Math 248B. Algebraic theory of q-expansions

1. Basic definitions

Let $X = X_1(N)$ for $N \geq 5$, viewed as a scheme over $\mathbf{Z}[1/N]$, and let $f : \mathbf{E} \rightarrow X$ be the universal generalized elliptic curve. In particular, $\mathbf{E}^{\mathrm{sm}} \rightarrow X$ has multiplicative fibers along the closed complement $X - Y$ of the open locus $Y = Y_1(N)$ over which $\mathbf{E}$ is smooth.

Remark 1.1. Let's explain how to think about $\mathbf{E}^{\mathrm{sm}}$ without reference to generalized elliptic curves. Since $X_{\mathbf{Q}}$ is Dedekind, it makes sense to speak of the Néron model over $X_{\mathbf{Q}}$ of the universal elliptic curve over $Y_{\mathbf{Q}}$. This Néron model turns out to be $\mathbf{E}^{\mathrm{sm}}$ (a proof requires some thought, due to component group issues in the non-proper fibers). However, since X is 2-dimensional there is no concept of "Néron model" over X.

Over X we can at least describe the relative identity component $(\mathbf{E}^{\mathrm{sm}})^0$: this is a semi-abelian scheme over X extending the universal elliptic curve over Y. This latter viewpoint in terms of multiplicative degeneration works well in the analytic theory, as well as in higher relative dimension.

Let $\omega = \omega_{\mathbf{E}/X}$ be the associated line bundle on X extending $(f_Y)_*(\Omega^1_{\mathbf{E}_Y/Y})$. In fancy terms, $\omega := f_*(K_{\mathbf{E}/X})$ where $K_{\mathbf{E}/X}$ on $\mathbf{E}$ is the (invertible) relative dualizing sheaf. In equivalent but more group-theoretic terms that carry over to the analytic and higher-dimensional settings, upon identifying $\Omega^1_{\mathbf{E}_Y/Y}$ with $e^*(\Omega^1_{\mathbf{E}_Y/Y}) = \mathrm{Lie}(\mathbf{E}_Y/Y)^\vee$ we have $\omega = e^*(\Omega^1_{\mathbf{E}^{\mathrm{sm}}/X}) = \mathrm{Lie}(\mathbf{E}^{\mathrm{sm}}/X)^\vee$. Recall from class that both viewpoints on ω have merits. For example, the relative dualizing sheaf viewpoint is needed in the construction of moduli spaces, such as to verify that "stable genus-1 curve" is an open condition on the base of a proper flat family of curves. On the other hand, the relative Lie algebra viewpoint is needed to make contact with formal groups, as is needed to affirm the consistency between the analytic theory of q-expansions (and spaces of modular forms) and the algebraic definitions of q-expansions and modular forms over $\mathbf{C}$ (to be reviewed below, and generalized significantly).

In class we proved that under the identification of $\Gamma(Y^{\mathrm{an}}_{\mathbf{C}}, (\omega^{\mathrm{an}}_{\mathbf{C}})^{\otimes k})$ with the huge vector space of holomorphic quasi-modular forms of weight k for $\Gamma_1(N)$ (i.e., without any growth condition at the cusps), the subspace $\Gamma(X^{\mathrm{an}}_{\mathbf{C}}, (\omega^{\mathrm{an}}_{\mathbf{C}})^{\otimes k})$ is carried isomorphically onto $M_k(\Gamma_1(N))$. By Serre's GAGA theorem, the natural analytification map

$$\Gamma(X_{\mathbf{C}}, \omega_{\mathbf{C}}^{\otimes k}) \rightarrow \Gamma(X^{\mathrm{an}}_{\mathbf{C}}, (\omega^{\mathrm{an}}_{\mathbf{C}})^{\otimes k}) = M_k(\Gamma_1(N))$$

is an isomorphism, and in class we proved that the classical notion of q-expansion at ∞ on the target agrees with the following purely algebraic definition of q-expansions over $\mathbf{C}$.

Fix a choice of $i = \sqrt{-1} \in \mathbf{C}$ and consider the generator $e^{2\pi i/N}$ of $\mu_N(\mathbf{C})$. Using the canonical closed subgroup inclusion of μ_N into $\mathrm{Tate}^{\mathrm{sm}}$ over $\mathbf{Z}[\![q]\!]$, by scalar extension to $\mathbf{C}[\![q]\!]$ we likewise identify μ_N as a closed $\mathbf{C}[\![q]\!]$-subgroup of $\mathrm{Tate}(q)^{\mathrm{sm}}$, so $e^{2\pi i/N}$ is thereby a $\Gamma_1(N)$-structure on the generalized elliptic curve $\mathrm{Tate}(q)$ over $\mathbf{C}[\![q]\!]$ (with 1-gon special fiber). This $\Gamma_1(N)$-structure defines a map $\mathrm{Spec}\,\mathbf{C}[\![q]\!] \rightarrow X$ *and* an identification of $\mathrm{Tate}(q)$ with the pullback of $\mathbf{E} \rightarrow X$ along this map. In this way, identify $(\mathrm{d}t/t)^{\otimes k}$ with an $\mathscr{O}^\wedge_{X_{\mathbf{C}},\infty}$-basis of the completed stalk $(\omega^{\otimes k}_{\mathbf{C},\infty})^\wedge$, and likewise identify $\mathscr{O}^\wedge_{X_{\mathbf{C}},\infty}$ with $\mathbf{C}[\![q]\!]$.

Definition 1.2. For $f \in \Gamma(X_{\mathbf{C}}, \omega_{\mathbf{C}}^{\otimes k})$, the *$q$-expansion* $f_\infty \in \mathbf{C}[\![q]\!]$ is the coefficient of f in the completed stalk

$$(\omega^{\otimes k}_{\mathbf{C},\infty})^\wedge = \mathbf{C}[\![q]\!](\frac{\mathrm{d}t}{t})^{\otimes k}$$

under the identifications defined above. That is, $f = f_\infty (\mathrm{d}t/t)^{\otimes k}$.

Remark 1.3. As stated, this definition depends on a choice of i (through the specified level structure $e^{2\pi i/N} \in \mu_N$): if we change the choice then the map $\mathrm{Spec}\,\mathbf{C}[\![q]\!] \rightarrow X_{\mathbf{C}}$ does *not* change (so the isomorphism $\mathscr{O}^\wedge_{X_{\mathbf{C}},\infty} \simeq \mathbf{C}[\![q]\!]$ is canonical) but the isomorphism of $\mathrm{Tate}(q)$ with the pullback of $\mathbf{E} \rightarrow X$ along this map changes by inversion. That in turn changes the identification of $\mathrm{d}t/t$ with a basis of $\omega^\wedge_{\mathbf{C},\infty}$ by a sign, and so changes the image of $(\mathrm{d}t/t)^{\otimes k}$ in this completed stalk by a sign of $(-1)^k$.

There are two other roles of a choice of i that arise when converting q-expansions into holomorphic functions on a half-plane. First, we have to choose a connected component of $\mathbf{C} - \mathbf{R}$ to which we pull back a holomorphic function on Δ^* via $\tau \mapsto q = e^{2\pi i\tau}$. Second, we have to choose a uniformization $\mathbf{C} \to \mathbf{C}^\times$ via $z \mapsto w = e^{2\pi iz}$ to convert global sections of globally nontrivial line bundles into global holomorphic functions. Note that $\mathrm{d}w/w = 2\pi i\mathrm{d}z$, so $(\mathrm{d}w/w)^{\otimes k} = (2\pi i)^k(\mathrm{d}z)^{\otimes k}$. This latter appearance of i has a sign change effect of $(-1)^k$, and exactly matches the choice used to pick $e^{2\pi i/N} \in \mu_N(\mathbf{C})$ above. The choice of i implicit in $q = e^{2\pi i\tau}$ (or in other words, the choice of connected component of $\mathbf{C} - \mathbf{R}$) is actually *invisible*, since modular forms viewed over $\mathbf{C} - \mathbf{R}$ satisfy a transformation law relative to subgroups of $\mathrm{GL}_2(\mathbf{Z})$. More specifically, for the connected moduli space of $\Gamma_1(N)$-structures we do have invariance under $\left(\begin{smallmatrix} -1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ (which swaps τ and $-\tau$).

Since Definition 1.2 computes the classical q-expansion correctly when $\mathbf{C}$ is viewed as a $\mathbf{Z}[1/N, \zeta_N]$-algebra via $\zeta_N \mapsto e^{2\pi i/N}$ (as we discussed in class), we are now motivated to make some algebraic definitions with rather general coefficients rings. These definitions are given in the next section.

The aim of this handout is to prove that q-expansions as defined algebraically below can detect the "field of definition", or even "ring of definition", of a holomorphic modular form relative to algebraic models of modular curves. For simplicity we will focus throughout on the case of $X_1(N)$. This modular curve has the virtue of being smooth and proper over $\mathbf{Z}[1/N]$ with *geometrically connected* fibers. This is the reason that in the algebraic theory it will suffice to work with a single cusp (just like in the analytic theory: analytic continuation is great on a *connected* complex manifold).

Some of the techniques below were introduced by Katz in his great paper "p-adic properties of modular curves and modular forms". Variants can be introduced to handle non-smooth modular curves in the presence of geometrically reduced but reducible fibers (such as when considering congruences modulo primes that divide the "level"). In such generality one needs to use q-expansions along a collection of cusps big enough so that each irreducible component of each fiber over the base ring meets some cusp under consideration.

2. Preparations with ∞

Let $\mathscr{O} = \mathbf{Z}[1/N, \zeta_N]$, and let A be an $\mathscr{O}$-algebra. (The most basic case of interest is when A is a subring of $\mathbf{C}$, or even a subfield of $\mathbf{C}$, with $\mathscr{O}$ embedded into $\mathbf{C}$ via $\zeta_N \mapsto e^{2\pi i/N}$ for a fixed choice of $i = \sqrt{-1} \in \mathbf{C}^\times$). Over $\mathbf{Z}[\![q]\!][1/N, \zeta_N] = \mathscr{O} \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ we have the $\Gamma_1(N)$-structure $\zeta_N \in \mu_N \hookrightarrow \mathrm{Tate}(q)^{\mathrm{sm}}_{\mathscr{O}\otimes_{\mathbf{Z}}\mathbf{Z}[\![q]\!]}$, and this defines a map

$$\mathrm{Spec}(\mathscr{O} \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]) \to X_{\mathscr{O}}$$

which lifts a section $\infty : \mathrm{Spec}\,\mathscr{O} \to X_{\mathscr{O}}$. Likewise after base change along $\mathscr{O} \to A$, we get a map

$$\mathrm{Spec}(A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]) \to X_A.$$

Setting $q = 0$ gives a section $\infty_A : \mathrm{Spec}\,A \to X_A$.

Remark 2.1. It can be proved that the map $\mathrm{Spec}(A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]) \to X_A$ identifies $A[\![q]\!]$ with the coordinate ring of the formal completion of X_A along ∞_A.

Definition 2.2. For any $\mathbf{Z}[1/N]$-algebra A, the A-module $M_k(\Gamma_1(N), A)$ is $\Gamma(X_A, \omega_A^{\otimes k})$. This is covariant in A in the evident manner.

For $A = \mathbf{C}$, the preceding definition recovers the classical space $M_k(\Gamma_1(N))$ via the GAGA isomorphism, as discussed above. More generally, for any A-algebra B there is a natural base change morphism

$$B \otimes_A M_k(\Gamma_1(N), A) \to M_k(\Gamma_1(N), B)$$

which is an isomorphism when B is A-flat. (The formation of global sections of a quasi-coherent sheaf on a quasi-compact and separated A-scheme commutes with flat base change on A.)

Example 2.3. Since $\mathbf{Z}[1/N] \to \mathbf{C}$ is flat, we have

$$\mathbf{C} \otimes_{\mathbf{Z}[1/N]} M_k(\Gamma_1(N), \mathbf{Z}[1/N]) \simeq M_k(\Gamma_1(N), \mathbf{C}).$$

Hence, $M_k(\Gamma_1(N), \mathbf{Z}[1/N])$ defines a canonical $\mathbf{Z}[1/N]$-structure on the space of holomorphic modular forms of weight k for $\Gamma_1(N)$. Is this the space of holomorphic forms f whose q-expansion coefficients lie in $\mathbf{Z}[1/N]$? Typically *no*. The reason is that although X is a $\mathbf{Z}[1/N]$-structure on the algebraization $X_{\mathbf{C}}$ of the classical modular curve, for the cusp ∞ we have only obtained it as a section over $\mathbf{Z}[1/N, \zeta_N]$ (with $\zeta_N \mapsto e^{2\pi i/N}$). We cannot expect to dig deeper than algebras over $\mathscr{O}$ based on the preceding definitions.

For our purposes, the main point is that for any $\mathscr{O}$-algebra A and global section $f \in M_k(\Gamma_1(N), A)$ of $\omega_A^{\otimes k}$ over X_A, its pullback over $\mathrm{Spec}(A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!])$ is identified with an element of

$$\omega_{\mathrm{Tate}(q)/\mathbf{Z}[\![q]\!]} \otimes_{\mathbf{Z}} A = (A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!])(\mathrm{d}t/t)^{\otimes k}.$$

This brings us to:

Definition 2.4. Using notation as above, the pullback of f has the form $f_\infty(\mathrm{d}t/t)^{\otimes k}$ for a unique $f_\infty \in A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$. This is the *$q$-expansion of f along ∞_A*.

Remark 2.5. Beware that an argument is required to justify that f_∞ determines f uniquely (as we know for $A = \mathbf{C}$ via the analytic theory). We will address this in Lemma 3.2.

The "justification" for the terminology is three-fold: the natural map $A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to A[\![q]\!]$ is *injective* (by Lemma 2.6 below), the image of f_∞ in $A[\![q]\!]$ is the multiplier coefficient when computing the pullback of f to the formal completion of X_A along ∞_A (by Remark 2.1), and for $A = \mathbf{C}$ we thereby recover Definition 1.2 (which we have seen in class is consistent with the analytic definition of q-expansions).

Lemma 2.6. *For any $\mathbf{Z}$-module M, the natural map $M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to M[\![q]\!]$ is injective.*

This map is generally *not* surjective when M is not finitely generated over $\mathbf{Z}$; e.g., for $M = \mathbf{Q}$ the image consists of power series of $\mathbf{Q}$ with bounded denominators.

Proof. Since an element of the tensor product is a finite sum of elementary tensors, and an injection $M' \to M$ induces an injection $M'[\![q]\!] \to M[\![q]\!]$, to prove the vanishing of the kernel we may replace M with a finitely generated $\mathbf{Z}$-submodule (namely, arising from an expression for a hypothetical element in the kernel as a sum of elementary tensors). Then by the good behavior with respect to *finite* direct sum decompositions in M we reduce to two cases: $M = \mathbf{Z}$ and $M = \mathbf{Z}/n\mathbf{Z}$. In both of these cases (and hence for the case of finitely generated M in general) the map of interest is clearly an isomorphism. ∎

In view of the injectivity of this lemma, when we speak of f_∞ it is no loss of information to work in $A[\![q]\!]$ rather than in $A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$. Nonetheless, the fact that the q-expansion in the $A[\![q]\!]$-sense actually lies in the subring $A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ tells us something rather non-obvious from the viewpoint of the analytic theory of q-expansions:

Proposition 2.7 (bounded denominators)**.** *Let $M' \hookrightarrow M$ be an injection of $\mathbf{Z}$-modules. Inside of $M[\![q]\!]$, the intersection of $M'[\![q]\!]$ and $M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is equal to $M' \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$.*

In particular, if f is a holomorphic cusp form for $\Gamma_1(N)$ and its q-expansion in $\mathbf{C}[\![q]\!]$ lies in $K[\![q]\!]$ for a number field K then it has bounded denominators (i.e., there exists a nonzero $c \in \mathscr{O}_K$ such that cf has q-expansion in $\mathscr{O}_K[\![q]\!]$.

Note that we really are getting bounded denominators at *all* primes, even those dividing N. For this it is crucial that we have the Tate curve available over $\mathbf{Z}[\![q]\!][1/N]$ and not just over the much larger ring $\mathbf{Z}[1/N][\![q]\!]$.

Proof. The second claim follows from the first by the compatibility of algebraic and analytic q-expansions over $\mathbf{C}$ (discussed in class) by taking $M' = K$ and $M = \mathbf{C}$. The point is that $K \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is the localization at $\mathscr{O}_K - \{0\}$ of $\mathscr{O}_K \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] = \mathscr{O}_K[\![q]\!]$ (equality since $\mathscr{O}_K$ is a finitely generated $\mathbf{Z}$-module, even finite free).

To prove the first claim, note that $\mathbf{Z} \to \mathbf{Z}[\![q]\!]$ is *flat*, so the diagram

$$0 \to M' \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to (M/M') \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to 0$$

is a short exact sequence. But by Lemma 2.6, the final term injects into $(M/M')[\![q]\!] = M[\![q]\!]/M'[\![q]\!]$. Thus, $M' \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is the set of elements in $M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ whose image in $M[\![q]\!]$ lies in $M'[\![q]\!]$, and that in turn is exactly the assertion we are trying to prove. ∎

3. Descending a ring of definition

Consider a map $A \to B$ of $\mathscr{O}$-algebras. (We do *not* assume these are $\mathbf{Z}$-flat, or even subrings of $\mathbf{C}$.) Since X is $\mathbf{Z}[1/N]$-flat, if $A \to B$ is injective then the pullback along the map $X_B \to X_A$ is injective for global sections of vector bundles. In particular, the natural pullback map

$$M_k(\Gamma_1(N), A) = \mathrm{H}^0(X_A, \omega_A^{\otimes k}) \to \mathrm{H}^0(X_B, \omega_B^{\otimes k}) = M_k(\Gamma_1(N), B)$$

is injective in such cases. In general there are respective q-expansion maps to $A[\![q]\!]$ and $B[\![q]\!]$ from the two sides, using the Tate curve construction described above, and upon unraveling the definitions it is not hard to verify the *commutativity* of the diagram

$$\begin{array}{ccc} \mathrm{H}^0(X_A, \omega_A^{\otimes k}) & \longrightarrow & \mathrm{H}^0(X_B, \omega_B^{\otimes k}) \\ \downarrow & & \downarrow \\ A[\![q]\!] & \longrightarrow & B[\![q]\!] \end{array} \tag{3.1}$$

where the vertical maps are the q-expansion maps and the bottom map is the natural map.

Example 3.1. Let $A = \mathbf{C}$ and $B = \mathbf{C}$ but let $A \to B$ be an abstract field automorphism over $\mathbf{Q}(\zeta_N)$. The top map in (3.1) then corresponds to the "global" $\mathrm{Aut}(\mathbf{C}/\mathbf{Q}(\zeta_N))$-action on $M_k(\Gamma_1(N)) = \mathbf{C} \otimes_{\mathbf{Q}(\zeta_N)} M_k(\Gamma_1(N), \mathbf{Q}(\zeta_N))$ provided by the $\mathbf{Q}(\zeta_N)$-structure $X_{\mathbf{Q}(\zeta_N)}$ on the algebraization $X_{\mathbf{C}}$ of the classical modular curve $X_{\Gamma_1(N)}$. The commutativity of (3.1) implies that this automorphism group action corresponds exactly to the action on q-expansion coefficients! Note that this consistency of actions is not a definition (though it also doesn't lie too deep).

As a consequence of the commutativity of (3.1), if $A \to B$ is injective then both horiztonal maps are injective. An important point is that the vertical maps are *always* injective (regardless of $A \to B$):

Lemma 3.2. *For any $\mathscr{O}$-algebra A, the q-expansion map*

$$\mathrm{H}^0(X_A, \omega_A^{\otimes k}) \to A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \hookrightarrow A[\![q]\!]$$

is injective.

Proof. We now use an elegant trick due to Katz, replacing rings with modules to acquire more variation. To be precise, since $X_A \to X_{\mathscr{O}}$ is an affine map, we can identify the left side with the global sections on $X_{\mathscr{O}}$ of the quasi-coherent pushforward sheaf $A \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k}$. Now the role of A as a ring has gone away: it only matters as an $\mathscr{O}$-module! That is, for any $\mathscr{O}$-module M we can use pullback along $\mathrm{Spec}\, \mathscr{O}[\![q]\!] \to X_{\mathscr{O}}$ to define a q-expansion map

$$\mathrm{H}^0(X_{\mathscr{O}}, M \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k}) \to M \otimes_{\mathscr{O}} (\mathscr{O} \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]) = M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$$

and when $M = A$ is an $\mathscr{O}$-algebra this construction recovers the earlier one (check!).

The formation of this construction for any M is functional in M and even compatible with the formation of direct limits in M. Hence, to prove the injectivity in general we may express M as a direct limit of finitely generated $\mathscr{O}$-submodules and thereby reduce the problem to the case when M is $\mathscr{O}$-finite (but possibly not free, nor even torsion-free).

If $0 \to M' \to M \to M'' \to 0$ is an exact sequence of $\mathscr{O}$-modules then

$$0 \to M' \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k} \to M \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k} \to M'' \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k} \to 0$$

is short exact as a sequence of quasi-coherent sheaves on $X_{\mathscr{O}}$ since $\omega_{\mathscr{O}}$ is a line bundle on the $\mathscr{O}$-flat $X_{\mathscr{O}}$. Thus, the diagram of global sections is left-exact. The flat scalar extension of $0 \to M' \to M \to M'' \to 0$ against the $\mathscr{O}$-flat $\mathscr{O} \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ also yields a left-exact (even short exact) sequence, so a simple diagram chase shows that the injectivity result for M' and M'' implies the same for M. Thus, to solve our problem for $\mathscr{O}$-finite M it suffices to separately treat the torsion-case and the torsion-free case (using $M' = M_{\mathrm{tor}}$ and $M'' = M/M_{\mathrm{tor}}$). But in the torsion-free case we can embed M into a finite free $\mathscr{O}$-module. Thus, by the compatibility with direct sums and the structure of finitely generated torsion $\mathscr{O}$-modules (and a few more applications of the preceding argument with short exact sequences, applied to a composition series for a finitely generated torsion $\mathscr{O}$-module), we finally reduce to two basic cases: $M = \mathscr{O}$ and $M = \mathscr{O}/\mathfrak{m}$ for a maximal ideal $\mathfrak{m}$ of $\mathscr{O}$.

In terms of our original problem, the above formal manipulations have reduced the problem to the special cases $A = \mathscr{O}$ and $A = \kappa = \mathscr{O}/\mathfrak{m}$ for a maximal ideal $\mathfrak{m}$ of $\mathscr{O}$. For the case $A = \mathscr{O}$, we can use the injection into the fraction field K to reduce to treating K. Hence, it suffices to prove that the two q-expansion maps

$$\mathrm{H}^0(X_K, \omega_K^{\otimes k}) \to K[\![q]\!],\ \ \mathrm{H}^0(X_\kappa, \omega_\kappa^{\otimes k}) \to \kappa[\![q]\!]$$

are injective. By the identification of $K[\![q]\!]$ with the completed local ring of X_K at the rational point $\infty \in X_K(K)$, the case $A = K$ reduces to the assertion that if the global section of a line bundle on a *connected* smooth variety vanishes in the completed stalk at some point (and hence vanishes in the actual stalk, by injectivity of completion for local noetherian rings) then the global section is zero. This in turn follows from the fact that a smooth connected variety is *integral*, so the space of global sections injects into the generic stalk, which in turn contains all other stalks.

The argument just used over K can be applied equally well over κ, due to X_κ being smooth and *connected*, provided that we can identify the q-expansion process with computing the completed stalk at ∞_κ. That is, the map $\mathrm{Spec}(\kappa[\![q]\!]) \to X_\kappa$ arising from the Tate curve construction induces a local κ-algebra map $\mathscr{O}^\wedge_{X_\kappa,\infty} \to \kappa[\![q]\!]$ between complete discrete valuation rings with residue field κ, and we wish to prove that this map is an isomorphism. Exactly as we saw in class over K, it suffices to prove surjectivity onto $\kappa[\![q]\!]/(q^2)$, which amounts to the Tate curve modulo q^2 (over κ) being a *non-trivial* equicharacteristic deformation of its 1-gon special fiber.

To prove the nontriviality, it suffices to find some "invariant" that vanishes for the trivial first-order deformation over κ but is nonzero for the Tate curve modulo q^2. The invariant is found in the annihilator ideal of the coherent sheaf $\Omega^2_{\mathrm{Tate}(q)/\kappa[\![q]\!]}$. This sheaf is physically supported at the singularity in the special fiber, so its annihilator ideal cuts out an infinitesimal structure supported at that point. We claim that this ideal contains the element q that is *nonzero* modulo q^2, contrasting with the elementary fact that for the trivial deformation the ideal has vanishing intersection with the base ring (even modulo q^2). To do the calculation of the annihilator ideal, observe that the formal singularity is $R = \kappa[\![q]\!][\![u, v]\!]/(uv - q)$, and

$$\widehat{\Omega}^1_{R/\kappa[\![q]\!]} = (R\mathrm{d}u \oplus R\mathrm{d}v)/(u\mathrm{d}v + v\mathrm{d}u),$$

so passing to second exterior powers gives the R-module $R/(u, v)$. This has R-annihilator (u, v), which contains the element $uv = q$ of $\kappa[\![q]\!]$ that is nonzero modulo q^2. ■

Theorem 3.3. *Let $A \to B$ be an injective map of $\mathscr{O}$-algebras, and choose $f \in \mathrm{H}^0(X_B, \omega_B^{\otimes k})$. If the q-expansion $f_\infty \in B[\![q]\!]$ lies in $A[\![q]\!]$ then f lies in the A-submodule $\mathrm{H}^0(X_A, \omega_A^{\otimes k})$.*

In particular, if f is a holomorphic modular form on $\Gamma_1(N)$ and its q-expansion coefficients lie in an $\mathscr{O}$-subalgebra R of $\mathbf{C}$ then $f \in \mathrm{H}^0(X_R, \omega_R^{\otimes k})$.

The most "popular" applications of the final part of the theorem are with R a number field, or localized integer ring thereof (all containing $\mathscr{O}$). We emphasize that this theorem does not make any assertion concerning $\mathbf{Z}[1/N]$-subalgebras of $\mathbf{C}$ not containing a primitive Nth root of unity.

Proof. By Proposition 2.7 and Lemma 3.2, it is equivalent to prove that the intersection

$$\mathrm{H}^0(X_B, \omega_B^{\otimes k}) \cap (A \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!])$$

inside of $B \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is equal to $\mathrm{H}^0(X_A, \omega_A^{\otimes k})$. Once again using the trick of viewing X_B and X_A as affine over $X_{\mathscr{O}}$, we can recast the problem relative to a general injection $M' \to M$ of $\mathscr{O}$-modules (in place of the injection $A \to B$ of $\mathscr{O}$-algebras): we claim that

$$\mathrm{H}^0(X_{\mathscr{O}}, M \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k}) \cap (M' \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!])$$

inside of $M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is equal to $\mathrm{H}^0(X_{\mathscr{O}}, M' \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k})$. (This generalization "makes sense", since the proof of Lemma 3.2 proved the injectivity of the q-expansion construction with coefficients in any $\mathscr{O}$-module, and $M' \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!] \to M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ is injective due to the $\mathbf{Z}$-flatness of $\mathbf{Z}[\![q]\!]$.)

An element of the intersection has q-expansion in $M \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$ that vanishes in $(M/M') \otimes_{\mathbf{Z}} \mathbf{Z}[\![q]\!]$, so by the functoriality of the q-expansion construction with $\mathscr{O}$-module coefficients and its *injectivity* in general (shown in the proof of Lemma 3.2), it follows that an element in the intersection of interest has *vanishing* image in $\mathrm{H}^0(X_{\mathscr{O}}, (M/M') \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k})$. But the diagram of quasi-coherent sheaves

$$0 \to M' \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k} \to M \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k} \to (M/M') \otimes_{\mathscr{O}} \omega_{\mathscr{O}}^{\otimes k}$$

on $X_{\mathscr{O}}$ is exact since $\omega_{\mathscr{O}}$ is a line bundle on the $\mathscr{O}$-flat $X_{\mathscr{O}}$, so the associated diagram of global sections is also exact. This gives the result. ■