K and CM type Φ, then it must be K-linearly isogenous to the CM torus $(\mathbf{R}\otimes_{\mathbf{Q}} K)_\Phi/\mathscr{O}_K$ with CM type Φ. Thus, every CM torus is isogenous (in a manner respecting the action of the underlying CM field) to one whose CM order $\mathrm{End}(X)\cap K$ is $\mathscr{O}_K$. Note that whereas the CM type determines a CM torus up to isogeny (respecting the action of the underlying CM field), the CM order is rather sensitive to the specific complex torus and it is not determined by the CM type.

Example 1.3.10. Let us consider the case of elliptic curves. What can be said about $\mathrm{End}(E)$ and $\mathrm{End}^0(E)$ for an elliptic curve E? Since $\mathrm{End}^0(E)$ injects into $\mathrm{End}_{\mathbf{C}}(\mathrm{T}_e(E)) = \mathbf{C}$, it follows that $\mathrm{End}^0(E)$ is a commutative field and so its $\mathbf{Q}$-degree is at most $2g = 2$. If $\mathrm{End}^0(E) = \mathbf{Q}$ then certainly the order $\mathrm{End}(E)$ is equal to $\mathbf{Z}$. Otherwise the preceding considerations show that $F = \mathrm{End}^0(E)$ must be an *imaginary quadratic* field. In this latter CM case the action by F on $\mathrm{T}_e(E)$ is given by an embedding $\varphi : F \hookrightarrow \mathbf{C}$ inducing an $\mathbf{R}$-algebra isomorphism $\mathbf{R}\otimes_{\mathbf{Q}} F \simeq \mathbf{C}$; this CM type $\{\varphi\}$ gives a *canonical* embedding of $F = \mathrm{End}^0(E)$ into $\mathbf{C}$; the canonicity is because this embedding is determined by the underlying 1-dimensional $\mathbf{C}$-structure on $\mathrm{T}_e(E)$ and the $\mathbf{C}$-linearity of the action by $\mathrm{End}^0(E)$ on this line. Thus, if $\mathscr{O} = \mathrm{End}(E)$ is the corresponding quadratic order in $\mathscr{O}_F$ then there is a (non-canonical) *$\mathscr{O}$-linear* isomorphism of elliptic curves

$$E \simeq (\mathbf{R}\otimes_{\mathbf{Q}} F)_{\{\varphi\}}/\mathfrak{a} \simeq \mathbf{C}/\varphi(\mathfrak{a})$$

for a fractional $\mathscr{O}$-ideal $\mathfrak{a} \subseteq F \subseteq \mathbf{C}$ satisfying $\mathrm{End}_{\mathbf{Z}}(\mathfrak{a}) \cap F = \mathscr{O}$. In practice, the phrase *CM elliptic curve* (over $\mathbf{C}$) may have two possible meanings: an elliptic curve E for which $\mathrm{End}^0(E)$ is an imaginary quadratic field F, or a CM torus (E, i) of dimension 1; the distinction is that in the latter case we are specifying an *abstract* field K together with an isomorphism $i : K \simeq \mathrm{End}^0(E)$. Passing to the conjugate CM type changes i but does not change E or $\mathrm{End}^0(E)$.

With more work it can be proved that, in the preceding notation, $\mathfrak{a}$ is an *invertible* $\mathscr{O}$-ideal, or in more intrinsic terms that $\mathrm{H}_1(E, \mathbf{Z})$ is an invertible $\mathrm{End}(E)$-module. In algebraic terms, if R is a Dedekind domain (such as $\mathbf{Z}$) and F'/F is a quadratic separable extension of its fraction field then for any 1-dimensional F'-vector space V' and rank-2 projective R-module $M' \subseteq V'$ such that $\mathscr{O} = \mathrm{End}_R(M') \cap F'$ (inside $\mathrm{End}_R(V')$) is strictly larger than R (and so is an order in the integral closure of R in F') then M' is an invertible $\mathscr{O}$-module. This is proved by a direct calculation using the explicit description of quadratic orders. See [S, 4.11, 5.4.2] for the method of proof in the special case $R = \mathbf{Z}$; this argument applies to any principal ideal domain in the role of $\mathbf{Z}$, and hence to any Dedekind domain R via localization to the case of a discrete valuation ring. Such invertibility even when the quadratic order is non-maximal is important in certain applications, such as in the theory of Heegner points.

Example 1.3.11. Let us explain what a Heegner point is from an analytic perspective. The word "point" is used because these structures correspond to certain points on modular curves, but for our present purposes we just describe what is being classified by such points.

Let K be an imaginary quadratic field with discriminant D and let $N \geq 1$ be an integer such that the prime factors of N are split in $\mathscr{O}_K$ (so $(N, D) = 1$). Let $\mathscr{O}$ be an order in K whose conductor f is coprime to ND. A *level-N Heegner point* (over $\mathbf{C}$) is an $\mathscr{O}$-linear isogeny $\pi : E' \to E$ between elliptic curves with CM order $\mathscr{O}$ such that $\ker\pi$ is cyclic with size N. (Since $(f, ND) = 1$, the CM order of E is automatically equal to that of E' if $\ker\pi$ is a submodule for the action of $\mathrm{End}(E')$.) The common CM type provides a *canonical* embedding of K into $\mathbf{C}$, and we may write $E' = V/\Lambda$ for an invertible $\mathscr{O}$-submodule $\Lambda \subseteq V$ that is a lattice (with $\mathscr{O}$ acting on V through the embedding of K into $\mathbf{C}$ provided by the CM type). The kernel $\ker\pi$ is an $\mathscr{O}$-module isomorphic to $\mathbf{Z}/N\mathbf{Z}$, so by the splitting hypothesis on the prime factors of N in $\mathscr{O}_K$ and the relative primality of the conductor of $\mathscr{O}$ and ND it follows that $\ker\pi$ is a free module of rank 1 over $\mathscr{O}/\mathfrak{n}$ for a unique invertible ideal $\mathfrak{n}$ of $\mathscr{O}$ satisfying $\mathfrak{n}\overline{\mathfrak{n}} = N\mathscr{O}$ with $(\mathfrak{n}, \overline{\mathfrak{n}}) = 1$. In terms of the geometric setup, $\mathfrak{n}$ is the annihilator ideal $\mathrm{Ann}_{\mathscr{O}}(\ker\pi)$. We may therefore uniquely identify E with V/Λ' for an $\mathscr{O}$-lattice Λ' in

V containing Λ with index N such that the given isogeny $E' \to E$ is the natural projection $V/\Lambda \to V/\Lambda'$. Hence, $\Lambda'/\Lambda = \ker \pi \simeq \mathscr{O}/\mathfrak{n}$, so $\Lambda' = \mathfrak{n}^{-1}\Lambda$ inside of V.

In the theory of CM elliptic curves one sometimes sees the notion of a *normalized* CM elliptic curve. This is an artificial and confusing concept that arises from the fact that in the classical theory all number fields were considered to be subfields of $\mathbf{C}$. More specifically, if E is an elliptic curve over $\mathbf{C}$ and K is an imaginary quadratic *subfield* of $\mathbf{C}$ equipped with an embedding $i : K \hookrightarrow \mathrm{End}^0(E)$ (so this is necessarily an isomorphism, by Example 1.3.10) then one gets *another* embedding of K into $\mathbf{C}$ through the CM type $\{\varphi\}$ that encodes the $\mathbf{C}$-linear action of $\mathrm{End}^0(E)$ on $\mathrm{T}_e(E)$. Now there arises the natural question: does φ agree with the way K is originally embedded into $\mathbf{C}$? If so, we say that the action is *normalized*. If the action is not normalized then composing it with complex conjugation on K will make it become normalized because there are only two embeddings of K into $\mathbf{C}$ and they are intertwined by complex conjugation; this amounts to passing to the conjugate CM type. Here is a non-normalized action:

Example 1.3.12. Let E be the elliptic curve $y^2 = x^3 - x$, and let $K/\mathbf{Q}$ be a splitting field of $z^2 + 1 = 0$. Let $\sqrt{-1} \in \mathscr{O}_K$ be a square root of -1, and pick an embedding of K into $\mathbf{C}$. The endomorphism $T : (x, y) \mapsto (-x, \varphi(\sqrt{-1})y)$ has square equal to $(x, y) \mapsto (x, -y) = [-1](x, y)$, so $\sqrt{-1} \mapsto T$ defines an embedding $i : \mathscr{O}_K \hookrightarrow \mathrm{End}(E)$ (and this is an equality, since $\mathscr{O}_K$ is the maximal order in K). Note that $-\sqrt{-1}$ acts as $-T$, so this $\mathscr{O}_K$-action i is independent of the choice of $\sqrt{-1}$ in K. Since y/x is a local coordinate at the identity and $T^*(y/x) = \varphi(-\sqrt{-1})(y/x)$, the CM type is $\overline{\varphi}$. Hence, if we consider K to be a subfield of $\mathbf{C}$ via φ then (E, i) is non-normalized.

The concept of a normalized action is only meaningful if one considers the imaginary quadratic field K as starting life inside of $\mathbf{C}$ prior to the consideration of its action on the elliptic curve. However, this is a tremendous mistake: one should always consider the quadratic field K as an "abstract" field without a preferred embedding into $\mathbf{C}$, and it is the *tangential representation* through the action on the elliptic curve (that is, the CM type) that should be taken to select the embedding of K into $\mathbf{C}$. With this modern point of view, the distinction between normalized and non-normalized actions never arises. The "abstract" ring $\mathrm{End}^0(E)$ may be considered to be the imaginary quadratic field associated to a CM elliptic curve E, and its image in $\mathrm{End}_{\mathbf{C}}(\mathrm{T}_e(E)) = \mathbf{C}$ is what *canonically* identifies $\mathrm{End}^0(E)$ with a subfield of $\mathbf{C}$. If we are interested in studying isomorphism classes of pairs (E, i) of CM elliptic curves with CM field equal to a fixed "abstract" imaginary quadratic field K then we will find both possible CM types arising for the same underlying elliptic curve E and CM field K, so it may be unwise to rule out half of the CM types arbitrarily.

Example 1.3.13. Let (E, i) be a CM elliptic curve with underlying CM field K. Let $X = E \times E$. Hence, via the specification of the isomorphism $i : K \simeq \mathrm{End}^0(E)$ we have $\mathrm{End}^0(X) \simeq \mathrm{Mat}_{2\times 2}(K)$. For *any* quadratic extension K'/K we may use a choice of K-basis for K' to embed K' into $\mathrm{Mat}_{2\times 2}(K)$ (through the K-linear multiplication action of K' on itself), and so we get an embedding $K' \hookrightarrow \mathrm{End}^0(X)$ of the quartic field K' into the endomorphism algebra of the 2-dimensional complex torus X. Such a field K' is totally complex (since it contains K) but it is often not CM. Observe that the action of K' on $\mathrm{T}_e(X)$ corresponds to the two embeddings of K' into $\mathbf{C}$ that lift the single embedding $i : K \hookrightarrow \mathrm{End}_{\mathbf{C}}(\mathrm{T}_e(E)) = \mathbf{C}$ that is the CM type of the given CM elliptic curve.

2. Analytic aspects of line bundles, duality, and pairings

In §1 we saw that some aspects of the theory of complex tori have a "linear algebra" flavor. For example, passage to the isogeny category is analogous to passing from $\mathbf{Z}$-lattices to rational vector spaces, and CM tori may be described in terms of uniformizations defined via CM types and fractional ideals for orders in CM fields. In this section we will study the structure of holomorphic line bundles on complex tori, and these will be described in terms of bilinear algebra data. This leads to a notion of duality for complex tori that is somewhat analogous to the concept of duality for vector spaces and lattices. We begin by introducing the notion of the dual complex torus in an *ad hoc* manner, and eventually we will see that this construction is a concrete model for an abstract universal object. The abstract viewpoint is what will be required for the correct notion of duality in the algebraic theory (for which the crutch of exponential uniformizations is not available, especially in positive characteristic).

2.1. **The dual torus.** Let X be a complex torus, and let $F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)/\Lambda$ be its double-coset uniformization as in (1.1.5). Equivalently, the quotient $V = F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$ contains Λ as a lattice via $\lambda\mapsto 1\otimes\lambda$, and V/Λ is the exponential uniformization of X. Let $F^{\dagger}\subseteq(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)^{\vee} = \mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\vee}$ be the annihilator of F, so concretely the subspace $F^{\dagger}$ is the dual $V^{\vee}$ to the quotient V of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$. Since $\mathbf{C}\otimes_{\mathbf{Z}}\mathbf{Z}(1)\simeq\mathbf{C}$ via multiplication, we may write $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\vee} = \mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\vee}(1)$ with $\Lambda^{\vee}(1) = \Lambda^{\vee}\otimes_{\mathbf{Z}}\mathbf{Z}(1)\simeq\mathrm{Hom}_{\mathbf{Z}}(\Lambda,\mathbf{Z}(1))$, and so for the lattice $\Lambda^{\dagger} = \Lambda^{\vee}(1)$ of $\mathbf{Z}(1)$-valued linear forms on Λ we may also say that $F^{\dagger}$ is a $\mathbf{C}$-subspace of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger}$. The reason for our interest in $\Lambda^{\dagger}$ rather than $\Lambda^{\vee}$ will become apparent later.

Theorem 2.1.1. *The map* $\Lambda^{\dagger}\to F^{\dagger}\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger})$ *is a lattice inclusion.*

The proof of the theorem works if we use $\Lambda^{\vee}$ in the role of $\Lambda^{\dagger}$.

Proof. We have to show that the $\mathbf{C}$-linear map $F^{\dagger}\oplus(\mathbf{R}\otimes_{\mathbf{Z}}\Lambda^{\dagger})\to\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger}$ is an isomorphism, or equivalently (for $\mathbf{R}$-dimension reasons) it is injective. An element $\ell\in\mathbf{R}\otimes_{\mathbf{Z}}\Lambda^{\dagger} = \mathrm{Hom}(\Lambda,\mathbf{R}(1))$ goes to the element in $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger} = \mathrm{Hom}(\Lambda,\mathbf{C}(1)) = \mathrm{Hom}(\Lambda,\mathbf{C})$ given by composing ℓ with the inclusion of $\mathbf{R}(1)$ into $\mathbf{C}$, so we have to prove that this composite is not induced by an element of $F^{\dagger}$, which is to say that its scalar extension $\ell_{\mathbf{C}}$ to a $\mathbf{C}$-linear map $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda\to\mathbf{C}$ does not kill F, except if $\ell = 0$. If $\ell_{\mathbf{C}}$ does kill F then $\ell_{\mathbf{C}}$ factors through the quotient $(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)/F$ that is $\mathbf{R}$-linearly spanned by Λ, so the image of $\ell_{\mathbf{C}}$ in $\mathbf{C}$ is the $\mathbf{R}$-span of $\ell(\Lambda)\subseteq\mathbf{R}(1)$. However, the image of the $\mathbf{C}$-linear map $\ell_{\mathbf{C}}$ is a $\mathbf{C}$-subspace of $\mathbf{C}$ and thus the containment inside $\mathbf{R}(1)$ forces it to be 0. ■

Definition 2.1.2. With notation as above, the *dual complex torus* $X^{\vee}$ is $F^{\dagger}\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger})/\Lambda^{\dagger}$.

In an evident manner, the dual complex torus $X^{\vee}$ has contravariant functorial dependence on X, and there is a canonical isomorphism $\iota_X : X\simeq X^{\vee\vee}$ induced by the evident double duality isomorphism $\Lambda^{\dagger\dagger} = \mathrm{Hom}(\mathrm{Hom}(\Lambda,\mathbf{Z}(1)),\mathbf{Z}(1))\simeq\Lambda$ (that carries $F^{\dagger\dagger}$ to F after extending scalars to $\mathbf{C}$); in particular, if $f : X\to X'$ is a map of complex tori then there is a natural induced map $f^{\vee} : X'^{\vee}\to X^{\vee}$. One also checks that $\iota_X^{\vee} = \iota_{X^{\vee}}^{-1}$, and that $f^{\vee\vee} = f$ via ι_X and $\iota_{X'}$ (this merely expresses the functoriality of ι_X in X).

Let us describe the exponential uniformization of the dual torus $X^{\vee}$ in terms of the exponential uniformization V/Λ of X, thereby recovering a formulation of the dual that is found in many books (such as [Mum]). We will find an intervention of the "non-holomorphic" operation of complex conjugation in this description, and so it is the double-coset formulation that is better-suited to generalizing this construction to the setting of analytically varying families of complex tori. Moreover, we shall see that the double-coset point of view is more convenient for the classification of holomorphic line bundles on a complex torus.

Under the decomposition $F\oplus\overline{F} = \mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ the subspace $\overline{F}$ maps isomorphically onto $V = F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$ and the quotient of $(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)^{\vee}\simeq\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger}$ modulo the annihilator $F^{\dagger}$ of the subspace F is identified with $F^{\vee}$. Since formation of the linear dual commutes with passage to the conjugate space, the $\mathbf{C}$-linear identification of $\overline{F}$ with V gives a $\mathbf{C}$-linear identification of $F^{\vee}$ with the conjugate-dual space $\overline{V}^{\vee}$. This conjugate-dual space may be ($\mathbf{C}$-linearly) viewed as the space of conjugate-linear functionals $\ell : V\to\mathbf{C}$ (that is, $\ell(c\cdot v) = \overline{c}\cdot\ell(v)$ for $c\in\mathbf{C}$ and $v\in V$) with $\mathbf{C}$-structure through the values of the functionals. One checks that the composite $\mathbf{C}$-linear isomorphism

$$(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda^{\dagger})/F^{\dagger}\simeq F^{\vee}\simeq\overline{V}^{\vee}$$

carries the lattice $\Lambda^{\dagger} = \Lambda^{\vee}(1) = \mathrm{Hom}(\Lambda,\mathbf{Z}(1))$ over to the $\mathbf{Z}$-submodule (and in fact, lattice) of semi-linear functionals $\ell : V\to\mathbf{C}$ whose imaginary component ℓ_{im} is $\mathbf{Z}(1)$-valued on Λ. Explicitly, for any $\lambda' : \Lambda\to\mathbf{Z}(1)$ we may form the $\mathbf{R}$-linear extension $\lambda'_{\mathbf{R}} : V = \mathbf{R}\otimes_{\mathbf{Z}}\Lambda\to\mathbf{R}(1)$ and we recover the $\mathbf{C}$-linear $\ell : V\to\mathbf{C}$ as $\ell(v) = -\sqrt{-1}\cdot\lambda'_{\mathbf{R}}(\sqrt{-1}\cdot v) + \lambda'_{\mathbf{R}}(v)$ for $v\in V$ and either $\sqrt{-1}\in\mathbf{C}$.

To summarize, we may say that the exponential uniformization of $X^{\vee}$ "is" $\overline{V}^{\vee}/\Lambda^{\vee}(1)$, where $\lambda'\in\Lambda^{\vee}(1)$ maps to the unique semilinear functional $V\to\mathbf{C}$ with imaginary component $\lambda'_{\mathbf{R}}$. In this language, if $f : X\to X'$ is a map of complex tori induced by a $\mathbf{C}$-linear map $T : V\to V'$ on tangent spaces at the identity then the dual map $f^{\vee} : X'^{\vee}\to X^{\vee}$ is induced by the conjugate-linear dual $\overline{T}^{\vee} : \overline{V}'^{\vee}\to\overline{V}^{\vee}$ on tangent spaces at the identity and on homology lattices the induced map $\mathrm{H}_1(f^{\vee})$ is the $\mathbf{Z}(1)$-dual of $\mathrm{H}_1(f)$.

Example 2.1.3. Let C be a compact connected Riemann surface and let X be its Jacobian as in Example 1.1.9. Let $\Lambda = \mathrm{H}_1(C, \mathbf{Z})$ and $F = \ker(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda \twoheadrightarrow W^\vee)$ with W the space of holomorphic 1-forms on C. The intersection pairing on cycles defines a canonical skew-symmetric perfect bilinear pairing

$$\mathrm{H}_1(C, \mathbf{Z}) \times \mathrm{H}_1(C, \mathbf{Z}) \to \mathbf{Z}(1)$$

given on a pair of smooth transverse loops by

$$(\sigma, \sigma') \mapsto \sigma.\sigma' = \sum_{x \in \sigma \cap \sigma'} \varepsilon_{i,x}(\sigma, \sigma') \cdot 2\pi i,$$

where $i = \sqrt{-1} \in \mathbf{C}$ and $\varepsilon_{i,x}(\sigma, \sigma') = \pm 1$ is the sign comparing two orientations: the i-orientation on $\mathrm{T}_x(C) \simeq \mathrm{T}_x(\sigma) \oplus \mathrm{T}_x(\sigma')$ and the product of the canonical orientations on the tangent lines to the oriented loops σ and σ' at x. Note that $\varepsilon_{i,x}(\sigma, \sigma') \cdot 2\pi i$ is independent of the choice of i (so this $\mathbf{Z}(1)$-valued intersection form does not depend on a choice of orientation for the underlying surface, in contrast with the $\mathbf{Z}$-valued intersection form used in elementary surface theory). This skew-symmetric pairing identifies Λ with $\Lambda^\dagger = \Lambda^\vee(1)$ via $\lambda \mapsto (\cdot).\lambda$, and the $\mathbf{C}$-scalar extension $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda \simeq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger$ carries F to $F^\dagger$. Consequently, this defines an isomorphism $\varphi_C : X \simeq X^\vee$. If we had used the other self-duality $\lambda \mapsto \lambda.(\cdot)$ that is negative to the one just considered, we would get the negated isomorphism $-\varphi_C$. We shall see in Example 3.1.4 why φ_C is better than $-\varphi_C$.

2.2. **Classification of line bundles.** Let $X = F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)/\Lambda$ be a complex torus with positive dimension. We wish to describe all holomorphic line bundles L over X. It is more convenient to give the description in terms of the double-coset uniformization rather than in terms of the exponential uniformization V/Λ. Moreover, we shall see that in this classification, due to Appell and Humbert, there is a "discrete" part and a "continuous" part. The dual torus $X^\vee$ will turn out to parameterize those isomorphism classes of $L \to X$ for which the "discrete" part of the classification data is trivial (in an appropriate sense).

Definition 2.2.1. Let Λ be a finite free $\mathbf{Z}$-module of rank $2g > 0$ and let $F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ be a $\mathbf{C}$-subspace of dimension g such that Λ maps to a lattice in $V = F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)$ via the natural map $\lambda \mapsto 1 \otimes \lambda$. An *Appell–Humbert datum* on (Λ, F) is a pair (H, α) where $H : (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda) \times (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda) \to \mathbf{C}$ is a $\mathbf{C}$-bilinear map and $\alpha : \Lambda \to \mathbf{C}^\times$ is a set-theoretic map such that:

1. H factors through the projection $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda \twoheadrightarrow V$ modulo F in the second variable,
2. $\Psi_H(w, w') = H(w, w') - H(w', w)$ has $\mathbf{Z}(1)$-valued restriction to $\Lambda \times \Lambda$,
3. $\alpha(\lambda + \lambda') = \alpha(\lambda)\alpha(\lambda')e^{\Psi_H(\lambda,\lambda')/2}$ for all $\lambda, \lambda' \in \Lambda$.

The order of subtraction in the definition of Ψ_H is required by the desire to have agreement between Riemann forms and Chern classes later in the theory; we will not address such compatibilities here. Note also that $(\lambda, \lambda') \mapsto e^{\Psi_H(\lambda,\lambda')/2}$ takes values in $\{\pm 1\}$.

Example 2.2.2. Let $B : V \times V \to \mathbf{C}$ be a symmetric bilinear pairing and $\ell \in V^\vee$ a linear functional. Let H_B denote the composition of B with the projection $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda \twoheadrightarrow V$ in both variables, and let $\alpha_\ell : \Lambda \to \mathbf{C}^\times$ be $e^{\ell|_\Lambda}$. The pair (H_B, α_ℓ) is obviously an Appell–Humbert datum. Note that such a pair uniquely determines B and ℓ, since if $\alpha_\ell = 1$ then $\ell(\Lambda) \subseteq \mathbf{Z}(1)$ and so the $\mathbf{C}$-subspace $\ell(V) = \ell(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda)$ in $\mathbf{C}$ is contained in $\mathbf{R}(1)$ (forcing it to vanish).

The set $\mathrm{AH}_{(\Lambda,F)}$ of all Appell–Humbert data for (Λ, F) forms a commutative group in an evident manner (componentwise in H and α), and the subset arising from pairs (B, ℓ) as above is a subgroup $\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$ that we shall call the *trivial* Appell–Humbert data. This subgroup of trivial data has a natural structure of Euclidean space over $\mathbf{C}$ (as it is the additive group of pairs (B, ℓ)).

The interest in Appell–Humbert data, and the special role of the trivial ones, is due to two facts: (i) any such datum (H, α) allows us to define a holomorphic line bundle $L(H, \alpha)$ on complex torus X, and (ii) this line bundle is trivial precisely when $(H, \alpha) \in \mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$. To explain how this works, write $X = V/\Lambda$ with $V = F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)$, so to give a holomorphic line bundle $L \to X$ it suffices to give a line bundle over V equipped with suitable Λ-action covering the Λ-translations on V. Consider the trivial line bundle $\mathbf{C} \times V \to V$. On this bundle we make $\lambda \in \Lambda$ act via

$$[\lambda] : (c, v) \mapsto (\alpha(\lambda) e^{H(\lambda,v)+H(\lambda,\lambda)/2} c, v + \lambda).$$

The cocycle condition that this action descends $\mathbf{C} \times V \to V$ to a holomorphic line bundle $L(H, \alpha) \to X$ is precisely the requirement on α in part (3) of Definition 2.2.1 (since $e^{\Psi_H(\lambda,\lambda')/2} = \pm 1$ for $\lambda, \lambda' \in \Lambda$).

Given two pairs (H, α) and (H', α') in $\mathrm{AH}_{(\Lambda,F)}$ we likewise get a unique isomorphism $L(H + H', \alpha\alpha') \simeq L(H, \alpha) \otimes L(H', \alpha')$ of line bundles over X descending the multiplication isomorphism $(\mathbf{C} \times V) \otimes (\mathbf{C} \times V) \simeq \mathbf{C} \times V$ of line bundles over V. Moreover, if $(H, \alpha) = (H_B, \alpha_\ell)$ is a trivial Appell–Humbert datum then $L(H, \alpha) \to X$ is canonically trivial (in terms of (H, α)) because the trivializing section $v \mapsto e^{B(v,v)/2+\ell(v)}$ of $\mathbf{C} \times V \to V$ descends to a trivializing section of $L(H, \alpha) \to X$. Hence, to each $(H, \alpha) \in \mathrm{AH}_{(\Lambda,F)}$ we have associated a line bundle on X whose isomorphism class only depends on (H, α) modulo trivial Appell–Humbert data. This construction exhausts the entire group $\mathrm{Pic}(X)$ of isomorphism classes of holomorphic line bundles on X (with group structure via tensor product):

Theorem 2.2.3 (Appell–Humbert)**.** *The natural map of groups* $\mathrm{AH}_{(\Lambda,F)}/\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)} \to \mathrm{Pic}(X)$ *induced by* $(H, \alpha) \mapsto L(H, \alpha)$ *is an isomorphism.*

Proof. See [Mum, Ch. I, §2] for a proof using different normalizations than ours (in the definition of an Appell–Humbert datum, and so also in the recipe for defining $L(H, \alpha)$). ■

Since Ψ_H is unaffected by replacing H with $H + H_B$ for any B as above (due to the symmetry of B), we conclude from the Appell–Humbert theorem that if $L \to X$ is a holomorphic line bundle and we choose (H, α) such that $L \simeq L(H, \alpha)$ then the skew-symmetric difference $\Psi_H : \Lambda \times \Lambda \to \mathbf{Z}(1)$ as in Definition 2.2.1(2) only depends on L and not on (H, α). Hence, we may denote it Ψ_L. This skew-symmetric $\mathbf{Z}(1)$-valued bilinear form on Λ is called the *Riemann form* of the holomorphic line bundle $L \to X$. By contstruction, the Riemann form makes F isotropic.

A $\mathbf{C}$-dimension count shows that every skew-symmetric $\mathbf{C}$-bilinear form on $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ that makes F isotropic has the form $(w, w') \mapsto H(w, w') - H(w', w)$ for some $\mathbf{C}$-bilinear form H on $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ that kills F in the second variable. Hence, for every $\mathbf{Z}(1)$-valued skew-symmetric bilinear form Ψ on Λ the subset $\mathrm{AH}^{\Psi}_{(\Lambda,F)}$ consisting of those (H, α) satisfying $\Psi_H = \Psi$ is a (possibly empty!) union of $\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$-cosets in the group $\mathrm{AH}_{(\Lambda,F)}$ and the elements of this union represent precisely those classes in $\mathrm{Pic}(X)$ with Riemann form Ψ.

Remark 2.2.4. By [Mum, pp. 36-37], if $g > 1$ then for "almost all" (Λ, F) the locus $\mathrm{AH}^{\Psi}_{(\Lambda,F)}$ is empty for every $\Psi \neq 0$. It is generally a subtle algebraic problem to determine those $\Psi \neq 0$ for which $\mathrm{AH}^{\Psi}_{(\Lambda,F)}$ is non-empty, and we will see later that in the algebraic theory over $\mathbf{C}$ such $\Psi \neq 0$ always exist. Put another way, it is a hard algebraic problem to determine if there is a nonzero $\mathbf{Z}(1)$-valued skew-symmetric form Ψ on Λ such that $\Psi_{\mathbf{C}}$ makes F isotropic.

Since the group $\mathrm{AH}_{(\Lambda,F)}$ consists of pairs (H, α) satisfying some simple linear-algebra and group-theoretic conditions, it is easy to write down explicit algebraic equations (say in terms of bases for Λ and F) that describe all such pairs and such that the group law on these pairs is given by algebraic expressions in the coordinates. More specifically, for each fixed Ψ the locus $\mathrm{AH}^{\Psi}_{(\Lambda,F)}$ is an open and closed subset (possibly empty) of $\mathrm{AH}_{(\Lambda,F)}$ and it is described by polynomial equations in a Euclidean space. In this way $\mathrm{AH}_{(\Lambda,F)}$ acquires a structure of Lie group with $\mathrm{AH}^{0}_{(\Lambda,F)}$ an open and closed subgroup that contains $\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$. But for any (H, α) such that $\Psi_H = 0$ it must be that H is symmetric, whence H also kills F in the first variable and so $H = H_B$ for a unique symmetric bilinear B on V. Thus, $\mathrm{AH}^{0}_{(\Lambda,F)}$ is the Lie group of pairs (B, α) for which B is a symmetric bilinear form on V and $\alpha : \Lambda \to \mathbf{C}^{\times}$ is a group homomorphism; this has the evident analytic structure as a product of two Lie groups: the Euclidean space of B's and the "algebraic torus" $\mathrm{Hom}(\Lambda, \mathbf{C}^{\times})$ of α's. Within here the subgroup of trivial Appell–Humbert pairs consists of those (B, α)'s for which α is an exponential $\alpha_\ell = e^{\ell|_\Lambda}$ for a (necessarily unique) $\ell \in V^{\vee}$. The problem of passing to an analytic quotient by $\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$ is entirely a problem with the α-aspect in the open and closed subgroup $\mathrm{AH}^{0}_{(\Lambda,F)}$ (as the Euclidean space of B's factors out as an analytic direct factor that is eliminated in the quotient).

Much as with Tate uniformizations in Example 1.1.12, we may use the canonical exact analytic exponential sequence $0 \to \mathbf{Z}(1) \to \mathbf{C} \to \mathbf{C}^\times \to 1$ to identify the Lie group $\mathrm{Hom}(\Lambda, \mathbf{C}^\times)$ with the analytic quotient

$$\mathrm{Hom}(\Lambda, \mathbf{C})/\,\mathrm{Hom}(\Lambda, \mathbf{Z}(1)) = (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\vee)/\Lambda^\vee(1) = (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)/\Lambda^\dagger.$$

In this way, the map $\ell \mapsto \alpha_\ell$ for $\ell \in V^\vee$ that provides the α-part of trivial Appell–Humbert pairs corresponds to the canonical mapping

$$V^\vee = F^\dagger \to (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)/\Lambda^\dagger$$

that is a closed embedding of Lie groups: it is the kernel of the submersion of Lie groups

$$(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)/\Lambda^\dagger \to F^\dagger \backslash (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)/\Lambda^\dagger = X^\vee$$

(that is obtained from passage to the quotient by the properly discontinuous action by the discrete $\Lambda^\dagger$ on source and target for the $\Lambda^\dagger$-equivariant quotient map of $\mathbf{C}$-vector spaces $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger \twoheadrightarrow F^\dagger \backslash (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)$). We may summarize our conclusions in a refined version of the Appell–Humbert theorem:

Theorem 2.2.5. *With its natural structure as a Euclidean space of pairs (B, ℓ), $\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$ is a closed subgroup of the Lie group $\mathrm{AH}_{(\Lambda,F)}$. The quotient Pic_X by this connected closed subgroup is naturally isomorphic to $\mathrm{Pic}(X)$ as a group and its connected components are precisely the open and closed subsets*

$$\mathrm{Pic}^\Psi_X = \mathrm{AH}^\Psi_{(\Lambda,F)}/\mathrm{AH}^{\mathrm{triv}}_{(\Lambda,F)}$$

that classify those line bundles with a fixed Riemann form Ψ and that are non-empty. The identity component Pic^0_X is naturally isomorphic to the dual complex torus $X^\vee$.

This theorem makes precise the sense in which the data of the Riemann form Ψ is the only discrete invariant of a holomorphic line bundle on a complex torus X: it classifies precisely the connected components of $\mathrm{Pic}(X)$ when this group is given an analytic structure as Pic_X via the preceding constructions. Let us work out what this means in a special case:

Example 2.2.6. Suppose that $X = E$ is an elliptic curve (with identity e). In this case we claim that Ψ_L "is" just the degree of L by another name. By the definition of the Riemann form, clearly Ψ_L depends additively on L. By writing down explicit Appell–Humbert data, one can show that the map $E \to \mathrm{Pic}_E$ defined by $x \mapsto \mathscr{O}_X((e) - (x))$ is analytic. By the theory of elliptic curves this is injective, so it follows by elementary analytic considerations in dimension 1 that this map must have image precisely equal to the compact 1-dimensional identity component Pic^0_E. Hence, Pic^0_E consists of exactly the line bundles of degree 0, whence Pic_E has infinite cyclic component group of rank 1 as desired. (There is another approach, somewhat less *ad hoc*: relate Riemann forms and first Chern classes, and using such Chern classes to provide the relation with degree for line bundles on a compact connected Riemann surface such as E.)

One can go further and build a "universal line bundle" $P_X \to X \times \mathrm{Pic}_X$ such that, among other things, $P_X|_{X \times \{\xi\}}$ represents the isomorphism class ξ for all $\xi \in \mathrm{Pic}_X = \mathrm{Pic}(X)$. To make this universality precise goes somewhat beyond our aims, and so we shall now focus our attention on an analogous universality result for the identity component Pic^0_X that we have seen is identified with the dual complex torus $X^\vee$. Roughly speaking, we have shown that $X^\vee$ classifies those holomorphic line bundles that can be put in a connected analytic family with the trivial line bundle (or, less geometrically, it classifies those $L \to X$ whose discrete invariant Ψ_L vanishes), and we want to give a precise meaning to this rough idea. (For example, what is a "connected analytic family" of holomorphic line bundles over X?) We require a new concept:

Definition 2.2.7. Let M be a compact connected complex manifold and let $m_0 \in M$ be a point. An *m_0-trivialized line bundle* on M is a pair (L, i) with $L \to M$ a holomorphic line bundle and $i : \mathbf{C} \simeq L(m_0)$ a trivialization of the fiber over m_0 (that is, a choice of basis of $L(m_0)$).

Example 2.2.8. Let $M = X \times X'$ with X and X' compact connected manifolds and $m_0 = (x_0, x'_0)$. In this case, a holomorphic line bundle $L \to M$ may be viewed as a "connected analytic family" of line bundles on X with parameter space X' in the sense that for all $x' \in X'$ we get a holomorphic line bundle $L_{x'} = L|_{X \times \{x'\}}$. (The holomorphicity of the "dependence on x'" is encoded in the holomorphicity of $L \to M$.) To give an m_0-trivialization is to give a basis in the fiber line $L(x_0, x'_0)$.

Observe that, in contrast with "bare" holomorphic line bundles, an m_0-trivialized pair (L, i) has no non-trivial automorphisms in the sense that if $\phi : L \simeq L$ is a bundle automorphism such that $\phi(m_0) : L(m_0) \simeq L(m_0)$ respects the chosen basis on this line (via i) then ϕ must be the identity. Indeed, as an automorphism of a holomorphic line bundle the map ϕ must be multiplication by a global non-vanishing holomorphic function, and since M is compact and connected such a function is a nonzero scalar. By working on the m_0-fiber, we see that this scalar is 1. It follows that if (L, i) and (L', i') are two m_0-trivialized line bundles on M then an isomorphism $\phi : L \simeq L'$ that respects the m_0-trivializations (in the sense that $\phi(m_0) : L(m_0) \simeq L'(m_0)$ respects the chosen basis in each fiber line) must be unique if it exists. This uniqueness aspect says that the notion of an m_0-rigidified line bundle allows no ambiguity when we speak of two such structures as being isomorphic (that is, we do not need to say what the isomorphism is because it is unique if it exists).

Example 2.2.9. Let $e^\vee \in X^\vee$ be the origin. We shall use the Appell–Humbert construction to build an $(e, e^\vee)$-rigidified line bundle P_X over the complex torus $X \times X^\vee$. We have the double-coset descriptions

$$X = F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)/\Lambda, \quad X^\vee = F^\dagger\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)/\Lambda^\dagger$$

with $\Lambda^\dagger = \mathrm{Hom}(\Lambda, \mathbf{Z}(1))$ and $F^\dagger$ equal to the annihilator of $F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$. Concretely, $F^\dagger = V^\vee$ as a subspace of $(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)^\vee$, where V is the quotient $F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)$. By the Appell–Humbert theorem, to give a line bundle on $X \times X^\vee$ is "the same" as to give an Appell–Humbert datum modulo trivial such data. We will give a canonical such datum, but we first need to write down a convenient double-coset description for the complex torus $X \times X^\vee$. By taking direct sums of the description for each factor, we have

$$X \times X^\vee = (F \oplus F^\dagger)\backslash(\mathbf{C} \otimes_{\mathbf{Z}} (\Lambda \oplus \Lambda^\dagger))/(\Lambda \oplus \Lambda^\dagger) \tag{2.2.1}$$

with $\Lambda^\dagger = \Lambda^\vee(1)$ and $F^\dagger$ dual to the quotient V of $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ modulo F.

Since the conjugate subspace $\overline{F} \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ projects isomorphically onto the quotient V it gives a natural $\mathbf{C}$-linear splitting

$$\mathbf{C} \otimes_{\mathbf{Z}} \Lambda = F \oplus V,$$

and we likewise have that the conjugate subspace to $F^\dagger = V^\vee$ projects isomorphically to the quotient $F^\vee$ of $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger = (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)^\vee$ so we have a natural $\mathbf{C}$-linear splitting $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger = V^\vee \oplus F^\vee$ (where $V^\vee$ is the subspace $F^\dagger$). Using these decompositions, we write typical elements $w \in \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ and $w' \in \mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger$ in the form $w = w_F + w_V$ and $w' = w'_{V^\vee} + w'_{F^\vee}$. The subspace $F \oplus F^\dagger$ in $\mathbf{C} \otimes_{\mathbf{Z}} (\Lambda \oplus \Lambda')$ (as in (2.2.1)) is precisely the subspace $F \oplus V^\vee$, so an object in $\mathrm{AH}_{(\Lambda \oplus \Lambda^\dagger, F \oplus F^\dagger)}$ is a pair of maps

$$H : ((\mathbf{C} \otimes_{\mathbf{Z}} \Lambda) \oplus (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)) \times ((\mathbf{C} \otimes_{\mathbf{Z}} \Lambda) \oplus (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger)) \to \mathbf{C}, \quad \alpha : \Lambda \oplus \Lambda^\dagger \to \mathbf{C}^\times$$

where H is a $\mathbf{C}$-bilinear map and α is a group homomorphism that together satisfy the Appell–Humbert conditions: H kills $F \oplus V^\vee$ in the second variable, the map

$$\Psi_H(w_1 + w'_1, w_2 + w'_2) = H(w_1 + w'_1, w_2 + w'_2) - H(w_2 + w'_2, w_1 + w'_1)$$

is $\mathbf{Z}(1)$-valued on $(\Lambda \oplus \Lambda^\dagger) \times (\Lambda \oplus \Lambda^\dagger)$, and $\alpha(\ell_1 + \ell_2) = \alpha(\ell_1)\alpha(\ell_2)e^{\Psi_H(\ell_1, \ell_2)/2}$ for $\ell_1, \ell_2 \in \Lambda \oplus \Lambda^\dagger$.

Define $H_X(w_1 + w'_1, w_2 + w'_2) = w'_{2,F^\vee}(w_1) - w'_{1,V^\vee}(w_2)$ and $\alpha_X(\lambda + \lambda') = e^{\lambda'(\lambda)/2} = \pm 1$ for $w_j \in \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$, $w'_j \in \mathbf{C} \otimes_{\mathbf{Z}} \Lambda^\dagger = (\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)^\vee$, $\lambda \in \Lambda$, and $\lambda' \in \Lambda^\dagger = \mathrm{Hom}(\Lambda, \mathbf{Z}(1))$. It is straightforward to check all of the above requirements on H and α, and in fact $\Psi_{H_X}(w_1 + w'_1, w_2 + w'_2) = w'_2(w_1) - w'_1(w_2)$. Thus, the line bundle $L(H_X, \alpha_X)$ on $X \times X^\vee$ has Riemann form Ψ on $(\Lambda \oplus \Lambda^\dagger) \times (\Lambda \oplus \Lambda^\dagger)$ given by

$$(\lambda_1 + \lambda'_1, \lambda_2 + \lambda'_2) \mapsto \Psi(\lambda_1 + \lambda'_1, \lambda_2 + \lambda'_2) = \lambda'_2(\lambda_1) - \lambda'_1(\lambda_2) \in \mathbf{Z}(1)$$

for $\lambda_j \in \Lambda$ and $\lambda'_j \in \Lambda^\vee(1)$. It can be shown by direct calculation that the associated line bundle $P_X = L(H_X, \alpha_X) \to X \times X^\vee$ has pullback over $X \times \{\xi\}$ that represents $\xi \in \mathrm{Pic}^0_X$ (for any choice of ξ), and more importantly that the pullbacks $P_X|_{X \times \{e^\vee\}}$ over X and $P_X|_{\{e\} \times X^\vee}$ over $X^\vee$ admit canonical trivializations i_X and $i_{X^\vee}$ that induce the same basis on the fiber over $(e, e^\vee)$. In particular, P_X is an $(e, e^\vee)$-trivialized line bundle over $X \times X^\vee$.

2.3. **Correspondences and universal properties.** The triple $(P_X, i_X, i_{X^\vee})$ constructed in Example 2.2.9 is equipped with extra trivialization structure as in the following definition:

Definition 2.3.1. Let (M, m_0) and (M', m_0') be compact connected complex manifolds endowed with marked points. A *correspondence* on $M \times M'$ between M and M' is a triple (L, i, i') with $L \to M \times M'$ a holomorphic line bundle and $i : \mathbf{C} \times M \simeq L|_{M \times \{m_0'\}}$ and $i' : \mathbf{C} \times M' \simeq L|_{\{m_0\} \times M'}$ trivializations over M and M' such that they induce the same basis vector on the fiber $L(m_0, m_0')$. If i and i' are understood from context, we may simply refer to L as a correspondence.

As with rigidified line bundles, a correspondence (L, i, i') admits no non-trivial automorphisms. Moreover, the notion of correspondence is symmetric in M and M' in the sense that if $s : M' \times M \simeq M \times M'$ is the flipping automorphism then $(s^*(L), i', i)$ is a correspondence on $M' \times M$ between M' and M. Also, if $f : M_1 \to M_2$ and $f' : M_1' \to M_2'$ are analytic maps respecting the base points then pullback along $f \times f'$ carries correspondences between M_2 and M_2' back to correspondences between M_1 and M_1'. The absence of non-trivial automorphisms for correspondences ensures that there is no ambiguity when we say that a correspondence between M_1 and M_1' "is" a pullback (along $f \times f'$) of one between M_2 and M_2'; it is not necessary to say what the meaning of "is" is.

The triple $(P_X, i_X, i_{X^\vee})$ is the *Poincaré correspondence* between X and $X^\vee$, and the line bundle P_X is the *Poincaré bundle* over $X \times X^\vee$. The Poincaré correspondence is characterized abstractly by the following universal property:

Theorem 2.3.2. *Let (X, e) be a complex torus and (M', m_0') be a pointed connected compact complex manifold. Let (L, i, i') be a correspondence on $X \times M'$ between (X, e) and (M', m_0'). There is a unique map of complex manifolds $f : M' \to X^\vee$ carrying m_0' to $e^\vee$ such that (L, i, i') is the $(1_X \times f)$-pullback of the Poincaré correspondence $(P_X, i_X, i_{X^\vee})$.*

In case $(M', m_0') = (X', e')$ is a complex torus, the map f is a map of complex tori since it respects the identity points.

Proof. In case (M', m_0') is required to be a complex torus, which is the case we need, this is the content of [Mum, Ch. II, §9], up to the use of different Appell–Humbert normalizations than ours. Set-theoretically, $f(m') \in X^\vee \subseteq \mathrm{Pic}(X)$ is the isomorphism class of the line bundle $L_{m'} = L|_{X \times \{m'\}}$ over X that sits in the "connected analytic family" $\{L_{m'}\}_{m' \in M'}$ with the trivial bundle $L_{m_0'}$ (trivialization provided by i). ■

Example 2.3.3. The universal property in Theorem 2.3.2 gives us a more conceptual way to understand the meaning of the double-duality isomorphism $\iota_X : X \simeq X^{\vee\vee}$ that was earlier constructed as an *ad hoc* isomorphism via the explicit double-coset analytic definition of the dual complex torus. The essential point is that Theorem 2.3.2 gives us a way to think about $X^\vee$ *independent* of an explicit construction. To be more specific, if we apply Theorem 2.3.2 with the complex torus $X^\vee$ in the role of X then the correspondence $(s^*(P_X), i_{X^\vee}, i_X)$ between $X^\vee$ and X must arise as the $(1_{X^\vee} \times f)$-pullback of the Poincaré correspondence between $X^\vee$ and $X^{\vee\vee}$ for a unique map of tori $f : X \to X^{\vee\vee}$. Set-theoretically, $f(x) \in X^{\vee\vee} \subseteq \mathrm{Pic}(X^\vee)$ is the isomorphism class of the line bundle $P_X|_{\{x\} \times X^\vee}$. If one looks closely at the analytic models for $X^\vee$ and P_X, it emerges that this map f is precisely ι_X (and in particular it is an isomorphism).

Example 2.3.4. Recall that if $f : X \to X'$ is a map of tori then by using the explicit analytic models for the dual tori we made an *ad hoc* dual map $f^\vee : X'^\vee \to X^\vee$. Now this map can be given a conceptual meaning: pullback by f defines a map of groups $f^* : \mathrm{Pic}(X') \to \mathrm{Pic}(X)$, and the restriction of this map to the subgroup $X'^\vee$ is precisely $f^\vee$. That is, for the isomorphism class $\xi' \in X'^\vee$ of some line bundle L' on X' with trivial discrete invariant, $f^\vee(\xi') \in X^\vee$ is the isomorphism class of the line bundle $f^*(L')$ on X. The verification of this conceptual description of the dual map rests on direct calculation with Appell–Humbert data.

Recall that via the double-duality identifications ι_X and $\iota_{X'}$, we have $f^{\vee\vee} = f$ (as this is obvious from the analytic viewpoint). In fact, this equality can also be deduced from the abstract approach via universal properties for the dual torus.)

The significance of Theorem 2.3.2, Example 2.3.3, and Example 2.3.4 cannot be overstated: they provide abstract ways to think about the dual torus, dual map, and double-duality isomorphism in purely geometric language (line bundles, *etc.*) *without* the intervention of analytic uniformizations. It is with this abstract geometric point of view, based on universal properties, that we shall be able to formulate analogous notions in the algebraic theory.

Remark 2.3.5. Analogues of Theorem 2.3.2 can be proved for more general compact connected pointed complex manifolds in the role of (X, e). For example, when we work with (C, x_0) for a compact connected Riemann surface C then there is a "universal correspondence" over $C \times \mathrm{Jac}(C)$ (rigified along x_0 and the origin of $\mathrm{Jac}(C)$). We will not address the construction here, but we emphasize that such a universal property for the Jacobian provides the mechanism by which one defines and constructs a Jacobian in the algebraic theory. The classical Abel-Jacobi theorem identifies the underlying group of $\mathrm{Jac}(C)$ with the group $\mathrm{Pic}^0(C)$ of isomorphism classes of degree-0 line bundles in $\mathrm{Pic}(C)$, but it is important that not only the group structure on $\mathrm{Jac}(C)$ be given an intrinsic meaning (as the group of isomorphism classes of degree-0 line bundles on C) but also the *analytic structure* on this group be given an intrinsic meaning. It is the viewpoint of universal correspondences that enables one to give the analytic structure on this group (via $\mathrm{Jac}(C)$) an intrinsic meaning (via a universal property for maps from compact connected complex manifolds endowed with a marked point).

Let us conclude our discussion of correspondences with an important class of self-correspondences. Our construction will satisfy certain symmetry properties, so we begin with a definition:

Definition 2.3.6. A correspondence (L, i, i') between a complex torus (X, e) and itself it *symmetric* if $s^*(L) \simeq L$ as correspondences (in particular, $i' = i$). A map $f : X \to X^\vee$ is *symmetric* if the dual map $f^\vee \circ \iota_X : X \simeq X^{\vee\vee} \to X^\vee$ is equal to f.

By the universal property of the dual torus, if L is a correspondence on $X \times X$ and $f : X \to X^\vee$ is the associated map of tori (so $L \simeq (1_X \times f)^*(P_X)$ as correspondences) then the correspondence $s^*(L)$ is classified by the dual mapping $f^\vee$. It follows that symmetry for f is equivalent to symmetry for L.

Example 2.3.7. Let $m : X \times X \to X$ be the multiplication mapping, and let L be a line bundle on X. The line bundle

$$\wedge(L) = m^*(L) \otimes p_1^*(L)^{-1} \otimes p_2^*(L)^{-1}$$

on $X \times X$ has pullbacks along $X \times e$ and $e \times X$ that are each identified with $L(e) \times (X \times X) \to X \times X$. Hence, upon picking a basis of the fiber $L(e)$ we obtain compatible trivializations i and i' for $\wedge(L)$ along $e \times X$ and $X \times e$, so we get a correspondence $(\wedge(L), i, i')$ on $X \times X$; we shall call this the *Mumford correspondence* associated to L (and $\wedge(L)$ is the *Mumford bundle* associated to L). If we change the choice of basis of $L(e)$ then both i and i' change by the same $\mathbf{C}^\times$-factor. Hence, the isomorphism class of this correspondence is independent of the choice of basis of the line $L(e)$. In what follows, everything we do will be independent of this choice and so we do not dwell on it any further.

It is clear by construction that the Mumford correspondence $\wedge(L)$ on $X \times X$ is symmetric. What is the associated symmetric map $\phi_L : X \to X^\vee$? By the description given in the proof of Theorem 2.3.2, $\phi_L(x) \in X^\vee$ is the isomorphism class of the line bundle

$$(\wedge(L))|_{X \times \{x\}} = t_x^*(L) \otimes L^{-1} \otimes (L(x)^\vee \times X) \simeq t_x^*(L) \otimes L^{-1}$$

on X (as $L(x)^\vee \simeq \mathbf{C}$), with $t_x : X \to X$ the additive translation map $y \mapsto y + x$. By symmetry of $\wedge(L)$, it follows that the map $\phi_L : X \to X^\vee$ sending x to $t_x^*(L) \otimes L^{-1}$ is symmetric. It is obvious from the definition that the Mumford correspondence has additive dependence on L, so $\phi_{L_1 \otimes L_2} = \phi_{L_1} + \phi_{L_2}$.

In the language of divisors, if $L = \mathscr{O}_X(D)$ is the line bundle of meromorphic functions with "poles no worse than D" then $t_x^*(L) = \mathscr{O}_X(t_x^{-1}(D)) = \mathscr{O}_X(t_{-x}(D))$. Hence, $\phi_{\mathscr{O}_X(D)}(x) = \mathscr{O}_X(t_{-x}(D) - D)$. The appearance of $-x$ rather than x in this formula, even in the case of elliptic curves, has led some to say that this definition of ϕ_L is "wrong" by a sign. We shall see in Example 3.1.6 that such objections based on the language of divisors are misleading: *every* symmetric map $X \to X^\vee$ has the form ϕ_L for a line bundle L, and if L is an *ample* line bundle (a "positivity" condition to be defined later) then $(1, \phi_L)^*(P_X) = \Delta_X^*(\wedge(L))$ is also ample on X, whereas $(1, -\phi_L)^*(P_X) = (1, \phi_{L^{-1}})^*(P_X) = \Delta_X^*(\wedge(L))^{-1}$ is not (with $\Delta_X : X \to X \times X$ the diagonal map). In this respect, ϕ_L is "better" than $-\phi_L$.

As a concrete example, in the case of an elliptic curve (E, e) if we take $L = \mathscr{O}_E(e)$ then $\phi_L : E \to E^\vee$ is the map $x \mapsto \mathscr{O}_E((-x) - (e))$ that is the negative of the "classical" auto-duality of an elliptic curve. (In particular, $\phi_{\mathscr{O}_E(e)}$ is an isomorphism.) However, it turns out that $(1, \phi_{\mathscr{O}_E(e)})^*(P_E) \simeq \mathscr{O}_E(2e)$, so this is the line bundle associated to an effective divisor, whereas $(1, -\phi_{\mathscr{O}_E(e)})^*(P_E) \simeq \mathscr{O}_E(-2e)$ is associated to an anti-effective divisor. It is the former that has good positivity properties, and this is the sense in which $\phi_{\mathscr{O}_E(e)}$ is better than the classical autoduality $-\phi_{\mathscr{O}_E(e)} = \phi_{\mathscr{O}_E(-e)}$.

Example 2.3.8. The Mumford correspondence is well-behaved with respect to pullback, in the sense that if $f : X \to X'$ be a map of tori and L' is a line bundle on X' then $(f \times f)^*(\wedge(L')) \simeq \wedge(f^*(L'))$ as correspondences on $X \times X$ (since f is compatible with the multiplication laws on X and X'). By using the isomorphism of correspondences $(1 \times \phi_{L'})^*(P_{X'}) \simeq \wedge(L')$ that uniquely characterizes $\phi_{L'}$, as well as the analogous characterization for $f^\vee$, one can deduce by diagram-chasing that $f^\vee \circ \phi_{L'} \circ f = \phi_{f^*(L')}$. This can also be seen by direct calculation: since $t_{f(x)} \circ f = f \circ t_x$ for any $x \in X$, we have

$$(f^\vee \circ \phi_{L'} \circ f)(x) = f^*(t_{f(x)}^*(L') \otimes L'^{-1}) \simeq t_x^*(f^*(L')) \otimes (f^*(L'))^{-1} = \phi_{f^*(L')}(x).$$

2.4. **Pairings and torsion.** Roughly speaking, a correspondence between complex tori is to be considered as analogous to a bilinear pairing $W \times W' \to k$ of finite-dimensional vector spaces over a field and the associated map to the dual torus is analogous to the associated linear map $W' \to W^\vee$. In this spirit, the bilinear evaluation pairing $B_W : W \times W^\vee \to k$ is an analogue of the Poincaré correspondence, and it even has a similar (obvious) universal property: for every bilinear pairing $B : W \times W' \to k$ there is a unique linear map $T_B : W' \to W^\vee$ such that $B(w, w') = B_W(w, T_B(w')) = (T_B(w'))(w)$. Can we strenthen this analogy by associating a genuine bilinear pairing to a correspondence between complex tori? If so, do the linear-algebra concepts of symmetric bilinear form and non-degenerate bilinear form have reasonable analogues for correspondences? We shall provide affirmative answers to these questions in two ways: we will first construct pairings on homology $\mathbf{Z}$-lattices, and then we will give analogues on torsion points. It is the formulation via torsion (and consequently an ℓ-adic torsion-free formulation) that will carry over to the algebraic theory.

Let (L, i, i') be a correspondence between complex tori X and X', and let Λ and Λ' be the associated homology lattices, so $\Lambda \oplus \Lambda'$ is the homology lattice for $X \times X'$. The line bundle L on $X \times X'$ has a skew-symmetric Riemann form

$$\Psi_L : (\Lambda \oplus \Lambda') \times (\Lambda \oplus \Lambda') \to \mathbf{Z}(1)$$

and by inspection of Appell–Humbert data this pairing makes both subgroups $\Lambda, \Lambda' \subseteq \Lambda \oplus \Lambda'$ isotropic, so Ψ_L is uniquely determined by its restriction $\psi_{(L,i,i')} = \Psi_L|_{(\Lambda \oplus \{0\}) \times (\{0\} \oplus \Lambda')}$ to $\Lambda \times \Lambda'$; that is, $\psi_{(L,i,i')}(\lambda, \lambda') = \Psi_L((\lambda, 0), (0, \lambda'))$. We call $\psi_{(L,i,i')}$ the *Riemann form* associated to the correspondence (L, i, i'), and we denote it ψ_L since it does not depend on the auxiliary choice of compatible trivializations i and i'. It must be emphasized that the Riemann form ψ_L is a pairing between Λ and Λ', so it is *not* to be confused with the Riemann form Ψ_L attached to the line bundle L on the complex torus $X \times X'$ (as this is an entirely different object, namely a skew-symmetric pairing between $\Lambda \oplus \Lambda'$ and itself). Beware that there is another natural pairing to consider, namely $\psi'_L(\lambda, \lambda') = \Psi_L((0, \lambda'), (\lambda, 0)) = -\psi_L(\lambda, \lambda')$. The superiority of ψ_L over ψ'_L will be explained soon.

The Riemann form ψ_L uniquely determines the correspondence (L, i, i') in the sense of the following theorem (whose proof is a calculation with Appell–Humbert data):

Theorem 2.4.1. *Let $V = \mathrm{T}_e(X)$ and $V' = \mathrm{T}_{e'}(X')$. The correspondence (L, i, i') on $X \times X'$ descends the trivial correspondence $\mathbf{C} \times (V \oplus V')$ over $V \oplus V'$ (equipped with the natural trivializations along the origins of V and V') via the $\Lambda \oplus \Lambda'$-action*

$$[\lambda + \lambda'] : (c, v + v') = (e^{\psi_L(\lambda_F, \lambda') - \psi_{L,\mathbf{C}}(v, \lambda') + \psi_{L,\mathbf{C}}(\lambda, v')} c, (v + \lambda) + (v' + \lambda'))$$

for $c \in \mathbf{C}$, $v \in V$, and $v' \in V'$ (with λ_F denoting the F-component of $\lambda \in \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ under the canonical splitting of the Hodge filtration, and the elements $v \in V$ and $v' \in V'$ likewise viewed in $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ and $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda'$ respectively via such splittings).

Example 2.4.2. Consider the Poincaré correspondence P_X between X and $X^\vee$, using the analytic model for $X^\vee$ as in our initial definition. Let $\Lambda = \mathrm{H}_1(X, \mathbf{Z})$, so from the analytic definition of $X^\vee$ we obtain an identification of $\mathrm{H}_1(X^\vee, \mathbf{Z})$ with $\Lambda^\dagger = \mathrm{Hom}(\Lambda, \mathbf{Z}(1))$. We thereby obtain a canonical bilinear pairing

$$\psi_{P_X} : \Lambda \times \mathrm{Hom}(\Lambda, \mathbf{Z}(1)) \to \mathbf{Z}(1),$$

so there naturally arises the question of what this canonical pairing is; it is indeed the evaluation pairing. In particular, if $f : X \to X'$ is a map of complex tori then the adjoint of $\mathrm{H}_1(f)$ with respect to ψ_{P_X} is the homology map $\mathrm{H}_1(f^\vee)$ of the dual morphism. The canonical pairing ψ_{P_X} between the homology lattices of X and $X^\vee$ is called the *Weil pairing.*

Now we can give two reasons why ψ_L is better than ψ'_L: Theorem 2.4.1 with ψ'_L has a sign in the exponential, and the comparison of ψ'_{P_X} with the evaluation pairing gets a sign.

Definition 2.4.3. If $f : X' \to X^\vee$ is a map of tori, then its *Riemann form* $\Psi_f : \Lambda \times \Lambda' \to \mathbf{Z}(1)$ is the Riemann form ψ_L associated to the correspondence $L = (1_X \times f)^*(P_X)$ that is "classified" by the map f.

Remark 2.4.4. The Riemann form Ψ_f is can be described in rather concrete terms. It may be viewed as a map $\Lambda' \to \Lambda^\vee(1)$, and as such it is *exactly* the homology map $\mathrm{H}_1(f)$ (due to Example 2.4.2). Thus, by Theorem 1.2.6, the pairing Ψ_f is non-degenerate if and only if f is an isogeny. There is also a uniqueness result: by Theorem 2.4.1, the Riemann form Ψ_f attached to a map $f : X' \to X^\vee$ and the Riemann form ψ_L attached to a correspondence L on $X \times X'$ are bilinear data that *uniquely determine* f and L respectively.

For any map $f : X' \to X^\vee$ we may identify the dual map $f^\vee$ with a map $X \simeq X^{\vee\vee} \to X'^\vee$, so $\Psi_{f^\vee}$ is identified with a bilinear pairing $\Lambda' \times \Lambda \to \mathbf{Z}(1)$. It therefore makes sense to compare $\Psi_f(\lambda, \lambda')$ and $\Psi_{f^\vee}(\lambda', \lambda)$ for $\lambda \in \Lambda$ and $\lambda' \in \Lambda'$. This comparison involves a sign,

$$\Psi_{f^\vee}(\lambda', \lambda) = -\Psi_f(\lambda, \lambda'),$$

essentially because Riemann forms are defined in terms of skew-symmetric pairings as in Definition 2.2.1(2). In the case $X' = X$, which is to say the case of self-correspondences of X, it follows that the self-pairing Ψ_f on Λ induced by a map $f : X \to X^\vee$ is skew-symmetric if and only if the map f is symmetric (in the sense of Definition 2.3.6). Equivalently, a self-correspondence L on $X \times X$ is symmetric if and only if its Riemann form ψ_L is a skew-symmetric form on Λ.

Example 2.4.5. If $f : X \to X^\vee$ is a map then $L_0 = (1, f)^*(P_X)$ is a line bundle on X, and since f is determined by the bilinear self-pairing Ψ_f on Λ it makes sense to ask for the computation of the skew-symmetric Riemann form $\Psi_{L_0} : \Lambda \times \Lambda \to \mathbf{Z}(1)$ in terms of Ψ_f. The answer is

$$\Psi_{L_0}(\lambda_1, \lambda_2) = \Psi_f(\lambda_1, \lambda_2) - \Psi_f(\lambda_2, \lambda_1),$$

so in the case that f is symmetric (or equivalently, Ψ_f is skew-symmetric) this collapses to the identity $\Psi_{L_0} = 2\Psi_f$.

Example 2.4.6. Let us return to the Mumford correspondences constructed in Example 2.3.7. Pick a line bundle L on X and let $f = \phi_L : X \to X^\vee$ be the map classifying the symmetric Mumford correspondence $\wedge(L)$. How is the skew-symmetric pairing $\Psi_f = \psi_{\wedge(L)}$ on Λ related to the skew-symmetric Riemann form Ψ_L on Λ? Unwinding the analytic construction of the Poincaré correspondence, one finds that $\Psi_{\phi_L} = \Psi_L$ as skew-symmetric bilinear pairings on Λ. Thus, ϕ_L is an isogeny if and only if Ψ_L is a non-degenerate pairing. Also, since f is determined by Ψ_f it follows that f is determined by the discrete invariant of L. For example, $\phi_L = 0$ if and only if $L \in \mathrm{Pic}^0_X$. We emphasize that the map ϕ_L only "knows" L through the discrete invariant Ψ_L (that is the skew-symmetric Riemann form of the symmetric map ϕ_L); changing L by tensoring against a member of Pic^0_X does not change the associated symmetric map from X to $X^\vee$. In §3 we shall see that there is much interest in certain cases for which the symmetric map ϕ_L is an isogeny.

We conclude our discussion of pairings by considering their effect on torsion. If (L, i, i') is a correspondence between tori X and X', and $f : X' \to X^\vee$ is the associated map, then we get a bilinear Riemann form

$$\Psi_f = \psi_L : \Lambda \times \Lambda' \to \mathbf{Z}(1)$$

that is non-degenerate if and only f is an isogeny, and when $X' = X$ this pairing is skew-symmetric if and only if f (or equivalently, L) is symmetric. Working modulo n we obtain a pairing

$$\Psi_{f,n} = \psi_{L,n} : \Lambda/n\Lambda \times \Lambda'/n\Lambda' \to \mathbf{Z}(1)/n\mathbf{Z}(1) \simeq \mu_n(\mathbf{C})$$

into the group of nth roots of unity in $\mathbf{C}$ for every positive integer n. We have canonical identifications $\Lambda/n\Lambda \simeq (1/n)\Lambda/\Lambda \simeq X[n]$ and $\Lambda'/n\Lambda' \simeq X'[n]$ with the n-torsion on X and X', so we may rephrase $\Psi_{f,n}$ as a bilinear pairing

$$(\cdot,\cdot)_{f,n} = (\cdot,\cdot)_{L,n} : X[n] \times X'[n] \to \mu_n(\mathbf{C}).$$

It is such torsion pairings (and not the $\mathbf{Z}$-lattice homology pairings) that will be rediscovered in the algebraic theory.

The are two cases of special interest: $X' = X^\vee$ with $L = P_X$ on $X \times X^\vee$, and $X' = X$ with L a symmetric correspondence on $X \times X$. In the case of the Poincaré correspondence, Example 2.4.2 shows that $(\cdot,\cdot)_{P_X,n}$ is a *perfect* pairing with respect to which the adjoint of the n-torsion map induced by some $g : X \to Y$ is induced by the dual map $g^\vee : Y^\vee \to X^\vee$. Moreover, under our initial analytic construction of $X^\vee$ using the homology lattice $\Lambda^\vee(1) = \mathrm{Hom}(\Lambda, \mathbf{Z}(1))$ that identifies $X^\vee[n]$ with $\mathrm{Hom}(X[n], \mu_n(\mathbf{C}))$, the pairing $(\cdot,\cdot)_{P_X,n}$ is the evaluation pairing between $X[n]$ and $\mathrm{Hom}(X[n], \mu_n(\mathbf{C}))$. This canonical pairing between $X[n]$ and $X^\vee[n]$ is called the *Weil n-torsion pairing*. In the case of a symmetric correspondence L on $X \times X$, or equivalently a symmetric map $f : X \to X^\vee$, the n-torsion pairings $(\cdot,\cdot)_{L,n} = (\cdot,\cdot)_{f,n}$ are skew-symmetric. All of these properties of the torsion pairings will hold (with proper formulation) in the algebraic theory.

3. Analytic aspects of polarizations and CM tori

The preceding sections provide many ways to encode information for complex tori in terms of linear algebra data. In particular, there is a good theory of duality that is related to pairings on homology lattices via Riemann forms. We saw that such pairings allow us to encode isogenies in terms of non-degenerate lattice pairings. However, whereas every finite-dimensional vector space is (non-canonically) isomorphic to its own dual, we have not given any reason to expect that $\mathrm{Hom}(X, X^\vee)$ is even nonzero for a complex torus X. More specifically, any nonzero finite-dimensional vector space admits a non-degenerate symmetric bilinear form so we can ask the analogous question: does there exist a symmetric isogeny $f : X \to X^\vee$? By Example 2.4.5, if there is a nonzero symmetric map f from X to $X^\vee$ then the line bundle $(1,f)^*(P_X)$ on X has Riemann form $2\Psi_f \neq 0$. It therefore follows from Remark 2.2.4 that if $g > 1$ then a generic complex torus of dimension g has no such nonzero maps, as it has vanishing Riemann form for all of its holomorphic line bundles.

In this section we shall focus our efforts on understanding the fundamental concept of *polarization* for a complex torus. There are several ways to view this notion, and it essentially amounts to the specification of a projective embedding of the complex torus. Those tori that admit projective embeddings are rather special within the analytic theory (if $g > 1$), and they have very nice properties. As but one striking example of a class of analytically-constructed complex tori that turn out to admit such embeddings, in Theorem 3.3.7 we shall see that CM tori necessarily admit projective embeddings; this is a property that generally fails if we allow arbitrary totally complex fields rather than CM fields in the definition of a CM torus. Moreover, we will see in §4 that the algebraic theory over $\mathbf{C}$ meets the analytic theory in precisely the class of complex tori that admit projective embeddings.

3.1. Polarizations.

Let us first review the special property of projective n-space $\mathbf{CP}^n$ in complex-analytic geometry (for $n \geq 1$). Following Grothendieck's convention, we interpret points $[a_0, \ldots, a_n]$ in $\mathbf{CP}^n$ as classifying nonzero linear functionals on $\mathbf{C}^{n+1}$ up to $\mathbf{C}^\times$-multiple (or, what comes to the same by passing to the kernel, points in $\mathbf{CP}^n$ classify hyperplanes in $\mathbf{C}^{n+1}$). Concretely, such a homogeneous $(n+1)$-tuple corresponds to the homothety class of the nonzero linear functional $\sum a_j t_j$ on $\mathbf{C}^{n+1}$, or its kernel hyperplane. Better yet (from the viewpoint of a definition that is well-suited to variation in analytic families), we may view points of $\mathbf{CP}^n$ as isomorphism classes of 1-dimensional quotients of $\mathbf{C}^{n+1}$.

Consider a complex manifold M and a holomorphic line bundle $L \to M$ endowed with an ordered $(n+1)$-tuple of global sections $s_0, \ldots, s_n$ that generate L in the sense that they never simultaneously vanish (*i.e.*, for each $m \in M$ some $s_j(m)$ is nonzero in the fiber line $L(m)$) and so they span all fibers of $L \to M$. These sections define a map of holomorphic vector bundles $\mathbf{C}^{n+1}\times M \to L$ over M given by $(a_j) \mapsto \sum a_j s_j(m)$ on m-fibers. The fibral non-vanishing condition implies that this is a surjection of vector bundles. More specifically, for each $m \in M$ we get a 1-dimensional quotient $L(m)$ of $\mathbf{C}^{n+1}$ via the quotient map $(a_j) \mapsto \sum a_j s_j(m)$. This gives a point in $\mathbf{CP}^n$ that is traditionally denoted $[s_0(m), \dots, s_n(m)]$. To be precise, for each $0 \leq j \leq n$ the open set $M_j \subseteq M$ where s_j is non-vanishing on fibers gives a map $M_j \to \mathbf{CP}^n$ into the standard chart where the jth homogeneous coordinate function is nonvanishing: over M_j we have $s_i = h_{ij}s_j$ for all i with holomorphic h_{ij} on M_j (so $h_{ii} = 1$), and $m \mapsto [h_{0j}(m), \dots, h_{nj}(m)]$ is an analytic map from M_j into the jth standard chart $U_j \subseteq \mathbf{CP}^n$. These maps agree on overlaps precisely because the s_i's are sections of a line bundle, and so we get a global analytic map $h : M \to \mathbf{CP}^n$.

This procedure can be reversed, and so it provides a universal property for $\mathbf{CP}^n$. To formulate such a property, recall that over $\mathbf{CP}^n$ there is a canonical line bundle $\mathscr{O}(1)$ endowed with an ordered $(n+1)$-tuple of global sections $(\sigma_0, \dots, \sigma_n)$ that are nowhere simultaneously vanishing. Explicitly, over the jth standard chart U_j the section σ_j is a generator and $\sigma_i = (t_i/t_j)\sigma_j$ with $[t_0, \dots, t_n]$ the standard homogeneous coordinates on $\mathbf{CP}^n$. The universal property is that if M is any complex manifold and $L \to M$ is a line bundle that is endowed with an ordered $(n+1)$-tuple of generating sections $s_0, \dots, s_n$ then there is a unique analytic map $h : M \to \mathbf{CP}^n$ and an isomorphism $h^*(\mathscr{O}(1)) \simeq L$ carrying $h^*(\sigma_j)$ to s_j for all j. (Such a bundle isomorphism is uniquely determined, so we do not need to give it a name.) Concretely, $h(m) \in \mathbf{CP}^n$ is the isomorphism class of the 1-dimensional quotient $L(m)$ of $\mathbf{C}^{n+1}$ defined by the $s_j(m)$'s. As a consequence, to give a map from a complex manifold to a projective space we must produce a line bundle with "lots" of holomorphic sections. It is an even stronger condition to ask that there be enough sections for this map to be a closed embedding into the projective space. The existence of such embeddings merits a name:

Definition 3.1.1. A line bundle L on a compact complex manifold M is *ample* if there exists a closed embedding $j : M \hookrightarrow \mathbf{CP}^n$ and an isomorphism $j^*(\mathscr{O}(1)) \simeq L^{\otimes r}$ for some n and some $r \geq 1$.

It may seem better to insist that $r = 1$, but for various reasons in the general theory it turns out that allowing $r \geq 1$ leads to a more useful notion.

Example 3.1.2. By using Veronese d-uple embeddings of a projective space into higher-dimensional projective spaces one sees that ampleness is preserved under tensor products on a fixed compact manifold M, and so if $\dim M > 0$ then the non-ampleness of the trivial line bundle implies that an ample line bundle on M has non-ample inverse. Likewise, by using Segre maps from a product of projective spaces into a higher-dimensional projective space one sees that if $L \to M$ and $L' \to M'$ are ample holomorphic line bundles on compact complex manifolds M and M' then $p_1^*(L) \otimes p_2^*(L')$ is ample on $M \times M'$ (where p_1, p_2 are the projections to the factors). It is a fundamental fact that if $f : M' \to M$ is a map between compact complex manifolds and it has finite fibers (*e.g.*, an isogeny of complex tori) then $f^*(L)$ is ample on M' for any ample line bundle L on M.

In the case of complex tori, ampleness is a concrete property: under what conditions on an Appell–Humbert datum (H, α) may we infer that $L(H, \alpha)$ has a high-dimensional space of global sections? There is a characterization of ampleness in terms of the Riemann form, as follows. Let $X = F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda)/\Lambda$ be a complex torus and let $L \to X$ be a holomorphic line bundle. The decomposition $F \oplus \overline{F} = \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ gives rise to a projection map $\mathbf{C} \otimes_{\mathbf{Z}} \Lambda \twoheadrightarrow F$ that carries Λ to a lattice in F and so restricts to an $\mathbf{R}$-linear isomorphism $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda \simeq F$. This puts a complex structure on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ (that is *conjugate* to the complex structure induced by $V = \mathrm{T}_e(X)$). The Riemann form $\Psi_L : \Lambda \times \Lambda \to \mathbf{Z}(1)$ is skew-symmetric, so its $\mathbf{R}$-linear extension $\Psi_{L,\mathbf{R}}$ is identified with a skew-symmetric form $F \times F \to \mathbf{R}(1)$. By unwinding how Ψ_L is defined, one finds upon studying the interaction of $\Psi_{L,\mathbf{R}}$ with the complex structure on F that $\Psi_{L,\mathbf{R}}$ is the imaginary component of a (unique) Hermitian form on F.

Theorem 3.1.3 (Lefschetz)**.** *The line bundle L is ample if and only if the unique Hermitian form on $F \simeq \mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ with imaginary component $\Psi_{L,\mathbf{R}}$ is positive-definite. In such cases Ψ_L is non-degenerate. In particular, whether or not L is ample depends only on the discrete invariant Ψ_L of L.*

This theorem explains why ampleness is considered to be a positivity condition. Beware that if we use the conjugate complex structure arising from V then the resulting Hermitian form with imaginary component

$\Psi_{L,\mathbf{R}}$ has real part with opposite sign to the one on F. Hence, to formulate the ampleness criterion in terms of a positive-definiteness property (rather than a negative-definiteness property) it is essential to use F rather than V. (Hence, this gives yet another illustration of why the viewpoint of Hodge structures is better than that of exponential uniformization when studying complex tori.)

Proof. Up to unwinding many changes in normalizations (and fixing a miraculously harmless sign error), this is [Mum, pp. 30–33]. ■

Example 3.1.4. Let C be a compact connected Riemann surface with genus $g > 0$, and let $X = \mathrm{Jac}(C)$ be its Jacobian. Consider the isomorphism $\varphi_C : X \simeq X^\vee$ built in Example 2.1.3; this rests on the map $\theta_C : \lambda \mapsto (\cdot).\lambda$ from $\Lambda = \mathrm{H}_1(C, \mathbf{Z})$ to $\Lambda^\vee(1)$. By Example 2.4.2 the map φ_C classifies a correspondence L over $X \times X$ whose Riemann form is the self-pairing

$$\Psi_C : \mathrm{H}_1(C, \mathbf{Z}) \times \mathrm{H}_1(C, \mathbf{Z}) \to \mathbf{Z}(1)$$

given by the intersection form $\Psi_C(\lambda, \lambda') = (\theta_C(\lambda'))(\lambda) = \lambda.\lambda'$. By skew-symmetry of Ψ_C, the isomorphism φ_C is symmetric. Moreover, Example 2.4.5 implies that this intersection form must also be the Riemann form of the line bundle $(1, \varphi_C)^*(P_X)$ on X.

A direct calculation shows the $\mathbf{R}$-scalar extension $\Psi_{C,\mathbf{R}}$ is the imaginary component of a positive-definite Hermitian form on $\mathrm{H}_1(C, \mathbf{R})$ when given the complex structure from F (conjugate to that of V), so by Lefschetz' theorem $(1, \varphi_C)^*(P_X)$ is ample on the Jacobian X. The line bundle $(1, -\varphi_C)^*(P_X)$ turns out to be the inverse of $(1, \varphi_C)^*(P_X)$, and so it is non-ample. In the special case that $C = E$ is an elliptic curve, so $X = E$ via Example 1.1.11 with x_0 equal to the identity, φ_C is equal to the map $\phi_{\mathscr{O}_E(e)}$ as in Example 2.3.7 and so the pullback $(1, \varphi_E)^*(P_E)$ is $\mathscr{O}_E(2e)$, which is indeed ample.

The map φ_C in the preceding example satisfies the conditions in the following fundamental definition:

Definition 3.1.5. A *polarization* on a complex torus X is a symmetric map $\phi : X \to X^\vee$ for which the line bundle $(1, \phi)^*(P_X)$ on X is ample.

Let $\phi : X \to X^\vee$ be a polarization. Since the line bundle $(1, \phi)^*(P_X)$ has Riemann form $2\Psi_\phi$ (by Example 2.4.5) and by ampleness this Riemann form must be non-degenerate, it follows that ϕ must be an isogeny. The degree $\deg(\phi)$ is called the *degree* of the polarization. This degree is always a square because $\ker(\phi)$ is the cokernel of the map $\mathrm{H}_1(\phi) : \Lambda \to \Lambda^\vee(1)$ that is associated to a non-degenerate skew-symmetric form Ψ_ϕ over $\mathbf{Z}$ (and so the order of the cokernel is the square of a Pfaffian). We say ϕ is a *principal polarization* when $\deg(\phi) = 1$. If we consider a symmetric isogeny ϕ to be analogous to a non-degenerate symmetric bilinear forms B (with the analogy made more precise by means of formation of the Riemann form Ψ_ϕ) then the formation of $(1, \phi)^*(P_X) = \Delta_X^*(1_X \times f)^*(P_X)$ is analogous to the formation of the associated quadratic form $w \mapsto B(w, w)$, whence the notion of polarization is roughly analogous to the concept of a positive-definite quadratic form.

Example 3.1.6. Let L be a line bundle on a complex torus X and consider the symmetric map $\phi_L : X \to X^\vee$ as in Example 2.3.7. What can we say about the line bundle $(1, \phi_L)^*(P_X)$ on X? Its Riemann form is $2\Psi_{\phi_L}$, and by Example 2.4.6 this equals $2\Psi_L$. Since the Riemann form determines ampleness via a positive-definiteness property in Lefschetz' theorem, and multiplication by 2 on a quadratic form over $\mathbf{R}$ does not affect its signature, we conclude that ϕ_L is a polarization if and only if L is ample. In general, an inspection of Appell–Humbert data shows that $n\phi_L = \phi_{L^{\otimes n}}$ for every integer n, so $n\phi_L$ has Riemann form $\Psi_{L^{\otimes n}}$ and with $n = -1$ we conclude that $-\phi_L$ is never a polarization when L is ample (and $X \neq 0$). For example, if $X = E$ is an elliptic curve then the map $\phi_{\mathscr{O}_E(e)}$ that is the negative of the classical autoduality is a principal polarization (and so the classical autoduality is not a polarization).

This procedure is exhaustive in the sense that it gives *all* polarizations (even though the definition of polarization made no mention of the ϕ_L-construction!). The reason for this is that *any* symmetric map of tori $f : X \to X^\vee$ has the form ϕ_L for some line bundle L on X. Indeed, the Riemann form Ψ_f is skew-symmetric on Λ by symmetry of f (Remark 2.3.8), and the symmetric correspondence $N = (1 \times f)^*(P_X)$ on $X \times X$ has Riemann form $\psi_N = \Psi_f$, so (by definition of ψ_N) the $\mathbf{C}$-linear extension of Ψ_f makes the $\mathbf{C}$-subspace $F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ be isotropic for the $\mathbf{C}$-linear extension of ψ_N. It therefore follows from the dimension-counting preceding Remark 2.2.4 that there exists an Appell–Humbert datum (H, α) for X such that $\Psi_H = \psi_N$, so $L = L(H, \alpha)$ has Riemann form ψ_N. But the symmetric correspondence $\wedge(L)$ satisfies $\psi_{\wedge(L)} = \Psi_L$ by Example 2.4.6, so $\psi_{\wedge(L)} = \psi_N$. Since a symmetric correspondence is uniquely determined by its Riemann form (Theorem 2.4.1), this gives an isomorphism of correspondences

$$(1 \times f)^*(P_X) = N \simeq \wedge(L) = (1 \times \phi_L)^*(P_X),$$

so by the universal property of P_X we conclude $f = \phi_L$ as desired.

Remark 3.1.7. In many references, such as [M1] and [GIT], one sees the concept of polarization *defined* as a map of the form ϕ_L for an ample L. This approach gives rise to annoying questions concerning sign conventions in the definition of ϕ_L. By using the definition we have given (which I learned from deJong) and then deducing *a posteriori* that it coincides with the "traditional" definition in terms of Mumford correspondences, we can be certain that it is our definition and not one resting on $-\phi_L$'s that is the correct one to use.

Example 3.1.8. In the case of an elliptic curve (E, e), we can make the polarizations rather explicit: for each positive integer d, the map $\phi_{\mathscr{O}_E(e)} \circ [d]_E : E \to E^\vee$ is the unique polarization of degree d^2. To prove this, first note that by Example 3.1.6 polarizations $\phi : E \to E^\vee$ are precisely maps ϕ_L for an ample line bundle L on E, and ϕ_L is uniquely determined by its Riemann form $\Psi_{\phi_L} = \Psi_L$ that is the unique discrete invariant of L. The only discrete invariant of a line bundle on an elliptic curve is its degree (Example 2.2.6), so if L has degree d then

$$\phi = \phi_{\mathscr{O}_E(de)} = \phi_{\mathscr{O}_E(e)^{\otimes d}} = [d]_{E^\vee} \circ \phi_{\mathscr{O}_E(e)} = \phi_{\mathscr{O}_E(e)} \circ [d]_E.$$

Since L is ample, we must have $d > 0$.

The Riemann form for $\phi_{\mathscr{O}_E(de)}$ is d times the Riemann form for $\phi_{\mathscr{O}_E(e)}$, so to describe the Riemann form of the unique polarization of degree d^2 we just have to describe it in the case $d = 1$. By Example 3.1.4, the principal polarization for E has Riemann form on $\mathrm{H}_1(E, \mathbf{Z})$ given by the $\mathbf{Z}(1)$-valued intersection form. In the classical terminology with $E = \mathbf{C}/(\mathbf{Z}\tau \oplus \mathbf{Z})$ and homology basis $\{[\tau], [1]\}$ we have $[1].[\tau] = 2\pi i_\tau$ with $i_\tau = \sqrt{-1} \in \mathbf{C} - \mathbf{R}$ lying in the connected component of τ in $\mathbf{C} - \mathbf{R}$. In terms of the ordered homology basis $\{[\tau], [1]\}$ (note the ordering!), the matrix for the unique polarization of degree d^2 is therefore

$$2\pi i_\tau \cdot \begin{pmatrix} 0 & -d \\ d & 0 \end{pmatrix}$$

for $d > 0$.

If t denotes the standard coordinate on $\mathbf{C}$ then the integration map π from $\mathrm{H}_1(E, \mathbf{C})$ onto the dual space of the global holomorphic 1-forms carries $[1]$ and $[\tau]$ to the functionals that send $\mathrm{d}t$ to 1 and τ respectively, so the nonzero cycle $c = \tau \otimes [1] - [\tau] \in \mathbf{C} \otimes \mathrm{H}_1(E, \mathbf{Z})$ spans the 1-dimensional kernel $F = \ker \pi$. The Hermitian form on F associated to the Riemann form (*i.e.*, the intersection form) on $\mathrm{H}_1(E, \mathbf{Z})$ via Lefschetz' theorem works out to be precisely the $\mathbf{C}$-bilinear intersection pairing $F \times \overline{F} \to \mathbf{C}$ induced by the intersection form on $\mathrm{H}_1(E, \mathbf{C})$, and direct calculation shows $c.\overline{c} = 2\pi(\tau - \overline{\tau})/i_\tau$. This is indeed positive, as we knew it had to be, since i_τ and $\tau - \overline{\tau}$ lie in the same component of the nonzero part of the imaginary axis.

Example 3.1.9. Let us push the preceding example a bit further. Using the canonical principal polarization $\phi_{\mathscr{O}_E(e)} : E \simeq E^\vee$ for an elliptic curve, we obtain canonical self-pairings

$$(\cdot, \cdot)_n : E[n] \times E[n] \simeq E[n] \times E^\vee[n] \to \mu_n(\mathbf{C})$$

that are perfect and skew-symmetric (since $\phi_{\mathscr{O}_E(e)}$ is a symmetric isomorphism). This is obtained by mod-n reduction of the skew-symmetric Riemann form $\Psi = \Psi_{\phi_{\mathscr{O}_E(e)}}$ that we computed in the preceding example. For any $\lambda, \lambda' \in \mathrm{H}_1(E, \mathbf{Z})$, we have $(\lambda/n, \lambda'/n) = e^{\Psi(\lambda, \lambda')/n}$. Thus, in terms of an analytic model $E = \mathbf{C}/(\mathbf{Z}\tau \oplus \mathbf{Z})$ we have $(1/n, \tau/n)_n = e^{2\pi i_\tau/n}$. In terms of the analytic Tate model $E = \mathbf{C}^\times/q^{\mathbf{Z}}$ with $0 < |q| < 1$, we have $(\zeta, q^{1/n})_n = \zeta$ for any $\zeta \in \mu_n(\mathbf{C}) = \mathbf{C}^\times[n]$ and any nth root of q (the point being that $|e^{2\pi i_\tau \tau}| < 1$). There is a lot of confusion in the literature concerning sign conventions for Weil self-pairings on torsion of elliptic curves. However, the way we have developed the theory avoids *ad hoc* questions concerning sign conventions in the definitions (and so shows that from the viewpoint of other approaches to the theory of Weil self-pairings there is a unique "correct" choice of such conventions).

3.2. **Siegel's construction.** The purpose of Siegel's construction is to describe *all* polarized complex tori; in the 1-dimensional case it recovers Example 3.1.8. Let (X, ϕ) be a polarized complex torus with dimension $g > 0$, so the Riemann form Ψ_ϕ is a non-degenerate symplectic form on $\Lambda = \mathrm{H}_1(X, \mathbf{Z})$ with values in $\mathbf{Z}(1)$. As such, it has discrete invariants given by elementary divisors $1 \leq d_1 | d_2 | \dots | d_g$. That is, by the structure theorem for sympectic spaces over a principal ideal domain, we can find a basis $\{\sigma_1, \dots, \sigma_{2g}\}$ of Λ with respect to which Ψ_ϕ is given by the matrix

$$\Psi_{\mathbf{d}} = 2\pi i \begin{pmatrix} 0 & -\mathbf{d} \\ \mathbf{d} & 0 \end{pmatrix}$$

where $\mathbf{d} = \mathrm{diag}(d_1, \dots, d_g)$ and $i = \sqrt{-1}$ is a fixed choice. We shall write $\Psi_{\mathbf{d},\mathbf{R}}$ to denote the $\mathbf{R}$-scalar extension of $\Psi_{\mathbf{d}}$. In terms of this fixed discrete data we seek to describe possibilities for X and ϕ.

Note first of all that $\{\sigma_{g+1}, \dots, \sigma_{2g}\}$ spans a maximal $\Psi_{\mathbf{d}}$-isotropic subspace of Λ over $\mathbf{Z}$, and hence it spans a maximal $\Psi_{\mathbf{d},\mathbf{R}}$-isotropic subspace of $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ over $\mathbf{R}$. We claim that this forces the g vectors σ_{g+j} to be a $\mathbf{C}$-basis of the g-dimensional V, or equivalently of the g-dimensional F (viewed as conjugate complex structures on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$). To prove this, it suffices to prove the vanishing of any $w' \in F$ that is orthogonal of the $\mathbf{C}$-span W of the σ_{g+j}'s in F when using the (positive-)definite Hermitian form on F as in Lefschetz' theorem. If $w' \neq 0$ then it must lie in $F - W$, yet it is also orthogonal to the $\mathbf{R}$-span of the σ_{g+j}'s with respect to the imaginary component $\Psi_{\mathbf{d},\mathbf{R}}$ of the Hermitian form. By maximal isotropicity it follows that w' is contained in the $\mathbf{R}$-span of the σ_{g+j}'s, a contradiction since $w \notin W$. This settles the $\mathbf{C}$-basis claim.

Arguing as in the derivation of (1.1.2) identifies V with $\mathbf{C}^g$ by using the $\mathbf{C}$-basis of σ_{g+j}'s, and so $\Lambda \subseteq V$ is identified with a lattice of the form $Z'(\mathbf{Z}^g) + \mathbf{Z}^g$ for some $Z' \in \mathrm{Mat}_{g \times g}(\mathbf{C})$ with invertible imaginary component. The condition that the skew-symmetric pairing $\Psi_{\mathbf{d}}$ on this lattice arises as a Riemann form Ψ_ϕ of a (symmetric) map of tori ϕ from $X = \mathbf{C}^g/(Z'(\mathbf{Z}^g) + \mathbf{Z}^g)$ to $X^\vee$ is precisely the condition that F is isotropic with respect to the $\mathbf{C}$-linear extension $\Psi_{\mathbf{d},\mathbf{C}}$. After some computation, this is precisely the condition that the matrix $Z'\mathbf{d}^{-1}$ is symmetric. With some more computation, the ampleness aspect via Lefschetz' theorem says exactly that with respect to our choice of $2\pi i$ the imaginary part of $Z'\mathbf{d}^{-1}$ is positive-definite (and hence invertible). Since $Z'_{\mathrm{im}} = (Z'\mathbf{d}^{-1})_{\mathrm{im}}\mathbf{d}$ with $\mathbf{d}$ invertible, upon renaming $Z'\mathbf{d}^{-1}$ as Z this proves:

Theorem 3.2.1 (Siegel)**.** *Fix a choice of $i = \sqrt{-1}$ and $1 \leq d_1 | d_2 | \dots | d_g$. Let $\mathbf{d} = \mathrm{diag}(d_1, \dots, d_g)$. Within the vector space of symmetric $g \times g$ matrices over $\mathbf{C}$, let $\mathfrak{h}_{g,i}$ be the open subset of those Z for which the symmetric real matrix $(Z - \overline{Z})/2\pi i$ is positive-definite.*

On the homology lattice of the complex torus $X_{Z,\mathbf{d}} = \mathbf{C}^g/(Z\mathbf{d}\mathbf{Z}^g + \mathbf{Z}^g)$, the skew-symmetric form having matrix

$$2\pi i \begin{pmatrix} 0 & -\mathbf{d} \\ \mathbf{d} & 0 \end{pmatrix}$$

with respect to the $\mathbf{Z}$-basis $\{Z\mathbf{d}(e_1), \dots, Z\mathbf{d}(e_g), e_1, \dots, e_g\}$ is the Riemann form $\Psi_{\phi_{Z,\mathbf{d}}}$ of a (unique) polarization $\phi_{Z,\mathbf{d}} : X_{Z,\mathbf{d}} \to X_{Z,\mathbf{d}}^\vee$. Moreover, every polarized complex torus (X, ϕ) for which the Riemann form Ψ_ϕ has invariant factors $\{d_1, \dots, d_g\}$ arises in this manner.

The polarized complex tori $(X_{Z,\mathbf{d}}, \phi_{Z,\mathbf{d}})$ in this theorem are endowed with more structure than a polarization with specified invariant factors. Indeed, these $X_{Z,\mathbf{d}}$'s are endowed with a homology basis with respect to which the Riemann form of the polarization is in a standard form. The $X_{Z,\mathbf{d}}$'s can also be put into a global analytic family that provides an analytic model for a Siegel modular variety, as we shall see in §7, where we will also study this construction from an arithmetic point of view.

The manifold $\mathfrak{h}_{g,i}$ is the *Siegel half-space* associated to (g, i). In the case $g = 1$ it is the connected component of $\mathbf{C} - \mathbf{R}$ containing the choice of i. In general it is a connected open cone in the space of symmetric $g \times g$ matrices over $\mathbf{C}$. (The topological structure of $\mathfrak{h}_{g,i}$ will be addressed in §7.1.)

Observe that to each $Z \in \mathfrak{h}_{g,i}$ we have associated a polarized complex torus $X_{Z,\mathbf{d}}$ (with polarization having Riemann form whose invariant factors are fixed), and we claim that the locus of Z's for which $X_{Z,\mathbf{d}}$ is a CM torus is dense in the analytic topology on $\mathfrak{h}_{g,i}$. The most classical case is $g = 1$, for which the fiber over $Z = \tau$ in a connected component $\mathfrak{h}$ of $\mathbf{C} - \mathbf{R}$ is $\mathbf{C}/(d\tau\mathbf{Z} \oplus \mathbf{Z})$ and hence is a CM elliptic curve if and only if τ is imaginary quadratic, in which case there is CM by an order in the imaginary quadratic field $K = \mathbf{Q}(\tau) \subseteq \mathbf{C}$. The imaginary quadratic points in $\mathfrak{h}$ are obviously dense, even if we fix the CM field. A good way to conceptually understand the density is to recall that there is a topological (and even real-analytic) isomorphism $\mathfrak{h} \simeq \mathrm{SL}_2(\mathbf{R})/\mathrm{SO}_2(\mathbf{R})$ under which the standard linear fractional action by $\mathrm{SL}_2(\mathbf{R})$ on $\mathfrak{h}$ is identified with left translation by $\mathrm{SL}_2(\mathbf{R})$ on the coset space, so by density of $\mathrm{SL}_2(\mathbf{Q})$ in $\mathrm{SL}_2(\mathbf{R})$ any $\mathrm{SL}_2(\mathbf{Q})$-orbit in $\mathfrak{h}$ is topologically dense. The relevance of this fact is that $\mathrm{SL}_2(\mathbf{Q})$-orbit of a CM point consists entirely of CM points with the same CM field. The points of such an orbit have varying CM orders in the CM field. (The infinite locus of CM points with a *fixed* CM order has *finite* image in any modular curve quotient of $\mathfrak{h}$, such as in the sense of Example 6.1.9, because the set of isomorphism classes of elliptic curves with a fixed CM order can be shown to be a torsor for the finite class group of the order.)

In the case $g > 1$, to get a similar density statement for CM tori with a fixed CM field K of degree $2g$ over $\mathbf{Q}$ and a fixed CM type Φ (but varying CM order within K) it turns out to be rather difficult to analyze the situation by a direct inspection of lattices of the form $Z\mathbf{d}(\mathbf{Z}^g) + \mathbf{Z}^g$ inside of $\mathbf{C}^g$. The difficulty is due to the fact that the Siegel description of polarized tori is not well-adapted to the way we have analytically described CM tori (as $(\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi/\mathfrak{a}$) and the Riemann form of a K-linear polarization on such tori (via the trace form in the case $\mathfrak{a} = \mathscr{O}_K$). We will return to this problem in §7.1, using an orbit argument with the symplectic group Sp_{2g} replacing SL_2 as used above.

3.3. **Endomorphisms and polarizations.** The presence of a polarization is a powerful tool for "digging holes" into a complex torus. For example, the category of polarizable complex tori has no interesting exact sequences, due to:

Theorem 3.3.1 (Poincaré reducibility theorem)**.** *If X is a polarizable complex torus and $X_0 \subseteq X$ is a complex subtorus then X_0 is polarizable and there exists an isogeny-complement: a complex subtorus $X_0' \subseteq X$ such that the map $X_0 \times X_0' \to X$ defined by addition is an isogeny.*

Proof. Let $\phi = \phi_L : X \to X^\vee$ be a polarization for an ample line bundle L, and let $j : X_0 \to X$ be the inclusion. By Example 2.3.8 the composite map $j^\vee \circ \phi \circ j$ is equal to ϕ_{j^*L}, and $j^*(L)$ an ample line bundle on X_0. Thus, $j^\vee \circ \phi \circ j$ is an isogeny from X_0 to $X_0^\vee$. In particular, X_0 has finite intersection with the identity component $X_0' = (\ker(j^\vee \circ \phi))^0$ that is a subtorus of X. The dual $j^\vee$ is surjective since j is an embedding (this can be checked on the level of rational homology), so certainly $\dim X_0' \geq \dim(X) - \dim(X_0)$. Since the map $X_0 \times X_0' \to X$ has finite kernel, the dimension inequality is forced to be an equality and so this addition map is an isogeny. ■

It follows from Poincaré's theorem that a polarizable complex torus X contains no nonzero proper complex subtori if and only if it is not isogenous to a product of two nonzero complex tori, so the non-existence of nonzero proper complex subtori is *equivalent* to $\mathrm{End}^0(X)$ being a division algebra (that is, all nonzero self-maps of X are surjective and hence isogenies). Such complex tori are called *simple*. We get the important corollary:

Corollary 3.3.2. *Every polarizable complex torus X is isogenous to a product $\prod X_j^{e_j}$ with the X_j's pair-wise non-isogeneous simple polarizable complex tori. The X_j's and e_j's are necessarily unique, and $D_j = \mathrm{End}^0(X_j)$ is a finite-dimensional division algebra over $\mathbf{Q}$. In particular, the category of such tori is semisimple, all objects have finite length, and $\mathrm{End}^0(X) = \prod_j \mathrm{Mat}_{e_j \times e_j}(D_j)$ is a semisimple $\mathbf{Q}$-algebra.*

Remark 3.3.3. We call these X_j's the *simple isogeny factors* of X, and $\prod X_j^{e_j}$ is the *semisimple decomposition* of X.

An *automorphism* ξ of a polarized complex torus (X, ϕ) is an automorphism ξ of X such that the Riemann form

$$\Psi_\phi : \Lambda \times \Lambda \to \mathbf{Z}(1)$$

is invariant under applying ξ to both factors, or equivalently $\xi^\dagger \xi = 1$ with $f^\dagger = \phi^{-1} \circ f^\vee \circ \phi$ on $\mathrm{End}^0(X)$. The automorphism group of (X, ϕ) is contained in the discrete subgroup $\mathrm{GL}(\Lambda) \subseteq \mathrm{GL}(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda)$, and in view of

Lefschetz' theorem it is contained in the compact orthogonal group for a certain positive-definite quadratic form on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$. This yields an important finiteness property of polarized complex tori:

Theorem 3.3.4. *The automorphism group of a polarized complex torus is finite.*

Remark 3.3.5. An analogous finiteness result holds in the algebraic theory, but it is considerably harder to prove.

Now we turn to the interaction of polarizations and CM-structures. This is the first place where CM fields (rather than general totally complex fields) play a crucial role in the theory of CM tori. Let us start with a lemma.

Lemma 3.3.6. *Let K be a CM field and (X, i) a CM torus with CM field K. Endow the dual torus $X^\vee$ with a structure of CM torus via the embedding $i^\vee : K \hookrightarrow \mathrm{End}^0(X^\vee)$ defined by $i^\vee(c) = i(\overline{c})^\vee$ for all $c \in K$, where $c \mapsto \overline{c}$ is the complex conjugation on K.*

The CM torus $(X^\vee, i^\vee)$ has the same CM type as (X, i).

Proof. Let Φ be a CM type for X. By the definition of $i^\vee$, it suffices to show that letting K act on $X^\vee$ through duality yields the conjugate CM type $\overline{\Phi}$. The double-coset uniformization for $X^\vee$ shows that its tangent space is $\mathbf{C}$-dual to a direct summand of $\mathbf{C} \otimes_{\mathbf{Z}} \mathrm{H}_1(X, \mathbf{Z})$ that is "complementary" to the quotient $\mathrm{T}_e(X)$. Since linear duality does not change eigencharacters and Φ and $\overline{\Phi}$ have disjoint union equal to $\mathrm{Hom}(K, \mathbf{C})$, it suffices to prove that the set of eigencharacters for the K-action on $\mathbf{C} \otimes_{\mathbf{Q}} \mathrm{H}_1(X, \mathbf{Q})$ is precisely the set of embeddings of K into $\mathbf{C}$ without repetition. The K-action is through a structure of 1-dimensional K-vector space on $\mathrm{H}_1(X, \mathbf{Q})$, and so our problem is solved by the trivial observation

$$\mathbf{C} \otimes_{\mathbf{Q}} K \simeq \prod_{\varphi : K \to \mathbf{C}} \mathbf{C}$$

via $a \otimes x \mapsto (a\varphi(x))_\varphi$ with φ running through the set of *all* embeddings of K into $\mathbf{C}$. ■

Since the CM type determines the K-linear isogeny class of a CM torus with CM by K, by following the procedure in the lemma we know that (X, i) and $(X^\vee, i^\vee)$ are K-linearly isogenous for any CM torus (X, i) with CM field K. In fact, we can find such an isogeny with better properties:

Theorem 3.3.7. *Let (X, i) be a CM torus with CM field K. There exists a K-linear polarization $X \to X^\vee$.*

Proof. By Example 2.3.8, if $f : X_1 \to X_2$ is an isogeny of complex tori and $\phi : X_2 \to X_2^\vee$ is a polarization of the form ϕ_{L_2} for an ample line bundle on L_2 then $f^\vee \circ \phi \circ f = \phi_{f^*(L)}$ and so this is a polarization on X_1 because $f^*(L)$ is ample (due f being an isogeny; see Example 3.1.2). Since all possible pairs (X, i) for a fixed K and fixed CM type are K-linearly isogenous, it therefore follows that we only need to treat a *single* such pair for each K and each CM type Φ for K. We choose $X = (\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi / \mathscr{O}_K$ with its canonical $\mathscr{O}_K$-action (via Φ).

Since K/K_0 is a totally imaginary quadratic extension of a totally real field, there exists $c \in K$ such that $c^2 = c_0 \in K_0$ with $K = K_0(c)$ and necessarily $\varphi(c) \in \mathbf{C}^\times$ lies on the imaginary axis for all $\varphi \in \Phi$ since K_0 is totally real and c_0 must be totally negative. These φ's exhaust the set of archimedean places of K *without* repetition, so we can use weak approximation to improve the choice of c in order that all $\varphi(c)$ lie in the same component of the punctured imaginary axis $\mathbf{R}(1) - \{0\}$. In other words, for $i = \sqrt{-1}$ in this component, $\varphi(c)/i > 0$ for all $\varphi \in \Phi$. By scaling, we can take $c \in \mathscr{O}_K$. Letting $z \mapsto \overline{z}$ be the CM involution on K, $\Psi(z, z') = 2\pi i \mathrm{Tr}_{K/\mathbf{Q}}(cz\overline{z'})$ is a non-degenerate $\mathbf{Z}(1)$-valued bilinear form on $\mathscr{O}_K$. It is skew-symmetric because $\overline{c} = -c$ and the trace is invariant under the CM involution. (Note also that Ψ is independent of the choice of i.) Direct calculation (see [Mum, pp. 212-213] for the details, using a different normalization) shows that this arises as the imaginary component of a Hermitian form on $(\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi$, thereby providing the required isotropicity condition on $\Psi_{\mathbf{C}}$ to ensure that $\Psi = \Psi_\phi$ is the Riemann form of a map $\phi : X \to X^\vee$. The skew-symmetry and non-degeneracy of Ψ imply that ϕ is a symmetric isogeny. With a bit of calculation, the positivity of the $\varphi(c)/i$'s ensures the required positive-definiteness in Lefschetz' theorem to confirm that ϕ is a polarization.

The K-linearity of ϕ is equivalent to the property that the K-action on the rational homology for X has its Ψ_ϕ-adjoint action given by the CM involution. That is, we need to show $\Psi(az, z') = \Psi(z, \overline{a}z')$ for $a \in K$. This is obvious from the definition of Ψ. ■

We can restate the K-linearity property of a polarization $X \to X^\vee$ in another way, by using the following important concept:

Definition 3.3.8. If X is a complex torus and $\phi : X \to X^\vee$ is a polarization then the involution $f^\dagger = \phi^{-1} \circ f^\vee \circ \phi$ on $\mathrm{End}^0(X)$ is the *Rosati involution.*

Remark 3.3.9. The map $f \mapsto f^\dagger$ is an anti-homomorphism of the endomorphism algebra because dualizing is additive and anti-multiplicative, and it is an involution because ϕ is symmetric. Also, it depends on the choice of ϕ in general (except if $\dim X = 1$, in which case all polarizations are $\mathbf{Q}^\times_{>0}$-multiples of each other).

Since $f^\vee$ computes the adjoint to f with respect to the Weil pairing on $\mathbf{Z}$-homology lattices, the $\mathbf{Q}$-homology map $\mathrm{H}_1(f^\dagger)$ induced by the Rosati involution is the $(\Psi_\phi)_{\mathbf{Q}}$-adjoint of $\mathrm{H}_1(f)$. In other words, for Weil self-pairings induced by a polarization, the Rosati involution computes the adjoint. In the case of a K*-linear* polarization as in Theorem 3.3.7, this involution takes on a very concrete form for self-maps of X arising from the CM field K: it restricts to the CM involution on $K \subseteq \mathrm{End}^0(X)$.

Let us conclude our treatment of the basic analytic aspects of CM tori by explaining how a CM torus may be built from simpler pieces. First consider a generalization of Example 1.3.13, as follows. Let K'/K be an extension of CM fields with $n = [K' : K]$ and let (X, i) be a complex torus with CM by K and CM type Φ. The torus X^n may be endowed with a CM structure by K' via the embedding

$$K' \hookrightarrow \mathrm{Mat}_{n\times n}(K) \subseteq \mathrm{End}^0(X^n)$$

defined by a K-basis of K' and the K-linear action of K' on itself through multiplication. It is not hard to check that this identifies the tangent space of X^n with $K' \otimes_K \mathrm{T}_e(X)$ as a $K' \otimes_{\mathbf{Q}} \mathbf{C}$-module, so using the ring decomposition

$$K' \otimes_{\mathbf{Q}} \mathbf{C} \simeq K' \otimes_K (K \otimes_{\mathbf{Q}} \mathbf{C}) \simeq \prod_{\varphi:K\hookrightarrow\mathbf{C}} K' \otimes_{K,\varphi} \mathbf{C}$$

shows that the CM type Φ' on X^n consists of those embeddings $K' \hookrightarrow \mathbf{C}$ that lift the embeddings $K \hookrightarrow \mathbf{C}$ coming from Φ. This procedure can be reversed in the following sense (whose proof rests on knowledge of the possible endomorphism algebras of simple polarizable complex tori [Mum, §21]):

Theorem 3.3.10. *Let (X, i) be a CM complex torus with CM type Φ and CM field K. The semisimple decomposition of X is a power X_0^n of a single simple complex torus and $\mathrm{End}^0(X_0)$ is a CM field, so $\mathrm{End}^0(X_0)$ is the center of $\mathrm{End}^0(X)$ and the resulting inclusion $\mathrm{End}^0(X_0) \subseteq i(K)$ into the maximal commutative subfield $i(K) \subseteq \mathrm{End}^0(X)$ gives a canonical identification $i_0 : \mathrm{End}^0(X_0) \simeq K_0$ onto a CM subfield K_0 in K. The CM type Φ is the set of embeddings of K into $\mathbf{C}$ over the embeddings of K_0 into $\mathbf{C}$ determined by i_0.*

The complex torus X is non-simple if and only if there exists a proper CM subfield $K_1 \subseteq K$ and a CM type Φ_1 on K_1 such that Φ is the set of embeddings of K into $\mathbf{C}$ over the embeddings of K_1 into $\mathbf{C}$ given by Φ_1.

If we change the embedding $i : K \hookrightarrow \mathrm{End}^0(X)$ then the subfield $K_0 \subseteq K$ in Theorem 3.3.10 will usually change. It is a striking fact (proved in the same analysis of endomorphism algebras that feeds into the proof of Theorem 3.3.10) that if X is a simple complex torus for which the division algebra $\mathrm{End}^0(X)$ is a commutative field with degree $2\dim X$ then this totally complex field must be a CM field.

3.4. **Tate modules.** We have seen that the integral homology lattice is an extremely useful tool in the analytic study of complex tori. Although this lattice cannot be reconstructed algebraically, a mild variant on it can be. This so-called "Tate module" construction turns out to be as vital in the algebraic theory as the homology lattice is in the analytic theory. To conclude our discussion of the analytic theory, we shall explain the Tate module from an analytic point of view and translate some of our earlier results on lattice pairings into the language of Tate modules. This will motivate several important constructions in the algebraic theory.

Let X be a complex torus. The lattice $\Lambda = \mathrm{H}_1(X, \mathbf{Z})$ is related to X via the analytic exponential uniformization. However, its quotients $\Lambda/n\Lambda$ can be rediscovered without any analytic machinery: as we saw in Example 1.2.4, there is an isomorphism $\Lambda/n\Lambda \simeq (1/n)\Lambda/\Lambda \simeq X[n]$ onto the n-torsion in X for any nonzero integer n. Moreover, for any nonzero integer n' we have a commutative diagram:

$$\begin{array}{ccccc} \Lambda/nn'\Lambda & \xrightarrow{\simeq} & (1/nn')\Lambda/\Lambda & \xrightarrow{\simeq} & X[nn'] \\ \downarrow & & \downarrow{\scriptstyle n'} & & \downarrow{\scriptstyle n'} \\ \Lambda/n\Lambda & \xrightarrow[\simeq]{} & (1/n)\Lambda/\Lambda & \xrightarrow[\simeq]{} & X[n] \end{array}$$

in which the left column is reduction and the middle and right columns are induced by multiplication by n'. Taking n to run through powers of a prime ℓ, we may pass to the inverse limit to get an isomorphism

$$\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \Lambda \simeq \varprojlim \Lambda/\ell^m\Lambda \simeq \varprojlim X[\ell^m] = \mathrm{T}_\ell(X)$$

where the final equality is a definition of notation and the inverse system $\{X[\ell^m]\}$ has its transition maps given by multiplication by ℓ from each stage onto the next. We call $\mathrm{T}_\ell(X)$ the *ℓ-adic Tate module* of X; it is a finite free $\mathbf{Z}_\ell$-module of rank $2\dim(X)$, and it is functorial in X. Note that we can recover torsion-levels from the Tate module: the natural map $\mathrm{T}_\ell(X)/\ell^m\mathrm{T}_\ell(X) \to X[\ell^m]$ is an isomorphism for all $m \geq 1$. The associated $\mathbf{Q}_\ell$-vector space $\mathbf{Q}_\ell \otimes_{\mathbf{Z}_\ell} \mathrm{T}_\ell(X)$ is denoted $\mathrm{V}_\ell(X)$.

The natural maps $\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathrm{H}_1(X, \mathbf{Z}) \to \mathrm{H}_1(X, \mathbf{Z}_\ell)$ and $\mathbf{Q}_\ell \otimes_{\mathbf{Z}} \mathrm{H}_1(X, \mathbf{Z}) \to \mathrm{H}_1(X, \mathbf{Q}_\ell)$ are isomorphisms, so we can say that the $\mathbf{Z}_\ell$-module $\mathrm{T}_\ell(X)$ and the $\mathbf{Q}_\ell$-vector space $\mathrm{V}_\ell(X)$ that are build intrinsically from the torsion in X recover the ℓ-adic homology of X with coefficients in $\mathbf{Z}_\ell$ and $\mathbf{Q}_\ell$ respectively. The functoriality of these constructions provides natural maps

$$(3.4.1) \quad \mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathrm{Hom}(X, X') \to \mathrm{Hom}_{\mathbf{Z}_\ell}(\mathrm{T}_\ell(X), \mathrm{T}_\ell(X')), \;\; \mathbf{Q}_\ell \otimes_{\mathbf{Q}} \mathrm{Hom}^0(X, X') \to \mathrm{Hom}_{\mathbf{Q}_\ell}(\mathrm{V}_\ell(X), \mathrm{V}_\ell(X')).$$

These are nothing but the representation of torus maps in ℓ-adic homology, but formulated in a manner that has no dependence on algebraic topology or analysis. One consequence of the topological interpretation via homology lattices is:

Theorem 3.4.1. *The maps* (3.4.1) *are injective.*

Proof. We apply the flat extensions of scalars $\mathbf{Z} \to \mathbf{Z}_\ell$ and $\mathbf{Z} \to \mathbf{Q}_\ell$ to the injection

$$\mathrm{Hom}(X, X') \to \mathrm{End}_{\mathbf{Z}}(\mathrm{H}_1(X, \mathbf{Z}), \mathrm{H}_1(X', \mathbf{Z})).$$

■

Remark 3.4.2. This injectivity result holds in the algebraic theory, but by a different method of proof.

Tate modules provide an ℓ-adic version of the torsion-level pairings that we made at the end of §2.4, as follows. Passing to the inverse limit on the canonical Weil torsion pairings $X[\ell^m] \times X^\vee[\ell^m] \to \mu_{\ell^m}(\mathbf{C})$ gives a perfect $\mathbf{Z}_\ell$-bilinear pairing

$$(\cdot, \cdot)_{\ell, X} : \mathrm{T}_\ell(X) \times \mathrm{T}_\ell(X^\vee) \to \varprojlim \mu_{\ell^m}(\mathbf{C}) = \mathbf{Z}_\ell(1).$$

Although this pairing has been constructed using analytic methods (especially the analytic model for $X^\vee$), the torsion-level pairings do admit another means of construction in terms of the Poincaré bundle P_X, and in this way the ℓ-adic pairing admits a natural analogue in the algebraic theory.

Via our knowledge at torsion level, we see that for any map $f : X \to X'$ of complex tori the maps $\mathrm{T}_\ell(f)$ and $\mathrm{T}_\ell(f^\vee)$ are adjoint with respect to the ℓ-adic pairings. Likewise, if $\phi : X \to X^\vee$ is a map then we get an induced ℓ-adic self-pairing $(\cdot, \cdot)_{\phi, \ell}$ on $\mathrm{T}_\ell(X)$ via composition with $\mathrm{T}_\ell(\phi)$ into the second variable $\mathrm{T}_\ell(X^\vee)$ for $(\cdot, \cdot)_{\ell, X}$. By Remark 2.4.4 this is nothing more or less than the ℓ-adic scalar extension of the Riemann form Ψ_ϕ, and so we conclude that $(\cdot, \cdot)_{\phi, \ell}$ is skew-symmetric if and only if ϕ is a symmetric map. This symmetry criterion for ϕ in terms of skew-symmetry for the associated ℓ-adic pairing will carry over to the algebraic theory, using another method of proof.

Example 3.4.3. Let X be a complex torus of dimension $g > 0$ and let $F \subseteq \mathrm{End}^0(X)$ be a commutative subfield. By Theorem 3.4.1, we have an injection

$$\prod_{v|\ell} F_v = \mathbf{Q}_\ell \otimes_{\mathbf{Q}} F \subseteq \mathbf{Q}_\ell \otimes_{\mathbf{Q}} \mathrm{End}^0(X) \hookrightarrow \mathrm{End}_{\mathbf{Q}_\ell}(\mathrm{V}_\ell(X)),$$

so $\mathrm{V}_\ell(X)$ is a faithful module over $\prod_{v|\ell} F_v$. Hence, the associated decomposition $\mathrm{V}_\ell(X) \simeq \prod_{v|\ell} \mathrm{V}_v(X)$ into v-components has each component not equal to 0; that is, $\mathrm{V}_v(X)$ is a nonzero F_v-vector space. But $\mathrm{V}_\ell(X)$ has $\mathbf{Q}_\ell$-dimension $2g$, so

$$[F : \mathbf{Q}] = \sum_v [F_v : \mathbf{Q}_\ell] \leq \sum_v \dim_{\mathbf{Q}_\ell} \mathrm{V}_v(X) \leq \dim_{\mathbf{Q}_\ell} \mathrm{V}_\ell(X) = 2g.$$

Of course, this is just a fancy way of encoding our earlier proof of this dimension bound by considering $\mathrm{H}_1(X, \mathbf{Q})$ as an F-vector space. The advantage of the ℓ-adic argument is that it avoids any use of homology *once* one knows the basic structure of the torsion in X and one knows the injectivity in Theorem 3.4.1. These latter properties will carry over (by different proofs) in the algebraic theory, and so the dimension bound $[F : \mathbf{Q}] \leq 2g$ will also carry over. This argument is a typical illustration of how working with ℓ-adic Tate modules allows one to argue "as if" one is using homology.

Example 3.4.4. Suppose (X, i) is a CM torus with CM by the maximal order $\mathscr{O}_K$ in the CM field K. In this case $\mathrm{T}_\ell(X)$ is a faithful module over $\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathscr{O}_K \simeq \prod_{v|\ell} \mathscr{O}_{K,v}$, so for reasons of $\mathbf{Z}_\ell$-rank it is forced to be a free module of rank 1. That is, the decomposition $\mathrm{T}_\ell(X) = \prod_v \mathrm{T}_v(X)$ has each $\mathrm{T}_v(X)$ free of rank 1 over $\mathscr{O}_{K,v}$. Of course, in the analytic setting this description is silly since we can make the more refined statement that $\mathrm{H}_1(X, \mathbf{Z})$ is a torsion-free $\mathscr{O}_K$-module with $\mathbf{Z}$-rank equal to $2g = [\mathscr{O}_K : \mathbf{Z}]$, so it is an invertible module and thus free of rank 1 after we algebraically localize at ℓ. The importance of the the ℓ-adic conclusion is that it will carry over to the algebraic theory (and so provides a substitute for the properties of the integral homology as an $\mathscr{O}_K$-module).

Rather than distinguish a particular prime, we can also work with all ℓ at once. By using the primary decomposition for each $\mathbf{Z}/n\mathbf{Z}$, we see that the profinite completion $\widehat{\mathbf{Z}} = \varprojlim \mathbf{Z}/n\mathbf{Z}$ of $\mathbf{Z}$ naturally decomposes: $\widehat{\mathbf{Z}} \simeq \prod_\ell \mathbf{Z}_\ell$ as compact topological rings. Likewise, the *total Tate module* $\mathrm{T}_{\widehat{\mathbf{Z}}}(X) = \varprojlim X[n]$ (with the inverse system taken according to divisibility) is naturally identified with $\prod_\ell \mathrm{T}_\ell(X)$. In particular, $\mathrm{T}_{\widehat{\mathbf{Z}}}(X)$ is a free $\widehat{\mathbf{Z}}$-module of rank $2g$. We likewise have $\widehat{\mathbf{Z}}$-bilinear pairings (with values in $\widehat{\mathbf{Z}}(1) = \prod_\ell \mathbf{Z}_\ell(1)$), and so on. The ring $\widehat{\mathbf{Z}}$ is not a domain, and so in certain respects the ℓ-adic spaces present fewer algebraic complications. However, $\mathbf{Q} \otimes_{\mathbf{Z}} \widehat{\mathbf{Z}}$ is the ring $\mathbf{A}_{\mathrm{f}}$ of finite adeles of $\mathbf{Q}$, and so for adelic considerations it is convenient to work with the free rank-$2g$ module $\mathrm{V}_{\widehat{\mathbf{Z}}}(X) = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{T}_{\widehat{\mathbf{Z}}}(X)$ over $\mathbf{A}_{\mathrm{f}}$. For example, the adelic formulation of the Main Theorem of Complex Multiplication is expressed in terms of total Tate modules (or rather their $\mathbf{Q}$-scalar extensions as modules over the ring of finite adeles).

4. Basic algebraic definitions and examples

Having explored the analytic theory for a while, we now wish to see how it carries over to an algebraic setting. The algebraic analogue of a complex torus is called an abelian variety, and these exhibit a wide array of interesting geometric and arithmetic properties. In addition to carrying over many of the structures and concepts that arise in the analytic theory, we have to confront new phenomena related to the possibility of a base field that is not algebraically closed. This creates some technical problems, and by far the most serious technical problems to arise are connected with p-power torsion and inseparable isogenies in characteristic $p > 0$. These problems are best illustrated by considering the purely inseparable bijective pth-power map $\mathrm{GL}_1 \to \mathrm{GL}_1$ of degree p over an algebraically closed field of characteristic p.

Even if one's ultimate interest is to study Shimura varieties or abelian varieties over number fields, all of which are objects in characteristic 0, it is an unavoidable fact of life that most interesting number-theoretic properties require a good theory of reduction into positive characteristic; moreover, the tools developed to understand p-power torsion in characteristic p have grown into very useful techniques even for applications in characteristic 0. We shall avoid restriction on the characteristic whenever possible, and we refer the reader to [Mum] and [M1] for a serious treatment of the proofs of the results we shall discuss; we shall almost entirely ignore proofs, due to our minimal background assumptions in algebraic geometry.

We assume that the reader is familiar with the concept of an algebraic variety over an algebraically closed field k. For arithmetic purposes, we must allow k to be more general, such as a number field or finite field. In fact, it is also important to allow imperfect fields. For example, global function fields over finite fields are never perfect, yet they provide an excellent testing-ground for conjectures that one may hope to eventually prove in the case of number fields. Roughly speaking, to "do" algebraic geometry over a general field what one has to do is to systematically keep track of coefficients that define varieties and morphisms: it must be ensured that all constructions can be "done over the base field". There are several methods for doing this, such as using Galois theory in the case of a perfect base field. The only truly satisfying method is to use algebro-geometric foundations based on schemes, but we do not wish to use the language of schemes here (even though it is necessary to prove many results in the algebraic theory in a clean and intuitive manner). We therefore ask the reader to accept on faith that there is a systematic (and elegant!) way to keep track of coefficients to ensure one can take them to be within the base field throughout various constructions.

The viewpoint we shall use is that of working with points having coordinates in a fixed algebraically closed extension K of the base field k (with $k_s \subseteq K$ the separable closure of k in K), but the geometric objects and maps must be defined by equations whose coefficients are in the base field. For example, when we say "Let $L \to X$ be a line bundle over an variety over k" we mean that the gluing data (both maps and open affine overlaps) that describe X and Zariski-local trivializations for L are given in terms of absolutely reduced finitely generated algebras over k (and in particular, all polynomials that appear in the definitions have coefficients in k). We allow varieties to be reducible, but we require them to be quasi-compact. More precisely, a k-variety X is required to be absolutely reduced (that is, coordinate rings of affine opens are finitely generated k-algebras that acquire no nonzero nilpotents after any extension of the base field), to have irreducible components that are absolutely irreducible (that is, remain irreducible after any extension of the base field), and to have a covering by finitely many affine opens. In algebraic terms, we are requiring that for each of the finitely many irreducible components X_i of X the field k is algebraically closed in the function field $k(X_i)$ and that the extension $k(X_i)/k$ has a separating transcendence basis over k (automatic if k is perfect); in ancient terminology, this says that $k(X_i)/k$ is a regular extension of fields. Note in particular that when we say "irreducible k-variety" we are requiring absolute irreducibility.

When we wish to emphasize the underlying set of a k-variety X and ignore its k-structure, we shall write $X(K)$ rather than X. If $L \subseteq K$ is an intermediate field over k, we write $X(L)$ to denote the set of points of X that are L-rational. The reader may take all of our varieties to be quasi-projective, though one should not focus too much attention on the quasi-projective embedding. Unless we say to the contrary, all varieties and all maps among them are understood to be "defined" over a fixed base field k. It is important to keep track of behavior with respect to change in the base field, but we will generally not address this aspect of the theory.

This naive perspective is adequate for giving an exposition of the theory over a field, though it is inadequate for proving anything serious (if one wishes to allow general base fields) and anyone who intends to study the modern subject in any depth (especially with number-theoretic inclinations) must learn the modern techniques of algebraic geometry.

4.1. **Maps and torsion.** The basic object of study is:

Definition 4.1.1. An *abelian variety* X over k is a complete and irreducible group variety over k. That is, X is a complete and irreducible variety over k that is endowed with a k-rational point $e \in X(k)$ and maps of varieties

$$m : X \times X \to X, \;\; \mathrm{inv} : X \to X$$

such that the standard group-law diagrams commute (or, equivalently, $X(K)$ becomes a group via this data, with e as the identity). If $\dim X = 1$, then X is called an *elliptic curve.* A *homomorphism* between abelian varieties is a map of varieties that respects the group laws.

The concept of completeness of a variety is a "compactness" property that can be defined in several ways, such as requiring that every rational map from a smooth projective curve to the given variety is an everywhere-defined morphism. In particular, projective varieties are complete. It is an important theorem of Weil that abelian varieties are necessarily projective (see [Mum, pp. 60-62]), and with more work it can even be deduced from this that *any* abstract group variety over k is quasi-projective.

The importance of allowing the definition of an abelian variety to be apparently more general (with projectivity relaxed to completeness) is because some abstract constructions of abelian varieties are well-suited to the verification of completeness, but not projectivity. For our expository purposes, the reader may take projectivity to be a requirement in the definition of an abelian variety (as this does not actually change the underlying concept).

By a method of proof entirely different from what we have seen in the analytic case, one gets basic results analogous to what we have seen in the analytic theory. For example:

Theorem 4.1.2. *Every abelian variety is smooth, and the group law on an abelian variety is commutative and determined by the identity section. Moreover, if (X, e) and (X', e') are abelian varieties over k and $f : X \to X'$ is a map of varieties such that $f(e) = e'$ then f is a homomorphism.*

Theorem 4.1.3. *An elliptic curve has genus 1, and any smooth projective curve C over k with genus 1 and a marked point $e \in C(k)$ admits a unique structure of abelian variety with identity e.*

Example 4.1.4. As a special case of Serre's GAGA theorem, algebraic and analytic geometry for smooth projective varieties is "the same" as for compact complex manifolds admitting a projective embedding. This includes an equivalence of categories for vector bundles. Hence, an abelian variety over $\mathbf{C}$ is "the same" as a polarizable complex torus. The finiteness results for Hom-groups and the structural results for torsion thereby carry over to the algebraic side over $\mathbf{C}$, as do some results for line bundles (provided the formulation is suitably algebraic).

The structure of torsion works out rather nicely in the algebraic theory, at least away from the characteristic. Using intersection-theoretic or cohomological methods, one obtains:

Theorem 4.1.5. *Let X be an abelian variety of dimension g, and let n be a positive integer. The multiplication map $[n]_X : X \to X$ is surjective with finite fibers, and it induces an extension $[n]_X^* : k(X) \to k(X)$ with degree n^{2g} on the level of function fields. This extension is inseparable if* $\operatorname{char}(k)|n$*, and if* $\operatorname{char}(k) \nmid n$ *then all fibers of $[n]_X$ have order n^{2g} and for $x \in X(k_s)$ all points in $[n]_X^{-1}(x)$ are k_s-rational. In particular, $X(k_s)$ is n-divisible.*

The surjectivity of $[n]_X$ for all $n \neq 0$ implies:

Corollary 4.1.6. *For abelian varieties X and X' over k, the* $\mathbf{Z}$*-module* $\operatorname{Hom}_k(X, X')$ *is torsion-free.*

Consider a positive integer n not divisible by the characteristic. The commutative group $X(k_s)[n]$ is finite with order n^{2g}, and for all $d|n$ the d-torsion subgroup $X(k_s)[d]$ has order d^{2g} with multiplication by n/d carries $X(k_s)[n]$ onto $X(k_s)[d]$. By the structure theorem for finite abelian groups, the only possibility consistent with this numerology is that $X(k_s)[n] \simeq (\mathbf{Z}/n\mathbf{Z})^{2g}$ for all n. Of course, $X(k)[n]$ may have size considerably smaller than n^{2g}, and it may even be trivial.

Example 4.1.7. Suppose that X is an abelian variety over a number field k. The groups $X(k_s)[n]$ and $X(\mathbf{C})[n]$ both have size $n^{2\dim X}$, and so they coincide. That is, all torsion in $X(\mathbf{C})$ is automatically "algebraic". This is a crucial fact in permitting the use of analytic techniques in the study of abelian varieties over number fields. More generally, if k'/k is an extension of fields and k is separably closed, then for any abelian variety X over k we have $X(k)[n] = X(k')[n]$ for any n not divisible by the characteristic.

Example 4.1.8. The Mordell–Weil theorem asserts that if X is an abelian variety over a global field k then $X(k)$ is a finitely generated group. (The proof of this theorem involves serious input from algebraic number theory.) As a special case, the torsion subgroup of $X(k)$ is finite. It is a theorem of Merel that if E is an elliptic curve over a number field k then the torsion subgroup of $E(k)$ is bounded in terms of $[k : \mathbf{Q}]$ (and earlier work of Mazur gave the list of all possible torsion subgroups in the case $k = \mathbf{Q}$, the largest being $\mathbf{Z}/2\mathbf{Z} \times \mathbf{Z}/8\mathbf{Z}$); no higher-dimensional analogue is known (or even conjectured). The analogous such "torsion boundedness" theorem for elliptic curves over global function fields is a triviality in comparison with the work of Mazur and Merel (as it rests on only essentially elementary geometric facts concerning modular curves).

As in the analytic theory, for the study of homomorphisms it is convenient to introduce the notion of isogeny:

Definition 4.1.9. An *isogeny* between abelian varieties is a surjective homomorphism with finite kernel.

Many properties of isogenies of complex tori carry over to the algebraic theory (though one has to change some of the proofs, and use a more sophisticated theory of quotients than in the analytic case). For example, a map between abelian varieties of the same dimension is an isogeny if and only if it is either surjective or has finite kernel. Also, a map $f : X \to X'$ between abelian varieties of the same dimension is an isogeny if and only if there exists a map $f' : X' \to X$ such that $f' \circ f = [n]_X$ for a nonzero integer n, in which case f' is necessarily an isogeny. Thus, as in the analytic case, it is reasonable to consider the property of a pair of abelian varieties being *isogenous* (though this notion is sensitive to the base field, as $\mathrm{Hom}_k(X, X')$ may increase if k is increased).

We define $\mathrm{Hom}^0_k(X, X') = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{Hom}_k(X, X')$, and this is the Hom-group in the *isogeny category* of abelian varieties over k. By Corollary 4.1.6, passage to the isogeny category over k is a faithful functor.

4.2. **Tate modules and applications.** For a prime $\ell \neq \mathrm{char}(k)$, consider the inverse system $\{X(k_s)[\ell^m]\}$ with multiplication by ℓ as transition map at each stage. This is structurally quite similar to what we saw in the analytic theory, and we may define:

Definition 4.2.1. The ℓ-adic Tate module $\mathrm{T}_\ell(X)$ is $\varprojlim X(k_s)[\ell^m]$. Its associated vector space is $\mathrm{V}_\ell(X) = \mathbf{Q}_\ell \otimes_{\mathbf{Z}_\ell} \mathrm{T}_\ell(X)$.

From the structure of the inverse system $\{X(k_s)[\ell^m]\}$ it is clear that $\mathrm{T}_\ell(X)$ is a finite free $\mathbf{Z}_\ell$-module with rank $2g$ ($g = \dim X$), and that the natural map $\mathrm{T}_\ell(X)/\ell^m\mathrm{T}_\ell(X) \to X(k_s)[\ell^m]$ is an isomorphism for all $m \geq 1$. Hence, $\mathrm{T}_\ell(X)$ "knows" all of the ℓ-power torsion in X. These torsion subgroups are also rather substantial inside of X:

Theorem 4.2.2. *Fix a prime $\ell \neq \mathrm{char}(k)$. A closed set in X that contains all $X(k_s)[\ell^m]$'s must coincide with X. That is, the ℓ-power torsion is Zariski-dense.*

If $k = \mathbf{C}$ then density for the classical analytic topology is clear via the exponential uniformization since $\mathbf{Z}[1/\ell] \subseteq \mathbf{R}$ is dense, and this analytic topological density is stronger than density for the Zariski topology.

Proof. By the second part of Example 4.1.7, we may assume that k is algebraically closed. Let $Z \subseteq X$ be the Zariski-closure of the ℓ-power torsion. By standard arguments as with topological groups, it is easy to see that Z must be a subgroup of X. Since the subgroup of ℓ-power torsion in X is a divisible group whereas the component group of Z must be finite (and so has finite ℓ-multiplicity), it follows that all such points are in the identity component of Z. Thus, Z is connected and hence irreducible and smooth (as it is a k-group), so Z is an abelian variety. But $Z(k)[\ell]$ therefore has order $\ell^{2\dim Z}$, yet $X(k)[\ell] \subseteq Z$ is a subgroup of order $\ell^{2\dim X}$, so this forces $\dim Z = \dim X$. Hence, $Z = X$. ∎

Since the group law on X is defined over k, there is a natural structure of discrete $\mathrm{Gal}(k_s/k)$-module on $X(k_s)$, and in particular for n not divisible by $\mathrm{char}(k)$ each $X(k_s)[n]$ has such a Galois action (factoring through the Galois group of a sufficiently large finite Galois extension of k that depends on n). By continuity of each action map $\mathrm{Gal}(k_s/k) \to \mathrm{Aut}(X(k_s)[n])$, we may pass to the inverse limit with n ranging through increasing powers of a prime $\ell \neq \mathrm{char}(k)$ to arrive at an action map

$$\rho_{X,\ell} : \mathrm{Gal}(k_s/k) \to \mathrm{Aut}_{\mathbf{Z}_\ell}(\mathrm{T}_\ell(X)) \simeq \mathrm{GL}_{2g}(\mathbf{Z}_\ell)$$

that is *continuous* for the Krull topology on $\mathrm{Gal}(k_s/k)$ and the natural ℓ-adic topology on the target. Concretely, what this says is that a pair of elements of $\mathrm{Gal}(k_s/k)$ that agree on a sufficiently large finite Galois extension of k have matrix images under $\rho_{X,\ell}$ that are congruent modulo a large power of ℓ.

Definition 4.2.3. The representation $\rho_{X,\ell}$ is the *ℓ-adic representation* attached to X.

It is sometimes convenient to work instead with the action by $\mathrm{Gal}(k_s/k)$ on $\mathrm{V}_\ell(X)$, which is to say that we work in $\mathrm{GL}_{2g}(\mathbf{Q}_\ell)$, and we shall also call this the ℓ-adic representation attached to X. The context should make the intended meaning clear, though usually we prefer to work over $\mathbf{Z}_\ell$ rather than $\mathbf{Q}_\ell$.

Example 4.2.4. If k is a finite field, say with size q, then there is a canonical *arithmetic Frobenius* element $\phi \in \mathrm{Gal}(k_s/k)$ defined by $\phi(x) = x^q$. The action $\rho_{X,\ell}(\phi)$ is extremely interesting, as it encodes a lot of information about X (such as its zeta function). Weil proved that the characteristic polynomial of $\rho_{X,\ell}(\phi)$ lies in $\mathbf{Z}[T]$ and is independent of ℓ. He also showed that the algebraic-integer roots of this polynomial have absolute value $q^{1/2}$ under all embeddings into $\mathbf{C}$. This is the famous *Riemann Hypothesis* for abelian varieties.

If $f : X \to X'$ is a map between abelian varieties over k then for all n not divisible by $\mathrm{char}(k)$ the induced maps $X(k_s)[n] \to X'(k_s)[n]$ are $\mathrm{Gal}(k_s/k)$-equivariant since f is "defined over k". Hence, passage to the inverse limit gives a $\mathbf{Z}_\ell[\mathrm{Gal}(k_s/k)]$-module map

$$\mathrm{T}_\ell(f) : \mathrm{T}_\ell(X) \to \mathrm{T}_\ell(X').$$

The same goes for the V_ℓ's. In particular, we get a natural $\mathbf{Z}_\ell$-module map

$$\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathrm{Hom}_k(X, X') \to \mathrm{Hom}_{\mathbf{Z}_\ell[\mathrm{Gal}(k_s/k)]}(\mathrm{T}_\ell(X), \mathrm{T}_\ell(X')) \subseteq \mathrm{Mat}_{2g' \times 2g}(\mathbf{Z}_\ell) \tag{4.2.1}$$

with $g = \dim X$ and $g' = \dim X'$.

In §5 we will discuss the algebraic theory of polarizations (providing a "positivity" structure in the theory), and together with Corollary 4.1.6 and arguments with ℓ-adic Tate modules (used as a substitute for homology lattices) one obtains an important theorem of Weil that is analogous to results obtained very easily from the exponential uniformization in the analytic theory:

Theorem 4.2.5 (Weil)**.** *The map* (4.2.1) *is injective for all* $\ell \neq \mathrm{char}(k)$ *and* $\mathrm{Hom}_k(X, X')$ *is a finite free* $\mathbf{Z}$*-module. In particular, it has* $\mathbf{Z}$*-rank at most* $4gg'$ *and* $\mathrm{T}_\ell(X)$ *is a faithful module over* $\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathrm{End}_k(X)$.

Remark 4.2.6. Since $\mathrm{Hom}_k(X, X')$ naturally injects into the Hom-group $\mathrm{Hom}_K(X_K, X'_K)$ over K, the $\mathbf{Z}$-finiteness problem is one over algebraically closed base fields; that is, it is a geometric problem and not an arithmetic one.

Example 4.2.7. If $R \subseteq \mathrm{End}_k(X)$ is a commutative domain then it is $\mathbf{Z}$-finite and hence is an order in a number field F. By Theorem 4.2.5, the $\mathbf{Q}_\ell$-vector space $\mathrm{V}_\ell(X)$ with dimension $2\dim X$ is a faithful module over the ring $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$, and so the argument in Example 3.4.3 may be used *verbatim* to get the bound $[F : \mathbf{Q}] \leq 2\dim X$.

Let us go further and show that, as in the algebraic theory, if $[F : \mathbf{Q}] = 2\dim X$ then F must be its own centralizer in $\mathrm{End}_k^0(X)$. It is equivalent to check that $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$ is its own centralizer in $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} \mathrm{End}_k^0(X)$, so by the injectivity of (4.2.1) it suffices to prove more generally that for any injection of $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$ into $\mathrm{Mat}_{2g \times 2g}(\mathbf{Q}_\ell)$ the image is its own centralizer. But such an injection is a structure of faithful $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$-module on $\mathbf{Q}_\ell^{2g}$, and $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$ is a product of extensions $F_v/\mathbf{Q}_\ell$ with degrees adding up to $2g$, so for $\mathbf{Q}_\ell$-dimension reasons this module structure must be free of rank 1. Hence, all module endomorphisms are given by multiplication by an element of $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F$, as desired.

The preceding example motivates interest in the case $[F : \mathbf{Q}] = 2\dim X$. This gives rise to the notion of *CM abelian variety* and *CM type*; these are analogues of notions we have seen in the analytic theory. Let us consider a special case:

Example 4.2.8. Suppose that the ring of integers $\mathscr{O}_F$ of a number field F with degree $2g$ over $\mathbf{Q}$ acts on a g-dimensional abelian variety X over k. Since $\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathscr{O}$ is a finite product of discrete valuation rings, over which all torsion-free finite modules are free, we can argue as in Example 4.2.7 but with $\mathbf{Z}_\ell$ replacing $\mathbf{Q}_\ell$ to conclude that $\mathrm{T}_\ell(X)$ is a free module of rank 1 over $\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathscr{O}$. The ℓ-adic representation makes the action of $\mathrm{Gal}(k_s/k)$ on $\mathrm{T}_\ell(X)$ be linear for this rank-1 module structure, and so it can be expressed as a *continuous* map

$$\rho_{X,\ell} : \mathrm{Gal}(k_s/k) \to (\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathscr{O}_F)^\times \simeq \prod_{v|\ell} \mathscr{O}_{F,v}^\times.$$

Of course, even if we are merely given a commutative subfield $F \subseteq \mathrm{End}_k^0(X)$ with $[F : \mathbf{Q}] = 2g$ but perhaps F meets $\mathrm{End}_k(X)$ is a non-maximal order we can still conclude that $\mathrm{V}_\ell(X)$ is free of rank 1 over $\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F = \prod_{v|\ell} F_v$, so the ℓ-adic representation may be expressed as a *continuous* map

$$\rho_{X,\ell} : \mathrm{Gal}(k_s/k) \to (\mathbf{Q}_\ell \otimes_{\mathbf{Q}} F)^\times \simeq \prod_{v|\ell} F_v^\times.$$

By compactness of the source, this map lands inside of the product of the unit groups of the valuation rings at the places of F over ℓ.

Note in particular that such abelian varieties give rise to *abelian* Galois representations

Just as we introduced ℓ-adic Tate modules via inverse limits over powers of a fixed prime $\ell \neq \mathrm{char}(k)$, when $\mathrm{char}(k) = 0$ we may also form an inverse limit of $X[n]$'s over *all* positive integers n as in the analytic case:

Definition 4.2.9. If $\mathrm{char}(k) = 0$, the *total Tate module* of an abelian variety X over k is $\mathrm{T}_{\widehat{\mathbf{Z}}}(X) = \varprojlim X(k_s)[n]$. Also define the $\mathbf{A}_{\mathrm{f}}$-module $\mathrm{V}_{\widehat{\mathbf{Z}}}(X) = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{T}_{\widehat{\mathbf{Z}}}(X)$.

The total Tate module is a finite free $\widehat{\mathbf{Z}}$-module with rank $2g$, and it is canonically isomorphic to $\prod_\ell \mathrm{T}_\ell(X)$ (via primary decomposition for the $X(k_s)[n]$'s). Likewise, $\mathrm{V}_{\widehat{\mathbf{Z}}}(X)$ is a finite free module of rank $2g$ over the ring $\mathbf{A}_{\mathrm{f}}$ of finite adeles. There are associated continuous Galois representations into the topological groups of invertible $2g \times 2g$ matrices over $\widehat{\mathbf{Z}}$ (simply $\prod_\ell \rho_{X,\ell}$) and $\mathbf{A}_{\mathrm{f}}$, and for X as in Example 4.2.8 these are abelian representations with values in either the topological group $\mathbf{A}_{F,\mathrm{f}}^\times$ of finite ideles for F or its compact subgroup $\widehat{\mathscr{O}}_F^\times$ of unit finite ideles.

In the arithmetic study of CM abelian varieties over number fields, complex-analytic methods are extremely useful. To make a link between the analytic and arithmetic base fields, one has to invoke the following technical consequence of the $\mathbf{Z}$-finiteness property of Hom-groups and the Zariski-density of ℓ-power torsion:

Theorem 4.2.10. *Let X and X' be abelian varieties over k. There exists a finite separable extension k'/k such that $\mathrm{Hom}_{k'}(X_{k'}, X'_{k'}) = \mathrm{Hom}_{k''}(X_{k''}, X'_{k''})$ for any further extension field k''/k'. In particular, for any field L/k_s we have $\mathrm{Hom}_L(X_L, X'_L) = \mathrm{Hom}_{k'}(X_{k'}, X'_{k'})$.*

Nothing like this theorem is true for more general projective varieties, or even for abelian varieties if we consider maps that are not required to carry the origin to the origin.

Proof. By increasing L to be separably closed, we can assume it contains k_s. Since $\mathrm{Hom}_{k_s}(X_{k_s}, X'_{k_s})$ is finitely generated, and (by "chasing coefficients") each generator is defined over a finite subextension, we get a finite separable extension k'/k over which all of the generators are defined. Hence, $\mathrm{Hom}_{k_s}(X_{k_s}, X'_{k_s}) = \mathrm{Hom}_{k'}(X_{k'}, X'_{k'})$. The problem is therefore to prove that passage from k_s to L introduces no new maps.

Put another way, we may assume k is separably closed and in this case the problem is to show that any map $X_L \to X'_L$ as abelian varieties over an extension L/k is automatically defined over k. The key point is that for a fixed prime $\ell \neq \mathrm{char}(k)$ all ℓ-power torsion over L is already "defined" over k (as $X(L)[n]$ has size at most $n^{2\dim X}$ whenever n is not divisible by $\mathrm{char}(k)$, yet the subgroup $X(k)[n]$ is already this big since k is separably closed). By Theorem 4.2.2, within the abelian varieties X_L and X'_L over L the ℓ-power torsion subgroups are *functorial*, Zariski-dense, and consist entirely of k-rational points. This provides an enormous amount of rigidity that makes it possible to descend the given map from L down to k; the details require descent theory (especially to handle positive characteristic). ■

Example 4.2.11. Let X and X' be abelian varieties over a number field k, and suppose we construct an analytic map $f : X(\mathbf{C}) \to X'(\mathbf{C})$ between the associated complex tori. By GAGA, this arises from a unique map of abelian varieties $X_{\mathbf{C}} \to X'_{\mathbf{C}}$ over $\mathbf{C}$. By Theorem 4.2.10, the algebraization of f over $\mathbf{C}$ descends to a map $X_{k'} \to X'_{k'}$ over a finite extension k'/k. In fact, since $\mathrm{Hom}(X(\mathbf{C}), X'(\mathbf{C}))$ is finitely generated as a $\mathbf{Z}$-module we get more: there is a sufficiently large finite extension of k down to which *all* analytic maps of tori $X(\mathbf{C}) \to X'(\mathbf{C})$ descend. The ability to descend complex analytic maps down to a (possibly unknown) number field is a crucial construction technique in the theory of CM abelian varieties, and it also has important applications in the *construction* of abelian varieties over finite fields (the existence aspect of Honda–Tate theory).

Example 4.2.12. Let $\overline{\mathbf{Q}} \subseteq \mathbf{C}$ be the subfield of algebraic numbers. Consider the passage from abelian varieties over $\overline{\mathbf{Q}}$ to abelian varieties over $\mathbf{C}$ (or complex tori) via extension of the base field (or analytification over $\mathbf{C}$). Let us say that an abelian variety X over $\mathbf{C}$ is *algebraic* if it arises from an abelian variety over $\overline{\mathbf{Q}}$. Such a descent to $\overline{\mathbf{Q}}$ is *unique and functorial*. Equivalently, extending scalars from $\overline{\mathbf{Q}}$ to $\mathbf{C}$ is fully faithful (*i.e.*, it induces a bijection on Hom-groups). This is immediate from Theorem 4.2.10, and by Example 4.1.7 and the algebraic theory of quotients it follows that the property of being algebraic is even isogeny-invariant over $\mathbf{C}$. The same conclusions carry over (with the same proofs) when the extension $\mathbf{C}/\overline{\mathbf{Q}}$ is replaced with any other extension of a separably closed base field in the role of $\overline{\mathbf{Q}}$, except that for the isogeny-invariance aspect we have to restrict attention to isogenies of degree not divisible by the characteristic. (This degree restriction is a serious issue in the study of abelian varieties in positive characteristic: there are examples of elliptic curves E over $\mathbf{F}_p$ such that $E \times E$ acquires "new" p-isogenous quotients whenever we pass to a larger base field.)

Example 4.2.13. Let K be a CM field. By using the complex-analytic classification of complex tori with CM by K in terms of the CM type, it can be shown that CM tori are not only algebraic over $\mathbf{C}$ (as we know from the existence of polarizations for such tori) but they are even necessarily defined over $\overline{\mathbf{Q}}$ and hence over a *number field*. In the terminology of the preceding example, every CM torus (including the data of its CM-structure) is algebraic in the strong sense of being uniquely and functorially defined over $\overline{\mathbf{Q}}$.

Example 4.2.14. The descent-theory technique required in the proof of Theorem 4.2.10 also yields a variant that is very useful in characteristic p: if L/k is any purely inseparable extension of fields (such as a perfect closure of a field of positive characteristc) then $\mathrm{Hom}_k(X, X') = \mathrm{Hom}_L(X_L, X'_L)$ for any abelian varieties X and X' over k. Applying this to the case when L is the perfect closure of k, this reduces some construction problems for maps to the case when the base field is perfect; for example, this is a key ingredient in the construction of splitting sections in the proof of the Poincaré reducibility theorem over arbitrary base fields in positive characteristic.

4.3. **Duality and Weil pairing.** If (V, v_0) and (V', v'_0) are a pair of marked connected projective varieties then we can define the notion of a *correspondence* between them: a triple (L, i, i') with $L \to V \times V'$ a line bundle and i and i' trivializations along $V \times \{v'_0\}$ and $\{v_0\} \times V'$ that induce the same basis on the fiber $L(v_0, v'_0)$. As in the analytic theory, these admit no non-trivial automorphisms. (The content is that the only global functions on a connected projective k-variety are the elements of k.)

For a fixed (V, v_0) we can seek a *universal correspondence*: a correspondence (L, i, i') between (V, v_0) and some (X', x'_0) such that any corresponence between (V, v_0) and any other marked connected projective variety (V', v'_0) is obtained by pullback via a unique map $V' \to X'$ carrying v'_0 to x'_0. These do exist in some cases, and this provides important examples of abelian varieties:

Example 4.3.1. Let C be a smooth (irreducible) projective curve over k with genus g and equipped with a rational point $x_0 \in C(k)$. We have an analogue of Remark 2.3.5: there is a g-dimensional abelian variety $\mathrm{Jac}(C)$ called the *Jacobian* of C that is equipped with a universal correspondence on $C \times \mathrm{Jac}(C)$ (trivialized along x_0 and the origin of $\mathrm{Jac}(C)$). The construction of this Jacobian is considerably more difficult than in the analytic case, though if one tries to actually prove that the double-coset construction in the analytic case satisfies the property of being universal for correspondences containing a trivial fiber at the marked point then one will meet most of the difficulties that arise in the algebraic case. See [M2] for a treatment of the algebraic case.

Let (X, e) be an abelian variety of dimension $g > 0$ over k. There is an analogue of Theorem 2.3.2 (and, due to the absence of the Appell–Humbert classification in the algebraic theory, the proof is much more difficult; see [Mum, Ch. 2-3] and [M1, §10]):

Theorem 4.3.2. *There is an abelian variety $(X^\vee, e^\vee)$ equipped with a universal correspondence $(P_X, i, i^\vee)$ on $X \times X^\vee$ with compatible trivializations along e and $e^\vee$.*

We call this universal correspondence the *Poincaré correspondence*, and $P_X \to X \times X^\vee$ is called the *Poincaré line bundle*. Concretely, the points of $X^\vee$ classify isomorphism classes of line bundles on X (over variable extension fields of k) that can be put in a connected algebraic family with the trivial bundle; in particular, the origin $e^\vee \in X^\vee(k)$ corresponds to the trivial line bundle on X. We call $X^\vee$ the *dual abelian variety*, and it does have the same dimension as X.

The dual abelian variety, equipped with the Poincaré correspondence, has all of the properties we would expect in analogy with the analytic theory, such as analogues of Example 2.3.3 and Example 2.3.4. More specifically, if $f : X \to X'$ is a map between abelian varieties then $(f \times 1_{X'^\vee})^*(P_{X'})$ is a correspondence on $X \times X'^\vee$, so by the universal property of $X^\vee$ it has the form $(1_X \times f^\vee)^*(P_X)$ for a unique map of marked connected varieties $f^\vee : (X'^\vee, e'^\vee) \to (X^\vee, e^\vee)$. By Theorem 4.1.2, this is a map of abelian varieties, and it is called the *dual morphism*. Concretely, if we interpret points of the dual as representing isomorphism classes of line bundles then $f^\vee$ is simply the pullback operation $L' \mapsto f^*(L')$ along f. In this manner, the dual abelian variety is contravariantly functorial in the original abelian variety, as in the analytic case.

A close study of the theory of line bundles lying in connected families with the trivial bundle reveals the non-trivial fact that the formation of the dual morphism is additive: $(f_1+f_2)^\vee = f_1^\vee + f_2^\vee$ for $f_1, f_2 : X \rightrightarrows X'$. In particular, $[n]_X^\vee = [n]_{X^\vee}$ for any $n \in \mathbf{Z}$, so by the isogeny criterion for a map $f : X \to X'$ between abelian varieties of the same dimension (via factorization through $[n]_X$ for some nonzero integer n) we conclude that f is an isogeny if and only if $f^\vee$ is an isogeny. As one would expect from the analytic theory, for an isogeny f the degree of f is the same as that of its dual isogeny $f^\vee$; the proof is inspired by an observation in the analytic theory, namely that there is a duality between $\ker f$ and $\ker f^\vee$ in the analytic case, but the construction of such a duality in the algebraic is especially subtle if $\mathrm{char}(k)$ divides either of the two degrees.

Double duality also carries over to the algebraic theory, though again it is especially subtle in positive characteristic, and it expresses the inherent symmetry between X and $X^\vee$ via P_X:

Theorem 4.3.3. *If $s : X^\vee \times X \to X \times X^\vee$ is the flipping isomorphism then the map $\iota_X : X \to X^{\vee\vee}$ induced by the correspondence $(s^*(P_X), i^\vee, i)$ on $X^\vee \times X$ is an isomorphism, and $\iota_X^\vee = \iota_{X^\vee}^{-1}$.*

In the analytic theory, at the n-torsion level we built a natural perfect pairing $X[n] \times X^\vee[n] \to \mu_n(\mathbf{C})$ with dual-adjoint functoriality. These pairings also can be found in the algebraic theory, together with Galois-compatibility. To motivate the approach to an algebraic construction, let us revisit the analytic situation from another point of view. Let X be a complex torus, and choose two points $x \in X[n]$ and $x^\vee \in X^\vee[n]$. The point $x^\vee$ is the isomorphism class of a line bundle $L \to X$ such that $L^{\otimes n}$ is trivial. In particular, $n\Psi_L = \Psi_{L^{\otimes n}} = 0$, so $\Psi_L = 0$. This forces $\phi_L = 0$, so the Mumford correspondence $\wedge(L)$ is trivial over $X \times X$. That is, $m^*(L) \simeq p_1^*(L) \otimes p_2^*(L)$ over $X \times X$ (with $m : X \times X \to X$ denoting the multiplication map). Pulling this back along the diagonal gives $[2]_X^*(L) \simeq L^{\otimes 2}$. If we likewise write $[r]_X = m \circ (1 \times [r-1]_X)$ for any $r > 1$, by induction on r we get $[r]_X^*(L) \simeq L^{\otimes r}$ for all $r \geq 1$. Taking $r = n$ gives that $[n]_X^*(L)$ is trivial. Since $x \in X[n]$ we have a canonical isomorphism

$$t_x^*([n]_X^*(L)) \simeq [n]_X^*(t_{nx}^*(L)) \simeq [n]_X^*(L),$$

so for any trivializing section s of the trivial bundle $[n]_X^*(L)$ we get $t_x^*(s)$ as another such section. The ratio $s/t_x^*(s) \in \mathbf{C}^\times$ is clearly independent of s and it in fact equals the Weil n-torsion pairing $(x, x^\vee)_{P_X, n}$ as originally defined analytically.

This alternative description in terms of line bundles can be carried over to any abelian variety X over a field k (once one has available the algebraic Mumford correspondences, to be discussed in §5), where it gives *canonical* $\mathrm{Gal}(k_s/k)$-equivariant perfect pairings

$$(\cdot,\cdot)_{X,n} : X(k_s)[n] \times X^\vee(k_s)[n] \to \mu_n(k_s)$$

for all $n \geq 1$ not divisible by $\mathrm{char}(k)$; this is called the *Weil n-torsion pairing*. The canonical nature is two-fold: for any map $f : X \to X'$ between abelian varieties, the dual map $f^\vee : X'^\vee \to X^\vee$ induces the adjoint to f on n-torsion, and if n' is a multiple of n then $(\cdot,\cdot)_{X,n'}$ is compatible with $(\cdot,\cdot)_{X,n}$ via the maps from n'-torsion onto n-torsion using multiplication by n'/n. Hence, we can pass to the limit through powers of a prime $\ell \neq \mathrm{char}(k)$ to get a perfect Galois-equivariant ℓ-adic pairing

$$\mathrm{T}_\ell(X) \times \mathrm{T}_\ell(X^\vee) \to \varprojlim \mu_{\ell^m}(k_s) = \mathbf{Z}_\ell(1)$$

where the final equality is *notation*: it denotes a rank-1 free $\mathbf{Z}_\ell$-module on which $\mathrm{Gal}(k_s/k)$ acts through the ℓ-adic cyclotomic character. In the algebraic case over $\mathbf{C}$, this pairing coincides with the ℓ-adic scalar extension of the canonical pairing on homology lattices.

Example 4.3.4. If $\mathrm{char}(k) = 0$ then we can take the product of all ℓ-adic Weil pairings to get a perfect Galois-equivariant $\widehat{\mathbf{Z}}$-bilinear pairing

$$\mathrm{T}_{\widehat{\mathbf{Z}}}(X) \times \mathrm{T}_{\widehat{\mathbf{Z}}}(X^\vee) \to \varprojlim \mu_n(k_s) = \widehat{\mathbf{Z}}(1) = \prod_\ell \mathbf{Z}_\ell(1).$$

Extending scalars to $\mathbf{Q}$, this becomes a canonical perfect Galois-equivariant $\mathbf{A}_{\mathrm{f}}$-bilinear pairing

$$\mathrm{V}_{\widehat{\mathbf{Z}}}(X) \times \mathrm{V}_{\widehat{\mathbf{Z}}}(X^\vee) \to \varprojlim \mu_n(k_s) = \mathbf{A}_{\mathrm{f}}(1).$$

5. Algebraic theory of polarizations and endomorphisms

A crucial result in §4 is the $\mathbf{Z}$-module finiteness of Hom-groups in the algebraic case. The proof of this finiteness requires a "positivity" input, and this is provided by the theory of polarizations. Our present aim is to develop of understanding of the algebraic theory of polarizations, and to see how the ℓ-adic aspects of the algebraic theory provide a replacement for the Riemann-form pairings on homology lattices in the analytic theory. This further supports the principle that Tate modules are a good replacement for the homology lattice. It should also be emphasized that polarizations are not merely technical devices to prove theorems, but they are essential ingredients in the construction of good moduli functors for families of abelian varieties. Many Shimura varieties classify abelian varieties equipped with polarizations and other auxiliary discrete structure.

5.1. **The Mumford construction.** We have seen in §4.3 that there is a good algebraic theory of the dual abelian variety and dual morphisms. By using double duality, we can define the notion of a *symmetric map* $f : X \to X^\vee$: this is a map equal to its own dual in the sense that $f^\vee \circ \iota_X : X \to X^\vee$ is equal to f. The universal property of the Poincaré correspondence identifies $\mathrm{Hom}(X, X^\vee)$ with the group of self-correspondences on $X \times X$ (with tensor product as group law), under which to each f we associate the correspondence $(1_X \times f)^*(P_X)$ and to each correspondence (L, i, i') on $X \times X$ we associate the map $f : X \to X^\vee$ that "classifies" it in the sense of the universal property of the dual abelian variety (that is, $L \simeq (1_X \times f)^*(P_X)$ as correspondences on $X \times X$). In particular, if f is associated to (L, i, i') on $X \times X$ then $f^\vee \in \mathrm{Hom}(X, X^\vee)$ is associated to $(s^*(L), i', i)$ with $s : X \simeq X \simeq X \times X$ the flipping involution. Thus, f is symmetric in the sense just defined if and only if $i' = i$ and $s^*(L) \simeq L$ (which we call a *symmetric correspondence*).

In the algebraic case, a line bundle $L \to V$ on a complete k-variety V is *ample* if for some $r \geq 1$ we have $L^{\otimes r} \simeq j^*(\mathscr{O}(1))$ for a closed embedding $j : V \to \mathbf{P}^n_k$ into a projective space over k; the pair $(\mathbf{P}^n_k, \mathscr{O}(1))$ is universal as a line bundle equipped with an ordered $(n+1)$-tuple of generating sections over a k-variety, in analogy with the complex-analytic case. As in the analytic theory, ampleness on complete varieties is preserved under pullback along maps with finite fibers.

We may now copy the definition of polarization from the analytic theory:

Definition 5.1.1. Let X be an abelian variety. A *polarization* on X is a symmetric map $\phi : X \to X^\vee$ such that the line bundle $(1, \phi)^*(P_X)$ on X is ample.

In the language of correspondences, an equivalent condition is this: the data of a symmetric correspondence (L, i, i') on $X \times X$ such that $\Delta_X^*(L)$ is ample on X, where $\Delta_X : X \to X \times X$ is the diagonal. Indeed, the relationship between correspondences and elements of $\mathrm{Hom}(X, X^\vee)$ goes via the relation $L \simeq (1_X \times \phi)^*(P_X)$, and pullback along the diagonal then gives $\Delta_X^*(L) \simeq (1, \phi)^*(P_X)$ as line bundles over X.

By GAGA, in the algebraic case over $\mathbf{C}$ the above definition of polarization recovers the analytic notion of polarization. In the analytic theory we had a very concrete description of all polarizations, via Mumford correspondences with an ample line bundle. Let us formulate the Mumford construction in the algebraic setting, and then consider the extent to which it exhausts all examples. (We will get an affirmative answer in the "geometric" setting of an algebraically closed base field, but not otherwise in general.)

Definition 5.1.2. If $L \to X$ is a line bundle on an abelian variety X, the *Mumford correspondence* $(\wedge(L), i, i')$ is the line bundle

$$\wedge(L) = m^*(L) \otimes p_1^*(L)^{-1} \otimes p_2^*(L)^{-1}$$

on $X \times X$ equipped with trivializations i and i' along $X \times \{e\}$ and $\{e\} \times X$ via a common basis of $L(e)$. The associated map $\phi : X \to X^\vee$ for which $(1 \times \phi)^*(P_X) \simeq \wedge(L)$ as correspondences is denoted ϕ_L.

Remark 5.1.3. Ampleness (or not) is unaffected by change of the base field, so a map $\phi : X \to X^\vee$ is a polarization if and only if it becomes one after an extension of the base field.

By exactly the same calculation as in the analytic case, on points we have that for $x \in X(K)$ the point $\phi_L(x) \in X^\vee(K)$ is the isomorphism class of the line bundle $t_x^*(L_K) \otimes L_K^{-1}$ on X_K. Since the Mumford correspondence is visibly symmetric, the map ϕ_L is symmetric. Recall from Example 3.1.6 that in the analytic case (and so in the algebraic case over $\mathbf{C}$) *every* symmetric map has the form ϕ_L for some line bundle L. Is this true in the algebraic case? In an affirmative direction, one has:

Theorem 5.1.4. *If k is algebraically closed then every symmetric map $\phi : X \to X^\vee$ for an abelian variety X over k has the form ϕ_L for some line bundle L on X.*

Proof. This is proved in [Mum, §20, Thm. 2; §23, Thm. 3]. The only step in the proof that requires k to be algebraically closed in the analysis of the kernel in case ϕ is an isogeny (which is the essential case): one has to use the fact that for such a base field this kernel (taken in the appropriate scheme-theoretic sense if $\mathrm{char}(k) | \deg(\phi)$) can be filtered with successive quotients of a very concrete type. The case $\mathrm{char}(k) | \deg(\phi)$ requires input from the theory of finite commutative group schemes of p-power order in characteristic p. ∎

In the analytic theory, we saw that a map of the form ϕ_L is a polarization if and only if L is ample, since $(1, \phi_L)^*(P_X)$ was shown to have Riemann form $2\Psi_L$ and hence it is ample if and only if L is ample (because ampleness is determined by the Riemann form, by Lefschetz' theorem). A similar argument works in the algebraic case over any field: by [Mum, p. 59, Cor. 3], $(1, \phi_L)^*(P_X) = \Delta_X^*(\wedge(L)) \simeq L^{\otimes 2} \otimes L_0$ for some $L_0 \in \mathrm{Pic}(X)$ that comes from $X^\vee(k)$. In the algebraic theory it is shown that (as in the analytic theory) ampleness is unaffected by deformation, or in other words tensoring a line bundle against one classified by a point of the dual abelian variety does not affect the ampleness property. Since L is ample if and only if $L^{\otimes 2}$ is ample, and so if and only if $L^{\otimes 2} \otimes L_0$ is ample, it follows that in the algebraic case (as in the analytic case) the symmetric map ϕ_L is a polarization if and only if L is ample. This proves:

Corollary 5.1.5. *For algebraically closed k the Mumford construction gives all polarizations: they are precisely maps of the form ϕ_L for ample L.*

For a general k, Remark 5.1.3 implies that a map $\phi : X \to X^\vee$ is a polarization if and only if after extending the base field to an algebraic closure $\overline{k}/k$ the induced map $\phi_{\overline{k}} : X_{\overline{k}} \to X_{\overline{k}}^\vee$ acquires the form ϕ_L for some ample line bundle L on $X_{\overline{k}}$. (Equivalently, a symmetric correspondence on $X \times X$ is a polarization if and only if on $X_{\overline{k}} \times X_{\overline{k}}$ it is a Mumford correspondence arising from an ample line bundle on $X_{\overline{k}}$.) However, the map ϕ_L does not generally determine L; we saw this in the analytic theory, where ϕ_L was seen to be uniquely determined by its Riemann form $\Psi_{\phi_L} = \Psi_L$ that is the discrete invariant of L. Thus, even if $\phi_L = \phi_{\overline{k}}$ descends to a symmetric map $\phi : X \to X^\vee$, we cannot expect such an L on $X_{\overline{k}}$ to descend to X. That is, over a general field we cannot expect all polarizations of X to arise from Mumford correspondences for ample line bundles on X (nor even that all symmetric maps arise from Mumford correspondences for line bundles on X). The best one can say is the following refinement of Corollary 5.1.5:

Corollary 5.1.6. *If ϕ is a symmetric correspondence on an abelian variety X over a field k then there exists a finite separable extension k'/k such that $\phi_{k'} = \phi_{L'}$ in $\mathrm{Hom}_{k'}(X_{k'}, X^\vee_{k'})$ for some line bundle L' on $X_{k'}$, and if $n = [k' : k]$ then $n\phi = \phi_L$ for some line bundle L on X. If k is finite or separably closed then we may take $k' = k$.*

Proof. By using Corollary 5.1.5 and some general results of Grothendieck concerning Hom-schemes, it can be shown that there is a smooth irreducible projective k-variety T_ϕ that is an étale torsor for $X^\vee$ and whose k'-rational points over any field k'/k are in natural bijection with the set of isomorphism classes of line bundles L' on $X_{k'}$ for which $\phi_{L'} = \phi_{k'}$. It is a general fact that a smooth variety over a field always has a rational point over some finite separable extension, so this provides the desired finite separable k'/k (and in particular settles the case of separably closed k). In the language of Galois cohomology, if k_s/k is a separable closure then T_ϕ gives rise to a cohomology class $c_\phi \in \mathrm{H}^1(k_s/k, X^\vee(k_s))$ and for a subextension k'/k, the set $T_\phi(k')$ is non-empty if and only if the restriction $\mathrm{Res}_{k'}(c_\phi) \in \mathrm{H}^1(k_s/k', X^\vee(k_s))$ vanishes. Since $X^\vee$ is a smooth connected group variety, Lang's trick ensures the vanishing of $\mathrm{H}^1(k_s/k, X^\vee(k_s))$ when k is finite; this settles the case of finite fields.

In general, for any finite subextension k'/k the composite

$$\mathrm{H}^1(k_s/k, M) \overset{\mathrm{Res}}{\to} \mathrm{H}^1(k_s/k', M) \overset{\mathrm{Cor}}{\to} \mathrm{H}^1(k_s/k, M)$$

is multiplication by $[k' : k]$ for any discrete $\mathrm{Gal}(k_s/k)$-module M, so if we choose k'/k of degree n such that $T_\phi(k')$ is non-empty then $nc_\phi = 0$. By inspecting the construction of the class c_ϕ one finds that $Nc_\phi = c_{N\phi}$ for any integer N, so indeed $c_{n\phi} = 0$. That is, $n\phi = \phi_L$ for some line bundle L on X. ■

Over number fields there are explicit examples due to Poonen and Stoll [PS] in which one cannot take $k' = k$ in Corollary 5.1.6. More specifically, their work provides an obstruction to a polarization ϕ on an abelian variety X over a number field k having the form ϕ_L for an ample line bundle L on X, and this obstruction is encoded in a subtle arithmetic invariant of X, the Cassels–Tate pairing between the 2-torsion subgroups of the Tate–Shafarevich groups of X and $X^\vee$; in [PS] there are given several interesting examples of non-triviality of this obstruction.

Remark 5.1.7. Although one cannot expect in general that a polarization $\phi : X \to X^\vee$ over a field k has the form ϕ_L for a line bundle L, Corollary 5.1.6 provides a positive integer n for which $n\phi$ has such a form. Hence, by working in the free rank-1 $\mathbf{Z}$-module $(\mathbf{Q} \cdot \phi) \cap \mathrm{End}_k(X)$ within $\mathrm{End}^0_k(X)$ we see that the set of $q \in \mathbf{Q}$ such that $q\phi$ is a polarization arising from Mumford's construction is the set of positive integral multiples of some $q_0 > 0$. In the work of Weil and Shimura, the term "polarization" means the concept now called a *homogenous polarization*: the set of $\mathbf{Q}^\times_{>0}$-multiples of what we call a polarization, which is to say the "ample" half-line in a line $\mathbf{Q} \cdot \phi \subseteq \mathrm{End}^0_k(X)$ as above. We have just seen that in each such half-line there is a *uniquely determined* minimal element $\phi_0 = q_0\phi$ that has the form ϕ_{L_0} for an ample L_0 on X over k. Weil and Shimura defined a *basic polar divisor* attached to the homogenous polarization $\mathbf{Q}^\times_{>0}\phi$ to be any effective divisor D_0 on X for which $\phi_0 = \phi_{\mathscr{O}_X(D_0)}$ (by arguments in the proof of the "vanishing theorem" [Mum, p. 150], any ample line bundle on X has the form $\mathscr{O}_X(D)$ for an effective divisor D), and this concept eliminated the $\mathbf{Q}^\times$-ambiguities for Weil and Shimura. Note that an automorphism of X that preserves the homogenous polarization necessarily preserves the canonically determined ϕ_0, and so the automorphism group of a homogenously polarized abelian variety is identified with the automorphism group of a polarized abelian variety. This is a very fortuitous coincidence; see Remark 5.2.5.

Having explored the range of "universality" of the Mumford construction among all symmetric correspondences on an abelian variety over a general field, let us work out the Mumford construction concrete in some special cases.

Example 5.1.8. Let E be an elliptic curve over a field k and let $L = \mathscr{O}(D)$ be the line bundle associated to a divisor D on E. This is ample if and only if $d = \deg D > 0$, and the associated symmetric map ϕ_L is determined on k'-points (for any extension field k'/k) by the formula

$$\phi_L(a) = t_a^*(\mathscr{O}(D)) \otimes \mathscr{O}(-D) = \mathscr{O}(t_{-a}(D)) \otimes \mathscr{O}(-D) \simeq \mathscr{O}(t_{-a}(D) - D) \simeq \mathscr{O}([-da] - [0]) \simeq \mathscr{O}([0] - [da])$$

with the second to last isomorphism arising from the geometric description of the group law on an elliptic curve. This map is denoted ϕ_d, so $\phi_d = \phi_1 \circ [d]_E = [d]_{E^\vee} \circ \phi_1$ and in particular we see that for $d \neq 0$ the symmetric map ϕ_d is an isogeny with degree d^2 and ϕ_1 is the *unique* degree-1 polarization of E.

Note that the self-duality $E \simeq E^\vee$ used in books of Silverman, Katz–Mazur, and Serre is ϕ_{-1} and not ϕ_1, so the symmetric self-dualities ϕ used by those authors are "wrong" in the sense that $(1, \phi)^*(P_E)$ is *not* ample for $\phi = \phi_{-1}$. In fact, as we saw in the analytic theory, the maps $\phi_d = [d] \circ \phi_1$ are akin to negative-definite quadratic forms for $d < 0$. By adopting the conceptual "formula-free" approach to polarizations as has been used in these notes, with "$\phi_L(x) = t_x^*(L) \otimes L^{-1}$" derived by a calculation and not introduced as an *ad hoc* definition, we see that there is no doubt that ϕ_1 is the better self-duality to use for elliptic curves, and it is an artifact of the passage between divisors and line bundles on a curve (which has no analogue for higher-dimensional abelian varieties) that the unique degree-1 polarization of an elliptic curve has the description $a \mapsto \mathscr{O}([-a] - [0]) \simeq \mathscr{O}([0] - [a])$ in the language of divisors.

There is a similar result for Jacobians of smooth projective curves C of higher genus g. We omit an extensive discussion of this generalization, because it is most conveniently done by using the framework of correspondences for connected projective (or complete) varieties, not just abelian varieties, to pass between the symmetric correspondence

$$L = \mathscr{O}(\Delta_C) \otimes p_1^*(\mathscr{O}(x_0))^{-1} \otimes p_2^*(\mathscr{O}(x_0))^{-1}$$

on $C \times C$ and a suitable symmetric correspondence on $\mathrm{Jac}(C) \times \mathrm{Jac}(C)$, with $x_0 \in C(k)$ a rational point. We simply note that the key point is the calculation that (with I_{Δ_C} denoting the inverse bundle to $\mathscr{O}(\Delta_C)$) the line bundle

$$\Delta_C^*(L) \simeq (I_{\Delta_C}/I_{\Delta_C}^2)^{-1} \otimes \mathscr{O}(-x_0)^{\otimes 2} \simeq (\Omega^1_{C/k})^{-1} \otimes \mathscr{O}(-x_0)^{\otimes 2}$$

has degree $-(2g-2)-2 = -2g < 0$ for $g > 0$, and so it is anti-ample. Thus, it is really L^{-1} that corresponds to a polarization of $\mathrm{Jac}(C)$; this is the same sign discrepancy as in [M2, Lemma 6.9]. These annoying signs are entirely due to the translation between the languages of line bundles and divisors on curves.

Recall that in the analytic theory, we deduced that polarizations must be isogenies because positive-definite Hermitian forms are necessarily non-degenerate. In the algebraic case, the same conclusion holds but the proof is more difficult:

Theorem 5.1.9. *A polarization is an isogeny.*

Proof. The isogeny property may be checked over an algebraically closed extension of k, and so by the exhaustiveness of the Mumford construction over an algebraically closed base field it is enough to prove that ϕ_L is an isogeny for ample L. Since composition with an isogeny does not affect whether or not a map is an isogeny, and $\phi_L \circ [n]_X = \phi_{L^{\otimes n}}$, we may assume $L = j^*(\mathscr{O}(1))$ for a closed embedding $j : X \hookrightarrow \mathbf{P}^m_k$ for some $m \geq 1$. In this case $L = \mathscr{O}_X(D)$ for an *effective* divisor D (*e.g.*, a hyperplane slice of X under the embedding j), so the desired isogeny property follows from [Mum, pp. 60-61]. ■

As a consequence, we get the algebraic version of the Poincaré reducibility theorem over any field:

Theorem 5.1.10 (Poincaré)**.** *If X is an abelian variety over a field k and $X_0 \subseteq X$ is an abelian subvariety then there exists an isogeny-complement: an abelian subvariety $X_0' \subseteq X$ such that the map $X_0 \times X_0' \to X$ defined by addition is an isogeny.*

Proof. Pick an ample line bundle L on X over k, so $\phi = \phi_L : X \to X^\vee$ is a polarization. Let $j : X_0 \to X$ be the inclusion. As in the analytic case, the composite map $j^\vee \circ \phi \circ j$ is equal to ϕ_{j^*L} and $j^*(L)$ an ample line bundle on X_0, so by Theorem 5.1.9 the map $j^\vee \circ \phi \circ j$ is an isogeny from X_0 to $X_0^\vee$. In particular, X_0 has finite intersection with the identity component $X_0' = (\ker(j^\vee \circ \phi))^0$. Suppose for a moment that this identity component is a smooth k-subgroup of X, and so is an abelian subvariety. The dual $j^\vee$ is surjective since j is an embedding (this can be checked using exactness and perfectness properties of ℓ-adic Tate modules and pairings over $\mathbf{Q}_\ell$ to replace the homology arguments in the analytic case), so $\dim X_0' \geq \dim(X) - \dim(X_0)$. Since the map $X_0 \times X_0' \to X$ has finite kernel, the dimension inequality is forced to be an equality and so this addition map is an isogeny.

There remains the problem of X_0' being a smooth k-subgroup of X. If k is algebraically closed then this is no problem. If k is perfect then one can carry out a Galois descent from the algebraic closure (that is Galois over k). Hence, in general we have solved the problem over the perfect closure $k_{\rm p}/k$. But the isogeny-complement built over $k_{\rm p}$ can be expressed as the image of a projector in the endomorphism algebra (in the isogeny category) over $k_{\rm p}$, and by Example 4.2.14 this projector descends to k since $k_{\rm p}/k$ is purely inseparable. The image of a sufficiently divisible nonzero $\mathbf{Z}$-multiple of this projector provides the desired structure on X_0' inside of X. ■

As in the analytic case, we may now formally deduce everything in Corollary 3.3.2 in the algebraic case over a field, except for the finite-dimensional aspects for the End_k^0's over $\mathbf{Q}$. That is, there remains the serious gap in our knowledge that Theorem 4.2.5 has to be proved. (Nothing in our discussion of polarizations has yet required this theorem.) Theorem 4.2.5 is a geometric result, in the sense that it suffices to prove it over an algebraically closed base field. We refer the reader to [Mum, pp. 172-8] for a discussion of the proof, using Poincaré reducibility and a lot of cohomological input.

5.2. **ℓ-adic Riemann forms and the Rosati involution.** Let $\phi : X \to X^\vee$ be a symmetric morphism, so it is associated to a symmetric correspondence $L = (1 \times \phi)^*(P_X)$ on $X \times X$. Double-duality of abelian varieties is compatible with Cartier duality of torsion groups up to a minus sign [Oda, Cor. 1.3(*ii*)]. That is, when we use double duality of abelian varieties then (for n not divisible by $\mathrm{char}(k)$) the pairings

$$\langle\cdot,\cdot\rangle_{X,n} : X(k_s)[n] \times X^\vee(k_s)[n] \to \mu_n(k_s), \ \ \langle\cdot,\cdot\rangle_{X^\vee,n} : X^\vee(k_s)[n] \times X(k_s)[n] \to \mu_n(k_s)$$

are negative to each other (in the multiplicative sense) under flip of the factors. Hence, the self-pairing

$$e_{\phi,n} = \langle\cdot,\cdot\rangle_{X,n} \circ (1 \times \phi)$$

on $X(k_s)[n]$ is *skew-symmetric* when $\phi = \phi^\vee$ (as dual morphisms are adjoints with respect to the intrinsic Weil n-torsion pairings $\langle\cdot,\cdot\rangle_{X,n}$). In particular, if ϕ is an isogeny then we obtain a canonical non-degenerate pairing

$$e_{\phi,\ell^\infty} : \mathrm{T}_\ell(X) \times \mathrm{T}_\ell(X) \to \mathbf{Z}_\ell(1)$$

for any prime $\ell \neq \mathrm{char}(k)$ and this pairing is skew-symmetric if and only if ϕ is symmetric. This can be seen concretely in the complex-analytic case because the ℓ-adic pairing is the ℓ-adic scalar extension of the $\mathbf{Z}(1)$-valued lattice pairing that is the Riemann form Ψ_ϕ and we know in the analytic theory that a correspondence map ϕ is symmetric if and only if its Riemann form is skew-symmetric (as $\Psi_{\phi^\vee}(x,y) = -\Psi_\phi(x,y)$). The pairings e_{ϕ,ℓ^∞} are to be considered as ℓ-adic versions of Riemann forms in the algebraic theory.

Lemma 5.2.1. *If a symmetric isogeny $\phi : X \to X^\vee$ is a polarization then its degree* $\deg\phi > 0$ *is a perfect square.*

Proof. In the complex-analytic theory, this is proved by identifying the degree with a lattice index that is an absolute determinant of a skew-symmetric matrix (arising from the Riemann form), and the Pfaffian expresses such determinants as universal perfect squares. In the algebraic case, one can imitate this homology argument by using Tate modules (for analysis of the ℓ-part with $\ell \nmid \mathrm{char}(k)$) and Dieudonné modules (for analysis of the p-part if $\mathrm{char}(k) = p > 0$) to establish that the prime-factor multiplicities of the positive integer $\deg\phi$ are all even. An alternative algebraic approach that avoids ℓ-adic and p-adic skew-symmetric self-pairings is to extend scalars to an algebraically closed base field, where we can write $\phi = \phi_L$ for a line bundle L on X, and then use the "Riemann–Roch formula" $\deg\phi_L = \chi(L)^2$ [Mum, §16]. ■

Definition 5.2.2. A *principal polarization* of an abelian variety X over a field k is a polarization whose associated symmetric isogeny $\phi : X \to X^\vee$ is an isomorphism. A *principally polarized abelian variety* is a pair (X,ϕ) where X is an abelian variety and $\phi : X \simeq X^\vee$ is a principal polarization.

A principal polarization is a very special kind of self-duality: it is an isomorphism $\phi : X \simeq X^\vee$ that is not only symmetric but also gives rise to an ample pullback $(1,\phi)^*(P_X)$. The negative of a principal polarization is never a principal polarization (on a nonzero abelian variety), and under the dictionary of analogies between linear algebra and abelian varieties we consider a principal polarization to be the analogue of a positive definite quadratic form over $\mathbf{Z}$ with discriminant 1 (or, in other words, a positive-definite symmetric perfect bilinear pairing on a lattice over $\mathbf{Z}$). In general, an abelian variety over a field may not admit a principal polarization over the field. An abelian variety over an algebraically closed base field always admits an isogeny to a principally polarized abelian variety [Mum, Cor. 1, p. 234]. A similar result can be proved by explicit considerations with homology lattices for polarizable complex tori in the analytic theory, using the elementary fact that if B is a non-degenerate symmetric $\mathbf{Z}$-bilinear form on a lattice L and $\mathrm{disc}(B)$ is a square then there exists an isogenous lattice L' in $\mathbf{Q}\otimes_{\mathbf{Z}} L$ such that $B|_{L'\times L'}$ is $\mathbf{Z}$-valued with discriminant 1.

For any symmetric isogeny $\phi : X \to X^\vee$, the equality $\phi = \phi^\vee$ implies that the anti-automorphism of $\mathrm{End}^0_k(X)$ defined by

$$\lambda \mapsto \lambda^\dagger \stackrel{\mathrm{def}}{=} \phi^{-1} \circ \lambda^\vee \circ \phi$$

is an involution (that is, $\lambda^{\dagger\dagger} = \lambda$); this is the *Rosati involution* associated to ϕ. It is clear by definition that the Rosati involution is a $\mathbf{Q}$-algebra anti-automorphism of $\mathrm{End}^0_k(X)$ whose formation is compatible with extension of the base field. Since the ordinary dual morphism is adjoint for the intrinsic Weil pairings between X and $X^\vee$, it follows that the Rosati involution computes the adjoint morphism for the non-degenerate self-pairing e_{ϕ,ℓ^∞} on $\mathrm{T}_\ell(X)$ for any prime $\ell \neq \mathrm{char}(k)$.

We are especially interested in the Rosati involution attached to a symmetric isogeny ϕ that is a polarization, which is to say for which the line bundle $(1,\phi)^*P_X$ on X is ample. More precisely, a crucial ingredient in the proof of the Riemann Hypothesis for abelian varieties over finite fields is a positivity property for Rosati involutions associated to a polarization over an arbitrary field. Let us first recall a definition (and see [Mum, §19, Thm. 4] for the proof that it is $\mathbf{Z}$-valued and independent of ℓ):

Definition 5.2.3. Let X be an abelian variety over a field k. The linear form $\mathrm{Tr} : \mathrm{End}_k(X) \to \mathbf{Z}$ is induced by

$$\mathbf{Z}_\ell \otimes_{\mathbf{Z}} \mathrm{End}_k(X) \hookrightarrow \mathrm{End}_{\mathbf{Z}_\ell}(\mathrm{T}_\ell(X)) \stackrel{\mathrm{trace}}{\longrightarrow} \mathbf{Z}_\ell$$

for any prime $\ell \neq \mathrm{char}(k)$.

The scalar extension $\mathrm{End}^0_k(X) \to \mathbf{Q}$ of the trace will also be denote Tr. Note that it is compatible with extension of the base field. Also, by definition this trace is insensitive to order of composition of endomorphisms, and so if $\phi : X \to X^\vee$ is any symmetric isogeny and $\lambda \mapsto \lambda^\dagger$ is the associated Rosati involution then the $\mathbf{Q}$-bilinear form

$$[\lambda_1, \lambda_2]_\phi = \mathrm{Tr}(\lambda_1 \circ \lambda_2^\dagger)$$

on $\mathrm{End}^0_k(X)$ is symmetric. In the complex analytic case, the Rosati involution associated to a symmetric isogeny is essentially an adjoint with respect to a non-degenerate Hermitian form (arising from the Appell-Humbert datum for the line bundle $(1,\phi)^*(P_X)$), and so the associated symmetric $\mathbf{Q}$-bilinear trace form $[\cdot,\cdot]_\phi$ is positive-definite when ϕ is a polarization: this is because if (V,H) is a finite-dimensional Hermitian inner product space and $T : V \to V$ is a nonzero linear map then its composite TT^* against the H-adjoint is an H-normal operator that is positive-definite and hence has positive trace. By using a geometric technique, this positivity for $[\cdot,\cdot]_\phi$ in the case of polarizations on a complex torus can be proved over any ground field and thereby provides the positivity input in the proof of the Riemann Hypothesis for abelian varieties over finite fields:

Theorem 5.2.4. *Let X be an abelian variety over a field and let $\phi : X \to X^\vee$ be a symmetric isogeny. Let $\lambda \mapsto \lambda^\dagger$ be the associated Rosati involution on $\mathrm{End}^0_k(X)$. If ϕ is a polarization then the symmetric bilinear form $[\cdot,\cdot]_\phi$ on $\mathrm{End}^0_k(X)$ is positive-definite.*

Note that this positivity condition does not characterize polarizations among all symmetric isogenies, as the endomorphism algebra may be very small (*e.g.*, it might be $\mathbf{Q}$). Nonetheless, this theorem does give a certain positivity property for polarizations in the algebraic theory.

Proof. We have to prove that $[\lambda, \lambda]_\phi > 0$ for any nonzero $\lambda \in \mathrm{End}_k^0(X)$, and by scaling we may restrict attention to genuine endomorphisms $\lambda \in \mathrm{End}_k(X)$. Without loss of generality, we may suppose k is algebraically closed. Hence, $\phi = \phi_L$ for an ample line bundle L on X. By projectivity of X, we may write $L = \mathscr{O}_X(D)$ for an effective divisor D on X after perhaps replacing L with a higher tensor power (which we may certainly do). Letting $g = \dim X$, there is an explicit intersection-theory formula [Mum, §21]:

$$[\lambda, \lambda]_\phi = \frac{2g}{(D^g)} \cdot (D^{g-1}.\lambda^* D)$$

involving g-fold intersections among effective divisors. Such intersection numbers only depend on the linear equivalence class of the divisors, and so by general-position and Bertini-style arguments with very ample divisors we see that the intersection numbers are positive (since D is ample, and $\lambda^* D$ is an effective and non-empty divisor in X for D not containing $\lambda(X)$ because $\lambda(X)$ is a positive-dimensional closed subset of the projective variety X). ∎

Remark 5.2.5. A further application of integrality properties of characteristic polynomials is an important finiteness theorem that motivated Weil's interest in the concept of polarized abelian varieties: the automorphism group of a polarized abelian variety is *finite*. We refer the reader to [Mum, Thm. 5, p. 207] for this finiteness result. (This reference actually proves a stronger theorem, and yields the finiteness claim as an immediate consequence.) The analogous finiteness result in the analytic theory was obtained via a simple compactness argument in Theorem 3.3.4.)

6. Analytic models and algebraic descent

In the final two sections, we shift our viewpoint from the study of individual abelian varieties (or complex tori) to the study of *families* of such objects. We will be particularly interested in certain analytic manifolds and algebraic varieties (over arithmetically interesting fields) that parameterize such families in a universal manner. This is essentially the technique of moduli spaces, and it provides a good conceptual way to make sense of descending an algebraic structure over $\mathbf{C}$ to one over a number field. We also discuss the general technique of Shimura that gives a way to characterize such "descent of the base field" in the absence of a moduli-theoretic interpretation *provided* that we specify a Galois action at a suitable Zariski-dense set of points. The Main Theorem of Complex Multiplication provides guidance for how to describe Galois actions at CM points, which provide Zariski-dense loci in the $\mathbf{C}$-points of many Shimura varieties.

6.1. **Analytic and algebraic modular curves.** In §3.2 we saw how to build a natural analytic "family" of g-dimensional polarized complex tori with an i-oriented homology basis that standardizes the Riemann form of the polarization, using the complex manifold $\mathfrak{h}_{g,i}$ as a (set-theoretic) parameter space. In this section we shall focus our attention on the special case $g = 1$, in which case $\mathfrak{h}_{g,i}$ becomes the connected component $\mathfrak{h}_i$ of $\mathbf{C} - \mathbf{R}$ that contains a fixed choice of $2\pi i$. For each $z \in \mathfrak{h}_i$ we get an elliptic curve $E_z = \mathbf{C}/(\mathbf{Z}z \oplus \mathbf{Z})$. Specifying polarizations in the case of elliptic curves is redundant information, since an elliptic curve admits a unique polarization of each square degree. Hence, throughout our discussion we will no longer mention polarizations.

In what sense is the collection of E_z's an "analytic family" of elliptic curves parameterized by $\mathfrak{h}_i$? We certainly wish to make this better than a mere set-theoretic statement. There are (at least) three ways to make this precise, two of which are explicit and one of which is conceptual. We begin with the two explicit points of view. Letting $\Lambda_z = \mathbf{Z}z \oplus \mathbf{Z}$, the functions $g_2(\Lambda_z)$ and $g_3(\Lambda_z)$ in classical Weierstrass theory are analytic in z so

$$\mathscr{E}^{\mathrm{alg}} = \{([x, y, w], z) \in \mathbf{CP}^2 \times \mathfrak{h}_i \,|\, y^2 w = 4x^3 - g_2(\Lambda_z) x w^2 - g_3(\Lambda_z) w^3\}$$

is a closed complex submanifold in $\mathbf{CP}^2 \times \mathfrak{h}_i$ equipped with an analytic section $z \mapsto ([0, 1, 0], z)$ such that the projection $\mathscr{E}^{\mathrm{alg}} \to \mathfrak{h}_i$ is a proper submersion with fiber over each $z \in \mathfrak{h}_i$ recovering the classical Weierstrass model for $\mathbf{C}/\Lambda_z$. Another explicit point of view is to make the exponential uniformization move in a family. More precisely, consider the equivalence relation on $\mathbf{C} \times \mathfrak{h}_i$ under which $(t, z) \sim (t', z')$ if and only if $z = z'$ and $t - t' \in \Lambda_z$. One shows that the quotient $\mathscr{E}$ of $\mathbf{C} \times \mathfrak{h}_i$ modulo this equivalence relation admits a unique topology with respect to which the projection $\mathbf{C} \times \mathfrak{h}_i \to \mathscr{E}$ is a covering map, and a unique complex-analytic manifold structure with respect to which this projection is a local analytic isomorphism. The zero-section of $\mathbf{C} \times \mathfrak{h}_i$ composes with the projection to give an analytic section $e : \mathfrak{h}_i \to E$. By using the analyticity of the map

$$(t, z) \mapsto ([\wp_{\Lambda_z}(t), \wp'_{\Lambda_z}(t), 1], z)$$

on the open set

$$U = \{(t, z) \in \mathbf{C} \times \mathfrak{h}_i, \,|\, t \notin \Lambda_z\}$$

we get a unique analytic isomorphism $\mathscr{E} \simeq \mathscr{E}^{\mathrm{alg}}$ over $\mathfrak{h}_i$ that carries the "constant section" $[0, 1, 0]$ to the section e.

Finally, we give the conceptual point of view that unifies the two explicit (yet isomorphic) structures just built over $\mathfrak{h}_i$.

Definition 6.1.1. An *elliptic curve over a complex manifold* S (or a *family of elliptic curves parameterized by* S) is a proper submersion $\pi : E \to S$ of complex manifolds equipped with a holomorphic section $e : S \to E$ such that each compact Hausdorff fiber E_s is a connected Riemann surface with genus 1 (and so is uniquely an elliptic curve in the traditional sense with identity $e(s)$). A *morphism* between elliptic curves (E, e) and (E', e') over S is a map $f : E \to E'$ over S such that $f \circ e = e'$.

This definition aims to capture the intuitive idea of an analytically-varying family of elliptic curves with parameter space S. There are several ways in which one may seek to generalize this definition (*e.g.*, allow S to have singularities), but it is adequate for our present purposes to stick with the definition as just given. Both $\mathscr{E}$ and $\mathscr{E}^{\mathrm{alg}}$ built above are elliptic curves over $\mathfrak{h}_i$ and they are in fact isomorphic as such.

Example 6.1.2. If a and b are two analytic functions on S such that $a^3 - 27b^2$ is non-vanishing on S then

$$\{([x, y, w], s) \in \mathbf{CP}^2 \times S \,|\, y^2 w = 4x^3 - a(s)xw^2 - b(s)w^3\}$$

is an elliptic curve over S. In fact, *every* analytic family of elliptic curves parameterized by a complex manifold looks like this example *locally* over the base. The proof requires some serious input from analytic sheaf cohomology.

We shall need to use the notion of *pullback* for an elliptic curve $\pi : E \to S$ with respect to an analytic map $f : S' \to S$. Roughly speaking, the pullback should be the family of elliptic curves $\{E_{f(s')}\}_{s' \in S'}$ parameterized by the points of S'. Rigorously, it is

$$E' = \{(x, s') \in E \times S' \,|\, \pi(x) = f(s')\} \subseteq E \times S',$$

and one shows that E' is a closed submanifold of $E \times S'$ such that its natural projection π' to S' and section $e' = (e, \mathrm{id}_{S'})$ gives it a structure of elliptic curve over S' (with s'-fiber $E_{f(s')}$ as desired). In an evident manner, if $f' : S'' \to S'$ is a second map then the pullback of E' along f' is naturally isomorphic to the pullback of E along $f \circ f'$; in this sense, the formation of pullback is transitive.

It is a basic fact from differential geometry (Ehresmann's theorem) that any proper submersion π of C^∞ manifolds is a C^∞-fibration. In particular, the homologies at nearby fibers are *canonically* identified with each other. (Namely, over a small open ball B in the base over which $\pi^{-1}(B) \to B$ is a split fibration, $\mathrm{H}_1(\pi^{-1}(b), \mathbf{Z}) \to \mathrm{H}_1(\pi^{-1}(B), \mathbf{Z})$ is an isomorphism for all $b \in B$.) It therefore makes sense to speak of parallel transport of fibral homology along paths in the base S, and when S is connected this defines a representation of $\pi_1(S, s_0)$ on the fibral homology at $s_0 \in S$. If this representation is trivial (*e.g.*, if S is simply connected), then we can consistently identify the fibral homologies on *all* fibers in a manner that is compatible with parallel transport along paths in S.

Definition 6.1.3. An elliptic curve $E \to S$ over a complex manifold has *trivialized fibral homology* if there is given an ordered basis of $\{\sigma_s, \sigma'_s\}$ of $\mathrm{H}_1(E_s, \mathbf{Z})$ for all $s \in S$ such that both $s \mapsto \sigma_s$ and $s \mapsto \sigma'_s$ are compatible with parallel transport. This structure is *i-oriented* for $i = \sqrt{-1} \in \mathbf{C}$ if the intersection product $\sigma'_s.\sigma_s \in \mathbf{Z}(1)$ is equal to $2\pi i$ for all $s \in S$. (Note the order of the intersection pairing.)

Example 6.1.4. Consider the elliptic curve $\mathscr{E} \to \mathfrak{h}_i$. The fibral homology basis $\{[z], [1]\}$ for $\mathrm{H}_1(\mathbf{C}/\Lambda_z, \mathbf{Z}) = \mathrm{H}_1(\mathscr{E}_z, \mathbf{Z})$ is a trivialized fibral homology structure on $\mathscr{E} \to \mathfrak{h}_i$. It is also i-oriented, since $[1].[z] = 2\pi i_z = 2\pi i$ for $i_z = \sqrt{-1} \in \mathbf{C} - \mathbf{R}$ in the connected component of $z \in \mathfrak{h}_i$.

Example 6.1.5. If $E \to S$ is an elliptic curve with trivialized fibral homology and $f : S' \to S$ is an analytic map then the pullback elliptic curve $E' \to S'$ has trivialized fibral homology in an evident manner (using the identification $E'_{s'} = E_{f(s')}$ and some topological considerations).

An elliptic curve equipped with trivialized fibral homology has no non-trivial automorphisms. More precisely, an automorphism of an elliptic curve that preserves a chosen basis on fibral homologies must be the identity. Hence, when we speak of two such structures being isomorphic, there is exactly one such isomorphism and so we do not need to name it.

Example 6.1.4 is *universal* in the following sense:

Theorem 6.1.6. *Let $E \to S$ be an elliptic curve equipped with i-oriented trivialized fibral homology via the homology basis $\{\sigma_s, \sigma'_s\}$ on the s-fiber for all $s \in S$. There is a unique analytic map $S \to \mathfrak{h}_i$ along which the pullback of $\mathscr{E} \to \mathfrak{h}_i$ with its canonical i-oriented trivialized fibral homology is isomorphic to $E \to S$ with its given fibral homology trivialization.*

We emphasize that this theorem assigns *conceptual meaning* to the old idea of an upper half-space serving as a parameter space for elliptic curves. The actual structure being parameterized is not merely an elliptic curve, but rather an elliptic curve equipped with an i-oriented homology basis. (How does such a structure determine a preferred tangent vector at the origin, as is implicit in the standard Weierstrass uniformization $\mathbf{C}/\Lambda_z$?) In particular, we now have a way to think about this structure over $\mathfrak{h}_i$ in a conceptual manner that does not rest on the recipe of its explicit construction.

Proof. Set-theoretically, for each $s \in S$ the nonzero points σ_s, σ'_s in the lattice $\mathrm{H}_1(E_s, \mathbf{Z}) \subseteq \mathrm{T}_{e(s)}(E_s)$ are $\mathbf{C}$-linearly dependent. We may therefore write $\sigma_s = z(s)\sigma'_s$ with respect to this complex structure, using a unique $z(s) \in \mathbf{C}$. The $\mathbf{R}$-independence of σ_s and σ'_s forces $z(s) \in \mathbf{C} - \mathbf{R}$, and the i-orientation hypothesis forces $z(s) \in \mathfrak{h}_i$. The map $z : S \to \mathfrak{h}_i$ defined by $s \mapsto z(s)$ is the unique possibility that can work, and the hard part is to prove that this map really is analytic and that it does work. A rigorous justification of these plausible claims requires a lot of preparations, so we omit the details. ■

The preceding theorem is an example of a solution to an *analytic moduli problem*: we consider a certain kind of analytic structure relative to a base manifold (such as an elliptic curve equipped with i-oriented trivialized fibral homology) such that there is a reasonable notion of "pullback", and we seek such a structure that is *universal* in the sense that all instances of the structure are *uniquely* pullbacks of the universal one. The concept of a universal structure is an extremely important one throughout modern mathematics, and in a geometric setting the base space for such a structure is called a *fine moduli space*: one should imagine that the points of such a space are universal parameters for the structure under consideration. When working with fine moduli spaces in the setting of algebraic geometry this interpretation of points in the base space requires some care, due to issues related to field of rationality of points. This is easiest to understand with some specific examples, as we shall provide shortly over $\mathbf{Q}$.

The importance of universal objects is that one can use the universality to construct global analytic maps by "pure thought" without the interference of explicit calculations and analytic verifications; all of the hard work is put into the proof of universality (which we omitted in the case of the family we built over $\mathfrak{h}_i$). More precisely, there are often several (or more) explicit models for the *same* universal structure, and so explicit formulas may be an artifact of the particular model. An abstract viewpoint helps us to distinguish those properties that have conceptual meaning and those that depend on the choice of explicit model.

For example, consider

$$\gamma = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbf{Z}).$$

For any pair $(E \to S, (\sigma, \sigma'))$ consisting of an analytic family of elliptic curves and an i-oriented trivialized fibral homology, we get another such pair by replacing (σ, σ') with $(a\sigma + b\sigma', c\sigma + d\sigma')$; the property that $ad - bc \in \mathbf{Z}^\times$ be positive is required precisely to preserve the i-orientation condition on the fibral homology basis. Observe that this operation on pairs $(E, (\sigma, \sigma'))$ makes no reference to any elliptic curve over an upper half-plane. Applying it to the universal structure $\mathscr{E} \to \mathfrak{h}_i$, it follows from universality that there is a unique analytic map $[\gamma] : \mathfrak{h}_i \to \mathfrak{h}_i$ that pulls back the universal structure to the one we just made by manipulating the fibral homology basis in accordance with γ. We claim $[\gamma](z) = (az + b)/(cz + d)$, thereby revealing the intrinsic meaning of the classical linear fractional action of $\mathrm{SL}_2(\mathbf{Z})$ on $\mathfrak{h}_i$. Indeed, multiplication by $cz + d$ on $\mathbf{C}$ carries $\Lambda_{(az+b)/(cz+d)}$ to $\mathbf{Z}(az+b) \oplus \mathbf{Z}(cz+d) = \Lambda_z$ and so induces an isomorphism $\mathscr{E}_{(az+b)/(cz+d)} \simeq \mathscr{E}_z$ that carries $\sigma((az + b)/(cz + d))$ to $a\sigma(z) + b\sigma'(z)$ and carries $\sigma'((az + b)/(cz + d))$ to $c\sigma(z) + d\sigma'(z)$. This shows (by the intrinsic meaning of the points of $\mathfrak{h}_i$ in terms of the universal structure over it) that γ must be the asserted linear fractional action.

Let us now pose two more examples of analytic moduli problems, one of which admits a solution and the which does not. In contrast with the case of i-oriented trivialized fibral homology for analytic families of elliptic curves, the problems we now consider will mention only torsion and so will have reasonable analogues in the algebraic theory (whereas the notion of an i-oriented homology basis is manifestly topological and so has no good algebraic analogue).

Example 6.1.7. A $\Gamma_1(N)$-*structure* on an elliptic curve E is a point of exact order N. More generally, for an elliptic curve $\pi : E \to S$ over a complex manifold a $\Gamma_1(N)$-*structure* is a section $P : S \to E$ such that $P(s) \in E_s$ has exact order N for all $s \in S$. Loosely speaking, it is the specification of an analytically-varying point of exact order N in the fibers E_s. For example, the family $\mathscr{E} \to \mathfrak{h}_i$ admits many $\Gamma_1(N)$-structures, such as

$$z \mapsto ([\wp_{\Lambda_z}((rz + s)/N), \wp'_{\Lambda_z}((rz + s)/N), 1], z) \tag{6.1.1}$$

for $(r, s, N) = 1$.

Example 6.1.8. A $\Gamma_0(N)$-*structure* on an elliptic curve E is a cyclic subgroup $C \subseteq E$ with order N. In order to make this definition work in analytic families, we have to formulate the correct replacement for a cyclic subgroup of order N; we *do not* consider such a subgroup as $\mathbf{Z}/N\mathbf{Z}$, since this description specifies extra information consisting of a choice of generator and so puts us back in the setting of $\Gamma_1(N)$-structures. Let us define a *finite group* over a complex manifold S to be a finite covering space $G \to S$ endowed with a fibral group structure that is locally constant with respect to local splittings of the covering space. (That is, the canonical bijection between nearby fibers is a group isomorphism.) Such a structure is called commutative, cyclic of order N, *etc.*, when the fibers have this property.

A $\Gamma_0(N)$-*structure* on an elliptic curve $E \to S$ is a cyclic group $C \to S$ of order N equipped with a closed embedding $C \hookrightarrow E$ over S such that $C_s \to E_s$ is a group homomorphism for all $s \in S$. For example, any $\Gamma_1(N)$-structure P on an elliptic curve $E \to S$ gives rise to a $\Gamma_0(N)$-structure by taking the union of the subgroups $\langle P(s) \rangle \subseteq E_s$ for all $s \in S$; one can check that this union has a unique structure of finite group over S with respect to which its set-theoretic map to E is a closed embeding. Thus, for example, we get many $\Gamma_0(N)$-structures on $\mathscr{E} \to \mathfrak{h}_i$ by using (6.1.1) for varying pairs (r, s).

The notions of *pullback* for $\Gamma_1(N)$-structure and $\Gamma_0(N)$-structure are defined in an evident manner. Hence, it makes sense to inquire about the possible existence of universal structures of these types (as we vary the analytic family of elliptic curves).

Example 6.1.9. We can push Example 6.1.7 further. Let $\Gamma_1(N) \subseteq \mathrm{SL}_2(\mathbf{Z})$ be the subgroup

$$\Gamma_1(N) = \{\gamma \in \mathrm{SL}_2(\mathbf{Z}) \,|\, \gamma \equiv \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix} \bmod N\}.$$

Under the universal action by $\mathrm{SL}_2(\mathbf{Z})$ on $\mathscr{E} \to \mathfrak{h}_i$, the action by $\Gamma_1(N)$ preserves the $\Gamma_1(N)$-structure P given by Example 6.1.7 with $r = 0$ and $s = 1$. That is, if $\gamma \in \Gamma_1(N)$ then the unique isomorphism $\mathscr{E}_{\gamma(z)} \simeq \mathscr{E}_z$ as elliptic curves with i-oriented trivialized homology carries $P(\gamma(z)) = 1/N \bmod \Lambda_{\gamma(z)}$ to $P(z) = 1/N \bmod \Lambda_z$.

In this manner, we have an action by $\Gamma_1(N)$ on the $\Gamma_1(N)$-structure $(\mathscr{E} \to \mathfrak{h}_i, P)$ *ignoring* the fibral homology structure.

Provided that $N \geq 4$, the action by $\Gamma_1(N)$ on $\mathfrak{h}_i$ is a free action, and so we can pass to the quotient by this left action to get a $\Gamma_1(N)$-structure

$$(\mathscr{E}_1(N), P_1) \to Y_1(N)^{\mathrm{an}} := \Gamma_1(N)\backslash\mathfrak{h}_i.$$

This procedure does *not* work with $\Gamma_0(N)$ because the action by this group on $\mathfrak{h}_i$ is never free and so although the quotient $Y_0(N)^{\mathrm{an}} = \Gamma_0(N)\backslash\mathfrak{h}_i$ makes sense as an open Riemann surface the quotient by the $\Gamma_0(N)$-action on $\mathscr{E}$ turns out to have singularities and is not an elliptic curve over $Y_0(N)^{\mathrm{an}}$.

By building on the techniques used in the proof of Theorem 6.1.6 one can show that the preceding construction for $\Gamma_1(N)$ is universal (in the sense that all $\Gamma_1(N)$-structures are uniquely obtained from it via pullback):

Theorem 6.1.10. *For $N \geq 4$, the $\Gamma_1(N)$-structure $(\mathscr{E}_1(N), P_1) \to Y_1(N)^{\mathrm{an}}$ is universal among all $\Gamma_1(N)$-structures on elliptic curves over complex manifolds. For any $N \geq 1$, there is no universal object among $\Gamma_0(N)$-structures on elliptic curves over complex manifolds but the underlying set of $Y_0(N)^{\mathrm{an}}$ is in bijective correspondence with the set of isomorphism classes of $\Gamma_0(N)$-structures on ordinary elliptic curves. Moreover, for any $\Gamma_0(N)$-structure C on an elliptic curve $E \to S$ over a complex manifold S there is a functorially associated analytic map $S \to Y_0(N)^{\mathrm{an}}$ (sending $s \in S$ to the point classifying the isomorphism class of (E_s, C_s)) and $Y_0(N)^{\mathrm{an}}$ is initial among complex manifolds with this property.*

We may summarize this theorem as saying that $Y_1(N)^{\mathrm{an}}$ is a *fine moduli space* for the problem of classfying analytic families of $\Gamma_1(N)$-structures, whereas $Y_0(N)^{\mathrm{an}}$ is merely a *coarse moduli space* for the problem of classifying analytic families of $\Gamma_0(N)$-structures.

In the algebraic theory, it seems reasonable that we should be able to define a good notion of algebraic family of elliptic curves parameterized by a variety, and the notions of $\Gamma_1(N)$-structure and $\Gamma_0(N)$-structure on such a family if N is not divisible by the characteristic. The one serious technical algebraic issue that does not arise in the analytic setting is that we have to confront the problem of non-rational points on varieties. By imposing some restrictions that are unpleasant from the viewpoint of the general theory but adequate for our purposes, the basic definitions may be given as follows:

Definition 6.1.11. Let k be a field and V a smooth k-variety. An *elliptic curve over* V is a smooth k-variety E equipped with a map $\pi : E \to V$ and a section $e : V \to E$ such that each fiber $E_v = \pi^{-1}(v)$ of π is an elliptic curve with identity $e(v)$. For any $N \geq 1$ not divisible by $\mathrm{char}(k)$, a $\Gamma_1(N)$*-structure* on $E \to V$ is a section $P : V \to E$ such that for all k'/k and $v \in V(k')$ the point $P(v) \in E_v(k')$ has exact order N. (It suffices to work with k' an algebraic closure of k.) Likewise, a $\Gamma_0(N)$-structure on $E \to V$ is a closed subset $C \subseteq E$ that is a cyclic subgroup of order N on all fibers over $V(\overline{k})$ and that is finite étale over V.

Remark 6.1.12. The notion of a *finite étale* map in Definition 6.1.11 is an algebraic analogue of a finite covering space (and over $k = \mathbf{C}$ it is equivalent to the associated map on $\mathbf{C}$-points with the classical topology being a finite covering map).

A reasonable notion of *pullback* can be defined for elliptic curves and $\Gamma_1(N)$-structures and $\Gamma_0(N)$-structures (agreeing set-theoretically with the analytic picture when working with $\overline{k}$-fibers), so it is meaningful to ask if there are universal objects in each case. It is also meaningful to ask for a coarse moduli variety: one whose points rational over an algebraically closed field are in bijection with isomorphism classes of structures and which is initial among all smooth varieties that universally receive classifying maps from the base of any family for the structure being considered (such as $\Gamma_0(N)$-structures). Based on the analytic theory, the following result is not surprising:

Theorem 6.1.13. *For $N \geq 4$, within the category of $\mathbf{Q}$-varieties there exists a universal $\Gamma_1(N)$-structure $(E, P) \to Y_1(N)$ over a smooth absolutely irreducible affine curve $Y_1(N)$ over $\mathbf{Q}$, and this universal property persists after any extension of the base field. Moreover, the analytification of its scalar extension to $\mathbf{C}$ is a fine moduli space for $\Gamma_1(N)$-structures on varying elliptic curves in the complex-analytic category.*

For any $N \geq 1$, there is a smooth affine absolutely irreducible curve $Y_0(N)$ over $\mathbf{Q}$ that is a coarse moduli variety for the moduli problem of $\Gamma_0(N)$-structures on elliptic curves over smooth $\mathbf{Q}$-varieties, and this property persists after any extension of the base field. Moreover, the analytification of its scalar extension to $\mathbf{C}$ is a coarse moduli space for $\Gamma_0(N)$-structures on varying elliptic curves in the complex-analytic category.

Proof. This is most conceptually proved by using techniques of Grothendieck. We omit the details; see [KM] for an exhaustive discussion in rather greater generality. ■

6.2. **Characterization of algebraic models via Galois action.** In general, if k'/k is an extension of fields and X and Y are two algebraic varieties over k, then it may happen that the associated varieties $X_{k'}$ and $Y_{k'}$ over k' are isomorphic without X and Y being isomorphic over k. For example, if we consider smooth conics in the projective plane over k then such conics that have a k-rational point are k-isomorphic to the projective line over k but those without a k-rational point certainly are not. Hence, all smooth planar conics over an algebraically closed base field k are k-isomorphic, whereas $u^2+v^2+w^2$ over $\mathbf{Q}$ has no rational points and $uv = w^2$ has many rational points so these curves are not $\mathbf{Q}$-isomorphic.

These examples show that the following problem is non-trivial: given a variety V' over k', how can we encode the specification of a pair (V, i) consisting of a k-variety V and an isomorphism $i : V_{k'} \simeq V'$ of k'-varieties? We say such a V is a *descent* of V' down to k. If (V_1, i_1) and (V_2, i_2) are two descents of V' down to k (via k'-isomorphisms $i_j : V_{j/k'} \simeq V'$) then we say that they are *isomorphic* as descents if there is a k-isomorphism $f : V_1 \simeq V_2$ such that the induced map $f_{k'}$ over k' is compatible with the identifications of $V_{1/k'}$ and $V_{2/k'}$ with V'; that is, $i_2 \circ f_{k'} = i_1$. We shall see some elementary examples of non-isomorphic descents below. Note that the k'-isomorphism i is a crucial part of the data of a descent (V, i) of V' down to k.

Of course, we wish to consider descent not only for geometric objects, but also for maps between them. For example, if we are given k-varieties V and W that are descents of k'-varieties V' and W', how can we determine if a given k'-map $f' : V' \to W'$ arises from a k-map $f : V \to W$, and to what extent is such an f uniquely determined? Although the affine case is easier to work with, it is important for applications that we consider these problems for varieties that are not necessarily affine. Also, we have to allow for the possibility that k'/k is not a finite extension, nor even algebraic, since in the theory of Shimura varieties one typically has the case $k' = \mathbf{C}$ and k a number field. Note that in these cases k' is algebraically closed and the characteristic is 0.

The theory that we are going to explain in this section was developed by Weil, Shimura, and others, back in the days when algebraic geometry was concerned with objects over a field. The theory was later brilliantly generalized by Grothendieck within the framework as schemes, and Grothendieck's descent theory grew into one of the fundamental techniques in modern algebraic geometry. We shall give an exposition of the main definitions, examples, and results from the viewpoints of Weil and Shimura, but the reader should keep in mind that Grothendieck's approach provides a much clearer picture and gives stronger results, even in the setting over a field. (Strictly speaking, certain aspects of the theory of Shimura and Weil do not fit exactly into Grothendieck's theory, but the results of Shimura and Weil can nonetheless all be most clearly proved by using Grothendieck's point of view.) When working with descent one also wants to descend *properties* of objects and maps. That is, if $\mathbf{P}$ is a property of varieties (resp. maps between varieties), such as smoothness, irreducibility, or having dimension d (resp. being surjective, having smooth fibers, being a group law) then to what extent is the property $\mathbf{P}$ for a given variety over k (resp. map between k-varieties) equivalent to the same property after extending scalars to k'? We will not address this kind of question here, but in practice it is an important problem and Grothendieck's theory gives the most satisfying systematic approach to solving all such problems.

Example 6.2.1. Let E' be an elliptic curve over k'. We can ask for a descent of E' down to k as merely an algebraic curve, or also as an elliptic curve (that is, descend the group structure as well). If E is an elliptic curve over k that descends E' then $j(E') = j(E) \in k$. Hence, if $j(E') \in k'$ does not lie in k then we certainly cannot descend E' to k *as an elliptic curve*. However, this is a stronger requirement than mere descent as an algebraic curve. For example, let C be a smooth planar cubic over k with no k-rational point. (A famous example over $k = \mathbf{Q}$ is $3u^3 + 4v^3 + 5w^3 = 0$; the non-existence of $\mathbf{Q}$-points is due to Selmer.)

There is certainly a finite Galois extension k'/k for which $C(k')$ is non-empty, so the curve $E' = C_{k'}$ over k' is a smooth plane cubic with a k'-rational point; let $e' \in E'(k')$ be such a point. The pair (E', e') is an elliptic curve over k' and C is a descent of E' down to k as an algebraic curve but not as an elliptic curve.

If E' is an elliptic curve over k' such that $j(E') \in k$, then does E' necessarily descend to an elliptic curve over k, and if so then how many such descents are there? Rather than consider the question in general, we illustrate the subtleties with an example. Let $f \in k[x]$ be a separable cubic over a field k with $\mathrm{char}(k) \neq 2$ (say, $k = \mathbf{Q}$) and for each $d \in k^\times$ let E_d be the elliptic curve $dy^2 = f(x)$ over k. It can be proved that E_d and $E_{d'}$ are k-isomorphic as elliptic curves if and only if the ratio $d/d' \in k^\times$ is a square in k. Thus, if d is a nonsquare in k then E_d and E_1 become isomorphic over the quadratic extension $k(\sqrt{d})$ but they are not isomorphic over k. If k'/k is a compositum of many of the quadratic fields $k(\sqrt{d})$ for d ranging through a set of representatives of multiplicatively independent elements of $k^\times/(k^\times)^2$ then we get many pairwise non-isomorphic descents over k for a single elliptic curve over k'.

Suppose that k'/k is a Galois extension (maybe with infinite degree), and let $\Gamma = \mathrm{Gal}(k'/k)$. Let V be a k-variety, and let $V' = V_{k'}$ be the associated k'-variety, so V descends V' down to k. Let us show how V' acquires some interesting structure through the specification of its descent V down to k, somewhat in the spirit of Galois theory.

We begin by defining an "action" of Γ on geometric objects and maps between them, somewhat analogous to the way in which Galois groups act on elements of a field. For any k'-variety X' and any $\gamma \in \Gamma$, let X'^{γ} denote the k'-variety obtained from X' via the extension of the base field given by $\gamma : k' \simeq k'$. Concretely (at the level of open affine subvarieties), X'^{γ} is defined by the same equations over k' as is X' except that the k'-coefficients are replaced with their images under γ. For example, if X' is a hypersurface in $\mathbf{P}^n_{k'}$ defined by the vanishing of some absolutely irreducible degree-d homogeneous polynomial $g(t_0, \dots, t_n) = \sum c'_I t^I$ with $c'_I \in k'$ (and I ranging through all multi-indices $(i_0, \dots, i_n)$ of non-negative integers adding up to d) then X'^{γ} is the projective hypersurface defined by the vanishing of the polynomial $\gamma(g) = \sum \gamma(c'_I) t^I$. Similarly, if $f' : X'_1 \to X'_2$ is a k'-map between two k'-varieties, then we get a k'-map $f'^{\gamma} : X_1'^{\gamma} \to X_2'^{\gamma}$ by applying γ to the k'-coefficients that define f' between open affine subvarieties of X'_1 and X'_2. It is not difficult to check that the formation of f'^{γ} is compatible with composition in f'.

For $\gamma_1, \gamma_2 \in \Gamma$ there is a natural isomorphism $c_{\gamma_1,\gamma_2} : (V'^{\gamma_1})^{\gamma_2} \simeq V'^{\gamma_2\gamma_1}$ that expresses the transitivity of tensor products (note the order of multiplication on the target of c_{γ_1,γ_2}!) and this is "associative" with respect to a third element of Γ in an evident manner. In general one does not expect X'^{γ} to be k'-isomorphic to X' for $\gamma \neq 1$, but in certain situations there arise interesting isomorphisms of this type. More specifically, in the case of a k'-variety $V' = V_{k'}$ for a k-variety V we have k'-isomorphisms

$$[\gamma] : V' = V_{k'} \simeq (V_{k'})^{\gamma} = V'^{\gamma} \tag{6.2.1}$$

induced geometrically by applying γ to the coordinates of solutions of the equations over k that define the k-descent V of V'. (As in Galois theory, if we are studying solutions to polynomials over k then we can apply field automorphisms over k to rearrange the solutions). In terms of coordinate rings at the affine level, (6.2.1) says that if A is a k-algebra and $A' = k' \otimes_k A$ is the associated k'-algebra then there is an isomorphism of k'-algebras $k' \otimes_{\gamma,k'} A' \simeq A'$ induced by

$$k' \otimes_{\gamma,k'} A' = k' \otimes_{\gamma,k'} (k' \otimes_k A) \simeq k' \otimes_k A = A'$$

where the middle isomorphism is induced by associativity of tensor products (since $\gamma' : k' \simeq k'$ is a k-map). The *fundamental observation* is that the geometric isomorphisms $[\gamma]$ over k' satisfy a compatibility condition that relates composition among the $[\gamma]$'s and multiplication in the group Γ: the composite k'-isomorphism

$$V' \overset{[\gamma_2]}{\simeq} V'^{\gamma_2} \overset{[\gamma_1]^{\gamma_2}}{\simeq} (V'^{\gamma_1})^{\gamma_2} \overset{c_{\gamma_1,\gamma_2}}{\simeq} V'^{\gamma_2\gamma_1}$$

is equal to $[\gamma_2\gamma_1]$. This is the *cocycle condition*, and it is written more compactly as the identity:

$$[\gamma_2\gamma_1] = c_{\gamma_1,\gamma_2} \circ [\gamma_1]^{\gamma_2} \circ [\gamma_2].$$

The reason why the cocycle condition is so fundamental is that it only mentions the k'-variety V' and the k'-isomorphisms $[\gamma]$; it *does not mention* a k-descent V of V' (though we just showed how the structure of a k-descent provides such a collection of k'-isomorphisms). Let us now isolate this structure by giving it a name:

Definition 6.2.2. Let k'/k be a Galois extension of fields. A *Galois k'/k-descent datum* on a k'-variety V' is a collection of k'-isomorphisms $[\gamma] : V' \simeq V'^{\gamma}$ (for all $\gamma \in \mathrm{Gal}(k'/k)$) such that the cocycle condition is satisfied. If V'_1 and V'_2 are two k'-varieties endowed with respective Galois descent data $\{[\gamma]_1\}$ and $\{[\gamma]_2\}$, then a k'-map $f' : V'_1 \to V'_2$ *respects the descent data* if the diagram

$$\begin{array}{ccc} V_1'^{\gamma} & \xrightarrow{f'^{\gamma}} & V_2'^{\gamma} \\ {\scriptstyle [\gamma]_1}\uparrow\simeq & & \simeq\uparrow{\scriptstyle [\gamma]_2} \\ V'_1 & \xrightarrow[f']{} & V'_2 \end{array}$$

commutes for all γ.

Example 6.2.3. If V is a k-variety and $V' = V_{k'}$ then the Galois k'/k-descent datum $\{[\gamma]\}$ constructed above is called the *canonical* Galois descent datum on V' arising from V. If $f : V_1 \to V_2$ is a k-map of k-varieties and $f' : V'_1 \to V'_2$ is the induced k'-map of k'-varieties then f' respects the canonical Galois descent data on its source and target.

Example 6.2.4. Let $(V', \{[\gamma]\})$ be a k'-variety equipped with k'/k-descent data. For any field L/k' the descent data provides an action of $\mathrm{Aut}(L/k)$ on the set $V'(L)$, as follows. Pick $\sigma \in \mathrm{Aut}(L/k)$ and let $\overline{\sigma} = \sigma|_{k'} \in \Gamma$ be the induced automorphism on the Galois subextension k'/k. For any $v' \in V'(L)$, we get an induced point $v'^{\sigma} \in V'^{\overline{\sigma}}(L)$ by applying σ to the coordinates of v'. Applying the k'-isomorphism $[\overline{\sigma}]^{-1} : V'^{\overline{\sigma}} \simeq V'$ carries v'^{σ} to an L-rational point of V', denoted $\sigma(v')$. By using the cocycle condition, one checks that $v' \mapsto \sigma(v')$ is an action of $\mathrm{Aut}(L/k)$ on the set $V'(L)$. (That is, $\tau(\sigma(v')) = (\tau \circ \sigma)(v')$ for all $\tau, \sigma \in \mathrm{Aut}(L/k)$; note the order of composition!)

Likewise, there is also an action by Γ on the function field $k'(V')$ over the natural action on the constant subfield k'; by *function field* in the possibly reducible case we really mean the product of the function fields of the finitely many irreducible components of V'. We define the action of $\gamma \in \Gamma$ on $k'(V')$ via pullback by $[\gamma]$: this induces a k'-algebra isomorphism of function fields

$$k' \otimes_{\gamma, k'} k'(V') = k'(V'^{\gamma}) \simeq k'(V')$$

via $a' \otimes g' \mapsto a' \cdot [\gamma^{-1}]^*(g'^{\gamma})$ and the cocycle condition implies that assigning to each $g' \in k'(V')$ the image $[\gamma^{-1}]^*(g'^{\gamma})$ of $1 \otimes g'$ under this composite isomorphism defines an action of Γ on $k'(V')$ over the natural Γ-action on k'.

Let $\mathscr{C}_k$ be the category of k-varieties (with k-maps) and let $\mathscr{C}_{k'/k}$ be the category of pairs $(V', \{[\gamma]\})$ consisting of a k'-variety V' and a Galois k'/k-descent datum (with maps in $\mathscr{C}_{k'/k}$ given by k'-maps of varieties respecting the Galois descent data). We get a functor $D : \mathscr{C}_k \to \mathscr{C}_{k'/k}$ by assigning to any k-variety V (resp. k-map $f : V_1 \to V_2$ between two k-varieties) the k'-variety $V_{k'}$ equipped with its canonical Galois k'/k-descent datum (resp. the k'-map $f_{k'}$). There arise two natural questions. First, is the functor $D : \mathscr{C}_k \to \mathscr{C}_{k'/k}$ fully faithful? That is, for a pair of k-varieties V_1 and V_2, is the natural map $\mathrm{Hom}_k(V_1, V_2) \to \mathrm{Hom}_{k'}(V_{1/k'}, V_{2/k'})$ given by $f \mapsto f_{k'}$ a bijection onto the subset of k'-maps that respect the canonical Galois k'/k-descent data? We shall see below that the answer to this question is affirmative, so a k-variety V is *functorially determined* (as a k-variety) by the specification of its associated k'-variety $V' = V_{k'}$ equipped with canonical descent data as an object in $\mathscr{C}_{k'/k}$. That is, if V_1 and V_2 are k-varieties with associated canonical k'/k-descent data $\{[\gamma]_j\}$ on $V_{j/k'}$ and we are given a map

$$(V_{1/k'}, \{\gamma]_1\}) \to (V_{2/k'}, \{[\gamma]_2\})$$

in the category $\mathscr{C}_{k'/k}$ then this k'-map $f' : V_{1/k'} \to V_{2/k'}$ has the form $f_{k'}$ for a unique k-map $f : V_1 \to V_2$. In particular, if f' is an isomorphism respecting the descent data then it uniquely has the form $f_{k'}$ for an isomorphism $f : V_1 \simeq V_2$ of k-varieties.

To what extent is the functor D an equivalence of categories? That is, if we are given an object $(V', \{[\gamma]\})$ in $\mathscr{C}_{k'/k}$ then does there exist a k-variety V such that $D(V)$ in $\mathscr{C}_{k'/k}$ is isomorphic to $(V', \{[\gamma]\})$? (In other words, we seek a k-variety V and a k'-isomorphism $i : V_{k'} \simeq V'$ such that i carries the canonical Galois descent datum on $V_{k'}$ over to the given k'/k-descent data $\{[\gamma]\}$ on V'.) This is called the *descent problem* for $(V', \{[\gamma]\})$, and if the descent problem has an affirmative answer (in which case the answer is *functorially unique*, in view of the affirmative answer to the first question) then we say that the descent for $(V', \{[\gamma]\})$ down to k is *effective*.

There are two necessary conditions for effectivity, as follows. Choose an algebraic closure $\overline{k}/k$ containing k', and consider the $\mathrm{Aut}(\overline{k}/k)$-orbits in $V'(\overline{k})$ (via Example 6.2.4 with $L = \overline{k}$) and the $\mathrm{Gal}(k'/k)$-action on $k'(V')$ (as in Example 6.2.4). The *continuity condition* is the condition that all $\mathrm{Aut}(\overline{k}/k)$-orbits in $V'(\overline{k})$ are finite and that all $\mathrm{Gal}(k'/k)$-orbits on $k'(V')$ are finite, and the *affineness condition* is the condition that each orbit in $V'(\overline{k})$ is contained in an affine open k'-subvariety of V'.

The continuity condition expresses the fact that for any finitely generated k-algebra A any k-map $A \to \overline{k}$ has only finitely many conjugates under composition with $\mathrm{Aut}(\overline{k}/k)$ and every element of $k' \otimes_k A$ is a finite sum of elementary tensors. The affineness condition expresses the fact that any solution V to the descent problem is covered by open affine k-subvarieties. Of course, the continuity condition is trivially satisfied in case $[k' : k]$ is finite (though for application with Shimura varieties one wants to consider the case when k'/k is an algebraic closure of a number field). Note that if the continuity condition holds then the affineness property has an affirmative answer if more generally *any* finite set of points in V' is contained in an open affine, and this latter property is *always* satisfied if V' is quasi-projective over k' (as will be the case in all applications). These two necessary conditions for effectivity are also sufficient:

Theorem 6.2.5. *Let k'/k be a Galois extension of fields. The functor $D : \mathscr{C}_k \to \mathscr{C}_{k'/k}$ is fully faithful, and an object $(V', \{[\gamma]\})$ in $\mathscr{C}_{k'/k}$ has effective descent if and only if for an algebraic closure $\overline{k}/k'$ the $\mathrm{Aut}(\overline{k}/k)$-action induced on $V'(\overline{k})$ by the descent datum satisfies the continuity condition and the affineness condition.*

Proof. The proof is best understood within the framework of schemes, so we omit the details except to say a few words. First of all, the problem of descending maps is converted into a geometric problem of descending locally closed subsets (and then even closed subsets) by consideration of the graph. Suitable use of Galois theory and the Nullstellensatz then permits one to settle the problem of descent for locally closed subsets by descending radical ideals over open affines. This takes care of the full faithfulness aspect. The real difficulty in the proof of effectivity is to reduce to the affine case. It is this reduction to the affine case that uses the continuity and affineness conditions in order to cover V' by open affine k'-subvarieties U' that are stable under the descent datum (in the sense that $[\gamma]$ carries U' isomorphically to U'^{γ} for all γ). Grothendieck's techniques show that in such a situation it suffices to solve the problem for the U''s, and moreover in the case that k'/k is finite. Once we are in the affine case with k'/k finite, everything can be deduced from the classical Galois descent theorem for vector spaces: if W' is a k'-vector space (perhaps with infinite dimension) and if Γ is given an action on W' that is semi-linear over the natural action on the scalars (that is, for $c' \in k'$ and $w' \in W'$ we have $[\gamma](c'w') = \gamma(c')[\gamma(w')]$) then for the k-subspace W of Γ-invariant elements in W' the natural k'-linear Γ-equivariant map $k' \otimes_k W \to W'$ is an isomorphism. This classical theorem is proved by using the normal basis theorem in Galois theory. ■

Theorem 6.2.5 is a useful result, but for our purposes it is really just a source of intuition for the style of descent theorem that is required for the definition of canonical models in the sense of Shimura. The problem Shimura considered was to specify a structure of descent to a number field for an algebraic variety over $\mathbf{C}$. The effectivity aspect for geometric objects turns out to be rather subtle, but the full faithfulness aspect for maps (and so the *functorial uniqueness* of descent, if it exists) works out quite nicely, as we shall see. If we are given an algebraic variety V' over $\mathbf{C}$, how can we specify a descent of V' to a k-variety V over a number field $k \subseteq \mathbf{C}$? Suppose we are given a $\mathbf{C}$-isomorphism $V' \simeq V_{\mathbf{C}}$, so exactly as in the case of Galois extensions we get an "action" of $\Gamma = \mathrm{Aut}(\mathbf{C}/k)$ on V'. That is, for each $\gamma \in \Gamma$ there is an evident $\mathbf{C}$-isomorphism $[\gamma] : V' \simeq V'^{\gamma}$, and these satisfy the cocycle condition. In particular, exactly as in Example 6.2.4 we get an action of Γ on the set $V'(\mathbf{C})$. Since $\mathrm{Aut}(\mathbf{C}/k)$ acting on $\mathbf{C}$ preserves the subfield $\overline{\mathbf{Q}}$ of algebraic numbers, it follows from unwinding the definitions that the subset $V(\overline{\mathbf{Q}}) \subseteq V'(\mathbf{C})$ of algebraic points of V is Γ-stable. Moreover, it is a Zariski-dense subset of V'. This provides the key to generalizing Galois descent to the case of non-algebraic extensions. We first need to introduce a suitable generalization of the Galois condition:

Definition 6.2.6. A field extension k'/k is *quasi-Galois* if the subfield of $\mathrm{Aut}(k'/k)$-invariants in k' is k.

Of course, if k'/k is a Galois extension (with perhaps infinite degree) then it is quasi-Galois. Another important class of examples is:

Example 6.2.7. Let k be a perfect field and let k'/k be an algebraically closed extension of k. The examples to keep in mind are k a number field and $k' = \mathbf{C}$. We claim that k'/k is quasi-Galois. Let $\overline{k}$ be the algebraic closure of k in k', so $\overline{k}/k$ is an algebraic closure (as k' is algebraically closed). Abstractly, $k'/\overline{k}$ is an algebraic closure of a purely transcendental extension of $\overline{k}$ with some intrinsic cardinality, and so as such if $a \in k'$ is not in $\overline{k}$ then each of a and $a+1$ can be put into its own transcendence basis for k'/k and so we can use the abstract description of k' to make a $\overline{k}$-automorphism of k' that moves a to $a+1$. This shows that an $\mathrm{Aut}(k'/k)$-invariant element of k' must lie in $\overline{k}$. By Galois theory, since $\overline{k}$ is Galois over k (as k is perfect) the only $\mathrm{Aut}(\overline{k}/k)$-invariant elements of $\overline{k}$ are the elements of k. Hence, it is enough to prove that the natural map $\mathrm{Aut}(k'/k) \to \mathrm{Aut}(\overline{k}/k)$ is surjective. More generally, we shall prove that any automorphism of $\overline{k}$ extends to an automorphism of k'.

In fact, for any extension k'/k of fields with k' algebraically closed, we claim that every automorphism σ of k lifts to an automorphism of k'. Let $\{x_i\}$ be a transcendence basis of k'/k, so k' is an algebraic closure of the field $k(x_i)$. The automorphism σ lifts to an automorphism of the subfield $k(x_i)$ by action on the coefficients, so upon renaming $k(x_i)$ as k we are reduced to proving that an automorphism of a field lifts to an automorphism of an algebraic closure. This follows from the uniqueness of algebraic closures.

Let k'/k be a quasi-Galois extension of fields, with $\Gamma = \mathrm{Aut}(k'/k)$, and let $\overline{k} \subseteq k'$ be the algebraic closure of k in k'. (This is an algebraic closure of k in the abstract sense if k' is algebraically closed, which is the case of most interest in applications, though in general the notation is merely suggestive and $\overline{k}$ may not be algebraically closed.) By [S, Prop. 6.13], the algebraic extension $\overline{k}/k$ is Galois and Γ surjects onto $\mathrm{Gal}(\overline{k}/k)$. By [S, Prop. 6.14], for every finite extension k_0/k contained in $\overline{k}$ the extension k'/k_0 is quasi-Galois.

Exactly as in the Galois case, if V' is a k'-variety then we can define the concept of a *k'/k-descent datum* on V', namely a collection of k'-automorphisms $V' \simeq {V'}^{\gamma}$ that satisfy the cocycle condition, and for any k-variety V the associated k'-variety $V_{k'}$ is equipped with a canonical k'/k-descent datum. We may also use a k'/k-descent datum to define a natural action by Γ on the set $V'(k')$. If V' and W' are two k'-varieties equipped with k'/k-descent data, then we define the concept of a map of k'-varieties $f' : V' \to W'$ being *compatible with the descent data* in exactly the same manner as in the Galois case.

For our purposes, the most convenient formulation for a suitable theorem analogous to Theorem 6.2.5 will use another structure, as follows. Let V' be a k'-variety and let $S' \subseteq V'(k')$ be a Zariski-dense subset that is equipped with an action by $\mathrm{Gal}(\overline{k}/k)$ with finite orbits. For example, if V is a k-variety and $V' = V_{k'}$ then $S' = V(\overline{k})$ works if we use the natural action by Γ on $V'(k') = V(k')$. Also, if $f : V_1 \to V_2$ is a k-map of k-varieties then the induced map $f_{k'} : V_{1/k'} \to V_{2/k'}$ over k' is compatible with the canonical Γ-actions on the subsets $V_{j/k'}(k') = V_j(k')$ and carries $V_1(\overline{k})$ over into $V_2(\overline{k})$. Thus, if a map $f' : V_{1/k'} \to V_{2/k'}$ over k' is to have any hope of descending to a k-map $V_1 \to V_2$ then it must carry the Zariski-dense subset $V_1(\overline{k})$ into the Zariski-dense subset $V_2(\overline{k})$ and be compatible with the $\mathrm{Gal}(\overline{k}/k)$-actions on these subsets. We are led to the following variant on the definition of a Galois descent datum:

Definition 6.2.8. Let k'/k be a quasi-Galois extension of fields, with $\overline{k} \subseteq k'$ denoting the algebraic closure of k in k', and let V' be a k'-variety. A *Shimura k'/k-descent datum* on V' is a Zariski-dense subset $S' \subseteq V'$ consisting of k'-rational points and an action by $\mathrm{Gal}(\overline{k}/k)$ on S' with finite orbits. A *descent* of such a descent datum is a pair (V, i) consisting of a k-variety V and a k'-isomorphism $i : V_{k'} \simeq V'$ such that $S' \subseteq i(V(\overline{k}))$ and the given action by $\mathrm{Gal}(\overline{k}/k)$ on S' is compatible via i with the natural action by this group on $V(\overline{k})$.

This definition is a bit of a mouthful, so let us restate it informally: the idea is that we wish to specify a k-descent of $V_{k'}$ such that the points in S' become "algebraic" with respect to the descent V over k. Since

$S' \subseteq V'(k')$ and the descent structure identifies $V'(k')$ with $V_{k'}(k') = V(k')$, the algebraicity property for the points of S' with respect to V says exactly that such k'-rational points have coordinates with respect to V that lie in the algebraic closure $\overline{k}$ of k in k'. But there is a natural action by $\mathrm{Gal}(\overline{k}/k)$ on the subset $V(\overline{k}) \subseteq V(k') = V_{k'}(k') = V'(k')$, and so we can be greedy and ask for more: perhaps the subset S' identified inside of $V(\overline{k})$ is stable under this action and the resulting action on S' is given by some recipe that we specify in advance. Such an *a priori* recipe for this Galois action is precisely the content of the Shimura k'/k-descent datum, and the hope is that by specifying such a Galois action on a Zariski-dense set we have imposed enough structure to determine the k-descent V of V' uniquely (if it exists!). Indeed this works, as is recorded in the following variant on Theorem 6.2.5:

Theorem 6.2.9 (Shimura)**.** *Let V' be a k'-variety and $S' \subseteq V'(k')$ a Zariski-dense subset endowed with a Shimura k'/k-descent datum. If there exists a descent (V, i) of this descent datum then it is unique up to unique k-isomorphism. That is, if (V_1, i_1) and (V_2, i_2) are two such pairs then there is a unique k-isomorphism $f : V_1 \simeq V_2$ such that $i_2 \circ f_{k'} = i_1$.*

Moreover, if W' is another k'-variety equipped with a Zariski-dense subset $T' \subseteq W'$ consisting of k'-rational points and having a Shimura k'/k-descent datum, and if (W, j) is a descent of (W', T'), then a k'-map $V' \to W'$ carrying S' into T' and equivariant for the $\mathrm{Gal}(\overline{k}/k)$*-actions on these subsets uniquely descends to a k-map $f : V \to W$.*

The meaning of this theorem is that (up to the specification of the Zariski-dense subset) we can functorially determine a k-descent of a k'-variety by specifying a Zariski-dense subset of points that are to become algebraic points on the descent and on which the $\mathrm{Gal}(\overline{k}/k)$-action is given by a specified action. In applications to certain Shimura varieties, we will have $k' = \mathbf{C}$ and k a number field, and the set S' will consist of CM points with a Galois action specified by the Main Theorem of Complex Multiplication. Observe that we are not making any assertion in the direction of the hardest part of Theorem 6.2.5, namely the existence of effective descent for geometric objects. This is a very difficult problem; all we are doing is giving an analogue of the easier claims in Theorem 6.2.5 concerning the existence and uniqueness of descent for maps.

Proof. The uniqueness of descent up to unique isomorphism is a special case of the second assertion concerning unique descent of maps. Indeed, if (V_1, i_1) and (V_2, i_2) are two descents of a given Shimura descent datum (V', S') (with specified $\mathrm{Gal}(\overline{k}/k)$-action on $S' \subseteq V'(k')$) then the composite k'-isomorphism $V_{1/k'} \simeq V' \simeq V_{2/k'}$ and its inverse both satisfy the requirements for the second part of the theorem when we use the subsets of $V_1(k')$ and $V_2(k')$ coming from the Zariski-dense $S' \subseteq V'(k')$. Hence, if the second part of the theorem is proved then we can descend these composite k'-maps between $V_{1/k'}$ and $V_{2/k'}$ to k-maps $V_1 \to V_2$ and $V_2 \to V_1$ that must be inverse to each other by uniqueness (as these maps become inverse over k'). Hence, it suffices to verify the second part of the theorem.

Let V and W be k-varieties and let $f' : V_{k'} \to W_{k'}$ be a k'-map such that $f'(S') \subseteq T'$ for Zariski-dense subsets $S' \subseteq V_{k'}$ and $T' \subseteq W_{k'}$ coming from $V(\overline{k})$ and $W(\overline{k})$, and assume that $f' : S' \to T'$ is equivariant for the natural actions by $\Gamma = \mathrm{Aut}(k'/k)$ on $V(k')$ and $W(k')$. (Recall that by definition of a descent of a Shimura k'/k-descent datum, these Γ-actions do preserve S' and T', and moreover on these subsets such Γ-actions factor through given actions by $\mathrm{Gal}(\overline{k}/k)$ on each subset.) In such a situation, we need to prove that f' has the form $f' = f_{k'}$ for a unique k-map $f : V \to W$. Observe that the k-structures V and W provide k'-isomorphisms $[\gamma]_V : V' \simeq V'^{\gamma}$ and $[\gamma]_W : W' \simeq W'^{\gamma}$ for all $\gamma \in \Gamma$ such that the cocycle condition is satisfied. The first key observation is that $f'^{\gamma} = f'$ for all $\gamma \in \Gamma$, by which we really mean $[\gamma]_W^{-1} \circ f'^{\gamma} \circ [\gamma]_V = f'$. To verify such an equality of k'-maps it is enough to do so on a Zariski-dense subset, so it suffices to check on the subset $S' \subseteq V'(k') = V(k')$. Using the natural Γ-action on $V(k')$ and $W(k')$, for any $v' \in V(k')$ we have

$$([\gamma]_W^{-1} \circ f'^{\gamma} \circ [\gamma]_V)(v') = \gamma^{-1}(f'(\gamma(v')))$$

in $W(k')$. Hence, to show f'^{γ} and f' coincide on S' it is the same as to show $f'(\gamma(s')) = \gamma(f'(s'))$ in $W(k')$ for all $s' \in S'$. But this identity follows from the hypothesis of Γ-compatibility for $f' : S' \to T'$. Thus, we indeed have $f'^{\gamma} = f'$ for all $\gamma \in \Gamma$.

The Zariski-dense subsets have done their job, and we now ignore them by instead proving rather generally that if $f' : V_{k'} \to W_{k'}$ is a k'-map that satisfies $f'^{\gamma} = f'$ for all $\gamma \in \Gamma$ then f' uniquely descends to a k-map $f : V \to W$. Consider the graph map $(1, f') : V' \to V' \times W'$. This is an immersion onto a locally closed k'-subvariety, and if we let $Z' \subseteq V' \times W'$ denote this locally closed k'-subvariety then its projection to V' is a k'-isomorphism whose inverse (composed with projection to W') recovers the map f'. Moreover, Z' is stable under the natural k'/k-descent datum on $V' \times W' = (V \times W)_{k'}$ in the sense that the natural isomorphism

$$V' \times W' \simeq V'^{\gamma} \times W'^{\gamma} \simeq (V' \times W')^{\gamma}$$

carries Z' isomorphically to Z'^{γ} for all $\gamma \in \Gamma$. Hence, rather than try to descend k'-maps it suffices to descend locally closed k'-subvarieties that are stable under the descent data, *provided* it is known that a k-map that becomes an isomorphism over k' had to have been an isomorphism over k. The descendability of the isomorphism property is an elementary result in Grothendieck's descent theory. Our problem therefore is indeed reduced to that of descending locally closed k'-subvarieties.

To allow ourselves to work locally, it is convenient to formulate the problem more generally without the interference of cartesian products. Let P be a k-variety (such as $V \times W$ above) and let Z' be a locally closed k'-subvariety of $P' = P_{k'}$ such that Z' is carried isomorphically to Z'^{γ} under the natural k'-isomorphism $P' \simeq P'^{\gamma}$ for all $\gamma \in \Gamma = \mathrm{Aut}(k'/k)$. We claim that $Z' = Z_{k'}$ for a unique locally closed k-subvariety Z in P. In view of the uniqueness aspect of the claim, it suffices to solve the problem by working locally on P. Also, formation of Zariski-closure for a locally closed set commutes with extension of the base field, so the Zariski-closure of Z' in P' is a closed k'-subvariety that satisfies the same hypotheses as does Z' in P'. Hence, if we can solve the problem for closed k'-subvarieties then the closure of Z' in P' will uniquely descend to a closed k-subvariety of P, and upon replacing P with such a closed k-subvariety it will then suffice to uniquely descend open k-subvarieties (since any locally closed subvariety is open in its closure). To summarize, it is enough to solve the problem of existence and uniqueness of k'/k-descent separately for open subvarieties and closed subvarieties within P' that are stable under the natural k'/k-descent datum on $P' = P_{k'}$, provided that we prove the unique descents in both cases are open and closed respectively. Grothendieck's descent theory ensures that a locally closed immersion of k-varieties is an open immersion (resp. closed immersion) if and only if the induced k'-map has this property, so it is enough to focus on uniquely descending open subvarieties to open subvarieties and closed subvarieties to closed subvarieties.

Any open subvariety $U' \subseteq P'$ has closed complement that is possibly reducible but nonetheless admits an algebraic structure via the corresponding radical ideal on affines opens in P'. (This closed complement may not be a k-variety; the function fields of its irreducible components may not be regular extensions of k.) The uniqueness of this algebraic structure on the closed complement of U' implies that it is stable under the isomorphisms $[\gamma] : P' \simeq P'^{\gamma}$ whenever $[\gamma]$ carries U' isomorphically to U'^{γ} for all γ. We can therefore shift our attention to the more general problem of proving existence and uniqueness of descent for arbitrary closed subsets $Z' \subseteq P_{k'}$ equipped with the unique "reduced" structure through the associated radical ideal on affine opens in P'.

By working locally on P we may assume P is affine, say with coordinate ring A, and Z' corresponds to a radical ideal I' in $A' = k' \otimes_k A$. The stability condition on Z' within $P' = P_{k'}$ under the k'-isomorphisms $P' \simeq P'^{\gamma}$ is precisely the condition that $I' \subseteq A'$ is stable under the natural Γ-action on A' through action on the left tensor factor in $k' \otimes_k A$. An elementary argument with bases of vector spaces shows that if W_1, W_2 are subspaces of a vector space W over a field F and if F'/F is an extension field then

$$F' \otimes_F (W_1 \cap W_2) = (F' \otimes_F W_1) \cap (F' \otimes_F W_2)$$

inside of $F' \otimes_F W$, and that $W_1 \subseteq W_2$ inside of W if and only if $F' \otimes_F W_1 \subseteq F' \otimes_F W_2$ inside of $F' \otimes_F W$. It follows that if we can show $I' = k \otimes_k I$ as k'-subspaces of A' for a k-subspace $I \subseteq A$ then I is unique and it is necessarily an ideal in A.

Our problem is now reduced to a general question in linear algebra: if k'/k is a quasi-Galois extension of fields (a property that we have not yet used!) and if W is a k-vector space then for any k'-subspace $W_1' \subseteq W' = k' \otimes_k W$ that is stable under the action by $\Gamma = \mathrm{Aut}(k'/k)$ on W' has the form $k' \otimes_k W_1$ for a unique k-subspace $W_1 \subseteq W$. The uniqueness is immediate, for by choosing a k-basis of W that extends a

k-basis of W_1 we see via the equality ${k'}^{\Gamma} = k$ that $(W_1')^{\Gamma} = W_1$ inside of ${W'}^{\Gamma} = W$. Hence, the problem is one of existence, and we are led to define $W_1 = (W_1')^{\Gamma} \subseteq {W'}^{\Gamma} = W$ and to prove that the natural map $k' \otimes_k W_1 \to W_1'$ is an isomorphism. Injectivity of this map follows from the fact that $k' \otimes_k W \to W'$ is injective (even an isomorphism), so the problem is one of surjectivity. We argue by contradiction, using the classical argument of minimal expressions in elementary tensors. Suppose W_1' is larger than its k'-subspace $k' \otimes_k W_1$. Any element $w' \in W'$ not in $k' \otimes_k W_1$ must be nonzero and when viewed as an element of $W' = k' \otimes_k W$ it can be expressed as a finite sum $\sum_i c_i' \otimes w_i$ of elementary tensors. We consider such an element w' that can be expressed as such a sum with a minimal number of elementary tensors, and we shall deduce a contradiction.

Since w' must be nonzero, there are actually some elementary tensors that show up in a minimal expression $\sum_{i=1}^n c_i' \otimes w_i$ for w'. The w_i's must have k-linearly independent images in W/W_1, due to the minimality condition. It is clear that replacing w' with a ${k'}^{\times}$-multiple does not affect that it lies in $W_1' - k' \otimes_k W_1$ and that it satisfies the minimality property. Hence, we may assume $c_1' = 1$. For each $\gamma \in \Gamma$, the γ-stability of W' implies that W' contains the vector

$$w' - \gamma(w') = \sum_{i=1}^{n}(c_i' - \gamma(c_i')) \otimes w_i = \sum_{i=2}^{n}(c_i' - \gamma(c_i')) \otimes w_i.$$

This is a sum of at most $n-1$ elementary tensors, whence by minimality of n this difference must lie in $k' \otimes_k W_1$. But the w_i's are linearly independent modulo W_1, so we can find a k-basis of W that contains the w_i's and a k-basis of W_1. Hence, the only way a k'-linear combination of the $1 \otimes w_i$'s can lie in $k' \otimes_k W_1$ is if the linear combination is trivial. We conclude that all differences $c_i' - \gamma(c_i')$ in k' for $i > 1$ must vanish for all γ, so $c_2', \dots, c_n'$ lie in ${k'}^{\Gamma} = k$ (the quasi-Galois hypothesis!). Hence,

$$w' = 1 \otimes w_1 + \sum_{i=2}^{n} 1 \otimes (c_i' w_i) = 1 \otimes w$$

for $w = w_1 + \sum_{i>1} c_i' w_i \in W$. But $1 \otimes w$ is obviously Γ-invariant, so $w \in (W_1')^{\Gamma} = W_1$. This yields $w' = 1 \otimes w \in k' \otimes_k W_1$, contrary to the hypothesis on w'. ∎

Let us record another useful descent theorem that follows by methods rather similar to the ones just used:

Theorem 6.2.10. *Let k'/k be quasi-Galois, with $\overline{k}$ denoting the algebraic closure of k in k'. Let V be a k-variety, and let $W' \subseteq V_{k'}$ be a closed k'-subvariety. Assume that there is given a subset $T' \subseteq W'(k')$ whose image in $V_{k'}(k') = V(k')$ lies in $V(\overline{k})$, and assume that the action by $\mathrm{Gal}(\overline{k}/k)$ on $V(\overline{k})$ carries T' back into itself. There exists a unique closed k-subvariety $W \subseteq V$ such that $W_{k'} = W'$ inside of $V_{k'}$.*

In this theorem, it is crucial that we assume W' is closed rather than just locally closed in $V_{k'}$. For example, suppose W' is the open complement in $V_{k'}$ of a single k'-rational point v_0' that is not $\overline{k}$-rational. If the open W' in V' descends to a locally closed k-subvariety W in V then such a descent must be an open k-subvariety, and its closed complement has to be a single closed point that is necessarily a $\overline{k}$-point that descends v_0'. This contradicts how v_0' was chosen.

Proof. In view of the arguments with descent for locally closed subvarieties in the proof of Theorem 6.2.9, all we need to check is that under the natural isomorphism $V_{k'} \simeq V_{k'}^{\gamma}$ for all $\gamma \in \Gamma = \mathrm{Aut}(k'/k)$ the closed k'-subvariety W' in V' is carried into the closed k'-subvariety ${W'}^{\gamma}$ in ${V'}^{\gamma}$ (and then by the cocycle property for the isomorphisms $V_{k'} \simeq V_{k'}^{\gamma}$ the resulting k'-maps $W' \to {W'}^{\gamma}$ are forced to be k'-isomorphisms for all $\gamma \in \Gamma$). We have to prove that the preimage of ${W'}^{\gamma}$ in $V_{k'}$ contains W'. This preimage is a closed set and so by Zariski-denseness of T' in W' it is enough to prove that this preimage contains T'. That is, we want the isomorphism $V_{k'} \simeq V_{k'}^{\gamma}$ to carry T' into ${W'}^{\gamma}$ for all $\gamma \in \Gamma$. Working on the level of k'-points (since $T' \subseteq V(k')$), this will follow if the natural action of Γ on $V(k')$ carries $T' \subseteq V(\overline{k})$ into itself. But Γ acts on $V(\overline{k})$ through its quotient $\mathrm{Gal}(\overline{k}/k)$, and as such it is a hypothesis that the $\mathrm{Gal}(\overline{k}/k)$-action on $V(\overline{k})$ carries T' into itself. ∎

For applications to Shimura varieties it is convenient to prove a variant on Theorem 6.2.10 (typically applied to the extension $\mathbf{C}/k$ for a number field k):

Theorem 6.2.11. *Let k'/k be an extension of fields. Let V be a k-variety and $W' \subseteq V_{k'}$ a locally closed k'-subvariety. Let $\{k_i\}$ be a collection of finite separable extensions of k in k' such that $\cap k_i = k$. If W' descends (necessarily uniquely) to a locally closed k_i-subvariety $W_i \subseteq V_{k_i}$ for all i then $W' = W_{k'}$ for a unique locally closed k-subvariety of V, and $W_{k_i} = W_i$ inside of V_{k_i} for all i.*

Proof. By uniqueness of descent for locally closed subvarieties, if W' descends to a locally closed k-subvariety W in V then automatically $W_{k_i} = W_i$ inside of V_{k_i} for all i (as then both W_{k_i} and W_i descend W' to a locally closed k_i-subvariety of V_{k_i}). Thus, since finitely many of the k_i's already have intersection k (as all k_i are k-finite) it follows that we may replace the given collection of k_i's with a finite subcollection having intersection k. Also, we can replace k' with the subextension generated by the k_i's (over which the W_i's have the same scalar extension, necessarily the unique descent of W' to this subfield). In this way we can assume k'/k is finite separable.

The closure of W' in $V_{k'}$ is a closed k'-subvariety in which W' is open, and likewise for the closure of each W_i in V_{k_i}. Since formation of closure of a locally closed subvariety commutes with extension of the base field, the closures of the W_i's descend the closure of W'. Hence, it suffices to separately treat the case of closed subvarieties and then open subvarieties. In the case of open subvarieties, the closed complements (given their unique "reduced" algebraic structure) may not be subvarieties but the formation of such a reduced structure is compatible with *separable* extension of the base field (as such extension preserves the property of an algebra being reduced). Hence, by the separability of k'/k (and hence of all intermediate extensions) it follows that for the case of an open subvariety the formation of the reduced structure on the closed complement is compatible with the extensions of the base field that are being considered.

To summarize, it suffices to solve the descent problem in the case of arbitrary closed subsets given their unique "reduced" algebraic structure. Working locally on V we may assume V to be affine, and so our problem becomes to show that if A is a k-algebra and $I' \subseteq k' \otimes_k A$ is an ideal that arises by scalar extension from ideals $I_i \subseteq k_i \otimes_k A$ for all i then $I' = k' \otimes_k I$ for an ideal I of A. As in the proof of Theorem 6.2.9, it suffices to solve the analogous problem in linear algebra: if F'/F is a finite separable extension of fields and W is a vector space over F equipped with an F'-subspace $W_0' \subseteq F' \otimes_F W$ that arises as a scalar extension of a subspace of $F_i \otimes_F W$ for subfields $F_i \subseteq F$ of finite degree over F such that $\cap F_i = F$ then W_0' is a scalar extension of a subspace $W_0 \subseteq W$. We may increase F' to be finite Galois over F, and by elementary Galois descent the necessary and sufficient condition for an F'-subspace W_0' of $F' \otimes_F W$ to descend to a subspace of $L \otimes_F W$ for an intermediate extension L/F inside of F' is that W_0' be stable under the natural action of $\mathrm{Gal}(F'/L)$ on $F' \otimes_F W$ through the left tensor factor. The subgroups $\mathrm{Gal}(F'/F_i)$ generate $\mathrm{Gal}(F'/F)$ since $\cap F_i = F$, so W_0' is stable under the action by $\mathrm{Gal}(F'/F)$. Hence, W_0' descends to an F-subspace $W_0 \subseteq W$ as desired. ■

There is a final variant on these descent results that eliminates the intervention of an ambient variety over the base field. However, the conclusion is also weaker: rather than give existence criteria for a unique descent, we merely give a uniqueness criterion. In practice, this is useful only to ensure that certain constructions are uniquely characterized by abstract properties that are independent of the construction process.

Theorem 6.2.12. *Let k'/k be an extension of fields and let V' be a k'-variety. Let $\{k_i\}$ be a collection of intermediate fields that are finite separable over k and satisfy $\cap k_i = k$. Assume there is given a k_i-descent V_i of V' for each i (that is, we are provided with k'-isomorphisms $j_i : V_{i/k'} \simeq V'$ for all i). There exists at most one k-descent V of V' that is compatible with the V_i's. That is, if (V, j) and $(\widetilde{V}, \widetilde{j})$ are two k-descents of V' and there are given k_i-isomorphisms $V_{k_i} \simeq V_i$ and $W_{k_i} \simeq V_i$ respectively inducing $j_i^{-1} \circ j$ and $j_i^{-1} \circ \widetilde{j}$ over k' for all i then there are unique k-isomorphisms $V \simeq W$ as k-descents of the V_i's and V'.*

Proof. Since finitely many k_i's have intersection k, by the uniqueness aspect it is enough to treat the case when the collection of k_i's is finite. As in the preceding proof, we may replace k' with the subextension generated by the k_i's so that k'/k is finite separable, and then by increasing it a bit we may assume k'/k is finite Galois. By Theorem 6.2.5, a descent V_0 of V' down to an intermediate field k_0/k is uniquely determined by the specification of a collection of k'-isomorphisms $[\gamma] : V' \simeq V'^{\gamma}$ (for all $\gamma \in \mathrm{Gal}(k'/k_0)$) such that the cocycle condition is satisfied. (Note we do not need to consider the more subtle issue of when such Galois descent is effective, as we are only aiming to prove uniqueness for a specific kind of descent rather than an existence result.)

The V_i's are therefore characterized by a system of k'-isomorphisms $[\gamma]_i : V' \simeq V'^{\gamma}$ for all $\gamma \in \mathrm{Gal}(k'/k_i)$ with the cocycle condition satisfied. The hypothetical k-descent V is likewise characterized by a system of k'-isomorphisms $[\gamma] : V' \simeq V'^{\gamma}$ for all $\gamma \in \mathrm{Gal}(k'/k)$ with the cocycle condition satisfied, and the compatibility hypothesis between V and the V_i's is precisely the condition that $[\gamma]_i = [\gamma]$ for all $\gamma \in \mathrm{Gal}(k'/k_i) \subseteq \mathrm{Gal}(k'/k)$. The uniqueness problem for (V, j) is the assertion that this final condition uniquely determines the $[\gamma]$'s for *all* $\gamma \in \mathrm{Gal}(k'/k)$. But this is obvious: the full Galois group is generated by the subgroups $\mathrm{Gal}(k'/k_i)$ since $\cap k_i = k$ and the cocycle condition implies (by induction) that the $[\gamma]$'s are determined by the special case of γ's in a generating set for the group. ■

Corollary 6.2.13. *Let k'/k be quasi-Galois, let $\overline{k}/k$ be the algebraic closure of k in k', and let $\{k_i\}$ be a collection of subfields of finite degree over k such that $\cap k_i = k$. Let V' be a k'-variety, and suppose that for each i we are given a Zariski-dense subset $S_i \subseteq V'(k')$ equipped with an action ρ_i by $\mathrm{Gal}(\overline{k}/k_i)$ with finite orbits. Up to unique k-isomorphism there exists at most one k-descent V of V' such that $V(\overline{k}) \subseteq V(k') = V'(k')$ contains the S_i's with the $\mathrm{Gal}(\overline{k}/k_i)$-action on $V(\overline{k})$ inducing ρ_i on S_i for all i.*

The method of proof (as with those above) also provides functoriality results for descent of maps. We leave it to the interested reader to carry out such formulations and verifications.

Proof. By Theorem 6.2.9, the V_{k_i}'s are uniquely determined up to unique k_i-isomorphism. Theorem 6.2.12 therefore provides the uniqueness of V up to unique k-isomorphism. ■

6.3. **Heegner points and classical CM theory.** We wish to use the Main Theorem of CM for elliptic curves to describe the Galois theoretic action on a Zariski-dense set of CM points of the modular curves $Y_0(n)$ over $\mathbf{Q}$ that we discussed in §6.1. More precisely, for an infinite set of imaginary quadratic fields K (depending on n) and embeddings $\varphi : K \hookrightarrow \mathbf{C}$ we wish to exhibit an infinite set of points in $Y_0(n)(\overline{\mathbf{Q}})$ whose action under $\mathrm{Gal}(\overline{\mathbf{Q}}/K)$ can be described explicitly. It follows from Theorem 6.2.9 that such data will uniquely determined each $Y_0(n)_K$ over K, and so by Corollary 6.2.13 even the $\mathbf{Q}$-structure on $Y_0(n)$ is thereby uniquely determined.

Consider an imaginary quadratic field $K = \mathbf{Q}(\sqrt{D})$ with discriminant D such that $(n, D) = 1$ and the prime factors of n are split in K. By Chebotarev's density theorem, there are infinitely many such K and we may take D to be relatively prime to any desired nonzero integer. We fix an embedding $\varphi : K \to \mathbf{C}$. Example 1.3.11 provides Heegner points on $Y_0(n)$ with CM order (in $\mathscr{O}_K$) whose conductor relatively prime to nD. As we vary the conductor f we get an infinite set of points. Each such point considered as a $\mathbf{C}$-point on $Y_0(n)$ turns out to be an algebraic point with field of definition given by a so-called *ring class field* $H_{K,f}$ of K with conductor f; the ring class field is a certain canonical finite abelian extension provided by class field theory for K, equipped with a natural isomorphism from $\mathrm{Gal}(H_{K,f}/K)$ onto the finite class group of the quadratic order in $\mathscr{O}_K$ with conductor f. In the case $f = 1$ this ring class field is the Hilbert class field $H = H_K$ of K. The Main Theorem of CM for elliptic curves allows one to explicitly describe the action by $\mathrm{Gal}(H_{K,f}/K)$ on n-isogeny Heegner points for K with conductor f. We will give the details only in the case $f = 1$, largely to avoid technical complications inherent in the statement of the Main Theorem of CM for elliptic curves with non-maximal CM order; however, upon formulating the Main Theorem for non-maximal orders the method used below can be adapted to the case of any order.

We normalize the isomorphism $\mathrm{Gal}(H/K) \simeq \mathrm{Cl}_K$ from class field theory so that a maximal ideal $\mathfrak{p}$ of $\mathscr{O}_K$ is carried to the corresponding arithmetic Frobenius element (acting on $\mathscr{O}_H/\mathfrak{P}$ by raising to the $\mathbf{N}\mathfrak{p}$th-power for any prime $\mathfrak{P}$ of $\mathscr{O}_H$ over $\mathfrak{p}$). In general, for a fractional ideal $\mathfrak{b}$ of K we let $\sigma_{\mathfrak{b}}$ denote the corresponding element of $\mathrm{Gal}(H/K)$.

To fix our notation, we begin by recalling the statements of the Main Theorem of Complex Multiplication that we shall require. To keep matters concrete in the setting of elliptic curves, we shall use [L] as our reference for the development of the classical CM theory for elliptic curves (including arbitrary orders in imaginary quadratic fields). Unfortunately, the techniques in [L] rest on rather hands-on arguments with elliptic functions and Weierstrass equations, so they do not generalize to the higher-dimensional case (though many of the results do generalize, via more conceptual methods). In order to state the main theorems of CM for elliptic curves (with maximal CM order), we need to recall some general idelic notation.

Let K be a number field (only the imaginary quadratic case is relevant for us) and $\mathfrak{a}$ a fractional ideal of K. For any idele s of K, we let (s) denote the corresponding fractional ideal of K (with multiplicity $\mathrm{ord}_v(s_v)$ at each $v \nmid \infty$). We have a canonical decomposition

$$K/\mathfrak{a} \simeq \bigoplus_{v \nmid \infty} K_v/\mathfrak{a}_v,$$

and multiplication by s_v on the vth factor sets up an $\mathscr{O}_K$-linear isomorphism

$$s : K/\mathfrak{a} \simeq K/(s)\mathfrak{a}.$$

In the special case when s is a principal idele with generator α then this map is just the multiplication map

$$\alpha : K/\mathfrak{a} \simeq K/\alpha\mathfrak{a}.$$

Another important special case is when s_v is a unit for all $v \nmid \infty$. This is the case when the Artin symbol $(s|K) \in \mathrm{Gal}(K^{\mathrm{ab}}/K)$ acts *trivially* on the Hilbert class field of K (essentially by the idelic description of the ideal class group and the idelic formulation of class field theory for K). In this case, $(s)\mathfrak{a} = \mathfrak{a}$ but the map $s : K/\mathfrak{a} \to K/\mathfrak{a}$ is rather far from the identity in general (we will see that in the motivating imaginary quadratic case relative to an algebraic model for $\mathbf{C}/\mathfrak{a}$ over H, this action essentially corresponds to the action of $(s|K)$ on $(\mathbf{C}/\mathfrak{a})_{\mathrm{tors}}$, whose coordinates all lie in K^{ab}).

By [L, Ch. 10; §1, Thm. 2; §3, Thm. 5, Rem. 1(p. 133)]) we arrive at the following special case of the first main theorem of CM for elliptic curves:

Theorem 6.3.1. *Let $K \subseteq \mathbf{C}$ be an imaginary quadratic subfield of $\mathbf{C}$, and $\mathfrak{a}$ a fractional ideal of K. Let $H \subseteq \mathbf{C}$ be the Hilbert class field of K. Then*

$$j(\mathbf{C}/\mathfrak{a}) \in H$$

and this j-invariant generates H over K.

Moreover, for any fractional ideal $\mathfrak{b}$ of K, we have

$$\sigma_{\mathfrak{b}}(j(\mathfrak{a})) = j(\mathfrak{a}\mathfrak{b}^{-1}).$$

If we had decided to use Deligne's normalization for the Artin isomorphism (under which local uniformizers correspond to geometric Frobenius elements), there would be $\mathfrak{b}$ rather than $\mathfrak{b}^{-1}$ in the final identity of the theorem.

Corollary 6.3.2. *For any fractional ideal $\mathfrak{a}$ of K, the elliptic curve $\mathbf{C}/\mathfrak{a}$ admits a Weierstrass model over the subfield $H \subseteq \mathbf{C}$.*

Proof. It is a basic fact from the theory of elliptic curves that an elliptic curve over an algebraically closed field admits a Weierstrass equation whose coefficients lie in the subfield generated over the prime field by the j-invariant. ■

A special case of the second main theorem of CM is the following [L, Ch. 10, §2, Thm. 3]:

Theorem 6.3.3. *Let $K \subseteq \mathbf{C}$ be an imaginary quadratic field, and $H \subseteq \mathbf{C}$ its Hilbert class field. Let $\varphi : \mathbf{C}/\mathfrak{a} \simeq A$ be an analytic isomorphism onto a Weierstrass model over $\mathbf{C}$. Let $\sigma \in \mathrm{Aut}(\mathbf{C}/K)$ be a K-automorphism, and s an idele of K with $(s|K) = \sigma|_{K^{\mathrm{ab}}}$. Let A^σ denote the Weierstrass model obtained from applying σ to the coefficients, so the set-theoretic map $(x, y) \mapsto (\sigma(x), \sigma(y))$ induces an isomorphism of abelian groups $A \simeq A^\sigma$.*

Then there exists a unique analytic isomorphism $\psi_\sigma : \mathbf{C}/(s)^{-1}\mathfrak{a} \simeq A^\sigma$ such that the diagram

(6.3.1)
$$\begin{array}{ccccc} K/\mathfrak{a} & \longrightarrow & \mathbf{C}/\mathfrak{a} & \xrightarrow[\simeq]{\varphi} & A \\ {\scriptstyle s^{-1}}\downarrow & & & & \downarrow{\scriptstyle \sigma} \\ K/(s)^{-1}\mathfrak{a} & \longrightarrow & \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1} & \xrightarrow[\psi_\sigma]{\simeq} & A^\sigma \end{array}$$

commutes. Here, the maps $K/\Lambda \to \mathbf{C}/\Lambda$ are the canonical inclusions.

If we had used Deligne's normalization for the Artin isomorphism, then s^{-1} would be replaced with s in both appearances on the left side of (6.3.1).

The proof of this theorem in [L] involves a detailed study of Weber functions that make torsion explicit in terms of Weierstrass theory. The key issue is that $K/\mathfrak{a}$ is exactly the torsion subgroup of $\mathbf{C}/\mathfrak{a}$ (and likewise for the bottom row of (6.3.1) using $(s)^{-1}\mathfrak{a}$), so the commutativity of the diagram certainly requires an understanding of torsion in terms of analytic Weierstrass models. It is perhaps also worth stressing that the proof of Theorem 6.3.3 proceeds in a series of stages, where one first proves a weaker result for arithmetic Frobenius elements away from some "bad" finite set and then an indirect procedure extends things to the case of general σ.

Observe that uniqueness of ψ_σ is immediate from the denseness of $(\mathbf{C}/\Lambda)_{\rm tors}$ inside of $\mathbf{C}/\Lambda$. Note that there is no meaningful way to fill in a middle arrow in (6.3.1) making a commutative square on either the left or the right. Thus, it is only the rectangle (6.3.1) which really makes sense.

In the special case where we choose the Weierstrass model to have coefficients in H (which can be done, thanks to Corollary 6.3.2) and σ is the identity on H, then $A^\sigma = A$ and $\sigma|_{K^{\rm ab}} = (s|K)$ for an idele s which is a *unit* at all finite places. Thus, as we noted earlier, we then have $(s)^{-1}\mathfrak{a} = \mathfrak{a}$ for all fractional ideals $\mathfrak{a}$ of K. Putting this all together, we get the following consequence of the preceding theorem:

Corollary 6.3.4. *Let $K \subseteq \mathbf{C}$ be an imaginary quadratic subfield with Hilbert class field H. Let $\mathfrak{a}$ be a fractional ideal of K and $\varphi : \mathbf{C}/\mathfrak{a} \simeq A$ an analytic isomorphism onto a Weierstrass model with coefficients in H. For any $\sigma \in {\rm Aut}(\mathbf{C}/H)$ and idele s of K with $\sigma|_{K^{\rm ab}} = (s|K)$ and $s_v \in \mathscr{O}_v^\times$ for all $v \nmid \infty$, there is a unique analytic isomorphism $\psi_\sigma : \mathbf{C}/\mathfrak{a} \simeq A$ such that the diagram*

$$\begin{array}{ccccc} K/\mathfrak{a} & \longrightarrow & \mathbf{C}/\mathfrak{a} & \xrightarrow[\simeq]{\varphi} & A \\ {\scriptstyle s^{-1}}\downarrow & & & & \downarrow{\scriptstyle \sigma} \\ K/\mathfrak{a} & \longrightarrow & \mathbf{C}/\mathfrak{a} & \xrightarrow[\psi_\sigma]{\simeq} & A \end{array}$$

commutes, where the right vertical map is $(x, y) \mapsto (\sigma(x), \sigma(y))$ (which carries A to itself since σ acts as the identity on the coefficients in H).

Remark 6.3.5. Clearly $\psi_\sigma = \varphi \circ \xi_\sigma$ for an analytic automorphism ξ_σ of $\mathbf{C}/\mathfrak{a}$, which is to say $\xi_\sigma \in \mathscr{O}_K^\times \subseteq \mathbf{C}$.

This completes our overview of the classical CM theory of elliptic curves.

6.4. **Applications of CM theory to algebraicity.** We let K be an imaginary quadratic subfield of $\mathbf{C}$ and $\mathfrak{a}$ a fractional ideal of K. Let $\mathfrak{b}$ be a (non-zero) integral ideal of K, so the finite $\mathscr{O}_K$-module

$$(\mathbf{C}/\mathfrak{a})[\mathfrak{b}] \subseteq \mathbf{C}/\mathfrak{a}$$

makes sense, and is the kernel of the natural finite projection map

$$\mathbf{C}/\mathfrak{a} \to \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1}.$$

A special case of this situation is that described by a Heegner point. We want to prove the following theorem:

Theorem 6.4.1. *With notation as above, and $H \subseteq \mathbf{C}$ the Hilbert class field of K, the situation*

$$\mathbf{C}/\mathfrak{a} \to \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1}$$

can be "defined over H".

This can be interpreted in two (equivalent) ways. On the one hand, we can take this to mean that projective models (even Weierstrass models) can be found for the two elliptic curves using H-coefficients in such a way that the analytic projection map can be described (locally for the Zariski topology relative to H-coefficients) in terms of polynomial equations with H-coefficients. But clearly such an *ad hoc* point of view is inadequate for proving anything serious. In more sophisticated terms, we are claiming that there exists an isogeny $E \to E'$ of elliptic curves over H whose base change to $\mathbf{C}$ (and resulting analytification) is isomorphic to the given analytic isogeny.

In order to verify statements of this sort, we will use the following basic fact which is a special case of the descent discussion in §6.2.9.

Lemma 6.4.2. *Let F be a field of characteristic 0 and let F'/F an algebraically closed extension field. Let E be an elliptic curve over F and $\pi' : E_{/F'} \to E'_1$ an isogeny of elliptic curves over F'. Then there exists an isogeny*

$$\pi : E \to E_1$$

over F whose base change to F' is isomorphic to π' if and only if the finite subgroup

$$\ker(\pi')(F') \subseteq E(F')$$

is stable under the "coordinate-wise" action of $\mathrm{Aut}(F'/F)$ on $E(F')$.

Now we can prove Theorem 6.4.1:

Proof. By the algebraicity of the category of complex analytic elliptic curves, the analytic diagram

$$\mathbf{C}/\mathfrak{a} \to \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1}$$

can be realized as the analytification of an algebraic isogeny between elliptic curves over $\mathbf{C}$. Moreover, we have already noted that $\mathbf{C}/\mathfrak{a}$ may be realized by a Weierstrass equation E with coefficients in H. Let $\varphi : \mathbf{C}/\mathfrak{a} \simeq E(\mathbf{C})$ be an analytic isomorphism.

By the preceding lemma we just have to show that the subgroup $\varphi(\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a})$ in $E(\mathbf{C})$ is stable under the action of $\mathrm{Aut}(\mathbf{C}/H)$. Choose $\sigma \in \mathrm{Aut}(\mathbf{C}/H)$. In the notation of Corollary 6.3.2 and the subsequent remark, we have

$$\sigma(\varphi(\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a})) = \varphi(\xi_\sigma((\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a}))),$$

where ξ_σ is multiplication by some element in $\mathscr{O}_K^\times$. The right side is visibly $\varphi((\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a}))$, so we're done. ∎

We now use the arithmetic theory of modular curves, or more specifically the theory in §6.1 for realizing classical analytic modular curves as arising from algebraic curves over $\mathbf{Q}$. From now on, we write $Y_0(n)$ to denote the canonical algebraic curve over $\mathbf{Q}$ (whose analytified base change to $\mathbf{C}$ is naturally analytically identified with the classical analytic construction denoted $Y_0(n)^{\mathrm{an}}$). This is a smooth irreducible affine curve over $\mathbf{Q}$. It is an unfortunate fact of life that one cannot view $Y_0(n)$ as a fine moduli scheme, but the theory that makes this algebraic curve over $\mathbf{Q}$ provides the following:

Theorem 6.4.3. *Let F be a subfield of $\mathbf{C}$ and $E \to E'$ a cyclic n-isogeny of elliptic curves over F. Then the point on $Y_0(n)^{\mathrm{an}}$ corresponding to the analytic data $E(\mathbf{C}) \to E'(\mathbf{C})$ is induced by a point in $Y_0(n)(F)$.*

Combining this fact with Theorem 6.4.1, we obtain:

Corollary 6.4.4. *Let K be an imaginary quadratic field, and n a positive integer all of whose prime factors are totally split in K. Then all Heegner points on $Y_0(n)^{\mathrm{an}}$ arise from points in $Y_0(n)(H)$.*

Since $\mathrm{Gal}(H/K)$ acts on the set $Y_0(n)(H)$ (as $Y_0(n)$ is an algebraic curve over K, or even $\mathbf{Q}$), it makes sense (by Corollary 6.4.4) to ask how this group acts on Heegner points. Can the action be described in terms of the analytic data $([\mathfrak{a}], \mathfrak{n})$ which we have used to describe Heegner points? Let us see how it goes.

Theorem 6.4.5. *Let $\sigma \in \mathrm{Gal}(H/K)$ correspond to a fractional ideal class $[\mathfrak{b}]$ under the class field theory isomorphism $\mathrm{Gal}(H/K) \simeq \mathrm{Cl}_K$. Then as points in $Y_0(n)(H)$ we have*

$$\sigma([\mathfrak{a}], \mathfrak{n}) = ([\mathfrak{a}\mathfrak{b}^{-1}], \mathfrak{n}).$$

Of course, we could describe the equality in this theorem purely in terms of ideal classes rather than in terms of representative elements, but the distinction is minor and in practice we'll certainly be computing with representatives anyway. A more subtle point is to observe that the statement of the theorem certainly depends on how we define the Artin isomorphism of class field theory. We have made the convention to normalize the isomorphism to associated local uniformizers to arithmetic Frobenius elements because this is the convention used in [L] (to which we have referred for the statements of the main theorems of CM). If we had adopted Deligne's preference to associated local uniformizers with geometric Frobenius then the equality in the theorem would have $\mathfrak{b}$ rather than $\mathfrak{b}^{-1}$ on the right side.

Proof. We first choose an isogeny $\pi : E \to E'$ of elliptic curves over H which induces a point in $Y_0(n)(H) \subseteq Y_0(n)(\mathbf{C}) = Y_0(n)^{\mathrm{an}}$ that corresponds to the Heegner point data $\mathbf{C}/\mathfrak{a} \to \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1}$. We let

$$\varphi : \mathbf{C}/\mathfrak{a} \simeq E(\mathbf{C}),\ \ \varphi' : \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1} \simeq E'(\mathbf{C})$$

denote corresponding compatible analytic isomorphisms making the diagram

$$\begin{array}{ccc} \mathbf{C}/\mathfrak{a} & \xrightarrow{\varphi} & E(\mathbf{C}) \\ \downarrow & & \downarrow{\scriptstyle \pi_{\mathbf{C}}} \\ \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1} & \xrightarrow[\varphi']{} & E'(\mathbf{C}) \end{array}$$

commute, where the left column is the canonical projection.

Fix a lifting of σ to a K-automorphism of $\mathbf{C}$, again denoted σ, and choose an idele s of K for which $(s|K) = \sigma|_{K^{\mathrm{ab}}}$, so $(s) = \mathfrak{b}$ by the definition of $\mathfrak{b}$ (and by our convention for normalizing the Artin isomorphism). By Theorem 6.3.3, there are unique analytic isomorphism

$$\psi_\sigma : \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1} \simeq E^\sigma(\mathbf{C})$$

and

$$\psi'_\sigma : \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1} \simeq E'^\sigma(\mathbf{C})$$

which fit into respective commutative diagrams of groups

$$\begin{array}{ccccc} K/\mathfrak{a} & \longrightarrow & \mathbf{C}/\mathfrak{a} & \xrightarrow[\simeq]{\varphi} & E(\mathbf{C}) \\ {\scriptstyle s^{-1}}\downarrow & & & & \downarrow{\scriptstyle \sigma} \\ K/\mathfrak{a}\mathfrak{b}^{-1} & \longrightarrow & \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1} & \xrightarrow[\psi_\sigma]{\simeq} & E^\sigma(\mathbf{C}) \end{array}$$

and

$$\begin{array}{ccccc} K/\mathfrak{a}\mathfrak{n}^{-1} & \longrightarrow & \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1} & \xrightarrow[\simeq]{\varphi} & E'(\mathbf{C}) \\ {\scriptstyle s^{-1}}\downarrow & & & & \downarrow{\scriptstyle \sigma} \\ K/\mathfrak{a}\mathfrak{n}^{-1}\mathfrak{b}^{-1} & \longrightarrow & \mathbf{C}/\mathfrak{a}\mathfrak{n}^{-1}\mathfrak{b}^{-1} & \xrightarrow[\psi'_\sigma]{\simeq} & E'^\sigma(\mathbf{C}) \end{array}$$

Now consider the analytic diagram

(6.4.1)
$$\begin{array}{ccc} \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1} & \xrightarrow[\simeq]{\psi'_\sigma} & E^\sigma(\mathbf{C}) \\ \downarrow & & \downarrow{\scriptstyle (\pi^\sigma)_{\mathbf{C}}} \\ \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1}\mathfrak{n}^{-1} & \xrightarrow[\psi'_\sigma]{\simeq} & E'^\sigma(\mathbf{C}) \end{array}$$

where the left column in the canonical projection. Here is the crucial point: this diagram *commutes*. To check such commutativity, by *continuity* it suffices to compose everything back to the dense torsion subgroup

$K/\mathfrak{a}\mathfrak{b}^{-1} \subseteq \mathbf{C}/\mathfrak{a}\mathfrak{b}^{-1}$, and even to compose with the isomorphism

$$s^{-1} : K/\mathfrak{a} \simeq K/\mathfrak{a}\mathfrak{b}^{-1}.$$

But if we use the trivial commutativity of general diagrams

$$\begin{array}{ccc} K/\mathfrak{c} & \longrightarrow & \mathbf{C}/\mathfrak{c} \\ \downarrow & & \downarrow \\ K/\mathfrak{c}' & \longrightarrow & \mathbf{C}/\mathfrak{c}' \end{array}$$

(for fractional ideals $\mathfrak{c} \subseteq \mathfrak{c}'$) and

$$\begin{array}{ccc} K/\mathfrak{c} & \xrightarrow[\simeq]{s^{-1}} & K/(s)^{-1}\mathfrak{c} \\ \downarrow & & \downarrow \\ K/\mathfrak{c}' & \xrightarrow[s^{-1}]{\simeq} & K/(s)^{-1}\mathfrak{c}' \end{array}$$

(for fractional ideals $\mathfrak{c} \subseteq \mathfrak{c}'$), then everything fits together rather easily, thanks to the characterizing properties of ψ_σ and ψ'_σ when composed back to torsion. It is too painful for me to type the corresponding 3-dimensional diagram which puts everything together, so we leave it as a pleasant exercise to see how the jigsaw puzzle works.

With the commutativity of (6.4.1) settled, we look at what this diagram says! The left side of (6.4.1) is exactly the data corresponding to the Heegner point $([\mathfrak{a}\mathfrak{b}^{-1}], \mathfrak{n})$. The right side of (6.4.1) is exactly the analytification of the base change to $\mathbf{C}$ of the action of $\sigma \in \mathrm{Gal}(H/K)$ on the original H-model $E \to E'$ in $Y_0(n)(H)$ for our initial Heegner point $([\mathfrak{a}], \mathfrak{n})$. Hence, the commutativity of (6.4.1) with analytic isomorphisms across the horizontal directions says exactly that $\sigma \in \mathrm{Gal}(H/K)$ acting on $Y_0(n)(H)$ takes $([\mathfrak{a}], \mathfrak{n})$ over to $([\mathfrak{a}\mathfrak{b}^{-1}], \mathfrak{n})$, as desired. ■

7. Siegel moduli spaces as Shimura varieties

In §6 we saw some examples of 1-dimensional analytic moduli spaces for abelian varieties, and in the case of certain modular curves we used methods of Shimura (based on the Main Theorem of Complex Multiplication for elliptic curves) to characterize specific algebraic models for such curves over $\mathbf{Q}$. We now wish to turn our attention to a higher-dimensional example by revisiting Siegel's construction from §3.2. By passing to quotients by suitable discrete group actions (generalizing the procedure used for modular curves) we shall make analytic Siegel modular varieties. These are solutions to complex-analytic moduli problems that make sense in the algebraic theory, and in this way the Siegel modular varieties conceptually acquire an algebraic structure over $\mathbf{C}$ (that is moreover unique by work of Baily and Borel). The relevant moduli problems will make sense over number fields, and work of Mumford [GIT] provides models over over number fields that solve these moduli problems. Our aim will be to show how, via the Main Theorem of Complex Multiplication for abelian varieties, the models made by Mumford have Galois action at CM points that can be described within the analytic theory. The significance of Shimura's approach to canonical models is that it gives a way to characterize algebraic models over number fields in situations when there is no available moduli problem for abelian varieties. We will make essential use of the moduli-theoretic properties of Mumford's models in our proof that they are canonical models in the sense of Shimura.

7.1. **Analytic Siegel families and moduli spaces.** Fix $g \geq 1$ and positive integers $d_1, \dots, d_g$ satisfying $d_1|d_2|\dots|d_g$. Let $\mathbf{d}$ be the diagonal $g \times g$ matrix with jj-entry d_j for $1 \leq j \leq g$. Also fix a basis $2\pi i$ for $\mathbf{Z}(1)$. Let $\mathfrak{h}_{g,i}$ be the Siegel half-space given by the open complex manifold inside the vector space of symmetric $g \times g$ matrices consisting of those Z for which the real symmetric matrix $(Z - \overline{Z})/2\pi i$ is positive-definite (and hence invertible). For each such Z the imaginary component Z_{im} is invertible, and so $(Z\mathbf{d})_{\mathrm{im}} = Z_{\mathrm{im}}\mathbf{d}$ is also invertible. Hence, $X_{Z,\mathbf{d}} = \mathbf{C}^g/(Z\mathbf{d}\mathbf{Z}^g + \mathbf{Z}^g)$ is a complex torus. Recall from §3.2 that there is a unique polarization $\phi_{Z,\mathbf{d}} : X_{Z,\mathbf{d}} \to X_{Z,\mathbf{d}}^\vee$ whose Riemann form has associated matrix

$$2\pi i \begin{pmatrix} 0 & -\mathbf{d} \\ \mathbf{d} & 0 \end{pmatrix} \tag{7.1.1}$$

with respect to the ordered homology basis

$$\iota_{Z,\mathbf{d}} : \mathbf{Z}^{2g} \simeq \mathrm{H}_1(X_{Z,\mathbf{d}}, \mathbf{Z})$$

corresponding to $\{Z\mathbf{d}(e_1), \dots, Z\mathbf{d}(e_g), e_1, \dots, e_g\} \subseteq \mathbf{C}^g$, and that *every* pair (X, ϕ) consisting of a polarized g-dimensional complex torus whose skew-symmetric Riemann form has invariant factors $\{d_1, \dots, d_g\}$ arises via this construction. One can be more precise: if we consider triples (X, ϕ, ι) where X is a g-dimensional complex torus, $\phi : X \to X^\vee$ is a polarization, and $\iota : \mathbf{Z}^{2g} \simeq \mathrm{H}_1(X, \mathbf{Z})$ is an ordered homology basis with respect to which the Riemann form of the polarization has matrix (7.1.1) then (via an evident notion of isomorphism for such triples) we have that (X, ϕ, ι) is *uniquely* isomorphic to a triple $(X_{Z,\mathbf{d}}, \phi_{Z,\mathbf{d}}, \iota_{Z,\mathbf{d}})$ for a *unique* $Z \in \mathfrak{h}_{g,i}$. In other words, the set $\mathfrak{h}_{g,i}$ is in "natural" bijection with the set of isomorphism classes of such triples (X, ϕ, ι). But what does "natural" mean? We wish to explain Grothendieck's relative viewpoint that gives it a *useful* mathematical meaning, and moreover paves the way for algebraic and arithmetic reformulation.

We shall first build a global structure ("analytic family of complex tori") over $\mathfrak{h}_{g,i}$ that satisfies a universal property (recovering Example 6.1.4 in the case $g = 1$), and then we shall use the universal property to conceptually analyze some structures on $\mathfrak{h}_{g,i}$ that may be introduced via "bare hands" matrix manipulations but lead to messy calculations when studied in such a naive manner. An elementary example of this dichotomy in viewpoints is the classical linear fractional action of $\mathrm{SL}_2(\mathbf{Z})$ on a half-plane $\mathfrak{h} \subseteq \mathbf{C} - \mathbf{R}$ and the conceptual interpretation of this action provided in the discussion preceding Example 6.1.7. It is important to observe that the discussion there adopted the perspective of moving lattices within $\mathbf{C}$, and this gave no insight into the more general linear fractional action by $\mathrm{SL}_2(\mathbf{R})$ on $\mathfrak{h}$. A conceptual understanding of the meaning of this latter action requires the viewpoint of variation of complex structure (as in the discussion following Example 1.1.7), and we shall see this worked out in the case of arbitrary $g \geq 1$ in our study of $\mathfrak{h}_{g,i}$ with an action by the symplectic group $\mathrm{Sp}_{2g}(\mathbf{R})$ (recovering the classical situation with $\mathrm{SL}_2(\mathbf{R})$ and $\mathfrak{h}$ upon setting $g = 1$).

Before we construct a universal analytic family of complex tori over $\mathfrak{h}_{g,i}$, let us first define the general concept of an analytic family of tori (recovering Definition 6.1.1 if $g = 1$).

Definition 7.1.1. Let S be a complex manifold. A *complex torus over S* (or an *analytic family of complex tori parameterized by S*) is a proper submersion of complex manifolds $f : A \to S$ equipped with a section $e : S \to A$ such that each fiber $(A_s, e(s))$ admits a (necessarily unique) structure of complex torus and the group laws vary analytically in $s \in S$. This final condition means that the fiberwise inversion map $A \to A$ over S (restricting to inversion on each A_s) is an analytic map and for the locally closed complex submanifold

$$A \times_S A := \{(a, a') \in A \times A \,|\, f(a) = f(a')\} = \cup_{s \in S}(A_s \times A_s) \subseteq A \times A$$

the fiberwise group law $m : A \times_S A \to A$ over S (restricting to the group law on each A_s) is an analytic map.

If (A, e) and (A', e') are two analytic families of complex tori over S, a *map of S-tori* $f : A \to A'$ is an analytic map over S such that $f \circ e = e'$ (or, equivalently, such that $f_s : A_s \to A'_s$ is a map of Lie groups for all $s \in S$).

For any such (A, e) over S the fiber-dimension function $s \mapsto \dim A_s$ is locally constant on S, so we usually take it to be a constant $g \geq 1$. This case is called a complex torus *with relative dimension g* over S. Given such an (A, e), the exponential and double-coset uniformizations in the classical setting (as in §1) do admit analogues over S. One has to replace vector spaces with holomorphic vector bundles and lattices with local systems of finite free $\mathbf{Z}$-modules. The most natural way to develop such a classification is to first relativize classical Lie theory in families, but this is too lengthy a task to explain here so we leave it to the imagination of the interested reader to work out how this goes. This relativization of the classical theory leads to the most natural proof of Theorem 7.1.8 below.

Example 7.1.2. Fix $\mathbf{d}$ as above, and impose an equivalence relation on $\mathbf{C}^g \times \mathfrak{h}_{g,i}$ via the condition $(w, Z) \sim (w', Z')$ if and only if $Z = Z'$ and $w - w' \in \Lambda_{Z,\mathbf{d}} = Z\mathbf{d}\mathbf{Z}^g + \mathbf{Z}^g$. (Geometrically, on the fiber over each $Z \in \mathfrak{h}_{g,i}$ we are imposing the condition of congruence modulo the lattice $\Lambda_{Z,\mathbf{d}}$.) The quotient $\mathscr{A}_{g,\mathbf{d},i}$ of $\mathbf{C}^g \times \mathfrak{h}_{g,i}$ modulo this equivalence relation admits unique topology with respect to which the projection from $\mathbf{C}^g \times \mathfrak{h}_{g,i}$ onto $\mathscr{A}_{g,\mathbf{d},i}$ is a covering map and a unique complex-analytic manifold structure with respect to which this projection is a local analytic isomorphism. With this structure the map $\mathscr{A}_{g,\mathbf{d},i} \to \mathfrak{h}_{g,i}$ is a proper submersion, and composition of the zero-section of $\mathbf{C}^g \times \mathfrak{h}_{g,i} \to \mathfrak{h}_{g,i}$ with the projection to $\mathscr{A}_{g,\mathbf{d},i}$ defines an analytic section $e : \mathfrak{h}_{g,i} \to \mathscr{A}_{g,\mathbf{d},i}$ so that $(\mathscr{A}_{g,\mathbf{d},i}, e)$ is a complex torus over $\mathfrak{h}_{g,i}$. On the fiber over each $Z \in \mathfrak{h}_{g,i}$ this fibral torus is precisely $X_{Z,\mathbf{d}}$. The complex torus $\mathscr{A}_{g,\mathbf{d},i}$ over $\mathfrak{h}_{g,i}$ is a "gluing" of the $X_{Z,\mathbf{d}}$'s; it is called the *Siegel torus* (or *Siegel family*) *of type* $(\mathbf{d}, i)$ over the Siegel half-space $\mathfrak{h}_{g,i}$.

In the case $g = 1$ this recovers the family of elliptic curves $\mathscr{E} \to \mathfrak{h}_i = \mathfrak{h}_{1,i}$ from §6.1. There is no higher-dimensional analogue of the explicit "algebraic" model $\mathscr{E}^{\mathrm{alg}} \to \mathfrak{h}_i$ because the explicit nature of the Weierstrass theory for elliptic curves has no higher-dimensional analogue.

We want to glue the polarization and homology trivialization on the fibers $X_{Z,\mathbf{d}}$, so we first have to define notions of polarization and homology trivialization for analytic families of complex tori. There are several ways to do this, and we choose a method of definition that is well-suited to our expository purposes (but is not well-suited to giving natural proofs of subsequent assertions):

Definition 7.1.3. Let (A, e) be a complex torus over a complex manifold S. A *fibral homology trivialization* on (A, e) over S is a choice of ordered basis $\{\lambda_1(s), \ldots, \lambda_{2g}(s)\}$ of $\mathrm{H}_1(A_s, \mathbf{Z})$ for all $s \in S$ such that under the canonical isomorphism of homologies on nearby fibers A_s and $A_{s'}$ (via parallel transport along small paths in the base S) the homology class $\lambda_j(s)$ is transported to the homology class $\lambda_j(s')$ for all $1 \le j \le 2g$.

Example 7.1.4. Consider the Siegel torus $\mathscr{A}_{g,\mathbf{d},i} \to \mathfrak{h}_{g,i}$ of type $(\mathbf{d}, i)$. On the fiber over each $Z \in \mathfrak{h}_{g,i}$ we get the complex torus

$$X_{Z,\mathbf{d}} = \mathbf{C}^g/\Lambda_{Z,\mathbf{d}}$$

whose homology $\Lambda_{Z,\mathbf{d}} = Z\mathbf{d}\mathbf{Z}^g + \mathbf{Z}^g$ has trivialization $\{Z\mathbf{d}(e_1), \ldots, Z\mathbf{d}(e_g), e_1, \ldots, e_g\}$. This is a fibral homology trivialization in the sense of the preceding definition.

The notion of correspondence has an analogue in families by using proper maps in the role of compact manifolds and using sections in the role of base points:

Definition 7.1.5. Let $f : M \to S$ and $f' : M' \to S$ be proper submersions between complex manifolds, and assume that all fibers are connected. Let $m_0 : S \to M$ and $m_0' : S \to M$ be analytic sections, and let $M \times_S M'$ be the "fiber product" manifold: this is the locally closed submanifold in $M \times M'$ given by

$$M \times_S M' = \{(m, m') \in M \times M' \,|\, f(m) = f'(m')\} = \cup_{s \in S} M_s \times M'_s \hookrightarrow M \times M'.$$

A *relative correspondence* between (M, m_0) and (M', m_0') over S is a line bundle $L \to M \times_S M'$ equipped with trivializations $i : (1_M \times_S m_0')^*(L) \simeq \mathbf{C} \times M$ and $i' : (m_0 \times_S 1_{M'})^*(L) \simeq \mathbf{C} \times M'$ along the locally closed submanifolds $1_M \times_S m_0' : M \hookrightarrow M \times_S M'$ and $m_0 \times_S 1_{M'} : M' \hookrightarrow M \times_S M'$ such that $m_0^*(i) = {m_0'}^*(i')$ as trivializations of the line bundle $(m_0 \times_S m_0')^*(L) \to S$ over S.

This definition may look complicated, but all it is saying is that for the line bundle L over $M \times_S M'$ the pullbacks L_s on each $M_s \times M'_s$ (for $s \in S$) are endowed with trivializations making each L_s a correspondence between the compact connected marked manifolds $(M_s, m_0(s))$ and $(M'_s, m_0'(s))$ such that the trivializations over s-fibers "vary analytically" in s. In the "absolute" case when S is a single point, a relative correspondence is a correspondence in the earlier sense. It follows from the absolute case that relative correspondences admit no non-trivial automorphisms.

Recall that the notion of a correspondence between complex tori (with the origins as the marked points) is equivalent to a certain kind of $\mathbf{Z}(1)$-valued homology lattice pairing (the Riemann form, whose $\mathbf{C}$-scalar extension satisfies some orthogonality restrictions with respect to the complex structures). A similar dictionary works in the relative setting, but the justification requires a long digression into the relative version of the Appell-Humbert theorem so we shall just state the result and illustrate it with an example. If (A, e)

and (A', e') are complex tori over a complex manifold S, then we may define a *relative homology pairing* between (A, e) and (A', e') to be a collection of bilinear pairings $\psi = \{\psi_s\}_{s \in S}$ with each ψ_s a $\mathbf{Z}(1)$-valued pairing between the homology lattices of A_s and A'_s such that ψ_s has continuous dependence on s. That is, for nearby points s_1 and s_2 in S, the canonical isomorphisms $\mathrm{H}_1(A_{s_j}, \mathbf{Z}) \simeq \mathrm{H}_1(A'_{s_j}, \mathbf{Z})$ carry ψ_{s_1} to ψ_{s_2}. For example, if we are given a relative correspondence $L \to A \times_S A'$ between (A, e) and (A', e') over S then the collection of Riemann forms $\psi_L = \{\psi_{L_s}\}$ is a relative homology pairing (called the *relative Riemann form* of the relative correspondence). It can be shown that, conversely, if a relative homology pairing $\psi = \{\psi_s\}$ has each ψ_s arising as the Riemann form of a correspondence between $(A_s, e(s))$ and $(A'_s, e'(s))$ for all $s \in S$ then up to unique isomorphism there is a relative correspondence L between (A, e) and (A', e') for which $\psi_L = \psi$ (and so, for example, the s-fiber of L is the unique correspondence between $(A_s, e(s))$ and $(A'_s, e'(s))$ with Riemann form ψ_s).

Example 7.1.6. For the relative complex torus constructed in Example 7.1.4, consider the bilinear self-pairing on the Z-fiber homology given by the matrix (7.1.1) with respect to the fibral homology trivialization in Example 7.1.4. This is a relative homology pairing, and it is a relative Riemann form because on fibers it is the Riemann form of the correspondence $(1 \times \phi_{Z,\mathbf{d}})^*(P_{X_{Z,\mathbf{d}}})$ associated to the Z-fiber polarization $\phi_{Z,\mathbf{d}}$. Thus, this relative homology pairing is the relative Riemann form of a relative self-correspondence of the complex torus $\mathscr{A}_{Z,\mathbf{d},i}$ over $\mathfrak{h}_{g,i}$.

There is a good theory of the relative dual torus, essentially "gluing" the fibral dual tori into a global analytic family in a natural way, but we will not discuss it here; the theory of the relative dual torus is useful in the study of relative correspondences between analytic families of complex tori. We can use the notion of relative correspondence to define a relative notion of polarization for a complex torus over a complex manifold:

Definition 7.1.7. Let (A, e) be a complex torus over a complex manifold S. A (*relative*) *polarization* of A over S is a relative self-correspondence $L \to A \times_S A$ over S such that for each $s \in S$ the induced correspondence $L_s \to A_s \times A_s$ is a polarization; that is, it is symmetric and has pullback $\Delta^*_{A_s}(L_s) \to A_s$ that is an ample line bundle. The collection $\phi = \{\phi_s\}_{s \in S}$ of maps $\phi_s : A_s \to A_s^\vee$ associated to the L_s's uniquely determines the relative correspondence L and so is also referred to as the relative polarization. The locally constant function $s \mapsto \deg(\phi_s)$ on S is the *degree* of the polarization. A *principal relative polarization* is a relative polarization with degree 1.

A relative self-correspondence $L \to A \times_S A$ is a polarization if and only if its fibral Riemann forms ψ_{L_s} are skew-symmetric and satisfy a positivity condition (as in Lefschetz' theorem). For example, Example 7.1.6 is a polarization on $\mathscr{A}_{g,\mathbf{d},i}$ over $\mathfrak{h}_{g,i}$. The notion of polarization can be equivalently formulated in terms of a map of S-tori from A to its relative dual, but we will not address this aspect since we have omitted a discussion of the relative dual torus. In concrete terms, giving a formulation in terms of such a map amounts to making precise the notion of a collection of maps $\phi_s : A_s \to A_s^\vee$ "varying analytically in s".

In general, for an analytic family of polarized tori $\{(A_s, \phi_s)\}_{s \in S}$ with $\dim A_s = g > 0$ for all $s \in S$ the Riemann form ψ_s associated to each ϕ_s is a non-degenerate $\mathbf{Z}(1)$-valued skew-symmetric form on the $\mathbf{Z}$-lattice $\mathrm{H}_1(A_s, \mathbf{Z})$ of rank $2g$, so it has a collection of invariant factors $\{d_j(s)\}$ with positive integers $d_1(s)|\dots|d_g(s)$. One can check that $d_j(s)$ is locally constant in s, and so over each connected component of S the $d_j(s)$'s are constant. Hence, for the purpose of studying analytic families of polarized complex tori it is reasonable to *fix* such a g-tuple of d_j's and to consider only analytic families of g-dimensional polarized complex tori for which the Riemann form has these fixed invariant factors (in which case the polarization has constant degree $\prod d_j^2$). The Siegel torus of type $(\mathbf{d}, i)$ over $\mathfrak{h}_{g,i}$ is an example of such an analytic family.

Rather than focus on the matrix language as in (7.1.1), it is conceptually better to work with an abstract rank-$2g$ non-degenerate symplectic space (Λ, Ψ) over $\mathbf{Z}$ (with $\mathbf{Z}(1)$-valued Ψ) and to consider polarized complex tori whose fibral homologies equipped with their Riemann forms are continuously identified with (Λ, Ψ). We will sometimes avoid this abstract language and use the special cases (that are the general cases) given by the matrix language of (7.1.1); however, the abstract viewpoint eliminates the annoying intervention of bases for lattices and $\mathbf{Z}(1)$, and it makes certain properties of the Siegel half-space much easier to understand (as we shall see).

We have now assembled enough terminology so that we can formulate the universal property of the Siegel torus of type $(\mathbf{d}, i)$. Roughly speaking, it is the universal polarized complex torus whose fibral Riemann forms are continuously identified with (7.1.1). More precisely:

Theorem 7.1.8. *Let $2\pi i$ be a basis of $\mathbf{Z}(1)$ and fix positive integers $d_1|\dots|d_g$. Let A be a polarized complex torus with fiber-dimension $g > 0$ over a complex manifold S, and let ι be a fibral homology trivialization that gives each fibral Riemann form the matrix (7.1.1). For each $s \in S$, let $Z(s) \in \mathfrak{h}_{g,i}$ be the unique point for which A_s is isomorphic to $X_{Z(s),\mathbf{d}}$ as polarized tori equipped with homology trivializations. The map $f : S \to \mathfrak{h}_{g,i}$ defined by $s \mapsto Z(s)$ is analytic, and the pullback of $\mathscr{A}_{g,\mathbf{d},i} \to \mathfrak{h}_{g,i}$ along f is uniquely isomorphic to $A \to S$ as polarized complex tori over S equipped with fibral homology trivializations.*

Remark 7.1.9. The notion of *pullback* for complex tori is defined in a manner analogous to pullback for vector bundles and elliptic curves over S. Briefly, if $f : S' \to S$ is a map of complex manifolds and $M \to S$ is a submersion of complex manifolds then the *pullback* of $M \to S$ along f is the submersion of complex manifolds $M' \to S'$ given by the locally closed submanifold

$$M' = M \times_S S' := \{(m, s') \in M \times S' \,|\, m \in M_{f(s')}\}$$

in $M \times S'$. (This is a locally closed submanifold because $M \to S$ is a submersion.) Roughly speaking, if we think of $M \to S$ as a "family" $\{M_s\}_{s\in S}$ with parameter space S then the pullback is the "family" $\{M_{f(s')}\}_{s'\in S'}$ with parameter space S'. (Of course, such a set-theoretic parameterization is much weaker than the structure provided by the complex manifold M' equipped with its analytic projection to S'.)

7.2. **Siegel half-spaces as coset spaces.** In the case $g = 1$ and $d_1 = 1$, Theorem 7.1.8 asserts a universal property for a (principally polarized) elliptic curve over $\mathfrak{h}_{1,i}$ and thereby becomes exactly Theorem 6.1.6 (in view of the uniqueness of the principal polarization of an elliptic curve). We wish to consider quotients of the universal structure in Theorem 7.1.8 modulo sufficiently small subgroups in the symplectic group $\mathrm{Sp}_{2g}(\mathbf{Z})$, so first we should explain how $\mathrm{Sp}_{2g}(\mathbf{Z})$ acts on $\mathfrak{h}_{g,i}$. For topological purposes we shall first give a transitive real-analytic action by $\mathrm{Sp}_{2g}(\mathbf{R})$ on the real-analytic manifold $\mathfrak{h}_{g,i}$. We want to give the action a conceptual meaning independent of matrix manipulations, so we will use the viewpoint of complex structures and symplectic lattices. (We will also illustrate the abstract definitions with concrete matrix formulas.) Our strategy is to discover another interpretation of the Siegel half-space $\mathfrak{h}_{g,i}$ in terms of the language of variation of complex structure. This will provide a natural action by the real Lie group $\mathrm{Sp}_{2g}(\mathbf{R})$.

Let (Λ, Ψ) be a non-degenerate $\mathbf{Z}(1)$-valued symplectic space of rank $2g > 0$ over $\mathbf{Z}$. That is, Λ is a finite free $\mathbf{Z}$-module of rank $2g$ and Ψ is a non-degenerate $\mathbf{Z}(1)$-valued skew-symmetric bilinear form on Λ. Any such pair is isomorphic to a unique (7.1.1) on the lattice $\mathbf{Z}^{2g}$, and we shall refer to such coordinatized examples as the *standard types*. Recall from §1.1 that to give a complex structure on $\mathbf{R}\otimes_{\mathbf{Z}}\Lambda$ is the same as to give a g-dimensional $\mathbf{C}$-subspace $F \subseteq \mathbf{C}\otimes_{\mathbf{Z}}\Lambda$ such that Λ maps isomorphically onto a lattice in $F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)$ (or, equivalently, such that F and $\overline{F}$ are complementary $\mathbf{C}$-subspaces in $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$). The space of such F's is the Grassmann manifold of g-dimensional subspaces of $\mathbf{C}\otimes_{\mathbf{Z}}\Lambda$, and the condition on F that Λ maps isomorphically onto a lattice modulo F is an *open* condition in this Grassmannian. In this way the set of complex structures on $\mathbf{R}\otimes_{\mathbf{Z}}\Lambda$ acquires a natural structure of complex manifold (in fact, open submanifold of a Grassmann manifold); for example, if $g = 1$ this is the open submanifold $\mathbf{C} - \mathbf{R} = \mathbf{CP}^1 - \mathbf{RP}^1$ in $\mathbf{CP}^1$ (that is, it is a union of two classical half-planes).

We are not interested in arbitrary complex structures on $\mathbf{R}\otimes_{\mathbf{Z}}\Lambda$, but rather only those that interact in a certain way with the symplectic form Ψ. The $\mathbf{C}$-scalar extension

$$\Psi_{\mathbf{C}} : (\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)\times(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda) \to \mathbf{C}(1) = \mathbf{C}$$

is a skew-symmetric $\mathbf{C}$-bilinear form, and so for a given complex structure, or equivalently an F as above, we may consider the condition that F is isotropic with respect to $\Psi_{\mathbf{C}}$. (In the case $g = 1$ this condition is automatically satisfied because F is a line and all lines in a symplectic space are isotropic.) This is precisely the condition that the skew-symmetric Ψ is the Riemann form of a symmetric map $X \to X^\vee$ with $X = F\backslash(\mathbf{C}\otimes_{\mathbf{Z}}\Lambda)/\Lambda$, and so to encode the property of being a polarization we wish to demand the further requirement that the Hermitian form on F arising from the restriction of $\Psi_{\mathbf{C}}$ to a bilinear pairing $F\times\overline{F} \to \mathbf{C}$

is positive-definite. (That is, $\Psi_{\mathbf{C}}(\xi, \overline{\xi}) > 0$ for all nonzero $\xi \in F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$.) Such F satisfy a very important extra condition:

Lemma 7.2.1. *Let W be a $2g$-dimensional $\mathbf{R}$-vector space and let $\psi : W \times W \to \mathbf{R}(1)$ be a non-degenerate symplectic pairing. Let $F \subseteq \mathbf{C} \otimes_{\mathbf{R}} W$ be a g-dimensional subspace corresponding to a complex structure on W (that is, W maps isomorphically to the quotient $F\backslash(\mathbf{C} \otimes_{\mathbf{R}} W)$). If F is $\psi_{\mathbf{C}}$-isotropic and the Hermitian pairing $(\xi_1, \xi_2) \mapsto \psi_{\mathbf{C}}(\xi_1, \overline{\xi}_2)$ on F is positive-definite then for any maximal ψ-isotropic subspace $W_1 \subseteq W$ the natural map $\mathbf{C} \otimes_{\mathbf{R}} W_1 \to F\backslash(\mathbf{C} \otimes_{\mathbf{R}} W)$ is an isomorphism; that is, an $\mathbf{R}$-basis of W_1 is a $\mathbf{C}$-basis of $F\backslash(\mathbf{C} \otimes_{\mathbf{R}} W)$.*

The significance of this lemma is that a $\mathbf{Z}$-basis for *any* maximal Ψ-isotropic lattice in Λ is a $\mathbf{C}$-basis for the complex structure on $W = \mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ associated to F. In other words, if we pick a basis $\{\lambda_j\}$ of Λ with respect to which Ψ acquires the matrix form as in (7.1.1) then the $\mathbf{Z}$-basis $\{\lambda_{g+1}, \dots, \lambda_{2g}\}$ for a maximal Ψ-isotropic lattice in Λ (and hence for a maximal $\Psi_{\mathbf{R}}$-isotropic subspace W_1 of W) is *automatically* a $\mathbf{C}$-basis for the complex structure on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ determined by F. That is, by *fixing* such a basis of Λ and allowing F to *vary*, the map
$$\mathbf{C}^{2g} = \mathbf{C} \otimes_{\mathbf{Z}} \Lambda \twoheadrightarrow F\backslash(\mathbf{C} \otimes_{\mathbf{Z}} \Lambda) = \mathbf{C}^g$$
whose kernel is F can *always* be described as in (1.1.2), which is to say as a map $\begin{pmatrix} Z & 1_g \end{pmatrix}$ with $Z \in \mathrm{Mat}_{g \times g}(\mathbf{C})$ having Z_{im} invertible. Moreover, the isotropicity of F is precisely the condition that Z is symmetric and the positive-definiteness condition is that $Z_{\mathrm{im}}/(2\pi i)$ is positive-definite.

Proof. We have to prove that the g-dimensional $\mathbf{C}$-subspace $\mathbf{C} \otimes_{\mathbf{R}} W_1$ has vanishing intersection with the codimension-g subspace $F \subseteq \mathbf{C} \otimes_{\mathbf{R}} W$. Pick $\xi \in F \cap (\mathbf{C} \otimes_{\mathbf{R}} W_1)$, so $\overline{\xi} \in \overline{F} \cap (\mathbf{C} \otimes_{\mathbf{R}} W_1)$. By the ψ-isotropicity of W_1, we must have $\psi_{\mathbf{C}}(\xi, \overline{\xi}) = 0$. Hence, the definiteness property for the Hermitian pairing on F defined by $\psi_{\mathbf{C}}$ implies $\xi = 0$. ■

Applying Lemma 7.2.1 with $W_1 = \mathbf{R}\lambda_1 \oplus \dots \mathbf{R}\lambda_g$ (in the usual notation), we deduce the important fact that each $g \times g$ matrix $Z \in \mathfrak{h}_{g,i}$ is actually *invertible*. A local calculation also shows that within the complex manifold of complex structures on $W = \mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ (which we have identified with an open subset of a Grassmann manifold), the locus of F's that are $\Psi_{\mathbf{C}}$-isotropic is a closed submanifold. Within this closed submanifold, the positive-definiteness condition as in the preceding lemma is an open condition. We define the *abstract Siegel half-space* $\mathfrak{h}_{(\Lambda,\Psi)}$ to be the complex manifold obtained in this manner.

Example 7.2.2. In the case of the standard types, which is to say $\Lambda = \mathbf{Z}^{2g}$ and Ψ as in (7.1.1), $\mathfrak{h}_{(\Lambda,\Psi)}$ is identified (as a complex manifold, not just as a set) with the usual Siegel half-space $\mathfrak{h}_{g,i}$. Explicitly, to each $Z \in \mathfrak{h}_{g,i}$ we assign the subspace $F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda = \mathbf{C}^{2g}$ that is the kernel of the projection $\mathbf{C}^{2g} \twoheadrightarrow \mathbf{C}^g$ carrying λ_{j+g} to e_j and λ_j to $Z\mathbf{d}(e_j)$ for $1 \le j \le g$. In this quotient, the image of $\Lambda = \mathbf{Z}^{2g}$ is the lattice $Z\mathbf{d}(\mathbf{Z}^g) + \mathbf{Z}^g \subseteq \mathbf{C}^g$ on which the matrix for Ψ with respect to the ordered basis $\{Z\mathbf{d}(e_1), \dots, Z\mathbf{d}(e_g), e_1, \dots, e_g\}$ is as in (7.1.1).

Note that $\mathfrak{h}_{(\Lambda,\Psi)}$ only depends on the structure $(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda, \Psi_{\mathbf{R}})$ over $\mathbf{R}$ rather than on the $\mathbf{Z}$-structure (Λ, Ψ). Hence, on the underlying real-analytic manifold of $\mathfrak{h}_{(\Lambda,\Psi)}$ there is a canonical left action by the automorphism group of the symplectic space $(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda, \Psi_{\mathbf{R}})$ under which a symplectic automorphism g acts by carrying F to $g(F)$. However, to make a link with some classical formulas it is convenient to instead consider a natural action by the automorphism group G of the *dual* real symplectic space $(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda^\vee(1), \Psi_{\mathbf{R}}^\vee)$; the $\mathbf{R}(1)$-valued *dual symplectic form* $\Psi_{\mathbf{R}}^\vee$ is a special case of the general formation of a dual non-degenerate bilinear form $B^\vee$ (with $B^\vee(\ell, \ell') = B(w, w')$ if $\ell = B(\cdot, w)$ and $\ell' = B(\cdot, w')$). A symplectic automorphism $\gamma \in G$ of $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda^\vee(1)$ acts on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ via $(\gamma^{-1})^\vee$, with inversion used to make this a left action of G. For the standard type $(\mathbf{d}, i)$ with $d_j = 1$ for all j (the principally polarized case), the Lie group G is identified with the classical group $\mathrm{Sp}_{2g}(\mathbf{R})$ via the dual standard basis on the $\mathbf{Z}(1)$-dual lattice $\Lambda^\vee(1)$, and the automorphism group of $(\Lambda^\vee(1), \Psi^\vee)$ is identified with the discrete subgroup $\mathrm{Sp}_{2g}(\mathbf{Z})$. (Passing to the dual basis gives $\Psi^\vee$ the negative of the standard matrix in (7.1.1) with $\mathbf{d} = 1_g$, but such signs cancel out in conjugation calculations and so present no complications.) In the case of standard types $(\mathbf{d}, i)$ with $\mathbf{d} \neq 1_g$ (the non-principally polarized case) there is an identification of G with $\mathrm{Sp}_{2g}(\mathbf{R})$, but this identification involves an unpleasant intervention of the d_j's in the translation to the language of coordinatized symplectic groups over $\mathbf{R}$ (and especially over $\mathbf{Z}$).

Remark 7.2.3. There is a more natural motivation for introducing the symplectic group of the dual space, without reference to "rigging" the setup to get classical formulas, namely that Grassmann manifolds in the sense of Grothendieck classify isomorphism classes of quotient bundles rather than subbundles. To put the F's into a global family it is therefore natural to instead form the family of their duals as a holomorphic vector bundle of quotients of the complexified dual symplectic space.

Let us work out the action of G in the case of the standard types. This will show that the action is real-analytic, and it will also link up with some classical matrix formulas (now endowed with a conceptual meaning). To avoid annoying consideration with the d_j's in the formulas, we will restrict attention to the principally polarized case. For $Z \in \mathfrak{h}_{g,i}$ the associated $\Psi_{\mathbf{C}}$-isotropic subspace $F \subseteq \mathbf{C} \otimes_{\mathbf{Z}} \Lambda$ is the kernel of the map $\begin{pmatrix} Z & 1 \end{pmatrix} : \mathbf{C}^{2g} \twoheadrightarrow \mathbf{C}^g$ (with a symmetric Z for which $(Z - \overline{Z})/2\pi i$ is positive-definite), and for

$$\gamma = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \in \mathrm{Aut}(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda^\vee(1), \Psi_{\mathbf{R}}^\vee) \simeq \mathrm{Sp}_{2g}(\mathbf{R})$$

we therefore have that the subspace $(\gamma^{-1})^\vee(F)$ is the kernel of

$$\begin{pmatrix} Z & 1 \end{pmatrix} \begin{pmatrix} A & B \\ C & D \end{pmatrix}^{\mathrm{t}} = \begin{pmatrix} Z & 1 \end{pmatrix} \begin{pmatrix} A^{\mathrm{t}} & C^{\mathrm{t}} \\ B^{\mathrm{t}} & D^{\mathrm{t}} \end{pmatrix} = \begin{pmatrix} ZA^{\mathrm{t}} + B^{\mathrm{t}} & ZC^{\mathrm{t}} + D^{\mathrm{t}} \end{pmatrix}.$$

This must be the kernel of the matrix $\begin{pmatrix} Z' & 1 \end{pmatrix}$ for a unique $Z' \in \mathfrak{h}_{g,i}$, so it is necessary that $ZC^{\mathrm{t}} + D^{\mathrm{t}}$ is invertible (or equivalently, its transpose $CZ + D$ is invertible; such invertibility is not at all obvious from the viewpoint of matrix manipulation). Composition with its inverse on the target does not change the kernel and so gives

$$Z' = (ZC^{\mathrm{t}} + D^{\mathrm{t}})^{-1}(ZA^{\mathrm{t}} + B^{\mathrm{t}}).$$

Since $Z' \in \mathfrak{h}_{g,i}$ is symmetric, passing to the transpose yields the formula

$$Z' = (AZ + B)(CZ + D)^{-1}. \tag{7.2.1}$$

In other words, with the parameter Z we have recovered the classical linear fractional action of $\mathrm{Sp}_{2g}(\mathbf{R})$ on $\mathfrak{h}_{g,i}$ and we recognize that this Lie group is *not* the automorphism group of the given symplectic space but rather of the dual symplectic space.

One advantage of the conceptual approach to this classical action is that we can painlessly prove:

Theorem 7.2.4. *The action of* $\mathrm{Sp}_{2g}(\mathbf{R})$ *on* $\mathfrak{h}_{g,i}$ *is transitive with stabilizer subgroups given by the maximal compact subgroups. In particular,* $\mathfrak{h}_{g,i}$ *is connected and each* $\mathrm{Sp}_{2g}(\mathbf{Q})$*-orbit is dense.*

Proof. Any point in an abstract Siegel half-space corresponds to a complex structure on $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ with respect to which $\Psi_{\mathbf{R}}$ is the imaginary component of a positive-definite Hermitian form. But any two positive-definite Hermitian spaces with the same finite dimension are isomorphic, so for any two points in such a Siegel half-space we get an $\mathbf{R}$-linear automorphism of $\mathbf{R} \otimes_{\mathbf{Z}} \Lambda$ that carries one complex structure to the other and one Hermitian form to the other. In particular, this automorphism respects the formation of the common imaginary component $\Psi_{\mathbf{R}}$ of these Hermitian forms, which is to say that the automorphism we have built is an automorphism of the symplectic space $(\mathbf{R} \otimes_{\mathbf{Z}} \Lambda, \Psi_{\mathbf{R}})$ over $\mathbf{R}$. By working with the dual of the inverse automorphism in G, we obtain precisely the desired transitivity result. The stabilizer of a point is an automorphism of a positive-definite Hermitian space, which is to say a unitary group. The maximality of these compact stabilizers can be proved by brute force, but it is more elegant to obtain it as an obvious consequence of the general fact that any compact group acting continuously on a cone in a real Euclidean space has points with full isotropy group (proof: average via integration over the compact group).

The connectivity of $\mathfrak{h}_{g,i}$ now follows from transitivity of the real-analytic (hence continuous) action by $\mathrm{Sp}_{2g}(\mathbf{R})$ on $\mathfrak{h}_{g,i}$, together with the connectivity of this symplectic group. (To prove connectivity, either use that the symplectic group is generated by transvections, or that Sp_{2g} is a semisimple linear algebraic group).

The density of $\mathrm{Sp}_{2g}(\mathbf{Q})$-orbits is a consequence of the density of $H(\mathbf{Q})$ in $H(\mathbf{R})$ for any connected reductive linear algebraic group H over $\mathbf{Q}$ (use the Bruhat decomposition). ■

The following corollary is rather fundamental:

Corollary 7.2.5. *Fix a CM field K with degree $2g$ over $\mathbf{Q}$, and any CM type Φ for K. The locus of points $Z \in \mathfrak{h}_{g,i}$ such that the polarized torus fiber $(X_{Z,\mathbf{d}}, \phi_{Z,\mathbf{d}})$ admits CM by K and CM type Φ with respect to which the Rosati involution restricts to complex conjugation on K is a dense subset of $\mathfrak{h}_{g,i}$.*

In contrast with the case $g = 1$, for $g > 1$ it seems to be hopeless to easily characterize those Z for which the Z-fiber $X_{Z,\mathbf{d}}$ of the Siegel family of type $(\mathbf{d}, i)$ over $\mathfrak{h}_{g,i}$ (using any $\mathbf{d}$) admits a structure of CM torus (let alone with a specific CM field and CM type).

Proof. If Z and Z' in $\mathfrak{h}_{g,i}$ lie in the same $\mathrm{Sp}_{2g}(\mathbf{Q})$-orbit then the associated polarized torus fibers in the Siegel family are isogenous. (Think about the case $g = 1$ first.) Moreover, this isogeny respects the polarizations up to rational multiple. Hence, if $(X_{Z,\mathbf{d}}, \phi_{Z,\mathbf{d}})$ admits CM by K with CM type Φ such that the Rosati involution restricts to complex conjugation on K then all fibers of the Siegel family over the $\mathrm{Sp}_{2g}(\mathbf{Q})$-orbit of Z have this property. By Theorem 7.2.4, it therefore suffices to exhibit a single such point Z. That is, we seek a principally polarized CM torus with CM field K and CM type Φ such that the Rosati involution restricts to complex conjugation on K. The proof that any polarized complex torus is isogenous to a principally polarized torus shows likewise that if the given polarization is linear for a CM-structure then a principal polarization can be found with the same property. Hence, we can relax the polarization requirement to merely a polarization requirement. We have constructed the required polarization on $(\mathbf{R} \otimes_{\mathbf{Q}} K)_\Phi / \mathscr{O}_K$, so we are done. ■

7.3. **Algebraic Siegel modular varieties.** In the previous section, we undertook a detailed study of the Siegel families over Siegel half-spaces. These families are analytic in nature, and are characterized by analytic properties (especially homology trivializations) with no algebraic analogue. We now shall explain how passage to the quotient by the action of suitable subgroups of $\mathrm{Sp}_{2g}(\mathbf{Z})$ gives analytic families that are characterized by universal properties expressed entirely in the geometric language of polarizations and torsion-points. Such universal properties will admit a natural translation in terms of algebraic geometry, and this leads to algebraic models (so-called *Siegel modular varieties*) that generalize what we have seen for $g = 1$ via the theory of algebraic models for modular curves. To focus on the main ideas without bothersome technical distractions, we shall restrict our attention to fine moduli problems and the principally polarized case. This is already sufficient to illustrate the role of the Main Theorem of Complex Multiplication in the proof that such modular varieties are canonical models in the sense of Shimura. More specifically, in §7.4 it will be proved that for certain models of Siegel modular varieties over number fields the Galois-action at CM-points is described in a manner that generalizes what we have seen in our study of Heegner points on modular curves. Corollary 7.2.5 will ensure that the locus of CM-points is dense for the Zariski-topology on such varieties over $\mathbf{C}$ (as density will even hold for the analytic topology), so the framework of Shimura's descent theory will be applicable to characterize these algebraic models via the explicit Galois action at CM points.

Let $\mathscr{A}_{g,i}$ be the universal principally polarized complex torus over $\mathfrak{h}_{g,i}$; this is the case $d_j = 1$ for all j, and the universal property involves a fibral homology trivialization that identifies the Riemann form of the fibral polarizations with (7.1.1) for $\mathbf{d} = 1_g$. The action we have defined for $\mathrm{Sp}_{2g}(\mathbf{R})$ on $\mathfrak{h}_{g,i}$ does not lift to an action on $\mathscr{A}_{g,i}$ since it does not preserve the lattice $\Lambda = \mathbf{Z}^{2g}$ inside of $\mathbf{R} \otimes_{\mathbf{R}} \Lambda = \mathbf{R}^{2g}$. However, the subgroup $\mathrm{Sp}_{2g}(\mathbf{Z})$ does preserve this lattice and the action by this discrete subgroup does lift to an action on $\mathscr{A}_{g,i}$ covering the action on $\mathfrak{h}_{g,i}$. More conceptually, let (Λ, Ψ) be a perfect symplectic space of rank $2g$ over $\mathbf{Z}$ with a $\mathbf{Z}(1)$-valued Ψ (such as $\Lambda = \mathbf{Z}^{2g}$ and Ψ as in (7.1.1) with $\mathbf{d} = 1_g$), and let γ be an element in the automorphism group Γ of the dual symplectic lattice $(\Lambda^\vee(1), \Psi^\vee)$. We let γ act on the set of triples (A, ϕ, ι) consisting of a principally polarized complex torus (A, ϕ) and an isomorphism $\iota : \Lambda \simeq \mathrm{H}_1(A, \mathbf{Z})$ carrying Ψ to the Riemann form (or Weil self-pairing) of ϕ via the procedure

$$[\gamma](A, \phi, \iota) = (A, \phi, \iota \circ \gamma^\vee).$$

Note that this *is* a left action of Γ on the family $\mathscr{A}_{(\Lambda,\Psi)} \to \mathfrak{h}_{(\Lambda,\Psi)}$ over the abstract Siegel half-space that is universal among analytic families of principally polarized complex tori whose fibral homologies (as symplectic spaces over $\mathbf{Z}$) are continuously identified with (Λ, Ψ).

Let us now check that this action agrees with the action induced by that of the real Lie group G on $\mathfrak{h}_{(\Lambda,\Psi)}$ as built earlier via the identification of $\mathfrak{h}_{(\Lambda,\Psi)}$ as a parameter space for certain variations of complex structure. By a direct calculation, in the case of the standard $(1_g, i)$-type this second construction is an action of $\mathrm{Sp}_{2g}(\mathbf{Z})$ on $\mathfrak{h}_{g,i}$ and it is given by the explicit formula that carries Z to Z' where $\begin{pmatrix} Z' & 1_g \end{pmatrix}$ is a $\mathrm{GL}_g(\mathbf{C})$-translate of

$$\begin{pmatrix} Z & 1_g \end{pmatrix} \begin{pmatrix} A^{\mathrm{t}} & C^{\mathrm{t}} \\ B^{\mathrm{t}} & D^{\mathrm{t}} \end{pmatrix} = \begin{pmatrix} ZA^{\mathrm{t}} + B^{\mathrm{t}} & ZC^{\mathrm{t}} + D^{\mathrm{t}} \end{pmatrix},$$

which is to say $Z' = (ZC^{\mathrm{t}} + D^{\mathrm{t}})^{-1}(ZA^{\mathrm{t}} + B^{\mathrm{t}})$. As before, since Z' is symmetric we may pass to the transpose formula to recover the formula (7.2.1) (with A, B, C, and D now restricted to have entries in $\mathbf{Z}$). Hence, we have indeed recovered the action by the integral symplectic group, and we get more: a (unique) lifting of this to an action on the universal complex torus over the Siegel half-space, and the fact that this action is properly discontinuous (since the Siegel space is an orbit space for a real Lie group modulo a compact subgroup, and the overlap of a compact subgroup and a discrete subgroup is finite).

To formulate a structure on complex tori that is amenable to translation into the algebraic theory, we must work with trivializations of symplectic pairings at some torsion-level rather than at the level of homology lattices. This rests on the following definition:

Definition 7.3.1. Let $\{(A_s, \phi_s)\}_{s\in S}$ be an analytic family of principally polarized complex tori of dimension $g \geq 1$ parameterized by complex manifold S, and let n be a positive integer. Let L be a free $\mathbf{Z}/n\mathbf{Z}$-module with rank $2g$ and $\psi : L \times L \to R$ an alternating bilinear form with values in a free $\mathbf{Z}/n\mathbf{Z}$-module R with rank 1 such that ψ is perfect (*i.e.*, either of the two associated linear maps $L \to \mathrm{Hom}_{\mathbf{Z}/n\mathbf{Z}}(L, R)$ is an isomorphism).

An (L, ψ)*-structure* on $A \to S$ is a continuously varying family $\iota = \{\iota_s\}_{s\in S}$ of isomorphisms of symplectic spaces $\iota_s : (L, \psi) \simeq (A_s[n], e_{\phi_s,n})$, where $e_{\phi_s,n}$ is the $\mu_n(\mathbf{C})$-valued Weil n-torsion self-pairing on $A_s[n]$ induced by the symmetric isomorphism $\phi_s : A_s \simeq A_s^\vee$. That is, each ι_s is a pair of isomorphisms $L \simeq A_s[n]$ and $R \simeq \mu_n(\mathbf{C})$ of $\mathbf{Z}/n\mathbf{Z}$-modules carrying ψ to $e_{\phi_s,n}$.

Example 7.3.2. Let $L = (\mathbf{Z}/n\mathbf{Z})^{2g}$ and $R = \mathbf{Z}/n\mathbf{Z}$, and let ψ given by the standard matrix

$$\psi = \begin{pmatrix} 0 & -1_g \\ 1_g & 0 \end{pmatrix}.$$

In this case, an (L, ψ)-structure is called a *full level-n structure of type ζ* when the isomorphism $R \simeq \mu_n(\mathbf{C})$ carries $1 \bmod n\mathbf{Z} \in R$ to the primitive nth root of unity ζ in $\mathbf{C}$. On a principally polarized complex torus (A, ϕ) such a structure is precisely the data consisting of an ordered basis $\{\lambda_1, \ldots, \lambda_{2g}\}$ of $A[n]$ and a primitive nth root of unity $\zeta \in \mu_n(\mathbf{C})$ such that $e_{\phi,n}(\lambda_j, \lambda_{j'}) \in \mu_n(\mathbf{C})$ is trivial for $j' \neq j \pm g$ and it is equal to ζ (resp. ζ^{-1}) for $j' = j - g$ (resp. for $j' = j + g$).

Example 7.3.3. Let $L = (\mathbf{Z}/n\mathbf{Z})^g \times \mu_n(\mathbf{C})^g$ and $R = \mu_n(\mathbf{C})$, and let $\psi : L \times L \to R$ be the standard symplectic form that makes $(\mathbf{Z}/n\mathbf{Z})^g$ and $\mu_n(\mathbf{C})^g$ isotropic and restricts to the canonical bilinear pairing $\mu_n(\mathbf{C})^g \times (\mathbf{Z}/n\mathbf{Z})^g \to \mu_n(\mathbf{C})$ given by

$$((\zeta_1, \ldots, \zeta_g), (r_1, \ldots, r_g)) \mapsto \prod \zeta_j^{r_j}.$$

In this case, an (L, ψ)-structure is called a *μ-type level-n structure*. The interest in such structures is that they give rise to modular varieties defined over $\mathbf{Q}$ rather than over cyclotomic extensions of $\mathbf{Q}$. (See Remark 7.3.12.)

Let (Λ, Ψ) be a perfect symplectic space of rank $2g$ over $\mathbf{Z}$ with Ψ taking values in $\mathbf{Z}(1)$. The reduction $(\Lambda_n, \Psi_n) = (\Lambda/n\Lambda, \Psi \bmod n)$ is a perfect symplectic space over $\mathbf{Z}/n\mathbf{Z}$ with the symplectic form taking values in $\mathbf{Z}(1)/n\mathbf{Z}(1) = \mu_n(\mathbf{C})$. The universal principally polarized complex torus $\mathscr{A}_{(\Lambda,\Psi)} \to \mathfrak{h}_{(\Lambda,\Psi)}$ has a canonical

(Λ_n, Ψ_n)-structure that is invariant under the natural action on the family by the principal level-n congruence subgroup

$$\Gamma_{(\Lambda,\Psi)}(n) = \ker(\mathrm{Aut}(\Lambda^\vee(1), \Psi^\vee) \to \mathrm{Aut}(\Lambda_n^\vee(1), \Psi_n^\vee).$$

In the case of the standard $(1_g, i)$-type, this says that the classical principally polarized Siegel family $\mathscr{A}_{g,i} \to \mathfrak{h}_{g,i}$ has a full level-n structure of type $e^{2\pi i/n}$ that is invariant under the natural action on the family by the principal level-n congruence subgroup

$$\Gamma(n; g) = \ker(\mathrm{Sp}_{2g}(\mathbf{Z}) \to \mathrm{Sp}_{2g}(\mathbf{Z}/n\mathbf{Z})).$$

For $n \geq 3$ the properly discontinuous actions by $\Gamma_{(\Lambda,\Psi)}(n)$ on both $\mathscr{A}_{(\Lambda,\Psi)}$ and $\mathfrak{h}_{(\Lambda,\Psi)}$ are free, so we may pass to the quotient to get a new analytic family of complex tori equipped with a (Λ_n, Ψ_n)-structure. By smoothness of the symplectic group, it follows from the universal property of the analytic family over $\mathfrak{h}_{(\Lambda,\Psi)}$ that this structure modulo the principal level-n congruence subgroup is *universal* among principally polarized complex tori equipped with a (Λ_n, Ψ_n)-structure. We record this result in the classical setting where it has a concrete form:

Theorem 7.3.4. *Fix $n \geq 3$ and $g \geq 1$. The analytic family of principally polarized complex tori*

$$\Gamma(n; g)\backslash\mathscr{A}_{g,i} \to \Gamma(n; g)\backslash\mathfrak{h}_{g,i}$$

equipped with its canonical full level-n structure of type $e^{2\pi i/n}$ is universal. That is, any principally polarized complex torus equipped with a full level-n structure of type $e^{2\pi i/n}$ over a complex manifold S is uniquely isomorphic to a pullback of the family over $\Gamma(n; g)\backslash\mathfrak{h}_{g,i}$ via a unique analytic map $S \to \Gamma(n; g)\backslash\mathfrak{h}_{g,i}$.

We say that $\Gamma(n; g)\backslash\mathfrak{h}_{g,i}$ (equipped with the quotient principally polarized complex torus and full level-n structure over it) is a *fine moduli space* for the moduli problem of classifying analytic families of g-dimensional principally polarized complex tori equipped with a full level-n structure of type $e^{2\pi i/n}$. This is called the *analytic Siegel moduli space* of principally polarized complex tori of dimension g equipped with full level-n structure of type $e^{2\pi i/n}$. Note that this fine moduli space is *connected*, due to the connectivity of $\mathfrak{h}_{g,i}$.

Example 7.3.5. Consider the case when S is a point, so let (A, ϕ) be a principally polarized complex torus of dimension g and pick a full level-n structure of type $e^{2\pi i/n}$ on the n-torsion. The theorem implies that there exists $Z \in \mathfrak{h}_{g,i}$ and a unique isomorphism $(A, \phi) \simeq (X_{Z,1_g}, \phi_{Z,1_g})$ carrying the chosen basis of $A[n]$ to the basis of $X_{Z,1_g}[n] \simeq \mathrm{H}_1(X_{Z,1_g}, \mathbf{Z}/n\mathbf{Z})$ obtained by reduction of the standard basis of $\mathrm{H}_1(X_{Z,1_g}, \mathbf{Z})$, and that Z is unique up to the left action of $\Gamma(n; g)$ on $\mathfrak{h}_{g,i}$. In the case $g = 1$, this recovers the classical modular curve $Y(n) = \Gamma(n)\backslash\mathfrak{h}_i$ *equipped with* its universal elliptic curve (whose fiber over the orbit $\Gamma(n)z$ is $\mathbf{C}/(\mathbf{Z}z \oplus \mathbf{Z})$ equipped with the n-torsion basis $\{z/n, 1/n\}$, or equivalently is $\mathbf{C}^\times/q_z^{\mathbf{Z}}$ equipped with the n-torsion basis $\{q_{z/n}, e^{2\pi i/n}\}$ where $q_\tau = e^{2\pi i\tau}$ for $\tau \in \mathfrak{h}_i$).

Example 7.3.6. Let ζ be an arbitrary primitive nth root of unity in $\mathbf{C}$ with $n \geq 3$. Does there exist a universal principally polarized complex torus of relative dimension g equipped with a full level-n structure of type ζ? Indeed there does, and the underlying analytic family of polarized tori is again given by $\mathfrak{h}_{g,i}$; all we have to do is modify the level structure on the polarized torus family over this base space. More specifically, if ζ and ζ' are two primitive nth roots of unity in $\mathbf{C}$ and $j \in (\mathbf{Z}/n\mathbf{Z})^\times$ is the unique unit such that $\zeta' = \zeta^j$ then for any principally polarized complex torus $A \to S$ we may naturally pass between full level-n structures of types ζ and ζ' on $A \to S$ by multiplying the first g of the $2g$ basis vectors by j (or $1/j$). Thus, if we are given a universal family of one type then by applying this procedure to the universal n-torsion basis we obtain a universal family of the other type. As an illustration, in the case $g = 1$ we may view $Y(n)$ as a fine moduli space for elliptic curves equipped with an ordered n-torsion basis $\{P, Q\}$ having Weil n-torsion pairing $e_n(P, Q) = e^{-2\pi ij/n}$ by using the standard elliptic curve $\Gamma(n)\backslash\mathscr{E} \to Y(n)$ and the n-torsion basis $\{jz/n, 1/n\}$ in the z-fiber $\mathbf{C}/(\mathbf{Z}z \oplus \mathbf{Z})$. The case $g > 1$ over $\mathfrak{h}_{g,i}$ proceeds similarly.

For applications to certain Shimura curves, it is convenient to also consider the following example (that arises after "forgetting" an action on an abelian surface by an order in a quaternion division algebra):

Example 7.3.7. Let B be a quaternion division algebra over $\mathbf{Q}$ that is split at the infinite place, and let D be the product of its ramified primes. Pick a maximal order O in B, and choose $\rho \in O$ such that $\rho^2 = -D$. Let $x' = \mathrm{Trd}(x) - x$ denote the main involution of B. The pairing $B \times B \to \mathbf{Q}$ defined by $(x, y)_\rho = (1/D)\mathrm{Trd}(\rho x y')$ is a perfect $\mathbf{Z}$-valued skew-symmetric form on O.

Fix a primitive nth root of unity ζ in $\mathbf{C}$ and let $[x] \in O/nO$ denote the residue class of $x \in O$. The pairing $\psi_{\rho,n} : O/nO \times O/nO \to \mu_n(\mathbf{C})$ defined by $\psi_{\rho,n}([x],[y]) = \zeta^{(x,y)_\rho}$ is a perfect skew-symmetric pairing. It is natural to ask if there is a universal analytic family of principally polarized complex tori of dimension g equipped with an $(O/nO, \psi_{\rho,n})$-structure for $n \geq 3$. Since any two perfect symplectic spaces of the same rank over $\mathbf{Z}/n\mathbf{Z}$ are isomorphic, such a universal structure does exist and its underlying polarized complex torus is the universal principally polarized complex torus over $\Gamma(n; g)\backslash\mathfrak{h}_{2,i}$. (This problem becomes much more interesting when we impose extra structure on the abelian surface, namely an O-action and principal polarization whose Riemann form is $\pm 2\pi i(\cdot,\cdot)_\rho$.)

The notion of (L, ψ)-structure makes sense in the algebraic setting, and so we may formulate moduli problems and seek universal families in an algebraic context such that over $\mathbf{C}$ we may hope to recover the universal analytic structures (via analytification). Let us begin our discussion of these algebraic topics by giving some definitions; the definitions we provide are adequate for expository purposes, but not for proving the theorems that we shall state.

Definition 7.3.8. Let S be a smooth algebraic variety over a field k. An *abelian S-variety* (or *relative abelian variety over S*) is a smooth variety A equipped with a smooth proper surjection $\pi : A \to S$ and a section $e : S \to A$ such that for every field k'/k and $s \in S(k')$ the fiber $(A_s, e(s))$ has a (necessarily unique) structure of abelian variety over k'. We shall also call such a structure an *algebraic family of abelian varieties* $\{A_s\}$ parameterized by S.

Remark 7.3.9. The notions of *smoothness* and *properness* for maps of varieties as used in this definition are algebraic analogues of the notions of *submersion* and *properness* in the setting of complex manifolds. These are conditions that ensure each fiber has a natural structure of smooth complete variety, but it is too much of a digression to explain the definitions of these algebraic notions in detail.

There is a good notion of pullback for relative abelian varieties, and many properties of analytic families of complex tori have analogues for relative abelian varieties. For example, the function $s \mapsto \dim A_s$ is locally constant on S (and so we usually restrict attention to the case when it is a constant $g \geq 1$, also called a relative abelian variety with *relative dimension* g), and if n is an integer not divisible by the characteristic of the base field then the fibral n-torsion subgroups $A_s[n]$ naturally "glue" together to give a finite étale S-group $A[n] \to S$. This latter object is a structure that generalizes the Galois-module structure on n-torsion in the classical case of abelian varieties over a field that is not assumed to be separably closed.

The notions of relative correspondence and relative polarization in the analytic theory admit reasonable algebraic analogues, and we consider a relative polarization of a relative abelian variety $A \to S$ to be a collection of polarizations $\phi_s : A_s \to A_s^\vee$ of the fibers such that ϕ_s "depends algebraically on s". Making this notion precise requires more technique than we wish to discuss here, but it roughly says that the description of the ϕ_s's in suitable affine charts involves polynomials whose coefficients are algebraic functions on the base space S.

The preceding notion of (L, ψ)-structure has an algebraic analogue, but we restrict our attention to a special case for which the definition is easier to formulate.

Definition 7.3.10. Let $A \to S$ be a relative abelian variety with relative dimension $g \geq 1$ over a smooth variety S over a field k such that $\mathrm{char}(k) \nmid n$. A *full level-n structure* on $A \to S$ is an ordered $2g$-tuple of algebraic sections $P_j : S \to A$ for $1 \leq j \leq 2g$ such that $\{P_j(s)\}$ is a $\mathbf{Z}/n\mathbf{Z}$-basis of $A_s[n]$ for all $s \in S$.

If $A \to S$ is equipped with a principal polarization $\phi = \{\phi_s\}$ and if $\zeta \in \mu_n(k)$ is a primitive nth root of unity such that the matrix $(e_{\phi_s,n}(P_i, P_j))$ is equal to

$$\begin{pmatrix} 0 & \zeta^{-1} \cdot 1_g \\ \zeta \cdot 1_g & 0 \end{pmatrix}$$

for all $s \in S$ then this full level-n structure on A is *of type* ζ with respect to ϕ.

A fundamental result of Mumford [GIT, Ch. 6] is the existence of universal algebraic families of such structures over arbitrary $\mathbf{Z}[1/n]$-schemes for $n \geq 3$. We record his result over fields:

Theorem 7.3.11 (Mumford)**.** *Let k be a field with $\mathrm{char}(k) \nmid n$ and let $\zeta \in k$ be a primitive nth root of unity. Over k there exists a universal principally polarized relative abelian variety of relative dimension g equipped with a full level-n structure of type ζ over a smooth k-variety $M_{g,n,\zeta/k}$. This k-variety is irreducible, and for any extension field k'/k the family obtained by extension of the base field to k' is the universal family among k'-varieties. Moreover, for $k = \mathbf{C}$ the analytification of this universal family recovers the universal analytic family of type ζ.*

The variety $M_{g,n,\zeta/k}$ is called the *Siegel modular variety of level* n (and type ζ) over k. Its connectivity in characteristic 0 is proved by comparison with the analytic theory over $\mathbf{C}$, and its connectivity in positive characteristic requires deep work of Faltings and Chai on the compactification of Siegel moduli schemes over $\mathrm{Spec}\,\mathbf{Z}[1/n]$ (from which the connectivity may be inferred via the known connectivity in characteristic 0).

Remark 7.3.12*.* It is sometimes desired to have a model for the Siegel modular varieties over $\mathbf{Q}$. That is, for each $n \geq 3$ we seek an absolutely irreducible variety over $\mathbf{Q}$ that descends the modular variety classifying principally polarized g-dimensional abelian varieties with full level-n structure (of type ζ_n) over $\mathbf{Q}(\zeta_n)$. Briefly, one has to pose a twisted moduli problem that avoids the specification of a primitive nth root of unity. The idea is to work with the viewpoint of μ-type moduli problems, as in Example 7.3.3, and once the proper definitions are given the construction of the moduli space over $\mathbf{Q}$ is a consequence of a Galois-twisting argument that we omit.

Remark 7.3.13*.* For (L, ψ) as in Example 7.3.7, the notion of an (L, ψ)-structure has a reasonable definition in the relative algebraic setting when the base field k has characteristic not dividing n and a primitive nth root of unity ζ is chosen in k. Once such a definition is given, a straightforward argument provides a universal such algebraic structure with underlying principally polarized relative abelian variety given by Mumford's algebraic family over k. In particular, the moduli space in this case is a connected smooth k-variety.

Example 7.3.14*.* In the case $k \subseteq \mathbf{C}$ is the nth cyclotomic field and $\zeta \in k$ is a primitive nth root of unity, the Siegel modular variety of level n and type ζ is to be considered as an algebraic structure over k for the quotient $\Gamma(n; g)\backslash\mathfrak{h}_{g,i}$ (viewed as the analytic fine moduli space for classifying g-dimensional principally polarized complex tori with a full level-n structure of type ζ). The algebraic structure over $\mathbf{C}$ on this quotient turns out to be unique, even though it is non-compact, due to work of Baily and Borel. Let us define a *CM point* on this modular variety (over $\mathbf{C}$) to be a point classifying a principally polarized abelian variety that admits a structure of CM abelian variety of type (K, Φ) such that the Rosati involution restricts to complex conjugation on K. For a fixed (K, Φ), the locus of such points in the modular variety is Zariski-dense, due to two facts: (i) the identification of the analytification of the universal algebraic object over $\mathbf{C}$ with the universal analytic one over $\Gamma(n; g)\backslash\mathfrak{h}_{g,i}$, and (ii) the analytic denseness of such points in $\mathfrak{h}_{g,i}$ (Theorem 7.2.5).

The CM points on the modular variety are *algebraic*: they are rational over the subfield $\overline{\mathbf{Q}} \subseteq \mathbf{C}$ with respect to the $\overline{\mathbf{Q}}$-structure provided by the Mumford model. Indeed, by Example 4.1.7, Theorem 4.2.10, and Example 4.2.13 it follows that any CM torus equipped with a (principal) polarization and level structure over $\mathbf{C}$ may be uniquely descended (together with its polarization and level structure) to $\overline{\mathbf{Q}}$. Hence, the universal property of the modular variety over $\overline{\mathbf{Q}}$ ensures that such $\mathbf{C}$-points are algebraic.

As we let (K, Φ) vary with K containing the cyclotomic field k, the associated reflex fields K^* satisfy $\cap_{(K,\Phi)} kK^* = k$. Hence, by Corollary 6.2.13 we can uniquely determine the algebraic structure on the level-n Siegel modular variety over the cyclotomic field k if we can give a direct analytic description of the action by $\mathrm{Aut}(\mathbf{C}/kK^*)$ on CM points of type (K, Φ) for all (or even a sufficiently large finite set of) such pairs (K, Φ). This problem is taken up in §7.4 using the Main Theorem of Complex Multiplication for abelian varieties, and the method we use can be applied to many other modular varieties for polarized abelian varieties.

7.4. **Galois action at CM points.** In this final section, we recall the statement of the Main Theorem of Complex Multiplication for abelian varieties (in an intrinsic form that avoids unnecessary choices) and we use it to give an algebraic characterization of Siegel modular varieties in the manner suggested in Example 7.3.14.

Let $\overline{\mathbf{Q}}$ be an algebraic closure of $\mathbf{Q}$ and let (K,Φ) be a CM field of degree $2g$ over $\mathbf{Q}$ equipped with a CM type $\Phi \subseteq \mathrm{Hom}(K,\overline{\mathbf{Q}})$. Let (A,i,ϕ) be a CM abelian variety over $\overline{\mathbf{Q}}$ equipped with a K-linear polarization $\phi : A \to A^\vee$, where $A^\vee$ is made into a CM abelian variety of type (K,Φ) via the K-action $i^\vee : K \to \mathrm{End}^0(A^\vee)$ defined by $i^\vee(c) = i(\overline{c})^\vee$ (with $c \mapsto \overline{c}$ denoting complex conjugation on K). An equivalent formulation of the requirement on ϕ is that the Rosati involution on $\mathrm{End}^0(A)$ induced by ϕ restricts to complex conjugation on the $\mathbf{Q}$-subalgebra $i(K)$. For any automorphism σ of $\overline{\mathbf{Q}}$, we get another triple $(A^\sigma, i^\sigma, \phi^\sigma)$ through algebraic extension of scalars by $\sigma : \overline{\mathbf{Q}} \simeq \overline{\mathbf{Q}}$ (with $i^\sigma(c) = i(c)^\sigma$ for $c \in K$), and this has type $(K, \sigma \circ \Phi)$ (where σ acts on $\mathrm{Hom}(K,\overline{\mathbf{Q}})$ through composition on the target). Recall that the *reflex field* $K^* \subseteq \overline{\mathbf{Q}}$ is the fixed field of the open stabilizer of Φ under the left action of $\mathrm{Gal}(\overline{\mathbf{Q}}/\mathbf{Q})$ on $\mathrm{Hom}(K,\overline{\mathbf{Q}})$ via composition on the target field $\overline{\mathbf{Q}}$; this subfield of $\overline{\mathbf{Q}}$ depends on Φ in general. If σ fixes the reflex field $K^* \subseteq \overline{\mathbf{Q}}$ then (A^σ, i^σ) is again of type (K,Φ). (Note that whereas K is an *abstract* field not endowed with an embedding into $\overline{\mathbf{Q}}$, by definition the reflex field K^* does have such an embedding.)

Fix an embedding of $\overline{\mathbf{Q}}$ into $\mathbf{C}$, so it makes sense to speak of $\mathbf{C}$-points of an algebraic variety over $\overline{\mathbf{Q}}$, and let $\mathrm{N}_{\mathrm{ref}} : \mathbf{A}_{K^*}^\times \to \mathbf{A}_K^\times$ be the *idelic reflex norm* associated to the CM type (K,Φ). The problem that is solved by the Main Theorem of Complex Multiplication is to use class field theory and the reflex norm to give an analytic description of $(A^\sigma, i^\sigma, \phi^\sigma)$ in terms of an analytic description of (A,i,ϕ). The traditional formulation requires the non-canonical choices of an embedding of K into $\mathbf{C}$ and a basis for the free module $\mathrm{T}_e(A(\mathbf{C}))$ of rank 1 over the $\mathbf{C} \otimes_{\mathbf{Q}} K$-algebra $(\mathbf{R} \otimes_\Phi K)_\Phi$, as well as far-out stuff like the automorphism group of $\mathbf{C}$. We prefer to avoid such things in our formulation of the Main Theorem:

Theorem 7.4.1 (Shimura, Taniyama)**.** *Let V/Λ and V_σ/Λ_σ be the canonical analytic uniformizations of the complex tori $A(\mathbf{C})$ and $A^\sigma(\mathbf{C})$ respectively, and pick an idele $s \in \mathbf{A}_{K^*}^\times$ whose image $(s|K^*) \in$* $\mathrm{Gal}((K^*)^{\mathrm{ab}}/K^*)$ *under the Artin map is $\sigma|_{(K^*)^{\mathrm{ab}}}$. There is a unique $\mathbf{C} \otimes_{\mathbf{Q}} K$-linear isomorphism $V_\sigma \simeq V$ under which Λ_σ is carried to the lattice* $\mathrm{N}_{\mathrm{ref}}(1/s)\Lambda$ *in the* 1*-dimensional K-vector space $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda$, and under the resulting identification of $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda_\sigma$ with $\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda$ the diagram*

$$\begin{array}{ccccc} (\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda)/\Lambda & \xrightarrow{\simeq} & A(\mathbf{C})_{\mathrm{torsion}} & \xleftarrow{\simeq} & A(\overline{\mathbf{Q}})_{\mathrm{torsion}} \\ {\scriptstyle \mathrm{N}_{\mathrm{ref}}(1/s)}\downarrow{\scriptstyle\simeq} & & & & \downarrow{\scriptstyle\sigma} \\ (\mathbf{Q} \otimes_{\mathbf{Z}} \Lambda)/\mathrm{N}_\Phi(1/s)\Lambda & \xrightarrow[\simeq]{} & A^\sigma(\mathbf{C})_{\mathrm{torsion}} & \xleftarrow[\simeq]{} & A^\sigma(\overline{\mathbf{Q}})_{\mathrm{torsion}} \end{array} \tag{7.4.1}$$

commutes, where the right side is the Galois action on torsion points and the left side is idelic multiplication.

Moreover, the induced identification of rational homology lattices

$$\mathrm{H}_1(A^\sigma(\mathbf{C}),\mathbf{Q}) = \mathbf{Q} \otimes_{\mathbf{Z}} \Lambda_\sigma \simeq \mathbf{Q} \otimes_{\mathbf{Z}} \Lambda = \mathrm{H}_1(A(\mathbf{C}),\mathbf{Q})$$

carries the Riemann form Ψ_{ϕ^σ} to $q_{\sigma,s}\Psi_\phi$ with $q_{\sigma,s} \in \mathbf{Q}_{>0}^\times$ the unique positive rational generator of the fractional ideal in $\mathbf{Q}$ associated to the rational idele $\mathrm{N}_{K^*/\mathbf{Q}}(s)$.

Let $k \subseteq \mathbf{C}$ be the nth cyclotomic field and let $\zeta \in k$ be a primitive nth root of unity. We saw in Example 7.3.14 that to give a Galois-theoretic characterization of the level-n Siegel modular variety of type ζ over k it suffices to give an analytic description of the action of $\mathrm{Gal}(\overline{\mathbf{Q}}/kK^*)$ on CM points of every possible type (K,Φ) with $[K:\mathbf{Q}] = 2g$. Such a description is easily obtained from the theorem of Shimura and Taniyama, as follows. We pick a point (A,ϕ,ι) on the Siegel variety, where ι is the specification of a full level-n structure of type ζ on (A,ϕ), and we assume it admits a CM structure of type (K,Φ) with respect to which $\phi : A \simeq A^\vee$ is K-linear. Such structures are uniquely algebraic (over $\mathbf{C}$) and have a unique descent to $\overline{\mathbf{Q}}$. Finally, let V/Λ be the analytic exponential uniformization of $A(\mathbf{C})$ and let $\Psi_\phi : \Lambda \times \Lambda \to \mathbf{Z}(1)$ be the Riemann form corresponding to the polarization ϕ. The full level-n structure ι of type ζ on (A,ϕ) corresponds to a certain kind of embedding θ of $(\mathbf{Z}/n\mathbf{Z})^{2g}$ into the torsion subgroup $\Lambda_{\mathbf{Q}}/\Lambda$ of V/Λ, where $\Lambda_{\mathbf{Q}} = \mathbf{Q} \otimes_{\mathbf{Z}} \Lambda$. Our problem is to give an analytic description of the triple $(A^\sigma, \phi^\sigma, \iota^\sigma)$ for any $\sigma \in \mathrm{Gal}(\overline{\mathbf{Q}}/kK^*)$, where $K^* \subseteq \overline{\mathbf{Q}}$ is the reflex field of (K,Φ).

Here is the answer one obtains from Theorem 7.4.1. Let $s \in \mathbf{A}_{K^*}^{\times}$ be an idele such that $(s|K^*) = \sigma|_{(K^*)^{\mathrm{ab}}}$. The rational vector space $\Lambda_{\mathbf{Q}}$ is a 1-dimensional vector space over K, so finite ideles of K (such as $\mathrm{N}_{\mathrm{ref}}(1/s)$) may be "multiplied" against lattices (such as Λ) in this rational vector space. The complex torus $A^{\sigma}(\mathbf{C})$ is given by $V/\mathrm{N}_{\mathrm{ref}}(1/s)\Lambda$. The polarization ϕ^{σ} corresponds to the $\mathbf{Z}(1)$-valued Riemann form on $\mathrm{N}_{\mathrm{ref}}(1/s)\Lambda$ whose induced $\mathbf{Q}(1)$-valued pairing on $\Lambda_{\mathbf{Q}}$ is $q_{\sigma,s}\Psi_{\phi,\mathbf{Q}}$, where $q_{\sigma,s}$ is the unique positive rational generator of the fractional $\mathbf{Q}$-ideal associated to the $\mathbf{Q}$-idele $\mathrm{N}_{K^*/\mathbf{Q}}(s)$. The level-$n$ structure corresponds to the composite embedding

$$(\mathbf{Z}/n\mathbf{Z})^{2g} \overset{\theta}{\hookrightarrow} \Lambda_{\mathbf{Q}}/\Lambda \overset{\mathrm{N}_{\mathrm{ref}}(1/s)}{\simeq} \Lambda_{\mathbf{Q}}/\mathrm{N}_{\mathrm{ref}}(1/s)\Lambda = A^{\sigma}(\mathbf{C})_{\mathrm{torsion}} = A^{\sigma}(\overline{\mathbf{Q}})_{\mathrm{torsion}}$$

because (7.4.1) commutes. (This is a full level-n structure of type ζ because σ is assumed to act trivially on the nth cyclotomic field k inside of $\overline{\mathbf{Q}}$.)

References

[KM] N. Katz, B. Mazur, *Arithmetic moduli of elliptic curves*, Princeton Univ. Press, 1985.

[L] S. Lang, *Elliptic functions* (2nd ed.), Springer-Verlag, 1987.

[M1] J. Milne, "Abelian varieties," Ch. V in *Arithmetic geometry* (Cornell/Silverman ed.), Springer-Verlag, 1986.

[M2] J. Milne, "Jacobian varieties," Ch. VII in *Arithmetic geometry* (Cornell/Silverman ed.), Springer-Verlag, 1986.

[Mum] D. Mumford, *Abelian varieties*, Oxford University Press, Bombay, 1970.

[GIT] D. Mumford, *et al.*, *Geometric invariant theory* (3rd ed.), Springer-Verlag, 1994.

[Oda] T. Oda, *The first de Rham cohomology and Dieudonné modules*, Ann. Sci. École Norm. Sup. (4) **2** (1969), pp. 63–135.

[PS] B. Poonen, M. Stoll, *The Cassels–Tate pairing on polarized abelian varieties*, Annals of Math. 150 **2** (1999), pp. 1109–1149.

[S] G. Shimura, *Introduction to the arithmetic theory of automorphic forms*, Princeton Univ. Press, Princeton, 1970.

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA
E-mail address: bdconrad@umich.edu