It is finally time for the unitary group G_0 introduced in the first lecture to make its appearance. Recall that our basic object is a cohomological automorphic representation Π of $GL(n, F)$, satisfying several axioms that guarantee its descent to both unitary groups G and G_0. We henceforward assume that Π has cohomology with trivial coefficients. We fix a level subgroup $K_f \subset G_0(\mathbf{A}_f)$ such that $\pi'^{,K_f} \neq 0$. With some additional work we can even assume $\dim \pi'^{,K_f} = 1$ in practice, but this will not be necessary. Then the L-packet π' has non-zero intersection with the space of functions

$$M'(K_f, \mathbb{C}) = C(G_0(\mathbb{Q})\backslash G_0(\mathbf{A})/G_0(\mathbb{R}) \cdot K_f, \mathbb{C}) = M'(K_f, \mathbb{Z}) \otimes \mathbb{C}$$

where $K_f \subset G_0(\mathbf{A}_f)$ is a compact open subgroup and $M'(K_f, \mathbb{Z})$ is the free $\mathbb{Z}$-module of integer-valued functions. Fixing the level K_f guarantees that the module $M'(K_f, \mathbb{Z})$ is free of finite rank. We will impose the following conditions on K_f:

(1) $K_f = \prod_v K_v$ as v ranges over finite primes of F^+.
(2) At primes that remain inert in F/F^+, K_v is a hyperspecial maximal compact subgroup.
(3) At primes that ramify in F/F^+, we could assume K_v is a "very special" maximal compact subgroup, in the terminology of Labesse. However, this condition can be ignored, because we can always assume after a quadratic base change that F/F^+ is everywhere unramified.
(4) If v splits in F/F^+, and if $v \notin S$, then $K_v \simeq GL(n, \mathcal{O}_v)$.
(5) If $v \in S$, K_v is adapted to the situation. In practice, $S = Q \cup R$ where Q are the Taylor-Wiles primes, at which

$$K_v = \{g \in GL(n, \mathcal{O}_v) \mid g \equiv \begin{pmatrix} *_{n-1,n-1} & *_{n-1,1} \\ 0_{1,n-1} & 1 \end{pmatrix} \pmod{\mathfrak{m}_v}\}.$$

and R are the level-raising primes, at which

$$K_v = \{g \in GL(n, \mathcal{O}_v) \mid g \pmod{\mathfrak{m}_v} \text{ is upper-triangular unipotent } \}$$

We will not really be able to give a precise account of the final condition. In the present version of the article we can always use base change to assume K_v contains an Iwahori subgroup of $GL(n, \mathcal{O}_v)$, but it is unnecessarily restrictive to make this hypothesis.

For v split in F/F^+, $v \notin S$, the Hecke algebra $\mathcal{H}(K_v)$ of biinvariant $\mathbb{Z}$-valued functions functions $\varphi \in C_c(K_v\backslash G_{0,v}/K_v, \mathbb{Z})$ acts on $M(K_f, \mathbb{Z})$. The double coset space $K_v\backslash G_{0,v}/K_v$ is discrete. We can define the action as follows: for $f \in M(K_f, \mathbb{Z}), \varphi \in \mathcal{H}(K_v)$, we have

$$T(\varphi)f(h) = \int_{G_{0,v}} f(hg')\varphi(g')dg'$$

where dg' is the $\mathbb{Q}$-valued bi-invariant Haar measure on $G_{0,v}$ that gives volume 1 to K_v; since the integrand is right-invariant under K_v, this is a sum of values of a $\mathbb{Z}$-valued function with $\mathbb{Z}$ coefficients.

In order to relate modular forms to deformation algebras, we work systematically with $M(K_f, \mathcal{O}) = M(K_f, \mathbb{Z}) \otimes \mathcal{O}$, where $\mathcal{O}$ is an ℓ-adic integer ring, as in the earlier lectures, to be specified momentarily. We let $T(K_f)$ denote the subalgebra of $M(K_f, \mathcal{O})$ generated by the $T(\varphi)$, for $\varphi \in \mathcal{H}(K_v)$, v but not in S. This is a finite free commutative $\mathcal{O}$-subalgebra of $End(M(K_f, \mathcal{O}))$. The commutativity is a consequence of the commutativity of $\mathcal{H}(K_v)$ for each v, and this is part of the theory of the Satake transform that we have already seen. Since $T(K_f)$ is finite over $\mathcal{O}$, it is a semi-local ring, and we let $\mathfrak{m}_i \subset T(K_f)$ denote its maximal ideals, as i runs over a finite index set.

$M(K_f, \mathbb{C})[\pi'] = \pi'(K_f) \cap M(K_f, \mathbb{C}) \neq 0$. We assume that $\mathcal{O}$ is chosen so that $\pi'(K_f) \cap M(K_f, \mathcal{O}) \neq 0$, i.e., that $\pi'(K_f)$ contains functions with values in $\mathcal{O}$. This can be done first with $\mathcal{O}$ replaced by the integers in a finite extension of $\mathbb{Q}$, and then $\mathcal{O}$ can be defined as an ℓ-adic completion of this integer ring. Now $M(K_f, \mathcal{O})[\pi']$, with the obvious definition, is a $T(K_f)$-invariant submodule, precisely because $(\pi')^{K_f}$ is an eigenspace for all the Hecke operators, and indeed, the action of $T(K_f)$ on $(\pi')^{K_f}$ is given by a character $T(K_f) \to \mathcal{O}$ taking each $T(\varphi)$ to its eigenvalue on $(\pi')^{K_f}$. Since $\mathcal{O}$ is an integral domain, there is exactly one $\mathfrak{m}_i$, denote $\mathfrak{m}(\Pi)$, such that

$$M(K_f, \mathcal{O})[\pi'] \otimes_{T(K_f)} T(K_f)_{\mathfrak{m}(\Pi)} \neq 0.$$

Now we write

$$\mathbb{T}_\Pi = T(K_f)_{\mathfrak{m}(\Pi)}, \qquad M_\Pi = M(K_f, \mathcal{O}) \otimes_{T(K_f)} \mathbb{T}_\Pi.$$

By definition, M_Π is a $\mathbb{T}_\Pi$-module.

Lemma. *The algebra $\mathbb{T}_\Pi$ is reduced.*

This follows easily from the fact that $M(K_f, \mathbb{C})$ is a semisimple $T(K_f)$-module, which in turn comes from the fact that $T(K_f)$ is hermitian with respect to the L_2-inner product (Petersson norm).

The point of this construction is that $\mathcal{H}(K_v)$ is the same local Hecke algebra at v that acts on the cohomological automorphic forms on $GL(n, F)$ and on the middle-dimensional cohomology $H^{n-1}(Sh(G), \mathbb{Q}_\ell)$, where the avatar of Π is the L-packet we've denoted $\{\pi\}$. We let $\rho = \rho_{\ell,\pi}$ – which we can also write ρ_Π – be the n-dimensional Galois representation, and let $r = r_(\Pi) : G_{F^+} \to \mathcal{G}_n(\mathcal{O})$ be the corresponding homomorphism. We always assume

$$\bar{\rho}_\Pi isabsolutelyirreducible.$$

Our initial goal is to prove the existence of an isomorphism

$$R_{\bar{r}_\Pi, \emptyset} \xrightarrow{\sim} \mathbb{T}_\Pi.$$

As I mentioned at the end of the last lecture, this isomorphism implies that every lifting of $\bar{r}_\Pi$ in the category $Def_{\mathcal{D}}$ comes from some automorphic representation, not of G nor of $GL(n,F)$, but of the definite unitary group G_0. But then by base change we can deduce that every lifting of $\bar{r}_\Pi$ in fact comes from a cohomological automorphic representation of $GL(n,F)$.

Before we can construct the isomorphism $R = T$, as this is usually abbreviated, we need to construct a homomorphism. Now $\mathbb{T}_\Pi$ has been constructed to have the residue field k of $\mathcal{O}$, and therefore belongs to the category $\mathcal{C}_{\mathcal{O}}$. Thus a homomorphism from the universal deformation algebra R to T is equivalent to the existence of a lifting r_T from $\bar{r}_\Pi$ to $\mathcal{G}_n(T)$. We in fact construct a representation with values in $GL(n,T(K_f))$. For any irreducible representation τ of G_0 with $\tau^{K_f} subset M(K_f,\overline{\mathbb{Q}}_\ell)$, let $p_\tau : T(K_f) \to \overline{\mathbb{Q}}_\ell$ be the representation on τ^{K_f}.

Proposition. *(a) There is a representation*

$$\rho_{K_f} : G_F \to GL(n,T(K_f))$$

with the property that, for any irreducible representation τ of G_0 as above, admitting base change to a cuspidal cohomological automorphic representation $B(\tau)$ of $GL(n,F)$, the projection

$$\rho_\tau = p_\tau \circ \rho_{K_f} : \rho_{K_f} \to GL(n,\overline{\mathbb{Q}}_\ell)$$

is an l-adic representation of geometric type, unramified outside $S \cup S_\ell)$, with

$$L^{S\cup S_\ell)}(s,\rho_\tau) = L^{S\cup S_\ell)}(s,B(\tau)).$$

(b) Let $E(\Pi)$ be the set of τ as in (a) such that (loosely speaking) p_τ factors through $\mathbb{T}_\Pi \cap T(K_f)$. The projection

$$\rho_\Pi = \oplus_{\tau\in E(\Pi)}\rho_\tau : G_F \to GL(n,\oplus_{\tau\in E(\Pi)}\overline{\mathbb{Q}}_\ell)$$

is conjugate to a homomorphism $\rho_{\mathbb{T}_\Pi}$ with values in the subgroup $GL(n,\mathbb{T}_\Pi) \subset GL(n,\oplus_{\tau\in E(\Pi)}\overline{\mathbb{Q}}_\ell)$. The reduction of ρ_Π modulo the maximal ideal of $\mathbb{T}_\Pi$ is equivalent to $\bar{\rho}_\Pi$.

(c) The homomorphism $\rho_{\mathbb{T}_\Pi}$ extends to a homomorphism

$$r_{\mathbb{T}_\Pi} : G_{F^+} \to \mathcal{G}_n(T)$$

whose reduction modulo the maximal ideal is equivalent to $\bar{r}_\Pi$.

Assertion (a) is another way of expressing the theorem of Kottwitz and Clozel described in the second lecture. Assertion (b) is a consequence of a very useful theorem of Carayol, and depends on the hypothesis that $\bar{r}_\Pi$ be absolutely irreducible. The last claim in (b) is the assertion that all the Galois representations associated to $\tau \in E(\Pi)$ have the same reduction modulo ℓ. This is not difficult to show – it suffices to calculate traces of Frobenius modulo ℓ, and the point is that these are determined by the eigenvalues of the Hecke operators modulo ℓ, and completion at $\mathfrak{m}_\Pi$ simply picks out all the τ whose Hecke eigenvalues are congruent modulo ℓ. But it deserves to be emphasized, since this is precisely the basis of the theory of congruences of modular forms, as developed by Serre, Mazur, Katz, and Ribet in the 1970s, and extended in various directions by Hida, Wiles, Taylor, and now many others. Assertion (c) makes use of the Petersson pairing.

This is also the place to remark that the set of v that split in F/F^+, with a finite subset removed, suffice to determine any ℓ-adic representation of G_F, because their extensions to F have Dirichlet density 1. Practically every step of the argument makes reference to Chebotarev's density theorem, which in the present setting allows us to determine an ℓ-adic representation up to equivalence by the traces of Frobenius at primes v belonging to a set of Dirichlet density 1. This also suffices for the surjectivity assertion in the following Corollary:

Corollary. *There is a surjective homomorphism in $\mathcal{C}_\mathcal{O}$:*

$$\phi_\emptyset : R_{\bar{r}_\Pi,\emptyset} \to \mathbb{T}_\Pi.$$

To prove surjectivity we need to know that all the $T(\varphi)$ are in the image of the homomorphism for a set of generators $\varphi \in \mathcal{H}(K_v)$ for v split and not in S. Now the structure theory of $\mathcal{H}(K_v)$ gives a set of generators $T_{i,v}, i = 0, \dots n$, whose eigenvalues on τ^{K_v}, $\tau \in E(\Pi)$, are precisely the coefficients of the Euler factor at v of $L(s,\tau)$. Since τ is associated to some ρ_τ, these are the coefficients of the Euler factor at v of $L(s,\rho_\tau)$. This means that it suffices to show that the coefficients of the characteristic polynomial of $Frob_v$, acting on the universal lifting of $\bar{r}_\Pi$, lie in $R_{\bar{r}_\Pi,\emptyset}$, and this is a standard fact in representation theory.

To prove this surjection is an isomorphism, one applies the method of Taylor-Wiles, as simplified by Diamond and Fujiwara. This is where the sets of primes Q_N come in. For each N, one defines an algebra $\mathbb{T}_{\Pi,Q_N}$ acting on a module M_{Π,Q_N} of $\mathcal{O}$-valued modular forms, and admitting a surjective homomorphism

$$\phi_{Q_N} : R_{\bar{r}_\Pi,Q_N} \to \mathbb{T}_{\Pi,Q_N}.$$

The M_{Π,Q_N} is defined by analogy with M_Π, but for $v \in Q_N$ we replace K_v by $K_v^1 \subset GL(n,\mathcal{O}_v)$, defined to be the subgroup of matrices with final row $(0,0,\dots,0,1)$. For each $v \in Q_N$, there is a subgroup $D_v \subset R^\times_{\bar{r}\Pi,Q_N}$ isomorphic to the ℓ-torsion subgroup of $(\mathbb{Z}/Q_N\mathbb{Z})^\times$, which acts on M_{Π,Q_N} as the corresponding subgroup of K_v^0/K_v^1, where K_v^0 is defined like K_v^1 but with bottom row $(0,\dots,0,*)$. Let $D_{Q_N} = \prod_{v\in Q_N} D_v$.

These data satisfy the following relations, each of which conceals a hypothesis on our choice of Q_N that is allowed by the condition that the image of $\bar{\rho}$ be "big."

(1) The module M_{Π,Q_N} is a free $\mathbb{Z}_\ell[D_{Q_N}]$-module. (This is the crucial condition).

(2) Let $I_{Q_N} \subset \mathbb{Z}_\ell[D_{Q_N}]$ be the augmentation ideal. Then

$$M_{\Pi,Q_N}/I_{Q_N}M_{\Pi,Q_N} \xrightarrow{\sim} M_{\Pi,\emptyset}$$

as $\mathbb{T}_{\Pi,Q_N}$-modules.

The very rough idea is that a pigeonhole principle argument allows us to extract from the sequence of Q_N a subsequence such that the algebras R_{Π,Q_N} and $\mathbb{T}_{\Pi,Q_N}$, together with the modules M_{Π,Q_N}, all patch together, and in the end one has power series rings. The relative dimension over $\mathcal{O}$ of the power series ring $R_{\Pi,\infty}$ is given by the integer $r - n\sum_{v|\infty}\frac{c_v+1}{2}$ from the last lecture. On the other hand, the $\mathbb{Z}_\ell[D_{Q_N}]$ can be seen as truncated power series rings in r variables, and the freeness of M_{Π,Q_N} over $\mathbb{Z}_\ell[D_{Q_N}]$ for all N implies that the limit algebra $\mathbb{T}_\infty$ is a module over a power series ring in r variables, contained in the image of R_∞. Since the map $R_\infty \to \mathbb{T}_\infty$ is surjective, this implies that all the $c_v = -1$ and that the map is an isomorphism, and one deduces by reducing modulo augmentation ideals that the original map $R_\emptyset \to \mathbb{T}_\emptyset$ was also an isomorphism.

The original objective of my project with Taylor, begun in 1996, was to extend the theorem of Taylor-Wiles to automorphic forms in higher dimension, where the Galois representations are of dimension $n > 2$ in general. I remind you that the Taylor-Wiles theorem, which has been generalized in various directions by Diamond and Fujiwara, and more recently by Genestier-Tilouine for symplectic groups, asserts that, in favorable circumstances, if a mod ℓ representation $\bar{\rho}$ of $Gal(\bar{\mathbb{Q}}/E)$ lifts to an ℓ-adic representation which comes from automorphic forms on $GL(n)$ of a certain type, where E is totally real or CM field, then any lifting with minimal additional ramification also comes from automorphic forms. By "favorable circumstances" I mean, for example, that $\ell > n$, so we can apply Fontaine-Laffaille (nearly ordinary would also work) and $\bar{\rho}$ should be absolutely irreducible and with image not too small. The work of Wiles on $GL(2)$, and later Diamone, Skinner-Wiles and Breuil-Conrad-Diamond-Taylor, showed that the lifting theorem remained true under more general ramification conditions. This is the problem of "level raising" that involves a different range of techniques.

Let me remind you of the Fontaine-Mazur version of Langlands' conjectures. Let F (resp. E) be a number field, which we will assume totally real (resp. CM). For the moment we stick with F. Let S_ℓ denote the set of primes of F dividing ℓ. Following Fontaine and Mazur, we say an n-dimensional ℓ-adic representation ρ of $Gal(\bar{F}/F)$ is *of geometric type* if it is unramified outside the finite set $S\coprod S_\ell$ of primes of F, where $S \cap S_\ell = \emptyset$; and if at every $v \in S_\ell$ it has Fontaine's de Rham property, which means in particular that it allows us to associate a set of Hodge-Tate numbers $h^p(\rho)$ to ρ, with $n = \sum h^p(\rho)$ (varying with v in general).

Following Clozel, we define an automorphic representation π of $GL(n,\mathbb{Q})$ if π_∞ has integral infinitesimal character – which means it can be associated to a Hodge structure.

Conjecture. *(a) Let ρ be an irreducible n-dimensional ℓ-adic representation of G_F of geometric type. Then there is a cuspidal automorphic representation of $GL(n,\mathbb{Q})$ π_ρ of algebraic type associated to ρ, in the sense that $L(s,\pi_\rho) = L(s,\rho)$ where the former is the L-function associated by automorphic theory. In particular, $L(s,\rho)$ has an analytic continuation to an entire function satisfying the usual sort of functional equation.*

(b) Conversely, if π is an automorphic representation of $GL(n,\mathbb{Q})$ of algebraic type,

then there exists an ℓ-adic representation ρ_π of geometric type associated to π.

A theorem of Taylor-Wiles type is the following:

Prototype of generalized Taylor-Wiles theorem. *Suppose ρ is as in (a), and suppose π_ρ exists. Let ρ' be a second n-dimensional $\mathfrak{p}$-adic representation, with $\rho' \equiv \rho \pmod{\mathfrak{p}}$. Suppose moreover that*

(i) *ρ' satisfies a minimality condition, typically that ρ' is no more ramified than $\bar\rho := \rho \pmod{\ell}$ at primes in S or otherwise, apart from primes in S_ℓ.*
(ii) *Some more precise condition on the restriction of ρ' to $Gal(K_{\rho,q}/\mathbb{Q}_q)$, for $q \in S - S_\ell$ e.g. that ρ' and ρ have isomorphic restrictions to $Gal(K_{\rho',q}/\mathbb{Q}_q)$.*
(iii) *A specific Fontaine-type condition on the restriction of ρ' to $Gal(K_{\rho',v}/F_v)$ for $v \in S_\ell$; in practice, ρ and ρ' are assumed crystalline and have the same Hodge-Tate numbers).*
(iv) *Additional conditions, e.g.*
(a) *$\rho \pmod{\ell}$ is irreducible, even after restriction to $F(\zeta_\ell)^+$ and has big image,*
(b) *$\ell > n$ and is unramified in F.*

Then ρ' also satisfies the Fontaine-Mazur conjecture (a).

The original Taylor-Wiles theorem applies to $n = 2$ and for specific conditions (ii-iv). The generalization due to Wiles, and Taylor and his collaborators more generally,in the case $n = 2$ removes condition (i). I note that we always assume that all primes in S split completely in E/F.

Theorem A. (MH-Taylor). *(F totally real.) Let K be a finite extension of $\mathbb{Q}_\ell$, with residue field k. Suppose V_ℓ is an n-dimensional K-vector space and $\rho : Gal(\overline{\mathbb{Q}}/F) \to GL(n, K)$ is as in (a), and suppose π_ρ exists and is cohomological at infinity (this will be automatic). Suppose moreover that*

(1) Polarized of weight $n-1$: there is a Galois-invariant bilinear form

$$V_\ell \otimes V_\ell \to \mathbb{Q}(1-n)$$

, of parity $(-1)^{n-1}$;

(2) Regular: $h^p(\rho) \leq 1$ for all (p), and $h^p = h^{n-1-p}$ (in practice ρ will be pure of weight $n-1$ by construction)

(3) At some finite place v_0 of F π_ρ is either supercuspidal or Steinberg. In the supercuspidal case this implies that $\rho\mid_{G_{v_0}}$ is irreducible, and we assume this is still true of $\bar\rho$. In the Steinberg case there is a way around this.

Suppose ρ' is as above and satisfies conditions (i)-(iv) and is polarized ((2) and (3) are automatic). Then $\pi_{\rho'}$ also exists.

For E CM rather than totally real, we replace condition (1) by the hypothesis that $V_\ell^\vee \xrightarrow{\sim} V_\ell^c(1-n)$. The weight condition can be relaxed somewhat but I prefer not to go into it. In fact, we only consider CM fields of the form $E = F \cdot \mathcal{K}$ where $\mathcal{K}$ is imaginary quadratic, though it should be possible to treat more general cases by a descent argument.

The theorem is not proved by working with automorphic forms on $GL(n)$ but rather with definite unitary groups attached to division algebras over E ($= F \cdot \mathcal{K}$ if necessary). These are like the groups that occur in our book, except they are positive-definite everywhere, which allows us to construct Taylor-Wiles systems without difficulty, because the Hecke algebra modules in question are in 0-dimensional cohomology. Conditions (1) and (2) are unavoidable for cohomological representations. Indeed, there are reasons to believe representations of this type are generic, in the sense that they fit into the largest ℓ-adic analytic families (in the sense of Hida, Coleman, Mazur, etc.) with dense subsets of representations of geometric type. On the other hand, (3) is a symptom of our dependence on the current state of the stable trace formula, and is likely to be unnecessary in the near future. This turns out to be unfortunate, for reasons I can't explain. Note that (1) and (2) are conditions on ρ' whereas (3) is a condition on $\pi_{\rho'}$, but by the results of my book with Taylor (3) is in fact a condition on ρ'.

From now on I will work with E CM rather than F.

Theorem B. *Suppose ρ and π_ρ are as in the previous theorem and satisfy (1)-(3). Suppose ρ' is as in the Fontaine-Mazur conjecture, is polarized, and satisfies conditions (ii)-(iv) (i) is optional). Then $\pi_{\rho'}$ also exists.*

Here is an unenlightening definition of "big image". Let $\mathcal{O} = \mathcal{O}_K$. The polarization condition implies that ρ extends to a homomorphism

$$\tilde{\rho}: Gal(\overline{\mathbb{Q}}/F) \to \mathcal{G}_n(\mathcal{O}) := [GL(n,\mathcal{O}) \times GL(1,\mathcal{O})] \ltimes \{1, c\}$$

where $c(g,\mu) = (\mu^t g^{-1}, \mu)$ (more than one extension is in principle possible...) Let $\bar{\tilde{\rho}} := \tilde{\rho} \pmod{\ell}$, $H = Im(\bar{\tilde{\rho}}) \cap Gal(\overline{\mathbb{Q}}/\mathbb{Q}(\zeta_\ell)^+)$. We say H is "big" if

(i) $H^0(H, ad(\bar{\rho})) = H^1(H, ad(\bar{\rho})) = 0$;

(ii) For any irreducible submodule $W \subset \mathbf{A}(\bar{\rho})$ there exists $h \in H \cap GL(n,k)$ and $\alpha \in k$ such that the α-generalized eigenspace $V_{h,\alpha}$ of h is of dimension 1 and $\pi_{h,\alpha} \circ W \circ i_{h,\alpha} \neq 0$ where $\pi_{h,\alpha}$ and $i_{h,\alpha}$ are respectively the h-equivariant projection on $V_{h,\alpha}$ and the inclusion.

Surprisingly, this condition is not very restrictive. Of course if $H \subset GL(n,k)$ this is automatic, but it is also true when the image of H is relatively small.

The proof of Theorem B follows the pattern of the proof of Wiles' theorem and those that have followed. First one proves Theorem A by constructing appropriate Taylor-Wiles systems. One then relaxes the minimality hypothesis. Our approach uses a trick invented by Skinner-Wiles (when $n = 2$) to avoid the difficult level-lowering argument of Ribet.

More precisely, we let $R_\emptyset = R_\emptyset(\bar{\rho}, F)$ denote the deformation ring associated to the deformation problem implicit in conditions (i), (ii), (iii) (and the polarization condition (1)). In other words, we follow Mazur, Wiles, Ramakrishna, etc. and define a functor from Artinian $\mathcal{O}$-algebras A with residue field k to sets, classifying equivalence classes of liftings of $\bar{\rho}$ to maps $Gal(\overline{\mathbb{Q}}/F) \to \mathcal{G}_n(\mathcal{O})$ satisfying conditions (i), (ii), (iii). (For (iii) we use the Fontaine-Laffaille condition, which is one reason we need $\ell > n$ and unramified in F. At places $q \in S$ we introduce a new minimality condition when $\pi_{\rho,q}$ is Steinberg.). One proves this functor is representable, hence we have $R_\emptyset$.

We also need a Hecke algebra $T_\emptyset(\bar{\rho}, F) = T_\emptyset$ and a $\mathcal{O}$-module of automorphic forms $M_\emptyset$ on which $T_\emptyset$ acts. The module $M_\emptyset$ is contained in a module $\mathbb{M}$ of $\mathcal{O}$-valued modular forms on the definite unitary group $U(B)$ of the central division algebra B/E of dimension n^2 with involution of the second kind (i.e., restricting to c on the center E). Since this is a definite unitary group over the field F we can just take the module of functions $U(B)(F)\backslash U(B)(\mathbf{A}_f)/C$. Here $C = C_{\emptyset,S}$ is a compact open subgroup of $U(B)(\mathbf{A}_f)$ defined appropriately for ρ and S; it is hyperspecial locally at primes not in S (including S_ℓ). The theory of base change, plus the theory developed by Clozel and Kottwitz, and in our book, shows that automorphic forms of this type give rise to n-dimensional ℓ-adic representations of $Gal(\overline{\mathbb{Q}}/E)$ satisfying (1) and (2), with Hodge-Tate weights $0, 1, \ldots, n-1$, each with multiplicity one; and usually satisfying (3) as well (this is an additional condition we impose). These Galois representations arise in the cohomology of a Shimura variety attached to a different unitary group, but this is a story that doesn't concern us here. One can also obtain different regular polarized representations, but this requires additional notation.

At places of F that split in E we have $U(B)(F_w)$ is either $GL(n, F_w)$ or, for a finite set, the multiplicative group of a division algebra over F_w (an extra technical hypothesis). The Hecke algebra $\mathbb{T}$ is defined as the $\mathcal{O}$-subalgebra of $End(\mathbb{M})$ generated by the usual Hecke operators for $GL(n, F_w)$ at split places. The representation $\bar{\rho}$ corresponds to a maximal ideal $\mathfrak{m} \subset \mathbb{T}$, and we let $M_\emptyset$ and $T_\emptyset$ denote the localizations at $\mathfrak{m}$.

As in Wiles' theory, there is a map $R_\emptyset \to T_\emptyset$ which is automatically surjective. In particular, $M_\emptyset$ becomes naturally an $R_\emptyset$-module. The subscript $\emptyset$ refers to the fact that the liftings are minimal. The Taylor-Wiles method, as improved by Diamond and Fujiwara, replaces $R_\emptyset$ by certain deformation rings R_Q in which the minimality condition is slightly relaxed, and likewise replaces $C_{\emptyset,S}$ by slightly smaller compact open subgroups, and hence replaces $M_\emptyset$ by the slightly bigger M_Q. The remarkable fact is that the action of certain finite ℓ-groups guarantees that the M_Q grow rapidly, whereas a local Euler characteristic calculation in Galois cohomology guarantees that the R_Q grow no more rapidly than the M_Q, and using standard results in commutative algebra Diamond and Fujiwara conclude that $R_\emptyset \xrightarrow{\sim} T_\emptyset$ is an isomorphism of complete intersections and $M_\emptyset$ is a free $R_\emptyset$-module.

To obtain Theorem B from Theorem A Wiles' original approach was to eliminate condition (i), specifically to allow unrestricted ramification at primes in $S \cup \mathcal{R}$, where $\mathcal{R}$ is now the set of places where ρ' is ramified (not including primes dividing ℓ) together with the set of places where the ramification in π_ρ is not accounted for by $\bar{\rho}$. (This is not exactly right, but I can't go into this.) This seems too hard in general. However, by replacing E/F by E'/F' where F'/F is an appropriate solvable extension, and $E' = E \cdot F'$, one can arrange that π_ρ has the following property: at any v in $\mathcal{R}$, $\pi_{\rho,v}$ has a vector fixed by the Iwahori subgroup of $GL(n, F'_v)$. In other words, $\pi_{\rho,v}$ is the (unique) generic subquotient of some tamely ramified principal series representation. Moreover, we can also assume that all primes $v \in \mathcal{R}$ satisfy $q_v \equiv 1 \pmod{\ell}$. Our generalization of the Skinner-Wiles trick is the following lemma:

Lemma. *Over E', we may replace π_ρ by π_{ρ_1} with $\bar{\rho}_1 = \bar{\rho}$ but with $\rho_1(I_v)$ a finite group for all $v \in \mathcal{R}$.*

In other words, the local monodromy logarithms at $v \in \mathcal{R}$ vanish for ρ_1; this certainly need not have been true before. Enlarging F' even more, we may then assume that ρ_1 is unramified principal series at all $v \in \mathcal{R}$. Now we can define a deformation ring $R_{\mathcal{R}}$ – for representations of $Gal(\overline{\mathbb{Q}}/F')$ with residual representation equal to $\bar{\rho}$ – which allows unrestricted ramification at $v \in \mathcal{R}$. This is now really a case of level-raising, because we know that there is at least one lifting, namely ρ_1, which is unramified at $v \in \mathcal{R}$. The difficulty is now to define $M_{\mathcal{R}}$ (and therefore $T_{\mathcal{R}}$). We begin by taking $M^*_{\mathcal{R}}$ to be the module of all modular forms defined as before but with the local factor C_v of C at $v \in \mathcal{R}$ equal to $U_1(v^n)$ (define). There are operators $U_{i,v}$, $i = 1, \ldots, n$, generalizing the classical U-operators at v, and we let $M_{\mathcal{R}} = M^*_{\mathcal{R},\mathfrak{m},\mathfrak{n}}$ where

$$\mathfrak{n} = (U_{i,v}, v \in \mathcal{R}, 1 \leq i \leq n-1).$$