Towards generic modularity of higher theta series over global function fields and p-adic local fields

Chaim Avram (Avi) Bettigole Zeff

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2025

Abstract

Towards generic modularity of higher theta series over global function fields and p-adic local fields

Chaim Avram (Avi) Bettigole Zeff

Recent work of Feng–Yun–Zhang constructs higher theta series for unitary groups valued in cycles on the moduli stack of Hermitian shtukas, and proves using Fourier-theoretic methods that these theta series are modular after restriction to the generic fiber and on ℓ-adic cohomology. In the first part of this thesis, motivated by ideas from the relative Langlands program we generalize their construction to a natural class of reductive dual pairs, which contains almost all dual pairs for which such a construction should be expected, and prove generic modularity on ℓ-adic cohomology. In characteristic zero, no analogue of moduli stacks of shtukas with more than one leg exists globally, but such analogues do exist locally over p-adic fields, as defined by Scholze–Weinstein. In the second part of this thesis, we construct special cycles on these stacks, which under certain assumptions should be directly analogous to the higher theta series, and in this case prove generic modularity. Our main tool is a Fourier transform for cohomological correspondences, which we develop in the setting of p-adic geometry, building on ideas of Anschütz–Le Bras and using the power of recent six-functor formalisms due to Scholze and Heyer–Mann.

Table of Contents

Acknowledgements

First and foremost, I would like to thank my advisor Chao Li for his support, guidance, and endless patience. I have learned an incredible amount from him directly, and much of what I've learned from others or independently has nevertheless been inspired or facillitated by discussions with him.

I would also like to thank Yujie Xu for her mentorship and advice. The results of Chapter 2 of this thesis were obtained as the result of a collaboration with her, and I am grateful to her for working with me on them and agreeing to their inclusion here.

I am also grateful to Michael Harris, Dorian Goldfeld, and A. Raghuram for serving on my thesis defense committee and for their helpful questions and comments, and to Johan de Jong, Qiao He, David Ben-Zvi, Yiannis Sakellaridis, Nina Wawrów, Ryan Chen, and especially Tony Feng for a number of helpful discussions relating to the topic of this thesis. I would also like to thank the staff of the Columbia math department for creating a comfortable and welcoming working environment, and my students for suffering my foibles and helping me become a better teacher and thinker.

Finally, I am grateful to Nolan Reilly, Sonia Reilly, and Tristan McLaurin for their friendship and company over the better part of the last decade; Alex Xu and Neo Liang, for getting me into climbing; my more physically distant friends Maya Sankar, Alex LaGrassa, Nyra Yearwood, Leah Petit, Jessie Grosen, and everyone else I have undoubtedly shamefully forgotten to mention, may we see each other more often or at least continue to talk irregularly but at length; the members and organizers of SWC-UAW 2710, for both ensuring that I am able to survive in New York and have

reasonable working conditions and for creating a warm and welcoming community; my parents and brothers for their love, support, and occasional tact; and Mack, for more than I can say.

Chapter 1: Introduction and background

1.1 Background and motivation

The connection between theta series and L-functions goes back to Riemann's proof of the functional equation for the Riemann zeta function in [1] using the modularity of the Jacobi theta function

$$\theta(z) = \sum_{n \in \mathbb{Z}} q^{n^2} \tag{1.1}$$

where $q = e^{2\pi i z}$, in turn proved by Poisson as one of the first applications of Poisson summation. A century later, building on work of Siegel, Weil proved the Siegel–Weil formula relating more general theta series, constructed as generating series of representation numbers of quadratic or Hermitian forms, to a different analytic quantity: Siegel–Eisenstein series $E(s, z)$, or more precisely their central values in s. Equivalently, taking Fourier coefficients, the Siegel–Weil formula relates individual representation numbers to central values of Fourier coefficients of Siegel–Eisenstein series. In particular, as Siegel–Eisenstein series are modular forms in z, the theta series must be as well, generalizing Poisson's result; this can also be proven more directly via Poisson summation. The Siegel–Weil formula can in turn be used to relate the Petersson inner product of theta functions (more precisely, theta lifts) to special values of L-functions; this is the Rallis inner product formula [2, 3].

A more modern approach goes back to Kudla–Millson [4]: they replaced the numerical theta series by an analogue given by a generating series of cohomology classes of "special cycles" on a Shimura variety, and proved its modularity. Building on this work, Kudla [5, 6] and later Kudla–Rapoport [7, 8] introduced a program to study special cycles on Shimura varieties, as well as their integral models, including a generalization of the Siegel–Weil formula relating their degrees to

central values, resp. derivatives, of Fourier coefficients of Siegel–Eisenstein series, generalizing the Siegel–Weil formula, an (arithmetic) inner product formula, and the conjecture that the corresponding theta series are modular.

As the language of quadratic and Hermitian forms suggests, the natural setting for Kudla's program is related to groups of automorphisms of vector spaces equipped with a bilinear pairing of some type, and more precisely to pairs of reductive groups (G, H) called dual reductive pairs. We refer to Section 2.1 for details; loosely speaking, we think of a dual pair as a pair of reductive groups equipped with standard representations Std_G and Std_H such that the tensor product $W = \mathrm{Std}_G \otimes \mathrm{Std}_H$ is naturally equipped with a symplectic pairing, so that $G \times H$ can be viewed as a subgroup of $\mathrm{Sp}(W)$. We sometimes refer to W as the standard symplectic representation of (G, H). The main examples studied in Kudla's program are the case when G and H are both unitary groups and when one of G and H is an orthogonal group and the other is a symplectic group, which we refer to as the unitary and orthogonal cases respectively.

The language of dual pairs is useful to understand how to obtain the theta series a priori, analogous to (1.1), without first describing its Fourier coefficients: choosing a Lagrangian subspace $L \subset \mathrm{Std}_G$, inducing a Lagrangian subspace $U := L \otimes \mathrm{Std}_H \subset W$, there is a canonical representation of $\mathrm{Sp}(W)$ called the Weil representation[1] inducing a representation of $G \times H$ on Schwartz functions on U, possibly after base change, e.g. to the adeles. Fixing a Schwartz function Φ—in practice often the indicator function of a lattice, at least at finite places—the theta series is then given by

$$(g, h) \mapsto \sum_{x \in U} ((g, h) \cdot \Phi)(x). \tag{1.2}$$

Specializing to simple cases such as $G = \mathrm{Sp}(2) \simeq \mathrm{SL}(2)$ in the global number field case recovers a function on $\mathrm{SL}_2(\mathbb{A})$, which is relatively easy to see is an automorphic form; the dictionary between automorphic forms and modular forms matches this example with the Jacobi theta function (1.1). Replacing $\Phi(x)$ by a special cycle associated to x on a Shimura variety for H (or on its integral

[1] More precisely, it is a representation of the metaplectic group $\mathrm{Mp}(W)$, which is a cover of $\mathrm{Sp}(W)$; for each dual pair (G, H) of interest to us this cover splits after pullback to $G \times H$, so we still get a representation ω of $G \times H$.

model), we obtain Kudla's arithmetic theta series.

Many of the conjectures of Kudla's program have now been proven in many cases, including some in Kudla's original papers; in particular, the arithmetic Siegel–Weil formula is now known for Fourier coefficients corresponding to nonsingular matrices by the work of Li–Zhang for in the unitary case [9] and orthogonal case [10], and the arithmetic inner product formula in the unitary case by the work of Li–Liu [11, 12]. For orthogonal groups the modularity conjecture has been proven (on the generic fiber) by Zhang [13] and Bruinier–Westerhold-Raum [14], and for special cycles of codimension 1 on the integral model by [15] in the unitary case, as well as a wide variety of other cases due to many other authors.

Remark 1.1.1. One can take a different perspective on reductive dual pairs: the standard symplectic representation W with its action of $G \times H$ is an example of a hyperspherical variety as defined in [16], and hence fits into the framework of the relative Langlands program. In particular, there should be an automorphic form on $G \times H$ associated to the hyperspherical variety W, namely the period function of W, which looks formally similar to the definition of the theta series (1.2). This perspective is discussed in greater detail in Section 2.1; for the moment, we note only that in practice, the definition of the period depends on the auxiliary datum of a choice of polarization of the hyperspherical variety, although the resulting period function is expected to be independent of the choice. Such a polarization need not always exist; when it does, we say that the hyperspherical variety is polarizable. In the non-polarizable case, it is not yet clear how to define the period.

For our case of $G \times H$ acting on W, a polarization of W amounts to a G-stable Lagrangian subspace of W. When such a polarization exists, so when W is polarizable, we say that (G, H) is of type II; otherwise we say it is of type I. As the period can be well understood via the methods of [16] in the type II case, the more interesting case for us is when (G, H) is of type I, so W is non-polarizable. Nevertheless, when we can choose a Lagrangian $L \subset \mathrm{Std}_G$—not necessarily G-stable!—we can still define a theta series following (1.2), which we still hope to be independent of this choice of L. We elaborate on this below.

1.2 Main results

We now turn to the function field setting, and so fix a curve X over a finite field $k = \mathbb{F}_q$. We will at times want the additional data of an étale cover $X' \to X$ of degree either 1 or 2, and we write σ for the generator of $\mathrm{Aut}(X'/X)$.

In [17], Feng–Yun–Zhang introduce a "higher theta series," formed as a generating series of special cycles on stacks of unitary shtukas with r legs for each $r \geq 0$. For $r = 0$ and $r = 1$, this is analogous to the theta series and arithmetic theta series valued in cycles on Shimura varieties and their integral model respectively; for $r \geq 2$ there is no characteristic 0 analogue, due to the lack of a notion of "shtukas with multiple legs" over $\mathrm{Spec}\,\mathbb{Z}$. Feng–Yun–Zhang conjecture that the theta series is a modular form for the (anti-)unitary group, and prove in [18] that it satisfies a "higher Siegel–Weil formula" relating the degrees of (nonsingular) terms to higher central derivatives of a Siegel–Eisenstein series. In [19], they prove the modularity conjecture after restriction to the generic fiber ("generic modularity") and imposing certain additional technical restrictions, most prominently after passing from Chow groups to ℓ-adic cohomology, removed in [20] by Feng–Khan. In the function field setting, the theta series a priori depends on (1) a G-bundle $\mathcal{G}$ and (2) a Lagrangian subbundle $\mathcal{E} \subset \mathcal{G}$; the modularity conjecture states that it is in fact independent of $\mathcal{E}$.

Motivated by Kudla's program, our first goal is to generalize the construction and generic modularity of the higher theta series to other reductive dual pairs (G, H), beyond the unitary dual pair considered in [17, 19]. In particular as in Remark 1.1.1 from the point of view of the relative Langlands program the theta series can be understood for $r = 0$ as the period of $G \times H$ acting on its standard symplectic representation W, and we speculate that the $r > 0$ case should be thought of as a "period cycle" generalizing the period function. The modularity for $r = 0$ is then expected from the relative Langlands philosophy, as the period should be independent of auxiliary polarization-like data such as the choice of Lagrangian $\mathcal{E}$, and the modularity for general r can be thought of as a prediction of a conjectural "higher relative Langlands program."

The presence of the generating series rather than constructing the period cycle directly as a

special cycle ultimately comes from the Weil representation, which is necessary since the hyperspherical pair $(G \times H, W)$ is non-polarizable when (G, H) is the unitary dual pair; thus our main case of interest is when $(G \times H, W)$ is non-polarizable, which is to say (G, H) is a reductive dual pair of type I. Fixing a Lagrangian subspace $L \subset \mathrm{Std}_G$—not necessarily G-stable!—we let $P \subset G$ be the parabolic subgroup stabilizing L, so that Bun_P parametrizes G-bundles $\mathcal{G}$ together with Lagrangian subbundles $\mathcal{E} \subset \mathcal{G}$. We write m for the rank of L, so Std_G has rank $2m$; on the other side, we write n for the rank of Std_H. In this setting, we define for each $\mathcal{E}$ a special cycle $\mathcal{Z}^r_{\mathcal{E}} \to \mathrm{Sht}^r_H$ on the stack of shtukas for the group H, which decomposes as the disjoint union of substacks $\mathcal{Z}^r_{\mathcal{E}}(a)$ for a ranging over the space $\mathcal{A}_{\mathcal{E}}(k)$ of (skew-)Hermitian pairings on $\mathcal{E}$, depending on the dual pair (G, H). The pairing on $\mathcal{G}$ allows us to view $\mathcal{G}$ as an extension of $\mathcal{E}^*$ by $\mathcal{E}$ up to certain twists, giving rise to an Ext class $e_{\mathcal{G},\mathcal{E}}$, which we can pair with $a \in \mathcal{A}_{\mathcal{E}}(k)$ to obtain an element of $\mathbb{F}_q$. Choosing a nontrivial additive character $\psi : \mathbb{F}_q \to \overline{\mathbb{Q}}_\ell^\times$, we can then construct the higher theta series

$$\widetilde{Z}^r_m : \mathrm{Bun}_P(k) \to \mathrm{Ch}_*(\mathrm{Sht}^r_H)$$

as a generating series

$$\widetilde{Z}^r_m(\mathcal{G}, \mathcal{E}) = \chi(\det \mathcal{E}) q^{\frac{n}{2}(\deg \mathcal{E} - \deg \omega_X)} \sum_{a \in \mathcal{A}_{\mathcal{E}}(k)} \psi(\langle e_{\mathcal{G},\mathcal{E}}, a \rangle)[\mathcal{Z}^r_{\mathcal{E}}(a)],$$

where $[\mathcal{Z}^r_{\mathcal{E}}(a)]$ denotes the virtual cycle class, defined in Section 2.2.2.

We write $\widetilde{Z}^r_{m,\eta}$ for the image under the ℓ-adic cycle class and restriction of $\widetilde{Z}^r_m$ to (approximately) the generic fiber of Sht^r_H; see Section 2.3 for more details. Our main result is then the following.

Theorem 1.2.1 (Theorem 2.3.2)**.** *(a) If (G, H) is a type I reductive dual pair such that the symplectic representation W admits an H-stable Lagrangian subspace and is not of excluded type, then the higher theta series $\widetilde{Z}^r_{m,\eta}$ descends to $\mathrm{Bun}_G(k)$, i.e. $\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$.*

(b) *If (G, H) is of excluded type, then $\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$ and on the class of $\mathcal{E}$ in* $\pi_0 \operatorname{LGr}(m, 2m)$.

Here $\operatorname{LGr}(m, 2m)$ is the Lagrangian Grassmannian, parametrizing Lagrangian subspaces of the rank $2m$ vector space Std_G, and $\pi_0 \operatorname{LGr}(m, 2m)$ is its set of connected components. The pairs of excluded type are described in Remark 2.1.2; the fundamental example is $(G, H) = (\mathrm{O}(m, m), \mathrm{Sp}(n))$ where $\mathrm{O}(m, m)$ is the split orthogonal group (i.e. the automorphism group of a rank $2m$ orthogonal space with a rank m isotropic subspace) and m is even. The obstruction in this case to full generic modularity is made clear in Lemma 2.4.1 and Proposition 2.4.2; the issue is with a reduction to transverse Lagrangian subspaces, which does not completely work in the excluded case. After this restriction all methods work uniformly.

Our approach broadly follows [19], using the theory of derived and arithmetic Fourier transforms, modified as needed for our generality. A key technical observation is Lemma 2.4.3: although we allow general type I dual pairs, so the vector spaces with which we are concerned in general may be symmetric or anti-symmetric, Hermitian or skew-Hermitian, a certain key vector space V will always be Hermitian regardless of the choice of dual pair. This is key to our analysis, and is what allows the generality of Theorem 2.3.2.

Although our main concern is with the type I (non-polarizable) case, we can still make many of the analogous definitions and arguments in the type II case. As suggested above, here the generating series is replaced by a single class, up to normalization; we denote the resulting class as $[\mathcal{Z}^{r,\mathrm{norm}}_{\mathcal{E}}]$. Then we have the following result:

Theorem 1.2.2 (Theorem 2.6.2)**.** *For each $\mathcal{E} \in \mathrm{Bun}_G(k)$, we have $[\mathcal{Z}^{r,\mathrm{norm}}_{\mathcal{E}}] = [\mathcal{Z}^{r,\mathrm{norm}}_{\sigma^*\mathcal{E}^*}]$.*

Remark 1.2.3. Let $U \subset W$ be a fixed $G \times H$-stable Lagrangian subspace, which exists since (G, H) is type II. When $r = 0$, $[\mathcal{Z}^0_{\mathcal{E}}]$ is the function on $\mathrm{Bun}_H(k)$ sending $\mathcal{F}$ to $\#\operatorname{Hom}_D(\mathcal{E}, \mathcal{F})$, which as a function on $H(\mathbb{A})$ is given by

$$h \mapsto \sum_{u \in U(F)} \Phi(g^{-1}uh)$$

where $g \in G(\mathbb{A})$ is in the class of $\mathcal{E} \in \mathrm{Bun}_G(k) \simeq G(F)\backslash G(\mathbb{A})/G(\widehat{O})$ and Φ is the indicator

function of the standard lattice $U(\widehat{O})$ in $U(\mathbb{A})$. In particular, note that this agrees with the period function for the action of $G \times H$ on U as defined in [16, §10.3], and the involution $\mathcal{E} \mapsto \sigma^*\mathcal{E}^*$ corresponds to replacing U by its dual $\sigma^*U^* \subset W$. Therefore in this case Theorem 2.6.2 recovers [16, Example 10.9.1], which there is proven by Poisson summation. In fact our proof for $r = 0$ is essentially by Poisson summation as well, and for arbitrary r the proof can be viewed as a geometrization of Poisson summation.

The above work (Theorem 1.2.1 and Theorem 1.2.2) is joint with Yujie Xu; the author of this thesis is grateful to her for agreeing to its inclusion here.

We mentioned above that for $r \geq 2$ there is no characteristic 0 analogue of the moduli space of shtukas with r legs, and so we cannot hope to define a higher theta series. This is true globally, over number fields. Locally, however, over p-adic fields it is possible to define stacks of *local* shtukas with arbitrarily many legs: this was done in [21] by Scholze–Weinstein, using the language of p-adic geometry: perfectoid spaces, diamonds, and v-stacks. This motivates the following question: can we work analogously to the function field case to define higher theta series over p-adic local fields, and prove their (generic) modularity?

For concreteness, we restrict to the case where (G, H) is the unitary dual pair, as in [19], although we expect that the techniques used to prove the theorems above should allow us to work similarly with more general dual pairs. The analogues of skew-Hermitian vector bundles on the (absolute) curve are lattices $\mathfrak{G}$ inside of trivial G-torsors $\mathcal{G}$ on the Fargues–Fontaine curve, and so one might hope for a statement to the effect that the theta series, which should a priori depend on a choice of Lagrangian sublattice $\mathfrak{E} \subset \mathcal{E}$, is in fact independent of this choice. However, although morally the two settings are analogous, following the philosophy of Fargues [22] that the local Langlands program should be thought of as geometric Langlands on the Fargues–Fontaine curve, in practice the analogy is not as close as one might hope: for example, stacks of p-adic shtukas are no longer most naturally presented as an intersection of a Hecke correspondence with the graph of Frobenius, as the Frobenius is already incorporated into the Fargues–Fontaine curve. For our purposes, the most important obstruction is less technical: the theta series in the function field case

can be thought of as the convolution of a special cycle $[\mathcal{Z}^r_{\mathcal{E}}]$ with a Gaussian $\psi(\langle -, - \rangle)$. In the p-adic setting, the natural analogue of a Gaussian is no longer given by evaluating a character on a quadratic form but instead by the indicator function of a lattice, which looks quite different!

Nevertheless, in the special case $(\mathfrak{E}_2 \subset \mathcal{E}_2) = \sigma^*(\mathfrak{E}_1 \subset \mathcal{E}_1)^* := (\sigma^*\mathfrak{E}_1^* \subset \sigma^*\mathcal{E}_1^*)$ as in Theorem 1.2.2, the extension class $e_{\mathcal{G},\mathcal{E}}$ vanishes and so the Gaussian term in the theta series is trivial, so this difficulty does not arise. In this case, we can define a special cycle $[\mathcal{Z}^r_{\mathcal{E},\mathfrak{E},\eta,C}]$ on the stack $\mathrm{Sht}^r_{H,\eta,C}$ of p-adic shtukas with r legs after restriction to the generic fiber and base change to a geometric (rather than absolute) point $\operatorname{Spa} C$; this is a class in $H^{\mathrm{BM}}_{2(n-m)r}(\mathrm{Sht}^r_{H,\eta,C})$. Then we can prove the following result:

Theorem 1.2.4 (Theorem 3.6.3)**.** *For all $r \geq 0$,*

$$[\mathcal{Z}^r_{\mathcal{E},\mathfrak{E},\eta,C}] = [\mathcal{Z}^r_{\sigma^*\mathcal{E}^*,\sigma^*\mathfrak{E}^*,\eta,C}].$$

As in the type II case over function fields, it would be desirable to have a more general result beyond this "split case," at least conjecturally, but it is not entirely clear what the right formulation should be. The author hopes to return to this question in future work.

In the case $r = 0$, this class is a zero-cycle on the constant perfectoid space $H(F)/H(\mathcal{O}_F)$, while the pair $(\mathcal{E}, \mathfrak{E})$ is a point of the analogous space $M(F)/M(\mathcal{O}_F)$ for a Levi subgroup M of a parabolic in G.[2] Thus we can think of the special cycles as defining a function on $M(F)/M(\mathcal{O}_F) \times H(F)/H(\mathcal{O}_F)$. The duality property of Theorem 3.6.3 can then be understood as follows: embedding M into G, in general we might hope to get a function (the general theta series for $r = 0$) on $G(F)/G(\mathcal{O}_F) \times H(F)/H(\mathcal{O}_F)$ whose restriction to M is as above. Following the model of the function field case, we expect that this should be done by defining the function on $G(F)/P(\mathcal{O}_F)$ and then showing that it descends to $G(F)/G(\mathcal{O}_F)$. For the piece defined so far, all we can show is the invariance under the longest Weyl element, which on the pair $(\mathcal{E}, \mathfrak{E})$ acts by $(\mathcal{E}, \mathfrak{E}) \mapsto (\sigma^*\mathcal{E}^*, \sigma^*\mathfrak{E}^*)$, so the theorem gives the invariance in this case.

[2]This does not make sense as written, as H and M need not be defined over $\mathcal{O}_F$ but only over F; instead, one should replace $H(\mathcal{O}_F)$ and $M(\mathcal{O}_F)$ with certain compact open subgroups of $H(F)$ and $M(F)$ respectively.

Remark 1.2.5. This can also be viewed as an analogue of Remark 1.2.3 on the Fargues–Fontaine curve, though note that here we are in the type I case; the restrictions to the split case in Theorems 1.2.2 and 3.6.3 are for different reasons.

The broad approach is analogous to the techniques used in the function field case and in [19, 20], but as mentioned above there are a number of technical and substantive differences. We include in Table 1.1 a table of analogies and differences between the function field and p-adic cases, which may help clarify why on the one hand we should expect similar phenomena in the two cases and on the other some of the key difficulties in adapting the method from the function field setting to the p-adic setting.

Global function fields F	p**-adic local fields** F
curve X over a finite field, with $F = k(X)$	Fargues–Fontaine curve X_F
G-bundles on X	G-bundles on X_F
shtukas $\mathcal{E} \dashrightarrow \mathrm{Fr}^* \mathcal{E}$	p-adic shtukas $\mathcal{E} \dashrightarrow \mathcal{E}^b$ for $b \in B(G)$[3]
stack of shtukas is the intersection of the Hecke stack with the graph of Frobenius	stack of shtukas is a cover of a piece of the Hecke stack
{shtukas with zero legs} $\simeq$ {G-bundles} $\simeq G(F)\backslash G(\mathbb{A}_F)/G(\widehat{\mathcal{O}}_F)$	{shtukas with zero legs} $\simeq$ {lattices in the trivial $G(F)$-torsor} $\simeq G(F)/G(\mathcal{O}_F)$
locally free sheaf/perfect complex $\mapsto$ classical/derived vector bundle	locally free sheaf/perfect complex $\mapsto$ Banach–Colmez space
special cycles $\mathcal{Z}^r_{\mathcal{E}} \to \mathrm{Sht}^r_H$ are finite	special cycles $\mathcal{Z}^r_{\mathcal{E}} \to \mathrm{Sht}^r_H$ are not finite, but are proper for a suitable six-functor formalism

Table 1.1: Some relevant analogies and differences between global function fields and p-adic local fields

[3] Here $B(G)$ is the Kottwitz set of [23].

For example, rather than derived vector bundles, we use the Fourier transform on Banach–Colmez spaces developed in [24], and develop the Fourier transform of cohomological correspondences in this setting. In fact, we at times need more powerful results from [25], especially for an integral analogue needed to recover results on spaces of shtukas at finite level. We also need to modify the six-functor formalism used to work with $\mathbb{Q}_\ell$-coefficients, rather than torsion coefficients; we use the recent theory of Berkovich motives developed by Scholze in [26], which has all the properties we will need and is analogous to the motivic coefficients used in [20].

1.3 Outline of this thesis

In Chapter 2, we define special cycles on stacks of shtukas for reductive dual pairs, and show that they have well-behaved virtual fundamental classes. Using these classes, we define a higher theta series and conjecture its modularity. Finally, using derived and arithmetic Fourier analysis, we prove Theorem 1.2.1 and Theorem 1.2.2.

In Chapter 3, we develop the theory of the Fourier transform for cohomological correspondences over Banach–Colmez spaces and their integral analogues, and show that it allows us to prove duality results on stacks of p-adic shtukas. We then construct the relevant stacks, define our special cycles, and prove Theorem 1.2.4.

Chapter 2: Generic modularity for type I reductive dual pairs

2.1 Reductive dual pairs

Over any field F, let W be a finite-dimensional symplectic vector space. A reductive dual pair in $\mathrm{Sp}(W)$ is a pair of subgroups (G, H) of $\mathrm{Sp}(W)$ such that G is the centralizer of H, H is the centralizer of G, and the action of both G and H on W is completely reducible. As a symplectic vector space, W always admits Lagrangian subspaces; but there may or may not be Lagrangian subspaces invariant under both the action of G and H. If there does exist such a subspace, we say that (G, H) is type II; otherwise we say (G, H) is type I.

The study of reductive groups acting on symplectic vector spaces is closely related to the relative Langlands program, as developed in [16] building upon the work of [27, 28, 29, 30, 31, 32, 33, 34, 35] and many others. Broadly, the idea of the program is that where the Langlands program can be viewed as an identification between certain "automorphic" objects (or spaces, or categories) associated to a reductive group G over a field of interest (such as a global or local field for number-theoretic purposes, or an algebraically closed field for geometric Langlands) and "spectral" or "Galois" objects associated to the Langlands dual group $\check{G}$, the relative Langlands program is concerned with extending this relationship to certain morphisms between reductive groups. In the usual Langlands program, this is the study of Langlands functoriality. However, the relative Langlands program takes a more expansive view of what is meant by a "morphism": rather than a homomorphism of reductive groups $G \to H$, it should instead be a certain type of symplectic variety equipped with compatible actions of G and H. As this is equivalent to an action of $G \times H$, the study of such morphisms reduces to the study of these symplectic varieties with G-action, which in [16] are called hyperspherical varieties, thought of as morphisms from the trivial group to G or vice versa. Examples include $T^*(G/H)$ where $H \subset G$ is a spherical subgroup, or

suitable symplectic representations of G.

If a hyperspherical variety M for a reductive group G is thought of as a morphism from the trivial group, then functoriality suggests that it should produce an element of the automorphic object associated to G. Over a global field, which will be our principle area of concern in this thesis, this is to say that we should associate to M an automorphic form, inside the space of automorphic forms for G. This is the period function of M, and is closely related to classical periods of automorphic forms. On the spectral side, we similarly produce an object associated to M, which is called its L-function; the connection with special values of classical L-functions is somewhat more obscure, but is explored somewhat further in e.g. [36], and in many cases can be made explicit.

The key idea of relative Langlands duality is that there should be a duality $M \mapsto \check{M}$ between hyperspherical varieties for G and those for $\check{G}$, such that the period of M is identified with the L-function of $\check{M}$. This recovers many classical identities relating periods of automorphic forms to special values of L-functions, such as Hecke's period formula or Waldspurger's formula [37], among many others; we refer to [16, Table 1.5.1] for a more extensive (though still incomplete) list of examples. However, in this thesis we will generally only be concerned with the automorphic side and periods.

In practice, the period and L-function are defined to depend on certain additional data, namely a polarization of M, which for our purposes can be thought of as a G-invariant Lagrangian subvariety. In general, such a subvariety need not exist; when one does, we say that M is polarizable, and it is only under this assumption that the period and L-function are defined in general. Nevertheless, they can be shown to be independent of the choice of polarization, and it is expected that analogous objects can be defined in general.

Consider the case when G is defined over a global function field F over a finite field $k = \mathbb{F}_q$. Then automorphic forms on G are functions on $\mathrm{Bun}_G(k) \simeq G(F)\backslash G(\mathbb{A}_F)/G(\widehat{\mathcal{O}}_F)$, where $\mathbb{A}_F$ is the ring of adeles of F and $\widehat{\mathcal{O}}_F$ the integral adeles. In the polarized case $M = T^*X$ for a spherical G-variety X, the period function can be defined, up to normalization, as the left $G(F)$- and right

$G(\widehat{O}_F)$-invariant function on $G(\mathbb{A}_F)$ defined by

$$g \mapsto \sum_{x \in X(F)} (g \cdot \Phi)(x) \tag{2.1}$$

where Φ is the indicator function of $X(\widehat{O}_F) \subset X(\mathbb{A}_F)$ and the action of g is by precomposition (see [16, §10.3] for a more precise definition). Note that this is closely analogous to the theta series

$$(g, h) \mapsto \sum_{x \in U} ((g, h) \cdot \Phi)(x),$$

with the chosen Lagrangian in the standard symplectic representation U replaced by X and the group action by the Weil representation of $G \times H$. Thus we think of the choice of the polarization X as similar to the choice of a Lagrangian $U \subset W$ (or $L \subset \mathrm{Std}_G$), and of the period as analogous to the theta series for this choice, with independence of polarization analogous to modularity of the theta series.

When our hyperspherical pair is given by $G \times H$ acting on its standard symplectic representation W for a reductive dual pair (G, H), we can at least in some cases make this analogy concrete. It is not always possible to find $L \subset \mathrm{Std}_G$ which is not only a Lagrangian subspace but also a subrepresentation of G; when it is possible, $U = L \otimes \mathrm{Std}_H$ gives a genuine polarization of W, and so we can define the period straightforwardly. This is the case when the dual pair (G, H) is of type II, while the non-polarizable case is type I.

Remark 2.1.1. This philosophy motivates the main restrictions on the pair (G, H): as the period can be well understood via the methods of [16] in the type II case, the more interesting case for us is when (G, H) is of type I, so W is non-polarizable; however we discuss the type II case briefly in Section 2.6.1. Further it must be possible to find an H-stable Lagrangian subspace of W; if e.g. G is a symplectic group, then this is automatic, but if it is e.g. an orthogonal group then it must be split.

We now give a more explicit description of type I dual pairs in our case of interest; a reference

is [38, §3.1]. Let $k = \mathbb{F}_q$ be a finite field, X a smooth projective curve over k, and $F = k(X)$ its function field. We let D be a finite-dimensional division algebra over F with center F', an unramified field extension of F of degree at most 2, equipped with an involution $\sigma : D \to D$ inducing an automorphism $\sigma \in \mathrm{Gal}(F'/F)$ so that $F = (F')^{\sigma=1}$. We write $\nu : X' \to X$ for the étale cover of X corresponding to $F \hookrightarrow F'$.

Let $(W, \langle -, - \rangle)$ be a nondegenerate finite-dimensional symplectic vector space over F, and fix $\epsilon \in \{\pm 1\}$. We let W_1 be a right D-module equipped with a right (D, σ)-linear ϵ-Hermitian pairing $\langle -, - \rangle_1 : W_1 \times W_1 \to D$ and W_2 a left D-module equipped with a left (D, σ)-linear $(-\epsilon)$-Hermitian pairing $\langle -, - \rangle_2 : W_2 \times W_2 \to D$, i.e. for all $d, d' \in D$, $w_1, w_1' \in W_1$, $w_2, w_2' \in W_2$ we have

$$\begin{aligned} (\langle w_1', w_1 \rangle_1) &= \epsilon \langle w_1, w_1' \rangle_1, & \langle dw_1, d'w_1' \rangle &= \sigma(d) \langle w_1, w_1' \rangle_1 d', \\ \sigma(\langle w_2', w_2 \rangle_2) &= -\epsilon \langle w_2, w_2' \rangle_2, & \langle dw_2, d'w_2' \rangle &= d \langle w_2, w_2' \rangle_2 \sigma(d'). \end{aligned}$$

Let $\phi : W_1 \otimes_D W_2 \xrightarrow{\sim} W$ be an F-linear isomorphism such that

$$\langle \phi(w_1 \otimes w_2), \phi(w_1' \otimes w_2') \rangle = \mathrm{Tr}_{D/F}(\sigma(\langle w_1, w_1' \rangle_1) \langle w_2, w_2' \rangle_2),$$

using which we identify $W_1 \otimes_D W_2$ with W. Then

$$(G, H) = (\mathrm{Aut}_D(W_1, \langle -, - \rangle_1), \mathrm{Aut}_D(W_2, \langle -, - \rangle_2))$$

is a reductive dual pair of type I in $\mathrm{Sp}_F(W, \langle -, - \rangle)$, and all dual pairs of type I are of this form.

For example, suppose $D = F'$. If σ acts trivially on F', i.e. $F = F'$, and $\epsilon = 1$, then G is an orthogonal group and H a symplectic group. For $\epsilon = -1$, we exchange G and H, so G is a symplectic group and H an orthogonal group. If σ acts nontrivially on F', so F' is a quadratic extension of F, both sides are unitary groups.

We can also define similitude groups by allowing similitudes in the automorphism groups, i.e. $\widetilde{G} := \{(g, \lambda) \in \mathrm{GL}_D(W_1) | \langle gw, gw' \rangle_1 = \lambda \langle w, w' \rangle_i\}$ with $\lambda = \lambda(g) \in D$ giving the similitude

character $\widetilde{G} \xrightarrow{\lambda} D$, and similarly $\widetilde{H}$ as the similitude group for H. In particular, taking $D = F'$ an unramified quadratic extension of F, $\epsilon = -1$, and W_1 and W_2 of ranks $2m$ and n respectively recovers the case considered in [19], with $\widetilde{G} = \mathrm{GU}^-(2m)$ the anti-unitary similitude group[1] and $\widetilde{H} = \mathrm{GU}(n)$ the unitary similitude group.

We will not quite work in the generality of all type I (similitude) reductive dual pairs over G: we need to assume that W_1 is split, i.e. admits a submodule which is its own orthogonal complement. We fix such a submodule W_1', which we call a Lagrangian in W_1 (although W_1 need not be symplectic, so this differs slightly from standard terminology). Let $m = \mathrm{rank}_D W_1'$, so that $\mathrm{rank}_D W_1 = 2m$; we set $n = \mathrm{rank}_D W_2$. We write $l = \dim_{F'} D$, so that W_1' and W_2 have F'-dimensions ml and nl respectively, and we assume $n \geq m + 1 - \frac{1}{l}$.

Remark 2.1.2. If $\epsilon = 1$, $F = F'$, and m is even, then we will sometimes say (G, H) (or $(\widetilde{G}, \widetilde{H})$) is of excluded type, as for technical reasons the results we can prove for dual pairs of this type are slightly weaker than for most type I dual pairs.

Let $P \subset \widetilde{G}$ be the parabolic subgroup stabilizing W_1', and let $M \simeq \mathrm{Res}_{D/F} \mathrm{GL}_D(W_1')$ be the Levi subgroup of P. We denote elements of $\mathrm{Bun}_{\widetilde{G}}$ by $\mathcal{G}$, those of Bun_M by $\mathcal{E}$, and those of Bun_P by $(\mathcal{G}, \mathcal{E})$ via the projections to Bun_M and $\mathrm{Bun}_{\widetilde{G}}$, although strictly P-bundles comprise more data in general.[2] Each $\mathcal{G} \in \mathrm{Bun}_{\widetilde{G}}$ includes the datum of a (right) $\mathrm{GL}_D(W_1)$-torsor over X'[3] equipped with a right (D, σ)-linear ϵ-Hermitian pairing $\langle -, - \rangle_1 : \mathcal{G} \times \sigma^* \mathcal{G} \to \mathfrak{M}' \otimes \omega_{X'}$ for a (left) $D^\times$-torsor $\mathfrak{M}'$ on X' such that $\sigma^* \mathfrak{M}' = \mathfrak{M}'$. Equivalently, we have $\mathfrak{M}' = \nu^* \mathfrak{M}$ for some (left) $(D^\times)^{\sigma=1}$-torsor $\mathfrak{M}$ on X, so the pairing $\langle -, - \rangle_1$ can be taken to be valued in $\nu^*(\mathfrak{M} \otimes \omega_X)$. We often think of $\mathcal{G}$ and $\mathcal{E}$ as vector bundles with $\mathcal{O}_D$-structure for a maximal order $\mathcal{O}_D \subset D$ thought of as an $\mathcal{O}_X$-algebra, and by an abuse of notation sometimes write $\mathrm{Hom}_D(-, -)$ for morphisms of such objects.

[1] That is, the automorphism group of a rank $2m$ anti-Hermitian space, with similitudes; this is abstractly isomorphic to the unitary similitude group $\mathrm{GU}(2m)$, but we write it as GU^- to emphasize that it is half of a (similitude) reductive dual pair, with the other half $\mathrm{GU}(n)$.

[2] In particular a P-bundle $(\mathcal{G}, \mathcal{E})$ also includes the datum of an embedding $\mathcal{E} \hookrightarrow \mathcal{G}$ of D-modules, via the embeddings in $\mathrm{GL}_D(W_1)$ and $\mathrm{GL}_D(W_1')$, so that we think of $\mathcal{E}$ as a Lagrangian subbundle of $\mathcal{G}$.

[3] A priori, it should be an $\mathrm{GL}_D(W_1)$-torsor over X, where we view D as an F-algebra; but since $F \to D$ factors through F', we can equivalently think of it as an $\mathrm{GL}_D(W_1)$-torsor over X', where we think of D as an F'-algebra, and recover the version over X as the pushforward.

Restricting $\langle -, - \rangle_1$ to $\mathcal{E} \times \mathcal{G}/\mathcal{E}$ induces an isomorphism $\mathcal{G}/\mathcal{E} \simeq \sigma^*\mathcal{E}^\vee \otimes \nu^*\mathfrak{M}$, where $\mathcal{E}^\vee = \underline{\mathrm{Hom}}_D(\mathcal{E}, D \otimes \omega_{X'})$ is the Serre dual. Thus we have a short exact sequence

$$0 \to \mathcal{E} \to \mathcal{G} \to \sigma^*\mathcal{E}^\vee \otimes \nu^*\mathfrak{M} \to 0. \tag{2.2}$$

Note that when tensor products appear between two objects equipped with right and left D-module structures (e.g. $\sigma^*\mathcal{E}^\vee \otimes \nu^*\mathfrak{M}$ above), we implicitly mean the tensor product as D-modules.

Each $\mathcal{F} \in \mathrm{Bun}_{\widetilde{H}}$ is similarly a $\mathrm{GL}_D(W_2)$-torsor on X' equipped with a left (D, σ)-linear $(-\epsilon)$-Hermitian pairing $\langle -, - \rangle_W : \mathcal{F} \times \sigma^*\mathcal{F} \to \nu^*\mathcal{L} \otimes \omega_{X'}$ for a left $(D^\times)^{\sigma=1}$-torsor $\mathcal{L}$ on X. Via the F-action on $D^{\sigma=1}$ we can view $\mathfrak{M} \otimes \omega_X$ as a $D^{\sigma=1,\times}$-torsor, which we assume is isomorphic to $\mathcal{L}$, so that (2.2) becomes

$$0 \to \mathcal{E} \to \mathcal{G} \to \sigma^*\mathcal{E}^* \otimes \nu^*\mathcal{L} \to 0. \tag{2.3}$$

2.2 Shtukas, special cycles, and virtual fundamental classes

2.2.1 Stack of $\widetilde{H}$-shtukas and special cycles

Fix a similitude dual reductive pair $(\widetilde{G}, \widetilde{H})$ of type I as above, with W_1 split.

Definition 2.2.1. An $\widetilde{H}$-shtuka over an $\mathbb{F}_q$-scheme S is a tuple

$$((x_1, \dots, x_r), (\mathcal{F}_0, \dots, \mathcal{F}_r), (\mathcal{L}_0, \dots, \mathcal{L}_r), (f_1, \dots, f_r), (g_1, \dots, g_r), \varphi)$$

where

- $x_1, \dots, x_r : S \to X'$ are legs;
- $\mathcal{F}_0, \dots, \mathcal{F}_r$ are $\mathrm{GL}_D(W_2)$-torsors on $X'_S := X' \times_{\mathbb{F}_q} S$ together with a left (D, σ)-linear $(-\epsilon)$-Hermitian pairing $\mathcal{F}_i \otimes_{\mathcal{O}_{X'_S}} \mathcal{F}_i \to \nu^*\mathcal{L}_i \otimes_{\mathcal{O}_{X'_S}} \omega_{X'_S}$;
- $\mathcal{L}_0, \dots, \mathcal{L}_r$ are $(D^\times)^{\sigma=1}$-torsors on X;

- $\varphi : {}^{\tau}\mathcal{F}_r \xrightarrow{\sim} \mathcal{F}_0$ is a $(-\epsilon)$-Hermitian isomorphism, and there is an isomorphism $\psi : \mathcal{L}_r \xrightarrow{\sim} {}^{\tau}\mathcal{L}_0$ such that ${}^{\tau}\langle -, - \rangle_0 \circ (\varphi \otimes \varphi) = \psi \circ \langle -, - \rangle_r$;

- there are modifications

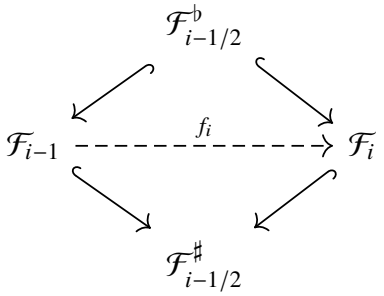

  such that the cokernels of the left inclusions are isomorphic to the pullback of a $D^{\times}$-torsor on X' along x_i and the cokernels of the right inclusions are isomorphic to the pullback of a $D^{\times}$-torsor on X' along $\sigma(x_i)$ and such that the maps to $\mathcal{F}^{\sharp}_{i-1/2}$ are dual to the maps from $\mathcal{F}^{\flat}_{i-1/2}$, i.e. the injections $\mathcal{F}_{i-1} \hookrightarrow \mathcal{F}^{\sharp}_{i-1/2}$ and $\mathcal{F}_i \hookrightarrow \mathcal{F}^{\sharp}_{i-1/2}$ are given by the compositions

$$\mathcal{F}_{i-1} \xrightarrow{\sim} \mathcal{F}^{\vee}_{i-1} \otimes \nu^* \mathcal{L}_{i-1} \to \mathcal{F}^{\flat\vee}_{i-1/2} \otimes \nu^* \mathcal{L}_{i-1},$$
$$\mathcal{F}_i \xrightarrow{\sim} \mathcal{F}^{\vee}_i \otimes \nu^* \mathcal{L}_i \to \mathcal{F}^{\flat}_{i-1/2} \otimes \nu^* \mathcal{L}_i$$

  respectively;

- $f_i : \mathcal{F}_{i-1} \dashrightarrow \mathcal{F}_i$ are isomorphisms on $X'_S \backslash \Gamma_{x_i}$ bounded by μ_i, which we take to be minuscule to recover the modifications as described above, and $g_i : \mathcal{L}_{i-1} \xrightarrow{\sim} \mathcal{L}_i$ are isomorphisms satisfying $\langle -, - \rangle_i \circ (f_i \otimes f_i) = \nu^* g_i \circ \langle -, - \rangle_{i-1}$, so the $\nu^* g_i$'s are the multiplier maps of the f_i's.

In particular, all the $\mathcal{L}_i$ are isomorphic to some $\mathcal{L} \boxtimes \mathcal{O}_S$ for some $(D^{\times})^{\sigma=1}$-torsor $\mathcal{L}$ on X, justifying ignoring the distinction between $\mathcal{L}_{i-1}$ and $\mathcal{L}_i$ in the second-to-last point. We fix minuscule cocharacters μ_i to bound the modifications as described above, and write $\mathrm{Sht}^r_{\widetilde{H}} = \mathrm{Sht}_{\widetilde{H},(\mu_1,\dots,\mu_r)}$ for the stack sending S to the groupoid of shtukas on S in the sense of Definition 2.2.1. It is equipped with a map to $\mathrm{Bun}_{(D^{\times})^{\sigma=1}}(k)$ recording $\mathcal{L}$, and we write $\mathrm{Sht}^r_{H,\mathcal{L}}$ for the fiber over $\mathcal{L}$.

Taking $D = F'$ an unramified quadratic extension of F and $\epsilon = -1$ recovers [17, Definition 2.9].

We can now define for each $\mathcal{E} \in \mathrm{Bun}_M(k)$ a stack $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}} \to \mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$, which we call the special cycle of $\mathcal{E}$ on $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$: an S-point of $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}$ lying over an S-point $\mathcal{F}_0 \dashrightarrow \mathcal{F}_1 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_r \xrightarrow{\sim} {}^\tau\mathcal{F}_0$ of $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$ is a commutative diagram

$$\begin{array}{ccccccccc} \mathcal{E}\boxtimes\mathcal{O}_S & = & \mathcal{E}\boxtimes\mathcal{O}_S & = & \cdots & = & \mathcal{E}\boxtimes\mathcal{O}_S & \xrightarrow{\sim} & {}^\tau\mathcal{E}\boxtimes\mathcal{O}_S \\ \downarrow{\scriptstyle t_0} & & \downarrow{\scriptstyle t_1} & & & & \downarrow{\scriptstyle t_r} & & \downarrow{\scriptstyle {}^\tau t_0} \\ \mathcal{F}_0 & \dashrightarrow & \mathcal{F}_1 & \dashrightarrow & \cdots & \dashrightarrow & \mathcal{F}_r & \xrightarrow{\sim} & {}^\tau\mathcal{F}_0 \end{array}, \tag{2.4}$$

where the vertical arrows are D-linear maps. Each $\mathcal{F}_i$ is equipped with a $\nu^*\mathcal{L}\otimes\omega_{X'}$-valued $(-\epsilon)$-Hermitian (D,σ)-linear pairing, inducing a $(-\epsilon)$-Hermitian isomorphism $h_{\mathcal{F}_i} : \mathcal{F}_i \xrightarrow{\sim} \mathcal{F}_i^\vee \otimes \nu^*\mathcal{L}_i$ (where $\mathcal{L}_i \simeq \mathcal{L}\boxtimes\mathcal{O}_S$), so for each i we can form the composition

$$a_i : \mathcal{E}\boxtimes\mathcal{O}_S \xrightarrow{t_i} \mathcal{F}_i \xrightarrow{h_{\mathcal{F}_i}} \mathcal{F}_i^\vee\otimes\nu^*\mathcal{L}_i \xrightarrow{t_i^\vee\otimes\nu^*\mathcal{L}_i} (\mathcal{E}^\vee\boxtimes\mathcal{O}_S)\otimes\nu^*\mathcal{L}_i \simeq (\mathcal{E}^\vee\otimes\nu^*\mathcal{L}_i)\boxtimes\mathcal{O}_S.$$

Since the diagram (2.4) commutes and the modifications are compatible with the pairings, the a_i must satisfy $a_0 = \cdots = a_r = {}^\tau a_0$ and so be given by $a\boxtimes\mathrm{id}_S$ for some D-linear map $a : \mathcal{E}\to\mathcal{E}^\vee\otimes\nu^*\mathcal{L}$ on X' satisfying $a^\vee\otimes\nu^*\mathcal{L} = -\epsilon a$. Therefore defining

$$\mathcal{A}_{\mathcal{E},\mathcal{L}}(k) = \{a : \mathcal{E}\to\mathcal{E}^\vee\otimes\nu^*\mathcal{L} \text{ s.t. } a^\vee\otimes\nu^*\mathcal{L} = -\epsilon a\}$$

we get a map $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}} \to \mathcal{A}_{\mathcal{E},\mathcal{L}}(k)$, and we write $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)$ for the fiber over a. This is finite over $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$:

Proposition 2.2.2. *For each $a \in \mathcal{A}_{\mathcal{E},\mathcal{L}}(k)$, the map $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a) \to \mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$ is finite.*

Proof. The proof of [18, Proposition 7.5], which is the special case when the reductive dual pair is the unitary pair and $\mathcal{L}$ is trivial, generalizes straightforwardly to our setting. □

In the following subsection, we will define a virtual fundamental class $[\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)]$, which will

allow us to define the higher theta series for $(\widetilde{G}, \widetilde{H})$.

2.2.2 Derived Hitchin stacks

Our goal in this section is to define a derived stack whose classical truncation agrees with $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}$ but which is quasi-smooth of the expected dimension, so that its fundamental class defines a well-behaved virtual fundamental class for $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}$. We will construct this stack as the stack of shtukas over the derived Hitchin stack; that is, we will construct a derived stack $\mathscr{M}$ which we call the derived Hitchin stack and a self-correspondence $\mathrm{Hk}^r_{\mathscr{M}}$ whose intersection with the graph of Frobenius gives a stack $\mathrm{Sht}^r_{\mathscr{M}}$ with the desired properties.

Most of this subsection will follow §3 and §5 of [17], to which we refer for most definitions and notations. However our conventions differ slightly: we write M_n for $\mathrm{Res}_{D/F}\,\mathrm{GL}_n(D)$, which we can view as a group scheme over X, so M_m is the group M of Section 2.1, and we say that an n-framing of a gerbe $\mathscr{G}$ over X is a map $\mathscr{G} \to \mathrm{B}M_n$ and an n-framed gerbe is a gerbe $\mathscr{G}$ over X together with an n-framing. (In the case $D = F'$ a quadratic extension of F we recover [17, §3.3].) If X is a proper k-scheme and $\mathscr{G}$ a gerbe over X, we also write $\mathrm{Sect}(X, \mathscr{G})$ for the derived mapping stack $\mathscr{M}ap_X(X, \mathscr{G})$ denoted $\mathscr{S}ect(X, \mathscr{G})$ in §5 of *loc. cit.*, which on the level of 0-cells sends an animated k-algebra $R_\bullet$ to the anima of morphisms $X_{R_\bullet} = X \times_{\mathrm{Spec}\,k} \mathrm{Spec}\,R_\bullet \to \mathscr{G}$ over X, i.e. sections of the structure map $\mathscr{G} \to X$.

Let $\mathrm{B}\widetilde{H}_{\mathcal{L}}$ be the gerbe sending a scheme $s : S \to X$ to the groupoid of $\widetilde{H}$-bundles on $S' := S \times_X X'$ with similitude bundle $s^*\mathcal{L}$, i.e. locally free sheaves of left D-modules $\mathcal{F}$ of rank n on S' equipped with a left (D, σ)-linear isomorphism $h : \mathcal{F} \xrightarrow{\sim} \mathcal{F}^* \otimes s^*(\mathcal{L} \otimes \omega_X)$, so that $\mathrm{Bun}_{\mathrm{B}\widetilde{H}_{\mathcal{L}}} = \mathrm{Bun}_{\widetilde{H},\mathcal{L}}$. Forgetting the data of h gives an n-framing $\mathrm{B}\widetilde{H}_{\mathcal{L}} \to \mathrm{B}M_n$.

We define the standard representation $\mathrm{Std}(m, n)$ of $\mathrm{B}M_m \times_X \mathrm{B}M_n$ over X to be the stack sending $s : S \to X$ to the groupoid of tuples $(\mathcal{E}, \mathcal{F}, t)$ where $(\mathcal{E}, \mathcal{F}) \in (\mathrm{B}M_m \times_X \mathrm{B}M_n)(S)$ is the underlying S-point and $t : \mathcal{E} \to \mathcal{F}$ is a D-linear map. Forgetting t gives a map $\omega_{m,n} : \mathrm{Std}(m, n) \to \mathrm{B}M_m \times_X \mathrm{B}M_n$. More generally if $\mathrm{B}H$ is an m-framed gerbe and $\mathrm{B}H'$ an n-framed gerbe, then we define

$\mathrm{Std}(\mathrm{B}H, \mathrm{B}H')$ via the Cartesian diagram

$$\begin{array}{ccc} \mathrm{Std}(\mathrm{B}H, \mathrm{B}H') & \longrightarrow & \mathrm{Std}(m,n) \\ \downarrow{\scriptstyle \mu} & & \downarrow{\scriptstyle \omega_{m,n}} \\ \mathrm{B}H \times_X \mathrm{B}H' & \longrightarrow & \mathrm{B}M_m \times_X \mathrm{B}M_n \end{array} , \tag{2.5}$$

with the bottom arrow given by the framings. We define the derived Hitchin stack to be the derived mapping stack

$$\mathscr{M}_{H,H'} = \mathrm{Sect}(X, \mathrm{Std}(\mathrm{B}H, \mathrm{B}H')),$$

which via the left vertical map in (2.5) carries a natural map

$$\pi : \mathscr{M}_{H,H'} \to \mathrm{Sect}(X, \mathrm{B}H \times_X \mathrm{B}H') = \mathrm{Bun}_{\mathrm{B}H} \times \mathrm{Bun}_{\mathrm{B}H'} .$$

Proposition 2.2.3. *For any animated ring $R_\bullet$, at an $R_\bullet$-point $(\mathcal{E}, \mathcal{F}, t)$ of $\mathscr{M}_{H,H'}$ the tangent complex of π is naturally isomorphic to*

$$R\mathrm{pr}_*(\mathcal{H}om_{X_{R_\bullet}}(\mathrm{V}(\mathcal{E}), \mathrm{V}(\mathcal{F})))$$

where $\mathrm{pr} : X_{R_\bullet} \to \mathrm{Spec}\, R_\bullet$ *is the projection map. In particular, π is quasi-smooth, so is $\mathscr{M}_{H,H'}$ is quasi-smooth as well.*

Proof. This follows from [17, Corollary 5.10] with $\mathscr{Y} = \mathrm{Std}(\mathrm{B}H, \mathrm{B}H')$, $\mathscr{Y}' = \mathrm{B}H \times_X \mathrm{B}H'$, and μ as in (2.5). □

Our next goal is to define a Hecke correspondence over this stack. We specialize to the case of interest: set $\mathrm{B}H = \mathrm{B}M_m$ and $\mathrm{B}H' = \mathrm{B}\widetilde{H}_{\mathcal{L}}$.

First, we define $\mathscr{H}_{\mathrm{B}M_m} \to X \times X'$ to be the functor sending a k-scheme S over $X \times X'$ to the groupoid of tuples $(\mathcal{F}^\flat_{1/2}, \mathcal{F}_0, \mathcal{F}_1)$ of D-modules of rank m on $S' = S \times_X X'$, together with injective maps $\mathcal{F}_0 \leftarrow \mathcal{F}^\flat_{1/2} \to \mathcal{F}_1$ with cokernels given by rank 1 D-modules supported on $S' \times_{X' \times X', \Delta} X'$ and $S' \times_{X' \times X', (\mathrm{id}, \sigma)} X'$ respectively. Projecting to the $\mathcal{F}_i$ gives two maps $\mathscr{H}_{\mathrm{B}M_m} \to \mathrm{B}M_m \times X'$ of gerbes

over $X \times X'$, and after applying $\mathrm{Sect}(X' \times X/X', -/X)$ and taking classical truncations these maps recover $h_i \times p_X : \mathrm{Hk}^1_{M_m} \to \mathrm{Bun}_{M_m} \times X'$ for $i = 0, 1$. We also have a map $h^\flat_{1/2} : \mathscr{H}_{M_m} \to \mathrm{B}M_m$ by projecting to $\mathcal{F}^\flat_{1/2}$. For $D = F'$ an unramified quadratic extension of F this recovers [17, Definition 5.18].

Similarly, we define $\mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}} \to X \times X'$ to send a k-scheme S to the analogous data with M_m-bundles replaced by $\widetilde{H}$-bundles with similitude bundle $\mathcal{L}$. The n-framing of $\mathrm{B}\widetilde{H}_\mathcal{L}$ induces a map $\mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}} \to \mathscr{H}_{\mathrm{B}M_n}$, and we again have two maps $\mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}} \to \mathrm{B}\widetilde{H}_\mathcal{L} \times X'$ given by projecting to the $\widetilde{H}$-bundles $\mathcal{F}_0$ and $\mathcal{F}_1$. This recovers $\mathscr{H}_{\mathrm{U}(n),\mathcal{L}}$ as in [17, Definition 5.14].

We define the standard representation

$$\mathrm{Std}(\mathrm{B}M_m, \mathscr{H}_{\mathrm{B}M_n}) := \mathrm{Std}(m, n) \times_{\mathrm{B}M_n, h^\flat_{1/2}} \mathscr{H}_{\mathrm{B}M_n}$$

and thence

$$\mathrm{Std}(\mathrm{B}M_m, \mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}}) = \mathrm{Std}(\mathrm{B}M_m, \mathscr{H}_{\mathrm{B}M_n}) \times_{\mathrm{B}M_n} \mathrm{B}\widetilde{H}_\mathcal{L}.$$

Then we set

$$\mathrm{Hk}^1_{\mathscr{M}_{M_m,\widetilde{H}_\mathcal{L}}} = \mathrm{Sect}(X \times X'/X', \mathrm{Std}(\mathrm{B}M_m, \mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}})).$$

The two projections $p_0, p_1 : \mathscr{H}_{\mathrm{B}\widetilde{H}_\mathcal{L}} \to \mathrm{B}\widetilde{H}_\mathcal{L} \times X'$ induce two projections $\mathrm{Hk}^1_{\mathscr{M}_{M_m,\widetilde{H}_\mathcal{L}}} \to \mathscr{M}_{M_m,\widetilde{H}_\mathcal{L}}$. Abbreviating $\mathscr{M} = \mathscr{M}_{M_m,\widetilde{H}_\mathcal{L}}$, we write

$$\mathrm{Hk}^r_\mathscr{M} = \mathrm{Hk}^1_\mathscr{M} \times_\mathscr{M} \mathrm{Hk}^1_\mathscr{M} \times_\mathscr{M} \cdots \times_\mathscr{M} \mathrm{Hk}^1_\mathscr{M}$$

where on each factor of $\mathrm{Hk}^1_\mathscr{M}$ the left map to $\mathscr{M}$ is given by p_0 and the right map is given by p_1, so that e.g. $\mathrm{Hk}^2_\mathscr{M}$ parametrizes pairs of points $(A, B) \in \mathrm{Hk}^1_\mathscr{M} \times \mathrm{Hk}^1_\mathscr{M}$ such that $p_1(A) = p_2(B)$. There is a natural map $\pi : \mathrm{Hk}^r_\mathscr{M} \to \mathrm{Bun}_{M_m} \times \mathrm{Hk}^r_{\widetilde{H},\mathcal{L}} = \mathrm{Bun}_M \times \mathrm{Hk}_{\widetilde{H},\mathcal{L}}$, whose tangent complex at $(\mathcal{E}, (x_i), (\mathcal{F}_i), (t_i)) \in \mathrm{Hk}^r_\mathscr{M}(R_\bullet)$ we can compute as in [17, Lemma 5.31]:

Proposition 2.2.4. *For any animated ring $R_\bullet$, at an $R_\bullet$-point $(\mathcal{E}, (x_i), (\mathcal{F}_i), (t_i))$ of $\mathrm{Hk}^r_\mathscr{M}$ the tan-*

gent complex of π is naturally isomorphic to

$$R\mathrm{pr}_*\left(\bigoplus_{i=1}^{r}\mathcal{H}om(\mathrm{V}(\mathcal{E}),\mathrm{V}(\mathcal{F}^\flat_{i-1/2}))\to\bigoplus_{i=1}^{r-1}\mathcal{H}om(\mathrm{V}(\mathcal{E}),\mathrm{V}(\mathcal{F}_i))\right)$$

where $\mathrm{pr}: X_{R_\bullet}\to \mathrm{Spec}\, R_\bullet$ *is the projection and the differential is induced by the map*

$$\bigoplus_{i=1}^{r}\mathcal{F}^\flat_{i-1/2}\to\bigoplus_{i=1}^{r-1}\mathcal{F}_i$$

sending $(v_{1/2},\dots,v_{r-1/2})\mapsto(v_{1/2}-v_{3/2},v_{3/2}-v_{5/2},\dots,v_{r-3/2}-v_{r-1/2})$. *In particular,* π *is quasi-smooth of relative dimension* $-mlr$*, so* $\mathrm{Hk}^r_{\mathscr{M}}$ *is quasi-smooth.*

The proof is identical to the proof of [17, Lemma 5.31].

Therefore (following [17, §5.8]), if

$$\mathrm{Sht}^r_{\mathscr{M}}=\mathrm{Hk}^r_{\mathscr{M}}\times_{\mathscr{M}\times\mathscr{M},(\mathrm{id},\mathrm{Frob})}\mathscr{M},$$

we conclude that the natural map $\mathrm{Sht}^r_{\mathscr{M}}\to\mathrm{Bun}_M(k)\times\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$ is quasi-smooth of (virtual) relative dimension $-mlr$ and so $\mathrm{Sht}^r_{\mathscr{M}}$ is quasi-smooth of (virtual) dimension $((n-m-1)l+1)r$,[4] and therefore has a well-defined fundamental class. The fiber of its classical truncation over $\{\mathcal{E}\}\times\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$ recovers $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}$, and so we can define the virtual fundamental class of $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}$ and its fibers $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)$ as the respective pullbacks of $[\mathrm{Sht}^r_{\mathscr{M}}]$. By Proposition 2.2.2, each $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)$ is finite over $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$, and so the pushforward of the virtual fundamental class is well-defined; by a mild abuse of notation, we write $[\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)]$ for the pushforward of the virtual fundamental class to $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$, so it can be thought of as the (virtual) cycle class of $\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)\to\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$.

[4] To compute the dimension, we also need to know the dimension of $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$; this follows from the computation of the dimension of the corresponding Hecke stack in Lemma 2.5.3.

2.3 The modularity conjecture

The short exact sequence (2.3) gives rise to an extension class $e_{\mathcal{G},\mathcal{E}} \in \mathrm{Ext}^1(\sigma^*\mathcal{E}^* \otimes \nu^*\mathcal{L}, \mathcal{E})$, which has a Serre duality pairing with $\mathrm{Hom}(\mathcal{E}, \sigma^*\mathcal{E}^\vee \otimes \nu^*\mathcal{L})$, so for each $a \in \mathcal{A}_{\mathcal{E},\mathcal{L}}(k)$ we get $\langle e_{\mathcal{G},\mathcal{E}}, a\rangle \in \mathbb{F}_q$.

Let $\eta_{F'/F} : \mathrm{Pic}_X(k) \to \overline{\mathbb{Q}}_\ell^\times$ be the character associated to the cover X'/X (which may be trivial), and let $\chi : \mathrm{Bun}_{D^\times}(k) \to \overline{\mathbb{Q}}_\ell^\times$ be any character (with respect to the group operation of taking tensor products over D) such that the restriction of χ along $\mathrm{Pic}_X(k) \to \mathrm{Pic}_{X'}(k) = \mathrm{Bun}_{\mathrm{GL}_1,X'}(k) \to \mathrm{Bun}_{D^\times}(k)$ is given by $\eta^{nl}_{F'/F}$.

Since D is a central division algebra over F', there exists some finite extension E/F' over which D is isomorphic to a product of matrix algebras after passing to which we can take the product of the determinants, which if we start with values in F' will be valued in $F'^\times$. This defines a homomorphism (the reduced norm) $D^\times \to F'^\times$. Similarly if we start by taking σ-invariants we get a reduced norm map $(D^\times)^{\sigma=1} \to F^\times$, and (composing with the usual determinant) $M = \mathrm{GL}_D(W_1') \to F'^\times = \mathrm{GL}_1$ over F', inducing $\det : \mathrm{Bun}_{(D^\times)^{\sigma=1}} \to \mathrm{Bun}_{\mathrm{GL}_1,X} = \mathrm{Pic}(X)$ over F and $\det : \mathrm{Bun}_M \to \mathrm{Bun}_{\mathrm{GL}_1,X'} = \mathrm{Pic}(X')$. In particular, we have well-defined degrees on each target, and so we can define $\deg \mathcal{L}$ and $\deg \mathcal{E}$ to be the degrees of their determinants. Note that these might not agree with the degrees of the underlying vector bundles.

Finally, fix a nontrivial additive character $\psi : \mathbb{F}_q \to \overline{\mathbb{Q}}_\ell^\times$, and write $[\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)]$ for the virtual cycle class from Section 2.2.2. Then we can define

$$\begin{aligned} \widetilde{Z}^r_{m,\mathcal{L}} : \mathrm{Bun}_{P,\omega_X^{-1}\otimes\mathcal{L}}(k) &\to \mathrm{Ch}^{mlr}(\mathrm{Sht}^r_{H,\mathcal{L}})_{\overline{\mathbb{Q}}_\ell} \\ (\mathcal{G},\mathcal{E}) &\mapsto \chi(\det\mathcal{E}) q^{\frac{n}{2}(\deg\mathcal{E}-\deg\mathcal{L}-\deg\omega_X)} \sum_{a\in\mathcal{A}_{\mathcal{E},\mathcal{L}}(k)} \psi(\langle e_{\mathcal{G},\mathcal{E}}, a\rangle)[\mathcal{Z}^r_{\mathcal{E},\mathcal{L}}(a)]. \end{aligned} \tag{2.6}$$

Letting $\mathcal{L}$ vary, we can more generally form the series

$$\begin{aligned} \widetilde{Z}^r_m : \mathrm{Bun}_P(k) &\to \mathrm{Ch}^{mlr}(\mathrm{Sht}^r_{\widetilde{H}})_{\overline{\mathbb{Q}}_\ell} \\ (\mathcal{G},\mathfrak{M},\mathcal{E}) &\mapsto \widetilde{Z}^r_{m,\mathfrak{M}\otimes\omega_X}(\mathcal{G},\mathcal{E}). \end{aligned} \tag{2.7}$$

We often omit $\mathfrak{M}$ from the notation, and implicitly let $\mathcal{L} = \mathfrak{M} \otimes \omega_X$ be the similitude bundle.

Conjecture 2.3.1 (Modularity for general type I dual pairs)**.** *The higher theta series $\widetilde{Z}^r_m$ is left-invariant under $G(F)$, i.e. $\widetilde{Z}^r_m$ is a modular form on $\widetilde{G}$ with values in* $\mathrm{Ch}^{mlr}(\mathrm{Sht}^r_{\widetilde{H}})_{\overline{\mathbb{Q}}_\ell}$.

Equivalently, Conjecture 2.3.1 states that $\widetilde{Z}^r_m(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$ and not on $\mathcal{E}$.

It is convenient to weaken Conjecture 2.3.1 in the following three ways. First, in order to more easily apply sheaf-theoretic tools such as the ℓ-adic Fourier transform, we pass to ℓ-adic cohomology via the cycle class map. Second, the target stack $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$ admits a presentation as a colimit of finite type open substacks $\mathrm{Sht}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ for Harder–Narasimhan polygons μ for $\widetilde{H}$. Note that the role of μ here is to bound the underlying $\widetilde{H}$-bundles $\mathcal{F}_i$, not to bound the modifications of the shtukas. By working with one stratum at a time, we can restrict our attention to finite-type stacks. Finally, let $\eta = \mathrm{Spec}\, F' \to X'$ be the generic point, and $\eta^r \to X'^r$ its r-fold self-fiber product over $\mathrm{Spec}\, k$ (note this is larger than the generic point of X'^r). Restricting to $\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}} \times_{X'^r} \eta^r$ allows us to make convenient assumptions about the support of certain torsion sheaves on X' introduced in Section 2.5.

Combining these modifications, we write

$$\widetilde{Z}^r_{m,\eta} : \mathrm{Bun}_P(k) \to \varprojlim_{\mu} H^{2mlr}(\mathrm{Sht}^{r,\leq\mu}_{\widetilde{H}} \times_{X'^r} \eta^r, \overline{\mathbb{Q}}_\ell) \tag{2.8}$$

for the composition of $\widetilde{Z}^r_m$ with the cycle class map and passing to the generic fiber and Harder–Narasimhan strata. Then our main result is that Conjecture 2.3.1 holds after replacing $\widetilde{Z}^r_m$ by $\widetilde{Z}^r_{m,\eta}$, i.e. $\widetilde{Z}^r_{m,\eta}$ is modular, up to mild modification for dual pairs of excluded type:

Theorem 2.3.2. *(a) If (G, H) is a type I reductive dual pair such that the symplectic representation W admits an H-stable Lagrangian subspace and is not of excluded type, then the higher theta series Z^r_η descends to* $\mathrm{Bun}_G(k)$*, i.e. $\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$.*

(b) If (G, H) is of excluded type, then $\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$ and on the class of $\mathcal{E}$ in $\pi_0 \mathrm{LGr}(m, 2m)$.

2.4 The case $r = 0$

Following [19], we first explain our argument in the case $r = 0$; in this case the virtual fundamental classes are not needed and Z^0_m and $Z^0_{m,\eta}$ are actually the same. Although the general case is more complicated, the idea of the proof is similar and in certain places will reduce back to the $r = 0$ case; in particular many of the modifications of the method used in *loc. cit.* needed to work in our generality will already become apparent here.

Lemma 2.4.1. *Let D be a finite-dimensional division algebra over a field F of characteristic different from 2 and V be a right D-module of rank $2m$ with a right (D, σ)-linear ϵ-Hermitian bilinear form as above. Let $L_1, L_2 \subset V$ be Lagrangian D-submodules, and if $\epsilon = 1$ and $F' = F$ then assume $\dim_D L_1 \cap L_2 \equiv m \pmod 2$. Then there exists a third Lagrangian L of V of rank m such that $L_1 \cap L = L \cap L_2 = 0$.*

Proof. Assume first that either $\epsilon = -1$ or $F' \neq F$, i.e. σ is nontrivial and so F' is a degree 2 extension of F. Then a mild modification of the proof of [19, Lemma 2.2.1] works: if $I = L_1 \cap L_2$, we have an orthogonal decomposition $V = V' \oplus I \oplus I^*$, with $L_i = L'_i \oplus I$ and each L'_i Lagrangian in V'. Then it suffices to prove that there exists a maximal isotropic submodule L' of V' such that $L'_1 \cap L' = L' \cap L'_2 = 0$; then we can take $L = L' \oplus I^*$. Via the bilinear form on V' and the decomposition $V' \simeq L'_1 \oplus L'_2$, we can identify L_2 with L_1^*, so choosing a basis for L'_1 gives a dual basis for L'_2 and a decomposition of V' into ϵ-Hermitian D-modules of rank 2. Therefore we can reduce to the case $m = 1$. When $\epsilon = -1$, every 1-dimensional D-submodule is isotropic; when $\epsilon = 1$ and $[F' : F] = 2$, since $k = \mathbb{F}_q$ contains at least 3 elements there are at least 4 isotropic 1-dimensional D-submodules. In either case, since any two distinct 1-dimensional D-submodules have trivial intersection, we can find a third isotropic submodule L' whose intersection with each L'_i is trivial as desired.

We therefore restrict our attention to the case where $\epsilon = 1$ and $F' = F$. In this case the Grassmannian of isotropic D-submodules of V of rank m has two isomorphic components, with L_1 and L_2 in the same component if and only if $\dim_D L_1 \cap L_2 \equiv m \pmod 2$. This explains the

assumption in the statement of the lemma: if $L_1 \cap L = L \cap L_2 = 0$, then either all of L_1, L_2, and L are in the same component, or L_1 and L_2 are in the same component and L is in the other, depending on the parity of m; either way, L_1 and L_2 must be in the same component for such an L to exist.

In the case $m = 1$ we have only two isotropic lines, one from each component, and so while any two must be either equal or have trivial intersection we cannot find a third isotropic line L' as needed for the above approach. Instead, let L be an isotropic D-submodule of V whose intersection with each L_i is trivial; the trivial submodule $\{0\}$ satisfies these assumptions, so such a submodule certainly exists, and we let L be the maximal such submodule. Let $L^\perp \subset V$ be the coisotropic submodule of vectors $v \in V$ such that $\langle v, l\rangle = 0$ for all $l \in L$. Since L is isotropic, $L^\perp/L$ is a nondegenerate orthogonal D-module; let $\pi : L^\perp \to L^\perp/L$ be the projection. Then each $\pi(L_i \cap L^\perp)$ is an isotropic submodule of $L^\perp/L$, hence of dimension at most $\frac{1}{2}\dim_D L^\perp/L$. If $\dim_D L^\perp/L > 2$, or equivalently if $\dim_D L < m-1$, then we can find an isotropic line $\ell \in L^\perp/L$ whose intersection with each $\pi(L_i \cap L^\perp)$ is trivial, and $\pi^{-1}(\ell)$ is an isotropic subspace of $L^\perp$ which strictly contains L and whose intersection with each L_i is trivial. Therefore either $\dim_D L = m-1$ or $\dim_D L = m$. If $\dim_D L = m$ then we are done, so assume $\dim_D L = m-1$.

There are exactly two Lagrangians containing L, given by the preimage under π of the two isotropic lines in the rank 2 D-module $L^\perp/L$; call these L' and L''. Since their intersection L has dimension $m-1$, they must be in different components of the Grassmannian, say L' in the same component as L_1 and L_2 and L'' in the opposite component. Each has intersection with L_i of dimension at most 1. Suppose first m is even. Then $\dim L' \cap L_i$ must be even, hence 0, and so $L_1 \cap L' = L' \cap L_2 = 0$. Similarly, if m is odd, then $\dim L'' \cap L_i$ must be even, hence 0, and so $L_1 \cap L'' = L'' \cap L_2 = 0$. □

For two different Lagrangian subbundles $\mathcal{E}_1, \mathcal{E}_2 \subset \mathcal{G}$ (i.e. elements of $\mathrm{Bun}_P(k)$ agreeing on the projection to $\mathrm{Bun}_{\widetilde{G}}(k)$), say that they are *transverse* if their intersection in $\mathcal{G}$ is the zero section.

Proposition 2.4.2 (Reduction to transverse Lagrangians)**.** *Suppose that for any $\mathcal{G} \in \mathrm{Bun}_{\widetilde{G}}(k)$ and*

$\mathcal{E}_1, \mathcal{E}_2$ transverse Lagrangian subbundles as above we have

$$\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E}_1) = \widetilde{Z}^r_{m,\eta}(\mathcal{G}, \mathcal{E}_2).$$

Then Theorem 2.3.2 holds.

Proof. As the transversality condition is equivalent to the intersection of $\mathcal{E}_1$ and $\mathcal{E}_2$ on the generic fiber being trivial, applying Lemma 2.4.1 with $V = \mathcal{G}_\eta$ and $L_i = \mathcal{E}_\eta$ allows us to link any two Lagrangians by a third transverse to both, assuming $\dim L_1 \cap L_2 \equiv m \pmod 2$ if $\epsilon = 1$ and $F' = F$. If $\epsilon = 1$, $F' = F$, and $\dim L_1 \cap L_2 \not\equiv m \pmod 2$, then L_1 and L_2 are in different components in the Lagrangian Grassmannian. Assume m is odd, and choose a Lagrangian L transverse to L_1 (so in the opposite component to L_1, and the same one as L_2). By Lemma 2.4.1, we can now find a Lagrangian L' transverse to both L and L_2 (in the same component as L_1, and the opposite one to L and L_2), so we can link any two Lagrangians by either one or two intermediate Lagrangians, with any two adjacent Lagrangians transverse; together with the above, this shows that the conditions of the proposition imply Theorem 2.3.2 (a). When m is even, any two transverse Lagrangians lie in the same component of the Grassmannian, so we cannot link two Lagrangians in different components; but Lemma 2.4.1 allows us to link any two in the same component, which is enough for Theorem 2.3.2 (b). □

We can therefore safely restrict our attention to the case of transverse Lagrangians, and henceforth generally do so.

Let $\mathcal{E}_1, \mathcal{E}_2 \subset \mathcal{G}$ be transverse Lagrangian subbundles, which we can view as two elements $(\mathcal{G}, \mathcal{E}_i) \in \mathrm{Bun}_P(k)$. The short exact sequence (2.3) for $\mathcal{E}_1$ gives an injection $\mathcal{E}_1 \to \mathcal{G}$, and for $\mathcal{E}_2$ a surjection $\mathcal{G} \to \sigma^*\mathcal{E}_2^* \otimes \nu^*\mathcal{L}$. Composing these gives a map

$$b_{12} : \mathcal{E}_1 \to \mathcal{G} \to \sigma^*\mathcal{E}_2^* \otimes \nu^*\mathcal{L},$$

which by the transversality of $\mathcal{E}_1$ and $\mathcal{E}_2$ is generically full rank and in particular is injective with

cokernel Q_2 a torsion coherent sheaf on X', i.e. torsion $\mathcal{O}_{X'}$-modules or equivalently having support at finitely many points. Similarly, exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$, we obtain a map

$$b_{21} : \mathcal{E}_2 \to \mathcal{G} \to \sigma^*\mathcal{E}_1^* \otimes \nu^*\mathcal{L},$$

which is again injective with torsion cokernel Q_1 on X',[5] satisfying

$$\sigma^* b_{12}^* \otimes \nu^*\mathcal{L} = \epsilon b_{21}. \tag{2.9}$$

For a torsion sheaf T on X', we write $T^* = \underline{\mathrm{Hom}}(T, F'/\mathcal{O}_{X'})$; then taking cokernels (2.9) gives isomorphisms

$$\beta_{12} : Q_1 \xrightarrow{\sim} \sigma^* Q_2^* \otimes \nu^*\mathcal{L} \quad \text{and} \quad \beta_{21} : Q_2 \xrightarrow{\sim} \sigma^* Q_1^* \otimes \nu^*\mathcal{L}. \tag{2.10}$$

Let $\mathcal{G}^\sharp := (\sigma^*\mathcal{E}_2^* \oplus \sigma^*\mathcal{E}_1^*) \otimes \nu^*\mathcal{L}$. The natural map $\mathcal{G} \to \mathcal{G}^\sharp$ has kernel $\mathcal{E}_1 \cap \mathcal{E}_2 = 0$, so it is a subbundle of the same rank. The $\nu^*\mathcal{L}$-valued pairing on $\mathcal{G}$ extends rationally to a pairing $\mathcal{G}^\sharp \times \sigma^*\mathcal{G}^\sharp \to F' \otimes_{\mathcal{O}_{X'}} \nu^*\mathcal{L}$, which by restriction induces pairings $\sigma^*\mathcal{E}_1^* \times \mathcal{E}_2^* \to F' \otimes_{\mathcal{O}_{X'}} \nu^*\mathcal{L}$ and $\sigma^*\mathcal{E}_2^* \times \mathcal{E}_1^* \to F' \otimes \mathcal{O}_{X'}\nu^*\mathcal{L}$. Taking cokernels, these pairings induce pairings $\gamma_{12} : Q_1 \times \sigma^* Q_2 \to F'/\mathcal{O}_{X'} \otimes \nu^*\mathcal{L}$ and $\gamma_{21} : Q_2 \times \sigma^* Q_1 \to F'/\mathcal{O}_{X'} \otimes \nu^*\mathcal{L}$, which correspond to the isomorphisms β_{12} and β_{21} respectively. Since the pairing on $\mathcal{G}$ is ϵ-Hermitian, the one on $\mathcal{G}^\sharp$ is as well. Therefore for any open subspace $U \subset X'$ and sections s_1, s_2 of Q_1 and $\sigma^* Q_2$ respectively over U, we have $\gamma_{12}(s_1, s_2) = \epsilon\sigma^*\gamma_{21}(\sigma^* s_2, \sigma^* s_1)$. Therefore

$$\sigma^*\beta_{12}^* \otimes \nu^*\mathcal{L} = \epsilon\beta_{21}. \tag{2.11}$$

Similarly to the Q_i, the cokernel Q of $\mathcal{G} \hookrightarrow \mathcal{G}^\sharp$ is a torsion sheaf, and as in [19, §2.3.3] the induced maps $Q \to Q_i$ are isomorphisms. Composing these isomorphisms with those of (2.10)

[5]Note that since the maps b_{12} and b_{21} are D-linear, giving each Q_i the structure of a torsion coherent sheaf of D-modules in a suitable way. The various isomorphisms between variations on these sheaves are constructed from D-linear maps and so respect this structure.

gives

$$\begin{aligned} h_{12} : Q \xrightarrow{\sim} Q_1 \xrightarrow{\beta_{12}} \sigma^* Q_2^* \otimes \nu^* \mathcal{L} \xrightarrow{\sim} \sigma^* Q^* \otimes \nu^* \mathcal{L}, \\ h_{21} : Q \xrightarrow{\sim} Q_2 \xrightarrow{\beta_{21}} \sigma^* Q_1 \otimes \nu^* \mathcal{L} \xrightarrow{\sim} \sigma^* Q^* \otimes \nu^* \mathcal{L}, \end{aligned} \tag{2.12}$$

which therefore satisfy

$$\sigma^* h_{12}^* \otimes \nu^* \mathcal{L} = \epsilon h_{21}. \tag{2.13}$$

This brings us to the following key lemma:

Lemma 2.4.3. *With notation as above,* $h_{12} = -h_{21}$.

Note that the statement is independent of ϵ or any of the other data associated to the dual pair $(\widetilde{G}, \widetilde{H})$ other than h_{12} and h_{21} themselves: it is identical to [19, Lemma 2.3.3]. This is critical: it will allow us to use the theory of Fourier analysis on a Hermitian (or symmetric) vector space, rather than potentially symplectic or skew-Hermitian spaces. The proof is nearly the same as in *loc. cit.*, but we include it due to the importance of the lemma.

Proof. Consider the pairing $\gamma : (Q_1 \oplus Q_2) \times (\sigma^* Q_1 \oplus \sigma^* Q_2) \to F'/\mathcal{O}_{X'} \otimes \nu^* \mathcal{L}$ given by

$$\gamma(s_1, s_2, s_1', s_2') = \gamma_{12}(s_1, s_2') + \gamma_{21}(s_2, s_1').$$

In particular the restrictions to $Q_i \times \sigma^* Q_i$ are trivial. Each of h_{12} and h_{21} are derived from β_{12} and β_{21} respectively by identifying Q_1 and Q_2 with Q in appropriate ways, and so modulo such identifications the claim $h_{12} = -h_{21}$ amounts to a symmetry of γ. More precisely, let $\iota_i : Q_i \xrightarrow{\sim} Q$ be the fixed isomorphism. Then the pairing $Q \times \sigma^* Q \to F'/\mathcal{O}_{X'} \otimes \nu^* \mathcal{L}$ associated to h_{12} is

$$(s, s') \mapsto \gamma_{12}(\iota_1^{-1}(s), \iota_2^{-1}(s'))$$

and likewise the pairing associated to h_{21} is

$$(s, s') \mapsto \gamma_{21}(\iota_2^{-1}(s), \iota_1^{-1}(s')).$$

Thus the claim $h_{12} = -h_{21}$ is equivalent to

$$\gamma_{12}(\iota_1^{-1}(s), \iota_2^{-1}(s')) = -\gamma_{21}(\iota_2^{-1}(s), \iota_1^{-1}(s')). \tag{2.14}$$

Recall that we have a surjection $\mathcal{G}^\sharp \to Q_1 \oplus Q_2$, with kernel $\mathcal{E}_1 \oplus \mathcal{E}_2$. The pairing γ fits into a commutative diagram

$$\begin{array}{ccccc} \mathcal{G} \times \sigma^*\mathcal{G} & \longrightarrow & \mathcal{G}^\sharp \times \sigma^*\mathcal{G}^\sharp & \longrightarrow & (Q_1 \oplus Q_2) \times (\sigma^*Q_1 \oplus \sigma^*Q_2) \\ \downarrow & & \downarrow & & \downarrow{\scriptstyle \gamma} \\ \nu^*\mathcal{L} & \longrightarrow & F' \otimes_{O_{X'}} \nu^*\mathcal{L} & \longrightarrow & F'/O_{X'} \otimes \nu^*\mathcal{L} \end{array}, \tag{2.15}$$

so the image $\overline{\mathcal{G}} \subset Q_1 \oplus Q_2$ of $\mathcal{G}$ is isotropic under γ.

Via the inclusion $\overline{\mathcal{G}} \subset Q_1 \oplus Q_2$, we can think of $\overline{\mathcal{G}}$ as a relation in $Q_1 \times Q_2$. The projection $\overline{\mathcal{G}} \to Q_1$ can be identified with $\operatorname{coker}(\mathcal{E}_1 \oplus \mathcal{E}_2 \to \mathcal{G}) \simeq \operatorname{coker}(\mathcal{E}_2 \to \operatorname{coker}(\mathcal{E}_1 \to \mathcal{G})) = \operatorname{coker}(\mathcal{E}_2 \to \sigma^*\mathcal{E}_1^* \otimes \nu^*\mathcal{L}) = \operatorname{coker}(b_{21}) = Q_1$ and similarly for Q_2. In particular, both projections $\overline{\mathcal{G}} \to Q_i$ are isomorphisms, so we can think of $\overline{\mathcal{G}} \hookrightarrow Q_1 \oplus Q_2$ as the graph of an isomorphism $\varphi : Q_1 \xrightarrow{\sim} Q_2$. In other words,

$$\overline{\mathcal{G}} = \{(s_1, \varphi(s_1))\} \tag{2.16}$$

as s_1 varies over sections of Q_1 over any open subspace $U \subset X'$. Since $Q = \mathcal{G}^\sharp/\mathcal{G} = (Q_1 \oplus Q_2)/\overline{\mathcal{G}}$, we have $\iota_1(s_1) + \iota_2(s_2) = 0$ for all $(s_1, s_2) \in \overline{\mathcal{G}} \subset Q_1 \oplus Q_2$, so via (2.16) we have

$$\iota_1(s_1) + \iota_2(\varphi(s_1)) = 0,$$

i.e. $\iota_2^{-1}(\iota_1(s_1)) = -\varphi(s_1)$ for all $s_1 \in Q_1$. Therefore we have

$$\varphi = -\iota_2^{-1} \circ \iota_1. \tag{2.17}$$

Since $\overline{\mathcal{G}}$ is isotropic for γ, we have

$$\gamma(s_1, \varphi(s_1), s_1', \varphi(s_1')) = \gamma_{12}(s_1, \varphi(s_1')) + \gamma_{21}(\varphi(s_1), s_1') = 0. \tag{2.18}$$

By (2.17), equation (2.18) is equivalent to

$$-\gamma_{12}(s_1, \iota_2^{-1}(\iota_1(s_1'))) - \gamma_{21}(\iota_2^{-1}(\iota_1(s_1)), s_1') = 0,$$

which after writing $s := \iota_1(s_1)$, $s' := \iota_1(s_1')$ gives (2.14). □

Fix $\mathcal{F} \in \mathrm{Bun}_{\widetilde{H},\mathcal{L}}(k)$, i.e. a vector bundle of rank n on X' equipped with an isomorphism $h_{\mathcal{F}} : \mathcal{F} \xrightarrow{\sim} \sigma^*\mathcal{F}^\vee \otimes \nu^*\mathcal{L}$ such that $\sigma^* h_{\mathcal{F}}^\vee = -\epsilon h_{\mathcal{F}}$. We define a vector space

$$V = \mathrm{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \sigma^* Q).$$

Since $\sigma^* h_{12}^* \otimes \nu^*\mathcal{L} = \epsilon h_{21} = -\epsilon h_{12}$, it follows that h_{12} is induced from a $(-\epsilon)$-Hermitian form $Q \times \sigma^* Q \to F'/\mathcal{O}_{X'} \otimes \nu^*\mathcal{L}$, and similarly for h_{21}. As a corollary, V is naturally equipped with a (D, σ)-linear Hermitian pairing, independent of the value of ϵ. It therefore carries a quadratic form $\mathfrak{q}_{12}$, and similarly its negative $\mathfrak{q}_{21} = -\mathfrak{q}_{12}$.

For each D-linear map $t : \mathcal{E}_i \to \mathcal{F}$, we can define

$$s = s(t) : \mathcal{F}^* \otimes \nu^*\mathcal{L} \xrightarrow{t^* \otimes \nu^*\mathcal{L}} \mathcal{E}_i^* \otimes \nu^*\mathcal{L} \to Q_i \xrightarrow{\sim} Q,$$

which gives a pair of maps $f_i : \mathrm{Hom}(\mathcal{E}_i, \mathcal{F}) \to V$ for $i = 1, 2$. Imitating the proof of Lemma 2.3.4 of [19], we find that for $t : \mathcal{E}_1 \to \mathcal{F}$ with image $a(t) \in \mathcal{A}_{\mathcal{E}_1,\mathcal{L}}(k)$ and $s(t) = f_1(t) \in V$ we have

$$\langle a(t), e_{\mathcal{G},\mathcal{E}_1} \rangle = \mathfrak{q}_{12}(s(t)), \tag{2.19}$$

and similarly for $t : \mathcal{E}_2 \to \mathcal{F}$

$$\langle a(t), e_{\mathcal{G},\mathcal{E}_2} \rangle = \mathfrak{q}_{21}(s(t)). \tag{2.20}$$

We can now explain the proof of the $r = 0$ case of Theorem 2.3.2. For each $a \in \mathcal{A}_{\mathcal{E},\mathcal{L}}(k)$, $\mathcal{Z}^0_{\mathcal{E},\mathcal{L}}(a)$ is a finite discrete space over $\mathrm{Sht}^0_{\widetilde{H},\mathcal{L}} = \mathrm{Bun}_{\widetilde{H},\mathcal{L}}(k)$, with class in $\mathrm{Ch}_0(\mathrm{Bun}_{\widetilde{H},\mathcal{L}}(k))$ given by its degree over each fiber; the same is true of the theta series $\widetilde{Z}^r_{m,\mathcal{L}}(\mathcal{G},\mathcal{E})$ (after base change to $\overline{\mathbb{Q}}_\ell$), so it suffices to compute the degrees of its fibers at each $\mathcal{F} \in \mathrm{Bun}_{\widetilde{H},\mathcal{L}}(k)$.

By definition,

$$\begin{aligned}
\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_1)_{\mathcal{F}} &= \chi(\det \mathcal{E}_1) q^{\frac{n}{2}(\deg \mathcal{E}_1 - \deg \mathcal{L} - \deg \omega_X)} \sum_{a \in \mathcal{A}_{\mathcal{E}_1,\mathcal{L}}(k)} \psi(\langle a, e_{\mathcal{G},\mathcal{E}_1}\rangle)[\mathcal{Z}^0_{\mathcal{E}_1,\mathcal{L}}(a)]_{\mathcal{F}} \\
&= \chi(\det \mathcal{E}_1) q^{\frac{n}{2}(\deg \mathcal{E}_1 - \deg \mathcal{L} - \deg \omega_X)} \sum_{t \in \mathrm{Hom}(\mathcal{E}_1,\mathcal{F})} \psi(\langle a(t), e_{\mathcal{G},\mathcal{E}_1}\rangle \\
&= \chi(\det \mathcal{E}_1) q^{\frac{n}{2}(\deg \mathcal{E}_1 - \deg \mathcal{L} - \deg \omega_X)} \sum_{t \in \mathrm{Hom}(\mathcal{E}_1,\mathcal{F})} \psi(\mathfrak{q}_{12}(s(t))) \\
&= \chi(\det \mathcal{E}_1) q^{\frac{n}{2}(\deg \mathcal{E}_1 - \deg \mathcal{L} - \deg \omega_X)} \sum_{s \in V} \psi(\mathfrak{q}_{12}(s)) f_{1!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})}
\end{aligned}$$

where for any map of finite sets $f : A \to B$ and function $\gamma : A \to \overline{\mathbb{Q}}_\ell$ we write $(f_!\gamma)(b) = \sum_{a \in f^{-1}(b)} \gamma(a)$, and $\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})}$ denotes the constant function 1. For functions α, β on V, write

$$\langle \alpha, \beta \rangle_V = \sum_{s \in V} \alpha(s)\beta(s);$$

then we have

$$\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_1) = \chi(\det \mathcal{E}_1) q^{\frac{n}{2}(\deg \mathcal{E}_1 - \deg \mathcal{L} - \deg \omega_X)} \langle \mathfrak{q}_{12}^* \psi, f_{1!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})} \rangle. \tag{2.21}$$

Similarly, exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$ everywhere above, we obtain

$$\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_2) = \chi(\det \mathcal{E}_2) q^{\frac{n}{2}(\deg \mathcal{E}_2 - \deg \mathcal{L} - \deg \omega_X)} \langle \mathfrak{q}_{21}^* \psi, f_{2!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_2,\mathcal{F})} \rangle. \tag{2.22}$$

The Hermitian pairing $(-,-) : V \times \sigma^* V \to k$ on V can be viewed as an isomorphism of V with

its dual, which lets us define the Fourier transform of a function φ on V as follows:

$$\mathrm{FT}_V(\varphi)(s) = (-1)^{\dim V} \sum_{s' \in V} \varphi(s')\psi((s, s'))$$

where ψ is our fixed nontrivial character of k. This is exactly as in [19, §2.3.7], to which we refer the reader for an array of good properties of this Fourier transform. Of particular note is the Plancherel formula:

$$\langle \varphi_1, \varphi_2 \rangle_V = q^{-\dim V} \langle \mathrm{FT}_V([-1]^* \varphi_1), \mathrm{FT}_V(\varphi_2) \rangle_V. \tag{2.23}$$

Using (2.23), we will deduce the $r = 0$ case of Theorem 2.3.2 from the following proposition. Let D_Q be the divisor on X associated to the Hermitian torsion sheaf Q as in [19, §2.3.8].

Proposition 2.4.4. *The following formulas hold:*

(a) $\mathrm{FT}_V([-1]^* \mathfrak{q}_{12}^* \psi) = (-1)^{\dim V} q^{\frac{1}{2}\dim V} \eta_{F'/F}(D_Q)^{nl} (\mathfrak{q}_{21}/4)^* \psi$.

(b) $\mathrm{FT}_V(f_{1!} \mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1, \mathcal{F})}) = (-1)^{\dim V} q^{\chi(X, \mathcal{F} \otimes \mathcal{E}_2 \otimes \nu^* \mathcal{L}^*) + \dim V} f_{2!} \mathbf{1}_{\mathrm{Hom}(\mathcal{E}_2, \mathcal{F})}$.

(c) $\frac{1}{2} n \deg \mathcal{E}_1 + \frac{1}{2} \dim V + \chi(X, \mathcal{F} \otimes \mathcal{E}_2 \otimes \nu^* \mathcal{L}^*) = \frac{1}{2} n \deg \mathcal{E}_2$.

(d) $\chi(\det \mathcal{E}_1) \eta_{F'/F}(D_Q)^{nl} = \chi(\det \mathcal{E}_2)$.

Proof. (a) The argument is essentially the same as in the proof of [19, Lemma 2.3.8].

(b) We have an exact sequence

$$\begin{aligned} 0 \to \mathrm{Hom}(\mathcal{F}^* \otimes \nu^* \mathcal{L}, \mathcal{E}_2) \to (\mathrm{Hom}(\mathcal{F}^* \otimes \nu^* \mathcal{L}, \mathcal{E}_1^* \otimes \nu^* \mathcal{L}) \simeq \mathrm{Hom}(\mathcal{E}_1, \mathcal{F})) \\ \xrightarrow{f_1} (\mathrm{Hom}(\mathcal{F}^* \otimes \nu^* \mathcal{L}, Q_1) = V) \xrightarrow{g} \mathrm{Ext}^1(\mathcal{F}^* \otimes \nu^* \mathcal{L}, \mathcal{E}_2), \end{aligned}$$

from which it follows that

$$f_{1!} \mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1, \mathcal{F})} = q^{\dim \mathrm{Hom}(\mathcal{F}^* \otimes \nu^* \mathcal{L}, \mathcal{E}_2)} g^* \delta_{\mathrm{Ext}^1(\mathcal{F}^* \otimes \nu^* \mathcal{L}), \mathcal{E}_2)},$$

where δ_A for an abelian group A is the indicator function of $0 \in A$. Therefore

$$\begin{aligned}\mathrm{FT}_V(f_{1!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})}) &= q^{\dim \mathrm{Hom}(\mathcal{F}^*\otimes\nu^*\mathcal{L},\mathcal{E}_2)}\,\mathrm{FT}_V(g^*\delta_{\mathrm{Ext}^1(\mathcal{F}^*\otimes\nu^*\mathcal{L},\mathcal{E}_2)}) \\ &= (-1)^{\dim V} q^{\dim \chi(X,\mathcal{F}\otimes\mathcal{E}_2\otimes\nu^*\mathcal{L}^*)+\dim V} f_{2!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_2,\mathcal{F})}\end{aligned}$$

via Serre duality between $\mathrm{Hom}(\mathcal{E}_2,\mathcal{F}) \simeq \mathrm{Hom}(\mathcal{F}^*,\mathcal{E}_2^*)$ and $\mathrm{Ext}^1(\mathcal{F}^*\otimes\nu^*\mathcal{L},\mathcal{E}_2) \simeq \mathrm{Ext}^1(\mathcal{F}\otimes\omega_X^{-1},\mathcal{E}_2)$ (via $h_{\mathcal{F}}$).

(c) By Riemann–Roch $\chi(X',\widetilde{\mathcal{E}}_2\otimes\mathcal{F}\otimes\nu^*\mathcal{L}^*)$ is

$$m\deg\mathcal{F} + n\deg\widetilde{\mathcal{E}}_2 - mn[F':F]\deg\mathcal{L} - \frac{1}{2}mn\deg\omega_{X'}, \tag{2.24}$$

and the $(-\epsilon)$-Hermitian structure on $\mathcal{F}$ forces its degree to be $\frac{1}{2}n([F':F]\deg\mathcal{L}+\deg\omega_{X'})$. Therefore (2.24) is

$$\chi(X',\widetilde{\mathcal{E}}_2\otimes\mathcal{F}\otimes\nu^*\mathcal{L}^*) = n\deg\widetilde{\mathcal{E}}_2 - \frac{1}{2}mn[F':F]\deg\mathcal{L}. \tag{2.25}$$

Here all degrees are computed with respect to D, as in the discussion preceding (2.6). The short exact sequence

$$0 \to \sigma^*\widetilde{\mathcal{E}}_2 \to \widetilde{\mathcal{E}}_1^*\otimes\nu^*\mathcal{L} \to \sigma^*Q_1 \to 0$$

gives $-\deg\widetilde{\mathcal{E}}_1 + m[F':F]\deg\mathcal{L} = \deg\widetilde{\mathcal{E}}_2 + \deg\sigma^*Q_1$, so $\frac{1}{2}n(\deg\widetilde{\mathcal{E}}_1 - \deg\widetilde{\mathcal{E}}_2) = -n\deg\widetilde{\mathcal{E}}_2 + \frac{1}{2}n(m[F':F] - \deg\sigma^*Q_1) = -n\deg\widetilde{\mathcal{E}}_2 + \frac{1}{2}mn[F':F]\deg\mathcal{L} - \frac{1}{2}\dim V$, which by (2.25) is $-\chi(X',\widetilde{\mathcal{E}}_2\otimes\mathcal{F}\otimes\nu^*\mathcal{L}^*) - \frac{1}{2}\dim V$. Thus

$$\frac{1}{2}n(\deg\widetilde{\mathcal{E}}_1 - \deg\widetilde{\mathcal{E}}_2) = -\chi(X',\widetilde{\mathcal{E}}_2\otimes\mathcal{F}\otimes\nu^*\mathcal{L}^*) - \frac{1}{2}\dim V,$$

which rearranges to the claim.

(d) This is explained in [19, §2.3.9], replacing n in *loc. cit.* by nl.

□

Proof of Theorem 2.3.2 when $r = 0$. By (2.21) and (2.23), we have

$$\begin{aligned}\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_1)_{\mathcal{F}} &= \chi(\det\mathcal{E}_1)q^{\frac{n}{2}(\deg\mathcal{E}_1-\deg\mathcal{L}-\deg\omega_X)}\langle\mathfrak{q}_{12}^*\psi, f_{1!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})}\rangle_V \\ &= \chi(\det\mathcal{E}_1)q^{\frac{n}{2}(\deg\mathcal{E}_1-\deg\mathcal{L}-\deg\omega_X)-\dim V}\langle \mathrm{FT}_V([-1]^*\mathfrak{q}_{12}^*\psi), \mathrm{FT}_V(f_{1!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_1,\mathcal{F})})\rangle_V.\end{aligned}$$

By Proposition 2.4.4 (a) and (b), this is

$$\chi(\det\mathcal{E}_1)\eta_{F'/F}(D_Q)^{nl}q^{\frac{n}{2}(\deg\mathcal{E}_1-\deg\mathcal{L}-\deg\omega_X)+\frac{1}{2}\dim V+\chi(X,\mathcal{F}\otimes\mathcal{E}_2\otimes\nu^*\mathcal{L}^*)}\langle(\frac{1}{4}\mathfrak{q}_{21})^*\psi, f_{2!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_2,\mathcal{F})}\rangle$$

which by Proposition 2.4.4 (c) and (d) is

$$\chi(\det\mathcal{E}_2)q^{\frac{n}{2}(\deg\mathcal{E}_2-\deg\mathcal{L}-\deg\omega_X)}\langle\mathfrak{q}_{21}^*\psi, f_{2!}\mathbf{1}_{\mathrm{Hom}(\mathcal{E}_2,\mathcal{F})}\rangle.$$

Finally by (2.21) this is $\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_2)_{\mathcal{F}}$, so letting $\mathcal{F}$ vary we find $\widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_1) = \widetilde{Z}^0_{m,\mathcal{L}}(\mathcal{G},\mathcal{E}_2)$ whenever $\mathcal{E}_1$ and $\mathcal{E}_2$ are transverse. Finally by Proposition 2.4.2 this implies the $r = 0$ case of Theorem 2.3.2. □

2.5 Generic modularity

In Section 2.4, we studied for each $\mathcal{F} \in \mathrm{Bun}_{\widetilde{\mathcal{H}},\mathcal{L}}(k) \simeq \mathrm{Sht}^0_{\widetilde{\mathcal{H}},\mathcal{L}}$ an $\mathbb{F}_q$-vector space $V = \mathrm{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \sigma^*Q)$. If we allowed $\mathcal{F}$ to vary, this would give a family of $\mathbb{F}_q$-vector spaces over $\mathrm{Sht}^0_{\widetilde{\mathcal{H}},\mathcal{L}}$. For general r, we will construct an object $\mathrm{Sht}^r_V \to \mathrm{Sht}^r_{\widetilde{\mathcal{H}},\mathcal{L}}$, which can loosely be thought of as a family of $\mathbb{F}_q$-vector spaces and which has similar properties to V; we will then relate the special cycles $\mathcal{Z}^r_{\mathcal{E}_1}$ and $\mathcal{Z}^r_{\mathcal{E}_2}$ to Sht^r_V, generalizing the $r = 0$ case above, and use a version of Fourier duality to prove Theorem 2.3.2. Here we need the tools of derived and arithmetic Fourier analysis introduced in [19].

One complication is that to ensure the objects arising in this setup are well-behaved, we need to impose some assumptions on our Lagrangians $\mathcal{E}_i$. To avoid this issue, following [19], we observe

that for any choice of Lagrangians $\widetilde{\mathcal{E}}_i$ there exist $\mathcal{E}_i \subset \widetilde{\mathcal{E}}_i$ with torsion cokernel which do satisfy our assumptions; and with a little care we can show that the same geometric arguments can be used to show the result for $\widetilde{\mathcal{E}}_i$, not just $\mathcal{E}_i$.

2.5.1 Some derived vector bundles on $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$

By Proposition 2.4.2, we can restrict to the case of transverse sub-D-modules $\widetilde{\mathcal{E}}_1, \widetilde{\mathcal{E}}_2 \subset \mathcal{G}$, where we recall that by D-modules we mean vector bundles with $\mathcal{O}_D$-actions, and similarly below we write $\mathrm{Hom}_D(-,-)$ for morphisms in this category. We write b_{12}, b_{21}, Q_1, Q_2, Q, etc. for the notations of Section 2.4 with $\mathcal{E}_i$ replaced by $\widetilde{\mathcal{E}}_i$, and fix a Harder–Narasimhan polygon μ for $\widetilde{H}$. Let $\mathcal{E}_i \hookrightarrow \widetilde{\mathcal{E}}_i$ be sub-D-modules with cokernel $\mathcal{T}_i$ supported on a finite subset of X',[6] satisfying:

(a) the supports of Q, $\mathcal{T}_1$, and $\mathcal{T}_2$ in X' are disjoint; and

(b) for every $\widetilde{H}$-bundle $\mathcal{F} \in \mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}(\overline{k})$ we have

$$\mathrm{Ext}^1_{X'_{\overline{k}}}(\mathcal{F}^*, \mathcal{E}_i^*) = 0,$$

which by Serre duality is equivalent to

$$\mathrm{Hom}_{X'_{\overline{k}}}(\mathcal{F}^*, \mathcal{E}_i) = 0.$$

Both conditions can be ensured by taking $\mathcal{E}_i = \widetilde{E}_i \otimes \mathcal{O}(-D_i)$ for sufficiently large divisors D_i with support disjoint from that of Q.

Using assumption (a) and the transversality of $\widetilde{\mathcal{E}}_1$ and $\widetilde{\mathcal{E}}_2$, the composite map

$$\sigma^*\mathcal{E}_1 \hookrightarrow \sigma^*\widetilde{\mathcal{E}}_1 \xrightarrow{\sigma^* b_{12}} \widetilde{\mathcal{E}}_2^* \otimes \nu^*\mathcal{L} \hookrightarrow \mathcal{E}_2^* \otimes \nu^*\mathcal{L}$$

[6] Again, the $\mathcal{T}_i$ inherit a D-module structure (as torsion coherent sheaves), so $\widetilde{Q}_i$ will as well.

is injective with cokernel given by

$$\widetilde{Q}_2 := \sigma^* Q_2 \oplus \sigma^* \mathcal{T}_1 \oplus \mathcal{T}_2^* \otimes \nu^* \mathcal{L}. \tag{2.26}$$

Similarly,

$$\sigma^* \mathcal{E}_2 \hookrightarrow \sigma^* \widetilde{\mathcal{E}}_2 \xrightarrow{\sigma^* b_{21}} \widetilde{\mathcal{E}}_1^* \otimes \nu^* \mathcal{L} \hookrightarrow \mathcal{E}_1^* \otimes \nu^* \mathcal{L}$$

is injective with cokernel

$$\widetilde{Q}_1 := \sigma^* Q_1 \oplus \sigma^* \mathcal{T}_2 \oplus \mathcal{T}_1^* \otimes \nu^* \mathcal{L}. \tag{2.27}$$

The isomorphism $\beta_{12} : Q_1 \simeq \sigma^* Q_2^* \otimes \nu^* \mathcal{L}$ induces

$$\widetilde{Q}_1 \simeq Q_2^* \otimes \nu^* \mathcal{L} \oplus \mathcal{T}_2 \oplus \mathcal{T}_1^* \otimes \nu^* \mathcal{L} = \sigma^* \widetilde{Q}_2^* \otimes \nu^* \mathcal{L},$$

and similarly $\widetilde{Q}_2 \simeq \sigma^* \widetilde{Q}_1^* \otimes \nu^* \mathcal{L}$. Let $X^\circ = X - \nu(|Q| \cup |\mathcal{T}_1| \cup |\mathcal{T}_2|) = X - \nu(|\widetilde{Q}_1|) = X - \nu(|\widetilde{Q}_2|)$, where by $|\mathcal{T}|$ for any torsion coherent sheaf $\mathcal{T}$ we mean its set-theoretic support, and for any stack $\mathcal{Y} \to X'^r$ write $\mathcal{Y}^\circ = \mathcal{Y} \times_{X'^r} (X'^\circ)^r$, so e.g. $X'^\circ = \nu^{-1}(X^\circ)$.

The short exact sequence defining $\widetilde{Q}_1$ induces an exact triangle of sheaves on X'

$$\mathcal{E}_1^* \otimes \nu^* \mathcal{L} \to \widetilde{Q}_1 \to \sigma^* \mathcal{E}_2[1], \tag{2.28}$$

and we correspondingly define three perfect complexes on $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$ by taking $R\,\mathrm{Hom}(\mathcal{F} \otimes \nu^* \mathcal{L}, -)$.

Definition 2.5.1. Let $\mathcal{F}_{\mathrm{univ}}$ be the universal $\widetilde{H}$-bundle on $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$.

- Let $\mathcal{U} = \underline{R\,\mathrm{Hom}}_D(\mathcal{F}_{\mathrm{univ}}^*, \mathcal{E}_1^*)$ be the complex on $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$ whose pullback along any map $\mathrm{Spec}\, R \to \mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$, i.e. a $\widetilde{H}$-bundle $\mathcal{F}$ on $X'_R = X' \times \mathrm{Spec}\, R$, is $R\,\mathrm{Hom}_D(\mathcal{F}^*, \mathcal{E}_1^* \boxtimes R)$ as an animated R-module, i.e. derived D-linear maps of D-modules over X'_R. Let $U = \mathrm{Tot}_{\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}}(\mathcal{U})$ be the associated derived vector bundle over $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$ (as in [19, §6.1.1]).
- Let $\mathcal{V} = \underline{R\,\mathrm{Hom}}_D(\mathcal{F}_{\mathrm{univ}}^* \otimes \nu^* \mathcal{L}, \widetilde{Q}_1)$ be the complex on $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}^{\leq \mu}$ whose pullback along any

$\operatorname{Spec} R \to \operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ corresponding to $\mathcal{F}$ on X'_R is $R\operatorname{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \widetilde{Q}_1 \boxtimes R)$ as an animated R-module, and let $V = \operatorname{Tot}_{\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathcal{V})$ be the associated derived vector bundle over $\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$.

- Let $\mathcal{W} = \underline{R\operatorname{Hom}}_D(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^*\mathcal{E}_2[1])$ be the complex on $\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ whose pullback along any $\operatorname{Spec} R \to \operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ corresponding to $\mathcal{F}$ on X'_R is $R\operatorname{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \mathcal{E}_2[1] \boxtimes R)$ as an animated R-module, and let $W = \operatorname{Tot}_{\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathcal{W})$ be the associated derived vector bundle over $\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$.

Remark 2.5.2. Note that the complexes $\underline{R\operatorname{Hom}}_D(-,-)$ are well-defined as derived sheaves of abelian groups on $\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$, but do not in general carry D-actions, at least a priori. However, via the forgetful functor they do carry $\mathcal{O}_{\operatorname{Bun}_{\widetilde{H},\mathcal{L}}}$-actions and so form perfect complexes of quasicoherent sheaves, whence it makes sense to discuss the associated derived vector bundles.

The exact triangle (2.28) induces a Cartesian square of derived vector bundles

$$\begin{array}{ccc} U & \longrightarrow & V \\ \downarrow & & \downarrow \\ \operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}} & \longrightarrow & W \end{array} \tag{2.29}$$

where the bottom map is the zero section. By (b) and the fact that $\widetilde{Q}_1$ is torsion we see that in fact these are all classical vector bundles.

2.5.2 Analogues over the Hecke stack

Let $\operatorname{Hk}^r_{\widetilde{H},\mathcal{L}}$ be the stack parametrizing the same data as in Definition 2.2.1 without φ or ψ, with maps $h_i : \operatorname{Hk}^r_{\widetilde{H},\mathcal{L}} \to \operatorname{Bun}_{\widetilde{H},\mathcal{L}}$ for $0 \leq i \leq r$. Write $\operatorname{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \subset \operatorname{Hk}^r_{\widetilde{H},\mathcal{L}}$ for $\bigcap_{i=0}^r h_i^{-1}(\operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}})$.

Lemma 2.5.3. *Each map $h_i : \operatorname{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \to \operatorname{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ is smooth of relative dimension $((n-1)l+1)r$.*

Proof. As in [18, Lemma 6.9], the fibers of the map $(h_i, p_{X'}) : \operatorname{Hk}^1_{\widetilde{H},\mathcal{L}} \to \operatorname{Bun}_{\widetilde{H},\mathcal{L}} \times X'$ are determined by the choice of $\mathcal{F}^\flat_{1/2}$, which is the choice of a sub-D-module of corank 1. These are parametrized by a smooth space of dimension $(n-1)l$ (namely D-projective $n-1$-space), and so the further projection to $\operatorname{Bun}_{\widetilde{H},\mathcal{L}}$ is smooth of relative dimension $(n-1)l+1$ (since X' is a smooth

curve). Therefore each $\mathrm{Hk}^r_{\widetilde{H},\mathcal{L}} \to \mathrm{Bun}_{\widetilde{H},\mathcal{L}}$ has relative dimension $((n-1)l+1)r$. Since our h_i is the base change of such a morphism, it has the same relative dimension. □

For each S-point $(\mathcal{F}_0, \ldots, \mathcal{F}_r; \mathcal{F}^\flat_{1/2}, \ldots, \mathcal{F}^\flat_{r-1/2})$ of $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$, we define $\mathcal{F}^\flat_\bullet$ on $X' \times S$ to be the perfect complex

$$(\mathcal{F}^\flat_{1/2} \oplus \cdots \oplus \mathcal{F}^\flat_{r-1/2}) \to (\mathcal{F}_1 \oplus \ldots \oplus \mathcal{F}_{r-1})$$

in degrees 0 and 1, where the differential sends

$$(s_{1/2}, \ldots, s_{r-1/2}) \mapsto (s_{1/2} - s_{3/2}, \ldots, s_{r-3/2} - s_{r-1/2})$$

after viewing $\mathcal{F}^\flat_{i-1/2}$ as subsheaves of $\mathcal{F}_{i-1}$ and $\mathcal{F}_i$ via the given embeddings. Letting the point of $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}(S)$ and S vary, this gives a perfect complex $\mathcal{F}^\flat_{\bullet,\mathrm{univ}}$ on $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$.

Definition 2.5.4. With $\mathcal{F}^\flat_{\bullet,\mathrm{univ}}$ as above,

- let $\mathcal{U}^\flat_{\mathrm{Hk}} = \underline{R\,\mathrm{Hom}}_D(\mathcal{F}^{\flat *}_{\bullet,\mathrm{univ}}, \mathcal{E}_1^*)$ and $\mathrm{Hk}^\flat_U = \mathrm{Tot}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathcal{U}^\flat_{\mathrm{Hk}})$ be the associated derived vector bundle over $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$;[7]
- let $\mathcal{V}^\flat_{\mathrm{Hk}} = \underline{R\,\mathrm{Hom}}_D(\mathcal{F}^{\flat *}_{\bullet,\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{Q}_1)$ and $\mathrm{Hk}^\flat_V = \mathrm{Tot}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathcal{U}^\flat_{\mathrm{Hk}})$;
- and let $\mathcal{W}^\flat_{\mathrm{Hk}} = \underline{R\,\mathrm{Hom}}_D(\mathcal{F}^{\flat *}_{\bullet,\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^*\mathcal{E}_2[1])$ and $\mathrm{Hk}^\flat_W = \mathrm{Tot}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathcal{W}^\flat_{\mathrm{Hk}})$.

The triangle (2.28) again induces an exact triangle of perfect complexes on $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$

$$\mathcal{U}^\flat_{\mathrm{Hk}} \to \mathcal{V}^\flat_{\mathrm{Hk}} \to \mathcal{W}^\flat_{\mathrm{Hk}}$$

[7] Similar to [19, Remark 9.1.2], $\mathrm{Hk}^\flat_U$ is isomorphic to the derived pullback of the derived Hecke stack $\mathrm{Hk}^r_{\mathscr{M}} \to \mathrm{Bun}_M \times \mathrm{Hk}^r_{\widetilde{H},\mathcal{L}}$ along $\{\mathcal{E}_1\} \times \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \to \mathrm{Bun}_M \times \mathrm{Hk}^r_{\widetilde{H},\mathcal{L}}$.

and thus a derived Cartesian square of derived vector bundles over $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{U} & \xrightarrow{f} & \mathrm{Hk}^{\flat}_{V} \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle g} \\ \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} & \xrightarrow{z} & \mathrm{Hk}^{\flat}_{W} \end{array} \tag{2.30}$$

where the bottom map z is again the zero section.

For each $0 \leq i \leq r-1$, the projection $\mathcal{F}^{\flat}_{\bullet} \to \mathcal{F}^{\flat}_{i+1/2}$ together with the inclusion $\mathcal{F}_{i+1/2} \hookrightarrow \mathcal{F}_i$ gives a map $\mathcal{F}^{\flat}_{\bullet} \to \mathcal{F}_i$; similarly for $1 \leq i \leq r$ the composition $\mathcal{F}^{\flat}_{\bullet} \to \mathcal{F}^{\flat}_{i-1/2} \hookrightarrow \mathcal{F}_i$ likewise gives such a map. For $1 \leq i \leq r-1$, these constructions agree. Dualizing gives maps

$$\mathcal{F}^{*}_{i} \to \mathcal{F}^{\flat *}_{\bullet}$$

for each i. We can view $\mathcal{F}^{*}_{i}$ as arising from the universal bundle $\mathcal{F}_{\mathrm{univ}}$ on $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ via pullback along the ith projection $h_i : \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \to \mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$. Since the perfect complexes and derived vector bundles of Definitions 2.5.1 and 2.5.4 respectively are defined by taking derived homomorphisms from $\mathcal{F}^{*}_{\mathrm{univ}}$ and $\mathcal{F}^{\flat *}_{\bullet,\mathrm{univ}}$, these induce maps

$$\mathcal{U}^{\flat}_{\mathrm{Hk}} \to h_i^{*}\mathcal{U}, \qquad \mathcal{V}^{\flat}_{\mathrm{Hk}} \to h_i^{*}\mathcal{V}, \qquad \mathcal{W}^{\flat}_{\mathrm{Hk}} \to h_i^{*}\mathcal{W}.$$

Writing $\widetilde{U}_i$, $\widetilde{V}_i$ and $\widetilde{W}_i$ for the derived vector bundles on $\mathrm{Hk}^{r}_{\widetilde{H},\mathcal{L}}$ given by the pullbacks along h_i of U, V, and W respectively, or equivalently the derived vector bundles associated to the complexes $h_i^{*}\mathcal{U}$, $h_i^{*}\mathcal{V}$, and $h_i^{*}\mathcal{W}$, we obtain maps

$$a_i : \mathrm{Hk}^{\flat}_{U} \to U, \qquad b_i : \mathrm{Hk}^{\flat}_{V} \to V, \qquad c_i : \mathrm{Hk}^{\flat}_{W} \to W.$$

Lemma 2.5.5. *For each* $0 \leq i \leq r$,

(a) each $\widetilde{a}_i$ *is a quasi-smooth closed embedding, and* a_i *is quasi-smooth;*

(b) each $\widetilde{b}_i$ and b_i is quasi-smooth and separated (and so is representable in derived schemes);

(c) each $\widetilde{c}_i$ is a smooth vector bundle, and each c_i is smooth.

Proof. The proofs of the properties of the $\widetilde{a}_i$, $\widetilde{b}_i$, and $\widetilde{c}_i$ are identical to those of [19, Lemma 9.1.3]. Since $h_i : \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \to \mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$ is smooth by Lemma 2.5.3, so is the projection $\widetilde{U}_i \to U$, and likewise for U replaced by V and W. Therefore the claimed properties of a_i, b_i, and c_i follow from the respective properties of $\widetilde{a}_i$, $\widetilde{b}_i$, and $\widetilde{c}_i$. □

Corollary 2.5.6. *The diagram*

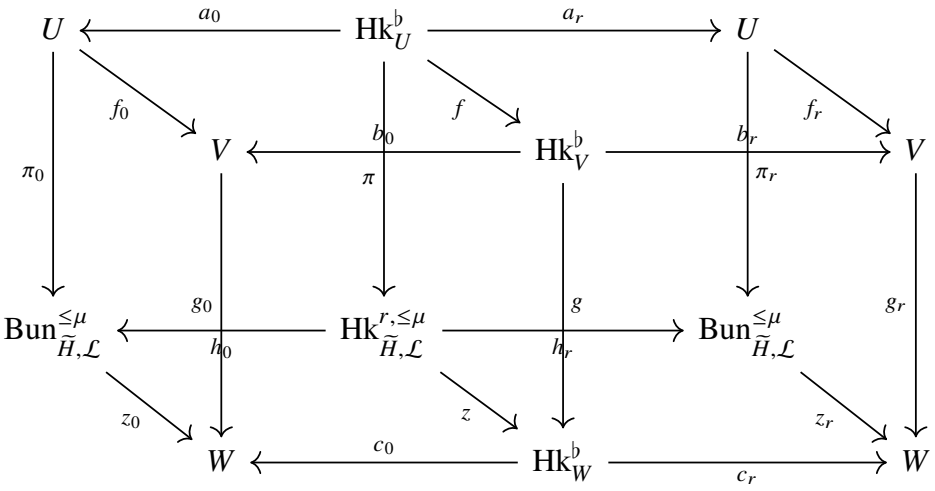

satisfies the conditions of [19, §5.1.1].

With Lemma 2.5.5 established, the proof is formally identical to the proof of Corollary 9.1.5 of [19].

Lemma 2.5.7. *For each $0 \leq i \leq r$, we have $d(\widetilde{a}_i) + d(\widetilde{c}_i) = d(\widetilde{b}_i)$.*

Here d denotes the virtual relative dimension, i.e. the Euler characteristic of the tangent complex.

Proof. By construction $\widetilde{a}_i$ is a map between derived vector bundles over $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ and so has relative

dimension

$$\begin{aligned} d(\widetilde{a}_i) &= \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{U} - \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} U \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}},h_i} \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} \\ &= \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{U} - \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} U. \end{aligned}$$

Similarly

$$\begin{aligned} d(\widetilde{b}_i) &= \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{V} - \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V, \\ d(\widetilde{c}_i) &= \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{W} - \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} W, \end{aligned}$$

and by (2.29) and (2.30) respectively we have

$$\begin{aligned} \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} U + \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} W &= \operatorname{rank}_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V, \\ \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{U} + \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{W} &= \operatorname{rank}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^{\flat}_{V}, \end{aligned}$$

whence the claim follows. □

2.5.3 Duality

Exchanging the roles of $\mathcal{E}_1$ and $\mathcal{E}_2$, parallel to (2.28) we have an exact triangle

$$\mathcal{E}_2^* \otimes \nu^* \mathcal{L} \to \widetilde{Q}_2 \to \mathcal{E}_1[1] \tag{2.31}$$

and correspondingly define $\mathcal{U}^\perp$, $\mathcal{V}'$, and $\mathcal{W}^\perp$ with derived vector bundles $U^\perp$, V', and $W^\perp$ respectively parallel to Definition 2.5.1, with the roles of $\mathcal{E}_1$ and $\mathcal{E}_2$ (resp. $\widetilde{Q}_1$ and $\widetilde{Q}_2$) switched. As for

Definition 2.5.1, these will in fact be classical vector bundles, so the derived Cartesian square

$$\begin{array}{ccc} U^{\perp} & \longrightarrow & V' \\ \downarrow & & \downarrow \\ \mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}} & \longrightarrow & W \end{array} \tag{2.32}$$

can be thought of as a short exact sequence of classical vector bundles

$$0 \to U^{\perp} \to V' \to W^{\perp} \to 0,$$

so we can think of $U^{\perp}$ as a subbundle of V' and $W^{\perp}$ as a quotient.

Lemma 2.5.8. *Serre duality identifies V' with the dual vector bundle of V over $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$. Under this duality, U is dual to $W^{\perp}$ and W is dual to $U^{\perp}$, so that* (2.29) *and* (2.32) *are dual fiber squares.*

Proof. For an R-point $\mathcal{F}$ of $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$, the pullback of $\mathcal{V}$ along this point is

$$R\,\mathrm{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \widetilde{Q}_1 \boxtimes R) \simeq R\,\mathrm{Hom}_D(\widetilde{Q}_1^*[-1] \otimes \nu^*\mathcal{L} \boxtimes R, \mathcal{F}), \tag{2.33}$$

which by Serre duality is R-dual to $R\,\mathrm{Hom}_D(\mathcal{F}, \widetilde{Q}_1^* \otimes \nu^*\mathcal{L} \otimes \omega_X \boxtimes R)$. Applying the isomorphism $h_{\mathcal{F}} : \mathcal{F} \xrightarrow{\sim} \mathcal{F}^{\vee} \otimes \nu^*\mathcal{L}$, both sides of (2.33) are isomorphic to

$$\begin{aligned} R\,\mathrm{Hom}_D(\mathcal{F}^{\vee} \otimes \nu^*\mathcal{L}, \widetilde{Q}_1^* \otimes \nu^*\mathcal{L} \otimes \omega_X \boxtimes R) &\simeq R\,\mathrm{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \widetilde{Q}_1^* \otimes \nu^*\mathcal{L} \boxtimes R) \\ &\simeq R\,\mathrm{Hom}_D(\mathcal{F}^* \otimes \nu^*\mathcal{L}, \widetilde{Q}_2) \end{aligned}$$

(bearing in mind that in fact everything here is concentrated in degree 0). This is the value of $\mathcal{V}'$ on $\mathcal{F}$, and as everything is functorial in $\mathcal{F}$ this gives our desired duality between V and V'. Under this duality, the exact triangle

$$R\,\mathrm{Hom}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{E}_1^*) \to R\,\mathrm{Hom}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{Q}_1) \to R\,\mathrm{Hom}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \mathcal{E}_2[1])$$

is dual to the exact triangle

$$R\operatorname{Hom}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{E}^*_2) \to R\operatorname{Hom}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{Q}_2) \to R\operatorname{Hom}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \mathcal{E}_1[1]),$$

which is the desired duality on the level of perfect complexes; taking associated derived vector bundles gives the duality of fiber squares. □

We would like to have an analogous duality between Hecke stacks. First, though, we need to introduce auxiliary Hecke stacks $\mathrm{Hk}^\sharp_{\_}$. One can give an explicit definition involving a perfect complex $\mathcal{F}^\sharp_{\bullet,\mathrm{univ}}$, completely parallel to the definition of the stacks $\mathrm{Hk}^\flat_{\_}$ using the prefect complex $\mathcal{F}^\flat_{\bullet,\mathrm{univ}}$, but a more parsimonious definition is to define them as the derived vector bundles over $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ fitting into the the derived Cartesian squares

$$\begin{array}{ccc} \mathrm{Hk}^\flat_U & \xrightarrow{\widetilde{a}_0} & \widetilde{U}_0 \\ \downarrow{\scriptstyle \widetilde{a}_r} & & \downarrow{\scriptstyle \widetilde{a}'_r} \\ \widetilde{U}_r & \xrightarrow{\widetilde{a}'_0} & \mathrm{Hk}^\sharp_U \end{array}, \tag{2.34}$$

$$\begin{array}{ccc} \mathrm{Hk}^\flat_V & \xrightarrow{\widetilde{b}_0} & \widetilde{V}_0 \\ \downarrow{\scriptstyle \widetilde{b}_r} & & \downarrow{\scriptstyle \widetilde{b}'_r} \\ \widetilde{V}_r & \xrightarrow{\widetilde{b}'_0} & \mathrm{Hk}^\sharp_V \end{array}, \tag{2.35}$$

$$\begin{array}{ccc} \mathrm{Hk}^\flat_W & \xrightarrow{\widetilde{c}_0} & \widetilde{W}_0 \\ \downarrow{\scriptstyle \widetilde{c}_r} & & \downarrow{\scriptstyle \widetilde{c}'_r} \\ \widetilde{W}_r & \xrightarrow{\widetilde{c}'_0} & \mathrm{Hk}^\sharp_W \end{array}, \tag{2.36}$$

so that the $\mathrm{Hk}^\sharp_{\_}$ can be defined as pushouts. We define $\mathrm{Hk}^\flat_{U^\perp}$, $\mathrm{Hk}^\flat_{V'}$, and $\mathrm{Hk}^\flat_{W^\perp}$ parallel to Defini-

tion 2.5.4, fitting into a derived Cartesian square

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{U^\perp} & \xrightarrow{f^\perp} & \mathrm{Hk}^{\flat}_{V'} \\ \downarrow{\scriptstyle \pi^\perp} & & \downarrow{\scriptstyle g^\perp} \\ \mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}} & \xrightarrow{z^\perp_{\mathrm{Hk}}} & \mathrm{Hk}^{\flat}_{W^\perp} \end{array}, \tag{2.37}$$

and then similarly define the $\mathrm{Hk}^{\sharp}_{\_}$-stacks by

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{U^\perp} & \xrightarrow{\widetilde{a}^\perp_0} & \widetilde{U}^\perp_0 \\ \downarrow{\scriptstyle \widetilde{a}^\perp_r} & & \downarrow{\scriptstyle \widetilde{a}'^\perp_r} \\ \widetilde{U}^\perp_r & \xrightarrow{\widetilde{a}'^\perp_0} & \mathrm{Hk}^{\sharp}_{U^\perp} \end{array}, \tag{2.38}$$

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{V'} & \xrightarrow{\widetilde{\beta}_0} & \widetilde{V}'_0 \\ \downarrow{\scriptstyle \widetilde{\beta}_r} & & \downarrow{\scriptstyle \widetilde{\beta}'_r} \\ \widetilde{V}'_r & \xrightarrow{\widetilde{\beta}'_0} & \mathrm{Hk}^{\sharp}_{V'} \end{array}, \tag{2.39}$$

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{W^\perp} & \xrightarrow{\widetilde{c}^\perp_0} & \widetilde{W}^\perp_0 \\ \downarrow{\scriptstyle \widetilde{c}^\perp_r} & & \downarrow{\scriptstyle \widetilde{c}'^\perp_r} \\ \widetilde{W}^\perp_r & \xrightarrow{\widetilde{c}'^\perp_0} & \mathrm{Hk}^{\sharp}_{W^\perp} \end{array} \tag{2.40}$$

as above.

Lemma 2.5.9. *Fourier duality of derived vector bundles on* $\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ *gives identifications*

(a) $\mathrm{Hk}^{\sharp}_{U^\perp} \simeq \widehat{\mathrm{Hk}^{\flat}_{W}}$ *and* $\mathrm{Hk}^{\flat}_{U^\perp} \simeq \widehat{\mathrm{Hk}^{\sharp}_{W}}$*;*

(b) $\mathrm{Hk}^{\sharp}_{V'} \simeq \widehat{\mathrm{Hk}^{\flat}_{V}}$ *and* $\mathrm{Hk}^{\flat}_{V'} \simeq \widehat{\mathrm{Hk}^{\sharp}_{V}}$*;*

(c) $\mathrm{Hk}^{\sharp}_{W^\perp} \simeq \widehat{\mathrm{Hk}^{\flat}_{U}}$ *and* $\mathrm{Hk}^{\flat}_{W^\perp} \simeq \widehat{\mathrm{Hk}^{\sharp}_{U}}$.

Under these identifications, the fiber squares (2.34) *and* (2.40)*,* (2.35) *and* (2.39)*, and* (2.36) *and* (2.38) *are Fourier dual pairs. In particular*

$$\widetilde{\beta}_0 = \widehat{\widetilde{b}'_r}, \qquad \widetilde{\beta}_r = \widehat{\widetilde{b}'_0}, \qquad \widetilde{\beta}'_r = \widehat{\widetilde{b}}_0, \qquad \widetilde{\beta}'_0 = \widehat{\widetilde{b}}_r.$$

This is just the analogue of Lemma 2.5.8 over $\mathrm{Hk}^r_{\widetilde{H},\mathcal{L}}$, and the proof is the same translated to the Hecke setting.

Recall from Lemma 2.5.5 that $a_r : \mathrm{Hk}^\flat_U \to U$ is a quasi-smooth map and so its relative fundamental class defines a cohomological correspondence

$$\mathfrak{c}_U = [a_r] \in \mathrm{Corr}_{\mathrm{Hk}^\flat_U}(\mathbb{Q}_{\ell,U}, \mathbb{Q}_{\ell,U\langle -d(a_r)\rangle})$$

supported on the correspondence $U \xleftarrow{a_0} \mathrm{Hk}^\flat_U \xrightarrow{a_r} U$. Dually we have a quasi-smooth map $a_r : \mathrm{Hk}^\flat_{U^\perp} \to U^\perp$ defining a cohomological correspondence

$$\mathfrak{c}_{U^\perp} = [a_r^\perp] \in \mathrm{Corr}_{\mathrm{Hk}^\flat_{U^\perp}}(\mathbb{Q}_{\ell,U^\perp}, \mathbb{Q}_{\ell,U^\perp\langle -d(a_r^\perp)\rangle})$$

supported on $U^\perp \xleftarrow{a_0^\perp} \mathrm{Hk}^\flat_{U^\perp} \xrightarrow{a_r^\perp} U^\perp$. By Corollary 2.5.6 we can push forward these correspondences along the maps of correspondences

$$\begin{array}{ccccc} U & \xleftarrow{a_0} & \mathrm{Hk}^\flat_U & \xrightarrow{a_r} & U \\ \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_r} \\ V & \xleftarrow{b_0} & \mathrm{Hk}^\flat_V & \xrightarrow{b_r} & V \end{array}, \qquad \begin{array}{ccccc} U^\perp & \xleftarrow{a_0^\perp} & \mathrm{Hk}^\flat_{U^\perp} & \xrightarrow{a_r^\perp} & U^\perp \\ \downarrow{\scriptstyle f_0^\perp} & & \downarrow{\scriptstyle f^\perp} & & \downarrow{\scriptstyle f_r^\perp} \\ V' & \xleftarrow{b_0'} & \mathrm{Hk}^\flat_{V'} & \xrightarrow{b_r'} & V' \end{array}$$

to get

$$f_!\mathfrak{c}_U \in \mathrm{Corr}_{\mathrm{Hk}^\flat_V}(f_{0!}\mathbb{Q}_{\ell,U}, f_{r!}\mathbb{Q}_{\ell,U\langle -d(a_r)\rangle})$$

and

$$f_!^\perp\mathfrak{c}_{U^\perp} \in \mathrm{Corr}_{\mathrm{Hk}^\flat_{V'}}(f_{0!}^\perp\mathbb{Q}_{\ell,U^\perp}, f_{r!}^\perp\mathbb{Q}_{\ell,U^\perp\langle -d(a_r^\perp)\rangle}).$$

We would like to understand

$$\mathrm{FT}_{\mathrm{Hk}^\flat_V}(f_!\mathfrak{c}_U) \in \mathrm{Corr}_{\mathrm{Hk}^\flat_{V'}}(\mathrm{FT}_V(f_{0!}\mathbb{Q}_{\ell,U}), \mathrm{FT}_V(f_{r!}\mathbb{Q}_{\ell,U})_{[d(\widetilde{b}_0)+d(\widetilde{b}_r)](d(\widetilde{b}_0))\langle -d(a_r)\rangle}).$$

Since $d(\widetilde{b}_0) = d(\widetilde{b}_r)$, this is equivalently an element of

$$\mathrm{Corr}_{\mathrm{Hk}^{\flat}_{V'}}(\mathrm{FT}_V(f_{0!}\mathbb{Q}_{\ell,U}), \mathrm{FT}_V(f_{r!}\mathbb{Q}_{\ell,U})\langle d(\widetilde{b}_r)-d(a_r)\rangle).$$

Via the Hermitian structure on V, we can look at the orthogonal complement of a subbundle I of V, which is naturally identified with the Fourier dual of V/I; for $I = U$, Lemma 2.5.8 identifies the dual of $V/U = W$ with $U^{\perp}$, so that as in [19] we have

$$\mathrm{FT}_V(f_{i!}\mathbb{Q}_{\ell,U}) \simeq f_{i!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}[\mathrm{rank}\, V]\langle -\mathrm{rank}\, U\rangle}$$

(see the discussion preceding Theorem 9.4.1 there, as well as §6.2.4 and Example 6.2.3). By Lemma 2.5.7, we have $d(\widetilde{b}_r) - d(\widetilde{a}_r) = d(\widetilde{c}_r)$, and by Lemma 2.5.9 $\widetilde{c}_r$ is dual to $\widetilde{a}_0'^{\perp}$, which has the same relative dimension as its pullback $\widetilde{a}_0^{\perp}$. Therefore $d(\widetilde{b}_r) - d(a_r) = d(\widetilde{b}_r) - d(\widetilde{a}_r) - d(h_r) = -d(\widetilde{a}_0^{\perp}) - d(h_r) = -d(a_0^{\perp}) = -d(a_r^{\perp})$. Thus we can view $\mathrm{FT}_{\mathrm{Hk}^{\flat}_V}(f_!\mathfrak{c}_U)$ as an element of

$$\mathrm{Corr}_{\mathrm{Hk}^{\flat}_{V'}}(f_{0!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}[\mathrm{rank}\, V]\langle -\mathrm{rank}\, U\rangle}, f_{r!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}[\mathrm{rank}\, V]\langle -\mathrm{rank}\, U-d(a_r^{\perp})\rangle})$$

or equivalently

$$\mathbb{T}_{[\mathrm{rank}\, V-2\,\mathrm{rank}\, U](-\mathrm{rank}\, U)}\,\mathrm{Corr}_{\mathrm{Hk}^{\flat}_{V'}}(f_{0!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}}, f_{r!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}\langle -d(a_r^{\perp})\rangle}).$$

Up to shift and twist, this is the same space as the dual correspondence $f_!^{\perp}\mathfrak{c}_{U^{\perp}}$, so we might wonder if we can relate these two elements of this space. An affirmative answer is given by the following theorem:

Theorem 2.5.10. *We have*

$$\mathbb{T}_{[2\,\mathrm{rank}\, U-\mathrm{rank}\, V](\mathrm{rank}\, U)}\,\mathrm{FT}_{\mathrm{Hk}^{\flat}_V}(f_!\mathfrak{c}_U) = f_!^{\perp}\mathfrak{c}_{U^{\perp}} \in \mathrm{Corr}_{\mathrm{Hk}^{\flat}_{V'}}(f_{0!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}}, f_{r!}^{\perp}\mathbb{Q}_{\ell,U^{\perp}\langle -d(a_r^{\perp})\rangle}).$$

Proof. Let $\mathfrak{s} = [h_r] : h_0^*\overline{\mathbb{Q}}_{\ell} \to h_r^!\overline{\mathbb{Q}}_{\ell}$, and note that $\mathfrak{c}_U = \pi^*\mathfrak{s}$, $\mathfrak{c}_{U^{\perp}} = \pi^{\perp *}\mathfrak{s}$ (completely parallel to

[19, Lemma 9.4.2]), and $\mathrm{FT}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathfrak{s}) = \mathfrak{s}$ since this is the trivial vector bundle. Therefore by [19, Proposition 7.2.4] we have

$$\begin{aligned}
\mathbb{T}_{[2\operatorname{rank}U-\operatorname{rank}V](\operatorname{rank}U)}\,\mathrm{FT}_{\mathrm{Hk}^{\flat}_{V}}(f_!\mathfrak{c}_U) &= \mathbb{T}_{[2\operatorname{rank}U-\operatorname{rank}V](\operatorname{rank}U)}\,\mathrm{FT}_{\mathrm{Hk}^{\flat}_{V}}(f_!\pi^*\mathfrak{s}) \\
&= \mathbb{T}_{[\operatorname{rank}U](\operatorname{rank}U)}g^{\perp *}\,\mathrm{FT}_{\mathrm{Hk}^{\flat}_{U}}(\pi^*\mathfrak{s}) \\
&= g^{\perp *}z^{\perp}_{\mathrm{Hk}!}\,\mathrm{FT}_{\mathrm{Hk}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}}(\mathfrak{s}) \\
&= g^{\perp *}z^{\perp}_{\mathrm{Hk}!}\mathfrak{s}.
\end{aligned}$$

Finally by Corollary 2.5.6 we can apply [19, Theorem 5.1.3] to see that the right-hand side is $f_!^{\perp}\pi^{\perp *}\mathfrak{s} = f_!^{\perp}\mathfrak{c}_{U^{\perp}}$ as claimed. □

We would like to take shtuka traces of both sides of Theorem 2.5.10 to obtain an arithmetic Fourier duality result relating the virtual fundamental classes of our special cycles, or approximations of them. There is an immediate issue, however: $\mathrm{Sht}^r_V = \mathrm{Sht}(\mathrm{Hk}^{\flat}_V)$ is not in general a relative $\mathbb{F}_q$-vector space over $\mathrm{Sht}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$, and so one cannot apply the arithmetic Fourier theory of [19, §8]. Restricting to $\mathrm{Sht}^{r,\circ}_V$, however, the problem disappears:

Lemma 2.5.11. *For any $0 \leq i \leq r$, the restrictions $\widetilde{b}_i^{\circ} : \mathrm{Hk}^{\flat,\circ}_V \to V \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}},h_i} \mathrm{Hk}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$ and $\widetilde{b}_i'^{\circ} : V \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}},h_i} \mathrm{Hk}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}} \to \mathrm{Hk}^{\sharp\circ}_V$ are isomorphisms.*

Proof. Choose an R-point $\mathcal{F}_\star = (\mathcal{F}_0, \ldots, \mathcal{F}_r; \mathcal{F}^{\flat}_{1/2}, \ldots, \mathcal{F}^{\flat}_{r-1/2}; \mathcal{F}^{\sharp}_{1/2}, \ldots, \mathcal{F}^{\sharp}_{r-1/2})$. The map $\widetilde{b}_i$ (or rather its pullback along this R-point) is induced by the injection $\mathcal{F}_i^* \hookrightarrow \mathcal{F}_\bullet^{\flat *}$ with cokernel supported on the legs of $\mathcal{F}_\star$, which by assumption are disjoint from the support of $\widetilde{Q}_1$ and so this map induces an isomorphism

$$R\operatorname{Hom}_D(\mathcal{F}_\bullet^{\flat *} \otimes \nu^*\mathcal{L}, \widetilde{Q}_1 \boxtimes R) \xrightarrow{\sim} R\operatorname{Hom}_D(\mathcal{F}_i^*, \widetilde{Q}_1 \boxtimes R),$$

so varying R gives the desired isomorphism. The argument for $\widetilde{b}_i'$ is similar. □

Corollary 2.5.12. *The projection* $\mathrm{Sht}_V^{r,\circ} \to \mathrm{Sht}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$ *is a relative finite-dimensional* $\mathbb{F}_q$*-vector space, i.e. an étale locally free* $\mathbb{F}_q$*-vector space bundle of finite rank.*

Proof. By Lemma 2.5.11, $\mathrm{Hk}_V^{\flat,\circ}$ is a classical vector bundle over $\mathrm{Hk}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$ since V is classical over $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$, so the result follows from [19, Lemma 8.3.1]. □

This enables us to use [19, Theorem 8.3.2] to apply the duality results of Lemma 2.5.9 to $\mathrm{Sht}_V^{r,\circ}$.

Corollary 2.5.13. *We have*

$$\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}_V^{r,\circ}}(\mathrm{Sht}(f)_!^\circ[\mathrm{Sht}_U^{r,\circ}]) = (-1)^{\mathrm{rank}\, V} q^{\mathrm{rank}\, U} \cdot \mathrm{Sht}(f^\perp)_!^\circ[\mathrm{Sht}_{U^\perp}^{r,\circ}]$$

in $H^{\mathrm{BM}}_{2((n-m-1)l+1)r}(\mathrm{Sht}_{V'}^{r,\circ})$.

Proof. Since U and $U^\perp$ are classical vector bundles over $\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}$, they are in particular smooth and so by [19, Corollary 4.7.3] we have

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U) = [\mathrm{Sht}^r_U] \in H^{\mathrm{BM}}_{2d(a_r)}(\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}) \quad \text{and} \quad \mathrm{Tr}(\mathfrak{c}_{U^\perp}) = [\mathrm{Sht}^r_{U^\perp}] \in H^{\mathrm{BM}}_{2d(a_r^\perp)}(\mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}),$$

where $\mathrm{Tr}^{\mathrm{Sht}}(-)$ is the shtuka trace, as defined in [19, §4.2]. Since f, f_0, f_r, and their $-^\perp$ versions are all closed embeddings, they are all proper, and therefore by Proposition 4.5.1 in *loc. cit.* we have

$$\mathrm{Tr}^{\mathrm{Sht}}(f_!\mathfrak{c}_U) = \mathrm{Sht}(f)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U) = \mathrm{Sht}(f)_![\mathrm{Sht}^r_U].$$

Similarly we have

$$\mathrm{Tr}^{\mathrm{Sht}}(f_!^\perp\mathfrak{c}_{U^\perp}) = \mathrm{Sht}(f^\perp)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{U^\perp}) = \mathrm{Sht}(f^\perp)_![\mathrm{Sht}^r_{U^\perp}],$$

and the same holds after restriction to $(X'^\circ)^r$. Now we apply [19, Theorem 8.3.2] with $E = V$,

$C_E = \mathrm{Hk}_V^{\flat,\circ}$, and $\mathfrak{c} = (f_! \mathfrak{c}_U)|_{\mathrm{Hk}_V^{\flat,\circ}}$ together with Theorem 2.5.10 to get

$$\begin{aligned}\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}_V^{r,\circ}}(\mathrm{Sht}(f)^\circ_! [\mathrm{Sht}_U^{r,\circ}]) &= \mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}_V^{r,\circ}}(\mathrm{Tr}^{\mathrm{Sht}}(f_! \mathfrak{c}_U)|_{\mathrm{Sht}_V^{r,\circ}}) \\ &= \mathrm{Tr}^{\mathrm{Sht}}(\mathrm{FT}_{\mathrm{Hk}_V^\flat}(f_! \mathfrak{c}_U))|_{\mathrm{Sht}_{V'}^{r,\circ}} \\ &= \mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{[\mathrm{rank}\, V - 2\,\mathrm{rank}\, U](-\mathrm{rank}\, U)} f_!^\perp \mathfrak{c}_{U^\perp})|_{\mathrm{Sht}_{V'}^{r,\circ}}.\end{aligned}$$

Applying equation (4.2.6) from *loc. cit.* to calculate the shtuka trace of the shift and twist, we find that the right-hand side is

$$(-1)^{\mathrm{rank}\, V} q^{\mathrm{rank}\, U}\, \mathrm{Tr}^{\mathrm{Sht}}(f_!^\perp \mathfrak{c}_{U^\perp})|_{\mathrm{Sht}_{V'}^{r,\circ}} = (-1)^{\mathrm{rank}\, V} q^{\mathrm{rank}\, U}\, \mathrm{Sht}(f^\perp)^\circ_! [\mathrm{Sht}_{U^\perp}^{r,\circ}] \tag{2.41}$$

as desired.

To conclude, we observe that although the class (2.41) is a priori an element of $H^{\mathrm{BM}}_{2d(a_r^\perp)}(\mathrm{Sht}_{V'}^{r,\circ})$, we can compute $d(a_r^\perp) = ((n-m-1)l+1)r$ and so it is in the expected degree. Indeed, we have $a_r^\perp = h_r \circ \widetilde{a}_r^\perp$, so it suffices to compute the relative dimensions of h_r and $\widetilde{a}_r^\perp$. By Lemma 2.5.3, the relative dimension of h_r is $((n-1)l+1)r$. By Lemma 2.5.5 (or rather its equivalent under exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$), $\widetilde{a}_r^\perp$ is a quasi-smooth closed embedding of derived vector bundles with cokernel the derived vector bundle associated to $\underline{R\,\mathrm{Hom}}_D(T_{r,\mathrm{univ}}, \mathcal{E}_2^*)$ for $T_{r,\mathrm{univ}}$ a torsion sheaf of D-modules of degree r (as a D-module). Therefore this cokernel has virtual D-rank $-\deg_D T_{r,\mathrm{univ}} \cdot \mathrm{rank}_D\, \mathcal{E}_2^* = -mr$ and therefore has virtual rank $-mlr$ as a derived vector bundle on X'. Therefore in total $d(a_r^\perp) = ((n-1)l+1)r - mlr = ((n-m-1)l+1)r$. □

2.5.4 Theta series and modularity

It remains to deduce the modularity result from Corollary 2.5.13. We start by observing that Sht_U^r can be reinterpreted in more familiar language: it is just a derived version of $\mathcal{Z}^{r,\leq\mu}_{\mathcal{E}_1,\mathcal{L}}$. More precisely, it is the pullback to $\mathrm{Sht}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ of the (derived) fiber of $\mathrm{Sht}^r_{\mathscr{M}}$ over $\{\mathcal{E}_1\} \times \mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$, and $\mathcal{Z}^{r,\leq\mu}_{\mathcal{E}_1,\mathcal{L}}$ is its classical truncation. Therefore $[\mathrm{Sht}_U^r] = [\mathrm{Sht}^r_{\mathscr{M}}]|_{\mathcal{Z}^{r,\leq\mu}_{\mathcal{E}_1,\mathcal{L}}}$, which by definition is (after

pushforward to the base) the $\leq \mu$-restriction of the virtual fundamental class $[\mathcal{Z}^r_{\mathcal{E}_1,\mathcal{L}}]$.

This is not quite the right thing for modularity, though: we want to study the theta series, rather than just the special cycle, and to use the initial bundles $\widetilde{\mathcal{E}}_i$ rather than their sub-D-modules $\mathcal{E}_i$. We will obtain these modifications by multiplying the special cycle $\mathrm{Sht}(f)^\circ_!\,[\mathrm{Sht}^{r,\circ}_U]$ by certain functions on $\mathrm{Sht}^{r,\circ}_V$, and then with everything in place we can prove Theorem 2.3.2 via a geometric generalization of the proof of the $r = 0$ case appearing in Section 2.4.

The decompositions (2.26), (2.27)

$$\widetilde{Q}_1 = \sigma^* Q_1 \oplus \mathcal{T}_1^* \otimes \nu^*\mathcal{L} \oplus \sigma^*\mathcal{T}_2, \qquad \widetilde{Q}_2 = \sigma^* Q_2 \oplus \sigma^*\mathcal{T}_1 \oplus \mathcal{T}_2^* \otimes \nu^*\mathcal{L}$$

induce isomorphisms

$$\begin{aligned}\mathcal{V} &= \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{Q}_1) \\ &= \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^* Q_1) \oplus \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{T}_1^*) \oplus \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^*\mathcal{T}_2)\end{aligned}$$

and thus

$$V = V^{(0)} \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V^{(1)} \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V^{(2)} \tag{2.42}$$

where $V^{(i)}$ are the vector bundles associated to each factor. Similarly we have a decomposition

$$V' = V'^{(0)} \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V'^{(1)} \times_{\mathrm{Bun}^{\leq\mu}_{\widetilde{H},\mathcal{L}}} V'^{(2)} \tag{2.43}$$

where $V'^{(0)}$, $V'^{(1)}$, and $V'^{(2)}$ are the derived vector bundles corresponding to $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^* Q_2)$, $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \sigma^*\mathcal{T}_1)$, and $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{T}_2^*)$ respectively, and the duality of Lemma 2.5.8 preserves these decompositions, i.e. $\widehat{V^{(i)}} = V'^{(i)}$. The pairings $h_{12} = -h_{21}$ from Lemma 2.4.3 induce two isomorphisms $V'^{(0)} = \widehat{V^{(0)}} \xrightarrow{\sim} V^{(0)}$, which are the negatives of each other.

We similarly get decompositions

$$\mathrm{Hk}^\flat_V = \mathrm{Hk}^\flat_{V^{(0)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^\flat_{V^{(1)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^\flat_{V^{(1)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}}, \tag{2.44}$$

$$\mathrm{Hk}^\flat_{V'} = \mathrm{Hk}^\flat_{V'^{(0)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^\flat_{V'^{(1)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Hk}^\flat_{V'^{(1)}} \times_{\mathrm{Hk}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}}. \tag{2.45}$$

Using Lemma 2.5.11 to avoid passing to $\mathrm{Hk}^\sharp_-$-stacks, by Lemma 2.5.9 these are again dual vector bundles on each factor after restriction to $\mathrm{Hk}^{r,\le\mu,\circ}_{\widetilde{H},\mathcal{L}}$. The $(-\epsilon)$-Hermitian structures h_{12} and h_{21} again induce two isomorphisms $\mathrm{Hk}^\flat_{V^{(0)}} \simeq \mathrm{Hk}^\flat_{V'^{(0)}}$ (without restriction), which are again negatives of each other. Finally passing to shtuka stacks we have decompositions

$$\mathrm{Sht}^r_V = \mathrm{Sht}^r_{V^{(0)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Sht}^r_{V^{(1)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Sht}^r_{V^{(1)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}}, \tag{2.46}$$

$$\mathrm{Sht}^r_{V'} = \mathrm{Sht}^r_{V'^{(0)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Sht}^r_{V'^{(1)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \mathrm{Sht}^r_{V'^{(1)}} \times_{\mathrm{Sht}^{r,\le\mu}_{\widetilde{H},\mathcal{L}}} \tag{2.47}$$

which are dual relative $\mathbb{F}_q$-vector spaces after restriction to $\mathrm{Sht}^{r,\le\mu,\circ}_{\widetilde{H},\mathcal{L}}$, and again we have two isomorphisms $\mathrm{Sht}^r_{V^{(0)}} \simeq \mathrm{Sht}^r_{V'^{(0)}}$ which are negatives of each other. Note that, as in Section 2.4, the $(-\epsilon)$-Hermitian structures h_{12} and h_{21} on Q together with that on $\mathcal{F}_{\mathrm{univ}}$ gives rise to a Hermitian structure on $\mathrm{Sht}^r_{V^{(0)}}$, i.e. two quadratic forms $\mathfrak{q}_{12} = -\mathfrak{q}_{21} : \mathrm{Sht}^r_{V^{(0)}} \to \underline{k}$. Recall that we have a fixed additive character $\psi : k \to \overline{\mathbb{Q}}^\times_\ell$; then we can pull it back to a function $\mathfrak{q}_{12}^*\psi$ on $\mathrm{Sht}^r_{V^{(0)}}$, which functions as the Gaussian.

On the other two components $\mathrm{Sht}^r_{V^{(1)}}$ and $\mathrm{Sht}^r_{V^{(2)}}$, we want to define two further functions. There are two natural candidates on any relative $\mathbb{F}_q$-vector space, given by the constant function $\mathbf{1}_{\mathrm{Sht}^r_{V^{(i)}}}$ and the characteristic function of the zero section $\delta_{\mathrm{Sht}^r_{V^{(i)}}}$, and similarly on $\mathrm{Sht}^r_{V'^{(i)}}$. By a mild abuse of notation we use the same notations for the pullbacks to $\mathrm{Sht}^{r,\circ}_V$ and $\mathrm{Sht}^{r,\circ}_{V'}$ respectively. Recall we can associate to the torsion sheaf Q a divisor D_Q on X.

Lemma 2.5.14. *Let $d^{(i)}$ be the rank of $\mathrm{Sht}^{r,\circ}_{V^{(i)}}$ as a relative $\mathbb{F}_q$-vector space over $\mathrm{Sht}^{r,\le\mu,\circ}_{\widetilde{H},\mathcal{L}}$. Then we have*

(a) $\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V^{(i)}}}(\delta_{\mathrm{Sht}^{r,\circ}_{V^{(i)}}}) = (-1)^{d^{(i)}}\mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'^{(i)}}}$;

(b) $\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V^{(i)}}}(\mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V^{(i)}}}) = (-1)^{d^{(i)}} q^{d^{(i)}} \delta_{\mathrm{Sht}^{r,\circ}_{V'^{(i)}}}$;

(c) $\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V^{(0)}}}(\mathfrak{q}_{12}^{*}\psi) = (-1)^{d^{(0)}} q^{d^{(0)}/2} \eta_{F'/F}(D_Q)^{nl} [1/2]^{*} \mathfrak{q}_{12}^{*} [-1]^{*} \psi$.

Proof. As in the proof of the analogous Lemma 10.2.3 of [19] (see also the proof of Lemma 8.2.4 there), we can reduce to the case where $\mathrm{Sht}_{V^{(0)}} \to \mathrm{Sht}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$ is split, in which case we just have the usual Fourier transform over finite fields. Then parts (a) and (b) follow from Example 2.3.7 of *loc. cit.*, while part (c) follows from Proposition 2.4.4 (a). □

Corollary 2.5.15. *Let* $d = d^{(0)} + d^{(1)} + d^{(2)}$ *be the rank of* $\mathrm{Sht}^{r,\circ}_{V}$ *over* $\mathrm{Sht}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$. *Then*

$$\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V}}(\mathfrak{q}_{12}^{*}\psi \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V^{(1)}}} \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V^{(2)}}}) = (-1)^{d} q^{d^{(2)}+\frac{1}{2}d^{(0)}} \eta_{F'/F}(D_Q)^{nl} \cdot ([1/2]^{*}\mathfrak{q}_{12}^{*}[-1]^{*}\psi \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'^{(1)}}} \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V'^{(2)}}})$$

as functions on $\mathrm{Sht}^{r,\circ}_{V'}$.

Proof. This is the product of (a), (b), and (c) from Lemma 2.5.14. □

Our strategy will be to relate the higher theta series to the product of the functions appearing in Corollary 2.5.15 with $\mathrm{Sht}(f)^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U}]$ as in Corollary 2.5.13, so that we can apply these duality results to the theta series itself. First, we need to see how to pass from $\mathcal{E}_i$ to $\widetilde{\mathcal{E}}_i$, in terms of the functions above.

The argument is identical to that appearing in [19, §10.2.2] with the pair $(\mathrm{GL}(m)', \mathrm{U}(n))$ replaced by $(M, \widetilde{H}_{\mathcal{L}})$. In particular, letting $N_{\mathcal{T}_i}$ be the derived vector bundle (again, a posteriori classical) associated to $\underline{R\,\mathrm{Hom}}(\mathcal{F}^{*}_{\mathrm{univ}}, \mathcal{T}_i^{*})$ on $\mathrm{Bun}_{\widetilde{H},\mathcal{L}}$ with Hecke version $\mathrm{Hk}^{\flat}_{N_{\mathcal{T}_i}}$ and $\mathrm{Sht}^{r}_{N_{\mathcal{T}_i}}$ constructed as usual, we have an exact triangle

$$\underline{R\,\mathrm{Hom}}(\mathcal{F}^{*}_{\mathrm{univ}}, \widetilde{\mathcal{E}}_1^{*}) \to \underline{R\,\mathrm{Hom}}(\mathcal{F}^{*}_{\mathrm{univ}}, \mathcal{E}_1^{*}) \to \underline{R\,\mathrm{Hom}}(\mathcal{F}^{*}_{\mathrm{univ}}, \mathcal{T}_1^{*}) \tag{2.48}$$

inducing a derived Cartesian square

$$\begin{array}{ccc} \mathcal{Z}^r_{\widetilde{\mathcal{E}}_i,\mathcal{L}} & \longrightarrow & \mathcal{Z}^r_{\mathcal{E}_i,\mathcal{L}} \\ \downarrow & & \downarrow \\ \mathrm{Sht}^r_{\widetilde{H},\mathcal{L}} & \longrightarrow & \mathrm{Sht}^r_{N_{\mathcal{T}_i}} \end{array} \tag{2.49}$$

where the bottom map is the zero section. Restricting to $\mathrm{Sht}^{r,\leq\mu}_{\widetilde{H},\mathcal{L}}$, $\mathrm{Sht}^r_{N_{\mathcal{T}_1}}$ becomes $\mathrm{Sht}^r_{V^{(1)}}$, while $\mathrm{Sht}^r_{N_{\mathcal{T}_2}}$ becomes $\mathrm{Sht}^r_{V'^{(2)}}$; restricting $\mathcal{Z}^r_{\mathcal{E}_i,\mathcal{L}}$ gives Sht^r_U for $i = 1$ and $\mathrm{Sht}^r_{U^\perp}$ for $i = 2$.

Write $c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_i,\mathcal{L}}$ for the pushforward to $\mathrm{Sht}^{r,\circ}_V$ of the pullback of the virtual fundamental class of $\mathcal{Z}^r_{\widetilde{\mathcal{E}}_i,\mathcal{L}}$ along $\mathrm{Sht}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}} \to \mathrm{Sht}^r_{\widetilde{H},\mathcal{L}}$.

Corollary 2.5.16. *We have*

$$\mathrm{Sht}(f)^\circ_!\left[\mathrm{Sht}^{r,\circ}_U\right] \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V^{(1)}}} = c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}$$

and

$$\mathrm{Sht}(f^\perp)^\circ_!\left[\mathrm{Sht}^{r,\circ}_{U^\perp}\right] \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V'^{(2)}}} = c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_2,\mathcal{L}}.$$

Indeed, it follows from (2.49) that restricting $c^{r,\leq\mu,\circ}_{\mathcal{E}_1,\mathcal{L}} = \mathrm{Sht}(f)^\circ_!\left[\mathrm{Sht}^{r,\circ}_U\right]$ along the zero section of $\mathrm{Sht}^{r,\circ}_{V^{(1)}}$ recovers $c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}$ as a class in $H^{\mathrm{BM}}_{2d(a_r)}(\mathrm{Sht}^{r,\circ}_V)$, and similarly restricting $c^{r,\leq\mu,\circ}_{\mathcal{E}_2,\mathcal{L}} = \mathrm{Sht}(f^\perp)^\circ_!\left[\mathrm{Sht}^{r,\circ}_{U^\perp}\right]$ along the zero section of $\mathrm{Sht}^{r,\circ}_{V'^{(2)}}$ gives $c^{r,\leq\mu,\circ}_{\widetilde{E}_2,\mathcal{L}}$ as a class in $H^{\mathrm{BM}}_{2d(a_r^\perp)}(\mathrm{Sht}^{r,\circ}_{V'})$.

As above, we have a decomposition of the classes $c^{r,\leq\mu,\circ}_{\mathcal{E},\mathcal{L}}$ into $c^{r,\leq\mu,\circ}_{\mathcal{E},\mathcal{L}}(a)$ for $a \in \mathcal{A}_{\mathcal{E},\mathcal{L}}(k)$. Let $p^\circ : \mathrm{Sht}^{r,\circ}_V \to \mathrm{Sht}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$ and $\widehat{p}^\circ : \mathrm{Sht}^{r,\circ}_{V'} \to \mathrm{Sht}^{r,\leq\mu,\circ}_{\widetilde{H},\mathcal{L}}$ be the projections; recall that these are relative $\mathbb{F}_q$-vector spaces, so in particular finite étale, and note that $p^\circ_! c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}} = [\mathcal{Z}^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}]$ and $\widehat{p}^\circ_! c^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_2,\mathcal{L}} = [\mathcal{Z}^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_2,\mathcal{L}}]$. Then Corollary 2.5.16 implies the following result.

Corollary 2.5.17. *Let η be the generic point of X', and $\widetilde{Z}^r_{m,\eta}$ as in* (2.8)*. Then $\widetilde{Z}^r_{m,\eta}(\mathcal{G}, \widetilde{\mathcal{E}}_1)$ is the base change along $\eta^r \to X'^r$ of*

$$\chi(\det \widetilde{\mathcal{E}}_1) q^{\frac{n}{2}(\deg \widetilde{\mathcal{E}}_1 - \deg \mathcal{L} - \deg \omega_X)} p^\circ_!(\mathrm{Sht}(f)^\circ_!\left[\mathrm{Sht}^{r,\circ}_U\right] \cdot (\mathfrak{q}_{12}^* \psi \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V^{(1)}}} \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V^{(2)}}}))$$

and $\widetilde{Z}^r_{m,\eta}(\mathcal{G},\widetilde{\mathcal{E}}_2)$ *is the base change of*

$$\chi(\det\widetilde{\mathcal{E}}_2)q^{\frac{n}{2}(\deg\widetilde{\mathcal{E}}_2-\deg\mathcal{L}-\deg\omega_X)}\widehat{p}^{\circ}_!(\mathrm{Sht}(f^{\perp})^{\circ}_![\mathrm{Sht}^{r,\circ}_{U^{\perp}}]\cdot(\mathfrak{q}_{21}^*\psi\cdot\mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'(1)}}\cdot\delta_{\mathrm{Sht}^{r,\circ}_{V'(2)}})).$$

Proof. By Corollary 2.5.16, we have

$$p^{\circ}_!(\mathrm{Sht}(f)^{\circ}_![\mathrm{Sht}^{r,\circ}_U]\cdot\delta_{\mathrm{Sht}^{r,\circ}_{V(1)}})=\sum_{a\in\mathcal{A}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}(k)}[\mathcal{Z}^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}(a)].$$

Composing the map to $\mathcal{A}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}(k)$ with the pairing with $e_{\mathcal{G},\widetilde{\mathcal{E}}_1}$ gives a function on Sht^r_U which by (2.19) agrees with $\mathfrak{q}_{12}\circ\mathrm{Sht}(f)$ (as usual reducing to the split case). Therefore $\mathfrak{q}_{12}^*\psi$ agrees with $\psi(\langle e_{\mathcal{G},\widetilde{\mathcal{E}}_1},a\rangle)$ and so

$$\begin{aligned}p^{\circ}_!(\mathrm{Sht}(f)^{\circ}_![\mathrm{Sht}^{r,\circ}_U]\cdot\mathfrak{q}_{12}^*\psi\cdot\delta_{\mathrm{Sht}^{r,\circ}_{V(1)}})&=p^{\circ}_!(\mathrm{Sht}(f)^{\circ}_![\mathrm{Sht}^{r,\circ}_U]\cdot\mathfrak{q}_{12}^*\psi\cdot\delta_{\mathrm{Sht}^{r,\circ}_{V(1)}}\cdot\mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V(2)}})\\&=\sum_{a\in\mathcal{A}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}(k)}\psi(\langle e_{\mathcal{G},\widetilde{\mathcal{E}}_1},a\rangle)[\mathcal{Z}^{r,\leq\mu,\circ}_{\widetilde{\mathcal{E}}_1,\mathcal{L}}(a)].\end{aligned}$$

After base change to η^r and inserting the leading factors, this is precisely $\widetilde{Z}^r_{m,\eta}(\mathcal{G},\widetilde{\mathcal{E}}_1)$. The proof of the second claim is similar, using (2.20) in place of (2.19). □

We are now ready to conclude. We recall that by Proposition 2.4.2, to prove Theorem 2.3.2 it suffices to show that the values of the theta series at any two transverse Lagrangians agree.

Theorem 2.5.18. *For any two transverse Lagrangian sub-D-modules* $\widetilde{\mathcal{E}}_1,\widetilde{\mathcal{E}}_2\subset\mathcal{G}$*, we have*

$$\widetilde{Z}^r_{m,\eta}(\mathcal{G},\widetilde{\mathcal{E}}_1)=\widetilde{Z}^r_{m,\eta}(\mathcal{G},\widetilde{\mathcal{E}}_2).$$

Proof. By Corollary 2.5.17, it suffices to prove that

$$\chi(\det\widetilde{\mathcal{E}}_1)q^{\frac{n}{2}(\deg\widetilde{\mathcal{E}}_1-\deg\mathcal{L}-\deg\omega_X)}p^{\circ}_!(\mathrm{Sht}(f)^{\circ}_![\mathrm{Sht}^{r,\circ}_U]\cdot(\mathfrak{q}_{12}^*\psi\cdot\delta_{\mathrm{Sht}^{r,\circ}_{V(1)}}\cdot\mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V(2)}})) \tag{2.50}$$

and

$$\chi(\det \widetilde{\mathcal{E}}_2) q^{\frac{n}{2}(\deg \widetilde{\mathcal{E}}_2 - \deg \mathcal{L} - \deg \omega_X)} \widehat{p}^{\circ}_{!}(\mathrm{Sht}(f^{\perp})^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U^{\perp}}] \cdot (\mathfrak{q}_{21}^{*}\psi \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'(1)}} \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V'(2)}})) \tag{2.51}$$

agree.

By the Plancherel formula and near-involutivity of the arithmetic Fourier transform [19, Lemmas 8.2.1 and 8.2.4], we have

$$\begin{aligned} &p^{\circ}_{!}(\mathrm{Sht}(f)^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U}] \cdot (\mathfrak{q}_{12}^{*}\psi \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V(1)}} \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V(2)}})) \\ &\qquad = q^{-d} \widehat{p}_{!}(\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V}}(\mathrm{Sht}(f)^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U}]) \cdot [-1]^{*} \mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}^{r,\circ}_{V}}(\mathfrak{q}_{12}^{*}\psi \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V(1)}} \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V(2)}})), \end{aligned}$$

which by Corollary 2.5.13 and Corollary 2.5.15 (noting that $\operatorname{rank} V = d$) is

$$\eta_{F'/F}(D_Q)^{nl} q^{-d+\operatorname{rank} U + d^{(2)} + \frac{1}{2}d^{(0)}} \cdot \widehat{p}_{!}(\mathrm{Sht}(f^{\perp})^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U^{\perp}}] \cdot ([1/2]^{*}\mathfrak{q}_{12}^{*}[-1]^{*}\psi \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'(1)}} \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V'(2)}})).$$

As $\mathrm{Sht}(f^{\perp})^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U^{\perp}}]$ is invariant under the scaling action of $\mathbb{F}_q^{\times}$, we can eliminate the $[1/2]^{*}$ on the right. Since $\mathfrak{q}_{12} = -\mathfrak{q}_{21}$, $\mathfrak{q}_{12}^{*}[-1]^{*} = \mathfrak{q}_{21}^{*}$ and so in all we have shown that (2.50) is

$$\begin{aligned} &\chi(\det \widetilde{\mathcal{E}}_1) \eta_{F'/F}(D_Q)^{nl} q^{\frac{n}{2}(\deg \widetilde{\mathcal{E}}_1 - \deg \mathcal{L} - \deg \omega_X) - d + \operatorname{rank} U + d^{(2)} + \frac{1}{2}d^{(0)}} \\ &\qquad \cdot \widehat{p}_{!}(\mathrm{Sht}(f^{\perp})^{\circ}_{!}[\mathrm{Sht}^{r,\circ}_{U^{\perp}}] \cdot (\mathfrak{q}_{12}^{*}\psi \cdot \mathbf{1}_{\mathrm{Sht}^{r,\circ}_{V'(1)}} \cdot \delta_{\mathrm{Sht}^{r,\circ}_{V'(2)}})) \end{aligned}$$

This clearly matches (2.51) up to signs and powers of q; so what remains is to show that

$$\frac{n}{2}(\deg \widetilde{\mathcal{E}}_1 - \deg \mathcal{L} - \deg \omega_X) - d + \operatorname{rank} U + d^{(2)} + \frac{1}{2}d^{(0)} = \frac{n}{2}(\deg \widetilde{\mathcal{E}}_2 - \deg \mathcal{L} - \deg \omega_X),$$

or equivalently

$$\frac{n}{2}(\deg \widetilde{\mathcal{E}}_1 - \deg \widetilde{\mathcal{E}}_2) = \frac{1}{2}d^{(0)} + d^{(1)} - \operatorname{rank} U, \tag{2.52}$$

and

$$\chi(\det \widetilde{\mathcal{E}}_1)\eta_{F'/F}(D_Q)^{nl} = \chi(\det \widetilde{\mathcal{E}}_2). \tag{2.53}$$

The latter is Proposition 2.4.4 (d). For the former, note that the exact triangle (2.48) inducing (2.49) gives an equality

$$\begin{aligned}\operatorname{rank} U &= \operatorname{rank} \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{E}^*_1)\\ &= \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{\mathcal{E}}^*_1) + \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{T}^*_1)\\ &= \operatorname{rank} \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{\mathcal{E}}^*_1) + d^{(1)},\end{aligned}$$

while the exact triangle

$$\underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{\mathcal{E}}_2) \to \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{\mathcal{E}}^*_1) \to \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, Q_1)$$

gives

$$\begin{aligned}\operatorname{rank} \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{\mathcal{E}}^*_1) &= \operatorname{rank} \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, \widetilde{\mathcal{E}}_2) + \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}} \otimes \nu^*\mathcal{L}, Q_1)\\ &= \chi(X', \mathcal{F}_{\mathrm{univ}} \otimes \widetilde{\mathcal{E}}_2 \otimes \nu^*\mathcal{L}^*) + d^{(0)}.\end{aligned}$$

Finally, we computed in Proposition 2.4.4 (c) (in slightly different language)

$$\chi(X', \mathcal{F}_{\mathrm{univ}} \otimes \widetilde{\mathcal{E}}_2 \otimes \nu^*\mathcal{L}^*) = \frac{n}{2}(\deg \widetilde{\mathcal{E}}_2 - \widetilde{\mathcal{E}}_1) - \frac{1}{2}d^{(0)}.$$

Putting everything together we find

$$\begin{aligned}\frac{1}{2}n(\deg \widetilde{\mathcal{E}}_1 - \deg \widetilde{\mathcal{E}}_2) &= -\chi(X', \mathcal{F}_{\mathrm{univ}} \otimes \widetilde{\mathcal{E}}_2 \otimes \nu^*\mathcal{L}^*) - \frac{1}{2}d^{(0)}\\ &= -\operatorname{rank} \underline{R\operatorname{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{\mathcal{E}}^*_1) + \frac{1}{2}d^{(0)}\\ &= -\operatorname{rank} U + d^{(1)} + \frac{1}{2}d^{(0)},\end{aligned}$$

giving (2.52). □

As an immediate consequence of Theorem 2.5.18 and Proposition 2.4.2, Theorem 2.3.2 holds.

2.6 Addenda

2.6.1 The type II case

Although generally in this chapter we have been concerned with type I dual pairs, which are more mysterious from the perspective of the relative Langlands program due to their non-polarizability, we briefly investigate the analogue of the higher theta series in the type II case.

Let F, D, F', and W be as in Section 2.1, and let W_1 be a right D-module, W_2 a left D-module, and $\phi : W_1 \otimes_D W_2 \hookrightarrow W$ is an F-linear injection with image a Lagrangian subspace of W. Then $(G, H) = (\mathrm{GL}_D(W_1), \mathrm{GL}_D(W_2))$ form a reductive dual pair of type II in $\mathrm{Sp}_F(W, \langle -, - \rangle)$.[8,9] Let $m = \mathrm{rank}_D W_1$ and $n = \mathrm{rank}_D W_2$, and assume $m \leq n$.

Definition 2.6.1. An H-shtuka over an $\mathbb{F}_q$-scheme S is a tuple

$$((x_1, \ldots, x_r), (\mathcal{F}_0, \ldots, \mathcal{F}_r), (f_1, \ldots, f_r), \varphi)$$

where

- $x_1, \ldots, x_r : S \to X'$ are legs;
- $\mathcal{F}_0, \ldots, \mathcal{F}_r$ are $H = \mathrm{GL}_D(W_2)$-torsors on $X'_S = X' \times_{\mathbb{F}_q} S$;
- $\varphi : {}^{\tau}\mathcal{F}_r \xrightarrow{\sim} \mathcal{F}_0$ is a D-linear isomorphism;
- there are modifications

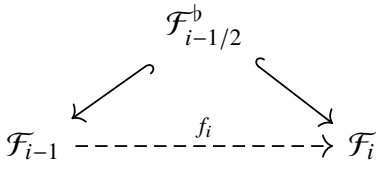

[8] See e.g. [38, §4.1]; all dual pairs of type II are of this form.

[9] We neglect the similitude case in this section for simplicity, or equivalently restrict to the case where the similitude line bundle is trivial.

such that the cokernel of the left inclusion is isomorphic to the pullback of a $D^\times$-torsor on X' along x_i and the cokernel of the right inclusion is isomorphic to the pullback of a $D^\times$-torsor on X' along $\sigma(x_i)$;

- $f_i : \mathcal{F}_{i-1} \to \mathcal{F}_i$ are isomorphisms on $X'_S \backslash \Gamma_{x_i}$ bounded by μ_i, which we take to be minuscule to recover the modifications as described above.

We define Sht^r_H to be the stack over $\mathbb{F}_q$ sending S to the groupoid of shtukas over S in the sense of Definition 2.6.1 (compare [17, Definition 2.4], which is the special case where $D = F'$ is an étale double cover of F).

For each $\mathcal{E} \in \mathrm{Bun}_G(k)$, we define a special cycle $\mathcal{Z}^r_{\mathcal{E}} \to \mathrm{Sht}^r_H$ with S-points lying over a fixed S-point of Sht^r_H given by commutative diagrams

$$\begin{array}{ccccccccc} \mathcal{E} \boxtimes \mathcal{O}_S & = & \mathcal{E} \boxtimes \mathcal{O}_S & = & \cdots & = & \mathcal{E} \boxtimes \mathcal{O}_S & \xrightarrow{\sim} & {}^\tau\mathcal{E} \boxtimes \mathcal{O}_S \\ \downarrow{\scriptstyle t_0} & & \downarrow{\scriptstyle t_1} & & & & \downarrow{\scriptstyle t_r} & & \downarrow{\scriptstyle \tau t_0} \\ \mathcal{F}_0 & \dashrightarrow & \mathcal{F}_1 & \dashrightarrow & \cdots & \dashrightarrow & \mathcal{F}_r & \xrightarrow{\sim} & {}^\tau\mathcal{F}_0 \end{array} , \tag{2.54}$$

where the vertical arrows are D-linear maps.

In Section 2.3, we needed to modify the cycle $\mathcal{Z}^r_{\mathcal{E}}$ by convolving with a Gaussian: the cycle itself depends only on the datum of the "auxiliary" M-bundle $\mathcal{E}$ rather than the "intrinsic" datum $\mathcal{G}$. More abstractly, this boils down to the fact that we cannot find a Lagrangian subspace W' of W fixed by the actions of both groups but only one fixed by a subgroup $M \times \widetilde{H}$, so the W'-period needs to be promoted to a period for the full action via the Weil representation. In our situation no such difficulty arises: $\mathcal{Z}^r_{\mathcal{E}}$ already depends only on a G-bundle $\mathcal{E}$, and for $r = 0$ it recovers the period function of $G \times H$ acting on the symplectic representation W with respect to the Lagrangian $W' := \phi(W_1 \otimes_D W_2) \subset W$.

From the perspective of relative Langlands, there is then one remaining ambiguity: the choice of the Lagrangian W'. We would like to think of the period cycle $\mathcal{Z}^r_{\mathcal{E}}$ as associated just to the datum of $(G \times H, W)$, but instead the period appears to depend on the $G \times H$-representation W' inside $W \simeq W' \oplus W'^*$. Thus we are led to predict that replacing the Lagrangian W' by the alternative

$G \times H$-stable Lagrangian W'^* should give the same period cycles, after suitable normalization.

Generalizing [16, (10.22)], if $\omega_X^{1/2}$ is a square root of the canonical bundle then, as $[\mathcal{Z}^0_{\mathcal{E}}]$ is the function on $\mathrm{Bun}_H(k)$ sending $\mathcal{F}$ to $q^{\dim \mathrm{Hom}_D(\mathcal{E},\mathcal{F})}$ we expect that the normalized period function should be

$$\begin{aligned}(\mathcal{E},\mathcal{F}) &\mapsto q^{\dim H^0(\mathcal{H}om_D(\mathcal{E},\mathcal{F})\otimes\omega_X^{1/2})-\frac{1}{2}\deg(\mathcal{H}om_D(\mathcal{E},\mathcal{F})\otimes\omega_X^{1/2})}\\ &= [\mathcal{Z}^0_{\mathcal{E}}]_{\mathcal{F}\otimes\omega_X^{1/2}}\cdot q^{\frac{n}{2}\deg\mathcal{E}-\frac{m}{2}\deg(\mathcal{F}\otimes\omega_X^{1/2})}.\end{aligned}$$ 10

The action of $\mathrm{Pic}_X(k)$ on Sht^r_H gives a map $-\otimes\omega_X^{1/2}:\mathrm{Sht}^r_H\to\mathrm{Sht}^r_H$, and we let $\mathcal{Z}^{r,\mathrm{norm}}_{\mathcal{E}}$ be the pullback along this map of $q^{\frac{n}{2}\deg\mathcal{E}-\frac{m}{2}\deg(-)}\cdot\mathcal{Z}^r_{\mathcal{E}}$ where $\deg(-)$ denotes the degree function on Sht^r_H, given by composing the projection to Bun_H with the reduced norm to $\mathrm{Bun}_{\mathrm{GL}_1}$ and taking the degree there:

$$\deg:\mathrm{Sht}^r_H\to\mathrm{Bun}_H\to\mathrm{Bun}_{\mathrm{GL}_1}\to\mathbb{Z}. \tag{2.55}$$

Writing $W\simeq(W_1\oplus\sigma^*W_1^*)\otimes_D W_2$, the essential ambiguity is up to replacing W_1 by $W_1^\dagger:=\sigma^*W_1^*$ (see however Remark 2.6.3). On G-torsors, this amounts to replacing $\mathcal{E}$ by $\sigma^*\mathcal{E}^*$; on H-torsors, the objects are fixed but the determinant map is replaced by its inverse, so in particular an H-bundle $\mathcal{F}$ in this context is taken to have degree $-\deg\mathcal{F}$ in sense of (2.55).[11] Therefore the desired formula becomes:

Theorem 2.6.2. *The normalized period cycle satisfies*

$$[\mathcal{Z}^{r,\mathrm{norm}}_{\mathcal{E}}]=[\mathcal{Z}^{r,\mathrm{norm}}_{\sigma^*\mathcal{E}^*}]$$

in $\varprojlim_\mu H^{\mathrm{BM}}_{2(([F':F](n-1)-m)l+1)r}(\mathrm{Sht}^{r,\leq\mu}_H)$. *Equivalently,*

$$q^{\frac{n}{2}\deg\mathcal{E}-\frac{m}{2}\deg(-)}[\mathcal{Z}^r_{\mathcal{E}}]=q^{\frac{n}{2}\deg\mathcal{E}^*+\frac{m}{2}\deg(-)}[\mathcal{Z}^r_{\sigma^*\mathcal{E}^*}]. \tag{2.56}$$

[11] This is a somewhat odd-seeming convention, but seems to be necessary for the desired symmetry; it may be related to the choice of the spin structure being interchanged under this duality.

Here again the classes $[\mathcal{Z}^r_{\mathcal{E}}]$ and $[\mathcal{Z}^r_{\sigma^*\mathcal{E}^*}]$ and their normalized versions should be interpreted as virtual fundamental classes as in Section 2.2.2. Indeed, Section 2.2.2 applies directly to this situation as well, replacing the Levi subgroup M there by our group G and taking trivial similitude bundle, so we can define $[\mathcal{Z}^r_{\mathcal{E}}]$ to be the pullback of the virtual fundamental class of $\mathrm{Sht}^r_{\mathscr{M}}$ to the fiber over $\{\mathcal{E}\} \times \mathrm{Sht}^r_H$. These cycles are therefore again in codimension mlr; by combining the proofs of [17, Lemma 2.6] and Lemma 2.5.3, $\mathrm{Sht}^r_H \to X''^r \times \mathrm{Bun}_H$ is smooth of relative dimension $[F' : F](n-1)lr$, so that the virtual fundamental classes live in degree $2(([F' : F](n-1)-m)l+1)r$.

Proof. In the case $r = 0$, $[\mathcal{Z}^0_{\mathcal{E}}]$ is the function on $\mathrm{Bun}_H(k)$ sending $\mathcal{F}$ to $\#\mathrm{Hom}_D(\mathcal{E}, \mathcal{F})$, which as a function on $H(\mathbb{A})$ is given by

$$h \mapsto \sum_{w' \in W'(F)} \Phi(g^{-1}w'h)$$

where $g \in G(\mathbb{A})$ is in the class of $\mathcal{E} \in \mathrm{Bun}_G(k) \simeq G(F)\backslash G(\mathbb{A})/G(\widehat{O})$ and Φ is the indicator function of the standard lattice $W'(\widehat{O})$ in $W'(\mathbb{A})$. In particular, note that this agrees with the period function as defined in [16, §10.3]. The Fourier transform on W_1 with respect to the self-dual measure and a fixed additive character $\psi : \mathbb{F}_q \to \overline{\mathbb{Q}}_\ell^\times$ gives an isomorphism between Schwartz functions on $W_1 \otimes W_2$ and on $W_1^\dagger \otimes W_2$, and since Φ is the indicator function of a compact self-dual subgroup it is fixed by the Fourier transform in that $\mathrm{FT}(\Phi) = \Phi^\dagger$, the indicator function of $(W_1^\dagger \otimes W_2)(\widehat{O})$; therefore by Poisson summation

$$[\mathcal{Z}^0_{\mathcal{E}}]_{\mathcal{F}} = \sum_{w' \in W'} \Phi(g^{-1}w'h) = \mathrm{Vol}(g^{-1}W') \sum_{w'^\dagger \in W_1^\dagger \otimes_D W_2} \Phi^\dagger(gw'^\dagger h) = q^{-n \deg \mathcal{E}}[\mathcal{Z}^0_{\sigma^*\mathcal{E}^*}]_{\mathcal{F}},$$

agreeing with the prediction (2.56) (this is the argument of [16, Example 10.9.1]). We briefly explain how to promote this argument to the $r > 0$ case, similar to Section 2.5.

In general, we write $\widetilde{\mathcal{E}}$ in place of $\mathcal{E}$, and let $\mathcal{E}_1 \subset \widetilde{\mathcal{E}}$ and $\mathcal{E}_2 \subset \sigma^*\widetilde{\mathcal{E}}^*$ be subbundle with torsion cokernels $\mathcal{T}_1$ and $\mathcal{T}_2$ respectively with disjoint support, such that $\mathrm{Ext}^1_{X'_{\overline{k}}}(\mathcal{F}^*, \mathcal{E}_i^*) = 0$ for every H-bundle $\mathcal{F} \in \mathrm{Bun}_H^{\leq \mu}(\overline{k})$ and $i \in \{1, 2\}$. Let $U \to \mathrm{Bun}_H^{\leq \mu}$ for a Harder–Narasimhan polygon μ be the (a priori) derived vector bundle corresponding to the locally free sheaf $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{E}_1^*)$, and

$W \to \mathrm{Bun}_H^{\leq\mu}$ corresponding to $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \sigma^*\mathcal{E}_2[1])$; exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$ we get $U^\perp$ and $W^\perp$. In this situation, the composite

$$\sigma^*\mathcal{E}_2 \to \widetilde{\mathcal{E}}^* \to \mathcal{E}_1^*$$

is injective with cokernel $\widetilde{Q}_1 := \mathcal{T}_1^* \oplus \sigma^*\mathcal{T}_2$, and likewise

$$\sigma^*\mathcal{E}_1 \to \sigma^*\widetilde{\mathcal{E}} \to \mathcal{E}_2^*$$

is injective with cokernel $\widetilde{Q}_2 := \sigma^*\mathcal{T}_1 \oplus \mathcal{T}_2$. Then we let V correspond to the locally free sheaf $\underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \widetilde{Q}_1) \simeq \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{T}_1^*) \oplus \underline{R\,\mathrm{Hom}}(\mathcal{F}^*_{\mathrm{univ}}, \sigma^*\mathcal{T}_2)$, inducing a decomposition $V \simeq V^{(1)} \times_{\mathrm{Bun}_H^{\leq\mu}} V^{(2)}$. We similarly define $V' \simeq V'^{(1)} \times_{\mathrm{Bun}_H^{\leq\mu}} V'^{(2)}$. As in Lemma 2.5.8, we have $\widehat{U} \simeq W^\perp$, $\widehat{V} \simeq V'$ (preserving the decompositions), and $\widehat{W} \simeq U^\perp$.

We can form Hecke and shtuka stacks as in Section 2.5, and find that $\mathrm{Hk}_U^\flat \to \mathrm{Hk}_H^{r,\leq\mu}$ has dual derived vector space given by $\mathrm{Hk}^\sharp_{W^\perp}$, $\mathrm{Hk}_V^\flat \to \mathrm{Hk}_H^{r,\leq\mu}$ has dual $\mathrm{Hk}^\sharp_{V'}$, and $\mathrm{Hk}_W^\flat \to \mathrm{Hk}_H^{r,\leq\mu}$ has dual $\mathrm{Hk}^\sharp_{U^\perp}$. The correspondence $U \xleftarrow{a_0} \mathrm{Hk}_U^\flat \xrightarrow{a_r} U$ given by 0th and rth projections supports a cohomological correspondence given by the relative fundamental class of a_r, which we denote by $\mathfrak{c}_U$. Its pushforward along $f : \mathrm{Hk}_U^\flat \to \mathrm{Hk}_V^\flat$ satisfies

$$\mathbb{T}_{[2\operatorname{rank} U - \operatorname{rank} V](\operatorname{rank} U)}\,\mathrm{FT}_{\mathrm{Hk}_V^\flat}(f_!\mathfrak{c}_U) = f_!^\perp \mathfrak{c}_{U^\perp},$$

with the analogous notations on the $-^\perp$ side, proven just as for Theorem 2.5.10. Taking shtuka traces analogously gives

$$\mathrm{FT}^{\mathrm{arith}}_{\mathrm{Sht}_V^{r,\circ}}(\mathrm{Sht}(f)^\circ_!\,[\mathrm{Sht}_U^{r,\circ}] = (-1)^{\operatorname{rank} V} q^{\operatorname{rank} U} \cdot \mathrm{Sht}(f^\perp)^\circ_!\,[\mathrm{Sht}_{U^\perp}^{r,\circ}].$$

We can then conclude as in the proof of Theorem 2.5.18 with $V^{(0)}$ trivial that

$$[\mathcal{Z}^r_{\mathcal{E}}] q^{-d^{(1)} + \operatorname{rank} U} [\mathcal{Z}^r_{\sigma^*\mathcal{E}^*}]. \tag{2.57}$$

Here we have used that Sht_U^r is the derived version of $\mathcal{Z}_{\mathcal{E}}^r$ and the fiber of Sht_U^r over $\mathrm{Sht}_{V^{(2)}}^r \to \mathrm{Sht}_V$, and similarly on the dual side.

Since $-d^{(1)} + \operatorname{rank} U = \chi(\mathcal{F}_{\mathrm{univ}} \otimes^{\mathbb{L}} \widetilde{\mathcal{E}}^*) = -n \deg \widetilde{\mathcal{E}} + m \deg(-) - \frac{1}{2}mn \deg \omega_X$, replacing each $\mathcal{F}$ by $\mathcal{F} \otimes \omega_X^{1/2}$ gives $-n \deg \widetilde{\mathcal{E}} + m \deg(-)$. Therefore, taking the limit over μ, (2.57) rearranges to (2.56). □

Remark 2.6.3. Theorem 2.6.2 is analogous to the special case of Theorem 2.3.2 in the type I case where, in the language of Section 2.5, $\widetilde{\mathcal{E}}_2 = \sigma^* \widetilde{\mathcal{E}}_1^*$. It would be preferable to have a more general modularity statement in the type II case, at least conjecturally, analogous to Conjecture 2.3.1. Outside of the split case, however, it is not completely clear what the statement should be: one needs to make the action of $\mathrm{Sp}(W)$ on the Lagrangian bundle $\mathcal{E}$ explicit, and understand the resulting transformation of the leading terms. We expect that this should be possible, and once formulated it should be possible to prove a version of this conjecture using a similar adaptation of the methods of Section 2.5 after similar restrictions to those described in Section 2.3, in particular on the generic fiber.

2.6.2 Other symplectic representations

Now we relax the conditions on the pair (G, H): suppose $\rho : G \times H \to \mathrm{Sp}_{F_0}(W)$ is a symplectic representation representation,[12] equipped with a fixed lift to the metaplectic group of W. In principle, this should be enough to define a period function for the pair $(G \times H, W)$, without any polarizability conditions. However, thus far we only know how to build periods for representations in the case of the tensor product of the standard representations of reductive dual pairs. Thus we assume that ρ factors as a tensor product $\rho_1 \otimes_D \rho_2$ over some F_0-division algebra D where $\rho_1 : G \to G' \to \mathrm{GL}_D(W_1)$, $\rho_2 : H \to H' \to \mathrm{GL}_D(W_2)$, $W_1 \otimes_D W_2 \simeq W$, and (G', H') is a reductive dual pair. We say that ρ is of type I (resp. II) if (G', H') is.

In the type I case, let W_1' be a subspace of W_1 such that $W_1' \otimes_D W_2$ is a Lagrangian subspace of

[12] We again neglect similitude factors in this section, but it should be straightforward to include them as in Section 2.1.

W, and let $P = MN$ be the parabolic in G fixing W_1' via ρ_1 and similarly $Q = LR$ be the parabolic in G' fixing W'; we write $\rho_1' : M \to L$ for the restriction of ρ_1 to M on W_1'. This gives maps

$$\mathrm{Bun}_{\rho_1'} : \mathrm{Bun}_M \to \mathrm{Bun}_L, \qquad \mathrm{Bun}_{\rho_1|_P} : \mathrm{Bun}_P \to \mathrm{Bun}_Q,$$
$$\mathrm{Bun}_{\rho_1} : \mathrm{Bun}_G \to \mathrm{Bun}_{G'}, \qquad \mathrm{Sht}^r_{\rho_2} : \mathrm{Sht}^r_H \to \mathrm{Sht}^r_{H'},$$

where $\mathrm{Sht}^r_{H'}$ is defined as in Definition 2.2.1 (with trivial similitude bundle) and Sht^r_H as in e.g. [39]. This lets us define $\widetilde{Z}^r(\mathcal{G}, \mathcal{E}) \in \mathrm{Ch}^{mlr}(\mathrm{Sht}^r_H)_{\overline{\mathbb{Q}}_\ell}$ for each $(\mathcal{G}, \mathcal{E}) \in \mathrm{Bun}_P(k)$ as the pullback of $\widetilde{Z}^r_m(\mathrm{Bun}_{\rho_1|_P}(\mathcal{G}, \mathcal{E})) \in \mathrm{Ch}^{mlr}(\mathrm{Sht}^r_{H'})$ where $m = \mathrm{rank}_D W_1'$.[13] Thus modularity for $(G \times H, W)$, i.e. independence of $\mathcal{E}$, reduces to modularity for dual reductive pairs of type I:

Proposition 2.6.4. *With notation as above, Conjecture 2.3.1 implies that $\widetilde{Z}^r(\mathcal{G}, \mathcal{E})$ is independent of $\mathcal{E}$.*

Proof. Indeed, assuming Conjecture 2.3.1 we have

$$\widetilde{Z}^r(\mathcal{G}, \mathcal{E}_1) = \mathrm{Sht}^{r*}_{\rho_2} \widetilde{Z}^r_m(\mathcal{G}, \mathcal{E}_1) = \mathrm{Sht}^{r*}_{\rho_2} \widetilde{Z}^r_m(\mathcal{G}, \mathcal{E}_2) = \widetilde{Z}^r(\mathcal{G}, \mathcal{E}_1).$$

□

Defining $\widetilde{Z}^r_\eta$ as in Section 2.3 using $\widetilde{Z}^r$ in place of $\widetilde{Z}^r_m$, so passing to ℓ-adic cohomology, the limit over Harder–Narasimhan strata, and the generic fiber, we obtain a generic modularity statement from Theorem 2.3.2:

Corollary 2.6.5. *(a) If (G, H) is a type I reductive dual pair such that the symplectic representation W admits an H-stable Lagrangian subspace and is not of excluded type, then the higher theta series Z^r_η descends to $\mathrm{Bun}_G(k)$, i.e. $\widetilde{Z}^r_\eta(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$.*

(b) If (G, H) is of excluded type, then $\widetilde{Z}^r_\eta(\mathcal{G}, \mathcal{E})$ depends only on $\mathcal{G}$ and on the class of $\mathcal{E}$ in $\pi_0 \mathrm{LGr}(m, 2m)$.

[13] It should be possible to give an equivalent definition in terms of special cycles (defined via morphisms on the associated bundles through ρ) similar to (2.4), which we leave to the reader.

In particular, the modularity of the period function is given by the $r = 0$ case of Theorem 2.3.2.

Similarly, in the type II case ρ induces a map

$$\mathrm{Bun}_G \times \mathrm{Sht}^r_H \to \mathrm{Bun}_{G'} \times \mathrm{Sht}^r_{H'}$$

and we define the normalized period to be the pullback of that for (G', H'), so the independence of the choice of Lagrangian subspace W' reduces to that for (G', H'), so we can interchange $\mathcal{E}$ and $\sigma^*\mathcal{E}^*$ as in Theorem 2.6.2.

Remark 2.6.6. In the special case where (G, H) is itself a reductive dual pair (of either type), in which case we could directly define the theta series via the Weil representation on Schwartz functions on the standard representation of $M \times H$, it may not agree with the one defined here by pullback. For example, if the standard representation of G has rank $m' \neq m$, then these classes in $\mathrm{Ch}^*(\mathrm{Sht}^r_H)$ will not even have the same codimension. However this is a feature rather than a bug: viewed as a period, the theta series should depend not just on a group $G \times H$ but an action of this product on a hyperspherical vector space, i.e. a symplectic representation,[14] so that the cycle $\widetilde{Z}^r : \mathrm{Bun}_G(k) \to \mathrm{Ch}^{mlr}(\mathrm{Sht}^r_H)_{\overline{\mathbb{Q}}_\ell}$ is the period cycle associated to the hyperspherical pair $(G \times H, W)$ for the action of $G \times H$ on W via ρ through $G' \times H'$.

[14]Further we should restrict to "anomaly free" hyperspherical varieties, i.e. in this case a symplectic representation lifting to the metaplectic cover; see [16, §5.1] for details. In particular, polarized representations are automatically anomaly free.

Chapter 3: Split generic p-adic modularity

3.1 Sheaf theory

In what follows we will often refer to a (derived) category of ℓ-adic sheaves $D(S, \overline{\mathbb{Q}}_\ell)$ on a v-stack S, equipped with a six-functor formalism. Much of what we do is independent of the particular choice, other than that we require certain maps to be cohomologically smooth, etc.; at certain points when we refer [24] we need to assume that the six-functor formalism also satisfies excision.[1] For concreteness, we use Scholze's recent theory of Berkovich motives [26] with ℓ-adic coefficients, using the functor a'^* from §12 of *loc. cit.* to work in terms of v-stacks as in [41] (i.e. "over-convergent motivic sheaves"), so $D(S, \overline{\mathbb{Q}}_\ell) := \mathcal{D}^{\mathrm{oc}}_{\mathrm{mot}}(S, \overline{\mathbb{Q}}_\ell)$. This has the advantage that as a motivic cohomology theory the Borel–Moore homology groups $H^{\mathrm{BM}}_{2*}(X) = \mathrm{Hom}(\mathbb{Q}_\ell(*)[2*], \pi_*\pi^!\mathbb{Q}_\ell)$ for $\pi : X \to S$ deserve to be called Chow groups, as in [20, §3.5]. However we will rarely need to unwind any of the definitions. The key properties of this formalism that we need are that it gives a six-functor formalism with $\mathbb{Q}_\ell$-coefficients on small v-stacks; it satisfies excision; and it is compatible with the six-functor formalism for $\mathbb{Q}_\ell$-sheaves on condensed anima of [42], as discussed in the proof of [41, Proposition 2.1].

We also at times refer to the more recent work [25], which builds upon [24], in particular for the functor σ in Section 3.3.2. This is primarily for the integral analogue of the Banach–Colmez functor $\mathcal{E} \mapsto \mathcal{BC}(\mathcal{E})$, which in their language can be understood as the rationalization $\sigma_{\mathbb{Q}_p}$. The integral setting in turn is needed to pass to stacks of shtukas at finite level.[2]

We sometimes write $\mathbf{1}$ for the unit object $\mathbb{Q}_\ell$, and $-\langle d\rangle$ or $-_{\langle d\rangle}$ for $-(d)[2d]$, where (d) is the

[1] In *loc. cit.* the authors work with the formalism of $\mathbb{Z}/n\mathbb{Z}$-sheaves for $p \nmid n$, as developed in [40], but the results go through for any formalism satisfying excision and in which ℓ-cohomologically smooth maps are cohomologically smooth, proper maps are cohomologically proper, and $f_!$ is defined whenever f is smooth-locally nice.

[2] It might be possible to prove an analogous result at infinite level without the integral theory; but the descent to finite level does not seem directly accessible.

Tate twist. We will say that a morphism $f : X \to Y$ is cohomologically quasi-smooth if there is a natural transformation $f^* \to f^!_{\langle d \rangle}$ for some d. This is the necessary condition to define pullback on Borel–Moore homology. For this it suffices to have a map $\mathbf{1} \simeq f^*\mathbf{1} \to f^!\mathbf{1}_{\langle d \rangle}$, as in [43, Remark 3.7]. We will often refer to the notions of pullable or pushable squares as in [19, Definition 3.1.1]; replacing quasi-smooth morphisms by cohomologically quasi-smooth ones, we call the resulting notion cohomologically pullable squares.[3]

In [20], at various points we need to check that certain motivic sheaves are universally strongly locally acyclic (USLA) or geometric. In particular these conditions are necessary to form the trace of cohomological correspondences: the key fact is that if X_0, X_1 are derived Artin stacks over S and $\mathcal{K}_0 \in D(X_0)$ is USLA over S and $\mathcal{K}_1 \in D(X_1)$ is geometric, then the natural map

$$\mathbb{D}_{X_0/S}(\mathcal{K}_0) \boxtimes_S \mathcal{K}_1 \to \underline{R\,\mathrm{Hom}}_{X_0 \times_S X_1}(\mathrm{pr}_0^* \mathcal{K}_0, \mathrm{pr}_1^! \mathcal{K}_1) \tag{3.1}$$

is an isomorphism (Proposition 3.8.5 in *loc. cit.*). This is used to define the trace, which also implicitly uses that when $\mathcal{K}$ on $X \to S$ is geometric and USLA over S, the object $(X, \mathcal{A})$ in the category of correspondences over S is dualizable (Proposition 5.2.3).

In the p-adic context, there is a good analogue of the USLA condition, namely universally locally acyclic (ULA) sheaves as in [44, §IV.2], which satisfy similar conditions to (3.1) as in e.g. Theorem IV.2.23 in *loc. cit.*[4] In the setting of Berkovich motives, we can abstract out this condition: we replace the notion of USLA or ULA sheaves with suave sheaves, again with respect to the morphism $X \to S$, as in e.g. [42, Definition 4.4.1]; in the setup of (3.1), if $\mathcal{K}_0$ is suave over S and $X_1 \to S$ is suave (see Definition 4.5.1) then by Lemmas 4.4.5, 4.4.9, and 4.4.17 in *loc. cit.* (3.1) holds. Note that there is no condition on $\mathcal{K}_1$, though we now incorporate an additional condition on X_1 (which will be satisfied in all relevant situations).

We also need to check dualizability. As $\mathcal{K}$ is dualizable if and only if it is id-suave (Example

[3] The choice to write "cohomologically pullable" but not "cohomologically pushable" is really just illustrative; we only ever use the cohomological behavior, and sometimes only prove that morphisms are cohomologically proper rather than geometrically.

[4] Strictly speaking what is shown there is a special case of (3.1), but it is not difficult to obtain the generality above; note that there is no geometricity condition.

4.4.3 in *loc. cit.*), in practice all our sheaves of interest will arise via pushforward and pullback of the unit object, and so to check dualizability it will suffice to observe that suave objects are preserved under suave pullback and prim pushforward (Lemma 4.5.16 in *loc. cit.*), and so we will need to check only that the relevant maps are suave or prim. We note that cohomologically smooth maps are suave and that the unit is f-suave if f is suave, which will be our main tool for this condition.

Subject to replacing USLA and geometricity conditions with suaveness (and dualizability, as needed) as described above, the proofs of [20, Theorems 5.4.2, 5.4.3] apply without modification, and we cite them without further discussion of the difference in the setups.

3.2 Intersections of correspondences

Let $c = (c_0, c_1) : C \to X \times Y$ be a correspondence of v-stacks over a base v-stack S. In the rest of this subsection, products refer to fiber products over S if not otherwise specified. There are at least two notions of products of correspondences; to disambiguate them we will denote them by $\circ$ and $\cap$, and refer to them as composition and intersection.

The first, as the name suggests, is the composition of correspondences: if $d = (d_0, d_1) : D \to Y \times Z$ is another correspondence, the composite correspondence from X to Z is given by $C \times_Y D$, with map to X given by the projection onto C composed with c_0 and map to Z given by the projection to D composed with d_1. We call this correspondence $D \circ C$.

On the other hand, if D is again a correspondence $d = (d_0, d_1) : D \to X \times Y$, then we can take its intersection with C over $X \times Y$: $C \cap D = C \times_{X \times Y} D$, with the obvious projection to $X \times Y$. For example, let Δ denote the diagonal correspondence $X \to X \times X$ (by a mild abuse of notation); then for any $C \to X \times X$, $C \cap \Delta$ is the space of fixed points $\mathrm{Fix}(C)$.

Being the composition of morphisms in a category, the composition product $\circ$ satisfies a number of good properties. For example, if $X = Z$, A and B are sheaves on X and Y respectively, and $u : c_0^* A \to c_1^! B$ and $v : d_0^* B \to d_1^! A$ are cohomological correspondences supported on C and D respectively, then one can form the composite cohomological correspondence $v \circ u$ from A to

itself supported on $D \circ C$ (see e.g. [45, Definition 4.3.4]), whose trace then gives an element of $H^0(\mathrm{Fix}(D \circ C), K_{\mathrm{Fix}(D\circ C)})$, giving a pairing on cohomological correspondences $(u, v) \mapsto \mathrm{tr}(v \circ u)$.

For the special case $X = Y$ and $D = \Delta$, we can understand cohomology classes on $C \cap \Delta = \mathrm{Fix}(C)$ as traces of cohomological correspondences supported on C. Slightly more generally, when $d = (\mathrm{id}, f) : D = X \to X \times X$ is the graph of an endomorphism $f : X \to X$, we can understand $C \cap \Gamma_f$ as the f-twisted fixed points of C (for example when f is the Frobenius of X we recover $\mathrm{Sht}(C)$), and cohomology classes on it as f-twisted traces of cohomological correspondences supported on C. It is not immediately clear however how to extend this interpretation to arbitrary D.

Fortunately, we can relate the intersection product to the composition product. For a correspondence $c = (c_0, c_1) : C \to X \times Y$, write $C^\vee$ for the same correspondence but with the roles of X and Y switched, i.e. $c^\vee = (c_1, c_0) : C^\vee = C \to Y \times X$. If $D \to X \times Y$ is another correspondence, we obtain a self-correspondence of X as $D^\vee \circ C$. By the definition of the composition product, this has T-points given by pairs $s : T \to C$ and $t : T \to D$ together with identifications $c_1 \circ s = d_1 \circ t$. As $D^\vee \circ C$ is a self-correspondence of X, we can take its stack of fixed points, whose T-points are then pairs (s, t) as above together with an identification $c_1 \circ s = d_1 \circ t$ and another $c_0 \circ s = d_0 \circ t$. But this is precisely the definition of $C \cap D$; in other words,

$$C \cap D = \mathrm{Fix}(D^\vee \circ C).^5 \tag{3.2}$$

In particular, we get a pairing sending cohomological correspondences $u : c_0^* A \to c_1^! B$ and $v^\vee : d_1^* B \to d_0^! A$ to $\mathrm{tr}(v^\vee \circ u) \in H^0(C \cap D, K_{C\cap D})$.

Note that although the right-hand side of (3.2) appears to depend on the order of C and D, the left-hand side is independent of it: $C \cap D \simeq D \cap C$ as a correspondence over $X \times Y$. For example, when $X = Y$ and $D = \Delta$, $D^\vee = \Delta$ and $D^\vee \circ C = \Delta \circ C = C$, so we recover the identity $C \cap \Delta = \mathrm{Fix}(C)$, and similarly $\Delta \cap C = \mathrm{Fix}(C^\vee \circ \Delta)$, so we recover the identity $\mathrm{Fix}(C) \simeq \mathrm{Fix}(C^\vee)$

[5] Note that although this identification is compatible with the maps to $X \times Y$, this may be misleading: when $X = Y$, $C \cap D \to X \times X$ does not necessarily factor through the diagonal, as the map to the second factor does not come formally from the fixed point construction but from the particular construction of $D^\vee \circ C$.

over X. Thus we can also replace both C and D by their duals without changing the fixed points, so e.g. $C \cap D = \mathrm{Fix}(D \circ C^\vee) = \mathrm{Fix}(C \circ D^\vee)$. For example, for $D = \Gamma_f$ as above, $C \circ \Gamma_f^\vee$ is the correspondence $c^{(f)} = (f \circ c_0, c_1) : C^{(f)} = C \to X \times X$, recovering the Frobenius-twisted version $C^{(1)}$ of C when f is the Frobenius.[6] In that case there is a canonical choice of cohomological correspondence $\mathrm{Weil}_A^\vee : \mathrm{Fr}^* A \to \mathrm{id}^! A = A$ supported on $\Gamma_{\mathrm{Fr}}^\vee$, given by the Weil structure on A, and so evaluating our pairing on $\mathrm{Weil}_A^\vee$ gives a map from cohomological correspondences $u : c_0^* A \to c_1^! A$ to $\mathrm{tr}(u \circ \mathrm{Weil}_A^\vee) \in H^0(\mathrm{Fix}(C^{(1)}), K_{\mathrm{Fix}(C^{(1)})}) = H^0(\mathrm{Sht}(C), K_{\mathrm{Sht}(C)})$. This is the shtuka trace of [19, §4.2]. Indeed, more generally following loc. cit. for correspondences $C, D \to X \times Y$ and sheaves A, B supported on X and Y respectively, given cohomological correspondences $u : c_0^* A \to c_1^! B$ and $v^\vee : d_1^* B \to d_0^! A_{\langle -i \rangle}$ we can form $\mathrm{tr}(v^\vee \circ u) \in H_{2i}^{\mathrm{BM}}(D \cap C)$.

For a commutative diagram

$$\begin{array}{ccccccccc} X & \longleftarrow & C & \longrightarrow & Y & \longleftarrow & D & \longrightarrow & X \\ \downarrow & & \downarrow{\scriptstyle f} & & \downarrow & & \downarrow{\scriptstyle g} & & \downarrow \\ X' & \longleftarrow & C' & \longrightarrow & Y' & \longleftarrow & D' & \longrightarrow & X' \end{array},$$

with the maps $X \to X'$ agreeing, we get a morphism of correspondences $h : D^\vee \circ C \to D'^\vee \circ C'$, giving on fixed points a map

$$\mathrm{Fix}(h) : C \cap D \to C' \cap D'.$$

For u and $v^\vee$ as above, what we have called $\mathrm{tr}(v^\vee \circ u)$ is written in [20] as $\langle u, v^\vee \rangle$; they show that under suitable conditions, $\mathrm{Fix}(h)$ is proper and

$$\mathrm{Fix}(h)_* \langle u, v^\vee \rangle = \langle f_! u, g_! v^\vee \rangle.$$

Similarly under suitable pullback conditions $\mathrm{Fix}(h)$ will be smooth and

$$\mathrm{Fix}(h)^! \langle u', v'^\vee \rangle = \langle f^* u', g^* v'^\vee \rangle.$$

[6] We could also choose a different presentation, e.g. $C \cap \Gamma_f = \mathrm{Fix}(\Gamma_f^\vee \circ C)$, but the correspondence $\Gamma_f^\vee \circ C$ will be of a slightly different form although its fixed points will agree.

3.3 Fourier transforms

3.3.1 Fourier transform over stacks in F-vector spaces

We briefly recall some notions from [24]. In particular, we will use the notion of a stack in F-vector spaces, which (over a fixed base S) is equivalent to a complex v-sheaves of F-vector spaces on S with cohomology concentrated in degrees 0 and -1, with e.g. the stack $\underline{F}$ corresponding to $F[0]$ and $*/\underline{F}$ to $F[1]$. For such a stack $V \to S$, we write $\widehat{V} \to S$ for the dual $\widehat{V} = \underline{\mathrm{Hom}}(V, S/\underline{F})$, the internal hom object in the category of stacks in F-vector spaces; if V corresponds to the complex $\mathcal{V}$, then $\widehat{V}$ corresponds to $\widehat{\mathcal{V}} = \tau_{\leq 0}\underline{R\,\mathrm{Hom}}(\mathcal{V}, F[1])$. If the natural map $V \to \widehat{\widehat{V}}$ is an isomorphism, we say that V is dualizable. We tend to identify the sheaf $\mathcal{V}$ and its associated stack V.

For $\mathcal{E}$ a vector bundle on X_S, we write $\mathcal{BC}(\mathcal{E})$ for the sheaf over S sending $T \to S$ to $R\Gamma(X_T, \mathcal{E}|_{X_T})$. This is a complex of F-vector spaces on S concentrated in degrees 0 and -1, and so we can think of it as a stack in F-vector spaces. In fact one can make the same definitions for any perfect complex $\mathcal{E}$ on X_S, so long as the resulting cohomology is concentrated in the same degrees.[7]

Suppose $p : V \to S$ and $\widehat{p} : \widehat{V} \to S$ are representable in Artin v-stacks and smooth-locally nice, so that a six-functor formalism is available for them. Fix a nontrivial character $\psi : F \to \overline{\mathbb{Q}}_\ell^\times$, and write $\mathcal{L}_\psi$ for the corresponding rank 1 local system on $S/\underline{F}$. We have a natural pairing $\mathrm{ev}_V : V \times_S \widehat{V} \to S/\underline{F}$ and smooth-locally nice projections $\pi : V \times_S V \to V$, $\widehat{\pi} : V \times_S \widehat{V} \to \widehat{V}$, which lets us define the Fourier transform

$$\mathrm{FT} = \mathrm{FT}_V^\psi := \widehat{\pi}_!(\pi^*(-) \otimes \mathrm{ev}_V^* \mathcal{L}_\psi) : D(V, \overline{\mathbb{Q}}_\ell) \to D(\widehat{V}, \overline{\mathbb{Q}}_\ell).$$

[7] Our conventions for Banach–Colmez spaces are slightly different than those of [24]: for example if $\mathcal{E}$ has only negative slopes then the $\mathcal{BC}(\mathcal{E})$ of loc. cit. would be $\mathcal{BC}(\mathcal{E}[1])$ in our notation.

When V is dualizable, we obtain similarly

$$\mathrm{FT}^{\psi}_{\widehat{V}} := \pi_!(\widehat{\pi}^*(-) \otimes \mathrm{ev}^*_{\widehat{V}}\psi) : D(\widehat{V}, \overline{\mathbb{Q}}_\ell) \to D(V, \overline{\mathbb{Q}}_\ell),$$

and if further after decomposing S into open and closed substacks FT^{ψ}_{V} and $\mathrm{FT}^{\psi^{-1}}_{\widehat{V}}$ are inverse up to shift and twist we say that V is very nice (following [24]). When $\mathrm{FT}^{\psi^{-1}}_{\widehat{V}} \circ \mathrm{FT}^{\psi}_{V} \simeq (-d)[-2d]$ for $d \in \mathbb{Z}$ a locally constant function on S (depending on V), we say that V is very nice of Fourier dimension d, in which case $\widehat{V}$ is as well. By Propositions 3.34 and 3.35 of *loc. cit.*, replacing the ℓ-cohomological smoothness results for Banach–Colmez spaces from [44, Propositions II.2.5, II.3.5] with [41, Propositions 2.2, 2.3], the Banach–Colmez spaces $\mathcal{BC}(\mathcal{E})$ for any vector bundle $\mathcal{E}$ (or in fact suitable perfect complex) are very nice, and examining their proofs we see that in fact the Fourier dimension of $\mathcal{BC}(\mathcal{E})$ is its cohomological dimension over S, justifying the terminology. The proofs also show that the structure maps of Banach–Colmez spaces are prim, which will be useful later.

Let $f : V \to V'$ be a morphism of stacks in F-vector spaces over S, with V and V' very nice of Fourier dimensions d and d' respectively, and write $\widehat{f} : \widehat{V}' \to \widehat{V}$ for its dual; write $d(f) = d - d'$, and similarly $d(\widehat{f}) = d' - d = -d(f)$. Then by [24, Proposition 3.31] we have

$$\mathrm{FT}^{\psi}_{V'} \circ f_! \simeq \widehat{f}^* \circ \mathrm{FT}^{\psi}_{V} \tag{3.3}$$

(keeping track of the shift, we see that it is actually trivial). By the (near-)involutivity of the Fourier transform, it follows that

$$f_! \circ \mathrm{FT}^{\psi^{-1}}_{\widehat{V}} \simeq \mathrm{FT}^{\psi^{-1}}_{\widehat{V}'} \circ \widehat{f}^*(-d(f))[-2d(f)]$$

and so replacing f by $\widehat{f}$ and ψ by ψ^{-1} we find

$$\mathrm{FT}^{\psi}_{V} \circ f^* \simeq \widehat{f_!} \circ \mathrm{FT}_{V'}(-d(f))[-2d(f)]. \tag{3.4}$$

Suppose

$$\begin{array}{ccc} C^\flat & \xrightarrow{c_0} & V_0 \\ \downarrow{\scriptstyle c_1} & & \downarrow{\scriptstyle c_1'} \\ V_1 & \xrightarrow{c_0'} & C^\sharp \end{array}$$

is a Cartesian and co-Cartesian diagram where each object is a very nice stack in F-vector spaces and each map c_i, c_i' is a morphism in this category. We refer to such a diagram as a correspondence of F-stacks from V_0 to V_1. The duality $V \mapsto \widehat{V}$ applied to this diagram results in a dual correspondence

$$\begin{array}{ccc} \widehat{C^\sharp} & \xrightarrow{\widehat{c_1'}} & \widehat{V}_0 \\ \downarrow{\scriptstyle \widehat{c_0'}} & & \downarrow{\scriptstyle \widehat{c_0}} \\ \widehat{V}_1 & \xrightarrow{\widehat{c_1}} & \widehat{C^\flat} \end{array}.$$

If $\mathcal{K}_i \in D(V_i)$ are sheaves and $u : c_0^*\mathcal{K}_0 \to c_1^!\mathcal{K}_1$, or equivalently $u' : c_{1!}c_0^*\mathcal{K}_0 \to \mathcal{K}_1$, is a cohomological correspondence supported on $C^\flat$, then by (3.3), (3.4) its Fourier transform

$$\mathrm{FT}_{C^\flat}(u') : \mathrm{FT}_{V_1}(c_{1!}c_0^*\mathcal{K}_0) \simeq \widehat{c}_1^*\widehat{c}_{0!}\,\mathrm{FT}_{V_0}(\mathcal{K}_0)(-d(c_0))[-2d(c_0)] \simeq \widehat{c}_{0!}'\widehat{c}_1'^*\,\mathrm{FT}_{V_0}(\mathcal{K}_0) \to \mathrm{FT}_{V_1}(\mathcal{K}_1)$$

where d_i is the Fourier dimension of V_i over S and $d^\flat$ is that of $C^\flat$. By adjunction this is equivalent to a cohomological correspondence

$$\widehat{c}_1'^*\,\mathrm{FT}_{V_0}(\mathcal{K}_0) \to \widehat{c}_0'^!\,\mathrm{FT}_{V_1}(\mathcal{K}_1)(d(c_0))(2d(c_0))$$

supported on $\widehat{C^\sharp}$, which we denote as $\mathrm{FT}_{C^\flat}(u)$. Thus we get a Fourier transform on cohomological correspondences

$$\mathrm{FT}_{C^\flat} : \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \to \mathrm{Corr}_{\widehat{C^\sharp}}(\mathrm{FT}(\mathcal{K}_0), \mathrm{FT}(\mathcal{K}_1)(d(c_0))[2d(c_0)],$$

analogous to [19, (7.1.2)] up to a different normalization of the Fourier transform. The Fourier transform is compatible with base change, so if instead we have a correspondence $X \leftarrow C \to Y$

and $V_0 \to X$, $V_1 \to Y$, and $C^\flat \to C$ very nice stacks in F-vector spaces such that the maps $C^\flat \to \widetilde{V}_0 = V_0 \times_X C$ and $C^\flat \to \widetilde{V}_1 \times_Y C$ are morphisms of stacks in F-vector spaces then we can define the Fourier transform of cohomological correspondences supported on $V_0 \leftarrow C^\flat \to V_1$ as in [19, §7.1.2].

Consider the dual maps of correspondences of F-stacks

$$\begin{array}{ccccc} V_0 & \xleftarrow[c_0]{} & C^\flat & \xrightarrow{c_1} & V_1 \\ \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f^\flat} & & \downarrow{\scriptstyle f_1} \\ W_0 & \xleftarrow[d_0]{} & D^\flat & \xrightarrow{d_1} & W_1 \end{array}, \qquad \begin{array}{ccccc} \widehat{W_0} & \xleftarrow[\widehat{d_1'}]{} & \widehat{D^\sharp} & \xrightarrow{\widehat{d_0'}} & \widehat{W}_1 \\ \downarrow{\scriptstyle \widehat{f_0}} & & \downarrow{\scriptstyle \widehat{f^\sharp}} & & \downarrow{\scriptstyle \widehat{f_1}} \\ \widehat{V}_0 & \xleftarrow[\widehat{c_1'}]{} & \widehat{C^\sharp} & \xrightarrow{\widehat{c_0'}} & \widehat{V}_1 \end{array}.$$

We assume that all the stacks are (cohomologically) smooth over the base.

Proposition 3.3.1. *With notation as above, suppose that $f^\flat$ is left pushable and $\widehat{f^\sharp}$ is cohomologically right pullable of defect δ.*[8] *Then $\delta = d(c_0) - d(d_0)$ and we have an identification*

$$\widehat{f^\sharp}^* \circ \mathrm{FT}_{C^\flat} \simeq \mathrm{FT}_{D^\flat} \circ f_!^\flat \tag{3.5}$$

of morphisms

$$\mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \to \mathrm{Corr}_{\widehat{D^\sharp}}(\widehat{f_0^*}\,\mathrm{FT}(\mathcal{K}_0), \widehat{f_1^*}\,\mathrm{FT}(\mathcal{K}_1)(d(d_0))[2d(d_0)])$$

naturally in $\mathcal{K}_i \in D(V_i, \overline{\mathbb{Q}}_\ell)$.

This is the analogue of Proposition 7.2.3 (1) in [19]. There, the result follows from Propositions 6.3.3 and 6.4.2, analogues of which we record below.

Proposition 3.3.2. *Let*

$$\begin{array}{ccc} B & \xrightarrow{g'} & A \\ \downarrow{\scriptstyle f'} & & \downarrow{\scriptstyle f} \\ D & \xrightarrow{g} & C \end{array} \quad \textit{and} \quad \begin{array}{ccc} \widehat{C} & \xrightarrow{\widehat{f}} & \widehat{A} \\ \downarrow{\scriptstyle \widehat{g}} & & \downarrow{\scriptstyle \widehat{g}'} \\ \widehat{D} & \xrightarrow{\widehat{f}'} & \widehat{B} \end{array}$$

[8] That is, the natural map $a : \widehat{D^\sharp} \to \widehat{C^\sharp} \times_{\widehat{V}_1} \widehat{W}_1$ is cohomologically quasi-smooth, i.e. carries a natural transformation $a^* \to a^!(-\delta)[-2\delta]$.

be dual Cartesian diagrams of (cohomologically) smooth very nice stacks in F-vector spaces. The base change isomorphisms and those of (3.3)*,* (3.4) *fit into a commutative diagram*

$$\begin{array}{ccc} \widehat{g}_! \widehat{f}^* \operatorname{FT}_A & \longrightarrow & \operatorname{FT}_D g^* f_! \langle d(g) \rangle \\ \downarrow & & \downarrow \\ \widehat{f}'^* \widehat{g}'_! \operatorname{FT}_A & \longrightarrow & \operatorname{FT}_D f'_! g'^* \langle d(g) \rangle \end{array} .$$

Note that unlike in [19, Proposition 6.3.3], we do not need to assume global presentability, as we can always factor a map $\phi : V \to V'$ of smooth stacks in F-vector spaces as $V \to \widetilde{V} \to V'$ with the first map a closed immersion and the second smooth. Indeed, we can reduce to the case where $V = S$ and $\phi = z_{V'}$ is the zero section: if the claim holds for $z_{V'}$, then the pullback

$$\begin{array}{ccccc} V & \longrightarrow & V \times_S \widetilde{V}' & \longrightarrow & V \times_S V' \\ \downarrow & & \downarrow & & \downarrow \\ S & \longrightarrow & \widetilde{V}' & \longrightarrow & V' \end{array}$$

along $\phi - \mathrm{id} : V \times_S V' \to V'$ gives such a factorization of the graph $V \to V \times_S V'$ of ϕ, and since the projection $V \times_S V' \to V'$ is smooth $V \to V \times_S \widetilde{V}' \to V'$ gives a factorization as claimed. To prove the claim for $z_{V'}$, we observe that when (the associated complex to) V' is concentrated in degree 0, $z_{V'}$ is a closed immersion, so it suffices to prove by induction that if $z_{V'} : S \to V'$ admits such a factorization, then so does $z_{V'/F} : S \to V'/F$. But this is immediate by $z_{V'/F} : S \to V' \to V'/F$, since the last map is smooth.

Following the proof of Proposition A.4.3 in [19], the existence of factorizations as above allows us to assume that each of f and g (or equivalently their duals) is either (cohomologically) smooth or a closed embedding (Lemma A.4.2 in *loc. cit.*). The proof of Lemma A.4.2 is then formal and applies in our situation without modification, except that we replace $\mathbb{A}^1$ by S/F and we at times need to check more possibilities, since e.g. f being a closed embedding is not necessarily equivalent to $\widehat{f}$ being smooth. The only place where this is an issue is that we need to separately check the case where g is smooth and $\widehat{g}$ is a closed embedding and the case where g and $\widehat{g}$ are both

smooth, rather than combining them; but this is similar to the case explained in *loc. cit.*

Next, suppose that $f : V \to V'$ is a map of very nice smooth stacks in F-vector spaces, and so comes equipped with a natural transformation $f^* \langle d(f) \rangle \to f^!$. By adjunction, we obtain a transformation $f_! f^* \langle d(f) \rangle \to \mathrm{id}$, and applying $\mathrm{FT}_{V'}$ and (3.3), (3.4) we get a transformation $\widehat{f^*} \widehat{f_!} \to \mathrm{id}$, so by adjunction $\widehat{f_!} \to \widehat{f_*}$.

Proposition 3.3.3. *Suppose that $\mathcal{K} \in D(V')$ is of the form $g_! \mathbf{1} \langle i \rangle$ for some integer i and map of F-vector spaces $g : V'' \to V'$, or generated by objects of this form.*[9] *Then the diagram*

$$\begin{array}{ccc} \widehat{f_!}\,\mathrm{FT}_{V'}\,\mathcal{K} & \longrightarrow & \widehat{f_*}\,\mathrm{FT}_{V'}\,\mathcal{K} \\ \downarrow & & \downarrow \\ \mathrm{FT}_V\, f^*\mathcal{K}\,\langle d(f) \rangle & \longrightarrow & \mathrm{FT}_V\, f^!\mathcal{K} \end{array}$$

commutes.

Here the left vertical arrow is the isomorphism of (3.4), and the right vertical arrow is adjoint to (3.3).

Proof. It suffices to show the result on the unit sheaf $\mathbf{1}$, i.e. that the image under FT_V of the morphism $f^* \mathbf{1} \langle d(f) \rangle \simeq \mathbf{1} \langle d(f) \rangle \to f^! \mathbf{1}$ is the map

$$\mathrm{FT}_V(f^* \mathbf{1} \langle d(f) \rangle) \simeq \widehat{f_!}\,\mathrm{FT}_{V'}(\mathbf{1}) \simeq \widehat{f_!} z_{\widehat{V'}!} \mathbf{1} \langle -d(V') \rangle \to \widehat{f_*} z_{\widehat{V'}!} \mathbf{1} \langle -d(V') \rangle \to \mathrm{FT}_V(f^! \mathbf{1}). \tag{3.6}$$

The compatibility with pushforwards can then be worked out using the base change compatibility of Proposition 3.3.2, which straightforwardly gives the result for the sheaves of the desired form. We argue that this holds, following §6.4.3 in *loc. cit.*

Factoring f as $V \xrightarrow{i} V'' \xrightarrow{p} V'$ with i a closed immersion and p a smooth surjection, it suffices to show the claim above with each of i and p in place of f, and similarly for $\widehat{f}$. So assume first that f is a smooth surjection, so that $f^! \mathbf{1} \simeq \mathbf{1} \langle d(f) \rangle$. By the compatibility of our category of

[9] This is close to the notion of geometric motives used in [20], but we require g to be a linear rather than smooth map, which makes it a more restrictive notion; however all of the relevant sheaves will satisfy this property.

(overconvergent) motives with the six-functor formalism of [46] and Theorem 4.4.29 and Footnote 2 in *loc. cit.*, we can take this isomorphism to be the identity. If $\widehat{f}$ is a closed immersion, then we can take $\widehat{f_!} \to \widehat{f_*}$ the identity, as in e.g. the proof of [42, Lemma 4.6.4 (ii)]. Thus this follows as the Fourier transform takes the identity to the identity.

By choosing the factorization of $\widehat{f}$ in general suitably we can assume that $\widehat{V}''$ is nice, and in particular has zero section a closed immersion; so if $\widehat{f}$ is a smooth surjection, we can assume that $\widehat{V}'$ has zero section a closed immersion. Since $\widehat{f}$ is a linear smooth surjection, it can be viewed as the quotient map for the quotient of $\widehat{V}'$ by $\ker \widehat{f}$, a subspace stack of $\widehat{V}'$ whose inclusion is a closed immersion; the pushout of this inclusion along $\ker \widehat{f} \to *$ is the zero section $z_{\widehat{V}}$, so this is also a closed immersion. Then

$$\widehat{f}_! z_{\widehat{V}'!}\mathbf{1} = z_{\widehat{V}!}\mathbf{1} = z_{\widehat{V}*}\mathbf{1} = \widehat{f}_* z_{\widehat{V}'*}\mathbf{1} = \widehat{f}_* z_{\widehat{V}'!}\mathbf{1}$$

and so the central map $\widehat{f}_! z_{\widehat{V}'!}\mathbf{1}\langle -d(V')\rangle \to \widehat{f}_* z_{\widehat{V}'!}\mathbf{1}\langle -d(V')\rangle$ in (3.6) is the identity as above.

Next, suppose that f is a closed immersion. Since everything is compatible with base change, we can reduce to the case where S is a point, and so choose a smooth splitting $\pi : V' \to V$ of S. Since $\pi \circ f = \mathrm{id}$, the map

$$\mathrm{id}^* \mathbf{1} = f^*\pi^*\mathbf{1} \to f^*\pi^!\mathbf{1}\langle -d(\pi)\rangle \to f^!\pi^!\mathbf{1}\langle -d(\pi) - d(f)\rangle = \mathrm{id}^! \mathbf{1}$$

is the identity, and since π is smooth as above we can take the first map to be the identity; so the second must be as well. If $\widehat{f}$ is a closed immersion, we can conclude that the central map of (3.6) is the identity as above. If $\widehat{f}$ is smooth, then as above we can assume that $z_{\widehat{V}'}$ and thus $z_{\widehat{V}}$ are closed immersions, so

$$z_{\widehat{V}!}\mathbf{1} = \widehat{f}_! z_{\widehat{V}'!}\mathbf{1} \to \widehat{f}_* z_{\widehat{V}'!}\mathbf{1} = \widehat{f}_* z_{\widehat{V}'*}\mathbf{1} = z_{\widehat{V}*}\mathbf{1}$$

is the identity, and so in particular the map in the middle must be as well. Thus in either case since the Fourier transform takes the identity to the identity the claim follows. □

Similarly, we have the following, using (3.4) in place of (3.3):

Proposition 3.3.4. *If $f^\flat$ is cohomologically right pullable of defect δ and $\widehat{f^\sharp}$ is left pushable, then $\delta = d(C^\flat) - d(D^\flat) - d(V_1) + d(W_1)$ and we have an identification*

$$\widehat{f^\sharp_!} \circ \mathrm{FT}_{D^\flat} \simeq \mathbb{T}_{(d(f_0))[2d(f_0)]} \mathrm{FT}_{C^\flat} \circ f^{\flat *}$$

of morphisms

$$\mathrm{Corr}_{D^\flat}(\mathcal{K}_0, \mathcal{K}_1) \to \mathrm{Corr}_{\widehat{C^\sharp}}(\widehat{f_{0!}} \mathrm{FT}(\mathcal{K}_0), \widehat{f_{1!}} \mathrm{FT}(\mathcal{K}_1)(d(d_0))[2d(d_0)])$$

naturally in $\mathcal{K}_i \in D(W_i, \overline{\mathbb{Q}}_\ell)$.

Here $\mathbb{T}_{(m)[n]}$ is the shift operator

$$\mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \xrightarrow{\sim} \mathrm{Corr}_{C^\flat}(\mathcal{K}_0(m)[n], \mathcal{K}_1(m)[n]).$$

The proof is similar to that of Proposition 3.3.1; alternatively it can be deduced from it via involutivity.

Note that for the purposes of evaluating the trace pairing $\langle -, - \rangle$, we sometimes need to apply a shift or twist to align the sheaves so that we can form the composition of cohomological correspondences; the shift and twist are determined uniquely by the degrees of the sheaves, so we freely drop these from the notation.[10]

3.3.2 Fourier transform over $\mathcal{O}_F$-submodule and quotient stacks

To discuss the objects of interest, we turn to the language of [25]. They construct a six-functor formalism $\mathcal{D}_{\mathrm{FF}}(-, \mathbb{Z}_p)$ which on suitable perfectoid spaces S in characteristic p can be understood as "modified solid quasi-coherent sheaves" on the un-quotiented Fargues–Fontaine curve $\mathcal{Y}_{[0,\infty)}$ over $S/\varphi^{\mathbb{Z}}$. Replacing $\mathcal{Y}_{[0,\infty)}$ by $\mathcal{Y}_{(0,\infty)}$ gives another six-functor formalism $\mathcal{D}_{\mathrm{FF}}(-, \mathbb{Q}_p)$, which

[10] An alternative would be to incorporate an index keeping track of the shifts as in [20, §5].

for suitable S can be understood as modified solid quasi-coherent sheaves on X_S. Restriction gives a natural map $\mathcal{D}_{\mathrm{FF}}(-,\mathbb{Z}_p) \to \mathcal{D}_{\mathrm{FF}}(-,\mathbb{Q}_p)$, which we denote by $\mathcal{F} \mapsto \mathcal{F}[1/p]$.

In loc. cit., Anschütz–Le Bras–Mann define a natural transformation

$$\sigma = \sigma_{\mathbb{Z}_p} : \mathcal{D}_{\mathrm{FF}}(-,\mathbb{Z}_p) \to \mathcal{D}((-)_v,\mathbb{Z}_p)$$

and analogously

$$\sigma_{\mathbb{Q}_p} : \mathcal{D}_{\mathrm{FF}}(-,\mathbb{Q}_p) \to \mathcal{D}((-)_v,\mathbb{Q}_p),$$

compatibly with restriction. The rational version $\sigma_{\mathbb{Q}_p}$ denoted by $R\tau_*$ in [24]; applied to vector bundles (or flat coherent sheaves) on the Fargues–Fontaine curve, after passing through the equivalence with stacks in F-vector spaces $\sigma_{\mathbb{Q}_p}$ can be understood as taking Banach–Colmez spaces. By Remark 4.2.6 (iii) and the proof of Lemma 5.3.1 in [25], in all of the above we can replace $\mathbb{Q}_p$ by any finite extension F and $\mathbb{Z}_p$ by $\mathcal{O}_F$, which we henceforth do without comment.

Let $\mathcal{L}$ be a complex of v-sheaves of $\mathcal{O}_F$-modules with cohomology concentrated in degrees 0 and -1, and set $\mathcal{V} = \mathcal{L}[1/p]$. There is a canonical map (by adjunction) of sheaves of $\mathcal{O}_F$-modules $\mathcal{L} \to \mathcal{V}$; if it is injective, we say that $\mathcal{L}$ is torsion-free. Conversely if $\mathcal{V} = \mathcal{L}[1/p] = 0$, say that $\mathcal{L}$ is torsion. In particular $\mathcal{V}/\mathcal{L} = \mathrm{coker}(\mathcal{L} \to \mathcal{V})$ is torsion.

In the above setup, $\mathcal{V}$ has an associated stack in F-vector spaces V over S, and $\mathcal{L}$ can then be viewed as a v-stack $p_L : L \to S$ equipped with a map $i_L : L \to V$. Suppose L is torsion-free, so that i_L is injective; we will then say that L is an $\mathcal{O}_F$-submodule stack of V. Then we can associate to $\mathcal{V}/\mathcal{L}$ the stack $V/L := \mathrm{cofib}(i_L)$, which is a v-stack over S with an $\mathcal{O}_F$-action and quotient map $q_L : V \to V/L$.

We consider the duality on complexes of $\mathcal{O}_F$-sheaves given by $\underline{R\,\mathrm{Hom}}(-, F/\mathcal{O}_F[1])$, given on stacks by $\widehat{-} := \underline{\mathrm{Hom}}(-, */(F/\mathcal{O}_F))$. Restricted to stacks in F-vector spaces, this recovers the duality from §3.3.1. If

$$\mathcal{L} \to \mathcal{V} \to \mathcal{V}/\mathcal{L}$$

is the exact triangle in complexes of v-sheaves of $\mathbb{Z}_p$-modules associated to an $\mathcal{O}_F$-submodule

stack $L \to V$, taking duals we obtain an exact triangle

$$\widehat{\mathcal{V}/\mathcal{L}} \to \widehat{\mathcal{V}} \to \widehat{\mathcal{L}}$$

provided $R^2 \operatorname{Hom}(\mathcal{V}/\mathcal{L}, F/\mathcal{O}_F) = 0$. In particular $\widehat{V/L} \to \widehat{V}$ is an $\mathcal{O}_F$-submodule stack.

Let L be an $\mathcal{O}_F$-submodule stack of V. Let $\mathrm{ev}_L : L \times \widehat{L} \to */(F/\mathcal{O}_F)$ be the evaluation map and $\pi_L : L \times \widehat{L} \to L$, $\pi_{\widehat{L}} \to \widehat{L}$ be the projections. We define the Fourier transform

$$\mathrm{FT} = \mathrm{FT}_L^{\psi} = \pi_{\widehat{L}!}(\pi_L^*(-) \otimes \mathrm{ev}_L^* \mathcal{L}_\psi) : D(L, \overline{\mathbb{Q}}_\ell) \to D(\widehat{L}, \overline{\mathbb{Q}}_\ell).$$

Similarly, we can define

$$\mathrm{FT} = \mathrm{FT}_{V/L}^{\psi} = \pi_{\widehat{V/L}!}(\pi_{V/L}^*(-) \otimes \mathrm{ev}_{V/L}^* \mathcal{L}_\psi : D(V/L, \overline{\mathbb{Q}}_\ell) \to D(\widehat{V/L}, \overline{\mathbb{Q}}_\ell),$$

with analogous notation. We assume henceforth that $L \to V$ satisfies $R^2 \operatorname{Hom}(\mathcal{V}/\mathcal{L}, F/\mathcal{O}_F) = 0$ and that $\widehat{\widehat{L}} \simeq L$.

We check some functoriality properties of this integral Fourier transform. If $L \to V$ and $L' \to V'$ are $\mathcal{O}_F$-submodule stacks and $f : L \to L'$ is a morphism such that $f[1/p] : V \to V'$ is a morphism of stacks in F-vector spaces, the dual $\widehat{f[1/p]} : \widehat{V}' \to \widehat{V}$ induces $\widehat{f} : \widehat{L}' \to \widehat{L}$. We have

$$\mathrm{FT}_{L'}^{\psi} \circ f_! \simeq \widehat{f}^* \circ \mathrm{FT}_L^{\psi}; \tag{3.7}$$

this is proven in essentially the same way as for the classical ℓ-adic Fourier transform, analogous to (3.3). The same holds on the torsion side:

$$\mathrm{FT}_{V'/L'}^{\psi} \circ f_! \simeq \widehat{f}^* \circ \mathrm{FT}_{V/L}^{\psi}\,. \tag{3.8}$$

For the analogue of (3.4), we again need to pass through some version of (shifted) involutivity.

Consider the Fourier transform of the unit sheaf

$$\mathrm{FT}^{\psi}_{\widehat{L}}(\mathbf{1}) \simeq \pi_{L!}\,\mathrm{ev}^{*}_{\widehat{L}}\,\mathcal{L}_{\psi}.$$

As $\widehat{L} \simeq \widehat{V}/(\widehat{V/L})$, we have $p_L^* p_{\widehat{V/L}!}\mathbf{1} \otimes \pi_{L!} \simeq \pi_{L!}(- \otimes (q_{\widehat{V/L}!} \times \mathrm{id}_L)_!\mathbf{1}) \simeq i_L^* i_{L!}\pi_{L!}(q_{\widehat{V/L}!} \times \mathrm{id}_L)_!(q_{\widehat{V/L}!} \times \mathrm{id}_L)^* \simeq i_L^*\pi_{V!}(\mathrm{id}_{\widehat{V}} \times i_L)_!$ and so

$$p_L^* p_{\widehat{V/L}!}\mathbf{1} \otimes \mathrm{FT}^{\psi}_{\widehat{L}}(\mathbf{1}) \simeq i_L^*\,\mathrm{FT}^{\psi}_{\widehat{V}}(\mathbf{1}) \simeq i_L^* z_{V!}\mathbf{1}(-d(V))[-2d(V)]$$

by (3.4), where $z_V : S \to V$ is the zero section and $d(V)$ is the Fourier dimension. We will later see that the $p_{\widehat{V/L}!}\mathbf{1}$ term is actually just a shift and twist, which simplifies the expression, but for now we leave it as above. Just as in the proof of [19, Lemma A.2.4], this lets us identify the kernel sheaf of $p_L^* p_{\widehat{V/L}!}\mathbf{1} \otimes \mathrm{FT}^{\psi^{-1}}_{\widehat{L}} \circ \mathrm{FT}^{\psi}_{L}$ with $\Delta^-_!\mathbf{1}(-d(V))[-2d(V)]$, where $\Delta^- : L \to L \times L$ is the anti-diagonal. Thus we have an identification of natural transforms

$$p_L^* p_{\widehat{V/L}!}\mathbf{1} \otimes \mathrm{FT}^{\psi^{-1}}_{\widehat{L}} \circ \mathrm{FT}^{\psi}_{L} \simeq [-1]^*(-d(V))[-2d(V)]. \tag{3.9}$$

If we assume that FT^{ψ}_{L} is an equivalence, then it follows that its inverse is

$$p_L^* p_{\widehat{V/L}!}\mathbf{1} \otimes [-1]^*\,\mathrm{FT}^{\psi^{-1}}_{\widehat{L}}(d(V))[2d(V)].$$

In this case, together with the assumptions above, we say that L is very nice, as well as V. In all the cases in which we will want to apply our Fourier analysis this will be true. In particular for L obtained via σ from lattices inside torsors on the Fargues–Fontaine curve, arguing as in [24, Propositions 3.34, 3.35] shows that L is very nice.

As a corollary, when L is very nice we can deduce the following analogue of (3.4), under the assumptions above on $L \to V$ (and similarly on $L' \to V'$) and those on f from the discussion

around (3.7): applying (3.9) to (3.7) and (3.8) we find

$$\begin{aligned} p_{L'}^* p_{\widehat{V/L}!}\mathbf{1} \otimes [-1]^* f_! \circ \mathrm{FT}_{\widehat{L}}^{\psi^{-1}} \langle -d(V')\rangle &\simeq p_{L'}^*(p_{\widehat{V/L}!}\mathbf{1} \otimes p_{\widehat{V'/L'}!}\mathbf{1}) \otimes \mathrm{FT}_{\widehat{L'}}^{\psi^{-1}} \circ \mathrm{FT}_{L'}^{\psi} \circ f_! \circ \mathrm{FT}_{\widehat{L}}^{\psi^{-1}} \\ &\simeq p_{L'}^*(p_{\widehat{V/L}!}\mathbf{1} \otimes p_{\widehat{V'/L'}!}\mathbf{1}) \otimes \mathrm{FT}_{\widehat{L'}}^{\psi^{-1}} \circ \widehat{f}^* \circ \mathrm{FT}_{L}^{\psi} \circ \mathrm{FT}_{\widehat{L}}^{\psi^{-1}} \\ &\simeq p_{L'}^* p_{\widehat{V'/L'}!}\mathbf{1}\, \mathrm{FT}_{\widehat{L'}}^{\psi^{-1}} \circ \widehat{f}^*[-1]^* \langle -d(V)\rangle \end{aligned}$$

and so

$$p_{L'}^* p_{\widehat{V/L}!}\mathbf{1} \otimes f_! \circ \mathrm{FT}_{\widehat{L}}^{\psi^{-1}} \simeq p_{L'}^* p_{\widehat{V'/L'}!}\mathbf{1} \otimes \mathrm{FT}_{\widehat{L'}}^{\psi^{-1}}\, \widehat{f}^* \langle -d(f[1/p])\rangle\,,$$

or replacing f with $\widehat{f}$ and ψ with ψ^{-1}

$$p_{\widehat{L}}^* p_{\widehat{V/L}!}\mathbf{1} \otimes \mathrm{FT}_{L}^{\psi} \circ f^* \simeq p_{\widehat{L}}^* p_{\widehat{V'/L'}!}\mathbf{1} \otimes \widehat{f_!} \circ \mathrm{FT}_{L'} \langle -d(f[1/p])\rangle\,. \tag{3.10}$$

Suppose we have a pair of O_F-submodule stacks $L_0 \to V_0$, $L_1 \to V_1$, and consider the correspondence $V_0 \leftarrow L_0 \times L_1 \to V_1$. This fits into the Cartesian and co-Cartesian diagram

$$\begin{array}{ccc} L_0 \times L_1 & \xrightarrow{c_0} & V_0 \\ \downarrow{\scriptstyle c_1} & & \downarrow{\scriptstyle c_1'} \\ V_1 & \xrightarrow{c_0'} & V_0/L_0 \times V_1/L_1 \end{array}\,.$$

Taking duals as above gives the dual diagram

$$\begin{array}{ccc} \widehat{V_0/L_0} \times \widehat{V_1/L_1} & \xrightarrow{\widehat{c_1'}} & \widehat{V}_0 \\ \downarrow{\scriptstyle \widehat{c_0'}} & & \downarrow{\scriptstyle \widehat{c}_0} \\ \widehat{V}_1 & \xrightarrow{\widehat{c}_1} & \widehat{L}_0 \times \widehat{L}_1 \end{array}\,.$$

Let $\mathcal{K}_i \in D(V_i, \overline{\mathbb{Q}}_\ell)$ and let $u : c_0^*\mathcal{K}_0 \to c_1^!\mathcal{K}_1$, or equivalently $u' : c_{1!}c_0^*\mathcal{K}_0 \to \mathcal{K}_1$, be a cohomological correspondence supported on $L_0 \times L_1$. By (3.7) and (3.10), its Fourier transform

gives a map

$$\begin{aligned} p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}!}\mathbf{1}\otimes \mathrm{FT}_{V_1}(c_{1!}c_0^*\mathcal{K}_0) &\simeq p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}!}\mathbf{1}\otimes \widehat{c}_1^*\,\mathrm{FT}_{L_0\times L_1}(c_0^*\mathcal{K}_0)\\ &\simeq \widehat{c}_1^*\widehat{c}_{0!}\,\mathrm{FT}_{V_0}(\mathcal{K}_0)(-d(V_1))[-2d(V_1)]\\ &\to p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}!}\mathbf{1}\otimes \mathrm{FT}_{V_1}(\mathcal{K}_1), \end{aligned}$$

or equivalently by base change and adjunction

$$\mathrm{FT}_{L_0\times L_1}(u) : \widehat{c}_0'^*\,\mathrm{FT}_{V_0}(\mathcal{K}_0) \to p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}!}\mathbf{1}\otimes \widehat{c}_1'^!\,\mathrm{FT}_{V_1}(\mathcal{K}_1)\,\langle d(V_1)\rangle\,.$$

Thus we get a Fourier transform

$$\mathrm{FT}_{L_0\times L_1} : \mathrm{Corr}_{L_0\times L_1}(\mathcal{K}_0,\mathcal{K}_1) \to \mathrm{Corr}_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}}(\mathrm{FT}(\mathcal{K}_0), p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}!}\mathbf{1}\otimes \mathrm{FT}(\mathcal{K}_1)\,\langle d(V_1)\rangle).$$

Similarly to the discussion in §3.3.1, the functoriality relations (3.7), (3.10) induces an identification

$$\widehat{f^\sharp}^*\circ \mathrm{FT}_{L_0\times L_1} \simeq \mathrm{FT}_{L_0'\times L_1'}\circ f^\flat_! \tag{3.11}$$

for $f^\flat : L_0\times L_1 \to L_0'\times L_1'$ left pushable compatibly with $f_i : V_i \to V_i'$ and $f^\sharp$ the induced map $V_0/L_0\times V_1/L_1 \to V_0'/L_0'\times V_1'/L_1'$. Consider the map $i^\flat : L_0\times L_1 \to V_0\times V_1$ of correspondences of $\mathcal{O}_F$-submodule stacks between V_0 and V_1, with dual map $\widehat{i^\sharp} : \widehat{V_0/L_0}\times\widehat{V_1/L_1} \to \widehat{V}_0\times\widehat{V}_1$. By (3.11), for $u\in \mathrm{Corr}_{L_0\times L_1}(\mathcal{K}_0,\mathcal{K}_1)$ we have $\mathrm{FT}_{V_0\times V_1}(i_!u) \simeq \widehat{i^\sharp}^*\,\mathrm{FT}(u)$. The left-hand side is a cohomological correspondence between $\mathrm{FT}_{V_0}(\mathcal{K}_0)$ and $\mathrm{FT}_{V_1}(\mathcal{K}_1)(d(V_1))[2d(V_1)]$, while the right-hand side is between $\mathrm{FT}_{V_0}(\mathcal{K}_0)$ and $p^*_{\widehat{V}_1} p_{\widehat{V_0/L_0}\times\widehat{V_0/L_1}!}\mathbf{1}\otimes\mathrm{FT}_{V_1}(\mathcal{K}_1)(d(V_1)-d(\widehat{i^\sharp}))[2(d(V_1)-d(\widehat{i^\sharp}))]$, so we can identify $p_{\widehat{V_0/L_0}\times\widehat{V_0/L_1}!}\mathbf{1}$ with $\mathbf{1}(d(\widehat{i^\sharp}))[2d(\widehat{i^\sharp})] = \mathbf{1}(d(i^\flat))[2d(i^\flat)]$. In particular taking one of the L_i to be all of V_i we find that

$$p_{\widehat{V/L}!}\mathbf{1} \simeq \mathbf{1}(d(i_L))[2d(i_L)] \tag{3.12}$$

as claimed; so very nice $\mathcal{O}_F$-submodule stacks are prim over the base. Therefore we can rewrite the Fourier transform as

$$\mathrm{FT}_{L_0\times L_1} : \mathrm{Corr}_{L_0\times L_1}(\mathcal{K}_0, \mathcal{K}_1) \to \mathrm{Corr}_{\widehat{V_0/L_0}\times\widehat{V_1/L_1}}(\mathrm{FT}(\mathcal{K}_0), \mathrm{FT}(\mathcal{K}_1)(d(c_0))[2d(c_0)]).$$

For $f^\flat$ right pullable we similarly have

$$\widehat{f^\sharp}_! \circ \mathrm{FT}_{L_0'\times L_1'} \simeq \mathbb{T}_{(d(c_0')-d(c_0))[2(d(c_0')-d(c_0))]}\, \mathrm{FT}_{L_0\times L_1} \circ f^{\flat *}, \tag{3.13}$$

where to line up the shifts and twists we have used the above identification of the pushforward with the twist.

3.4 Hecke and shtuka stacks

We now suppose that p is an odd prime, F a finite extension of $\mathbb{Q}_p$ as above, and F'/F a quadratic unramified extension, with residue fields $\mathbb{F}_{q'}/\mathbb{F}_q$. Fix a vector space V over F' of rank $2m$ with a nondegenerate skew-Hermitian pairing $\langle -, - \rangle : V \times V \to F'$ with respect to the quadratic extension F'/F, and let G be the similitude automorphism group of $(V, \langle -, - \rangle)$, viewed as a reductive group over F; that is, for an F-algebra R,

$$G(R) = \{g \in \mathrm{GL}(V \otimes_F R) | \langle gu, gv \rangle = \lambda \langle u, v \rangle \forall u, v \in V \otimes_F R\}$$

where $\lambda = \lambda(g)$ depends on g but not on u and v. Fix an $\mathcal{O}_{F'}$-lattice $\Lambda \subset V$, and let $J \subset G(F)$ be the compact open subgroup stabilizing Λ. Choose a Lagrangian subspace $L \subset V$, and let $P \subset G$ be the parabolic subgroup stabilizing L with Levi subgroup $M \simeq \mathrm{Res}_{F'/F}\,\mathrm{GL}(m) \subset P$.

On the other side, let W be a F'/F-Hermitian vector space of rank n and H be its unitary similitude group. We again fix a lattice $\Xi \subset W$ and let $K_1 \subset H(F)$ be the compact open subgroup stabilizing Ξ. We define the following v-stacks:

Definition 3.4.1. (a) Let Bun_H be the stack over $\mathrm{Spd}\,\mathbb{F}_q$ sending a perfectoid space S over

$\mathbb{F}_q$ to the groupoid of H-bundles on the relative Fargues–Fontaine curve $X_S = X_{F,S} = (S\dot{\times}_{\operatorname{Spa}\mathbb{F}_q} \operatorname{Spa} F)/\varphi^{\mathbb{Z}}$ where φ is the Frobenius on S and $S\dot{\times}_{\operatorname{Spa}\mathbb{F}_q} \operatorname{Spa} F$ is defined as in [21]. Explicitly, a H-bundle on X_S can be thought of as a vector bundle $\mathcal{F}$ on $X'_S = X_{F',S}$ together with a line bundle $\mathcal{L}$ on X_S and a Hermitian pairing $\mathcal{F} \otimes \mathcal{F} \to \nu^*\mathcal{L}$, where $\nu : X'_S \to X_S$ is the associated double cover.

(b) Let Hk^1_H be the stack over $\mathrm{Bun}_H \times \mathrm{Bun}_H$ sending S to a pair of H-bundles $(\mathcal{F}_0, \mathcal{F}_1)$ on X_S, i.e. Hermitian vector bundles on X'_S as above, together with a vector bundle $\mathcal{F}^\flat_{1/2}$ on X'_S and injections $j_i : \mathcal{F}^\flat_{1/2} \hookrightarrow \mathcal{F}_i$ for $i = 0, 1$ such that $\operatorname{coker}(j_i)$ is a locally free sheaf supported on a degree 1 divisor D_i of X'_S, such that the automorphism σ of X'_S over X_S interchanges D_0 and D_1, and such that away from D_0 and D_1 the composite isomorphism $\mathcal{F}_0|_{X_S \setminus D_0 \cup D_1} \xrightarrow{\sim} \mathcal{F}_1|_{X_S \setminus D_0 \cup D_1}$ is compatible with the Hermitian pairings on both sides.[11] In particular, if $\mathcal{L}_i$ is the similitude line bundle of $\mathcal{F}_i$, then the modification induces an isomorphism $\mathcal{L}_0 \simeq \mathcal{L}_1$, so we fix a single similitude bundle $\mathcal{L}$.[12] Write $p_i : \mathrm{Hk}^1_H \to \mathrm{Bun}_H$ for $i = 0, 1$ for the projections $(\mathcal{F}_0 \leftarrow \mathcal{F}^\flat_{1/2} \to \mathcal{F}_1) \mapsto \mathcal{F}_i$, and for $r \geq 0$ set

$$\mathrm{Hk}^r_H = \mathrm{Hk}^1_H \times_{p_1, \mathrm{Bun}_H, p_0} \cdots \times_{p_1, \mathrm{Bun}_H, p_0} \mathrm{Hk}^1_H,$$

so that Hk^r_H parametrizes chains

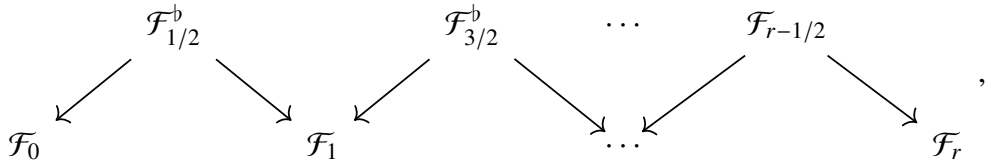

which we sometimes write more informally as $(\mathcal{F}_0 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_r)$. The divisor D_0 corresponds to an untilt $S^\sharp_0$ of S over $\mathcal{O}_{F'}$, so Hk^1_H carries a map p_X to the ν-sheaf $\operatorname{Spd}\mathcal{O}_{F'}$ parametrizing untilts over $\mathcal{O}_{F'}$; since $D_1 = \sigma(D_0)$ this fully describes the locations of the

[11] More properly, this should be $\mathrm{Hk}^1_{H,\mu}$ for minuscule μ bounding modifications to be of this form, with similar stacks for any cocharacter μ; but we will work with modifications of this type throughout and so disregard the distinction.

[12] We will often implicitly regard $\mathcal{L}$ as fixed, amounting to taking the component of the relevant stacks over each $\mathcal{L}$.

modifications. Thus there is a corresponding map $p_X : \mathrm{Hk}_H^r \to (\operatorname{Spd} \mathcal{O}_{F'})^r$. Similarly, the projections to $\mathcal{F}_i$ give maps $p_i : \mathrm{Hk}_H^r \to \mathrm{Bun}_H$ for $0 \le i \le r$.

(c) For any $b \in B(H)_{\text{basic}}$ and perfectoid space S over $\mathbb{F}_q$, let $\mathcal{F}^b$ denote the H-bundle on X_S given by the pullback of the H-isocrystal associated to b along the map $X_S \to \operatorname{Spa} \breve{F}/\varphi^{\mathbb{Z}}$; for example for $b = 1$, $\mathcal{F}^1$ is the trivial H-bundle. Here $B(H)$ is the Kottwitz set [23]. This can be viewed as a morphism $\operatorname{Spd} \mathbb{F}_q \to \mathrm{Bun}_H$, whose image is the locus Bun_H^b. We let $\mathrm{Sht}_{H,b,\infty}^r = \mathrm{Hk}_H^r \times_{p_0 \times p_r, \mathrm{Bun}_H \times \mathrm{Bun}_H} \operatorname{Spd} \mathbb{F}_q$ where the map $\operatorname{Spd} \mathbb{F}_q \to \mathrm{Bun}_H \times \mathrm{Bun}_H$ is given by the product of the maps associated to $\mathcal{F}^1$ and $\mathcal{F}^b$; thus $\mathrm{Sht}_{H,b,\infty}^r$ parametrizes chains $(\mathcal{F}^1 \dashrightarrow \mathcal{F}_1 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_{r-1} \dashrightarrow \mathcal{F}^b)$. We have an action of $H(F)$ on $\mathrm{Sht}_{H,b,\infty}^r$ as the automorphisms of $\mathcal{F}^1$, and similarly an action of the inner form $H_b(F)$ as the automorphisms of $\mathcal{F}^b$; for $K_1 \subset H(F)$ a compact open subgroup as above $K_b \subset H_b(F)$ another compact open subgroup, $K = K_1 \times K_b$, we set $\mathrm{Sht}_{H,b,K}^r = \mathrm{Sht}_{H,b,\infty}^r / K$. By [21, Theorem 22.5.2], we can equivalently think of $\mathcal{F}^1$ as the trivial H-torsor on S, and similarly for $\mathcal{F}^b$ as the trivial H_b-torsor; so $\mathrm{Sht}_{H,b,K}^r$ can be thought of as classifying chains $(\mathcal{F}^1 \dashrightarrow \mathcal{F}_1 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_{r-1} \dashrightarrow \mathcal{F}^b)$ as above together with a K_1-torsor $\mathfrak{P}$ inside the trivial H-torsor on S and a K_b-torsor $\mathbb{R}$ inside the trivial H_b-torsor on S. The similitude bundle of $\mathcal{F}^1$ is trivial, so so must that of each $\mathcal{F}_i$ be; thus the data of the similitude bundle is just the choice of $\mathcal{O}_F$-lattice in the trivial line bundle.

Remark 3.4.2. The stack of shtukas of Definition 3.4.1 is not quite the same as that of [21, Definition 23.1.1]; it is defined over the Fargues–Fontaine curve X_S, after quotienting by Frobenius, rather than over $Y_S = S \dot{\times} \operatorname{Spa} F$, and uses iterated modifications. More substantially, it involves integral data at both $\mathcal{F}_0 \simeq \mathcal{F}^1$ and $\mathcal{F}_r \simeq \mathcal{F}^b$, rather than just at $\mathcal{F}_0$; equivalently we quotient the infinite level stack of shtukas by subgroups of both the $H(F)$- and $H_b(F)$-actions rather than just one. The term "symmetrized stack of shtukas" is already taken (see [47]), but we might view this version as "symmetrized" with respect to the involution of local shtuka data $(G, b, \{\mu_i\}) \mapsto (G_b, b^{-1}, \{\mu_i^{-1}\})$ (cf. [21, Corollary 23.3.2]).

Despite Remark 3.4.2, we can still use [21, p. 23.1.4] to conclude that our stacks of shtukas are

small v-stacks, and in fact locally spatial diamonds, as they are a pro-étale quotient of the infinite-level stacks of shtukas in *loc. cit.* Taking suitable quotients and letting $\mathcal{F}_0$ and $\mathcal{F}_r$ vary, we see that our Hecke stacks are small v-stacks as well, as is Bun_H by [44, Proposition III.1.3].

Remark 3.4.3. Consider the correspondence $(p_0, p_1) : \mathrm{Hk}_H^1 \to \mathrm{Bun}_H \times \mathrm{Bun}_H$ of Definition 3.4.1 (b). Note that Hk_H^r is the r-fold self-composition of Hk_H^1 as a self-correspondence of Bun_H; thus we think of it as a correspondence $(p_0, p_r) : \mathrm{Hk}_H^r \to \mathrm{Bun}_H \times \mathrm{Bun}_H$. Fix $b \in B(H)_{\mathrm{basic}}$, and consider the correspondence $j = (j^1, j^b) : * \to \mathrm{Bun}_H \times \mathrm{Bun}_H$ given by the pair of H-bundles $(\mathcal{F}^1, \mathcal{F}^b)$, with image $\mathrm{Bun}_H^1 \times \mathrm{Bun}_H^b$. Then $\mathrm{Sht}_{H,b,\infty}^r = \mathrm{Hk}_H^r \cap j$. By (3.2), this is $\mathrm{Fix}(j^\vee \circ \mathrm{Hk}_H^r)$. Noting that j^1 factors as $* \to *\/H(F) \simeq \mathrm{Bun}_H^1 \hookrightarrow \mathrm{Bun}_H$, with the first map a quotient and the second a locally closed embedding, a cohomological correspondence $\mu : j^{b*}\mathbf{1} = \mathbf{1} \to j^{1!}\mathbf{1}$ is equivalent to a choice of Haar measure on $H(F)$. With étale coefficients, this (as well as the more general discussion below) follows from [45, Example 4.2.2]; the passage to motivic coefficients is explained briefly in the proof of Proposition 2.1 in [41] and in greater detail in [42, Example 5.3.21, Proposition 5.4.4]. Fixing such a Haar measure, we obtain a trace map $u \mapsto \mathrm{tr}(\mu \circ u)$ from self-cohomological correspondences of the constant sheaf (up to shift and twist) supported on Hk_H^r to $H_*^{\mathrm{BM}}(\mathrm{Sht}_{H,b,\infty}^r)$. If we replace $j^1 : * \to *\/H(F) \hookrightarrow \mathrm{Bun}_H$ by $j_K^1 : *\/K_1 \to *\/H(F) \hookrightarrow \mathrm{Bun}_H$ and similarly j^b by $j_K^b : * \to *\/K_b \to *\/H_b(F) \hookrightarrow \mathrm{Bun}_H$, we obtain the correspondence $j_K = j_K^1 \times j_K^b : *\/K \to \mathrm{Bun}_H \times \mathrm{Bun}_H$ whose intersection with Hk_H^r gives the finite-level shtuka space $\mathrm{Sht}_{H,b,K}^r$, with the Haar measure μ replaced by the induced Haar measure on $H(F)/K_1$ together with a functional on the module of Haar measures on $K_b(F)$, which since K_b is compact we can take to be evaluation on the constant function 1.

3.5 Relative vector spaces and duality

In this section we assume for simplicity that $\mathcal{L}$ is trivial, though there should be no difficulty in generalizing to arbitrary $\mathcal{L}$.

Recall that we defined the compact open subgroup $J \subset G(F)$ as the stabilizer of an $\mathcal{O}_{F'}$-lattice $\Lambda \subset V$. Let $\breve{F}$ denote the maximal unramified extension of F, and $\breve{F}' = \breve{F} \otimes_F F'$ be a degree

2 extension, possibly split. Base changing to $\breve{F}$ gives a subgroup $\breve{J}$ of $G(\breve{F})$ as the stabilizer of $\Lambda \otimes_{O_F} O_{\breve{F}}$. For every $a \in B(G)_{\text{basic}}$, we obtain a compact open subgroup J_a of the corresponding inner form $G_a(F)$ defined as $J_a = \breve{J} \cap G_a(F) \subset G(\breve{F})$, recovering $J = J_1$.

Similarly, we can define twisted forms $P_a \subset G_a$ of the parabolic $P \subset G$, equipped with compact open subgroups $J_a \cap P_a(F) \subset P_a(F)$. The idea is that we should broaden the scope of the modularity statement to include all basic a simultaneously. That is: we will a priori define the theta series as a function on $\bigsqcup_a \mathrm{Sht}^0_{P_a, J_a \cap P_a(F)} = \bigsqcup_a P_a(F)/(J_a \cap P_a(F))$, valued in cycles on the stack of H-shtukas, and show (under suitable restrictions) that it descends to $\bigsqcup_a \mathrm{Sht}_{G_a, J_a} = \bigsqcup_a G_a(F)/J_a$. Restricting to the $a = 1$ component will then recover the expected statement of modularity.

Let C be an algebraically closed nonarchimedean field of characteristic p, so that $\operatorname{Spa} C$ is a geometric point; our test perfectoid spaces S will all be over C. Equivalently, we can view our various v-stacks from Definition 3.4.1 as base changed to C.

Remark 3.5.1. It would be preferable, at least aesthetically, to work over the "absolute point" $\operatorname{Spd} \mathbb{F}_q$. Here, though, the Fargues–Fontaine curve is an incarnation of the p-adic field F, which is too rigid for our needs: e.g. any injection of finite rank F-modules is an isomorphism.

Let $\mathcal{G} \in \mathrm{Bun}_G(C)$ be the G-bundle corresponding to some $a \in B(G)_{\text{basic}}$, and let $\widetilde{\mathcal{E}}_1, \widetilde{\mathcal{E}}_2 \subset \mathcal{G}$ be transverse Lagrangian subbundles, which will be M-bundles on the Fargues–Fontaine curve corresponding to the induced class in $B(M)_{\text{basic}}$, which by an abuse of notation we also denote as a; these can be viewed as trivial G_a-torsors, resp. M_a-torsors. Let $\mathfrak{G} \subset \mathcal{G}$, $\widetilde{\mathfrak{E}}_1 \subset \widetilde{\mathcal{E}}_1$, and $\widetilde{\mathfrak{E}}_2 \subset \widetilde{\mathcal{E}}_2$ be respectively J_a- and $J_a \cap M_a$-subtorsors of the trivial G_a- (M_a-) torsor. For each $i \in \{1, 2\}$ let $\mathcal{E}_i \hookrightarrow \widetilde{\mathcal{E}}_i$ be a subbundle with cokernel $\mathcal{T}_i$ a torsion sheaf on the Fargues–Fontaine curve X_C. Note that $\mathcal{E}_i$ can again be viewed as a trivial $M_{a'}$-torsor on the Fargues–Fontaine curve for some class $a' \in B(M)_{\text{basic}}$. We correspondingly let $\mathfrak{E}_i \subset \mathcal{E}_i$ be $J_a \cap M_a(F)$-subtorsor, fitting into the

commutative diagram

$$\begin{array}{ccc} \mathfrak{E}_i & \longrightarrow & \widetilde{\mathfrak{E}}_i \\ \downarrow & & \downarrow \\ \mathcal{E}_i & \longrightarrow & \widetilde{\mathcal{E}}_i \end{array},$$

and write $\mathfrak{T}_i$ for the cokernel of $\mathfrak{E}_i \to \widetilde{\mathfrak{E}}_i$.

Let

$$b_{12} : \widetilde{\mathcal{E}}_1 \to \mathcal{G} \to \sigma^* \widetilde{\mathcal{E}}_2^*, \qquad b_{21} : \widetilde{\mathcal{E}}_2 \to \mathcal{G} \to \sigma^* \widetilde{\mathcal{E}}_1^* \tag{3.14}$$

be the composite maps, injective by the transversality of the $\widetilde{\mathcal{E}}_i$, with cokernels Q_2 and Q_1 respectively, $\mathcal{G}^\sharp = \sigma^* \widetilde{\mathcal{E}}_1^* \oplus \sigma^* \widetilde{\mathcal{E}}_2^*$, and Q the cokernel of the natural injection $\mathcal{G} \to \mathcal{G}^\sharp$, which is isomorphic to both Q_1 and Q_2. We assume the the supports of Q and of the $\mathcal{T}_i$ are disjoint. Writing $\widetilde{Q}_1 = Q^* \oplus \mathcal{T}_1^* \oplus \sigma^* \mathcal{T}_2$, where $T^* = \underline{\mathrm{Ext}}^1(T, \mathcal{O})$ for a torsion sheaf T, and $\widetilde{Q}_2 = \sigma^* Q \oplus \sigma^* \mathcal{T}_1 \oplus \mathcal{T}_2^*$, we obtain short exact sequences

$$0 \to \sigma^* \mathcal{E}_2 \to \mathcal{E}_1^* \to \widetilde{Q}_1 \to 0 \quad \text{and} \quad 0 \to \sigma^* \mathcal{E}_1 \to \mathcal{E}_2^* \to \widetilde{Q}_2 \to 0. \tag{3.15}$$

Analogously, we have maps

$$\mathfrak{b}_{12} : \widetilde{\mathfrak{E}}_1 \to \mathfrak{G} \to \sigma^* \widetilde{\mathfrak{E}}_2^*, \qquad \mathfrak{b}_{21} : \widetilde{\mathfrak{E}}_2 \to \mathfrak{G} \to \sigma^* \widetilde{\mathfrak{E}}_1$$

with cokernels $\mathfrak{Q}_2$ and $\mathfrak{Q}_1$ respectively, $\mathfrak{G}^\sharp = \sigma^* \widetilde{\mathfrak{E}}_1^* \oplus \sigma^* \widetilde{\mathfrak{E}}_2^*$, and $\mathfrak{Q}$ the cokernel of the natural injection $\mathfrak{G} \to \mathfrak{G}^\sharp$, isomorphic to each of the $\mathfrak{Q}_i$. Parallel to (3.15), we obtain short exact sequences

$$0 \to \sigma^* \mathfrak{E}_2 \to \mathfrak{E}_1^* \to \widetilde{\mathfrak{Q}}_1 \to 0 \quad \text{and} \quad 0 \to \sigma^* \mathfrak{E}_1 \to \mathfrak{E}_2^* \to \widetilde{\mathfrak{Q}}_2 \to 0 \tag{3.16}$$

where $\widetilde{\mathfrak{Q}}_i$ are defined analogously to $\widetilde{Q}_i$, replacing Q by $\mathfrak{Q}$ and $\mathcal{T}_i$ by $\mathfrak{T}_i$.

Let U be the stack in F-vector spaces over Bun_H corresponding to the complex of F-vector

spaces $\mathcal{U}_{\mathcal{F}}$ with fiber over $\mathcal{F} \in \mathrm{Bun}_H(S)$ given by

$$\mathbb{H}\mathrm{om}(\mathcal{E}_1, \mathcal{F}) \simeq \mathbb{H}\mathrm{om}(\mathcal{F}^*, \mathcal{E}_1^*) := \mathcal{BC}(\underline{R\,\mathrm{Hom}}(\mathcal{F}^*, \mathcal{E}_1^*)).$$

Here we write $\mathcal{E}_1^*$ for its pullback to X'_S by an abuse of notation. We similarly define $V \to \mathrm{Bun}_H$ as the stack corresponding to the complex $\mathcal{V}$ whose fiber over an S-point $\mathcal{F}$ of Bun_H is $\mathcal{V}_{\mathcal{F}} = \mathbb{H}\mathrm{om}(\mathcal{F}^*, \widetilde{Q}_1)$, and $W \to \mathrm{Bun}_H$ corresponding to the complex $\mathcal{W}$ whose fiber over $\mathcal{F}$ corresponds to $\mathbb{E}\mathrm{xt}^1(\sigma^*\mathcal{E}_2^*, \mathcal{F}) \simeq \mathbb{E}\mathrm{xt}^1(\mathcal{F}^*, \sigma^*\mathcal{E}_2)$. We assume that each of the corresponding sheaves of F-vector spaces is concentrated in degree 0; this is automatic for V since $\widetilde{Q}_1$ is torsion, and for U and V is really an assumption on the $\mathcal{E}_i$, which can be guaranteed by taking the degrees of the torsion sheaves $\mathcal{T}_i$ sufficiently large.[13] In particular, we can even assume that each $\mathcal{T}_i$ is supported at a fixed divisor on the Fargues–Fontaine curve, which will be useful later. We write $U^\perp$, $V^\perp$, and $W^\perp$ respectively for the analogous spaces with the roles of $\mathcal{E}_1$ and $\mathcal{E}_2$ exchanged, replacing $\widetilde{Q}_1$ by $\widetilde{Q}_2$. The short exact sequences (3.15) induce short exact sequences of stacks in F-vector spaces

$$0 \to U \to V \to W \to 0 \quad \text{and} \quad 0 \to U^\perp \to V^\perp \to W^\perp \to 0. \tag{3.17}$$

Proposition 3.5.2. *As stacks in F-vector spaces over Bun_H, we have*

$$\widehat{U} \simeq W^\perp, \qquad \widehat{V} \simeq V^\perp, \qquad \textit{and} \qquad \widehat{W} \simeq U^\perp,$$

so that Fourier duality exchanges the short exact sequences of (3.17)*.*

Proof. This essentially follows from [24, Corollary 3.10, Example 3.16]. We work out the first case; the others are similar. We have $\widehat{U} = \mathcal{BC}(\widehat{\mathcal{E}_1^* \otimes \mathcal{F}}) \simeq \mathcal{BC}(\mathcal{E}_1 \otimes \mathcal{F}^*[1])$ (bearing in mind the difference in convention), which is the same thing as $\mathbb{E}\mathrm{xt}^1(\mathcal{F}, \mathcal{E}_1)$. This is close to the definition of $W^\perp$, but we would want to replace $\mathcal{F}$ by $\sigma^*\mathcal{F}^*$; however $\mathcal{F}$ comes equipped with an isomorphism $h_{\mathcal{F}} : \mathcal{F} \simeq \sigma^*\mathcal{F}^*$ given by its Hermitian structure, and so we can identify $\widehat{U}$ with $W^\perp$. □

[13] This is the essential reason to allow a to vary; otherwise we cannot twist the $\mathcal{E}_i$ as we need to to allow these assumptions.

Following [19, §9.1.3], given an S-point $\mathcal{F}_0 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_r$ of Hk_H^r, let $\mathcal{F}_\bullet^\flat$ be the perfect complex on X'_S in degrees 0 and 1

$$\bigoplus_{i=1}^{r} \mathcal{F}_{i-1/2}^\flat \to \bigoplus_{i=1}^{r-1} \mathcal{F}_r,$$

with the differential given by

$$(s_{1/2}, \dots, s_{r-1/2}) \mapsto (s_{1/2} - s_{3/2}, \dots, s_{r-3/2} - s_{r-1/2}),$$

and let $\mathcal{F}_\bullet^{\flat *}$ be its linear dual, the cone of the dual map

$$\bigoplus_{i=1}^{r} \mathcal{F}_r^* \to \bigoplus_{i=1}^{r} \mathcal{F}_{i-1/2}^{\flat *}.$$

For each $1 \le i \le r-1$, either projecting to $\mathcal{F}_{i-1/2}^\flat$ and applying the inclusion $\mathcal{F}_{i-1/2}^\flat \hookrightarrow \mathcal{F}_i$ or projecting to $\mathcal{F}_{i+1/2}^\flat$ and applying $\mathcal{F}_{i+1/2}^\flat \to \mathcal{F}_i$ gives the same map $\mathcal{F}_\bullet^\flat \to \mathcal{F}_i$; for $i = r$ only the former applies, while for $i = 0$ only the latter, so for all $0 \le i \le r$ we get a well-defined map $\mathcal{F}_\bullet^\flat \to \mathcal{F}_i$, or dually $\mathcal{F}_i^* \to \mathcal{F}_\bullet^{\flat *}$. We define $\mathcal{U}_{\mathrm{Hk}}^\flat$, $\mathcal{V}_{\mathrm{Hk}}^\flat$, and $\mathcal{W}_{\mathrm{Hk}}^\flat$ to be the sheaves over Hk_H^r of complexes of F-modules whose pullback along an S-point $\mathcal{F}_0 \dashrightarrow \cdots \dashrightarrow \mathcal{F}_r$ is $\mathbb{H}\mathrm{om}(\mathcal{F}_\bullet^{\flat *}, \mathcal{E}_1^*)$, $\mathbb{H}\mathrm{om}(\mathcal{F}_\bullet^{\flat *}, \widetilde{Q}_1)$, and $\mathbb{H}\mathrm{om}(\mathcal{F}_\bullet^{\flat *}, \sigma^* \mathcal{E}_2[1])$ respectively. Exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$ similarly gives complexes $\mathcal{U}_{\mathrm{Hk}}^{\perp,\flat}$, $\mathcal{V}_{\mathrm{Hk}}^{\perp,\flat}$, and $\mathcal{W}_{\mathrm{Hk}}^{\perp,\flat}$, with the short exact sequences (3.15) inducing

$$0 \to \mathcal{U}_{\mathrm{Hk}}^\flat \to \mathcal{V}_{\mathrm{Hk}}^\flat \to \mathcal{W}_{\mathrm{Hk}}^\flat \to 0 \quad \text{and} \quad 0 \to \mathcal{U}_{\mathrm{Hk}}^{\perp,\flat} \to \mathcal{V}_{\mathrm{Hk}}^{\perp,\flat} \to \mathcal{W}_{\mathrm{Hk}}^{\perp,\flat} \to 0. \tag{3.18}$$

Writing $\widetilde{\mathcal{U}}_i$ for the pullback of $\mathcal{U} \to \mathrm{Bun}_H$ along $p_i : \mathrm{Hk}_H^r \to \mathrm{Bun}_H$ and similarly for $\mathcal{U}$, $\mathcal{V}$, $\mathcal{W}$, $\mathcal{U}^\perp$, $\mathcal{V}^\perp$, and $\mathcal{W}^\perp$, for each $? \in \{\mathcal{U}, \mathcal{V}, \mathcal{W}, \mathcal{U}^\perp, \mathcal{V}^\perp, \mathcal{W}^\perp\}$ we let $?_{\mathrm{Hk}}^\sharp$ be the sheaf on Hk_H^r of

complexes of F-vector spaces fitting into the Cartesian and co-Cartesian diagram

$$\begin{array}{ccc} ?^{\flat}_{\mathrm{Hk}} & \longrightarrow & \widetilde{?}_0 \\ \downarrow & & \downarrow \\ \widetilde{?}_r & \longrightarrow & ?^{\sharp}_{\mathrm{Hk}} \end{array}.$$

Proposition 3.5.3. *There is a distinguished triangle*

$$\mathcal{K}_i \to \mathcal{W}^{\flat}_{\mathrm{Hk}} \xrightarrow{\widetilde{c}_i} \widetilde{\mathcal{W}}_i$$

with $\mathcal{K}_i$ a perfect complex of F-vector spaces concentrated in degree 0.

In particular, since $\widetilde{\mathcal{W}}_i$ is the pullback of $\mathcal{W}$ it is concentrated in degree 0, so so is $\mathcal{W}^{\flat}_{\mathrm{Hk}}$.

Proof. Let T_i be the cokernel of the map $\mathcal{F}_i^* \to \mathcal{F}_\bullet^{\flat *}$, so that we have a distinguished triangle

$$\mathcal{K}_i := \mathbb{H}\mathrm{om}(T_i, \sigma^*\mathcal{E}_2[1]) \to \mathcal{W}^{\flat}_{\mathrm{Hk}} = \mathbb{H}\mathrm{om}(\mathcal{F}_\bullet^{\flat *}, \sigma^*\mathcal{E}_2[1]) \to \widetilde{\mathcal{W}}_i = \mathbb{H}\mathrm{om}(\mathcal{F}_i^*, \sigma^*\mathcal{E}_2[1]).$$

As T_i is torsion and $\sigma^*\mathcal{E}_2$ locally free, $\underline{R\,\mathrm{Hom}}_{X'_S}(T_i, \sigma^*\mathcal{E}_2[1])$ is a locally free sheaf concentrated in degree 0, so its Banach–Colmez sheaf $\mathbb{H}\mathrm{om}(T_i, \sigma^*\mathcal{E}_2[1]) = \sigma(\underline{R\,\mathrm{Hom}}_{X'_S}(T_i, \sigma^*\mathcal{E}_2[1]))$ is concentrated in degree 0. □

We would like to study the corresponding stacks in F-vector spaces; but a priori, it is not clear that these complexes are concentrated in degrees 0 and -1. However, we make the following restriction: let $(\mathrm{Spd}\,\mathcal{O}_F)^\circ$ be the v-sheaf sending S to the space of untilts $S^\sharp$ of S over $\mathcal{O}_F$ such that the resulting divisor $S^\sharp \to Y_{[0,\infty),S,\mathcal{O}_F}$ is disjoint from the support of either $\widetilde{Q}_i$ (or rather their image under $Y_{[0,\infty),S,\mathcal{O}_{F'}} \to Y_{[0,\infty),S,\mathcal{O}_F}$), and $(\mathrm{Spd}\,\mathcal{O}'_F)^\circ$ the preimage of $(\mathrm{Spd}\,\mathcal{O}_F)^\circ$. Write $\mathrm{Hk}^{r,\circ}_H$ for the preimage of $((\mathrm{Spd}\,\mathcal{O}_{F'})^\circ)^r$ under the map p_X, and write $-^\circ$ for the pullback along $\mathrm{Hk}^{r,\circ}_H \to \mathrm{Hk}^r_H$.

Just as in [19, Lemma 10.1.3], we find that after restriction to $\mathrm{Hk}^{r,\circ}_H$ the maps $\mathcal{V}^{\flat,\circ}_{\mathrm{Hk}} \to \widetilde{\mathcal{V}}^\circ_i \to$

$\mathcal{V}^{\sharp,\circ}_{\mathrm{Hk}}$ are isomorphisms. We therefore have a commutative diagram

$$\begin{array}{ccccc}
\mathcal{U}^{\flat,\circ}_{\mathrm{Hk}} & \xrightarrow{f^\circ} & \mathcal{V}^{\flat,\circ}_{\mathrm{Hk}} & \xrightarrow{g^\circ} & \mathcal{W}^{\flat,\circ}_{\mathrm{Hk}} \\
\downarrow{\scriptstyle \widetilde{a}_i^\circ} & & \downarrow{\scriptstyle \widetilde{b}_i^\circ} & & \downarrow{\scriptstyle \widetilde{c}_i} \\
\widetilde{\mathcal{U}}_i^\circ & \xrightarrow{\widetilde{f}_i} & \widetilde{\mathcal{V}}_i^\circ & \xrightarrow{\widetilde{g}_i} & \widetilde{\mathcal{W}}_i^\circ \\
\downarrow & & \downarrow & & \downarrow \\
\mathrm{coker}(\widetilde{a}_i^\circ) & \longrightarrow & 0 & \longrightarrow & \mathcal{K}_i^\circ[1]
\end{array} \tag{3.19}$$

with all columns and the top two rows distinguished triangles, and so by the (triangulated) nine lemma the bottom row is distinguished as well, i.e. $\mathcal{K}_i^\circ \xrightarrow{\sim} \mathrm{coker}(\widetilde{a}_i^\circ)$. By Proposition 3.5.3, $\mathcal{K}_i^\circ$ is concentrated in degree 0, so so is $\mathrm{coker}(\widetilde{a}_i^\circ)$ as a perfect complex of F-vector spaces. Therefore $\mathcal{U}^{\flat,\circ}_{\mathrm{Hk}}$ is concentrated in degree 0 as well.

In the rest of this note, we always work over the locus $\mathrm{Hk}^{r,\circ}_H$, and so drop the $-^\circ$ from the notation. Ultimately as the support of the $\widetilde{Q}_i$ varies this will lead to the restriction to generic, rather than integral, modularity.

Remark 3.5.4. Interestingly, this observation works in the function field case as well; for generic modularity as in [19], it may be possible to avoid derived Fourier analysis entirely.

For each ? as above we write $\mathrm{Hk}^\flat_?$ and $\mathrm{Hk}^\sharp_?$ for the corresponding stacks in F-vector spaces over Hk^r_H. By an analogous argument to Proposition 3.5.2, we have the following:

Proposition 3.5.5. *As stacks in F-vector spaces over Hk^r_H, we have*

$$\widehat{\mathrm{Hk}}^\flat_U \simeq \mathrm{Hk}^\sharp_{W^\perp}, \qquad \widehat{\mathrm{Hk}}^\flat_V \simeq \mathrm{Hk}^\sharp_{V^\perp}, \qquad \widehat{\mathrm{Hk}}^\flat_W \simeq \mathrm{Hk}^\sharp_{U^\perp}$$

and

$$\widehat{\mathrm{Hk}}^\sharp_U \simeq \mathrm{Hk}^\flat_{W^\perp}, \qquad \widehat{\mathrm{Hk}}^\sharp_V \simeq \mathrm{Hk}^\flat_{V^\perp}, \qquad \widehat{\mathrm{Hk}}^\sharp_W \simeq \mathrm{Hk}^\flat_{U^\perp},$$

so that the diagrams

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{\widetilde{a}_0} & \widetilde{U}_0 \\ \downarrow{\scriptstyle \widetilde{a}_r} & & \downarrow{\scriptstyle \widetilde{a}'_r} \\ \widetilde{U}_r & \xrightarrow{\widetilde{a}'_0} & \mathrm{Hk}_U^\sharp \end{array}, \qquad \begin{array}{ccc} \mathrm{Hk}_V^\flat & \xrightarrow{\widetilde{b}_0} & \widetilde{V}_0 \\ \downarrow{\scriptstyle \widetilde{b}_r} & & \downarrow{\scriptstyle \widetilde{b}'_r} \\ \widetilde{V}_r & \xrightarrow{\widetilde{b}'_0} & \mathrm{Hk}_V^\sharp \end{array}, \qquad \begin{array}{ccc} \mathrm{Hk}_W^\flat & \xrightarrow{\widetilde{c}_0} & \widetilde{W}_0 \\ \downarrow{\scriptstyle \widetilde{c}_r} & & \downarrow{\scriptstyle \widetilde{c}'_r} \\ \widetilde{W}_r & \xrightarrow{\widetilde{c}'_0} & \mathrm{Hk}_W^\sharp \end{array}$$

are dual to

$$\begin{array}{ccc} \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{\widetilde{a}_0^\perp} & \widetilde{U^\perp}_0 \\ \downarrow{\scriptstyle \widetilde{a}_r^\perp} & & \downarrow{\scriptstyle \widetilde{a}_r'^\perp} \\ \widetilde{U}_r^\perp & \xrightarrow{\widetilde{a}_0'^\perp} & \mathrm{Hk}_{U^\perp}^\sharp \end{array}, \qquad \begin{array}{ccc} \mathrm{Hk}_{V^\perp}^\flat & \xrightarrow{\widetilde{b}_0^\perp} & \widetilde{V^\perp}_0 \\ \downarrow{\scriptstyle \widetilde{b}_r^\perp} & & \downarrow{\scriptstyle \widetilde{b}_r'^\perp} \\ \widetilde{V}_r^\perp & \xrightarrow{\widetilde{b}_0'^\perp} & \mathrm{Hk}_{V^\perp}^\sharp \end{array}, \qquad \begin{array}{ccc} \mathrm{Hk}_{W^\perp}^\flat & \xrightarrow{\widetilde{c}_0^\perp} & \widetilde{W^\perp}_0 \\ \downarrow{\scriptstyle \widetilde{c}_r^\perp} & & \downarrow{\scriptstyle \widetilde{c}_r'^\perp} \\ \widetilde{W}_r^\perp & \xrightarrow{\widetilde{c}_0'^\perp} & \mathrm{Hk}_{W^\perp}^\sharp \end{array}$$

respectively.

We next want to define self-correspondences $j^?$ for $? \in \{U, V, W, U^\perp, V^\perp, W^\perp\}$ fitting into maps of correspondences

$$\begin{array}{ccccc} ? & \longleftarrow & j^? & \longrightarrow & ? \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{Bun}_H & \longleftarrow & *\!/K & \longrightarrow & \mathrm{Bun}_H \end{array} .$$

For this, we need to incorporate the integral structures $\widetilde{\mathfrak{E}}_i$ and $\mathfrak{G}$. Let $(\mathfrak{P}, \mathfrak{R}) \in (*/K)(S) \simeq (*/K_1)(S) \times (*/K_b)(S)$ be an S-point of $*/K$ lying over $(\mathcal{F}, \mathcal{F}') \in \mathrm{Bun}_H$, with $\mathcal{F}$ trivial and $\mathcal{F}'$ isomorphic to $\mathcal{F}^b$. Each space $?$ can be understood as parametrizing over $\mathcal{F}$ maps $\mathcal{F}^* \to A$ for some perfect complex A on the Fargues–Fontaine curve, given by $A = \mathcal{E}_1^*$ for $? = U$, $A = \widetilde{Q}_1$ for $? = V$, and $A = \sigma^*\mathcal{E}_2[1]$ for $? = W$, and similarly exchanging $\mathcal{E}_1$ and $\mathcal{E}_2$ for the $-^\perp$ versions. From our discussion above, we have integral versions of all of these: $\mathfrak{E}_1^* \to \mathcal{E}_1^*$, $\widetilde{\mathfrak{Q}}_1 \to \widetilde{Q}_1$, and $\sigma^*\mathfrak{E}_2[1] \to \sigma^*\mathcal{E}_2[1]$.[14] On the other hand, either $\mathfrak{P}$ or $\mathfrak{R}$ could be viewed as an integral replacement for $\mathcal{F} \in \mathrm{Bun}_H$, depending on whether $\mathcal{F}$ is isomorphic to $\mathcal{F}^1$ or $\mathcal{F}^b$. Thus we define $\mathfrak{U}^1$ as the stack over $*/K_1$ sending $\mathfrak{P}$ to $\mathbb{H}\mathrm{om}(\mathfrak{P}^*, \mathfrak{E}_1^*)$ and $\mathfrak{U}^b$ over $*/K_b$ sending $\mathfrak{R}$ to

[14] Note that the $-^*$ operation in e.g. $\mathfrak{E}_1^*$ is linear duality of sheaves of O_F-modules on the Fargues–Fontaine curve, not the Fourier duality discussed above.

$\mathbb{H}\mathrm{om}(\mathfrak{R}^*, \mathfrak{E}_1^*)$; and similarly $\mathfrak{V}^1$, $\mathfrak{V}^b$, $\mathfrak{W}^1$, $\mathfrak{W}^b$ and $-^\perp$ versions, using the integral versions of A and either $\mathfrak{P}$ or $\mathfrak{R}$. We can view for example $\mathfrak{U}^1$ as parametrizing those maps $\mathcal{F}^* \to \mathcal{E}_1^*$ whose restriction to $\mathfrak{P}^*$ lands in $\mathfrak{E}_1^*$, and similarly for the other such stacks. Then we can define $j^U = \mathfrak{U}^1 \times \mathfrak{U}^b \to U \times U$, with the maps given by forgetting the integral structure, and similarly $j^V = \mathfrak{V}^1 \times \mathfrak{V}^b$, $j^W = \mathfrak{W}^1 \times \mathfrak{W}^b$, and $-^\perp$ versions; these give morphisms of correspondences to $j : */K \to \mathrm{Bun}_H \times \mathrm{Bun}_H$ as above. The short exact sequences (3.16) induce short exact sequences

$$0 \to \mathfrak{U}^c \to \mathfrak{V}^c \to \mathfrak{W}^c \to 0, \tag{3.20}$$

$$0 \to \mathfrak{U}^{\perp,c} \to \mathfrak{V}^{\perp,c} \to \mathfrak{W}^{\perp,c} \to 0 \tag{3.21}$$

for $c \in \{1, b\}$, which upon inverting p recover (3.17) base changed along $j_K^c : */K_c \to \mathrm{Bun}_H^c$; for each $? \in \{U, V, W, U^\perp, V^\perp, W^\perp\}$ we denote this base change by $?^c$. By the nine lemma, we obtain short exact sequences

$$0 \to U^c/\mathfrak{U}^c \to V^c/\mathfrak{V}^c \to W^c/\mathfrak{V}^c \tag{3.22}$$

and

$$0 \to U^{\perp,c}/\mathfrak{U}^{\perp,c} \to V^{\perp,c}/\mathfrak{V}^{\perp,c} \to W^{\perp,c}/\mathfrak{V}^{\perp,c}. \tag{3.23}$$

Proposition 3.5.6. *We have*

$$\widehat{\mathfrak{U}}^c \simeq W^{\perp,c}/\mathfrak{W}^{\perp,c}, \qquad \widehat{\mathfrak{V}}^c \simeq V^{\perp,c}/\mathfrak{V}^{\perp,c}, \qquad \textit{and} \qquad \widehat{\mathfrak{W}}^c \simeq U^{\perp,c}/\mathfrak{U}^{\perp,c},$$

so that Fourier duality exchanges (3.20) *with* (3.23)*.*

Proof. This is similar to Proposition 3.5.2; again we work out only the first case with $c = 1$, with the others similar. For A, B O_F-submodule stacks of $A[1/p]$ and $B[1/p]$ respectively, we have $\widehat{\mathbb{H}\mathrm{om}(A, B)} \simeq \mathbb{H}\mathrm{om}(A^* \otimes B, F/O_F[1]) \simeq \mathbb{H}\mathrm{om}(A^*, \widehat{B})$. Applied to the case $A = \mathfrak{P}^*$ and $B = \mathfrak{E}_1^*$ and making use of the Hermitian structure $\mathfrak{P} \simeq \sigma^*\mathfrak{P}^*$, we obtain $\widehat{\mathbb{H}\mathrm{om}(\mathfrak{P}^*, \mathfrak{E}_1^*)} \simeq \mathbb{H}\mathrm{om}(\mathfrak{P}, \mathcal{E}_1/\mathfrak{E}_1[1]) \simeq \mathbb{H}\mathrm{om}(\mathfrak{P}^*, \sigma^*\mathcal{E}_1/\mathfrak{E}_1[1])$, which is the cofiber of $\mathbb{H}\mathrm{om}(\mathfrak{P}^*, \sigma^*\mathfrak{E}_1[1]) \to$

$\mathbb{H}\mathrm{om}(\mathfrak{P}^*, \sigma^*\mathcal{E}_1[1]) \simeq \mathbb{H}\mathrm{om}(\mathcal{F}^*, \sigma^*\mathcal{E}_1[1])$ (since any map $\mathfrak{P}^* \to \sigma^*\mathcal{E}_1[1]$ factors through $\mathcal{F}^* = \mathfrak{P}^*[1/p]$). Thus $\widehat{\mathfrak{U}}^1$ is the cofiber of $\mathfrak{W}^{\perp,1} \to W^{\perp,1}$ as claimed. □

We now take the point of view that the key objects are pairs of correspondences $X \leftarrow C \to Y \leftarrow D \to X$, from which we can form the intersection $C \times_{X\times Y} D = \mathrm{Fix}(D \circ C)$. In particular we have the "base" example

$$\mathrm{Bun}_H \xleftarrow{h_0} \mathrm{Hk}^r_H \xrightarrow{h_r} \mathrm{Bun}_H \xleftarrow{j^\flat_K} */K \xrightarrow{j^1_K} \mathrm{Bun}_H, \tag{3.24}$$

whose intersection recover $\mathrm{Sht}^r_{H,b,K}$. In addition for each $? \in \{U, V, W, U^\perp, V^\perp, W^\perp\}$ we have the example

$$? \leftarrow \mathrm{Hk}^\flat_? \to ? \leftarrow j^? \to ?,$$

equipped with maps to the base example

$$\begin{array}{ccccccccc} ? & \longleftarrow & \mathrm{Hk}^\flat_? & \longrightarrow & ? & \longleftarrow & j^? & \longrightarrow & ? \\ \downarrow{\scriptstyle p_{?0}} & & \downarrow{\scriptstyle p_?} & & \downarrow{\scriptstyle p_{?r}} & & \downarrow{\scriptstyle q_?} & & \downarrow{\scriptstyle p_{?0}} \\ \mathrm{Bun}_H & \longleftarrow & \mathrm{Hk}^r_H & \longrightarrow & \mathrm{Bun}_H & \longleftarrow & */K & \longrightarrow & \mathrm{Bun}_H \end{array}.$$

We sometimes write $\pi = p_U$ by analogy with the notation of [19]. We write $\mathrm{Sht}(p_?) : \mathrm{Sht}^r_? := \mathrm{Hk}^\flat_? \times_{\mathrm{Bun}_H \times \mathrm{Bun}_H} j^? \to \mathrm{Sht}^r_{H,b,K}$ for the induced map on fixed points.

By Proposition 3.5.6, the $j^?$ fit into the formalism of the $\mathcal{O}_F$-linear Fourier transform: e.g. the Cartesian diagrams

$$\begin{array}{ccc} j^U & \longrightarrow & U^{1,b}_1 \\ \downarrow & & \downarrow \\ U^{1,b}_b & \longrightarrow & U^1/\mathfrak{U}^1 \times U^b/\mathfrak{U}^b \end{array}, \qquad \begin{array}{ccc} j^{W^\perp} & \longrightarrow & W^{\perp,1,b}_1 \\ \downarrow & & \downarrow \\ W^{\perp,1,b}_b & \longrightarrow & W^{\perp,1}/\mathfrak{W}^{\perp,1} \times W^{\perp,b}/\mathfrak{W}^{\perp,b} \end{array}$$

are Fourier dual, where $?^{1,b}_c$ denotes the base change of U to $*/K \to */K_c \to \mathrm{Bun}_H$, and so the Fourier transform relates cohomological correspondences supported on j^U and on $j^{W^\perp}$.

The short exact sequences (3.17) can be rephrased as Cartesian (and co-Cartesian) diagrams

$$\begin{array}{ccc} U & \longrightarrow & V \\ \downarrow & & \downarrow \\ \mathrm{Bun}_H & \longrightarrow & W \end{array}, \qquad \begin{array}{ccc} U^\perp & \longrightarrow & V^\perp \\ \downarrow & & \downarrow \\ \mathrm{Bun}_H & \longrightarrow & W^\perp \end{array}$$

where the bottom map is the zero section. We similarly have versions of these diagrams for Hecke stacks

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{f} & \mathrm{Hk}_V^\flat \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle g} \\ \mathrm{Bun}_H & \xrightarrow{z} & \mathrm{Hk}_W^\flat \end{array}, \qquad \begin{array}{ccc} \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{f^\perp} & \mathrm{Hk}_{V^\perp}^\flat \\ \downarrow{\scriptstyle \pi^\perp} & & \downarrow{\scriptstyle g^\perp} \\ \mathrm{Hk}_H^r & \xrightarrow{z^\perp} & \mathrm{Hk}_{W^\perp}^\flat \end{array},$$

which we think of as Cartesian squares of correspondences. Likewise for the correspondences $j^?$ we obtain from (3.20), (3.21) Cartesian squares of correspondences

$$\begin{array}{ccc} j^U & \xrightarrow{\phi} & j^V \\ \downarrow{\scriptstyle q_U} & & \downarrow{\scriptstyle \gamma} \\ */K & \xrightarrow{\zeta} & j^W \end{array}, \qquad \begin{array}{ccc} j^{U^\perp} & \xrightarrow{\phi^\perp} & j^{V^\perp} \\ \downarrow{\scriptstyle q_{U^\perp}} & & \downarrow{\scriptstyle \gamma^\perp} \\ */K & \xrightarrow{\zeta^\perp} & j^{W^\perp} \end{array}.$$

We collect some geometric properties of the objects we have defined which will be useful below. We write $a_i = h_i \circ \widetilde{a}_i : \mathrm{Hk}_U^\flat \to \widetilde{U}_i$, and similarly for b_i, c_i, and the $-^\perp$-versions, and bear in mind that we are restricting to $\mathrm{Hk}_H^{r,\circ}$ though it is omitted from the notation.

Proposition 3.5.7. *(a) For each $0 \le i \le r$, $\widetilde{a}_i$ and $\widetilde{a}_i^\perp$ are cohomologically quasi-smooth closed immersions, b_i and $b_i^\perp$ are isomorphisms, and c_i and $c_i^\perp$ are smooth and separated.*

(b) For each $0 \le i \le r$, the squares

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{a_i} & U \\ \downarrow{\scriptstyle \pi} & & \downarrow \\ \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \end{array} \quad \textit{and} \quad \begin{array}{ccc} \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{a_i^\perp} & U \\ \downarrow{\scriptstyle \pi^\perp} & & \downarrow \\ \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \end{array}$$

are pushable and cohomologically pullable.

(c) *The structure maps* $U \to \mathrm{Bun}_H$ *and* $U^\perp \to \mathrm{Bun}_H$ *are cohomologically smooth and prim, and* $q_U : \mathfrak{U}^1 \times \mathfrak{U}^b \to */K$ *and* $q_{U^\perp} : \mathfrak{U}^{\perp,1} \times \mathfrak{U}^{\perp,b} \to */K$ *are cohomologically quasi-smooth and proper.*

(d) *For each* $i \in \{0, r\}$*, the squares*

$$\begin{array}{ccc} j^U & \longrightarrow & U \\ \downarrow{\scriptstyle q_U} & & \downarrow{\scriptstyle p_{U,i}} \\ */K & \longrightarrow & \mathrm{Bun}_H \end{array} \quad \text{and} \quad \begin{array}{ccc} j^{U^\perp} & \longrightarrow & U^\perp \\ \downarrow{\scriptstyle q_{U^\perp}} & & \downarrow{\scriptstyle p_{U^\perp,i}} \\ */K & \longrightarrow & \mathrm{Bun}_H \end{array}$$

are pushable and cohomologically pullable.

(e) *For* $0 \le i \le r$*, the squares*

$$\begin{array}{ccc} \mathrm{Hk}_V^\flat & \xrightarrow{b_i} & V \\ \downarrow{\scriptstyle \rho} & & \downarrow{\scriptstyle \rho_i} \\ \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathrm{Hk}_{V^\perp}^\flat & \xrightarrow{b_i^\perp} & V^\perp \\ \downarrow{\scriptstyle \rho^\perp} & & \downarrow{\scriptstyle \rho_i^\perp} \\ \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \end{array}$$

are pushable.

(f) *For* $c \in \{1, b\}$*, the squares*

$$\begin{array}{ccc} j^V & \longrightarrow & V \\ \downarrow{\scriptstyle r} & & \downarrow \\ */K & \xrightarrow{j_K^c} & \mathrm{Bun}_H \end{array} \quad \text{and} \quad \begin{array}{ccc} j^{V^\perp} & \longrightarrow & V^\perp \\ \downarrow{\scriptstyle r^\perp} & & \downarrow \\ */K & \xrightarrow{j_K^c} & \mathrm{Bun}_H \end{array}$$

are pushable.

(g) *For each* $0 \le i \le r$*, the squares*

$$\begin{array}{ccc} \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle z_V} & & \downarrow{\scriptstyle z_{V,i}} \\ \mathrm{Hk}_V^\flat & \xrightarrow{b_i} & V \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle z_{V^\perp}} & & \downarrow{\scriptstyle z_{V^\perp,i}} \\ \mathrm{Hk}_{V^\perp}^\flat & \xrightarrow{b_i^\perp} & V^\perp \end{array}$$

are cohomologically pullable.

(h) For $c \in \{1, b\}$, the squares

$$\begin{array}{ccc} */K & \xrightarrow{j_K^c} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle \zeta} & & \downarrow \\ j^V & \longrightarrow & V \end{array} \qquad \textit{and} \qquad \begin{array}{ccc} */K & \xrightarrow{j_K^c} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle \zeta^\perp} & & \downarrow \\ j^{V^\perp} & \longrightarrow & V^\perp \end{array}$$

are cohomologically pullable.

(i) For $0 \leq i \leq r$, the squares

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{a_i} & U \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_i} \\ \mathrm{Hk}_V^\flat & \xrightarrow{b_i} & V \end{array} \qquad \textit{and} \qquad \begin{array}{ccc} \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{a_i^\perp} & U^\perp \\ \downarrow{\scriptstyle f^\perp} & & \downarrow{\scriptstyle f_i^\perp} \\ \mathrm{Hk}_{V^\perp}^\flat & \xrightarrow{b_i^\perp} & V^\perp \end{array}$$

are pushable.

(j) For $0 \leq i \leq r$, the squares

$$\begin{array}{ccc} \mathrm{Hk}_V^\flat & \xrightarrow{b_i} & V \\ \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle g_i} \\ \mathrm{Hk}_W^\flat & \xrightarrow{c_i} & W \end{array} \qquad \textit{and} \qquad \begin{array}{ccc} \mathrm{Hk}_{V^\perp}^\flat & \xrightarrow{b_i^\perp} & V^\perp \\ \downarrow{\scriptstyle g^\perp} & & \downarrow{\scriptstyle g_i^\perp} \\ \mathrm{Hk}_{W^\perp}^\flat & \xrightarrow{c_i^\perp} & W^\perp \end{array}$$

are cohomologically pullable.

(k) For $0 \leq i \leq r$, the squares

$$\begin{array}{ccc} \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle z} & & \downarrow{\scriptstyle z_i} \\ \mathrm{Hk}_W^\flat & \xrightarrow{c_i} & W \end{array} \qquad \textit{and} \qquad \begin{array}{ccc} \mathrm{Hk}_H^r & \xrightarrow{h_i} & \mathrm{Bun}_H \\ \downarrow{\scriptstyle z^\perp} & & \downarrow{\scriptstyle z_i^\perp} \\ \mathrm{Hk}_{W^\perp}^\flat & \xrightarrow{c_i^\perp} & W^\perp \end{array}$$

are pushable.

(l) For $i \in \{0, r\}$, the squares

$$\begin{array}{ccc} j^U & \longrightarrow & U \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle f_i} \\ j^V & \longrightarrow & V \end{array} \quad \text{and} \quad \begin{array}{ccc} j^{U^\perp} & \longrightarrow & U^\perp \\ \downarrow{\scriptstyle \phi^\perp} & & \downarrow{\scriptstyle f_i^\perp} \\ j^{V^\perp} & \longrightarrow & V^\perp \end{array}$$

are pushable.

(m) For $i \in \{0, r\}$, the squares

$$\begin{array}{ccc} j^V & \longrightarrow & V \\ \downarrow{\scriptstyle \gamma} & & \downarrow{\scriptstyle g_i} \\ j^W & \longrightarrow & W \end{array} \quad \text{and} \quad \begin{array}{ccc} j^{V^\perp} & \longrightarrow & V^\perp \\ \downarrow{\scriptstyle \gamma^\perp} & & \downarrow{\scriptstyle g_i^\perp} \\ j^{W^\perp} & \longrightarrow & W^\perp \end{array}$$

are cohomologically pullable.

Proof. (a) All three properties follow from (3.19) and the observation above that all the sheaves involved are concentrated in degree 0: $\widetilde{a}_i$ is an inclusion of smooth stacks in F-vector spaces, hence a cohomologically quasi-smooth closed immersion; $\widetilde{b}_i$ is an isomorphism (keeping in mind the restriction to $\mathrm{Hk}_H^{r,\circ}$); and $\widetilde{c}_i$ presents $\mathrm{Hk}_W^\flat$ as a smooth stack in F-vector spaces over $\widetilde{W}_i$, and its diagonal $\mathrm{Hk}_W^\flat \to \mathrm{Hk}_W^\flat \times_{\widetilde{W}_i} \mathrm{Hk}_W^\flat$ is the pullback of the zero section of $\mathrm{Hk}_W^\flat$ along the subtraction map, hence proper since $\mathrm{Hk}_W^\flat$ is concentrated in degree 0. The $-^\perp$-side is analogous.

(b) It suffices to prove the claim after base changing everything to Hk_H^r, so we replace Bun_H by Hk_H^r and U and $U^\perp$ by $\widetilde{U}_i$ and $\widetilde{U}_i^\perp$ respectively. Then both squares are pushable since $\widetilde{a}_i$ and $\widetilde{a}_i^\perp$ are closed immersions by (a) and $\mathrm{id} : \mathrm{Hk}_H^r \to \mathrm{Hk}_H^r$ is separated, and cohomologically pullable since $\widetilde{a}_i$ and $\widetilde{a}_i^\perp$ are cohomologically quasi-smooth and id is smooth.

(c) We explain the proof for U; the version for $U^\perp$ is the same. For the cohomologically (quasi-)smoothness claims, the first follows from the smoothness of Banach–Colmez spaces of vector bundles and the assumption that $\mathcal{U}$ is concentrated in degree 0, and the second since

for each $c \in \{1, b\}$ the factor $\mathfrak{W}^{\perp,c}$ is an O_F-submodule stack of $W^{\perp,c} := W^{\perp} \times_{\mathrm{Bun}_H} */K_c$ with and therefore for some d we have $\widehat{W^{\perp,c}/\mathfrak{W}}^{\perp,c} \simeq \mathfrak{U}^c$ by Proposition 3.5.6, so for some d we have $q_{U!}\mathbf{1} \simeq \mathbf{1}\langle d\rangle$ from (3.12). By adjunction we obtain a map $\mathbf{1} \to q_U^!\mathbf{1}\langle d\rangle$ for some d, i.e. q_U is cohomologically quasi-smooth.

The primness of $U \to \mathrm{Bun}_H$ is automatic since it is a Banach–Colmez space, as mentioned in Section 3.3.1. For the properness of q_U, by [42, Lemma 4.6.3] it suffices to show the claim after pullback along the cover $* \to */K$. The result is an O_F-submodule stack of a stack in F-vector spaces concentrated in degree 0, hence an O_F-local system which on the point is trivial, and so is proper by Lemma 4.8.2 (ii) in *loc. cit*.

(d) Cohomological pullability follows immediately from (c). For pushability, it suffices to show the claim after base change to $*/K$, so that the bottom maps become isomorphisms and the top maps are the product of the projections from $\mathfrak{U}^c$ to a point for $c \in \{1, b\}$, which is proper by a similar argument to that in (c), and the inclusion $\mathfrak{U}^{c'} \to U^{c'}$, which is a closed immersion hence proper, and similarly for the $-^{\perp}$ versions.

(e) Again base changing to Hk_H^r, this follows from the fact that $\widetilde{b}_i$ and $\widetilde{b}_i^{\perp}$ are isomorphisms, hence proper, and id is separated.

(f) By a similar argument to in (c), r is proper.

(g) This follows from the fact that $\widetilde{b}_i$ and $\widetilde{b}_i^{\perp}$ are isomorphisms and id is cohomologically quasi-smooth.

(h) Say $c = 1$; then the pullback $j^V \times_V \mathrm{Bun}_H$ is $\mathfrak{W}^b \times */K_1$, with the map $*/K$ given by the base change of the zero section of $\mathfrak{W}^b$. As in (a) this is a cohomologically quasi-smooth closed immersion, so in particular the square is pullable. The same holds exchanging 1 and b for $c = b$, and similarly for the $-^{\perp}$-square.

(i) This follows from (a), since $\widetilde{a}_i$ and $\widetilde{a}_i^{\perp}$ are proper and $\widetilde{b}_i$ and $\widetilde{b}_i^{\perp}$ are isomorphisms, hence separated.

(j) This follows from (a), since $\widetilde{b}_i$, $\widetilde{c}_i$, $\widetilde{b}_i^\perp$, and $\widetilde{c}_i^\perp$ are all smooth.

(k) This follows since after base change to Hk_H^r, id is proper and $\widetilde{c}_i$ and $\widetilde{c}_i^\perp$ are separated.

(l) By a similar argument to in (a), f_i, $f_i^\perp$, ϕ, and $\phi^\perp$ are closed immersions, hence proper.

(m) Since g_i and $g_i^\perp$ are surjections of smooth stacks of F-vector spaces in degree 0, they are themselves smooth stacks in F-vector spaces; so it suffices to see that γ and $\gamma^\perp$ are cohomologically quasi-smooth. Just as for $\widetilde{c}_i$ and $\widetilde{c}_i^\perp$, $j^W \to */K$ and $j^{W^\perp} \to */K$ are cohomologically smooth, and so the cohomological quasi-smoothness of r and $r^\perp$ from (f) yield that of γ and $\gamma^\perp$.

□

Remark 3.5.8. In fact, the proof of Proposition 3.5.7 (b) implies that the relative cohomological dimensions of π and $\pi^\perp$ are given by $d(a_r)$ and $d(a_r^\perp)$ respectively, which by (3.19) and its $-^\perp$ analogue are given by the negatives of the dimensions of the stack-theoretic realizations of $\mathcal{K}_r = \mathbb{H}\mathrm{om}(T_r, \sigma^*\mathcal{E}_2[1])$ and $\mathcal{K}_r^\perp := \mathbb{H}\mathrm{om}(T_r, \sigma^*\mathcal{E}_1[1])$ respectively. Since T_r is torsion and, by the restriction to $(\mathrm{Spd}\, \mathcal{O}_F)^\circ$, we can assume it is supported away from $p = 0$, by [24, Corollary 3.10, §4.2] if $\mathcal{E}_2 \simeq \mathcal{O}^{\oplus m}$ is trivial the associated Banach–Colmez space for $\mathcal{K}_r$ is an extension of affine spaces over untilts, up to diamondization, and since T_r has degree r the resulting space has dimension mr, and similarly for $\mathcal{K}_r^\perp$. In general, we allow the $\mathcal{E}_i$ to vary over different isomorphism classes, corresponding to letting $a \in B(G)_{\mathrm{basic}}$ vary; but ultimately we will restrict to a trivial, after passing to $\widetilde{\mathcal{E}}_i$, and so the analogous maps $\widetilde{\pi} : \mathrm{Hk}_{\widetilde{U}}^\flat \to \mathrm{Hk}_H^r$ and $\widetilde{\pi}^\perp : \mathrm{Hk}_{\widetilde{U}^\perp}^\flat \to \mathrm{Hk}_H^r$ will both have relative cohomological dimensions $-mr$.

Let $\mathfrak{s} : h_0^*\mathbb{Q}_\ell \to h_r^!\mathbb{Q}_{\ell\langle -d(h_r)\rangle}$ be the cohomological correspondence supported on Hk_H^r induced by the smoothness of h_r. As $p_U = \pi$ is cohomologically right pullable by Proposition 3.5.7 (b), we can pull back $\mathfrak{s}$ to $\mathfrak{c}_U := \pi^*\mathfrak{s}$, supported on $\mathrm{Hk}_U^\flat$.

As in Remark 3.4.3, choosing a cohomological self-correspondence of the unit sheaf supported on j is equivalent to choosing a Haar measure on $H(F)/K_1$; make such a choice, and denote the

resulting cohomological correspondence by μ. Then $q_U : j^U \to */K$ is likewise cohomologically right pullable by Proposition 3.5.7 (d), and we can pull back μ to $\mu_U := q_U^* \mu$.

We note that the base change theorem [19, Theorem 5.1.3] and its proof apply without modification in our setting, other than replacing quasi-smoothness and pullability by their cohomological variants.

Similarly define $\mathfrak{c}_{U^\perp}$ and $\mu_{U^\perp}$. We have the following Fourier duality property:

Theorem 3.5.9. *We have:*

(1)

$$\mathrm{FT}_{\mathrm{Hk}_V^b}(f_! \mathfrak{c}_U) = \mathbb{T}_{(-d(U_0)[-2d(U_0)])} f_!^\perp \mathfrak{c}_{U^\perp}.$$

(2)

$$\mathrm{FT}_{j^V}(\phi_! \mu_U) = \mathbb{T}_{(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))[2(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))]} \phi_!^\perp \mu_{U^\perp}.$$

Proof. (1) Indeed, writing $\mathfrak{c}_U = \pi^* \mathfrak{s}$, by Proposition 3.3.1, Proposition 3.3.4 we have

$$\begin{aligned}
\mathrm{FT}_{\mathrm{Hk}_V^b}(f_! \pi^* \mathfrak{s}) &= g^{\perp *} \mathrm{FT}_{\mathrm{Hk}_U^b}(\pi^* \mathfrak{s}) \\
&= \mathbb{T}_{(-d(\pi_0)[-2d(\pi_0)])} g^{\perp *} z_!^\perp \mathrm{FT}_{\mathrm{Hk}_H^r}(\mathfrak{s}) \\
&= \mathbb{T}_{(-d(U_0))[-2d(U_0)]} f_!^\perp \pi^{\perp *} \mathfrak{s} \\
&= \mathbb{T}_{(-d(U_0))[-2d(U_0)]} f_!^\perp \mathfrak{c}_{U^\perp}
\end{aligned}$$

by base change, using Proposition 3.5.7 and the preceding discussion to verify the conditions of [19, Theorem 5.1.3].

(2) Similarly, using (3.11) and (3.13) (and Proposition 3.5.7 for base change),

$$\begin{aligned}\mathrm{FT}_{j^\vee}(\phi_! q_U^* \mu) &= \gamma^{\perp *}\,\mathrm{FT}_{j^U}(q_U^*\mu)\\ &= \mathbb{T}_{(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))[2(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))]}\gamma^{\perp *}\zeta_!^{\perp}\,\mathrm{FT}_j(\mu)\\ &= \mathbb{T}_{(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))[2(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))]}\gamma^{\perp *}\phi_!^{\perp} q_{U^\perp}^*\mu\\ &= \mathbb{T}_{(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))[2(d(\mathfrak{U}^1)+d(\mathfrak{U}^b)-d(U_0))]}\gamma^{\perp *}\phi_!^{\perp}\mu_{U^\perp}.\end{aligned}$$

□

3.6 Proof of split generic modularity

We first need to formulate a version of the modularity conjecture in our setting, for which we restrict to the simplest case, where $\widetilde{\mathcal{E}}_1 =: \widetilde{\mathcal{E}}$, $\widetilde{\mathcal{E}}_2 = \sigma^*\widetilde{\mathcal{E}}^*$, and $\mathcal{G} = \mathcal{G}^\sharp = \widetilde{\mathcal{E}} \oplus \sigma^*\widetilde{\mathcal{E}}^*$, and similarly on lattices.

By [20, Theorem 5.4.3], using Proposition 3.5.7 (b), (c), (d) (together with the discussion of Section 3.1) to verify the conditions, the intersection $\mathrm{Sht}(\widetilde{\pi}) : \mathrm{Hk}^\flat_{\widetilde{U}} \cap j^{\widetilde{U}} \to \mathrm{Hk}^r_H \cap j = \mathrm{Sht}^r_{H,b,K}$ is cohomologically quasi-smooth, where $\widetilde{U}$ is the same as U but with $\mathcal{E}_1$ replaced by $\widetilde{\mathcal{E}}_1 = \widetilde{\mathcal{E}}$, and similarly for $\mathrm{Hk}^\flat_{\widetilde{U}}$ and $j^{\widetilde{U}}$, and the results from Proposition 3.5.7 adapt straightforwardly.[15] We can therefore define

$$[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}] = [\mathrm{Hk}^\flat_{\widetilde{U}} \cap j^{\widetilde{U}}] := \mathrm{Sht}(\widetilde{\pi})^![\mathrm{Sht}^r_{H,b,K}] \tag{3.25}$$

the (virtual) fundamental class, where $[\mathrm{Sht}^r_{H,b,K}] = \langle \mathfrak{s}, \mu\rangle$ is the fundamental class of the stack of shtukas. By Remark 3.5.8, the cohomological dimension of $\widetilde{\pi}$ is $-mr$. Working similarly to Lemma 2.5.3, replacing the one-dimensional smooth curve X' with $\mathrm{Spd}\, O_F$, which is cohomologically smooth by [44, Corollary IV.2.34] and of cohomological dimension 1 by the proof there (see also [26, Proposition 2.1] for the result after quotienting by Frobenius with explicitly motivic coefficients), the dimension of $h_r : \mathrm{Hk}^r_H \to \mathrm{Bun}_H$ is nr, and so the class $[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}] = [\mathrm{Hk}^\flat_{\widetilde{U}} \cap j^{\widetilde{U}}]$ is an

[15] We apologize for the conflict of notation with $\widetilde{U}_i = U \times_{\mathrm{Bun}_H} \mathrm{Hk}^r_H$; we never write the latter without a subscript, and it does not occur in this section.

element of $H^{\mathrm{BM}}_{2(n-m)r}(\mathrm{Sht}^r_{H,b,K}) = H^{-2(n-m)r}(\mathrm{Sht}^r_{H,b,K}, s^!\mathbb{Q}_\ell(-(n-m)r))$ where $s : \mathrm{Sht}^r_{H,b,K} \to *$ is the structure map.

Conjecture 3.6.1. *For all $r \geq 0$,*

$$[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}] = [\mathcal{Z}^r_{\sigma^*\widetilde{\mathcal{E}}^*,\sigma^*\widetilde{\mathfrak{E}}^*}].$$

Remark 3.6.2. More generally, we would like to define a "refined" theta series for any Lagrangian $\mathfrak{E} \subset \mathfrak{G}$ by decomposing the cycle $[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}]$ into finite pieces $[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}(a)]$ indexed by Hermitian maps $a : \mathcal{E} \to \sigma^*\mathcal{E}^*$, and considering the generating series of these pieces parallel to [17], which we would then expect to be independent of $\mathfrak{E}$; Conjecture 3.6.1 is the special case $\mathfrak{G} \simeq \mathfrak{E} \oplus \sigma^*\mathfrak{E}^*$ of this expectation. In principle there is no difficulty in making such a definition; however without further modification the Fourier-theoretic method used here does not work in this generality, even in the case $r = 0$, due to the lack of a suitable p-adic Gaussian, and it is not entirely clear what the "right" definition should be.

The main result of this chapter is that after restricting to the generic fiber, Conjecture 3.6.1 holds. We make this precise: for each $t \in \mathbb{R}_{\geq 0}$, let $(\mathrm{Spd}\, \mathcal{O}_{F'})_{>t}$ be the subsheaf of $\mathrm{Spd}\, \mathcal{O}_{F'}$ whose S-points are the untilts of S over $\mathcal{O}_{F'}$, equivalently degree 1 divisors $S^\sharp$ on the curve $Y_{[0,\infty),S}$ over $\mathcal{O}_{F'}$, such that $S^\sharp \to Y_{[0,\infty),S}$ factors through $Y_{(t,\infty),S}$; so for example $(\mathrm{Spd}\, \mathcal{O}_{F'})_{>0} = \mathrm{Spd}\, F'$. We fix some $t > 0$, and abbreviate $\eta = (\mathrm{Spd}\, \mathcal{O}_{F'})_{>t} \to \mathrm{Spd}\, \mathcal{O}_{F'}$. In addition, since we restrict to $S \to \mathrm{Spa}\, C$, we can study only the base change of our cycles to C; so we let $[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}},\eta,C}]$ be the base change of $[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}}}]$ to $\mathrm{Sht}_{H,b,K,\eta,C} := \mathrm{Sht}^r_{H,b,K} \times_{(\mathrm{Spd}\, \mathcal{O}_{F'})^r} \eta^r \times \mathrm{Spa}\, C$.

Theorem 3.6.3. *For all $r \geq 0$,*

$$[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}},\eta,C}] = \mathcal{Z}^r_{\sigma^*\widetilde{\mathcal{E}}^*,\sigma^*\widetilde{\mathfrak{E}}^*,\eta,C}]$$

in $H^{\mathrm{BM}}_{2(n-m)r}(\mathrm{Sht}_{H,b,K,\eta,C})$.

As observed above, we can fix the support of the $\mathcal{T}_i$; we choose them to have support on the

locus $(\operatorname{Spd} \mathcal{O}_{F'})_{\leq t} = \operatorname{Spd} \mathcal{O}_{F'} - (\operatorname{Spd} \mathcal{O}_{F'})_{>t}$, so that the restriction to $(\operatorname{Spd} \mathcal{O}_{F'})_{>t}$ ensures that we are over $\mathrm{Hk}_H^{r,\circ}$ everywhere.

Under our assumptions, b_{12} and b_{21} are isomorphisms, and so $Q_1 \simeq Q_2 \simeq Q$ are trivial; hence $\widetilde{Q}_1 = \mathcal{T}_1^* \oplus \sigma^* \mathcal{T}_2$ and $\widetilde{Q}_2 = \sigma^* \mathcal{T}_1 \oplus \mathcal{T}_2^*$. This gives a decomposition

$$V \simeq V^{(1)} \times_{\mathrm{Bun}_H} V^{(2)}$$

where $V^{(1)} = \mathbb{H}\mathrm{om}(\mathcal{F}_{\mathrm{univ}}^*, \mathcal{T}_1^*)$ and $V^{(2)} = \mathbb{H}\mathrm{om}(\mathcal{F}_{\mathrm{univ}}^*, \sigma^* \mathcal{T}_2)$, which we can think of as a direct sum of stacks in F-vector spaces, and similarly on Hecke stacks. The short exact sequence

$$0 \to \widetilde{\mathcal{E}}^* \to \mathcal{E}_1^* \to \mathcal{T}_1^* \to 0$$

induces a Cartesian diagram

$$\begin{array}{ccc} \widetilde{U} & \longrightarrow & U \\ \downarrow & & \downarrow \\ \mathrm{Bun}_H & \longrightarrow & V^{(1)} \end{array} .$$

Thus we have a commutative diagram

$$\begin{array}{ccc} \mathrm{Hk}_{\widetilde{U}}^{\flat} & \xrightarrow{i} & \mathrm{Hk}_U^{\flat} \\ \downarrow{\scriptstyle \widetilde{f}} & & \downarrow{\scriptstyle f} \\ \mathrm{Hk}_{V^{(2)}}^{\flat} & \xrightarrow{\widetilde{z}^{(1)}} & \mathrm{Hk}_V^{\flat} \\ \downarrow{\scriptstyle \rho^{(2)}} & & \downarrow{\scriptstyle \mathrm{id} \times \rho^{(2)}} \\ \mathrm{Hk}_H^r & \xrightarrow{z^{(1)}} & \mathrm{Hk}_{V^{(1)}}^{\flat} \end{array} \tag{3.26}$$

with all squares Cartesian. Using Proposition 3.5.7 and observing that (e) and (g) apply after replacing V with $V^{(i)}$, we observe that each square satisfies the hypotheses for base change: we are left only to verify the cohomological pullability of i. We do this by studying the dual: the map $\mathrm{Hk}_{\widetilde{U}}^{\flat} \to \mathrm{Hk}_U^{\flat} \times_U \widetilde{U}$ is cohomologically quasi-smooth if the dual map $\mathrm{Hk}_{W^\perp}^{\sharp} \sqcup_{W^\perp} \widetilde{W}^\perp \to \mathrm{Hk}_{\widetilde{W}^\perp}^{\sharp}$ is

(cohomologically) proper, or equivalently if $\mathrm{Hk}^{\flat}_{\widetilde{W}^{\perp}} \to \mathrm{Hk}^{\flat}_{W^{\perp}} \times_{W^{\perp}} \widetilde{W}^{\perp}$ is, i.e. if the square

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{\widetilde{W}^{\perp}} & \longrightarrow & \widetilde{W}^{\perp} \\ \downarrow & & \downarrow \\ \mathrm{Hk}^{\flat}_{W^{\perp}} & \longrightarrow & W^{\perp} \end{array}$$

is pushable. But this is immediate since $W^{\perp} \to \widetilde{W}^{\perp}$ is a closed immersion of stacks in F-vector spaces concentrated in degree 0, induced by $\mathcal{E} \to \widetilde{\mathcal{E}}$.

Dually, we have a decomposition

$$\widehat{V} \simeq \widehat{V}^{(1)} \times_{\mathrm{Bun}_H} \widehat{V}^{(2)}$$

where $\widehat{V}^{(1)} = \mathbb{H}\mathrm{om}(\mathcal{F}^*_{\mathrm{univ}}, \sigma^*\mathcal{T}_1)$ and $\widehat{V}^{(2)} = \mathbb{H}\mathrm{om}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{T}_2^*)$. The short exact sequence

$$0 \to \widetilde{\mathcal{E}}^* \to \mathcal{E}_2^* \to \mathcal{T}_2^* \to 0$$

similarly yields a Cartesian diagram

$$\begin{array}{ccc} \widetilde{U}^{\perp} & \longrightarrow & U^{\perp} \\ \downarrow & & \downarrow \\ \mathrm{Bun}_H & \longrightarrow & \widehat{V}^{(2)} \end{array}$$

where $\widetilde{U}^{\perp}$ is $U^{\perp}$ with $\mathcal{E}_2$ replaced by $\widetilde{\mathcal{E}}_2 = \sigma^*\widetilde{\mathcal{E}}^*$. The analogous construction on Hecke stacks yields a commutative diagram

$$\begin{array}{ccc} \mathrm{Hk}^{\flat}_{\widetilde{U}^{\perp}} & \xrightarrow{i^{\perp}} & \mathrm{Hk}^{\flat}_{U^{\perp}} \\ \downarrow{\scriptstyle \widetilde{f}^{\perp}} & & \downarrow{\scriptstyle f^{\perp}} \\ \mathrm{Hk}^{\flat}_{\widehat{V}^{(1)}} & \xrightarrow{\widetilde{z}^{\perp(2)}} & \mathrm{Hk}^{\flat}_{\widehat{V}} \\ \downarrow{\scriptstyle \widehat{\rho}^{(1)}} & & \downarrow{\scriptstyle \mathrm{id}\times\widehat{\rho}^{(1)}} \\ \mathrm{Hk}^{r}_{H} & \xrightarrow{z^{\perp(2)}} & \mathrm{Hk}^{\flat}_{V^{(2)}} \end{array} \tag{3.27}$$

with both squares Cartesian and the bottom square dual to that of (3.26). Again, observing that Proposition 3.5.7 (f), (h) apply after replacing V with $V^{(i)}$, we have base change for cohomological correspondences along each square, and $i^\perp$ is cohomologically pullable by a similar argument.

Let $\widetilde{\pi} : \mathrm{Hk}^\flat_{\widetilde{U}} \to \mathrm{Hk}^r_H$ and $\widetilde{\pi}^\perp : \mathrm{Hk}^\flat_{\widetilde{U}^\perp}$ be the structure maps. Then, freely using the pushability and pullability results from Proposition 3.5.7, by Theorem 3.5.9 we have

$$\begin{aligned}
\widetilde{\pi}_! \mathfrak{c}_{\widetilde{U}} &= \widetilde{\pi}_! i^* \mathfrak{c}_U \\
&= \rho^{(2)}_! \widetilde{f}_! i^* \mathfrak{c}_U \\
&= \rho^{(2)}_! \widetilde{z}^{(1)*} f_! \mathfrak{c}_U \\
&= \mathrm{FT}_{\mathrm{Hk}^r_H}(\rho^{(2)}_! \widetilde{z}^{(1)*} f_! \mathfrak{c}_U) \\
&= \mathbb{T}_{(d(\rho_0^{(1)}))[2d(\rho_0^{(1)})]} z^{\perp(2)*} (\mathrm{id} \times \widehat{\rho}^{(1)})_! \, \mathrm{FT}(f_! \mathfrak{c}_U) \\
&= \mathbb{T}_{(d(\rho_0^{(1)})-d(U_0))[2(d(\rho_0^{(1)}-d(U_0))-d(U_0))]} z^{\perp(2)*} (\mathrm{id} \times \widehat{\rho}^{(1)})_! f^\perp_! \mathfrak{c}_{U^\perp} \\
&= \mathbb{T}_{(d(\rho_0^{(1)})-d(U_0))[2(d(\rho_0^{(1)})-d(U_0))]} \widehat{\rho}^{(1)}_! \widetilde{f}^\perp_! i^{\perp *} \mathfrak{c}_{U^\perp} \\
&= \mathbb{T}_{(d(\rho_0^{(1)})-d(U_0))[2(d(\rho_0^{(1)})-d(U_0))]} \widetilde{\pi}^\perp_! i^{\perp *} \mathfrak{c}_{U^\perp} \\
&= \mathbb{T}_{(d(\rho_0^{(1)})-d(U_0))[2(d(\rho_0^{(1)})-d(U_0))]} \widetilde{\pi}^\perp_! \mathfrak{c}_{\widetilde{U}^\perp}.
\end{aligned} \tag{3.28}$$

Working analogously throughout with our integral versions, we obtain a pair of dual commutative diagrams

$$\begin{array}{ccc}
j^{\widetilde{U}} & \xrightarrow{\iota} & j^U \\
\downarrow{\scriptstyle \widetilde{\phi}} & & \downarrow{\scriptstyle \phi} \\
j^{V^{(2)}} & \xrightarrow{\widetilde{\zeta}^{(1)}} & j^V \\
\downarrow{\scriptstyle r^{(2)}} & & \downarrow{\scriptstyle \mathrm{id} \times r^{(2)}} \\
*/K & \xrightarrow{\zeta^{(1)}} & j^{V^{(1)}}
\end{array} \quad , \qquad
\begin{array}{ccc}
j^{\widetilde{U}^\perp} & \xrightarrow{\iota^\perp} & j^{U^\perp} \\
\downarrow{\scriptstyle \widetilde{\phi}^\perp} & & \downarrow{\scriptstyle \phi^\perp} \\
j^{\widehat{V}^{(1)}} & \xrightarrow{\widetilde{\zeta}^{\perp(2)}} & j^{\widehat{V}} \\
\downarrow{\scriptstyle \widehat{r}^{(1)}} & & \downarrow{\scriptstyle \mathrm{id} \times \widehat{r}^{(1)}} \\
*/K & \xrightarrow{\zeta^{\perp(2)}} & j^{\widehat{V}^{(2)}}
\end{array} \quad , \tag{3.29}$$

with ι and $\iota^\perp$ cohomologically pullable by a similar argument as for i. As above we have

$$\begin{aligned}
q_{\widetilde{U}!}\mu_{\widetilde{U}} &= q_{\widetilde{U}!}\iota^*\mu_U \\
&= r_!^{(2)}\widetilde{\phi}_!\iota^*\mu_U \\
&= r_!^{(2)}\widetilde{\zeta}^{(1)*}\phi_!\mu_U \\
&= \mathrm{FT}_{*/K}(r_!^{(2)}\widetilde{\zeta}^{(1)*}\phi_!\mu_U) \\
&= \mathbb{T}_{\left\langle -d(j^{V^{(1)}})+d(V^{(1)})\right\rangle}\zeta^{\perp(2)*}(\mathrm{id}\times\widehat{r}^{(1)})_!\,\mathrm{FT}_{j^V}(\phi_!\mu_U) \\
&= \mathbb{T}_{\left\langle -d(j^{V^{(1)}})+d(V^{(1)})+d(j^U)-d(U_0)\right\rangle}\zeta^{\perp(2)*}(\mathrm{id}\times\widehat{r}^{(1)})_!\phi_!^\perp\mu_{U^\perp} \\
&= \mathbb{T}_{\left\langle -d(j^{V^{(1)}})+d(V^{(1)})+d(j^U)-d(U_0)\right\rangle}\widehat{r}_!^{(1)}\widetilde{\phi}_!^\perp\iota^{\perp*}\mu_{U^\perp} \\
&= \mathbb{T}_{\left\langle -d(j^{V^{(1)}})+d(V^{(1)})+d(j^U)-d(U_0)\right\rangle}q_{\widetilde{U}^\perp!}\iota^{\perp*}\mu_{U^\perp} \\
&= \mathbb{T}_{\left\langle -d(j^{V^{(1)}})+d(V^{(1)})+d(j^U)-d(U_0)\right\rangle}q_{\widetilde{U}^\perp!}\mu_{\widetilde{U}^\perp}.
\end{aligned} \tag{3.30}$$

We can now prove Theorem 3.6.3. Combining [20, Theorems 5.4.2 and 5.4.3], using Proposition 3.5.7 (b), (c), (d), with (3.28) and (3.30), we have (applying shifts as needed to align the terms)

$$\begin{aligned}
[\mathcal{Z}^r_{\widetilde{\mathcal{E}},\widetilde{\mathfrak{E}},\eta,C}] &= \mathrm{Sht}(\widetilde{\pi}^\circ)_*\,\mathrm{Sht}(\widetilde{\pi}^\circ)^!\langle\mathfrak{s},\mu\rangle \\
&= \mathrm{Sht}(\widetilde{\pi}^\circ)_*\langle\mathfrak{c}_{\widetilde{U}},\mu_{\widetilde{U}}\rangle \\
&= \langle\widetilde{\pi}_!\mathfrak{c}_{\widetilde{U}}, q_{\widetilde{U}!}\mu_{\widetilde{U}}\rangle \\
&= \langle\widetilde{\pi}_!^\perp\mathfrak{c}_{\widetilde{U}^\perp}, q_{\widetilde{U}^\perp!}\mu_{\widetilde{U}^\perp}\rangle \\
&= [\mathcal{Z}^r_{\sigma^*\widetilde{\mathcal{E}}^*,\sigma^*\widetilde{\mathfrak{E}}^*,\eta,C}].
\end{aligned}$$

References

[1] B. Riemann, "Ueber die Anzahl der Primzahlen unter einer gegebenen Grosse," *Ges. Math. Werke und Wissenschaftlicher Nachlaß*, vol. 2, no. 145-155, p. 2, 1859.

[2] S. Rallis, "Injectivity properties of liftings associated to Weil representations," *Compositio Math.*, vol. 52, no. 2, pp. 139–169, 1984.

[3] J.-S. Li, "Non-vanishing theorems for the cohomology of certain arithmetic quotients," 1992.

[4] S. S. Kudla and J. J. Millson, "Intersection numbers of cycles on locally symmetric spaces and Fourier coefficients of holomorphic modular forms in several complex variables," *Inst. Hautes Études Sci. Publ. Math.*, no. 71, pp. 121–172, 1990.

[5] S. S. Kudla, "Central derivatives of Eisenstein series and height pairings," *Ann. of Math. (2)*, vol. 146, no. 3, pp. 545–646, 1997.

[6] S. S. Kudla, "Special cycles and derivatives of Eisenstein series," in *Heegner points and Rankin L-series*, ser. Math. Sci. Res. Inst. Publ. Vol. 49, Cambridge Univ. Press, Cambridge, 2004, pp. 243–270, ISBN: 0-521-83659-X.

[7] S. Kudla and M. Rapoport, "Special cycles on unitary Shimura varieties I. Unramified local theory," *Invent. Math.*, vol. 184, no. 3, pp. 629–682, 2011.

[8] S. Kudla and M. Rapoport, "Special cycles on unitary Shimura varieties II: Global theory," *J. Reine Angew. Math.*, vol. 697, pp. 91–157, 2014.

[9] Q. Lu and W. Zheng, "Categorical traces and a relative Lefschetz-Verdier formula," *Forum Math. Sigma*, vol. 10, Paper No. e10, 24, 2022.

[10] C. Li and W. Zhang, "On the arithmetic Siegel-Weil formula for GSpin Shimura varieties," *Invent. Math.*, vol. 228, no. 3, pp. 1353–1460, 2022.

[11] C. Li and Y. Liu, "Chow groups and L-derivatives of automorphic motives for unitary groups," *Ann. of Math. (2)*, vol. 194, no. 3, pp. 817–901, 2021.

[12] C. Li and Y. Liu, "Chow groups and L-derivatives of automorphic motives for unitary groups, II.," *Forum Math. Pi*, vol. 10, Paper No. e5, 71, 2022.

[13] W. Zhang, "Modularity of generating functions of special cycles on Shimura varieties," Ph.D. dissertation, Columbia University New York, 2009.

[14] J. H. Bruinier and M. Westerholt-Raum, "Kudla's modularity conjecture and formal Fourier-Jacobi series," *Forum Math. Pi*, vol. 3, e7, 30, 2015.

[15] J. H. Bruinier, B. Howard, S. S. Kudla, M. Rapoport, and T. Yang, "Modularity of generating series of divisors on unitary Shimura varieties," *Astérisque*, no. 421, pp. 7–125, 2020.

[16] D. Ben-Zvi, Y. Sakellaridis, and A. Venkatesh, "Relative Langlands duality," *arXiv preprint arXiv:2409.04677*, 2024.

[17] T. Feng, Z. Yun, and W. Zhang, "Higher theta series for unitary groups over function fields," *arXiv preprint arXiv:2110.07001*, 2021.

[18] T. Feng, Z. Yun, and W. Zhang, "Higher Siegel-Weil formula for unitary groups: The non-singular terms," *Invent. Math.*, vol. 235, no. 2, pp. 569–668, 2024.

[19] T. Feng, Z. Yun, and W. Zhang, "Modularity of higher theta series I: Cohomology of the generic fiber," *arXiv preprint arXiv:2308.10979*, 2023.

[20] T. Feng and A. Khan, "Modularity of higher theta series II: Chow group of the generic fiber," *arXiv preprint arXiv:2403.19711*, 2024.

[21] P. Scholze and J. Weinstein, *Berkeley lectures on p-adic geometry* (Annals of Mathematics Studies). Princeton University Press, Princeton, NJ, 2020, vol. 207, pp. x+250, ISBN: 978-0-691-20209-9; 978-0-691-20208-2; 978-0-691-20215-0.

[22] L. Fargues, "Geometrization of the local Langlands correspondence: an overview," *arXiv preprint arXiv:1602.00999*, 2016.

[23] R. E. Kottwitz, "Isocrystals with additional structure," *Compositio Math.*, vol. 56, no. 2, pp. 201–220, 1985.

[24] J. Anschütz and A.-C. Le Bras, "A Fourier Transform for Banach-Colmez spaces," *arXiv preprint arXiv:2111.11116*, 2021.

[25] J. Anschütz, A.-C. L. Bras, and L. Mann, "A 6-functor formalism for solid quasi-coherent sheaves on the Fargues–Fontaine curve," *arXiv preprint arXiv:2412.20968*, 2024.

[26] P. Scholze, "Berkovich motives," *arXiv preprint arXiv:2412.03382*, 2024.

[27] D. Gaitsgory and D. Nadler, "Spherical varieties and Langlands duality," *Mosc. Math. J.*, vol. 10, no. 1, pp. 65–137, 271, 2010.

[28] Y. Sakellaridis and A. Venkatesh, "Periods and harmonic analysis on spherical varieties," *Astérisque*, no. 396, pp. viii+360, 2017.

[29] Y. Sakellaridis, "Spherical functions on spherical varieties," *Amer. J. Math.*, vol. 135, no. 5, pp. 1291–1381, 2013.

[30] Y. Sakellaridis, "On the unramified spectrum of spherical varieties over p-adic fields," *Compos. Math.*, vol. 144, no. 4, pp. 978–1016, 2008.

[31] H. Jacquet and S. Rallis, "Symplectic periods," *J. Reine Angew. Math.*, vol. 423, pp. 175–197, 1992.

[32] H. Jacquet, E. Lapid, and J. Rogawski, "Periods of automorphic forms," *J. Amer. Math. Soc.*, vol. 12, no. 1, pp. 173–240, 1999.

[33] A. Ichino and T. Ikeda, "On the periods of automorphic forms on special orthogonal groups and the Gross-Prasad conjecture," *Geom. Funct. Anal.*, vol. 19, no. 5, pp. 1378–1425, 2010.

[34] B. H. Gross and D. Prasad, "On the decomposition of a representation of SO_n when restricted to SO_{n-1}," *Canad. J. Math.*, vol. 44, no. 5, pp. 974–1002, 1992.

[35] W. T. Gan, B. H. Gross, and D. Prasad, "Symplectic local root numbers, central critical L values, and restriction problems in the representation theory of classical groups," in 346, Sur les conjectures de Gross et Prasad. I, 2012, pp. 1–109, ISBN: 978-2-85629-348-5.

[36] T. Feng and J. Wang, "Geometric Langlands Duality for Periods," *Geom. Funct. Anal.*, vol. 35, no. 2, pp. 463–541, 2025.

[37] J.-L. Waldspurger, "Sur les valeurs de certaines fonctions L automorphes en leur centre de symétrie," *Compositio Math.*, vol. 54, no. 2, pp. 173–242, 1985.

[38] S. W. Shin, "Geometric reductive dual pairs and a mod p theta correspondence," *preprint*, 2011.

[39] U. Hartl and Y. Xu, "Uniformizing the moduli stacks of global G-shtukas II," *arXiv preprint arXiv:2309.17441*, 2023.

[40] P. Scholze, "Étale cohomology of diamonds," *arXiv preprint arXiv:1709.07343*, 2017.

[41] P. Scholze, "Geometrization of the local Langlands correspondence, motivically," *arXiv preprint arXiv:2501.07944*, 2025.

[42] C. Heyer and L. Mann, "6-Functor Formalisms and Smooth Representations," *arXiv e-prints*, arXiv–2410, 2024.

[43] A. Khan, "Virtual fundamental classes of derived stacks I," *arXiv e-prints*, arXiv–1909, 2019.

[44] L. Fargues and P. Scholze, "Geometrization of the local Langlands correspondence," *arXiv preprint arXiv:2102.13459*, 2021.

[45] D. Hansen, T. Kaletha, and J. Weinstein, "On the Kottwitz conjecture for local shtuka spaces," *Forum Math. Pi*, vol. 10, Paper No. e13, 79, 2022.

[46] J. Ayoub, M. Gallauer, and A. Vezzani, "The six-functor formalism for rigid analytic motives," *Forum Math. Sigma*, vol. 10, Paper No. e61, 182, 2022.

[47] S. D. Li-Huerta, "The plectic conjecture over local fields," *arXiv preprint arXiv:2208.09962*, 2022.