The modularity theorem

Avi Zeff

Seminar on Fermat's last theorem

The goal of today's talk is to go into slightly more detail about how the modularity theorem is actually proved. Since the proof is a hundred-page paper assuming a tremendous amount of background, we will not give anything approaching a full proof; but we will try to sketch some of the ideas that go into it.

Recall that the modularity theorem states that every (semistable) elliptic curve E over $\mathbb{Q}$ corresponds to a Hecke eigenform f, with the property that if we write out the q-expansion

$$f = \sum_{n \geq 1} a_n q^n$$

(so a_n is the eigenvalue of the nth Hecke operator) then $\#E(\mathbb{F}_p) = p + 1 - a_p$ for each prime number p (not dividing the level of the form f). This can equivalently be phrased as an equality of L-functions $L(f, s) = L(E, s)$.

We've noted before that an elliptic curve being modular can be viewed as "really" a property of its associated Galois representations $T_p(E) = \varprojlim_n E[p^n]$. Note that there's some ambiguity: we can choose the prime p! So part of the claim is: if $T_p(E)$ is modular, then so is E, and therefore so is each $T_\ell(E)$ for $\ell \neq p$ as well.

We'll approach our goal to prove that any semistable elliptic curve E over $\mathbb{Q}$ is modular via two steps. First, we'll discuss a "modularity lifting theorem," which lets us use a simpler criterion: it says that (under some technical conditions) if the mod p representation $E[p]$ (which we recall is a rank 2 vector space over $\mathbb{F}_p$) is modular, then so is $T_p(E)$. Then we'll look at the possibilities for $T_p(E)$. In particular since we can use any prime p, we'll choose explicitly $p = 3$, which is essentially the simplest case; this will often work, and when it doesn't we'll see that we can get around the obstruction by also using $p = 5$. (This is sort of like an induction argument: the induction step is via the modularity lifting theorem, and the base case is via some more classical results together with the 3-5 trick.)

The modularity lifting theorem is often expressed as an identity $\mathcal{R} = \mathbb{T}$. Before we can say anything about this, we need to answer two questions: what is $\mathcal{R}$? and what is $\mathbb{T}$? In the course of answering these questions, we'll see that we can actually find a natural map $\mathcal{R} \to \mathbb{T}$, so the main issue will be to show that this is an isomorphism. We'll skip over how this is done, though we may return to it at the end if there's time.

We'll start with the ring $\mathcal{R}$. This is a very interesting object, defined as the universal deformation ring of a Galois representation; so we first of all want to say what this means.

1 Deformations of Galois representations

Let G be a group. We are used to studying representations of G over $\mathbb{C}$, i.e. complex vector spaces V together with an action of G, i.e. a group homomorphism $G \to \mathrm{GL}(V)$; but we've also sometimes seen representations over finite fields $\mathbb{F}_p$ such as the p-torsion points $E[p]$ or even over rings which are not fields such as $\mathbb{Z}_p$ in the case of the Tate module $T_p(E)$.

More generally, let's say that a representation of G over an arbitrary (commutative) ring R is an R-module M together with a G-action on M, compatible with the R-module structure; equivalently this is a group homomorphism $\rho : G \to \operatorname{Aut}_R(M)$. When R is a field, M is a vector space, which is particularly nice; even when R is not a field, the most common situation of interest is for M to be a free module, as for example in the case of $T_p(E)$ over $\mathbb{Z}_p$. We'll write $\operatorname{Rep}_R(G)$ for the set of representations of G over R, elements of which we'll often abbreviate to ρ. (In fact, this naturally has the structure of a category, but we generally won't need this; you can often think of functors when we talk about maps between sets (or categories) of representations.)

Now, let's say we have two rings R and S, and a ring homomorphism $R \to S$. Suppose we have a representation of G over R, i.e. an R-module M with G-action; let's imagine for concreteness that M is actually free, so $M \simeq R^{\oplus r}$, so the G-action is given by a homomorphism $\rho : G \to \mathrm{GL}_r(R)$. The ring homomorphism $R \to S$ induces a group homomorphism $\mathrm{GL}_r(R) \to \mathrm{GL}_r(S)$ (by taking the ring map on each entry of the matrices), so composing gives a map

$$G \xrightarrow{\rho} \mathrm{GL}_r(R) \to \mathrm{GL}_r(S),$$

which is to say a G-action on the free S-module $S^{\oplus r}$. Thus we've produced a (free) G-representation over S from a (free) G-representation over R; in fact a similar operation works for non-free representations as well, so we get a map

$$\operatorname{Rep}_R(G) \to \operatorname{Rep}_S(G).$$

Consider for example a ring like $R = \mathbb{C}[T]$, the ring of polynomials in one variable over T, and take $S = \mathbb{C}$. In principle, we could have a lot of maps $\mathbb{C} \to \mathbb{C}$, so maps $\mathbb{C}[T] \to \mathbb{C}$ are even more complicated; but we'll say for simplicity that these have to be maps of $\mathbb{C}$-algebras, i.e. they have to be the identity on the constant polynomials $\mathbb{C} \subset \mathbb{C}[T]$. Therefore such maps are completely determined by the image of T, which can be any complex number; that is, there is a natural bijection between homomorphisms of $\mathbb{C}$-algebras

$$f : \mathbb{C}[T] \to \mathbb{C}$$

and complex numbers $f(T)$.

Suppose we have a representation ρ of G over $\mathbb{C}[T]$, i.e. a $\mathbb{C}[T]$-module M with G-action. Then the construction above gives, for every map $f : \mathbb{C}[T] \to \mathbb{C}$, a $\mathbb{C}$-representation of G, which we might write as ρ_f. Since the maps f correspond to complex numbers, this gives us a whole family of complex representations of G, parametrized by complex numbers; to understand this whole family, we could instead think about the original representation ρ.

We can also do this in more unfamiliar contexts. For example, suppose we have a representation ρ on a module M over the integers $\mathbb{Z}$. Then by reducing modulo p, we obtain a $\mathbb{F}_p$-representation M/pM for every prime p, which we can again think of as a family of representations $\overline{\rho}_p$ parametrized by the set of prime numbers and induced by the original representation of ρ over $\mathbb{Z}$.

A related construction is: say we have a representation of G over $\mathbb{Z}/p^n\mathbb{Z}$ for some integer n. Reducing modulo p gives a representation over $\mathbb{F}_p$, and in fact reducing modulo p^m for any $m \leq n$ gives a representation over $\mathbb{Z}/p^m\mathbb{Z}$. This suggests a question: if we're given a representation of G over $\mathbb{F}_p$, when does it come from a representation over $\mathbb{Z}/p^n$, and for which n? How high can we go?

More generally, if we're concerned with *all* representations of G, we could hope that they *all* arise as a single family in some way like this: that is, we'd look for some ring R and representation ρ over R such that every representation over any ring S is given by applying this construction to some map $R \to S$. In this case we'd call ρ the universal representation of G, and R its universal representation ring.

This is usually too much to hope for. A good example of why comes from the case above where we looked at representations modulo p: if we just study representations of G over $\mathbb{F}_p$, we may still have many different ones, and if $\overline{\rho}_1$ and $\overline{\rho}_2$ are two different $\mathbb{F}_p$-representations then they can't come from the reduction modulo p of the same $\mathbb{Z}$-representation. In principle they could still come from a more complicated ring with multiple maps to $\mathbb{F}_p$, which sometimes happens, but in general we still get a lot of different "classes" of representations.

Thus the best we can really hope for in this situation is that all representations with the same reduction modulo p should come from the same "universal" representation. To make this a little more precise, we fix some base ring k (such as $\mathbb{F}_p$, or more generally typically a field), and consider the class of rings S equipped with maps $S \to k$ (which we often require to be surjective). If indeed $S \to k$ is a surjection onto a field, its kernel is a maximal ideal, which will sometimes be relevant.

Now, fix a G-representation $\overline{\rho}_0$ over k. For any representation ρ over S, the map $S \to k$ induces a representation $\overline{\rho}$ over k; we're only interested in representations ρ such that $\overline{\rho} \simeq \overline{\rho}_0$, i.e. whose reduction to k is given by the fixed representation $\overline{\rho}_0$. Such representations ρ are often called *deformations* of $\overline{\rho}_0$ to S, since the idea is often that S is some "enlargement" of k so that ρ is an extension of $\overline{\rho}_0$ to this larger space, which recovers $\overline{\rho}_0$ when we specialize back to k.

The hope is that all of these ρ really will come from some representation ρ^{univ} of G over some ring $R^{\text{univ}}_{\overline{\rho}_0}$ by the process above; that is, a deformation ρ of $\overline{\rho}_0$ to S is equivalent to a map $R^{\text{univ}}_{\overline{\rho}_0} \to S$, with ρ obtained via the construction above along this map from the representation ρ^{univ}. In this case ρ^{univ} is called the universal deformation, and $R^{\text{univ}}_{\overline{\rho}_0}$ the universal deformation ring of $\overline{\rho}_0$.

There are a lot of technical questions about when this is true that we will not get into; but the upshot is that in most cases of interest it will be true. This is not quite as miraculous as it sounds; it just means that we can translate the question of studying G-representations with a fixed reduction to k into the question of studying ring homomorphisms out of the ring $R^{\text{univ}}_{\overline{\rho}_0}$, which is a priori not necessarily any easier since we don't know anything about this ring. However, it turns out that there are ways to study it, so this will end up being a powerful tool.

The case that we're interested in is when G is the absolute Galois group $G = \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ of the rational numbers, and the representations that we're most interested in are Tate modules of elliptic curves $T_p(E)$. In this case, we know what the reduction modulo p looks like: it's $E[p]$, a rank 2 representation of G over $\mathbb{F}_p$, so $T_p(E)$ is a particular deformation of $E[p]$ to the ring $S = \mathbb{Z}_p$. We fix a particular mod p representation $\overline{\rho}$, so we're interested in elliptic curves E with $E[p] \simeq \overline{\rho}$ (of course, by varying $\overline{\rho}$ we can get all elliptic curves in one of these sets). The deformation $T_p(E)$ of $\overline{\rho}$ is then equivalent to a map $R^{\text{univ}}_{\overline{\rho}} \to \mathbb{Z}_p$. Since $\mathbb{Z}_p$

isn't a field, we can't describe this map as the quotient by a maximal ideal; but there's a natural inclusion $\mathbb{Z}_p \subset \mathbb{Q}_p$ giving a map $R_{\overline{\rho}}^{\text{univ}} \to \mathbb{Q}_p$ (corresponding to the $\mathbb{Q}_p$-representation $T_p(E) \otimes_{\mathbb{Z}_p} \mathbb{Q}_p$), which does correspond to a maximal ideal. Although there may be many maps $R_{\overline{\rho}}^{\text{univ}} \to \mathbb{Z}_p$, the composite maps $R_{\overline{\rho}}^{\text{univ}} \to \mathbb{F}_p$ should all agree, since they all correspond to the same $\mathbb{F}_p$-representation $\overline{\rho}$.

In fact, in general the universal deformation ring tells us much more than just about the lifts $T_p(E)$. For example, while the $T_p(E)$ are lifts all the way to $\mathbb{Z}_p$, or equivalently compatible systems of lifts to every $\mathbb{Z}/p^n\mathbb{Z}$, we could have representations $\overline{\rho}$ that don't lift all the way: for example, perhaps they lift to $\mathbb{Z}/p^3\mathbb{Z}$, but not further, or perhaps they don't even lift past $\mathbb{F}_p$. These properties are then reflected in the universal deformation ring $R_{\overline{\rho}}^{\text{univ}}$: if it is an $\mathbb{F}_p$-algebra, i.e. $p = 0$ in $R_{\overline{\rho}}^{\text{univ}}$, then it cannot have any maps to rings in which $p \neq 0$, so the deformation does not lift at all; if $p^2 \neq 0$ but $p^3 = 0$ in the universal deformation ring, then there exist lifts modulo p^3 but not higher. In the case $\overline{\rho} = E[p]$, the existence of the lift $T_p(E)$ means that $p^n \neq 0$ for all n in the universal deformation ring, i.e. p is not nilpotent.

One can also impose other conditions on the deformations ρ of $\overline{\rho}$ parametrized by maps from $R_{\overline{\rho}}^{\text{univ}}$. In the case of interest $G = \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, recall that there are various properties that Galois representations can have, e.g. being unramified at various points; one of the things we will require is that our representations be unramified at primes not dividing a fixed positive integer N (which will be our level), as well as more complicated conditions at primes dividing N whose details we won't get into. We'll write the universal deformation ring for such representations deforming our fixed $\mathbb{F}_p$-representation $\overline{\rho}$ as $\mathcal{R}$. It is a $\mathbb{Z}_p$-algebra, concretely of the form $\mathcal{R} \simeq \mathbb{Z}_p[[x_1, \ldots, x_n]]/I$ for some ideal I; and all of the Galois representations (over any ring S) with which we will be concerned can be thought of equivalently as ring homomorphisms $\mathcal{R} \to S$.

2 The Hecke algebra

We've seen Hecke operators before, which act on the space of modular forms. Given two Hecke operators, we can compose them, which we think of as a kind of multiplication; we've seen before that this multiplication is commutative. We can also add two Hecke operators, with $T_1 + T_2$ given by the operator $f \mapsto T_1(f) + T_2(f)$; this is compatible with the multiplication, and makes the space of Hecke operators into a commutative ring $\mathbb{T}$.

Given a Hecke eigenform f (of weight 2, as all our eigenforms will be) and a Hecke operator T, we can associate to them the Hecke eigenvalue a_T, i.e. $T(f) = a_T f$. If the eigenform f has rational coefficients (in its q-expansion), which is the kind we care about, then it defines a map

$$\begin{aligned} \mathbb{T} &\to \mathbb{Q} \\ T &\mapsto a_T \end{aligned}$$

which one can show is a ring homomorphism. Since the target is a field, the kernel is a maximal ideal $\mathfrak{m}_f$. More generally, we can identify Hecke eigenforms (with rational coefficients) with maps $\mathbb{T} \to \mathbb{Q}$, or equivalently with maximal ideals of $\mathbb{T}$.

We've mentioned before that there is a construction due to Eichler–Shimura that associates to each Hecke eigenform with rational coefficients f an elliptic curve E_f over $\mathbb{Q}$, such that for each prime p not dividing the level of f we have $\#E_f(\mathbb{F}_p) = p + 1 - a_p$. The modularity theorem can be viewed as the statement that this gives a bijection.

In our language, it's more convenient to work with rings like $\mathbb{Z}_p$ than $\mathbb{Q}$, so we ask instead for a map $\mathbb{T} \to \mathbb{Z}_p$, which (away from the level N) should be equivalent. (Really we should ask for some completion $\widehat{\mathbb{T}}$ of $\mathbb{T}$ to make everything p-local and complete, which we'll largely ignore for the moment.) Quotienting by the resulting maximal ideal gives some field over $\mathbb{F}_p$.

Given such an f and its associated elliptic curve E_f, we get the Tate module $T_p(E_f)$, which is a Galois representation over $\mathbb{Z}_p$. If $E_f[p] \simeq \overline{\rho}$, then this is the same thing as a map $\mathcal{R} \to \mathbb{Z}_p$; assume this for the moment. Then we can think of Eichler–Shimura as a machine that takes in maps $\widehat{\mathbb{T}} \to \mathbb{Z}_p$ and produces a map $\mathcal{R} \to \mathbb{Z}_p$, or similarly for any $\mathbb{Z}_p$-algebra S.

In general, a suitable (i.e. functorial in S) machine

$$\mathrm{Hom}(A, S) \to \mathrm{Hom}(B, S)$$

is actually the same thing as a ring homomorphism $B \to A$. One direction is easy: given a map $B \to A$, we can turn a map $A \to S$ into a map $B \to S$ by composition. The other direction requires some category theory and we omit it, but we're only interested in this special case: the Eichler–Shimura machine can be interpreted as a ring homomorphism

$$\mathcal{R} \to \widehat{\mathbb{T}}.$$

Let's return to our assumption that $E_f[p] \simeq \overline{\rho}$, since (for a fixed $\overline{\rho}$) this can't possibly be true for all E_f and thus for all f, so we can't hope that this map could be an isomorphism. This has to do with the completion issues we elided: if we reduce $\mathbb{T}$ modulo the maximal ideal $\mathfrak{m}_f$, we obtain a field $\mathbb{T}/\mathfrak{m}_f$ of characteristic p. The fixed $\mathbb{F}_p$-representation $\overline{\rho}$ then induces a Galois representation over this field. On the other hand, the map $\mathcal{R} \to \widehat{\mathbb{T}} \to \mathbb{T}/\mathfrak{m}_f$ should also correspond to a Galois representation over $\mathbb{T}/\mathfrak{m}_f$, so we need these to agree. This is achieved by completing $\mathbb{T}$ at the maximal ideal $\mathfrak{m}_f$, which corresponds to taking only the eigenforms f whose corresponding elliptic curves E_f satisfy $E_f[p] \simeq \overline{\rho}$. Thus we really do have a map

$$\mathcal{R} \to \widehat{\mathbb{T}},$$

and it is not unreasonable to hope it may actually be an isomorphism. Note however that this requires there to exist some eigenform f such that $E_f[p] \simeq \overline{\rho}$, i.e. that the mod p representation $\overline{\rho}$ is itself modular.

Suppose that $\overline{\rho}$ is modular and that the associated map $\mathcal{R} \to \widehat{\mathbb{T}}$ is an isomorphism. What would this mean? Well, we can go the other direction: for every deformation of $\overline{\rho}$ to a $\mathbb{Z}_p$-algebra S, we can find a corresponding map $\widehat{\mathbb{T}} \to S$ giving an eigenform over S. In particular, given an arbitrary deformation ρ of $\overline{\rho}$ to $\mathbb{Z}_p$, this corresponds to a map $\mathcal{R} \simeq \widehat{\mathbb{T}} \to \mathbb{Z}_p$ and hence to an eigenform over $\mathbb{Z}_p$; the completion map $\mathbb{T} \to \widehat{\mathbb{T}}$ induces $\mathbb{T} \to \mathbb{Z}_p$ which factors through $\mathbb{T} \to \mathbb{Z}_{(p)}$ and hence gives an eigenform f with rational coefficients such that $\rho \simeq T_p(E_f)$ via the Eichler–Shimura construction. Thus if we can show that for every (semistable) elliptic curve E over $\mathbb{Q}$, the resulting $\mathbb{F}_p$-representation $E[p]$ is modular and that the resulting map $\mathcal{R} \to \widehat{\mathbb{T}}$ is an isomorphism, then we can conclude that E must be modular as well; that is, we will have proven the modularity theorem.

The statement that for $\overline{\rho}$ modular, $\mathcal{R} \to \widehat{\mathbb{T}}$ is an isomorphism is the Wiles–Taylor modularity lifting theorem: if the reduction of the Galois representation of E modulo p is modular, then this lets us "lift" the modularity up to recover the modularity of E. This is the technical heart of their proof. We may attempt to say a few words at the end of the class about how it works, but for now we'll leave it alone. The main remaining thing we want to discuss is the "base case": how do we know if $E[p]$ is modular?

3 Mod p modularity

We want to analyze what happens to Galois representations modulo p, and since we're free to choose p we'll start by choosing a convenient one.

The smallest prime is 2, but this is often too small and strange things tend to happen at $p = 2$; so the next-simplest case is $p = 3$. This turns out to be a happy medium; part of this is that $\mathrm{GL}_2(\mathbb{F}_3)$ is a solvable group, while already $\mathrm{GL}_2(\mathbb{F}_5)$ is not solvable so many things are harder at $p = 5$.

Fix a semistable elliptic curve E over $\mathbb{Q}$. We sometimes write ρ for its Galois representation $T_p(E)$, and $\overline{\rho}$ for its modulo p representation $E[p]$. At the prime $p = 3$, if the representation $\overline{\rho}$ is irreducible, then it is known by work of Langlands–Tunnell that $\overline{\rho}$ must be modular, i.e. there must be some eigenform f with $E_f[p] \simeq \overline{\rho}$ (note we do not necessarily need to have $E_f \simeq E$! so the modularity lifting theorem is necessary). So the remaining case is when $\overline{\rho}$ is reducible.

In this case, we can't say much, so we simply try the next-simplest prime, $p = 5$. Now we put the prime in a subscript for clarity, i.e. ρ_3, ρ_5, $\overline{\rho}_3$, $\overline{\rho}_5$. We can assume that $\overline{\rho}_3$ is reducible by the above, so there are two cases: either $\overline{\rho}_5$ is reducible or it is irreducible. One can actually show that if $\overline{\rho}_3$ and $\overline{\rho}_5$ are both reducible, then the original curve E cannot be semistable, so we can focus on the case where $\overline{\rho}_3$ is reducible and $\overline{\rho}_5$ is irreducible.

We no longer have a theorem like Langlands–Tunnell at $p = 5$, so the fact that $\overline{\rho}_5$ is irreducible doesn't let us conclude just yet. However, one can do something called the "3-5 trick": one can find another elliptic curve E' such that the corresponding mod 5 representations are isomorphic, $\overline{\rho}_5 \simeq \overline{\rho}'_5$, and the mod 3 representation $\overline{\rho}'_3$ is irreducible. Then by Langlands–Tunnell, it follows that $\overline{\rho}'_3$ is modular, so by modularity lifting E' is modular, so all of its Tate modules $\rho'_p = T_p(E')$ are modular for all primes p and in particular ρ'_5 and hence $\overline{\rho}'_5$ is modular. Therefore $\overline{\rho}_5 \simeq \overline{\rho}'_5$ is also modular, and therefore the original curve E is modular.

4 Patching

The proof of the $\mathcal{R} = \mathbb{T}$ theorem is very complicated, but we'll try to say something about really just one of the key ideas involved: patching. The idea is this: for certain sets Q of "Taylor–Wiles primes," one gets a modified version of the universal deformation ring $\mathcal{R}_Q$

and Hecke algebra $\mathbb{T}_Q$ allowing ramification at Q, fitting into commutative diagrams

$$\begin{tikzcd}
\mathcal{R}_Q \arrow[r] \arrow[d] & \mathbb{T}_Q \arrow[d] \\
\mathcal{R} \arrow[r] & \mathbb{T}
\end{tikzcd}$$

where all the maps are surjective (we can think of $\mathbb{T}$ as the quotient of $\mathcal{R}$ parametrizing the modular deformations). Each of the maps $\mathcal{R}_{\mathbb{Q}} \to \mathbb{T}_{\mathbb{Q}}$ may be rather poorly behaved, as may both sides. However, what we'd like to do is take some sort of inverse limit over the different sets Q to get a map at "infinite level" $\mathcal{R}_\infty \to \mathbb{T}_\infty$, where now in the limit the bad behavior has been somehow "smoothed out." This is analogous to how although each ring $\mathbb{Z}/p^n\mathbb{Z}$ has rather bad technical properties (other than for $n = 1$), since they have lots of zero divisors, the inverse limit $\mathbb{Z}_p$ is an integral domain with very good properties. Indeed, we mentioned before that $\mathcal{R} \simeq \mathbb{Z}_p[[x_1, \ldots, x_n]]/I$ for some ideal I; similarly at level Q we have $\mathcal{R} \simeq \mathbb{Z}_p[[x_1, \ldots, x_n]]/I_Q$ for some ideal I_Q, and in the limit the ideals vanish and we just get $\mathcal{R}_\infty \simeq \mathbb{Z}_p[[x_1, \ldots, x_n]]$.

In particular, if we take the module M of modular forms (suitably completed) over the algebra of Hecke operators $\widehat{\mathbb{T}}$, the map $\mathcal{R} \to \mathbb{T}$ makes M an $\mathcal{R}$-module. If we could show that M is a finite rank free $\mathcal{R}$-module, then in particular we could find some vector $v \in M$ such that $- \cdot v : \mathcal{R} \to M$ is an embedding of $\mathcal{R}$-modules. On the other hand the $\mathcal{R}$-action factors through the surjection $\mathcal{R} \to \widehat{\mathbb{T}}$, so so does this map $\mathcal{R} \to M$; since the latter is injective, this is only possible if $\mathcal{R} \to \widehat{\mathbb{T}}$ is actually an isomorphism. However, it is too hard to show that M is free over $\mathcal{R}$. Instead, we find a version M_Q for each Q, consisting roughly of modular forms of level $NQ = N \cdot \prod_{\ell \in Q} \ell$, and again take the limit to get M_∞ over $\widehat{\mathbb{T}}_\infty$. Now the formal properties of the rings $\mathcal{R}_\infty$ and $\widehat{\mathbb{T}}_\infty$ and the map between them ensure that M_∞ will be free as an $\mathcal{R}_\infty$-module, and so the above argument implies that $\mathcal{R}_\infty \simeq \widehat{\mathbb{T}}_\infty$. Both sides are given by quotienting by the same ideal, so we can then deduce that back at the base level $\mathcal{R} \simeq \widehat{\mathbb{T}}$.

There is a problem, however: when we formed $\mathbb{Z}_p$ as the inverse limit of the $\mathbb{Z}/p^n\mathbb{Z}$, we thought of its elements as compatible systems of elements of the $\mathbb{Z}/p^n\mathbb{Z}$ along the natural maps $\mathbb{Z}/p^{n+1}\mathbb{Z} \to \mathbb{Z}/p^n\mathbb{Z}$ given by reduction modulo p^n. Here, however, we have no natural maps between the various rings and modules at level Q: there is a priori no real relationship between modular forms of level NQ and level NQ'.

Nevertheless, via the notion of patching data Taylor–Wiles showed that one can choose sequences of Q's and maps between them such that it is possible to take the inverse limits and get the desired behavior. This involves making a lot of non-canonical choices, and it is quite surprising that it seems to give "the right thing."