J. Inst. Math. Jussieu (2026), **25**(3), 1441–1486
doi:10.1017/S1474748025101503

p-ADIC SHEAVES ON CLASSIFYING STACKS, AND THE p-ADIC JACQUET-LANGLANDS CORRESPONDENCE

DAVID HANSEN[1] AND LUCAS MANN[2]

[1] *National University of Singapore, Singapore*
(dhansen@nus.edu.sg)
[2] *University of Muenster, Germany*
(mann.lucas@uni-muenster.de)

(Received 4 January 2024; revised 3 July 2025; accepted 2 December 2025;
first published online 30 March 2026)

Abstract We establish several new properties of the p-adic Jacquet-Langlands functor defined by Scholze in terms of the cohomology of the Lubin-Tate tower. In particular, we reprove Scholze's basic finiteness theorems, prove a duality theorem, and show a kind of partial Künneth formula. Using these results, we deduce bounds on Gelfand-Kirillov dimension, together with some new vanishing and nonvanishing results.

Our key new tool is the six functor formalism with solid almost $\mathcal{O}^+/p$-coefficients developed recently by the second author [Man22]. One major point of this paper is to extend the domain of validity of the !-functor formalism developed in [Man22] to allow certain 'stacky' maps. In the language of this extended formalism, we show that if G is a p-adic Lie group, the structure map of the classifying small v-stack $B\underline{G}$ is p-cohomologically smooth.

1 Introduction **1442**
1.1. Prologue 1442
1.2. Context and duality 1443
1.3. The stacky 6-functor formalism 1446
1.4. Applications to representation theory 1447

Keywords: p-adic Jacquet-Langlands correspondence; Lubin-Tate tower; Gelfand-Kirillov dimension

Mathematics Subject Classification: Primary 14G45
Secondary 14G35

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (https://creativecommons.org/licenses/by-nc-nd/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use and/or adaptation of the article.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

1.5. A word on notation 1450
2 Geometric preparations **1450**
2.1. Recollections on the solid six functors 1450
2.2. Six functors for stacky morphisms 1452
2.3. Cohomologically smooth maps 1455
3 Representation-theoretic preparations **1457**
3.1. Smooth representations 1458
3.2. Sheaves on classifying stacks 1461
3.3. Admissible representations 1469
4 Duality and applications **1473**
4.1. The main result 1473
4.2. Applications to Gelfand-Kirillov dimension 1476
4.3. The case of GL_2 1478
5 A partial Künneth formula **1480**
References **1485**

1. Introduction

1.1. Prologue

This paper has two goals. Our immediate goal is to prove some new properties of the p-adic Jacquet-Langlands functor defined by Scholze. Unlike the classical Jacquet-Langlands correspondence, which is defined by a simple character identity, Scholze's functor is defined inexplicitly, and has turned out to be rather difficult to study. Scholze himself proved a basic finiteness theorem, and established a local-global compatibility result in some cases. We reprove Scholze's finiteness theorem, and prove a duality theorem and a kind of 'partial' Künneth formula. As consequences, we deduce bounds on Gelfand-Kirillov dimension and prove some new vanishing and nonvanishing results for Scholze's functor.

To explain our second goal, let us briefly recall some recent history. The study of the classical local Langlands correspondence has recently been given new life through the geometrization program proposed by Fargues [Far16]. This program has already seen spectacular success, most notably in the work of Fargues-Scholze [FS21]. A key technical ingredient in this program is a sufficiently flexible six-functor formalism for ℓ-adic étale cohomology of diamonds and v-stacks. This formalism was developed by Scholze [Sch18a], and then extended further in [GHW22].

Our second goal is to lay some technical foundations for the geometrization of the p-adic Langlands correspondence. In particular, we extend the p-adic six functor formalism developed in [Man22] to allow certain stacky maps, in parallel with the extension provided in [GHW22], and show that classifying stacks of p-adic Lie groups fit perfectly into this formalism. This stacky six-functor machinery allows for an elegant formulation of the proof of our main results. With it in hand, one can also begin a meaningful study of p-adic sheaves on Bun_G.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Remark 1.1. We prove p-cohomological smoothness of Bun_G in [ABM24], building on the ideas in this paper, and there is some further discussion there on how the mod p Langlands should behave. Moreover, we have recently made advances in the ℓ-adic program and hope that some of the ideas there can be carried over to the p-adic world using the formalism built here.

1.2. Context and duality

Let us begin by recalling the classical Jacquet-Langlands correspondence. Fix a finite extension $F/\mathbf{Q}_p$ and an integer $n>1$. Set $G=\mathrm{GL}_n(F)$, and let D/F be the central division algebra of invariant $1/n$.

Theorem 1.2 (Jacquet-Langlands, Deligne-Kazhdan-Vignéras, Rogawski). *There is a natural bijection*

$$\mathrm{JL}:\mathrm{Irr}_{L^2}(G)\xrightarrow{\sim}\mathrm{Irr}(D^\times)$$

from the irreducible essentially square-integrable complex representations of G to the irreducible representations of $D^\times$, characterized by the requirement that $\chi_{\mathrm{JL}(\pi)}(g')=(-1)^{n-1}\chi_\pi(g)$ for all pairs of matching elliptic elements $g\in G$ and $g'\in D^\times$.

This bijection is compatible with twisting and passing to contragredient representations.

A very similar statement holds with $\overline{\mathbf{F}_\ell}$-coefficients instead of complex coefficients, cf. [Dat12].

In recent years, a p-adic version of the Langlands program has emerged, many aspects of which are still conjectural. In the context of this program, a candidate for a p-adic Jacquet-Langlands functor has been proposed by Scholze [Sch18b]. Let us begin by recalling the definition of Scholze's functor. Set $C=\widehat{\overline{F}}$, and let $\mathcal{M}_\infty$ be the infinite level Lubin-Tate space over $\mathrm{Spa}C$. Recall that this is a perfectoid space with commuting continuous actions of G and $D^\times$, and with a natural surjective period map $\pi_{\mathrm{GH}}:\mathcal{M}_\infty\to\mathbf{P}^{n-1}_C$. The period map π_{GH} is a G-torsor, and is $D^\times$-equivariant for a suitable action of $D^\times$ on projective space.

Using this geometry, Scholze defines an exact functor, denoted $\pi\mapsto\mathcal{F}_\pi$, from the category $\mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G)$ of smooth G-representations towards the category of étale $\mathbf{F}_p$-sheaves on $\mathbf{P}^{n-1}_C$. This functor is defined by a simple and explicit recipe: given any étale map $U\to\mathbf{P}^{n-1}_C$, one defines $\mathcal{F}_\pi(U)=\mathcal{C}_G(|U\times_{\mathbf{P}^{n-1}_C}\mathcal{M}_\infty|,\pi)$. Since π_{GH} is $D^\times$-equivariant, the sheaf $\mathcal{F}_\pi$ is also $D^\times$-equivariant, and the étale cohomology groups

$$\mathcal{J}^i(\pi):=H^i(\mathbf{P}^{n-1}_C,\mathcal{F}_\pi)$$

are naturally smooth $D^\times$-representations which also carry a natural action of the Weil group W_F. By general vanishing theorems (e.g. using [Hub96, Corollary 2.8.3]), it is easy to see that $\mathcal{J}^i(\pi)=0$ for all $i>2n-2$, and it is also easy to compute $\mathcal{J}^0(\pi)$ explicitly. However, these are the only obvious properties of $\mathcal{J}^i$. One of the main results of [Sch18b] is a highly nonobvious finiteness theorem: if π is an *admissible* G-representation, then also each $\mathcal{J}^i(\pi)$ is an admissible $D^\times$-representation. As a key step in the proof of this finiteness result, Scholze also proved a form of the primitive comparison theorem: if π is

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

admissible, the natural map

$$H^i(\mathbf{P}^{n-1}_C,\mathcal{F}_\pi)\otimes\mathcal{O}_C/p\to H^i(\mathbf{P}^{n-1}_C,\mathcal{F}_\pi\otimes\mathcal{O}^+/p)$$

is an almost isomorphism.

Consideration of simple examples shows that in general, $\mathcal{J}^i(\pi)$ can be nonzero for multiple values of i. As such, we prefer to package the individual cohomologies $\mathcal{J}^i$ into a single operation on the derived category. More precisely, by the exactness of the association $\pi\mapsto\mathcal{F}_\pi$, it naturally upgrades to a functor $\mathcal{D}(G):=\mathcal{D}(\mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G))\to\mathcal{D}(\mathbf{P}^{n-1}_{C,\text{ét}},\mathbf{F}_p)$, and we write $\mathcal{J}$ for the composite functor

$$\mathcal{D}(G)\overset{V\mapsto\mathcal{F}_V}{\longrightarrow}\mathcal{D}(\mathbf{P}^{n-1}_{C,\text{ét}},\mathbf{F}_p)\overset{R\Gamma}{\to}\mathcal{D}(D^\times).$$

By the finiteness theorems mentioned above, the functor $\mathcal{J}$ preserves bounded complexes with admissible cohomology.

Now, one would guess that $\mathcal{J}$ intertwines some natural duality operation on G-representations and $D^\times$-representations. However, in stark contrast with the case of complex or ℓ-adic representations, the naive functor of passing to smooth duals on mod-p representations of p-adic Lie groups is essentially useless. To remedy this, Kohlhaase has introduced a family of 'derived' smooth duality functors $\mathcal{S}^i_H$ on the smooth mod-p representations of any p-adic Lie group H, which can be packaged into a single endofunctor $\mathcal{S}_H:\mathcal{D}(H)\to\mathcal{D}(H)$. This functor has excellent properties, which we review in §3.1 below.

Our first main result can now be stated as follows.

Theorem 1.3. *There is a natural equivalence $\mathcal{S}_{D^\times}\circ\mathcal{J}\cong(\mathcal{J}\circ\mathcal{S}_G)[2n-2](n-1)$ of functors $\mathcal{D}_{\mathrm{adm}}(G)\to\mathcal{D}_{\mathrm{adm}}(D^\times)$.*

In other words, Scholze's functor $\mathcal{J}$ intertwines Kohlhaase's derived duality on G-representations and $D^\times$-representations, up to an explicit shift and twist. This result was conjectured by the first author, after observing that for certain V's arising from the completed cohomology of Shimura varieties, both $\mathcal{S}_{D^\times}\circ\mathcal{J}(V)$ and $(\mathcal{J}\circ\mathcal{S}_G)(V)[2n-2](n-1)$ can be computed by hand, and turn out to be abstractly isomorphic.

In the remainder of this section, we give a detailed sketch of the proof of Theorem 1.3. The first key ingredient is a more conceptual interpretation of Kohlhaase's derived duality functor.

Proposition 1.4. *For any p-adic Lie group G, the functor $\mathcal{S}_G(-):\mathcal{D}(G)\to\mathcal{D}(G)$ is naturally isomorphic to the functor $\underline{\mathrm{Hom}}_G(-,\mathbf{1})$. Here $\underline{\mathrm{Hom}}_G$ denotes the natural internal Hom in $\mathcal{D}(G)$.*

Be aware that in this paper, $\underline{\mathrm{Hom}}$ will always denote the *derived* internal hom (which is often denoted by $R\underline{\mathrm{Hom}}$ in the literature) – see Section 1.5 for more remarks on that choice. In the case of mod-ℓ coefficients with $\ell\neq p$, $\underline{\mathrm{Hom}}_G(-,\mathbf{1})$ is actually t-exact and therefore restricts to a functor on the hearts, where it coincides with the usual functor of smooth duality (Proposition 3.5). In this light, Proposition 1.4 puts Kohlhaase's duality on the same footing as classical smooth duality. We also note that Proposition 1.4 was independently discovered by Schneider-Sorensen [SS22].

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Next, we invoke a slightly different formula for the functor $\mathcal{J}$. For this, we briefly return to the geometry of the Lubin-Tate tower. Along with the surjective period map $\pi_{\mathrm{GH}} : \mathcal{M}_\infty \to \mathbf{P}_C^{n-1}$ mentioned earlier, there is a second period map $\pi_{\mathrm{HT}} : \mathcal{M}_\infty \to \Omega_C^{n-1}$, presenting $\mathcal{M}_\infty$ as a $D^\times$-torsor over Drinfeld space. Both period maps are $G \times D^\times$-equivariant, with G (resp. $D^\times$) acting trivially on $\mathbf{P}_C^{n-1}$ (resp. on Ω_C^{n-1}). In particular, we can pass to the stack quotient X of $\mathcal{M}_\infty$ by the whole $G \times D^\times$-action, obtaining a small v-stack X which admits two natural presentations

$$X \cong [\mathbf{P}_C^{n-1}/\underline{D^\times}] \cong [\Omega_C^{n-1}/\underline{G}].$$

We think of this stack as sitting in a diagram

$$\begin{array}{ccc} & X & \\ {\scriptstyle f}\swarrow & & \searrow{\scriptstyle g} \\ B\underline{G} = [\mathrm{Spa}\, C/\underline{G}] & & B\underline{D^\times} = [\mathrm{Spa}\, C/\underline{D^\times}] \end{array}$$

where f is smooth, and g is smooth and proper, with both maps being of pure relative dimension $n-1$. Let us call this **the fundamental diagram**. Now, for any small v-stack X, the formalism in [Sch18a] defines an associated category $\mathcal{D}_{\text{ét}}(X, \mathbf{F}_p)$ of étale $\mathbf{F}_p$-sheaves. We will see below (Proposition 3.9) that there is a natural identification $\mathcal{D}(H) = \mathcal{D}_{\text{ét}}(B\underline{H}, \mathbf{F}_p)$ of symmetric monoidal categories for any p-adic Lie group H. Under this identification, it is not hard to show that $\mathcal{J}$ is simply given by the operation $g_{\text{ét}*} f_{\text{ét}}^*$ (where by $g_{\text{ét}*}$ we denote the *derived* pushforward, often denoted $Rg_{\text{ét}*}$ – cf. Section 1.5). At this point, it is heuristically clear how the argument should go: f is smooth and g is proper smooth, so the operation $g_{\text{ét}*} f_{\text{ét}}^*$ should commute with internal hom towards the monoidal unit, up to an explicit shift and twist. However, it is entirely unclear how to make this heuristic rigorous, because the !-functors associated with the maps f and g are conspicuously absent from the $\mathcal{D}_{\text{ét}}$-formalism (Of course, this absence is a special feature of the $\ell = p$ setting of the present paper; one of the main points of [Sch18a] is to construct the !-functors in the $\mathcal{D}_{\text{ét}}$-formalism when the coefficient ring is killed by an integer prime to p.).

The key new tool which lets us circumvent this issue is the formalism of *solid almost $\mathcal{O}^+/p$-sheaves* recently developed by the second author [Man22]. This is a theory which attaches to any small v-stack X over $\mathrm{Spa}\, C$ a symmetric monoidal stable ∞-category $\mathcal{D}_\square^a(\mathcal{O}_X^+/p)$, whose objects are (to first approximation) complexes of solid quasicoherent almost $\mathcal{O}_X^+/p$-sheaves on the étale site of X. There is also a variant $\mathcal{D}_\square^a(\mathcal{O}_X^+/p)^\varphi$ where the objects are equipped with a suitable φ-module structure, which comes with an evident forgetful functor $\mathcal{D}_\square^a(\mathcal{O}_X^+/p)^\varphi \to \mathcal{D}_\square^a(\mathcal{O}_X^+/p)$. Crucially, the sheaf theories $X \mapsto \mathcal{D}_\square^a(\mathcal{O}_X^+/p)$ and $\mathcal{D}_\square^a(\mathcal{O}_X^+/p)^\varphi$ underlie *full* six functor formalisms, and all six functors are compatible with the operation of forgetting the φ-module structure. In particular, for any map of small v-stacks $f : X \to Y$ over $\mathrm{Spa}\, C$, we have functors f_* and f^*, and if f is a so-called bdcs ('p-**b**ounde**d**, locally **c**ompactifiable, and representable in locally **s**patial diamonds'.) map we also have functors $f_!$ and $f^!$ with all expected properties. Moreover, there is a notion of p-cohomologically smooth maps, for which $f^!$ is an invertible twist of f^*. Finally,

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

if H is a p-adic Lie group, there is a natural identification $\mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{H}}/p) = \mathcal{D}^{\mathrm{sm},a}_\square(H,\mathcal{O}_C/p)$, where the right-hand side denotes the derived ∞-category of smooth H-representations on complexes of solid almost $\mathcal{O}_C/p$-modules (and similarly with φ-structures).

Returning to the fundamental diagram, we can define a solid variant $\mathcal{J}_\square$ of the functor $\mathcal{J}$ by the formula

$$\mathcal{J}_\square = g_* f^* : \mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{G}}/p)^\varphi \to \mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{D^\times}}/p)^\varphi.$$

Using the formalism developed in [Man22], it is entirely straightforward to prove a solid variant of our goal, i.e. an isomorphism $\mathcal{S}_{D^\times,\square} \circ \mathcal{J}_\square \cong (\mathcal{J}_\square \circ \mathcal{S}_{G,\square})[2n-2](n-1)$, where (for any p-adic Lie group H) $\mathcal{S}_{H,\square}$ denotes the internal hom towards the monoidal unit in $\mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{H}}/p)^\varphi$. However, we have potentially lost sight of our goal of understanding duality for the functor $\mathcal{J}$.

At this point we use another key feature of the formalism developed in [Man22], namely the *p-torsion Riemann-Hilbert correspondence*, which gives a fully faithful embedding $-\otimes\mathcal{O}^{+a}_X/p : \mathcal{D}_{\text{ét}}(X,\mathbf{F}_p)^\dagger \hookrightarrow \mathcal{D}^a_\square(\mathcal{O}^+_X/p)^\varphi$ for any small v-stack X. Here the superscripted $\dagger$ denotes the restriction to the full subcategory of overconvergent sheaves. This embedding comes with a right adjoint $M \to M^\varphi$ which is (morally) the functor of φ-invariants. Using the basic properties of these functors, it is formal to check that $\mathcal{J}$ can be recovered from $\mathcal{J}_\square$, in the sense that $\mathcal{J}(A) \cong \mathcal{J}_\square(A\otimes\mathcal{O}^{+a}_{B\underline{G}}/p)^\varphi$ for any $A \in \mathcal{D}(G) = \mathcal{D}_{\text{ét}}(B\underline{G},\mathbf{F}_p)^\dagger$. We now apply a key new result, Proposition 3.26, which gives a practical criterion for an object $M \in \mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{H}}/p)^\varphi$ to arise from a bounded admissible complex in $\mathcal{D}(H)$ via the Riemann-Hilbert embedding, and shows that in this case the solid dual $\mathcal{S}_{H,\square}$ is the Riemann-Hilbert embedding of the smooth dual $\mathcal{S}_H$. It turns out that in order to apply this criterion to $\mathcal{J}_\square$, we only need to know two things:

(1) $\mathcal{J}_\square$ preserves discrete objects, and

(2) for any admissible $A \in \mathcal{D}(G)$, $\mathcal{J}_\square(A\otimes\mathcal{O}^{+a}_{B\underline{G}}/p)$ is isomorphic to its solid double dual.

We now conclude by noting that (1) is formal, and (2) is a consequence of the solid duality theorem. For the details of this argument, see Theorem 4.3 and its proof.

Let us emphasize that in the course of running this argument, we will reprove Scholze's finiteness and primitive comparison theorems for $\mathcal{J}$. Just as in Scholze's proof, the key point here is the properness of $\mathbf{P}^{n-1}_C$, or equivalently the properness of the map g. However, Scholze exploits this properness in a very indirect way, by a Cartan-Serre style argument with shrinking covers and an elaborate use of spectral sequences. For us, the properness of g is used entirely through the identification $g_! = g_*$.

1.3. The stacky 6-functor formalism

In the process of proving our main results we also extend the 6-functor formalism from [Man22] to 'stacky' maps and study the behavior of this extended 6-functor formalism in the setting of classifying stacks. More concretely, we say that a map $f : Y \to X$ of small v-stacks is *p-fine* if it admits a cover $g : Z \to Y$ such that g admits universal p-codescent and $f \circ g$ is bdcs (see Definition 2.5). Using the results on abstract 6-functor formalisms in [Man22], it is formal that the 6-functor formalism extends to p-fine maps, i.e. to any

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

p-fine map f we get associated functors $f_!$ and $f^!$ with all the expected properties. The typical example of a map satisfying universal p-codescent is a p-cohomologically smooth cover, but it is important to not restrict to smooth covers in the definition of p-fine maps.

Now, if G is a profinite group which has finite p-cohomological dimension, we prove that the structure map $f\colon */G \to *$ is p-fine and that $f_! = f_*$ (see Corollary 3.12). The key point here is to show that the obvious cover $* \to */G$, which is not p-cohomologically smooth, nevertheless admits universal p-codescent (see Lemma 3.11); the argument for this crucially uses Mathew's notion of descendability. This result easily implies that for all *locally* profinite groups G which *locally* have finite p-cohomological dimension, the map $*/G \to *$ is p-fine. Moreover, we provide a precise criterion on G determining whether or not $*/G \to *$ is p-cohomologically smooth (see Theorem 3.16). If this is the case then we say that G is *virtually p-Poincaré*. Using classical results of Lazard [Laz65], we show that every p-adic Lie group G is virtually p-Poincaré (see Theorem 3.18), so that for these groups the associated map $*/G \to *$ is p-cohomologically smooth. This applies to groups like $\mathrm{GL}_n(\mathbf{Q}_p)$ and $D^\times$. With these results in hand, we deduce a p-adic version of the ℓ-adic smoothness criterion [Sch18a, Proposition 24.2]:

Proposition 1.5 (see Corollary 3.17). *Let G be a locally profinite group which is virtually p-Poincaré (e.g. G could be any p-adic Lie group). Suppose G acts on a small v-stack Y and $f\colon Y \to X$ is a G-equivariant separated bdcs map of small v-stacks, where G acts trivially on X. Then the induced map $Y/G \to X$ is p-fine, and it is p-cohomologically smooth as soon as f is p-cohomologically smooth.*

Note that in the ℓ-adic case, one proves [Sch18a, Proposition 24.2] directly, and then deduces that classifying stacks of p-adic Lie groups are ℓ-cohomologically smooth. Such a direct argument seems very unlikely in the p-adic setting, due to the nonexistence of Haar measures.

The above results on the stacky 6-functor formalism and classifying stacks are useful for proving the crucial Proposition 3.26. However, all we really need in that proof is the more classical statement that smooth group cohomology of p-adic Lie groups satisfies Poincaré duality (as proved by Lazard). The real power of the stacky 6-functor formalism will reveal itself when we apply it to Bun_G in subsequent papers.

1.4. Applications to representation theory

One of our main motivations for proving the duality theorem is that Kohlhaase's duality functors are intimately related to the *Gelfand-Kirillov dimension* $\dim_G \pi$ of an admissible representation π. Recall that this is a nonnegative integer measuring the asymptotic size of π; for instance, $\dim_G \pi = 0$ exactly when $\dim_{\mathbf{F}_p} \pi < \infty$.

Computing the GK dimension of admissible representations arising in nature is generally quite difficult. However, our duality theorem gives a powerful tool for understanding how the functor $\mathcal{J}$ interacts with Gelfand-Kirillov dimension. Our first application in this direction is the following result, which gives a quantitative form of Scholze's finiteness theorem.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Theorem 1.6. *Let π be an admissible smooth $G = \mathrm{GL}_n(F)$-representation, and let $N_\pi \geq 0$ be the least integer such that $\mathcal{J}^i(\pi) = 0$ for all $i > N_\pi$. Then*

$$\dim_{D^\times} \mathcal{J}^i(\pi) \leq \dim_G(\pi) + N_\pi$$

for all i. In particular, since $N_\pi \leq 2n-2$, we have

$$\dim_{D^\times} \mathcal{J}^i(\pi) \leq \dim_G(\pi) + 2n - 2$$

for all π and all i.

Note that the possible increase of dimension by $2n-2$ is quite small compared to the maximal Gelfand-Kirillov dimension in this setting, which is $\dim G = \dim D^\times = n^2[F:\mathbf{Q}_p]$. Nevertheless, we conjecture that Theorem 1.6 is not optimal, and that in fact one always has an inequality

$$\dim_{D^\times} \mathcal{J}^i(\pi) \leq \dim_G(\pi),$$

cf. Conjecture 4.7 below. This conjecture in turn implies that $\mathcal{J}^{2n-2}(\pi)$ is always a finite-dimensional vector space, which can be computed explicitly (Proposition 4.9).

Under a stronger assumption on π, we can also prove a bound in the other direction, namely that for some i the association $\pi \rightsquigarrow \mathcal{J}^i(\pi)$ cannot decrease the GK dimension by very much. Here is a precise statement.

Theorem 1.7. *Let π be an admissible smooth G-representation which is Cohen-Macaulay in the sense that $\mathcal{S}^i_G(\pi)$ is nonzero for a single value of i. Then either $\mathcal{J}(\pi) = 0$ or*

$$\max_i \dim_{D^\times} \mathcal{J}^i(\pi) \geq \dim_G \pi - (2n-2).$$

We note that the Cohen-Macaulay condition is often satisfied in practice, and is preserved by parabolic induction.

The duality theorem is even more useful when combined with nontrivial vanishing theorems for $\mathcal{J}$. To illustrate this, we now specialize to the case $n = 2$, but with $F/\mathbf{Q}_p$ still arbitrary. Set $d = [F:\mathbf{Q}_p]$. Let $B \subset G = \mathrm{GL}_2(F)$ be the upper-triangular Borel subgroup, with $T = F^\times \times F^\times \subset B$ the diagonal maximal torus. For any smooth characters $\chi_1, \chi_2 : F^\times \to \mathbf{F}_p^\times$, we have the usual parabolic induction $\pi(\chi_1,\chi_2) = \mathrm{Ind}_B^G(\omega\chi_1 \otimes \chi_2)$, where ω is the mod-p cyclotomic character. We assume that $\chi_1/\chi_2 \neq \omega^{\pm 1}$. Then $\pi(\chi_1,\chi_2)$ and $\pi(\chi_2,\chi_1)$ are irreducible, and are exchanged by the functor $\mathcal{S}^d_G(-) \otimes (\chi_1\chi_2\omega \circ \det)$. Moreover, $\mathcal{S}^i_G(-)$ applied to either representation vanishes for all $i \neq d$.

By irreducibility, it is easy to see that $\mathcal{J}^0(\pi(\chi_1,\chi_2)) = 0$, and likewise with the characters exchanged. Much less obviously, the main theorem of [JLH21] shows also that $\mathcal{J}^2(\pi(\chi_1,\chi_2)) = 0$. With these vanishing results in mind, we define

$$\tau(\chi_1,\chi_2) = \mathcal{J}^1(\pi(\chi_1,\chi_2)).$$

Our duality theorem then implies the following dichotomy.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Theorem 1.8. *Notation and assumptions as above, exactly one of the following is true.*

i. $\tau(\chi_1,\chi_2)=\tau(\chi_2,\chi_1)=0$.

ii. Both $\tau(\chi_1,\chi_2)$ *and* $\tau(\chi_2,\chi_1)$ *are nonzero, and are exchanged by the functor* $\mathcal{S}^d_{D^\times}(-)\otimes(\chi_1\chi_2\omega\circ\mathrm{Nm})$. *In particular, they are both of GK dimension d. Moreover,* $\mathcal{S}^i_{D^\times}(-)$ *applied to either representation vanishes for all* $i\neq d$.

We conjecture that in fact scenario i. in this theorem *never* occurs. When $F=\mathbf{Q}_p$ and $p>3$, this follows from work of Paškūnas [Paš22], who showed by global methods that $\tau(\chi_1,\chi_2)$ and $\tau(\chi_2,\chi_1)$ cannot *both* vanish. Combining his results with the previous theorem yields the following corollary.

Corollary 1.9. *Maintain the notation and assumptions above, and assume moreover that* $F=\mathbf{Q}_p$ *and* $p>3$. *Then* $\tau(\chi_1,\chi_2)$ *and* $\tau(\chi_2,\chi_1)$ *are nonzero, of GK dimension one, and are exchanged by the functor* $\mathcal{S}^1_{D^\times}(-)\otimes(\chi_1\chi_2\omega\circ\mathrm{Nm})$. *Moreover,* $\mathcal{S}^0_{D^\times}(-)$ *applied to either representation vanishes.*

We note that the final vanishing statement here has a more concrete interpretation: it simply means that $\tau(\chi_1,\chi_2)$ does not admit any nonzero H-stable quotient σ with $\dim_{\mathbf{F}_p}\sigma<\infty$, for any open subgroup $H\subset D^\times$ (and similarly for $\tau(\chi_2,\chi_1)$). Intuitively, this means that $\tau(\chi_1,\chi_2)$ is 'maximally non-semisimple'. It would be very interesting to find an explicit construction of one-dimensional admissible representations of $D^\times$ with this property. We also note that a similar but more complicated theorem holds in the setting of *reducible* principal series representations of $\mathrm{GL}_2(\mathbf{Q}_p)$; see §4.3.1 for details.

Our last result is an intriguing alternative formula for $\mathcal{J}(\pi)$ in the case where π is parabolically induced. To state this result, we momentarily change our notation slightly, to emphasize the dependence of the Jacquet-Langlands functor on the chosen integer n. More precisely, we write $\mathcal{J}_n$ for the Jacquet-Langlands functor from $\mathcal{D}(\mathrm{GL}_n(F))$ towards $\mathcal{D}(D^\times_{1/n})$. Choose an integer $0<d<n$, and let $P_{n-d,d}\subset\mathrm{GL}_n(F)$ be the usual upper-triangular parabolic with Levi $\mathrm{GL}_{n-d}(F)\times\mathrm{GL}_d(F)$. If $\sigma_1\boxtimes\sigma_2$ is any smooth representation of the Levi, we write $\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1\boxtimes\sigma_2)$ for the usual parabolic induction to a smooth representation of $\mathrm{GL}_n(F)$. In §5 below, we define a certain auxiliary spatial diamond $W_{n,d}$ over $\mathrm{Spd}\,C$, which comes equipped with a natural $D^\times_{1/n}$-action as well as natural $D^\times_{1/n}$-equivariant maps $h_1:W_{n,d}\to B\underline{\mathrm{GL}_{n-d}(F)}$ and $h_2:W_{n,d}\to B\underline{D^\times_{1/d}}$. In the special case $d=1$, $W_{n,1}$ is exactly the quotient $\mathcal{M}_P=\mathcal{M}_\infty/\underline{P_{n-1,1}}$ studied in [JLH21].

Theorem 1.10. *Notation as above, there is a natural isomorphism*

$$\mathcal{J}_n(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1\boxtimes\sigma_2))\cong R\Gamma(W_{n,d},h_{1,\text{ét}}^*\sigma_1\otimes h_{2,\text{ét}}^*\mathcal{J}_d(\sigma_2)).$$

We like to think of this as a 'partial' Künneth formula. Although this theorem is logically independent of our other main results, the key point in the proof is still the new machinery provided by [Man22].

Just as in [JLH21], this alternative formula leads to new vanishing results. Here we state the simplest such corollary.

Corollary 1.11. *Let* π *be an irreducible admissible representation of* $\mathrm{GL}_3(F)$ *which is not supersingular, and whose underlying vector space is infinite-dimensional. Then* $\mathcal{J}^4(\pi)=0$.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

1.5. A word on notation

This paper depends heavily on the mathematics developed in [Man22]. However, since DH is less derived than LM, there are some small notational differences between [Man22] and the present paper. In particular, we will write $R\mathrm{Hom}$ for the derived external Hom in any sheaf theory, and $R\Gamma$ for derived global sections. These are notated as Hom and Γ in [Man22]. On the other hand, we follow [Man22] in writing $\underline{\mathrm{Hom}}$ for the derived internal Hom, f_* for derived pushforward, etc. We hope this does not cause any confusion.

We also adapt some general terminology from [Man22]: In a derived setting, we use the word 'static' to refer to objects that are concentrated in degree 0 (e.g. a static object in $\mathcal{D}(\mathrm{Ab})$ is the same as an abelian group). In a condensed setting we use the word 'discrete' to indicate that the underlying condensed set comes from a discrete topological space. We sometimes use a subscript ω to refer to the subcategory of discrete objects. For example, if $\mathcal{D}_\square(\mathbf{Z})$ denotes the derived category of solid abelian groups, then $\mathcal{D}_\square(\mathbf{Z})_\omega = \mathcal{D}(\mathrm{Ab})$ is the derived category of abelian groups.

2. Geometric preparations

In this section we recall the p-adic 6-functor formalism constructed in [Man22] and extend it to a context where it can be applied to classifying stacks of p-adic Lie groups. In particular we extend the 6-functor formalism to incorporate shriek functors along 'stacky' maps, similar to the stacky extension of the ℓ-adic formalism provided in [GHW22].

2.1. Recollections on the solid six functors

We recall the p-adic 6-functor formalism on small v-stacks defined in [Man22]. Let us first introduce the correct category of geometric objects to work with. In order to make sense of the sheaf $\mathcal{O}_X^+/\pi$ on a small v-stack, we need to choose an 'untilt' $X^\sharp$ of X (this is automatically chosen if X comes from an adic space over $\mathbf{Z}_p$) and a pseudouniformizer π on $X^\sharp$. To make things more precise, we define an *untilted small v-stack* $X^\sharp$ to be a small v-stack X together with a chosen map $X \to \mathrm{Spd}\,\mathbf{Z}_p$. For every map $Y \to X$ from a perfectoid space Y to X, the untilt $X^\sharp$ induces an untilt $Y^\sharp$ (in the sense of perfectoid spaces) of Y. The untilted small v-stack $X^\sharp$ comes equipped with a structure sheaf $\mathcal{O}_{X^\sharp}$ and an integral structure sheaf $\mathcal{O}_{X^\sharp}^+ \subset \mathcal{O}_{X^\sharp}$ on X_{v} such that on every affinoid perfectoid $Y \in X_{\mathrm{v}}$ we have $\mathcal{O}_{X^\sharp}(Y) = \mathcal{O}_{Y^\sharp}(Y^\sharp)$ (and similarly for $\mathcal{O}_{X^\sharp}^+$). A *pseudouniformizer* on $X^\sharp$ is a sheaf of ideals $\pi \subset \mathcal{O}_{X^\sharp}^+$ such that for every affinoid perfectoid $Y = \mathrm{Spa}(B, B^+) \in X_{\mathrm{v}}$ the restriction of π to Y is a principal ideal and generated by a pseudouniformizer in $B^{\sharp +}$. We denote by $\mathrm{vStack}_\pi^\sharp$ the category of pairs $(X^\sharp, \pi)$, where $X^\sharp$ is an untilted small v-stack and π is a pseudouniformizer on $X^\sharp$. A map $(Y^\sharp, \pi_Y) \to (X^\sharp, \pi_X)$ in $\mathrm{vStack}_\pi^\sharp$ is a map of small v-stacks $Y \to X$ over $\mathrm{Spd}\,\mathbf{Z}_p$ such that $\pi_X|_Y = \pi_Y$.

The first main result of [Man22] is the following:

Theorem 2.1 (see [Man22, Theorem 3.1.27]). *There is a unique hypercomplete v-sheaf*

$$(\mathrm{vStack}_\pi^\sharp)^{\mathrm{op}} \to \mathcal{C}at_\infty^\otimes, \qquad (X^\sharp, \pi) \mapsto \mathcal{D}_\square^a(\mathcal{O}_{X^\sharp}^+/\pi)$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

such that for every untilted affinoid perfectoid space $X^\sharp = \mathrm{Spa}(A^\sharp, A^{\sharp+})$ *with pseudouniformizer* $\pi \in A^{\sharp+}$*, where* X *is weakly of perfectly finite type over a totally disconnected space, we have*

$$\mathcal{D}^a_\square(\mathcal{O}^+_{X^\sharp}/\pi) = \mathcal{D}^a_\square(A^{\sharp+}/\pi).$$

Moreover, the pullback functors in this v-sheaf are induced by the solid base-change functors.

Here $\mathcal{C}at^\otimes_\infty$ denotes the ∞-category of symmetric monoidal ∞-categories and $\mathcal{D}^a_\square(A^{\sharp+}/\pi)$ denotes the derived ∞-category of *solid almost* $A^{\sharp+}/\pi$*-modules*. The term 'almost' refers to the fact that we care about everything only up to $\mathfrak{m}_{A^\sharp}$-torsion. The term 'solid' was introduced by Clausen-Scholze in [Sch19] and roughly means 'complete topological'. See [Man22, §2.2, §2.9] for the precise definition of $\mathcal{D}^a_\square(A^{\sharp+}/\pi)$.

In the following we abstract a little further in order to simplify notation: We denote vStack_Λ the category of pairs (X,Λ), where X is a small v-stack and Λ is a system of integral torsion coefficients on X (see [Man22, Definition 3.2.10]). The main example for Λ is $\Lambda = \mathcal{O}^+_{X^\sharp}/\pi$ for an untilt $X^\sharp$ of X and a pseudouniformizer π on $X^\sharp$ – the reader is invited to always assume this case and replace vStack_Λ by $\mathrm{vStack}^\sharp_\pi$ in the following (note that these two categories are not precisely the same, as the latter captures an untilt while the former only captures the relevant sheaf of coefficients). To every $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ one can then define the ∞-category $\mathcal{D}^a_\square(X,\Lambda)$ in the same way as in Theorem 2.1.

The second main result of [Man22] is the construction of a 6-functor formalism for the ∞-categories $\mathcal{D}^a_\square(\mathcal{O}^+_X/\pi)$. In order to formulate it, let us recall the definition of an abstract 6-functor formalism from [Man22, §A.5]: Suppose we are given an ∞-category $\mathcal{C}$ of geometric objects and a collection of edges E in $\mathcal{C}$ which contains all isomorphisms and is stable under composition and pullback. To the pair $(\mathcal{C},E)$ one can then associate an ∞-operad $\mathrm{Corr}(\mathcal{C})^\otimes_{E,all}$ which roughly has the following shape. The underlying ∞-category $\mathrm{Corr}(\mathcal{C})_{E,all}$ has as objects the same objects as $\mathcal{C}$ and as morphisms the diagrams of the form $Y \leftarrow Y' \to X$ such that $[Y' \to X] \in E$, where the composition of morphisms is given by the obvious pullback diagrams. The operadic structure on $\mathrm{Corr}(\mathcal{C})_{E,all}$ is roughly given by $X \otimes Y = X \times Y$ (if $\mathcal{C}$ admits products then this is literally true and $\mathrm{Corr}(\mathcal{C})_{E,all}$ is a symmetric monoidal ∞-category). With this notation at hand, a *pre-6-functor formalism* on $(\mathcal{C},E)$ is defined to be a map of ∞-operads $\mathcal{D}\colon \mathrm{Corr}(\mathcal{C})^\otimes_{E,all} \to \mathcal{C}at^\times_\infty$, where $\mathcal{C}at^\times_\infty$ denotes the cartesian structure on $\mathcal{C}at_\infty$. Such a functor $\mathcal{D}$ encodes the following data: For every $X \in \mathcal{C}$ we get a symmetric monoidal ∞-category $\mathcal{D}(X)$, for every morphism $f\colon Y \to X$ in $\mathcal{C}$ we get a symmetric monoidal functor $f^*\colon \mathcal{D}(X) \to \mathcal{D}(Y)$ and for every morphism $[f\colon Y \to X] \in E$ we get a functor $f_!\colon \mathcal{D}(Y) \to \mathcal{D}(X)$ such that the functors $\otimes$, f^* and $f_!$ satisfy the usual compatibilities. We say that $\mathcal{D}$ is a *6-functor formalism* if the functors $\otimes$, f^* and $f_!$ admit right adjoints.

In the setting of the desired p-adic 6-functor formalism, we take $\mathcal{C} = \mathrm{vStack}_\Lambda$ and $E = bdcs$, the class of those maps $f\colon Y \to X$ of small v-stacks which are p-bounded (this roughly means that the map has finite p-cohomological dimension, see [Man22, §3.5]), locally compactifiable and representable in locally spatial diamonds. The main result of [Man22] now reads as follows:

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Theorem 2.2 (see [Man22, Theorem 3.6.12]). *There is a 6-functor formalism*

$$\mathcal{D}^a_{\square}\colon \mathrm{Corr}(\mathrm{vStack}_\Lambda)^{\otimes}_{bdcs,all} \to \mathcal{C}at^{\times}_{\infty}, \qquad (X,\Lambda) \mapsto \mathcal{D}^a_{\square}(X,\Lambda)$$

with the following properties:

(i) *The symmetric monoidal ∞-categories $\mathcal{D}^a_{\square}(X,\Lambda)$ and pullback functors f^* coincide with the version in Theorem 2.1.*

(ii) *For every bdcs morphism $f\colon Y \to X$ in vStack_Λ the functor $f_!\colon \mathcal{D}^a_{\square}(Y,\Lambda) \to \mathcal{D}^a_{\square}(X,\Lambda)$ has the following properties: If f is proper then $f_! = f_*$ is right adjoint to f^*. If f is étale then $f_!$ is left adjoint to f^*.*

2.2. Six functors for stacky morphisms

In order to efficiently apply the above 6-functor formalism to representation theory, i.e. sheaves on classifying stacks, it is very useful to generalize the definition of the shriek functors $f_!$ and $f^!$ to a larger class of maps f. Namely, we need to allow certain 'stacky' maps f as well, similar to what is done in [GHW22].

Definition 2.3. We say that a bdcs map $f\colon Y \to X$ of small v-stacks admits *universal p-codescent* if it satisfies the following property: Given any $(X',\Lambda) \in \mathrm{vStack}_\Lambda$ which admits a map (in vStack_Λ) to some strictly totally disconnected perfectoid space and given any map $X' \to X$ with base-change $f'\colon Y' \to X'$ and Čech nerve $Y'_\bullet \to X'$ the natural functor

$$\mathcal{D}^{a!}_{\square}(X',\Lambda) \xrightarrow{\sim} \varprojlim_{n\in\Delta} \mathcal{D}^{a!}_{\square}(Y'_n,\Lambda)$$

is an equivalence. Here $\mathcal{D}^{a!}_{\square}$ denotes the functor $(Z,\Lambda) \mapsto \mathcal{D}^a_{\square}(Z,\Lambda)$, $f \mapsto f^!$.

Remark 2.4. The assumption that (X',Λ) admits a map to some strictly totally disconnected perfectoid space is not necessary for the abstract theory of universal p-codescent, but it gives us more flexibility when showing that a map has this property, see the proof of Lemma 3.11.

Definition 2.5. A map $f\colon Y \to X$ of small v-stacks is called *p-fine* if there is a map $g\colon Z \to Y$ such that $f \circ g$ is bdcs and g is bdcs and admits universal p-codescent.

Lemma 2.6.

(i) *The condition of being p-fine is étale local on source and target.*

(ii) *The collection of p-fine maps is stable under composition and base-change.*

(iii) *Every bdcs map is p-fine. In particular every étale map and every map of rigid varieties over a fixed base field is p-fine.*

(iv) *Let $f\colon Y \to X$ and $g\colon Z \to Y$ be maps of small v-stacks. If f and $f \circ g$ are p-fine, then so is g.*

Proof. It follows from [LZ17, Lemma 3.1.2.(4)] that bdcs maps admitting universal p-codescent are stable under composition. Part (ii) of the claim follows immediately and part (iii) is obvious (see [Man22, Example 3.6.11] for the second claim in (iii)). Part (i) follows from the observation that étale maps admit universal p-codescent.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

It remains to prove (iv), so let f and g be given as in the claim. Choose a bdcs map $Y' \to Y$ admitting universal p-codescent such that the composed map $Y' \to X$ is bdcs. Denote $g' \colon Z' \to Y'$ the base-change of g along $Y' \to Y$. Since the map $Z' \to Z$ is bdcs, it follows from (ii) that the composed map $Z' \to X$ is p-fine. We can therefore choose a bdcs map $Z'' \to Z'$ admitting universal p-codescent such that $Z'' \to X$ is bdcs. Note that the map $Z'' \to X$ factors as $Z'' \to Y' \to X$ and since $Y' \to X$ is bdcs it follows from [Man22, Lemma 3.6.10.(iv)] that $Z'' \to Y'$ is bdcs. In particular $Z'' \to Y$ is bdcs, so the map $Z'' \to Z$ shows that g is p-fine. □

It is completely formal that the p-adic 6-functor formalism extends to all p-fine maps (in fact a similar extension can be performed in any 6-functor formalism):

Proposition 2.7. *The 6-functor formalism in Theorem 2.2 extends uniquely to a 6-functor formalism*

$$\mathcal{D}^a_{\square} \colon \mathrm{Corr}(\mathrm{vStack}_\Lambda)_{pfine,all} \to \mathcal{C}at_\infty, \qquad (X,\Lambda) \mapsto \mathcal{D}^a_{\square}(X,\Lambda),$$

where $pfine$ denotes the class of p-fine maps in vStack_Λ.

Proof. Let $\mathcal{C} \subset \mathrm{vStack}_\Lambda$ be the full subcategory spanned by those (X,Λ) which admit a map (in vStack_Λ) to some strictly totally disconnected space. Then $\mathcal{C}$ is a basis of vStack_Λ, hence by [Man22, Proposition A.5.16] and [Man22, Proposition A.3.11. (i)] it is enough to construct the desired extension of the 6-functor formalism on $\mathcal{C}$, i.e. we need to construct a 6-functor formalism $\mathcal{D}^a_{\square} \colon \mathrm{Corr}(\mathcal{C})_{pfine,all} \to \mathcal{C}at_\infty$ extending the one from Theorem 2.2. We now apply [Man22, Proposition A.5.14] with $E = bdcs$, $E' = pfine$ and $S \subset E$ being the subset of those maps which admit universal p-codescent. Condition (a) requires $\mathcal{D}^a_{\square}(X,\Lambda)$ to be presentable, which can be achieved by temporarily restricting to κ-condensed objects (as in the proof of [Man22, Theorem 3.6.12]). Condition (b) follows from the definition of universal p-codescent, condition (c) follows from the definition of p-fine maps and condition (d) follows from [Man22, Lemma 3.6.10] (see [Man22, Remark A.5.15.(ii)]. □

Remark 2.8. More concretely, the shriek functor $f_! \colon \mathcal{D}^a_{\square}(Y,\Lambda) \to \mathcal{D}^a_{\square}(X,\Lambda)$ along a p-fine map $f \colon Y \to X$ in vStack_Λ is computed as follows. Via v-descent we can reduce to the case that (X,Λ) admits a map to some strictly totally disconnected space. Now choose a bdcs map $g \colon Z \to Y$ such that g admits universal p-codescent and $f \circ g$ is bdcs. Let $g_\bullet \colon Z_\bullet \to Y$ denote the Čech nerve of g. It follows from the p-codescent property of g that for every $\mathcal{M} \in \mathcal{D}^a_{\square}(Y,\Lambda)$ the natural map $\varinjlim_{n\in\Delta} g_{n!}g_n^!\mathcal{M} \xrightarrow{\sim} \mathcal{M}$ is an isomorphism. Consequently we have

$$f_!\mathcal{M} = f_! \varinjlim_{n\in\Delta} g_{n!}g_n^!\mathcal{M} = \varinjlim_{n\in\Delta} (f\circ g_n)_! g_n^!\mathcal{M}.$$

Note that the maps $f \circ g_n \colon Z_n \to X$ are bdcs, hence $(f \circ g_n)_!$ was already defined in Theorem 2.2.

In order to make use of Proposition 2.7 we need a good amount of examples of bdcs maps admitting universal p-codescent. Arguably the most important class of such maps

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

is formed by the p-cohomologically smooth maps (see [Man22, Definition 3.8.1]). The following is an analog of [GHW22, Proposition 3.16]:

Lemma 2.9. *Every bdcs and p-cohomologically smooth map of small v-stacks admits universal p-codescent.*

Proof. Let $f\colon Y \to X$ be a bdcs and p-cohomologically smooth map in vStack_Λ with Čech nerve $f_\bullet\colon Y_\bullet \to X$ (i.e. $Y_0 = Y$, $Y_1 = Y \times_X Y$ and so on); we need to show that the natural functor $f_\bullet^!\colon \mathcal{D}_\square^{a!}(X,\Lambda) \xrightarrow{\sim} \varprojlim_{n\in\Delta} \mathcal{D}_\square^{a!}(Y_n,\Lambda)$ is an isomorphism. This functor admits a left adjoint

$$f_{\bullet!}\colon \varprojlim_{n\in\Delta} \mathcal{D}_\square^{a!}(Y_n,\Lambda) \to \mathcal{D}_\square^{a!}(X,\Lambda), \qquad \mathcal{M}_\bullet \mapsto \varinjlim_{n\in\Delta} f_{n!}\mathcal{M}_n.$$

Note that $f_\bullet^!$ is conservative: This follows from the fact that $f_0^!$ is conservative, which in turn is implied by the fact that $f_0^!$ is just a twist of f_0^* by definition of the p-cohomological smoothness of f_0. Thus in order to show that $f_\bullet^!$ is an equivalence, it is enough to show that $f_{\bullet!}$ is fully faithful (cf. [Lur18, Corollary 03UZ]), i.e. for given $\mathcal{M}_\bullet \in \varprojlim_{n\in\Delta} \mathcal{D}_\square^{a!}(Y_n,\Lambda)$ we need to see that the natural map $f_\bullet^! f_{\bullet!}\mathcal{M}_\bullet \xrightarrow{\sim} \mathcal{M}_\bullet$ is an isomorphism. This boils down to showing that the map $f_0^! f_{\bullet!}\mathcal{M}_\bullet \to \mathcal{M}_0$ is an isomorphism. Since f_0 is p-cohomologically smooth, $f_0^!$ preserves small colimits, so we end up with the claim that the natural map

$$\varinjlim_{n\in\Delta} f_0^! f_{n!}\mathcal{M}_n \xrightarrow{\sim} \mathcal{M}_0$$

is an isomorphism. Let us denote by $f'_\bullet\colon Y'_\bullet \to Y_0$ the base-change of $Y_\bullet \to X$ along $f_0\colon Y_0 \to X$. Denote $g_\bullet\colon Y'_\bullet \to Y_\bullet$ the natural projections. Then for every $n \geq 0$ the natural morphism $f'_{n!} g_n^! \mathcal{M}_n \xrightarrow{\sim} f_0^! f_{n!}\mathcal{M}_n$ is an isomorphism (by p-cohomological smoothness of f_0 and g_n this reduces to proper base-change). On the other hand $g_\bullet^!\mathcal{M}_\bullet$ defines an object of $\varprojlim_{n\in\Delta} \mathcal{D}_\square^{a!}(Y'_n,\Lambda)$ so the above claimed isomorphism reduces to showing that the natural functor $\mathcal{D}_\square^{a!}(Y_0,\Lambda) \xrightarrow{\sim} \varprojlim_{n\in\Delta} \mathcal{D}_\square^{a!}(Y'_n,\Lambda)$ is an equivalence. This means that we can replace $Y_\bullet \to X$ by $Y'_\bullet \to Y_0$. But the latter augmented simplicial object is split, hence the desired equivalence of ∞-categories is formal. □

Let $\mathrm{vStack}^\sharp_{\pi|p} \subset \mathrm{vStack}^\sharp_\pi$ be the full subcategory spanned by those pairs $(X^\sharp,\pi)$ such that $\pi|p$. Recall the definition of the ∞-category $\mathcal{D}_\square^a(\mathcal{O}^+_{X^\sharp}/\pi)^\varphi$ from [Man22, Definition 3.9.10. (b)]: Roughly, this is the ∞-category of pairs $(\mathcal{M},\varphi_\mathcal{M})$ where $\mathcal{M} \in \mathcal{D}_\square^a(\mathcal{O}^+_{X^\sharp}/\pi)$ and $\varphi_\mathcal{M}$ is an isomorphism $\varphi_\mathcal{M}\colon \varphi^*\mathcal{M} \xrightarrow{\sim} \mathcal{M}$; here φ denotes the Frobenius $x \mapsto x^p$ on $\mathcal{O}^+_{X^\sharp}/\pi$. The above constructed 6-functor formalism extends to φ-modules:

Proposition 2.10. *The 6-functor formalism from Proposition 2.7 extends to a 6-functor formalism*

$$\mathcal{D}_\square^a(-)^\varphi\colon \mathrm{Corr}(\mathrm{vStack}_\Lambda)_{pfine,all} \to \mathcal{C}at_\infty, \qquad (X,\Lambda) \mapsto \mathcal{D}_\square^a(X,\Lambda)^\varphi.$$

The forgetful functor $\mathcal{D}_\square^a(X,\Lambda)^\varphi \to \mathcal{D}_\square^a(X,\Lambda)$ preserves all small limits and colimits and commutes with the 6 functors $\otimes$, $\underline{\mathrm{Hom}}$, f^*, f_*, $f_!$ and $f^!$.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proof. By [Man22, Proposition 3.9.13] the same is true for the 6-functor formalism on bdcs maps. To get this result also on p-fine maps, we note that the same extension procedure as in the proof of Proposition 2.7 applies verbatim to the φ-module case so that we get the desired 6-functor formalism. By [Man22, Proposition 3.9.13] and functoriality of Kan extensions the forgetful functor preserves small limits and colimits and commutes with all the 6 functors except possibly with $f^!$. To see that it also commutes with $f^!$, note first that by the same argument as in the proof of [Man22, Lemma 3.9.12. (v)] we can reduce to the case that the target of $f\colon Y \to X$ is a strictly totally disconnected space. Choose also a bdcs cover $g\colon Z \twoheadrightarrow Y$ such that g admits universal p-codescent and $f \circ g$ is bdcs. Then by universal p-codescent the functor $g^!$ is conservative, hence to check that $f^!$ commutes with the forgetful functor it is enough to check that $g^! f^!$ commutes with the forgetful functor. This follows from the bdcs case. □

In the following we will often omit the explicit mention of untilt and pseudouniformizer and simply write $X \in \mathrm{vStack}^{\sharp}_{\pi|\mathrm{p}}$. The main reason for introducing φ-modules is that they have a tight connection to actual étale $\mathbf{F}_p$-sheaves: The p-torsion Riemann-Hilbert correspondence (see [Man22, Theorem 3.9.23]) tells us that for $X \in \mathrm{vStack}^{\sharp}_{\pi|p}$ there is a natural fully faithful embedding

$$\mathcal{D}_{\text{ét}}(X,\mathbf{F}_p)^{\mathrm{oc}} \hookrightarrow \mathcal{D}^a_{\square}(\mathcal{O}_X^+/\pi)^{\varphi}, \qquad \mathcal{F} \mapsto \mathcal{F} \otimes_{\mathbf{F}_p} \mathcal{O}_X^{+a}/\pi$$

which induces an equivalence of dualizable objects on both sides. Here $\mathcal{D}_{\text{ét}}(X,\mathbf{F}_p)^{\mathrm{oc}}$ denotes the (derived) ∞-category of overconvergent $\mathbf{F}_p$-sheaves. The Riemann-Hilbert functor admits a right adjoint $(-)^{\varphi}\colon \mathcal{D}^a_{\square}(\mathcal{O}_X^+/\pi)^{\varphi} \to \mathcal{D}_{\text{ét}}(X,\mathbf{F}_p)^{\mathrm{oc}}$.

2.3. Cohomologically smooth maps

The notion of p-cohomologically smooth maps of small v-stacks generalizes in a straightforward way from bdcs maps to p-fine maps:

Definition 2.11. A p-fine map $f\colon Y \to X$ of small v-stacks is called *p-cohomologically smooth* if for every map $X' \to X$ from a strictly totally disconnected perfectoid space $X' = \mathrm{Spa}(A', A'^+)$ there is a pseudouniformizer $\pi \in A'^+$ such that the map $f'\colon Y' := Y \times_X X' \to X'$ satisfies the following property: The natural morphism $f'^* \otimes f'^!(\mathcal{O}_{X'}^{+a}/\pi) \xrightarrow{\sim} f'^!$ is an isomorphism of functors $\mathcal{D}^a_{\square}(\mathcal{O}_{X'}^+/\pi) \to \mathcal{D}^a_{\square}(\mathcal{O}_{Y'}^+/\pi)$ and $f'^!(\mathcal{O}_{X'}^{+a}/\pi)$ is invertible.

The basic results on bdcs and p-cohomologically smooth maps extend to the p-fine setting, as follows.

Proposition 2.12. (i) *Let $f\colon Y \to X$ be a p-fine and p-cohomologically smooth map in* $\mathrm{vStack}_{\Lambda}$. *Then the natural morphism*

$$f^* \otimes f^! \Lambda^a \xrightarrow{\sim} f^!$$

is an isomorphism of functors $\mathcal{D}^a_{\square}(X,\Lambda) \to \mathcal{D}^a_{\square}(Y,\Lambda)$ and $f^!\Lambda^a$ is invertible.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

(ii) *Let*

$$\begin{array}{ccc} Y' & \xrightarrow{g'} & X' \\ \downarrow{\scriptstyle f'} & & \downarrow{\scriptstyle f} \\ X' & \xrightarrow{g} & X \end{array}$$

be a cartesian square in vStack_Λ.

(a) *Assume that f is p-fine and that either f or g is p-fine and p-cohomologically smooth. Then the natural morphism*

$$g'^* f^! \xrightarrow{\sim} f'^! g^*$$

is an isomorphism of functors $\mathcal{D}^a_\square(X,\Lambda) \to \mathcal{D}^a_\square(Y',\Lambda)$.

(b) *Assume that g is p-fine and p-cohomologically smooth. Then the natural morphism*

$$g^* f_* \xrightarrow{\sim} f'_* g'^*$$

is an isomorphism of functors $\mathcal{D}^a_\square(Y,\Lambda) \to \mathcal{D}^a_\square(X',\Lambda)$.

(iii) *Let $f\colon Y \to X$ be a p-fine and p-cohomologically smooth map in* vStack_Λ. *Then for all $\mathcal{M},\mathcal{N} \in \mathcal{D}^a_\square(X,\Lambda)$ the natural morphism*

$$f^* \underline{\mathrm{Hom}}(\mathcal{M},\mathcal{N}) \xrightarrow{\sim} \underline{\mathrm{Hom}}(f^*\mathcal{M}, f^*\mathcal{N})$$

is an isomorphism.

Proof. We first note that part (i) is formal if X is strictly totally disconnected and follows for general X once we have shown (ii).(a) (see the proof of [Man22, Proposition 3.8.4. (i)]). Moreover, once we have shown (ii).(a), also (iii) is formal (see the proof of [Man22, Proposition 3.8.7]) and (ii).(b) is formal (see the proof of [Man22, Proposition 3.8.6.(ii)]. Therefore, the whole claim boils down to showing (ii).(a). By the proof of [Man22, Proposition 3.8.6.(iii)] the case that g is p-cohomologically smooth follows formally from the case that f is p-cohomologically smooth. Thus from now on we can assume that we are given a cartesian square as in (ii) with f being p-fine and p-cohomologically smooth and we need to show that the morphism $g'^* f^! \xrightarrow{\sim} f'^! g^*$ is an isomorphism. As in the proof of [Man22, Proposition 3.8.4.(ii)] we can formally reduce to the case that X and X' are strictly totally disconnected. Then by (i) (which holds unconditionally on (ii).(a) in the strictly totally disconnected case) we are reduced to showing that the natural map $\alpha\colon g'^* f^! \Lambda^a \xrightarrow{\sim} f'^! \Lambda^a$ is an isomorphism. We claim that g'_* is conservative on bounded objects (in the sense of [Man22, Definition 3.3.1. (a)]): Pick any cover $Z' \twoheadrightarrow X'$ such that the composition $Z \to X$ is bdcs (this can be done by the definition of p-fine maps). Then Z' is a locally spatial diamond, so we can find a pro-étale cover of Z' by totally disconnected spaces Z_i. Letting $W_i := Z_i \times_{X'} Y' = Z_i \times_X X'$ and denoting $g'_i\colon W_i \to Z_i$ the base-change of g', we note that each W_i is a p-bounded affinoid perfectoid space. Hence by [Man22, Theorem 3.5.21] each g'_i is a forgetful functor of module categories and in particular conservative. Using also bounded base-change (see [Man22, Proposition 3.3.7]) we deduce that g' is conservative on bounded objects. Now note that $f^! \Lambda^a$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

is bounded: If $\Lambda = \mathcal{O}_X^+/\pi$ for some $\pi|p$ then this follows from the Riemann-Hilbert correspondence (see [Man22, Theorem 3.9.23]); for general Λ one can make the usual reductions. Similarly $f'^!\Lambda^a$ is bounded. Thus altogether we deduce that it is enough to show that $g'_*\alpha\colon g'_*g'^*f^!\Lambda^a \xrightarrow{\sim} g'_*f'^!\Lambda^a$ is an isomorphism. The remaining proof of [Man22, Proposition 3.8.4.(ii)] applies verbatim. □

Lemma 2.13.

(i) *The condition of being p-fine and p-cohomologically smooth is analytically local on both source and target.*

(ii) *Among p-fine maps, the condition of being p-cohomologically smooth is v-local on the target and p-cohomologically smooth local on the source.*

(iii) *p-fine and p-cohomologically smooth maps are stable under composition and base-change.*

Proof. This is formal, see [Man22, Lemma 3.8.5]. □

3. Representation-theoretic preparations

In this section, we apply the 6-functor formalism to study representations of p-adic Lie groups. Since a number of categories will appear in what follows, we begin by introducing them. For the moment, we allow G to be any locally profinite group with locally finite p-cohomological dimension (in the sense of [SW00, p. 4.1]).

First of all, we will need the ∞-category of solid G-representations $\mathcal{D}_\square(G,\mathbf{F}_p) = \mathcal{D}_\square(\mathbf{F}_p)^{BG}$, as defined in [Man22, Definition 3.4.3]. By definition it is the derived ∞-category of the category of continuous G-representations on solid $\mathbf{F}_p$-vector spaces (the latter being defined as the category of modules over the analytic ring $\mathbf{F}_{p\square}[G]$ – one can equivalently describe this category in terms of solid $\mathbf{F}_p$-vector spaces V together with a map of condensed sets $G \times V \to V$ satisfying the unit, associativity and a certain $\mathbf{F}_p$-linearity condition). We also have the derived category $\mathcal{D}^{\mathrm{sm}}_\square(G,\mathbf{F}_p)$ of smooth G-representations on solid $\mathbf{F}_p$-vector spaces [Man22, Definition 3.4.11]. There is a natural functor $\mathcal{D}^{\mathrm{sm}}_\square(G,\mathbf{F}_p) \to \mathcal{D}_\square(G,\mathbf{F}_p)$, which is not fully faithful. Most classically, we have the usual derived category $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$ of the abelian category of smooth G-representations on discrete $\mathbf{F}_p$-vector spaces. There is a natural identification $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) = \mathcal{D}^{\mathrm{sm}}_\square(G,\mathbf{F}_p)_\omega$, identifying $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$ as the full subcategory of $\mathcal{D}^{\mathrm{sm}}_\square(G,\mathbf{F}_p)$ consisting of those smooth representations whose underlying $\mathbf{F}_p$-vector space is discrete (in the topological, i.e. condensed, sense), which we will use without comment. We also note that the composite functor $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) \to \mathcal{D}_\square(G,\mathbf{F}_p)$ is fully faithful.

Now fix an algebraically closed nonarchimedean field C of residue characteristic p, and a pseudouniformizer $\pi|p$. Let $B\underline{G} = \mathrm{Spd}(C)/\underline{G}$ denote the classifying stack of G. We then have the ∞-category $\mathcal{D}^a_\square(\mathcal{O}^+_{B\underline{G}}/\pi)^\varphi$, which (by [Man22, Lemma 3.4.26.(ii)]) can be identified with the ∞-category of smooth G-representations on solid almost φ-modules over $\mathcal{O}_C^+/\pi$. This category can be directly accessed by the 6-functor formalism, as we will see. On the other hand, Proposition 3.9 gives a natural identification $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) = \mathcal{D}_{\text{ét}}(B\underline{G},\mathbf{F}_p)$. The Riemann-Hilbert functor then gives a natural fully faithful symmetric

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

monoidal embedding

$$\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) \to \mathcal{D}^a_{\square}(\mathcal{O}^+_{B\underline{G}}/\pi)^{\varphi},$$

which factors through the full subcategory $\mathcal{D}^a_{\square}(\mathcal{O}^+_{B\underline{G}}/\pi)^{\varphi}_{\omega}$ of discrete objects.

3.1. Smooth representations

We continue to fix a locally profinite group G of locally bounded p-cohomological dimension. In this section, we recall some results on the derived ∞-category $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$. However, to place our results in context, we momentarily allow our coefficients to be any field k (still assuming that G has locally bounded char$k$-cohomological dimension). Thus, let $\mathrm{Rep}^{\mathrm{sm}}_k(G)$ denote the category of smooth G-representations on k-vector spaces, and let $\mathcal{D}^{\mathrm{sm}}(G,k)$ denote its derived ∞-category. When k is clear from context, we just write $\mathcal{D}(G) = \mathcal{D}^{\mathrm{sm}}(G,k)$. We sometimes write $\mathbf{1} \in \mathrm{Rep}^{\mathrm{sm}}_k(G)$ for k with the trivial G-action.

The assertions in the following proposition are well-known.

Proposition 3.1.

1. $\mathrm{Rep}^{\mathrm{sm}}_k(G)$ *is a Grothendieck abelian category. In particular, unbounded complexes have K-injective resolutions.*
2. *The derived category* $\mathcal{D}^{\mathrm{sm}}(G,k)$ *is left-complete and has a natural symmetric monoidal structure given by k-linear tensor product of cochain complexes.*
3. *If* $H \subset G$ *is an open subgroup, the restriction functor* $r = r^H_G : \mathcal{D}^{\mathrm{sm}}(G,k) \to \mathcal{D}^{\mathrm{sm}}(H,k)$ *is t-exact and symmetric monoidal.*

Let $\underline{\mathrm{Hom}}_G(A,B)$ be the internal hom in $\mathcal{D}^{\mathrm{sm}}(G,k)$, defined as a formal adjoint satisfying $\mathrm{Hom}_{\mathcal{D}(G)}(C \otimes_k A, B) = \mathrm{Hom}_{\mathcal{D}(G)}(C, \underline{\mathrm{Hom}}_G(A,B))$.

Lemma 3.2. *If* $H \subset G$ *is an open subgroup, then there is a natural isomorphism* $r\underline{\mathrm{Hom}}_G(A,B) \cong \underline{\mathrm{Hom}}_H(rA,rB)$ *for any* $A,B \in \mathcal{D}(G)$. *Here* $r : \mathcal{D}(G) \to \mathcal{D}(H)$ *denotes the restriction functor.*

Proof. The functor r has a left adjoint, given by the (t-exact) functor $\mathrm{ind}^G_H(-) := - \otimes_{k[H]} k[G]$ of compact induction. We then compute

$$\begin{aligned}
\mathrm{Hom}_{\mathcal{D}(H)}(C, r\underline{\mathrm{Hom}}_G(A,B)) &\cong \mathrm{Hom}_{\mathcal{D}(G)}(\mathrm{ind}^G_H C, \underline{\mathrm{Hom}}_G(A,B)) \\
&\cong \mathrm{Hom}_{\mathcal{D}(G)}(A \otimes \mathrm{ind}^G_H C, B) \\
&\cong \mathrm{Hom}_{\mathcal{D}(G)}(\mathrm{ind}^G_H(rA \otimes C), B) \\
&\cong \mathrm{Hom}_{\mathcal{D}(H)}(rA \otimes C, rB) \\
&\cong \mathrm{Hom}_{\mathcal{D}(H)}(C, \underline{\mathrm{Hom}}_H(rA,rB))
\end{aligned}$$

naturally in $C \in \mathcal{D}(H)$, so the result follows from Yoneda. □

If $A \in \mathcal{D}(G)$ is any object and $H \subset G$ is any open subgroup, we write

$$R\Gamma_{\mathrm{sm}}(H,A) = R\mathrm{Hom}_{\mathcal{D}(H)}(\mathbf{1}, rA) = R\mathrm{Hom}_{\mathcal{D}(G)}(\mathrm{ind}^G_H \mathbf{1}, A).$$

To actually compute this, let $A \xrightarrow{\sim} I^{\bullet}$ be a K-injective resolution by injectives in $\mathrm{Rep}^{\mathrm{sm}}_k(G)$; then $R\Gamma_{\mathrm{sm}}(H,A) \cong (I^{\bullet})^H$. When H is a p-adic Lie group, $R\Gamma_{\mathrm{sm}}(H,-)$ coincides with continuous group cohomology, by the main result in [Fus21].

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Lemma 3.3. *For any $A \in \mathcal{D}(G)$, we have $A \cong \mathrm{colim}_{H \subset G\,\mathrm{open}} R\Gamma_{\mathrm{sm}}(H,A)$.*

Proof. If $A \xrightarrow{\sim} I^\bullet$ is a K-injective resolution by injectives, then $R\Gamma_{\mathrm{sm}}(H,A) \cong (I^\bullet)^H$ for any H; here we use that restriction along $\mathrm{Rep}_k^{\mathrm{sm}}(G) \to \mathrm{Rep}_k^{\mathrm{sm}}(H)$ preserves K-injective complexes of injectives. But then

$$\mathrm{colim}_{H \subset G\,\mathrm{open}} R\Gamma_{\mathrm{sm}}(H,A) \cong \mathrm{colim}_{H \subset G\,\mathrm{open}} (I^\bullet)^H \cong I^\bullet \cong A,$$

using the fact that each I^n is a smooth representation. □

Lemma 3.4. *For any $A,B \in \mathcal{D}(G)$, we have $\underline{\mathrm{Hom}}_G(A,B) \cong \mathrm{colim}_H R\mathrm{Hom}_{\mathcal{D}(H)}(A,B)$.*

Here (and in what follows) we drop the restriction functors r from the notation.

Proof. There are natural isomorphisms

$$\begin{aligned} \mathrm{colim}_H R\mathrm{Hom}_{\mathcal{D}(H)}(A,B) &\cong \mathrm{colim}_H R\Gamma_{\mathrm{sm}}(H, \underline{\mathrm{Hom}}_H(A,B)) \\ &\cong \mathrm{colim}_H R\Gamma_{\mathrm{sm}}(H, \underline{\mathrm{Hom}}_G(A,B)) \\ &\cong \underline{\mathrm{Hom}}_G(A,B), \end{aligned}$$

where the second line follows from Lemma 3.2, and the third line follows from Lemma 3.3. □

Proposition 3.5. *When G is locally pro-p and $\mathrm{char} k \neq p$, $\underline{\mathrm{Hom}}_G(-,\mathbf{1})$ is the usual exact smooth duality functor, i.e. it sends any $\pi \in \mathrm{Rep}_k^{\mathrm{sm}}(G)$ to the smooth vectors in $\mathrm{Hom}_k(\pi,k)$.*

Proof. Easy, using the previous lemma together with the fact that $\mathbf{1}$ is an injective object of $\mathrm{Rep}_k^{\mathrm{sm}}(H)$ for any pro-p open compact $H \subset G$. (Use that any injective map $V \to W$ in $\mathrm{Rep}_k^{\mathrm{sm}}(H)$ admits an H-equivariant retraction, by averaging any k-linear retraction against a k-valued Haar measure.) □

When G is a p-adic Lie group and $k = \mathbf{F}_p$, the naive smooth duality functor on $\mathrm{Rep}_{\mathbf{F}_p}^{\mathrm{sm}}(G)$ is rather useless, because the functor of passage to smooth vectors is no longer exact. To remedy this, Kohlhaase has introduced a family of derived smooth duality functors $\mathcal{S}_G^i(-)$ on $\mathrm{Rep}_{\mathbf{F}_p}^{\mathrm{sm}}(G)$, which are essentially the right-derived functors of the functor $V \mapsto (V^*)^{\mathrm{sm}}$. More precisely, Kohlhaase defines

$$\mathcal{S}_G^i(V) = \mathrm{colim}_{H \subset G\,\mathrm{open\,compact}} \mathrm{Ext}^i_{\mathbf{F}_p[[H]]}(k, V^*).$$

Here $V^* = \mathrm{Hom}_{\mathbf{F}_p}(V,\mathbf{F}_p)$ denotes the Pontryagin dual of V with its natural pseudocompact $\mathbf{F}_p$-module structure; this is naturally a pseudocompact $\mathbf{F}_p[[H]]$-module for any open compact $H \subset G$. We prefer to package the $\mathcal{S}^i$'s into a single functor on the derived category. In other words, we consider the endofunctor of $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$ defined by

$$V \mapsto \mathcal{S}_G(V) = \mathrm{colim}_{H \subset G\,\mathrm{open\,compact}} R\mathrm{Hom}_{\mathbf{F}_p[[H]]}(\mathbf{F}_p, V^*),$$

where $V^* \in \mathcal{D}(\mathbf{F}_p[[H]] - \mathrm{mod})$ is computed termwise. Here $R\mathrm{Hom}_{\mathbf{F}_p[[H]]}$ denotes $R\mathrm{Hom}$ in the derived category of $\mathbf{F}_p[[H]]$-modules. Of course, $H^i(\mathcal{S}(V)) = \mathcal{S}^i(V)$ when V is concentrated in degree zero.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

The following proposition was noted independently by Schneider-Sorensen [SS22].

Proposition 3.6. *As endofunctors of $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$, the functor $\underline{\mathrm{Hom}}_G(-,\mathbf{1})$ coincides with Kohlhaase's derived smooth duality functor $\mathcal{S}_G(-)$.*

Proof. As noted above, $\mathcal{S}_G(V)$ is defined by the formula

$$\mathcal{S}_G(V) = \mathrm{colim}_{H\subset G\,\mathrm{open\,compact}} R\mathrm{Hom}_{\mathbf{F}_p[[H]]}(\mathbf{F}_p, V^*).$$

To actually compute this, choose a uniformly powerful open pro-p subgroup $G_0 \subset G$. By work of Lazard [Laz65, Ch. V, Section 2], we can choose a finite length resolution $P_\bullet \to \mathbf{F}_p \to 0$ of $\mathbf{F}_p$ by finite free $\mathbf{F}_p[[G_0]]$-modules. Then we can express $\mathcal{S}_G(-)$ more concretely as

$$\mathcal{S}_G(V) = \mathrm{colim}_{H\subset G_0\,\mathrm{open}} \mathrm{Hom}_{\mathbf{F}_p[[H]]}(P_\bullet, V^*).$$

Set $I^n = \mathrm{Hom}^{\mathrm{cts}}_{\mathbf{F}_p}(P_n,\mathbf{F}_p)$, so $0 \to \mathbf{F}_p \to I^\bullet$ is a finite resolution of $\mathbf{1}$ by injective objects in $\mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G_0)$. Since Pontryagin duality gives an exact anti-equivalence between smooth H-representations over $\mathbf{F}_p$ and pseudocompact $\mathbf{F}_p[[H]]$-modules, we have natural identifications

$$\mathrm{Hom}_{\mathbf{F}_p[[H]]}(P_\bullet, V^*) \cong \mathrm{Hom}_{\mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(H)}(V, I^\bullet) \cong R\mathrm{Hom}_{\mathcal{D}(H)}(V,\mathbf{1}).$$

Putting things together, we get

$$\begin{aligned}\mathcal{S}_G(V) &\cong \mathrm{colim}_{H\subset G_0\,\mathrm{open}} \mathrm{Hom}_{\mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(H)}(V, I^\bullet)\\ &\cong \mathrm{colim}_{H\subset G_0\,\mathrm{open}} R\mathrm{Hom}_{\mathcal{D}(H)}(V,\mathbf{1})\\ &\cong \underline{\mathrm{Hom}}_G(V,\mathbf{1}),\end{aligned}$$

where the last line follows by Lemma 3.4. □

In the remainder of this subsection, we continue to assume that G is a p-adic Lie group. Recall that the *Gelfand-Kirillov (GK) dimension* of an admissible smooth representation $V \in \mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G)$ is the quantity

$$\dim_G V = \dim G \cdot \lim_{n\to\infty} \frac{\log \dim_{\mathbf{F}_p} V^{K_n}}{\log [K_n : K_0]} = \lim_{n\to\infty} \frac{\log \dim_{\mathbf{F}_p} V^{K_n}}{n \log p}.$$

Here $K_0 \subset G$ is a fixed choice of a uniformly powerful open pro-p subgroup, and $K_n = K^{p^n} \subset K_0$ is the closed normal subgroup generated by the p^nth powers of elements of K. (One automatically has that $[K_n : K_0] = p^{n\cdot \dim G}$, which explains the second equality above.) By [EP20, Proposition 2.18], this limit exists independently of the choice of K_0, and defines an integer in the interval $[0, \dim G]$. Moreover, $\dim_G V = 0$ iff $\dim_{\mathbf{F}_p} V < \infty$. Additionally, if $0 \to V' \to V \to V'' \to 0$ is an exact sequence of admissible smooth representations, then $\dim_G V = \max(\dim_G V', \dim_G V'')$. In particular, any subquotient W of a given V satisfies $\dim_G W \leq \dim_G V$.

A key point for us is that $\dim_G$ is closely related to the behavior of the functors $\mathcal{S}^i_G$.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proposition 3.7.

(i) *If V is an admissible smooth G-representation, then each $\mathcal{S}_G^i(V)$ is admissible smooth, with $\dim_G \mathcal{S}_G^i(V) \leq i$ for all i. Moreover, if $d = \dim_G V$ then $\dim_G \mathcal{S}_G^d(V) = d$ and $\mathcal{S}_G^i(V) = 0$ for all $i > d$.*

(ii) *For any $A \in \mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p)$ whose cohomologies $H^i(A)$ are all admissible, the complex $\mathcal{S}_G(A)$ has admissible cohomologies, and the biduality map $A \to (\mathcal{S}_G \circ \mathcal{S}_G)(A)$ is an isomorphism.*

(iii) *If V is an admissible smooth representation with $\dim_G V = 0$, so $\dim_{\mathbf{F}_p} V < \infty$, then $\mathcal{S}(V) \cong \mathcal{S}^0(V)$ is the naive linear dual of V.*

Proof. The first part is essentially a summary of results from [Koh17]. The second part is Pontryagin-dual to the biduality discussed in the proof of [Koh17, Theorem 3.5]; see also [SS22, Proposition 3.4]. The third part is immediate from the first part. □

Note that part (i) here implies that $\dim_G V = \sup\{i \mid \mathcal{S}_G^i(V) \neq 0\}$. We also observe that $\mathcal{S}^0(-)$ can be completely understood.

Proposition 3.8. *Let V be any admissible smooth G-representation. Then $\mathcal{S}_G^0(V) \neq 0$ if and only if V admits a nonzero G-stable quotient $V \twoheadrightarrow W$ with $\dim_{\mathbf{F}_p} W < \infty$.*

Proof. 'If': Given such a $V \to W$, we get a nonzero map $\mathcal{S}(W) \to \mathcal{S}(V)$. But the previous proposition shows that $\mathcal{S}(W) = \tau^{\leq 0}\mathcal{S}(W) = \mathcal{S}^0(W)$, so we get a nonzero map $\mathcal{S}^0(W) \to \mathcal{S}(V)$ which necessarily factors over a nonzero map $\mathcal{S}^0(W) \to \mathcal{S}^0(V)$.

'Only if': Suppose $\mathcal{S}^0(V) \neq 0$. By the previous proposition, this is zero-dimensional as a G-representation, so $\mathcal{S}(\mathcal{S}^0(V)) \cong \mathcal{S}^0(\mathcal{S}^0(V))$. In particular, applying $\mathcal{S}$ to the map $\mathcal{S}^0(V) = \tau^{\leq 0}\mathcal{S}(V) \to \mathcal{S}(V)$ and using biduality, we get a nonzero map $V \cong \mathcal{S}(\mathcal{S}(V)) \to \mathcal{S}(\mathcal{S}^0(V)) \cong \mathcal{S}^0(\mathcal{S}^0(V))$, and $\mathcal{S}^0(\mathcal{S}^0(V))$ is finite-dimensional over $\mathbf{F}_p$. □

We conclude this section by giving a geometric interpretation of $\mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p)$. Let C be any algebraically closed nonarchimedean field of residue characteristic p.

Proposition 3.9. *Notation as above, there are natural symmetric monoidal equivalences*

$$\mathcal{D}_{\text{ét}}([\mathrm{Spd}(C)/\underline{G}], \mathbf{F}_p) \cong \mathcal{D}([\mathrm{Spd}(C)/\underline{G}]_{\text{ét}}, \mathbf{F}_p) \cong \mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p).$$

In particular, the first two categories are left-complete and canonically independent of C.

Proof. This follows from the proof of Theorem V.1.1 in [FS21], noting that that $[\mathrm{Spd}(C)/\underline{G}]_{\text{ét}}$ locally has uniformly bounded p-cohomological dimension by Lazard-Serre [Laz65, Ch. V, Section 2]. □

3.2. Sheaves on classifying stacks

We now study the stacky 6-functor formalism introduced above in the context of the classifying stack $*/G$ associated to a locally profinite group G. Recall that if G has locally finite p-cohomological dimension then for every totally disconnected space $X = \mathrm{Spa}(A, A^+)$ with pseudouniformizer π the ∞-category $\mathcal{D}_{\square}^a(\mathcal{O}_{X/G}^+/\pi)$ can be identified with the ∞-category of smooth G-representations on solid almost A^+/π-modules (see

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

[Man22, Lemma 3.4.26.(ii)]). In the following we will show that if G has locally finite p-cohomological dimension then $*/G \to *$ is p-fine, and we will deduce a precise criterion for when the map $*/G \to *$ is p-cohomologically smooth. In particular, this will hold in all practical situations, see Theorem 3.18 below.

To start things off, we first translate the finite p-cohomological dimension condition into a vanishing of Ext-groups. Also recall the definition of the p-cohomological dimension $\mathrm{cd}_p G$ from [Man22, Definition 3.4.20] (see [Man22, Proposition 3.4.22] to relate this to more classical versions).

Proposition 3.10. *Let G be a profinite group and let $M,N \in \mathcal{D}_\square(\mathbf{F}_p)^{BG}$ be static continuous G-representations on solid $\mathbf{F}_p$-vector spaces. If M is discrete then*

$$\mathrm{Ext}^i_G(M,N) = 0 \qquad \text{for } i > \mathrm{cd}_p G.$$

Proof. The ∞-category $\mathcal{D}_\square(\mathbf{F}_p)^{BG}$ comes naturally equipped with a symmetric monoidal structure such that the forgetful functor $q^*\colon \mathcal{D}_\square(\mathbf{F}_p)^{BG} \to \mathcal{D}_\square(\mathbf{F}_p)$ is symmetric monoidal (e.g. one can construct $\mathcal{D}_\square(\mathbf{F}_p)^{BG}$ as certain sheaves on the classifying stack $*/G$ in a suitable geometric setup). The forgetful functor q^* has a left adjoint $q_\natural\colon \mathcal{D}_\square(\mathbf{F}_p) \to \mathcal{D}_\square(\mathbf{F}_p)^{BG}$ which is given by $q_\natural(M) = \mathbf{F}_{p\square}[G] \otimes_{\mathbf{F}_{p\square}} M$. It is straightforward to check that $q_\natural$ satisfies the projection formula, i.e. that for all $M \in \mathcal{D}_\square(\mathbf{F}_p)^{BG}$ and $N \in \mathcal{D}_\square(\mathbf{F}_p)$ the natural morphism $q_\natural(N \otimes q^*M) \xrightarrow{\sim} q_\natural N \otimes M$ is an isomorphism. The symmetric monoidal structure on $\mathcal{D}_\square(\mathbf{F}_p)^{BG}$ is closed, i.e. admits internal hom's $\underline{\mathrm{Hom}}$ (this is clear if we restrict to κ-condensed objects for some cardinal κ and then follows for unrestricted condensed objects by the following computation). It follows formally from the projection formula for $q_\natural$ that for all $M,N \in \mathcal{D}_\square(\mathbf{F}_p)^{BG}$ the natural morphism

$$q^*\underline{\mathrm{Hom}}(M,N) \xrightarrow{\sim} \underline{\mathrm{Hom}}(q^*M, q^*N)$$

is an isomorphism – in other words, the internal hom in $\mathcal{D}_\square(\mathbf{F}_p)^{BG}$ is computed on the underlying solid $\mathbf{F}_p$-vector spaces. In particular, if M and N are static and M is discrete (and thus $q^*M \cong \bigoplus_I \mathbf{F}_p$ for some set I) then $\underline{\mathrm{Hom}}(M,N)$ is static. Therefore, from

$$\mathrm{Ext}^i_G(M,N) = \pi_0 R\mathrm{Hom}(M, N[i]) = \pi_0 \mathrm{Hom}(\mathbf{F}_p, \underline{\mathrm{Hom}}(M,N[i])) = H^i(G, \underline{\mathrm{Hom}}(M,N))$$

we immediately deduce the claim. □

We will now show that $*/G \to *$ is p-fine. This reduces to showing that the canonical cover $* \to */G$ admits universal p-codescent, which will be established by adapting Mathew's ideas on descendable algebras [Mat16].

Lemma 3.11. *Let G be a profinite group with $\mathrm{cd}_p G < \infty$. Then the map $* \to */G$ of small v-stacks is bdcs and admits universal p-codescent.*

Proof. It is clear that the map $* \to */G$ is bdcs (e.g. the pullback along itself is $G \to *$, which is pro-étale). Now pick any $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ admitting a map to a strictly totally disconnected space, pick any map $X \to */G$ and denote $q\colon Y \to X$ the base-change of $* \to */G$. We need to show that q admits p-codescent, i.e. if $q_\bullet\colon Y_\bullet \to X$ denotes the Čech nerve of q then the natural functor

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

$$\mathcal{D}^{a!}_{\square}(X,\Lambda) \xrightarrow{\sim} \varprojlim_{n\in\Delta} \mathcal{D}^{a!}_{\square}(Y_n,\Lambda)$$

is an equivalence. To see this we employ Lurie's Beck-Chevalley condition, see [Lur, Corollary 4.7.5.3]. Part (2) of this condition is easily satisfied: Since all the maps in the diagram $Y_\bullet \to X$ are proper, the required left adjointability for the upper shriek functors reduces to proper base-change (by passing to left adjoints). Thus in order to apply Lurie's result it remains to verify the following:

(a) The functor $q^!\colon \mathcal{D}^a_{\square}(X,\Lambda) \to \mathcal{D}^a_{\square}(Y,\Lambda)$ is conservative.

(b) The functor $q^!$ preserves geometric realizations of $q^!$-split simplicial objects in $\mathcal{D}^a_{\square}(X,\Lambda)$.

To prove (a) and (b) we apply ideas by Mathew [Mat16] (which were also used extensively in the setup of the p-adic 6-functor formalism, cf. [Man22, §2.6]). More concretely, we will prove the following claim:

(∗) Let $\langle q_*q^!\rangle \subset \mathrm{Fun}(\mathcal{D}^a_{\square}(X,\Lambda),\mathcal{D}^a_{\square}(X,\Lambda))$ be the full subcategory generated by the endofunctor $q_*q^!$ under compositions, retracts and finite (co)limits. Then $\langle q_*q^!\rangle$ contains the identity functor.

Let us first see how (∗) implies (a) and (b), so assume that (∗) holds for the moment. To prove (a), let $\mathcal{M} \to \mathcal{N}$ be a given map in $\mathcal{D}^a_{\square}(X,\Lambda)$ such that $q^!\mathcal{M} \xrightarrow{\sim} q^!\mathcal{N}$ is an isomorphism. Then also $q_*q^!\mathcal{M} \xrightarrow{\sim} q_*q^!\mathcal{N}$ is an isomorphism and one deduces easily that $F(\mathcal{M}) \xrightarrow{\sim} F(\mathcal{N})$ is an isomorphism for all $F \in \langle q_*q^!\rangle$; taking $F = \mathrm{id}$ shows that $\mathcal{M} \xrightarrow{\sim} \mathcal{N}$ is an isomorphism, as desired. To prove (b), let $\mathcal{M}_\bullet$ be a $q^!$-split simplicial object in $\mathcal{D}^a_{\square}(X,\Lambda)$. Then $q_*q^!\mathcal{M}_\bullet$ is a constant Ind-object (i.e. lies in the essential image of $\mathcal{D}^a_{\square}(X,\Lambda) \hookrightarrow \mathrm{Ind}(\mathcal{D}^a_{\square}(X,\Lambda))$) and it follows immediately that the same is true for $F(\mathcal{M}_\bullet)$ for all $F \in \langle q_*q^!\rangle$. Taking $F = \mathrm{id}$ shows that $\mathcal{M}_\bullet$ is a constant Ind-object, so that its colimit commutes with any exact functor; in particular it commutes with $q^!$.

It now remains to verify (∗). We compute (using [Man22, Corollary 3.6.17. (i)])

$$q_*q^! = q_*\underline{\mathrm{Hom}}(\Lambda^a,q^!) = \underline{\mathrm{Hom}}(q_!\Lambda^a,-) = \underline{\mathrm{Hom}}(q_*\Lambda^a,-).$$

Therefore (∗) can be reduced to showing the following: Let $\langle q_*\Lambda^a\rangle \subset \mathcal{D}^a_{\square}(X,\Lambda)$ be the full subcategory generated by $q_*\Lambda^a$ under tensor product, retracts and finite (co)limits; then $\langle q_*\Lambda^a\rangle$ contains the tensor unit Λ^a. By [Mat16, Proposition 3.20] this is equivalent to the following claim:

(∗′) The algebra object $q_*\Lambda^a \in \mathcal{D}^a_{\square}(X,\Lambda)$ admits descent.

It remains to prove that (∗′) is true. We first perform some reductions. By assumption on (X,Λ) there is a map $(X,\Lambda) \to (Z,\Lambda)$ with $Z = \mathrm{Spa}(A,A^+)$ being strictly totally disconnected. We thus get an induced map $X \to Z \times */G = Z/G$ such that q is the pullback of the projection $q_Z\colon Z \to Z/G$. The induced pullback functor $\mathcal{D}^a_{\square}(Z/G,\Lambda) \to \mathcal{D}^a_{\square}(X,\Lambda)$ is exact, symmetric monoidal and sends $q_{Z*}\Lambda^a$ to $q_*\Lambda^a$ (by proper base-change). It is therefore enough to prove that (∗') holds for q_Z in place of q, so from now on we assume $X = Z/G$ and hence $Y = Z$.

Given a morphism $\Lambda' \to \Lambda$ of integral torsion coefficients on Z, the induced base-change functor $\mathcal{D}^a_{\square}(Z,\Lambda') \to \mathcal{D}^a_{\square}(Z,\Lambda)$ preserves tensor products, retracts and finite (co)limits

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

and sends $q_*\Lambda'^a$ to $q_*\Lambda^a$ (by [Man22, Proposition 3.6.16]), hence if $(*')$ holds for Λ' then it also holds for Λ. By definition of integral torsion coefficients we can thus reduce to the case that $\Lambda = \mathcal{O}^{+a}_{Z^\sharp}/\pi$ for some untilt $Z^\sharp$ of Z and some pseudouniformizer π on $Z^\sharp$. Pick some pseudouniformizer ϖ on $Z^\sharp$ such that $\varpi|p$ and some integer $n > 0$ such that $\pi|\varpi^n$. By [Mat16, Proposition 3.35] the map of discrete rings $A^{\sharp+}/\pi \to A^{\sharp+}/\varpi$ admits descent, which after pullback along $Z/G \to Z$ implies that the map $\mathcal{O}^{+a}_{Z^\sharp/G}/\pi \to \mathcal{O}^{+a}_{Z^\sharp/G}/\varpi$ in $\mathcal{D}^a_\square(\mathcal{O}^+_{Z^\sharp/G}/\pi)$ admits descent. Therefore, if we can show that $(*')$ holds for $\Lambda = \mathcal{O}^+_{Z^\sharp/G}/\varpi$ then by [Mat16, Proposition 3.24] and [BS17, Lemma 11.10] it also follows for $\Lambda = \mathcal{O}^+_{Z^\sharp/G}/\pi$. Thus from now on we can assume that $\pi = \varpi$, i.e. we now have $\pi|p$.

Now recall the Riemann-Hilbert functor $RH\colon \mathcal{D}_{\text{ét}}(Z/G,\mathbf{F}_p)^{\text{oc}} \to \mathcal{D}^a_\square(\mathcal{O}^+_{Z^\sharp/G}/\pi)$ from [Man22, Definition 3.9.21. (a)] (where we ignore the φ-module structure). Then RH is symmetric monoidal and exact and satisfies $RH(q_*\mathbf{F}_p) = q_*(\mathcal{O}^{+a}_{Z^\sharp}/\pi)$. To verify the latter, note that this identity can be checked after applying q^*, where we get (using base-change)

$$\begin{aligned} q^*RH(q_*\mathbf{F}_p) = RH(q^*q_*\mathbf{F}_p) = RH(\mathcal{C}(G,\mathbf{F}_p)) = \mathcal{C}(G,\mathbf{F}_p)\otimes_{\mathbf{F}_p} A^{\sharp+,a}/\pi = \mathcal{C}(G,A^{\sharp+,a}/\pi) \\ = q^*q_*(A^{\sharp+,a}/\pi) \end{aligned}$$

(here we implicitly use the identity $\mathcal{D}^a_\square(\mathcal{O}^+_{Z^\sharp}/\pi) = \mathcal{D}^a_\square(A^{\sharp+}/\pi)$). All in all this means that it is enough to show that $q_*\mathbf{F}_p \in \mathcal{D}_{\text{ét}}(Z/G,\mathbf{F}_p)^{\text{oc}}$ is descendable. It follows easily from [Man22, Lemma 3.9.18.(iii)] that $\mathcal{D}_{\text{ét}}(Z/G,\mathbf{F}_p)^{\text{oc}} = \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p(Z))_\omega$ is the derived ∞-category of discrete smooth G-representations on $\mathbf{F}_p(Z)$-modules (use the same argument as in the proof of [Man22, Lemma 3.4.26.(iii)]). Our goal is to show that the G-representation $\mathcal{C}(G,\mathbf{F}_p(Z))$ admits descent in this ∞-category. By considering base-change along $\mathbf{F}_p \to \mathbf{F}_p(Z)$ we reduce the claim to the following:

$(*'')$ The algebra object $\mathcal{C}(G,\mathbf{F}_p) \in \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)_\omega$ admits descent.

Note that $\mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)_\omega$ is equivalent to the full subcategory of $\mathcal{D}_\square(\mathbf{F}_p)^{BG}$ spanned by the discrete G-representations (e.g. both ∞-categories can be constructed using the same descent diagram), hence by Proposition 3.10 it follows that for all static $M,N \in \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)_\omega$ we have $\pi_0\operatorname{Hom}(M,N[i]) = 0$ for $i > d := \operatorname{cd}_p G$. Now consider the triangle

$$\mathbf{F}_p \to \operatorname{Tot}_d(\mathcal{C}(G,\mathbf{F}_p)^{\otimes\bullet+1}) \to X$$

in $\mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)_\omega$, where Tot_d denotes the d-truncated totalization of a cosimplicial object (i.e. the limit over $\Delta_{\le d}$). We know that $\operatorname{Tot}(\mathcal{C}(G,\mathbf{F}_p)^{\otimes\bullet+1}) = \mathbf{F}_p$ (by descent) and by the dual version of [Lur, Proposition 1.2.4.5] we have $\pi_i\operatorname{Tot}_d = \pi_i\operatorname{Tot}$ for $i > -d$. It follows that X is concentrated in homological degrees $\le -d$, so that by the above vanishing of hom's the connecting map $X \to \mathbf{F}_p[1]$ must be zero. It follows that $\operatorname{Tot}_d(\mathcal{C}(G,\mathbf{F}_p)^{\otimes\bullet+1}) \cong \mathbf{F}_p \oplus X$, which immediately implies that $\mathcal{C}(G,\mathbf{F}_p)$ admits descent. □

Corollary 3.12. *Let G be a profinite group with $\operatorname{cd}_p G < \infty$. Then:*

(i) *The map $*/G \to *$ is p-fine.*

(ii) *For every $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ with associated map $f\colon X/G \to X$ there is a natural isomorphism $f_! = f_*$ of functors $\mathcal{D}^a_\square(X/G,\Lambda) \to \mathcal{D}^a_\square(X,\Lambda)$.*

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proof. Part (i) follows from Lemma 3.11 by using the factorization $\mathrm{id} = [* \to */G \to *]$. To prove (ii) let (X,Λ) be given and assume first that (X,Λ) admits a map to some strictly totally disconnected space. Let $q\colon X \to X/G$ denote the projection and let $q_\bullet\colon X_\bullet \to X/G$ be the Čech nerve of q. By Lemma 3.11 we have $\mathcal{D}^{a!}_\square(X/G,\Lambda) = \varprojlim_{n\in\Delta} \mathcal{D}^{a!}_\square(X_n,\Lambda)$, which means that for all $\mathcal{M} \in \mathcal{D}^a_\square(X/G,\Lambda)$ the natural map $\varinjlim_{n\in\Delta} q_{n!}q_n^!\mathcal{M} \xrightarrow{\sim} \mathcal{M}$ is an isomorphism. Since both f_* and $f_!$ preserve small colimits (for f_* this follows from $\mathrm{cd}_p G < \infty$) and all q_n and $f\circ q_n$ are proper we deduce

$$f_* = f_* \varinjlim_n q_{n!}q_n^! = \varinjlim_n (f\circ q_n)_* q_n^! = \varinjlim_n (f\circ q_n)_! q_n^! = f_!.$$

The case of general (X,Λ) can be deduced via descent (more precisely, do the above computation on sections of the cocartesian fibrations over Δ). □

Having established that $*/G \to *$ is p-fine in many cases, we now determine criteria for when this map is p-cohomologically smooth. When G has finite p-cohomological dimension, the precise answer is captured by the following definition.

Definition 3.13. Let G be a profinite group with $\mathrm{cd}_p G < \infty$. We say that G is *p-Poincaré of dimension d* if it satisfies the following conditions:

(i) $H^i(G,M)$ is finite for all $i \geq 0$ and all G-representations M on finite $\mathbf{F}_p$-vector spaces.

(ii) The solid $\mathbf{F}_p$-vector space $R\Gamma(G,\mathbf{F}_{p\square}[G])$ is invertible and concentrated in cohomological degree d (i.e. it is isomorphic to $\mathbf{F}_p[-d]$).

Note that condition (i) is equivalent to the finiteness of $H^i(H,\mathbf{F}_p)$ for all $i \geq 0$ and all open subgroups $H \subseteq G$. Note also that if G is p-Poincaré of dimension d, then so is every open subgroup of G.

Proposition 3.14. *Let G be a profinite group with $\mathrm{cd}_p G < \infty$.*

(i) *G satisfies condition (i) of Definition 3.13 if and only if for all strictly totally disconnected spaces $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ with associated map $f\colon X/G \to X$ the morphism $f^* \otimes f^!\Lambda^a \xrightarrow{\sim} f^!$ is an equivalence of functors $\mathcal{D}^a_\square(X,\Lambda) \to \mathcal{D}^a_\square(X/G,\Lambda)$.*

(ii) *G is a p-Poincaré group (of dimension d) if and only if the map $*/G \to *$ is p-cohomologically smooth (of pure dimension $d/2$).*

Proof. Fix a strictly totally disconnected space $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ with associated map $f\colon X/G \to X$. Then $\mathcal{D}^a_\square(X,\Lambda) = \mathcal{D}^a_\square(\Lambda(X))$ is the ∞-category of solid almost $\Lambda(X)$-modules, while $\mathcal{D}^a_\square(X/G,\Lambda)$ is the ∞-category of smooth G-representations on such modules (see [Man22, Lemma 3.4.26]). Moreover, f_* computes smooth G-cohomology and f^* associates to every $M \in \mathcal{D}^a_\square(\Lambda(X))$ the trivial G-representation on M.

We first prove (i). By [Man22, Lemma 3.4.19] the collection of functors $R\Gamma_{\mathrm{sm}}(H,-)\colon \mathcal{D}^a_\square(X/G,\Lambda) \to \mathcal{D}^a_\square(X,\Lambda)$ for open compact subgroups H is conservative. Therefore the natural map $f^* \otimes f^!\Lambda^a \to f^!$ is an equivalence if and only if for all compact open subgroups $H \subset G$ the map $R\Gamma_{\mathrm{sm}}(H, f^* \otimes f^!\Lambda^a) \to R\Gamma_{\mathrm{sm}}(H, f^!)$ is an equivalence of endofunctors of

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

$\mathcal{D}^a_\square(X,\Lambda)$. Let us denote $f_H\colon */H \to *$ the projection, which factors over the finite étale projection $*/H \to */G$. Then we see that $f^* \otimes f^!\Lambda^a \to f^!$ is an equivalence if and only if for all compact open subgroups $H \subset G$ the map

$$f_{H*}(f_H^* \otimes f_H^!\Lambda^a) \to f_{H*}f_H^!$$

is an equivalence. By Corollary 3.12 we have $f_{H*} = f_{H!}$, so in particular f_{H*} satisfies the projection formula. We compute

$$f_{H*}(f_H^* \otimes f_H^!\Lambda^a) = - \otimes f_{H*}f_H^!\Lambda^a = - \otimes \underline{\mathrm{Hom}}(f_{H*}\Lambda^a,\Lambda^a),$$
$$f_{H*}f_H^! = f_{H*}\underline{\mathrm{Hom}}(\Lambda^a, f_H^!) = \underline{\mathrm{Hom}}(f_{H!}\Lambda^a,-) = \underline{\mathrm{Hom}}(f_{H*}\Lambda^a,-).$$

It follows that $f^* \otimes f^!\Lambda^a \to f^!$ is an equivalence if and only if $R\Gamma_{\mathrm{sm}}(H,\Lambda(X)^a) = f_{H*}\Lambda^a$ is dualizable in $\mathcal{D}^a_\square(X,\Lambda)$ for all open compact $H \subset G$. Now assume that G satisfies condition (i) of Definition 3.13. Then $R\Gamma_{\mathrm{sm}}(H,\mathbf{F}_p)$ is dualizable in $\mathcal{D}_\square(\mathbf{F}_p)$ for all H. If $p = 0$ in $\Lambda(X)$ then the same follows for $R\Gamma_{\mathrm{sm}}(H,\Lambda(X)^a) = R\Gamma_{\mathrm{sm}}(H,\mathbf{F}_p) \otimes_{\mathbf{F}_p} \Lambda(X)^a$. For general Λ we can argue as in the proof of [Man22, Lemma 3.8.4. (i)]. Conversely, assume that $f^* \otimes f^!\Lambda^a \xrightarrow{\sim} f^!$ is an isomorphism for all (X,Λ). Pick $X = \mathrm{Spa}C$ for some algebraically closed field over $\mathbf{Q}_p$ and $\Lambda = \mathcal{O}_C/p$. Then the Riemann-Hilbert correspondence from [Man22, Theorem 3.9.23] implies that the base-change functor $RH\colon \mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) \hookrightarrow \mathcal{D}^{\mathrm{sm}}(G,\mathcal{O}_C^a/p)^\varphi$ is fully faithful and induces an equivalence of dualizable objects. On the other hand, this base-change functor sends $R\Gamma_{\mathrm{sm}}(H,\mathbf{F}_p)$ to $R\Gamma_{\mathrm{sm}}(H,\mathcal{O}_C^a/p)$, hence it follows from the above that $R\Gamma_{\mathrm{sm}}(H,\mathbf{F}_p)$ is dualizable for all open compact subgroups $H \subset G$. But every G-representation M on a finite static $\mathbf{F}_p$-vector space is a finite colimit of the generating representations $\mathbf{F}_p[G/H]$, hence also $R\Gamma_{\mathrm{sm}}(G,M)$ is dualizable. In particular all $H^i(G,M)$ are finite (here we implicitly use $R\Gamma(G,M) = R\Gamma_{\mathrm{sm}}(G,M)$ since M is discrete).

We now prove (ii), so from now on we can assume that G satisfies the equivalent conditions of (i). We then need to see that G is p-Poincaré if and only if for every strictly totally disconnected $(X,\Lambda) \in \mathrm{vStack}_\Lambda$ with $\Lambda = \mathcal{O}^+_{X^\flat}/\pi$ for some pseudouniformizer π on $X^\flat$ the object $f^!\Lambda^a \in \mathcal{D}^a_\square(X/G,\Lambda)$ is invertible. Given (X,Λ), let us denote $q\colon X \to X/G$ the projection. Since invertibility satisfies descent, $f^!\Lambda^a$ is invertible if and only if $q^*f^!\Lambda^a$ is invertible. For every compact open subgroup $H \subset G$ we denote $f_H\colon */H \to *$ and $q_H\colon */H \to */G$ the natural maps and we compute (using [Man22, Proposition 3.3.9.(iii)] and Corollary 3.12)

$$q^*f^!\Lambda^a = f_*q_*q^*f^!\Lambda^a = \varinjlim_H f_*q_{H*}q_H^*f^!\Lambda^a = \varinjlim_H f_{H*}f_H^!\Lambda^a = \varinjlim_H \underline{\mathrm{Hom}}(f_{H*}\Lambda^a,\Lambda^a).$$

By the proof of (i) we know that $f_{H*}\Lambda^a = R\Gamma(H,\Lambda(X)^a)$ is dualizable in $\mathcal{D}^a_\square(X,\Lambda)$ and similarly $R\Gamma(H,\mathbf{F}_p)$ is dualizable in $\mathcal{D}_\square(\mathbf{F}_p)$. It follows that

$$\underline{\mathrm{Hom}}(f_{H*}\Lambda^a,\Lambda^a) = \underline{\mathrm{Hom}}(\Gamma(H,\mathbf{F}_p),\mathbf{F}_p) \otimes_{\mathbf{F}_p} \Lambda(X)^a.$$

Therefore we obtain

$$q^*f^!\Lambda^a = \varinjlim_H \underline{\mathrm{Hom}}(f_{H*}\Lambda^a,\Lambda^a) = (\varinjlim_H \underline{\mathrm{Hom}}(\Gamma(H,\mathbf{F}_p),\mathbf{F}_p)) \otimes_{\mathbf{F}_p} \Lambda(X)^a$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

We have $R\Gamma(H,\mathbf{F}_p) = \bigoplus_i H^i(H,\mathbf{F}_p)[-i]$ and since there are no higher Ext-groups between discrete $\mathbf{F}_p$-vector spaces we even have a decomposition of diagrams $(R\Gamma(H,\mathbf{F}_p)_H = \bigoplus_i (H^i(H,\mathbf{F}_p))_H$. On the other hand, we know that if $(M_i)_i$ is any cofiltered system of finite discrete $\mathbf{F}_p$-vector spaces then $\varprojlim_i M_i$ is concentrated in degree 0 and $\underline{\mathrm{Hom}}(\varprojlim_i M_i,\mathbf{F}_p) = \varinjlim_i \underline{\mathrm{Hom}}(M_i,\mathbf{F}_p)$. Altogether this implies

$$\varinjlim_H \underline{\mathrm{Hom}}(R\Gamma(H,\mathbf{F}_p),\mathbf{F}_p) = \underline{\mathrm{Hom}}(\varprojlim_H R\Gamma(H,\mathbf{F}_p),\mathbf{F}_p) = \underline{\mathrm{Hom}}(\varprojlim_H \Gamma(G,\mathbf{F}_p[G/H]),\mathbf{F}_p) =$$
$$= \underline{\mathrm{Hom}}(R\Gamma(G,\mathbf{F}_{p\square}[G]),\mathbf{F}_p).$$

We deduce:

$$q^* f^! \Lambda^a = \underline{\mathrm{Hom}}(R\Gamma(G,\mathbf{F}_{p\square}[G]),\mathbf{F}_p) \otimes_{\mathbf{F}_p} \Lambda(X)^a.$$

In particular, if G is p-Poincaré then $q^* f^! \Lambda^a$ is invertible. Conversely, assume that $*/G \to *$ is p-cohomologically smooth. Choosing $X = \mathrm{Spa}\, C$ for some algebraically closed field C, the Riemann-Hilbert correspondence (see [Man22, Theorem 3.9.23]) tells us that the invertibility of $q^* f^! \Lambda^a$ implies that $\underline{\mathrm{Hom}}(R\Gamma(G,\mathbf{F}_{p\square}[G]),\mathbf{F}_p)$ is invertible. But then also $R\Gamma(G,\mathbf{F}_{p\square}[G])$ is invertible and thus G is p-Poincaré. □

So far we have only dealt with profinite groups of bounded p-cohomological dimension. However, the above results generalize in a straightforward way to *locally* profinite groups satisfying a much weaker cohomological dimension condition. The precise definition is as follows.

Definition 3.15. Let G be a locally profinite group.

(a) We say that G has *locally finite p-cohomological dimension* if there is some open subgroup $H \subseteq G$ such that $\mathrm{cd}_p H < \infty$.

(b) We say that G is *virtually p-Poincaré (of dimension d)* if there is an open compact subgroup $H \subset G$ such that $\mathrm{cd}_p H < \infty$ and H is p-Poincaré of dimension d in the sense of Definition 3.13.

Theorem 3.16. *Let G be a locally profinite group which has locally finite p-cohomological dimension. Then the map $*/G \to *$ is p-fine, and it is p-cohomologically smooth (of pure dimension $d/2$) if and only if G is virtually p-Poincaré (of dimension d).*

Proof. This follows immediately from Proposition 3.14 by observing that being p-fine and being p-cohomologically smooth can both be checked étale locally on the source (note that if $H \subseteq G$ is an open subgroup then the map $*/H \twoheadrightarrow */G$ is an étale cover). □

As a formal consequence of Theorem 3.16 we obtain the following p-adic analog of [Sch18a, Proposition 24.2]:

Corollary 3.17. *Let G be a locally profinite group of locally finite p-cohomological dimension acting on a small v-stack Y and let $f: Y \to X$ be a G-equivariant separated bdcs map of small v-stacks, where G acts trivially on X.*

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

(i) *The map* $f' \colon Y/G \to X$ *is* p*-fine.*

(ii) *If* f *is* p*-cohomologically smooth and* G *is virtually* p*-Poincaré, then* f' *is* p*-cohomologically smooth.*

Proof. Factor f' as $Y/G \to X/G \to X$. We claim that the first map, i.e. $f/G \colon Y/G \to X/G$, is bdcs: Base-changing this map along the pro-étale cover $X \twoheadrightarrow X/G$ gives f, hence the map $Y/G \to X/G$ is separated by [Sch18a, Proposition 10.11.(ii)], compactifiable by [Sch18a, Proposition 22.3.(iii)] (note that f is automatically compactifiable by [Sch18a, Proposition 22.3. (v)]), p-bounded by [Man22, Lemma 3.5.10.(ii)] and representable in locally spatial diamonds by [Sch18a, Proposition 13.4.(iv)]. Since $X/G \to X$ is p-fine by Corollary 3.12 we deduce that f' is p-fine. This proves (i).

To prove (ii), assume in addition that f is p-cohomologically smooth and G is virtually p-Poincaré. Then it follows from [Man22, Lemma 3.8.5.(ii)] that f/G is p-cohomologically smooth, and it follows from Theorem 3.16 that the map $X/G \to X$ is p-cohomologically smooth. Together this implies that f' is p-cohomologically smooth, as desired. □

We end this subsection by providing examples of virtual p-Poincaré groups. The key result for applications is the following theorem.

Theorem 3.18. *Any* p*-adic Lie group is virtually* p*-Poincaré.*

By the local nature of Definition 3.15 and standard structural results on p-adic Lie groups, this is an immediate consequence of the following theorem, which is what we shall actually prove.

Theorem 3.19. *Any torsion-free compact* p*-adic analytic group is* p*-Poincaré.*

The proof relies on deep results of Lazard [Laz65], which are in some sense the only 'non-formal' input we need. To apply Lazard's results, we need some general comparison between condensed group cohomology and classical continuous group cohomology, which takes into account a natural topological structure on the latter. To state the relevant result, fix a locally profinite group G. Let V be a continuous representation of G on a Hausdorff locally compact $\mathbf{F}_p$-module. Then $\underline{V}$ is a solid $\mathbf{F}_p$-vector space with an action by the condensed group G, i.e. an $\mathbf{F}_{p\square}[G]$-module. Thus we may form the condensed group cohomology

$$R\Gamma(G,\underline{V}) = \underline{\mathrm{Hom}}_{\mathbf{F}_{p\square}[G]}(\mathbf{F}_p,\underline{V}) \in \mathcal{D}_\square(\mathbf{F}_p)$$

as considered above.

On the other hand, the classical continuous group cohomology $R\Gamma_{\mathrm{cts}}(G,V)$ is by definition computed by the usual complex

$$0 \to V \to \mathcal{C}(G,V) \to \mathcal{C}(G^2,V) \to \cdots$$

of inhomogeneous cochains. Giving each $\mathcal{C}(G^n,V)$ the compact-open topology, this is naturally a complex of Hausdorff locally compact $\mathbf{F}_p$-modules. In particular, passing termwise to the condensation, we get a complex

$$0 \to \underline{V} \to \underline{\mathcal{C}(G,V)} \to \underline{\mathcal{C}(G^2,V)} \to \cdots$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

of solid $\mathbf{F}_p$-vector spaces. Let $\underline{R\Gamma_{\mathrm{cts}}(G,V)}$ denote the image of this complex in $\mathcal{D}_{\square}(\mathbf{F}_p)$.

Lemma 3.20. *There is a natural isomorphism*

$$R\Gamma(G,\underline{V}) \simeq \underline{R\Gamma_{\mathrm{cts}}(G,V)}$$

in $\mathcal{D}_{\square}(\mathbf{F}_p)$.

Proof of Theorem 3.19. Let G be as in the statement of the theorem. Then condition (i) in Definition 3.13 is classical, see, e.g. [SW00, Theorems 5.1.1 and 5.1.5]. For condition (ii), recall that $R\Gamma_{\mathrm{cts}}(G,\mathbf{F}_p[[G]]) \simeq \mathbf{F}_p[-\dim G]$ by [Laz65, Ch. V, Sec. 2] and [SW00, Theorem 5.1.9]. Since $\underline{\mathbf{F}_p[[G]]} = \mathbf{F}_{p\square}[G]$, the conclusion now follows from Lemma 3.20. □

Proof of Lemma 3.20. Arguing as in the proof of [Man22, Proposition 3.4.6], one sees that $R\Gamma(G,\underline{V})$ is computed as the cohomology of the complex

$$0 \to \underline{V} \to \underline{\mathrm{Hom}}(\underline{G},\underline{V}) \to \underline{\mathrm{Hom}}(\underline{G^2},\underline{V}) \to \cdots$$

where the internal homs are taken in condensed sets. By the definition of $\underline{R\Gamma_{\mathrm{cts}}(G,V)}$, it now suffices to show that $\underline{\mathrm{Hom}}(\underline{S},\underline{V})$ identifies with $\underline{\mathcal{C}(S,V)}$ functorially in locally profinite S. This is a special case of the next lemma. □

Lemma 3.21. *Let X be any locally compact Hausdorff space, and let Y be any compactly generated T1 topological space. Equip the space of continuous maps $\mathcal{C}(X,Y)$ with the compact-open topology. Then $\mathcal{C}(X,Y)$ is compactly generated T1, and $\underline{\mathcal{C}(X,Y)} = \underline{\mathrm{Hom}}(\underline{X},\underline{Y})$, where $\underline{\mathrm{Hom}}$ denotes the internal hom in condensed sets.*

Proof. The claim regarding the topology on $\mathcal{C}(X,Y)$ is easy. For the remaining claim, let T be any profinite space. We then compute that

$$\begin{aligned}\mathrm{Hom}_{\mathrm{Cond}}(\underline{T},\underline{\mathrm{Hom}}(\underline{X},\underline{Y})) &= \mathrm{Hom}_{\mathrm{Cond}}(\underline{T}\times\underline{X},\underline{Y})\\ &= \mathrm{Hom}_{\mathrm{Top}}(T\times X,Y)\\ &= \mathrm{Hom}_{\mathrm{Top}}(T,\mathcal{C}(X,Y))\\ &= \mathrm{Hom}_{\mathrm{Cond}}(\underline{T},\underline{\mathcal{C}(X,Y)})\end{aligned}$$

naturally in T,X,Y. The first equality is immediate from the universal property of $\underline{\mathrm{Hom}}$. The second equality follows from the fact that on compactly generated T1 spaces, $S \mapsto \underline{S}$ is fully faithful and commutes with finite products. The fourth equality follows similarly, using the first claim in the lemma. The third equality follows from the assumption that X is locally compact Hausdorff, which is exactly the condition needed to ensure that $\mathcal{C}(X,-)$ with its compact-open topology is right-adjoint to $-\times X$. Since T is arbitrary, the result now follows from Yoneda. □

3.3. Admissible representations

Given a locally profinite group G, an important subclass of all smooth G-representations are the admissible ones. Classically, a smooth G-representation on an $\mathbf{F}_p$-vector space V

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

is called admissible if for a basis of compact open subgroups $H \subseteq G$ the space V^H of H-invariants is finite. In the condensed derived world the correct analogue of this definition seems to be the following:

Definition 3.22. Let G be a locally profinite group which has locally finite p-cohomological dimension. A smooth G-representation $V \in \mathcal{D}_{\square}^{\mathrm{sm}}(G,\mathbf{F}_p)$ is called *admissible* if for a basis of compact open subgroups $H \subseteq G$ the smooth group cohomology $R\Gamma_{\mathrm{sm}}(H,V) \in \mathcal{D}_{\square}(\mathbf{F}_p)$ is perfect.

Remarks 3.23.

(i) An admissible G-representation $V \in \mathcal{D}_{\square}^{\mathrm{sm}}(G,\mathbf{F}_p)$ is automatically discrete because all $R\Gamma_{\mathrm{sm}}(H,V)$ are discrete by assumption and hence so is $V = \varinjlim_H R\Gamma_{\mathrm{sm}}(H,V)$ (cf. [Man22, Lemma 3.4.19]). Thus V admits a representation as a complex of classical smooth G-representations.

(ii) Suppose that G is a p-adic Lie group. Then by [SS22, Corollary 4.12] a smooth representation $V \in \mathcal{D}_{\square}^{\mathrm{sm}}(G,\mathbf{F}_p)$ is admissible if and only if it is bounded and discrete and all $\pi_n V$ are admissible smooth representations in the classical sense.

Now suppose we are additionally given an algebraically closed perfectoid field C with pseudouniformizer $\pi|p$ and we consider the classifying stack $B\underline{G} = \mathrm{Spd}(C)/\underline{G}$. Then $\mathcal{D}_{\square}^{a}(\mathcal{O}^{+}_{B\underline{G}}/\pi)^{\varphi}$ is the ∞-category of smooth G-representations on solid almost φ-modules over $\mathcal{O}_C^a/\pi$ and by the p-torsion Riemann-Hilbert correspondence this ∞-category contains the ∞-category $\mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$ of smooth G-representations as a full subcategory. We make the following definition:

Definition 3.24. Let G be a locally profinite group which has locally finite p-cohomological dimension. Let C be an algebraically closed perfectoid field with pseudouniformizer $\pi \in \mathcal{O}_C$ such that $\pi|p$. A smooth G-representation $M \in \mathcal{D}_{\square}^{\mathrm{sm},a}(G,\mathcal{O}_C/\pi)^{\varphi}$ is called *admissible* if it comes from an admissible $\mathbf{F}_p$-representation via the embedding

$$-\otimes_{\mathbf{F}_p} \mathcal{O}_C^a/\pi \colon \mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p) \hookrightarrow \mathcal{D}_{\square}^{\mathrm{sm},a}(G,\mathcal{O}_C/\pi)^{\varphi}$$

obtained from the Riemann-Hilbert correspondence on the classifying stack $B\underline{G}$.

One might also consider defining M to be admissible if the smooth group cohomology groups $R\Gamma_{\mathrm{sm}}(H,M)$ are perfect for a basis of open subgroups $H \subseteq G$. In fact, this easily turns out to be equivalent to Definition 3.24, as we will see below. If G is p-Poincaré then there is also a characterization of admissible representations in terms of their behavior under duality. To show this, we need the following characterization of perfect (equivalently, dualizable) objects in terms of their duals:

Lemma 3.25. *Let K be a perfectoid field with pseudouniformizer $\pi \in \mathcal{O}_K$. Then a φ-module $M \in \mathcal{D}_{\square}^{a}(\mathcal{O}_K/\pi)^{\varphi}$ is dualizable if and only if M is bounded, both M and $M^{\vee}$ are discrete and the natural map $M \to M^{\vee\vee}$ is an isomorphism.*

Proof. First assume that M is dualizable. Then by [Man22, Proposition 3.7.5] both M and $M^{\vee}$ are discrete. It follows from [Man22, Lemmas 3.9.2, 3.9.3] that M is bounded. The isomorphism $M \xrightarrow{\sim} M^{\vee\vee}$ is clear. This proves the 'only if' part of the claim.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

We now prove the 'if' part, so assume that M is bounded, M and $M^\vee$ are discrete and $M \xrightarrow{\sim} M^{\vee\vee}$ is an isomorphism. We want to show that M is dualizable, for which we can assume that M is concentrated in homological degree ≥ 0. Now consider $M_0 := \pi_0 M$, which is a static discrete $\mathcal{O}_K^a$-module. From the fact that $M^\vee$ is discrete it follows that $\pi_0 M_0^\vee = \pi_0 M^\vee$ is discrete. After choosing a surjection $\bigoplus_I \mathcal{O}_K^a/\pi \twoheadrightarrow M_0$ we get an injection $\pi_0 M_0^\vee \hookrightarrow \prod_I \mathcal{O}_K^a/\pi$. Now apply $\pi_0(-)_*$ to get a submodule

$$M_0' := \pi_0(\pi_0 M_0^\vee)_* \subseteq \prod_I A,$$

where we abbreviate $A := \pi_0(\mathcal{O}_K^a/\pi)_*$. Note that M_0' is equipped with the subspace topology since it is a kernel of a map $\prod_I A \to \prod_J A$. Moreover, the fact that $\pi_0 M_0^\vee$ is discrete implies that for every $\varepsilon \in \mathfrak{m}_K$ the submodule $T_\varepsilon \subset M_0'$ of ε-torsion in M_0' is open. In particular there is an open subset $U_\varepsilon \subset \prod_I A$ containing 0 such that $U_\varepsilon \cap M_0' \subseteq T_\varepsilon$. In fact we can choose U_ε to be an element of the standard basis for the product topology, i.e. $U_\varepsilon = U_\varepsilon' \times \prod_{I \setminus I_\varepsilon} A$ for some finite subset $I_\varepsilon \subset I$ and some open subset $U_\varepsilon' \subset \prod_{I_\varepsilon} A$. This means that if $(a_i)_i, (b_i)_i \in \prod_I A$ lie in M_0' and $a_i = b_i$ for $i \in I_\varepsilon$ then $\varepsilon(a_i)_i = \varepsilon(b_i)_i$. Now let $M_{0,\varepsilon}' \subset \prod_{I_\varepsilon} A$ be the image of M_0' under the projection $\prod_I A \twoheadrightarrow \prod_{I_\varepsilon} A$ (equipped with the subspace topology). Then there is a natural projection map $f_\varepsilon' \colon M_0' \to M_{0,\varepsilon}'$. We can also construct a map $g_\varepsilon' \colon M_{0,\varepsilon}' \to M_0'$ as follows: Given $a \in M_{0,\varepsilon}'$ we pick any $a' \in M_0'$ such that $f_\varepsilon'(a') = a$. Then we set $g_\varepsilon'(a) := \varepsilon a'$, which by the above observation is independent of the choice of a'. Note that $f_\varepsilon' g_\varepsilon' = \varepsilon$ and $g_\varepsilon' f_\varepsilon' = \varepsilon$, hence we have $M_0' \approx_\varepsilon M_{0,\varepsilon}'$. Going back to the almost world we now see that

$$\pi_0 M_0^\vee \approx_\varepsilon M_{0,\varepsilon}'^a, \qquad \text{with } M_{0,\varepsilon}'^a \subseteq \prod_{i=1}^{n_\varepsilon} \mathcal{O}_K^a/\pi \text{ a discrete submodule,}$$

where we took the liberty of naming $I_\varepsilon = \{1, \dots, n_\varepsilon\}$. By [Man22, Lemma 3.7.15] the discrete dualization functor $N \mapsto (N^\vee)_\omega$ is t-exact on discrete N. By applying this functor to $M^\vee$ it follows from the isomorphism $M \xrightarrow{\sim} M^{\vee\vee}$ that also $M_0 \xrightarrow{\sim} ((\pi_0 M_0^\vee)^\vee)_\omega$ is an isomorphism. Letting $M_{0,\varepsilon} := ((M_{0,\varepsilon}'^a)^\vee)_\omega$ we deduce that $M_0 \approx_\varepsilon M_{0,\varepsilon}$ and there is a surjection $\bigoplus_{i=1}^{n_\varepsilon} \mathcal{O}_K^a/\pi \twoheadrightarrow M_{0,\varepsilon}$. As this works for all ε, we deduce that M_0 is almost finitely generated as an $\mathcal{O}_K^a/\pi$-module. Then it is also almost finitely generated as an $\mathcal{O}_K^a$-module (where we equip $\mathcal{O}_K^a$ with the discrete topology in this proof), hence by [Sch13, Theorem 2.5] we have

$$M_0 \approx (\mathcal{O}_K^a/\pi)^r \oplus \mathcal{O}_K^a/I_{\gamma_1} \oplus \mathcal{O}_K^a/I_{\gamma_2} \oplus \dots$$

for a series of real numbers $1 > \gamma_1 \geq \gamma_2 \geq \dots \geq 0$ converging to 0. Here for a real number γ we denote by $I_\gamma \subset \mathcal{O}_K$ the ideal generated by $\pi^{\gamma'}$ for all $\gamma' > \gamma$ (see [Sch13, Example 2.4. (i)]). The φ-module structure $\varphi^* M \xrightarrow{\sim} M$ induces an isomorphism $\pi_0 \varphi^* M_0 \xrightarrow{\sim} M_0$. On the other hand for $\gamma < 1$ we have

$$\pi_0 \varphi^*(\mathcal{O}_K^a/I_\gamma) = \begin{cases} \mathcal{O}_K^a/\pi, & \text{if } \gamma \geq 1/p, \\ \mathcal{O}_K^a/I_{p\gamma}, & \text{else,} \end{cases}$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

as $\mathcal{O}_K^a/\pi$-modules. This forces all $\gamma_i = 0$, i.e. we have $M_0 \approx (\mathcal{O}_K^a/\pi)^r$. It follows that M_0 is dualizable: We need to check that the natural map $\alpha\colon M_0 \otimes M_0^\vee \to \underline{\mathrm{Hom}}(M_0, M_0)$ is an isomorphism but since the same statement holds for $(\mathcal{O}_K^a/\pi)^r$ in place of M_0 one deduces easily that both kernel and cokernel of $\pi_i\alpha$ are killed by ε for all $i \in \mathbf{Z}$ and $\varepsilon \in \mathfrak{m}_K$ (cf. the proof of [Man22, Lemma 3.7.17]). In particular M_0 is flat (which can also be seen directly), so that $\pi_0\varphi^* M_0 = \varphi^* M$ and therefore the φ-module structure on M_0 restricts to a φ-module structure on M_0. Since the conditions on M are preserved under finite limits, we reduce the claimed dualizability of M to that of $\tau_{\geq 1}M$, which satisfies the same assumptions as M. Since M is bounded, we can conclude by induction. □

As promised, here come the three equivalent ways of characterizing admissible representations on a classifying stack. The third characterization is very useful for showing that admissible representations are stable under various of the 6 operations, and plays a key role in our analysis of the Jacquet-Langlands functor.

Proposition 3.26. *Let C be an algebraically closed perfectoid field with pseudouniformizer $\pi \in \mathcal{O}_C$ such that $\pi | p$ and let G be a locally profinite p-Poincaré group. Then for $M \in \mathcal{D}_\square^{\mathrm{sm},a}(G, \mathcal{O}_C/\pi)^\varphi$ the following are equivalent:*

(i) *M is bounded and admissible.*

(ii) *M is bounded and for a basis of compact open subgroups $H \subseteq G$, $\Gamma_{\mathrm{sm}}(H, M)$ is dualizable in $\mathcal{D}_\square^a(\mathcal{O}_C/\pi)$.*

(iii) *Both M and $M^\vee$ are bounded and discrete and the natural morphism $M \xrightarrow{\sim} M^{\vee\vee}$ is an isomorphism.*

If this is the case then $M^\vee$ is bounded and admissible and this dual can unambiguously be computed in any of the following ∞-categories: $\mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p)$, $\mathcal{D}_\square^{\mathrm{sm}}(G, \mathbf{F}_p)$, $\mathcal{D}_\square^{\mathrm{sm},a}(G, \mathcal{O}_C/\pi)_\omega^\varphi$, $\mathcal{D}_\square^{\mathrm{sm},a}(G, \mathcal{O}_C/\pi)^\varphi$.

Proof. Note first that (i) implies (ii) because $R\Gamma_{\mathrm{sm}}(H, V) \otimes_{\mathbf{F}_p} \mathcal{O}_C^a/\pi = R\Gamma_{\mathrm{sm}}(H, V \otimes_{\mathbf{F}_p} \mathcal{O}_C^a/\pi)$ for all $V \in \mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p)$ (e.g. by the explicit computation of smooth group cohomology in [Man22, Corollary 3.4.17]). Conversely, assume that (ii) is satisfied. Then M is discrete because $M = \varinjlim_H R\Gamma_{\mathrm{sm}}(H, M)$, where on the right-hand side H ranges through all compact open subgroups of G (see [Man22, Lemma 3.4.19]). We claim that M lies in $\mathcal{D}^{\mathrm{sm}}(G, \mathbf{F}_p)$, i.e. that the natural morphism $M^\varphi \otimes_{\mathbf{F}_p} \mathcal{O}_C^a \xrightarrow{\sim} M$ is an isomorphism. By the colimit property we just used the collection of functors $R\Gamma_{\mathrm{sm}}(H, -)$, where H ranges through a basis of compact open subgroups of G, is conservative, hence it is enough to show that $R\Gamma_{\mathrm{sm}}(H, M^\varphi \otimes_{\mathbf{F}_p} \mathcal{O}_C^a) = R\Gamma_{\mathrm{sm}}(H, M)$ for all H in the basis. Noting that that $R\Gamma_{\mathrm{sm}}(H, -)$ commutes with $(-)^\varphi$ (which is formal) and with $- \otimes_{\mathbf{F}_p} \mathcal{O}_C^a/\pi$ (as used above), the claim boils down to showing that for $M_H := R\Gamma_{\mathrm{sm}}(H, M)$ we have $M_H^\varphi \otimes_{\mathbf{F}_p} \mathcal{O}_C^a/\pi \xrightarrow{\sim} M_H$. But by assumption M_H is dualizable, hence the desired equivalence follows from the Riemann-Hilbert correspondence (see [Man22, Theorem 3.9.23]). From the Riemann-Hilbert correspondence we also deduce that $R\Gamma_{\mathrm{sm}}(H, M^\varphi)$ is dualizable, hence M^φ is admissible, as desired.

We have now finished the proof that (i) and (ii) are equivalent. To prove that they are also equivalent to (iii) we make the following computation. Let $X = \mathrm{Spa}C$, $\Lambda = \mathcal{O}_C/\pi$,

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

$\omega_G := f^!\Lambda^a$ and for every compact open subgroup $H \subset G$ let $f_H \colon X/H \to X$ be the natural map. Then from Corollary 3.12 and Proposition 3.14 we get:

$$\begin{aligned} f_{H*}M^\vee &= f_{H*}\underline{\mathrm{Hom}}(M,\Lambda^a) = f_{H*}\underline{\mathrm{Hom}}(M\otimes\omega_G,\omega_G) = \underline{\mathrm{Hom}}(f_{H*}(M\otimes\omega_G),\Lambda^a) \\ &= (f_{H*}(M\otimes\omega_G))^\vee. \end{aligned}$$

Note also that for small enough H the action of H on ω_G is trivial, so that $-\otimes\omega_G$ is just a shift in the ∞-category of H-representations. Consequently for small enough H the functor $R\Gamma_{\mathrm{sm}}(H,-)$ preserves duals up to shift. Since the functors $R\Gamma_{\mathrm{sm}}(H,-)$ are also conservative (for H varying in a basis of compact open subgroups of G) it follows that (iii) is equivalent to the claim that M and $M^\vee$ are bounded and for a basis H of compact open subgroups of G we have that both $M_H := R\Gamma_{\mathrm{sm}}(H,M)$ and $M_H^\vee$ are discrete and the natural morphism $M_H \xrightarrow{\sim} M_H^{\vee\vee}$ is an isomorphism. By Lemma 3.25 this implies (ii). Moreover, if (ii) is satisfied then all the conditions in (iii) follow immediately except possibly that $M^\vee$ is also bounded. To see that $M^\vee$ is indeed bounded, note first that $M^\vee = \varinjlim_H R\Gamma_{\mathrm{sm}}(H,M^\vee)$ (see [Man22, Lemma 3.4.19]) and by the above computation the right-hand side is (up to a constant shift) isomorphic to $\varinjlim_H M_H^\vee$. On the other hand, the discrete dualization functor $N \mapsto (N^\vee)_\omega$ is t-exact on discrete objects in $\mathcal{D}^a_\square(\mathcal{O}_C/\pi)$ (see [Man22, Lemma 3.7.15]). In particular, if M_H is concentrated in homological degrees $[a,b]$ then $M_H^\vee$ is concentrated in homological degrees $[-b,-a]$. The fact that M is bounded implies that we can choose a and b independently of H, which then implies that $M^\vee$ is also bounded.

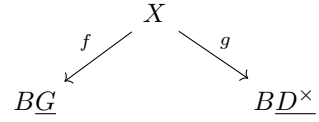
It remains to prove the last part of the claim, so assume that M is admissible and bounded. It follows immediately from (iii) that also $M^\vee$ (computed in $\mathcal{D}^{\mathrm{sm},a}_\square(G,\mathcal{O}_C/\pi)$) is admissible (and bounded), so in particular it is discrete and comes from an $\mathbf{F}_p$-representation via the Riemann-Hilbert correspondence. It is thus formal that this dual can be computed in any of the claimed ∞-categories. □

Remark 3.27. It follows from the proof that in Lemma 3.25 it is enough to require that there is *some* isomorphism $M \cong M^{\vee\vee}$ instead of insisting that the natural map $M \to M^{\vee\vee}$ is an isomorphism. It follows that in condition (iii) of Proposition 3.26 it is enough to require the existence of *some* isomorphism $M \cong M^{\vee\vee}$.

4. Duality and applications

4.1. The main result

Fix a finite extension $F/\mathbf{Q}_p$, let $G := \mathrm{GL}_n(F)$, let D be the central division algebra over F of invariant $1/n$ and let C be the completed algebraic closure of F. In the following we will work with small v-stacks over $\mathrm{Spd}\,C$, so in particular the final object $*$ will denote $\mathrm{Spd}\,C$ and $B\underline{H} = \mathrm{Spd}\,C/\underline{H}$ for any locally profinite group H. Recall the fundamental diagram from the introduction:



https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Here X is obtained from the infinite level Lubin-Tate tower by modding out the commuting actions of G and $D^\times$, so that in particular we have $X \cong [\mathbf{P}_C^{n-1}/\underline{D^\times}] \cong [\Omega_C^{n-1}/\underline{G}]$, where $\Omega_C^{n-1} = \mathbf{P}_C^{n-1} \setminus \mathbf{P}_C^{n-1}(\mathbf{Q}_p)$ denotes Drinfeld space. It follows that f and g are smooth and that g is proper. Recall also that for every locally profinite group H which has locally finite p-cohomological dimension (e.g. $H = G$ or $H = D^\times$) we have

$$\mathcal{D}_{\text{ét}}(B\underline{H},\mathbf{F}_p) = \mathcal{D}^{\text{sm}}(H,\mathbf{F}_p),$$
$$\mathcal{D}_\square^a(\mathcal{O}^+_{B\underline{H}}/p)^\varphi = \mathcal{D}_\square^{\text{sm},a}(H,\mathcal{O}_C/p)^\varphi,$$

and the upper ∞-category embedds fully faithfully into the lower one via the Riemann-Hilbert functor. As explained in the introduction, we can now define the Jacquet-Langlands functor as follows.

Definition 4.1. In the above setup, we define the *p-adic Jacquet-Langlands functors*

$$\mathcal{J}\colon \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p) \to \mathcal{D}^{\text{sm}}(D^\times,\mathbf{F}_p), \qquad A \mapsto g_{\text{ét},*}f_{\text{ét}}^* A,$$
$$\mathcal{J}_\square\colon \mathcal{D}_\square^{\text{sm},a}(G,\mathcal{O}_C/p)^\varphi \to \mathcal{D}_\square^{\text{sm},a}(D^\times,\mathcal{O}_C/p)^\varphi, \qquad A \mapsto g_* f^* A,$$

where we implicitly use the identifications of sheaves on the classifying stack with smooth representations discussed above.

The functor $\mathcal{J}_\square$ is defined in the setting of a full 6-functor formalism and should thus be seen as the 'actual' Jacquet-Langlands functor, whereas $\mathcal{J}$ is a more naively defined shadow of $\mathcal{J}_\square$. In fact, it is completely formal that for all $A \in \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)$ we have

$$\mathcal{J}(A) = (\mathcal{J}_\square(A \otimes \mathcal{O}^{+a}_{B\underline{G}}/p))^\varphi. \tag{1}$$

In other words, $\mathcal{J}(A)$ remembers the φ-invariants of the Jacquet-Langlands correspondent $\mathcal{J}_\square(A \otimes \mathcal{O}^+_{B\underline{G}}/p)$ of A. It turns out that if A is admissible then in fact these φ-invariants capture the whole Jacquet-Langlands correspondent. In fact, the Jacquet-Langlands correspondent of an admissible representation is itself admissible and hence lies in the image of the Riemann-Hilbert functor, as is shown in the following result.

Theorem 4.2. *In the above setup, let $A \in \mathcal{D}^{\text{sm}}(G,\mathbf{F}_p)$ be admissible. Then:*

(i) *$\mathcal{J}(A) \in \mathcal{D}^{\text{sm}}(D^\times,\mathbf{F}_p)$ is admissible.*

(ii) *The natural primitive comparison map*

$$\mathcal{J}(A) \otimes \mathcal{O}^{+a}_{B\underline{D^\times}}/p \xrightarrow{\sim} \mathcal{J}_\square(A \otimes \mathcal{O}^{+a}_{B\underline{G}}/p)$$

induced by Eq. (1) is an isomorphism.

(iii) *There is a natural isomorphism*

$$(\mathcal{S}_{D^\times} \circ \mathcal{J})(A) = (\mathcal{J} \circ \mathcal{S}_G)(A)[2n-2](n-1)$$

in $\mathcal{D}^{\text{sm}}(D^\times,\mathbf{F}_p)$.

Recall that Scholze already proved (i) and (ii). We will prove (i)–(iii) simultaneously, independently of Scholze's results. In fact, we will see that all (i)–(iii) follow formally from the following unconditional duality for $\mathcal{J}_\square$:

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Lemma 4.3. *In the above setup, there is a natural isomorphism*

$$(-)^\vee \circ \mathcal{J}_\square = (\mathcal{J}_\square \circ (-)^\vee)[2n-2](n-1)$$

of functors $\mathcal{D}_\square^{\mathrm{sm},a}(G,\mathcal{O}_C/p)^\varphi \to \mathcal{D}_\square^{\mathrm{sm},a}(D^\times,\mathcal{O}_C/p)^\varphi$. *Here* $(-)^\vee$ *denotes the internal hom towards the monoidal unit.*

Proof. For every $A \in \mathcal{D}_\square^{\mathrm{sm},a}(G,\mathcal{O}_C/p)^\varphi$ we get the following chain of natural equivalences (using that g is proper and smooth and that f is smooth):

$$\begin{aligned}\underline{\mathrm{Hom}}(g_* f^* A, \mathcal{O}^{+a}_{B\underline{D^\times}}/p) &= \underline{\mathrm{Hom}}(g_! f^* A, \mathcal{O}^{+a}_{B\underline{D^\times}}/p)\\ &= g_* \underline{\mathrm{Hom}}(f^* A, g^!(\mathcal{O}^{+a}_{B\underline{D^\times}}/p))\\ &= g_* \underline{\mathrm{Hom}}(f^* A, \mathcal{O}^{+a}_X/p)[2n-2](n-1)\\ &= g_* f^* \underline{\mathrm{Hom}}(A, \mathcal{O}^{+a}_{B\underline{G}}/p)[2n-2](n-1).\end{aligned}$$

In the last step we used that by the smoothness of f, f^* commutes with internal hom (see [Man22, Proposition 3.8.7]. □

Proof of Theorem 4.2. Fix an admissible $A \in \mathcal{D}^{\mathrm{sm}}(G,\mathbf{F}_p)$ as in the claim and let us denote $A' := A \otimes \mathcal{O}^{+a}_{B\underline{G}}/p$ the corresponding object in $\mathcal{D}_\square^{\mathrm{sm},a}(G,\mathcal{O}_C/p)^\varphi$ given by the Riemann-Hilbert embedding. Then A' is admissible (by definition) and hence by Proposition 3.26 both A' and $A'^\vee$ are bounded and discrete and the natural map $A' \xrightarrow{\sim} A'^{\vee\vee}$ is an isomorphism. Clearly $\mathcal{J}_\square$ preserves bounded and discrete objects (this is true for f^* like for any pullback and it follows for g_* from the fact that g is proper and p-bounded). Thus it follows from Lemma 4.3 that both $\mathcal{J}_\square(A')$ and $\mathcal{J}_\square(A')^\vee$ are bounded and discrete. Moreover, Lemma 4.3 also implies that there is an isomorphism $\mathcal{J}_\square(A') \cong \mathcal{J}_\square(A')^{\vee\vee}$, so altogether it follows from Proposition 3.26 (and Remark 3.27) that $\mathcal{J}_\square(A')$ is admissible. This immediately implies (i). It also implies (ii) because we now know that $\mathcal{J}_\square(A')$ lies in the image of the Riemann-Hilbert functor. Now (iii) follows from Lemma 4.3 and the last part of Proposition 3.26. □

Remark 4.4. In general, both sides of the duality isomorphism are somewhat intractable. However, in certain favorable situations it reduces to a more concrete statement. We record one useful instance of this in the following proposition. Recall from Proposition 3.7 that for any admissible smooth $\pi \in \mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G)$, the derived dual $\mathcal{S}_G^{\dim_G \pi}(\pi)$ is always nonzero. Following [Koh17], we say π is Cohen-Macaulay if $\mathcal{S}_G^i(\pi) = 0$ for all $i \neq \dim_G \pi$.

Proposition 4.5. *Let* $\pi \in \mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(G)$ *be an admissible representation which is Cohen-Macaulay, and such that there is an integer* c *with* $\mathcal{J}^i(\pi) = 0$ *for all* $i \neq c$. *Then there are isomorphisms*

$$\mathcal{S}^i_{D^\times}(\mathcal{J}^c(\pi)) \simeq \mathcal{J}^{i+2n-2-c-\dim_G \pi}(\check{\pi})$$

for all i, *where* $\check{\pi} = \mathcal{S}_G^{\dim_G \pi}(\pi)$ *is the unique nonzero derived dual of* π.

Proof. Evaluate both sides of the duality isomorphism $\mathcal{S}_{D^\times} \circ \mathcal{J}) \cong (\mathcal{J} \circ \mathcal{S}_G)[2n-2](n-1)$ on π. Since $\mathcal{J}(\pi) = \mathcal{J}^c(\pi)[-c]$ by assumption, the left-hand side evaluates to

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

$\mathcal{S}_{D^\times}(\mathcal{J}^c(\pi)[-c]) \cong \mathcal{S}_{D^\times}(\mathcal{J}^c(\pi))[c]$. Next, since $\mathcal{S}_G(\pi) = \check{\pi}[-\dim_G \pi]$ by assumption, the right-hand side simplifies to $\mathcal{J}(\check{\pi}[-\dim_G \pi])[2n-2](n-1) \cong \mathcal{J}(\check{\pi})[2n-2-\dim_G \pi](n-1)$. Equating both sides and rearranging the shifts, we get $\mathcal{S}_{D^\times}(\mathcal{J}^c(\pi)) = \mathcal{J}(\check{\pi})[2n-2-\dim_G \pi - c](n-1)$. Passing to cohomology on both sides (and ignoring the Tate twists), we get the result. □

Remark 4.6. We explain how to show a variant of our duality isomorphism in the setting of $\mathbf{Z}/p^n$ or $\mathbf{Z}_p$-coefficients. We explain the first case; the second follows from a limit argument. First some preliminary observations. Let H be any p-adic Lie group. One easily checks that Proposition 3.9 extends to the setting of $\mathbf{Z}/p^n$-coefficients, yielding a symmetric monoidal equivalence $\mathcal{D}^{\mathrm{sm}}(H,\mathbf{Z}/p^n) \cong \mathcal{D}_{\text{ét}}(B\underline{H},\mathbf{Z}/p^n)$. On these categories we again have a duality functor $\mathcal{S}_H(-)$, given by internal hom towards the monoidal unit. We say $A \in \mathcal{D}^{\mathrm{sm}}(H,\mathbf{Z}/p^n)$ is admissible if $R\Gamma_{\mathrm{sm}}(H,A)$ is a perfect complex of $\mathbf{Z}/p^n$-modules for a basis of open subgroups H. Admissible complexes are automatically bounded, and are preserved by duality and $-\otimes^{\mathbf{L}}_{\mathbf{Z}/p^n}\mathbf{F}_p$, and the latter functor is conservative on admissible complexes. Conversely, if A is given and $A\otimes^{\mathbf{L}}_{\mathbf{Z}/p^n}\mathbf{F}_p \in \mathcal{D}^{\mathrm{sm}}(H,\mathbf{F}_p)$ is admissible, then A is necessarily admissible.

With this in mind, we *define* $\mathcal{J}: \mathcal{D}^{\mathrm{sm}}(G,\mathbf{Z}/p^n) \to \mathcal{D}^{\mathrm{sm}}(D^\times,\mathbf{Z}/p^n)$ as $\mathcal{J} = g_{\text{ét}*}f^*_{\text{ét}}$. This is compatible with the $\mathbf{F}_p$ version of $\mathcal{J}$ via $-\otimes^{\mathbf{L}}_{\mathbf{Z}/p^n}\mathbf{F}_p$. Using this observation and the final statement in the previous paragraph together with the finiteness properties of $\mathcal{J}$ with $\mathbf{F}_p$-coefficients, it is easy to see that $\mathcal{J}$ preserves admissibility in this setting. Now, the key observation is that there is a natural map $\mathcal{J}(A)\otimes\mathcal{J}(\mathcal{S}_G(A)) \to \mathbf{Z}/p^n[2-2n](1-n)$, defined as the composition

$$\begin{aligned}\mathcal{J}(A)\otimes\mathcal{J}(\mathcal{S}_G(A)) &\to \mathcal{J}(A\otimes\mathcal{S}_G(A))\\ &\to \mathcal{J}(\mathbf{Z}/p^n)\\ &\to \tau^{\geq 2n-2}\mathcal{J}(\mathbf{Z}/p^n)\\ &\cong \mathbf{Z}/p^n[2-2n](1-n).\end{aligned}$$

Here the first arrow is obtained from the lax monoidal structure of $\mathcal{J}$, the second arrow is induced by the evaluation map $A\otimes\mathcal{S}_G(A) \to \mathbf{Z}/p^n$, and the final isomorphism is induced by the usual computation of the étale cohomology of $\mathbf{P}^{n-1}$. By adjunction (and rearranging the shifts and twists), this map induces a map $\mathcal{J}(\mathcal{S}_G(A))[2n-2](n-1) \to \mathcal{S}_{D^\times}(\mathcal{J}(A))$. When A is admissible, the source and target of this map are also admissible, so to see it is an isomorphism, it suffices to show this after applying $-\otimes^{\mathbf{L}}_{\mathbf{Z}/p^n}\mathbf{F}_p$. But all operations here are compatible with reduction mod p, so we now conclude by invoking the duality isomorphism with $\mathbf{F}_p$-coefficients.

4.2. Applications to Gelfand-Kirillov dimension

In this section we prove Theorems 1.6 and 1.7, along with some related results. We will freely use the basic properties of dimension, in particular the results of Propositions 3.7-3.8 and the surrounding discussion.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proof of Theorem 1.6. Let $d = \dim_G(\pi)$, so $\mathcal{S}_G(\pi) \in \mathcal{D}^{\mathrm{sm}}(G)^{[0,d]}$. Since $\mathcal{J}(-)$ carries $\mathcal{D}^{\mathrm{sm}}(G)^{[a,b]}$ into $\mathcal{D}^{\mathrm{sm}}(D^\times)^{[a,b+2n-2]}$, we deduce that

$$(\mathcal{J} \circ \mathcal{S}_G)(\pi)[2n-2] \in \mathcal{D}^{\mathrm{sm}}(D^\times)^{\leq d}.$$

By the duality isomorphism $(\mathcal{S}_{D^\times} \circ \mathcal{J})(\pi) \cong (\mathcal{J} \circ \mathcal{S}_G)(\pi)[2n-2]$, we get that $M = (\mathcal{S}_{D^\times} \circ \mathcal{J})(\pi) \in \mathcal{D}^{\mathrm{sm}}(D^\times)^{\leq d}$. Now look at the fourth quadrant spectral sequence

$$E_2^{i,j} = \mathcal{S}^i_{D^\times}(\mathcal{J}^{-j}(\pi)) \Rightarrow H^{i+j}(M).$$

Let $c = \max_i \dim_{D^\times} \mathcal{J}^i(\pi)$, and let i_0 be the least i realizing this maximum. Then $E_2^{c,-i_0}$ is c-dimensional as an $D^\times$-representation, and as we turn the pages of the spectral sequence it only interacts with representations of dimension $< c$ via the incoming differentials. Therefore a nonzero subquotient survives to the E_∞-page, giving that $H^{c-i_0}(M) \neq 0$. But $M \in \mathcal{D}^{\mathrm{sm}}(D^\times)^{\leq d}$, so this implies that $c - i_0 \leq d$, and therefore $c \leq d + i_0 \leq d + N_\pi$, as desired. □

We conjecture the following strengthening of Theorem 1.6.

Conjecture 4.7. *For any $\pi \in \mathrm{Rep}^{\mathrm{adm}}_{\mathbf{F}_p}(G)$, we have $\dim_{D^\times} \mathcal{J}^i(\pi) \leq \dim_G \pi$ for all $i \geq 0$.*

In many cases, this can be proved. In particular, we have the following result, which can be proved by an easy adaptation of the proof of Theorem 1.6.

Proposition 4.8. *Suppose $\pi \in \mathrm{Rep}^{\mathrm{adm}}_{\mathbf{F}_p}(G)$ is such that for all $i > n-1$ and all $j \geq 0$, $\mathcal{J}^i(\pi) = 0$ and $\mathcal{J}^i(\mathcal{S}^j_G(\pi)) = 0$. Then Conjecture 4.7 is true for π.*

The hypotheses of this proposition are satisfied for any representation of the form $\pi = \mathrm{Ind}_P^G \sigma$, where P is any parabolic subgroup contained in the standard maximal parabolic $P_{n-1,1}$ and σ is any admissible smooth representation of the Levi quotient $P \twoheadrightarrow M$. To see this, note that the main theorem of [JLH21] shows that $\mathcal{J}^i \circ \mathrm{Ind}_P^G = 0$ for all $i > n-1$. It's then enough to observe that for all j we have $\mathcal{S}^j_G(\mathrm{Ind}_P^G \sigma) = \mathrm{Ind}_P^G \delta_P \mathcal{S}_M^{j-d}(\sigma)$ for a certain integer d and smooth character δ_P (use [Koh17, Theorem 4.7 and Corollary 5.3]), so the main theorem of [JLH21] again applies.

We also observe that Conjecture 4.7 implies that $\mathcal{J}^{2n-2}$ is very small, and can be explicitly computed.

Proposition 4.9. *Suppose that Conjecture 4.7 is true. Then for any $\pi \in \mathrm{Rep}^{\mathrm{adm}}_{\mathbf{F}_p}(G)$ admitting a finite-length filtration whose graded pieces have central characters, we have $\dim_{\mathbf{F}_p} \mathcal{J}^{2n-2}(\pi) < \infty$, and in fact there is an isomorphism*

$$\mathcal{J}^{2n-2}(\pi) = \mathcal{S}^0_{D^\times}(\mathcal{J}^0(\mathcal{S}^0_G(\pi))).$$

Note that the slightly strange hypothesis on π is satisfied for any irreducible admissible representation. The final isomorphism here implies (assuming Conjecture 4.7) that $\mathcal{J}^{2n-2}(\pi)$ is nonzero *exactly* when π admits a G-stable quotient $\pi \twoheadrightarrow \tau$ such that $\dim_{\mathbf{F}_p} \tau < \infty$ and $\tau^{\mathrm{SL}_n(F)} \neq 0$. We also note that the following proof shows that $\mathcal{J}^{2n-2}(\pi) = \mathcal{S}^0_{D^\times}(\mathcal{J}^0(\mathcal{S}^0_G(\pi)))$ whenever $\dim_{D^\times} \mathcal{J}^{2n-2}(\pi) = 0$, independently of Conjecture 4.7.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proof. Set $M = (\mathcal{J} \circ \mathcal{S}_G)(\pi)$. By the duality theorem, we get an E_2 spectral sequence

$$E_2^{i,j} = \mathcal{S}^i_{D^\times}(\mathcal{J}^{-j}(\pi)) \Rightarrow H^{i+j+2n-2}(M).$$

Suppose that $\dim_{\mathbf{F}_p} \mathcal{J}^{2n-2}(\pi) = \infty$, i.e. that $c := \dim_{D^\times} \mathcal{J}^{2n-2}(\pi) > 0$. Then $E_2^{c,-(2n-2)}$ is nonzero and c-dimensional, and a c-dimensional quotient of it survives to the E_∞-page, showing that $\dim_{D^\times} H^c(M) \geq c$. On the other hand, we have a second spectral sequence $\mathcal{J}^i(\mathcal{S}^j_G(\pi)) \Rightarrow H^{i+j}(M)$, showing that $H^c(M)$ has a finite filtration whose gradeds are subquotients of $\mathcal{J}^i(\mathcal{S}^{c-i}_G(\pi))$ for varying i. Since $\dim_G \mathcal{S}^{c-i}_G(\pi) \leq c-i$, Conjecture 4.7 then implies that $\dim_{D^\times} \mathcal{J}^i(\mathcal{S}^{c-i}_G(\pi)) \leq c-i < c$ for $i > 0$. But when $i = 0$, $\dim_{D^\times} \mathcal{J}^0(\mathcal{S}^c_G(\pi)) = 0$ by the explicit calculation of $\mathcal{J}^0$, using our assumption on π to reduce to the case of π admitting a central character. Therefore $H^c(M)$ has a finite filtration whose gradeds have dimension $< c$, so $\dim_{D^\times} H^c(M) < c$. This is a contradiction.

For the final statement, note that the two spectral sequences show (unconditionally) that

$$\mathcal{S}^0_{D^\times}(\mathcal{J}^{2n-2}(\pi)) = \mathcal{J}^0(\mathcal{S}^0_G(\pi))$$

for all $\pi \in \mathrm{Rep}^{\mathrm{adm}}_{\mathbf{F}_p}(G)$. If $\mathcal{J}^{2n-2}(\pi)$ is zero-dimensional, then $\mathcal{S}^0_{D^\times}(\mathcal{J}^{2n-2}(\pi)) = \mathcal{S}_{D^\times}(\mathcal{J}^{2n-2}(\pi))$ is the full derived dual of $\mathcal{J}^{2n-2}(\pi)$, so the result follows from biduality. □

Proof of Theorem 1.7. Let $\check{\pi}$ be the unique nonzero derived dual of π. The duality isomorphism then simplifies to an isomorphism

$$\mathcal{S}_{D^\times}(\mathcal{J}(\pi)) \cong \mathcal{J}(\check{\pi})[2n-2-\dim_G \pi].$$

Since $\mathcal{J}(\check{\pi})$ is concentrated in degrees ≥ 0, the right-hand side, and then also the left-hand side, is concentrated in degrees $\geq \dim_G \pi - (2n-2)$.

If $\mathcal{J}(\pi)$ is zero, there is nothing to prove. Otherwise, let $c = \max_i \dim_{D^\times} \mathcal{J}^i(\pi)$, and let i_0 be any i realizing this maximum. Arguing as in the proof of Theorem 1.6, we see that the left-hand side has nonzero cohomology in degree $c - i_0$. Therefore $c \geq \dim_G \pi - (2n-2) + i_0 \geq \dim_G \pi - (2n-2)$, giving the desired result. □

4.3. The case of GL_2

Now we specialize to the case $n = 2$, with $F/\mathbf{Q}_p$ still arbitrary. Set $d = [F : \mathbf{Q}_p]$. Let $B \subset G = \mathrm{GL}_2(F)$ be the upper-triangular Borel subgroup, with $T = F^\times \times F^\times \subset B$ the diagonal maximal torus. Let $\omega : F^\times \to \mathbf{F}_p^\times$ be the character defined as the composition

$$F^\times \xrightarrow{\mathrm{Nm}_{F/\mathbf{Q}_p}} \mathbf{Q}_p^\times \xrightarrow{x \mapsto x|x|} \mathbf{Z}_p^\times \xrightarrow{\mathrm{red}} \mathbf{F}_p^\times.$$

If $\chi_1 \otimes \chi_2 : T \to \mathbf{F}_p^\times$ is any smooth character, we can consider the smooth induction $\mathrm{Ind}_B^G(\chi_1 \otimes \chi_2)$. This is irreducible iff $\chi_1 \neq \chi_2$. Moreover, $\mathcal{S}^i_G(\mathrm{Ind}_B^G(\chi_1 \otimes \chi_2)) = 0$ for all $i \neq d$, and

$$\mathcal{S}^d_G(\mathrm{Ind}_B^G(\chi_1 \otimes \chi_2)) \cong \mathrm{Ind}_B^G(\omega\chi_1^{-1} \otimes \omega^{-1}\chi_2^{-1}).$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

In particular, if we set $\pi(\chi_1,\chi_2) = \mathrm{Ind}_B^G(\omega\chi_1 \otimes \chi_2)$, then

$$\mathcal{S}_G^d(\pi(\chi_1,\chi_2)) \otimes (\chi_1\chi_2\omega \circ \det) \cong \pi(\chi_2,\chi_1),$$

and moreover $\pi(\chi_1,\chi_2)$ and its dual $\mathcal{S}_G^d(\pi(\chi_1,\chi_2))$ are both irreducible iff $\chi_1\chi_2^{-1} \neq \omega^{\pm 1}$.

Assume then that χ_1,χ_2 is a pair of characters such that $\chi_1/\chi_2 \neq \omega^{\pm 1}$. As noted in the introduction, we then have $\mathcal{J}^i(\pi(\chi_1,\chi_2)) = 0$ for all $i \neq 1$, and likewise with χ_1 and χ_2 swapped. With this vanishing in mind, we set $\tau(\chi_1,\chi_2) = \mathcal{J}^1(\pi(\chi_1,\chi_2))$.

Theorem 4.10. *Notation and assumptions as above, exactly one of the following is true.*

i. $\tau(\chi_1,\chi_2) = \tau(\chi_2,\chi_1) = 0$.

ii. *Both* $\tau(\chi_1,\chi_2)$ *and* $\tau(\chi_2,\chi_1)$ *are nonzero, and are exchanged by the functor* $\mathcal{S}_{D^\times}^d(-) \otimes (\chi_1\chi_2\omega \circ \mathrm{Nm})$. *Moreover,* $\mathcal{S}_{D^\times}^i(-)$ *applied to either representation vanishes for all* $i \neq d$.

Proof. Evaluate the duality isomorphism $(\mathcal{S}_{D^\times} \circ \mathcal{J}) \cong (\mathcal{J} \circ \mathcal{S}_G)[2]$ on the representation $\pi(\chi_1,\chi_2)$. Since $\mathcal{J}(\pi(\chi_1,\chi_2)) = \tau(\chi_1,\chi_2)[-1]$, the left-hand side simplifies to $\mathcal{S}_{D^\times}(\tau(\chi_1,\chi_2))[1]$ by the results recalled above. On the other hand, $\mathcal{S}_G(\pi(\chi_1,\chi_2)) \cong \pi(\chi_2,\chi_1) \otimes (\chi_1\chi_2\omega \circ \det)^{-1}[-d]$, so the right-hand side simplifies to $\mathcal{J}(\pi(\chi_2,\chi_1)) \otimes (\chi_1\chi_2\omega \circ \mathrm{Nm})^{-1}[2-d] \cong \tau(\chi_2,\chi_1) \otimes (\chi_1\chi_2\omega \circ \mathrm{Nm})^{-1}[1-d]$. Thus, canceling the shifts, we get an isomorphism

$$\mathcal{S}_{D^\times}(\tau(\chi_1,\chi_2)) \cong \tau(\chi_2,\chi_1) \otimes (\chi_1\chi_2\omega \circ \mathrm{Nm})^{-1}[-d],$$

which immediately gives the desired result. □

4.3.1. A non-generic case. In this section we sketch some consequences of our duality theorem in the setting of *reducible* principal series representations of $\mathrm{GL}_2(\mathbf{Q}_p)$. Maintain the previous notation, but with $F = \mathbf{Q}_p$. The principal series representation $\pi(\omega^{-1},1)$ is reducible of length two, and sits in a short exact sequence

$$0 \to \mathbf{1} \to \pi(\omega^{-1},1) \to \mathrm{St} \to 0 \quad (\dagger)$$

where St is the mod p Steinberg representation of $\mathrm{GL}_2(\mathbf{Q}_p)$. To ease notation, set $\kappa = \mathcal{S}_G^1(\mathrm{St})$. Note that $\mathcal{S}_G(\mathbf{1}) = \mathbf{1}[0]$ and $\mathcal{S}_G(\pi(\omega^{-1},1)) = \pi(1,\omega^{-1})[-1]$, where $\pi(1,\omega^{-1})$ is an irreducible principal series. The associated long exact sequence for $\mathcal{S}_G^i(-)$ is now easy to analyze: the embedding $\mathcal{S}_G^0(\mathrm{St}) \hookrightarrow \mathcal{S}_G^0(\pi(\omega^{-1},1)) = 0$ implies that the long exact sequence degenerates to a short exact sequence

$$0 \to \mathbf{1} \to \kappa \to \pi(1,\omega^{-1}) \to 0 \quad (\ddagger)$$

and that $\mathcal{S}_G^i(\mathrm{St})$ for all $i \neq 1$. Biduality then implies that the previous sequence is non-split, since otherwise the nonvanishing of $\mathcal{S}_G^0(-)$ applied to the first term would contradict the vanishing of $\mathcal{S}_G^0(\kappa)$. Note that κ is the unique non-split extension of $\pi(1,\omega^{-1})$ by $\mathbf{1}$.

Using the main result in [Lud17] together with the explicit computation of $\mathcal{J}^0$, the following vanishing results are clear:

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

- $\mathcal{J}(\mathbf{1}) = R\Gamma(\mathbf{P}^1_C, \mathbf{F}_p) = \mathbf{1} \oplus \mathbf{1}[-2]$
- $\mathcal{J}^i(\pi(1,\omega^{-1})) = 0$ for all $i \neq 1$.
- $\mathcal{J}^i(\pi(\omega^{-1},1)) = 0$ for all $i \neq 0,1$, and $\mathcal{J}^0(\pi(\omega^{-1},1)) = \mathbf{1}$.
- $\mathcal{J}^i(\mathrm{St}) = 0$ for all $i \neq 1$.

The long exact sequence in $\mathcal{J}^i$'s associated with the sequence (†) then degenerates to a short exact sequence

$$0 \to \mathcal{J}^1(\pi(\omega^{-1},1)) \to \mathcal{J}^1(\mathrm{St}) \to \mathbf{1} \to 0 \quad (\dagger').$$

Next, observe that $\mathcal{J}^2$ is right-exact, so applying it to the exact sequence (‡) gives a surjection $\mathbf{1} = \mathcal{J}^2(\mathbf{1}) \twoheadrightarrow \mathcal{J}^2(\kappa)$. This shows that $\dim_{D^\times} \mathcal{J}^2(\kappa) = 0$. Using Proposition 4.9 and the remark immediately afterwards, we then deduce from the vanishing $\mathcal{S}^0_G(\kappa) = 0$ that in fact $\mathcal{J}^2(\kappa) = 0$. Note that Hu-Wang proved the vanishing $\mathcal{J}^2(\kappa) = 0$ by global methods [HW22, Corollary 8.19], while our argument is purely local. With this vanishing in hand, we see that the long exact sequence in $\mathcal{J}^i$'s associated with the sequence (‡) degenerates to a short exact sequence

$$0 \to \mathcal{J}^1(\kappa) \to \mathcal{J}^1(\pi(1,\omega^{-1})) \to \mathbf{1} \to 0 \quad (\ddagger').$$

Next, we observe that the representations St and $\pi(1,\omega^{-1})$ both satisfy the conditions of Proposition 4.5, with $c = 1$ and $\dim_G(-) = 1$ in both cases. Applying that proposition, we get duality isomorphisms

$$\mathcal{S}^i_{D^\times}(\mathcal{J}^1(\mathrm{St})) \cong \mathcal{J}^i(\kappa)$$

and

$$\mathcal{S}^i_{D^\times}(\mathcal{J}^1(\pi(1,\omega^{-1}))) \cong \mathcal{J}^i(\pi(\omega^{-1},1))$$

for all $i \geq 0$. When $i = 0$ these yield no new information. The vanishing of the right-hand sides for $i \geq 2$ implies that $\mathcal{J}^1(\mathrm{St})$ and $\mathcal{J}^1(\pi(1,\omega^{-1}))$ are of dimension ≤ 1; the exact sequences $(\dagger')$ and $(\ddagger')$ then imply that $\mathcal{J}^1(\kappa)$ and $\mathcal{J}^1(\pi(\omega^{-1},1))$ are also of dimension ≤ 1. Finally, in the case $i = 1$, the left-hand side of each isomorphism is either identically zero, or Cohen-Macaulay of dimension one. This yields a dichotomy: either

(1) $\mathcal{J}^1(\mathrm{St}) = \mathcal{J}^1(\pi(1,\omega^{-1})) = \mathbf{1}$ and $\mathcal{J}^1(\pi(\omega^{-1},1)) = \mathcal{J}^1(\kappa) = 0$, or

(2) all four representations $\mathcal{J}^1(\mathrm{St})$, $\mathcal{J}^1(\pi(1,\omega^{-1}))$, $\mathcal{J}^1(\pi(\omega^{-1},1))$, and $\mathcal{J}^1(\kappa)$ are one-dimensional, and the last two are Cohen-Macaulay.

Just as in the situation of Theorem 1.8, we don't currently see how to rule out (1) by a purely local argument. However, using [HW22, Proposition 8.18. (i)] and their exact sequence (8.13), it is easy to see that (1) does not occur. We leave the details to the interested reader.

5. A partial Künneth formula

In this section only, we change notation slightly, to emphasize the dependence of the Jacquet-Langlands functor on the integer n. More precisely, we write $\mathcal{J}_n$ for the Jacquet-Langlands functor from $\mathcal{D}(\mathrm{GL}_n(F))$ towards $\mathcal{D}(D^\times_{1/n})$.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

For any integer $0 < d < n$, let $P_{n-d,d} \subset \mathrm{GL}_n(F)$ be the usual upper-triangular parabolic with Levi $\mathrm{GL}_{n-d}(F) \times \mathrm{GL}_d(F)$. Let $\sigma_1 \boxtimes \sigma_2$ be a smooth representation of the Levi, and let $\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2)$ be the usual parabolic induction to a smooth representation of $\mathrm{GL}_n(F)$. The main result in this section is an alternative formula for $\mathcal{J}_n(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2))$, in terms of the cohomology of an auxiliary space with coefficients in a certain sheaf built from σ_1 and $\mathcal{J}_d(\sigma_2)$. This builds on the analysis begun by the first author in [JLH21, Appendix A]. When $d = 1$, this formula is implicit in [JLH21], as explained below.

To state the main result, we need to introduce the key auxiliary space. In what follows we will freely use some standard notation related to relative Fargues-Fontaine curves. In particular, for any perfectoid space $S/\mathrm{Spa}C$, we have the adic relative Fargues-Fontaine curve $\mathcal{X}_{S^\flat}$, which comes with a canonical closed immersion $\iota : S \hookrightarrow \mathcal{X}_{S^\flat}$ realizing S as a closed Cartier divisor. We write $\mathcal{O}(1/n)$ for the usual stable vector bundle of rank n and degree 1, which is defined functorially over any $\mathcal{X}_{S^\flat}$.

Definition 5.1. Let $W_{n,d}$ denote the moduli space over $\mathrm{Spd}C$ parametrizing quotients of $\mathcal{O}(1/n)$ which are isomorphic to $\mathcal{O}(1/d)$ after pullback to any geometric point.

Note that $W_{n,d}$ carries a natural $\underline{D_{1/n}^\times}$-action by bundle automorphisms (Recall that the automorphism sheaf $S \mapsto \mathrm{Aut}(\mathcal{O}(1/n)^{\oplus m}/\mathcal{X}_{S^\flat})$ is exactly the sheaf $\underline{\mathrm{GL}_m(D_{1/n}^\times)}$.). It is easy to see that $W_{n,d}$ is a spatial diamond, and we will see below that it is proper over $\mathrm{Spd}C$. It is also easy to see that $W_{n,d}$ can be defined equivalently as the moduli space of *sub*bundles of $\mathcal{O}(1/n)$ which are isomorphic to $\mathcal{O}^{n-d}$ after pullback to any geometric point. Trivializing the subbundle in this alternative description gives a canonical $\mathrm{GL}_{n-d}(F)$-torsor over $W_{n,d}$, corresponding to a canonical map $h_1 : W_{n,d} \to B\underline{\mathrm{GL}_{n-d}(F)}$. On the other hand, trivializing the quotient bundle in the original moduli problem gives a canonical $D_{1/d}^\times$-torsor over $W_{n,d}$, corresponding to a canonical map $h_2 : W_{n,d} \to B\underline{D_{1/d}^\times}$.

With these preparations in hand, the main result in this section can be stated as follows.

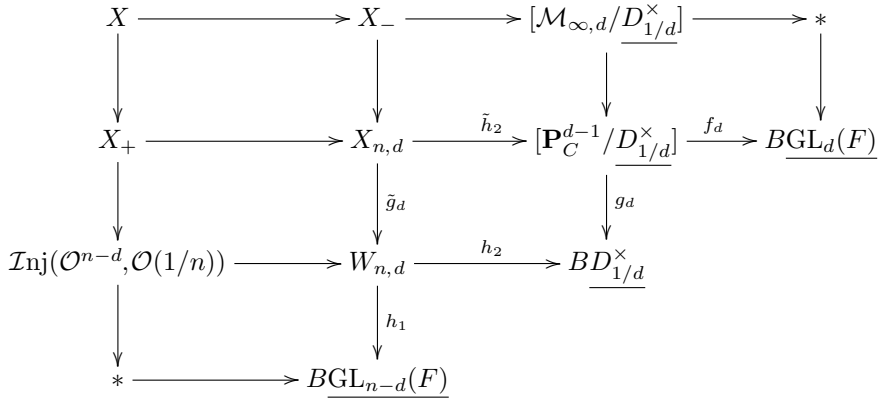
Theorem 5.2. *For any $\sigma_1 \in \mathcal{D}(\mathrm{GL}_{n-d}(F))$ and $\sigma_2 \in \mathcal{D}(\mathrm{GL}_d(F))$, there is a natural isomorphism*

$$\mathcal{J}_n(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2)) \cong R\Gamma(W_{n,d}, h_{1,\text{ét}}^*\sigma_1 \otimes h_{2,\text{ét}}^*\mathcal{J}_d(\sigma_2)).$$

When $d = 1$, this specializes to the isomorphism $\mathcal{J}_n(\mathrm{Ind}_{P_{n-1,1}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2)) \cong R\Gamma(\mathcal{M}_{P(F)}, \mathcal{F}_{\sigma_1 \boxtimes \sigma_2})$ implicit in the proof of [JLH21, Theorem 5.3.1] (where we freely use some notation from loc. cit.). Indeed, $\mathcal{M}_{P(F)} = X_{n,1} = W_{n,1}$ by the discussion below, and a direct calculation shows that $\mathcal{F}_{\sigma_1 \boxtimes \sigma_2} \cong h_1^*\sigma_1 \otimes h_2^*\sigma_2$, noting that $\mathcal{J}_1(\sigma_2) = \sigma_2$ has no effect.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

The proof relies on several auxiliary spaces, which we find helpful to collect into a large diagram



of small v-stacks, where every square is cartesian. Here $* = \mathrm{Spd}\, C$ for brevity, and $W_{n,d}$ and h_1, h_2 are as discussed above. Let us explain the parts of this diagram we have not yet discussed (We will not need every object in this diagram, but the reader might find it helpful to know that they exist.). Let $U \subset P_{n-d,d}$ be the unipotent radical, so $U \cdot \mathrm{GL}_{n-d}(F)$ and $U \cdot \mathrm{GL}_d(F)$ are both closed normal subgroups of $P_{n-d,d}$ with intersection U. Then by definition, $X = \mathcal{M}_\infty/\underline{U}$, $X_+ = \mathcal{M}_\infty/\underline{U \cdot \mathrm{GL}_d(F)}$, $X_- = \mathcal{M}_\infty/\underline{U \cdot \mathrm{GL}_{n-d}(F)}$, and $X_{n,d} = \mathcal{M}_\infty/\underline{P_{n-d,d}}$, whence the upper left square is cartesian. By [JLH21, Proposition A.3.3], the space $\underline{X_{n,d}}$ admits an explicit alternative description: it is the moduli space sending any perfectoid space S/C to (isomorphism classes of) diagrams $\mathcal{O}(1/n) \twoheadrightarrow \mathcal{E} \hookleftarrow \mathcal{F}$ of vector bundles over $\mathcal{X}_{S^\flat}$ such that $\mathcal{E} \simeq \mathcal{O}(1/d)$ and $\mathcal{F} \simeq \mathcal{O}^d$ at all geometric points, and such that $\mathrm{coker}(\mathcal{F} \to \mathcal{E}) \simeq \iota_* W$ for some rank one projective $\mathcal{O}_S$-module W, where $\iota : S \to \mathcal{X}_{S^\flat}$ is the usual divisor. The map $\tilde{h}_2$ corresponds to remembering only $\mathcal{E} \hookleftarrow \mathcal{F}$, while $\tilde{g}_d$ corresponds to remembering only $\mathcal{O}(1/n) \twoheadrightarrow \mathcal{E}$. The cartesian square spanned by g_d and h_2 is then a consequence of [JLH21, Proposition A.3.4]. Note that g_d is proper and p-bounded, since it comes from a proper map of rigid analytic spaces pro-étale-locally on the target. Note also that $X_{n,d} \to W_{n,d}$ is proper and surjective, while $X_{n,d} \to *$ is proper by [JLH21, Proposition A.3.2], so $W_{n,d} \to *$ is proper as claimed in the discussion preceding Theorem 5.2.

Lemma 5.3. *For any $M_1 \in \mathcal{D}^{\mathrm{sm},a}_\square(\mathrm{GL}_{n-d}(F), \mathcal{O}_C/p)^\varphi$ and $M_2 \in \mathcal{D}^{\mathrm{sm},a}_\square(\mathrm{GL}_d(F), \mathcal{O}_C/p)^\varphi$, there is a natural isomorphism*

$$\mathcal{J}_{n,\square}(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(M_1 \boxtimes M_2)) \cong R\Gamma_\square(X_{n,d}, \tilde{g}_d^* h_1^* M_1 \otimes \tilde{h}_2^* f_d^* M_2)$$

in $\mathcal{D}^{\mathrm{sm},a}_\square(D^\times, \mathcal{O}_C/p)^\varphi$.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Proof. Consider the commutative diagram

$$\begin{array}{ccccc}
X_{n,d} & & \longrightarrow & & * \\
\downarrow{\scriptstyle \tilde{\alpha}} & & & & \downarrow{\scriptstyle \alpha} \\
[X_{n,d}/\underline{D}^{\times}_{1/n}] & \xrightarrow{\overline{\pi}_{\mathrm{GH}}} & [\mathbf{P}^{n-1}_C/\underline{D}^{\times}_{1/n}] & \xrightarrow{g_n} & B\underline{D}^{\times}_{1/n} \\
\downarrow{\scriptstyle r} & & \downarrow{\scriptstyle f_n} & & \\
B\underline{P_{n-d,d}} & \xrightarrow{q} & B\underline{\mathrm{GL}_n(F)} & & \\
\downarrow{\scriptstyle s} & & & & \\
B\underline{\mathrm{GL}_{n-d}(F)} \times B\underline{\mathrm{GL}_d(F)} & & & &
\end{array}$$

where the square and rectangle are both cartesian. By definition, $\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(M_1 \boxtimes M_2) = q_* s^*(M_1 \boxtimes M_2)$ and $\mathcal{J}_{n,\square} = g_{n*} f_n^*$, so

$$\mathcal{J}_{n,\square}(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(M_1 \boxtimes M_2)) = g_{n*} f_n^* q_* s^*(M_1 \boxtimes M_2),$$

which by proper base-change transforms to

$$\begin{aligned} &= g_{n*}(\overline{\pi}_{\mathrm{GH}})_* r^* s^*(M_1 \boxtimes M_2) \\ &= R\Gamma_{\square}(X_{n,d}, r^* s^*(M_1 \boxtimes M_2)), \end{aligned}$$

where the $R\Gamma_{\square}$ on the last line denotes the $D^{\times}_{1/n}$-equivariant cohomology. To finish the proof it remains to identify the $D^{\times}_{1/n}$-equivariant sheaf $r^* s^*(M_1 \boxtimes M_2)$ on $X_{n,d}$. After forgetting the $D^{\times}_{1/n}$-equivariance of the sheaf, i.e. after applying $\tilde{\alpha}^*$, we get

$$\begin{aligned} \tilde{\alpha}^* r^* s^*(M_1 \boxtimes M_2) &= (s \circ r \circ \tilde{\alpha})^*(M_1 \boxtimes M_2) = ((h_1 \circ \tilde{g}_d) \times (f_d \circ \tilde{h}_2))^*(M_1 \boxtimes M_2) \\ &= \tilde{g}_d^* h_1^* M_1 \otimes \tilde{h}_2^* f_d^* M_2, \end{aligned}$$

using that $s \circ r \circ \tilde{\alpha} = (h_1 \circ \tilde{g}_d) \times (f_d \circ \tilde{h}_2)$. This gives the result. □

Proof of Theorem 5.2. By some straightforward manipulations with the Riemann-Hilbert correspondence, we easily reduce to proving the solid version of the theorem. In other words, we need to show that for any $M_1 \in \mathcal{D}_{\square}^{\mathrm{sm},a}(\mathrm{GL}_{n-d}(F), \mathcal{O}_C/p)^{\varphi}$ and $M_2 \in \mathcal{D}_{\square}^{\mathrm{sm},a}(\mathrm{GL}_d(F), \mathcal{O}_C/p)^{\varphi}$, there is a natural isomorphism

$$\mathcal{J}_{n,\square}(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(M_1 \boxtimes M_2)) \cong R\Gamma_{\square}(W_{n,d}, h_1^* M_1 \otimes h_2^* \mathcal{J}_{d,\square}(M_2)).$$

By the previous lemma, we already have a canonical isomorphism

$$\mathcal{J}_{n,\square}(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(M_1 \boxtimes M_2)) \cong R\Gamma_{\square}(X_{n,d}, \tilde{g}_d^* h_1^* M_1 \otimes \tilde{h}_2^* f_d^* M_2).$$

We then compute that

$$\begin{aligned} R\Gamma_{\square}(X_{n,d}, \tilde{g}_d^* h_1^* M_1 \otimes \tilde{h}_2^* f_d^* M_2) &\cong R\Gamma_{\square}(W_{n,d}, \tilde{g}_{d*}(\tilde{g}_d^* h_1^* M_1 \otimes \tilde{h}_2^* f_d^* M_2)) \\ &\cong R\Gamma_{\square}(W_{n,d}, h_1^* M_1 \otimes \tilde{g}_{d*} \tilde{h}_2^* f_d^* M_2)) \\ &\cong R\Gamma_{\square}(W_{n,d}, h_1^* M_1 \otimes h_2^* g_{d*} f_d^* M_2)). \end{aligned}$$

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

Here the first isomorphism is trivial, the second follows from the projection formula, and the third follows from proper base change. Since $g_{d*}f_d^*M_2 = \mathcal{J}_{d,\square}(M_2)$ by definition, this gives the desired result. □

By work of Johansson-Ludwig [JLH21, Theorem A], $W_{n,1}$ is a perfectoid space for all n. The proof of [JLH21, Theorem A] is a highly nontrivial argument using the p-adic geometry of global Shimura varieties. In the special case $n = 2$, the first author reproved this result by purely local methods [JLH21, §A.4]. However, the arguments in [JLH21, §A.4] easily generalize to show that $W_{n,n-1}$ is also a perfectoid space for all n. It thus seems natural to make the following conjecture.

Conjecture 5.4. *For all integers $0 < d < n$, $W_{n,d}$ is a perfectoid space.*

This conjecture implies a vanishing result.

Proposition 5.5. *Suppose $W_{n,d}$ is a perfectoid space. Then for any admissible smooth representations $\sigma_1 \in \mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(\mathrm{GL}_{n-d}(F))$ and $\sigma_2 \in \mathrm{Rep}^{\mathrm{sm}}_{\mathbf{F}_p}(\mathrm{GL}_d(F))$, $\mathcal{J}_n^i(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2)) = 0$ for all $i > d+n-2$.*

Note that $\mathcal{J}_n^i$ trivially vanishes in degrees $> 2n-2$, while this result gives vanishing in degrees $> d+n-2 = 2n-2-(n-d)$, so we improve on the trivial bound by an interval of length $n-d$. When $d = 1$, this specializes to the main vanishing theorem in [JLH21] (using of course the perfectoidness of $W_{n,1}$ proved in loc. cit.).

Proof. By the primitive comparison theorem plus some small additional arguments, we get an almost isomorphism

$$\mathcal{J}_n^i(\mathrm{Ind}_{P_{n-d,d}}^{\mathrm{GL}_n(F)}(\sigma_1 \boxtimes \sigma_2)) \otimes \mathcal{O}_C/p \cong H^i(W_{n,d}, h_1^*\sigma_1 \otimes h_2^*\mathcal{J}_d(\sigma_2) \otimes \mathcal{O}^+/p).$$

It thus suffices to show that $H^i(W_{n,d}, h_1^*\sigma_1 \otimes h_2^*\mathcal{J}_d(\sigma_2) \otimes \mathcal{O}^+/p)$ is almost zero in the claimed range of degrees. Recall that $\mathcal{J}_d(\sigma_2)$ is concentrated in degrees $\leq 2d-2$. Note also that $W_{n,d}$ is perfectoid by assumption, so the functor $H^i(W_{n,d}, - \otimes \mathcal{O}^+/p)$ on overconvergent étale $\mathbf{F}_p$-sheaves is almost zero for all $i > \dim W_{n,d} = n-d$, by the usual almost equality of étale and analytic cohomology. The evident spectral sequence then gives almost vanishing in degrees $> 2d-2+n-d = d+n-2$, whence the claim. □

When $n = 3$, $W_{3,1}$ and $W_{3,2}$ are both perfectoid spaces. Combining this observation with the previous proposition, the right-exactness of $\mathcal{J}_3^4$, and the classification of irreducible admissible representations of $\mathrm{GL}_3(F)$ [Her11, Theorem 1.1], we get the following vanishing result.

Corollary 5.6. *Let π be an irreducible admissible representation of $\mathrm{GL}_3(F)$ which is not supersingular, and whose underlying vector space is infinite-dimensional. Then $\mathcal{J}^4(\pi) = 0$.*

Acknowledgments. The idea that $\mathcal{J}$ should satisfy some form of duality was suggested to DH by Judith Ludwig during a visit to Heidelberg in January 2020. We heartily thank Judith for this key initial suggestion. We also sincerely thank Peter Scholze for some very helpful conversations about this material, and for suggesting that the

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

authors might profitably collaborate. Some of these results were discovered during the Oberwolfach workshop 'Nonarchimedean geometry and applications' in February 2022, and we sincerely thank the organizers for the opportunity to participate in this event. We also thank Gabriel Dospinescu, Eugen Hellmann, Yongquan Hu, Christian Johansson, Vytautas Paškūnas, Claus Sorensen, and Bogdan Zavyalov for some helpful conversations.

Competing interests. The authors confirm that they have no competing interests.

Data availability statement. Not applicable.

References

[ABM24] Anschütz J, Le Bras A-C and Mann L (2024) 'A 6-functor formalism for solid quasi-coherent sheaves on the Fargues-Fontaine curve'. https://arxiv.org/abs/2412.20968.

[BS17] Bhatt B and Scholze P (2017) Projectivity of the Witt vector affine Grassmannian. In: *Invent. Math.* **209**(2), 329–423. doi: 10.1007/s00222-016-0710-4.

[Dat12] Dat J-F (2012) Un cas simple de correspondance de Jacquet-Langlands modulo ℓ. *Proc. Lond. Math. Soc. (3)* **104**(4). With an appendix by Vignéras M-F, pp. 690–727. doi: 10.1112/plms/pdr043.

[EP20] Emerton M and Paškūnas V (2020) On the density of supercuspidal points of fixed regular weight in local deformation rings and global Hecke algebras. *J. É c. Polytech. Math.* **7**, 337–371. doi: 10.5802/jep.119.

[Far16] Fargues L (2016) 'Geometrization of the local Langlands correspondence: an overview'. https://webusers.imj-prg.fr/~laurent.fargues/Geometrization_review.pdf.

[FS21] Fargues L and Scholze P (2021) 'Geometrization of the local Langlands correspondence'. https://arxiv.org/abs/2102.13459.

[Fus21] Fust P (2021) 'Continuous cohomology and Ext-groups'. https://arxiv.org/abs/2106.04473.

[GHW22] Gulotta D, Hansen D and Weinstein J (2022) 'An enhanced six-functor formalism for diamonds and v-stacks'. http://www.davidrenshawhansen.com/enhanced.pdf.

[Her11] Herzig F (2011) 'The classification of irreducible admissible mod p representations of a p-adic GL_n'. *Invent. Math.* **186**(2), 373–434. doi: 10.1007/s00222-011-0321-z.

[Hub96] Huber R (1996) *É tale Cohomology of Rigid Analytic Varieties and Adic Spaces*. Aspects of Mathematics 30. Berlin: Vieweg+Teubner Verlag.

[HW22] Hu Y and Wang H (2022) 'On some mod p representations of quaternion algebra over $\mathbf{Q}p$'. https://arxiv.org/abs/2201.01464.

[JLH21] Johansson C, Ludwig J and Hansen D (2021) 'A quotient of the Lubin-Tate tower II'. *Math. Ann.* **380**(1–2), 43–89. doi: 10.1007/s00208-020-02104-3.

[Koh17] Kohlhaase J (2017) Smooth duality in natural characteristic. *Adv. Math.* **317**, 1–49. doi: 10.1016/j.aim.2017.06.038.

[Laz65] Lazard M (1965) Groupes analytiques p-adiques. *Inst. Hautes É tudes Sci. Publ. Math.* **26**, 389–603. http://www.numdam.org/item?id=PMIHES_1965__26__389_0.

[Lud17] Ludwig J (2017) A quotient of the Lubin-Tate tower. *Forum Math. Sigma* **5**, Paper No. e17, 41. doi: 10.1017/fms.2017.15.

[Lur] Lurie J. *Higher Algebra*.

[Lur18] Lurie J. *Kerodon*. https://kerodon.net. 2018.

[LZ17] Liu Y and Zheng W (2017) 'Enhanced six operations and base change theorem for higher Artin stacks'. https://arxiv.org/abs/1211.5948.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press

[Man22] Mann L (2022) 'A p-adic 6-functor formalism in rigid-analytic geometry'. https://arxiv.org/abs/2206.02022.

[Mat16] Mathew A (2016) 'The Galois group of a stable homotopy theory'. *Adv. Math.* **291**, 403–541. doi: 10.1016/j.aim.2015.12.017.

[Paš22] Paškūnas V (2022) 'On some consequences of a theorem of J. Ludwig'. *J. Inst. Math. Jussieu* **21**(3), 1067–1106. doi: 10.1017/S1474748020000547.

[Sch13] Scholze P (2013) 'p-adic Hodge theory for rigid-analytic varieties'. *Forum Math. Pi* **1**, e1.

[Sch18a] Scholze P (2018) 'Étale cohomology of diamonds'. https://www.math.uni-bonn.de/people/scholze/EtCohDiamonds.pdf.

[Sch18b] Scholze P (2018) 'On the p-adic cohomology of the Lubin-Tate tower'. *Ann. Sci. Éc. Norm. Supér. (4)* **51**(4). With an appendix by Rapoport M, 811–863. doi: 10.24033/asens.2367.

[Sch19] Scholze P (2019) *Lectures on Condensed Mathematics*. https://www.math.uni-bonn.de/people/scholze/Condensed.pdf.

[SS22] Schneider P and Sorensen C (2022) 'Duals in natural characteristic' https://arxiv.org/abs/2202.01800.

[SW00] Symonds P and Weigel T (2000) 'Cohomology of p-adic analytic groups'. In: *New Horizons in Pro-p Groups*. Vol. 184. Progr. Math. Boston, MA: Birkhäuser Boston, pp. 349–410.

https://doi.org/10.1017/S1474748025101503 Published online by Cambridge University Press