

WEIGHTS OF GALOIS REPRESENTATIONS ASSOCIATED TO HILBERT MODULAR FORMS

MICHAEL M. SCHEIN

ABSTRACT. Let F be a totally real field, $p \geq 3$ a rational prime unramified in F , and \mathfrak{p} a place of F over p . Let $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ be a two-dimensional mod p Galois representation which is assumed to be modular of some weight and whose restriction to a decomposition subgroup at \mathfrak{p} is irreducible. We specify a set of weights, determined by the restriction of ρ to inertia at \mathfrak{p} , which contains all the modular weights for ρ . This proves part of a conjecture of Diamond, Buzzard, and Jarvis, which provides an analogue of Serre’s epsilon conjecture for Hilbert modular forms mod p .

1. INTRODUCTION

Let F be a totally real field. To a Hilbert modular eigenform on F one can attach a mod p Galois representation $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ (see [Car2] and [Tay]). For every embedding $\iota : F \hookrightarrow \mathbb{R}$, there is a corresponding complex conjugation $c_\iota \in \text{Gal}(\overline{F}/F)$, and we call $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ totally odd if $\det(\rho(c_\iota)) = -1$ for all ι . There is a natural analogue of Serre’s conjecture:

Conjecture 1. Let $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ be a continuous, irreducible, totally odd representation. Then ρ is modular.

Diamond, Buzzard, and Jarvis ([BDJ]) have formulated an analogue of the strong Serre conjecture if p is unramified in F . This paper proves some cases of the “epsilon conjecture” in this context. We note that the author in [Sch1] has extended the conjecture to the case where p ramifies, and indicates there how to modify the argument of this paper to prove an analogous theorem in that case. We use the notion of modularity of Definition 1.2, which also views weights from a somewhat different perspective than the classical one of $[F : \mathbb{Q}]$ -tuples of integers.

Throughout this paper, we assume that $p \geq 3$ is unramified in F , and fix a prime ideal \mathfrak{p} of F dividing p . For any place v of F , we write \mathcal{O}_v for the completion of \mathcal{O}_F at v and k_v for the residue field. Let the cardinality of $k_{\mathfrak{p}}$ be $q = p^s$. Let $G_{\mathfrak{p}} \subset \text{Gal}(\overline{F}/F)$ be a decomposition subgroup for \mathfrak{p} , and $I_{\mathfrak{p}} \subset G_{\mathfrak{p}}$ the inertia subgroup.

Definition 1.1. A *weight* is an irreducible $\overline{\mathbb{F}}_p$ -representation of $\text{GL}_2(\mathcal{O}_F/\mathfrak{p}) = \prod_{v|p} \text{GL}_2(k_v)$.

By Prop. 1 of [BL], any weight has the form $\sigma = \bigotimes_{v|p} \sigma_v$, for

$$\sigma_v = \bigotimes_{\tau: k_v \hookrightarrow \overline{\mathbb{F}}_p} (\det^{w_\tau} \text{Sym}^{k_\tau - 2} k_v^2) \otimes_{\tau} \overline{\mathbb{F}}_p,$$

where $2 \leq k_\tau \leq p + 1$ and $0 \leq w_\tau \leq p - 1$, and the w_τ are not all $p - 1$.

Key words and phrases. Galois representations, Hilbert modular forms, Serre’s conjecture, Shimura curves. The author was supported by the NDSEG and NSF Graduate Research Fellowships.

Let B be a quaternion algebra over F . We shall always assume that B is split at exactly one real place of F and at all places over p . In other words, $B \otimes \mathbb{R} \simeq M_2(\mathbb{R}) \times \mathbb{H}^{[F:\mathbb{Q}]-1}$ and $B \otimes \mathbb{Q}_p \simeq M_2(F \otimes \mathbb{Q}_p)$. Consider the reductive group $G = \text{Res}_{F/\mathbb{Q}}(B^*)$ over \mathbb{Q} . If $F \neq \mathbb{Q}$, then for an open compact subgroup $U \subset G(\mathbb{A}^\infty)$ we can define a compact Riemann surface

$$(1) \quad M_U(\mathbb{C}) = G(\mathbb{Q}) \backslash G(\mathbb{A}^\infty) \times (\mathbb{C} - \mathbb{R}) / U.$$

This Riemann surface has a canonical model M_U over F . As U varies, the Shimura curves M_U/F form an inverse system of smooth complete algebraic curves. They are not in general geometrically connected. As in [BDJ], we let $\text{Pic}^0(M_U)$ denote the identity component of the relative Picard scheme of M_U/F , which parametrizes line bundles locally of degree zero. Then $\text{Pic}^0(M_U)$ is an abelian variety over F .

We say that an open compact $U \subset G(\mathbb{A}^\infty)$ is of type $(*)$ if $U = U_p \times U^p$, where $U_p = \ker((B \otimes \hat{\mathbb{Z}}_p^*) \rightarrow \text{GL}_2(\mathcal{O}_v) \rightarrow \text{GL}_2(\mathcal{O}_F/p))$ and $U^p \subset G(\mathbb{A}^{\infty,p})$. Let $V = \prod_{v|p} \text{GL}_2(\mathcal{O}_v) \times U^p$. If U^p is sufficiently small as in section 3.1, then M_U/M_V is a Galois cover with group $V/U = \text{GL}_2(\mathcal{O}_F/p)$. This gives an action of V/U on $\text{Pic}^0(M_U)$.

Definition 1.2. We say that a Galois representation $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$ is *modular of weight* σ if there exists a quaternion algebra B over F , split at one real place and all places over p , such that ρ is a Jordan-Hölder constituent of $(\text{Pic}^0(M_U(G))[p] \otimes_{\bar{\mathbb{F}}_p} \sigma)^{\text{GL}_2(\mathcal{O}_F/p)}$ for an open compact $U \subset (D \otimes \hat{\mathbb{Z}})^* \subset G(\mathbb{A}^\infty)$ of type $(*)$.

Given a continuous, irreducible, totally odd $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$, let $W(\rho)$ be the set of modular weights. Diamond, Buzzard, and Jarvis construct sets $W_v^?(\rho)$ for each $v|p$, such that $W_v^?(\rho)$ depends only on $\rho|_{I_v}$, and conjecture that $W(\rho) = \left\{ \sigma = \bigotimes_{v|p} \sigma_v : \sigma_v \in W_v^?(\rho) \right\}$.

If $\rho_{\mathfrak{p}} : G_{\mathfrak{p}} \rightarrow \text{GL}(V)$ is any continuous two-dimensional $\bar{\mathbb{F}}_p$ -representation, then the wild inertia $I'_{\mathfrak{p}}$ acts trivially on the semisimplification V^{ss} of V ; see [Edi2] for a proof. Therefore, $\rho^{ss}|_{I_{\mathfrak{p}}}$ factors through the tame inertia $I_{t,\mathfrak{p}} = I_{\mathfrak{p}}/I'_{\mathfrak{p}} \simeq \varprojlim \mathbb{F}_{p^r}^*$, and the action of $I_{t,\mathfrak{p}}$ is given by two characters $\phi, \phi' : I_{t,\mathfrak{p}} \rightarrow \bar{\mathbb{F}}_p^*$. Recall that a character $\mu : I_{t,\mathfrak{p}} \rightarrow \bar{\mathbb{F}}_p^*$ is said to be of level r if r is the smallest integer such that μ factors through $\mathbb{F}_{p^r}^*$. If μ is moreover induced by an embedding of fields $\mathbb{F}_{p^r} \hookrightarrow \bar{\mathbb{F}}_p$, then it is a fundamental character of level r ; there are r of these. If $\rho|_{G_{\mathfrak{p}}}$ is irreducible, then we see in Lemma 4.4 that ϕ and ϕ' have level $2s$ and $\phi' = \phi^q$.

In this case, $W_{\mathfrak{p}}^?(\rho)$ is defined as follows. Let $k'_{\mathfrak{p}}$ be a quadratic extension of $k_{\mathfrak{p}}$. Then each embedding $\tau : k_{\mathfrak{p}} \hookrightarrow \bar{\mathbb{F}}_p$ has two liftings to embeddings $\tau' : k'_{\mathfrak{p}} \hookrightarrow \bar{\mathbb{F}}_p$. For such an embedding τ (resp. τ') let $\psi_{\tau} : I_{t,\mathfrak{p}} \rightarrow \bar{\mathbb{F}}_p^*$ (resp. $\psi_{\tau'}$) be the corresponding fundamental character of level s (resp. level $2s$).

Then [BDJ] defines $\sigma_{\mathfrak{p}} = \bigotimes_{\tau} (\det^{w_{\tau}} \text{Sym}^{k_{\tau}-2} k_{\mathfrak{p}}^2) \otimes_{\tau} \bar{\mathbb{F}}_p \in W_{\mathfrak{p}}^?(\rho)$ if and only if for every τ it is possible to label one of the two liftings $\tilde{\tau}$ and the other one $\tilde{\tau}'$ in such a way that

$$\rho|_{I_{\mathfrak{p}}}^{ss} \sim \prod_{\tau} \psi_{\tau}^{w_{\tau}} \begin{pmatrix} \prod_{\tilde{\tau}} \psi_{\tilde{\tau}}^{k_{\tau}-1} & 0 \\ 0 & \prod_{\tilde{\tau}'} \psi_{\tilde{\tau}'}^{k_{\tau}-1} \end{pmatrix}.$$

The main theorem of this paper is essentially a proof of one direction of the conjecture of [BDJ] in this case. Our result is slightly weaker than the conjecture. To state it, we introduce a set of weights $W_{\mathfrak{p}}(\rho) \supset W_{\mathfrak{p}}^?(\rho)$, and the statement is:

Theorem 1.3. *Assume that p is unramified in F , and let $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$ be such that $\rho|_{G_p}$ is irreducible and ρ is modular for a weight of the form $\sigma_p \otimes (\otimes_{\beta \neq p} \sigma_\beta)$. Then $\sigma_p \in W_p(\rho)$.*

We define $W_p(\rho)$ as follows. Let $\tau_1, \tau_2, \dots, \tau_s$ be a labeling of the embeddings $k_p \hookrightarrow \bar{\mathbb{F}}_p$ such that $\tau_i = \tau_{i+1}^p$ for all $i \in \mathbb{Z}/s\mathbb{Z}$. Similarly let $\sigma_1, \dots, \sigma_{2s}$ be a labeling of the embeddings $k'_p \hookrightarrow \bar{\mathbb{F}}_p$ such that $\sigma_j = \sigma_{j+1}^p$ for all $j \in \mathbb{Z}/2s\mathbb{Z}$ and such that the restriction of σ_j to k_p is $\tau_{\pi(j)}$, where $\pi : \mathbb{Z}/2s\mathbb{Z} \rightarrow \mathbb{Z}/s\mathbb{Z}$ is the natural projection. Write ψ_j for ψ_{σ_j} and k_j for $k_{\tau_{\pi(j)}}$. Thus $\{\psi_{\tau_i}, \psi_{\tau'_i}\} = \{\psi_i, \psi_{i+s}\}$.

Now, let $\sigma_p = \otimes_{\tau} (\det^{w_\tau} \text{Sym}^{k_\tau - 2} k_p^2) \otimes_{\tau} \bar{\mathbb{F}}_p$ as before. Then $\sigma_p \in W_p(\rho)$ if and only if we have

$$(2) \quad \rho|_{I_p} \sim \prod_{\tau} \psi_{\tau}^{w_{\tau}} \begin{pmatrix} \prod_{j \in \mathbb{Z}/2s\mathbb{Z}} \psi_j^{m_j} & 0 \\ 0 & \prod_{j \in \mathbb{Z}/2s\mathbb{Z}} \psi_j^{m_{j+s}} \end{pmatrix},$$

such that for each $j \in \mathbb{Z}/2s\mathbb{Z}$ one of the following two conditions holds:

- (1) $\{m_j, m_{j+s}\} = \{k_j - 1, 0\}$ (note these are unordered sets).
- (2) $\{m_j, m_{j+s}\} = \{k_j - 2, 1\}$. Moreover, let $\alpha \in \{j, m + j\}$ be such that $m_\alpha = 1$. If $\{m_{j+1}, m_{j+1+s}\} = \{k_{j+1} - 1, 0\}$ (resp. $\{m_{j+1}, m_{j+1+s}\} = \{k_{j+1}, 1\}$), then $m_{\alpha+1} = 0$ (resp. $m_{\alpha+1} = 1$). Furthermore, there is some $l \in \mathbb{Z}/s\mathbb{Z}$ such that $k_l \neq p + 1$, and one of the following must hold:
 - (a) $k_j = p + 1$
 - (b) $[k_p : \mathbb{F}_p] \geq 3$ and $k_t \in \{p, p + 1\}$ for all $j \neq t \in \mathbb{Z}/s\mathbb{Z}$, and $k_t = p + 1$ for some t .

Clearly $W_p(\rho) \supset W_p^2(\rho)$. Moreover, if $2 \leq k_\tau \leq p$ for all τ , then condition (ii) never holds and $\sigma_p \in W_p(\rho)$ if and only if $\sigma_p \in W_p^2(\rho)$. Observe also that if $s = 1$ then $W_p(\rho) = W_p^2(\rho)$ for all ρ .

If $F = \mathbb{Q}$, then Theorem 1.3 is a result of Fontaine, proved in 1979 in letters to Serre. A proof was published in [Edi1], and our argument follows a similar method, albeit with more technical difficulties. From now on we assume $[F : \mathbb{Q}] > 1$, which allows us to use the results of [Car1].

As this work was completed, T. Gee proved many cases of the conjecture of [BDJ] by completely different methods [Gee2]. While his results apply more generally than ours, in the setting of this paper they are a subset of ours, as they apply only to weights with $3 \leq k_\tau \leq p - 1$ for all τ . Moreover, our argument can be extended to prove an analogous statement when p is ramified in F ([Sch1], Theorem 3.4), whereas Gee's techniques are inapplicable in that situation as the necessary modularity lifting theorems are unavailable at present.

Fix a representation $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$ satisfying the hypotheses of Theorem 1.3, and fix a suitable quaternion algebra B as in Definition 1.2. For every finite place v of F fix a uniformizer $\pi_v \in F_v$. We normalize the Artin reciprocity map so that uniformizers correspond to geometric Frobenius elements, and we write Frob_v for *arithmetic* Frobenius.

Unless stated otherwise, K will be the maximal unramified extension of F_p , and K'/K the totally ramified extension associated by class field theory to $\mathcal{O}_p^*/(1 + \pi_p \mathcal{O}_p)$. Let D and D' be the respective valuation rings of K and K' ; D is the Witt vectors of k_p . Let χ be the mod p cyclotomic character.

Our method is to construct a finite flat group scheme H/K such that $\text{Gal}(\bar{K}/K) \simeq I_p$ acts on $H(\bar{K})$ via the character ϕ . Suppose ρ is modular of weight σ . Choose a character $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \bar{\mathbb{F}}_p^*$ (see section 2.1 for definitions of the notations in this paragraph) such that σ_p is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_p)} \theta$. Using θ we find an appropriate piece \tilde{H} of $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p})}^{\text{bal}})[p^\infty]$ lifting H to D' . We compute the action of $\text{Gal}(K'/K)$ on $\text{cot}(\tilde{H}_{\bar{\mathbb{F}}_p})$, and Raynaud's theory of vector space schemes then

gives us a collection of possible ϕ ; this is precisely the collection of ϕ that are predicted by the conjecture of [BDJ] for some constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})}\theta$.

We then find the ϕ that appear in this way for all choices of θ such that $\sigma_{\mathfrak{p}}$ is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})}\theta$; they are exactly the ϕ in our definition of $W_{\mathfrak{p}}(\rho)$. Generically there are 2^s choices for θ , but if some of the parameters k_{τ} of $\sigma_{\mathfrak{p}}$ are $p+1$, there are fewer. This is the reason for condition (ii) above, which makes our result weaker than the one conjectured in [BDJ].

It is a pleasure to thank Richard Taylor for his very patient advice, explanations, and encouragement, Fred Diamond for describing his conjectures to us and for several fruitful conversations, and Kevin Buzzard for his comments on an earlier version of this work. This paper comes from part of the author's thesis [Sch2].

2. SHIMURA CURVES

2.1. Quaternionic Shimura curves. In this and the following two sections, we need not assume that p is unramified in F . We begin by defining some compact open subgroups of $G(\mathbb{A}^{\infty})$; we work with the Shimura curves with these level structures. Consider the following subgroups of $\text{GL}_2(\mathcal{O}_{\mathfrak{p}})$:

$$\begin{aligned} U^0(\mathfrak{p}) &= \text{GL}_2(\mathcal{O}_{\mathfrak{p}}) \\ U_0(\mathfrak{p}) &= \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}_2(\mathcal{O}_{\mathfrak{p}}) : c \in \mathfrak{p} \right\} \\ U_1(\mathfrak{p}) &= \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}_2(\mathcal{O}_{\mathfrak{p}}) : d-1, c \in \mathfrak{p} \right\} \\ U_1^{bal}(\mathfrak{p}) &= \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}_2(\mathcal{O}_{\mathfrak{p}}) : a-1, c, d-1 \in \mathfrak{p} \right\} \\ U^n(\mathfrak{p}) &= \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}_2(\mathcal{O}_{\mathfrak{p}}) : a-1, b, c, d-1 \in \mathfrak{p}^n \right\}, n \geq 1. \end{aligned}$$

We will denote by $\Gamma_0(\mathfrak{p})$ (resp. $\Gamma_1(\mathfrak{p})$, $\Gamma_1^{bal}(\mathfrak{p})$) the projection of $U_0(\mathfrak{p})$ (resp. $U_1(\mathfrak{p})$, $U_1^{bal}(\mathfrak{p})$) to $\text{GL}_2(k_{\mathfrak{p}})$. For $H \subset G(\mathbb{A}^{\infty, \mathfrak{p}})$, we define $K(n, H) = U^n(\mathfrak{p}) \times H$. We introduce the notations $M_{n, H} = M_{K(0, H)}$ and $M_{U_0(\mathfrak{p}), H} = M_{U_0(\mathfrak{p}) \times H}$. The Shimura curves $M_{U_1(\mathfrak{p}), H}$ and $M_{U_1^{bal}(\mathfrak{p}), H}$ are defined analogously. Finally, define $U_1(p) = \prod_{\mathfrak{p}|p} U_1(\mathfrak{p})$, and analogously for $U_1(p)^{bal}$. If the open compact $H \subset G(\mathbb{A}^{\infty, \mathfrak{p}})$ has the form $H = H_p \times H^p$ with $H_p \subset \prod_{v|p, v \neq \mathfrak{p}} \text{GL}_2(F_v)$ and $H^p \subset G(\mathbb{A}^{\infty, p})$, then we write $M_{U_1(p), H}$ for $M_{U_1(p) \times H^p}$, and analogously for $M_{U_1(p), H}^{bal}$.

Choose an order \mathcal{O} in the quaternion algebra B , and for every place v of F at which B splits, fix a ring isomorphism $B \otimes_F F_v \simeq M_2(F_v)$ such that $\mathcal{O} \otimes_{\mathcal{O}_F} \mathcal{O}_v$ is sent to $M_2(\mathcal{O}_v)$. Hence we get isomorphisms $(B \otimes F_v)^* \simeq \text{GL}_2(F_v)$.

Let $T \subset G$ be the torus $T = \text{Res}_{F/\mathbb{Q}} \mathbb{G}_m$, and let $\nu : G \rightarrow T$ be the reduced norm map. Then ([Car1], 1.2) the set of connected components of $M_K(\mathbb{C})$ are in bijection with the set $T(\mathbb{Q})^+ \backslash T(\mathbb{A}^{\infty}) / \nu(K)$, and ν induces the natural map $M_K \rightarrow \pi_0(M_K)$. Here $T(\mathbb{Q})^+$ denotes the totally positive elements of F^* , i. e. the ones which are mapped to $\mathbb{R}_{>0}$ by each embedding $F \hookrightarrow \mathbb{R}$.

For every $\gamma \in G(\mathbb{A}^{\infty})$ and every open compact $U \subset G(\mathbb{A}^{\infty})$, the action of γ by multiplication on the right induces an isomorphism $M_U(\mathbb{C}) \xrightarrow{\sim} M_{\gamma^{-1}U\gamma}(\mathbb{C})$. This action comes from a right action of $G(\mathbb{A}^{\infty})$ on the projective system of F -schemes M_U ([Car1], 1.1.2). Carayol proved that for sufficiently small H (where ‘‘sufficiently small’’ depends on n), the group $U^0(\mathfrak{p})/U^n(\mathfrak{p}) = \text{GL}_2(\mathcal{O}_{\mathfrak{p}}/\mathfrak{p}^n)$

acts freely on $M_{n,H}$ ([Car1], 1.4.1.3). Also let $\mathrm{GL}_2(\mathcal{O}_{\mathfrak{p}}/\mathfrak{p}^n)$ act from the right on $(\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$ by $v \cdot g = g^{-1}v$, for $g \in \mathrm{GL}_2(\mathcal{O}_{\mathfrak{p}}/\mathfrak{p}^n)$ and $v \in (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$. Carayol defined an $\mathcal{O}_{\mathfrak{p}}$ -module scheme $E_{n,H}$ over $M_{0,H}$ as follows:

$$E_{n,H} = (M_{n,H} \times (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2) / \mathrm{GL}_2(\mathcal{O}_{\mathfrak{p}}/\mathfrak{p}^n).$$

Note that $E_{n,H}$ is locally isomorphic to the constant $\mathcal{O}_{\mathfrak{p}}$ -module $(\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$, in the étale topology. If H is small enough that $E_{n,H}$ and $E_{n+1,H}$ are both defined, then $E_{n,H}$ is the \mathfrak{p}^n -torsion in $E_{n+1,H}$. Also, the $E_{n,H}$ are compatible with the maps $M_{0,H'} \rightarrow M_{0,H}$ for subgroups $H' \subset H$.

If H is sufficiently small, the Shimura curve $M_{0,H}$ over F has a (unique) smooth and proper integral model $\mathbf{M}_{0,H}$ over $\mathcal{O}_{(\mathfrak{p})} = \mathcal{O}_{\mathfrak{p}} \cap F$ ([Car1], 6.1). Furthermore, the $E_{n,H}$ lift to finite flat group schemes $\mathbf{E}_{n,H}$ of rank q^{2n} on $\mathbf{M}_{0,H}$, which retain the properties of $E_{n,H}$ listed above. Observe that only finitely many of the $\mathbf{E}_{n,H}$ are defined for any given H .

Now let H be small enough that $\mathbf{E}_{1,H}$ is defined on $\mathbf{M}_{0,H}$, and let S be any $\mathbf{M}_{0,H}$ -scheme. The pullback $\mathbf{E}_{1,H}|_S = \mathbf{E}_{1,H} \times_{\mathbf{M}_{0,H}} S$ is a finite locally free group scheme of rank q^2 over S equipped with an $\mathcal{O}_{\mathfrak{p}}$ -action. For $P \in \mathbf{E}_{1,H}(S)$, let $[P]$ be the image, viewed as a subscheme of $\mathbf{E}_{1,H}|_S$, and let \mathcal{I}_P be the ideal sheaf defining it. If J is a set of such sections, we write $\sum_{P \in J} [P]$ for the subscheme of $\mathbf{E}_{1,H}|_S$ defined by the product $\prod_{P \in J} \mathcal{I}_P$ of the ideal sheaves.

Definition 2.1. Let S be an $\mathbf{M}_{0,H}$ -scheme.

- (1) A *balanced $U_1(\mathfrak{p})$ -structure* on S is a pair (P, P') , such that there exists an fppf short exact sequence of $\mathcal{O}_{\mathfrak{p}}$ -modules on S :

$$0 \rightarrow \mathcal{K} \rightarrow \mathbf{E}_{1,H}|_S \rightarrow \mathcal{K}' \rightarrow 0$$

with \mathcal{K} and \mathcal{K}' finite locally free group schemes of rank q , such that $P \in \mathcal{K}(S)$ and $P' \in \mathcal{K}'(S)$, and P (resp. P') generates \mathcal{K} (resp. \mathcal{K}'). By this we mean that $\sum_{\alpha \in \mathcal{O}_{\mathfrak{p}}/\mathfrak{p}} [\alpha P^{(\prime)}] = \mathcal{K}^{(\prime)}$.

- (2) If H is small enough that $\mathbf{E}_{n,H}$ is defined, then a *Drinfeld basis of level n* on S is a homomorphism of $\mathcal{O}_{\mathfrak{p}}$ -modules $\varphi : (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2 \rightarrow \mathbf{E}_n|_S$ such that the subscheme $\sum_{\alpha \in (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2} [\varphi(\alpha)] \subset \mathbf{E}_n|_S$ coincides with $\mathbf{E}_n|_S$.

Each of the following functors associates a set to an $\mathbf{M}_{0,H}$ -scheme S .

$$\begin{aligned} \tilde{\mathcal{M}}_{U_1(\mathfrak{p}),H}^{bal} : S &\mapsto \{\text{balanced } U_1(\mathfrak{p})\text{-structures on } S\} \\ \tilde{\mathcal{M}}_{n,H} : S &\mapsto \{\text{Drinfeld bases of level } n \text{ on } S\} \end{aligned}$$

Carayol showed that the functor $\tilde{\mathcal{M}}_{n,H}$ is represented by a regular $\mathbf{M}_{0,H}$ -scheme $\mathbf{M}_{n,H}$ which is finite and flat over $\mathbf{M}_{0,H}$ and such that $\mathbf{M}_{n,H} \times_{\mathbf{M}_{0,H}} M_{0,H} = M_{n,H}$ ([Car1], 7.4). Jarvis constructed an integral model $\mathbf{M}_{U_1(\mathfrak{p}),H}$ of $M_{U_1(\mathfrak{p}),H}$, defined over $\mathcal{O}_{\mathfrak{p}}$; it is a regular scheme of dimension two, finite and flat over $\mathbf{M}_{0,H}$. Unfortunately, this model does not have semistable reduction, as can be seen from the description of its special fiber in [Jar2], 10.3. In the same way one can get an integral model of $M_{U_1(\mathfrak{p}),H}^{bal}$, but it suffers from the same disadvantages.

Let $F'_{\mathfrak{p}}/F_{\mathfrak{p}}$ be the totally ramified extension with Galois group $k_{\mathfrak{p}}^*$, and let $\mathcal{O}'_{\mathfrak{p}}$ be its ring of integers. By using the idea of canonical structures from Chapter 9 of [KM], one can get a model $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal}$ of $M_{U_1(\mathfrak{p}),H}^{bal}$ over $\mathrm{Spec}(\mathcal{O}_{\mathfrak{p}})$ that factors through $\mathrm{Spec}(\mathcal{O}'_{\mathfrak{p}})$ and does have semistable reduction over $\mathrm{Spec}(\mathcal{O}'_{\mathfrak{p}})$. Gee constructs such a curve in section 2 of [Gee1]; it is called $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal.can}$.

there. His argument follows Katz and Mazur and, though he works with unitary Shimura curves, transfers word for word to our quaternionic case.

Finally we describe the special fiber of this scheme. If S is a scheme over $\mathbf{M}_{0,H} \otimes_{\mathcal{O}_p} k_p$, then we have the absolute Frobenius morphism $\Phi : S \rightarrow S$, which acts by $x \mapsto x^q$ on the underlying affine rings. It induces the relative Frobenius $F : S \rightarrow S^{(q)}$, where $S^{(q)}$ is the base change of S by Φ .

We will be concerned with the Frobenius map $F : \mathbf{E}_{1,H}|_S \rightarrow \mathbf{E}_{1,H}|_{S^{(q)}}$. We also have the Verschiebung map $V : \mathbf{E}_{1,H}|_S^{(q)} \rightarrow \mathbf{E}_{1,H}|_S$, which is defined as the Cartier dual of:

$$F : (\mathbf{E}_{1,H}|_S)^\vee \rightarrow ((\mathbf{E}_{1,H}|_S)^\vee)^{(q)} = (\mathbf{E}_{1,H}|_{S^{(q)}})^\vee.$$

For each geometric point $x : \text{Spec } \overline{\mathbb{F}}_p \rightarrow \mathbf{M}_{0,H} \otimes k_p$ we get a \mathcal{O}_p -module as follows. Choose any lift y of x to a geometric point $y : \text{Spec } \overline{\mathbb{F}}_p \rightarrow \mathbf{M}_0 \otimes k_p$. Here \mathbf{M}_0 is the projective limit of the $\mathbf{M}_{0,H}$; since the $\mathbf{E}_{n,H}$ are compatible, it carries a p -divisible group \mathbf{E}_∞ . Now let $\mathbf{E}_\infty|_x$ be the pullback $\mathbf{E}_\infty \times_y \overline{\mathbb{F}}_p$; this is independent of the choice of y . It is a divisible height 2 formal \mathcal{O}_p -module over $\overline{\mathbb{F}}_p$. By Drinfeld's classification, $\mathbf{E}_\infty|_x$ is either $\Sigma_1 \times F_p/\mathcal{O}_p$ or Σ_2 , where Σ_h is the unique formal \mathcal{O}_p -module of height h . In the first case, we say that x is ordinary, and in the second that x is supersingular. Carayol ([Car1] 9.4.3) showed that there is a positive finite number of supersingular points. The following is essentially Theorem 2.18 of [Gee1].

Proposition 2.2. *The special fiber of $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal}$ consists of two smooth curves defined over k_p , which intersect transversally at the supersingular points. They are I , which parametrizes balanced $U_1(\mathfrak{p})$ -structures (P, P') where $\mathcal{K}^{(q)} = \ker V$ and E , which parametrizes structures with $\mathcal{K} = \ker F$.*

2.2. Hecke correspondences. Recall that $G(\mathbb{A}^\infty)$ acts on the projective system of M_U 's by right multiplication. For every $\gamma \in G(\mathbb{A}^\infty)$ we can define a Hecke correspondence $[U\gamma U]$ on M_U by the following diagram.

$$(3) \quad \begin{array}{ccc} & M_{\gamma U \gamma^{-1} \cap U} & \xrightarrow{\gamma} & M_{U \cap \gamma^{-1} U \gamma} \\ & \searrow \pi_1 & & \searrow \pi_2 \\ M_U & & & M_U \end{array}$$

Clearly these correspondences extend to the integral models \mathbf{M}_U . Now, let Δ be a set of left coset representatives of $U/(\gamma U \gamma^{-1} \cap U)$. Then clearly the correspondence $[U\gamma U]$ sends a class $xU \in M_U(\mathbb{C})$ (here $x \in G(\mathbb{Q}) \backslash G(\mathbb{A}^\infty) \times (\mathbb{C} \backslash \mathbb{R})$) to the set $\{x\delta\gamma U : \delta \in \Delta\} \subset M_U(\mathbb{C})$. It is easy to see that this is independent of the choice of Δ . Now observe that:

$$[U\gamma U] = \coprod_{\delta \in \Delta} \delta\gamma U.$$

We find that the correspondence $[U\gamma U]$ depends only on the double coset $[U\gamma U]$, which justifies the notation.

In practice we will always assume that $U = U_p \times H$, where U_p is one of the groups $U^n(\mathfrak{p})$, $U_0(\mathfrak{p})$, $U_1(\mathfrak{p})$, or $U_1^{bal}(\mathfrak{p})$. Let v be a finite place of F at which B splits. For $x \in \text{GL}_2(F_v)$, we denote by x_v the element of $\mathbf{G}(\mathbb{A}^\infty)$ which is trivial everywhere except at v , where it is x (recall that we have

fixed an identification of $(B \otimes F_v)^*$ with $\mathrm{GL}_2(F_v)$). Define the Hecke correspondences:

$$\begin{aligned} T_v(U) &= \left[U \begin{pmatrix} \pi_v & 0 \\ 0 & 1 \end{pmatrix}_v U \right] \\ S_v(U) &= \left[U \begin{pmatrix} \pi_v & 0 \\ 0 & \pi_v \end{pmatrix}_v U \right] \end{aligned}$$

Observe that the correspondence S_v is in fact an operator, since the element of $G(\mathbb{A}^\infty)$ in question is central. If $U_{\mathfrak{p}} \supset U^1(\mathfrak{p})$ and $\alpha, \beta \in k_{\mathfrak{p}}^*$, then we define the diamond operator

$$\langle \alpha, \beta \rangle(U) = \left[U \begin{pmatrix} \tilde{\alpha} & 0 \\ 0 & \tilde{\beta} \end{pmatrix}_{\mathfrak{p}} U \right]$$

where $\tilde{\alpha}, \tilde{\beta} \in \mathcal{O}_{\mathfrak{p}}^*$ are any lifts of α and β . Note that this is well-defined.

By decreasing the open compact U if necessary, we may assume that U is a product $\prod_v U_v$, with $U_v \in (B \otimes F_v)^*$. Let S be a finite collection of places of F containing all infinite places, all places dividing p , all places where B ramifies, and all places v such that $U_v \neq \mathrm{GL}_2(\mathcal{O}_v)$. We define \mathbf{T}_U to be the commutative \mathbb{Z} -subalgebra of $\mathrm{End}(\mathrm{Pic}^0(M_U))$ generated by $T_v(U)$ and $S_v(U)$ for all $v \notin S$, and $\tilde{\mathbf{T}}_U$ to be the subalgebra generated by \mathbf{T}_U and $\langle \alpha, \beta \rangle(U)$ for all $\alpha, \beta \in k_{\mathfrak{p}}^*$. Clearly \mathbf{T}_U and $\tilde{\mathbf{T}}_U$ are finite free over \mathbb{Z} .

Let $\mathbf{T} = \mathbf{T}_{U^1(p) \times H^p}$ and $\tilde{\mathbf{T}} = \tilde{\mathbf{T}}_{U^1(p) \times H^p}$. Then \mathbf{T} surjects onto $\mathbf{T}_{U_1(\mathfrak{p})^{bal} \times H}$ and $\mathbf{T}_{K(0, H)}$, so we can view \mathbf{T} as acting via correspondences on $M_{U_1(\mathfrak{p})^{bal}, H}$ and $M_{0, H}$, and similarly for $\tilde{\mathbf{T}}$.

2.3. An Eichler-Shimura relation. Let $U \subset G(\mathbb{A}^\infty)$ be one of the groups considered in section 2.2. Then the Hecke correspondences defined there commute with the $\mathrm{Gal}(\bar{F}/F)$ -action on \mathbf{M}_U , and these induce commuting actions of \mathbf{T} and of $\mathrm{Gal}(\bar{F}/F)$ on $\mathrm{Pic}^0(\mathbf{M}_U)$. In this section we prove an Eichler-Shimura relation.

First we briefly recall some facts about special fibers; see [Car1], §9 for the details. Let $F_{\mathfrak{p}}^{ab}$ and $F_{\mathfrak{p}}^{nr}$ be the maximal abelian and maximal unramified extensions of $F_{\mathfrak{p}}$, respectively, and let $F_{\mathfrak{p}}^{nr} \subset F_{\mathfrak{p}}^n \subset F_{\mathfrak{p}}^{ab}$ be the extension for which local class field theory identifies $\mathrm{Gal}(F_{\mathfrak{p}}^n/F_{\mathfrak{p}}^{nr})$ with $\mathcal{O}_{\mathfrak{p}}^*/(1 + \pi_{\mathfrak{p}}^n \mathcal{O}_{\mathfrak{p}})$. We denote by $\mathcal{O}_{\mathfrak{p}}^n$ the ring of integers of $F_{\mathfrak{p}}^n$. The action of the Weil group $W(F_{\mathfrak{p}}^{ab}/F_{\mathfrak{p}})$ on the ring $\mathcal{O}_{\mathfrak{p}}^n$ induces a right action of this group on the regular scheme $\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \mathcal{O}_{\mathfrak{p}}^n$, the normalization of $\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \mathcal{O}_{\mathfrak{p}}^n$. Write $\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p$ for the special fiber $(\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \mathcal{O}_{\mathfrak{p}}^n) \otimes_{\mathcal{O}_{\mathfrak{p}}^n} \bar{\mathbb{F}}_p$. Let $\mathbf{M} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p$ be the projective limit of these, with respect to both n and H ; it has commuting actions of $W(F_{\mathfrak{p}}^{ab}/F_{\mathfrak{p}})$ and $G(\mathbb{A}^\infty)$.

Let $\varphi_n : (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2 \rightarrow \mathbf{E}_{n, H}$ be the universal Drinfeld basis of level n on $\mathbf{M} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p$, and A_n be a line in the $(\mathcal{O}_{\mathfrak{p}}/\mathfrak{p}^n)$ -module $(\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$, i.e. a rank 1 direct summand. Then let $(\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p)_{A_n}$ be the closed smooth subscheme cut out by the condition $\varphi_n|_{A_n} = 0$; these are precisely the irreducible components of $\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p$. In particular, if A_n is always the line generated by $(1, 0) \in (\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$, then the $(\mathbf{M}_{n, H} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p)_{A_n}$ are clearly compatible, and the projective limit $(\mathbf{M} \otimes_{\mathcal{O}_{\mathfrak{p}}} \bar{\mathbb{F}}_p)_A$ is stable under the subgroup of $G(\mathbb{A}^\infty) \times W(F_{\mathfrak{p}}^{ab}/F_{\mathfrak{p}})$ consisting of elements whose $\mathrm{GL}_2(F_{\mathfrak{p}})$ -component is an upper triangular matrix. Denote the valuation of $F_{\mathfrak{p}}$ by v .

For $w \in W(F_{\mathfrak{p}}^{ab}/F_{\mathfrak{p}})$, let $[w] \in F_{\mathfrak{p}}^*$ be its image under the Artin reciprocity map. Then Carayol proved ([Car1], 10.3) that elements of $\mathrm{GL}_2(F_{\mathfrak{p}}) \times W(F_{\mathfrak{p}}^{ab}/F_{\mathfrak{p}})$ of the form

$$\left(\left(\begin{array}{cc} [w] & * \\ 0 & 1 \end{array} \right), w \right)$$

act on $(\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A$ by the $(v([w]^{-1}))$ -th power of absolute Frobenius.

Remark 2.3. Let A'_n be the subspace of $(\mathfrak{p}^{-n}/\mathcal{O}_{\mathfrak{p}})^2$ generated by $(0, 1)$. Then, by exactly the same argument as in [Car1], 10.4 we can show that

$$\left(\left(\begin{array}{cc} 1 & 0 \\ * & [w] \end{array} \right), w \right)$$

acts on $(\mathbf{M} \otimes \overline{\mathbb{F}}_p)_{A'}$ by the $(v([w]^{-1}))$ -th power of absolute Frobenius; in particular, it stabilizes points.

Now let $\overline{\Phi}$ denote absolute Frobenius on $(\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A$, which acts as $x \mapsto x^q$ on underlying affine rings. By Carayol's congruence relation, the action of $\overline{\Phi}$ on $(\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A$ is the same as that of

$$\left(\left(\begin{array}{cc} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & 1 \end{array} \right), \mathrm{Frob}_{\mathfrak{p}} \right).$$

Hence the same is true on cohomology $H^1((\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A, \mathbb{Q}_l)$ for any prime $l \neq p$, and so on $H^1(\mathbf{M}_{0,H} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l) = H^1((\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A, \mathbb{Q}_l)^{U_0(\mathfrak{p}) \times H}$. (Recall that only the upper triangular elements of $\mathrm{GL}_2(F_{\mathfrak{p}})$ act on $H^1((\mathbf{M} \otimes \overline{\mathbb{F}}_p)_A, \mathbb{Q}_l)$.) But note that $\overline{\Phi} = \Phi \times \mathrm{Frob}_{\mathfrak{p}}$, where Φ is absolute Frobenius on $\mathbf{M}_{0,H} \otimes_{\mathcal{O}_{\mathfrak{p}}} k_{\mathfrak{p}}$ and $\mathrm{Frob}_{\mathfrak{p}}$ acts on $\mathcal{O}_{\mathfrak{p}}^{nr}$ and hence $\overline{\mathbb{F}}_p$. Thus we have the following equality of correspondences on $H^1(\mathbf{M}_{0,H} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$:

$$\Phi \times 1 = \left(\begin{array}{cc} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & 1 \end{array} \right).$$

The right action of $G(\mathbb{A}^{\infty})$ on the projective system of Shimura curves \mathbf{M}_U induces a left action on the injective system of $H^1(\mathbf{M}_U \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$. By the diagram (3), we see that the operator $[U\gamma U]$ on cohomology is given by the following composition of morphisms, where tr is a trace map:

$$H^1(\mathbf{M}_U \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l) \xleftarrow{tr} H^1(\mathbf{M}_{\gamma U \gamma^{-1} \cap U} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l) \xleftarrow{\gamma^*} H^1(\mathbf{M}_{U \cap \gamma^{-1} U \gamma} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l) \xleftarrow{\pi_2^*} H^1(\mathbf{M}_U \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$$

Lemma 2.4. *On $H^1(\mathbf{M}_{0,H} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$ we have the following relation of correspondences:*

$$S_{\mathfrak{p}} \Phi^2 - T_{\mathfrak{p}} \Phi + q = 0.$$

Proof. To simplify notation, let $U = K(0, H)$ throughout this proof. Let $\gamma \in G(\mathbb{A}^{\infty})$ be the element which is $\left(\begin{array}{cc} \pi_{\mathfrak{p}} & 0 \\ 0 & 1 \end{array} \right)$ at \mathfrak{p} and 1 elsewhere. For each $\alpha \in k_{\mathfrak{p}}$, choose an $\tilde{\alpha} \in \mathcal{O}_{\mathfrak{p}}$ lifting it. Then:

$$U\gamma U = \prod_{\alpha \in k_{\mathfrak{p}}} \left(\begin{array}{cc} \pi_{\mathfrak{p}} & \tilde{\alpha} \\ 0 & 1 \end{array} \right) U \prod \left(\begin{array}{cc} 1 & 0 \\ 0 & \pi_{\mathfrak{p}} \end{array} \right) U.$$

Thus, on $H^1(\mathbf{M}_{0,H} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$ we have that $(S_{\mathfrak{p}}\Phi^2 - T_{\mathfrak{p}}\Phi + q)x =$

$$\begin{aligned} & \begin{pmatrix} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & \pi_{\mathfrak{p}} \end{pmatrix}^* x - \sum_{\alpha \in k_{\mathfrak{p}}} \begin{pmatrix} \pi_{\mathfrak{p}} & \tilde{\alpha} \\ 0 & 1 \end{pmatrix}^* \begin{pmatrix} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & 1 \end{pmatrix}^* x - \begin{pmatrix} 1 & 0 \\ 0 & \pi_{\mathfrak{p}} \end{pmatrix}^* \begin{pmatrix} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & 1 \end{pmatrix}^* x + qx = \\ & \begin{pmatrix} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & \pi_{\mathfrak{p}} \end{pmatrix}^* x - qx - \begin{pmatrix} \pi_{\mathfrak{p}}^{-1} & 0 \\ 0 & \pi_{\mathfrak{p}} \end{pmatrix}^* x + qx = 0 \end{aligned}$$

□

Remark 2.5. The action of $\Phi \times 1$ on $H^1(\mathbf{M}_{0,H} \otimes \overline{\mathbb{F}}_p, \mathbb{Q}_l)$ or $H^1(\mathbf{M}_{0,H} \otimes_{\mathcal{O}_{(p)}} \overline{F}, \mathbb{Q}_l)$ corresponds to the Galois action of geometric Frobenius: $1 \times \text{Frob}_p^{-1}$. Hence $S_{\mathfrak{p}}\text{Frob}_p^{-2} - T_{\mathfrak{p}}\text{Frob}_p^{-1} + q = 0$. Equivalently,

$$\text{Frob}_p^2 - (N\mathfrak{p})^{-1}T_{\mathfrak{p}}\text{Frob}_p + (N\mathfrak{p})^{-1}S_{\mathfrak{p}} = 0.$$

Proposition 2.6. *Let v be a place of F such that $v \notin S$ (see section 2.2 for the definition of S) and v lies over a rational prime $l \neq p$ which is unramified in F . Suppose also that B splits at all places of F above l ; clearly, almost all v satisfy these conditions. Then, the element $\text{Frob}_v^2 - T_v\text{Frob}_v + NvS_v$ of $\mathbf{T}[\text{Gal}(\overline{F}/F)]$ annihilates $\text{Pic}^0(\mathbf{M}_U)[p](\overline{F})$.*

Proof. We apply the previous lemma and remark with v in the role of \mathfrak{p} and p in the role of l . We can do this as the only assumptions made about \mathfrak{p} in [Car1] were that p is unramified in F and that B splits at all places of F above p . Since $v \notin S$, we have $U = \text{GL}_2(\mathcal{O}_v) \times U'$, where $U' \subset G(\mathbb{A}^{\infty,v})$. Recall that $\mathbf{M}_{0,H}$ was defined as the unique (up to canonical isomorphism) extension of the F -scheme $M_{0,H}$ to $\mathcal{O}_{(p)} = F \cap \mathcal{O}_{\mathfrak{p}}$. Similarly, there is an integral model $\mathbf{M}_{0,U'}^{(v)}$ of M_U over $\mathcal{O}_{(v)}$. Then by the previous remark we have the following relation on $H^1(\mathbf{M}_{0,U'}^{(v)} \otimes_{\mathcal{O}_{(v)}} \overline{F}, \mathbb{Q}_p)$:

$$\text{Frob}_v^2 - (Nv)^{-1}T_v\text{Frob}_v + (Nv)^{-1}S_v = 0.$$

Note that $\mathbf{M}_{0,H} \otimes \overline{F} \simeq \mathbf{M}_{0,U'}^{(v)} \otimes \overline{F}$. Since $H^1(M_{0,H} \otimes \overline{F}, \mathbb{Z}_p)$ is the cohomology of a proper smooth curve, it has no torsion and so injects into $H^1(M_{0,H} \otimes \overline{F}, \mathbb{Q}_p)$. Hence the above relation holds there as well. This group in turn surjects onto $H^1(M_{0,H} \otimes \overline{F}, \mathbb{F}_p)$. Now, $\text{Pic}^0(M_U)[p] = H^1(M_{0,H} \otimes \overline{F}, \mu_p)$, and on this space the action of Frob_v is twisted by (Nv) with respect to what we had earlier:

$$(Nv)^{-2}\text{Frob}_v^2 - (Nv)^{-2}T_v\text{Frob}_v + (Nv)^{-1}S_v = 0.$$

This completes the proof. □

2.4. A geometric Galois action. Let D, K, D', K' be as defined in the introduction; note that K' and D' are $F_{\mathfrak{p}}^1$ and $\mathcal{O}_{\mathfrak{p}}^1$, respectively. Each $\sigma \in \text{Gal}(K'/K)$ induces an automorphism $\text{id} \times \text{Spec}(\sigma)$ of $\mathbf{M}_{1,H} \otimes D'$ and hence of the normalization $\mathbf{M}_{1,H} \otimes_D D'$ and of its special fiber $\mathbf{M}_{1,H} \otimes \overline{\mathbb{F}}_p$. Similarly we get an automorphism of $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \otimes \overline{\mathbb{F}}_p$, i.e. of the special fiber of the normalization of $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \otimes_{\mathcal{O}'_{\mathfrak{p}}} D'$. Recall that this latter scheme has two smooth components, which we denoted by I and E .

Let $j : \text{Gal}(K'/K) \rightarrow \mathcal{O}_{\mathfrak{p}}^*/(1 + \mathfrak{p})$ be the isomorphism induced by the Artin reciprocity map. Then we have the following analogue of [Gro], Prop. 7.2:

Proposition 2.7. *Let $\sigma \in \text{Gal}(K'/K)$. Then σ acts on I and E by $\langle j(\sigma)^{-1}, 1 \rangle$ and $\langle 1, j(\sigma)^{-1} \rangle$, respectively.*

Proof. Let $\varphi : (\mathfrak{p}^{-1}/\mathcal{O}_{\mathfrak{p}})^2 \rightarrow \mathbf{E}_1|_T$ be a Drinfeld basis of level 1 on a $\mathbf{M}_{0,H}$ -scheme T . It induces a $U_1(\mathfrak{p})^{bal}$ -structure as follows: $P = \varphi(0, 1)$, and P' is the image in $\mathbf{E}_1|_T/\langle P \rangle$ of $\varphi(1, 0)$. This morphism of moduli problems is represented by the map $\mathbf{M}_{1,H} \rightarrow \mathbf{M}_{U_1(\mathfrak{p}),H}^{bal}$. In the notation of the previous section, $(\mathbf{M}_{1,H} \overline{\otimes} \overline{\mathbb{F}}_p)_{A'}$ is the subscheme defined by the condition $\phi_1|_{A'} = 0$. It clearly maps to the E component of $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} \overline{\mathbb{F}}_p$; indeed, if φ is zero on A'_1 , then in the corresponding $U_1(\mathfrak{p})^{bal}$ -structure $\langle P \rangle$ is connected and so is $\ker F$. Similarly, $(\mathbf{M}_{1,H} \overline{\otimes} \overline{\mathbb{F}}_p)_{A'}$ lands on the I component.

Now by [Car1], 10.3 and Remark 2.3 the actions of σ on $(\mathbf{M} \overline{\otimes} \overline{\mathbb{F}}_p)_A$ and $(\mathbf{M} \overline{\otimes} \overline{\mathbb{F}}_p)_{A'}$ are identical, respectively, to those of

$$\begin{pmatrix} j(\sigma)^{-1} & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ 0 & j(\sigma)^{-1} \end{pmatrix}.$$

Hence σ acts on I as $\langle j(\sigma)^{-1}, 1 \rangle$ and on E as $\langle 1, j(\sigma)^{-1} \rangle$. \square

Remark 2.8. This proposition can also be proved by passing to Carayol's unitary Shimura varieties \mathbf{M}'_U , which have a modular description. There one can carry out a similar argument to [Gro], Proposition 7.2. One uses the results of [Kas], 3.4 to define the analogue of the operator w_{ζ} .

Let J be the Néron model over D' of the curve $\text{Pic}^0(M_{U_1(\mathfrak{p}),H}^{bal}) \times_K K'$. Then $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} D')$ is isomorphic to the identity component J^0 of J by [BLR1], 9.5/4.

The right action of $\text{Gal}(K'/K)$ on $M_{U_1(\mathfrak{p}),H}^{bal} \times_K K'$, and hence on $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} D'$, that we just studied induces a left action of the same group on $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} D')$. Similarly, we have a right action of $\text{Gal}(K'/K)$ on $\text{Pic}^0(M_{U_1(\mathfrak{p}),H}^{bal}) \times_K K'$ and hence on its Néron model. By the above isomorphism, this induces a right action on $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} D')$.

Lemma 2.9. *Let $\sigma \in \text{Gal}(K'/K)$. Then the right action of σ on $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \overline{\otimes} D')$ is identical to the left action of σ^{-1} .*

Proof. For any regular curve X/K , there are two actions of $\text{Gal}(K'/K)$ on $\text{Pic}^0(X \times_K K')$: the obvious left action, and the right action coming from the isomorphism $\text{Pic}^0(X \times_K K') \simeq \text{Pic}^0(X) \times_K K'$. By the universal property of the Picard functor, the right action of σ on $\text{Pic}^0(X \times_K K')$ is identical to the left action of σ^{-1} , and the lemma follows immediately. \square

3. MOD p HILBERT MODULAR FORMS

3.1. Mod p Hilbert modular forms. From here on we assume that p is unramified in F . We first describe a general construction for lisse étale mod p sheaves associated to modular representations; cf. the p -adic sheaves of [Car2], 2.1.3 and [HT], III.2.

Consider a compact open $U \subset G(\mathbb{A}^{\infty})$, and let $U \supset U_1 \supset U_2 \supset \dots$ be a sequence of subgroups such that for each n , $M_{U_n} \rightarrow M_U$ is a finite étale Galois cover. Set $G = \varprojlim \text{Gal}(M_{U_n}/M_U)$. Let $\eta : G \rightarrow \text{GL}(W)$ be a representation of G on an $\overline{\mathbb{F}}_p$ -vector space W , and suppose that η is defined over a finite field $\mathbb{F} \subset \overline{\mathbb{F}}_p$.

We define a lisse étale sheaf \mathcal{F}_{η} on M_U as follows. Since η has finite image, it factors through $\text{Gal}(M_{U_N}/M_U)$ for some N . The sections of \mathcal{F}_{η} over an étale cover $Y \rightarrow M_U$ will be the functions

$$f : \pi_0(M_{U_N} \times_{M_U} Y) \rightarrow W$$

such that for all $\gamma \in \text{Gal}(M_{U_N}/M_U)$ and all $C \in \pi_0(M_{U_N} \times_{M_U} Y)$ the following relation holds:

$$f(C\gamma) = \eta(\gamma)^{-1}f(C).$$

In particular, if $M_V \rightarrow M_U$ is a finite étale Galois cover and η is an $\overline{\mathbb{F}}_p$ -representation of $\text{Gal}(M_V/M_U)$, then this construction produces a lisse étale sheaf \mathcal{F}_η on M_U .

Lemma 3.1. *For H^p sufficiently small, the finite étale cover $M_{U^1(p),H} \rightarrow M_{0,H}$ is Galois with Galois group $\text{GL}_2(\mathcal{O}_F/p)$.*

Proof. By [Car1], 1.4.1.1, with the observation that it suffices to vary H^p in that argument, for H^p sufficiently small $M_{1,H} \rightarrow M_{0,H}$ is a Galois cover with Galois group $K(0, H)/K(1, H)(K(0, H) \cap Z(\mathbb{Q}))$. Shrinking H^p if necessary, we can assume $(Z(\mathbb{Q}) \cap K(0, H))_{\mathfrak{p}} \subset \mathcal{O}_F^*$. By [Che], Théorème 1, there exists a rational integer a such that $(a, p) = 1$ and $\mathcal{O}_F^* \cap (1 + a\mathcal{O}_F) \subset 1 + \mathfrak{p}\mathcal{O}_F$. Now for all $v \notin S$, we require that $H_v \subset \{g \in \text{GL}_2(\mathcal{O}_v) : g \equiv I_2 \pmod{\pi_v^{\text{ord}_v(a)}}\}$. Then we have $(Z(\mathbb{Q}) \cap K(0, H))_{\mathfrak{p}} \subset U^1(\mathfrak{p})$. Repeating this process for all places of F dividing p , we obtain the lemma. \square

Assume from now on that H^p is sufficiently small to satisfy the previous lemma. Then for any modular representation $\xi : \text{GL}_2(\mathcal{O}_F/p) \rightarrow \text{GL}(W_\xi)$ we obtain by the construction above a lisse étale mod p sheaf \mathcal{F}_ξ on $M_{0,H}$.

Definition 3.2. Let σ be a weight. A *mod p Hilbert modular class of weight σ* is an element of $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$.

3.2. Hecke operators on cohomology. The Hecke correspondences on M_U introduced in Section 2.2 induce Hecke operators on the spaces $H^1(M_U \otimes \overline{F}, \mathcal{F}_\sigma)$. We will study them in this section. Set $\mathcal{F} = \mathcal{F}_\sigma$. The induced operator $[UxU]$ on cohomology is the composition of three maps:

$$(4) \quad H^i(M_U \otimes \overline{F}, \mathcal{F}) \rightarrow H^i(M_{xUx^{-1} \cap U} \otimes \overline{F}, x^*\mathcal{F}) \rightarrow H^i(M_{xUx^{-1} \cap U} \otimes \overline{F}, \mathcal{F}) \rightarrow H^i(M_U \otimes \overline{F}, \mathcal{F}).$$

The leftmost map is pullback by x , and the rightmost map is the trace. To give an explicit description of the map in the middle, we will construct the relevant map of sheaves $x^*\mathcal{F} \rightarrow \mathcal{F}$ on $M_{xUx^{-1} \cap U}$. By adjointness of x^* and x_* , this corresponds to a map $\mathcal{F} \rightarrow x_*\mathcal{F}$ on M_U .

Now let $Y \rightarrow M_U$ be an étale cover. To construct our map

$$(5) \quad \mathcal{F}(Y) \rightarrow x_*\mathcal{F}(Y) = \mathcal{F}(Y \times_{M_U, x} M_{xUx^{-1} \cap U}),$$

choose normal subgroups $W \subset U$ and $W' \subset xUx^{-1} \cap U$ such that W and W' trivialize \mathcal{F} . In other words, $\text{Gal}(M_W/M_U)$ (resp. $\text{Gal}(M_{W'}/M_{xUx^{-1} \cap U})$) acts on Λ/p via σ . It is easy to see that we can suppose $W' \subset xWx^{-1}$. Then we read off from the definition of the sheaf \mathcal{F} that

$$\mathcal{F}(Y) = \{f : \pi_0(Y \times_{M_U} M_W) \rightarrow \Lambda/p : \forall C \in \pi_0, \forall \gamma \in U, f(C\gamma) = \sigma(\gamma)^{-1}f(C)\}$$

and

$$\begin{aligned} \mathcal{F}(Y \times_{M_U, x} M_{xUx^{-1} \cap U}) = \\ \{F : \pi_0(Y \times_{M_U, x} M_{W'}) \rightarrow \Lambda/p : \forall C \in \pi_0, \forall \gamma' \in xUx^{-1} \cap U, F(C\gamma') = \sigma(\gamma')^{-1}F(C)\}. \end{aligned}$$

Observe that the conjugation map $x : M_{W'} \rightarrow M_{x^{-1}W'x} \rightarrow M_W$ induces a map on the base changes:

$$x : Y \times_{M_U, x} M_{W'} \rightarrow Y \times_{M_U} M_W.$$

Then the map in (5) is $f \mapsto F_f$, where $F_f(C) = xf(Cx)$.

3.3. Eisenstein ideals. We generalize the notion of Eisenstein ideals to F as in [Jar1]. Let $\mathfrak{d} \subset \mathcal{O}_F$ be an integral ideal, and set $A_{\mathfrak{d}} = \{a \in (\mathbb{A}_F^{\infty})^* : a - 1 \in \mathfrak{d}\}$. The Artin reciprocity map induces an isomorphism between $Cl(\mathfrak{d}) = \mathbb{A}_F^*/F^*A_{\mathfrak{d}}(F_{\infty}^*)^0$ and $\text{Gal}(F_{\mathfrak{d}}/F)$, where $F_{\mathfrak{d}}$ is the narrow ray class field associated to \mathfrak{d} . If v is a place of F , we will say that $v \equiv 1 \pmod{\mathfrak{d}}$ if $\pi_v \in F^*A_{\mathfrak{d}}$; observe that this definition is independent of the choice of uniformizer π_v .

Definition 3.3. If $\mathfrak{m} \subset \mathbf{T}$ is a maximal ideal, we will say that \mathfrak{m} is *Eisenstein* if there exists an integral ideal $\mathfrak{d} \subset \mathcal{O}_F$ such that for almost all places v of F such that $v \equiv 1 \pmod{\mathfrak{d}}$, we have $T_v - 2 \in \mathfrak{m}$ and $S_v - 1 \in \mathfrak{m}$.

Let $\mathfrak{m} \subset \mathbf{T}$ be a maximal ideal. If there exists a representation $\text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ such that for almost all places v of F , the characteristic polynomial of the action of Frob_v is $x^2 - T_v x + NvS_v$, then we will denote it by $\rho_{\mathfrak{m}}$.

Proposition 3.4. *Let $\mathfrak{m} \subset \mathbf{T}$ be a maximal ideal such that $H^1(M_U \otimes \overline{F}, \mathcal{F}_{\sigma})_{\mathfrak{m}} \neq 0$ for some weight σ . Suppose $\rho_{\mathfrak{m}}$ exists. Then $\rho_{\mathfrak{m}}$ is reducible if and only if \mathfrak{m} is Eisenstein.*

Proof. In the case $F = \mathbb{Q}$, this is Proposition 2 of [DT], and the proof is essentially the same. See also the comments in section 3 of [Jar1]. Let $\rho_{\mathfrak{m}}$ be reducible. Then its semisimplification has the form $\chi \oplus \chi^{-1} \det \rho_{\mathfrak{m}}$ for some character χ of $\text{Gal}(\overline{F}/F)$, hence of $\text{Gal}(\overline{F}/F)^{ab}$. Let \mathfrak{c} be the conductor of χ . Then for almost all $v \equiv 1 \pmod{\mathfrak{c}}$, we have $T_v \equiv \text{tr} \rho \equiv 1 + NvS_v \pmod{\mathfrak{m}}$. Observe that the Artin reciprocity map identifies $\pi_0(M_U) \simeq (\mathbb{A}_F^{\infty})^*/F^+\nu(U)$ with a finite quotient of $\text{Gal}(\overline{F}/F)^{ab}$. Hence $\pi_0(M_U)$ is a quotient of $Cl(\mathfrak{d})$ for some integral ideal $\mathfrak{d} \subset \mathcal{O}_F$; see [Neu] VI.1.8 for a proof of this fact. From the description of the action of the Hecke operators in section 3.2, it is clear that if $v \equiv 1 \pmod{\mathfrak{d}}$ and v does not divide p , then S_v acts like the identity on $H^1(M_U \otimes \overline{F}, \mathcal{F}_{\sigma})$, whence $S_v - 1 \in \text{Ann}_{\mathbf{T}} H^1(M_U \otimes \overline{F}, \mathcal{F}_{\sigma}) \subset \mathfrak{m}$. Now choose \mathfrak{d}' such that the narrow ray class field $F_{\mathfrak{d}'}$ contains $F_{\mathfrak{c}\mathfrak{d}}(\mu_p)$. Then for almost all $v \equiv 1 \pmod{\mathfrak{d}'}$ we see that $S_v - 1 \in \mathfrak{m}$ and $T_v - 2 \in \mathfrak{m}$ (since $Nv \equiv 1 \pmod{p}$).

The argument in the converse direction is identical to that in [DT]. \square

Proposition 3.5. *Let $A = H^0(M_U \otimes \overline{F}, \mathcal{F}_{\xi})$ for a sheaf \mathcal{F}_{ξ} as above. Let $\mathfrak{m} \subset \mathbf{T}$ be a maximal ideal such that \mathfrak{m} contains p and $A_{\mathfrak{m}} \neq 0$. Then \mathfrak{m} is Eisenstein.*

Proof. This is analogous to the proof of Lemma 3 in [DT]. Consider A as a $\mathbf{T}[\text{Gal}(\overline{F}/F)]$ -representation. It is clear from the construction of \mathcal{F}_{σ} that the action of $\text{Gal}(\overline{F}/F)$ on the semisimplification of $A \otimes \overline{\mathbb{F}}_p$ factors through $\pi_0(M_U) \simeq (\mathbb{A}_F^{\infty})^*/F^+\nu(U)$, hence through $Cl(\mathfrak{d})$, where \mathfrak{d} is as in the proof of Proposition 3.4. Therefore, $\text{Frob}_v - 1$ is nilpotent on A for all $v \equiv 1 \pmod{\mathfrak{d}}$. From the Eichler-Shimura relation (Proposition 2.6) we know that $T_v - 2 = (\text{Frob}_v - 1)(1 - \text{Frob}_v^{-1})$; since the two factors commute, $T_v - 2$ is also nilpotent on A .

From the description of the Hecke operators in section 3.2 we see that $S_v = 1$ on A for all $v \equiv 1 \pmod{\mathfrak{d}}$ such that $v \nmid p$. If $A_{\mathfrak{m}} \neq 0$, then $\text{Ann}(A) \subset \mathfrak{m}$, whence $T_v - 2 \in \mathfrak{m}$ and $S_v - 1 \in \mathfrak{m}$ for v as above. \square

Remark 3.6. Observe that the same result holds for $A = H^2(M_U \otimes \overline{F}, \mathcal{F}_{\xi})$ by Poincaré duality.

3.4. Lifting forms to “weight 2”. Suppose that $\sigma_{\mathfrak{p}}$ is an irreducible constituent of $\xi = \text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$, for a one-dimensional representation $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$. Indeed such a θ can always be found; generically there are 2^s of them. Let $I = \{\tau_1, \tau_2, \dots, \tau_s\}$ be the set of embeddings $k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p$. We will

assume they are labeled so that $\tau_{i-1} = \tau_i^p$ for $i \in \mathbb{Z}/s\mathbb{Z}$, and we will write w_i and k_i for w_{τ_i} and k_{τ_i} , respectively. Denote the set of all subsets of I by $\mathcal{P}(I)$. If $S \in \mathcal{P}(I)$, let ν_S be the characteristic function of $S^{(p)}$:

$$\nu_S(\tau_i) = \begin{cases} 1 & : \tau_{i+1} \in S \\ 0 & : \tau_{i+1} \notin S \end{cases}$$

Lemma 3.7. *Let $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p$ be a character of the form*

$$\begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{i \in \mathbb{Z}/s\mathbb{Z}} \tau_i(d)^{k_i-2}.$$

Then the irreducible constituents of the $\overline{\mathbb{F}}_p$ -representation $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$ are the following, where S runs over $\mathcal{P}(I)$:

$$\left(\bigotimes_{\tau \in S} \text{Sym}^{k_{\tau}-3+\nu_S(\tau)} k_{\mathfrak{p}}^2 \otimes_{\tau} \overline{\mathbb{F}}_p \right) \otimes \left(\bigotimes_{\tau \notin S} \det^{k_{\tau}-2+\nu_S(\tau)} \text{Sym}^{p+1-k_{\tau}-\nu_S(\tau)} k_{\mathfrak{p}}^2 \otimes_{\tau} \overline{\mathbb{F}}_p \right).$$

Here we have defined $\text{Sym}^{-1} k_{\mathfrak{p}}^2 = 0$.

This is Proposition 1.1 of [Dia]; the proof is a computation with Brauer characters. The unique irreducible submodule and quotient correspond to $S = I$ and $S = \emptyset$, respectively; since we do not use this fact in the paper, we do not prove it, but explicit maps can be found in [Sch2].

Remark 3.8. By twisting everything in Lemma 3.7 with characters $\prod_{\tau} \tau(\det^{w_{\tau}}) : \text{GL}_2(k_{\mathfrak{p}}) \rightarrow \overline{\mathbb{F}}_p^*$, we obtain the constituents of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$ for all characters $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p$. Observe that every $\sigma_{\mathfrak{p}}$ appears as a constituent of some θ , and generically of 2^s θ 's.

Definition 3.9. We will call $\sigma_{\mathfrak{p}}$ *untwisted* if it is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$ for some character $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ of the form

$$\theta : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{\tau : k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(d)^{k_{\tau}-2}.$$

We will also call characters θ of this form untwisted. Often we restrict to the case of untwisted $\sigma_{\mathfrak{p}}$, then use Corollary 3.21 to extend our results to the general case.

Let $\xi = \text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$. Since we are considering representations over $\overline{\mathbb{F}}_p$ of a group $(\text{GL}_2(k_{\mathfrak{p}}))$ whose order is divisible by p , Maschke's theorem fails, and $\sigma_{\mathfrak{p}}$ need not be a direct summand of ξ . However, there exists a composition series of submodules

$$0 = W_0 \subset W_1 \subset W_2 \subset \cdots \subset W_m = \xi$$

such that for some labeling $\{\sigma_1, \dots, \sigma_m\}$ of the irreducible constituents of ξ , we have $W_i/W_{i-1} \simeq \sigma_i$.

Definition 3.10. We call a Hecke eigenclass $f \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_{\sigma})$ *non-Eisenstein* if the corresponding maximal ideal $\mathfrak{m} \subset \mathbf{T}$ is non-Eisenstein.

Lemma 3.11. *Let $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$ be a weight. (Here $\sigma^{\mathfrak{p}}$ consists of the components of σ away from \mathfrak{p} .) Let $f \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_{\sigma})$ be a non-Eisenstein Hecke eigenclass. There exists an eigenclass $f' \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_{\xi \otimes \sigma^{\mathfrak{p}}})$ with the same system of eigenvalues.*

Proof. First of all, we drop the $\sigma^{\mathfrak{p}}$ from the notation; this will lead to no confusion. Then $f \in H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_\sigma)_{\mathfrak{m}} \neq 0$. Recall that $\{W_i\}$ is our filtration on $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$. For some $1 \leq i \leq m$ we have an exact sequence of representations

$$0 \rightarrow W_{i-1} \rightarrow W_i \rightarrow \sigma_{\mathfrak{p}} \rightarrow 0$$

Consider the following segment of the long exact cohomology sequence:

$$H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{W_i})_{\mathfrak{m}} \rightarrow H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_\sigma)_{\mathfrak{m}} \rightarrow H^2(M_{0,H} \otimes \bar{F}, \mathcal{F}_{W_{i-1}})_{\mathfrak{m}}.$$

By the remark after Proposition 3.5 the rightmost term is trivial. Since everything here clearly respects the action of \mathbf{T} , we can lift f to an element of $H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{W_i})_{\mathfrak{m}}$ with the same eigenvalues. Now we consider the sequence

$$0 \rightarrow W_i \rightarrow W_{i+1} \rightarrow W_{i+1}/W_i \rightarrow 0$$

and see that $H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{W_i})_{\mathfrak{m}}$ embeds in $H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{W_{i+1}})_{\mathfrak{m}}$. Iterating this argument, we see that it embeds in $H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_\xi)_{\mathfrak{m}}$. \square

For any place $v|p$, let $\mathbf{1}_v$ be the trivial representation of $\text{GL}_2(k_v)$, and let $\mathbf{1} = \bigotimes_{v|p} \mathbf{1}_v$ be the trivial weight; this corresponds to $k_\tau = 2$ and $w_\tau = 0$ for all τ .

Lemma 3.12. *Suppose that $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$ with $\sigma_{\mathfrak{p}}$ untwisted, and let $f \in H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_\sigma)$ be a non-Eisenstein eigenclass. There exists an eigenclass $f \in H^1(M_{U_1(\mathfrak{p}),H} \otimes \bar{F}, \mathcal{F}_{\mathbf{1}_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}})$ with the same system of eigenvalues.*

Proof. By Lemma 3.11 we can lift f to $f' \in H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{\xi \otimes \sigma^{\mathfrak{p}}})$. Let $B = \prod_{i=2}^r \text{GL}_2(\mathcal{O}_{\mathfrak{p}_i})$, and recall that $\xi \otimes \sigma^{\mathfrak{p}} = \text{Ind}_{\Gamma_0(\mathfrak{p})B}^{\text{GL}_2(k_{\mathfrak{p}})B} (\theta \otimes \sigma^{\mathfrak{p}})$. Consider the finite étale Galois cover $\pi : M_{U_0(\mathfrak{p}),H} \rightarrow M_{0,H}$; it is then a general fact from étale cohomology that for any lisse étale sheaf \mathcal{F} , we have $H^i(M_{0,H} \otimes \bar{F}, \pi_* \mathcal{F}) = H^i(M_{U_0(\mathfrak{p}),H} \otimes \bar{F}, \mathcal{F})$. Hence $H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_{\xi \otimes \sigma^{\mathfrak{p}}}) = H^1(M_{0,H} \otimes \bar{F}, \pi_* \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}}) = H^1(M_{U_0(\mathfrak{p}),H} \otimes \bar{F}, \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}})$. Statements of this form should be considered analogues of Shapiro's lemma in group cohomology.

Now observe that $\eta : M_{U_1(\mathfrak{p}),H} \rightarrow M_{U_0(\mathfrak{p}),H}$ is a Galois covering with group $G_{1,0} \simeq k_{\mathfrak{p}}^*$. Indeed, since $\nu(U_1(\mathfrak{p})) = \nu(U_0(\mathfrak{p}))$ we see that $M_{U_1(\mathfrak{p}),H}$ and $M_{U_0(\mathfrak{p}),H}$ have the same number of connected components, and each component of $M_{U_1(\mathfrak{p}),H}$ is a Galois cover of the corresponding component of $M_{U_0(\mathfrak{p}),H}$. It is a standard result that $H^1(M_{U_0(\mathfrak{p}),H} \otimes \bar{F}, \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}}) \simeq H^1(M_{U_1(\mathfrak{p}),H} \otimes \bar{F}, \eta^* \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}})^{G_{1,0}}$.

Finally we observe that $\eta^* \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}} = \mathcal{F}_{\theta|_{\Gamma_1(\mathfrak{p})} \otimes \sigma^{\mathfrak{p}}}$. Indeed, we have $\mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}} \simeq \eta_* \eta^* \mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}}^{G_{1,0}}$. Now $\eta_* \mathcal{F}_{\theta|_{\Gamma_1(\mathfrak{p})} \otimes \sigma^{\mathfrak{p}}} = \mathcal{F}_{\text{Ind}_{\Gamma_1(\mathfrak{p})}^{\Gamma_0(\mathfrak{p})} \theta|_{\Gamma_1(\mathfrak{p})} \otimes \sigma^{\mathfrak{p}}}$, and the $G_{1,0}$ -invariants of this sheaf clearly agree on stalks with $\mathcal{F}_{\theta \otimes \sigma^{\mathfrak{p}}}$. But we assumed θ to be untwisted, and so θ vanishes on $\Gamma_1(\mathfrak{p})$. Therefore, f indeed lifts to $H^1(M_{U_1(\mathfrak{p}),H} \otimes \bar{F}, \mathcal{F}_{\mathbf{1}_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}})$, and all of these manipulations are compatible with the Hecke action. \square

Remark 3.13. By exactly the same argument, we see that for arbitrary $\sigma_{\mathfrak{p}}$, $f \in H^1(M_{0,H} \otimes \bar{F}, \mathcal{F}_\sigma)$ can be lifted to an eigenform in $H^1(M_{U_1(\mathfrak{p}),H}^{\text{bal}} \otimes \bar{F}, \mathcal{F}_{\mathbf{1}_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}})$. This time, $\sigma_{\mathfrak{p}}$ is a constituent of θ of the form

$$\theta : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{\tau} \tau(ad)^{w_\tau} \tau(d)^{k_\tau - 2}$$

and one has to go up to $\Gamma_1^{\text{bal}}(\mathfrak{p})$ for θ to vanish.

Since \mathcal{F}_1 is the constant sheaf $\overline{\mathbb{F}}_p$, we can consider all the places above p at once in this argument to obtain the following:

Proposition 3.14. *Let σ be a weight such that for each $v|p$, the v -component σ_v is untwisted. Then a non-Eisenstein eigenclass $f \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$ can be lifted to an eigenclass $f \in H^1(M_{U_1(p),H} \otimes \overline{F}, \overline{\mathbb{F}}_p)$ with the same eigenvalues. For arbitrary σ , f can be lifted to $H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, \overline{\mathbb{F}}_p)$.*

3.5. Action of the diamond operators. While the diamond operators act trivially on $M_{U_0(p),H}$, this is not true on $M_{U_1(p),H}$ and $M_{U_1(p),H}^{bal}$.

Lemma 3.15. *Let $f \in H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, \mathcal{F}_{\mathbf{1}_p \otimes \sigma^p})$ be a lift of a class from $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$ as in the remark above. Recall that we have a fixed character $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ such that $\sigma_{\mathfrak{p}}$ is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$. Let $\alpha, \beta \in k_{\mathfrak{p}}^*$. Then*

$$\langle \alpha, \beta \rangle f = \theta^{-1} \left(\begin{pmatrix} \alpha & 0 \\ 0 & \beta \end{pmatrix} \right) f.$$

Proof. For some liftings $\tilde{\alpha}, \tilde{\beta} \in \mathcal{O}_{\mathfrak{p}}^*$ of α, β , let $\gamma \in G(\mathbb{A}^\infty)$ be the element such that $\gamma_v = 1$ for all $v \neq \mathfrak{p}$, and

$$\gamma_{\mathfrak{p}} = \begin{pmatrix} \tilde{\alpha} & 0 \\ 0 & \tilde{\beta} \end{pmatrix}.$$

From the description of the Hecke operators in section 3.2 we see that the operator $\langle \alpha, \beta \rangle$ on cohomology is induced by a map of sheaves $\mathcal{F}_{\overline{\mathbb{F}}_p \otimes \sigma^p} \rightarrow \gamma_* \mathcal{F}_{\overline{\mathbb{F}}_p \otimes \sigma^p}$ which sends a section $f \in \mathcal{F}_{\overline{\mathbb{F}}_p \otimes \sigma^p}(Y)$ to a section F_f , where $F_f(C) = \gamma f(C\gamma)$.

Since f comes from a section of the sheaf $\mathcal{F}_{\theta \otimes \sigma^p}$ on $M_{U_0(p),H}$, we see that $f(C\gamma) = \theta^{-1}(\gamma) f(C)$. On the other hand, $\text{GL}_2(F_{\mathfrak{p}})$ acts trivially on the underlying space of $\mathbf{1}_p \otimes \sigma^p$. Thus the map above is simply multiplication by $\theta^{-1}(\gamma) = \prod_{\tau} \tau(\alpha\beta)^{-w_{\tau}} \tau(\beta)^{2-k_{\tau}}$. \square

3.6. Galois representations.

Theorem 3.16. *Let $f \in H^1(M_{0,H} \times \overline{F}, \mathcal{F}_\sigma)$ be a Hilbert modular class which is an eigenvector of the Hecke algebra \mathbf{T} , and suppose the associated maximal ideal $\mathfrak{m} \subset \mathbf{T}$ is non-Eisenstein. Let $a_v \in \overline{\mathbb{F}}_p$ and $b_v \in \overline{\mathbb{F}}_p$ denote the eigenvalues of T_v and S_v , respectively, for places v where these operators exist. Then there is a representation $\overline{\rho}_f : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ such that if v is a place of F such that $H_v = \text{GL}_2(\mathcal{O}_v)$ (note that this criterion implies $v \nmid p$ and is satisfied by almost all v), then $\overline{\rho}_f$ is unramified at v , and*

- $\text{tr} \rho_f(\text{Frob}_v) = a_v$
- $\det \rho_f(\text{Frob}_v) = Nv \cdot b_v$.

Proof. We only sketch the proof. In section 3.4 we showed that f can be lifted to a Hilbert modular eigenclass $\tilde{f} \in H^1(M_{U_1(p),H}^{bal} \times \overline{F}, \overline{\mathbb{F}}_p)$ with the same system of eigenvalues. Since the constant sheaf $\overline{\mathbb{F}}_p$ corresponds to (the restriction to $\prod_{\mathfrak{p}|p} \Gamma_1^{bal}(\mathfrak{p})$ of) the trivial one-dimensional weight, our claim is analogous to the classical result that a modular form of weight k and level N can be lifted to a form of weight 2 and level Np ; see, for instance, [Gro] §9.

Since $H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, \overline{\mathbb{F}}_p) = H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, D) \otimes_D \overline{\mathbb{F}}_p$, by the Deligne-Serre lemma (Lemme 6.11 of [DS]) there exists $f' \in H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, D')$, for some discrete valuation ring $D' \supset D$, such

that the Hecke eigenvalues of \tilde{f} and f' have the same image in the residue field $\overline{\mathbb{F}}_p$. Clearly f' can be viewed as an element of $H^1(M_{U_1(p),H}^{bal} \otimes \overline{F}, \overline{\mathbb{Q}}_p)$.

Recall we have a fixed isomorphism $\overline{\mathbb{Q}}_p \xrightarrow{\sim} \mathbb{C}$. In Hida's notation, the constant sheaf $\overline{\mathbb{Q}}_p$ is the p -adic sheaf $\mathcal{L}(n, v, \mathbb{C})$ for $n = v = 0$ (which corresponds to $k = 2t$ and $w = t$, where t is the vector $(1, \dots, 1)$). By [Hid], Theorem 6.2, we have

$$H^1(M_{U_1(\mathfrak{p}),H}^{bal} \otimes \overline{F}, \overline{\mathbb{Q}}_{\mathfrak{p}}) \simeq S_{2t,t,\emptyset}(U_1(\mathfrak{p})^{bal} \times H, \mathbb{C}) \oplus S_{2t,t,I}(U_1(\mathfrak{p})^{bal} \times H, \mathbb{C}).$$

Here the two components on the right hand side are spaces of functions $f : G(\mathbb{Q}) \backslash G(\mathbb{A}) \rightarrow \mathbb{C}$ satisfying certain congruence properties, as defined on pp. 303-4 of [Hid]. Now by the Jacquet-Langlands theorem, there is a classical Hilbert eigenform of parallel weight 2 with the same eigenvalues as our eigenform $f \in S_{2t,t,I}(U_1(\mathfrak{p})^{bal} \times H, \mathbb{C})$, and for these one constructs Galois representations with the required properties as in [Car2] and [Tay]. \square

3.7. Modularity. We will now see that modularity in the sense of Definition 1.2 implies modularity in the intuitive sense that for a modular ρ there exists some mod p Hilbert modular eigenclass f such that $\rho \simeq \overline{\rho}_f$.

First we define a ‘‘diamond operator’’ map as follows. Map $(\mathbb{A}_F^\infty)^*$ to $(\mathbf{T}/\mathfrak{m})^*$ by sending an adèle a to $[U\gamma_a U]$. Here $\gamma_a \in G(\mathbb{A}^\infty)$ is defined as follows: for any place v of F , set $(\gamma_a)_v = a_v 1_v$, where 1_v is the identity of $(D \otimes F_v)^*$. Observe that if $a \in F^*$, then $\gamma_a \in G(\mathbb{Q})$ and γ_a is in the center of $G(\mathbb{A}^\infty)$, so a is in the kernel of our map. Composing with the Artin reciprocity map we obtain a ‘‘diamond operator’’ map:

$$\begin{aligned} \langle \cdot \rangle : \text{Gal}(\overline{F}/F) &\rightarrow (\mathbf{T}/\mathfrak{m})^* \\ \text{Frob}_v &\mapsto S_v^{-1} \end{aligned}$$

Proposition 3.17. *Suppose that $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ is irreducible, continuous, totally odd, and modular of weight σ (in the sense of Definition 1.2). Then there exists $f \in H^1(M_{0,H} \times \overline{F}, \mathcal{F}_\sigma)$ such that $\rho \simeq \overline{\rho}_f$.*

Proof. Consider the $\overline{\mathbb{F}}_p$ -vector space $V = (\text{Pic}^0(M_{1,H})[p] \otimes \sigma)^{\text{GL}_2(\mathcal{O}_F/p)}$, as a representation of the group $\text{Gal}(\overline{F}/F)$. We know that, for an appropriate H , ρ is a subrepresentation. Furthermore, V has an action of \mathbf{T} that commutes with the $\text{Gal}(\overline{F}/F)$ -action. This makes V into a \mathbf{T}/p -module. Since \mathbf{T} is finite free over \mathbb{Z} , we see that $\mathbf{T}/p = \bigoplus_{\mathfrak{m}} \mathbf{T}/\mathfrak{m}^{n_{\mathfrak{m}}}$, where \mathfrak{m} ranges over all maximal ideals of \mathbf{T} containing (p) and $n_{\mathfrak{m}} \in \mathbb{Z}$. Let $V_{\mathfrak{m}} = V \otimes_{\mathbf{T}} \mathbf{T}/\mathfrak{m}$. Then $V = \bigoplus_{\mathfrak{m}} V_{\mathfrak{m}}$.

Since this decomposition of V respects the $\text{Gal}(\overline{F}/F)$ -module structure and ρ is irreducible, we find that ρ appears in $V_{\mathfrak{m}}$ for some maximal ideal $\mathfrak{m} \subset \mathbf{T}$. In particular, $V_{\mathfrak{m}} \neq 0$. Also, there is a canonical isomorphism $\text{Pic}^0(M_{1,H})[p^\infty] \simeq H_{\text{ét}}^1(M_{1,H}, \mathbb{Z}_p(1))$ which is compatible with the Galois and Hecke actions. Taking p -torsion, we see that $\text{Pic}^0(M_{1,H})[p] \simeq H^1(M_{1,H} \otimes \overline{F}, \boldsymbol{\mu}_p)$. Therefore, we have a Hecke-equivariant (but not Galois-equivariant!) isomorphism:

$$V_{\mathfrak{m}} \simeq (H^1(M_{1,H} \otimes \overline{F}, \boldsymbol{\mu}_p) \otimes \sigma)_{\mathfrak{m}}^{\text{GL}_2(\mathcal{O}_F/p)} \simeq H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)_{\mathfrak{m}}.$$

Indeed, we will show that $(H^0(M_{1,H} \otimes \overline{F}, \boldsymbol{\mu}_p) \otimes \sigma)^{\text{GL}_2(\mathcal{O}_F/p)} \simeq H^0(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$, and by the Hochschild-Serre spectral sequence we will get a similar map for H^1 with Eisenstein kernel and cokernel. Now, $(H^0(M_{1,H} \otimes \overline{F}, \boldsymbol{\mu}_p) \otimes \sigma)$ consists of locally constant functions $f : (M_{1,H})_{\overline{F}} \rightarrow W_\sigma$, and $g \in \text{GL}_2(\mathcal{O}_F/p)$ acts by $(gf)(C) = \sigma(g)(f(Cg))$. Hence the $\text{GL}_2(\mathcal{O}_F/p)$ -invariants are the functions satisfying $f(Cg) = \sigma^{-1}(g)f(C)$, which are precisely the global sections of \mathcal{F}_σ on $M_{0,H}$.

By the Eichler-Shimura relation of Proposition 2.6, for almost all places v of F we have the following congruence on V and $V_{\mathfrak{m}}$:

$$(6) \quad \text{Frob}_v^2 - T_v \text{Frob}_v + (Nv)S_v \equiv 0 \pmod{\mathfrak{m}}.$$

Let $\alpha, \beta \in \overline{\mathbb{F}}_p$ be the roots of the polynomial $X^2 - T_v X + (Nv)S_v$, where we consider the coefficients as elements of $\mathbf{T}/\mathfrak{m} \subset \overline{\mathbb{F}}_p$. Consider a Jordan-Hölder decomposition of the $(\mathbf{T}/\mathfrak{m})[\text{Gal}(\overline{F}/F)]$ -module $V_{\mathfrak{m}}$. Let W be a constituent corresponding to ρ . Let $\lambda_1, \lambda_2, \dots, \lambda_{\dim W}$ be the roots of the characteristic polynomial of Frob_v on W . Clearly each λ_i is either α or β .

Let $\tilde{\chi} = \chi \langle \cdot \rangle^{-1}$, where χ is the mod p cyclotomic character. Observe that $\tilde{\chi}(\text{Frob}_v) = (Nv)S_v = \alpha\beta$ for v satisfying (6). Hence, the roots of the characteristic polynomial of Frob_v acting on the $(\mathbf{T}/\mathfrak{m})[\text{Gal}(\overline{F}/F)]$ -module $W^\vee \otimes \tilde{\chi}$ are precisely $\lambda_i^{-1}\alpha\beta$; note that this quantity is equal to β (resp. α) if $\lambda_i = \alpha$ (resp. $\lambda_i = \beta$). Finally we see that the characteristic polynomial of Frob_v acting on $W \oplus W^\vee \tilde{\chi}$ is $(X^2 - T_v X + (Nv)S_v)^{\dim W}$. Hence $\rho_{\mathfrak{m}}$ exists, and by the Brauer-Nesbitt theorem:

$$W \oplus W^\vee \tilde{\chi} \simeq (\rho_{\mathfrak{m}})^{\dim W}.$$

Since ρ is irreducible, we have that $\rho \simeq W$ and $\dim W = 2$. Hence $\rho \simeq \bar{\rho}_{\mathfrak{m}}$. In particular, $\bar{\rho}_{\mathfrak{m}}$ is irreducible, so by Proposition 3.4 \mathfrak{m} is non-Eisenstein. Let $0 \neq f \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)_{\mathfrak{m}}$; this space is non-zero since ρ is a constituent. Then by Theorem 3.16 $\bar{\rho}_f$ exists and is isomorphic to ρ . \square

Remark 3.18. For $f \in H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$, observe that $\bar{\rho}_f$ does *not* necessarily appear in the $\text{Gal}(\overline{F}/F)$ -representation $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$. Indeed, from the characteristic polynomial of Frob_v on $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$ in Remark 2.5 we find that $\bar{\rho}_f \otimes \chi$ appears in $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma)$, where χ is the cyclotomic character.

Proposition 3.19. *Let $\langle \cdot \rangle : \text{Gal}(\overline{F}/F) \rightarrow \overline{\mathbb{F}}_p^*$ be the map defined above. Then $(\det \rho) = \chi \langle \cdot \rangle^{-1}$.*

Proof. It suffices to verify the relation for uniformizers π_v at almost all finite places v of F ; note that the choice of π_v corresponds via the reciprocity map to (the inverse of) a choice of $\text{Frob}_v \in \text{Gal}(\overline{F}/F)$. But $(\det \rho)(\pi_v^{-1}) = \det \rho(\text{Frob}_v) = (Nv)S_v = \langle \pi_v \rangle \chi^{-1}(\pi_v)$. \square

3.8. Twisting by the determinant. Let $\eta = \bigotimes_{v|p} \left(\bigotimes_{\tau: k_v \hookrightarrow \overline{\mathbb{F}}_p} \det^{w_\tau} \otimes_\tau \overline{\mathbb{F}}_p \right)$ be a one-dimensional weight. We will show that twisting a weight by η simply twists the possible modular representations by a certain character. This will allow us to consider weights “up to twist” in section 5.

Let $\chi'_\eta : \prod_{v|p} k_v^* \rightarrow \overline{\mathbb{F}}_p^*$ be the character given by $\chi'_\eta|_{k_v} = \prod_{\tau: k_v \hookrightarrow \overline{\mathbb{F}}_p} \tau^{-w_\tau}$. Then χ'_η lifts to a character on $\prod_v \mathcal{O}_v^*$ and hence to one on \mathcal{O}_F^* . By Théorème 1 of [Che], there is an $a \in \mathbb{Z}$ such that $(p, a) = 1$ and $\mathcal{O}_F^* \cap (1 + a\mathcal{O}_F)$ lies in the kernel of χ'_η . Let $\mathcal{O}_{F,+}^*$ be the subgroup of totally positive elements of \mathcal{O}_F^* . The character $(\chi'_\eta)^{-1} : \mathcal{O}_{F,+}^* / (\mathcal{O}_{F,+}^* \cap (1 + a\mathcal{O}_F)) \rightarrow \overline{\mathbb{F}}_p^*$ extends to $\prod_{v|a} \mathcal{O}_v^* / (1 + \pi_v^{\text{ord}_v(a)} \mathcal{O}_v)$, and this in turn lifts to a character $\chi''_\eta : \prod_{v|a} \mathcal{O}_v^* \rightarrow \overline{\mathbb{F}}_p^*$. For $x \in \hat{\mathcal{O}}_F^*$, let $x_a = (x_v)_{v|a} \in \prod_{v|a} \mathcal{O}_v^*$. We now define a character $\chi_\eta : \hat{\mathcal{O}}_F^* / \overline{\mathcal{O}_{F,+}^*} \rightarrow \overline{\mathbb{F}}_p^*$, by $\chi_\eta(x) = \chi'_\eta(x_p) \chi''_\eta(x_a)$. Clearly χ_η is trivial on $\mathcal{O}_{F,+}^*$ and coincides with χ'_η on $\prod_{v|p} \mathcal{O}_v^*$.

Finally, observe that $\hat{\mathcal{O}}_F^* / \overline{\mathcal{O}_{F,+}^*} \subset \mathbb{A}_F^* / \overline{F^*(F_\infty^*)^0}$ is dense. Then χ_η extends to this larger group, and we view it via the reciprocity map as a character $\chi_\eta : \text{Gal}(\overline{F}/F) \rightarrow \overline{\mathbb{F}}_p^*$.

Let $V = K(0, H) \cap Z(\mathbb{Q})$. Decreasing H^p if necessary at the places dividing a , we may assume that $\nu(V) \subset \ker \chi_\eta$. Now $H^0(M_{0,H} \otimes \overline{F}, \mathcal{F}_\eta) = \{f : (\mathbb{A}_F^\infty)^* / F^* \nu(V) \rightarrow \Lambda/p : f(C\gamma) = \eta(\gamma)^{-1} f(C)\}$.

Thus $\chi_\eta \in H^0(M_{0,H} \otimes \overline{F}, \mathcal{F}_\eta)$. Furthermore, the cup product $f \mapsto f \cup \chi_\eta$ provides a map $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_\sigma) \rightarrow H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_{\sigma \otimes \eta})$.

Proposition 3.20. *If f is a Hecke eigenform corresponding to the Galois representation ρ , then $f \cup \chi_\eta$ is also a Hecke eigenform, and the corresponding Galois representation is $\rho \otimes \chi_\eta^{-1}$.*

Proof. We study the action of \mathbf{T} on χ_η . Suppose $v \notin S$ and let $\gamma_v \in G(\mathbb{A}^\infty)$ be the element with $T_v = [K(1, H)\gamma_v K(1, H)]$. As in (4), T_v acts on $H^1(M_{0,H} \otimes \overline{F}, \mathcal{F}_{\sigma \otimes \eta})$ via the composition of three maps. The middle map, or rather the corresponding map of sheaves $\mathcal{F}_{\sigma \otimes \eta} \rightarrow (\gamma_v)_* \mathcal{F}_{\sigma \otimes \eta}$ on $M_{0,H}$, sends $f \cup \chi_\eta$ to $F_{f \cup \chi_\eta}$, where $F_{f \cup \chi_\eta}(C) = \gamma_v f(C\gamma_v) \otimes \chi_\eta(C\gamma_v) = \gamma_v f(C\gamma_v) \otimes \chi_\eta(\pi_v)\chi_\eta(C)$. The third map is the trace, and since we can find a set of coset representatives of $\gamma_v K(1, H)\gamma_v^{-1}$ in $K(1, H)$ that lie in the kernel of the reduced norm ν , we see that $T_v(f \cup \chi_\eta) = T_v f \cup \chi_\eta(\pi_v)\chi_\eta = \lambda(T_v)\chi_\eta(\pi_v)(f \cup \chi_\eta)$. Similarly, $S_v(f \cup \chi_\eta) = \lambda(S_v)\chi_\eta(\pi_v)^2(f \cup \chi_\eta)$.

Therefore, $f \cup \chi_\eta$ is a Hecke eigenform. If ρ' is the Galois representation corresponding to $f \cup \chi_\eta$, then for all $v \notin S$,

$$\begin{aligned} \mathrm{tr} \rho'(\mathrm{Frob}_v) &= \lambda(T_v)\chi_\eta(\pi_v) = \lambda(T_v)\chi_\eta^{-1}(\pi_v^{-1}) = \mathrm{tr}(\rho \otimes \chi_\eta^{-1})(\mathrm{Frob}_v) \\ \det \rho'(\mathrm{Frob}_v) &= \lambda(S_v)\chi_\eta(\pi_v)^2 = \lambda(S_v)\chi_\eta^{-1}(\pi_v^{-1})^2 = \det(\rho \otimes \chi_\eta^{-1})(\mathrm{Frob}_v). \end{aligned}$$

Thus $\rho' = \rho \otimes \chi_\eta^{-1}$. □

Recall that for a Galois representation ρ , we denoted by $W(\rho)$ the set of modular weights.

Corollary 3.21. $W(\rho \otimes \chi_\eta^{-1}) = \{\sigma \otimes \eta : \sigma \in W(\rho)\}$.

4. FINITE FLAT GROUP SCHEMES

4.1. Review of vector space schemes. In this section we collect, without proof, facts about vector space schemes from [Ray] and section 5 of [Edi1]. Note that [Edi1] mostly treats the case $e = p - 1$, but the same arguments work to give the results stated here.

Let \mathbb{F} be a finite field with $w = p^r$ elements, where p is prime and $r \geq 1$. An \mathbb{F} -vector space scheme over a scheme S is a commutative group scheme G over S equipped with a homomorphism $\mathbb{F} \rightarrow \mathrm{End}_S(G)$.

Let D be a mixed characteristic strictly henselian discrete valuation ring with residue field k of characteristic p . Recall that *strictly* henselian means that k is separably closed. We will denote its field of fractions by K . Let $v : D \rightarrow \mathbb{Z}$ be the valuation and $e = v(p)$. Suppose that the base scheme S is a D -scheme.

Since D contains the $(w-1)$ th roots of unity, there are $w-1$ characters $\chi : \mathbb{F}^* \rightarrow \mu_{w-1}(D) \subset D^*$. We write \mathcal{I} for the augmentation ideal sheaf of G : $\mathcal{O}_G = \mathcal{O}_S \oplus \mathcal{I}$. For each character $\chi : \mathbb{F}^* \rightarrow D^*$ as above, we define the subsheaf

$$\mathcal{I}_\chi = \{f \in \mathcal{I} : \forall a \in \mathbb{F}^*, a^*(f) = \chi(a)f\}$$

Here a^* is the endomorphism of \mathcal{O}_G corresponding to scalar multiplication by a on G . It is easy to see that $\mathcal{I} = \bigoplus_\chi \mathcal{I}_\chi$.

Definition 4.1. We will say, following [Ray], that the \mathbb{F} -vector space scheme G/S satisfies condition $(**)$ if all $w-1$ sheaves \mathcal{I}_χ are invertible.

In the sequel we consider only schemes satisfying condition (**). If S is an integral scheme with fraction field of characteristic 0 and G/S is finite and locally free of rank w , then by [Ray], 1.2.2. G satisfies (**).

A character $\chi : \mathbb{F}^* \rightarrow \mu_{w-1}(D) \subset D^*$ is called fundamental if, when we define $\chi(0) = 0$, the composition $\mathbb{F} \xrightarrow{\chi} D \rightarrow k$ is a ring homomorphism. Clearly there are r fundamental characters. Pick one of them and call it χ_0 . Then all of them are $\chi_i = \chi_0^{p^i}$, where $i \in \mathbb{Z}/r\mathbb{Z}$.

Let G/K be any \mathbb{F} -vector space scheme satisfying condition (**). For all i let X_i be a generator of \mathcal{I}_{χ_i} . Since any character $\chi : \mathbb{F}^* \rightarrow D^*$ has the form $\chi = \chi_0^{\alpha_0} \chi_1^{\alpha_1} \cdots \chi_{r-1}^{\alpha_{r-1}}$ for $0 \leq \alpha_i \leq p-1$, it is easy to see that G is determined by r equations, one for each $i \in \mathbb{Z}/r\mathbb{Z}$:

$$(7) \quad X_i^p = \delta_i X_{i+1}, \delta_i \in K^*.$$

The map $(\delta_0, \delta_1, \dots, \delta_{r-1}) \mapsto (v(\delta_0), v(\delta_1), \dots, v(\delta_{r-1}))$ gives a bijection between isomorphism classes of such \mathbb{F} -vector space schemes over K and the set \mathbb{Z}^r/Λ , where

$$\Lambda = \mathbb{Z}(p, 0, \dots, 0, -1) \oplus \mathbb{Z}(-1, p, 0, 0, \dots, 0) \oplus \cdots \oplus \mathbb{Z}(0, \dots, 0, -1, p).$$

Observe that the set

$$\{(a_0, \dots, a_{r-1}) \in \mathbb{Z}^r : 0 \leq a_i \leq p-1, a_i \text{ not all } p-1\}$$

therefore corresponds bijectively with the isomorphism classes of G/K .

Similarly, an \mathbb{F} -vector space scheme G/D satisfying the condition (**) is given by equations of the form (7), where this time $\delta_i \in D$ and $0 \leq v(\delta_i) \leq e$. The second condition is necessary to define the group structure on G , i.e. to ensure that the inversion map is defined over D . The map $(\delta_0, \delta_1, \dots, \delta_{r-1}) \mapsto (v(\delta_0), v(\delta_1), \dots, v(\delta_{r-1}))$ is a bijection between the set of isomorphism classes of G/D and the set

$$\{(a_0, \dots, a_{r-1}) \in \mathbb{Z}^r \mid 0 \leq a_i \leq e\}.$$

If $e \geq p-1$, this tells us that every \mathbb{F} -vector space scheme G/K satisfying the condition (**) can be extended to D . However, if $e > p-1$ such extensions are not unique.

Now we specialize to our usual case of $D = W(k_p)$ and $K = F_{\mathfrak{p}}^{nr}$. Then D is a complete discrete valuation ring with $e = 1$ (since p was unramified in F) and residue field $\overline{\mathbb{F}}_p$. Let D' and K' be as defined in the introduction, let v' be their valuation, and observe that D' is a strictly henselian discrete valuation ring with $e' = p^s - 1$.

Let G/K be an \mathbb{F} -vector space scheme given by equations $X_i^p = \delta_i X_{i+1}$, and set $a_i = v(\delta_i)$. We can assume that $0 \leq a_i \leq p-1$. The pull-back $G_{K'} = G \times_K \text{Spec } K'$ is described by the same equations, with the δ_i now considered as elements of $(K')^*$. We have $v'(\delta_i) = (p^s - 1)v(\delta_i) = (p^s - 1)a_i$ for all i . Consider an extension $G_{D'}$ of $G_{K'}$ to D' ; as we remarked above, it is not unique if $s > 1$. If $G_{D'}$ is given by the equations $(X'_i)^p = \delta'_i X'_{i+1}$, set $a'_i = v'(\delta'_i)$. We have $0 \leq a'_i \leq p^s - 1$.

Now we study two Galois actions. First consider the action of $\text{Gal}(\overline{K}/K)$ on $G(\overline{K})$; it determines G as a finite group scheme by the equivalence of categories in [Wat], §6.3-4. Define the character $\psi_i : \mu_{w-1}(K) \rightarrow \mathbb{F}^*$ as the inverse of χ_i , for all $i \in \mathbb{Z}/r\mathbb{Z}$. Upon eliminating the variables X_1, \dots, X_{r-1} from the equations (7), we find that

$$X_0^w = \delta X_0,$$

where $\delta = \delta_0^{p^{r-1}} \delta_1^{p^{r-2}} \cdots \delta_{r-1}$. Then $G(\overline{K})$ can be identified with the solutions in \overline{K} of the above equation. If $a \in \mathbb{F}^*$, then multiplication by a on $G(\overline{K})$ corresponds to multiplication by $\chi_0(a)$ on these solutions. For $x \in G(\overline{K})$, clearly $G(\overline{K}) = \{\zeta x : \zeta \in \mu_{w-1}(K)\}$, so every element of $\text{Gal}(\overline{K}/K)$

acts on $G(\overline{K})$ by multiplication by some $\zeta \in \mu_{w-1}(K)$, which corresponds as above to multiplication by some $a \in \mathbb{F}^*$. Thus the action of $\text{Gal}(\overline{K}/K)$ is given by a character $\text{Gal}(\overline{K}/K) \rightarrow \mathbb{F}^*$. Since the only element of \mathbb{F}^* of p -power order is the identity, the wild inertia acts trivially. As the tame ramification of \overline{K}/K is identified with $\varprojlim \mu_{p^n-1}(K)$, the action is in fact given by a character $\phi : \mu_{w-1}(K) \rightarrow \mathbb{F}^*$. Furthermore,

$$(8) \quad \phi = \psi_0^{a_{r-1}} \psi_1^{a_0} \cdots \psi_{r-2}^{a_{r-2}}.$$

Let $\sigma \in \text{Gal}(K'/K)$; it induces automorphisms σ both of $\text{Spec } K'$ and of $\text{Spec } D'$. Denote by $g(\sigma)$ the automorphism $id_G \times \sigma$ of $G \times_K K'$ by $g(\sigma)$; it lifts to $G_{D'}$. We can now restrict to the special fiber to get an action of $\text{Gal}(K'/K)$ on $G_{\overline{\mathbb{F}}_p} = G_{D'} \times_{D'} \overline{\mathbb{F}}_p$ and hence on $\text{cot}(\overline{\mathbb{F}}_p)$, the cotangent space to $G_{\overline{\mathbb{F}}_p}$ at 0. This space has the presentation

$$(9) \quad \text{cot}(G_{\overline{\mathbb{F}}_p}) = \bigoplus_{i \in \mathbb{Z}/r\mathbb{Z}, a'_{i-1} > 0} \overline{\mathbb{F}}_p \cdot dX'_i.$$

One can define invariants b_i for each $i \in \mathbb{Z}/r\mathbb{Z}$ such that $g(\sigma)$ acts on the generator dX'_i , when it exists, by multiplication by (the image in the residue field $\overline{\mathbb{F}}_p$ of) $i_s(\sigma)^{-b_i}$. Here if K_n , with valuation v_n , is the totally ramified extension of K of degree $p^n - 1$, then $i_n : \text{Gal}(K_n/K) \xrightarrow{\sim} \mu_{p^n-1}(D)$ is defined as in [Ray], §3.1 by requiring that for all $\sigma \in \text{Gal}(K_n/K)$ and $x \in K_n$ there exists $u \in K'$ with $v'(u) \geq 1$ such that $\sigma(x) = x(i_n(\sigma))^{v_n(x)}(1+u)$. This gives the isomorphism $I_{t,p} \xrightarrow{\sim} \varprojlim_n \mu_{p^n-1}(K)$. The invariants a_i and b_i are related by the following fundamental equation, for all $i \in \mathbb{Z}/r\mathbb{Z}$:

$$(10) \quad a'_i = b_{i+1} - pb_i + (p^s - 1)a_i.$$

4.2. Torsion points on $\text{Pic}^0(\mathbf{M}_{U_1(p)^{bal,H}})$. We are finally ready to begin the proof of Theorem 1.3. Let $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ be modular of weight σ and suppose that $\rho|_{G_p}$ is irreducible. *Suppose that σ is untwisted.* We will see that this assumption involves no essential loss of generality. It implies that there exists a map $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ of the form

$$(11) \quad \theta : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(d)^{k_{\tau}-2}$$

such that σ is a Jordan-Hölder constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$. *Suppose that θ is non-trivial*, i.e. that the k_{τ} are not all 2 and not all $p+1$. The case of $\theta = 1$, which is easier, will be treated separately in section 5.4.

By Proposition 3.17 there exists a non-Eisenstein maximal ideal $\mathfrak{m} \subset \mathbf{T}$ and a Hilbert modular eigenform $f \in H^1(\mathbf{M}_{0,H} \times \overline{F}, \mathcal{F}_{\sigma})_{\mathfrak{m}}$ such that $\rho \simeq \overline{\rho}_f$. Then by Proposition 3.14 there exists $\tilde{f} \in H^1(\mathbf{M}_{U_1(p),H}^{bal} \times \overline{F}, \overline{\mathbb{F}}_p)_{\mathfrak{m}}$ with the same Hecke eigenvalues as f . By Lemma 3.15, this \tilde{f} is an eigenvector for the Hecke algebra $\tilde{\mathbf{T}}$ defined in section 2.2 and determines a map $\lambda : \tilde{\mathbf{T}} \rightarrow \overline{\mathbb{F}}_p$, which sends each Hecke operator to the corresponding eigenvalue. In particular, for every $\alpha, \beta \in k_{\mathfrak{p}}^*$ we have $\lambda(\langle \alpha, \beta \rangle) = \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(\beta)^{2-k_{\tau}}$. Clearly, the kernel of λ is a maximal ideal $\tilde{\mathfrak{m}} \subset \tilde{\mathbf{T}}$ containing \mathfrak{m} , and λ induces an embedding of fields $\tilde{\mathbf{T}}/\tilde{\mathfrak{m}} \subset \overline{\mathbb{F}}_p$. Set $\mathbb{F} = \tilde{\mathbf{T}}/\tilde{\mathfrak{m}}$. Increasing \mathbb{F} if necessary, we may assume that $\text{GL}_2(\mathbb{F})$ contains the image of ρ .

Since $\tilde{\mathbf{T}} \subset \text{End}(\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal}))$, $\tilde{\mathbf{T}}$ is a free \mathbb{Z} -module of finite rank and \mathbb{F} is a finite field. Also, we have $\tilde{\mathbf{T}}_p = \tilde{\mathbf{T}} \otimes_{\mathbb{Z}} \mathbb{Z}_p = \tilde{\mathbf{T}}_{\tilde{\mathfrak{m}}} \times \tilde{\mathbf{T}}'$, where $\tilde{\mathbf{T}}'$ is the product of the localizations of $\tilde{\mathbf{T}}$ at all the other maximal ideals above p . Now let $\varepsilon, \varepsilon' \in \tilde{\mathbf{T}}_p$ be idempotents such that $\varepsilon \tilde{\mathbf{T}}_p = \tilde{\mathbf{T}}_{\tilde{\mathfrak{m}}}$, $\varepsilon' \tilde{\mathbf{T}}_p = \tilde{\mathbf{T}}'$, and $\varepsilon + \varepsilon' = 1$. Consider the element $\eta = \frac{1}{q-1} \sum_{\alpha \in k_{\mathfrak{p}}^*} \langle \alpha, \alpha \rangle \in \tilde{\mathbf{T}}_p$. Clearly $\eta^2 = \eta$ and $\eta' = 1 - \eta$ is also an idempotent with $\eta\eta' = 0$. Since the θ is non-trivial, we see that $\lambda(\eta) = 0$, so $\eta \in \tilde{\mathfrak{m}}$. But η is an idempotent, so in fact $\eta \in \tilde{\mathfrak{m}}^n$ for all $n \geq 1$, whence $\varepsilon\eta = 0$. Thus, $\varepsilon\eta' = \varepsilon(1 - \eta) = \varepsilon$.

Recall from section 2 that $M_{U_1(\mathfrak{p}),H}^{bal}$ has semistable reduction over the field F'_p , where F'_p/F_p is a totally ramified extension with Galois group $k_{\mathfrak{p}}^*$. Note that $K' = F'_p F_p^{nr}$. Let J be the Néron model over D' of the Jacobian of $M_{U_1(\mathfrak{p}),H}^{bal}$, and denote its special fiber by $\bar{J} = J \times_{D'} \bar{\mathbb{F}}_p$. The identity component J^0 of J is isomorphic to $\text{Pic}^0(\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal}/D')$; see [BLR1], 9.5/4. Recall that we denoted the irreducible components of its special fiber by I and E . By [BLR1] 9.2/8, the pullbacks of invertible sheaves from $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \otimes_{D'} \bar{\mathbb{F}}_p$ to I and E give an exact sequence:

$$(12) \quad 0 \rightarrow T \rightarrow \bar{J}^0 \rightarrow J_I \oplus J_E \rightarrow 0$$

where J_I and J_E are the Jacobians of I and E , respectively, and T is a torus consisting of divisors of degree 0 supported on the (finitely many) supersingular points.

Now consider the p -divisible group $J^0[p^\infty] = (J^0[p^n])_n$ over D' . Observe that the $J^0[p^n]$ are quasi-finite and flat over D' . Since $J^0[p^\infty]$ is a $\tilde{\mathbf{T}} \otimes \mathbb{Z}_p$ -module, it decomposes as

$$J^0[p^\infty] = \varepsilon J^0[p^\infty] \oplus \varepsilon' J^0[p^\infty] = G \oplus G',$$

where we have defined the p -divisible groups G and G' by $G_n = \varepsilon J^0[p^n]$ and $G'_n = \varepsilon' J^0[p^n]$. Similarly there is a decomposition $J^0[p^\infty] = \eta J^0[p^\infty] \oplus \eta' J^0[p^\infty]$.

Observe that $\eta G = \eta \varepsilon J^0[p^\infty] = 0$. Recall that a geometric point of $\mathbf{M}_{U_1(\mathfrak{p}),H}^{bal} \otimes \bar{\mathbb{F}}_p$ has the form (x, P, P') , where x is the geometric point of $\mathbf{M}_{0,H} \otimes \bar{\mathbb{F}}_p$ lying below and (P, P') is a $U_1(\mathfrak{p})^{bal}$ -structure at x . It is easy to see from [Car1], 7.5 that the diamond operator $\langle \alpha, \beta \rangle$ sends (x, P, P') to $(x, \beta^{-1}P, \alpha^{-1}P')$. Indeed, the group $\text{GL}_2(\mathcal{O}_{\mathfrak{p}})$ acts on the right on Drinfeld bases of level 1 via its quotient $\text{GL}_2(k_{\mathfrak{p}})$ as follows: $(\varphi g)((x, y)) = \varphi((x, y)g^{-1})$, for $g \in \text{GL}_2(k_{\mathfrak{p}})$. In particular, $\langle \alpha, \beta \rangle$ stabilizes the supersingular points.

Let $J^0[p^\infty]^f = (J^0[p^n]^f)_n$ be the finite part of $J^0[p^\infty]$, i.e. the maximal subscheme which is finite over D' ; see [SGA], IX, 5.1. Since the diamond operators $\langle \alpha, \beta \rangle$, and hence η , act trivially on the torus $T \subset \bar{J}$, it follows by [SGA], IX 5.5.8 that they act trivially on the quotient $J^0[p^\infty]/J^0[p^\infty]^f$. But $\eta G = 0$, whence $G \subset J^0[p^\infty]^f$ and all the G_n are finite over D' .

Now we consider special fibers. Let $\bar{G} = G \times_{D'} \bar{\mathbb{F}}_p = \varepsilon \bar{J}^0[p^\infty]$. Taking p^∞ -torsion of the exact sequence (12) yields the following exact sequence of p -divisible groups over $\bar{\mathbb{F}}_p$:

$$0 \rightarrow T[p^\infty] \rightarrow \bar{J}^0[p^\infty] \rightarrow J_I[p^\infty] \oplus J_E[p^\infty] \rightarrow 0.$$

All these maps are compatible with the action of η and η' . Thus we get an exact sequence

$$0 \rightarrow \eta' T[p^\infty] \xrightarrow{\iota} \eta' \bar{J}^0[p^\infty] \xrightarrow{\pi} \eta' (J_I[p^\infty] \oplus J_E[p^\infty]) \rightarrow 0.$$

The only part of the exactness that is not obvious is that $\ker(\pi) \subset \text{im}(\iota)$. If $g \in \bar{J}^0[p^\infty]$ is such that $\pi(\eta'g) = 0$, then there is $t \in T[p^\infty]$ such that $\iota(t) = \eta'g$. But then $\iota(\eta't) = \eta'\iota(t) = (\eta')^2g = \eta'g$, proving the required exactness. Since η acts trivially on T , we have $\eta'T = 0$. Hence π is an

isomorphism. From the description above of the action of the $\langle \alpha, \beta \rangle$, we see that they preserve the decomposition $J_I \oplus J_E$, and hence so do η and η' . Thus we obtain an isomorphism

$$\eta' \bar{J}^0[p^\infty] \xrightarrow{\sim} \eta'(J_I[p^\infty] \oplus J_E[p^\infty]) = \eta' J_I[p^\infty] \oplus \eta' J_E[p^\infty].$$

Since $\varepsilon \eta' = \varepsilon$, we conclude that:

Lemma 4.2. *With the notations as above, we have a decomposition*

$$\bar{G} = \varepsilon \eta' \bar{J}^0[p^\infty] \simeq \varepsilon J_I[p^\infty] \oplus \varepsilon J_E[p^\infty] = \bar{G}_I \oplus \bar{G}_E.$$

Here the last equality is a definition; note that \bar{G}_I and \bar{G}_E are not defined as special fibers of group schemes over D' .

4.3. Relation between G and ρ . Consider the space \mathbb{F}^2 endowed with a $\text{Gal}(\bar{F}/F)$ -action given by $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\mathbb{F})$. This corresponds to a finite affine (étale) group scheme W over F , such that $W(\bar{F}) = \mathbb{F}^2$ (see the construction in [Sha], §3). Clearly W is an \mathbb{F} -vector space scheme.

Lemma 4.3. *There is an integer $d > 0$ such that $G[\mathfrak{m}]_{K'}$ is isomorphic to d copies of $W_{K'} = W \times_F K'$.*

Proof. By Theorem 1 of [BLR2], $G[\mathfrak{m}]_{K'}(\bar{F})$ is isomorphic to d copies of $W_{K'}(\bar{F})$ as $\mathbb{F}[\text{Gal}(\bar{F}/F)]$ -modules. Indeed, as we saw in the proof of Proposition 3.17, for almost all places v of F the characteristic polynomial of $\rho(\text{Frob}_v)$ is $X^2 - \lambda(T_v)X + Nv\lambda(S_v)$. By the Eichler-Shimura relation of Proposition 2.6, Frob_v satisfies the same polynomial on $\text{Pic}^0(M_{U_1(\mathfrak{p}), H}^{bal})(\bar{F})$, hence on $G[\mathfrak{m}]_{K'}$. Hence the hypotheses of [BLR2], Thm. 1 hold by the Chebotarev density theorem.

Now, $G[\mathfrak{m}]_{K'}$ and $\bigoplus_{i=1}^d W_{K'}$ are group schemes over a field of characteristic zero, hence étale, and have the same action of $\text{Gal}(\bar{K}/K')$, so they must be isomorphic. \square

4.4. Further decompositions. As we mentioned in the introduction (see [Edi2] for a proof), the pro- p group I'_p lies in the kernel of $(\rho|_{G_p})^{ss}$. Hence $(\rho)^{ss}|_{I'_p}$ factors through the tame inertia $I_{t,p} = I_p/I'_p$ and is given by two characters $\phi, \phi' : I_{t,p} \rightarrow \bar{\mathbb{F}}_p^*$.

Lemma 4.4. *If $\rho|_{G_p}$ is irreducible then ϕ and ϕ' are of level $2s$, with $\phi^q = \phi'$ and $(\phi')^q = \phi$.*

Proof. Since $I_p \trianglelefteq G_p$, $I'_p \trianglelefteq I_p$, and $I_{t,p}$ is abelian, we see that G_p/I_p acts on $I_{t,p}$ by conjugation. Observe that $G_p/I_p \simeq \text{Gal}(F_p^{nr}/F_p)$ is topologically generated by the element Φ which lifts the automorphism $x \mapsto x^q$ of the residue field \bar{k}_p of F_p^{nr} .

Now let $\sigma \in I_{t,p} = \varprojlim \text{Gal}(K_n/F_p^{nr})$. If π is any uniformizer of K_n , then $\sigma(\pi) = i_n(\sigma)\pi$ and:

$$\Phi \sigma \Phi^{-1}(\pi) = \Phi(i_n(\sigma)\Phi^{-1}(\pi)) = (i_n(\sigma))^q \pi.$$

Hence the representation $\rho' : I_{t,p} \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$ defined by $\rho'(\sigma) = \rho(\Phi \sigma \Phi^{-1})$ acts by the characters ϕ^q and $(\phi')^q$. Therefore $\{\phi, \phi'\} = \{\phi^q, (\phi')^q\}$, and since $q = p^s$ one of the following holds:

- (1) The characters ϕ and ϕ' are of level s . In this case, the action of G_p preserves the “ ϕ -eigenspace” of V , and so $\rho|_{G_p}$ is reducible.
- (2) The characters ϕ and ϕ' are of level $2s$, with $\phi' = \phi^q$ and $\phi = (\phi')^q$. Then $\rho|_{G_p}$ is irreducible.

Since we assumed that $\rho|_{G_p}$ was irreducible, we must be in the second case. \square

Thus under the hypotheses of Theorem 1.3, ϕ and ϕ' are both of level $2s$ (and not of level s). Let \mathbb{F}' be an extension of \mathbb{F} containing $\mathbb{F}_{p^{2s}}$. Then we have a decomposition $W_K \otimes_{\mathbb{F}} \mathbb{F}' = W_{\phi} \oplus W_{\phi'}$.

The left-hand side is the \mathbb{F}' -vector space scheme corresponding to the $\text{Gal}(\bar{K}/K)$ -module $(\mathbb{F}')^2 \otimes_{\mathbb{F}} \mathbb{F}'$ with the $\text{Gal}(\bar{K}/K)$ -action given by $\rho^{ss}|_{I_{\mathfrak{p}}}$. We define W_{ϕ} and $W_{\phi'}$ analogously; $\text{Gal}(\bar{K}/K) = I_{\mathfrak{p}}$ acts on them via the characters ϕ and ϕ' , respectively. Furthermore, let H_{ϕ} and $H_{\phi'}$ be the one-dimensional $\mathbb{F}_{p^{2s}}$ -vector space schemes over K corresponding to the characters ϕ and ϕ' . Then,

$$W_{\phi} = \bigoplus_{k=1}^r H_{\phi},$$

where $r = [\mathbb{F}' : \mathbb{F}]$, and there is an analogous decomposition of $W_{\phi'}$. These clearly base change to direct sum decompositions over K' . From the above and Lemma 4.3 it follows that

$$G[\tilde{\mathfrak{m}}]_{K'} \otimes_{\mathbb{F}} \mathbb{F}' = \bigoplus_{\substack{1 \leq j \leq d \\ 1 \leq k \leq r}} (H_{\phi,j,k})_{K'} \oplus (H_{\phi',j,k})_{K'}.$$

Here the $H_{\phi,j,k}$ and $H_{\phi',j,k}$ are copies of H_{ϕ} and $H_{\phi'}$, respectively. It follows from [Ray], 1.2.2 that they satisfy the condition $(**)$ as $\mathbb{F}_{p^{2s}}$ -vector space schemes. Since the absolute ramification degree of D' is $p^s - 1 \geq p - 1$, we know from Section 4.1 that all the vector space schemes in the above displayed formula extend to D' . If $s > 1$, these extensions are not unique.

Lemma 4.5. *Let H_1, H_2, \dots, H_{2dr} be any labeling of the set $\{(H_{\phi,j,k})_{K'}, (H_{\phi',j,k})_{K'} : 1 \leq j \leq d, 1 \leq k \leq r\}$. Then $G[\tilde{\mathfrak{m}}] \otimes_{\mathbb{F}} \mathbb{F}'$, as an $\mathbb{F}_{p^{2s}}$ -vector space scheme over D' , has a filtration:*

$$G[\tilde{\mathfrak{m}}] \otimes_{\mathbb{F}} \mathbb{F}' = V_0 \supset V_1 \supset \dots \supset V_{2dr} = 0$$

such that for every $1 \leq i \leq 2dr$, the quotient $\tilde{H}_i = V_{i-1}/V_i$ is an $\mathbb{F}_{p^{2s}}$ -vector space scheme over D' extending H_i . Furthermore, the action of $\text{Gal}(K'/K)$ preserves this filtration.

Proof. The image of the generic fiber $G[\tilde{\mathfrak{m}}]_{K'} = G[\tilde{\mathfrak{m}}] \times_{D'} K'$ in $G[\tilde{\mathfrak{m}}]$ is a dense open subset. Let $Q = \bigoplus_{i=2}^{2dr} H_i$ be the direct sum of all the components of $G[\tilde{\mathfrak{m}}]_{K'}$ except for H_1 , and let V_1 be the closure of its image in $G[\tilde{\mathfrak{m}}]$. This is clearly preserved by the continuous action of $\text{Gal}(K'/K)$; note that all the H_i are defined over K . By [Sha] §3, there exists a quotient group scheme V_0/V_1 , which is flat and finite over D' , and its generic fiber is clearly isomorphic to H_1 . Iterating this process, we construct all the elements of the filtration. \square

Taking special fibers, we get a filtration of schemes over $D'/\mathfrak{m}_{D'} D' = \bar{\mathbb{F}}_p$:

$$(13) \quad \bar{G}[\tilde{\mathfrak{m}}] \otimes \mathbb{F}' = \bar{V}_0 \supset \bar{V}_1 \supset \dots \supset \bar{V}_{2dr} = 0.$$

Observe that \bar{V}_{2dr-1} is the special fiber of the vector space scheme V_{2dr-1} over D' , whose generic fiber is in turn $(H_{2dr})_{K'}$. This is the setup of the end of section 4.1; $\text{cot}(\bar{V}_{2dr-1})$ is what we would call $\text{cot}((H_{2dr})_{\bar{\mathbb{F}}_p})$ in the notation of that section. Thus, $\text{cot}((H_{2dr})_{\bar{\mathbb{F}}_p})$ embeds into $\text{cot}(\bar{G})$.

5. PROOF OF THEOREM 1.3

5.1. Numerology of vector space schemes. We now suppose, without loss of generality, that in our chosen labeling H_1, \dots, H_{2dr} of the components of $G[\tilde{\mathfrak{m}}]_{K'}$ the last component H_{2dr} is one of the $H_{\phi,j,k}$. This is an $\mathbb{F}_{p^{2s}}$ -vector space scheme satisfying $(**)$, on which $\text{Gal}(\bar{K}/K) = I_{\mathfrak{p}}$ acts

by the character ϕ . The vector space scheme $\tilde{H}_{2dr} = V_{2dr-1}$ extends H_{2dr} to D' . We drop the subscript and write H for H_{2dr} .

Let $\sigma_0, \sigma_1, \dots, \sigma_{2s-1}$ be the embeddings of $\mathbb{F}_{p^{2s}}$ into $\overline{\mathbb{F}}_p$, labeled so that $\sigma_i^p = \sigma_{i-1}$ and the restriction of σ_i to \mathbb{F}_{p^s} is $\tau_{\pi(i)}$, where $\pi : \mathbb{Z}/2s\mathbb{Z} \rightarrow \mathbb{Z}/s\mathbb{Z}$ is the natural projection. Then let $\psi_i : I_{t,\mathfrak{p}} \rightarrow \overline{\mathbb{F}}_p^*$ be the fundamental character of level $2s$ corresponding to σ_i , and let $\chi_i : \mathbb{F}_{p^{2s}}^* \rightarrow D^*$ be the inverse mapping of ψ_i . Note that $\chi_i^p = \chi_{i+1}$, and we have precisely the setup of Section 4.1. Now let $a_0, a_1, \dots, a_{2s-1}$ be parameters corresponding to H , as defined in Section 4.1, and recall that $\phi = \psi_0^{a_{2s-1}} \psi_1^{a_0} \cdots \psi_{2s-1}^{a_{2s-2}}$. Similarly define a'_i and b_i , for $i \in \mathbb{Z}/2s\mathbb{Z}$.

Suppose that ρ is modular for a weight $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$ which is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$ for a non-trivial character $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ of the form

$$(14) \quad \theta : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(ad)^{w_{\tau}} \tau(d)^{k_{\tau}-2}.$$

Lemma 5.1. *As before, we write k_j for k_{τ_j} . Then,*

$$\phi\phi' = \prod_{i \in \mathbb{Z}/2s\mathbb{Z}} \psi_i^{2w_{\pi(i)} + k_{\pi(i)} - 1}.$$

Proof. Clearly, $\phi\phi'$ is a character of level s and factors through $\text{Gal}(K'/K)$. If $\chi_{\eta}^{-1} : \text{Gal}(\overline{F}/F) \rightarrow \overline{\mathbb{F}}_p^*$ is as in Corollary 3.21, then observe that $\chi_{\eta}^{-1}|_{I_{t,\mathfrak{p}}} = \prod_i \psi_i^{w_i}$. Hence it suffices to consider the case that $w_i = 0$ for all i .

This follows from Proposition 3.19. Let $\sigma \in \text{Gal}(\overline{K}/K)$, and suppose its image in $\text{Gal}(K'/K)$ corresponds to $j(\sigma) \in \mathcal{O}_{\mathfrak{p}}^*/(1 + \mathfrak{p})$ by the Artin reciprocity map. We identify $j(\sigma)$ with its image in $k_{\mathfrak{p}}^*$. Since $\mathcal{O}_{\mathfrak{p}}$ contains the $(q-1)$ -st roots of unity, we can identify $\mathcal{O}_{\mathfrak{p}}^*/(1 + \mathfrak{p})$ with $\mu_{q-1}(D)$. Then by Proposition 3.19 and Lemma 3.15 the following holds in $\mathbb{F}^* \subset \overline{\mathbb{F}}_p^*$:

$$\begin{aligned} \phi(\sigma)\phi'(\sigma) &= \det \rho(\sigma) = \langle \sigma \rangle^{-1} \chi(\sigma) = \lambda(\langle j(\sigma)^{-1} \rangle) N_{k_{\mathfrak{p}}/\mathbb{F}_p}(j(\sigma))^{-1} = \\ &= \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(j(\sigma))^{k_{\tau}-2} \tau(j(\sigma))^{-1} = \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(j(\sigma))^{k_{\tau}-1} = \prod_{i \in \mathbb{Z}/2s\mathbb{Z}} \sigma_i(j(\sigma))^{k_{\pi(i)}-1} = \prod_{i \in \mathbb{Z}/2s\mathbb{Z}} \psi_i(\sigma)^{k_{\pi(i)}-1}. \end{aligned}$$

The last equality is true because the maps $j : \text{Gal}(K'/K) \rightarrow \mathcal{O}_{\mathfrak{p}}^*/(1 + \mathfrak{p})$ and $i_s : \text{Gal}(K'/K) \rightarrow \mu_{q-1}(D)$ are the same by Proposition 3 of [Ser] (note that in that article the Artin reciprocity map is normalized so that uniformizers correspond to *arithmetic* Frobenius elements). \square

Lemma 5.2. *The cotangent space $\text{cot}(H_{\overline{\mathbb{F}}_p})$ does not vanish.*

Proof. Suppose that $\text{cot}(H_{\overline{\mathbb{F}}_p}) = 0$. We claim that $H_{\overline{\mathbb{F}}_p}$ is étale as a finite group scheme over $\overline{\mathbb{F}}_p$. It suffices to prove étaleness at the origin; by translation this will imply étaleness at any $\overline{\mathbb{F}}_p$ -point of $H_{\overline{\mathbb{F}}_p}$, and this is enough since the étale locus is open. Let \mathcal{O} be the complete local ring of $H_{\overline{\mathbb{F}}_p}$ at the origin, and let \mathfrak{n} be its maximal ideal. Since the cotangent space at the origin vanishes, $\mathfrak{n} = \mathfrak{n}^2$. But $\bigcap_n \mathfrak{n}^n = 0$ since \mathcal{O} is noetherian, so $\mathfrak{n} = 0$. Thus \mathcal{O} is a finite field extension of $\overline{\mathbb{F}}_p$, hence it is $\overline{\mathbb{F}}_p$ itself. This establishes that $H_{\overline{\mathbb{F}}_p}$ is étale.

Recall that a group scheme over D' is étale if and only if its special fiber is étale. Thus V_{2dr-1} is étale, so as a D' -scheme it is a direct sum of a finite number of copies of $\text{Spec } D'$; all finite étale

schemes over the strictly henselian ring D' are of this form ([Mil], I.4). Therefore, the generic fiber $H_{K'} = V_{2dr-1} \times_{D'} K'$ is a constant scheme over K' . Obviously, $\text{Gal}(\overline{K}/K')$ acts trivially on $H(\overline{K})$, and hence ϕ is a character of level s , giving a contradiction. \square

For each $j \in \mathbb{Z}/s\mathbb{Z}$, let $c_j = k_j - 2 + p(k_{j-1} - 2) + \cdots + p^{s-1}(k_{j+1} - 2)$. Since θ is non-trivial, $0 < c_j < p^s - 1$. We often write k_i and c_i instead of $k_{\pi(i)}$ and $c_{\pi(i)}$, for $i \in \mathbb{Z}/2s\mathbb{Z}$.

We temporarily return to the assumption that θ is untwisted.

Lemma 5.3. *Suppose that $\sigma_{\mathfrak{p}}$ is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$ for a character $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ of the form of (11). Then for all $i \in \mathbb{Z}/2s\mathbb{Z}$, we have $b_i \in \{0, c_{\pi(i)}\} \pmod{p^s - 1}$.*

Proof. By Proposition 2.7 and Lemma 2.9, we see that $\text{Gal}(K'/K)$ acts on J_I and J_E , and hence on $\text{cot}(\overline{G}_I)$ and $\text{cot}(\overline{G}_E)$, by $\langle j(\sigma)^{-1}, 1 \rangle^{-1} = \langle j(\sigma), 1 \rangle$ and $\langle 1, j(\sigma)^{-1} \rangle^{-1} = \langle 1, j(\sigma) \rangle$, respectively.

Recall the description of $\text{cot}(H_{\overline{\mathbb{F}}_p})$ from the end of the previous section. In particular, we see from (13) that $\text{cot}(H_{\overline{\mathbb{F}}_p}) \subset \text{cot}(\overline{G}[\mathfrak{m}])$. Since $\text{cot}(H_{\overline{\mathbb{F}}_p}) \neq 0$ by Lemma 5.2, there must be some $i \in \mathbb{Z}/2s\mathbb{Z}$ such that $a'_{i-1} > 0$ and dX'_i is a non-trivial element of $\text{cot}(H_{\overline{\mathbb{F}}_p})$. Now, $\text{cot}(\overline{G}[\mathfrak{m}]) = \text{cot}(\overline{G}_I[\mathfrak{m}]) \oplus \text{cot}(\overline{G}_E[\mathfrak{m}])$. By Lemma 3.15, in \mathbb{F}^* we have $\langle j(\sigma), 1 \rangle = 1$ and $\langle 1, j(\sigma) \rangle = \prod_{\tau: k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p} \tau(j(\sigma))^{2-k_{\tau}}$; observe that these are in fact contained in $\mathbb{F}_{p^{2s}}^*$. As we saw above, the maps j and i_s are the same.

Thus, σ can act on dX'_i either by the character 1 or by:

$$\chi_i \left(\prod_{j \in \mathbb{Z}/s\mathbb{Z}} \tau_j(i_s(\sigma))^{2-k_j} \right) = \chi_i(\tau_{\pi(i)}(i_s(\sigma))^{-c_{\pi(i)}}) = i_s(\sigma)^{-c_{\pi(i)}}.$$

But we know from the discussion preceding (9) that σ acts on dX'_i by $i_s(\sigma)^{-b_i}$. Thus we have shown that if $a'_{i-1} > 0$, then $b_i \in \{0, c_i\} \pmod{p^s - 1}$, and there exists at least one such i . Now, if $a'_{i-1} = 0$, then by the fundamental equation (10) we have $b_i \equiv pb_{i-1} \pmod{p^s - 1}$. The observation that $c_i \equiv pc_{i-1} \pmod{p^s - 1}$ completes the proof of the lemma. \square

We now apply (10) to determine the a_i and hence ϕ . Our analysis splits into four cases.

Case 1. $b_i = 0, b_{i+1} = c_{i+1}$. In this case, (10) tells us that

$$a'_i - (p^s - 1)a_i = b_{i+1} - pb_i = c_{i+1}.$$

Recall from Section 4.1 that since our b_i 's are only defined modulo $p^s - 1$, we can only determine the set of a_i 's modulo the lattice Λ . However, different sets of a_i equivalent modulo Λ give the same character ϕ , so this is enough for us. On the other hand, we know that $0 \leq a'_i \leq p^s - 1$. The latter condition implies that $a'_i = c_{i+1}$ and $a_i = 0$.

Case 2. $b_i = c_i, b_{i+1} = 0$. This time we see from (10) that:

$$\begin{aligned} a'_i - (p^s - 1)a_i &= -pc_i = -p^s(k_{i+1} - 2) - p^{s-1}(k_{i+2} - 2) - \cdots - p(k_i - 2) = \\ &= -(p^s - 1)(k_{i+1} - 2) - [p^{s-1}(k_{i+2} - 2) + \cdots + p(k_i - 2) + (k_{i+1} - 2)] = \\ &= [p^{s-1}(p + 1 - k_{i+2}) + \cdots + p(p + 1 - k_i) + (p + 1 - k_{i+1})] - (p^s - 1)(k_{i+1} - 1) \end{aligned}$$

Let β be the bracketed quantity in the bottom line of the displayed expression. Observe that $0 < \beta < p^s - 1$. We must have $a'_i = \beta$ and $a_i = k_{i+1} - 1$.

Case 3. $b_i = 0, b_{i+1} = 0$. From (10) we have $a'_i - (p^s - 1)a_i = 0$, which admits two solutions:

$$\begin{aligned} a'_i &= 0 & a_i &= 0 \\ a'_i &= p^s - 1 & a_i &= 1 \end{aligned}$$

Case 4. $b_i = c_i, b_{i+1} = c_{i+1}$. In this case we have:

$$a'_i - (p^s - 1)a_i = b_{i+1} - pb_i = c_{i+1} - pc_i = -(p^s - 1)(k_{i+1} - 2).$$

Therefore, there are again two possibilities in all cases, namely

$$\begin{aligned} a'_i &= 0 & a_i &= k_{i+1} - 2 \\ a'_i &= p^s - 1 & a_i &= k_{i+1} - 1 \end{aligned}$$

The following lemma shows that we may assume without loss of generality that $\{b_i, b_{i+s}\} = \{0, c_{\pi(i)}\}$ for each i ; this assumption will not change the set of possible characters ϕ that we obtain. We will write the collections of a_i 's as vectors: $\mathbf{a} = (a_0, a_1, \dots, a_{2s-1})$, and similarly for the b_i 's. We say that \mathbf{a} is associated to \mathbf{b} if these collections of a_i and b_i satisfy (10) for each i .

Lemma 5.4. *Let \mathbf{a} and \mathbf{b} be as above. Then $b_i \in \{0, c_i\}$ for each i , and \mathbf{a} is associated to \mathbf{b} . There exists $\hat{\mathbf{b}} = \{\hat{b}_0, \dots, \hat{b}_{2s-1}\}$ such that $\{\hat{b}_i, \hat{b}_{i+s}\} = \{0, c_i\}$ for each i and a vector $\hat{\mathbf{a}}$ associated to $\hat{\mathbf{b}}$ such that $\mathbf{a} - \hat{\mathbf{a}} \in \Lambda$.*

Proof. We prove this by brute force; see [Sch1], Lemma 3.2 for a more conceptual, but essentially equivalent, argument. Let $\mu_i = a_i + a_{i+s} - (k_{i+1} - 1)$. By Lemma 5.1, $\boldsymbol{\mu} = (\mu_0, \dots, \mu_{2s-1}) \in \Lambda$. Furthermore, the computations above show that, for all i , $|\mu_i| \leq k_{i+1} - 1 \leq p$. Hence each μ_i must be one of $0, \pm 1, \pm(p-1)$, or $\pm p$.

Case 1. Assume first that $\{b_i, b_{i+s}\} = \{0, c_i\}$ holds for at least one index i . If it is true for all i , we are done. Otherwise, for some i we have $\{b_{i+1}, b_{i+s+1}\} = \{0, c_{i+1}\}$, whereas $\{b_i, b_{i+s}\}$ is $\{0, 0\}$ or $\{c_i, c_i\}$. We will construct a vector $\tilde{\mathbf{b}}$ such that $\{\tilde{b}_k, \tilde{b}_{k+s}\} = \{0, c_k\}$ for all k such that $\{b_k, b_{k+s}\} = \{0, c_k\}$ and $\{\tilde{b}_i, \tilde{b}_{i+s}\} = \{0, c_i\}$, and an $\tilde{\mathbf{a}}$ associated to $\tilde{\mathbf{b}}$ such that $\mathbf{a} - \tilde{\mathbf{a}} \in \Lambda$. Iterating this process, we can construct $\tilde{\mathbf{b}}$ and $\tilde{\mathbf{a}}$ as required.

One considers a number of cases; we give an example of one. Assume that $b_{i+1} = 0$ and $b_{i+s+1} = c_{i+1}$, and consider the case $b_i = b_{i+s} = 0$. Then $a_i \in \{0, 1\}$ and $a_{i+s} = 0$. Observe that $\mu_{i+1} \in \{-1, 0, 1\}$, and hence $\mu_i \in \{0, \pm p\}$. Hence either $k_{i+1} = 2$ and $a_i = 1$, or $k_{i+1} = p + 1$ and $a_i = 0$.

Suppose that $b_{i-1} = c_{i-1}$. Then $a_{i-1} = k_i - 1$. Define $\tilde{b}_k = b_k$ for $k \neq i$ and $\tilde{b}_i = c_i$. Define $\tilde{a}_k = a_k$ for $k \neq i, i-1$. If $a_i = 1$, let $\tilde{a}_{i-1} = k_i - 1$ and $\tilde{a}_i = 1 = k_{i+1} - 1$; then $\tilde{\mathbf{a}} = \mathbf{a}$. If $a_i = 0$, let $\tilde{a}_i = p = k_{i+1} - 1$ and $\tilde{a}_{i-1} = k_i - 2$; clearly $\mathbf{a} - \tilde{\mathbf{a}} \in \Lambda$. In either case, $\tilde{\mathbf{a}}$ is associated to $\tilde{\mathbf{b}}$.

Case 2. Now suppose that $b_i = b_{i+s}$ for all i . Observe that for some i one of the following holds, since otherwise ϕ would be a character of level s : either $b_i = b_{i+s} = b_{i+1} = b_{i+s+1} = 0, a_i = 0, a_{i+s} = 1$ or $b_i = b_{i+s} = c_i, b_{i+1} = b_{i+s+1} = c_{i+1}, a_i = k_{i+1} - 1, a_{i+s} = k_{i+1} - 2$.

Let $\mu = \mu_{i+s} + p\mu_{i-1} + \dots + p^{s-1}\mu_{i+1}$. Since $\mu_k = \mu_{k+s}$, $\boldsymbol{\mu} \in \Lambda$ implies $(1+p^s)\mu \equiv 0 \pmod{p^{2s}-1}$, whence $\mu \equiv 0 \pmod{p^s-1}$. Also observe that $\mathbf{a} - \hat{\mathbf{a}} \in \Lambda$ if and only if $\alpha = (a_{i+s} - \hat{a}_{i+s}) + p(a_{i+s-1} - \hat{a}_{i+s-1}) + \dots + p^{2s-1}(a_{i+s+1} - \hat{a}_{i+s+1}) \equiv 0 \pmod{p^{2s}-1}$.

If $\mu = 0$, then for $k \in \{i+s+1, i+s+2, \dots, i+2s-1, i\}$ we set $\hat{b}_k = b_k$ and for the remaining k we switch: if $b_k = 0$ (resp. $b_k = c_k$), then $\hat{b}_k = c_k$ (resp. $\hat{b}_k = 0$). Then define $\hat{a}_k = a_k$ for $k \in \{i+s+1, \dots, i+2s-1\}$, $\hat{a}_k = (k_{k+1} - 1) - a_k$ for $k \in \{i+1, \dots, i+s-1\}$,

and $(\hat{a}_i, \hat{a}_{i+s}) = (0, k_{i+1} - 1)$ (resp. $(\hat{a}_i, \hat{a}_{i+s}) = (k_{i+1}, 0)$) if $a_i = 0$ (resp. $a_i = k_{i+1} - 1$). The construction for $\mu \neq 0$ is similar. \square

We will now assume that $\{b_i, b_{i+s}\} = \{0, c_i\}$ for all i . In this situation, the condition $\mu \in \Lambda$ implies that $\mu = 0$, so we have $a_i + a_{i+s} = k_{i+1} - 1$ for all i .

Proposition 5.5. *Let $\rho : \text{Gal}(\bar{F}/F) \rightarrow \text{GL}_2(\bar{\mathbb{F}}_p)$ be such that $\rho|_{G_p}$ is irreducible and ρ is modular of weight σ such that σ_p is a constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_p)} \theta$, where $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \bar{\mathbb{F}}_p^*$ as in (11) is untwisted and non-trivial. Then there exists a subset $S \subset I$ and a labeling $\{\tilde{\tau}, \tilde{\tau}'\}$ of the two liftings of $\tau : k_p \hookrightarrow \bar{\mathbb{F}}_p$ to $\mathbb{F}_{p^{2s}}$ for each τ , such that*

$$(15) \quad \rho|_{I_{t,p}} \sim \begin{pmatrix} \phi & 0 \\ 0 & \phi' \end{pmatrix}.$$

where $\phi' = \phi^q$ and

$$\phi = \prod_{\tau \in S} \psi_{\tilde{\tau}}^{k_{\tau}-2+\nu_S(\tau)} \prod_{\tau \notin S} \psi_{\tilde{\tau}}^p \psi_{\tilde{\tau}'}^{k_{\tau}-2+\nu_S(\tau)}.$$

Proof. Let $\Phi(\theta)$ be the set of all characters ϕ of the form in the above statement. Any $\phi \in \Phi(\theta)$ is specified by the data (S, ε_j) , where $S \subset I$ and for each $j \in \mathbb{Z}/s\mathbb{Z}$, ε_j is a bijection of sets $\varepsilon_j : \pi^{-1}(j) = \{j, j+s\} \rightarrow \{\psi_{\tilde{\tau}_j}, \psi_{\tilde{\tau}'_j}\}$ (think of the elements of the latter set as symbols). As usual, we write ε_i instead of $\varepsilon_{\pi(i)}$ for $i \in \mathbb{Z}/2s\mathbb{Z}$. Conversely, such data give rise to the character $\phi = \prod_{i \in \mathbb{Z}/2s\mathbb{Z}} \psi_i^{m_i}$, where

$$m_i = \begin{cases} k_i - 2 + \nu_S(\tau_i) & : \tau_i \in S, \varepsilon_i(i) = \psi_{\tilde{\tau}_i} \\ 0 & : \tau_i \in S, \varepsilon_i(i) = \psi_{\tilde{\tau}'_i} \\ p & : \tau_i \notin S, \varepsilon_i(i) = \psi_{\tilde{\tau}_i} \\ k_i - 2 + \nu_S(\tau_i) & : \tau_i \notin S, \varepsilon_i(i) = \psi_{\tilde{\tau}'_i} \end{cases}$$

It is clear that every $\phi \in \Phi(\theta)$ is described in this way, although different data may produce the same ϕ . The bijections ε_i correspond to the choice of which of σ_i, σ_{i+s} is $\tilde{\tau}_i$ and which is $\tilde{\tau}'_i$.

Similarly, let $\Omega(\theta)$ be the set of ϕ satisfying the conditions imposed by (10) via the computations above. Any $\phi \in \Omega(\theta)$ is specified by the data (S', δ_j) , where $S' \subset I$ and for each $j \in \mathbb{Z}/s\mathbb{Z}$, δ_j is a bijection of sets $\delta_j : \pi^{-1}(j) = \{j, j+s\} \rightarrow \{0, c_{j-1}\}$. Such data give rise to $\phi = \psi_0^{a_{2s-1}} \psi_1^{a_0} \cdots \psi_{2s-1}^{a_{2s-2}}$, where the a_i are determined by the following rules:

$$a_{i-1} = \begin{cases} 0 & : \delta_i(i) = 0, \delta_{i+1}(i+1) = c_i \\ k_i - 1 & : \delta_i(i) = c_{i-1}, \delta_{i+1}(i+1) = 0 \\ 0 & : \delta_i(i) = \delta_{i+1}(i+1) = 0, \tau_{i+1} \in S' \\ 1 & : \delta_i(i) = \delta_{i+1}(i+1) = 0, \tau_{i+1} \notin S' \\ k_i - 1 & : \delta_i(i) = c_{i-1}, \delta_{i+1}(i+1) = c_i, \tau_{i+1} \in S' \\ k_i - 2 & : \delta_i(i) = c_{i-1}, \delta_{i+1}(i+1) = c_i, \tau_{i+1} \notin S' \end{cases}$$

In other words, $\delta_i(i) = 0$ and $\delta_i(i) = c_{i-1}$ correspond to $b_{i-1} = 0$ and $b_{i-1} = c_{i-1}$, respectively, and S' decides which of the two options we choose in Cases 3 and 4. By virtue of Lemma 5.4 and the remark after it, each $\phi \in \Omega(\theta)$ may be described (perhaps not uniquely) in this way.

Now, let $\lambda_i : \{\psi_{\tilde{\tau}_i}, \psi_{\tilde{\tau}'_i}\} \rightarrow \{c_{i-1}, 0\}$ be the bijection which sends $\psi_{\tilde{\tau}_i}$ to c_{i-1} and $\psi_{\tilde{\tau}'_i}$ to 0. Let μ_i be the other bijection between the same sets. Consider the map $\omega : \Phi(\theta) \rightarrow \Omega(\theta)$ which sends the data (S, ε_i) to the data $\omega(S, \varepsilon_i) = (S', \delta_i)$, where $S' = S$ and

$$\delta_i = \begin{cases} \lambda_i \circ \varepsilon_i & : \tau_i \in S \\ \mu_i \circ \varepsilon_i & : \tau_i \notin S \end{cases}$$

Obviously ω gives a bijection between data defining elements of $\Phi(\theta)$ and data defining elements of $\Omega(\theta)$. One verifies that (S, ε_i) and $\omega(S, \varepsilon_i)$ give rise to the same ϕ , and hence $\Phi(\theta) = \Omega(\theta)$. \square

Remark 5.6. The statement of Proposition 5.5 has the same form as that of Theorem 1.3, which we wish to prove. If $3 \leq k_\tau \leq p$ for all τ , then by Lemma 3.7 $Ind_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})}\theta$ has a non-zero irreducible constituent for every $S \subset I$. It is easy to see that the union of the sets of ϕ 's predicted by the conjecture of [BDJ] for all of these constituents is precisely $\Phi(\theta)$. Thus, Proposition 5.5 is the ‘‘best possible result.’’ On the other hand, if $k_\tau = 2$ or $k_\tau = p + 1$ for some τ , then for some S we do not get a constituent, and $\Phi(\theta)$ is larger than the set predicted by the conjecture.

Corollary 5.7. *Retain the hypotheses of Proposition 5.5, except that $\sigma_{\mathfrak{p}}$ is now a constituent of $Ind_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})}\theta$ for an arbitrary $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ as in (14). Then there is a subset $S \subset I$ and a labeling $\{\tilde{\tau}, \tilde{\tau}'\}$ of the two liftings of $\tau : k_{\mathfrak{p}} \hookrightarrow \overline{\mathbb{F}}_p$ to $\mathbb{F}_{p^{2s}}$ for each τ , such that $\rho|_{I_{\mathfrak{p}}}$ has the form (15), with*

$$(16) \quad \phi = \prod_{\tau} \psi_{\tilde{\tau}}^{w_{\tau}} \psi_{\tilde{\tau}'_i}^{w_{\tau}} \left(\prod_{\tau \in S} \psi_{\tilde{\tau}}^{k_{\tau}-2+\nu_S(\tau)} \prod_{\tau \notin S} \psi_{\tilde{\tau}}^p \psi_{\tilde{\tau}'_i}^{k_{\tau}-2+\nu_S(\tau)} \right).$$

Proof. This is almost immediate from Proposition 5.5. Indeed, suppose that ρ is modular of weight $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$, with $\sigma_{\mathfrak{p}}$ as in the statement. Let η be a one-dimensional weight such that $\eta_{\mathfrak{p}} = \otimes_{\tau} (\det^{-w_{\tau}} \otimes_{\tau} \overline{\mathbb{F}}_p)$. By Corollary 3.21, ρ is modular of weight σ if and only if $\rho \otimes \chi_{\eta}^{-1}$ is modular of weight $\sigma \otimes \eta$. Also, $\sigma_{\mathfrak{p}} \otimes \eta_{\mathfrak{p}}$ is a constituent of $Ind_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})}\tilde{\theta}$, where

$$\tilde{\theta} : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{j \in \mathbb{Z}/s\mathbb{Z}} \tau_j(d)^{k_j-2}.$$

Proposition 5.5 applies to $\rho \otimes \chi_{\eta}^{-1}$. The Artin reciprocity map identifies $(I_{\mathfrak{p}})^{ab}$ with $\mathcal{O}_{\mathbb{F}_p}^*$, on which χ_{η} acts by $\prod_{\tau} \psi_{\tilde{\tau}}^{w_{\tau}} = \prod_{\tau} (\psi_{\tilde{\tau}} \psi_{\tilde{\tau}'_i})^{w_{\tau}}$. Hence $\rho = (\rho \otimes \chi_{\eta}^{-1}) \otimes \chi_{\eta}$ has the required form. \square

As before, we denote by $\Phi(\theta)$ the set of all characters ϕ of the form (16). For any $\sigma_{\mathfrak{p}}$, let $\Theta(\sigma_{\mathfrak{p}})$ be the set of $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ such that $\sigma_{\mathfrak{p}}$ is an irreducible constituent of $Ind_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})}\theta$. The previous argument shows that if the statement of Theorem 1.3 holds for some weight σ , then it holds for $\sigma \otimes \eta$ if η is any one-dimensional weight. Thus we assume from now on that $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$, where

$$\sigma_{\mathfrak{p}} = \bigotimes_{i \in \mathbb{Z}/s\mathbb{Z}} (\mathrm{Sym}^{k_i-2} k_{\mathfrak{p}}^2 \otimes_{\tau_i} \overline{\mathbb{F}}_p).$$

It is easy to see that $\Theta(\sigma_{\mathfrak{p}})$ contains some θ factoring through the determinant if and only if the k_{τ} are all 2 or all $p + 1$.

5.2. The cases $s = 1$ and $s = 2$ and the general method. Suppose that $s = 1$, i.e. $k_{\mathfrak{p}} = \mathbb{F}_p$; for instance, this will happen if p is completely split in F . In this case $I = \{\tau\}$ is a singleton, and $\sigma_{\mathfrak{p}} = \text{Sym}^{k-2} k_{\mathfrak{p}}^2 \otimes \overline{\mathbb{F}}_p$, with $k \neq 2, p+1$. In this case Theorem 1.3 is immediate from Proposition 5.5.

Now assume $s \geq 2$. The proofs in the several cases below all use the strategy outlined below. Let $\Phi(\sigma)$ be the set of characters ϕ that appear in $(\rho|_{I_{t,\mathfrak{p}}})^{\text{ss}}$ for some Galois representation ρ which is modular of weight σ and such that $\rho|_{G_{\mathfrak{p}}}$ is irreducible. We observe that $\Phi(\sigma) \subset \bigcap_{\theta \in \Theta(\sigma_{\mathfrak{p}})} \Phi(\theta)$ and proceed to compute this intersection; it will essentially be our bound on $\Phi(\sigma)$.

By Lemma 3.7, for each $S \subset I$ there is at most one θ_S such that $\sigma_{\mathfrak{p}}$ is the S -constituent of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$. Indeed, θ_S does not exist if and only if there is some $i \in \mathbb{Z}/s\mathbb{Z}$ such that $k_i = p+1$ and one of the following two conditions holds: either $\tau_i \in S$ and $\tau_{i+1} \notin S$, or $\tau_i \notin S$ and $\tau_{i+1} \in S$. Otherwise θ_S exists and is the following:

$$\theta_S : \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \prod_{\tau \in S} \tau(ad)^{p-1} \tau(d)^{k_{\tau}-1-\nu_S(\tau)} \prod_{\tau \notin S} \tau(ad)^{k_{\tau}-2} \tau(d)^{p+1-k_{\tau}-\nu_S(\tau)}.$$

By assumption the k_{τ} are not all 2 and not all $p+1$. Then if θ_S exists, it does not factor through the determinant, and by Corollary 5.7 $\Phi(\theta_S)$ consists of the following, where T runs over $\mathcal{P}(I)$:

$$(17) \quad \prod_{\substack{\tau \in T \\ \tau \in S}} \psi_{\tilde{\tau}}^{k_{\tau}+p-1-\nu_S(\tau)} \psi_{\tilde{\tau}'}^{p-\nu_T(\tau)} \prod_{\substack{\tau \in T \\ \tau \notin S}} \psi_{\tilde{\tau}}^{p-\nu_S(\tau)} \psi_{\tilde{\tau}'}^{k_{\tau}-1-\nu_T(\tau)} \prod_{\substack{\tau \notin T \\ \tau \in S}} \psi_{\tilde{\tau}}^{p-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{k_{\tau}-1-\nu_S(\tau)} \prod_{\substack{\tau \notin T \\ \tau \notin S}} \psi_{\tilde{\tau}}^{k_{\tau}-1-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{-\nu_S(\tau)}$$

Remark 5.8. Observe that $\bigcap_S \Phi(\theta_S)$ contains all ϕ of the form $\prod_{\tau} \psi_{\tilde{\tau}}^{k_{\tau}-1} \psi_{\tilde{\tau}'}^0$; here the intersection is over all S such that θ_S exists. Indeed, for each such S consider $T = S$. The corresponding ϕ are

$$\prod_{\tau \in S} \psi_{\tilde{\tau}}^{k_{\tau}+p-1-\nu_S(\tau)} \psi_{\tilde{\tau}'}^{p-\nu_S(\tau)} \prod_{\tau \notin S} \psi_{\tilde{\tau}}^{k_{\tau}-1-\nu_S(\tau)} \psi_{\tilde{\tau}'}^{-\nu_S(\tau)} = \prod_{\tau \in S} \psi_{\tilde{\tau}}^{k_{\tau}-1} \psi_{\tilde{\tau}'}^0 \prod_{\tau \notin S} \psi_{\tilde{\tau}}^{k_{\tau}-1} \psi_{\tilde{\tau}'}^0.$$

By (17), $\Phi(\theta_I)$ consists of all ϕ of the form $\prod_{\tau \in T} \psi_{\tilde{\tau}}^{k_{\tau}-2+\nu_T(\tau)} \psi_{\tilde{\tau}'}^0 \prod_{\tau \notin T} \psi_{\tilde{\tau}}^{k_{\tau}-2+\nu_T(\tau)} \psi_{\tilde{\tau}'}^p$. Hence if $\phi \in \bigcap_S \Phi(\theta_S)$, then in particular $\phi \in \Phi(\theta_I)$ and so $\phi = \prod_{j \in \mathbb{Z}/2s\mathbb{Z}} \psi_j^{m_j}$, where for each j the set $\{m_j, m_{j+s}\}$ is either $\{k_j - 1, 0\}$ or $\{k_j - 2, 1\}$. In the latter case, if $\alpha \in \{j, j+s\}$ is such that $m_{\alpha} = 1$, then if $\{m_{j+1}, m_{j+1+s}\} = \{k_{j+1} - 1, 0\}$ (resp. $\{m_{j+1}, m_{j+1+s}\} = \{k_{j+1} - 2, 1\}$), we see that $m_{\alpha+1} = 0$ (resp. $m_{\alpha+1} = 1$).

Now suppose $s = 2$. We resolve this case by explicit computation. Let $I = \{\tau_1, \tau_2\}$ and suppose that k_1 and k_2 are not both 2 and not both $p+1$. We just saw that $\bigcap_S \Phi(\theta_S)$ contains all ϕ of the form $\psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^0 \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^0$. Suppose first that $2 \leq k_1, k_2 \leq p$. Then θ_S exists for all S , and by (17):

$$\Phi(\theta_I) = \left\{ \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^0 \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^0, \psi_{\tilde{\tau}_1}^{k_1-2} \psi_{\tilde{\tau}_1'}^0 \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^p, \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^p \psi_{\tilde{\tau}_2}^{k_2-2} \psi_{\tilde{\tau}_2'}^0, \psi_{\tilde{\tau}_1}^{k_1-2} \psi_{\tilde{\tau}_1'}^p \psi_{\tilde{\tau}_2}^{k_2-2} \psi_{\tilde{\tau}_2'}^p \right\}$$

where T runs over $\mathcal{P}(I)$, and $\{\tilde{\tau}_1, \tilde{\tau}_1'\}$ and $\{\tilde{\tau}_2, \tilde{\tau}_2'\}$ run over labelings of $\{\sigma_1, \sigma_3\}$ and $\{\sigma_0, \sigma_2\}$, respectively. Similarly,

$$\Phi(\theta_{\{\tau_1\}}) = \left\{ \psi_{\tilde{\tau}_1}^{k_1+p-1} \psi_{\tilde{\tau}_1'}^{p-1} \psi_{\tilde{\tau}_2}^{k_2-2} \psi_{\tilde{\tau}_2'}^{p-1}, \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^0 \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^0, \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^{p-1} \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^{p-1}, \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^p \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^{-1} \right\}.$$

It is now clear that $\Phi(\sigma) \subset \bigcap_S \Phi(\theta_S) = \Phi(\theta_I) \cap \Phi(\theta_{\{\tau_1\}}) = \left\{ \psi_{\tilde{\tau}_1}^{k_1-1} \psi_{\tilde{\tau}_1'}^0 \psi_{\tilde{\tau}_2}^{k_2-1} \psi_{\tilde{\tau}_2'}^0 \right\}$, as desired.

Now suppose $k_1 = p + 1$. Then θ_S exists only if $S = I$ or $S = \emptyset$. It is easy to see from (17) that $\Phi(\theta_\emptyset) = \Phi(\theta_I)$. Hence $\bigcap_S \Phi(\theta_S) = \Phi(\theta_I)$, which is larger than we want. The following result, whose proof is independent of everything we have done here, is Corollary 8.4 of [Sch2].

Proposition 5.9. *Let $\theta : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$, and let $\rho : \text{Gal}(\overline{F}/F) \rightarrow \text{GL}_2(\overline{\mathbb{F}}_p)$ be such that $\rho|_{G_{\mathfrak{p}}}$ is irreducible and ρ is modular of weight $\sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$, where $\sigma_{\mathfrak{p}}$ is an irreducible quotient of $V = \text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta$, i.e. the constituent corresponding to $S = \emptyset$ in the notation of Lemma 3.7. Then there exists a constituent $\sigma'_{\mathfrak{p}} \neq \sigma_{\mathfrak{p}}$ of V such that ρ is modular of weight $\sigma'_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$.*

Hence if ρ is modular of weight $\sigma_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$, where $\sigma_{\mathfrak{p}} = (\text{Sym}^{p-1} k_{\mathfrak{p}}^2 \otimes_{\tau_1} \overline{\mathbb{F}}_p) \otimes (\text{Sym}^{k_2-2} k_{\mathfrak{p}}^2 \otimes_{\tau_2} \overline{\mathbb{F}}_p)$, then it is also modular of weight $\sigma'_{\mathfrak{p}} \otimes \sigma^{\mathfrak{p}}$, where $\sigma'_{\mathfrak{p}}$ is one of the other constituents of $\text{Ind}_{\Gamma_0(\mathfrak{p})}^{\text{GL}_2(k_{\mathfrak{p}})} \theta_\emptyset$. By Lemma 3.7, there are only two possibilities for $\sigma'_{\mathfrak{p}}$:

$$\begin{aligned} & (\det^{p-1} \text{Sym}^0 k_{\mathfrak{p}}^2 \otimes_{\tau_1} \overline{\mathbb{F}}_p) \otimes (\det^{k_2-2} \text{Sym}^{p+1-k_2} k_{\mathfrak{p}}^2 \otimes_{\tau_2} \overline{\mathbb{F}}_p) \\ & (\text{Sym}^{p-2} k_{\mathfrak{p}}^2 \otimes_{\tau_1} \overline{\mathbb{F}}_p) \otimes (\det^{k_2-1} \text{Sym}^{p-k_2} k_{\mathfrak{p}}^2 \otimes_{\tau_2} \overline{\mathbb{F}}_p) \end{aligned}$$

If $k_2 \neq 2$, then all exponents of symmetric powers in these weights are strictly less than $p - 1$, so we know that Theorem 1.3 holds for them. Thus ϕ must have either the form $\prod \psi_{\tau_1}^{p-1} \psi_{\tau_1}^p \psi_{\tau_2}^{k_2-2} \psi_{\tau_2}^p$ or the form $\prod \psi_{\tau_1}^{p-1} \psi_{\tau_1}^0 \psi_{\tau_2}^{k_2-1} \psi_{\tau_2}^0$ and it is easy to see that once we get these characters into the form of the statement of Theorem 1.3, we cannot have $\{m_2, m_4\} = \{k_2 - 2, 1\}$. If $k_2 = 2$, then $\{k_2 - 2, 1\} = \{k_2 - 1, 0\}$.

Remark 5.10. This proves Theorem 1.3 when $s = 2$ and k_1, k_2 are not both 2 or $p + 1$. Observe that the method of intersecting the $\Phi(\theta_S)$ broke down when $k_1 = p + 1$, and we had to rely on Proposition 5.9. A similar problem occurs in general when $k_l \in \{p, p + 1\}$ for all $l \neq j$, when we are trying to rule out $\{m_j, m_{j+s}\} = \{k_j - 2, 1\}$. This is the reason for the condition (2b) in the definition of $W_{\mathfrak{p}}(\rho)$.

5.3. The general case.

Proposition 5.11. *Suppose that $2 \leq k_\tau \leq p$ for all τ , and that the k_τ are not all 2. Then the statement of Theorem 1.3 holds for σ .*

Proof. We may assume $s \geq 3$. As before, we compute $\bigcap_S \Phi(\theta_S)$. Since $k_\tau \neq p + 1$ for all τ , it follows that θ_S exists for all $S \subset I$. By Remark 5.8, $\bigcap_S \Phi(\theta_S)$ contains all ϕ of the form $\prod_\tau \psi_{\tilde{\tau}}^{k_\tau-1} \psi_{\tilde{\tau}'}^0$. We claim it contains nothing else. Indeed, suppose that $\phi = \prod_j \psi_j^{m_j} \in \bigcap_S \Phi(\theta_S)$ is a character not of this form. Then $\phi \in \Phi(\theta_I)$, so there is some i such that $k_i \neq 2$ and $\{m_i, m_{i+s}\} = \{k_i - 2, 1\}$.

Let $S = \{\tau_i\}$. We assert that $\phi \notin \Phi(\theta_S)$, which will contradict $\phi \in \bigcap_{\theta \in \Theta(\sigma_{\mathfrak{p}})} \Phi(\theta)$. For every $T \subset I$, we claim ϕ is not one of the elements of $\Phi(\theta_S)$ corresponding to T in (17) above. These are:

$$\phi = \prod_{\substack{\tau \in T \\ \tau \neq \tau_i}} \psi_{\tilde{\tau}}^{k_\tau-1-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{p-\nu_S(\tau)} \prod_{\substack{\tau \notin T \\ \tau \neq \tau_i}} \psi_{\tilde{\tau}}^{k_\tau-1-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{-\nu_S(\tau)} \times \begin{cases} \psi_{\tilde{\tau}_i}^{k_i+p-1} \psi_{\tilde{\tau}_i'}^{p-\nu_T(\tau_i)} & : \tau_i \in T \\ \psi_{\tilde{\tau}_i}^{k_i-1} \psi_{\tilde{\tau}_i'}^{p-\nu_T(\tau_i)} & : \tau_i \notin T \end{cases}$$

Suppose that $\tilde{\tau}_i = \sigma_\alpha$, where $\alpha \in \{i, i + s\}$. There is a unique way to write ϕ such that all the exponents are between 0 and $p - 1$. It is clear that this redistribution of exponents can cause the residue mod p of each exponent to change by at most one. Since $3 \leq k_i \leq p$ and $\nu_S(\tau) = 0$

for $\tau \neq \tau_{i-1}$, we see that we must have $\tau_{i-1} \notin T$ and $\tilde{\tau}'_{i-1} = \sigma_{\alpha+s-1}$, and that the exponents of $\psi_{\alpha+1}, \psi_{\alpha+2}, \dots, \psi_{\alpha+s-2}$ in the displayed expression above must all be 0. This can only happen if for all $k \in \{i+1, \dots, i+s-2\}$ (note this set is non-empty since $s > 2$) we have $\tau_k \notin T$. So in particular $\nu_T(\tau_i) = 0$, whence the exponent of $\psi_{\alpha+s}$ is p . But $\psi_{\alpha+s-1}^{-1} \psi_{\alpha+s}^p = 1$, so after the redistribution of exponents, the exponent of ψ_α will not decrease after all. Hence $\phi \notin \Phi(\theta_S)$. \square

Finally we consider the general case, where $2 \leq k_i \leq p+1$ for all i , but the k_i are not all $p+1$ and not all 2. Suppose that $\phi \in \Phi(\theta_I)$, but ϕ does not have the form $\prod_\tau \psi_{\tilde{\tau}}^{k_\tau-1} \psi_{\tilde{\tau}'}^0$. As we saw above, there is some τ_i such that $\{m_i, m_{i+s}\} = \{k_i - 2, 1\}$. Suppose $k_i \neq p+1$. Now there may be more than one way to “redistribute” the exponents of ϕ , since we allow exponents to be p . If $k_{i-1} \neq p+1$, then set $S = \{\tau_i\}$. Then θ_S exists and the previous argument works to show that $\phi \notin \Phi(\theta_S)$.

Alternatively, suppose that there exists some $1 \leq m \leq s-1$ such that $k_{i-m} \neq p, p+1$. Then we may set $S = \{\tau_{i-m+1}, \dots, \tau_{i-1}, \tau_i\}$ and use a variant of the above argument. Observe that θ_S exists, and $\Phi(\theta_S)$ consists of:

$$\begin{aligned} \phi = & \prod_{\substack{\tau \in T \\ \tau \notin S}} \psi_{\tilde{\tau}}^{k_\tau-1-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{p-\nu_S(\tau)} \prod_{\substack{\tau \notin T \\ \tau \notin S}} \psi_{\tilde{\tau}}^{k_\tau-1-\nu_T(\tau)} \psi_{\tilde{\tau}'}^{-\nu_S(\tau)} \times \\ & \prod_{\substack{\tau \in T \\ \tau \in S \setminus \{\tau_i\}}} \psi_{\tilde{\tau}}^{k_\tau+p-2} \psi_{\tilde{\tau}'}^{p-\nu_T(\tau)} \prod_{\substack{\tau \notin T \\ \tau \in S \setminus \{\tau_i\}}} \psi_{\tilde{\tau}}^{k_\tau-2} \psi_{\tilde{\tau}'}^{p-\nu_T(\tau)} \times \begin{cases} \psi_{\tilde{\tau}_i}^{k_i+p-1} \psi_{\tilde{\tau}'_i}^{p-\nu_T(\tau_i)} & : \tau_i \in T \\ \psi_{\tilde{\tau}_i}^{k_i-1} \psi_{\tilde{\tau}'_i}^{p-\nu_T(\tau_i)} & : \tau_i \notin T \end{cases} \end{aligned}$$

As above, let $\tilde{\tau}_i = \sigma_\alpha$, and we see that we must have $\sigma_{\alpha+s-m} = \tilde{\tau}'_{i-m}$ (with exponent -1 in the above expression) and that all the exponents between them, i.e. the exponents of $\sigma_{\alpha+1}, \dots, \sigma_{\alpha+s-m-1}$ must be 0. When we redistribute exponents to bring ϕ into the form of (2), if the exponent of σ_α turns into $m_\alpha = k_\tau - 2$, then the exponent -1 of $\sigma_{\alpha+s-m}$ will turn into $p-1$. But then $p-1$ must be either $k_{i-m} - 1$ or $k_{i-m} - 2$, whence k_{i-m} is either p or $p+1$, which we assumed not to happen. Hence $\phi \notin \Phi(\theta_S)$.

The only case that has not been treated (for non-trivial θ) is when $k_i \neq p+1$, but $k_j \in \{p, p+1\}$ for all $j \neq i$. This is the reason for condition (2b) in the statement of Theorem 1.3.

5.4. The case of trivial θ . Let $\mathbf{1} : \Gamma_0(\mathfrak{p}) \rightarrow \overline{\mathbb{F}}_p^*$ be the trivial character. Then we see from Lemma 3.7 that the $\mathrm{GL}_2(k_{\mathfrak{p}})$ -module $\mathrm{Ind}_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})} \mathbf{1}$ has only two Jordan-Hölder constituents:

$$0 \rightarrow \bigotimes_{\tau \in I} \mathrm{Sym}^0 k_{\mathfrak{p}}^2 \otimes_{\tau} \overline{\mathbb{F}}_p \rightarrow \mathrm{Ind}_{\Gamma_0(\mathfrak{p})}^{\mathrm{GL}_2(k_{\mathfrak{p}})} \mathbf{1} \rightarrow \bigotimes_{\tau \in I} \mathrm{Sym}^{p-1} k_{\mathfrak{p}}^2 \otimes_{\tau} \overline{\mathbb{F}}_p \rightarrow 0.$$

We write $\sigma_{\mathfrak{p}}$ and $\sigma'_{\mathfrak{p}}$ for the submodule and quotient above, respectively. Suppose that $\rho : \mathrm{Gal}(\overline{F}/F) \rightarrow \mathrm{GL}_2(\overline{\mathbb{F}}_p)$ is modular of weight $\sigma = \sigma_{\mathfrak{p}} \otimes \sigma'_{\mathfrak{p}}$, where $\sigma_{\mathfrak{p}}$ is as above. The argument is similar to the one for non-trivial θ , but is much easier. There exists an eigenclass $f \in H^1(M_{0,H} \otimes \overline{F}, \overline{\mathbb{F}}_p)_{\mathfrak{m}}$ corresponding to ρ , where $\mathfrak{m} \subset \mathbf{T}$ is the appropriate maximal ideal. As before, we have $\mathbf{T}_{\mathfrak{p}} = \mathbf{T}_{\mathfrak{m}} \times \mathbf{T}'$; let ε be an idempotent such that $\varepsilon \mathbf{T}_{\mathfrak{p}} = \mathbf{T}_{\mathfrak{m}}$.

Recall that $\mathbf{M}_{0,H}$ has good reduction. Let J be the Néron model over D of $\mathrm{Pic}^0(M_{0,H})$. Define a p -divisible group G/D by $G_n = J^0[p^n]$; the G_n are finite and flat. Then we see as in Lemma 4.3 that $G[\mathfrak{m}]_K$ is isomorphic to some number of copies of $W_K = W \times_F K$, where W is the \mathbf{T}/\mathfrak{m} -vector

space scheme over F corresponding to ρ^{ss} . We find a $\mathbb{F}_{p^{2s}}$ -vector space subscheme $V \subset G[\mathfrak{m}] \otimes_{\mathbb{F}} \mathbb{F}'$ satisfying condition (**), such that $\text{Gal}(\overline{K}/K)$ acts on V_K by the character ϕ .

Clearly $\text{Gal}(K'/K)$ acts trivially on $\text{Pic}^0(\mathbf{M}_{0,H}) \otimes D$ and hence on the special fiber of $G[\mathfrak{m}]_{D'}$ and on its cotangent space. Thus $b_i = 0$ for all i . But then $a_i \in \{0, 1\}$ for all i , and from Lemma 5.1 we see that $(\phi\phi^q)|_{I_p} = \chi|_{I_p}$. Hence ϕ must have the form $\phi = \prod_{\tau} \psi_{\tau}^1 \psi_{\tau'}^0$.

Now suppose that ρ is modular of weight $\sigma' = \sigma'_p \otimes (\sigma')^p$, where σ' is as defined above. It follows from Proposition 5.9 that ρ is also modular of weight $\sigma_p \otimes (\sigma')^p$, and so ϕ must be of the form above. But this is equivalent to ϕ having the form $\phi = \prod_{\tau} \psi_{\tau}^p \psi_{\tau'}^0$.

The proof of Theorem 1.3 is now complete.

REFERENCES

- [SGA] *Groupes de monodromie en géométrie algébrique. I.* Springer-Verlag, Berlin, 1972. Séminaire de Géométrie Algébrique du Bois-Marie 1967–1969 (SGA 7 I), Dirigé par A. Grothendieck. Avec la collaboration de M. Raynaud et D. S. Rim, Lecture Notes in Mathematics, Vol. 288.
- [BL] L. Barthel and R. Livné. Irreducible modular representations of GL_2 of a local field. *Duke Math. J.* **75**(1994), 261–292.
- [BLR1] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. *Néron models*, volume 21 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1990.
- [BLR2] Nigel Boston, Hendrik W. Lenstra, Jr., and Kenneth A. Ribet. Quotients of group rings arising from two-dimensional representations. *C. R. Acad. Sci. Paris Sér. I Math.* **312**(1991), 323–328.
- [BDJ] Kevin Buzzard, Fred Diamond, and Frazer Jarvis. On Serre’s conjecture for mod l Galois representations over totally real fields. *Preprint* (2005).
- [Car1] Henri Carayol. Sur la mauvaise réduction des courbes de Shimura. *Compositio Math.* **59**(1986), 151–230.
- [Car2] Henri Carayol. Sur les représentations l -adiques associées aux formes modulaires de Hilbert. *Ann. Sci. École Norm. Sup. (4)* **19**(1986), 409–468.
- [Che] Claude Chevalley. Deux théorèmes d’arithmétique. *J. Math. Soc. Japan* **3**(1951), 36–44.
- [DS] Pierre Deligne and Jean-Pierre Serre. Formes modulaires de poids 1. *Ann. Sci. École Norm. Sup. (4)* **7**(1974), 507–530 (1975).
- [Dia] Fred Diamond. A correspondence between representations of local Galois groups and Lie-type groups. *Preprint* (2005).
- [DT] Fred Diamond and Richard Taylor. Nonoptimal levels of mod l modular representations. *Invent. Math.* **115**(1994), 435–462.
- [Edi1] Bas Edixhoven. The weight in Serre’s conjectures on modular forms. *Invent. Math.* **109**(1992), 563–594.
- [Edi2] Bas Edixhoven. Serre’s conjecture. In *Modular forms and Fermat’s last theorem (Boston, MA, 1995)*, pages 209–242. Springer, New York, 1997.
- [Gee1] Toby Stephen Gee. Companion forms over totally real fields. *Preprint* (2004).
- [Gee2] Toby Stephen Gee. On weights of mod p Hilbert modular forms. *Preprint* (2005).
- [Gro] Benedict H. Gross. A tameness criterion for Galois representations associated to modular forms (mod p). *Duke Math. J.* **61**(1990), 445–517.
- [HT] Michael Harris and Richard Taylor. *The geometry and cohomology of some simple Shimura varieties*, volume 151 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2001. With an appendix by Vladimir G. Berkovich.
- [Hid] Haruzo Hida. On p -adic Hecke algebras for GL_2 over totally real fields. *Ann. of Math. (2)* **128**(1988), 295–384.
- [Jar1] Frazer Jarvis. Level lowering for modular mod l representations over totally real fields. *Math. Ann.* **313**(1999), 141–160.
- [Jar2] Frazer Jarvis. Mazur’s principle for totally real fields of odd degree. *Compositio Math.* **116**(1999), 39–79.
- [Kas] Payman L. Kassaei. \mathcal{P} -adic modular forms over Shimura curves over totally real fields. *Compos. Math.* **140**(2004), 359–395.

- [KM] Nicholas M. Katz and Barry Mazur. *Arithmetic moduli of elliptic curves*, volume 108 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1985.
- [Mil] J. S. Milne. *Étale cohomology*, volume 33 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1980.
- [Neu] Jürgen Neukirch. *Algebraic number theory*, volume 322 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.
- [Ray] Michel Raynaud. Schémas en groupes de type (p, \dots, p) . *Bull. Soc. Math. France* **102**(1974), 241–280.
- [Sch1] Michael M. Schein. Weights in Serre’s conjecture for Hilbert modular forms: the ramified case. *Preprint* (2006).
- [Sch2] Michael M. Schein. *Weights of Galois representations associated to Hilbert modular forms*. PhD thesis, Harvard University, 2006.
- [Ser] Jean-Pierre Serre. Propriétés galoisiennes des points d’ordre fini des courbes elliptiques. *Invent. Math.* **15**(1972), 259–331.
- [Sha] Stephen S. Shatz. Group schemes, formal groups, and p -divisible groups. In *Arithmetic geometry (Storrs, Conn., 1984)*, pages 29–78. Springer, New York, 1986.
- [Tay] Richard Taylor. On Galois representations associated to Hilbert modular forms. *Invent. Math.* **98**(1989), 265–280.
- [Wat] William C. Waterhouse. *Introduction to affine group schemes*, volume 66 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1979.

E-mail address: `mschein@math.harvard.edu`

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, 1 OXFORD STREET, CAMBRIDGE, MA 02138 USA