SERRE WEIGHTS FOR LOCALLY REDUCIBLE TWO-DIMENSIONAL GALOIS REPRESENTATIONS

FRED DIAMOND AND DAVID SAVITT

Abstract. Let F be a totally real field, and v a place of F dividing an odd prime p. We study the weight part of Serre's conjecture for continuous totally odd representations $\overline{\rho}: G_F \to \mathrm{GL}_2(\overline{\mathbb{F}}_p)$ that are reducible locally at v. Let W be the set of predicted Serre weights for the semisimplification of $\overline{\rho}|_{G_{F_v}}$. We prove that when $\overline{\rho}|_{G_{F_v}}$ is generic, the Serre weights in W for which $\overline{\rho}$ is modular are exactly the ones that are predicted (assuming that $\overline{\rho}$ is modular). We also determine precisely which subsets of W arise as predicted weights when $\overline{\rho}|_{G_{F_v}}$ varies with fixed generic semisimplification.

Introduction

Let F be a totally real field and $\overline{\rho} : G_F \to \mathrm{GL}_2(\overline{\mathbb{F}}_p)$ a continuous totally odd representation. Suppose that $\overline{\rho}$ is automorphic in the sense that it arises as the reduction of a p-adic representation of G_F associated to a cuspidal Hilbert modular eigenform, or equivalently to a cuspidal holomorphic automorphic representation of $\mathrm{GL}_2(\mathbb{A}_F)$.

The weight part of Serre's Conjecture in this context was formulated in increasing generality by Buzzard, Jarvis and one of the authors [5], Schein [28] and Barnet–Lamb, Gee and Geraghty [1] (see also [14]). The structure of the statement is as follows: Let v be a prime of F dividing p, and let k denote its residue field. A *Serre weight* is then an irreducible representation of $\mathrm{GL}_2(k)$ over $\overline{\mathbb{F}}_p$. One can then define what it means for $\overline{\rho}$ to be modular of a given (Serre) weight, depending *a priori* on the choice of a suitable quaternion algebra over F, and we let $W^v_{\mathrm{mod}}(\overline{\rho})$ denote the set of weights at v for which $\overline{\rho}$ is modular. On the other hand one can define a set of weights $W_{\mathrm{expl}}(\overline{\rho})$ that depends only on $\overline{\rho}_v = \overline{\rho}|_{G_{F_v}}$, and the conjecture states that $W^v_{\mathrm{mod}}(\overline{\rho}) = W_{\mathrm{expl}}(\overline{\rho}_v)$.

A series of papers by Gee and coauthors [15, 20, 18, 1, 17, 16] proves the following, under mild technical hypotheses on $\overline{\rho}$:

- $W^v_{\mathrm{mod}}(\overline{\rho})$ depends only on $\overline{\rho}_v$;
- $W_{\mathrm{expl}}(\overline{\rho}_v) \subseteq W^v_{\mathrm{mod}}(\overline{\rho})$;
- $W_{\mathrm{expl}}(\overline{\rho}_v) = W^v_{\mathrm{mod}}(\overline{\rho})$ if F_v is unramified or totally ramified over $\mathbb{Q}_p$.

In this paper we study the reverse inclusion $W^v_{\mathrm{mod}}(\overline{\rho}) \subseteq W_{\mathrm{expl}}(\overline{\rho}_v)$ when $\overline{\rho}_v$ is reducible and F_v is an arbitrary finite extension of $\mathbb{Q}_p$ (i.e. not necessarily either unramified or totally ramified). One often refers to this inclusion as the problem

2010 *Mathematics Subject Classification.* 11F80.

Key words and phrases. Galois representations, Serre's conjecture, Breuil modules.

The first author was partially supported by Leverhulme Trust RPG-2012-530; the second author was partially supported by NSF grant DMS-0901049 and NSF CAREER grant DMS-1054032.

of "weight elimination," since one wishes to eliminate weights not in $W_{\mathrm{expl}}(\overline{\rho}_v)$ as possible weights for $\overline{\rho}$.

Suppose that $\overline{\rho}_v$ has the form

$$\begin{pmatrix} \chi_2 & * \\ 0 & \chi_1 \end{pmatrix}.$$

Then the set $W_{\mathrm{expl}}(\overline{\rho}_v)$ is a subset of $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$ which depends on the associated extension class $c_{\overline{\rho}_v} \in H^1(G_{F_v}, \overline{\mathbb{F}}_p(\chi_2\chi_1^{-1}))$. Assume that $\overline{\rho}$ satisfies the hypotheses of [16], as well as a certain genericity hypothesis (a condition on $\chi_2\chi_1^{-1}$; see Definition 3.5 for a precise statement). Then our main global result is the following.

Theorem A. *$W_{\mathrm{expl}}(\overline{\rho}_v) = W^v_{\mathrm{mod}}(\overline{\rho}) \cap W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$.*

In other words we prove under these hypotheses that weight elimination, and so also the weight part of Serre's Conjecture, holds for weights in $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$.

While the set $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$ is completely explicit, the dependence of $W_{\mathrm{expl}}(\overline{\rho}_v)$ on the extension class is given in terms of the existence of reducible crystalline lifts of $\overline{\rho}_v$ with prescribed Hodge–Tate weights. In particular it is not clear which subsets of $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$ arise as the extension class $c_{\overline{\rho}_v}$ varies. Another purpose of the paper is to address this question, which we resolve in the case where $\overline{\rho}_v$ is generic. These local results indicate a structure on the sets $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$. This structure should reflect properties of a mod p local Langlands correspondence in this context, in the sense that the set $W_{\mathrm{expl}}(\overline{\rho}_v)$ is expected to determine the $\mathrm{GL}_2(\mathcal{O}_{F_v})$-socle of $\pi(\overline{\rho}_v)$, the $\mathrm{GL}_2(F_v)$-representation associated to $\overline{\rho}_v$ by that correspondence.

To simplify the statement slightly for the introduction, we assume (in addition to genericity) that the restriction of $\chi_2\chi_1^{-1}$ to the inertia subgroup of G_{F_v} is not the cyclotomic character or its inverse. In particular this implies that $H^1(G_{F_v}, \overline{\mathbb{F}}_p(\chi_2\chi_1^{-1}))$ has dimension $[F_v : \mathbb{Q}_p] = ef$, where $f = [k : \mathbb{F}_p]$ and e is the absolute ramification degree of F_v. (We remark that this notation differs slightly from the notation in the body of the paper, where the ramification degree of F_v will be e'.) We shall define a partition of $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$ into subsets W_a indexed by the elements $a = (a_0, a_1, \ldots, a_{f-1})$ of $A = \{0, 1, \ldots, e\}^f$, and a subspace $L_a \subseteq H^1(G_{F_v}, \overline{\mathbb{F}}_p(\chi_2\chi_1^{-1}))$ of codimension $\sum_{i=0}^{f-1} a_i$ for each $a \in A$. We give the set A the usual (product) partial ordering.

Our main local result is the following.

Theorem B. *Suppose that $\sigma \in W_a$. Then $\sigma \in W_{\mathrm{expl}}(\overline{\rho}_v)$ if and only if $c_{\overline{\rho}_v} \in L_a$. Moreover there exists $b \in A$ (depending on $\overline{\rho}_v$) such that*

$$W_{\mathrm{expl}}(\overline{\rho}_v) = \coprod_{a \le b} W_a.$$

In other words the weights come in packets, where the packets arise in a hierarchy compatible with the partial ordering on A. In connection with the hypothetical mod p local Langlands correspondence mentioned above, Theorem B is consistent with the possibility that the associated $\mathrm{GL}_2(F_v)$-representation $\pi(\overline{\rho}_v)$ is equipped with an increasing filtration of length $[F_v : \mathbb{Q}_p] + 1$ such that $\mathrm{gr}^{\bullet}(\pi(\overline{\rho}_v)) \cong \pi(\overline{\rho}_v^{\mathrm{ss}})$ and $\mathrm{gr}^m(\pi(\overline{\rho}_v))$ has $\mathrm{GL}_2(\mathcal{O}_{F_v})$-socle consisting of the weights in the union of the W_a with $\sum_{i=0}^{f-1} a_i = m$ (cf. [4, Thm. 19.9]).

We now briefly indicate how our constructions and proofs proceed. The set W_a is defined using the reduction of a certain tamely ramified principal series type θ_a, and the space L_a is defined using Breuil modules with descent data corresponding to θ_a. In the first three sections of the paper, we show that the spaces L_a have the (co-)dimension claimed above, and that they satisfy $L_a \cap L_{a'} = L_{a''}$ where $a''_i = \max\{a_i, a'_i\}$. Section 1 contains a general analysis of the extensions of rank one Breuil modules. In Section 2 we define and study the extension spaces L_a, and in Section 3 we describe our subsets W_a of $W_{\mathrm{expl}}(\overline{\rho}_v^{\mathrm{ss}})$.

Having done the local analysis, the strategy for the proving the main results is similar to that of Gee, Liu and one of the authors [18] in the totally ramified case; in particular global arguments play a role in proving the local results. More precisely in Section 4 we prove the following three conditions are equivalent for each weight $\mu \in W_a$:

(1) $\mu \in W_{\mathrm{expl}}(\overline{\rho}_v)$;
(2) $\mu \in W^v_{\mathrm{mod}}(\overline{\rho})$;
(3) $c_{\overline{\rho}_v} \in L_a$.

The implication (1) $\Rightarrow$ (2) is already proved by Gee and Kisin [16], and (2) $\Rightarrow$ (3) is proved by showing that $\overline{\rho}_v$ has a potentially Barsotti–Tate lift of type θ_a. Having now proved that (1) $\Rightarrow$ (3), one deduces that L_a contains the relevant spaces of extensions with reducible crystalline lifts; equality follows on comparing dimensions, and this gives (3) $\Rightarrow$ (1).

The reason our results are not as definitive as those of [18] is that in the totally ramified case there is a tight connection between being modular of some Serre weight and having a potentially Barsotti-Tate lift of a certain type: in the totally ramified case the reduction mod p of the principal series type θ_a has at most two Jordan–Hölder factors, while in general it can have many more.

In fact, when F_v is allowed to be arbitrary some sort of hypothesis along the lines of genericity is necessary, in the sense that there exist F_v, χ_1, χ_2, and μ such that the subset of $H^1(G_{F_v}, \overline{\mathbb{F}}_p(\chi_2\chi_1^{-1}))$ corresponding to $\overline{\rho}_v$ with $\mu \in W_{\mathrm{expl}}(\overline{\rho}_v)$ is not equal to L_a for any choice of a. We give an example of this phenomenon in Section 5.

Finally, we must point out that some time after this paper was written, Gee, Liu, and the second author [19] announced a proof that $W_{\mathrm{expl}}(\overline{\rho}_v) = W^v_{\mathrm{mod}}(\overline{\rho})$ in general, thus improving on our Theorem A (by rather different methods). The arguments in [19] are entirely local, and depend on an extension to the ramified case of the p-adic Hodge theoretic results proved in [17] in the unramified case.

Acknowledgments. This collaboration was initiated during the Galois Trimester at the Institute Henri Poincaré in 2010, and the authors benefited from further opportunities for discussion during the LMS-EPSRC Durham Symposium (2011) and workshops at the Institute for Advanced Study (2011) and the Fields Institute (2012). Much of the research was carried out during a visit by DS to King's partly funded by an LMS grant in 2012, and a visit by FD to the Hausdorff Institute in 2013 was valuable in completing the work. The authors are grateful to these institutions and to the organizers of the above programs for their support, as well as to the anonymous referee for his or her comments and suggestions.

Notation and conventions. If M is a field, we let G_M denote its absolute Galois group. If M is a global field and v is a place of M, let M_v denote the completion of M at v. If M is a finite extension of $\mathbb{Q}_p$ for some p, we let M_0 denote the maximal unramified extension of $\mathbb{Q}_p$ contained in M, and we write I_M for the inertia subgroup of G_M.

Let p be an odd prime number. Let $K \supseteq L$ be finite extensions of $\mathbb{Q}_p$ such that K/L is a tame Galois extension. (These may be regarded as fixed, although at certain points in the paper we will make a specific choice for K.) Assume further that π is a uniformiser of $\mathcal{O}_K$ with the property that $\pi^{e(K/L)} \in L$, where $e(K/L)$ is the ramification index of the extension K/L. Let e, f and e', f' be the absolute ramification and inertial degrees of K and L respectively, and denote their residue fields by k and ℓ. From Section 1.3 onwards, $e(K/L)$ will always be divisible by $p^{f'} - 1$, and from Section 2.2 onwards we will have $f = f'$ and $e(K/L) = p^f - 1$. Write $\eta : \mathrm{Gal}(K/L) \to \mathcal{O}_K^\times$ for the function sending $g \mapsto g(\pi)/\pi$, and let $\overline{\eta} : \mathrm{Gal}(K/L) \to k^\times$ be the reduction of η modulo the maximal ideal of $\mathcal{O}_K$.

Our representations of G_L will have coefficients in $\overline{\mathbb{Q}}_p$, a fixed algebraic closure of $\mathbb{Q}_p$ whose residue field we denote $\overline{\mathbb{F}}_p$. Let E be a finite extension of $\mathbb{Q}_p$ contained in $\overline{\mathbb{Q}}_p$ and containing the image of every embedding of K into $\overline{\mathbb{Q}}_p$. Let $\mathcal{O}_E$ be the ring of integers in E, with uniformiser ϖ and residue field $k_E \subset \overline{\mathbb{F}}_p$. Note in particular that there exist f embeddings of k into k_E.

We write $\mathrm{Art}_L \colon L^\times \to W_L^{\mathrm{ab}}$ for the isomorphism of local class field theory, normalised so that uniformisers correspond to geometric Frobenius elements. For each $\sigma \in \mathrm{Hom}(\ell, \overline{\mathbb{F}}_p)$ we define the fundamental character ω_σ corresponding to σ to be the composite

$$I_L \longrightarrow \mathcal{O}_L^\times \longrightarrow \ell^\times \xrightarrow{\ \sigma\ } \overline{\mathbb{F}}_p^\times,$$

where the map $I_L \to \mathcal{O}_L^\times$ is induced by the restriction of Art_L^{-1}. Let ϵ denote the p-adic cyclotomic character and $\overline{\epsilon}$ the mod p cyclotomic character, so that $\prod_{\sigma \in \mathrm{Hom}(\ell, \overline{\mathbb{F}}_p)} \omega_\sigma^{e'} = \overline{\epsilon}$. We will often identify characters $I_L \to \overline{\mathbb{F}}_p^\times$ with characters $\ell^\times \to \overline{\mathbb{F}}_p^\times$ via the Artin map, as above, and similarly for their Teichmüller lifts.

Fix an embedding $\sigma_0 : k \hookrightarrow k_E$, and recursively define $\sigma_i : k \hookrightarrow k_E$ for all $i \in \mathbb{Z}$ so that $\sigma_{i+1}^p = \sigma_i$. We write ω_i for $\omega_{\sigma_i|_\ell}$. With these normalizations, if K/L is totally ramified of degree $e(K/L) = p^{f'} - 1$ then $\omega_i = (\sigma_i \circ \overline{\eta})|_{I_L}$.

We normalize Hodge–Tate weights so that all Hodge–Tate weights of the cyclotomic character are equal to 1. (See Definition 3.2 for further discussion of our conventions regarding Hodge–Tate weights.)

1. Extensions of Breuil modules

In the paper [2], Breuil classifies p-torsion finite flat group schemes over $\mathcal{O}_K$ in terms of semilinear-algebraic objects that have come to be known as Breuil modules. This classification has proved to be immensely useful, in part because Breuil modules are often amenable to explicit computation. In this section we make a careful study of the extensions between Breuil modules of rank one with coefficients and descent data. Many of these results are familiar, but the statements that we need are somewhat more general than those in the existing literature (*cf.* [3, 25, 6, 9]).

1.1. **Review of rank one Breuil modules.** We let ϕ denote the endomorphism of $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$ obtained by k_E-linearly extending the pth power map on $k[u]/u^{ep}$.

Define an action of $\mathrm{Gal}(K/L)$ on $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$ by the formula $g((a \otimes 1)u^i) = (g(a)\overline{\eta}(g)^i \otimes 1)u^i$, extended k_E-linearly.

Definition 1.1. The category of *Breuil modules with k_E-coefficients and generic fibre descent data from K to L*, denoted BrMod^K_{L,k_E}, is the category whose objects are quadruples $(\mathcal{M}, \mathrm{Fil}^1 \mathcal{M}, \phi_1, \{\widehat{g}\})$ where:

- $\mathcal{M}$ is a finitely generated free $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$-module,
- $\mathrm{Fil}^1 \mathcal{M}$ is a $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$-submodule of $\mathcal{M}$ containing $u^e\mathcal{M}$.
- $\phi_1 : \mathrm{Fil}^1 \mathcal{M} \to \mathcal{M}$ is a ϕ-semilinear map whose image generates $\mathcal{M}$ as a $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$-module,
- the maps $\widehat{g} : \mathcal{M} \to \mathcal{M}$ for each $g \in \mathrm{Gal}(K/L)$ are additive bijections that preserve $\mathrm{Fil}^1 \mathcal{M}$, commute with the ϕ_1-, and k_E-actions, and satisfy $\widehat{g}_1 \circ \widehat{g}_2 = \widehat{g_1 \circ g_2}$ for all $g_1, g_2 \in \mathrm{Gal}(K/L)$. Furthermore $\widehat{1}$ is the identity, and if $a \in (k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$, $m \in \mathcal{M}$ then $\widehat{g}(am) = g(a)\widehat{g}(m)$.

We will usually write $\mathcal{M}$ in place of $(\mathcal{M}, \mathrm{Fil}^1 \mathcal{M}, \phi_1, \{\widehat{g}\})$. A morphism $f : \mathcal{M} \to \mathcal{M}'$ in BrMod^K_{L,k_E} is a $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$-module homomorphism with $f(\mathrm{Fil}^1 \mathcal{M}) \subseteq \mathrm{Fil}^1 \mathcal{M}'$ that commutes with ϕ_1 and the descent data.

The category BrMod^K_{L,k_E} is equivalent to the category of finite flat group schemes over $\mathcal{O}_K$ together with a k_E-action and descent data on the generic fibre from K to L (see [2, 27]). This equivalence depends on the choice of uniformiser π. The covariant functor $T^L_{\mathrm{st},2}$ defined immediately before Lemma 4.9 of [26] associates to each object $\mathcal{M}$ of BrMod^K_{L,k_E} a k_E-representation of G_L, which we refer to as the *generic fibre* of $\mathcal{M}$.

Notation 1.2. We let $e_i \in k \otimes_{\mathbb{F}_p} k_E$ denote the idempotent satisfying $(x \otimes 1)e_i = (1 \otimes \sigma_i(x))e_i$ for all $x \in k$. Observe that $\phi(e_i) = e_{i+1}$. We adopt the convention that if $m_0, \dots, m_{f-1}$ are elements of some $(k \otimes k_E)$-module, then $\underline{m}$ denotes the sum $\sum_{i=0}^{f-1} m_i e_i$, as well as any inferrable variations of this notation: for instance if $r_0, \dots, r_{f-1}$ are integers then $u^{\underline{r}}$ denotes $\sum_{i=0}^{f-1} u^{r_i} e_i$. Conversely for any element written $\underline{a}$, we set $a_i = e_i\underline{a}$. When $\underline{a} \in (k \otimes k_E)[u]/u^{ep}$ we will generally identify a_i with its preimage in $k_E[u]/u^{ep}$ under the the map $k_E[u]/u^{ep} \simeq e_i((k \otimes k_E)[u]/u^{ep})$ sending $x \mapsto e_i x$.

The rank one objects of BrMod^K_{L,k_E} are classified as follows.

Lemma 1.3. *Every rank one object of* BrMod^K_{L,k_E} *has the form:*

- $\mathcal{M} = ((k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}) \cdot m$,
- $\mathrm{Fil}^1 \mathcal{M} = u^{\underline{r}}\mathcal{M}$,
- $\phi_1(u^{\underline{r}}m) = \underline{a}m$ *for some* $\underline{a} \in (k \otimes_{\mathbb{F}_p} k_E)^\times$*, and*
- $\widehat{g}(m) = (\overline{\eta}(g)^{\underline{c}} \otimes 1)m$ *for all* $g \in \mathrm{Gal}(K/L)$,

where $r_i \in \{0 \dots, e\}$ *and* $c_i \in \mathbb{Z}/(e(K/L))$ *are sequences that satisfy* $c_{i+1} \equiv p(c_i + r_i) \pmod{e(K/L)}$*, and the sequences* r_i, c_i, a_i *are each periodic with period dividing* f'.

Proof. This is a special case of [27, Thm. 3.5]. In the notation of that item we have $D = f'$ because of our assumption that k embeds into k_E, and the periodicity of the sequence a_i is equivalent to $\underline{a} \in (\ell \otimes_{\mathbb{F}_p} k_E)^\times$. □

Notation 1.4. We will denote a rank one Breuil module as in Lemma 1.3 by $\mathcal{M}(\underline{r}, \underline{a}, \underline{c})$, or else (for reasons of typographical aesthetics) by $\mathcal{M}(r, a, c)$.

We wish to consider maps between rank one Breuil modules, but before we do so, we note the following elementary lemma.

Lemma 1.5. *Let* $\mathrm{Gal}(K/L)$ *act on* $(k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$ *by* $g \cdot x = (\overline{\eta}(g)^{\underline{w}} \otimes 1)g(x)$, *where* $g(x)$ *denotes the usual action and* $\{w_i\}$ *is a sequence of integers that is periodic with period dividing* f'. *The* $\mathrm{Gal}(K/L)$*-invariants of this action are the elements* $\underline{x} \in (k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$ *such that each nonzero term of* x_i *has degree congruent to* $-w_i \pmod{e(K/L)}$, *and the sequence* x_i *is periodic with period dividing* f'.

Proof. There exists $g \in \mathrm{Gal}(K/L)$ such that $\overline{\eta}(g) = 1$ and the image of g generates $\mathrm{Gal}(k/\ell)$; since $g(e_i) = e_{i+f'}$, the equality $g \cdot \underline{x} = \underline{x}$ shows that x_i is periodic with period dividing f'. Consideration of the inertia group $I(K/L)$ gives the conditions on the degrees of nonzero terms. □

The following lemma is standard, but its setting is slightly more general than that of existing statements in the literature (*cf.* [25, Lem. 6.1], [9, Prop. 2.5]).

Lemma 1.6. *Let* $\mathcal{M} = \mathcal{M}(r, a, c)$ *and* $\mathcal{N} = \mathcal{M}(s, b, d)$ *be rank one Breuil modules as above. Define* $\alpha_i = p(p^{f-1}r_i + \cdots + r_{i+f-1})/(p^f - 1)$ *and* $\beta_i = p(p^{f-1}s_i + \cdots + s_{i+f-1})/(p^f - 1)$ *for all* i. *There exists a nonzero map* $\mathcal{M} \to \mathcal{N}$ *if and only if*

- $\beta_i - \alpha_i \in \mathbb{Z}_{\geq 0}$ *for all* i,
- $\beta_i - \alpha_i \equiv c_i - d_i \pmod{e(K/L)}$ *for all* i, *and*
- $\prod_{i=0}^{f'-1} a_i = \prod_{i=0}^{f'-1} b_i$.

Proof. A nonzero morphism $\mathcal{M} \to \mathcal{N}$ must have the form $m \mapsto \underline{\delta} u^{\underline{z}} n$ for some integers $z_i \geq 0$ and some $\underline{\delta} \in ((k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep})^\times$. For this map to preserve the filtrations, it is necessary and sufficient that $r_i + z_i \geq s_i$ for all i. For the map to commute with ϕ_1 it is necessary and sufficient that

$$\phi(\underline{\delta} u^{(\underline{z}+\underline{r}-\underline{s})})\underline{b} = \underline{\delta}\underline{a}u^{\underline{z}}.$$

It follows from this equation that $\underline{\delta} \in (k \otimes_{\mathbb{F}_p} k_E)^\times$, that $z_{i+1} = p(z_i + r_i - s_i)$ for all i, and that $\phi(\underline{\delta})/\underline{\delta} = \underline{a}/\underline{b}$. The unique solution to the system of equations for the z_i's is precisely $z_i = \beta_i - \alpha_i$ for all i. Note that the positivity of z_{i+1} is equivalent to the condition $r_i + z_i \geq s_i$.

For the map to commute with descent data, it is necessary and sufficient that $g(\underline{\delta}) = (\overline{\eta}(g)^{\underline{c}-\underline{z}-\underline{d}} \otimes 1)\underline{\delta}$ for all $g \in \mathrm{Gal}(K/L)$. By Lemma 1.5, and recalling that $\underline{\delta}$ has no non-constant terms, this is satisfied if and only if $z_i \equiv c_i - d_i \pmod{e(K/L)}$ for all i and the sequence $\delta_i \in k_E$ is periodic with period dividing f'. Finally, it is easy to check that there exists $\underline{\delta} \in (k \otimes_{\mathbb{F}_p} k_E)^\times$ with $\phi(\underline{\delta})/\underline{\delta} = \underline{a}/\underline{b}$ and having the necessary periodicity if and only if $\prod_{i=0}^{f'-1} a_i = \prod_{i=0}^{f'-1} b_i$. □

Remark 1.7. Suppose that $e(K/L)$ is divisible by $p^{f'} - 1$. By [27, Rem. 3.6] it is then automatic that the α_i and β_i of the preceding lemma are integers. Combining Lemma 1.6 with [21, Cor. 4.3] we see in this case that there exists a nonzero map $\mathcal{M} \to \mathcal{N}$ if and only if $T^L_{\mathrm{st},2}(\mathcal{M}) \simeq T^L_{\mathrm{st},2}(\mathcal{N})$ and $\beta_i \geq \alpha_i$ for all i.

We will use the notation $\alpha_i = p(p^{f-1}r_i + \cdots + r_{i+f-1})/(p^f - 1)$ throughout the paper, and similarly for β_i. Let us write $\mathrm{Nm}(\underline{a}) = \prod_{i=0}^{f'-1} a_i \in k_E$. The following is immediate from (the proof of) Lemma 1.3.

Corollary 1.8. *We have* $\mathcal{M}(r, a, c) \simeq \mathcal{M}(r', a', c')$ *if and only if* $r_i = r'_i$ *for all* i, $c_i = c'_i$ *for all* i, *and* $\mathrm{Nm}(\underline{a}) = \mathrm{Nm}(\underline{a}')$.

The following proposition is again standard, but slightly more general than the versions in the existing literature ([9, Prop. 2.6], [6, Prop. 5.6]).

Proposition 1.9. *Let $\mathcal{M} = \mathcal{M}(r, a, c)$ and $\mathcal{N} = \mathcal{M}(s, b, d)$ be rank one Breuil modules as above. There exists a rank one Breuil module $\mathcal{P}$ and a pair of nonzero maps $\mathcal{M} \to \mathcal{P}$ and $\mathcal{N} \to \mathcal{P}$ if and only if*

- $\beta_i - \alpha_i \in \mathbb{Z}$ *for all* i,
- $\beta_i - \alpha_i \equiv c_i - d_i \pmod{e(K/L)}$ *for all* i, *and*
- $\prod_{i=0}^{f'-1} a_i = \prod_{i=0}^{f'-1} b_i$.

In fact it is possible to take $\mathcal{P} = \mathcal{M}(t, a, v)$ such that if $\gamma_i = p(p^{f-1}t_i + \cdots + t_{i+f-1})/(p^f - 1)$ then $\gamma_i = \max(\alpha_i, \beta_i)$.

Proof. It follows directly from Lemma 1.6 that the listed conditions are necessary. For sufficiency, we follow the argument of [6, Prop. 5.6]. Define $\gamma_i = \max(\alpha_i, \beta_i)$, $n_i = \frac{1}{p}\max(0, \beta_i - \alpha_i)$, $t_i = r_i + pn_i - n_{i+1}$, and $v_i \equiv c_i + (\alpha_i - \gamma_i) \pmod{e(K/L)}$. Observe that n_i and $\alpha_i - \gamma_i$ are integers, so that t_i is an integer and v_i is well-defined. An argument identical to the one at *loc. cit.* shows that $t_i \in [0, e]$, and easy calculations show that $\gamma_i = p(p^{f-1}t_i + \cdots + t_{i+f-1})/(p^f - 1)$ and $v_{i+1} \equiv p(v_i + t_i) \pmod{e(K/L)}$. Thus $\mathcal{P} = \mathcal{M}(t, a, v)$ is a Breuil module with the property given in the last sentence of the proposition, and two applications of Lemma 1.6 show that there exist nonzero maps $\mathcal{M} \to \mathcal{P}$ and $\mathcal{N} \to \mathcal{P}$. (For the latter, note that $\gamma_i - \alpha_i \equiv c_i - v_i \pmod{e(K/L)}$, and together with our other hypotheses this implies that $\gamma_i - \beta_i \equiv d_i - v_i \pmod{e(K/L)}$.) $\square$

Corollary 1.10. *The conditions in Proposition 1.9 give necessary and sufficient conditions that $T^L_{\mathrm{st},2}(\mathcal{M}) \simeq T^L_{\mathrm{st},2}(\mathcal{N})$.*

Proof. Suppose that there exists $\mathcal{P}$ as in Proposition 1.9. Since the kernels of the maps produced by Lemma 1.6 do not contain any free $k[u]/u^{ep}$-submodules, it follows from [25, Prop. 8.3] that they induce isomorphisms $T^L_{\mathrm{st},2}(\mathcal{M}) \simeq T^L_{\mathrm{st},2}(\mathcal{P})$ and $T^L_{\mathrm{st},2}(\mathcal{N}) \simeq T^L_{\mathrm{st},2}(\mathcal{P})$.

Conversely, suppose $T^L_{\mathrm{st},2}(\mathcal{M}) \simeq T^L_{\mathrm{st},2}(\mathcal{N})$. Let $\mathcal{M}, \mathcal{N}$ correspond to the rank one k_E-vector space schemes $\mathcal{G}, \mathcal{H}$ with generic fibre descent data. By a theorem of Raynaud [24, Prop. 2.2.2, Cor 2.2.3] there exists a maximal rank one k_E-vector space scheme $\mathcal{G}'$ with nonzero maps $\mathcal{G}' \to \mathcal{G}$, $\mathcal{G}' \to \mathcal{H}$, and $\mathcal{G}'$ obtains generic fibre descent data by a scheme-theoretic closure argument as in [3, Prop. 4.1.3]. Then we can take $\mathcal{P}$ to be the Breuil module corresponding to $\mathcal{G}'$. $\square$

1.2. **Extensions of rank one Breuil modules.** We now describe the extensions between the rank one objects of BrMod^K_{L,k_E}. The main result is analogous to [3, Lem. 5.2.2], [25, Thm. 7.5] and [9, Thm. 3.9], and since the proof is substantively the same as the proofs given at those references, we will omit some details of the argument.

Theorem 1.11. *Let $\mathcal{M}, \mathcal{N}$ be rank one Breuil modules, with notation as in Section 1.1. Each extension of $\mathcal{M}$ by $\mathcal{N}$ is isomorphic to precisely one of the form*

- $\mathcal{P} = ((k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}) \cdot m + ((k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}) \cdot n$,
- $\mathrm{Fil}^1 \mathcal{P} = \langle u^{\underline{s}} n, u^{\underline{r}} m + \underline{h} n \rangle$,
- $\phi_1(u^{\underline{s}} n) = \underline{b} n$ *and* $\phi_1(u^{\underline{r}} m + \underline{h} n) = \underline{a} m$,
- $\widehat{g}(n) = (\overline{\eta}(g)^{\underline{d}} \otimes 1) n$ *and* $\widehat{g}(m) = (\overline{\eta}(g)^{\underline{c}} \otimes 1) m$ *for all* $g \in \mathrm{Gal}(K/L)$,

in which each $h_i \in k_E[u]/u^{ep}$ is a polynomial such that:

- *h_i is divisible by $u^{r_i+s_i-e}$,*
- *the sequence h_i is periodic with period dividing f',*
- *each nonzero term of h_i has degree congruent to $r_i+c_i-d_i \pmod{e(K/L)}$, and*
- *$\deg(h_i) < s_i$, except that when there exists a nonzero morphism $\mathcal{M} \to \mathcal{N}$, the polynomials h_i for $f' \mid i$ may also have a term of degree $r_0+\beta_0-\alpha_0$ in common.*

In particular the dimension of $\mathrm{Ext}^1(\mathcal{M},\mathcal{N})$ is given by the formula

$$\delta + \sum_{i=0}^{f'-1} \#\left\{j \in [\max(0, r_i+s_i-e), s_i) : j \equiv r_i+c_i-d_i \pmod{e(K/L)}\right\}$$

where $\delta = 1$ if there exists a map $\mathcal{M} \to \mathcal{N}$ and $\delta = 0$ otherwise.

Proof. Let $\mathcal{P}$ be any extension of $\mathcal{M}$ by $\mathcal{N}$. Then $\mathrm{Fil}^1 \mathcal{P} = \langle u^{\underline{s}} n, u^{\underline{r}} m + \underline{h} n\rangle$ for some $\underline{h}$ and some lift m of the given generator of $\mathcal{M}$, and $\phi(u^{\underline{r}} m + \underline{h} n) = \underline{a} m + \delta n$ for some δ. Replacing m with $m + \delta \underline{a}^{-1} n$ and suitably altering $\underline{h}$ shows that we can take $\delta = 0$. The condition that each h_i is divisible by $u^{r_i+s_i-e}$ is necessary and sufficient to ensure that $\mathrm{Fil}^1 \mathcal{P} \supset u^e \mathcal{P}$, so that the first three conditions given in the statement of the theorem define a Breuil module (without descent data). One checks straightforwardly that replacing m with $m + \underline{a}^{-1}\phi(\underline{t}) n$ for any $\underline{t} \in (k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$ preserves the shape of $\mathcal{P}$ while replacing $\underline{h}$ with $h - u^{\underline{r}}(\underline{ba^{-1}})\phi(\underline{t}) + u^{\underline{s}}\underline{t}$, and that these are precisely the changes of m that preserve the shape of $\mathcal{P}$.

Now the descent data on $\mathcal{P}$ must have the shape

$$\widehat{g}(m) = (\overline{\eta}(g)^{\underline{c}} \otimes 1) m + A_g n$$

for some collection of elements $A_g \in (k \otimes_{\mathbb{F}_p} k_E)[u]/u^{ep}$. The condition that $\widehat{hg} = \widehat{h} \circ \widehat{g}$, evaluated at m, implies that the function $g \mapsto (\overline{\eta}(g)^{-\underline{c}} \otimes 1) A_g$ is a cocycle in the cohomology group $H^1(\mathrm{Gal}(K/L), (k \otimes k_E)[u]/u^{ep})$ in which the action of $\mathrm{Gal}(K/L)$ on $(k \otimes k_E)[u]/u^{ep}$ is given by $g \cdot x = (\overline{\eta}(g)^{\underline{d}-\underline{c}} \otimes 1) g(x)$, where $g(x)$ is the usual action. This cohomology group is trivial since $\mathrm{Gal}(K/L)$ is assumed to have order prime to p, so that $(\overline{\eta}(g)^{-\underline{c}} \otimes 1) A_g$ is the coboundary of some element v. The relation $\phi_1 \circ \widehat{g} = \widehat{g} \circ \phi_1$ applied to $u^{\underline{r}} m + \underline{h} n$ implies that $A_g n$ lies in the image of ϕ_1, so that all nonzero terms in each A_g have degree divisible by p; it follows that we can take v to have the same property. One computes that replacing m with $m + \underline{a}^{-1}\phi(\underline{t}) n$ changes $(\overline{\eta}(g)^{-\underline{c}} \otimes 1) A_g$ by the coboundary of $\underline{a}^{-1}\phi(\underline{t})$, and choosing $\underline{t}$ so that $\underline{a}^{-1}\phi(\underline{t}) = -v$ allows us to take $A_g = 0$ for all g.

Thus our extension $\mathcal{P}$ has the shape as in the theorem, and it remains to investigate the possibilities for $\underline{h}$. In order that the given shape of $\mathcal{P}$ actually defines a Breuil module with descent data, it is necessary and sufficient that $u^{r_i+s_i-e}$ divides each h_i, and that the relation $\phi_1 \circ \widehat{g} = \widehat{g} \circ \phi_1$ is well-defined and satisfied when evaluated at $u^{\underline{r}} m + \underline{h} n$. A direct calculation shows that the latter condition is equivalent to the condition that $u^{e+\underline{s}}$ divides

$$(\overline{\eta}^{\underline{d}}(g) \otimes 1) g(\underline{h}) - (\overline{\eta}^{\underline{r}+\underline{c}}(g) \otimes 1) \underline{h}$$

for all $g \in \mathrm{Gal}(K/L)$, or equivalently that the remainder of $\underline{h}$ upon division by $u^{e+\underline{s}}$ is invariant under the action of Lemma 1.5 with $\underline{w} = \underline{r} + \underline{c} - \underline{d}$. From that lemma, we deduce that any term of h_i of degree $D < e + s_i$ must satisfy $D \equiv r_i + c_i - d_i$

(mod $e(K/L)$), and that such terms occur periodically with period dividing f'. Let $V \subseteq (k \otimes k_E)[u]/u^{ep}$ be the space of elements $\underline{h}$ satisfying the conditions in the previous sentence and with each h_i divisible by $u^{\max(0,r_i+s_i-e)}$.

Now let us examine the changes-of-variable $m \rightsquigarrow m + \underline{a}^{-1}\phi(\underline{t})n$ that preserve the shape of $\mathcal{P}$ (but may change $\underline{h}$). From the argument two paragraphs earlier, we see that such a change of variables preserves the shape of the descent data precisely when the coboundary of $\underline{a}^{-1}\phi(\underline{t})$ is trivial, or in other words precisely when $\phi(\underline{t}) = g \cdot \phi(\underline{t})$ under the $\mathrm{Gal}(K/L)$-action of that paragraph. Thus $\underline{t}$ may have arbitrary terms of degree at least e (since $\phi(u^e) = 0$), while by Lemma 1.5 the nonzero terms of t_i of degree $D < e$ must have $D \equiv p^{-1}(c_{i+1} - d_{i+1}) \pmod{e(K/L)}$, and these terms must occur periodically with period dividing f'. We say that a choice of $\underline{t}$ with these properties is *allowable*.

Recall from the beginning of the proof that replacing m with $m + \underline{a}^{-1}\phi(\underline{t})n$ has the effect of replacing $\underline{h}$ with $\underline{h}' = \underline{h} - u^{\underline{r}}(\underline{ba^{-1}})\phi(\underline{t}) + u^{\underline{s}}\underline{t}$. Let $U \subseteq (k \otimes k_E)[u]/u^{ep}$ be the space of allowable choices of $\underline{t}$, and $\Upsilon : U \to V$ the map that sends $\underline{t}$ to $u^{\underline{r}}(\underline{ba^{-1}})\phi(\underline{t}) - u^{\underline{s}}\underline{t}$. The above discussion shows that $\mathrm{Ext}^1(\mathcal{M}, \mathcal{N}) \simeq \mathrm{coker}(\Upsilon)$. We use this isomorphism to compute $\dim_{k_E} \mathrm{Ext}^1(\mathcal{M}, \mathcal{N})$. Let $y_i = \#\{j \in [\max(0, r_i + s_i - e), s_i) : j \equiv r_i + c_i - d_i \pmod{e(K/L)}\}$. One calculates directly from their definitions that

$$\dim_{k_E} U = e'f' + ef(p-1), \qquad \dim_{k_E} V = e'f' + ef(p-1) + \sum_{i=0}^{f-1} y_i - \sum_{i=0}^{f-1} s_i.$$

Suppose that $\underline{t} \in \ker(\Upsilon)$, i.e. that $u^{\underline{r}}(\underline{ba^{-1}})\phi(\underline{t}) = u^{\underline{s}}\underline{t}$. Observe (e.g. by comparing with the proof of Lemma 1.6) that this is precisely the condition required for the map $\mathcal{M} \to \mathcal{N}$ defined by $m \mapsto \phi(\underline{t})n$ to be a map of Breuil modules. If there are no such nonzero maps (i.e. if $\delta = 0$, with δ as in the statement of the Theorem), then $\ker(\Upsilon) = \{\underline{t} \in U : u^{\underline{s}}t = 0\}$ and so $\ker(\Upsilon)$ has dimension $\sum_i s_i$. If instead there exists a nonzero map $\mathcal{M} \to \mathcal{N}$ (i.e. if $\delta = 1$), then since that map must be unique up to scaling, we see that $u^{\underline{s}}\underline{t}$ is unique up to scaling and $\ker(\Upsilon)$ has dimension $1 + \sum_i s_i$. In either case $\dim_{k_E} \ker(\Upsilon) = \delta + \sum_i s_i$. Finally we calculate that $\mathrm{coker}(\Upsilon)$ has dimension

$$\dim_{k_E} V - \dim_{k_E} U + \dim_{k_E} \ker(\Upsilon) = \delta + \sum_{i=0}^{f-1} y_i.$$

Now let $W' \subseteq V$ be the space of elements $\underline{h}$ satisfying the conditions given in the statement of the theorem, and $W \subseteq W'$ the subspace of elements $\underline{h}$ for which the coefficient of degree $r_0 + \beta_0 - \alpha_0$ in h_0 is zero. (Thus $W \subsetneq W'$ if and only if $\delta = 1$, in which case W'/W has k_E-dimension 1.) It is easy to verify that $\dim_{k_E} W' = \delta + \sum_i y_i$, and so to complete the proof of the Theorem it suffices to show that $W' \cap \mathrm{im}(\Upsilon) = 0$. When $\delta = 1$, a straightforward computation (using the fact that $\mathrm{Nm}(\underline{a}) = \mathrm{Nm}(\underline{b})$ in this case) shows that if $\underline{h} \in \mathrm{im}(\Upsilon)$ then the coefficients ξ_i of degree $r_i + \beta_i - \alpha_i$ in h_i for $i = 0, \ldots, f' - 1$ satisfy the linear relation $\sum_{i=0}^{f-1} (a_0 \cdots a_i)(b_0 \cdots b_i)^{-1} \xi_i = 0$. If in addition we have $\underline{h} \in W'$ (so that $\xi_i = 0$ for $i \not\equiv 0 \pmod{f'}$) then $\xi_0 = 0$ and $\underline{h} \in W$. We are therefore reduced in all cases to showing that $W \cap \mathrm{im}(\Upsilon) = 0$.

Let $\pi_W : V \to W$ be the projection map that kills each term of h_i of degree at least s_i. Observe that we may write

$$\Upsilon(\underline{t}) = u^{\underline{s}}\Phi(\underline{t}) + \pi_W(\Upsilon(\underline{t}))$$

where $\Phi(\underline{t}) = u^{\underline{r}-\underline{s}}(\underline{ba^{-1}})\phi(\underline{t}) - \underline{t}$, with terms of negative degree in $\Phi(\underline{t})$ understood to be zero. To finish the argument we must show that if $u^{\underline{s}}\Phi(\underline{t}) = 0$ then $\pi_W(\Upsilon(\underline{t})) = 0$.

Observe that the defining formula for Φ also gives a well-defined map $\overline{\Phi} \in \mathrm{End}((k \otimes k_E)[u]/u^e)$. Fix an integer $v_i \in [0, e)$ and recursively define $v_j = (r_j - s_j) + pv_{j-1}$ for $j > i$. Since $u^{v_j}e_j$ and $(b_{j+1}/a_{j+1})u^{v_{j+1}}e_{j+1}$ are congruent modulo the image of $\overline{\Phi}$ (where the e_j's are the idempotents defined in 1.2), it follows that $u^{v_j}e_j \in \mathrm{im}(\overline{\Phi})$ except possibly if the sequence $\{v_j\}$ lies entirely within the interval $[0, e)$. In the latter case the sequence $\{v_j\}$ must be periodic, indeed with period dividing f', and one computes that $v_j = p^{-1}(\beta_{j+1} - \alpha_{j+1})$ for all j. Then one checks that $u^{v_j}e_j$ and $\mathrm{Nm}(\underline{ba^{-1}})u^{v_j}e_j$ are congruent modulo $\mathrm{im}(\overline{\Phi})$; so unless $\mathrm{Nm}(\underline{a}) = \mathrm{Nm}(\underline{b})$ we again have $u^{v_j}e_j \in \mathrm{im}(\overline{\Phi})$. We conclude that $\overline{\Phi}$ is surjective (hence bijective) unless $\mathrm{Nm}(\underline{a}) = \mathrm{Nm}(\underline{b})$ and $p^{-1}(\beta_i - \alpha_i) \in \{0, \ldots, e-1\}$ for all i, in which case the image of $\overline{\Phi}$ has codimension at most 1; and in all cases we conclude that $\ker(\overline{\Phi}) = \ker(\Upsilon') + u^e(k \otimes k_E)[u]/u^{ep}$, where Υ' is the endomorphism of $(k \otimes k_E)[u]/u^{ep}$ given by the same defining formula as Υ.

Now if $u^s\Phi(\underline{t}) = 0$ then $\overline{\Phi}(\underline{t}) = 0$, so $\underline{t} \in \ker(\Upsilon') + u^e(k \otimes k_E)[u]/u^{ep}$; finally

$$\Upsilon(\underline{t}) = \Upsilon'(\underline{t}) \in \Upsilon'(u^e(k \otimes k_E)[u]/u^{ep}) \subseteq u^e(k \otimes k_E)[u]/u^{ep},$$

and it follows that $\pi_W(\Upsilon(\underline{t})) = 0$. □

Remark 1.12. We have seen in the proof of Theorem 1.11 that $\mathcal{P}$ as in the first set of bullet points of Theorem 1.11 is a well-defined Breuil module provided that

- h_i is divisible by $u^{r_i+s_i-e}$,
- nonzero terms of h_i of degree less than $e + s_i$ have degree congruent to $r_i + c_i - d_i \pmod{e(K/L)}$, and occur periodically (for i) with period dividing f'.

We will denote this Breuil module by $\mathcal{P}(r, a, c; s, b, d; h)$.

1.3. **Comparison of extension classes.** We assume for the remainder of this paper that $e(K/L)$ is divisible by $p^{f'} - 1$, so that in particular Remark 1.7 is in force. We fix characters $\chi_1, \chi_2 : G_L \to k_E^\times$ and suppose that $\mathcal{M} = \mathcal{M}(r, a, c)$ and $\mathcal{N} = \mathcal{M}(s, b, d)$ are rank one Breuil modules whose generic fibres are χ_1, χ_2 respectively. The following lemma is [21, Cor. 4.3].

Lemma 1.13. *Set $\mathcal{M} = \mathcal{M}(r, a, c)$ and write $\lambda = \mathrm{Nm}(\underline{a})^{-1}$. Then $T^K_{\mathrm{st},2}(\mathcal{M}) = (\sigma_i \circ \overline{\eta}^{c_i+\alpha_i}) \cdot \mathrm{ur}_\lambda$, where ur_λ is the unramified character of G_L sending an arithmetic Frobenius element to λ.*

The character χ_1 and the sequence of r_i's determine $\mathcal{M}$ up to isomorphism (*cf.* Corollaries 1.8 and 1.10), and similarly for $\mathcal{N}$; moreover, one checks from Lemmas 1.3 and 1.13 that given χ_1 and $r_0, \ldots, r_{f-1}$ such that $\alpha_i \in \mathbb{Z}$ for some (hence all) i, there exists $\mathcal{M}(r, a, c)$ with generic fibre χ_1. In the remainder of this section, we compare extension classes in $\mathrm{Ext}^1(\mathcal{M}, \mathcal{N})$ with extension classes in $\mathrm{Ext}^1(\mathcal{M}', \mathcal{N}')$ where $\mathcal{M}', \mathcal{N}'$ are certain other Breuil modules with the same generic fibres as $\mathcal{M}, \mathcal{N}$ respectively; our treatment follows the treatment of the case $f = 1$ in Section 5.2 of [18].

Proposition 1.14. *The Breuil module $\mathcal{P} = \mathcal{P}(r,a,c;s,b,d;h)$ has the same generic fibre as $\mathcal{P}^\dagger = \mathcal{P}(0,a,c^\dagger;e,b,d^\dagger;u^{\underline{\delta}}h)$ where*

- $c_i^\dagger = c_i + \alpha_i$,
- $d_i^\dagger = d_i + \beta_i - ep/(p-1)$, *and*
- $\delta_i = ep/(p-1) - \beta_i + \alpha_i - r_i$.

Proof. Consider the Breuil module $\mathcal{P}^\ddagger = \mathcal{P}(r,a,c;e,b,d^\dagger;u^{\underline{\delta}^\ddagger}h)$ where $\delta_i^\ddagger = ep/(p-1) - \beta_i$. It is elementary from Remark 1.12 that both $\mathcal{P}^\dagger$ and $\mathcal{P}^\ddagger$ are well-defined (note that $\delta_i, \delta_i^\ddagger \geq 0$ and that $e(K/L)$ divides e); the key point of the calculation is that $\beta_i - s_i = \beta_{i+1}/p \leq e/(p-1)$, whence $e - s_i \leq ep/(p-1) - \beta_i$. Let $m^\dagger, n^\dagger$ and $m^\ddagger, n^\ddagger$ denote the standard basis elements for $\mathcal{P}^\dagger, \mathcal{P}^\ddagger$ respectively. One checks without difficulty that there is a map $f^\ddagger : \mathcal{P} \to \mathcal{P}^\ddagger$ sending

$$m \mapsto m^\ddagger, \qquad n \mapsto u^{ep/(p-1)-\underline{\beta}}n^\ddagger$$

as well as a map $f^\dagger : \mathcal{P}^\dagger \to \mathcal{P}^\ddagger$ sending

$$m^\dagger \mapsto u^{\underline{\alpha}}m^\ddagger, \qquad n^\dagger \mapsto n^\ddagger.$$

Since $\ker(f^\dagger), \ker(f^\ddagger)$ do not contain any free $k[u]/u^{ep}$-submodules, it follows from [25, Prop. 8.3] that $T^L_{\mathrm{st},2}(f^\dagger)$ and $T^L_{\mathrm{st},2}(f^\ddagger)$ are isomorphisms. □

Note that while the extension classes $\mathrm{Ext}^1_{k_E[G_L]}(\chi_1,\chi_2)$ realized by $\mathcal{P}$ and $\mathcal{P}^\dagger$ in Proposition 1.14 may not coincide, they differ by at most multiplication by a k_E-scalar, since the maps $f^\dagger$ and $f^\ddagger$ induce k_E-isomorphisms on the one-dimensional sub and quotient characters.

Definition 1.15. Let $L(\mathcal{M},\mathcal{N}) \subseteq \mathrm{Ext}^1_{k_E[G_L]}(\chi_1,\chi_2)$ denote the subspace consisting of extension classes of the form $T^L_{\mathrm{st},2}(\mathcal{P})$ for $\mathcal{P} \in \mathrm{Ext}^1(\mathcal{M},\mathcal{N})$.

The following proposition gives a criterion for one space of extensions $L(\mathcal{M},\mathcal{N})$ to be contained in another.

Proposition 1.16. *Suppose that $\mathcal{M} = \mathcal{M}(r,a,c)$ and $\mathcal{M}' = \mathcal{M}(r',a',c')$ have generic fibre χ_1 while $\mathcal{N} = \mathcal{M}(s,b,d)$ and $\mathcal{N}' = \mathcal{M}(s',b',d')$ have generic fibre χ_2. If there exist nonzero maps $\mathcal{M} \to \mathcal{M}'$ and $\mathcal{N}' \to \mathcal{N}$ then $L(\mathcal{M}',\mathcal{N}') \subseteq L(\mathcal{M},\mathcal{N})$.*

Proof. We show more generally that the conclusion holds provided that

$$\max(\alpha_{i+1}/p - \beta_i, \alpha_i - \beta_{i+1}/p - e) \leq \max(\alpha'_{i+1}/p - \beta'_i, \alpha'_i - \beta'_{i+1}/p - e)$$

for all i. (This inequality is easily checked when there exist maps $\mathcal{M} \to \mathcal{M}'$ and $\mathcal{N}' \to \mathcal{N}$, because $\alpha_i \leq \alpha'_i$ and $\beta'_i \leq \beta_i$ for all i in this case.)

By Corollaries 1.8 and 1.10 we may suppose without loss of generality that $a = a'$ and $b = b'$. Suppose that $\mathcal{P}' = \mathcal{P}(r',a,c';s',b,d';h')$. The given inequality is equivalent to

$$(\beta_i - \alpha_i + r_i) - (\beta'_i - \alpha'_i + r'_i) + \max(0, r'_i + s'_i - e) \geq \max(0, r_i + s_i - e)$$

which is precisely the condition that is required to make the assigments $\underline{h} = u^{(\underline{\beta}-\underline{\beta}')-(\underline{\alpha}-\underline{\alpha}')+(\underline{r}-\underline{r}')}\underline{h}'$ and $\mathcal{P} = \mathcal{P}(r,a,c;s,b,d;h)$ well-defined. Then $\mathcal{P}$ and $\mathcal{P}'$ both have the same generic fibre as the extension $\mathcal{P}^\dagger$ of Proposition 1.14, and so the generic fibre of $\mathcal{P}'$ is also in $L(\mathcal{M},\mathcal{N})$. □

We remark that Proposition 1.16 should also follow from a scheme-theoretic closure argument, but we give the above argument for the sake of expedience (we will need Proposition 1.14 again in Section 2.2).

2. Models of principal series type

We retain the notation and setting of the previous section; in particular recall that we have a running assumption that $e(K/L)$ is divisible by $p^{f'}-1$. Fix a pair of characters $\chi_1, \chi_2 : G_L \to k_E^\times$.

Recall that a two-dimensional *Galois type* is (the isomorphism class of) a representation $\tau : I_L \to \mathrm{GL}_2(\overline{\mathbb{Z}}_p)$ that extends to a representation of G_L and whose kernel is open. We say that τ is a *principal series type* if $\tau \simeq \lambda \oplus \lambda'$ where λ, λ' both extend to representations of G_L.

In this section we use the results of Section 1 to associate to the triple (χ_1, χ_2, τ) a subspace $L(\chi_1, \chi_2, \tau) \subseteq \mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2)$. We will see that $L(\chi_1, \chi_2, \tau)$ contains every extension of χ_1 by χ_2 that arises as the reduction mod p of a potentially Barsotti-Tate representation of type τ; in fact we will think of $L(\chi_1, \chi_2, \tau)$ as a finite flat avatar for the collection of such extensions. In Section 2.1 we define the set $L(\chi_1, \chi_2, \tau)$ and prove that it is a vector space (provided that it is nonempty). In Section 2.2 we restrict to the main local setting of our paper and study the spaces $L(\chi_1, \chi_2, \tau)$ in detail in that setting; for instance we compute the dimension of these spaces in many cases.

2.1. **Maximal and minimal models of type τ.** Raynaud [24] shows that if one fixes a finite flat p-torsion group scheme G over K, then the set of finite flat group schemes over $\mathcal{O}_K$ with generic fibre G has the structure of a lattice; in particular it possesses maximal and minimal elements. This has proved to be a valuable observation, and variants of it have recurred in numerous contexts (see [3, Lem. 4.1.2], [25, §8], [7, §3.3], and [19, §5.3] to name a few). Let τ be a principal series type. In this subsection we introduce the notion of a model of type τ (see Definition 2.2 below) and prove the existence of maximal and minimal models of type τ.

Definition 2.1. Write $\chi = \chi_1\chi_2$. If $\mathcal{M} = \mathcal{M}(r, a, c)$ has generic fibre χ_1, define the χ-*dual* of $\mathcal{M}$ to be the unique Breuil module $\mathcal{M}^\vee_\chi = \mathcal{M}(s, b, d)$ with generic fibre χ_2 such that $r_i + s_i = e$ for all i. The existence of $\mathcal{M}^\vee_\chi$ is implied by the paragraph following Lemma 1.13.

If $\tau \simeq \lambda \oplus \lambda'$ is a principal series type, we let $\overline{\lambda}, \overline{\lambda}'$ denote the reductions of λ, λ' modulo the maximal ideal of $\overline{\mathbb{Z}}_p$; we will usually abuse notation and write λ, λ' where we mean $\overline{\lambda}, \overline{\lambda}'$.

Definition 2.2. Let $\tau \simeq \lambda \oplus \lambda'$ be a principal series type. We say that $\mathcal{M}(r, a, c)$ is a *model of type* τ if $\sigma_i \circ \overline{\eta}^{c_i} \in \{\lambda, \lambda'\}$ for all i. Note that if $(\chi_1\chi_2)|_{I_{G_L}} = \lambda\lambda'\overline{\epsilon}$ and $\mathcal{M}(r, a, c)$ is a model of type τ with generic fibre χ_1, then its χ-dual $\mathcal{M}^\vee_\chi = \mathcal{M}(s, b, d)$ is a model of type τ with generic fibre χ_2, and moreover $\{\sigma_i \circ \overline{\eta}^{c_i}, \sigma_i \circ \overline{\eta}^{d_i}\} = \{\lambda, \lambda'\}$ for all i.

Definition 2.3. We define

$$L(\chi_1, \chi_2, \tau) = \cup_{\mathcal{M},\mathcal{N}} L(\mathcal{M}, \mathcal{N})$$

as $\mathcal{M}, \mathcal{N}$ range over all pairs of models of type τ with generic fibre χ_1, χ_2 respectively, and such that $\{\sigma_i \circ \overline{\eta}^{c_i}, \sigma_i \circ \overline{\eta}^{d_i}\} = \{\lambda, \lambda'\}$ for all i.

It follows, for instance from [21, Cor. 5.2], that $L(\chi_1, \chi_2, \tau)$ contains all extensions of χ_1 by χ_2 that arise as the reduction mod p of a potentially Barsotti-Tate representation of G_L of type τ. Note that if $L(\chi_1, \chi_2, \tau) \neq \varnothing$ then $(\chi_1\chi_2)|_{I_{G_L}} = \lambda\lambda'\overline{\epsilon}$.

Proposition 2.4. *Let $\mathcal{S}$ be the set of all $\mathcal{M}(r,a,c)$ of type $\lambda \oplus \lambda'$ with generic fibre χ. If $\mathcal{S}$ is nonempty, then it has a minimal and a maximal element; that is, there are Breuil modules $\mathcal{M}_-, \mathcal{M}_+ \in \mathcal{S}$ such that for any $\mathcal{M} \in \mathcal{S}$ there exist nonzero maps $\mathcal{M}_- \to \mathcal{M}$ and $\mathcal{M} \to \mathcal{M}_+$.*

Proof. By duality it suffices to prove the existence of $\mathcal{M}_+$. For this, since $\mathcal{S}$ is finite, it is enough to prove that any $\mathcal{M}, \mathcal{N} \in \mathcal{S}$ have an upper bound in $\mathcal{S}$, i.e. that there exists $\mathcal{P} \in \mathcal{S}$ together with nonzero maps $\mathcal{M} \to \mathcal{P}$ and $\mathcal{N} \to \mathcal{P}$.

Since $\mathcal{M}, \mathcal{N}$ have the same generic fibre, the conditions of Proposition 1.9 are satisfied, and we can form $\mathcal{P} = \mathcal{P}(t,a,v)$ as in the last sentence of the proposition. Note that if $\gamma_i = \alpha_i$ then $v_i = c_i$, while if $\gamma_i = \beta_i$ then $v_i = d_i$ (see the last sentence of the proof of Proposition 1.9, for instance). Thus $\sigma_i \circ \overline{\eta}^{v_i} \in \{\sigma_i \circ \overline{\eta}^{c_i}, \sigma_i \circ \overline{\eta}^{d_i}\} \subseteq \{\lambda, \lambda'\}$, and we conclude that $\mathcal{P} \in \mathcal{S}$. □

Remark 2.5. An argument identical to the above can be used to prove a much more general statement. Namely, we can fix sets $S_i \subseteq \{\sigma_i \circ \overline{\eta}^c : c \in \mathbb{Z}\}$ for each i, and consider the set $\mathcal{S}$ of Breuil modules $\mathcal{M}(r,a,c)$ with generic fibre χ such that $\sigma_i \circ \overline{\eta}^{c_i} \in S_i$ for all i; then if $\mathcal{S}$ is nonempty, it has a maximal and a minimal element.

Corollary 2.6. *If $L(\chi_1, \chi_2, \tau)$ is nonempty, then it is a vector space.*

Proof. Suppose that $L(\chi_1, \chi_2, \tau)$ is nonempty. By Proposition 2.4 there exists a minimal model $\mathcal{M}$ of type τ with generic fibre χ_1. It follows easily that $\mathcal{M}_\chi^\vee$ must be the maximal model of type τ with generic fibre χ_2. Proposition 1.16 implies that $L(\chi_1, \chi_2, \tau) = L(\mathcal{M}, \mathcal{M}_\chi^\vee)$, and the lemma follows. □

2.2. **The local setting.** For remainder of the paper we suppose that K/L is totally ramified of degree $p^{f'} - 1$, so that $K = L(\pi)$, $f = f'$, and $e(K/L) = p^f - 1$. Recall that in this setting we have $\omega_i = (\sigma_i \circ \overline{\eta})|_{I_L}$. The characters ω_i form a fundamental system of characters of niveau f, and we write

$$\lambda = \prod_{i=0}^{f-1} \omega_i^{\nu_i}, \qquad \lambda' = \prod_{i=0}^{f-1} \omega_i^{\nu_i'}$$

with $\nu_i, \nu_i' \in [0, p-1]$ for all i; when either λ or λ' is trivial we require $\nu_i = p-1$ for all i or $\nu_i' = p-1$ for all i, respectively. Write $\lambda'/\lambda = \omega_0^\delta$ and define integers $\delta_i \in [0, p-1]$ by $\lambda'/\lambda = \prod_{i=0}^{f-1} \omega_i^{\delta_i}$, with not all δ_i equal to $p-1$. Let $[p^i\delta]$ be the unique integer in the interval $[0, e(K/L)-1]$ congruent to $p^i\delta \pmod{e(K/L)}$.

From the equality $\prod_{i=0}^{f-1} \omega_i^{\delta_i+\nu_i} = \prod_{i=0}^{f-1} \omega_i^{\nu_i'}$ together with our bounds on the δ_i, ν_i, ν_i', it follows that there exists a unique collection of integers $\gamma_i \in \{0,1\}$ such that $\nu_i' = \delta_i + \nu_i - p\gamma_{i-1} + \gamma_i$. We write $C = \{i : \gamma_i = 1\}$ (the symbol C here stands for "carries").

With the above notation, we prove the following.

Proposition 2.7. *Suppose that $\mathcal{M} = \mathcal{M}(r,a,c)$ is a model of type $\lambda \oplus \lambda'$. Let $J = \{i : \sigma_i \circ \overline{\eta}^{c_i} \neq \lambda'\} \subseteq \{0,\ldots,f-1\}$. Define x_i by the formula*

$$r_i = \begin{cases} x_i e(K/L) & \text{if } i, i+1 \in J \text{ or } i, i+1 \notin J \\ x_i e(K/L) + [p^i \delta] & \text{if } i \in J,\ i+1 \notin J,\ i \notin C \\ x_i e(K/L) + (e(K/L) - [p^i\delta]) & \text{if } i \notin J,\ i+1 \in J,\ i \in C \\ x_i e(K/L) - [p^i\delta] & \text{if } i \notin J,\ i+1 \in J,\ i \notin C \\ x_i e(K/L) - (e(K/L) - [p^i\delta]) & \text{if } i \in J,\ i+1 \notin J,\ i \in C. \end{cases}$$

Then each x_i is an integer in the interval $[0, e']$, and if $\lambda \neq \lambda'$ then $x_i \neq e'$ in the second and third cases, while $x_i \neq 0$ in the fourth and fifth cases. Moreover the generic fibre of $\mathcal{M}$, on inertia, is equal to

$$\prod_{i \in J} \omega_i^{\nu_i} \prod_{i \notin J} \omega_i^{\nu_i'} \prod_{i=0}^{f-1} \omega_i^{x_i}.$$

Remark 2.8. Note that $J = \{i : \sigma_i \circ \overline{\eta}^{c_i} = \lambda\}$ unless $\lambda = \lambda'$, in which case $J = \varnothing$. The special case of Proposition 2.7 where $\lambda = 1$ is given in [30, §2.2.1]. The proof in [30, §2.2.1] is essentially the same as the one we give here, but the statement of the result when $\lambda = 1$ is somewhat simpler because $i \in C$ for all i.

Proof of Proposition 2.7. The case $\lambda = \lambda'$ is straightforward (note that $\delta = 0$, while $J = \varnothing$). Assume for the rest of the proof that $\lambda \neq \lambda'$. According to the definition of J we have $p^{f-i} c_i \equiv \sum_{j=0}^{f-1} p^{f-j} \nu_j$ if $i \in J$ and $p^{f-i} c_i \equiv \sum_{j=0}^{f-1} p^{f-j} \nu_j'$ if $i \notin J$. From the congruence $c_{i+1} \equiv p(c_i + r_i) \pmod{e(K/L)}$ together with the definitions preceding the statement of the Proposition, it follows that there exist integers y_i so that

$$r_i = \begin{cases} y_i e(K/L) & \text{if } i, i+1 \in J \text{ or } i, i+1 \notin J \\ y_i e(K/L) + [p^i\delta] & \text{if } i \in J,\ i+1 \notin J \\ y_i e(K/L) - [p^i\delta] & \text{if } i \notin J,\ i+1 \in J. \end{cases}$$

Since $r_i \in [0, e]$ for all i we have in particular that $y_i \in [0, e']$, with $y_i \neq e'$ if $i \in J, i+1 \notin J$ and $y_i \neq 0$ if $i \notin J, i+1 \in J$.

From this formula for the r_i's we calculate that

$$\sum_{i=0}^{f-1} p^{f-i} r_i = \sum_{i=0}^{f-1} p^{f-i} y_i e(K/L) + \sum_{i \in J, i+1 \notin J} p^{f-i}[p^i\delta] - \sum_{i \notin J, i+1 \in J} p^{f-i}[p^i\delta]. \tag{2.9}$$

Moreover we have $[p^i\delta] = \delta_i + p^{f-1}\delta_{i+1} + \cdots + p\delta_{i-1}$. Suppose that $0 \in J$. Let us compute the coefficient of δ_j on the right-hand side of (2.9). We see that $p^{f-i}[p^i\delta]$ contains a term of the form $p^{f-j}\delta_j$ if $i \geq j$ and $p^{2f-j}\delta_j$ if $i < j$. If $j \in J$ then the number of elements $i \in [0, j-1]$ such that $i \in J, i+1 \notin J$ is equal to the number of elements $i \in [0, j-1]$ such that $i \notin J, i+1 \in J$, and similarly for the interval $[j, f-1]$. It follows in this case that δ_j does not appear on the right-hand side of (2.9). If $j \notin J$ then instead the contribution of δ_j to the right-hand side of (2.9) is $(p^{2f-j} - p^{f-j})\delta_j$. We conclude that

$$\alpha_0 = \frac{1}{p^f - 1} \sum_{i=0}^{f-1} p^{f-i} r_i = \sum_{i=0}^{f-1} p^{f-i} y_i + \sum_{i \notin J} p^{f-i} \delta_i$$

and applying Lemma 1.13 we find that the generic fibre of $\mathcal{M}$, on inertia, is equal to $\prod_{i=0}^{f-1}\omega_i^{\nu_i}\prod_{i\notin J}\omega_i^{\delta_i}\prod_{i=0}^{f-1}\omega_i^{y_i}$, which rearranges to

$$\prod_{i\in J}\omega_i^{\nu_i}\prod_{i\notin J}\omega_i^{\nu_i'}\prod_{i\notin J}\omega_i^{p\gamma_{i-1}-\gamma_i}\prod_{i=0}^{f-1}\omega_i^{y_i} \tag{2.10}$$

by substituting for δ_i using the defining formula for the γ_i's. An analogous calculation in the case $0\notin J$ yields the formula

$$\prod_{i\in J}\omega_i^{\nu_i}\prod_{i\notin J}\omega_i^{\nu_i'}\prod_{i\in J}\omega_i^{-p\gamma_{i-1}+\gamma_i}\prod_{i=0}^{f-1}\omega_i^{y_i}$$

But $\prod_{i\notin J}\omega_i^{p\delta_{i-1}-\delta_i}=\prod_{i\in J}\omega_i^{-p\delta_{i-1}+\delta_i}$ since $\prod_{i=0}^{f-1}\omega_i^{p\delta_{i-1}-\delta_i}=1$, so in fact the formula (2.10) is valid in all cases. From the definition of the set C we can rewrite (2.10) as

$$\prod_{i\in J}\omega_i^{\nu_i}\prod_{i\notin J}\omega_i^{\nu_i'}\prod_{i\in C,i+1\notin J}\omega_i\prod_{i\in C,i\notin J}\omega_i^{-1}\prod_{i=0}^{f-1}\omega_i^{y_i}.$$

Now observe that with x_i as in the statement of the Proposition we have

$$x_i=\begin{cases} y_i+1 & \text{if } i\in J, i+1\notin J, i\in C\\ y_i-1 & \text{if } i\notin J, i+1\in J, i\in C\\ y_i & \text{otherwise,}\end{cases}$$

and the rest of the Proposition follows. □

Definition 2.11. If $x_0,\dots,x_{f-1}$ are integers in the interval $[0,e']$ with $x_i\neq e'$ whenever $i\in J, i+1\notin J, i\notin C$ or $i\notin J, i+1\in J, i\in C$, and $x_i\neq 0$ whenever $i\in J, i+1\notin J, i\in C$ or $i\notin J, i+1\in J, i\notin C$, we say that the x_i's are *allowable* for J. (Properly speaking we should say that they are allowable for J and C, but C will remain fixed in any calculation.) Observe that for every choice of J together with a collection of x_i's that are allowable for J, there exists a model of type $\lambda\oplus\lambda'$ as in Proposition 2.7 that possesses those invariants.

Proposition 2.12. *Suppose that $\nu_i'\in[p-1-e',p-1]$ and $\nu_i\le\nu_i'$ for all i.*

(1) There exists a model of type $\lambda\oplus\lambda'$ with trivial generic fibre.
(2) The minimal model of type $\lambda\oplus\lambda'$ with trivial generic fibre is $\mathcal{M}=\mathcal{M}(r,1,c)$ with $r_i=e(K/L)(p-1-\nu_i')$ and $c_i=\sum_{j=0}^{f-1}\nu_{i-j}'p^j$ for all i. In the notation of Proposition 2.7 we have $J=\varnothing$ and $x_i=p-1-\nu_i'$ for all i.

Proof. If the generic fibre of $\mathcal{M}$ is trivial, by Proposition 2.7 we can write

$$\prod_{i=0}^{f-1}\omega_i^{x_i}=\prod_{i\in J}\omega_i^{p-1-\nu_i}\prod_{i\notin J}\omega_i^{p-1-\nu_i'} \tag{2.13}$$

and conversely if this identity holds for some choice of J and allowable x_i's then taking $\underline{a}=1$ gives a model of type $\lambda\oplus\lambda'$ with trivial generic fibre. If we take $J=\varnothing$, then the integers $x_i=p-1-\nu_i'\in[0,e']$ are automatically allowable; this proves (1), and since $J=\varnothing$ we have $\sigma_i\circ\overline{\eta}^{c_i}=\lambda'$ for all i, which implies $c_i=\sum_{j=0}^{f-1}\nu_{i-j}'p^j$. It remains to show that the Breuil module $\mathcal{M}'$ corresponding to this data is actually the minimal model.

First suppose that there exists a choice of J and x_i's so that both sides of (2.13) are trivial. Since $\nu_i, \nu_i' \in [0, p-1]$, this implies that at least one of ν_i, ν_i' is $p-1$ for all i, or at least one of ν_i, ν_i' is 0 for all i. Since $\nu_i \le \nu_i'$, the latter would imply $\nu_i = 0$ for all i; but this contradicts our convention that $\nu_i = p-1$ for all i when $\lambda = 1$. So the former must hold, and we have $\nu_i' = p-1$ for all i. Then $J = \varnothing$ and $x_i = p - 1 - \nu_i' = 0$ for all i evidently gives a minimal model.

Now suppose it is never the case that both sides of (2.13) are trivial. Fix J and integers $x_i \ge 0$ so that (2.13) is satisfied (with $J = \varnothing$ if $\lambda = \lambda'$), define integers r_i by the formulas in the statement of Proposition 2.7, and then define $\alpha_i = \frac{1}{p^f-1}\sum_{j=0}^{f-1} p^{f-j} r_{i+j}$ as usual. (Any model of type $\lambda \oplus \lambda'$ with trivial generic fibre yields such data, with the x_i's allowable for J; however, note that in the argument that follows we do *not* assume that the x_i's are allowable for J.) To deduce that $\mathcal{M}'$ is the minimal model, it suffices by (the dual of) Proposition 1.9 to show that unless $J = \varnothing$ and $x_i = p - 1 - \nu_i'$ for all i, we must have $\alpha_i > \alpha_i'$ for some i, where the $\alpha_i' = \sum_{j=0}^{f-1} p^{f-j}(p - 1 - \nu_{i+j}')$ are the corresponding constants for $\mathcal{M}'$.

First suppose that $x_i \ge p$ for some i. Replacing x_i with $x_i - p$ and x_{i-1} with $x_{i-1} + 1$ leaves the truth of (2.13) unchanged, leaves α_j unchanged for all $j \ne i$, and replaces α_i with $\alpha_i - pe(K/L)$. By iterating this "carrying" operation we can reduce to the case where $x_i \le p-1$ for all i. In that case, since both sides of (2.13) are assumed to be nontrivial we must actually have

$$x_i = \begin{cases} p - 1 - \nu_i & \text{if } i \in J \\ p - 1 - \nu_i' & \text{if } i \notin J. \end{cases}$$

The claim in the case $J = \varnothing$ is now immediate, so suppose that $J \ne \varnothing$, and indeed suppose without loss of generality that $0 \in J$. Note that since $\nu_i' \ge \nu_i$ for all i the set C is empty, and the proof of Proposition 2.7 shows that α_0 is equal to $\sum_{i=0}^{f-1} p^{f-i} x_i + \sum_{i \notin J} p^{f-i}\delta_i$. An inequality $\alpha_0 \le \alpha_0'$, or equivalently

$$\sum_{i=0}^{f-1} p^{f-i} x_i + \sum_{i \notin J} p^{f-i}\delta_i \le \sum_{i=0}^{f-1} p^{f-i}(p - 1 - \nu_i'),$$

would imply that $\nu_i = \nu_i'$ for all $i \in J$, and $\delta_i = 0$ for all $i \notin J$. But $\delta_i = 0$ implies $\nu_i = \nu_i'$; so in fact we would have $\nu_i = \nu_i'$ for all i, contradicting that $\lambda \ne \lambda'$ when $J \ne \varnothing$. Therefore $\alpha_0 > \alpha_0'$. □

Corollary 2.14. *Let τ be a type as in Proposition 2.12. Write $\nu_i' = (p-1-e') + \mu_i$ for all i. If $\chi|_{I_{G_L}} = \lambda\lambda'\bar{\epsilon}$ then*

$$\dim_{k_E} L(1, \chi, \tau) \le \delta + \sum_{i=0}^{f-1} \mu_i$$

where $\delta = 1$ if $\chi = 1$ and $\delta = 0$ otherwise.

Proof. Let $\mathcal{M}$ be the minimal model of type τ with trivial generic fibre, as described by Proposition 2.12(2). By the proof of Corollary 2.6 we have $L(1, \chi, \tau) = L(\mathcal{M}, \mathcal{M}_\chi^\vee)$. We compute an upper bound on the dimension of $L(\mathcal{M}, \mathcal{M}_\chi^\vee)$ using Theorem 1.11. Since $J = \varnothing$ we have $r_i = (p - 1 - \nu_i')e(K/L)$ for all i, and

$$s_i = e - r_i = e(K/L)e' - r_i = e(K/L)\mu_i.$$

Thus the ith term in the dimension formula in Theorem 1.11 is μ_i. □

We will now use Proposition 1.14 to compare the spaces $L(1,\chi,\tau)$ as τ varies, at least in certain cases.

Proposition 2.15. *Let $\tau \simeq \lambda \oplus \lambda'$ be a type as in Proposition 2.12, and suppose further that $\nu_i + \nu'_i \geq p-1$ for all i, and $\chi \neq 1$. The space $L(1,\chi,\tau)$ is the set of extension classes of generic fibres of Breuil modules of the form $\mathcal{P}(0,1,0;e,b,d^\dagger;h)$, where $d^\dagger_i = \sum_{j=0}^{f-1}(\nu_{i-j}+\nu'_{i-j}-(p-1))p^j$, each h_i is a polynomial whose only nonzero terms have degree $t(p^f-1)-[\sum_{j=0}^{f-1}(\nu_{i-j}+\nu'_{i-j}-(p-1))p^j]$ with $p-1-\nu'_i < t \leq e'$, and $\mathrm{Nm}(\underline{b})^{-1}$ gives the unramified part of χ as in Lemma 1.13.*

Proof. Let $\mathcal{M}$ be the minimal model of Proposition 2.12. Then $\mathcal{M}^\vee_\chi = \mathcal{M}(s,b,d)$ with $s_i = e(L/K)\mu_i$, $d_i = \sum_{j=0}^{f-1}\nu_{i-j}p^j$, and $\underline{b}$ as in the statement of the proposition. By Theorem 1.11, classes in $\mathrm{Ext}^1(\mathcal{M},\mathcal{N})$ have h_i with terms of degree $m(p^f-1)+\sum_{j=0}^{f-1}(\nu'_{i-j}-\nu_{i-j})p^j$ with $0 \leq m < \mu_i$ (note that the hypotheses of Proposition 2.12 ensure that $\nu'_{i-j}-\nu_{i-j}$ are non-negative and not all zero).

Now compute that the $\underline{\delta}$ of Proposition 1.14 has

$$\delta_i = ep/(p-1) - \beta_i + \alpha_i - r_i = (p^f-1)(p-1-\nu'_i) + 2\sum_{j=0}^{f-1}(p-1-\nu'_{i-j})p^j,$$

and so the terms of the Breuil module $\mathcal{P}^\dagger$ of Proposition 1.14 have degree

$$m(p^f-1) + \sum_{j=0}^{f-1}(\nu'_{i-j}-\nu_{i-j})p^j + \delta_i = t(p^f-1) - \sum_{j=0}^{f-1}(\nu_{i-j}+\nu'_{i-j}-(p-1))p^j$$

where $t = p-1-\nu'_i+m+1$. When $\nu'_i = \nu_i = p-1$ for all i (i.e. in the unique case where $[\sum_{j=0}^{f-1}(\nu_{i-j}+\nu'_{i-j}-(p-1))p^j]$ and $\sum_{j=0}^{f-1}(\nu_{i-j}+\nu'_{i-j}-(p-1))p^j$ are different) note that there is a change of basis parameter $\underline{t}$ as in the proof of Theorem 1.11 with $\underline{t} \in (k\otimes k_E)^\times$ that exchanges the terms of degree 0 in the h_i's for terms of degree e'. One easily checks that $c^\dagger$ and $d^\dagger$ are as claimed, completing the proof. □

Remark 2.16. The Breuil modules $\mathcal{P}$ of Proposition 2.15 are usually in the canonical form of Theorem 1.11; the exception is that if $\lambda\lambda' = 1$ then we have terms of degree e' in h_i instead of terms of degree 0. However, as we have seen in the preceding argument, these are equivalent by a change of basis parameter $\underline{t}$ as in the proof of Theorem 1.11.

Corollary 2.17. *For any $\tau, \dot{\tau}$ as in Proposition 2.15 and $\chi \neq 1$ with $\chi|_{I_{G_L}} = \lambda\lambda'\overline{\epsilon}$, we have*

(1) $\dim_{k_E} L(1,\chi,\tau) = \sum_{i=0}^{f-1}\mu_i$,
(2) $L(1,\chi,\tau) \cap L(1,\chi,\dot{\tau}) = L(1,\chi,\ddot{\tau})$ where the type $\ddot{\tau}$ has $\ddot{\nu}_i = \max(\nu_i,\dot{\nu}_i)$ and $\ddot{\nu}'_i = \min(\nu'_i,\dot{\nu}'_i)$ (with the inferrable notation).

Proof. Let $\mathcal{M}_0 = \mathcal{M}(0,1,0)$ and $\mathcal{N}_0 = \mathcal{M}(e,b,d^\dagger)$, with b and $d^\dagger$ as in the statement of Proposition 2.15. The map $\mathrm{Ext}^1(\mathcal{M}_0,\mathcal{N}_0) \to \mathrm{Ext}^1_{k_E[G_L]}(1,\chi)$ is injective; for instance, this follows from the dimension calculation in Theorem 1.11 together with the fact that the map is surjective except in the case of cyclotomic χ when the image is the set of peu ramifiées classes. Now the result follows from Proposition 2.15 and Corollary 2.14, together with Remark 2.16 in the case where $\chi|_{I_{G_L}} = \overline{\epsilon}|_{I_{G_L}}$. □

3. Weights and types

3.1. **Serre weights.** We maintain the notation from the preceding section. In particular, L is a finite extension of $\mathbb{Q}_p$ of absolute ramification degree e', K/L is totally ramified of degree $p^{f'}-1$, so that $K = L(\pi)$, $f = f'$, and $e(K/L) = p^f - 1$. We continue to assume that the residue field is k of degree f over $\mathbb{F}_p$, and that $\sigma_i : k \hookrightarrow k_E$ are embeddings satisfying $\sigma_i = \sigma_{i+1}^p$ for $i = 0, \ldots, f-1$, taking indices modulo f.

Let $\overline{\rho} : G_L \to \mathrm{GL}_2(k_E)$ be a reducible representation, so that

$$\overline{\rho} \simeq \begin{pmatrix} \chi_2 & * \\ 0 & \chi_1 \end{pmatrix}$$

for some characters $\chi_1, \chi_2 : G_L \to k_E^\times$. In particular if $\overline{\rho}$ is decomposable, then we choose an ordering of the characters. The ordered pair of characters (χ_1, χ_2) will be fixed throughout the section.

Recall that a *Serre weight* in our context is an isomorphism class of absolutely irreducible representations of $\mathrm{GL}_2(k)$ in characteristic p. These are all defined over k_E and have the form

$$\mu_{m,n} := \bigotimes_{i=0}^{f-1} \left(\det{}^{m_i} \otimes \mathrm{Sym}^{n_i} k^2\right) \otimes_{k,\sigma_i} k_E$$

where $m = (m_0, \ldots, m_{f-1})$ and $n = (n_0, \ldots, n_{f-1})$ are f-tuples of integers satisfying $0 \le n_i \le p-1$ for all i. The representations $\mu_{m,n}$ and $\mu_{m',n'}$ are isomorphic if and only if $n = n'$ and $\sum_{i=0}^{f-1} m_i p^{f-i} \equiv \sum_{i=0}^{f-1} m'_i p^{f-i} \pmod{p^f - 1}$.

A set of predicted Serre weights for $\overline{\rho}$ is defined by Barnet-Lamb, Gee and Geraghty in [1, Def. 4.1.14] (building on [5, 28, 14], and following [13]). In order to give the definition, we use the notion of a Hodge–Tate module.

Definition 3.1. A *Hodge–Tate module of rank d (for L over E)* is an isomorphism class of filtered free $(L \otimes_{\mathbb{Q}_p} E)$-modules of rank d, i.e. of objects $(V, \mathrm{Fil}^\bullet)$, where V is a free $(L \otimes_{\mathbb{Q}_p} E)$-module of rank d and for $i \in \mathbb{Z}$, $\mathrm{Fil}^i V$ is a (not necessarily free) $(L \otimes_{\mathbb{Q}_p} E)$-submodule such that $\mathrm{Fil}^j V \subseteq \mathrm{Fil}^i V$ if $i \le j$, $\mathrm{Fil}^i V = V$ for $i \ll 0$ and $\mathrm{Fil}^i V = 0$ for $i \gg 0$.

Recall that we are assuming that E contains all the embeddings of L in $\overline{\mathbb{Q}}_p$, so to give a Hodge–Tate module of rank d is equivalent to giving, for each $\sigma : L \hookrightarrow \overline{\mathbb{Q}}_p$, a d-tuple of integers $(w_{\sigma,1}, \ldots, w_{\sigma,d})$ with $w_{\sigma,1} \le w_{\sigma,2} \le \cdots \le w_{\sigma,d}$. For consistency with our conventions we normalize this correspondence so that $(V, \mathrm{Fil}^\bullet)$ corresponds to the d-tuples $(w_{\sigma,1}, \ldots, w_{\sigma,d})$ defined by

$$-w_{\sigma,r} = \max\{\, w \,:\, r \le \dim_E \mathrm{Fil}^w(V \otimes_{L \otimes_{\mathbb{Q}_p} E} E)\,\},$$

where the tensor product is relative to the projection $L \otimes_{\mathbb{Q}_p} E \to E$ defined by $x \otimes y \mapsto \sigma(x)y$.

Definition 3.2. We refer to the d-tuple $(w_{\sigma,1}, w_{\sigma,2}, \ldots, w_{\sigma,d})$ as the *σ-labelled Hodge–Tate weights* of $(V, \mathrm{Fil}^\bullet)$. We say that $(V, \mathrm{Fil}^\bullet)$ is a *lift* of the Serre weight $\mu_{m,n}$ if $d = 2$ and for each $i = 0, \ldots, f-1$ there is an embedding $\tilde{\sigma}_i : L \hookrightarrow E$ lifting σ_i such that:

- $(V, \mathrm{Fil}^\bullet)$ has $\tilde{\sigma}_i$-labelled Hodge–Tate weights $(m_i, m_i + n_i + 1)$;
- for each $\sigma \neq \tilde{\sigma}_i$ lifting σ_i, $(V, \mathrm{Fil}^\bullet)$ has σ-labelled Hodge–Tate weights $(0, 1)$.

Recall that if $\rho : G_L \to \mathrm{GL}_d(E)$ is crystalline, then $D_{\mathrm{dR}}(\rho)$ has the structure of a filtered free $(L \otimes_{\mathbb{Q}_p} E)$-module of rank d as in Definition 3.1. We then define the Hodge–Tate module and σ-labelled Hodge–Tate weights of ρ to be those of $D_{\mathrm{dR}}(\rho)$.

Definition 3.3. We say that μ is a *predicted Serre weight* for $\overline{\rho}$ if, enlarging E if necessary, $\overline{\rho}$ has a reducible crystalline lift ρ whose Hodge–Tate type is a lift of μ. We then define $W_{\mathrm{expl}}(\overline{\rho})$ to be the set of predicted Serre weights for $\overline{\rho}$.

It is immediate from the definition that $W_{\mathrm{expl}}(\overline{\rho}) \subset W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$; moreover it follows from the description of reductions of crystalline characters that $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$ is precisely the set of Serre weights for $\overline{\rho}^{\mathrm{ss}}$ predicted by Schein in [28] (see [1, Lemma 4.1.22]). Recall that this is the set of $\mu_{m,n}$ such that

$$(3.4)\qquad \begin{array}{rcl} \chi_2|_{I_{G_L}} & = & \prod_{i\in J}\omega_i^{m_i+n_i+e'-d_i}\prod_{i\notin J}\omega_i^{m_i+e'-d_i} \\ \text{and}\quad \chi_1|_{I_{G_L}} & = & \prod_{i\in J}\omega_i^{m_i+d_i}\prod_{i\notin J}\omega_i^{m_i+n_i+d_i} \end{array}$$

for some $J \subseteq \{0,\ldots,f-1\}$ and integers d_i for $i = 0,\ldots,f-1$ satisfying $0 \le d_i \le e'-1$ if $i \in J$ and $1 \le d_i \le e'$ if $i \notin J$. Thus $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$ is indeed "explicit," as the notation is presumably meant to indicate; however $W_{\mathrm{expl}}(\overline{\rho})$ is less so since it is defined in terms of reductions of extensions of crystalline characters.

3.2. **A partition by types.** We fix $\overline{\rho}$ as in §3.1 and let $W' = W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$. The aim of this section is to define a partition of W' under the following hypothesis on $\overline{\rho}$:

Definition 3.5. We say that $\overline{\rho}$ is *generic* if $\chi_1^{-1}\chi_2|_{I_{G_L}} = \prod_{i=1}^{f}\omega_i^{b_i+e'}$ for some integers b_i satisfying

$$e' \le b_i + e' \le p-1-e'.$$

We assume for the remainder of the paper that $\overline{\rho}$ is generic, so that we have integers b_i as above.[1] Note in particular that this implies that $e' \le (p-1)/2$. We also write $\chi_1|_{I_{G_L}} = \prod_{i=1}^{f}\omega_i^{c_i}$ for some integers c_i.

Suppose that $\mu_{m,n} \in W'$, with J and $d = (d_0,\ldots,d_{f-1})$ as in (3.4). Then n satisfies the congruence:

$$\sum_{i=0}^{f-1}(b_i+2d_i)p^{f-i} \equiv \sum_{i\in J} n_i p^{f-i} - \sum_{i\notin J} n_i p^{f-i} \pmod{p^f-1}.$$

One easily sees that given J and d, there is a unique such n unless

$$\sum_{i=0}^{f-1}(b_i+2d_i)p^{f-i} \equiv \sum_{i\in J}(p-1)p^{f-i} \pmod{p^f-1}.$$

The genericity hypothesis implies that $0 \le b_i + 2d_i < p-1$ if $i \in J$, and $0 < b_i + 2d_i \le p-1$ if $i \notin J$, so we see that n is unique unless either $b = d = (0,\ldots,0)$ and $J = \{0,\ldots,f-1\}$, or $b = (p-1-2e',\ldots,p-1-2e')$, $d = (e',\ldots,e')$ and $J = \varnothing$ (and so in particular unless $\chi_1^{-1}\chi_2|_{I_{G_L}} = \overline{\epsilon}|_{I_{G_L}}^{\pm 1}$). It follows that aside from

[1] It appears to us that it should be possible to replace this genericity hypothesis with a somewhat weaker hypothesis and still prove the main results of this paper (using the methods of this paper). Indeed no such hypothesis was needed in the totally ramified case [18]. On the other hand the discussion in Section 5 below shows that with these methods one cannot expect to remove the genericity hypothesis entirely even for $L = \mathbb{Q}_{p^2}$, and so we have to some extent favored cleaner combinatorics over optimizing the genericity hypothesis.

these two exceptional cases, there is a unique $\mu_{m,n}$ for each pair (J, d), and one checks that it is given by:

$$(3.6)\quad \begin{array}{lll} m_i = c_i + p - 1 - d_i, & n_i = b_i + 2d_i, & \text{if } i \in J \text{ and } i+1 \in J; \\ m_i = c_i + p - 1 - d_i, & n_i = b_i + 2d_i + 1, & \text{if } i \in J \text{ and } i+1 \notin J; \\ m_i = c_i + b_i + d_i - 1, & n_i = p - b_i - 2d_i, & \text{if } i \notin J \text{ and } i+1 \in J; \\ m_i = c_i + b_i + d_i, & n_i = p - 1 - b_i - 2d_i, & \text{if } i \notin J \text{ and } i+1 \notin J. \end{array}$$

We let $\mu(J, d)$ denote the weight $\mu_{m,n}$ with m, n defined by (3.6). In the two exceptional cases, we obtain in addition to $\mu(J, d)$ the weight $\mu'(J, d)$ defined as follows: if $b = d = (0, \dots, 0)$ and $J = \{0, \dots, f-1\}$, then $\mu'(J, d) = \mu_{m,n}$ where $m_i = c_i$ and $n_i = p-1$ for all i, and if $b = (p-1-2e', \dots, p-1-2e')$, $d = (e', \dots, e')$ and $J = \varnothing$, then $\mu'(J, d) = \mu_{m,n}$ where $m_i = c_i - e'$ and $n_i = p-1$ for all i.

We let W denote the subset of W' consisting of the $\mu(J, d)$. Note also that for (m, n) as in (3.6), we always have $n_i < p - 1$ for all i. It follows that the additional weights $\mu'(J, d)$ (when they occur) are not in W. Note also that both additional weights arise if $b = (0, \dots, 0)$ and $e' = (p-1)/2$, but comparing values of m shows they are distinct from each other. Moreover the following lemma shows that the weights $\mu(J, d)$ are distinct.

Lemma 3.7. *Suppose $J, J' \subseteq S$ and that $d = (d_0, \dots, d_{f-1})$ and $d' = (d'_0, \dots, d'_{f-1})$ are f-tuples of integers satisfying $0 \le d_i \le e' - 1$ if $i \in J$, $1 \le d_i \le e'$ if $i \notin J$, $0 \le d'_i \le e' - 1$ if $i \in J'$ and $1 \le d'_i \le e'$ if $i \notin J'$. If $\mu(J, d)$ is isomorphic to $\mu(J', d')$, then $J = J'$ and $d = d'$.*

Proof. Write $\mu(J, d) = \mu_{m,n}$ and $\mu(J', d') = \mu_{m',n'}$ with (m, n) and (m', n') as in (3.6). Twisting by χ_1^{-1}, we may suppose that $c_i = 0$ for all i. Then $0 \le m_i \le p - 1$ for all i, and $m_i > 0$ for some i, so that $0 < \sum_{i=1}^{f} m_i p^{f-i} \le p^f - 1$. Since the same is true for m', we must have $m_i = m'_i$ for all i. We claim that $J = J'$. Indeed if not, then without loss of generality there is some $i \in J$ such that $i \notin J'$, but then

$$m_i = p - 1 - d_i \ge p - e' > b_i + d_i \ge m'_i,$$

giving a contradiction. Since $J = J'$ and $m = m'$, it follows immediately that $d = d'$. □

We will now define partitions of W and W' into subsets indexed by A, where A is the set of f-tuples $a = (a_0, a_1, \dots, a_{f-1})$ with $0 \le a_i \le e'$ for all i. For $a \in A$, we let τ_a denote the (at most) tamely ramified principal series inertial type

$$\tau_a := \prod_{i=0}^{f-1} \tilde{\omega}_i^{c_i + b_i + a_i} \oplus \prod_{i=0}^{f-1} \tilde{\omega}_i^{c_i - a_i},$$

where $\tilde{\omega}_i$ denotes the Teichmüller lift of ω_i.

If τ is a principal series type, we let θ_τ denote the $\mathrm{GL}_2(\mathcal{O}_L)$-type associated to τ by the inertial local Langlands correspondence, viewed as a representation of $\mathrm{GL}_2(k)$. If $\tau = \tau_a$ then we write θ_a for θ_τ; so if τ_a is non-scalar then explicitly

$$\theta_a = \mathrm{Ind}_B^{\mathrm{GL}_2(k)} \left(\prod_{i=0}^{f-1} \tilde{\omega}_i^{c_i + b_i + a_i} \otimes \prod_{i=0}^{f-1} \tilde{\omega}_i^{c_i - a_i} \right)$$

where B is the subgroup of upper-triangular matrices in $\mathrm{GL}_2(k)$, $\psi_1 \otimes \psi_2$ denotes the character of B sending $\begin{pmatrix} x & * \\ 0 & y \end{pmatrix}$ to $\psi_1(x)\psi_2(y)$, and we recall that we are identifying characters $k^\times = \ell^\times \to \overline{\mathbb{Q}}_p^\times$ with characters $I_L \to \overline{\mathbb{Q}}_p^\times$ via the local Artin map with its geometric normalisation. Note that if τ_a is scalar, then $\sum_{i=0}^{f-1}(b_i+2a_i)p^{f-i} \equiv 0 \pmod{p^f-1}$, which occurs only if $a=b=(0,\dots,0)$, or $a=(e',\dots,e')$ and $b=(p-1-2e',\dots,p-1-2e')$. In this case we let

$$\theta_{\tau_a} = \theta_a = \det \circ \prod_{i=0}^{f-1} \tilde{\omega}_i^{c_i+b_i+a_i}, \quad \text{and} \quad \theta'_{\tau_a} = \theta'_a = \theta_a \otimes \mathrm{Ind}_B^{\mathrm{GL}_2(k)} \mathbf{1}.$$

We then define

$$\begin{array}{rcl} W_a & := & \{\, \mu \in W' \,:\, \mu \text{ is a Jordan–Hölder constituent of } \overline{\theta}_a \,\}, \\ \text{and } W'_a & := & \{\, \mu \in W' \,:\, \mu \text{ is a Jordan–Hölder constituent of } \overline{\theta}'_a \,\}. \end{array}.$$

We will see shortly that W_a is in fact contained in W. Note that $W'_a = W_a$ unless $a=b=(0,\dots,0)$ in which case $W_a = \{\mu(J,d)\}$ and $W'_a = \{\mu(J,d), \mu'(J,d)\}$ with $J=\{0,\dots,f-1\}$ and $d=(0,\dots,0)$, or $a=(e',\dots,e')$ and $b=(p-1-2e',\dots,p-1-2e')$ in which case $W_a=\{\mu(J,d)\}$ and $W'_a=\{\mu(J,d),\mu'(J,d)\}$ with $J=\varnothing$ and $d=(e',\dots,e')$.

Proposition 3.8.

(1) W (resp. W') is the disjoint union of the W_a (resp. W'_a) for $a \in A$.
(2) $|W_a| = 2^{f-\delta_a}$ where $\delta_a = |\{\, i \in \{0,\dots,f-1\} : a_i = 0 \text{ or } e' \,\}|$.
(3) If $\mu(J,d)$ or $\mu'(J,d) \in W'_a$, then $\sum_{i=0}^{f-1} a_i = \sum_{i=0}^{f-1} d_i$.

Proof. First note that to prove the proposition, we may twist $\overline{\rho}$ so as to assume $c_i = 0$ for $i=0,\dots,f-1$.

To prove (1), we must show that for each (J,d) as in the definition of W, there is a unique $a \in A$ such that $\mu(J,d)$ is a Jordan–Hölder constituent of $\overline{\theta}_a$. For this we use the explicit description of $\overline{\theta}_a^{\mathrm{ss}}$ given for example in [11, Prop. 1.1]. In particular the Jordan–Hölder constituents are of the form $\nu(a,J')$ for certain subsets $J' \subseteq \{0,\dots,f-1\}$, where $\nu(a,J') = \mu_{m',n'}$ with $m'=(m'_0,\dots,m'_{f-1})$ and $n'=(n'_0,\dots,n'_{f-1})$ defined by

$$\begin{array}{llll} m'_i = p-1-a_i, & n'_i = b_i+2a_i, & \text{if } i \in J' \text{ and } i+1 \in J'; \\ m'_i = p-a_i, & n'_i = b_i+2a_i-1, & \text{if } i \in J' \text{ and } i+1 \notin J'; \\ m'_i = b_i+a_i, & n'_i = p-2-b_i-2a_i, & \text{if } i \notin J' \text{ and } i+1 \in J'; \\ m'_i = b_i+a_i, & n'_i = p-1-b_i-2a_i, & \text{if } i \notin J' \text{ and } i+1 \notin J'. \end{array} \tag{3.9}$$

The constituents are then precisely the $\nu(a,J')$ for those J' such that $n'_i \geq 0$ for all i, except in the case that τ_a is scalar, in which case there is only one Jordan–Hölder constituent, namely $\nu(a,J')$ with $J'=\{0,\dots,f-1\}$ (resp. $J'=\varnothing$) if $a=b=(0,\dots,0)$ (resp. $a=(e',\dots,e')$ and $b=(p-1-2e',\dots,p-1-2e')$).

Suppose now that $\nu(a,J') = \mu_{m',n'} \simeq \mu_{m,n} = \mu(J,d)$. Note that $0 \leq m'_i \leq p$ for $i=0,1,\dots,f-1$; we will rule out the possibility that $m'_i = p$ for some i. Indeed if $m'_i = p$, then we must have $a_i = 0$, $i \in J'$ and $i+1 \notin J'$. It follows that $m'_{i+1} = b_{i+1} + a_{i+1} \leq p-2$, so that

$$0 < \sum_{j=0}^{f-1} m'_{i-j} p^j < p^f - 1.$$

Since $0 \le m_i \le p-1$ for all i (and not all 0), and

$$\sum_{j=0}^{f-1} m_{i-j}p^j \equiv \sum_{j=0}^{f-1} m'_{i-j}p^j \pmod{p^f-1},$$

we see that the sums are equal. Therefore $m_i \equiv m'_i \equiv 0 \pmod p$, so in fact $m_i = 0$. The definition of m_i then implies that $b_i = 0$, giving $n'_i = -1$, a contradiction. Note also that if $m'_i = 0$ for all i, then $a = b = (0, \dots, 0)$ and $J' = \varnothing$, which is also impossible. Since

$$\sum_{i=1}^{f} m_i p^{f-i} \equiv \sum_{i=1}^{f} m'_i p^{f-i} \pmod{p^f-1}$$

and both sums are between 1 and p^f-1 (inclusive), it follows that $(m,n) = (m',n')$.

Next we show that $J' = J$. If $i \in J$ and $i \notin J'$ for some i, then

$$m'_i = a_i + b_i \le p-1-e' < p-1-d_i = m_i,$$

giving a contradiction. If $i \notin J$ and $i \in J'$ for some i, then the inequalities

$$m_i \le b_i + d_i \le p-1-e' \le p-1-a_i \le m'_i$$

must all be equalities, which implies that $i+1 \notin J$, $i+1 \in J'$, $b_i = p-1-2e'$ and $a_i = e'$. Iterating gives $J' = \{0, \dots, f-1\}$, $b = (p-1-2e', \dots, p-1-2e')$ and $a = (e', \dots, e')$, which is impossible.

Having shown that $J' = J$, it follows that a is determined by the equation

$$a_i = \begin{cases} d_i + 1, & \text{if } i \in J,\ i+1 \notin J, \\ d_i - 1, & \text{if } i \notin J,\ i+1 \in J, \\ d_i, & \text{otherwise.} \end{cases} \tag{3.10}$$

As indeed $(m',n') = (m,n)$ in this case (as well as $n_i \ne -1$ and $a_i \in [0,e']$ for all i), this gives the assertion for W. The assertion for W' follows upon checking that when $b = (0, \dots, 0)$ (resp. $b = (p-1-2e', \dots, p-1-2e')$) the constituent $\mu'(J,d)$ with $J = \{0, \dots, f-1\}$ and $d = (0, \dots, 0)$ (resp. $J = \varnothing$ and $d = (e', \dots, e')$) is not contained in W_a with $a \ne (0, \dots, 0)$ (resp. $a \ne (e', \dots, e')$).

To prove (2) we fix a and determine the $J \subseteq \{0, \dots, f-1\}$ for which (3.10) holds for some d as in the definition of W. The condition that $0 \le d_i \le e'-1$ if $i \in J$ and $1 \le d_i \le e'$ if $i \notin J$ translates into the condition that $0 \le a_i \le e'-1$ if $i+1 \in J$, and $1 \le a_i \le e'$ if $i+1 \notin J$. Therefore the only restrictions on J are that $i+1 \in J$ if $a_i = 0$, and that $i+1 \notin J$ if $a_i = e'$. The number of such J is $2^{f-\delta_a}$ as required.

Part (3) in the case $\mu(J,d) \in W'_a$ is immediate from (3.10) on noting that there are the same number of i satisfying $i \in J$, $i+1 \notin J$ as there are satisfying $i \notin J$, $i+1 \in J$. The formula in the case $\mu'(J,d) \in W'_a$ is immediate from the definitions. □

Remark 3.11. We remark that Schein [29, Prop. 3.2] also gives a decomposition of W' into subsets which are typically of cardinality 2^f, but it is visibly different from ours; for example, it is a decomposition into $(e')^f$ subsets rather than $(e'+1)^f$, and if $f = 1$, then the subsets are constituents of the reduction of a supercuspidal rather than principal series type.

Recall that $L(\chi_1, \chi_2, \tau)$ denotes the set of all extensions of χ_1 by χ_2 that arise as the generic fibre of a model of type τ. We translate Corollary 2.17 into the notation of this section.

Theorem 3.12. *For any $a, a' \in A$, we have*

(1) $\dim_{k_E} L(\chi_1, \chi_2, \tau_a) = \sum_{i=0}^{f-1}(e' - a_i)$,

(2) $L(\chi_1, \chi_2, \tau_a) \cap L(\chi_1, \chi_2, \tau_{a'}) = L(\chi_1, \chi_2, \tau_{a''})$ *where* $a''_i = \max(a_i, a'_i)$.

Proof. Reduce to the case of $\chi_1 = 1$ by twisting. Our genericity hypothesis rules out $\chi_2 = 1$. Note that for the type τ_a we have $\nu'_i = p - 1 - a_i$ and $\nu_i = a_i + b_i$, except that when $a = b = 0$ we (by convention) have $\nu_i = p - 1$ for all i. The conditions $\nu'_i \in [p-1-e', p-1]$, $\nu'_i \geq \nu_i$, and $\nu'_i + \nu_i \geq p - 1$ are all easily checked. Now $\mu_i = e' - a_i$ and (in the notation of Corollary 2.17) if $(\tau, \dot\tau) = (\tau_a, \tau_{a'})$ then $\ddot\tau = \tau_{a''}$, as desired. □

4. The main results

4.1. **The global setting.** Let F be a totally real field and $\overline{\rho} : G_F \to \mathrm{GL}_2(k_E)$ a continuous representation. We suppose that $\overline{\rho}$ is automorphic in the sense that it arises as the reduction of a p-adic representation of G_F associated to a cuspidal Hilbert modular eigenform of some weight and level, or equivalently to a cuspidal holomorphic automorphic representation of $\mathrm{GL}_2(\mathbb{A}_F)$. We fix a place v of F dividing p, and we let $L = F_v$, so that $k_v = k = \ell$ in what follows.

Let D be a quaternion algebra over F satisfying the following hypotheses:

- D is split at all primes dividing p;
- D is split at at most one infinite place of F;
- If w is a finite place of F at which D is ramified, then $\overline{\rho}|_{G_{F_w}}$ is either irreducible, or equivalent to a representation of the form $\psi \otimes \begin{pmatrix} \overline{\epsilon} & * \\ 0 & 1 \end{pmatrix}$ for some character $\psi : G_F \to k_E^\times$.

We let r denote the number of infinite places of F at which D is split (so $r = 0$ or 1), and if $r = 1$ we let ξ denote that infinite place and fix an isomorphism $D_\xi \simeq M_2(\mathbb{R})$. We also fix a maximal order $\mathcal{O}_D$ of D and an isomorphism $\mathcal{O}_{D_v} \simeq M_2(\mathcal{O}_L)$.

Remark 4.1. The hypothesis that D is split at all primes dividing p is made only to be able to invoke the results of [16] on the weight part of Serre's Conjecture. We expect however that the proofs of the required variants of their results, and hence the main results of this paper, carry over if we only require that D is split at v, without specifying the behavior of D at the other primes dividing p.

For any open compact subgroup U of $D_f^\times = (D \otimes \widehat{\mathbb{Z}})^\times$, we let X_U denote the associated Shimura variety of dimension r:

$$X_U = D^\times \backslash ((\mathfrak{H}^\pm)^r \times D_f^\times)/U,$$

where if $r = 1$ then $D^\times$ acts on $\mathfrak{H}^\pm = \mathbb{C} - \mathbb{R}$ via the isomorphism $D_\xi^\times \simeq \mathrm{GL}_2(\mathbb{R})$, and we let $S^D(U) = H^r(X_U, k_E)$. (Recall that r and ξ are defined just before Remark 4.1.) Let Σ_U denote the set of all finite places w of F such that (i) w does not divide p, (ii) D is split at w, (iii) U contains $\mathcal{O}_{D_w}^\times$, and (iv) $\overline{\rho}$ is unramified at w. Then $S^D(U)$ is equipped with the commuting action of Hecke operators T_w and S_w for all $w \in \Sigma_U$, hence with an action of the polynomial algebra over k_E generated by variables T_w and S_w for $w \in \Sigma_U$. We denote this algebra by $\mathbb{T}^{\Sigma_U}$, and let $\mathfrak{m}_{\overline{\rho}}^{\Sigma_U}$ denote the kernel of the k_E-algebra homomorphism $\mathbb{T}^{\Sigma_U} \to k_E$ defined by

$$T_w \mapsto \mathrm{Nm}(w)\,\mathrm{Trace}(\overline{\rho}(\mathrm{Frob}_w)); \qquad S_w \mapsto \mathrm{Nm}(w)\det(\overline{\rho}(\mathrm{Frob}_w))$$

for $w \in \Sigma^U$. We let $S^D(U)[\mathfrak{m}_{\overline{\rho}}^{\Sigma_U}]$ denote that set of $x \in S^D(U)$ such that $Tx = 0$ for all $T \in \mathfrak{m}_{\overline{\rho}}^{\Sigma_U}$.

Now let U_v denote the kernel of the map $\mathcal{O}_{D_v}^\times \to \mathrm{GL}_2(k)$ defined by composing the restriction of our fixed $\mathcal{O}_{D_v} \simeq M_2(\mathcal{O}_L)$ with reduction mod v. If $U = U_v U^v$ for some open compact $U^v \subseteq \ker(D_f^\times \to D_v^\times)$, then the natural right action of $\mathcal{O}_{D,v}^\times$ on X_U induces a left action of $\mathrm{GL}_2(k)$ on $S^D(U)$ which commutes with that of $\mathbb{T}^{\Sigma_U}$.

Definition 4.2. If μ is an irreducible representation of $\mathrm{GL}_2(k)$ over k_E, then we say that $\overline{\rho}$ is *modular of weight μ with respect to D and v* if

$$\mathrm{Hom}_{k_E[\mathrm{GL}_2(k)]}(\mu, S^D(U)[\mathfrak{m}_{\overline{\rho}}^{\Sigma_U}]) \neq 0$$

for some open compact subgroup $U = U_v U^v$ as above. We let $W_{\mathrm{mod}}^{D,v}(\overline{\rho})$ denote the set of Serre weights for which $\overline{\rho}$ is modular with respect to D and v.

The weight part of Serre's Conjecture for $\overline{\rho}$ (at v, with respect to D) states that

$$W_{\mathrm{mod}}^{D,v}(\overline{\rho}) = W_{\mathrm{expl}}(\overline{\rho}|_{G_L}), \tag{4.3}$$

where $W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$ is the set of predicted Serre weights as in [1, Def. 4.1.14], recalled in Definition 3.3 above in the case that $\overline{\rho}|_{G_L}$ is reducible.

One of the inclusions in (4.3) is proved under mild technical hypotheses by Gee and Kisin in [16], building on [15, 20, 18, 1, 17]. More precisely, we have the following result (*cf.* [16, Def. 5.5.3, Cor. 5.5.4]):

Theorem 4.4. *Suppose that $p > 2$, $\overline{\rho}|_{G_{F(\zeta_p)}}$ is irreducible, and if $p = 5$, then the projective image of $\overline{\rho}|_{G_{F(\zeta_5)}}$ is not isomorphic to A_5. Then the following hold:*

(1) $W_{\mathrm{mod}}^{D,v}(\overline{\rho})$ depends only on $\overline{\rho}|_{G_L}$;
(2) $W_{\mathrm{expl}}(\overline{\rho}|_{G_L}) \subseteq W_{\mathrm{mod}}^{D,v}(\overline{\rho})$;
(3) $W_{\mathrm{expl}}(\overline{\rho}|_{G_L}) = W_{\mathrm{mod}}^{D,v}(\overline{\rho})$ if L is unramified or totally ramified over $\mathbb{Q}_p$.

In particular the theorem ensures (under its hypotheses) that $W_{\mathrm{mod}}^{D,v}(\overline{\rho})$ is independent of the choice of D, which we henceforth suppress from the notation.

We will now restrict to the case where $\overline{\rho}|_{G_L}$ is reducible and generic (see Definition 3.5). Our main global result is the following:

Theorem 4.5. *If $\overline{\rho}$ is as in Theorem 4.4 and $\overline{\rho}|_{G_L}$ is reducible and generic, then*

$$W_{\mathrm{expl}}(\overline{\rho}|_{G_L}) = W_{\mathrm{mod}}^{v}(\overline{\rho}) \cap W_{\mathrm{expl}}(\overline{\rho}|_{G_L}^{\mathrm{ss}}).$$

In other words we prove the weight part of Serre's Conjecture holds in this case for weights in $W_{\mathrm{expl}}(\overline{\rho}|_{G_L}^{\mathrm{ss}})$. We will prove this theorem in §4.3 along with our main results in the local setting stated in §4.2.

Remark 4.6. When $p = 3$, the hypothesis that $\overline{\rho}$ is generic implies that $e' = 1$. Since the weight part of Serre's conjecture in the unramified case (i.e. the Buzzard–Diamond–Jarvis conjecture) is already known in full [1, 17, 16], Theorem 4.5 provides new information only when $p > 3$.

4.2. **The local setting.** We will now revert to the setting of §3, where $\overline{\rho} : G_L \to \mathrm{GL}_2(k_E)$ is a reducible representation, written as

$$\overline{\rho} \simeq \begin{pmatrix} \chi_2 & * \\ 0 & \chi_1 \end{pmatrix};$$

moreover we assume $\overline{\rho}$ is generic.

Suppose now that μ is a Serre weight in W in the notation of §3.2. Recall that W is a subset of $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$ with complement of cardinality at most 2, and that $\mu = \mu(J,d)$ for some (J,d) satisfying (3.4), where $J \subseteq \{0,\ldots,f-1\}$ and $d = (d_0,\ldots,d_{f-1})$ with $0 \le d_i \le e'-1$ if $i \in J$, $1 \le d_i \le e'$ if $i \notin J$.

Now choose a lift $\tilde{\sigma}_i : L \hookrightarrow E$ of σ_i for each $i \in \{0,\ldots,f-1\}$ and a subset $\tilde{J} \subseteq \{\sigma : L \hookrightarrow E\}$ such that

- $\tilde{\sigma}_i \in \tilde{J}$ if and only if $i \in J$, and
- $\{\, \sigma \in \tilde{J} : \sigma \text{ is a lift of } \sigma_i \,\}$ has cardinality $e' - d_i$.

Choose also a crystalline character $\tilde{\chi}_1 : G_L \to E^\times$ lifting χ_1 whose Hodge–Tate module V_1 has σ-labelled weights:

- 1, if $\sigma \notin \tilde{J}$ and $\sigma \notin \{\, \tilde{\sigma}_i : i = 0,\ldots,f-1 \,\}$;
- 0, if $\sigma \in \tilde{J}$ and $\sigma \notin \{\, \tilde{\sigma}_i : i = 0,\ldots,f-1 \,\}$;
- $m_i + n_i + 1$, if $\sigma = \tilde{\sigma}_i \notin J$;
- m_i, if $\sigma = \tilde{\sigma}_i \in J$.

That such a crystalline character exists follows for example from Lubin–Tate theory, or from [10, Prop. B.3]; moreover such a character is unique up to an unramified twist with trivial reduction. Similarly, let $\tilde{\chi}_2 : G_L \to E^\times$ be a lift of χ_2 whose Hodge–Tate module V_2 has σ-labelled weights:

- 0, if $\sigma \notin \tilde{J}$ and $\sigma \notin \{\, \tilde{\sigma}_i : i = 0,\ldots,f-1 \,\}$;
- 1, if $\sigma \in \tilde{J}$ and $\sigma \notin \{\, \tilde{\sigma}_i : i = 0,\ldots,f-1 \,\}$;
- m_i, if $\sigma = \tilde{\sigma}_i \notin J$;
- $m_i + n_i + 1$, if $\sigma = \tilde{\sigma}_i \in J$.

Note that $V_1 \oplus V_2$ is a Hodge–Tate module lifting μ.

We let $L_{\mathrm{cris},E}(\tilde{\chi}_1,\tilde{\chi}_2)$ denote the subspace of $\mathrm{Ext}^1_{E[G_L]}(\tilde{\chi}_1,\tilde{\chi}_2)$ corresponding to the set of extensions which are crystalline. We let $L_{\mathrm{cris},\mathcal{O}_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ denote the preimage of $L_{\mathrm{cris},E}(\tilde{\chi}_1,\tilde{\chi}_2)$ in $\mathrm{Ext}^1_{\mathcal{O}_E[G_L]}(\tilde{\chi}_1,\tilde{\chi}_2)$, and let $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ denote the image of $L_{\mathrm{cris},\mathcal{O}_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ in $\mathrm{Ext}^1_{k_E[G_L]}(\chi_1,\chi_2)$. Thus $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ consists of the set of extensions arising as reductions of crystalline representations of the form $\begin{pmatrix} \tilde{\chi}_2 & * \\ 0 & \tilde{\chi}_1 \end{pmatrix}$.

Recall that we have defined a partition of W into subsets W_a indexed by f-tuples $(a_0,a_1,\ldots,a_{f-1})$ with $0 \le a_i \le e'$ for all i (Proposition 3.8). Our main result comparing reductions of crystalline and potentially Barsotti–Tate extensions is the following.

Theorem 4.7. *If $\mu \in W_a$, then*

$$L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2) = L(\chi_1,\chi_2,\tau_a).$$

Remark 4.8. Note in particular that not only is $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ independent of the weight $\mu \in W_a$, but also of the various choices of lifts $\tilde{\sigma}_i$, $\tilde{J}$, $\tilde{\chi}_1$ and $\tilde{\chi}_2$.

Next we state the main result on the possible forms of $W_{\mathrm{expl}}(\overline{\rho})$. Recall that the partition of W into the W_a extends to a partition of $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$ into subsets W'_a defined in §3.2. To treat the case that $\overline{\rho}$ is equivalent to a representation of the form $\chi_1 \otimes \begin{pmatrix} \overline{\epsilon} & * \\ 0 & 1 \end{pmatrix}$, recall that such a representation is *très ramifiée* if the splitting field of its projective image is *not* of the form $L(\alpha_1^{1/p},\ldots,\alpha_s^{1/p})$ for some $\alpha_1,\ldots,\alpha_s \in \mathcal{O}_L^\times$.

Theorem 4.9. *We have*

$$W_{\mathrm{expl}}(\overline{\rho}) = \coprod_{a \le a^{\max}} W'_a$$

for some $a^{\max} \in A$ *depending on* $\overline{\rho}$*, unless* $\overline{\rho}$ *is très ramifiée, in which case* $W_{\mathrm{expl}}(\overline{\rho}) = \{\mu_{m,n}\}$ *where* $\chi_1|_{I_L} = \prod_{i=0}^{f-1} \omega_i^{m_i}$ *and* $n = (p-1, \dots, p-1)$*.*

Remark 4.10. It will also be clear from the proof that given any pair of characters χ_1, χ_2 such that $\chi_1 \oplus \chi_2$ is generic, every element of A arises as $a^{\max}$ for some peu ramifiée extension of χ_1 by χ_2. The theorem therefore completely determines the possible values of $W_{\mathrm{expl}}(\overline{\rho})$ for generic $\overline{\rho}$. As indicated in the footnote after Definition 3.5, we expect a similar description to be valid under hypotheses weaker than genericity, but not in full generality; see Section 5.

4.3. **The proofs.** In this section we will prove Theorems 4.5, 4.7 and 4.9, but first we note the following lemma.

Lemma 4.11. *Suppose that* $\overline{\rho} : G_F \to \mathrm{GL}_2(k_E)$ *is as in Theorem 4.4 with* $\overline{\rho}|_{G_L} \simeq \begin{pmatrix} \chi_2 & * \\ 0 & \chi_1 \end{pmatrix}$*, and that* τ *is a principal series type. If* $W^v_{\mathrm{mod}}(\overline{\rho})$ *contains a Jordan–Hölder factor of* $\overline{\theta}_\tau$*, then the extension defined by* $\overline{\rho}|_{G_L}$ *is in* $L(\chi_1, \chi_2, \tau)$*.*

Proof. By [5, Prop. 2.10] (stated there only for $r = 1$ and p unramified in F, but the proof carries over to our setting; see also the proof of [20, Lem. 3.4]), we have (replacing E by an extension E' if necessary) that $\overline{\rho} \simeq \overline{\rho}_\pi$ for some cuspidal holomorphic automorphic representation π of $\mathrm{GL}_2(\mathbb{A}_F)$ such that π_∞ is holomorphic of weight $(2, \dots, 2)$ with trivial central character and π_v has K-type θ_τ. (Note that our normalizations for the local Langlands correspondence differ from those of [5]; in our case $I(\psi_1 \otimes \psi_2)$ corresponds to $|\cdot|^{1/2}(\psi_1 \oplus \psi_2)$.)

Local–global compatibility at v of the Langlands correspondence (see the Corollary in the introduction to [22]) therefore implies that $\rho_\pi|_{G_L}$ is potentially Barsotti–Tate with associated Weil–Deligne representation of type τ. (Note that when τ is scalar, by definition θ_τ is not a twist of the Steinberg representation.) It follows from [21, Cor. 5.2] that $\overline{\rho}|_{G_L}$ is the generic fibre of a model of type τ in the sense of Definition 2.2, and hence that its associated extension class is in $L(\chi_1, \chi_2, \tau)$. Note furthermore that replacing E by E' has the effect of replacing $L(\chi_1, \chi_2, \tau)$ by $L(\chi_1, \chi_2, \tau) \otimes_{k_E} k_{E'}$ (for example as an application of Theorem 1.11), so that the conclusion holds without having extended scalars. ☐

Proof of Theorem 4.7. Since $\tilde{\chi}_1$ and $\tilde{\chi}_2$ are distinct characters, it follows from [23, Prop. 1.24(2)] that

$$\dim_E L_{\mathrm{cris},E}(\tilde{\chi}_1, \tilde{\chi}_2) = \dim_E(V/\operatorname{Fil}^0(V)),$$

where V is the Hodge–Tate module $\mathrm{Hom}_E(V_1, V_2)$. Note that $\dim_E(V/\operatorname{Fil}^0(V))$ is simply the number of $\sigma : L \hookrightarrow E$ such that the σ-labelled Hodge–Tate weight of V_2 is greater than that of V_1, which is the case if and only if $\sigma \in \tilde{J}$. Therefore

$$\dim_E L_{\mathrm{cris},E}(\tilde{\chi}_1, \tilde{\chi}_2) = |\tilde{J}| = \sum_{i=0}^{f-1} (e' - d_i).$$

The genericity hypothesis implies that $\chi_1 \neq \chi_2$, from which it follows that $\mathrm{Ext}^1_{\mathcal{O}_E[G_L]}(\tilde{\chi}_1, \tilde{\chi}_2)$ is torsion-free. Therefore $L_{\mathrm{cris},\mathcal{O}_E}(\tilde{\chi}_1, \tilde{\chi}_2)$ is torsion-free over $\mathcal{O}_E$

of rank $\dim_E L_{\mathrm{cris},E}(\tilde{\chi}_1, \tilde{\chi}_2)$, and it follows that

$$\dim_{k_E} L_{\mathrm{cris},k_E}(\tilde{\chi}_1, \tilde{\chi}_2) = \sum_{i=0}^{f-1}(e' - d_i).$$

By Proposition 3.8(3) and Theorem 3.12(1), this is the same as the dimension of $L(\chi_1, \chi_2, \tau_a)$, so it suffices to prove that

$$L_{\mathrm{cris},k_E}(\chi_1, \chi_2) \subseteq L(\chi_1, \chi_2, \tau_a).$$

Moreover since these subspaces of $\mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2)$ are well-behaved under extension of scalars, we may enlarge E in order to prove the inclusion.

Suppose now that we are given a representation $\overline{\varrho} : G_L \to \mathrm{GL}_2(k_E)$ giving rise to an extension class in $L_{\mathrm{cris},k_E}(\chi_1, \chi_2)$. By [16, Cor. A.3] we have that $\overline{\varrho} \simeq \overline{\rho}|_{G_L}$ for some totally real field F, representation $\overline{\rho} : G_F \to \mathrm{GL}_2(k_E)$ and embedding $F \hookrightarrow L$ such that Theorem 4.4 applies (enlarging E if necessary). Since $\mu \in W_{\mathrm{expl}}(\overline{\varrho})$, Theorem 4.4 implies that $\mu \in W^v_{\mathrm{mod}}(\overline{\rho})$, and hence Lemma 4.11 implies that the extension defined by $\overline{\varrho}$ is in $L(\chi_1, \chi_2, \tau_a)$. □

Proof of Theorem 4.9. From Theorem 4.7, Remark 4.8, and the definitions of $W_{\mathrm{expl}}(\overline{\rho})$ and $L_{\mathrm{cris},k_E}(\tilde{\chi}_1, \tilde{\chi}_2)$, we see that if $\mu \in W_a$, then $\mu \in W_{\mathrm{expl}}(\overline{\rho})$ if and only if the extension class associated to $\overline{\rho}$ is in $L(\chi_1, \chi_2, \tau_a)$. Let $A_{\overline{\rho}}$ denote the set of $a \in A$ for which this holds, so that

$$W_{\mathrm{expl}}(\overline{\rho}) \cap W = \coprod_{a \in A_{\overline{\rho}}} W_a.$$

By Theorem 3.12(1), we have $\dim_{k_E} L(\chi_1, \chi_2, \tau_{(0,\ldots,0)}) = e'f$. If $\chi_2 \neq \chi_1\overline{\epsilon}$, then this is the same as the dimension of

$$\mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2) \simeq H^1(G_L, \chi_1^{-1}\chi_2),$$

so we have that $(0, \ldots, 0) \in A_{\overline{\rho}}$, and in particular $A_{\overline{\rho}}$ is nonempty. In the case that $\chi_2 = \chi_1\overline{\epsilon}$, we have the isomorphism

$$\mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2) \simeq L^\times/(L^\times)^p \otimes k_E$$

of Kummer theory. Note that the genericity hypothesis implies that $\zeta_p \notin L$, so these spaces have dimension $e'f + 1$. The subspace $\mathcal{O}_L^\times/(\mathcal{O}_L^\times)^p \otimes k_E$ has dimension $e'f$, and the corresponding classes arise as generic fibres of models of type $\tau_{(0,\ldots,0)} = (\chi_1 \oplus \chi_1)|_{I_L}$. To see this, twist by χ_1^{-1} to reduce to the case where $\tau_{(0,\ldots,0)}$ is trivial, and then apply [12, Prop. 8.2] (or more precisely the analogous statement with L in place of $\mathbb{Q}_p$, which follows by the same proof). Therefore $L(\chi_1, \chi_2, \tau_{(0,\ldots,0)})$ corresponds to $\mathcal{O}_L^\times/(\mathcal{O}_L^\times)^p \otimes k_E$, and $(0, \ldots, 0) \in A_{\overline{\rho}}$ if and only if $\overline{\rho}$ is not très ramifiée.

Suppose now that $\overline{\rho}$ is not très ramifiée. In particular $A_{\overline{\rho}}$ is nonempty and Theorem 3.12(2) implies that

$$\bigcap_{a \in A_{\overline{\rho}}} L(\chi_1, \chi_2, \tau_a) = L(\chi_1, \chi_2, \tau_{a^{\max}})$$

where $a_i^{\max} = \max_{a \in A_{\overline{\rho}}}\{a_i\}$. Moreover $a \in A_{\overline{\rho}}$ if and only $a \leq a^{\max}$, so

$$W_{\mathrm{expl}}(\overline{\rho}) \cap W = \coprod_{a \leq a^{\max}} W_a.$$

On the other hand if $\overline{\rho}$ is très ramifiée, then we see that $W_{\mathrm{expl}}(\overline{\rho}) \cap W = \varnothing$.

To complete the proof of the theorem, we must treat the two possible additional weights $\mu'(J,d)$ arising when $\chi_1^{-1}\chi_2|_{I_L} = \overline{\epsilon}|_{I_L}^{\pm 1}$.

Note that the dimension calculations at the beginning of the proof of Theorem 4.7 apply equally with $n = (0,\dots,0)$ replaced by $n = (p-1,\dots,p-1)$. In the case $\chi_1|_{I_L} = \chi_2\overline{\epsilon}|_{I_L}$, $J = \varnothing$ and $d = (e',\dots,e')$, this gives $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2) = \{0\}$, so that

$$\mu'(J,d) \in W_{\mathrm{expl}}(\overline{\rho}) \quad \Leftrightarrow \quad \overline{\rho} \text{ splits} \quad \Leftrightarrow \quad \mu(J,d) \in W_{\mathrm{expl}}(\overline{\rho}).$$

In the case $\chi_2|_{I_L} = \chi_1\overline{\epsilon}|_{I_L}$, $J = \{0,\dots,f-1\}$ and $d = (0,\dots,0)$, we find that

$$\dim_{k_E} L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2) = e'f,$$

and we must show that $\mu'(J,d) \in W_{\mathrm{expl}}(\overline{\rho})$. If $\chi_1 \neq \chi_2\overline{\epsilon}$, then this holds since

$$L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2) = \mathrm{Ext}^1_{k_E[G_L]}(\chi_1,\chi_2).$$

If $\chi_1 = \chi_2\overline{\epsilon}$, then we must show that every class in

$$\mathrm{Ext}^1_{k_E[G_L]}(\chi_1,\chi_2) \simeq H^1(G_L, k_E(\overline{\epsilon}))$$

is in the codimension one subspace $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ for some choice of lifts $\tilde{\chi}_1$, $\tilde{\chi}_2$ as in §4.2 with $n = (p-1,\dots,p-1)$, enlarging E if necessary. This follows by exactly the same proof as that of [18, Prop. 5.2.9]. □

Proof of Theorem 4.5. We must show that if $\mu \in W' \cap W^v_{\mathrm{mod}}(\overline{\rho})$, then $\mu \in W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$. Suppose first that $\mu \in W_a$ for some $a \in A$. By Lemma 4.11, we have that the extension class associated to $\overline{\rho}|_{G_L}$ is in $L(\chi_1,\chi_2,\tau_a)$, hence in $L_{\mathrm{cris},k_E}(\tilde{\chi}_1,\tilde{\chi}_2)$ by Theorem 4.7, and therefore in $W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$.

Now we must deal with the two exceptional weights. If $\chi_2|_{I_L} = \chi_1\overline{\epsilon}|_{I_L}$, then Theorem 4.9 implies that $\mu'(J,d) \in W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$, where $J = \{0,\dots,f-1\}$ and $d = (0,\dots,0)$. Finally suppose that $\chi_1|_{I_L} = \chi_2\overline{\epsilon}|_{I_L}$ and that $\mu'(J,d) \in W^v_{\mathrm{mod}}(\overline{\rho})$, where $J = \varnothing$ and $d = (p-1,\dots,p-1)$. In this case the same argument as in the proof of Lemma 4.11 shows that $\overline{\rho} \simeq \overline{\rho}_\pi$ for some cuspidal holomorphic automorphic representation π of $\mathrm{GL}_2(\mathbb{A}_F)$ such that π_∞ is holomorphic of weight $(2,\dots,2)$ with trivial central character and $\psi \otimes \pi_v$ has a vector invariant under $U_0(v)$, where $\psi = [\chi_2]^{-1} \circ \det$ and $[\chi_2]$ denotes the Teichmüller lift of χ_2. Therefore $\psi \otimes \pi_v$ is either an unramified principal series, or an unramified twist of the Steinberg representation. If $\psi \otimes \pi_v$ is unramified, then in fact $\overline{\rho}$ is modular of weight $\mu(J,d)$, so it follows from the cases already proved that $\mu(J,d) \in W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$, and hence $\mu'(J,d) \in W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$ by Theorem 4.9. (In fact we see from the proof of Theorem 4.9 that in this case $\overline{\rho}|_{G_L}$ is split.) If $\psi \otimes \pi_v$ is an unramified twist of the Steinberg representation, then local-global compatibility at v gives that $\rho_\pi|_{G_L}$ is an unramified twist of a representation of the form $[\chi_2] \otimes \begin{pmatrix} \epsilon & * \\ 0 & 1 \end{pmatrix}$. Since $\chi_1|_{I_L} \neq \chi_2|_{I_L}$, it follows that $\overline{\rho}|_{G_L}$ is split, and hence that $\mu'(J,d) \in W_{\mathrm{expl}}(\overline{\rho}|_{G_L})$ in this case as well. □

5. A remark on the genericity hypothesis

In this section we show that the genericity hypothesis on $\overline{\rho}$ is, in general, necessary in order for our arguments in the proof of Theorem A to go through. That is, we give an example of a field L, characters χ_1, $\chi_2 : G_L \to \overline{\mathbb{F}}_p^\times$, and a weight μ such that the subset $L_{\mathrm{cris}} \subseteq H^1(G_L, \overline{\mathbb{F}}_p(\chi_2\chi_1^{-1}))$ corresponding to the representations

$\overline{\rho}$ with $\mu \in W_{\mathrm{expl}}(\overline{\rho})$ is not equal to the space $L(\chi_1, \chi_2, \tau)$ for any principal series type τ such that μ is a Jordan–Hölder constituent of $\overline{\theta}_\tau$.

Let $L = \mathbb{Q}_{p^2}$ be the unramified quadratic extension of $\mathbb{Q}_p$, so that $f = 2$ and $e' = 1$. Take χ_1 to be trivial, and $\chi_2 = \chi$ to be any extension to G_L of $\omega_0^{p-1}\omega_1^b$, where $b \in [1, p-2]$ is an integer. Observe that the weight $\mu_{m,n}$ with

$$m = (p-1, b-1), \qquad n = (p-1, p-b-1)$$

lies in $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$. One checks that $J = \{0\}$ is the only subset $J \subseteq \{0,1\}$ such that

$$\prod_{i \in J} \omega_i^{m_i + d_i} \prod_{i \notin J} \omega_i^{m_i + n_i + d_i} = 1$$

with $d_i = 0$ if $i \in J$ and $d_i = 1$ otherwise. It follows as in the proof of Theorem 4.7 that $\dim_{k_E} L_{\mathrm{cris},k_E}(\tilde{\chi}_1, \tilde{\chi}_2) = (1-1) + (1-0) = 1$, where $\tilde{\chi}_1$ and $\tilde{\chi}_2$ are defined as in §2.2. By [8, Rem. 7.13], the spaces $L_{\mathrm{cris},k_E}(\tilde{\chi}_1, \tilde{\chi}_2)$ are independent of the choice of $\tilde{\chi}_1$, $\tilde{\chi}_2$, so in fact $\dim_{k_E} L_{\mathrm{cris}} = 1$.

On the other hand, one checks (e.g. from [11, Prop. 1.1]) that there is exactly one principal series type τ such that μ is a Jordan–Hölder constituent of $\overline{\theta}_\tau$, namely

$$\tau \simeq \omega_0^{p-2}\omega_1^{p-1} \oplus \omega_0^{p-1}\omega_1^{b-1}.$$

The weight $\mu_{m',n'}$ with $m' = (0,0)$ and $n' = (p-2, b-1)$ is also a Jordan-Hölder constituent of $\overline{\theta}_\tau$ as well as an element of $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$; moreover the space $L_{\mathrm{cris},k_E}(\tilde{\chi}'_1, \tilde{\chi}'_2)$ corresponding to $\mu_{m',n'}$ has dimension 2, hence is equal to $\mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2)$. As in the proof of Theorem 4.7, we have that $L_{\mathrm{cris},k_E}(\tilde{\chi}'_1, \tilde{\chi}'_2) \subseteq L(\chi_1, \chi_2, \tau)$, so that $L(\chi_1, \chi_2, \tau) = \mathrm{Ext}^1_{k_E[G_L]}(\chi_1, \chi_2)$ properly contains L_{cris}.

We remark however that in the case $f = 2$ and $e' = 1$, Corollary 5.13 and Theorem 7.12 of [8] yield a partition of $W_{\mathrm{expl}}(\overline{\rho}^{\mathrm{ss}})$ into subsets W'_a for $a \in \{0,1\}^2$ such that Theorem 4.9 holds even without the genericity hypothesis. Indeed in the example above (with $\chi_1^{-1}\chi_2|_{I_L} = \omega_0^{p-1}\omega_1^b$ for some $b \in [1, p-2]$), one even has that each W'_a is a singleton exactly as in the generic case. On the other hand if $\chi_1^{-1}\chi_2|_{I_L} = \omega_1^b$ for some $b \in [1, p-1]$, then one of the two subsets W'_a with $a_0 + a_1 = 1$ must be empty, while the other has cardinality two.

References

[1] Tom Barnet-Lamb, Toby Gee, and David Geraghty. Serre weights for rank two unitary groups. *Math. Ann.*, page to appear, 2013.

[2] Christophe Breuil. Groupes p-divisibles, groupes finis et modules filtrés. *Ann. of Math. (2)*, 152(2):489–549, 2000.

[3] Christophe Breuil, Brian Conrad, Fred Diamond, and Richard Taylor. On the modularity of elliptic curves over **Q**. *J.A.M.S.*, 14:843–939, 2001.

[4] Christophe Breuil and Vytautas Paškūnas. Towards a modulo p Langlands correspondence for GL_2. *Mem. Amer. Math. Soc.*, 216(1016):vi+114, 2012.

[5] Kevin Buzzard, Fred Diamond, and Frazer Jarvis. On Serre's conjecture for mod l Galois representations over totally real fields. *Duke Math. J.*, 155(1):105–161, 2010.

[6] Xavier Caruso. Schémas en groupes et poids de Diamond–Serre. arXiv:0705.1213.

[7] Xavier Caruso and Tong Liu. Quasi-semi-stable representations. *Bull. Soc. Math. France*, 137(2):185–223, 2009.

[8] Seunghwan Chang and Fred Diamond. Extensions of rank one (ϕ, Γ)-modules and crystalline representations. *Compos. Math.*, 147(2):375–427, 2011.

[9] Chuangxun Cheng. Rank two Breuil modules: basic structures. *J. Number Theory*, 132(11):2379–2396, 2012.

[10] Brian Conrad. Lifting global representations with local properties. preprint, 2011.

[11] Fred Diamond. A correspondence between representations of local Galois groups and Lie-type groups. In *L-functions and Galois representations*, pages 187–206. Cambridge University Press, Cambridge, 2007.
[12] Bas Edixhoven. The weight in Serre's conjectures on modular forms. *Invent. math.*, 109(3):563–594, 1992.
[13] Matthew Emerton, Toby Gee, Florian Herzig, and David Savitt. General Serre weight conjectures. in preparation, 2013.
[14] Toby Gee. Automorphic lifts of prescribed types. *Math. Ann.*, 350(1):107–144, 2011.
[15] Toby Gee. On the weights of mod p Hilbert modular forms. *Inventiones Mathematicae*, 184:1–46, 2011. 10.1007/s00222-010-0284-5.
[16] Toby Gee and Mark Kisin. The Breuil-Mézard Conjecture for potentially Barsotti-Tate representations. http://arxiv.org/abs/1208.3179.
[17] Toby Gee, Tong Liu, and David Savitt. The Buzzard-Diamond-Jarvis conjecture for unitary groups. preprint, 2012.
[18] Toby Gee, Tong Liu, and David Savitt. Crystalline extensions and the weight part of Serre's conjecture. *Algebra and Number Theory*, 6(7):1537–1559, 2012.
[19] Toby Gee, Tong Liu, and David Savitt. The weight part of serre's conjecture for GL(2). 2013. preprint.
[20] Toby Gee and David Savitt. Serre weights for mod p Hilbert modular forms: the totally ramified case. *J. Reine Angew. Math.*, 660:1–26, 2011.
[21] Toby Gee and David Savitt. Serre weights for quaternion algebras. *Compos. Math.*, 147(4):1059–1086, 2011.
[22] Mark Kisin. Potentially semi-stable deformation rings. *J. Amer. Math. Soc.*, 21(2):513–546, 2008.
[23] Jan Nekovář. On p-adic height pairings. In *Séminaire de Théorie des Nombres, Paris, 1990–91*, volume 108 of *Progr. Math.*, pages 127–202. Birkhäuser Boston, Boston, MA, 1993.
[24] Michel Raynaud. Schémas en groupes de type $(p, p, \ldots, p)$. *Bull. Soc. Math. France*, 102:241–280, 1974.
[25] David Savitt. Modularity of some potentially Barsotti-Tate Galois representations. *Compos. Math.*, 140(1):31–63, 2004.
[26] David Savitt. On a conjecture of Conrad, Diamond, and Taylor. *Duke Math. J.*, 128(1):141–197, 2005.
[27] David Savitt. Breuil modules for Raynaud schemes. *J. Number Theory*, 128:2939–2950, 2008.
[28] Michael M. Schein. Weights in Serre's conjecture for Hilbert modular forms: the ramified case. *Israel J. Math.*, 166:369–391, 2008.
[29] Michael M. Schein. Reduction modulo p of cuspidal representations and weights in Serre's conjecture. *Bull. Lond. Math. Soc.*, 41(1):147–154, 2009.
[30] Ryan Smith. Serre weights: the partially ramified case. University of Arizona Ph. D. thesis, 2012.

E-mail address: fred.diamond@kcl.ac.uk

Department of Mathematics, King's College London

E-mail address: savitt@math.arizona.edu

Department of Mathematics, University of Arizona