MODULI STACKS OF (φ,Γ)-MODULES: ERRATA

MATTHEW EMERTON AND TOBY GEE

The following errata and clarifications are for the published version of [EG22]. All references are to [EG22] unless otherwise specified.

We begin with some minor corrections.

- In the notation section (and throughout the book), the residue field of K is k.
- In Remark 2.1.13, $K_0(\zeta_{p^\infty})$ should be replaced by $(K_0)_{\mathrm{cyc}}$, and the claim that this condition is equivalent to K being abelian over $\mathbf{Q}_p$ should be deleted.

  The same applies in Definition 3.2.3 and Remark 3.2.4, and the discussion between them. In particular we should define $K^{\mathrm{basic}} := K \cap (K_0)_{\mathrm{cyc}}$.

  (We would like to thank Dat Pham for pointing this out. This issue was also explained to us by Léo Poyeton some years previously, but we unfortunately confused $K_0(\zeta_{p^\infty})$ and $(K_0)_{\mathrm{cyc}}$ in our writeup.)
- We should explicitly assume throughout that the coefficient field E contains K. (This is implicitly assumed at several points, but not always explicitly asserted.)
- The definition of the universal unramified character in §5.3 is incorrect if $k \neq \mathbf{F}_p$; instead, we need to set $\varphi(v) = a'v$ where $a' \in (\mathbf{F} \otimes_k A)^\times$ has norm a. See Dat Pham's [Pha22, §2]. (Note that this construction is one place that uses the assumption mentioned in the previous point.)
- In Hypothesis 7.3.1, we should take $b \geq peh/(p-1)$ rather than $b \geq eh/(p-1)$, because in the following paragraph, the valuation of T' should be given by $pv(p)/(p-1)$ rather than $v(p)/(p-1)$.

  In the second paragraph of the proof of Lemma 7.3.5, the image of T' should be $\zeta_{p^s} - 1$, rather than $\zeta_{p^{s+1}} - 1$. Thus the image of T coincides with the image of the trace of $\zeta_{p^s} - 1$, which is in $K_{\mathrm{cyc},s}$ and thus fixed by $G_{K_{\mathrm{cyc},s}}$.

  (This was again pointed out to us by Dat Pham.)
- Strictly speaking, the proof of Theorem 8.6.2 is incomplete (but the gap is easily filled); see the discussion between the statement and proof of [CEGS22, Thm. 7.6] for a complete proof.
- In the first paragraph of the proof of Theorem 6.3.2, the representation $\overline{\rho}$ (in the deformation ring $R_{\overline{\rho}}^{\mathrm{crys},\underline{\lambda}}$) is equal to $\overline{\rho}_d$. (Thanks to Matteo Tamiozzo for this.)
- Contrary to the claim in Section 5.3, it is not always possible to choose the characters $\psi_{\underline{n}}$ in such a way that if $\underline{n}$, $\underline{n}'$ have $(\psi_{\underline{n}}\psi_{\underline{n}'}^{-1})|_{I_K} = \overline{\epsilon}|_{I_K}$, then in fact $\psi_{\underline{n}}\psi_{\underline{n}'}^{-1} = \overline{\epsilon}$. The only place that this assumption was used was (implicitly) in the definitions of the characters $\omega_{\underline{k},i}$ before Definition 5.5.11; accordingly, we simply replace the $\omega_{\underline{k},i}$ by unramified twists in such a way

that if $k_{\overline{\sigma},i} - k_{\overline{\sigma},i+1} = p-1$ for all $\overline{\sigma}$, then $\omega_{\underline{k},i} = \omega_{\underline{k},i+1}$. (Again, we thank Dat Pham for pointing this out to us.)

We would like to thank Dat Pham for pointing out that it is not clear that the proof of Theorem 5.5.12 is complete. More precisely, without making additional arguments, it is not obvious that our constructions cover all of $\mathcal{X}_{d,\mathrm{red}}(\overline{\mathbf{F}}_p)$. Since the proof of Theorem 5.5.12 is quite involved, rather than attempt to describe the additional arguments needed, we instead give a slightly different and more streamlined complete proof.

In particular, we allow ourselves to refer forwards to Theorem 6.5.1, which strengthens some of the conclusions of Theorem 5.5.12, showing that the stacks $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ in the statement of Theorem 5.5.12 are empty. This requires some justification, because the proof of Theorem 6.5.1 relies on Theorem 5.5.12! However, the proof of Theorem 5.5.12 is by induction on d, and the proof of Theorem 6.5.1 for any particular d only uses Theorem 5.5.12 for that value of d, so we are free to assume in the proof of Theorem 5.5.12 that $\mathcal{X}^{\mathrm{small}}_{d',\mathrm{red},\overline{\mathbf{F}}_p}$ is empty for each $d' < d$. Accordingly, we have also removed $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ from the statement of the following theorem.

Theorem (Theorem 5.5.12)**.**

(1) The Ind-algebraic stack $\mathcal{X}_{d,\mathrm{red}}$ is an algebraic stack, of finite presentation over $\mathbf{F}$.

(2) $\mathcal{X}_{d,\mathrm{red}}$ is equidimensional of dimension $[K:\mathbf{Q}_p]d(d-1)/2$. We can write $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$ as a union of irreducible components $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$, where each $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ is generically maximally nonsplit of niveau one and weight $\underline{k}$, and the corresponding eigenvalue morphism is dominant (i.e. has dense image in $(\mathbf{G}_m)^d_{\underline{k}}$).

(3) If we fix an irreducible representation $\overline{\alpha}: G_K \to \mathrm{GL}_a(\overline{\mathbf{F}}_p)$ (for some $a \geq 1$), then the locus of $\overline{\rho}$ in $\mathcal{X}_{d,\mathrm{red}}(\overline{\mathbf{F}}_p)$ for which $\dim \mathrm{Hom}_{G_K}(\overline{\rho},\overline{\alpha}) \geq r$ (for any $r \geq 1$) is (either empty, or) of dimension at most

$$[K:\mathbf{Q}_p]d(d-1)/2 - \lceil r\big((a^2+1)r - a\big)/2\rceil.$$

(4) If we fix an irreducible representation $\overline{\alpha}: G_K \to \mathrm{GL}_a(\overline{\mathbf{F}}_p)$ (for some $a \geq 1$), then the locus of $\overline{\rho}$ in $\mathcal{X}_{d,\mathrm{red}}(\overline{\mathbf{F}}_p)$ for which $\dim \mathrm{Ext}^2_{G_K}(\overline{\alpha},\overline{\rho}) \geq r$ is of dimension at most

$$[K:\mathbf{Q}_p]d(d-1)/2 - r.$$

Proof. Recall that a closed immersion of reduced algebraic stacks that are locally of finite type over $\mathbf{F}_p$ which is surjective on finite-type points is necessarily an isomorphism. As recalled above, $\mathcal{X}_{d,\mathrm{red}}$ is an inductive limit of such stacks (indeed, by Lemma A.9, we have $\mathcal{X}_{d,\mathrm{red}} = \varinjlim \mathcal{X}^a_{d,h,s,\mathrm{red}}$, where the $\mathcal{X}^a_{d,h,s}$ are as in Section 3.4), and so if we produce closed algebraic substacks $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ and $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ of $\mathcal{X}_d$, the union of whose $\overline{\mathbf{F}}_p$-points exhausts those of $\mathcal{X}_{d,\mathrm{red}}$, then $\mathcal{X}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ will in fact be an algebraic stack which is the union of its closed substacks $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ and $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$. Thus (1) is an immediate consequence of (2) (where the "union" statement in (2) is now to be understood on the level of $\overline{\mathbf{F}}_p$-points). In fact, it suffices to construct the closed substacks $\mathcal{X}^{\underline{k}}_{d,\mathrm{red}}$, and to show that the remaining $\overline{\mathbf{F}}_p$-points of $\mathcal{X}_{d,\mathrm{red}}$ are contained in a closed substack $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$, which is a finite union of finite type algebraic substacks of dimension less than $[K : \mathbf{Q}_p]d(d-1)/2$; indeed, this obviously suffices to prove (1), and it also suffices to prove (2) by an application of Theorem 6.5.1 . (See the discussion before the theorem for an explanation of why this is not a circular argument.)

Claim (4) follows from (3) (with $\overline{\alpha}$ replaced by $\overline{\alpha} \otimes \overline{\epsilon}$) by Tate local duality and the easily verified inequality

$$\lceil r\big((a^2+1)r - a\big)/2\rceil \geq r.$$

Thus it is enough to prove (2) and (3), which we do simultaneously by induction on d.

As recalled in Remark 5.3.4 , there are up to twist by unramified characters only finitely many irreducible $\overline{\mathbf{F}}_p$-representations of G_K of any fixed dimension. Accordingly, we let $\{\overline{\alpha}_i\}$ be a finite set of irreducible continuous representations $\overline{\alpha}_i : G_K \to \mathrm{GL}_{d_i}(\overline{\mathbf{F}}_p)$, such that any irreducible continuous representation of G_K over $\overline{\mathbf{F}}_p$ of dimension at most d arises as an unramified twist of exactly one of the $\overline{\alpha}_i$. We let the 1-dimensional representations in this set be the characters $\psi_{\underline{n}}$ defined in Section 5.3 .

Each $\overline{\alpha}_i$ corresponds to a finite type point of $\mathcal{X}_{d_i,\mathrm{red}}$, whose associated residual gerbe is a substack of $\mathcal{X}_{d_i,\mathrm{red}}$ of dimension -1: the morphism $\operatorname{Spec}\overline{\mathbf{F}}_p \to \mathcal{X}_{d_i,\mathrm{red}}$ corresponding to $\overline{\alpha}_i$ factors through a monomorphism $[\operatorname{Spec}\overline{\mathbf{F}}_p/\mathbf{G}_m] \to \mathcal{X}_{d_i,\mathrm{red}}$. It follows from Lemma 5.3.2 that for each $\overline{\alpha}_i$ there is an irreducible closed zero-dimensional algebraic substack of $\mathcal{X}_{d_i,\mathrm{red}}$ of finite presentation over $\overline{\mathbf{F}}_p$ which contains a dense open substack whose $\overline{\mathbf{F}}_p$-points are the unramified twists of $\overline{\alpha}_i$.

In particular, if $d = 1$, then we let $\mathcal{X}^{\underline{k}}_{1,\mathrm{red}}$ be the zero-dimensional stack constructed in the previous paragraph; this satisfies the required properties by definition, so (2) holds when $d = 1$. For (3), note that if $r > 0$ then we must have $a = 1$, and then the locus where $\operatorname{Hom}_{G_K}(\overline{\rho}, \overline{\alpha})$ is non-zero (equivalently, 1-dimensional) is exactly the closed substack of dimension -1 corresponding to $\overline{\alpha}$, so the required bound holds.

We now begin the inductive proof of (2) and (3) for $d > 1$. In fact, it will be helpful to simultaneously prove additional statements (2') and (4'), which we begin by stating. (It is trivial to verify that the discussion above also proves (2') and (4') when $d = 1$.)

(2') is as follows: for each $\underline{k}$, there is a closed irreducible algebraic substack $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ of $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$ of finite presentation over $\overline{\mathbf{F}}_p$ and dimension $[K : \mathbf{Q}_p]d(d-1)/2 - 1$, which is generically maximally nonsplit of niveau 1 and weight $\underline{k}$, and furthermore has the properties that the corresponding character ν_1 is trivial, and that $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ is obtained from $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ by unramified twisting. (The assumption that there is a dense open substack of $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red}}$ whose $\overline{\mathbf{F}}_p$-points correspond to representations which are of niveau 1 and are maximally nonsplit implies in particular that the corresponding family is twistable).

We now turn to (4'). For each character $\overline{\alpha} : G_K \to \overline{\mathbf{F}}_p^\times$, and each $\underline{k}$, let $\mathcal{X}_{\overline{\alpha}}$ be the locally closed substack consisting of those $\overline{\rho}$ in $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}(\overline{\mathbf{F}}_p)$ for which $\dim \operatorname{Ext}^2_{G_K}(\overline{\alpha}, \overline{\rho}) \neq 0$. Then by (4), the dimension of $\mathcal{X}_{\overline{\alpha}}$ is at most $[K : \mathbf{Q}_p]d(d-1)/2 - 1$. Then (4') is as follows: if $\mathcal{X}_{\overline{\alpha}}$ has dimension $[K : \mathbf{Q}_p]d(d-1)/2 - 1$, then after replacing $\overline{\alpha}$ by an unramified twist, $\mathcal{X}_{\overline{\alpha}}$ contains a dense open substack

$\mathcal{U}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ of $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$, and the complement $\mathcal{X}_{\overline{\alpha}} \setminus \mathcal{U}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ has dimension at most $[K:\mathbf{Q}_p]d(d-1)/2-2$.

We now prove the inductive step, so we assume that (2), (2'), (3) and (4') hold in dimension less than d. Let $\underline{k}_{d-1}$ be the Serre weight in dimension $(d-1)$ obtained by deleting the first entry in $\underline{k}$. By hypothesis there are corresponding (dense) open substacks $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red}}$ and $\mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ of $\mathcal{X}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red}}$ and $\mathcal{X}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ respectively, which are maximally nonsplit of weight $\underline{k}$. By our inductive construction, we can and do assume that $\mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ is obtained from $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red}}$ by unramified twisting. (These substacks are of course not uniquely determined, and we will allow ourselves to further shrink them in the below.)

Set $\overline{\alpha} := \overline{\epsilon}^{1-d}\omega_{\underline{k},1}$. If $k_{\overline{\sigma},1}-k_{\overline{\sigma},2}=p-1$ for all $\overline{\sigma}$ then we say that we are in the *très ramifiée case*, and we let $\mathcal{U}$ denote $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$; otherwise, we let $\mathcal{U}$ denote $\mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$. In the latter case, it follows from our inductive hypothesis (more precisely, from the assumption that (2) and (3) hold in dimension $(d-1)$), together with Tate local duality, that after possibly replacing $\mathcal{U}$ with an open substack we can and do assume that for each $\overline{\mathbf{F}}_p$-point u of $\mathcal{U}$, we have $\mathrm{Ext}^2_{G_K}(\overline{\alpha},\bar{r}_u)=0$, where $\bar{r}_u$ is the $(d-1)$-dimensional representation corresponding to u. If we are in the très ramifiée case then it follows similarly that after replacing $\mathcal{U}$ with an open substack we can and do assume that for each $\overline{\mathbf{F}}_p$-point u of $\mathcal{U}$, $\mathrm{Ext}^2_{G_K}(\overline{\alpha},\bar{r}_u)$ is 1-dimensional. (Note that the condition that $\mathrm{Ext}^2_{G_K}(\overline{\alpha},\bar{r}_u)$ is 1-dimensional is *a priori* only locally closed, but by hypothesis we can pass to an open substack where the dimension is at least 1, and the condition that the dimension is exactly 1 is then open.)

We let T be an irreducible scheme which smoothly covers $\mathcal{U}$, and we let $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ be the irreducible closed substack of $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$ constructed as the scheme-theoretic image of V in the notation of Proposition 5.4.4 . Part (2) of that proposition, together with the inductive hypothesis, implies that $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ has the claimed dimension. (Note in particular that if $K=\mathbf{Q}_p$, then condition (2d) of Proposition 5.4.4 holds. Indeed, if we are in the très ramifiée case then condition (2d)(iii) holds, and otherwise the inductive hypothesis that the image of the eigenvalue morphism is dense in $(\mathbf{G}_m)^{d-1}_{\underline{k}_{d-1}}$ implies that condition (2d)(i) holds after possibly further shrinking $\mathcal{U}$.) We then let $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ be the substack obtained from $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red}}$ by twisting by unramified characters, which has the claimed dimension by Lemma 5.3.2 .

By construction, we see that the stacks $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ and $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ satisfy the properties required of them by (2) and (2') (other, of course, than the claim that they exhaust the irreducible components of $\mathcal{X}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$); except that in the très ramifiée case, we need to check that $\mathcal{X}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ (and consequently $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$) is generically maximally nonsplit of niveau 1 and weight $\underline{k}$. More precisely, in this case we need to check the condition that the final extension is not just nonsplit, but is (generically) très ramifiée. This follows from Proposition 5.4.4 (3). Indeed by our inductive assumption that $\mathcal{X}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ is generically maximally nonsplit of niveau 1 and weight $\underline{k}_{d-1}$, and the image of the eigenvalue morphism is dense, we see that we are in case (3)(b)(i) unless $\overline{\epsilon}$ is trivial and furthermore we have $k_{\overline{\sigma}_2} - k_{\overline{\sigma}_3} = p-1$ for all $\overline{\sigma}$, in which case we satisfy (3)(b)(ii).

To complete the proof of (2), we need to construct $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$. Let $\mathcal{X}^{\mathrm{small}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ denote the union of the closed substacks

$$\mathcal{X}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p} \setminus \mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}.$$

By Proposition 5.4.4 , Tate local duality, upper semi-continuity of the fibre dimension, and the inductive hypothesis, we see that for each $1 \le i \le d$, each $\overline{\alpha}_i$ (of dimension a_i, say), and each $s \ge 0$ there is a finitely presented closed algebraic substack $\mathcal{X}_{s,\overline{\alpha}_i,\overline{\mathbf{F}}_p}$ of $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$, whose $\overline{\mathbf{F}}_p$-points contain all the representations of the form $0 \to \overline{\rho}_{d-a_i} \to \overline{\rho} \to \overline{\alpha}_i \to 0$ for which $\dim_{\overline{\mathbf{F}}_p} \mathrm{Ext}^2_{G_K}(\overline{\alpha}_i, \overline{\rho}_{d-a_i}) = s$, and whose dimension is at most

$$[K : \mathbf{Q}_p](d-a_i)(d-a_i-1)/2 - \lceil s(({a_i}^2+1)s - a_i)/2 \rceil + [K : \mathbf{Q}_p]a_i(d-a_i) + s - 1.$$

Furthermore, if $a = 1$, then the locus where $\overline{\rho}_{d-1}$ is an $\overline{\mathbf{F}}_p$-point of $\mathcal{X}^{\mathrm{small}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ is of dimension strictly less than this. (Here we use (4'), which shows in particular that the locus in $\mathcal{X}^{\mathrm{small}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ of points where $\dim_{\overline{\mathbf{F}}_p} \mathrm{Ext}^2_{G_K}(\overline{\alpha}_i, \overline{\rho}_{d-1}) = 1$ has dimension at most $[K : \mathbf{Q}_p](d-1)(d-2)/2 - 2$.)

These stacks are only nonzero for finitely many values of s. For fixed a_i, we see that as a function of s, this quantity is maximised by $s = 0$, as well as by $s = 1$ when $a_i = 1$. (To see this, we have to maximise the quantity $s - \lceil s(({a_i}^2+1)s - a_i)/2 \rceil$. Suppose firstly that $a_i > 1$. Then if $s = 0$ we have 0, while if $s > 0$ we have $s - \lceil s(({a_i}^2+1)s - a_i)/2 \rceil \le s - s(({a_i}^2+1)s - a_i)/2 \le s - s({a_i}^2+1-a_i)/2 \le s - 3s/2 < 0$. Meanwhile if $a_i = 1$, then for $s = 0$ we have 0, for $s = 1$ we have $1 - \lceil 1/2 \rceil = 1$, while for $s > 1$ we have $s - \lceil s(2s-1)/2 \rceil \le s - s(2s-1)/2 \le s - 3s/2 < 0$.) It follows that as a function of a_i the bound is maximised at $a_i = 1$ and $s = 0$ or 1, when it is equal to $[K : \mathbf{Q}_p]d(d-1)/2 - 1$, and it is otherwise strictly smaller.

By Lemma 5.3.2 , it follows that the locus in $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$ of representations of the form $0 \to \overline{\rho}_{d-a_i} \to \overline{\rho} \to \overline{\alpha}' \to 0$, with $\overline{\alpha}'$ an unramified twist of $\overline{\alpha}_i$ for which $\dim_{\overline{\mathbf{F}}_p} \mathrm{Ext}^2_{G_K}(\overline{\alpha}', \overline{\rho}_{d-a_i}) = s$, is of dimension at most $[K : \mathbf{Q}_p]d(d-1)/2$, with equality holding only if $a_i = 1$ and $s = 0$ or 1.

Putting this together, we claim that (2) holds in dimension d if we take $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ to be the union of the twists by unramified characters of the substacks $\mathcal{X}_{s,\overline{\alpha}_i}$ for which $\dim \overline{\alpha}_i > 1$ or $s > 1$, together with the union of the twists by unramified characters of the substacks of the $\mathcal{X}_{s,\overline{\alpha}_i,\overline{\mathbf{F}}_p}$ for which $\dim \overline{\alpha}_i = 1$, $s = 0$ or 1, and $\overline{\rho}_{d-1}$ is an $\overline{\mathbf{F}}_p$-point of $\mathcal{X}^{\mathrm{small}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$. Indeed, by construction, $\mathcal{X}^{\mathrm{small}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ is certainly a finite union of closed substacks of dimension less than $[K : \mathbf{Q}_p]d(d-1)/2$, and in order to see that it exhausts the remaining points of $\mathcal{X}_{d,\mathrm{red},\overline{\mathbf{F}}_p}(\overline{\mathbf{F}}_p)$, we only need to verify that if $\underline{k}$ is in the très ramifiée case, and $\overline{\rho}$ corresponds to a point of $\mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ with $\dim_{\overline{\mathbf{F}}_p} \mathrm{Ext}^2_{G_K}(\overline{\alpha}, \overline{\rho}) = 1$, then every extension of $\overline{\alpha}$ by $\overline{\rho}$ is contained in $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$. By construction, it is enough to check that such a $\overline{\rho}$ is automatically a point of $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$. But $\mathcal{U}^{\underline{k}_{d-1}}_{d-1,\mathrm{red},\overline{\mathbf{F}}_p}$ is obtained from $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red}}$ by unramified twisting, and since $\overline{\rho}$ is maximally nonsplit, the hypothesis that $\mathrm{Ext}^2_{G_K}(\overline{\alpha}, \overline{\rho}) \neq 0$ forces $\overline{\rho}$ to be a point of $\mathcal{U}^{\underline{k}_{d-1},\mathrm{fixed}}_{d-1,\mathrm{red}}$. This completes the proof of (2).

We now prove (3) in dimension d. In the case $r=0$, there is nothing to prove. If $\dim \mathrm{Hom}_{G_K}(\overline{\rho},\overline{\alpha}) \geq r \geq 1$, then we may place $\overline{\rho}$ in a short exact sequence

$$0 \to \overline{\theta} \to \overline{\rho} \to \overline{\alpha}^{\oplus r} \to 0,$$

where $\overline{\theta}$ is of dimension $d-ra<d$. We may apply part (2) so as to find that $\mathcal{X}_{d-ar,\mathrm{red},\overline{\mathbf{F}}_p}$ has dimension at most $[K:\mathbf{Q}_p](d-ar)(d-ar-1)/2$. Let $\mathcal{U}_s$ be the locally closed substack of $\mathcal{X}_{d-ar,\mathrm{red},\overline{\mathbf{F}}_p}$ over which $\dim H^2(G_K,\overline{\theta}\otimes\overline{\alpha}^\vee)=s$; by the inductive hypothesis, this locus has dimension at most $[K:\mathbf{Q}_p](d-ar)(d-ar-1)/2-s\big((a^2+1)s-a\big)/2$, and over this locus we may construct a universal family of extensions

$$0 \to \overline{\theta} \to \overline{\rho}_{\mathcal{U}_s} \to \overline{\alpha}^{\oplus r} \to 0.$$

The locus of $\overline{\rho}$ we are interested in is contained in the scheme-theoretic image of this family in $(\mathcal{X}_{d,\mathrm{red}})_{\overline{\mathbf{F}}_p}$, and Proposition 5.4.4 shows that this scheme-theoretic image has dimension bounded above by

$$[K:\mathbf{Q}_p](d-ar)(d-ar-1)/2-s((a^2+1)s-a)/2+r([K:\mathbf{Q}_p]a(d-ar)+s)-r^2$$

$$\begin{aligned} &=[K:\mathbf{Q}_p]d(d-1)/2-(r(a^2+1)r-a)/2-(r-s)^2/2-(ar(ar-1))/2 \\ &\leq[K:\mathbf{Q}_p]d(d-1)/2-(r(a^2+1)r-a)/2. \end{aligned}$$

Since this conclusion holds for each of the finitely many values of s (and since the dimension is an integer, allowing us to take the floor of this upper bound), we have proved (3).

Finally, we prove (4') in dimension d. Assume that $\mathcal{X}_{\overline{\alpha}}$ has dimension $[K:\mathbf{Q}_p]d(d-1)/2-1$. By Tate local duality, the condition $\mathrm{Ext}^2_{G_K}(\overline{\alpha},\overline{\rho})\neq 0$ is equivalent to $\mathrm{Hom}_{G_K}(\overline{\rho},\overline{\alpha}\otimes\overline{\epsilon})\neq 0$; so if $\overline{\rho}$ is a point of $\mathcal{X}_{\overline{\alpha}}$, then $\overline{\rho}$ admits $\overline{\alpha}\otimes\overline{\epsilon}$ as a Jordan–Hölder factor. It follows immediately that $\mathcal{X}_{\overline{\alpha}}$ is essentially twistable, so that the substack of $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ given by unramified twists of $\mathcal{X}_{\overline{\alpha}}$ has dimension $[K:\mathbf{Q}_p]d(d-1)/2$, so is dense in $\mathcal{X}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$. In particular, it contains a dense open substack of $\mathcal{U}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$.

Recalling that $\mathcal{U}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ is maximally nonsplit, it follows that after replacing $\overline{\alpha}$ by an unramified twist, and possibly shrinking $\mathcal{U}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$, the stack $\mathcal{X}_{\overline{\alpha}}$ contains $\mathcal{U}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$. Furthermore if $\mathcal{X}'_{\overline{\alpha}} := \mathcal{X}_{\overline{\alpha}}\setminus\mathcal{U}^{\underline{k}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$ had dimension $[K:\mathbf{Q}_p]d(d-1)/2-1$ then we could repeat this argument with $\mathcal{X}_{\overline{\alpha}}$ replaced by $\mathcal{X}'_{\overline{\alpha}}$, and conclude that $\mathcal{X}'_{\overline{\alpha}}$ contains a dense open substack of an unramified twist of $\mathcal{U}^{\underline{k},\mathrm{fixed}}_{d,\mathrm{red},\overline{\mathbf{F}}_p}$, which is a contradiction. This completes the proof of the theorem. □

References

[CEGS22] Ana Caraiani, Matthew Emerton, Toby Gee, and David Savitt, *The geometric Breuil–Mézard conjecture for two-dimensional potentially Barsotti–Tate Galois representations*, 2022.

[EG22] Matthew Emerton and Toby Gee, *Moduli stacks of étale (φ,Γ)-modules and the existence of crystalline lifts*, Annals of Math. Studies (to appear) (2022).

[Pha22] Dat Pham, *The Emerton–Gee stack of rank one (φ,Γ)-modules*, 2022.

Email address: emerton@math.uchicago.edu

Department of Mathematics, University of Chicago, 5734 S. University Ave., Chicago, IL 60637

Email address: toby.gee@imperial.ac.uk

Department of Mathematics, Imperial College London, London SW7 2AZ, UK