MODULARITY OF HIGHER THETA SERIES III: PROOF OF THE MODULARITY CONJECTURE

TONY FENG, ZHIWEI YUN, AND WEI ZHANG

ABSTRACT. We prove the Modularity Conjecture for higher theta series on moduli stacks of Hermitian shtukas. For general linear shtukas, we establish a more refined phenomenon that we call *supermodularity*. As a key input, we prove the Trace Conjecture for Hitchin stacks of low corank, realizing virtual fundamental classes of special cycles as categorical traces.

CONTENTS

1. Introduction 1
2. Methods 5
3. Notation 5

Part 1. The Trace Conjecture 7
4. Dimension bounds on special cycles in low corank 7
5. Proof of the Trace Conjecture in low corank 17

Part 2. General tools 24
6. Motivic derived Fourier transform 24
7. Cohomological correspondences 30
8. A trace formula for the sheaf-cycle correspondence 36

Part 3. The Modularity Conjecture 38
9. The genetic pattern of modularity 38
10. Proof of the Modularity Conjecture 51

Part 4. Supermodularity for general linear groups 59
11. Supermodularity 59
12. The genetic pattern of supermodularity 62
13. The proof of supermodularity 69
References 76

1. INTRODUCTION

1.1. **The main theorem.** Let $k = \mathbf{F}_q$ be a finite field of characteristic $p > 2$, let X be a smooth, proper, geometrically connected curve over k, and let $\nu\colon X' \to X$ be an étale double cover. In [FYZ25], the authors constructed *higher theta series* attached to this data. More precisely, for every *rank* $n \geq 1$, every *corank* $m \in [1, n]$, and every $r = 0, 1, \ldots$, the higher theta series $\widetilde{Z}_m^{n,r}$ is a certain Fourier series with coefficients in $\mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{U(n)})$, the Chow group of $r(n-m)$-dimensional cycle classes on the moduli stack $\mathrm{Sht}^r_{U(n)}$ of rank n Hermitian shtukas with r legs.

When $r = 0$, $\widetilde{Z}_m^{n,r}$ specializes to the classical theta functions. When $r = 1$, the construction may be viewed as a function field analogue of arithmetic theta series in the sense of Kudla [Kud04]. The adjective "higher" refers to the phenomena established in [FYZ24, FHM25] that for general r, the higher

theta series are connected to the rth derivatives of automorphic L-functions. While the $r = 0$ and $r = 1$ cases of this story have parallels over number fields, the existence of a story for higher derivatives is a distinctive feature of the function field setting.

Classical theta series and arithmetic theta series enjoy more symmetries than are evident in their definition. This is formulated precisely as *modularity* of the theta series: they descend to automorphic functions. For arithmetic theta series, the strongest formulations of modularity are still conjectural in most cases (and even the formulations may depend on conjectural ingredients). The main result of the earlier work [FYZ25] was the formulation of an analogous Modularity Conjecture for higher theta series. A priori, $\widetilde{Z}_m^{n,r}$ is a function on pairs $(\mathcal{G}, \mathcal{E})$ where $\mathcal{G}$ is a rank $2m$ skew-Hermitian bundle on X' and $\mathcal{E} \subset \mathcal{G}$ is a Lagrangian subbundle. The Modularity Conjecture predicts that $\widetilde{Z}_m^{n,r}$ is independent of the choice of $\mathcal{E} \subset \mathcal{G}$, and our main result is a proof of this conjecture.

Theorem 1.1.1. *The Modularity Conjecture* [FYZ25, Conjecture 4.15] *holds true.*

While this result is formulated for unitary groups and "signature $(1, n-1)$", the method of proof is very general and we expect it to work uniformly for all dual reductive pairs; see Zeff's thesis [Zef25] for initial steps toward such a general formalism.

1.2. **Prior work.** We may contextualize the Modularity Conjecture of [FYZ25] in terms of the analogous problem of *modularity of arithmetic theta series*, laid out in Kudla's MSRI article [Kud04]. It is useful to arrange the expectations of [Kud04] into a hierarchy of successively stronger modularity statements, summarized in Table 1.

Tier	**Number fields**	**Function fields**
Generic-fiber cohomology	Kudla–Millson [KM90].	Unitary groups [FYZ23].
Generic-fiber Chow groups	Orthogonal Shimura varieties: Borcherds in codimension one [Bor99] (Yuan–Zhang–Zhang [YZZ09] over totally real fields), in general Zhang [Zha09], Bruinier–Westerholt-Raum [BWR15].	Unitary groups [FK24].
Integral/arithmetic Chow groups	Unitary divisor case [BHK$^+$20]; Orthogonal Shimura varieties [HM25].	Unitary groups (this paper).

TABLE 1. Three tiers of modularity problems for (arithmetic) theta series and their prior progress.

(1) *Modularity in (Betti) cohomology of the generic fiber* was established for arithmetic theta series in the work of Kudla–Millson [KM90]. The work [FYZ23] established modularity of higher theta series in the ℓ-adic cohomology of the generic fiber, for unitary groups. Work in progress of Xu–Zeff will generalize this to more groups.

(2) *Modularity in the Chow group of the generic fiber* is subtler. For orthogonal Shimura varieties, this problem was raised by Kudla in [Kud97, Kud04]. Borcherds proved the divisor case over $\mathbf{Q}$ using Borcherds products [Bor99] and Yuan–Zhang–Zhang [YZZ09] proved the divisor case over any totally real number field; Zhang proved the higher-codimension case conditionally on a convergence assertion [Zha09]; and Bruinier–Westerholt-Raum supplied the convergence, completing the proof of Kudla's conjecture for orthogonal Shimura varieties [BWR15, BR25]. Pollack has given an alternative proof of the automatic convergence theorem [Pol24]; Raum has recently proved the analog in the unitary case over a general totally real number field in [Rau26]. There are also conditional extensions, assuming Beilinson–Bloch type conjectures, in work of Kudla and Maeda [Kud21, Mae21]. In related compactified or integral settings, Chow-valued modularity has been proved for divisors on unitary Shimura varieties by Bruinier–Howard–Kudla–Rapoport–Yang [BHK$^+$20] and for special zero-cycles on toroidal compactifications by Bruinier–Rosu–Zemel [BRZ24]. In the function field setting, the corresponding generic-fiber Chow result for unitary groups was proved in [FK24].

(3) *Modularity in the arithmetic Chow group of an integral model* is the strongest form envisioned in the Kudla program. In the number field setting, it requires not only integral extensions of the special cycles, but also Green currents. Kudla–Rapoport constructed the basic unitary special cycles for nonsingular coefficients [KR11, KR14]; see also the book of Kudla–Rapoport–Yang for low-dimensional cases [KRY06]. The divisor case for unitary Shimura varieties was proved in [BHK+20]. For integral models of orthogonal Shimura varieties, Howard–Madapusi proved Chow-valued modularity, though not Arakelov Chow-valued modularity [HM25]. In the higher codimension case, a detailed construction of the generating series valued in the arithmetic Chow group of an integral model of unitary Shimura varieties can be found in [Che24a, §3] and [Che24b, §3]; the modularity of these generating series is still conjectural. For function fields, [FYZ25] proposes the formulation of the Modularity Conjecture in this strength, and this paper provides the proof.

Despite the resemblances of the statements, the proofs of modularity (outside the case $r = 0$) are completely different in the number field and function field settings. We refer to the Introductions of [FYZ23] and [FK24] for discussion and comparison. Our proof of Theorem 1.1.1 can be regarded as an improvement on the strategy of [FYZ23, FK24], and will be discussed further in §1.5.

1.3. **Supermodularity for general linear groups.** When the double cover is split, $X' = X \sqcup X$, the unitary groups become general linear groups, and we prove a more refined phenomenon that we call *supermodularity*.

To formulate it, we first recall that [FYZ25] already extended the Modularity Conjecture to general linear groups. In the unitary case, we may view the higher theta series $\widetilde{Z}_m^{n,r}$ as a function on $\mathrm{Bun}_{P_m}(\mathbf{F}_q)$, where P_m is the Siegel parabolic subgroup of a rank $2m$ unitary group. In the general linear case, the Modularity Conjecture admits parabolic refinements as formulated in [FYZ25, §4.9]. Fix a decomposition $m = m_1 + m_2$, and let $P_{(m_1,m_2)} \subset \mathrm{GL}(m)$ be the corresponding standard parabolic (with the convention that $P_{(m,0)} = P_{(0,m)} = \mathrm{GL}(m)$). We can define a higher theta series $\widetilde{Z}^{\mu}_{m_1,m_2}$ (where μ is a sequence of signs ± 1 recording the leg types; see §3.1) on $\mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$, so that it depends a priori on a rank m bundle $\mathcal{G}$, a rank m_1 subbundle $\mathcal{E}_1 \subset \mathcal{G}$, and the resulting extension class. The Modularity Conjecture for general linear groups [FYZ25, Conjecture 4.24] asserts that the value depends on the subbundle $\mathcal{E}_1$ only through $\mathcal{G}$, i.e., that $\widetilde{Z}^{\mu}_{m_1,m_2}$ descends to $\mathrm{Bun}_m(\mathbf{F}_q)$.

Supermodularity expresses an even stronger modularity property of the higher theta series: it is essentially independent *even of the parabolic* $P_{(m_1,m_2)}$. Since in the extreme case $(m_1, m_2) = (m, 0)$ the higher theta series visibly reduces to the special cycle class $[\mathcal{Z}^{\mu}_{\mathcal{G},0}]$ attached to $\mathcal{G}$ alone, supermodularity expresses all the $\widetilde{Z}^{\mu}_{m_1,m_2}$ directly in terms of $[\mathcal{Z}^{\mu}_{\mathcal{G},0}]$. The precise result is formulated in Theorem 11.1.1.

Supermodularity is a phenomenon special to general linear groups, and therefore has no precedent for Shimura varieties over number fields. We nonetheless expect it to have a manifestation over p-adic local fields.

1.4. **Applications.** We list various applications (some potential, some already realized) of Theorem 1.1.1, for which the earlier generic modularity results of [FYZ23, FK24] are not sufficient.

- (*Arithmetic theta lifting*) In Kudla's program, the modularity of (arithmetic) theta series is needed for (arithmetic) theta lifting. In [FHM25], Theorem 1.1.1 is used to construct a higher theta lifting, and to prove an incarnation of a *higher arithmetic inner product formula*. Specifically, [FHM25, Corollary 1.2.1] is conditional on Theorem 1.1.1.
- (*Arithmetic intersection theory*) Modularity of arithmetic theta series was a key input to the results of [SSTT22], which established exceptional jumps of Picard ranks of K3 surfaces over number fields, using arithmetic intersection theory. Ruofan Jiang and Wenqing Wei are pursuing arithmetic intersection problems that rely on Theorem 1.1.1. In order to control the intersection theory, it is crucial to have the full integral modularity.
- (*Local modularity*) Yet another direction of application is towards local problems on special cycles. For this, it is again necessary to have modularity that holds integrally (not just over the generic fiber). Zhang [Zha21] applied modularity of arithmetic theta series to prove the Arithmetic

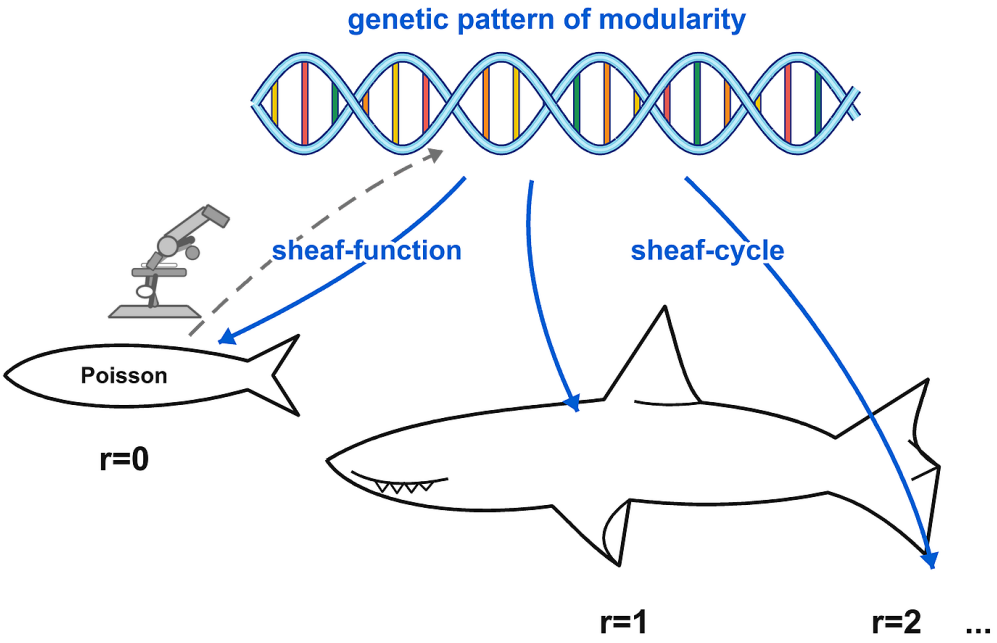

FIGURE 1. Cartoon of the proof of the Modularity Conjecture. By studying the $r = 0$ specimen (Poisson), we sequence the "genetic pattern of modularity". We promote this to a sheaf-theoretic level using (motivic) *derived Fourier analysis*, and then extract the modularity of all higher theta series via the *sheaf-cycle correspondence*.

Fundamental Lemma (AFL), and Luo [Luo25] used it to prove the Kudla–Rapoport Conjecture at places of maximal parahoric level structure.

1.5. **On the proof.** At a high level, the proof is illustrated in Figure 1. The philosophy is to find the "right" proof of modularity for the classical theta series ($r = 0$). Such a proof can be seen as an incarnation of the Poisson summation argument. It should then be promoted to the sheaf-theoretic level via the classical sheaf-function correspondence. We then apply an elaboration of that correspondence, called the *sheaf-cycle correspondence* in [FYZ23], to extract modularity of higher theta series.

Carrying out this strategy requires much foundational technology developed in [FYZ23] and [FK24], including the *derived algebraic geometry* approach to virtual fundamental classes of special cycles, the (motivic) *derived Fourier transform* to promote finite Fourier analysis to the sheaf-theoretic level, and the (motivic) sheaf-cycle correspondence to extract Chow cycles via categorical traces. These are described in more detail in the Introductions of [FYZ23, FK24]. We focus on explaining the novelties of the present paper, which allow us to prove the full Modularity Conjecture.

1.5.1. *The Trace Conjecture.* The *Trace Conjecture* of [FH25, Conjecture 9.4.2] is a fundamental formula that unites three different constructions of virtual fundamental classes for special cycles:

- the classical explicit approach present in [FYZ24, FYZ25] and all classical works in the Kudla program,
- the derived algebraic geometry approach introduced in [FYZ25] (with Shimura variety counterpart appearing in [Mad22]), and
- the (categorical) trace approach introduced in [FYZ23, FK24].

Each of these approaches possesses unique advantages, and all of them are exploited in an essential way here for the proof of the Modularity Conjecture.

A key new input in this paper is a proof of the Trace Conjecture in the low-corank range $m \leq n/3$ (Theorem 5.1.3). We use this to prove the Modularity Conjecture in the same range, and then deduce modularity for all coranks $m \leq n$ by an embedding trick. Both the proof of the Trace Conjecture in low corank and the embedding trick are insensitive to the group-theoretic setup, and should apply with little change to orthogonal and symplectic dual pairs once the relevant statements are formulated.

The previous generic modularity results implicitly used that the Trace Conjecture is easy to prove after restricting away from a point of X', by adding level structure to "resolve" the Hitchin stacks. In fact, at several junctures the earlier works [FYZ23, FK24] use this trick of adding level structure away from the legs to kill geometric difficulties. This trick is not available if one wants to prove the full Modularity Conjecture, where the legs are unrestricted.

Roughly speaking, [FYZ25] gave an elementary computation of the virtual fundamental classes of special cycles, expressing them as a linear combination of terms that are products of a "core" fundamental class and a Chern class (representing excess intersection). Part of the proof of the Trace Conjecture in the low corank range is a new dimension estimate of the "core" part (Theorem 4.1.1), showing that their dimensions do not exceed their "expected dimensions" when $m \leq n/3$. This principle was not anticipated in the analogous theory for Shimura varieties when the corank m exceeds 1, because it concerns a part of the special cycles which have no analogous counterpart in the context of Shimura varieties.

1.5.2. *The trace formula for Hitchin stacks.* Another source of progress comes from an improved understanding of the sheaf-cycle correspondence in the geometric situations appearing in this problem. The sheaf-cycle correspondence can be viewed as a "higher" trace operation producing Chow-valued traces. However, functoriality of the formation of trace requires "nice" geometric properties. This is already visible for the usual sheaf-function correspondence, where, for example, the Grothendieck–Lefschetz trace formula requires properness assumptions in general, which are not satisfied by the Hitchin stacks that govern the special cycles. Yet in §8 we establish a certain "trace formula" which does apply to these Hitchin stacks.

1.5.3. *Fourier duality.* Let us also mention a major departure of the Fourier duality argument in this paper from earlier works. In [FYZ23] and [FK24], it was necessary to add level structure along an uncontrolled finite subset of X. For example, this was necessary in order to define the *arithmetic Fourier transform* in [FYZ23, FK24]. Adding such level structure is useful for "taming" the derived aspect of the geometry, but forces one to restrict to the generic locus. In order to avoid such a restriction, we develop a new Fourier duality approach that interfaces well with derived geometry, and in particular bypasses the arithmetic Fourier transform altogether. A more substantive description of how our approach improves on the one in [FYZ23] is given in §9.2, after the necessary notation has been introduced.

1.6. **Acknowledgments.** T.F. was supported by the NSF (grants DMS-2302520 and DMS-2441922), the Simons Foundation, and the Alfred P. Sloan Foundation. Z.Y. was supported by the Simons Foundation. W.Z. was supported by the NSF (grant DMS-2401548) and the Simons Foundation.

2. Methods

The proof of Theorem 4.1.1 was found with assistance from GPT 5.5 Pro. Specifically, GPT 5.5 Pro was prompted with the statement of the theorem and produced essentially the same line of argument documented here (which in turn closely follows the arguments in [FYZ24, §9] and [FYZ25, §8]). A more detailed technical documentation of this contribution is given in Remark 4.1.2. No other results here were found with AI assistance. The paper was proofread and revised with assistance from various AI tools including Codex, Claude Code, Refine.ink, and Gemini via Google's Paper Assistant Tool (PAT). The figures and tables were produced by Codex.

3. Notation

Notation is as in [FK24] and [FYZ23].

3.1. **Main protagonists.** Let $k = \mathbf{F}_q$ be a finite field of characteristic $p > 2$, let X be a smooth, projective, geometrically connected curve over k of genus g, and let $\nu\colon X' \to X$ be an étale double cover, possibly split, with involution σ.

For a non-negative integer n, we denote by Bun_n the moduli stack of rank n vector bundles on X. For a non-negative integer r and $\mu = (\mu_1, \dots, \mu_r)$ a sequence with $\mu_i \in \{\pm 1\}$ and $\sum \mu_i = 0$, we denote by Sht_n^μ the moduli stack of rank n shtukas on X with simple modifications dictated by μ, defined in [YZ17, §5.1].

Fix a non-trivial additive character $\psi_0\colon \mathbf{F}_p \to \overline{\mathbf{Q}}^\times$, and write

$$\psi := \psi_0 \circ \mathrm{Tr}_{\mathbf{F}_q/\mathbf{F}_p}\colon \mathbf{F}_q \to \overline{\mathbf{Q}}^\times.$$

3.2. **Derived Artin stacks.** Unless noted otherwise, we always work in the category of derived Artin stacks. Hence when we say "Cartesian square" we mean what might be called "derived Cartesian square" (sometimes we keep the adjective "derived" for emphasis).

For a derived Artin stack A, we denote by A_{cl} its classical truncation.

By definition, a map of derived Artin stacks is a *closed embedding* or *proper* if the induced map of classical truncations has this property. For example, the inclusion of the classical truncation $A_{\mathrm{cl}} \hookrightarrow A$ is a closed embedding.

We say a map of derived Artin stacks is *schematic* if it is representable in derived schemes.

3.3. **Cotangent complexes.** For a map $f\colon A \to B$ of derived Artin stacks, we denote by $\mathbf{L}_f := \mathbf{L}_{A/B} \in \mathrm{Perf}(A)$ the relative cotangent complex.

Let $f\colon A \to B$ be a map of derived Artin stacks that is locally finitely presented on classical truncations. The map f is *étale* if the relative cotangent complex $\mathbf{L}_f$ vanishes (i.e., is isomorphic to $0 \in \mathrm{Perf}(A)$). The map f is *smooth* (resp. *quasi-smooth*) if the relative cotangent complex $\mathbf{L}_f$ is perfect of tor-amplitude $[0, \infty)$ (resp. $[-1, \infty)$). We use cohomological indexing for these tor-amplitude intervals; thus vector bundles lie in degree 0, positive degrees record "stacky directions", and negative directions record "derived directions".

Note that unlike properness, these properties cannot in general be detected on classical truncations. Moreover, while a smooth map is also smooth on classical truncations, quasi-smoothness is typically destroyed by classical truncation. Informally speaking, the ubiquity of quasi-smoothness[1] (in contrast to its classical counterpart, the property of being a local complete intersection) underpins the significance of derived algebraic geometry in our work.

When $\mathbf{L}_f$ is perfect, we write $d(f)$ for its virtual rank (or Euler characteristic), and call it the *relative dimension* of f. This integer is locally constant on A; for disconnected sources, all Chow degrees, shifts, and Tate twists involving $d(f)$ are read componentwise.

3.4. **Derived motives.** For a derived Artin stack A, we write $\mathcal{D}_{\mathrm{mot}}(A; \overline{\mathbf{Q}})$ for the derived category of motives on A with $\overline{\mathbf{Q}}$-coefficients in the sense of [FK24, §3.1], and $\mathcal{D}_{\mathrm{mot,gm}}(A; \overline{\mathbf{Q}}) \subset \mathcal{D}_{\mathrm{mot}}(A; \overline{\mathbf{Q}})$ for the full subcategory of geometric motives in the sense of [FK24, §3.3].

We denote the monoidal unit of $\mathcal{D}_{\mathrm{mot}}(A; \overline{\mathbf{Q}})$ by $\overline{\mathbf{Q}}_A$.

We use cohomological indexing. For $\mathcal{K} \in \mathcal{D}_{\mathrm{mot}}(A; \overline{\mathbf{Q}})$ and $i \in \mathbf{Z}$, we write $\mathcal{K}_{[i]}$ for the cohomological shift, $\mathcal{K}_{(i)}$ for Tate twist, and $\mathcal{K}_{\langle i \rangle} := \mathcal{K}_{[2i](i)} \in \mathcal{D}_{\mathrm{mot}}(A; \overline{\mathbf{Q}})$ for the indicated shift and Tate twist.

3.5. **Chow groups.** Let A be a derived Artin stack. We define the *Chow cohomology groups* of A (with $\overline{\mathbf{Q}}$-coefficients) by

$$\mathrm{CH}^i(A) := \mathrm{H}^{2i}(A; \overline{\mathbf{Q}}_A(i)) \cong \mathrm{H}^0(A; \overline{\mathbf{Q}}_A\langle i \rangle), \quad \text{for } i \in \mathbf{Z}.$$

Now suppose that A is locally of finite type over a field $\mathbf{F}$, and let $\pi\colon A \to \mathrm{Spec}\,(\mathbf{F})$ be the structural morphism. We define the *Chow groups* of A (with $\overline{\mathbf{Q}}$-coefficients) by

$$\mathrm{CH}_i(A) := \mathrm{H}^{-2i}(A; \pi^! \overline{\mathbf{Q}}_{\mathrm{Spec}\,(\mathbf{F})}(-i)) \cong \mathrm{H}^0(A; \pi^! \overline{\mathbf{Q}}_{\mathrm{Spec}\,(\mathbf{F})}\langle -i \rangle), \quad \text{for } i \in \mathbf{Z}$$

[1] This is sometimes called the "Hidden smoothness philosophy", and was anticipated by Deligne, Drinfeld, and Kontsevich.

This definition agrees with the classical definition of Chow groups with $\overline{\mathbf{Q}}$-coefficients *under assumptions* that A is "reasonable": see the discussion of [FK24, §3.5]. If A is a derived Artin stack, then by the derived invariance of $\mathcal{D}_{\mathrm{mot}}(A;\overline{\mathbf{Q}})$, the inclusion of the classical truncation $A_{\mathrm{cl}} \hookrightarrow A$ induces isomorphisms $\mathrm{CH}_i(A_{\mathrm{cl}}) \cong \mathrm{CH}_i(A)$.

More generally, for a morphism $f\colon A \to B$ of derived Artin stacks over a field $\mathbf{F}$ that is locally of finite type, we define the *relative Chow groups of* f to be

$$\mathrm{CH}_i(A/B) := \mathrm{H}^{-2i}(A; f^!\overline{\mathbf{Q}}_{B(-i)}) \cong \mathrm{H}^0(A; f^!\overline{\mathbf{Q}}_{B\langle -i\rangle}), \quad \text{for } i \in \mathbf{Z}.$$

For the identity morphism of A, these definitions give a canonical identification

$$\mathrm{CH}_i(A/A) \cong \mathrm{CH}^{-i}(A).$$

Functorial properties.

- The Chow groups are covariantly functorial with respect to proper morphisms. That is, if $f\colon A \to B$ is a proper morphism of derived Artin stacks, then we have pushforward maps
$$f_!\colon \mathrm{CH}_i(A) \to \mathrm{CH}_i(B) \tag{3.5.1}$$
and more generally $f_!\colon \mathrm{CH}_i(A/C) \to \mathrm{CH}_i(B/C)$ if f is defined over some C.
- Let $f\colon A \to B$ be a quasi-smooth map of derived Artin stacks, of relative dimension $d(f)$. There are (virtual) Gysin pullback maps
$$f^*\colon \mathrm{CH}_i(B) \to \mathrm{CH}_{i+d(f)}(A), \tag{3.5.2}$$
and more generally $f^*: \mathrm{CH}_i(B/C) \to \mathrm{CH}_{i+d(f)}(A/C)$ if B is defined over C. These are functorial and satisfy a base change formula with respect to proper pushforwards.

The maps (3.5.2) are induced by a natural transformation

$$[f]\colon f^* \to f^!\langle -d(f)\rangle \tag{3.5.3}$$

called the *relative fundamental class* (or *Gysin* natural transformation), constructed in [Kha19, §3]. It satisfies various natural compatibilities detailed in [Kha19, §3.2] or [FYZ23, §3.4].

Example 3.5.1 (Derived fundamental classes)**.** In particular, one has a relative (virtual) fundamental class

$$[A/B] \in \mathrm{CH}_{d(f)}(A/B) \tag{3.5.4}$$

defined as the Gysin pullback of the unit in $\mathrm{CH}_0(B/B) \cong \mathrm{CH}^0(B)$. Equivalently, it is determined by the morphism

$$\overline{\mathbf{Q}}_A \cong f^*\overline{\mathbf{Q}}_B \xrightarrow{\mathrm{gys}_f} f^!\overline{\mathbf{Q}}_{B\langle -d(f)\rangle} \tag{3.5.5}$$

obtained by evaluating the Gysin transformation (3.5.3) on $\overline{\mathbf{Q}}_B$. When $B = \mathrm{pt} := \mathrm{Spec}\ \mathbf{F}$, we abbreviate $[A] := [A/B]$ and call it the (derived) *fundamental class* of A.

3.6. **Derived vector bundles.** Let A be a derived Artin stack.

Given a perfect complex $\mathcal{E}$ on A, we denote by $\mathrm{Tot}(\mathcal{E})$ the derived stack of sections of $\mathcal{E}$, as in [FYZ23, §6.1.1]. We refer to $\mathrm{Tot}(\mathcal{E})$ as the *derived vector bundle* associated with $\mathcal{E}$. In terms of the functor of points, $\mathrm{Tot}(\mathcal{E})$ is the derived stack over A sending an A-scheme $u : T \to A$ to the mapping space $\mathrm{Map}_{\mathrm{QCoh}(T)}(\mathcal{O}_T, u^*\mathcal{E})$.

Throughout we use calligraphic letters such as $\mathcal{E}$ for perfect complexes, and Roman letters such as E for the corresponding total spaces.

We denote by $\mathcal{E}^*$ the dual perfect complex to $\mathcal{E}$. When $\mathcal{E}$ is a perfect complex on $X \times A$ (or on $X' \times A$, with $\omega_{X'}$ in place of ω_X), we denote by $\mathcal{E}^\vee := \mathcal{E}^* \otimes \omega_X$ the *Serre dual* perfect complex, with ω_X implicitly pulled back from X. We will denote the dual derived vector bundle to $E = \mathrm{Tot}(\mathcal{E})$ by $\widehat{E} = \mathrm{Tot}(\mathcal{E}^*)$.

Part 1. The Trace Conjecture

4. Dimension bounds on special cycles in low corank

In this section we prove a key dimension bound on the "injective part" of the special cycles defined in [FYZ25], in low corank. This will be a key input to the proof of the Trace Conjecture in low corank.

4.1. **Setup and statement.** Let $k = \mathbb{F}_q$ have odd characteristic, let X be a smooth proper geometrically connected curve over k, and let $\nu : X' \to X$ be an étale quadratic cover with involution σ. Throughout Part I, the cover $X' \to X$ is assumed connected; the split case is treated separately in Theorem 13.2.1 of Part 4. Fix integers $1 \le m \le n$, an integer $r \ge 0$, and a rank-m vector bundle $\mathcal{E}$ on X'.

4.1.1. *Hermitian bundles.* We recall the definitions needed in this section. For a vector bundle $\mathcal{F}$ on $X' \times S$, write $\mathcal{F}^\vee := \mathcal{H}om_{X'\times S}(\mathcal{F}, \omega_{X'} \boxtimes \mathcal{O}_S)$. A rank n Hermitian, or unitary, bundle is a pair $(\mathcal{F}, h)$ where $\mathcal{F}$ is a rank n vector bundle on $X' \times S$ and $h\colon \mathcal{F} \xrightarrow{\sim} \sigma^*\mathcal{F}^\vee$ satisfies $\sigma^* h^\vee = h$; see [FYZ24, Definition 6.1]. The stack of such bundles is denoted $\mathrm{Bun}_{U(n)}$.

4.1.2. *Hecke correspondences.* The Hecke stack $\mathrm{Hk}^r_{U(n)}$ is the stack of r successive unitary modifications from [FYZ24, Definition 6.3]. Its S-points consist of legs $x'_1, \dots, x'_r \in X'(S)$, Hermitian bundles $(\mathcal{F}_i, h_i)$ for $0 \le i \le r$, and isomorphisms

$$f_i\colon \mathcal{F}_{i-1}|_{X'\times S - (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} \xrightarrow{\sim} \mathcal{F}_i|_{X'\times S - (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} \quad (1 \le i \le r) \tag{4.1.1}$$

compatible with the Hermitian structures. The modification f_i is required to be lower of length 1 at x'_i and upper of length 1 at $\sigma(x'_i)$ in the terminology of [FYZ24, Definition 6.5]: equivalently, there is an intermediate vector bundle $\mathcal{F}^\flat_{i-1/2}$ with injections $\mathcal{F}^\flat_{i-1/2} \hookrightarrow \mathcal{F}_{i-1}$ and $\mathcal{F}^\flat_{i-1/2} \hookrightarrow \mathcal{F}_i$ whose cokernels are line bundles supported on the graphs of x'_i and $\sigma(x'_i)$, respectively.

4.1.3. *Hermitian shtukas.* The stack of rank n Hermitian shtukas is defined as in [FYZ24, Definition 6.6] by the Cartesian square

$$\begin{array}{ccc} \mathrm{Sht}^r_{U(n)} & \longrightarrow & \mathrm{Hk}^r_{U(n)} \\ \downarrow & & \downarrow{\scriptstyle (\mathrm{pr}_0, \mathrm{pr}_r)} \\ \mathrm{Bun}_{U(n)} & \xrightarrow[(\mathrm{Id}, \mathrm{Frob})]{} & \mathrm{Bun}_{U(n)} \times \mathrm{Bun}_{U(n)}\,. \end{array} \tag{4.1.2}$$

We abbreviate an S-point of $\mathrm{Sht}^r_{U(n)}$ as

$$(\mathcal{F}_0, h_0) \dashrightarrow (\mathcal{F}_1, h_1) \dashrightarrow \cdots \dashrightarrow (\mathcal{F}_r, h_r) \xrightarrow{\varphi} (\mathrm{Frob}_S^* \mathcal{F}_0, \mathrm{Frob}_S^* h_0). \tag{4.1.3}$$

4.1.4. *Special cycles.* The special cycle $\mathcal{Z}^r_{\mathcal{E}}$ is defined in [FYZ24, Definition 7.1]. Its S-points are points of $\mathrm{Sht}^r_{U(n)}(S)$ together with maps

$$\begin{array}{ccccccccc} \mathcal{E}\boxtimes\mathcal{O}_S & = & \mathcal{E}\boxtimes\mathcal{O}_S & = & \cdots & = & \mathcal{E}\boxtimes\mathcal{O}_S & \xrightarrow{\sim} & \mathrm{Frob}_S^*(\mathcal{E}\boxtimes\mathcal{O}_S) \\ \downarrow{\scriptstyle t_0} & & \downarrow{\scriptstyle t_1} & & \downarrow & & \downarrow{\scriptstyle t_r} & & \downarrow{\scriptstyle \mathrm{Frob}_S^* t_0} \\ \mathcal{F}_0 & \xrightarrow[f_1]{} & \mathcal{F}_1 & \dashrightarrow & \cdots & \xrightarrow[f_r]{} & \mathcal{F}_r & \xrightarrow[\varphi]{} & \mathrm{Frob}_S^*\mathcal{F}_0. \end{array} \tag{4.1.4}$$

such that φ identifies t_r with $\mathrm{Frob}_S^* t_0$, and for each $1 \le i \le r$ the maps t_{i-1} and t_i agree under the generic isomorphism f_i in (4.1.1). For this compatibility, the composite maps

$$a_t\colon \mathcal{E}\boxtimes\mathcal{O}_S \xrightarrow{t_i} \mathcal{F}_i \xrightarrow{h_i} \sigma^*\mathcal{F}_i^\vee \xrightarrow{\sigma^* t_i^\vee} \sigma^*\mathcal{E}^\vee \boxtimes \mathcal{O}_S, \tag{4.1.5}$$

are independent of i, and invariant under Frob_S. This implies that they must lie in the set $\mathcal{A}_{\mathcal{E}}(k)$ of Hermitian maps $\mathcal{E} \to \sigma^*\mathcal{E}^\vee$. This induces an open-closed decomposition

$$\mathcal{Z}^r_{\mathcal{E}} = \coprod_{a \in \mathcal{A}_{\mathcal{E}}(k)} \mathcal{Z}^r_{\mathcal{E}}(a).$$

Finally, $\mathcal{Z}^{r,\circ}_{\mathcal{E}} \subset \mathcal{Z}^r_{\mathcal{E}}$ denotes the open substack where the maps in (4.1.4) are injective after pullback to every geometric fiber of the test scheme. In the notation of [FYZ24, Definition 7.4], it restricts on the coefficient-a summand to $\mathcal{Z}^r_{\mathcal{E}}(a)^\circ$. The virtual class $[\mathcal{Z}^{r,\circ}_{\mathcal{E}}]$ is defined in [FYZ25, Definition 4.4], with an equivalent quotient formulation in [FYZ25, Definition 4.5], but the present section only uses the underlying geometric locus $\mathcal{Z}^{r,\circ}_{\mathcal{E}}$.

The main result of this section is the following theorem.

Theorem 4.1.1. *If $m \leq n/3$, then for all $\mathcal{E}$ and r we have*

$$\dim \mathcal{Z}_{\mathcal{E}}^{r,\circ} \leq (n-m)r. \tag{4.1.6}$$

Remark 4.1.2 (Description of AI assistance)**.** Theorem 4.1.1 is a similar result to [FYZ24, Proposition 9.5] which treated the case $\operatorname{rank}\mathcal{E} = 1$, but the proof — which was obtained with the aid of GPT 5.5 Pro — carries some additional subtleties. The main technical difference with the proof of *loc. cit.* is that when $\operatorname{rank}\mathcal{E} > 1$, the saturation of $\mathcal{E}$ in an ambient bundle $\mathcal{F}_i$ not necessarily given by twisting with a line bundle.

Now we describe the contribution of GPT 5.5 Pro in more detail. The first author queried GPT 5.5 Pro whether a bound of the form (4.1.6) was true for $m \leq \epsilon n$ for some $\epsilon > 0$. We knew from examples coming from parabolic shtukas that necessarily $\epsilon \leq 1/3$. GPT 5.5 Pro returned a response concluding that $\epsilon = 1/3$, which gave us confidence that it might be correct. We partially read this response, and explain how it relates to our eventual argument, as written below.

- GPT 5.5 Pro's response contained a passage very similar to the material of §4.2, although it gave a different argument based on Quot schemes, which we did not read.
- GPT 5.5 Pro's response contained the material of 4.3, and in fact we rewrote that section starting from its output. This material was already known to the authors, and is a straightforward generalization of the analysis in [FYZ24, §9.2].
- The key new trick in GPT 5.5 Pro's response (from the authors' perspective) was to use Lemma 4.4.4 to kill the dimensions coming from the possible saturations of $\mathcal{E}$ in the $\mathcal{F}_i$.

The argument below was written by the authors, except for the figures, which were generated by Codex.

4.2. **The saturation stratification.**

Definition 4.2.1. For a geometric point $(\mathcal{E}, \mathcal{F}_\bullet, t_\bullet)$ of $\mathcal{Z}_{\mathcal{E}}^{r,\circ}$, for each $i = 0, \ldots, r$ let

$$\mathcal{P}_i := \operatorname{Sat}_{\mathcal{F}_i}(t_i(\mathcal{E})) \quad \text{be the saturation of } t_i(\mathcal{E}) \text{ in } \mathcal{F}_i. \tag{4.2.1}$$

Then $\mathcal{P}_i \subset \mathcal{F}_i$ is a rank-m subbundle, and $\mathcal{F}_i/\mathcal{P}_i$ is locally free.

Lemma 4.2.2. *Let X' be a smooth projective curve over a field k. Let K/k be a field extension. Let T be a reduced finite-type K-scheme. Let $\mathcal{E}$ and $\mathcal{F}$ be vector bundles of ranks m and n, respectively, on $X'_T := X' \times_k T$, and let*

$$t : \mathcal{E} \longrightarrow \mathcal{F}$$

be a morphism which is injective after pullback to all geometric points of T. Then there is a finite locally closed stratification

$$T = \coprod_{\alpha=1}^{N} T_\alpha$$

such that after pullback to each T_α, the fiberwise saturation of $t(\mathcal{E}|_{T_\alpha})$ in $\mathcal{F}|_{T_\alpha}$ is represented by a rank-m subbundle $\mathcal{P}_\alpha \subset \mathcal{F}|_{T_\alpha}$ with locally free quotient. Moreover, $\mathcal{P}_\alpha/t(\mathcal{E}|_{T_\alpha})$ is finite flat over T_α.

Proof. By Noetherian induction, it suffices to produce a non-empty open subset $U \subset T$ and a rank m saturated sub-bundle $\mathcal{P}_U \subset \mathcal{F}|_U$ such that the following properties hold:

(i) $t(\mathcal{E}|_U) \subset \mathcal{P}_U \subset \mathcal{F}|_U$,
(ii) After pullback to any geometric point of U, $\mathcal{P}_U$ is the saturation of $t(\mathcal{E}|_U)$.
(iii) $\mathcal{P}_U/t(\mathcal{E}|_U)$ is finite flat over U.

Step 1. Choose a generic point η of an irreducible component of T. The saturation of $t(\mathcal{E}|_\eta)$ in $\mathcal{F}|_\eta$ is a rank m sub-bundle that we denote $\mathcal{P}_\eta$. Choose an open neighborhood U of η contained in that irreducible component. Since X'_η is a smooth curve and the quotient $\mathcal{F}_\eta/\mathcal{P}_\eta$ is torsion-free on X'_η, it is locally free of rank $n-m$. Spreading out the exact sequence

$$0 \to \mathcal{P}_\eta \to \mathcal{F}_\eta \to (\mathcal{F}_\eta/\mathcal{P}_\eta) \to 0$$

and shrinking U if necessary, we obtain a rank-m subbundle $\mathcal{P}_U \subset \mathcal{F}|_U$ such that $\mathcal{F}|_U/\mathcal{P}_U$ is locally free.

We next argue that it is possible to shrink U to arrange that $t(\mathcal{E}|_U) \subset \mathcal{P}_U$. The composite map $\mathcal{E}|_U \to \mathcal{F}|_U/\mathcal{P}_U$ vanishes at η. The support of its image is therefore a closed subset of X'_U missing the

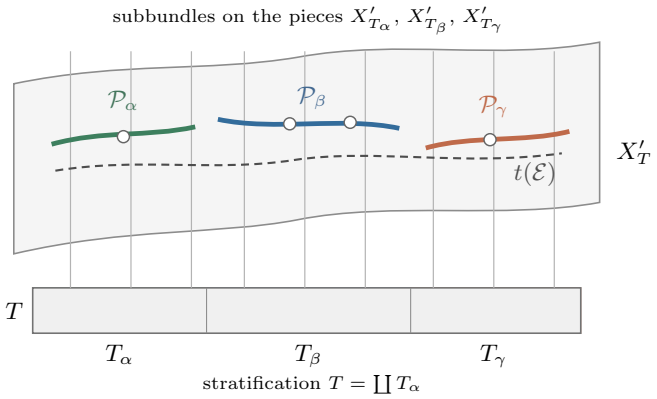

FIGURE 2. Schematic picture of Lemma 4.2.2. After stratifying the base T, the fiberwise saturation of $t(\mathcal{E})$ in $\mathcal{F}$ is represented over each stratum T_α by a rank-m subbundle $\mathcal{P}_\alpha \subset \mathcal{F}|_{X'_{T_\alpha}}$.

fiber over η. By properness of X', the image of this closed subset in U is a closed subset missing η. Therefore, after replacing U by an open subset, still containing η, we may assume that $\mathcal{E}|_U \to \mathcal{F}|_U/\mathcal{P}_U$ vanishes, so $\mathcal{P}_U$ contains $t(\mathcal{E}|_U)$.

Step 2. Consider the quotient $\mathcal{Q}_U := \mathcal{P}_U/t(\mathcal{E}|_U)$. Its support is closed in X'_U, hence is proper over U. Since $\mathcal{P}_\eta/t_\eta(\mathcal{E}_\eta)$ has finite support on X'_η, after shrinking U to a further open subset containing η, we may further assume that $\mathcal{Q}_U$ has support which is finite over U.

We claim that $\mathcal{P}_U$ represents the fiberwise saturated image over any geometric point $u \to U$. Since $\mathcal{F}_U/\mathcal{P}_U$ and $\mathcal{P}_U$ are locally free, $(\mathcal{F}_U/\mathcal{P}_U)_u \cong \mathcal{F}_u/\mathcal{P}_u$ is torsion-free on the smooth curve X'_u. Since $\mathcal{Q}_U$ is finite over U, $\mathcal{Q}_u = (\mathcal{P}_U/t(\mathcal{E}|_U))_u \cong (\mathcal{P}_u/t_u(\mathcal{E}_u))$ has finite support on X'_u, hence is torsion. Therefore $\mathcal{P}_u/t_u(\mathcal{E}_u)$ is precisely the torsion subsheaf of $\mathcal{F}_u/t_u(\mathcal{E}_u)$, so $\mathcal{P}_u$ is the saturation of $t_u(\mathcal{E}_u)$.

Step 3. We have an exact sequence

$$0 \to \mathcal{P}_U/t(\mathcal{E}|_U) \to (\mathcal{F}|_U)/t(\mathcal{E}|_U) \to (\mathcal{F}|_U)/\mathcal{P}_U \to 0.$$

Thanks to the assumption that the map $\mathcal{E}|_U \to \mathcal{F}|_U$ is injective after restriction to all geometric points of U, [Sta20, Tag 00ME] applies to show that $(\mathcal{F}|_U/t(\mathcal{E}|_U))$ is flat over U. The right-hand term is locally free on X'_U, hence flat over U. Therefore $\mathcal{P}_U/t(\mathcal{E}|_U)$ is flat over U. Since its support is finite over U by Step 2, $\mathcal{P}_U/t(\mathcal{E}|_U)$ is finite flat over U, as required. □

4.3. **Framed Hitchin stacks.** Fix $1 \le m \le n$. Let $\mathcal{N}$ be the stack over k whose R-points for a commutative k-algebra R form the groupoid of triples $(\mathcal{P}, \mathcal{F}, h)$, where:

- $(\mathcal{F}, h)$ is a rank n Hermitian bundle on X'_R, and
- $\mathcal{P} \subset \mathcal{F}$ is a rank m sub-bundle such that $\mathcal{F}/\mathcal{P}$ is locally free over X'_R.

This is a framed version of the Hitchin stack recalled in §5.1.

Let $\mathrm{Hk}^r_{\mathcal{N}}$ be the Hecke stack with R-points the groupoid of $\mathcal{F}_\bullet \in \mathrm{Hk}^r_{U(n)}(R)$ together with saturated sub-bundles $\mathcal{P}_i \subset \mathcal{F}_i$ such that the isomorphisms

$$\mathcal{F}_i|_{X'_R \setminus (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} \cong \mathcal{F}_{i-1}|_{X'_R \setminus (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} \qquad i = 0, \dots, r$$

identify

$$\mathcal{P}_i|_{X'_R \setminus (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} = \mathcal{P}_{i-1}|_{X'_R \setminus (\Gamma_{x'_i} \cup \Gamma_{\sigma x'_i})} \qquad i = 0, \dots, r.$$

For $0 \le i \le r$, let $\mathrm{pr}_i \colon \mathrm{Hk}^r_{\mathcal{N}} \to \mathcal{N}$ be the map recording $(\mathcal{P}_i, \mathcal{F}_i, h_i)$. Let $\mathrm{pr}_{\mathrm{legs}} \colon \mathrm{Hk}^r_{\mathcal{N}} \to (X')^r$ be the map recording the legs.

Proposition 4.3.1. *Over $\mathcal{N} \times X'$, let x be the universal point of X' and set $y := \sigma(x)$. Over the universal projective bundle $\mathbb{P}(\mathcal{F}^{\mathrm{univ}}_x) \cong \mathbb{P}((\mathcal{F}^{\mathrm{univ}}_y)^\vee) \to \mathcal{N} \times X'$ let $L \subset \mathcal{F}^{\mathrm{univ}}_x$ be the universal line; L is equivalent to the datum of the universal hyperplane $H = L^\perp \subset \mathcal{F}^{\mathrm{univ}}_y$, where orthogonal complements are taken with*

respect to the given Hermitian form. On each ambient projective bundle in the table, L and H denote the pullbacks of these universal objects. Define the four relative parameter spaces over $\mathcal{N} \times X'$ as follows:

type	ambient projective bundle	parameter space
0	$\mathbb{P}((\mathcal{F}_y^{\mathrm{univ}}/\mathcal{P}_y^{\mathrm{univ}})^\vee)$	$H^\perp \not\subset \mathcal{P}_x^{\mathrm{univ}}$
$+$	$\mathbb{P}(\mathcal{F}_y^{\mathrm{univ},\vee})$	$H \not\supset \mathcal{P}_y^{\mathrm{univ}},\ H^\perp \not\subset \mathcal{P}_x^{\mathrm{univ}}$
$-$	$\mathbb{P}(\mathcal{P}_x^{\mathrm{univ}})$	$L \subset (\mathcal{P}_y^{\mathrm{univ}})^\perp$
$\pm$	$\mathbb{P}(\mathcal{P}_x^{\mathrm{univ}})$	$L \not\subset (\mathcal{P}_y^{\mathrm{univ}})^\perp$

(In types $0, -, \pm$, the ambient projective bundle already enforces one of the closed fiber conditions, so the parameter space column records the non-automatic condition.) Note that the loci in type $0, +, \pm$ are open in their ambient projective bundles, while the type $-$ locus is closed.

Then there is a stratification of $\mathrm{Hk}^1_{\mathcal{N}}$ into loci ${}_?\mathrm{Hk}^1_{\mathcal{N}}$ with $? \in \{0, +, -, \pm\}$ such that the morphism

$$(\mathrm{pr}_1, \mathrm{pr}_{X'}) \colon {}_?\mathrm{Hk}^1_{\mathcal{N}} \longrightarrow \mathcal{N} \times X',$$

where $\mathrm{pr}_{X'}$ records the leg, realizes the stratum ${}_?\mathrm{Hk}^1_{\mathcal{N}}$ as the corresponding parameter space.

Proof. To complete $(\mathcal{P}, \mathcal{F}, h, x) \in \mathcal{N} \times X'(R)$ to a point of $\mathrm{Hk}^1_{\mathcal{N}}(R)$, we need to construct the diagram

$$\begin{array}{ccccc} \mathcal{F}' & \hookleftarrow & \mathcal{F}^\flat & \hookrightarrow & \mathcal{F} \\ \uparrow & & & & \uparrow \\ \mathcal{P}' & & & & \mathcal{P} \end{array}$$

where $\mathcal{F}'$ is related to $\mathcal{F}$ by a lower modification of $\mathcal{F}$ along y, followed by the Hermitian-dual upper modification at x. Let $H \subset \mathcal{F}_y$ be the hyperplane defining the lower modification, so that

$$\mathcal{F}^\flat = \ker(\mathcal{F} \to \mathcal{F}_y/H), \qquad \mathcal{F}' = \mathcal{F}^\flat + \mathcal{I}_x^{-1} L.$$

Then $\mathcal{F}'$ carries the Hermitian structure induced by h.

Given $(\mathcal{P}, \mathcal{F}, h, x)$ and H, or equivalently $L = H^\perp$, at most one saturated rank-m subbundle $\mathcal{P}' \subset \mathcal{F}'$ agrees with $\mathcal{P}$ away from $\Gamma_x \cup \Gamma_y$. Indeed, if $\mathcal{P}'$ and $\mathcal{P}''$ both satisfy this property, then $\mathcal{P}'' \to \mathcal{F}' \to \mathcal{F}'/\mathcal{P}'$ vanishes off the relative Cartier divisor $\Gamma_x \cup \Gamma_y$. Thus $\mathcal{P}'' \subset \mathcal{P}'$, and symmetry gives equality.

To classify the modification at geometric points, assume for the moment that R is a field. The Hermitian structure gives the perfect pairing

$$\langle\ ,\ \rangle_{x,y} : \mathcal{F}_x \times \mathcal{F}_y \longrightarrow \omega_{X'}|_y. \tag{4.3.1}$$

A hyperplane $H \subset \mathcal{F}_y$ determines the lower modification at y, and the orthogonal line $H^\perp \subset \mathcal{F}_x$ determines its Hermitian-dual upper modification at x. Thus $\mathcal{P}$ and $\mathcal{P}'$ agree away from x and y, and at each of these points the change has length zero or one. At y we have

$$\mathcal{F}^\flat = \ker(\mathcal{F} \twoheadrightarrow \mathcal{F}_y/H).$$

while the generic identification of $\mathcal{F}'$ with $\mathcal{F}^\flat$ gives

$$\mathcal{F}' = \mathcal{F}^\flat + \mathcal{I}_x^{-1} H^\perp$$

where $\mathcal{I}_x$ is the ideal sheaf of the graph of x. The two conditions $H \supset \mathcal{P}_y$ and $H^\perp \subset \mathcal{P}_x$ therefore determine the modification type, with the precise fiber changes recorded in (4.3.2).

fiber condition	sub-bundle modification
$H \supset \mathcal{P}_y$	$\mathcal{P}'_y = \mathcal{P}_y$
$H \not\supset \mathcal{P}_y$	$\mathrm{length}_y(\mathcal{P}/\mathcal{P}') = 1$
$H^\perp \not\subset \mathcal{P}_x$	$\mathcal{P}'_x = \mathcal{P}_x$
$H^\perp \subset \mathcal{P}_x$	$\mathrm{length}_x(\mathcal{P}'/\mathcal{P}) = 1.$

(4.3.2)

The four combinations give the following modification types.

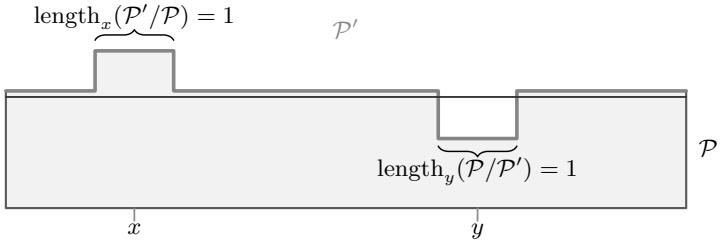

FIGURE 3. Cartoon of a modification between $\mathcal{P}$ and $\mathcal{P}'$. The bump at x indicates $\text{length}_x(\mathcal{P}'/\mathcal{P}) = 1$, while the dip at y indicates $\text{length}_y(\mathcal{P}/\mathcal{P}') = 1$.

type	condition at y	condition at x	sub-bundle modification
0	$H \supset \mathcal{P}_y$	$H^\perp \not\subset \mathcal{P}_x$	
$+$	$H \not\supset \mathcal{P}_y$	$H^\perp \not\subset \mathcal{P}_x$	
$-$	$H \supset \mathcal{P}_y$	$H^\perp \subset \mathcal{P}_x$	
$\pm$	$H \not\supset \mathcal{P}_y$	$H^\perp \subset \mathcal{P}_x$	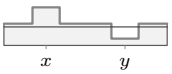

These are illustrated in Figure 4.

For an arbitrary test ring R, take $? \in \{0, +, -, \pm\}$ and let $\mathcal{P}$ and $H = L^\perp$ satisfy the corresponding parameter-space conditions. Define $\mathcal{P}'$ by the following formula:

type	$\mathcal{P}'$
0	$\mathcal{P}$
$+$	$\ker(\mathcal{P} \to \mathcal{F}_y/H)$
$-$	$\mathcal{P} + \mathcal{I}_x^{-1}L$
$\pm$	$\ker(\mathcal{P} \to \mathcal{F}_y/H) + \mathcal{I}_x^{-1}L.$

Each formula gives a saturated rank-m subbundle $\mathcal{P}' \subset \mathcal{F}'$ with locally free quotient $\mathcal{F}'/\mathcal{P}'$.

At every geometric point, exactly one of these four types occurs. Hence the substacks ${}_?\,\mathrm{Hk}^1_{\mathcal{N}}$ stratify $\mathrm{Hk}^1_{\mathcal{N}}$. Their geometry is read from the defining fiber conditions as follows.

- In type 0, H is a hyperplane in $\mathcal{F}_y$ containing $\mathcal{P}_y$ and such that $H^\perp \not\subset \mathcal{P}_x$. Imposing the first condition gives the projective bundle of hyperplanes in $\mathcal{F}_y$ containing $\mathcal{P}_y$, namely $\mathbb{P}((\mathcal{F}_y/\mathcal{P}_y)^\vee)$, and the second condition is open.
- In type $+$, H is a hyperplane in $\mathcal{F}_y$ satisfying the two open conditions $H \not\supset \mathcal{P}_y$ and $H^\perp \not\subset \mathcal{P}_x$.
- To understand the type $-$ and $\pm$ spaces, it is more convenient to choose L first. It must lie in the projective bundle of lines in $\mathcal{P}_x$, namely $\mathbb{P}(\mathcal{P}_x)$. Then the type $-$ parameter space is cut out by the closed condition $L \subset \mathcal{P}_y^\perp$, while the type $\pm$ parameter space is its open complement.

This completes the proof. □

Definition 4.3.2. Let $\tau\colon \{1, \dots, r\} \to \{0, +, -, \pm\}$ be a labeling of the modifications by the four possible types. Equivalently, we may think of τ as a tuple $(I_0, I_+, I_-, I_\pm)$ forming a partition of $\{1, \dots, r\}$, with $I_?$ designating the indices of type ?.

We define ${}_\tau\,\mathrm{Hk}^r_{\mathcal{N}}$ to be the locally closed stratum of $\mathrm{Hk}^r_{\mathcal{N}}$ where the modifications in $I_?$ lie in the stratum of type $? \in \{0, +, -, \pm\}$.

Corollary 4.3.3. *Let $\tau = (I_0, I_+, I_-, I_\pm)$ be a partition of $\{1, \dots, r\}$. For every geometric point $w \in {}_\tau\,\mathrm{Hk}^r_{\mathcal{N}}$ there exist*

- *a scheme U,*

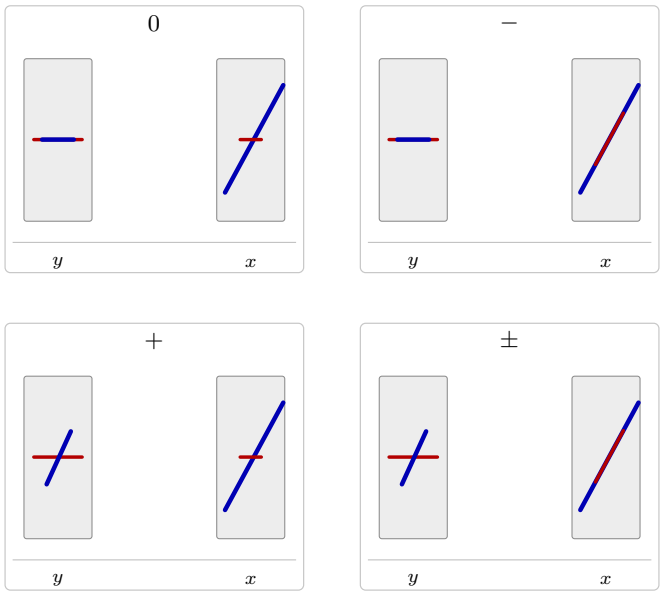

FIGURE 4. Cartoon of the four possible modification types. The grey rectangles depict the fibers of $\mathcal{F}_i$ at x and y. The blue line depicts the fibers of $\mathcal{P}$ at x and y. The red line over y depicts H and the red line over x depicts $H^\perp$.

- *a smooth morphism* $q : U \longrightarrow {}_\tau \mathrm{Hk}^r_{\mathcal{N}}$ *whose image contains* w,
- *a locally closed immersion* $U \hookrightarrow \overline{U}$, *and*
- *a smooth morphism* $\overline{U} \longrightarrow \mathcal{N} \times (X')^r$ *extending* $(\mathrm{pr}_r, \mathrm{pr}_{\mathrm{legs}}) \circ q$, *of relative dimension* $D_\tau + e$, *where*

$$D_\tau := \sum_{i=1}^{r} d_{\tau(i)} = (n-m-1)|I_0| + (n-1)|I_+| + (m-1)|I_-| + (m-1)|I_\pm|, \tag{4.3.3}$$

and e *is the relative dimension of* q *on the chosen component of* U.

Proof. The result is trivial if $r = 0$, so we henceforth assume that $r \geq 1$.

For $1 \leq a \leq r$, let $C_{[a,r]}$ denote the stratum of chains

$$(\mathcal{P}_{a-1}, \mathcal{F}_{a-1}) \dashrightarrow \cdots \dashrightarrow (\mathcal{P}_r, \mathcal{F}_r)$$

whose i-th modification has type $\tau(i)$. We prove, by descending induction on a, the following statement: after pullback to a smooth scheme chart $U \to C_{[a,r]}$ of relative dimension e, every point of U has a neighborhood locally closed in a scheme smooth over $\mathcal{N} \times (X')^{r-a+1}$ (via the last projection pr_r times the legs $x_a, \ldots, x_r$) of relative dimension $\sum_{i=a}^r d_{\tau(i)} + e$. This is the assertion of the corollary when $a = 1$.

For each type $? \in \{0, +, -, \pm\}$, write

$$B_? \longrightarrow \mathcal{N} \times X'$$

for the ambient projective bundle appearing in Proposition 4.3.1. Its relative dimension is $d_?$ where

$$d_0 = n - m - 1, \qquad d_+ = n - 1, \qquad d_- = d_\pm = m - 1.$$

The base case $a = r$ is Proposition 4.3.1: the one-step stratum becomes, after pullback to a smooth scheme chart of $\mathcal{N} \times X'$, locally closed in the corresponding ambient projective bundle $B_{\tau(r)}$. The chart dimension contributes the additional e.

Suppose the assertion has been proved for $C_{[a+1,r]}$. Fix a point in $C_{[a,r]}$, and choose a point of a smooth scheme chart $U \to C_{[a+1,r]}$ above its image, together with a locally closed immersion $U \hookrightarrow \overline{U}$, where $\overline{U}$ is smooth over $\mathcal{N} \times (X')^{r-a}$ of relative dimension $\sum_{i=a+1}^r d_{\tau(i)} + e$. Put

$$\widetilde{U} := C_{[a,r]} \times_{C_{[a+1,r]}} U.$$

The morphism $\widetilde{U} \to C_{[a,r]}$ is smooth of relative dimension e. The map $\mathrm{pr}_a : C_{[a+1,r]} \to \mathcal{N}$ pulls the ambient projective bundle $B_{\tau(a)}$ back to a projective bundle

$$Q \longrightarrow U \times X'$$

of relative dimension $d_{\tau(a)}$. By base change from Proposition 4.3.1, $\widetilde{U}$ is locally closed in Q.

Choose a Zariski open neighborhood $V \subset U \times X'$ of the chosen point on which the projective bundle Q is trivial. Since $U \hookrightarrow \overline{U}$ is locally closed, after shrinking around its image we may write $V = \overline{V} \cap (U \times X')$ for an open $\overline{V} \subset \overline{U} \times X'$. The trivialization identifies $Q|_V$ with $\mathbb{P}^{d_{\tau(a)}}_V$, so $\widetilde{U}|_V$ is locally closed in $\mathbb{P}^{d_{\tau(a)}}_{\overline{V}}$. This projective-space bundle is smooth over $\mathcal{N} \times (X')^{r-a+1}$ of relative dimension

$$d_{\tau(a)} + \sum_{i=a+1}^{r} d_{\tau(i)} + e.$$

Thus $\widetilde{U}|_V \to C_{[a,r]}$ and its embedding in $\mathbb{P}^{d_{\tau(a)}}_{\overline{V}}$ provide the required chart to prove the induction step. □

4.4. **Framed shtukas.** We define the stack of framed shtukas $\mathrm{Sht}^r_{\mathcal{N}}$ as the fibered product

$$\begin{array}{ccc} \mathrm{Sht}^r_{\mathcal{N}} & \longrightarrow & \mathrm{Hk}^r_{\mathcal{N}} \\ \downarrow & & \downarrow{\scriptstyle (\mathrm{pr}_0, \mathrm{pr}_r)} \\ \mathcal{N} & \xrightarrow{(\mathrm{Id},\mathrm{Frob})} & \mathcal{N} \times \mathcal{N} \end{array} \tag{4.4.1}$$

For $\tau \colon \{1, \dots, r\} \to \{0, +, -, \pm\}$ we define ${}_\tau\mathrm{Sht}^r_{\mathcal{N}}$ as the fibered product

$$\begin{array}{ccc} {}_\tau\mathrm{Sht}^r_{\mathcal{N}} & \longrightarrow & {}_\tau\mathrm{Hk}^r_{\mathcal{N}} \\ \downarrow & & \downarrow{\scriptstyle (\mathrm{pr}_0, \mathrm{pr}_r)} \\ \mathcal{N} & \xrightarrow{(\mathrm{Id},\mathrm{Frob})} & \mathcal{N} \times \mathcal{N} \end{array} \tag{4.4.2}$$

As τ varies over all such maps the ${}_\tau\mathrm{Hk}^r_{\mathcal{N}}$ form a locally closed stratification of $\mathrm{Hk}^r_{\mathcal{N}}$, hence the ${}_\tau\mathrm{Sht}^r_{\mathcal{N}}$ form a locally closed stratification of $\mathrm{Sht}^r_{\mathcal{N}}$.

Lemma 4.4.1. *Let $\tau = (I_0, I_+, I_-, I_\pm)$. If ${}_\tau\mathrm{Sht}^r_{\mathcal{N}}$ is non-empty, then $|I_+| = |I_-|$.*

Proof. Suppose ${}_\tau\mathrm{Sht}^r_{\mathcal{N}}$ has a geometric point $(x_\bullet, \mathcal{P}_\bullet, \mathcal{F}_\bullet)$. Then $\mathcal{P}_r \cong (\mathrm{Id}_{X'} \times \mathrm{Frob}_{\bar{k}})^* \mathcal{P}_0$, so their total degrees on X' agree. On the other hand, we have

$$\deg \mathcal{P}_0 - \deg \mathcal{P}_r = |I_-| - |I_+|.$$

Here $\mathrm{Frob}_{\bar{k}}$ is the q-Frobenius of the geometric test point in (4.4.1). Since $\mathcal{P}_r \cong \mathrm{Frob}^* \mathcal{P}_0$, we have $\deg \mathcal{P}_r = \deg \mathcal{P}_0$ hence if ${}_\tau\mathrm{Sht}^r_{\mathcal{N}}$ is non-empty then we must have $|I_-| = |I_+|$. □

The following lemma generalizes [FYZ24, Lemma 9.3] slightly: compared to *loc. cit.* the smooth envelope is allowed to contain U only as a locally closed subscheme, and we include an auxiliary T-factor.

Lemma 4.4.2. *Fix an algebraic closure $\overline{k}$ of k. For a $\overline{k}$-scheme Z, write $Z^{(1)} := Z \times_{\mathrm{Spec}\ \overline{k}, \mathrm{Frob}_q} \mathrm{Spec}\ \overline{k}$ for its q-Frobenius twist and $\mathrm{Frob} : Z \to Z^{(1)}$ for the relative q-Frobenius. Let W, Z, T be finite-type schemes over $\overline{k}$, and let*

$$(h_0, h_1, h_T) : W \longrightarrow Z \times Z^{(1)} \times T$$

be a morphism. Suppose that every point $w \in W$ has an open neighborhood $U \subset W$ for which there exist a finite-type scheme $\overline{U}$, a locally closed immersion $j : U \hookrightarrow \overline{U}$, and a smooth morphism

$$\overline{h} = (\overline{h}_1, \overline{h}_T) : \overline{U} \longrightarrow Z^{(1)} \times T$$

of relative dimension d, such that $\overline{h} \circ j = (h_1, h_T)|_U$. Define V by the Cartesian square

$$\begin{array}{ccc} V & \longrightarrow & W \\ \downarrow & & \downarrow{\scriptstyle (h_0, h_1)} \\ Z & \xrightarrow[(\mathrm{Id},\mathrm{Frob})]{} & Z \times Z^{(1)}. \end{array} \tag{4.4.3}$$

Then every fiber of $V \to T$ has dimension at most d.

Proof. Fix a geometric point $v \in V_t$, with image $w \in W$ and $t \in T$. It is enough to prove $\dim_v V_t \le d$. Choose $U \hookrightarrow \overline{U}$ at w as in the hypothesis and replace W by U. We may then work with a single smooth morphism

$$\overline{W} \longrightarrow Z^{(1)} \times T$$

with W locally closed in $\overline{W}$. For the purpose of proving the Lemma, we may make the following reductions:

- Base change to a geometric point Spec $\Omega \to T$ and replace $\overline{k}$ by Ω, reducing to the case where $T = \text{Spec } \Omega$.
- Replace Z by an affine neighborhood Z' of $h_0(w)$, and replace $\overline{W}$ by $\overline{h}_1^{-1}((Z')^{(1)})$ and W by $h_0^{-1}(Z') \cap \overline{h}_1^{-1}((Z')^{(1)})$, reducing to the case where Z is affine.

Since Z is affine, we may write

$$Z = \text{Spec } R, \qquad \text{where} \qquad R = \Omega[x_1, \dots, x_l]/I.$$

Then

$$Z^{(1)} = \text{Spec } R^{(1)}, \qquad \text{where} \qquad R^{(1)} \cong \Omega[\xi_1, \dots, \xi_l]/I^{(1)},$$

where $\xi_i = 1 \otimes x_i$ and $\text{Frob}^*(\xi_i) = x_i^q$. After shrinking around w, we may assume that $\overline{W}$ is affine and $W \subset \overline{W}$ is closed. The standard local form of the smooth morphism $\overline{h}_1 : \overline{W} \to Z^{(1)}$ of relative dimension d implies that there is a presentation

$$\overline{W} = \text{Spec } \frac{\Omega[\xi_1, \dots, \xi_l, y_1, \dots, y_{s+d}]_\Delta}{(I^{(1)}(\xi), r_1, \dots, r_s)}$$

for some element Δ, with

$$J := \det \left(\frac{\partial r_a}{\partial y_b} \right)_{1 \le a,b \le s}$$

invertible on $\overline{W}$.

Let $\bar{f}_i = h_0^* x_i \in \Gamma(W, \mathcal{O}_W)$. Because W is closed in $\overline{W}$, we may choose lifts

$$f_i \in \Omega[\xi_1, \dots, \xi_l, y_1, \dots, y_{s+d}]_\Delta$$

of the functions $\bar{f}_i$. Set

$$g_i := \xi_i - f_i^q, \qquad 1 \le i \le l.$$

By (4.4.3), the equations $\xi_i = \bar{f}_i^q$ cut out V on W. Hence, on this affine neighborhood, V is a closed subscheme of

$$U_{\text{amb}} := \text{Spec } \frac{\Omega[\xi_1, \dots, \xi_l, y_1, \dots, y_{s+d}]_\Delta}{(g_1, \dots, g_l, r_1, \dots, r_s)}.$$

Apart from the leading term ξ_i, each g_i depends on the variables through qth powers. In characteristic p, we have $\frac{\partial g_i}{\partial \xi_j} = \delta_{ij}$ and $\frac{\partial g_i}{\partial y_b} = 0$, so the Jacobian matrix for the equations $g_1, \dots, g_l, r_1, \dots, r_s$, using the column order

$$\xi_1, \dots, \xi_l \,;\, y_1, \dots, y_s \,;\, y_{s+1}, \dots, y_{s+d},$$

has the block form

$$\begin{pmatrix} \text{Id}_l & 0 & 0 \\ * & \left(\frac{\partial r_a}{\partial y_b}\right)_{1 \le a,b \le s} & * \end{pmatrix}.$$

The matrix has rank $l + s$, so U_{amb} is smooth of dimension $(l + s + d) - (l + s) = d$. Since V is locally a closed subscheme of U_{amb}, this gives $\dim_v V_t \le d$, proving the lemma. □

Corollary 4.4.3 (Stratum bound)**.** *Let $\tau = (I_0, I_+, I_-, I_\pm)$. Then we have*

$$\dim \left({}_\tau \text{Sht}^r_{\mathcal{N}} \right) \le (n - m)|I_0| + n|I_+| + m|I_-| + m|I_\pm|. \tag{4.4.4}$$

Proof. Let D_τ be the integer defined in (4.3.3), and let $T := (X')^r$. We first show that every geometric fiber of ${}_\tau\operatorname{Sht}^r_{\mathcal{N}} \longrightarrow T$ has dimension at most D_τ.

We may base change to $\overline{k}$. The assertion is local on T. Given a geometric point of T, choose a closed point $v \in |X|$ whose inverse image in X' is disjoint from all of its coordinates, and replace T by the open neighborhood

$$T_v := (X' \setminus \nu^{-1}(v))^r.$$

As v varies, these opens cover T.

It is also enough to work on finite-type truncations. Choose a Frobenius-stable truncation datum μ, consisting of a Harder–Narasimhan bound for the $\mathcal{F}_i$'s together with a lower bound for the locally constant integers $\deg \mathcal{P}_i$, and restrict to the locus on which all the framed bundles

$$(\mathcal{P}_i, \mathcal{F}_i), \qquad 0 \le i \le r,$$

are bounded by μ. As μ varies, these loci exhaust ${}_\tau\operatorname{Sht}^r_{\mathcal{N}}$.

Impose a sufficiently deep level structure at v, as in [FYZ24, proof of Proposition 9.1(2)]. Denote the resulting finite-type scheme covering the corresponding truncation of $\mathcal{N}$ by Z, and denote the corresponding level Hecke stratum by W. Since the legs lie in T_v, the level structure transports through every modification, and we obtain morphisms

$$h_0 : W \longrightarrow Z, \qquad h_r : W \longrightarrow Z^{(1)}, \qquad h_T : W \longrightarrow T_v.$$

Let V be defined by the Cartesian square

$$\begin{array}{ccc} V & \longrightarrow & W \\ \downarrow & & \downarrow{\scriptstyle (h_0,h_r)} \\ Z & \xrightarrow[(\mathrm{Id},\mathrm{Frob})]{} & Z \times Z^{(1)}. \end{array}$$

Corollary 4.3.3 implies that every point of W has an open neighborhood locally closed in a finite-type scheme smooth over $Z^{(1)} \times T_v$ of relative dimension D_τ. Lemma 4.4.2 therefore shows that every fiber of $V \longrightarrow T_v$ has dimension at most D_τ.

Since forgetting level structure is a finite étale map, it does not affect fiber dimensions. It follows that every geometric fiber of the map ${}_\tau\operatorname{Sht}^r_{\mathcal{N}} \longrightarrow T$ has dimension at most D_τ.

Since $\dim T = r$, we conclude that

$$\dim({}_\tau\operatorname{Sht}^r_{\mathcal{N}}) \le D_\tau + r = (n-m)|I_0| + n|I_+| + m|I_-| + m|I_\pm|,$$

as claimed. □

4.4.1. *Special cycles.* We prove Theorem 4.1.1. Since the coefficient map lands in the finite discrete set $\mathcal{A}_{\mathcal{E}}(k)$, the injective locus $\mathcal{Z}^{r,\circ}_{\mathcal{E}}$ is the finite disjoint union of the open-and-closed substacks $\mathcal{Z}^r_{\mathcal{E}}(a)^\circ$. It suffices to prove that if $m \le n/3$, then $(\mathcal{Z}^r_{\mathcal{E}}(a)^\circ)_{\mathrm{red}}$ is exhausted by finite type open substacks with dimension $\le r(n-m)$. We may thus impose a Harder–Narasimhan truncation and pass to a finite type open substack. Then we may apply Lemma 4.2.2 to the restriction of the universal family $\{t_\bullet \colon \mathcal{E} \hookrightarrow \mathcal{F}_\bullet\}$ (on a finite cover by a scheme) to obtain finitely many reduced locally closed strata where the saturation $\mathcal{P}_\bullet \subset \mathcal{F}_\bullet$ exists.

On each such reduced stratum, the saturations are compatible with the chain of modifications, thanks to the compatibility condition between each pair t_{i-1} and t_i. Thus the universal data on the stratum define a morphism to $\operatorname{Hk}^r_{\mathcal{N}}$, and the shtuka condition promotes it to a morphism to $\operatorname{Sht}^r_{\mathcal{N}}$. We now refine by pulling back the locally closed stratification $\{{}_\tau\operatorname{Hk}^r_{\mathcal{N}}\}_\tau$; on each refined stratum the type $\tau(i) \in \{0, +, -, \pm\}$ is constant for every leg. It suffices to show that each such refined stratum has dimension $\le r(n-m)$.

Pick any such stratum $\mathcal{S}$, and let $\tau = (I_0, I_+, I_-, I_\pm)$ be its type. Then we obtain a map $\mathcal{S} \to {}_\tau\operatorname{Sht}^r_{\mathcal{N}}$.

Lemma 4.4.4. *The map*

$$\mathcal{S} \to {}_\tau\operatorname{Sht}^r_{\mathcal{N}} \tag{4.4.5}$$

is quasi-finite.

Proof. Consider the composition

$$\mathcal{S} \to {}_\tau \operatorname{Sht}^r_{\mathcal{N}} \to \operatorname{Sht}^r_{U(n)}.$$

The composition $\mathcal{S} \to \operatorname{Sht}^r_{U(n)}$ is obtained from the coefficient-a injective-locus map $\mathcal{Z}^r_{\mathcal{E}}(a)^\circ \to \operatorname{Sht}^r_{U(n)}$ by restricting to a locally closed stratum. Since that map is quasi-finite by [FYZ24, Proposition 7.5], the composition $\mathcal{S} \to \operatorname{Sht}^r_{U(n)}$ is quasi-finite. Because this composition factors through the map (4.4.5), every fiber of (4.4.5) is contained in a fiber of $\mathcal{S} \to \operatorname{Sht}^r_{U(n)}$. Hence (4.4.5) is quasi-finite. □

Completion of the proof of Theorem 4.1.1. Thanks to Lemma 4.4.4, and after pulling back along smooth charts of the target so that quasi-finite maps do not increase dimension, it suffices to show that if ${}_\tau \operatorname{Sht}^r_{\mathcal{N}}$ is non-empty, then $\dim({}_\tau \operatorname{Sht}^r_{\mathcal{N}}) \le r(n-m)$. Let $\tau = (I_0, I_+, I_-, I_\pm)$ and

$$u := |I_+| = |I_-|, \qquad b := |I_\pm|$$

where the equality $|I_+| = |I_-|$ is Lemma 4.4.1.

Corollary 4.4.3 gives

$$\dim({}_\tau \operatorname{Sht}^r_{\mathcal{N}}) \le (n-m)|I_0| + (n+m)u + mb. \tag{4.4.6}$$

Since $r = |I_0| + 2u + b$, we have

$$\begin{aligned}(n-m)r - \big((n-m)|I_0| + (n+m)u + mb\big) &= (n-m)(|I_0| + 2u + b) - (n-m)|I_0| - (n+m)u - mb \\ &= (n-3m)u + (n-2m)b. \end{aligned} \tag{4.4.7}$$

The integers u and b are nonnegative. The hypothesis $3m \le n$ implies $n - 3m \ge 0$, and also $n - 2m = (n-3m) + m \ge 0$. Hence (4.4.7) is nonnegative, giving

$$(n-m)|I_0| + (n+m)u + mb \le (n-m)r.$$

Together with (4.4.6), this proves $\dim({}_\tau \operatorname{Sht}^r_{\mathcal{N}}) \le (n-m)r$. □

5. Proof of the Trace Conjecture in low corank

In [FK24, §4–6], the *motivic sheaf-cycle correspondence* is developed, refining the ℓ-adic version developed in [FYZ23, §3–5]. We use the notation and terminology of [FK24], including:

- The notion of *cohomological correspondences* between motivic sheaves.
- The notion of *pushable/pullable* squares and *pushable/pullable* maps of correspondences.
- The definition of *trace* of a self-cohomological correspondence between *geometric* motivic sheaves.
- The shtuka-twisted trace.

The purpose of this section is to prove the low-corank form of the Trace Conjecture, Theorem 5.1.3, which realizes the derived fundamental classes of special cycles as the shtuka-twisted traces of certain cohomological correspondences. Notably, we expect that the argument will generalize well to other settings, such as orthosymplectic groups.

5.1. **Formulation of the Trace Conjecture.** Let $\mathcal{M} := \mathcal{M}_{H_1,H_2}$ be a derived Hitchin stack as in [FYZ25, Definition 5.23]. Here BH_1 is a smooth m-framed gerbe over X and BH_2 is a smooth n-framed gerbe of unitary type over X. Let $\operatorname{Hk}^r_{\mathcal{M}}$ be the derived Hecke stack for $\mathcal{M}$ as defined in [FYZ25, §5.6]. Then we have a correspondence

$$\mathcal{M} \xleftarrow{h_0} \operatorname{Hk}^r_{\mathcal{M}} \xrightarrow{h_r} \mathcal{M}$$

Lemma 5.1.1. *The map $h_r\colon \operatorname{Hk}^r_{\mathcal{M}} \to \mathcal{M}$ is quasi-smooth.*

Proof. By [FYZ25, Proposition 5.35], $\mathcal{M}$ is the derived vector bundle over $\operatorname{Bun}_{H_1} \times \operatorname{Bun}_{H_2}$ associated to a certain perfect complex $\mathcal{K}$, and $\operatorname{Hk}^r_{\mathcal{M}}$ is the derived vector bundle over $\operatorname{Bun}_{H_1} \times \operatorname{Hk}^r_{H_2}$ associated to a certain perfect complex $\mathcal{K}^\flat$. Let $\widetilde{\mathcal{M}} := \operatorname{Tot}_{\operatorname{Bun}_{H_1} \times \operatorname{Hk}^r_{H_2}}(h_r^*\mathcal{K})$ be the base change of $\mathcal{M}$ to $\operatorname{Bun}_{H_1} \times \operatorname{Hk}^r_{H_2}$. Since the map $\operatorname{Hk}^r_{H_2} \to \operatorname{Bun}_{H_2}$ is smooth by [FYZ25, Lemma 2.6], the map $\widetilde{\mathcal{M}} \to \mathcal{M}$ is smooth. Hence it suffices to show that the map $\operatorname{Hk}^r_{\mathcal{M}} \to \widetilde{\mathcal{M}}$ is quasi-smooth. This follows from the observation that the cone of $\mathcal{K}^\flat \to h_r^*\mathcal{K}$ is represented by a locally free sheaf, since it arises as the relative cohomology of X' with coefficients in a torsion coherent sheaf, as explained in the proof of [FYZ23, Lemma 9.1.3(1)].

Equivalently, $\mathrm{Hk}^r_{\mathcal{M}} \to \widetilde{\mathcal{M}}$ is the derived base change of the zero section of the vector bundle associated to this cone, so it is quasi-smooth. □

Consequently, the map of correspondences

$$\begin{array}{ccccc} \mathcal{M} & \xleftarrow{h_0} & \mathrm{Hk}^r_{\mathcal{M}} & \xrightarrow{h_r} & \mathcal{M} \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{pt} & \longleftarrow & \mathrm{pt} & \longrightarrow & \mathrm{pt} \end{array}$$

is right pullable, so the pullback $\mathrm{Corr}_{\mathrm{pt}}(\overline{\mathbf{Q}}_{\mathrm{pt}}, \overline{\mathbf{Q}}_{\mathrm{pt}}) \to \mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}}}(\overline{\mathbf{Q}}_{\mathcal{M}}, \overline{\mathbf{Q}}_{\mathcal{M}\langle -d(h_r)\rangle})$ exists. We define $\mathfrak{c}_{\mathcal{M}} \in \mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}}}(\overline{\mathbf{Q}}_{\mathcal{M}}, \overline{\mathbf{Q}}_{\mathcal{M}\langle -d(h_r)\rangle})$ to be the pullback of the identity cohomological correspondence $\mathfrak{c}_{\mathrm{pt}} \in \mathrm{Corr}_{\mathrm{pt}}(\overline{\mathbf{Q}}_{\mathrm{pt}}, \overline{\mathbf{Q}}_{\mathrm{pt}})$. More explicitly, $\mathfrak{c}_{\mathcal{M}}$ is the composition

$$h_0^*\overline{\mathbf{Q}}_{\mathcal{M}} \cong \overline{\mathbf{Q}}_{\mathrm{Hk}^r_{\mathcal{M}}} \cong h_r^*\overline{\mathbf{Q}}_{\mathcal{M}} \xrightarrow{\mathrm{gys}_{h_r}} h_r^!\overline{\mathbf{Q}}_{\mathcal{M}\langle -d(h_r)\rangle}.$$

Conjecture 5.1.2 (The Trace Conjecture for Hitchin spaces, [FH25, §9.4])**.** *For any $r \geq 0$, and for any $m \in [1, n]$, we have*

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}}) = [\mathrm{Sht}^r_{\mathcal{M}}] \in \mathrm{CH}_{d(h_r)}(\mathrm{Sht}^r_{\mathcal{M}}). \tag{5.1.1}$$

Theorem 5.1.3 (Trace Conjecture in low corank)**.** *For any $r \geq 0$ and for $m \leq n/3$,* (5.1.1) *holds true.*

5.2. **Localized version.** We formulate an equivalent version of Conjecture 5.1.2 which is localized over a point of Bun_{H_1}. The stack $\mathcal{M}$ comes equipped with maps to Bun_{H_1}, such that

$$\mathrm{Hk}^r_{\mathcal{M}} \xrightarrow{h_i} \mathcal{M} \to \mathrm{Bun}_{H_1} \text{ is independent of } i \in \{0, \dots, r\}. \tag{5.2.1}$$

Let $\mathcal{E} \in \mathrm{Bun}_{H_1}(\mathbf{F}_q)$, let $\mathcal{M}_{\mathcal{E}}$ be the derived fiber of $\mathcal{M}$ over $\mathcal{E}$, and let $\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}$ be the derived fiber of $\mathrm{Hk}^r_{\mathcal{M}}$ over $\mathcal{E}$.

By Lemma 5.1.1, the base changed map $h_r\colon \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}} \to \mathcal{M}_{\mathcal{E}}$ is quasi-smooth, so we can define

$$\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}} \in \mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}}(\overline{\mathbf{Q}}_{\mathcal{M}_{\mathcal{E}}}, \overline{\mathbf{Q}}_{\mathcal{M}_{\mathcal{E}}\langle -d(h_r)\rangle})$$

to be the pullback of the identity correspondence in $\mathrm{Corr}_{\mathrm{pt}}(\overline{\mathbf{Q}}_{\mathrm{pt}}, \overline{\mathbf{Q}}_{\mathrm{pt}})$.

Theorem 5.2.1. *For any $r \geq 0$, and for $m \leq n/3$, we have*

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}) = [\mathcal{Z}^r_{\mathcal{E}}] \in \mathrm{CH}_{d(h_r)}(\mathcal{Z}^r_{\mathcal{E}}). \tag{5.2.2}$$

Lemma 5.2.2. *For all $1 \leq m \leq n$, the identity* (5.1.1) *is equivalent to the collection of identities* (5.2.2) *for all $\mathcal{E} \in \mathrm{Bun}_{H_1}(\mathbf{F}_q)$.*

Proof. Note that $\mathrm{Sht}^r_{\mathcal{M}}$ is fibered over the fixed-point stack of Bun_{H_1}. For $\mathcal{E} \in \mathrm{Bun}_{H_1}(\mathbf{F}_q)$, put $\overline{\mathcal{E}} := [\{\mathcal{E}\}/\operatorname{Aut}(\mathcal{E})]$ and let $\Gamma_{\mathcal{E}} := \operatorname{Aut}(\mathcal{E})(\mathbf{F}_q)$. The corresponding open-and-closed summand of $\mathrm{Sht}^r_{\mathcal{M}}$ is

$$\mathcal{Z}^r_{\overline{\mathcal{E}}} := [\mathcal{Z}^r_{\mathcal{E}}/\Gamma_{\mathcal{E}}],$$

where the quotient is by the finite group of Frobenius-fixed automorphisms of $\mathcal{E}$. This gives the decomposition

$$[\mathrm{Sht}^r_{\mathcal{M}}] = \sum_{\mathcal{E} \in \mathrm{Bun}_{H_1}(\mathbf{F}_q)} [\mathcal{Z}^r_{\overline{\mathcal{E}}}] \in \mathrm{CH}_{d(h_r)}(\mathrm{Sht}^r_{\mathcal{M}}). \tag{5.2.3}$$

The RHS is an infinite sum, but over each connected component of $\mathrm{Sht}^r_{\mathcal{M}}$ only finitely many summands are supported. Equivalently, only finitely many summands are supported on any finite-type truncation.

Consider the locally closed embedding $\overline{\mathcal{E}} \hookrightarrow \mathrm{Bun}_{H_1}$. We write $\mathcal{M}_{\overline{\mathcal{E}}} := [\mathcal{M}_{\mathcal{E}}/\operatorname{Aut}(\mathcal{E})]$ for the stack-theoretic quotient by the full automorphism group scheme of $\mathcal{E}$, and let $\mathfrak{c}_{\mathcal{M}_{\overline{\mathcal{E}}}}$ be the induced pullback of $\mathfrak{c}_{\mathcal{M}}$. By (5.2.1), $\mathrm{Hk}^r_{\mathcal{M}}$ *stabilizes* (cf. [FK24, Definition 7.2.1]) $\mathcal{M}_{\overline{\mathcal{E}}}$, so its Frobenius twist is *contracting* near $\mathcal{M}_{\overline{\mathcal{E}}}$ (cf. [FK24, Example 7.2.2]). Hence [FK24, Theorem 7.5.1] implies that

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}})|_{\mathcal{Z}^r_{\overline{\mathcal{E}}}} = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\overline{\mathcal{E}}}}) \in \mathrm{CH}_{d(h_r)}(\mathcal{Z}^r_{\overline{\mathcal{E}}}). \tag{5.2.4}$$

Let $q_{\mathcal{E}}\colon \mathcal{Z}^r_{\mathcal{E}} \to \mathcal{Z}^r_{\overline{\mathcal{E}}}$ be the natural finite étale $\Gamma_{\mathcal{E}}$-torsor. The smooth $\mathrm{Aut}(\mathcal{E})$-torsor $\mathcal{M}_{\mathcal{E}} \to \mathcal{M}_{\overline{\mathcal{E}}}$ induces a Cartesian map of correspondence data by (5.2.1), which induces $q_{\mathcal{E}}$ by formation of shtuka spaces. Compatibility of trace with smooth pullbacks, [FK24, Proposition 6.2.2], gives

$$q_{\mathcal{E}}^* \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\overline{\mathcal{E}}}}) = \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}) \in \mathrm{CH}_{d(h_r)}(\mathcal{Z}^r_{\mathcal{E}}). \tag{5.2.5}$$

Moreover $q_{\mathcal{E}}^*[\mathcal{Z}^r_{\overline{\mathcal{E}}}] = [\mathcal{Z}^r_{\mathcal{E}}]$, and with $\overline{\mathbf{Q}}$-coefficients the transfer identity $q_{\mathcal{E},!}q_{\mathcal{E}}^* = |\Gamma_{\mathcal{E}}|\,\mathrm{id}$ makes $q_{\mathcal{E}}^*$ injective. Hence

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\overline{\mathcal{E}}}}) = [\mathcal{Z}^r_{\overline{\mathcal{E}}}] \iff \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}) = [\mathcal{Z}^r_{\mathcal{E}}]. \tag{5.2.6}$$

If (5.1.1) holds, then we obtain (5.2.2) by restricting to $\mathcal{Z}^r_{\overline{\mathcal{E}}}$ and then pulling back along $q_{\mathcal{E}}$, using (5.2.4) and (5.2.5) to compute the restrictions. Conversely, if (5.2.2) holds for every $\mathcal{E} \in \mathrm{Bun}_{H_1}(\mathbf{F}_q)$, then (5.2.6) gives the desired equality on every open-and-closed summand $\mathcal{Z}^r_{\overline{\mathcal{E}}}$; summing these identities by (5.2.3) gives (5.1.1).

□

Henceforth we will focus on proving Theorem 5.2.1. (In any case, this is the version that we will actually use later.)

Remark 5.2.3. The split-cover refinement needed for supermodularity is stated separately as Theorem 13.2.1 in Part 4.

5.3. **The injective locus.** Recall that $\mathcal{M}_{\mathcal{E}}$ parametrizes an H_2-bundle $\mathcal{F}$ along with a map $t\colon \mathcal{E} \to \mathcal{F}$. Let $\mathcal{M}^\circ_{\mathcal{E}} \subset \mathcal{M}_{\mathcal{E}}$ be the locus where the map t is injective (as a map of coherent sheaves) after restriction to all geometric points of the test scheme. This is the fiber over $\{\mathcal{E}\}$ of [FYZ25, Definition 3.4]. This is an open subset preserved by the correspondence $\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}$, in the sense that the restriction gives a map of correspondences

$$\begin{array}{ccccc} \mathcal{M}^\circ_{\mathcal{E}} & \xleftarrow{h_0} & \mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}}} & \xrightarrow{h_r} & \mathcal{M}^\circ_{\mathcal{E}} \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{E}} & \xleftarrow{h_0} & \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}} & \xrightarrow{h_r} & \mathcal{M}_{\mathcal{E}} \end{array}$$

with both squares Cartesian, hence pushable and pullable. In particular, we have the pullback cohomological correspondence $\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}}} \in \mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}}}}(\overline{\mathbf{Q}}_{\mathcal{M}^\circ_{\mathcal{E}}}, \overline{\mathbf{Q}}_{\mathcal{M}^\circ_{\mathcal{E}}}\langle -d(h_r)\rangle)$ of $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}$.

Proposition 5.3.1. *For $m \le n/3$, we have*

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}}}) = [\mathcal{Z}^{r,\circ}_{\mathcal{E}}] \in \mathrm{CH}_{r(n-m)}(\mathcal{Z}^{r,\circ}_{\mathcal{E}}). \tag{5.3.1}$$

Proof. We can write $\mathcal{Z}^{r,\circ}_{\mathcal{E}}$ as the increasing union of Harder–Narasimhan truncations $\mathcal{Z}^{r,\circ,\le\nu}_{\mathcal{E}}$. Here, for any of the fibers $\mathcal{M}_\star$ below, the bounded Hecke correspondence means

$$\mathrm{Hk}^{r,\le\nu}_{\mathcal{M}_\star} := \bigcap_{i=0}^{r} h_i^{-1}(\mathcal{M}^{\le\nu}_\star),$$

and $\mathcal{Z}^{r,\le\nu}_\star$ is its shtuka fixed-point stack. Thus $\mathfrak{c}_{\mathcal{M}^{\le\nu}_\star}$ denotes the open restriction of $\mathfrak{c}_{\mathcal{M}_\star}$ to this bounded correspondence; a single Hecke modification need not preserve an HN bound, but these bounded open correspondences exhaust $\mathrm{Hk}^r_{\mathcal{M}_\star}$ as ν varies. The argument below uses the identification

$$\mathrm{CH}_{r(n-m)}(\mathcal{Z}^{r,\circ}_{\mathcal{E}}) = \varprojlim_{\nu} \mathrm{CH}_{r(n-m)}(\mathcal{Z}^{r,\circ,\le\nu}_{\mathcal{E}}), \tag{5.3.2}$$

so it suffices to show (5.3.1) after pullback to every Harder–Narasimhan truncation.

Fix the HN truncation ν. We first choose a closed point $x_0 \in |X|$ as follows. By Theorem 4.1.1, $\dim \mathcal{Z}^{r,\circ,\le\nu}_{\mathcal{E}} \le r(n-m)$, and the finite-type truncation has finitely many irreducible components of dimension $r(n-m)$. For each such component and each leg projection to X', either the projection is dominant or its image is a single point of X'. Therefore, there is a finite subset $B \subset |X'|$ of points that can occur as the image of a non-dominant leg projection from an $r(n-m)$-dimensional component. Choose x_0 such that $\nu^{-1}(x_0) \cap B = \varnothing$; choosing x_0 on X makes every divisor $D = dx_0$ and $\nu^* D$ defined

over $\mathbf{F}_q$ and σ-invariant. Given this x_0, we may choose $d \gg 0$ so that for $\widetilde{\mathcal{E}} := \mathcal{E}(-\nu^* D)$, $\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu}$ is smooth, cf. [FYZ23, §10.1.1], so then [FK24, Corollary 6.4.1] applies to the bounded correspondence above to yield

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu}}) = [\mathcal{Z}_{\widetilde{\mathcal{E}}}^{r,\leq \nu}] \in \mathrm{CH}_*(\mathcal{Z}_{\widetilde{\mathcal{E}}}^{r,\leq \nu}).$$

Then compatibility of trace with open restrictions, the open-embedding case of [FK24, Proposition 6.2.2], implies that

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu,\circ}}) = [\mathcal{Z}_{\widetilde{\mathcal{E}}}^{r,\circ,\leq \nu}] \in \mathrm{CH}_*(\mathcal{Z}_{\widetilde{\mathcal{E}}}^{r,\circ,\leq \nu}).$$

Let $X_0' := X' \setminus \nu^{-1}(x_0)$. We observe that $\mathcal{M}_{\mathcal{E}}^{\circ} \to \mathcal{M}_{\widetilde{\mathcal{E}}}^{\circ}$ is a closed embedding stabilized by $\mathrm{Hk}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\circ}}^{r}|_{(X_0')^r}$. Indeed, over $(X_0')^r$ the modifications are isomorphisms near $\nu^* D$, so extension of the maps from $\widetilde{\mathcal{E}}$ to $\mathcal{E}$ across $\nu^* D$ is a closed condition at every point of the chain. Hence $\mathfrak{c}^{(1)}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu,\circ}}|_{(X_0')^r}$ is contracting near $\mathcal{M}_{\mathcal{E}}^{\circ,\leq \nu}$ by [FK24, Example 7.2.2]. Furthermore, the derived vector bundle description of the Hecke stacks identifies $\mathrm{Hk}_{\mathcal{M}_{\mathcal{E}}^{\circ}}^{r}|_{(X_0')^r}$ with the derived pullback of $\mathrm{Hk}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\circ}}^{r}|_{(X_0')^r}$ along $\mathcal{M}_{\mathcal{E}}^{\circ} \hookrightarrow \mathcal{M}_{\widetilde{\mathcal{E}}}^{\circ}$, and Gysin base change for this derived Cartesian square identifies $\mathfrak{c}^{(1)}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu,\circ}}|_{(X_0')^r}$ with the restriction of $\mathfrak{c}^{(1)}_{\mathcal{M}_{\widetilde{\mathcal{E}}}^{\leq \nu,\circ}}|_{(X_0')^r}$ in the sense of [FK24, Construction 7.2.3]. Hence [FK24, Theorem 7.5.1] applies to yield

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}^{\circ,\leq \nu}}|_{(X_0')^r}) = [\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}|_{(X_0')^r}] \in \mathrm{CH}_{r(n-m)}(\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}|_{(X_0')^r}). \tag{5.3.3}$$

The choice of x_0 above ensures that no $r(n-m)$-dimensional component is contained in the boundary

$$\partial \mathcal{Z} := \mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu} \setminus \left(\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}|_{(X_0')^r}\right).$$

Thus, we have $\dim \partial \mathcal{Z} < r(n-m)$ and hence $\mathrm{CH}_{r(n-m)}(\partial \mathcal{Z}) = 0$. By the excision sequence, this implies that the restriction map

$$\mathrm{CH}_{r(n-m)}(\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}) \to \mathrm{CH}_{r(n-m)}(\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}|_{(X_0')^r})$$

is an isomorphism. Therefore, the desired equality (5.3.1) can be checked after restriction to $\mathcal{Z}_{\mathcal{E}}^{r,\circ,\leq \nu}|_{(X_0')^r}$. After such restriction, it becomes (5.3.3), so we are done. □

5.4. **Twisting by Chow cohomology.** Let C be a derived Artin stack. We use the Chow cohomology groups $\mathrm{CH}^i(C)$ defined in §3.5. Note that the ith Chern class of a vector bundle is valued in $\mathrm{CH}^i(C)$.

Suppose $L \in \mathcal{D}_{\mathrm{mot}}(C; \overline{\mathbf{Q}})$ is dualizable, and $c \in \mathrm{Hom}_C(L, L\langle i \rangle)$. Then we may form the *trace* of c as an element of Chow cohomology, $\mathrm{Tr}(c) \in \mathrm{CH}^i(C)$.

Example 5.4.1. Note that if L is $\otimes$-invertible, then we have

$$c \in \mathrm{Hom}_C(L, L\langle i \rangle) \cong \mathrm{Hom}_C(L \otimes L^\vee, L \otimes L^\vee \langle i \rangle) \cong \mathrm{Hom}_C(\overline{\mathbf{Q}}_C, \overline{\mathbf{Q}}_C \langle i \rangle) = \mathrm{CH}^i(C).$$

In this case, $\mathrm{Tr}(c)$ agrees with the image of c in $\mathrm{CH}^i(C)$.

A cohomological correspondence $\mathfrak{c} \in \mathrm{Corr}_C(\mathcal{K}_0, \mathcal{K}_1\langle -d \rangle)$ can be twisted by $c \in \mathrm{Hom}_C(L, L\langle i \rangle)$ to obtain $\mathfrak{c} \otimes c \in \mathrm{Corr}_C(\mathcal{K}_0 \otimes L, (\mathcal{K}_1 \otimes L)\langle -d+i \rangle)$. Recall also that Chow groups form a module over Chow cohomology.

Lemma 5.4.2. *Let $A \leftarrow C \to A$ be a correspondence of derived Artin stacks, and let $\rho_C \colon \mathrm{Fix}(C) \to C$ be the natural map. Let $\mathcal{K} \in \mathcal{D}_{\mathrm{mot,gm}}(A; \overline{\mathbf{Q}})$ and $\mathfrak{c} \in \mathrm{Corr}_C(\mathcal{K}, \mathcal{K}\langle j \rangle)$ for any $j \in \mathbf{Z}$. Let $c \in \mathrm{CH}^i(C) = \mathrm{Hom}_C(\overline{\mathbf{Q}}_C, \overline{\mathbf{Q}}_C\langle i \rangle)$, viewed as an endomorphism of the unit object on C. Then we have*

$$\mathrm{Tr}_C(\mathfrak{c} \otimes c) = \mathrm{Tr}_C(\mathfrak{c}) \cap \rho_C^* c \in \mathrm{CH}_*(\mathrm{Fix}(C)).$$

Proof. Clear from the definition. □

Example 5.4.3. Taking $L = \overline{\mathbf{Q}}_C$ in Example 5.4.1, and applying Lemma 5.4.2, we learn that cohomological correspondences can be tensored with Chow cohomology, and that formation of trace is linear over such twisting.

5.5. **Degeneracy stratification.** For each saturated subbundle $\mathcal{K} \subset \mathcal{E}$, there is a closed substack defined by the condition that $t\colon \mathcal{E} \to \mathcal{F}$ kills $\mathcal{K}$, equivalently where t factors through $\mathcal{E}/\mathcal{K}$. Inside it, let $\mathcal{M}_{\mathcal{E}}[\mathcal{K}]$ be the open substack where the induced map $\mathcal{E}/\mathcal{K} \to \mathcal{F}$ is injective as a map of coherent sheaves after restriction to every geometric point of the test scheme, as for $\mathcal{M}^\circ_{\mathcal{E}}$ above. This closed substack is stable under $\mathrm{Hk}^r_{\mathcal{M}}$.

Thus, for each $\mathcal{K}$ we have a map of correspondences

$$\begin{array}{ccccc} \mathcal{M}_{\mathcal{E}}[\mathcal{K}] & \longleftarrow & \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]} & \longrightarrow & \mathcal{M}_{\mathcal{E}}[\mathcal{K}] \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{E}} & \longleftarrow & \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}} & \longrightarrow & \mathcal{M}_{\mathcal{E}} \end{array} \tag{5.5.1}$$

where both squares are Cartesian. Notice that we have obvious maps $\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} \to \mathcal{M}_{\mathcal{E}}[\mathcal{K}]$ and $\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \to \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}$, which are each isomorphisms on classical truncations. We will use nil-invariance for motivic Chow groups and cohomological correspondences to identify the corresponding groups for these derived stacks having the same classical truncation. The commutative square

$$\begin{array}{ccc} \mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} & \longrightarrow & \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} \\ \downarrow & & \downarrow \\ \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}} & \longrightarrow & \mathcal{M}_{\mathcal{E}} \end{array} \tag{5.5.2}$$

is not Cartesian (unless $\mathcal{K} = 0$), but it is Cartesian on classical truncations: on classical points, $t_i(\mathcal{K})$ is generically zero and hence is zero as a subsheaf of the vector bundle $\mathcal{F}_i$. Furthermore, both horizontal maps are quasi-smooth, by Lemma 5.1.1.

Lemma 5.5.1. *The commutative square* (5.5.2) *is an* excess intersection square *in the sense of* [Kha19, Proposition 3.15]*, with excess bundle* $\mathcal{V}^r_{\mathcal{K}}$ *of* [FYZ25, §5.8.1].

Proof. We have already observed that the horizontal maps in (5.5.2) are quasi-smooth and the square is Cartesian on classical truncations. Let $f\colon \mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \to \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}$ be the left vertical map in (5.5.2). It remains to identify the (derived) fiber of the map

$$\mathbf{T}_{\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}/\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \to f^*\mathbf{T}_{\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}/\mathcal{M}_{\mathcal{E}}} \tag{5.5.3}$$

with $\mathcal{V}^r_{\mathcal{K}}[-2]$. Consider the commutative diagram

$$\begin{array}{ccccc} \mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} & \longrightarrow & \mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}} & \longrightarrow & \mathrm{Hk}^r_{H_2} \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} & \longrightarrow & \mathcal{M}_{\mathcal{E}} & \longrightarrow & \mathrm{Bun}_{H_2} \end{array}$$

This induces a commutative diagram of perfect complexes on $\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}$,

$$\begin{array}{ccc} \mathbf{T}_{\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}/\mathrm{Hk}^r_{H_2}} & \longrightarrow & f^*\mathbf{T}_{\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}/\mathrm{Hk}^r_{H_2}} \\ \downarrow & & \downarrow \\ \mathbf{T}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}/\mathrm{Bun}_{H_2}} & \longrightarrow & f^*\mathbf{T}_{\mathcal{M}_{\mathcal{E}}/\mathrm{Bun}_{H_2}} \end{array} \tag{5.5.4}$$

The derived fibers along the columns of (5.5.4) differ from the two terms in (5.5.3) by the same smooth summand pulled back from $\mathbf{T}_{\mathrm{Hk}^r_{H_2}/\mathrm{Bun}_{H_2}}$. The comparison is the identity on this common summand, so the derived fiber of (5.5.3) is obtained by cancelling it and then interchanging the order of forming derived fibers in (5.5.4): first taking the derived fibers along the rows, and then the derived fiber along the column.

Referring to the notation of [FYZ25, §5.8], [FYZ25, Proposition 5.35] implies that $\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}} \to \mathrm{Hk}^r_{H_2}$ is the derived vector bundle associated to $\underline{\mathrm{RHom}_{X'}(\mathcal{E}, \mathcal{F}^\flat_\star)}$, and $\mathcal{M}_\mathcal{E} \to \mathrm{Bun}_{H_2}$ is the derived vector bundle associated to $\underline{\mathrm{RHom}_{X'}(\mathcal{E}, \mathcal{F})}$, and similarly for $\mathcal{E}$ replaced by $\mathcal{E}/\mathcal{K}$. Hence the derived fibers along the rows of (5.5.4) are

$$\begin{array}{c} \underline{\mathrm{RHom}_{X'}(\mathcal{K}, \mathcal{F}^\flat_\star)}[-1] \\ \downarrow \\ \underline{\mathrm{RHom}_{X'}(\mathcal{K}, \mathcal{F})}[-1] \end{array}$$

and then forming the derived fiber along the column gives $\mathcal{V}^r_\mathcal{K}[-2]$, as desired.

□

Corollary 5.5.2. *We have*

$$\mathfrak{c}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} = \mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \otimes c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K}) \in \mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}}(\overline{\mathbf{Q}}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}, \overline{\mathbf{Q}}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}\langle -d(h_r)\rangle).$$

Proof. We can factor the diagram (5.5.2) through the right half of (5.5.1), obtaining a commutative square

$$\begin{array}{ccc} \mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} & \longrightarrow & \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} \\ \downarrow & & \downarrow \\ \mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} & \longrightarrow & \mathcal{M}_\mathcal{E}[\mathcal{K}] \end{array} \tag{5.5.5}$$

in which the vertical maps are isomorphisms on classical truncations. Since (5.5.2) factors through the derived Cartesian right square of (5.5.1), Lemma 5.5.1 shows that (5.5.5) is also an excess intersection square, with excess bundle $\mathcal{V}^r_\mathcal{K}$. Then the excess intersection formula [FYZ25, §6.1.12] applies to yield

$$[\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} / \mathcal{M}_\mathcal{E}[\mathcal{K}]] = c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K})[\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} / \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}] \in \mathrm{CH}_{d(h_r)}(\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} / \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}).$$

Under the isomorphism

$$\mathrm{Corr}_{\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}}(\overline{\mathbf{Q}}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}, \overline{\mathbf{Q}}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}\langle -d(h_r)\rangle) \cong \mathrm{CH}_{d(h_r)}(\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} / \mathcal{M}_\mathcal{E}[\mathcal{K}])$$

which is compatible with Gysin base change along the derived Cartesian squares above, the cohomological correspondence $\mathfrak{c}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}$ maps to $[\mathrm{Hk}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} / \mathcal{M}_\mathcal{E}[\mathcal{K}]]$. Under the nil-invariance identification with the top row, the correspondence $\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \otimes c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K})$ maps to $c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K}) \cap [\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} / \mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}]$. The excess formula identifies these two classes, completing the proof. □

Remark 5.5.3. The natural map $\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} \to \mathcal{M}_\mathcal{E}[\mathcal{K}]$ is an isomorphism on classical truncations: an injective map $\mathcal{E}/\mathcal{K} \to \mathcal{F}$ is the same thing as a map $\mathcal{E} \to \mathcal{F}$ with kernel exactly $\mathcal{K}$. Thus on classical truncations we have $\mathrm{Sht}^r_{\mathcal{M}_\mathcal{E}[\mathcal{K}]} \cong \mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}$. This is the identification used in Proposition 5.5.4.

Proposition 5.5.4. *Assume that* $\mathrm{rank}(\mathcal{E}/\mathcal{K}) \leq n/3$. *Let* $\rho_\mathcal{K}$ *be the natural map from* $\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}$ *to* $\mathrm{Hk}^r_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}$. *Then we have*

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}) = \rho^*_\mathcal{K} c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K}) \cap [\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}] \in \mathrm{CH}_*(\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}).$$

Proof. By Corollary 5.5.2 and then Lemma 5.4.2, we have

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_\mathcal{E}[\mathcal{K}]}) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}} \otimes c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K})) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}) \cap \rho^*_\mathcal{K} c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_\mathcal{K}) \in \mathrm{CH}_*(\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}).$$

If $\mathcal{E}/\mathcal{K} = 0$, then $\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}} = \mathrm{Bun}_{H_2}$ and $\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}} = \mathrm{Sht}^r_{H_2}$. The corresponding cohomological correspondence is the canonical Hecke correspondence, so [FK24, Corollary 6.4.1] gives $\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}) = [\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}]$. Hence assume that $\mathrm{rank}(\mathcal{E}/\mathcal{K}) > 0$. Then we may apply Proposition 5.3.1, which identifies

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}^\circ_{\mathcal{E}/\mathcal{K}}}) = [\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}] \in \mathrm{CH}_*(\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}).$$

□

5.6. **Completion of the proof.** Let $\mathcal{M}_{\geq d} \subset \mathcal{M}_{\mathcal{E}}$ be the closed substack where $\operatorname{rank}(\ker(t\colon \mathcal{E} \to \mathcal{F})) \geq d$ and let $\mathcal{M}_{\leq d} \subset \mathcal{M}_{\mathcal{E}}$ be the open substack where $\operatorname{rank}(\ker(t\colon \mathcal{E} \to \mathcal{F})) \leq d$. These are the usual determinantal rank loci for the map t on geometric fibers. They intersect in the locally closed substack $\mathcal{M}_d \hookrightarrow \mathcal{M}_{\mathcal{E}}$.

Furthermore, $j\colon \mathcal{M}_d \hookrightarrow \mathcal{M}_{\geq d}$ is an open embedding, while $i\colon \mathcal{M}_{\geq d+1} \hookrightarrow \mathcal{M}_{\geq d}$ is a closed embedding, and both are stable under $\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}$: along a Hecke chain the kernels $\ker(t_i) \subset \mathcal{E}$ have the same generic fiber, hence the same saturation and the same rank. Thus the pullbacks i^* and j^* exist on cohomological correspondences.

Denote by $\mathrm{Sht}^r_{\mathcal{M}_d}$ and $\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}$ the corresponding moduli stacks of shtukas. By [FK24, Proposition 7.5.4], the maps $\mathrm{Sht}^r_{\mathcal{M}_d} \to \mathrm{Sht}^r_{\mathcal{M}_{\mathcal{E}}} = \mathcal{Z}^r_{\mathcal{E}}$ are open-closed for all d.

For any stable locally closed substack $\mathcal{N} \subset \mathcal{M}_{\mathcal{E}}$ (in this subsection, $\mathcal{N}$ will be one of $\mathcal{M}_d$, $\mathcal{M}_{\geq d}$, $\mathcal{M}_{\geq d+1}$, or $\mathcal{M}_{\mathcal{E}}[\mathcal{K}]$), write $\mathrm{Hk}^r_{\mathcal{N}}$ for the derived preimage of $\mathcal{N}$ under the right leg h_r of $\mathrm{Hk}^r_{\mathcal{M}_{\mathcal{E}}}$; by stability, the left leg factors through $\mathcal{N}$. We write $\mathfrak{c}_{\mathcal{N}}$ for the pullback of $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}$ to $\mathrm{Hk}^r_{\mathcal{N}}$.

For $d < m$, decompose $\mathcal{M}_{\geq d} = \mathcal{M}_d \sqcup \mathcal{M}_{\geq d+1}$ and compute in $\mathrm{CH}_*(\mathrm{Sht}^r_{\mathcal{M}_{\geq d}})$ as follows.

$$\begin{aligned}
\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\geq d}}) &\overset{(1)}{=} \mathrm{Tr}^{\mathrm{Sht}}(j_! j^* \mathfrak{c}_{\mathcal{M}_{\geq d}}) + \mathrm{Tr}^{\mathrm{Sht}}(i_* i^* \mathfrak{c}_{\mathcal{M}_{\geq d}}) \\
&\overset{(2)}{=} \mathrm{Tr}^{\mathrm{Sht}}(j_! \mathfrak{c}_{\mathcal{M}_d}) + \mathrm{Tr}^{\mathrm{Sht}}(i_* \mathfrak{c}_{\mathcal{M}_{\geq d+1}}) \\
&\overset{(3)}{=} \mathrm{Sht}(j)_!(\mathrm{Tr}^{\mathrm{Sht}}(j_! \mathfrak{c}_{\mathcal{M}_d})|_{\mathrm{Sht}^r_{\mathcal{M}_d}}) + \mathrm{Sht}(i)_!(\mathrm{Tr}^{\mathrm{Sht}}(j_! \mathfrak{c}_{\mathcal{M}_d})|_{\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}}) \\
&\qquad + \mathrm{Sht}(j)_!(\mathrm{Tr}^{\mathrm{Sht}}(i_* \mathfrak{c}_{\mathcal{M}_{\geq d+1}})|_{\mathrm{Sht}^r_{\mathcal{M}_d}}) + \mathrm{Sht}(i)_!(\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\geq d+1}})|_{\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}}). \qquad (5.6.1)
\end{aligned}$$

The three numbered equalities are justified as follows.

(1) Equality (1) is trace additivity; the proof of [Jin24, Lemma 5.2.7] applies verbatim. Both legs of the Frobenius-twisted correspondence preserve the two strata, so $j_!, j^*, i_*, i^*$ act on the restricted correspondences.

(2) Equality (2) uses the derived Cartesian restriction squares

$$\begin{array}{ccc} \mathrm{Hk}^r_{\mathcal{M}_d} & \longrightarrow & \mathcal{M}_d \\ \downarrow & & \downarrow j \\ \mathrm{Hk}^r_{\mathcal{M}_{\geq d}} & \longrightarrow & \mathcal{M}_{\geq d} \end{array} \qquad \begin{array}{ccc} \mathrm{Hk}^r_{\mathcal{M}_{\geq d+1}} & \longrightarrow & \mathcal{M}_{\geq d+1} \\ \downarrow & & \downarrow i \\ \mathrm{Hk}^r_{\mathcal{M}_{\geq d}} & \longrightarrow & \mathcal{M}_{\geq d} \end{array}$$

which are Cartesian by the derived-preimage convention for $\mathrm{Hk}^r_{\mathcal{N}}$ and identify $j^* \mathfrak{c}_{\mathcal{M}_{\geq d}} = \mathfrak{c}_{\mathcal{M}_d}$ and $i^* \mathfrak{c}_{\mathcal{M}_{\geq d}} = \mathfrak{c}_{\mathcal{M}_{\geq d+1}}$.

(3) Equality (3) uses compatibility of $\mathrm{Tr}^{\mathrm{Sht}}$ with proper pushforward [FK24, Proposition 6.2.1]; both $\mathrm{Sht}(j)$ and $\mathrm{Sht}(i)$ are open-and-closed embeddings. It also uses the decomposition

$$\mathrm{Sht}^r_{\mathcal{M}_{\geq d}} = \mathrm{Sht}^r_{\mathcal{M}_d} \coprod \mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}.$$

We have

$$\mathrm{Sht}^r_{\mathcal{M}_d} = \coprod_{\substack{\mathcal{K} \subset \mathcal{E},\ \mathbf{F}_q\text{-rational} \\ \operatorname{rank}(\mathcal{K}) = d}} \mathrm{Sht}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}. \qquad (5.6.2)$$

The four terms in the last line of (5.6.1) simplify as follows.

- The third summand vanishes because $\mathrm{Tr}^{\mathrm{Sht}}(i_* \mathfrak{c}_{\mathcal{M}_{\geq d+1}}) = \mathrm{Sht}(i)_! \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\geq d+1}})$ is supported on the disjoint open-and-closed summand $\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}$.
- For the first summand, compatibility of trace with pullback along open embeddings gives

$$\mathrm{Tr}^{\mathrm{Sht}}(j_! \mathfrak{c}_{\mathcal{M}_d})|_{\mathrm{Sht}^r_{\mathcal{M}_d}} = \mathrm{Tr}^{\mathrm{Sht}}(j^* j_! \mathfrak{c}_{\mathcal{M}_d}) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_d}).$$

The decomposition (5.6.2) then gives

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_d}) = \sum_{\substack{\mathcal{K} \subset \mathcal{E},\ \mathbf{F}_q\text{-rational} \\ \operatorname{rank}(\mathcal{K}) = d}} (i_{\mathcal{K}})_! i_{\mathcal{K}}^* \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_d}) \qquad (5.6.3)$$

where $i_{\mathcal{K}}$ is the (open and closed) inclusion of $\mathrm{Sht}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}$ into $\mathrm{Sht}^r_{\mathcal{M}_d}$.

The correspondence $\mathrm{Hk}^r_{\mathcal{M}_d}$ stabilizes $\mathcal{M}_{\mathcal{E}}[\mathcal{K}]$. In this Cartesian setting, the restriction of [FK24, Construction 7.2.3] agrees with the pullback defining $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}$. Hence [FK24, Example 7.2.2] and [FK24, Theorem 7.5.1] give

$$i_{\mathcal{K}}^* \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_d}) = \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}).$$

Proposition 5.5.4 identifies this class as

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]}) =: \alpha_{\mathcal{K}} = \rho_{\mathcal{K}}^* c_{\mathrm{rank}(\mathcal{K})r}(\mathcal{V}^r_{\mathcal{K}}) \cap [\mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}}].$$

Substituting these observations into (5.6.3) gives

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_d}) = \sum_{\substack{\mathcal{K}\subset\subset\mathcal{E},\ \mathbf{F}_q\text{-rational}\\ \mathrm{rank}(\mathcal{K})=d}} (i_{\mathcal{K}})_! \alpha_{\mathcal{K}}.$$

- Since $\mathrm{Hk}^r_{\mathcal{M}_{\geq d}}$ stabilizes $\mathcal{M}_{\geq d+1}$, $(\mathrm{Hk}^r_{\mathcal{M}_{\geq d}})^{(1)}$ is contracting near $\mathcal{M}_{\geq d+1}$ by [FK24, Example 7.2.2], and then [FK24, Theorem 7.5.1] shows that

$$\operatorname{Tr}^{\mathrm{Sht}}(j_!\, \mathfrak{c}_{\mathcal{M}_d})|_{\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}} = 0.$$

Write $i_{\mathcal{K},\geq d} := \mathrm{Sht}(j) \circ i_{\mathcal{K}} \colon \mathrm{Sht}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]} \to \mathrm{Sht}^r_{\mathcal{M}_{\geq d}}$. Equation (5.6.1) therefore becomes

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\geq d}}) = \left(\sum_{\substack{\mathcal{K}\subset\subset\mathcal{E},\ \mathbf{F}_q\text{-rational}\\ \mathrm{rank}(\mathcal{K})=d}} (i_{\mathcal{K},\geq d})_! \alpha_{\mathcal{K}} \right) + (\mathrm{Sht}(i)_!\, \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\geq d+1}})|_{\mathrm{Sht}^r_{\mathcal{M}_{\geq d+1}}}) \tag{5.6.4}$$

When $d = m$, we have $\mathcal{M}_{\geq m} = \mathcal{M}_m$ and $\mathcal{M}_{\geq m+1} = \varnothing$, so (5.6.4) holds without its second summand. The identities for $0 \leq d \leq m$ telescope after pushing each class $\alpha_{\mathcal{K}}$ to $\mathrm{CH}_*(\mathcal{Z}^r_{\mathcal{E}})$ along the open-closed map

$$\iota_{\mathcal{K}} \colon \mathcal{Z}^{r,\circ}_{\mathcal{E}/\mathcal{K}} \cong \mathrm{Sht}^r_{\mathcal{M}_{\mathcal{E}}[\mathcal{K}]} \longrightarrow \mathcal{Z}^r_{\mathcal{E}},$$

we obtain

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}_{\mathcal{E}}}) = \sum_{\substack{\mathcal{K}\subset\subset\mathcal{E}\\ \mathcal{K}\ \mathbf{F}_q\text{-rational}}} (\iota_{\mathcal{K}})_! \alpha_{\mathcal{K}},$$

which agrees with $[\mathcal{Z}^r_{\mathcal{E}}]$ by [FYZ25, Theorem 6.6]. This concludes the proof of Theorem 5.2.1. □

Part 2. General tools

6. Motivic derived Fourier transform

In this section we develop a refined version of the homogeneous motivic Fourier transform introduced in [FK24]. Our version can be applied to inhomogeneous sheaves, and agrees with the homogeneous version of *loc. cit.* on homogeneous motivic sheaves. There are two reasons we need this generalization:

- We will need to apply the Fourier transform to inhomogeneous motivic sheaves.
- We will use the fact that the refined Fourier transform is given by a kernel which behaves like a "local system" (is trivial after a finite étale cover), which is not true of the homogeneous Fourier transform.

6.1. **Motivic Artin–Schreier sheaves.** Let $p = \operatorname{char} \mathbf{F}_q$. We recall the additive-character convention fixed in the Notation section: $\psi_0 \colon \mathbf{F}_p \to \overline{\mathbf{Q}}^\times$ is the character used to split the p-Artin–Schreier cover, and $\psi = \psi_0 \circ \mathrm{Tr}_{\mathbf{F}_q/\mathbf{F}_p}$ is the resulting character of $\mathbf{F}_q$. Let

$$f \colon \mathbf{A}^1_{\mathbf{F}_q} \longrightarrow \mathbf{A}^1_{\mathbf{F}_q}, \qquad u \longmapsto u^p - u, \tag{6.1.1}$$

be the Artin–Schreier morphism. It is a finite étale $\mathbf{F}_p$-torsor. With respect to this action, set

$$e_{\psi_0} := \frac{1}{p} \sum_{a \in \mathbf{F}_p} \psi_0(a)^{-1}[a] \in \overline{\mathbf{Q}}[\mathbf{F}_p]. \tag{6.1.2}$$

This is the primitive idempotent projecting to the ψ_0-isotypic summand. The motivic Artin–Schreier sheaf attached to ψ_0 is

$$\mathrm{AS}_{\psi_0} := e_{\psi_0}\big(f_!\overline{\mathbf{Q}}_{\mathbf{A}^1_{\mathbf{F}_q}}\big) \in \mathcal{D}_{\mathrm{mot}}(\mathbf{A}^1_{\mathbf{F}_q};\overline{\mathbf{Q}}). \tag{6.1.3}$$

Lemma 6.1.1. *Let $m\colon \mathbf{A}^1_{\mathbf{F}_q}\times\mathbf{A}^1_{\mathbf{F}_q}\to\mathbf{A}^1_{\mathbf{F}_q}$ denote the addition map. Then there is a canonical isomorphism*

$$m^*\,\mathrm{AS}_{\psi_0} \xrightarrow{\sim} \mathrm{pr}_1^*\,\mathrm{AS}_{\psi_0}\otimes\mathrm{pr}_2^*\,\mathrm{AS}_{\psi_0}\,. \tag{6.1.4}$$

Together with the canonical trivialization at the zero section, this makes AS_{ψ_0} into a character sheaf on the additive group $\mathbf{G}_a$.

Proof. Let us temporarily denote the Artin–Schreier cover as $f\colon \widetilde{\mathbf{A}}^1_{\mathbf{F}_q}\to\mathbf{A}^1_{\mathbf{F}_q}$. Since f is finite étale, $f_!=f_*$, and $f_!\overline{\mathbf{Q}}_{\widetilde{\mathbf{A}}^1_{\mathbf{F}_q}}$ carries the regular action of the deck group $\mathbf{F}_p$.[2] The addition map $m\colon \mathbf{A}^1_{\mathbf{F}_q}\times\mathbf{A}^1_{\mathbf{F}_q}\to\mathbf{A}^1_{\mathbf{F}_q}$ lifts to a map $\widetilde{m}\colon \widetilde{\mathbf{A}}^1_{\mathbf{F}_q}\times\widetilde{\mathbf{A}}^1_{\mathbf{F}_q}\to\widetilde{\mathbf{A}}^1_{\mathbf{F}_q}$, because

$$(u+v)^p-(u+v)=(u^p-u)+(v^p-v).$$

This covering map is equivariant for the addition homomorphism $\mathbf{F}_p\times\mathbf{F}_p\to\mathbf{F}_p$ of deck groups. The fiber of this homomorphism is the anti-diagonal $\{(a,-a)\}$, on which the character $\psi_0\boxtimes\psi_0$ is trivial. Therefore the $(\psi_0\boxtimes\psi_0)$-isotypic summand of $\mathrm{pr}_1^*\,f_!\overline{\mathbf{Q}}_{\widetilde{\mathbf{A}}^1}\otimes\mathrm{pr}_2^*\,f_!\overline{\mathbf{Q}}_{\widetilde{\mathbf{A}}^1}$ descends canonically to the ψ_0-isotypic summand of $m^*f_!\overline{\mathbf{Q}}_{\widetilde{\mathbf{A}}^1}$.

The addition map $(a,b)\mapsto a+b$ of $\mathbf{F}_p$-groups induces a map $\overline{\mathbf{Q}}[\mathbf{F}_p]\otimes_{\overline{\mathbf{Q}}}\overline{\mathbf{Q}}[\mathbf{F}_p]\to\overline{\mathbf{Q}}[\mathbf{F}_p]$ carrying

$$e_{\psi_0}\boxtimes e_{\psi_0}\longmapsto e_{\psi_0}.$$

This gives the isomorphism (6.1.4). The associativity and unit compatibilities follow from the associativity and unit of addition on the Artin–Schreier torsor, together with the canonical trivialization of the torsor over 0. □

For an $\mathbf{F}_q$-stack S, let $\mathrm{AS}_{\psi,S}$ denote the pullback of AS_{ψ_0} to $\mathbf{A}^1_S$. The construction of (6.1.4) is compatible with arbitrary base change in S.

6.2. **Motivic derived Fourier transform.** The notion of *derived vector bundles* is recalled in [FYZ23, §6.1]. Let $E\to S$ be a derived vector bundle over a derived Artin stack S over $\mathbf{F}_q$, and let $\widehat{E}\to S$ be its dual. Write

$$\langle\,,\,\rangle_E\colon \widehat{E}\times_S E\longrightarrow\mathbf{A}^1_S \tag{6.2.1}$$

for the evaluation morphism, and let $\mathrm{pr}_{\widehat{E}}$ and pr_E denote the two projections from $\widehat{E}\times_S E$. Let $d(E/S)$ be the virtual relative dimension of $E\to S$. It is locally constant on S; whenever S is disconnected, all shifts and Tate twists involving $d(E/S)$ are understood componentwise. The motivic Fourier transform associated with ψ is the $\overline{\mathbf{Q}}$-linear functor $\mathcal{D}_{\mathrm{mot}}(E;\overline{\mathbf{Q}})\to\mathcal{D}_{\mathrm{mot}}(\widehat{E};\overline{\mathbf{Q}})$ given by

$$\mathrm{FT}^{\psi}_E(\mathcal{K}) := \mathrm{pr}_{\widehat{E},!}\big(\mathrm{pr}_E^*\,\mathcal{K}\otimes\langle\,,\,\rangle_E^*\,\mathrm{AS}_{\psi,S}\big)[d(E/S)]. \tag{6.2.2}$$

This is the motivic analogue of the Fourier–Deligne transform. The shift in (6.2.2) is the Deligne–Laumon normalization. We use FT for this derived motivic Fourier transform throughout.

6.3. **Formal properties of the motivic Fourier transform.** We now record the motivic derived analogues of the formal derived Fourier analysis used below. The proofs of all statements in this subsection are deferred to the forthcoming work of Tong Zhou [Zho26]. They are formulated for the motivic categories $\mathcal{D}_{\mathrm{mot}}((-);\overline{\mathbf{Q}})$ and for the Artin–Schreier kernel $\mathrm{AS}_{\psi,S}$ fixed above. The motivic six-operation formalism for $\mathcal{D}_{\mathrm{mot}}((-);\overline{\mathbf{Q}})$ on derived Artin stacks is the one recalled in [FK24, §3.1], ultimately based on Khan's extension of motivic homotopy theory and its six operations to derived Artin stacks [Kha19, Appendix A]. In the rest of this section we write $\mathrm{FT}_E=\mathrm{FT}^{\psi}_E$, unless the additive character must be displayed.

[2] We fix once and for all the *pullback* convention for this action: $a\in\mathbf{F}_p$ acts via σ_a^*, where $\sigma_a(u)=u+a$ is the deck transformation. With this convention the image of the e_{ψ_0}-summand under any ℓ-adic realization is the Artin–Schreier local system with trace function $x\mapsto\psi_0(x)$.

Theorem 6.3.1 (Formal properties)**.** *Let $E \to S$ be a derived vector bundle over a derived Artin stack over $\mathbf{F}_q$, let $\widehat{E} \to S$ be its dual, and put $d = d(E/S)$. The functor FT_E of* (6.2.2) *satisfies the following canonical compatibilities.*

(1) For any morphism $h\colon \widetilde{S} \to S$, write $\widetilde{E} = E \times_S \widetilde{S}$, and let $h^E\colon \widetilde{E} \to E$ and $h^{\widehat{E}}\colon \widehat{\widetilde{E}} \to \widehat{E}$ be the induced maps. There are canonical natural isomorphisms

$$\mathrm{FT}_{\widetilde{E}} \circ (h^E)^* \xrightarrow{\sim} (h^{\widehat{E}})^* \circ \mathrm{FT}_E, \qquad \mathrm{FT}_{\widetilde{E}} \circ (h^E)^! \xrightarrow{\sim} (h^{\widehat{E}})^! \circ \mathrm{FT}_E, \tag{6.3.1}$$

and

$$\mathrm{FT}_E \circ h^E_! \xrightarrow{\sim} h^{\widehat{E}}_! \circ \mathrm{FT}_{\widetilde{E}}, \qquad \mathrm{FT}_E \circ h^E_* \xrightarrow{\sim} h^{\widehat{E}}_* \circ \mathrm{FT}_{\widetilde{E}}\,. \tag{6.3.2}$$

(2) There is a canonical involutivity isomorphism

$$\mathrm{FT}^{\psi}_{\widehat{E}} \circ \mathrm{FT}^{\psi}_E \xrightarrow{\sim} [-1]^*_E(-d), \tag{6.3.3}$$

where $[-1]_E\colon E \to E$ denotes fiberwise multiplication by -1.

(3) Let $f\colon E' \to E$ be a linear map of derived vector bundles over S, and let $\widehat{f}\colon \widehat{E} \to \widehat{E}'$ be its dual. Put $d' = d(E'/S)$. There are canonical natural isomorphisms

$$\begin{aligned} &\widehat{f}^* \circ \mathrm{FT}_{E'} \xrightarrow{\sim} \mathrm{FT}_E \circ f_{![d'-d]}, \qquad &&\widehat{f}^! \circ \mathrm{FT}_{E'} \xrightarrow{\sim} \mathrm{FT}_E \circ f_{*[d-d'](d-d')}, \\ &\mathrm{FT}_{E'} \circ f^* \xrightarrow{\sim} \widehat{f}_! \circ \mathrm{FT}_{E\,[d-d'](d-d')}, \qquad &&\mathrm{FT}_{E'} \circ f^! \xrightarrow{\sim} \widehat{f}_* \circ \mathrm{FT}_{E\,[d'-d]}. \end{aligned} \tag{6.3.4}$$

These isomorphisms intertwine the adjunctions $(f_!, f^!)$ and $(\widehat{f}^*, \widehat{f}_*)$, and the adjunctions (f^*, f_*) and $(\widehat{f}_!, \widehat{f}^!)$, with the shifts and Tate twists displayed in (6.3.4)*.*

(4) If $z_E\colon S \to E$ and $z_{\widehat{E}}\colon S \to \widehat{E}$ are the zero sections, set $\delta_E := z_{E!}\overline{\mathbf{Q}}_S$ and $\delta_{\widehat{E}} := z_{\widehat{E}!}\overline{\mathbf{Q}}_S$. Then

$$\mathrm{FT}_E(\delta_E) \xrightarrow{\sim} \overline{\mathbf{Q}}_{\widehat{E}[d]}, \qquad \mathrm{FT}_E(\overline{\mathbf{Q}}_E) \xrightarrow{\sim} \delta_{\widehat{E}[-d](-d)}. \tag{6.3.5}$$

(5) For $\mathcal{K}_0, \mathcal{K}_1 \in \mathcal{D}_{\mathrm{mot}}(E; \overline{\mathbf{Q}})$, with convolution taken along the additive group structure on E/S, there are canonical natural isomorphisms

$$\mathrm{FT}_E(\mathcal{K}_0 \star \mathcal{K}_1) \xrightarrow{\sim} \mathrm{FT}_E(\mathcal{K}_0) \otimes \mathrm{FT}_E(\mathcal{K}_1)_{[-d]}, \tag{6.3.6}$$

and

$$\mathrm{FT}_E(\mathcal{K}_0 \otimes \mathcal{K}_1) \xrightarrow{\sim} \mathrm{FT}_E(\mathcal{K}_0) \star \mathrm{FT}_E(\mathcal{K}_1)_{[d](d)}, \tag{6.3.7}$$

where the convolution on the right side of (6.3.7) *is taken on $\widehat{E}/S$.*

(6) Let $\mathbf{D}_E$ and $\mathbf{D}_{\widehat{E}}$ denote Verdier duality on E and $\widehat{E}$. There is a canonical natural isomorphism

$$\mathbf{D}_{\widehat{E}}\big(\mathrm{FT}^{\psi}_E(\mathcal{K})\big) \xrightarrow{\sim} \mathrm{FT}^{\psi^{-1}}_E\big(\mathbf{D}_E(\mathcal{K})\big)_{(d)}. \tag{6.3.8}$$

(7) For $\mathcal{K}_1, \mathcal{K}_2 \in \mathcal{D}_{\mathrm{mot}}(E; \overline{\mathbf{Q}})$, there is a canonical Plancherel isomorphism

$$\pi_{\widehat{E}!}\big(\mathrm{FT}_E(\mathcal{K}_1) \otimes \mathrm{FT}_E(\mathcal{K}_2)\big) \xrightarrow{\sim} \pi_{E!}\big(\mathcal{K}_1 \otimes [-1]^*_E\mathcal{K}_2\big)_{(-d)}. \tag{6.3.9}$$

Proposition 6.3.2 (Compatibility with proper base change)**.** *Consider a Cartesian square of globally presented derived vector bundles over a common base, together with its Fourier-dual Cartesian square,*

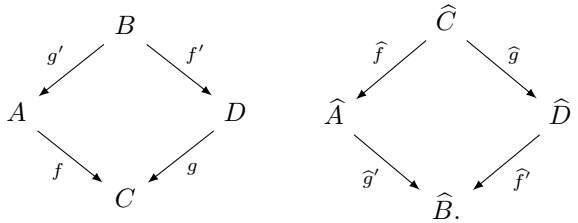

(6.3.10)

Put $d_f = d(f)$ and $d_g = d(g)$. Then the diagram of functors $\mathcal{D}_{\mathrm{mot}}(A;\overline{\mathbf{Q}}) \to \mathcal{D}_{\mathrm{mot}}(\widehat{D};\overline{\mathbf{Q}})$

$$\begin{array}{ccc} \widehat{g}_!\widehat{f}^*\,\mathrm{FT}_A & \xrightarrow{\ \sim\ } & \mathrm{FT}_D\; g^* f_{![d_f+d_g](d_g)} \\ \sim\downarrow & & \downarrow\sim \\ (\widehat{f}')^*\widehat{g}'_!\,\mathrm{FT}_A & \xrightarrow[\ \sim\ ]{} & \mathrm{FT}_D\; f'_!(g')^*{}_{[d_f+d_g](d_g)} \end{array} \tag{6.3.11}$$

commutes. Here the right vertical arrow is the Beck–Chevalley equivalence $g^* f_! \simeq f'_!(g')^*$ for the first square in (6.3.10)*, after applying FT_D and the displayed shifts and twists. The left vertical arrow is the Beck–Chevalley equivalence $\widehat{g}_!\widehat{f}^* \simeq (\widehat{f}')^*\widehat{g}'_!$ for the Fourier-dual square. The horizontal arrows are the Fourier functoriality isomorphisms of* (6.3.4)*, and the vertical arrows are the proper base-change isomorphisms for* (6.3.10) *and its Fourier-dual square.*

Proposition 6.3.3 (Fourier transform of the Gysin map)**.** *Let $f\colon E' \to E$ be a globally presented quasi-smooth linear map of derived vector bundles over S, and let $\widehat{f}\colon \widehat{E} \to \widehat{E}'$ be its dual. The quasi-smoothness of f implies that $\widehat{f}$ is separated and representable in derived schemes, by* [FYZ23, Lemma 6.1.5]*, so that the forget-supports transformation $\mathrm{can}(\widehat{f})\colon \widehat{f}_! \to \widehat{f}_*$ is defined. Then the diagram of functors $\mathcal{D}_{\mathrm{mot}}(E;\overline{\mathbf{Q}}) \to \mathcal{D}_{\mathrm{mot}}(\widehat{E}';\overline{\mathbf{Q}})$*

$$\begin{array}{ccc} \widehat{f}_!\,\mathrm{FT}_E & \xrightarrow{\ \mathrm{can}(\widehat{f})\ } & \widehat{f}_*\,\mathrm{FT}_E \\ \sim\downarrow & & \downarrow\sim \\ \mathrm{FT}_{E'}\, f^*{}_{[d(f)](d(f))} & \xrightarrow[\ [f]\ ]{} & \mathrm{FT}_{E'}\, f^!{}_{[-d(f)]} \end{array} \tag{6.3.12}$$

commutes, where $[f]$ denotes the natural transformation induced by the relative fundamental class of f, and $\mathrm{can}(\widehat{f})\colon \widehat{f}_! \to \widehat{f}_*$ is the forget-supports map.

6.4. **Self-duality of Gaussians.** Recall our standing assumption that $p \neq 2$. Throughout this subsection we write FT_E for FT_E^{ψ}.

Construction 6.4.1 (Induced Gaussian kernels)**.** Let $E \to S$ be a vector bundle of virtual relative dimension $d(E/S)$, and let

$$h_E\colon E \xrightarrow{\sim} \widehat{E}$$

be a symmetric self-duality; that is, under the canonical biduality identification $E \cong \widehat{\widehat{E}}$, one has $h_E^{\vee} = h_E$. Define

$$\beta_{h_E}\colon E \xrightarrow{(h_E,\mathrm{Id})} \widehat{E} \times_S E \xrightarrow{\langle\,,\,\rangle_E} \mathbf{A}^1_S. \tag{6.4.1}$$

The symmetry of h_E implies the following identity of morphisms $E \times_S E \to \mathbf{A}^1_S$:

$$\beta_{h_E}(e_1+e_2) - \beta_{h_E}(e_1) - \beta_{h_E}(e_2) = 2\langle h_E(e_1), e_2\rangle_E. \tag{6.4.2}$$

Thus β_{h_E} is the quadratic function associated with the symmetric bilinear form defined by h_E, with the normalization responsible for the multiplication-by-2 map $[2]_{\widehat{E}}$ in (6.4.5). The self-duality also determines the dual quadratic function

$$\beta^{\vee}_{h_E}\colon \widehat{E} \xrightarrow{(\mathrm{Id},h_E^{-1})} \widehat{E} \times_S E \xrightarrow{\langle\,,\,\rangle_E} \mathbf{A}^1_S. \tag{6.4.3}$$

One has $h_E^*\beta^{\vee}_{h_E} = \beta_{h_E}$. Define the associated Gaussian object by

$$\mathscr{G}_{h_E} := \beta^*_{h_E}\,\mathrm{AS}_{\psi,S} \in \mathcal{D}_{\mathrm{mot}}(E;\overline{\mathbf{Q}}). \tag{6.4.4}$$

Via h_E, the positive-dual Gaussian $(\beta^{\vee}_{h_E})^*\,\mathrm{AS}_{\psi,S}$ on $\widehat{E}$ identifies canonically with $\mathscr{G}_{h_E}$ on E. On the other hand, $h_E^*(-\beta^{\vee}_{h_E}) = -\beta_{h_E}$, so upon using h_E to identify $\widehat{E}$ with E, the Fourier-dual Gaussian is $(-\beta_{h_E})^*\,\mathrm{AS}_{\psi,S}$, which need not be isomorphic to $\mathscr{G}_{h_E}$.

Lemma 6.4.2 (Fourier transform of Gaussian sheaves)**.** *With the notation of Construction 6.4.1, there is a canonical isomorphism*

$$[2]_{\widehat{E}}^{*}\,\mathrm{FT}_E(\mathscr{G}_{h_E}) \xrightarrow{\sim} (-\beta_{h_E}^{\vee})^{*}\,\mathrm{AS}_{\psi,S} \otimes \pi_{\widehat{E}}^{*}\big(\pi_{E!}\mathscr{G}_{h_E}[d(E/S)]\big), \tag{6.4.5}$$

where $[2]_{\widehat{E}}\colon \widehat{E} \to \widehat{E}$ *is fiberwise multiplication by* 2*, and* $\pi_E\colon E \to S$ *and* $\pi_{\widehat{E}}\colon \widehat{E} \to S$ *are the structure morphisms.*

Proof. Consider the Cartesian square

$$\begin{array}{ccc} \widehat{E} \times_S E & \xrightarrow{[2]_{\widehat{E}} \times \mathrm{Id}_E} & \widehat{E} \times_S E \\ {\scriptstyle \mathrm{pr}_{\widehat{E}}}\downarrow & & \downarrow{\scriptstyle \mathrm{pr}_{\widehat{E}}} \\ \widehat{E} & \xrightarrow[[2]_{\widehat{E}}]{} & \widehat{E} \end{array} \tag{6.4.6}$$

Base change for the square (6.4.6), together with the definition (6.2.2) and the equality $\mathrm{pr}_E \circ ([2]_{\widehat{E}} \times \mathrm{Id}_E) = \mathrm{pr}_E$, gives a canonical identification

$$[2]_{\widehat{E}}^{*}\,\mathrm{FT}_E(\mathscr{G}_{h_E}) \cong \mathrm{pr}_{\widehat{E},!}\Big(\mathrm{pr}_E^{*}\mathscr{G}_{h_E} \otimes \big(\big\langle [2]_{\widehat{E}} \circ \mathrm{pr}_{\widehat{E}}, \mathrm{pr}_E \big\rangle_E\big)^{*}\,\mathrm{AS}_{\psi,S}\Big)[d(E/S)]. \tag{6.4.7}$$

Let

$$\tau_{h_E}\colon \widehat{E} \times_S E \longrightarrow \widehat{E} \times_S E, \qquad (\eta, e) \longmapsto \big(\eta, e + h_E^{-1}(\eta)\big) \tag{6.4.8}$$

be translation in the E-coordinate by $h_E^{-1}(\eta)$. It is an automorphism over $\widehat{E}$. The polarization identity (6.4.2) implies the following identity of morphisms $\widehat{E} \times_S E \to \mathbf{A}^1_S$:

$$\beta_{h_E} \circ \mathrm{pr}_E + \big\langle [2]_{\widehat{E}} \circ \mathrm{pr}_{\widehat{E}}, \mathrm{pr}_E \big\rangle_E = \beta_{h_E} \circ \mathrm{pr}_E \circ \tau_{h_E} - \beta_{h_E}^{\vee} \circ \mathrm{pr}_{\widehat{E}}\,. \tag{6.4.9}$$

Indeed, for $(\eta, e) \in \widehat{E} \times_S E$ and $v = h_E^{-1}(\eta)$, the right hand side of (6.4.9) is $\beta_{h_E}(e+v) - \beta_{h_E}(v) = \beta_{h_E}(e) + 2\langle h_E(e), v\rangle_E$, which equals $\beta_{h_E}(e) + \langle 2\eta, e\rangle_E$ by symmetry of h_E. Applying the Artin–Schreier additivity isomorphism (6.1.4) to (6.4.9) yields a canonical isomorphism

$$\mathrm{pr}_E^{*}\mathscr{G}_{h_E} \otimes \big(\big\langle [2]_{\widehat{E}} \circ \mathrm{pr}_{\widehat{E}}, \mathrm{pr}_E \big\rangle_E\big)^{*}\,\mathrm{AS}_{\psi,S} \cong \tau_{h_E}^{*}\,\mathrm{pr}_E^{*}\mathscr{G}_{h_E} \otimes \mathrm{pr}_{\widehat{E}}^{*}\big((-\beta_{h_E}^{\vee})^{*}\,\mathrm{AS}_{\psi,S}\big). \tag{6.4.10}$$

Substituting (6.4.10) into (6.4.7) and applying the projection formula for $\mathrm{pr}_{\widehat{E}}$ gives

$$[2]_{\widehat{E}}^{*}\,\mathrm{FT}_E(\mathscr{G}_{h_E}) \cong (-\beta_{h_E}^{\vee})^{*}\,\mathrm{AS}_{\psi,S} \otimes \mathrm{pr}_{\widehat{E},!}\big(\tau_{h_E}^{*}\,\mathrm{pr}_E^{*}\mathscr{G}_{h_E}\big)[d(E/S)]. \tag{6.4.11}$$

Because τ_{h_E} is an automorphism over $\widehat{E}$, the canonical isomorphism $\tau_{h_E,!}\tau_{h_E}^{*} \cong \mathrm{Id}$ gives

$$\mathrm{pr}_{\widehat{E},!}\big(\tau_{h_E}^{*}\,\mathrm{pr}_E^{*}\mathscr{G}_{h_E}\big) \cong \mathrm{pr}_{\widehat{E},!}\,\mathrm{pr}_E^{*}\mathscr{G}_{h_E}. \tag{6.4.12}$$

Finally, base change for the Cartesian square

$$\begin{array}{ccc} \widehat{E} \times_S E & \xrightarrow{\mathrm{pr}_E} & E \\ {\scriptstyle \mathrm{pr}_{\widehat{E}}}\downarrow & & \downarrow{\scriptstyle \pi_E} \\ \widehat{E} & \xrightarrow[\pi_{\widehat{E}}]{} & S \end{array}$$

identifies

$$\mathrm{pr}_{\widehat{E},!}\,\mathrm{pr}_E^{*}\mathscr{G}_{h_E} \cong \pi_{\widehat{E}}^{*}\pi_{E!}\mathscr{G}_{h_E}. \tag{6.4.13}$$

Combining (6.4.11), (6.4.12), and (6.4.13) gives

$$[2]_{\widehat{E}}^{*}\,\mathrm{FT}_E(\mathscr{G}_{h_E}) \cong (-\beta_{h_E}^{\vee})^{*}\,\mathrm{AS}_{\psi,S} \otimes \pi_{\widehat{E}}^{*}\big(\pi_{E!}\mathscr{G}_{h_E}[d(E/S)]\big),$$

which is the desired isomorphism (6.4.5). □

Lemma 6.4.3 (Invertibility of the relative Gauss cohomology)**.** *Maintain the notation of Construction 6.4.1, and assume in addition that $E \to S$ is a vector bundle in the classical sense (i.e., the total space of a locally free sheaf). Set*

$$\mathfrak{G}(E, h_E) := \pi_{E!}\mathscr{G}_{h_E}[d(E/S)] \in \mathcal{D}_{\mathrm{mot}}(S; \overline{\mathbf{Q}}). \tag{6.4.14}$$

Then $\mathfrak{G}(E, h_E)$ is tensor-invertible. More precisely, there is a canonical isomorphism

$$\mathfrak{G}(E, h_E) \otimes \mathfrak{G}(E, -h_E)(d(E/S)) \xrightarrow{\sim} \overline{\mathbf{Q}}_S. \tag{6.4.15}$$

Proof. Apply Lemma 6.4.2 to h_E and to $-h_E$. Since $p \neq 2$, the map $[2]_{\widehat{E}}$ is an automorphism over S, and Lemma 6.4.2 gives

$$[2]_{\widehat{E}}^* \,\mathrm{FT}_E(\mathscr{G}_{h_E}) \cong (-\beta_{h_E}^\vee)^* \,\mathrm{AS}_{\psi,S} \otimes \pi_{\widehat{E}}^* \mathfrak{G}(E, h_E),$$
$$[2]_{\widehat{E}}^* \,\mathrm{FT}_E(\mathscr{G}_{-h_E}) \cong (-\beta_{-h_E}^\vee)^* \,\mathrm{AS}_{\psi,S} \otimes \pi_{\widehat{E}}^* \mathfrak{G}(E, -h_E).$$

Here $\beta_{-h_E}^\vee = -\beta_{h_E}^\vee$, so

$$(-\beta_{h_E}^\vee)^* \,\mathrm{AS}_{\psi,S} \otimes (-\beta_{-h_E}^\vee)^* \,\mathrm{AS}_{\psi,S} \cong \overline{\mathbf{Q}}_{\widehat{E}}$$

canonically. (Under the identification $h_E \colon E \cong \widehat{E}$ these Fourier-dual Gaussians are $\mathscr{G}_{-h_E}$ and $\mathscr{G}_{h_E}$, respectively, though we will not need this.) Use the Plancherel isomorphism in (6.3.9) with $\mathcal{K}_1 = \mathscr{G}_{h_E}$ and $\mathcal{K}_2 = \mathscr{G}_{-h_E}$. Because $\beta_{-h_E} \circ [-1]_E = \beta_{-h_E}$, one has $[-1]_E^* \mathscr{G}_{-h_E} = \mathscr{G}_{-h_E}$, and the tensor product $\mathscr{G}_{h_E} \otimes \mathscr{G}_{-h_E}$ is canonically $\overline{\mathbf{Q}}_E$. Thus the right side of Plancherel is

$$\pi_{E!}\overline{\mathbf{Q}}_E(-d(E/S)).$$

The left side, using the two displayed Fourier-transform identifications (after inserting $[2]_{\widehat{E}}^*$, which does not change $\pi_{\widehat{E}!}$ since $[2]_{\widehat{E}}$ is an automorphism over S), is

$$\mathfrak{G}(E, h_E) \otimes \mathfrak{G}(E, -h_E) \otimes \pi_{\widehat{E}!}\overline{\mathbf{Q}}_{\widehat{E}}.$$

Since E is a vector bundle in the classical sense, of rank $d(E/S)$, we have

$$\pi_{E!}\overline{\mathbf{Q}}_E \cong \pi_{\widehat{E}!}\overline{\mathbf{Q}}_{\widehat{E}} \cong \overline{\mathbf{Q}}_S[-2d(E/S)](-d(E/S)).$$

Cancelling this common factor leaves the canonical isomorphism

$$\mathfrak{G}(E, h_E) \otimes \mathfrak{G}(E, -h_E)(d(E/S)) \cong \overline{\mathbf{Q}}_S,$$

which is (6.4.15). □

Lemma 6.4.4. *Assume the hypotheses of Lemma 6.4.3 and that $d = d(E/S)$ is constant. Then there are tensor-invertible objects $\mathfrak{g}_d \in \mathcal{D}_{\mathrm{mot}}(\mathrm{Spec}\ \mathbf{F}_q; \overline{\mathbf{Q}})$ and $\varepsilon(E, h_E) \in \mathcal{D}_{\mathrm{mot}}(S; \overline{\mathbf{Q}})$ such that*

- *There is a finite étale surjection $u \colon S' \to S$ such that $u^*\varepsilon(E, h_E) \cong \overline{\mathbf{Q}}_{S'}$;*
- *There is an isomorphism*

$$\mathfrak{G}(E, h_E) \cong \mathfrak{g}_d|_S \otimes \varepsilon(E, h_E).$$

Proof. The case $d = 0$ is immediate. Suppose $d > 0$, and set

$$q_d(x_1, \ldots, x_d) := x_1^2 + \cdots + x_d^2, \qquad \mathfrak{g}_d := \pi_{d!} q_d^* \,\mathrm{AS}_{\psi_0}[d], \qquad \pi_d \colon \mathbf{A}^d_{\mathbf{F}_q} \to \mathrm{Spec}\ \mathbf{F}_q.$$

Lemma 6.4.3 shows that $\mathfrak{g}_d$ is tensor-invertible. Define

$$\varepsilon(E, h_E) := \mathfrak{G}(E, h_E) \otimes (\mathfrak{g}_d|_S)^{-1}.$$

Let

$$\pi_P \colon P := \mathrm{Isom}_S\big((\mathcal{O}_S^d, q_d), (E, \beta_{h_E})\big) \longrightarrow S.$$

Because $p \neq 2$, the quadratic form β_{h_E} is étale-locally isometric to q_d; hence P is an $\mathrm{O}_d = \mathrm{Aut}(q_d)$-torsor. The universal isometry over P induces

$$\pi_P^* \mathfrak{G}(E, h_E) \xrightarrow{\sim} \mathfrak{g}_d|_P, \tag{6.4.16}$$

and therefore $\pi_P^* \varepsilon(E, h_E) \cong \overline{\mathbf{Q}}_P$.

We have two pullbacks of (6.4.16) to $P \times_S P$, and under the isomorphism $P \times_S P \cong P \times \mathrm{O}_d$, the discrepancy between them is the action of O_d on $\mathfrak{g}_d$. Since $\mathfrak{g}_d$ is tensor-invertible, this action gives a class

in $\mathrm{Aut}_{\mathcal{D}_{\mathrm{mot}}(\mathrm{O}_d;\overline{\mathbf{Q}})}(\overline{\mathbf{Q}}_{\mathrm{O}_d}) \cong \mathrm{H}^0(\mathrm{O}_d, \overline{\mathbf{Q}}_{\mathrm{O}_d})^\times$. Its restriction to SO_d is constant because SO_d is connected, and its value at the identity is 1. Hence $\varepsilon(E, h_E)$ is trivialized by the finite étale cover $P/\mathrm{SO}_d \to S$. □

7. Cohomological correspondences

The notion of *cohomological correspondence* is reviewed in [FYZ23, §4.1.1]. Here we develop some supplementary technical material about cohomological correspondences that will be needed later.

7.1. **Tensoring cohomological correspondences.** There is an operation of tensoring a cohomological correspondence with a map, defined formally as follows. Here $A_0 \xleftarrow{a_0} C^\flat \xrightarrow{a_1} A_1$ is a correspondence of derived Artin stacks, and $\mathcal{K}_i, \mathcal{L}_i \in \mathcal{D}_{\mathrm{mot}}(A_i; \overline{\mathbf{Q}})$. Consider a cohomological correspondence

$$a_0^*\mathcal{K}_0 \xrightarrow{\mathfrak{c}} a_1^!\mathcal{K}_1 \in \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1).$$

Suppose we also have a map

$$u : a_0^*\mathcal{L}_0 \xrightarrow{\sim} a_1^*\mathcal{L}_1.$$

Definition 7.1.1. The tensor product $\mathfrak{c} \otimes u \in \mathrm{Corr}_{C^\flat}(\mathcal{K}_0 \otimes \mathcal{L}_0, \mathcal{K}_1 \otimes \mathcal{L}_1)$ is the composite

$$a_0^*(\mathcal{K}_0 \otimes \mathcal{L}_0) = a_0^*\mathcal{K}_0 \otimes a_0^*\mathcal{L}_0 \xrightarrow{\mathfrak{c}\otimes u} a_1^!\mathcal{K}_1 \otimes a_1^*\mathcal{L}_1 \xrightarrow{\diamond} a_1^!(\mathcal{K}_1 \otimes \mathcal{L}_1).$$

Here the arrow $\diamond$ is the base change natural transformation $f^!(-) \otimes f^*(-) \to f^!(- \otimes -)$, applied with $f = a_1$. For a morphism $f\colon A \to B$, this is the Beck–Chevalley morphism called the "pull-pull" base change transformation $\diamond$ in [FYZ23, §3.5], obtained as the mate of the proper base change isomorphism for the Cartesian square

$$\begin{array}{ccc} A & \xrightarrow{(\mathrm{Id}_A, f)} & A \times B \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f\times \mathrm{Id}_B} \\ B & \xrightarrow[\Delta_B]{} & B \times B \end{array} \tag{7.1.1}$$

We are particularly interested in the case $\mathcal{K}_0 = \mathcal{K}_1$ and $\mathcal{L}_0 = \mathcal{L}_1$.

7.2. **Compositions.** We will use compositions of ordinary and cohomological correspondences. Some of this material appears in [FHM25, §8].

7.2.1. *Composing correspondences.* Two correspondences are composable if the codomain of one is the domain of the other:

$$\begin{array}{ccccccc} & C_{01} & & & C_{12} & & \\ {\scriptstyle a_0}\swarrow & & \searrow{\scriptstyle a_1} & {\scriptstyle c_1}\swarrow & & \searrow{\scriptstyle c_2} & \\ A_0 & & A_1 & & A_2. & & \end{array} \tag{7.2.1}$$

In this case, let C_{02} be the fibered product

$$\begin{array}{ccccccc} & & C_{02} & & & & \\ & {\scriptstyle \mathrm{pr}_{01}}\swarrow & & \searrow{\scriptstyle \mathrm{pr}_{12}} & & & \\ & C_{01} & & & C_{12} & & \\ {\scriptstyle a_0}\swarrow & & \searrow{\scriptstyle a_1} & {\scriptstyle c_1}\swarrow & & \searrow{\scriptstyle c_2} & \\ A_0 & & A_1 & & A_2. & & \end{array} \tag{7.2.2}$$

The composite correspondence is

$$A_0 \xleftarrow{p_0} C_{02} \xrightarrow{p_2} A_2, \qquad p_0 = a_0 \,\mathrm{pr}_{01}, \quad p_2 = c_2 \,\mathrm{pr}_{12}\,. \tag{7.2.3}$$

We will apply this to a 1-legged Hecke correspondence of the form

$$Y \xleftarrow{h_0} \mathrm{Hk}^1_Y \xrightarrow{h_1} Y.$$

For $r \geq 1$, the r-fold composition of such a correspondence is

$$\mathrm{Hk}_Y^r = \underbrace{\mathrm{Hk}_Y^1 \times_{h_1,Y,h_0} \mathrm{Hk}_Y^1 \times_{h_1,Y,h_0} \cdots \times_{h_1,Y,h_0} \mathrm{Hk}_Y^1}_{r \text{ factors}}, \qquad Y \xleftarrow{h_0} \mathrm{Hk}_Y^r \xrightarrow{h_r} Y. \tag{7.2.4}$$

For $r = 0$, we use the convention $\mathrm{Hk}_Y^0 = Y$.

7.2.2. *Composing cohomological correspondences.* Next we explain how to compose cohomological correspondences.

Definition 7.2.1. Let

$$\mathfrak{c}_{01} \colon a_0^* \mathcal{K}_0 \longrightarrow a_1^! \mathcal{K}_1, \qquad \mathfrak{c}_{12} \colon c_1^* \mathcal{K}_1 \longrightarrow c_2^! \mathcal{K}_2$$

be cohomological correspondences on (7.2.1). Their composite $\mathfrak{c}_{12} \circ \mathfrak{c}_{01} \in \mathrm{Corr}_{C_{02}}(\mathcal{K}_0, \mathcal{K}_2)$ is defined to be the morphism on C_{02}

$$\begin{aligned} p_0^* \mathcal{K}_0 = \mathrm{pr}_{01}^* \, a_0^* \mathcal{K}_0 &\xrightarrow{\mathrm{pr}_{01}^* \, \mathfrak{c}_{01}} \mathrm{pr}_{01}^* \, a_1^! \mathcal{K}_1 \\ &\xrightarrow{\diamond} \mathrm{pr}_{12}^! \, c_1^* \mathcal{K}_1 \xrightarrow{\mathrm{pr}_{12}^! \, \mathfrak{c}_{12}} \mathrm{pr}_{12}^! \, c_2^! \mathcal{K}_2 = p_2^! \mathcal{K}_2, \end{aligned} \tag{7.2.5}$$

where $\diamond$ is the pull-pull Beck–Chevalley transformation of [FYZ23, §3.5] for the pullback square in (7.2.2).

Composition of cohomological correspondences is compatible with tensoring, in the following sense.

Lemma 7.2.2. *With the notation of Definition 7.2.1, suppose we are given* $u_{01} \colon a_0^* \mathcal{L}_0 \xrightarrow{\sim} a_1^* \mathcal{L}_1$ *and* $u_{12} \colon c_1^* \mathcal{L}_1 \xrightarrow{\sim} c_2^* \mathcal{L}_2$. *Let* $u_{02} \colon p_0^* \mathcal{L}_0 \xrightarrow{\sim} p_2^* \mathcal{L}_2$ *be the composite*

$$p_0^* \mathcal{L}_0 = \mathrm{pr}_{01}^* \, a_0^* \mathcal{L}_0 \xrightarrow{\mathrm{pr}_{01}^* \, u_{01}} \mathrm{pr}_{01}^* \, a_1^* \mathcal{L}_1 = \mathrm{pr}_{12}^* \, c_1^* \mathcal{L}_1 \xrightarrow{\mathrm{pr}_{12}^* \, u_{12}} \mathrm{pr}_{12}^* \, c_2^* \mathcal{L}_2 = p_2^* \mathcal{L}_2. \tag{7.2.6}$$

Then we have

$$(\mathfrak{c}_{12} \otimes u_{12}) \circ (\mathfrak{c}_{01} \otimes u_{01}) = (\mathfrak{c}_{12} \circ \mathfrak{c}_{01}) \otimes u_{02} \in \mathrm{Corr}_{C_{02}}(\mathcal{K}_0 \otimes \mathcal{L}_0, \mathcal{K}_2 \otimes \mathcal{L}_2). \tag{7.2.7}$$

Proof. After pulling Definition 7.1.1 back to C_{02}, the tensoring square for $\mathfrak{c}_{01}$ is

$$\begin{array}{ccc} \mathrm{pr}_{01}^* \, a_0^* (\mathcal{K}_0 \otimes \mathcal{L}_0) & \xrightarrow{\mathrm{pr}_{01}^* \, \mathfrak{c}_{01} \otimes \mathrm{pr}_{01}^* \, u_{01}} & \mathrm{pr}_{01}^* \big(a_1^! \mathcal{K}_1 \otimes a_1^* \mathcal{L}_1 \big) \\ \| & & \downarrow \diamond \\ \mathrm{pr}_{01}^* \, a_0^* \mathcal{K}_0 \otimes \mathrm{pr}_{01}^* \, a_0^* \mathcal{L}_0 & \xrightarrow[\mathrm{pr}_{01}^* (\mathfrak{c}_{01} \otimes u_{01})]{} & \mathrm{pr}_{01}^* \, a_1^! (\mathcal{K}_1 \otimes \mathcal{L}_1) \end{array}$$

Compatibility of tensoring with the pull-pull map in (7.2.5) amounts to the commutativity of

$$\begin{array}{ccc} \mathrm{pr}_{01}^* \, a_1^! \mathcal{K}_1 \otimes \mathrm{pr}_{01}^* \, a_1^* \mathcal{L}_1 & \xrightarrow{\mathrm{pr}_{01}^* \diamond_\otimes} & \mathrm{pr}_{01}^* \, a_1^! (\mathcal{K}_1 \otimes \mathcal{L}_1) \\ {\scriptstyle \diamond_{\mathrm{pp}} \otimes \mathrm{Id}} \downarrow & & \downarrow {\scriptstyle \diamond_{\mathrm{pp}}} \\ \mathrm{pr}_{12}^! \, c_1^* \mathcal{K}_1 \otimes \mathrm{pr}_{12}^* \, c_1^* \mathcal{L}_1 & \xrightarrow[\diamond_\otimes]{} & \mathrm{pr}_{12}^! \, c_1^* (\mathcal{K}_1 \otimes \mathcal{L}_1) \end{array}$$

Here $\diamond_{\mathrm{pp}}$ is the pull-pull map for (7.2.2), and $\diamond_\otimes$ is the tensoring Beck–Chevalley map $f^!(-) \otimes f^*(-) \to f^!(- \otimes -)$. The morphism $\diamond_{\mathrm{pp}} \otimes \mathrm{Id}$ uses the identity $\mathrm{pr}_{01}^* \, a_1^* \mathcal{L}_1 = \mathrm{pr}_{12}^* \, c_1^* \mathcal{L}_1$. Applying $\mathrm{pr}_{12}^!$ to the tensoring square for $\mathfrak{c}_{12}$ gives

$$\begin{array}{ccc} \mathrm{pr}_{12}^! \, c_1^* \mathcal{K}_1 \otimes \mathrm{pr}_{12}^* \, c_1^* \mathcal{L}_1^{\mathrm{pr}_{12}^! \, \mathfrak{c}_{12} \otimes \mathrm{pr}_{12}^* \, u_{12}} & \mathrm{pr}_{12}^! \, c_2^! \mathcal{K}_2 \otimes \mathrm{pr}_{12}^* \, c_2^* \mathcal{L}_2 \\ {\scriptstyle \diamond_\otimes} \downarrow & \downarrow {\scriptstyle \diamond_\otimes} \\ \mathrm{pr}_{12}^! \, c_1^* (\mathcal{K}_1 \otimes \mathcal{L}_1) & \xrightarrow[\mathrm{pr}_{12}^! (\mathfrak{c}_{12} \otimes u_{12})]{} \mathrm{pr}_{12}^! \, c_2^! (\mathcal{K}_2 \otimes \mathcal{L}_2) \end{array}$$

The morphism $\diamond_\otimes$ is the composite of the tensoring map for pr_{12} with $\mathrm{pr}_{12}^!$ applied to the tensoring map for c_2. The tensoring squares for $\mathfrak{c}_{01}$ and $\mathfrak{c}_{12}$ commute by definition and naturality. In the square involving $\diamond_{\mathrm{pp}}$, the two paths correspond to the two decompositions of the Cartesian rectangle obtained by concatenating the graph square of a_1 with (7.2.2). Compatibility of base-change transformations with concatenation shows that they agree, as in [FYZ23, §§3.4–3.5] and (for the motivic setting) [FK24, §4.2].

Consequently, the same composite is expressed either as $(\mathfrak{c}_{12}\otimes u_{12})\circ(\mathfrak{c}_{01}\otimes u_{01})$ or, after forming $\mathfrak{c}_{12}\circ\mathfrak{c}_{01}$ and u_{02} as in (7.2.6), as $(\mathfrak{c}_{12}\circ\mathfrak{c}_{01})\otimes u_{02}$. This proves (7.2.7). □

7.3. **Cohomological co-correspondences.** Let

$$\begin{array}{ccccc} & & C^\flat & & \\ & \swarrow^{a_0} & & \searrow^{a_1} & \\ A_0 & & & & A_1 \\ & \searrow_{a_1'} & & \swarrow_{a_0'} & \\ & & C^\sharp & & \end{array} \tag{7.3.1}$$

be a commutative diagram of derived Artin stacks which is Cartesian on the underlying reduced stacks. (This weaker condition suffices for the proper base-change step $a_{1!}a_0^* \cong (a_0')^* a_{1!}'$ in the construction of (7.3.2), by topological invariance of $\mathcal{D}_{\mathrm{mot}}((-);\overline{\mathbf{Q}})$, cf. [FK24, Remark 4.2.4]: the comparison map from $C^\flat$ to the derived fiber product $A_0\times_{C^\sharp}A_1$ induces an equivalence of motivic categories, under which the exchange isomorphism for the honest Cartesian square transports to this diamond.) For $\mathcal{K}_0\in\mathcal{D}_{\mathrm{mot}}(A_0;\overline{\mathbf{Q}})$ and $\mathcal{K}_1\in\mathcal{D}_{\mathrm{mot}}(A_1;\overline{\mathbf{Q}})$, this Cartesian diamond gives the standard correspondence/co-correspondence conversion (cf. [FYZ23, 4.8.1])

$$\gamma_C\colon \mathrm{Corr}_{C^\flat}(\mathcal{K}_0,\mathcal{K}_1)\xrightarrow{\sim}\mathrm{CoCorr}_{C^\sharp}(\mathcal{K}_0,\mathcal{K}_1). \tag{7.3.2}$$

It is the isomorphism obtained from adjunctions and proper base change:

$$\begin{aligned}\mathrm{Hom}_{C^\flat}(a_0^*\mathcal{K}_0,a_1^!\mathcal{K}_1)&\cong\mathrm{Hom}_{A_1}(a_{1!}a_0^*\mathcal{K}_0,\mathcal{K}_1)\\&\cong\mathrm{Hom}_{A_1}((a_0')^*a_{1!}'\mathcal{K}_0,\mathcal{K}_1)\\&\cong\mathrm{Hom}_{C^\sharp}(a_{1!}'\mathcal{K}_0,a_{0*}'\mathcal{K}_1).\end{aligned}$$

We write γ_C^{-1} for the inverse operation, converting a cohomological co-correspondence on $C^\sharp$ into the corresponding cohomological correspondence on $C^\flat$.

7.4. **Fourier transform of cohomological correspondences.** In this subsection all Fourier transforms are the motivic derived Fourier transforms of §6; we write $\mathrm{FT}_E=\mathrm{FT}_E^\psi$. The statements below are the motivic analogues of the cohomological-correspondence Fourier analysis in [FYZ23, §7]. We do not repeat the proofs: the arguments of [FYZ23, §7] carry over verbatim with motivic sheaves, using the formalism of §6.

We first treat correspondences of derived vector bundles over a fixed base. Assume that

$$\begin{array}{ccccc} & & C^\flat & & \\ & \swarrow^{a_0} & & \searrow^{a_1} & \\ A_0 & & & & A_1 \\ & \searrow_{a_1'} & & \swarrow_{a_0'} & \\ & & C^\sharp & & \end{array} \tag{7.4.1}$$

is a Cartesian diagram of derived vector bundles over a common base S. Its dual diagram is

$$\begin{array}{ccccc} & & \widehat{C^\sharp} & & \\ & \swarrow^{\widehat{a_1'}} & & \searrow^{\widehat{a_0'}} & \\ \widehat{A}_0 & & & & \widehat{A}_1 \\ & \searrow_{\widehat{a}_0} & & \swarrow_{\widehat{a}_1} & \\ & & \widehat{C^\flat}. & & \end{array} \tag{7.4.2}$$

For $\mathcal{K}_i \in \mathcal{D}_{\mathrm{mot}}(A_i; \overline{\mathbf{Q}})$, put $\widehat{\mathcal{K}}_i := \mathrm{FT}_{A_i}(\mathcal{K}_i)$. The Fourier transform of a cohomological correspondence first gives a co-correspondence on $\widehat{C^\flat}$:

$$\mathrm{FT}'_{C^\flat} \colon \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \xrightarrow{\sim} \mathrm{CoCorr}_{\widehat{C^\flat}}\big(\widehat{\mathcal{K}}_0, \widehat{\mathcal{K}}_{1[d(a_0)+d(a_1)](d(a_0))}\big). \tag{7.4.3}$$

Applying the correspondence/co-correspondence conversion $\gamma_{\widehat{C}}^{-1}$ for the dual Cartesian diamond (7.4.2) gives the same-base Fourier transform of cohomological correspondences

$$\mathrm{FT}_{C^\flat} := \gamma_{\widehat{C}}^{-1} \circ \mathrm{FT}'_{C^\flat} \colon \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \xrightarrow{\sim} \mathrm{Corr}_{\widehat{C^\sharp}}\big(\widehat{\mathcal{K}}_0, \widehat{\mathcal{K}}_{1[d(a_0)+d(a_1)](d(a_0))}\big). \tag{7.4.4}$$

We also need the varying-base form. Suppose that

$$\begin{array}{ccccc} A_0 & \xleftarrow{p_0} & C^\flat & \xrightarrow{p_1} & A_1 \\ \downarrow & & \downarrow & & \downarrow \\ S_0 & \xleftarrow{h_0} & S_C & \xrightarrow{h_1} & S_1 \end{array}$$

is a map of correspondences in which $A_i \to S_i$ and $C^\flat \to S_C$ are derived vector bundles, and p_0, p_1 are linear. Let $\widetilde{A}_i := A_i \times_{S_i, h_i} S_C$, and let $\widetilde{p}_i \colon C^\flat \to \widetilde{A}_i$ be the induced maps. Let $C^\sharp$ be the pushout of $\widetilde{A}_0 \xleftarrow{\widetilde{p}_0} C^\flat \xrightarrow{\widetilde{p}_1} \widetilde{A}_1$ in derived vector bundles over S_C; then the diamond

$$\begin{array}{ccccc} & & C^\flat & & \\ & \swarrow^{\widetilde{p}_0} & & \searrow^{\widetilde{p}_1} & \\ \widetilde{A}_0 & & & & \widetilde{A}_1 \\ & \searrow & & \swarrow & \\ & & C^\sharp & & \end{array} \tag{7.4.5}$$

is Cartesian. Let $h_i^{\widehat{A}} \colon \widehat{\widetilde{A}_i} \to \widehat{A}_i$ be the base-change maps on the dual derived vector bundles. The varying-base motivic derived Fourier transform is the canonical isomorphism

$$\mathrm{FT}_{C^\flat} \colon \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) \xrightarrow{\sim} \mathrm{Corr}_{\widehat{C^\sharp}}\big((h_0^{\widehat{A}})^* \mathrm{FT}_{A_0}(\mathcal{K}_0), (h_1^{\widehat{A}})^! \mathrm{FT}_{A_1}(\mathcal{K}_1)_{[d(\widetilde{p}_0)+d(\widetilde{p}_1)](d(\widetilde{p}_0))}\big). \tag{7.4.6}$$

Following [FYZ23, §7], this isomorphism is *defined* as the composite of the following three steps. First, one uses the identification

$$\mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) = \mathrm{Corr}_{C^\flat}\big((h_0^A)^* \mathcal{K}_0, (h_1^A)^! \mathcal{K}_1\big)$$

of correspondence groups over S_C, where $h_i^A \colon \widetilde{A}_i \to A_i$ denote the projections. Second, one applies the transform (7.4.4) for the diamond (7.4.5). Third, one uses the base-change compatibilities $\mathrm{FT}_{\widetilde{A}_i} \circ (h_i^A)^* \cong (h_i^{\widehat{A}})^* \circ \mathrm{FT}_{A_i}$ and $\mathrm{FT}_{\widetilde{A}_i} \circ (h_i^A)^! \cong (h_i^{\widehat{A}})^! \circ \mathrm{FT}_{A_i}$, up to the indicated shifts and twists, of Theorem 6.3.1(1). When the base-change maps $h_i^{\widehat{A}}$ are clear from context, we suppress them from the notation.

Proposition 7.4.1 (Push-pull compatibility)**.** *Consider a morphism of varying-base linear correspondences*

$$\begin{array}{ccccc} A_0 & \xleftarrow{p_0} & C^\flat & \xrightarrow{p_1} & A_1 \\ {\scriptstyle f_0}\downarrow & & \downarrow{\scriptstyle f^\flat} & & \downarrow{\scriptstyle f_1} \\ B_0 & \xleftarrow[q_0]{} & D^\flat & \xrightarrow[q_1]{} & B_1 \end{array}$$

over the same base correspondence $S_0 \xleftarrow{h_0} S_C \xrightarrow{h_1} S_1$. Let $\widetilde{p}_i$ and $\widetilde{q}_i$ denote the correspondence maps after base change to S_C, and let $C^\sharp, D^\sharp$ be the pushouts used in (7.4.5)*. The morphism of spans induces a map $f^\sharp \colon C^\sharp \to D^\sharp$. We denote its dual by $\widehat{f}^\sharp \colon \widehat{D^\sharp} \to \widehat{C^\sharp}$, and similarly write $\widehat{f}_i \colon \widehat{B}_i \to \widehat{A}_i$ for the dual of f_i. Assume the induced diagram and its Fourier-dual diagram are globally presented.*

If $f^\flat$ is left pushable, then for $\mathcal{K}_i \in \mathcal{D}_{\mathrm{mot}}(A_i; \overline{\mathbf{Q}})$ the diagram

$$\begin{array}{ccc} \mathrm{Corr}_{C^\flat}(\mathcal{K}_0, \mathcal{K}_1) & \xrightarrow{\quad \mathrm{FT}_{C^\flat} \quad} & \mathrm{Corr}_{\widehat{C^\sharp}}\big(\mathrm{FT}_{A_0}\,\mathcal{K}_0, \mathrm{FT}_{A_1}\,\mathcal{K}_{1[d(\widetilde{p}_0)+d(\widetilde{p}_1)](d(\widetilde{p}_0))}\big) \\ \Big\downarrow{\scriptstyle f_!^\flat} & & \Big\downarrow{\scriptstyle (\widehat{f}^\sharp)^*} \\ \mathrm{Corr}_{D^\flat}(f_{0!}\mathcal{K}_0, f_{1!}\mathcal{K}_1) & \xrightarrow{\mathbb{T}_{[d(f_0)]}\,\mathrm{FT}_{D^\flat}} & \mathrm{Corr}_{\widehat{D^\sharp}}\big((\widehat{f}_0)^*\,\mathrm{FT}_{A_0}\,\mathcal{K}_0, (\widehat{f}_1)^*\,\mathrm{FT}_{A_1}\,\mathcal{K}_{1[d(\widetilde{q}_0)+d(\widetilde{q}_1)+d(f_0)-d(f_1)](d(\widetilde{q}_0))}\big) \end{array} \tag{7.4.7}$$

commutes. If $f^\flat$ is right pullable, then, setting $\delta_{f^\flat} := d(\widetilde{p}_1) - d(\widetilde{q}_1) = d(f^\flat) - d(f_1)$ for the defect of the right square (recall that the pullback of cohomological correspondences along a right pullable map carries the twist $\langle -\delta_{f^\flat}\rangle$, cf. [FYZ23, §4.4]*), for $\mathcal{K}_i \in \mathcal{D}_{\mathrm{mot}}(B_i; \overline{\mathbf{Q}})$ the diagram*

$$\begin{array}{ccc} \mathrm{Corr}_{D^\flat}(\mathcal{K}_0, \mathcal{K}_1) & \xrightarrow{\quad \mathrm{FT}_{D^\flat} \quad} & \mathrm{Corr}_{\widehat{D^\sharp}}\big(\mathrm{FT}_{B_0}\,\mathcal{K}_0, \mathrm{FT}_{B_1}\,\mathcal{K}_{1[d(\widetilde{q}_0)+d(\widetilde{q}_1)](d(\widetilde{q}_0))}\big) \\ \Big\downarrow{\scriptstyle (f^\flat)^*} & & \Big\downarrow{\scriptstyle (\widehat{f}^\sharp)_!} \\ \mathrm{Corr}_{C^\flat}(f_0^*\mathcal{K}_0, f_1^*\mathcal{K}_1\langle -\delta_{f^\flat}\rangle) & \xrightarrow[\mathbb{T}_{[d(f_0)](d(f_0))}\,\mathrm{FT}_{C^\flat}]{} & \mathrm{Corr}_{\widehat{C^\sharp}}\big((\widehat{f}_0)_!\,\mathrm{FT}_{B_0}\,\mathcal{K}_0, (\widehat{f}_1)_!\,\mathrm{FT}_{B_1}\,\mathcal{K}_{1[d(\widetilde{q}_0)+d(\widetilde{q}_1)](d(\widetilde{q}_0))}\big) \end{array} \tag{7.4.8}$$

commutes.

Next we investigate compatibility with compositions. Assume that the two correspondences in (7.2.1) are the upper halves of globally presented Cartesian diamonds of derived vector bundles

$$\begin{array}{ccccc} & & C_{01}^\flat & & \\ & {\scriptstyle a_0}\swarrow & & \searrow{\scriptstyle a_1} & \\ A_0 & & & & A_1 \\ & {\scriptstyle a_1'}\searrow & & \swarrow{\scriptstyle a_0'} & \\ & & C_{01}^\sharp & & \end{array} \qquad \begin{array}{ccccc} & & C_{12}^\flat & & \\ & {\scriptstyle c_1}\swarrow & & \searrow{\scriptstyle c_2} & \\ A_1 & & & & A_2 \\ & {\scriptstyle c_2'}\searrow & & \swarrow{\scriptstyle c_1'} & \\ & & C_{12}^\sharp. & & \end{array} \tag{7.4.9}$$

Here $C_{01}^\flat = C_{01}$ and $C_{12}^\flat = C_{12}$. Let $C_{02}^\flat = C_{01}^\flat \times_{A_1} C_{12}^\flat$, and let $C_{02}^\sharp$ be the lower term of the Cartesian diamond for the composite correspondence; equivalently, in global perfect-complex presentations, $C_{02}^\sharp$ is the pushout of $C_{01}^\sharp \leftarrow A_1 \to C_{12}^\sharp$. Thus the Fourier-dual correspondence is

$$\widehat{A}_0 \longleftarrow \widehat{C_{02}^\sharp} \cong \widehat{C_{01}^\sharp} \times_{\widehat{A}_1} \widehat{C_{12}^\sharp} \longrightarrow \widehat{A}_2. \tag{7.4.10}$$

Write $\widehat{\mathcal{K}}_i = \mathrm{FT}_{A_i}(\mathcal{K}_i)$, and set

$$\begin{aligned} \sigma_{01} &= d(a_0) + d(a_1), \quad & \tau_{01} &= d(a_0), \\ \sigma_{12} &= d(c_1) + d(c_2), \quad & \tau_{12} &= d(c_1), \\ \sigma_{02} &= d(p_0) + d(p_2), \quad & \tau_{02} &= d(p_0). \end{aligned}$$

Base change and additivity of virtual relative dimension give $\sigma_{02} = \sigma_{01} + \sigma_{12}$ and $\tau_{02} = \tau_{01} + \tau_{12}$. For $ij =\in \{01, 12, 02\}$, let $\mathrm{FT}_{C_{ij}^\flat}$ denote the shifted-and-twisted Fourier transform of cohomological correspondences of (7.4.4), so that

$$\mathrm{FT}_{C_{ij}^\flat}(\mathfrak{c}_{ij}) \in \mathrm{Corr}_{\widehat{C_{ij}^\sharp}}\big(\widehat{\mathcal{K}}_i, \widehat{\mathcal{K}}_{j[\sigma_{ij}](\tau_{ij})}\big).$$

Lemma 7.4.2. *In the globally presented setup* (7.4.9)*, Fourier transform is compatible with composition of cohomological correspondences. More precisely,*

$$\mathrm{FT}_{C_{02}^\flat}(\mathfrak{c}_{12} \circ \mathfrak{c}_{01}) = \big(\mathrm{FT}_{C_{12}^\flat}(\mathfrak{c}_{12})_{[\sigma_{01}](\tau_{01})}\big) \circ \mathrm{FT}_{C_{01}^\flat}(\mathfrak{c}_{01}) \in \mathrm{Corr}_{\widehat{C_{02}^\sharp}}\big(\widehat{\mathcal{K}}_0, \widehat{\mathcal{K}}_{2[\sigma_{02}](\tau_{02})}\big). \tag{7.4.11}$$

Proof. By Definition 7.2.1, the composition $\mathfrak{c}_{12}\circ\mathfrak{c}_{01}$ is obtained by pulling $\mathfrak{c}_{01}$ and $\mathfrak{c}_{12}$ to $C^\flat_{02}=C^\flat_{01}\times_{A_1}C^\flat_{12}$, inserting the pull-pull Beck–Chevalley transformation for the Cartesian square (7.2.2), and then composing. The two diamonds in (7.4.9) identify the Fourier-dual lower term of the composite with the fiber product $\widehat{C^\sharp_{01}}\times_{\widehat{A}_1}\widehat{C^\sharp_{12}}$, as displayed in (7.4.10).

The only non-formal input is the compatibility of this Beck–Chevalley transformation with Fourier transform. Applying Proposition 6.3.2 to the Cartesian square (7.2.2), using the two Cartesian diamonds in (7.4.9), shows that the Fourier transform of the pull-pull map used in the definition of $\mathfrak{c}_{12}\circ\mathfrak{c}_{01}$ is the pull-pull Beck–Chevalley map for the dual Cartesian square in (7.4.10). The correspondence/co-correspondence conversions $\gamma_{\widehat{C}}^{-1}$ are defined by adjunction and proper base change, so they are compatible with this identified Beck–Chevalley map.

After this identification, expanding the definition of $\mathrm{FT}_{C^\flat_{ij}}$ in (7.4.4) gives (7.4.11). The source and target twists agree because base change in (7.4.9) gives $\sigma_{02}=\sigma_{01}+\sigma_{12}$ and $\tau_{02}=\tau_{01}+\tau_{12}$; this is exactly the additional $[\sigma_{01}](\tau_{01})$ applied to $\mathrm{FT}_{C^\flat_{12}}(\mathfrak{c}_{12})$ in (7.4.11). □

7.5. **Tensor-convolution duality.** Given a commutative group stack A/S with addition operation $+\colon A\times_S A\to A$, we can define the convolution of $\mathcal{K},\mathcal{L}\in\mathcal{D}_{\mathrm{mot}}(A;\overline{\mathbf{Q}})$ as

$$\mathcal{K}\star\mathcal{L}:=+_!(\mathcal{K}\boxtimes_S\mathcal{L})\in\mathcal{D}_{\mathrm{mot}}(A;\overline{\mathbf{Q}}).$$

We next define the Fourier-dual operation to tensoring. Assume now that the diagram (7.3.1) is of commutative group stacks and homomorphisms over a common base S. Consider a cohomological co-correspondence

$$a'_{1!}\mathcal{K}_0\xrightarrow{\mathfrak{c}}a'_{0*}\mathcal{K}_1\in\mathrm{CoCorr}_{C^\sharp}(\mathcal{K}_0,\mathcal{K}_1).\tag{7.5.1}$$

Suppose we are given an isomorphism

$$v:a'_{1!}\mathcal{L}_0\xrightarrow{\sim}a'_{0!}\mathcal{L}_1.$$

Definition 7.5.1. We define the convolution $\mathfrak{c}\star v\in\mathrm{CoCorr}_{C^\sharp}(\mathcal{K}_0\star\mathcal{L}_0,\mathcal{K}_1\star\mathcal{L}_1)$ to be the composite

$$a'_{1!}(\mathcal{K}_0\star\mathcal{L}_0)\cong(a'_{1!}\mathcal{K}_0)\star(a'_{1!}\mathcal{L}_0)\xrightarrow{\mathfrak{c}\star v}(a'_{0*}\mathcal{K}_1)\star(a'_{0!}\mathcal{L}_1)\to a'_{0*}(\mathcal{K}_1\star\mathcal{L}_1).\tag{7.5.2}$$

Here the first isomorphism is the canonical convolution compatibility for the homomorphism a'_1. The last map is the Beck–Chevalley morphism, called the "push-push" base change natural transformation $\diamond$ in [FYZ23, §3.3], obtained as the mate of the proper base change isomorphism for the Cartesian square

$$\begin{array}{ccc} A\times_S A & \xrightarrow{+} & A\\ {\scriptstyle f\times\mathrm{Id}}\downarrow & & \downarrow{\scriptstyle f}\\ B\times_S A & \xrightarrow{+} & B\end{array}$$

with $f=a'_0$, where the bottom $+$ is the action map $(b,a)\mapsto b+f(a)$. This form of the square places the $*$-pushforward on the first factor, matching the middle term of (7.5.2); it is the Fourier-dual of the tensoring square (7.1.1).

Lemma 7.5.2. *Below, for $\mathcal{K}\in\mathcal{D}_{\mathrm{mot}}(A;\overline{\mathbf{Q}})$ we abbreviate $\widehat{\mathcal{K}}:=\mathrm{FT}_A(\mathcal{K})$.*

Assume that the diagrams (7.4.1) *and* (7.1.1) *with $f=a_1$ are globally presented in the sense of* [FYZ23, §6.3.1] *(so that motivic Fourier transform is compatible with proper base change, by Proposition 6.3.2). Put $r_i:=d(A_i/S)$ and $\delta_i:=d(a_i)$ for $i=0,1$.*

Let

$$a_0^*\mathcal{K}_0\xrightarrow{\mathfrak{c}}a_1^!\mathcal{K}_1\in\mathrm{Corr}_{C^\flat}(\mathcal{K}_0,\mathcal{K}_1),\qquad u:a_0^*\mathcal{L}_0\xrightarrow{\sim}a_1^*\mathcal{L}_1.$$

Denote their Fourier transforms[3] *by*

$$\widehat{\mathfrak{c}}:=\mathrm{FT}'_{C^\flat}(\mathfrak{c})\in\mathrm{CoCorr}_{\widehat{C^\flat}}\big(\widehat{\mathcal{K}}_0,\widehat{\mathcal{K}}_1[\delta_0+\delta_1](\delta_0)\big),$$
$$\widehat{u}:\widehat{a}_{0!}\widehat{\mathcal{L}}_0\xrightarrow{\sim}\widehat{a}_{1!}\widehat{\mathcal{L}}_1[\delta_0-\delta_1](\delta_0-\delta_1),$$

[3]See (7.4.3) and (7.4.4) for the motivic derived Fourier transform of cohomological correspondences.

where $\widehat{a}_i : \widehat{A}_i \to \widehat{C}^\flat$ *denotes the map dual to* a_i*. Under the Fourier-dual tensor/convolution identification*

$$\mathrm{FT}_{A_i}(\mathcal{K}_i \otimes \mathcal{L}_i) \cong \widehat{\mathcal{K}}_i \star \widehat{\mathcal{L}}_{i[r_i](r_i)}$$

from (6.3.7)*, we have*

$$\mathbb{T}_{[-r_0](-r_0)}\,\mathrm{FT}'_{C^\flat}(\mathfrak{c} \otimes u) = \widehat{\mathfrak{c}} \star \widehat{u} \in \mathrm{CoCorr}_{\widehat{C^\flat}}\big(\widehat{\mathcal{K}}_0 \star \widehat{\mathcal{L}}_0, \widehat{\mathcal{K}}_1 \star \widehat{\mathcal{L}}_{1[2\delta_0](2\delta_0-\delta_1)}\big).$$

Proof. By definition, $\mathfrak{c} \otimes u$ is given by the composition

$$\mathfrak{c} \otimes u \colon a_0^*(\mathcal{K}_0 \otimes \mathcal{L}_0) \to a_1^!\mathcal{K}_1 \otimes a_1^*\mathcal{L}_1 \xrightarrow{\diamond} a_1^!(\mathcal{K}_1 \otimes \mathcal{L}_1), \tag{7.5.3}$$

where $\diamond$ is the Beck–Chevalley transformation mated to proper base change for the Cartesian diagram (7.1.1). After applying $\mathrm{FT}'_{C^\flat}$, using the preceding tensor/convolution identification, and then applying the shift-twist operator $\mathbb{T}_{[-r_0](-r_0)}$, the first arrow in (7.5.3) becomes

$$(\widehat{a}_{0!}\widehat{\mathcal{K}}_0) \star (\widehat{a}_{0!}\widehat{\mathcal{L}}_0) \to (\widehat{a}_{1*}\widehat{\mathcal{K}}_{1[\delta_0+\delta_1](\delta_0)}) \star (\widehat{a}_{1!}\widehat{\mathcal{L}}_{1[\delta_0-\delta_1](\delta_0-\delta_1)}).$$

This is the first arrow in the definition (7.5.2) of $\widehat{\mathfrak{c}} \star \widehat{u}$, applied to the dual diamond (7.4.2) under the dictionary $a'_1 = \widehat{a}_0$, $a'_0 = \widehat{a}_1$, $C^\sharp = \widehat{C^\flat}$.

The Fourier transform of the second arrow in (7.5.3) is the Beck–Chevalley transformation mated to proper base change for the dual Cartesian diagram. Since we have assumed that the relevant diagrams are globally presented, Proposition 6.3.2 implies that the motivic Fourier transform carries the proper base change isomorphism for a Cartesian square to the proper base change isomorphism for the Fourier-dual square. This implies that the motivic Fourier transform of the second arrow in (7.5.3) is the second arrow in the definition (7.5.2) of $\widehat{\mathfrak{c}} \star \widehat{u}$, completing the proof. □

8. A trace formula for the sheaf-cycle correspondence

In this section we prove a technical result establishing compatibility of pushforwards with formation of traces for certain *nonproper* maps, the proper case being in [FK24].

8.1. **Formulation.** Let

$$S \xleftarrow{h_0} S^\flat \xrightarrow{h_1} S$$

be a correspondence of derived Artin stacks locally of finite type over $\mathbf{F}_q$. Let $\mathcal{E}$ on S and $\mathcal{E}^\flat$ on $S^\flat$ be *coconnective* perfect complexes, that is, perfect complexes of tor-amplitude in $[0, \infty)$ in our cohomological indexing. Recall from [FYZ23, §6.1] that $\mathrm{Tot}(\mathcal{E})$ is the functor $T \mapsto \mathrm{Map}_{\mathrm{QCoh}(T)}(\mathcal{O}_T, u^*\mathcal{E})$, and set $E := \mathrm{Tot}_S(\mathcal{E})$ and $E^\flat := \mathrm{Tot}_{S^\flat}(\mathcal{E}^\flat)$. By [FYZ23, Lemma 6.1.5], these are relative affine derived schemes over their bases. They are classical affine schemes when the complexes are locally free in degree 0; in general, their classical truncations are the associated abelian cones. Their zero sections are therefore closed embeddings.

For each $i \in \{0, 1\}$, suppose we are given a map $\mathcal{E}^\flat \to h_i^*\mathcal{E}$ whose derived fiber has tor-amplitude in $[1, \infty)$. By [FYZ23, Lemma 6.1.5], the induced map of total spaces $E^\flat \to h_i^*E$ is then a closed embedding of derived vector bundles. This implies that the map of correspondences

$$\begin{array}{ccccc} E & \xleftarrow{a_0} & E^\flat & \xrightarrow{a_1} & E \\ \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f^\flat} & & \downarrow{\scriptstyle f_1} \\ S & \xleftarrow{h_0} & S^\flat & \xrightarrow{h_1} & S \end{array} \tag{8.1.1}$$

is pushable, since its left comparison map is the closed embedding $E^\flat \hookrightarrow h_0^*E$. Consequently

$$f_! \colon \mathrm{Corr}_{E^\flat}(\mathcal{K}_0, \mathcal{K}_1) \to \mathrm{Corr}_{S^\flat}(f_{0!}\mathcal{K}_0, f_{1!}\mathcal{K}_1)$$

is defined.

The derived pushforwards below are along schematic maps locally of finite type and therefore preserve geometricity of motivic sheaves (cf. [FK24, §3.3]). The required compatibility is the following.

Proposition 8.1.1. *Assume the setup above, and let* $\mathfrak{c} \in \mathrm{Corr}_{E^\flat}(\mathcal{K}, \mathcal{K}_{\langle -i\rangle})$ *with* $\mathcal{K} \in \mathcal{D}_{\mathrm{mot,gm}}(E; \overline{\mathbf{Q}})$. *Then the induced map* $\mathrm{Sht}(f)\colon \mathrm{Sht}(E^\flat) \to \mathrm{Sht}(S^\flat)$ *is proper, and we have*

$$\mathrm{Tr}^{\mathrm{Sht}}(f_!(\mathfrak{c})) = \mathrm{Sht}(f)_! \,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}) \in \mathrm{CH}_i(\mathrm{Sht}(S^\flat)).$$

Traces are formed using the natural Weil structures (cf. [FK24, §6.4.3]*).*

Remark 8.1.2. This does *not* fall under the general compatibility results of [FK24, §6.2], the point being that the vertical maps in (8.1.1) are not proper. Proposition 8.1.1 can be thought of as a (relative) version of the Grothendieck–Lefschetz trace formula for certain types of non-proper spaces.

8.2. **Proof of Proposition 8.1.1.** The strategy is to compactify the map, and show that the boundary does not contribute to the trace.

8.2.1. *Projective completions.* The relevant projective completion is the following.

Lemma 8.2.1. *Let* T *be a derived Artin stack locally of finite type over* $\mathbf{F}_q$, *let* $\mathcal{F}$ *be a coconnective perfect complex on* T, *and let* $F := \mathrm{Tot}_T(\mathcal{F})$. *Define*

$$\mathbb{P}(F \oplus \mathcal{O}) \;:=\; \big(\mathrm{Tot}_T(\mathcal{F} \oplus \mathcal{O}) \smallsetminus z_{F\oplus\mathcal{O}}\big)/\mathbf{G}_m,$$

where $z_{F\oplus\mathcal{O}}$ *denotes the zero section (a closed embedding, by coconnectivity) and* $\mathbf{G}_m$ *acts by fiberwise scaling. We likewise write* z_F *for the zero section of* F. *Then:*

1. *The classical truncation of* $\mathbb{P}(F\oplus\mathcal{O})$ *is the projectivized abelian cone* $\mathrm{Proj}_{T_{\mathrm{cl}}} \mathrm{Sym}(\mathcal{Q}\oplus\mathcal{O})$, *where* $\mathcal{Q} := H^0(\mathcal{F}^*)$ *is coherent. In particular* $\mathbb{P}(F\oplus\mathcal{O}) \to T$ *is proper.*
2. *There is a stratification*
$$\mathbb{P}(F\oplus\mathcal{O}) = F \sqcup \mathbb{P}(F),$$
where $F \hookrightarrow \mathbb{P}(F\oplus\mathcal{O})$ *is the open locus on which the* $\mathcal{O}$*-coordinate is invertible, and the boundary identifies with* $\mathbb{P}(F) := (\mathrm{Tot}_T(\mathcal{F}) \smallsetminus z_F)/\mathbf{G}_m$.

Proof. (1) The fact that the classical truncation of $\mathbb{P}(F\oplus\mathcal{O})$ is the projectivized abelian cone $\mathrm{Proj}_{T_{\mathrm{cl}}} \mathrm{Sym}(\mathcal{Q}\oplus\mathcal{O})$ is clear from the definition. Properness is detected on classical truncations, so $\mathbb{P}(F\oplus\mathcal{O}) \to T$ is proper.

(2) The projection $\mathcal{F}\oplus\mathcal{O} \to \mathcal{O}$ gives a $\mathbf{G}_m$-equivariant fiberwise-linear map $\mathrm{Tot}_T(\mathcal{F}\oplus\mathcal{O}) \to \mathbb{A}^1_T$, hence a section of $\mathcal{O}(1)$ on the quotient. Its derived fiber over $z_{\mathbb{A}^1_T}\colon T \to \mathbb{A}^1_T$ is $\mathrm{Tot}_T(\mathcal{F})$, so its derived vanishing locus is $(\mathrm{Tot}_T(\mathcal{F}) \smallsetminus z_F)/\mathbf{G}_m = \mathbb{P}(F)$. On the complement, the $\mathcal{O}$-coordinate is invertible; rescaling it to 1 identifies the complement with $\mathcal{F}$. □

Set $\overline{E} := \mathbb{P}(E\oplus\mathcal{O})$ and $\overline{E^\flat} := \mathbb{P}(E^\flat\oplus\mathcal{O})$. By Lemma 8.2.1, these are proper over their bases, with open interiors $\jmath_E\colon E \hookrightarrow \overline{E}$ and $\jmath^\flat\colon E^\flat \hookrightarrow \overline{E^\flat}$ and boundaries $\mathbb{P}(E)$ and $\mathbb{P}(E^\flat)$.

The maps $(\mathcal{E}^\flat \to h_i^*\mathcal{E}) \oplus \mathrm{Id}_{\mathcal{O}}$ induce closed embeddings $\overline{E^\flat} \to h_i^*\overline{E}$. Since each has the form (linear) $\oplus$ $\mathrm{Id}_{\mathcal{O}}$, it preserves the interior and boundary strata. Let $\overline{a}_i\colon \overline{E^\flat} \to \overline{E}$ be the composite with the projection. The pullback of the boundary $\mathbb{P}(E) \subset \overline{E}$ pulls back under $\overline{a}_i$ to $\mathbb{P}(E^\flat)$, so the boundary is stabilized scheme-theoretically, as required in [FK24, Definition 7.2.1].

On the complementary opens, we have $\overline{a}_i^{-1}(E) = E^\flat$ as derived open substacks. Hence (8.1.1) extends to the compactified map

$$\begin{array}{ccccc} \overline{E} & \xleftarrow{\overline{a}_0} & \overline{E^\flat} & \xrightarrow{\overline{a}_1} & \overline{E} \\ \downarrow{\scriptstyle \overline{f}_0} & & \downarrow{\scriptstyle \overline{f}^\flat} & & \downarrow{\scriptstyle \overline{f}_1} \\ S & \xleftarrow{h_0} & S^\flat & \xrightarrow{h_1} & S \end{array} \tag{8.2.1}$$

All vertical maps in (8.2.1) are proper, and its comparison maps $\overline{E^\flat} \to h_i^*\overline{E}$ are closed embeddings; the map is therefore proper and pushable. The boundary strata form a correspondence $\mathbb{P}(E) \xleftarrow{\overline{a}_0} \mathbb{P}(E^\flat) \xrightarrow{\overline{a}_1} \mathbb{P}(E)$, which is the restriction of (8.2.1) in the sense of [FK24, Construction 7.2.3]. The open inclusions define a map of correspondences $\jmath$ from the top row of (8.1.1) to that of (8.2.1). Its squares are derived Cartesian, being preimages of open substacks, so $\jmath$ is pushable. Since (8.1.1) is the composite of $\jmath$ with (8.2.1), compatibility with composition of pushable maps (cf. [FYZ23, §3.2, especially Proposition 3.2.3]) gives

$$f_!\ \mathfrak{c} = \overline{f}_!\,(\jmath_!\ \mathfrak{c}). \tag{8.2.2}$$

8.2.2. *Decomposition of the shtuka stacks.* Both legs of (8.2.1) preserve the stratification $\overline{E} = \mathbb{P}(E) \sqcup E$. Since the boundary is stabilized scheme-theoretically, the Frobenius twist of the compactified self-correspondence is contracting near $\mathbb{P}(E)$ by [FK24, Example 7.2.2]. It follows from [FK24, Proposition 7.5.4] that $\mathrm{Sht}(\mathbb{P}(E^\flat)) \to \mathrm{Sht}(\overline{E^\flat})$ is an open-closed embedding on reduced substacks. Its complement is $\mathrm{Sht}(E^\flat)$: the underlying point of any shtuka lies in exactly one of the two strata, each stable under both legs. Thus we have

$$\mathrm{Sht}(\overline{E^\flat}) = \mathrm{Sht}(E^\flat) \sqcup \mathrm{Sht}(\mathbb{P}(E^\flat)) \tag{8.2.3}$$

on reduced substacks. Let $\jmath$ and $\imath$ denote the maps of correspondences

$$\begin{array}{ccccc} E & \xleftarrow{a_0} & E^\flat & \xrightarrow{a_1} & E \\ {\scriptstyle \jmath_0}\downarrow & & \downarrow{\scriptstyle \jmath^\flat} & & \downarrow{\scriptstyle \jmath_1} \\ \overline{E} & \xleftarrow[\overline{a}_0]{} & \overline{E^\flat} & \xrightarrow[\overline{a}_1]{} & \overline{E} \end{array} \qquad \begin{array}{ccccc} \mathbb{P}(E) & \longleftarrow & \mathbb{P}(E^\flat) & \longrightarrow & \mathbb{P}(E) \\ {\scriptstyle \imath_0}\downarrow & & \downarrow{\scriptstyle \imath^\flat} & & \downarrow{\scriptstyle \imath_1} \\ \overline{E} & \xleftarrow[\overline{a}_0]{} & \overline{E^\flat} & \xrightarrow[\overline{a}_1]{} & \overline{E} \end{array}$$

given by the interior and boundary inclusions; in particular $\jmath_0 = \jmath_1 = \jmath_E$, and $\imath_0 = \imath_1$ is the inclusion $\mathbb{P}(E) \hookrightarrow \overline{E}$. Write $\mathrm{Sht}(\jmath)$ and $\mathrm{Sht}(\imath)$ for the two open-closed embeddings. Furthermore, the morphism

$$\mathrm{Sht}(\overline{f})\colon\ \mathrm{Sht}(\overline{E^\flat}) \to \mathrm{Sht}(S^\flat)$$

is proper by [FK24, Proposition 6.2.1] applied to the Frobenius twist of (8.2.1). Hence $\mathrm{Sht}(f) = \mathrm{Sht}(\overline{f}) \circ \mathrm{Sht}(\jmath)$ is proper, being the composition of an open-closed embedding with a proper map; in particular $\mathrm{Sht}(f)_!$ is defined on Chow groups, as asserted in Proposition 8.1.1.

8.2.3. *Conclusion.* By (8.2.2) and the compatibility of the trace with proper pushforward, [FK24, Proposition 6.2.1], applied to the proper map of correspondences (8.2.1), we have

$$\mathrm{Tr}^{\mathrm{Sht}}(f_!\ \mathfrak{c}) \ =\ \mathrm{Tr}^{\mathrm{Sht}}\big(\overline{f}_!(\jmath_!\ \mathfrak{c})\big) \ =\ \mathrm{Sht}(\overline{f})_!\ \mathrm{Tr}^{\mathrm{Sht}}(\jmath_!\ \mathfrak{c}). \tag{8.2.4}$$

It remains to show that

$$\mathrm{Tr}^{\mathrm{Sht}}(\jmath_!\ \mathfrak{c}) \ =\ \mathrm{Sht}(\jmath)_!\ \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}) \ \in\ \mathrm{CH}_i(\mathrm{Sht}(\overline{E^\flat})). \tag{8.2.5}$$

By the open-closed decomposition (8.2.3), it suffices to compare the restrictions to its two components. On $\mathrm{Sht}(E^\flat)$, the open case of [FK24, Proposition 6.2.2] identifies the restriction of the left side with $\mathrm{Tr}^{\mathrm{Sht}}(\jmath^*\jmath_!\ \mathfrak{c}) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c})$, which is also the restriction of the right side. On $\mathrm{Sht}(\mathbb{P}(E^\flat))$, the Frobenius twist is contracting near the boundary, and the restriction of $\jmath_!\ \mathfrak{c}$ to the boundary correspondence in the sense of [FK24, Construction 7.2.3] is a cohomological correspondence for $\imath_0^*\jmath_{E!}\mathcal{K} = 0$. Hence [FK24, Corollary 7.5.7] implies that the restriction of $\mathrm{Tr}^{\mathrm{Sht}}(\jmath_!\ \mathfrak{c})$ vanishes, while the right side vanishes there for reasons of support. This proves (8.2.5).

Combining (8.2.4) and (8.2.5) with the equality $\mathrm{Sht}(\overline{f})_!\ \mathrm{Sht}(\jmath)_! = \mathrm{Sht}(f)_!$ completes the proof. □

Part 3. The Modularity Conjecture

9. The genetic pattern of modularity

In this section we establish certain sheaf-theoretic results which can be viewed as the "genetics of modularity" from Figure 1. The results here are essentially agnostic about the ground field.

9.1. **The bestiary of derived vector bundles.** Let $\mathcal{G} \in \mathrm{Bun}_{GU^-(2m)}(k)$. Let $\mathcal{E}_1, \mathcal{E}_2 \subset \mathcal{G}$ be transverse Lagrangians in the sense of [FYZ23, §2]. As discussed in [FYZ23, §2, §9.1], the transversality assumption implies that we have short exact sequences

$$\begin{aligned} 0 \to \sigma^*\mathcal{E}_2 \to \mathcal{E}_1^* \to Q_1 \to 0, \\ 0 \to \sigma^*\mathcal{E}_1 \to \mathcal{E}_2^* \to Q_2 \to 0 \end{aligned} \tag{9.1.1}$$

where Q_1 and $Q_2 \cong \sigma^*Q_1^* := \sigma^*\mathcal{E}xt^1_{X'}(Q_1, \mathcal{O}_{X'})$ are torsion coherent sheaves on X'.[4]

[4]Our Q_i agrees with the $\widetilde{Q}_i$ of [FYZ23, §9.1], i.e. with σ^* applied to the Q_i of [FYZ23, §2].

Fix a Harder–Narasimhan truncation μ for $\mathrm{Bun}_{U(n)}$ and write $S = \mathrm{Bun}_{U(n)}^{\leq \mu}$. We let Hk_S^r be the corresponding Harder–Narasimhan truncation of the Hecke stack [FYZ23, §9.1.2]. This fits into a correspondence diagram

$$S \xleftarrow{h_0} \mathrm{Hk}_S^r \xrightarrow{h_r} S.$$

From this setup we constructed a constellation of moduli spaces in [FYZ23, §9.1,§9.2], and we will use the same notation here.

(1) For $i \in \{0, r\}$: we have derived vector bundles U_i, V_i, W_i over S; and their respective dual derived vector bundles $W_i^\perp = \widehat{U}_i, V_i' = \widehat{V}_i, U_i^\perp = \widehat{W}_i$.
(2) For $i \in \{0, r\}$: pulling back U_i, V_i, W_i to Hk_S^r via h_i gives derived vector bundles $\widetilde{U}_i, \widetilde{V}_i, \widetilde{W}_i$. Similarly, we have $\widetilde{W}_i^\perp, \widetilde{V}_i', \widetilde{U}_i^\perp$.
(3) The Hecke stacks
$$\mathrm{Hk}_U^\flat, \mathrm{Hk}_U^\sharp, \mathrm{Hk}_V^\flat, \mathrm{Hk}_V^\sharp, \mathrm{Hk}_W^\flat, \mathrm{Hk}_W^\sharp,$$
which are derived vector bundles over Hk_S^r; and their respective dual derived vector bundles
$$\mathrm{Hk}_{W^\perp}^\sharp, \mathrm{Hk}_{W^\perp}^\flat, \mathrm{Hk}_{\widehat{V}}^\sharp, \mathrm{Hk}_{\widehat{V}}^\flat, \mathrm{Hk}_{U^\perp}^\sharp, \mathrm{Hk}_{U^\perp}^\flat .$$

We briefly recall the definitions of these objects.

9.1.1. *U, V, W and their duals.* Let $\mathcal{F}_{\mathrm{univ}}$ denote the universal Hermitian bundle on $X' \times S$. The six base vector bundles are the total spaces of the following perfect complexes on S:

$$\mathcal{U} = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, \mathcal{E}_1^*)}, \quad \mathcal{V} = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, Q_1)}, \quad \mathcal{W} = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, \sigma^*\mathcal{E}_{2[1]})},$$
$$\mathcal{U}^\perp = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, \mathcal{E}_2^*)}, \quad \mathcal{V}' = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, Q_2)}, \quad \mathcal{W}^\perp = \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}^*, \sigma^*\mathcal{E}_{1[1]})}.$$

Then $U = \mathrm{Tot}_S(\mathcal{U})$, $V = \mathrm{Tot}_S(\mathcal{V})$, $W = \mathrm{Tot}_S(\mathcal{W})$, and similarly for the perpendicular collection. For $i = 0, r$ we write U_i, V_i, W_i for the corresponding copies over the ith copy S_i of S, and we write $\widetilde{U}_i, \widetilde{V}_i, \widetilde{W}_i$ for their pullbacks along $h_i \colon \mathrm{Hk}_S^r \to S_i$. The exact sequences (9.1.1) induce exact triangles in $\mathrm{Perf}(S)$,

$$\mathcal{U} \to \mathcal{V} \to \mathcal{W}, \qquad \mathcal{U}^\perp \to \mathcal{V}' \to \mathcal{W}^\perp. \tag{9.1.2}$$

Passing to total spaces, this induces the following pair of derived Cartesian squares over S:

$$\begin{array}{ccc} U & \xrightarrow{f} & V \\ {\scriptstyle \pi_U}\downarrow & & \downarrow{\scriptstyle g} \\ S & \xrightarrow[z_W]{} & W \end{array} \qquad\qquad \begin{array}{ccc} U^\perp & \xrightarrow{\widehat{g}} & \widehat{V} \\ {\scriptstyle \widehat{z}_W}\downarrow & & \downarrow{\scriptstyle \widehat{f}} \\ S & \xrightarrow[\widehat{\pi}_U]{} & W^\perp \end{array} \tag{9.1.3}$$

Here z_W and $\widehat{\pi}_U = z_{W^\perp}$ denote the zero sections of their respective derived vector bundles, while π_U and $\widehat{z}_W = \pi_{U^\perp}$ is the structure morphism $U^\perp \to S$. From (9.1.3) we have corresponding derived Cartesian squares (U_i, V_i, S_i, W_i) and $(U_i^\perp, \widehat{V}_i, S_i, W_i^\perp)$ for $i = 0, r$. Here $\widehat{V}_i$ denotes the Serre-dual vector bundle to V_i; relative Serre duality identifies $\mathrm{Tot}_S(\mathcal{V}')$ with $\widehat{V}$, identifies $U^\perp$ with $\widehat{W}$, and identifies $W^\perp$ with $\widehat{U}$.

9.1.2. *Hecke correspondences.* We next recall the Hecke versions. Over Hk_S^r there are two universal complexes on $X' \times \mathrm{Hk}_S^r$. For an R-point $(\mathcal{F}_\star) \in \mathrm{Hk}_S^r(R)$, the complex $\mathcal{F}_\bullet^\flat$ on X'_R is the two-term complex in degrees 0 and 1

$$\left(\mathcal{F}_{1/2}^\flat \oplus \cdots \oplus \mathcal{F}_{r-1/2}^\flat\right) \longrightarrow \left(\mathcal{F}_1 \oplus \cdots \oplus \mathcal{F}_{r-1}\right),$$

where the differential maps $(s_{1/2}, \ldots, s_{r-1/2})$ to $(s_{1/2} - s_{3/2}, \ldots, s_{r-3/2} - s_{r-1/2})$. The sheaf $\mathcal{F}_\bullet^\sharp$ fits into the exact triangle

$$\mathcal{F}_\bullet^\flat \to \mathcal{F}_0 \oplus \mathcal{F}_r \to \mathcal{F}_\bullet^\sharp. \tag{9.1.4}$$

The map $\mathcal{F}_\bullet^\flat \to \mathcal{F}_0 \oplus \mathcal{F}_r$ has components induced by the inclusions $\mathcal{F}_{1/2}^\flat \subset \mathcal{F}_0$ and $\mathcal{F}_{r-1/2}^\flat \subset \mathcal{F}_r$, with the sign conventions of [FYZ23, §9.2]. However, when $r = 0$, these are interpreted according to the following convention: $\mathrm{Hk}_S^0 = S$ with the identity correspondence, $\mathcal{F}_\bullet^\flat := \mathcal{F}_0$ (for $r = 0$ the displayed two-term complex would be the zero complex), and (9.1.4) is defined to be the split triangle $\mathcal{F}_0 \xrightarrow{(1,1)} \mathcal{F}_0 \oplus \mathcal{F}_0 \to \mathcal{F}_0$, so that $\mathcal{F}_\bullet^\sharp = \mathcal{F}_0$.

Replacing $\mathcal{F}^*_{\mathrm{univ}}$ in the definitions of $\mathcal{U},\mathcal{V},\mathcal{W}$ by $(\mathcal{F}^\flat_\bullet)^*$ gives perfect complexes $\mathcal{U}^\flat_{\mathrm{Hk}},\mathcal{V}^\flat_{\mathrm{Hk}},\mathcal{W}^\flat_{\mathrm{Hk}}$ on Hk^r_S, and hence derived vector bundles $\mathrm{Hk}^\flat_U,\mathrm{Hk}^\flat_V,\mathrm{Hk}^\flat_W$. Replacing $(\mathcal{F}^\flat_\bullet)^*$ by $(\mathcal{F}^\sharp_\bullet)^*$ gives $\mathrm{Hk}^\sharp_U,\mathrm{Hk}^\sharp_V,\mathrm{Hk}^\sharp_W$. Applying the same construction to the second exact sequence in (9.1.1) gives $\mathrm{Hk}^{\flat/\sharp}_{U^\perp}$, $\mathrm{Hk}^{\flat/\sharp}_{\widehat{V}}$, and $\mathrm{Hk}^{\flat/\sharp}_{W^\perp}$. These fit into derived Cartesian squares

$$\begin{array}{ccc}\mathrm{Hk}^\flat_U & \xrightarrow{f} & \mathrm{Hk}^\flat_V\\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle g}\\ \mathrm{Hk}^r_S & \xrightarrow[z_{\mathrm{Hk}^\flat_W}]{} & \mathrm{Hk}^\flat_W\end{array}\qquad\qquad \begin{array}{ccc}\mathrm{Hk}^\flat_{U^\perp} & \xrightarrow{\widehat{g}} & \mathrm{Hk}^\flat_{\widehat{V}}\\ {\scriptstyle\widehat{z}}\downarrow & & \downarrow{\scriptstyle\widehat{f}}\\ \mathrm{Hk}^r_S & \xrightarrow[\widehat{\pi}]{} & \mathrm{Hk}^\flat_{W^\perp}\end{array} \tag{9.1.5}$$

The analogous $\sharp$-squares are obtained from the same construction using $(\mathcal{F}^\sharp_\bullet)^*$. (In the second square of (9.1.5), the maps with $\widehat{()}$ are the Fourier duals of the corresponding $\sharp$-maps; for instance $\widehat{g}=\widehat{g^\sharp}$ under the identification $\widehat{\mathrm{Hk}^\sharp_W}=\mathrm{Hk}^\flat_{U^\perp}$.)

For $i\in\{0,r\}$, write $h^U_i,h^V_i,h^W_i,h^{U^\perp}_i,h^{\widehat{V}}_i,h^{W^\perp}_i$ for the corresponding base-change projections to the bundles. We use the maps

$$\begin{aligned} a_i &:= h^U_i\circ\widetilde{a}_i, \qquad & b_i &:= h^V_i\circ\widetilde{b}_i, \qquad & c_i &:= h^W_i\circ\widetilde{c}_i,\\ a^\perp_i &:= h^{U^\perp}_i\circ\widetilde{a}^\perp_i, \qquad & \beta_i &:= h^{\widehat{V}}_i\circ\widetilde{\beta}_i, \qquad & c^\perp_i &:= h^{W^\perp}_i\circ\widetilde{c}^\perp_i. \end{aligned}$$

The exact triangle (9.1.4) induces three derived Cartesian squares of derived vector bundles over Hk^r_S:

$$\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_U & & \\
 & {\scriptstyle\widetilde{a}_0}\swarrow & & \searrow{\scriptstyle\widetilde{a}_r} & \\
\widetilde{U}_0 & & & & \widetilde{U}_r\\
 & {\scriptstyle\widetilde{a}'_r}\searrow & & \swarrow{\scriptstyle\widetilde{a}'_0} & \\
 & & \mathrm{Hk}^\sharp_U & &
\end{array}\qquad
\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_V & & \\
 & {\scriptstyle\widetilde{b}_0}\swarrow & & \searrow{\scriptstyle\widetilde{b}_r} & \\
\widetilde{V}_0 & & & & \widetilde{V}_r\\
 & {\scriptstyle\widetilde{b}'_r}\searrow & & \swarrow{\scriptstyle\widetilde{b}'_0} & \\
 & & \mathrm{Hk}^\sharp_V & &
\end{array}\qquad
\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_W & & \\
 & {\scriptstyle\widetilde{c}_0}\swarrow & & \searrow{\scriptstyle\widetilde{c}_r} & \\
\widetilde{W}_0 & & & & \widetilde{W}_r\\
 & {\scriptstyle\widetilde{c}'_r}\searrow & & \swarrow{\scriptstyle\widetilde{c}'_0} & \\
 & & \mathrm{Hk}^\sharp_W & &
\end{array} \tag{9.1.6}$$

Switching the roles of $\mathcal{E}_1$ (resp. Q_1) and $\mathcal{E}_2$ (resp. Q_2), the exact triangle (9.1.4) induces three (derived) Cartesian squares of derived vector bundles over Hk^r_S:

$$\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_{U^\perp} & & \\
 & {\scriptstyle\widetilde{a}^\perp_0}\swarrow & & \searrow{\scriptstyle\widetilde{a}^\perp_r} & \\
\widetilde{U}^\perp_0 & & & & \widetilde{U}^\perp_r\\
 & {\scriptstyle(\widetilde{a}'_r)^\perp}\searrow & & \swarrow{\scriptstyle(\widetilde{a}'_0)^\perp} & \\
 & & \mathrm{Hk}^\sharp_{U^\perp} & &
\end{array}\qquad
\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_{\widehat{V}} & & \\
 & {\scriptstyle\widetilde{\beta}_0}\swarrow & & \searrow{\scriptstyle\widetilde{\beta}_r} & \\
\widehat{\widetilde{V}}_0 & & & & \widehat{\widetilde{V}}_r\\
 & {\scriptstyle\widetilde{\beta}'_r}\searrow & & \swarrow{\scriptstyle\widetilde{\beta}'_0} & \\
 & & \mathrm{Hk}^\sharp_{\widehat{V}} & &
\end{array}\qquad
\begin{array}{ccccc}
 & & \mathrm{Hk}^\flat_{W^\perp} & & \\
 & {\scriptstyle\widetilde{c}^\perp_0}\swarrow & & \searrow{\scriptstyle\widetilde{c}^\perp_r} & \\
\widetilde{W}^\perp_0 & & & & \widetilde{W}^\perp_r\\
 & {\scriptstyle(\widetilde{c}'_r)^\perp}\searrow & & \swarrow{\scriptstyle(\widetilde{c}'_0)^\perp} & \\
 & & \mathrm{Hk}^\sharp_{W^\perp} & &
\end{array} \tag{9.1.7}$$

The construction is summarized by the following pair of diagrams:

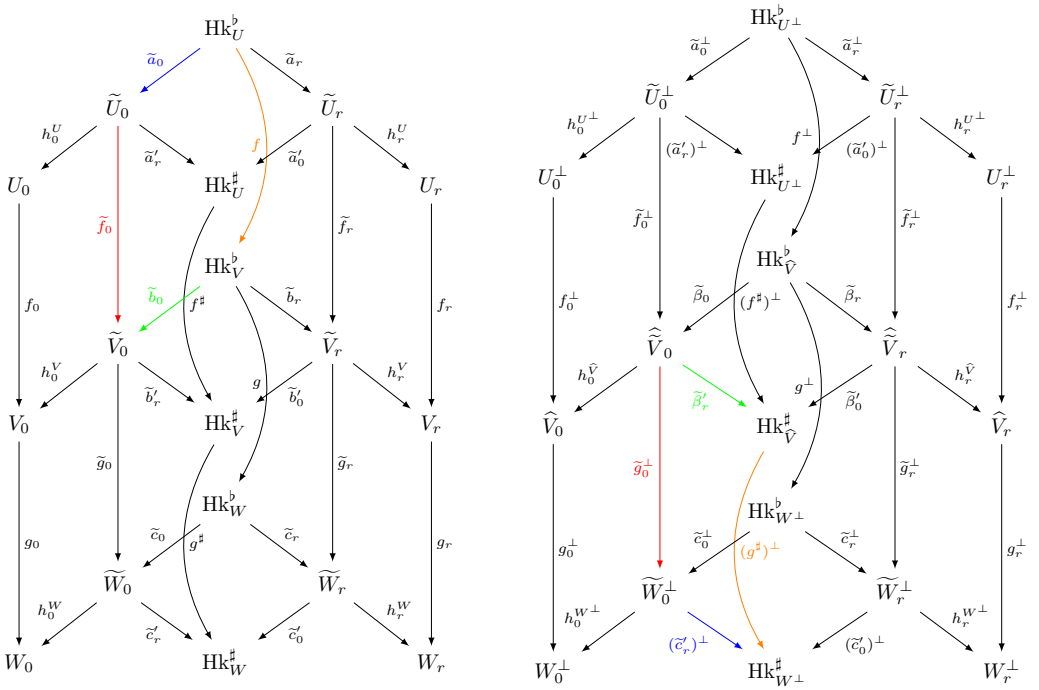

(9.1.8)

In each diagram:

- The maps in the columns are base changes (along the maps h_i, where appropriate) of the morphisms in the derived zero-fiber squares (9.1.3) and (9.1.5).
- The three diamonds in the middle are derived Cartesian.
- The four parallelograms on the left and right sides are derived Cartesian.

The diagram on the right is dual to the diagram on the left. The duality exchanges U with $W^\perp$, V with $\widehat{V}$, and W with $U^\perp$, and exchanges $\flat$ and $\sharp$ superscripts. Examples of dual morphisms are colored with the same color. By [FYZ23, §9.3], each of these diagrams is globally presented, so that the global-presentation formalism applies to them.

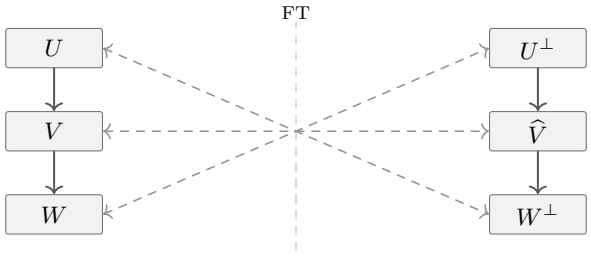

FIGURE 5. Cartoon of the dual pair of diagrams in (9.1.8). Fourier duality reverses the towers by exchanging U with $W^\perp$, V with $\widehat{V}$, and W with $U^\perp$.

9.2. **Departure from** [FYZ23]**.** The bestiary setup in §9.1 is identical to the one from [FYZ23, §9.1–9.2]. At this point, our proof will diverge in a crucial way from [FYZ23]. Let us indicate the shape of this departure, leaving some precise definitions for later.

The Trace Conjecture implies that

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U) = [\mathcal{Z}^r_{\mathcal{E}_1}], \qquad \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{U^\perp}) = [\mathcal{Z}^r_{\mathcal{E}_2}].$$

Therefore, where the quadratic function β is defined in §9.4, the trace

$$\mathrm{Tr}^{\mathrm{Sht}}\big(\pi_!(\mathfrak{c}_U \otimes \beta^* \,\mathrm{AS}_{\psi,S})\big)$$

is essentially (i.e., up to signs and powers of q) the higher theta series $\widetilde{Z}_m^{n,r}(\mathcal{G}, \mathcal{E}_1)$, while the trace

$$\mathrm{Tr}^{\mathrm{Sht}}\big(\pi_!^{\perp}(\mathfrak{c}_{U^\perp} \otimes (-\beta)^* \,\mathrm{AS}_{\psi,S})\big)$$

is essentially $\widetilde{Z}_m^{n,r}(\mathcal{G}, \mathcal{E}_2)$.

The strategy of [FYZ23] is to relate $\mathfrak{c}_U \otimes \beta^* \,\mathrm{AS}_{\psi,S}$ with $\mathfrak{c}_{U^\perp} \otimes (-\beta)^* \,\mathrm{AS}_{\psi,S}$ by pushing them both forward to the correspondence

$$V_0 \leftarrow \mathrm{Hk}_V^\flat \to V_r \tag{9.2.1}$$

and comparing them there. This comparison turns out to ultimately be a manifestation of the Plancherel formula.

However, this strategy has a key defect: in order to effect an actual comparison in $\mathrm{CH}_*(\mathrm{Sht}_S^r)$, we would have to push forward from (9.2.1) to

$$S \leftarrow \mathrm{Hk}_S^r \to S. \tag{9.2.2}$$

It turns out that the map from (9.2.1) to (9.2.2) is not pushable, so this is not possible. This is due to a certain "wildness" of the geometry of $\mathrm{Hk}_V^\flat$ when the legs meet the support of the torsion sheaf Q_1 defining V. On the complement of this locus, the map is in fact pushable, which is why this approach can work there. However, as the transverse pair $(\mathcal{E}_1, \mathcal{E}_2)$ varies, the support of the torsion sheaf Q_1 varies uncontrollably, so one is left only with automorphy over the generic fiber.

The key new insight here is that the relation between $\mathfrak{c}_U \otimes \beta^* \,\mathrm{AS}_{\psi,S}$ and $\mathfrak{c}_{U^\perp} \otimes (-\beta)^* \,\mathrm{AS}_{\psi,S}$ is already baked into their genetics: they are both *pulled back* from a pair of essentially self-dual cohomological correspondences on (9.2.1). From this starting point, the desired comparison can be effected by Fourier duality, without trying to push forward to (9.2.1).

9.3. **Geometric properties.** We establish some needed geometric properties of the maps in (9.1.8).

We use the terminology of *pushable* and *pullable* squares [FYZ23, §3.1 and §4.2]. A commutative square induces a comparison map α to the fibered product

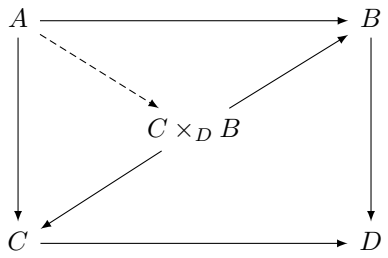

(9.3.1)

and we say that it is pushable if the (dashed) comparison map is proper, and pullable if it is quasi-smooth. We say that a map of correspondences

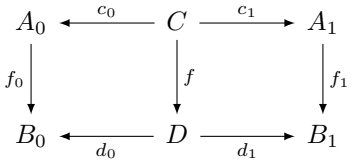

(9.3.2)

is *pushable* if the left square is pushable, and *pullable* if the right square is pullable. Following [FYZ23, §3.1], we call the virtual relative dimension of the comparison map of a pullable square its *defect*. Pullback of cohomological correspondences along a pullable map of correspondences carries a twist by the defect, which we track at each use.

Lemma 9.3.1. *(1) The maps of correspondences*

$$\begin{array}{ccccc} U_0 & \xleftarrow{a_0} & \mathrm{Hk}_U^\flat & \xrightarrow{a_r} & U_r \\ \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_r} \\ V_0 & \xleftarrow{b_0} & \mathrm{Hk}_V^\flat & \xrightarrow{b_r} & V_r \end{array} \qquad \begin{array}{ccccc} U_0^\perp & \xleftarrow{a_0^\perp} & \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{a_r^\perp} & U_r^\perp \\ \downarrow{\scriptstyle f_0^\perp} & & \downarrow{\scriptstyle f^\perp} & & \downarrow{\scriptstyle f_r^\perp} \\ \widehat{V}_0 & \xleftarrow{\beta_0} & \mathrm{Hk}_{\widehat{V}}^\flat & \xrightarrow{\beta_r} & \widehat{V}_r \end{array}$$

are pullable.

(2) The maps of correspondences

$$\begin{array}{ccccc} U_0 & \xleftarrow{a_0} & \mathrm{Hk}_U^\flat & \xrightarrow{a_r} & U_r \\ \downarrow{\scriptstyle \pi_0} & & \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi_r} \\ S_0 & \xleftarrow{h_0} & \mathrm{Hk}_S^r & \xrightarrow{h_r} & S_r \end{array} \qquad \begin{array}{ccccc} U_0^\perp & \xleftarrow{a_0^\perp} & \mathrm{Hk}_{U^\perp}^\flat & \xrightarrow{a_r^\perp} & U_r^\perp \\ \downarrow{\scriptstyle \pi_0^\perp} & & \downarrow{\scriptstyle \pi^\perp} & & \downarrow{\scriptstyle \pi_r^\perp} \\ S_0 & \xleftarrow{h_0} & \mathrm{Hk}_S^r & \xrightarrow{h_r} & S_r \end{array}$$

are pushable.

(3) The maps of correspondences

$$\begin{array}{ccccc} V_0 & \xleftarrow{b_0} & \mathrm{Hk}_V^\flat & \xrightarrow{b_r} & V_r \\ \downarrow{\scriptstyle g_0} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle g_r} \\ W_0 & \xleftarrow{c_0} & \mathrm{Hk}_W^\flat & \xrightarrow{c_r} & W_r \end{array} \qquad \begin{array}{ccccc} \widehat{V}_0 & \xleftarrow{\beta_0} & \mathrm{Hk}_{\widehat{V}}^\flat & \xrightarrow{\beta_r} & \widehat{V}_r \\ \downarrow{\scriptstyle g_0^\perp} & & \downarrow{\scriptstyle g^\perp} & & \downarrow{\scriptstyle g_r^\perp} \\ W_0^\perp & \xleftarrow{c_0^\perp} & \mathrm{Hk}_{W^\perp}^\flat & \xrightarrow{c_r^\perp} & W_r^\perp \end{array}$$

are pushable.

(4) The maps of correspondences

$$\begin{array}{ccccc} S_0 & \xleftarrow{h_0} & \mathrm{Hk}_S^r & \xrightarrow{h_r} & S_r \\ \downarrow{\scriptstyle z_{W_0}} & & \downarrow{\scriptstyle z_{\mathrm{Hk}_W^\flat}} & & \downarrow{\scriptstyle z_{W_r}} \\ W_0 & \xleftarrow{c_0} & \mathrm{Hk}_W^\flat & \xrightarrow{c_r} & W_r \end{array} \qquad \begin{array}{ccccc} S_0 & \xleftarrow{h_0} & \mathrm{Hk}_S^r & \xrightarrow{h_r} & S_r \\ \downarrow{\scriptstyle z_{W_0^\perp}} & & \downarrow{\scriptstyle z_{\mathrm{Hk}_{W^\perp}^\flat}} & & \downarrow{\scriptstyle z_{W_r^\perp}} \\ W_0^\perp & \xleftarrow{c_0^\perp} & \mathrm{Hk}_{W^\perp}^\flat & \xrightarrow{c_r^\perp} & W_r^\perp \end{array}$$

are pullable.

Proof. (4) We focus first on the left map of correspondences. By definition, we need to show pullability of the right square, in other words the map from Hk_S^r to the fibered product of $\mathrm{Hk}_W^\flat$ and S_r over W_r is quasi-smooth. The canonical base-change isomorphism

$$\mathrm{Hk}_W^\flat \times_{W_r} S_r \cong \mathrm{Hk}_W^\flat \times_{\widetilde{W}_r} \mathrm{Hk}_S^r$$

identifies its comparison map with the comparison map of the diagram

$$\begin{array}{ccc} \mathrm{Hk}_S^r & \xrightarrow{\mathrm{Id}} & \mathrm{Hk}_S^r \\ {\scriptstyle z_{\mathrm{Hk}_W^\flat}}\downarrow & & \downarrow{\scriptstyle z_{\widetilde{W}_r}} \\ \mathrm{Hk}_W^\flat & \xrightarrow[\widetilde{c}_r]{} & \widetilde{W}_r \end{array} \tag{9.3.3}$$

It therefore suffices to show that (9.3.3) is right pullable. The morphism $\widetilde{c}_r : \mathrm{Hk}_W^\flat \to \widetilde{W}_r$ is smooth (a vector bundle projection) by [FYZ23, Lemma 9.1.3(3)], while $\mathrm{Id}_{\mathrm{Hk}_S^r}$ is smooth. Hence (9.3.3) is pullable by [FYZ23, Example 3.1.3]. The same argument applies to the other map of correspondences, finishing the proof of (4).

We next deduce (1) from (4). By symmetry it suffices to handle the left map of correspondences. We similarly reduce to pullability of the diagram

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{\widetilde{a}_r} & \widetilde{U}_r \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle \widetilde{f}_r} \\ \mathrm{Hk}_V^\flat & \xrightarrow[\widetilde{b}_r]{} & \widetilde{V}_r. \end{array} \tag{9.3.4}$$

We have derived Cartesian squares

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{\pi} & \mathrm{Hk}_S^r \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle z_{\mathrm{Hk}_W^\flat}} \\ \mathrm{Hk}_V^\flat & \xrightarrow[g]{} & \mathrm{Hk}_W^\flat \end{array} \qquad \begin{array}{ccc} \widetilde{U}_r & \longrightarrow & \mathrm{Hk}_S^r \\ {\scriptstyle \widetilde{f}_r}\downarrow & & \downarrow{\scriptstyle z_{\widetilde{W}_r}} \\ \widetilde{V}_r & \xrightarrow[\widetilde{g}_r]{} & \widetilde{W}_r. \end{array} \tag{9.3.5}$$

The first square in (9.3.5) identifies $\mathrm{Hk}_U^\flat$ with $\mathrm{Hk}_V^\flat \times_{\mathrm{Hk}_W^\flat} \mathrm{Hk}_S^r$, and the second identifies $\widetilde{U}_r$ with $\widetilde{V}_r \times_{\widetilde{W}_r} \mathrm{Hk}_S^r$. Hence the comparison map of (9.3.4) is the derived base change of the comparison map of (9.3.3), and quasi-smoothness is stable under derived base change.

(3) follows from (1) by duality: the first map of correspondences in (1) is dual to the second map in (3), and the second map in (1) is dual to the first map in (3). Indeed, dualization exchanges the $\flat$- and $\sharp$-vertices. It also exchanges pullability and pushability by [FYZ23, Lemma 7.2.1] in its varying-base form (see the end of [FYZ23, §7.2]).

(2) is deduced from (3) by base change, similarly to how (1) was deduced from (4). Indeed, for the left square in (2), the comparison map $\widetilde{a}_0 \colon \mathrm{Hk}_U^\flat \to \widetilde{U}_0$ is the top horizontal arrow in the Cartesian square

$$\begin{array}{ccc} \mathrm{Hk}_U^\flat & \xrightarrow{\widetilde{a}_0} & \widetilde{U}_0 \\ {\scriptstyle f}\downarrow & & \downarrow \\ \mathrm{Hk}_V^\flat & \longrightarrow & \mathrm{Hk}_W^\flat \times_{W_0} V_0. \end{array} \tag{9.3.6}$$

The lower horizontal arrow in (9.3.6) is the left-square comparison map for (3), so $\widetilde{a}_0$ is its derived base change along $z_{\mathrm{Hk}_W^\flat}$. We conclude using that properness is stable under derived base change. □

In the proof of Lemma 9.3.1, we showed that the map of correspondences

$$\begin{array}{ccccc} \widetilde{U}_0 & \xleftarrow{\widetilde{a}_0} & \mathrm{Hk}_U^\flat & \xrightarrow{\widetilde{a}_r} & \widetilde{U}_r \\ \downarrow & & \downarrow & & \downarrow \\ \widetilde{V}_0 & \xleftarrow{\widetilde{b}_0} & \mathrm{Hk}_V^\flat & \xrightarrow{\widetilde{b}_r} & \widetilde{V}_r \end{array}$$

is pullable.

Lemma 9.3.2. *The map of correspondences*

$$\begin{array}{ccccc} \widetilde{V}_0 & \xleftarrow{\widetilde{b}_0} & \mathrm{Hk}_V^\flat & \xrightarrow{\widetilde{b}_r} & \widetilde{V}_r \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{Hk}_S^r & \xleftarrow{\mathrm{Id}} & \mathrm{Hk}_S^r & \xrightarrow{\mathrm{Id}} & \mathrm{Hk}_S^r \end{array}$$

is pullable. Consequently, the map of correspondences

$$\begin{array}{ccccc} \widetilde{U}_0 & \xleftarrow{\widetilde{a}_0} & \mathrm{Hk}_U^\flat & \xrightarrow{\widetilde{a}_r} & \widetilde{U}_r \\ \downarrow & & \downarrow & & \downarrow \\ \mathrm{Hk}_S^r & \xleftarrow{\mathrm{Id}} & \mathrm{Hk}_S^r & \xrightarrow{\mathrm{Id}} & \mathrm{Hk}_S^r \end{array}$$

is pullable. The same statements hold after replacing $\widetilde{U}_i$ (resp. $\widetilde{V}_i$) with U_i (resp. V_i) and the bottom left (resp. bottom right) Hk^r_S *with* S.

Proof. The last sentence follows immediately from the definition of pullability, as in the proof of Lemma 9.3.1.

The second assertion follows from stability of pullable maps under composition (two-out-of-three for pullable squares, [FYZ23, §3.1]). We now focus on proving the first statement. This amounts to showing that the square

$$\begin{array}{ccc} \mathrm{Hk}^\flat_V & \xrightarrow{\widetilde{b}_r} & \widetilde{V}_r \\ \downarrow & & \downarrow \\ \mathrm{Hk}^r_S & \xrightarrow[\mathrm{Id}]{} & \mathrm{Hk}^r_S \end{array}$$

is pullable, which is equivalent to $\widetilde{b}_r$ being quasi-smooth. This holds by [FYZ23, Lemma 9.1.3(2)]. □

Lemma 9.3.3. *For each i, the maps*

$$\widetilde{b}_i\colon \mathrm{Hk}^\flat_V \longrightarrow \widetilde{V}_i$$

of (9.1.8) *are quasi-smooth of virtual relative dimension* 0.

Proof. The quasi-smoothness is [FYZ23, Lemma 9.1.3(2)]. For the relative dimension, recall that over an R-point $(\mathcal{F}_\star)$ of Hk^r_S, the map $\widetilde{b}_i$ is induced by the map $\mathcal{F}_i^* \longrightarrow \mathcal{F}_\bullet^{\flat *}$. Write $\mathcal{F}^\flat_{\mathrm{univ},\bullet}$ for the universal complex whose pullback at $(\mathcal{F}_\star)$ is $\mathcal{F}^\flat_\bullet$, and define the universal perfect complex

$$T_{i,\mathrm{univ}} := \mathrm{cofib}\big(h_i^*\mathcal{F}^*_{\mathrm{univ}} \longrightarrow \mathcal{F}^{\flat *}_{\mathrm{univ},\bullet}\big) \quad \text{in } \mathrm{Perf}(X' \times \mathrm{Hk}^r_S),$$

and let T_i be its pullback to X'_R. At a classical R-point of Hk^r_S, T_i is a torsion sheaf supported at the legs. The relative tangent complex of $\widetilde{b}_i$ is the pullback from Hk^r_S of $\underline{\mathrm{RHom}(T_{i,\mathrm{univ}}, Q_1)}$. The virtual relative dimension may be checked at geometric points. There Riemann–Roch gives $\chi_{\mathrm{Eul}}(X', \mathrm{RHom}(T_i, Q_1)) = 0$, because both T_i and Q_1 are torsion. □

Lemma 9.3.4. *Assume $r = 1$. For $i = 0, r$, the maps*

$$\widetilde{b}_i\colon \mathrm{Hk}^\flat_V \longrightarrow \widetilde{V}_i$$

are LCI morphisms of virtual relative dimension 0.

Proof. When $r = 1$, the complex $\mathcal{F}^\flat_\bullet$ in (9.1.4) is the vector bundle $\mathcal{F}^\flat_{1/2}$ placed in degree 0. Thus $(\mathcal{F}^\flat_\bullet)^* = (\mathcal{F}^\flat_{1/2})^*$ is also a vector bundle. Since Q_1 is torsion, the relative Hom complexes

$$\underline{\mathrm{RHom}\big((\mathcal{F}^\flat_\bullet)^*, Q_1\big)} \quad \text{and} \quad \underline{\mathrm{RHom}(\mathcal{F}_i^*, Q_1)}$$

on Hk^1_S are represented by locally free coherent sheaves in degree 0. Hence $\mathrm{Hk}^\flat_V$ and $\widetilde{V}_i$ are ordinary vector bundles over Hk^1_S, and $\widetilde{b}_i$ is a morphism of vector bundles over Hk^1_S. A map of vector bundles over the same base is always LCI, since it can be factored as the composition of a regular embedding into its graph, followed by a smooth projection. Both bundles here are locally free of the same rank $n \cdot \mathrm{length}(Q_1)$, so the virtual relative dimension is 0. □

9.4. **Gaussians.** Abbreviate $\widetilde{S} := \mathrm{Hk}^r_S$, $V^\flat := \mathrm{Hk}^\flat_V$ and $V^\sharp := \mathrm{Hk}^\sharp_V$.

For $i = 0, r$, let $\widetilde{\pi}_i\colon \widetilde{V}_i \to \widetilde{S}$ be the projection to the base of the derived vector bundle, and let $h_i^V : \widetilde{V}_i \to V_i \cong V$ be the base change of $h_i\colon \widetilde{S} \to S$.

We have a derived Cartesian diagram of vector bundles over $\widetilde{S}$.

$$\begin{array}{ccccc} & & V^\flat & & \\ & \swarrow^{\widetilde{b}_0} & & \searrow^{\widetilde{b}_r} & \\ \widetilde{V}_0 & & & & \widetilde{V}_r \\ & \searrow_{\widetilde{b}'_r} & & \swarrow_{\widetilde{b}'_0} & \\ & & V^\sharp. & & \end{array} \tag{9.4.1}$$

Following [FYZ23, §2], the torsion sheaf Q_1 carries a canonical Hermitian structure. In the notation of *loc. cit.* one has $Q_1 = \sigma^* Q$, and the pullback of the Hermitian structure $h_{12}\colon Q \xrightarrow{\sim} \sigma^* Q^*$ is

$$\sigma^* h_{12}\colon Q_1 = \sigma^* Q \xrightarrow{\sim} Q^* \cong \sigma^* Q_1^*, \qquad Q^* := \mathcal{E}xt^1_{X'}(Q, \mathcal{O}_{X'}).$$

We continue to denote this transported form on Q_1 by h_{12}. (There are two natural choices h_{12} and h_{21}, related by $h_{12} = -h_{21}$ [FYZ23, §2]; we fix the normalization h_{12} throughout, following [FYZ23].) Tensoring h_{12} with the Hermitian structure of $\mathcal{F}_{\mathrm{univ}}$ and applying relative Serre duality along $X' \times S \to S$ produces a symmetric self-duality

$$h_V\colon V \xrightarrow{\sim} \widehat{V},$$

whose associated quadratic function on $\mathbf{F}_q$-points is the quadratic form $\mathfrak{q}_{12}$ of [FYZ23, §2].

We single out the features that we will use.

- The symmetric self-duality $h_V\colon V \xrightarrow{\sim} \widehat{V}$ pulls back to symmetric self-dualities
$$h_{V,i}\colon \widetilde{V}_i \xrightarrow{\sim} \widehat{\widetilde{V}_i}$$
over $\widetilde{S}$, for $i = 0, r$.
- There is an identification $V^\flat \xrightarrow{\sim} \widehat{V^\sharp}$. Thus, the diagram (9.4.1) is self-dual.
- With respect to these identifications, $\widetilde{b}'_r$ is the dual morphism to $\widetilde{b}_0$, and $\widetilde{b}'_0$ is the dual morphism to $\widetilde{b}_r$. (These three assertions are established in [FYZ23, §9.1–9.2], transported along the self-duality h_V defined above.)
- The diagram (9.4.1), the tensoring square (7.1.1) with $f = \widetilde{b}_r$, and their Fourier-dual diagrams are globally presented in the sense of [FYZ23, Definition 6.3.1].
- The morphism $\widetilde{b}_r$ is quasi-smooth of virtual relative dimension 0 by Lemma 9.3.3, and $\widetilde{b}'_0$ is separated, so that the relative derived fundamental class $[\widetilde{b}_r]\colon \widetilde{b}_r^* \to \widetilde{b}_r^!$ and the forget-supports map $\mathrm{can}(\widetilde{b}'_0)\colon \widetilde{b}'_{0!} \to \widetilde{b}'_{0*}$ are defined.

Write $d(V/S)$ for the virtual relative dimension of $V \to S$, equivalently of either $\widetilde{\pi}_i\colon \widetilde{V}_i \to \widetilde{S}$.

9.4.1. *Gaussian sheaves.* Let $\mathrm{AS}_{\psi_0} \in \mathcal{D}_{\mathrm{mot}}(\mathbf{A}^1_{\mathbf{F}_q}; \overline{\mathbf{Q}})$ be the motivic Artin–Schreier sheaf of (6.1.3), and let $\mathrm{AS}_{\psi,S}$ and $\mathrm{AS}_{\psi,\widetilde{S}}$ denote its pullbacks to $\mathbf{A}^1_S$ and $\mathbf{A}^1_{\widetilde{S}}$. Applying Construction 6.4.1 to h_V gives

$$\beta := \beta_{h_V}\colon V \longrightarrow \mathbf{A}^1_S, \qquad \beta^\vee := \beta^\vee_{h_V}\colon \widehat{V} \longrightarrow \mathbf{A}^1_S.$$

Set

$$\mathscr{G}_\beta := \beta^* \,\mathrm{AS}_{\psi,S} \in \mathcal{D}_{\mathrm{mot}}(V; \overline{\mathbf{Q}}), \qquad \mathscr{G}_{-\beta} := (-\beta)^* \,\mathrm{AS}_{\psi,S} \in \mathcal{D}_{\mathrm{mot}}(V; \overline{\mathbf{Q}}).$$

For $i = 0, r$, let

$$\beta_i := (h_i^V)^* \beta\colon \widetilde{V}_i \longrightarrow \mathbf{A}^1_{\widetilde{S}}.$$

Equivalently, β_i is the quadratic function induced by the pulled-back self-duality $h_{V,i}$. Let $\beta_i^\vee$ denote the corresponding dual quadratic function supplied by Construction 6.4.1.

Lemma 9.4.1. *We have*

$$\beta_0 \circ \widetilde{b}_0 = \beta_r \circ \widetilde{b}_r \quad \textit{as morphisms } V^\flat \to \mathbf{A}^1_{\widetilde{S}}. \tag{9.4.2}$$

Proof. Indeed, writing $\iota_\flat\colon V^\flat \cong \widehat{V^\sharp}$ for the identification above and $\iota_\sharp\colon V^\sharp \cong \widehat{V^\flat}$ for its dual, the duality assertions give

$$\widehat{\widetilde{b}_0} \circ h_{V,0} = \iota_\sharp \circ \widetilde{b}'_r \qquad \text{and} \qquad \widehat{\widetilde{b}_r} \circ h_{V,r} = \iota_\sharp \circ \widetilde{b}'_0.$$

Hence

$$\beta_0(\widetilde{b}_0 x) = \langle \iota_\sharp\, \widetilde{b}'_r\, \widetilde{b}_0\, x,\ x\rangle, \qquad \text{and} \qquad \beta_r(\widetilde{b}_r x) = \langle \iota_\sharp\, \widetilde{b}'_0\, \widetilde{b}_r\, x,\ x\rangle,$$

which agree by the commutativity of the diagram (9.4.1). □

Define

$$\mathscr{G}_{\beta,i} := \beta_i^*\, \mathrm{AS}_{\psi,\widetilde{S}} \in \mathcal{D}_{\mathrm{mot}}(\widetilde{V}_i; \overline{\mathbf{Q}}), \qquad \mathscr{G}_{-\beta,i} := (-\beta_i)^*\, \mathrm{AS}_{\psi,\widetilde{S}} \in \mathcal{D}_{\mathrm{mot}}(\widetilde{V}_i; \overline{\mathbf{Q}}). \tag{9.4.3}$$

The equality $h_{V,i}^* \beta_i^\vee = \beta_i$ identifies the positive-dual Gaussian with $\mathscr{G}_{\beta,i}$, while $h_{V,i}^*(-\beta_i^\vee) = -\beta_i$, so the transported Fourier-dual Gaussian is $\mathscr{G}_{-\beta,i}$.

9.4.2. *Cohomological correspondences.* Pulling $\mathrm{AS}_{\psi,\widetilde{S}}$ back along (9.4.2) gives

$$u_\beta\colon \widetilde{b}_0^* \mathscr{G}_{\beta,0} \xrightarrow{\sim} \widetilde{b}_r^* \mathscr{G}_{\beta,r}. \tag{9.4.4}$$

Applying the same construction to $(-\beta_0) \circ \widetilde{b}_0 = (-\beta_r) \circ \widetilde{b}_r$ gives

$$u_{-\beta}\colon \widetilde{b}_0^* \mathscr{G}_{-\beta,0} \xrightarrow{\sim} \widetilde{b}_r^* \mathscr{G}_{-\beta,r}. \tag{9.4.5}$$

By Lemma 9.3.2, we may define a cohomological correspondence $\mathfrak{c}_{\widetilde{V}} \in \mathrm{Corr}_{V^\flat}(\overline{\mathbf{Q}}_{\widetilde{V}_0}, \overline{\mathbf{Q}}_{\widetilde{V}_r})$ as the pullback of the identity cohomological correspondence on $\widetilde{S}$. Concretely, $\mathfrak{c}_{\widetilde{V}}$ is given explicitly by the composition

$$\mathfrak{c}_{\widetilde{V}}\colon \widetilde{b}_0^* \overline{\mathbf{Q}}_{\widetilde{V}_0} \cong \widetilde{b}_r^* \overline{\mathbf{Q}}_{\widetilde{V}_r} \xrightarrow{[\widetilde{b}_r]} \widetilde{b}_r^! \overline{\mathbf{Q}}_{\widetilde{V}_r}. \tag{9.4.6}$$

Tensoring (9.4.6) with (9.4.4) gives

$$\mathfrak{c}_{\widetilde{V}} \otimes u_\beta \in \mathrm{Corr}_{V^\flat}\big(\mathscr{G}_{\beta,0}, \mathscr{G}_{\beta,r}\big), \tag{9.4.7}$$

and tensoring with (9.4.5) gives

$$\mathfrak{c}_{\widetilde{V}} \otimes u_{-\beta} \in \mathrm{Corr}_{V^\flat}\big(\mathscr{G}_{-\beta,0}, \mathscr{G}_{-\beta,r}\big).$$

Both tensor products are in the sense of Definition 7.1.1.

9.4.3. *Fourier transform of Gaussians.* Let $\pi\colon V \to S$ be the structure morphism. Set

$$\mathfrak{G}_Q := \pi_! \mathscr{G}_{\beta[d(V/S)]} \in \mathcal{D}_{\mathrm{mot}}(S; \overline{\mathbf{Q}}), \qquad \mathfrak{G}_{Q,i} := h_i^* \mathfrak{G}_Q \in \mathcal{D}_{\mathrm{mot}}(\widetilde{S}; \overline{\mathbf{Q}}) \quad (0 \le i \le r). \tag{9.4.8}$$

By the definition of V in §9.1, it is the derived vector bundle associated to the perfect complex $\mathcal{V} = \underline{\mathrm{RHom}(\mathcal{F}^*_{\mathrm{univ}}, Q_1)}$. Since $\mathcal{F}^*_{\mathrm{univ}}$ is locally free and Q_1 is a fixed torsion sheaf on X', this is a locally free sheaf in degree 0, so V is a vector bundle in the classical sense. Hence Lemma 6.4.3 implies that $\mathfrak{G}_Q$, and with it each $\mathfrak{G}_{Q,i}$, is tensor-invertible.

Base change for $\widetilde{V}_i = V \times_{S,h_i} \widetilde{S}$, followed by Lemma 6.4.2, gives a canonical isomorphism

$$[2]^*_{\widetilde{V}_i} h_{V,i}^*\, \mathrm{FT}_{\widetilde{V}_i}(\mathscr{G}_{\beta,i}) \xrightarrow{\sim} \mathscr{G}_{-\beta,i} \otimes \widetilde{\pi}_i^* \mathfrak{G}_{Q,i}. \tag{9.4.9}$$

The scalar automorphism [2] acts on $\widetilde{V}_0, \widetilde{V}_r, V^\flat$, and $V^\sharp$. We write $[2]^*_V$ for the induced scalar pullback on the corresponding sheaves, correspondences, and co-correspondences.

We define a normalized Fourier transform of cohomological correspondences

$$\begin{aligned} [2]^*_V \mathbb{T}_{[-d(V/S)](-d(V/S))}\, \mathrm{FT}_{V^\flat}\colon\ & \mathrm{Corr}_{V^\flat}(\mathscr{G}_{\beta,0}, \mathscr{G}_{\beta,r}) \\ & \longrightarrow \mathrm{Corr}_{V^\flat}\Big((\mathscr{G}_{-\beta,0} \otimes \widetilde{\pi}_0^* \mathfrak{G}_{Q,0})_{[-d(V/S)](-d(V/S))}, \\ & \qquad\qquad (\mathscr{G}_{-\beta,r} \otimes \widetilde{\pi}_r^* \mathfrak{G}_{Q,r})_{[-d(V/S)](-d(V/S))} \Big), \end{aligned} \tag{9.4.10}$$

the composite of the Fourier transform (7.4.4) for (9.4.1) (whose legs have virtual relative dimension 0 by Lemma 9.3.3, so that no shift or twist appears), the identification $V^\flat \cong \widehat{V^\sharp}$, (9.4.9), and the displayed common shift and Tate twist.

Lemma 9.4.2. *Assume $r = 1$. For each sign $\epsilon = \pm 1$, with $u_{+\beta} := u_\beta$, the morphism*

$$[\widetilde{b}_r]_{\epsilon\beta} \colon \widetilde{b}_r^* \mathscr{G}_{\epsilon\beta,r} \longrightarrow \widetilde{b}_r^! \mathscr{G}_{\epsilon\beta,r}$$

obtained from the relative derived fundamental class $[\widetilde{b}_r]$ is an isomorphism. Consequently, after applying the isomorphisms $u_{\epsilon\beta} \colon \widetilde{b}_0^* \mathscr{G}_{\epsilon\beta,0} \xrightarrow{\sim} \widetilde{b}_r^* \mathscr{G}_{\epsilon\beta,r}$, there are canonical identifications

$$\mathrm{Corr}_{V^\flat}(\mathscr{G}_{\epsilon\beta,0}, \mathscr{G}_{\epsilon\beta,r}) = \mathrm{Hom}_{V^\flat}(\widetilde{b}_0^* \mathscr{G}_{\epsilon\beta,0}, \widetilde{b}_r^! \mathscr{G}_{\epsilon\beta,r}) \cong \mathrm{End}_{V^\flat}(\widetilde{b}_0^* \mathscr{G}_{\epsilon\beta,0}) \cong \mathrm{H}^0(V^\flat, \overline{\mathbf{Q}}_{V^\flat}). \tag{9.4.11}$$

Proof. By Lemma 9.3.3, the morphism $\widetilde{b}_r$ is quasi-smooth of virtual relative dimension 0. Thus the relative derived fundamental class has the form $[\widetilde{b}_r] \colon \widetilde{b}_r^* \to \widetilde{b}_r^!$ without any additional shift or Tate twist.

Let $p_{\epsilon,r} \colon \widetilde{V}_{\epsilon,r} \to \widetilde{V}_r$ be the finite étale Artin–Schreier cover obtained by pulling back (6.1.1) along the composite of $\epsilon\beta_r \colon \widetilde{V}_r \to \mathbf{A}^1_{\widetilde{S}}$ with the projection $\mathbf{A}^1_{\widetilde{S}} \to \mathbf{A}^1_{\mathbf{F}_q}$, with $\epsilon = \pm 1$. Then $\mathscr{G}_{\epsilon\beta,r} = e_{\psi_0}\, p_{\epsilon,r!} \overline{\mathbf{Q}}_{\widetilde{V}_{\epsilon,r}}$ where e_{ψ_0} is the idempotent of $\overline{\mathbf{Q}}[\mathbf{F}_p]$ defined in (6.1.2). Form the Cartesian square

$$\begin{array}{ccc} \widetilde{V}_\epsilon^\flat & \xrightarrow{\widetilde{b}_{r,\epsilon}} & \widetilde{V}_{\epsilon,r} \\ {\scriptstyle q_\epsilon}\downarrow & & \downarrow{\scriptstyle p_{\epsilon,r}} \\ V^\flat & \xrightarrow[\widetilde{b}_r]{} & \widetilde{V}_r. \end{array}$$

Since $p_{\epsilon,r}$ is finite étale, proper base change gives isomorphisms

$$\widetilde{b}_r^* \mathscr{G}_{\epsilon\beta,r} \cong e_{\psi_0}\, q_{\epsilon!} \overline{\mathbf{Q}}_{\widetilde{V}_\epsilon^\flat}, \qquad \widetilde{b}_r^! \mathscr{G}_{\epsilon\beta,r} \cong e_{\psi_0}\, q_{\epsilon!} \widetilde{b}_{r,\epsilon}^! \overline{\mathbf{Q}}_{\widetilde{V}_{\epsilon,r}}.$$

By Lemma 9.3.4, the map $\widetilde{b}_{r,\epsilon}$ is a classical LCI morphism of virtual relative dimension 0. Absolute purity for rational Beilinson motives [CD19, Theorem 14.4.1] implies that $\widetilde{b}_{r,\epsilon}^! \overline{\mathbf{Q}}_{\widetilde{V}_{\epsilon,r}} \cong \overline{\mathbf{Q}}_{\widetilde{V}_\epsilon^\flat}$, and then that $\widetilde{b}_r^! \mathscr{G}_{\epsilon\beta,r} \cong \widetilde{b}_r^* \mathscr{G}_{\epsilon\beta,r} \cong \widetilde{b}_0^* \mathscr{G}_{\epsilon\beta,0}$. Since $\widetilde{b}_0^* \mathscr{G}_{\epsilon\beta,0}$ is $\otimes$-invertible, the result follows. □

Lemma 9.4.3. *Let*

$$\begin{aligned} X^\circ &:= X - \nu(|Q_1| \cup |Q_2|), \\ \mathrm{Hk}_S^{1,\circ} &:= \mathrm{Hk}_S^1 \times_X X^\circ, \\ V^{\flat,\circ} &:= V^\flat \times_{\mathrm{Hk}_S^1} \mathrm{Hk}_S^{1,\circ}, \\ V^{\sharp,\circ} &:= V^\sharp \times_{\mathrm{Hk}_S^1} \mathrm{Hk}_S^{1,\circ}. \end{aligned}$$

For $i = 0, 1$, let $\widetilde{V}_i := V \times_{S,h_i} \mathrm{Hk}_S^1$, and let $\widetilde{b}_i \colon V^\flat \to \widetilde{V}_i$ be the map over Hk_S^1. Then the restrictions $\widetilde{b}_i^\circ \colon V^{\flat,\circ} \xrightarrow{\sim} \widetilde{V}_i^\circ$ are isomorphisms.

Proof. We prove the assertion for $\widetilde{b}_i^\circ$; the proof for $\widetilde{b}_i'^\circ$ is identical, applied to the map $(\mathcal{F}_\bullet^\sharp)^* \to \mathcal{F}_i^*$ obtained by dualizing (9.1.4), whose cokernel is again torsion supported on $\Gamma_{x'} \cup \Gamma_{\sigma x'}$. Over an R-point $(\mathcal{F}_\star)$ of $\mathrm{Hk}_S^{1,\circ}$, the map $\widetilde{b}_i$ is induced by the morphism

$$\mathfrak{p}_i^* \colon \mathcal{F}_i^* \longrightarrow \mathcal{F}_\bullet^{\flat *}$$

from the one-leg local model. The cokernel of $\mathfrak{p}_i^*$ is a torsion sheaf supported set-theoretically on $\Gamma_{x'} \cup \Gamma_{\sigma x'}$, where x' denotes the leg of $(\mathcal{F}_\star)$: for $i = 0$ the support lies on $\Gamma_{x'}$, and for $i = 1$ on $\Gamma_{\sigma x'}$. Since the fiber product defining $\mathrm{Hk}_S^{1,\circ}$ is taken over X, the image of the leg in X lies in X°, so both x' and $\sigma(x')$ are disjoint from $|Q_1| \cup |Q_2|$ (note that this set is σ-stable, since $Q_2 \cong \sigma^* Q_1^*$). Therefore $\mathfrak{p}_i^*$ is an isomorphism in an open neighborhood of $|Q_1| \times \mathrm{Spec}\, R \subset X'_R$, and hence induces an isomorphism

$$\mathrm{RHom}_{X'_R}(\mathcal{F}_\bullet^{\flat *}, Q_1 \otimes R) \xrightarrow{\sim} \mathrm{RHom}_{X'_R}(\mathcal{F}_i^*, Q_1 \otimes R).$$

This identification is functorial in R, so it identifies the corresponding total spaces over $\mathrm{Hk}_S^{1,\circ}$. □

Definition 9.4.4. On $V^{\flat,\circ}$, Lemma 9.4.3 and (9.4.2) give an isometry

$$\widetilde{b}_1^\circ \circ (\widetilde{b}_0^\circ)^{-1} \colon (\widetilde{V}_0^\circ, \beta_0|_{\widetilde{V}_0^\circ}) \xrightarrow{\sim} (\widetilde{V}_1^\circ, \beta_1|_{\widetilde{V}_1^\circ})$$

We denote by

$$\xi^\circ\colon\ \widetilde{b}_0^*\widetilde{\pi}_0^*\mathfrak{G}_{Q,0}\Big|_{V^{\flat,\circ}} \xrightarrow{\sim} \widetilde{b}_1^*\widetilde{\pi}_1^*\mathfrak{G}_{Q,1}\Big|_{V^{\flat,\circ}} \tag{9.4.12}$$

the induced transport isomorphism of relative Gauss factors in $\mathcal{D}_{\mathrm{mot}}(V^{\flat,\circ};\overline{\mathbf{Q}})$.

Proposition 9.4.5. *(1) Assume $r=1$. There is a unique isomorphism $\xi_1\colon \widetilde{b}_0^*\widetilde{\pi}_0^*\mathfrak{G}_{Q,0}\xrightarrow{\sim}\widetilde{b}_1^*\widetilde{\pi}_1^*\mathfrak{G}_{Q,1}$ in $\mathcal{D}_{\mathrm{mot}}(V^\flat;\overline{\mathbf{Q}})$ whose restriction to $V^{\flat,\circ}$ is ξ°. Moreover, with respect to the normalized map* (9.4.10)*, we have*

$$\begin{aligned}&[2]_V^*\big(\mathbb{T}_{[-d(V/S)](-d(V/S))}\,\mathrm{FT}_{V^\flat}(\mathfrak{c}_{\widetilde{V}}\otimes u_\beta)\big) = (\mathfrak{c}_{\widetilde{V}}\otimes u_{-\beta})\otimes \xi_1[-d(V/S)](-d(V/S)) \\ &\in \mathrm{Corr}_{V^\flat}\big((\mathscr{G}_{-\beta,0}\otimes\widetilde{\pi}_0^*\mathfrak{G}_{Q,0})[-d(V/S)](-d(V/S)),\ (\mathscr{G}_{-\beta,1}\otimes\widetilde{\pi}_1^*\mathfrak{G}_{Q,1})[-d(V/S)](-d(V/S))\big),\end{aligned} \tag{9.4.13}$$

the tensor operation being that of Definition 7.1.1.

(2) For general r, the identity (9.4.13) *holds with ξ_1 replaced by the isomorphism $\xi_r\colon \widetilde{b}_0^*\widetilde{\pi}_0^*\mathfrak{G}_{Q,0}\xrightarrow{\sim}\widetilde{b}_r^*\widetilde{\pi}_r^*\mathfrak{G}_{Q,r}$ defined as the composite of the pullbacks of ξ_1 along the factorization of $(\widetilde{V}_0\leftarrow V^\flat\to\widetilde{V}_r)$ into one-leg correspondences (and $\xi_0:=\mathrm{Id}$), as depicted in* (9.4.14).

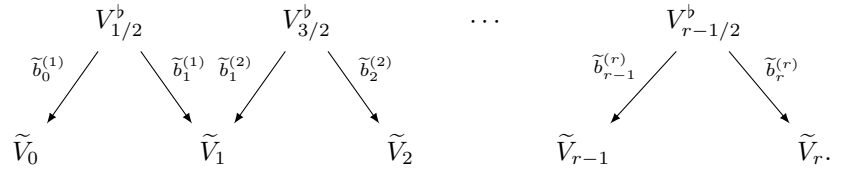

(9.4.14)

Proof. If $r=0$, both sides are the identity correspondence by the zero-leg convention, and $\xi_0=\mathrm{Id}$, so the assertion is immediate. Hence assume $r>0$.

First we will reduce to the case $r=1$. For each $1\le i\le r$, let $\rho_i\colon\widetilde{S}\to\mathrm{Hk}^1_S$ be the map recording the i-th one-step modification, and pull the one-leg identity back along ρ_i. This gives the correspondence

$$\widetilde{V}_{i-1}\leftarrow V^\flat_{i-1/2}\to\widetilde{V}_i$$

It is globally presented over the common base $\widetilde{S}$, being the base change of the globally presented one-leg diamond. By (9.4.14), these correspondences compose to $V^\flat$. Base change and functoriality of relative fundamental classes [Kha19, Theorems 3.12–3.13], together with Definition 7.2.1, identify $\mathfrak{c}_{\widetilde{V}}$ with the corresponding composite of the one-leg cohomological correspondences defined for $r=1$. Lemma 7.2.2 shows that the cohomological correspondence $\mathfrak{c}_{\widetilde{V}}\otimes u_\beta$ also factors as the composite of the $r=1$ case, while Lemma 7.4.2 shows that Fourier transform is compatible with compositions of cohomological correspondences. Thus we reduce to the case $r=1$, which we now assume.

By Lemma 6.4.3, each $\mathfrak{G}_{Q,i}$ is tensor-invertible; set

$$\mathcal{L}:=\big(\widetilde{b}_0^*\widetilde{\pi}_0^*\mathfrak{G}_{Q,0}\big)^{-1}\otimes\widetilde{b}_1^*\widetilde{\pi}_1^*\mathfrak{G}_{Q,1}\in\mathcal{D}_{\mathrm{mot}}(V^\flat;\overline{\mathbf{Q}}).$$

The object $\mathcal{L}$ is tensor-invertible. Tensor the identifications of Lemma 9.4.2 for $\epsilon=-1$ with $\widetilde{b}_i^*\widetilde{\pi}_i^*\mathfrak{G}_{Q,i}$ and apply the common shift and Tate twist in (9.4.10). This identifies the target group of (9.4.10) with $\mathrm{H}^0(V^\flat,\mathcal{L})$, sending s to $(\mathfrak{c}_{\widetilde{V}}\otimes u_{-\beta})\otimes s[-d(V/S)](-d(V/S))$; here $\mathfrak{c}_{\widetilde{V}}\otimes u_{-\beta}$ corresponds to 1 under (9.4.11). There is therefore a unique $s\in\mathrm{H}^0(V^\flat,\mathcal{L})$ such that

$$[2]_V^*\big(\mathbb{T}_{[-d(V/S)](-d(V/S))}\,\mathrm{FT}_{V^\flat}(\mathfrak{c}_{\widetilde{V}}\otimes u_\beta)\big)=(\mathfrak{c}_{\widetilde{V}}\otimes u_{-\beta})\otimes s[-d(V/S)](-d(V/S)).$$

We must show that s is an isomorphism restricting on $V^{\flat,\circ}$ to transport along the isometries of Lemma 9.4.3. Let $\pi^\flat\colon V^\flat\to\widetilde{S}=\mathrm{Hk}^1_S$ be the structure morphism; then $\pi^\flat=\widetilde{\pi}_i\circ\widetilde{b}_i$ for $i=0,1$. Lemma 6.4.4, applied to (V,h_V), gives

$$\mathfrak{G}_Q\cong\mathfrak{g}_{d(V/S)}\big|_S\otimes\varepsilon_Q,$$

where ε_Q is tensor-invertible and becomes trivial after a finite étale base change. Cancelling the constant Gauss factors in the definition of $\mathcal{L}$ gives

$$\mathcal{L}\cong(\pi^\flat)^*\big(h_0^*\varepsilon_Q^{-1}\otimes h_1^*\varepsilon_Q\big).$$

The RHS is likewise trivial after a finite étale base change. Moreover, $V^{\flat,\circ}$ is dense in every connected component, since each component of Hk^1_S dominates X through the smooth leg map and hence meets the dense open $\mathrm{Hk}^{1,\circ}_S$. It follows that restriction to $V^{\flat,\circ}$ is injective on $\mathrm{H}^0(-,\mathcal{L})$.

On $V^{\flat,\circ}$, the maps $\widetilde{b}_i^\circ$ identify the correspondence with the graph of an isometry of quadratic bundles. Naturality of (9.4.9) then identifies the restriction of the left side with $(\mathfrak{c}_{\widetilde{V}} \otimes u_{-\beta})|_{V^{\flat,\circ}}$ tensored with the tautological transport isomorphism ξ°. Thus $s|_{V^{\flat,\circ}} = \xi^\circ$. This is an isomorphism, and the finite étale local triviality used above implies that s is an isomorphism. Taking $\xi_1 := s$ completes the proof. □

9.4.4. *The fundamental identity.* We now pass from the $\widetilde{S}$-relative correspondence $\widetilde{V}_0 \leftarrow V^\flat \to \widetilde{V}_r$ to the varying-base correspondence

$$V_0 \longleftarrow V^\flat \longrightarrow V_r.$$

Recall that

$$b_i := h_i^V \circ \widetilde{b}_i \colon V^\flat \longrightarrow V_i \qquad (i = 0, r).$$

The maps $h_i \colon \widetilde{S} \to S$ are smooth, schematic, and separated by [FYZ24, Lemma 6.9]; hence their base changes $h_i^V \colon \widetilde{V}_i \to V_i$ are smooth of virtual relative dimension $d(h_i)$. By Lemma 9.3.3, the maps $\widetilde{b}_i$ are quasi-smooth of virtual relative dimension 0. Thus $b_i = h_i^V \circ \widetilde{b}_i$ is quasi-smooth of virtual relative dimension $d(h_i)$.

For $\mathcal{K}_0 \in \mathcal{D}_{\mathrm{mot}}(V_0; \overline{\mathbf{Q}})$ and $\mathcal{K}_r \in \mathcal{D}_{\mathrm{mot}}(V_r; \overline{\mathbf{Q}})$, we use the following cohomological correspondence group with varying base:

$$\begin{aligned} \mathrm{Corr}_{V^\flat}(\mathcal{K}_0, \mathcal{K}_r) &:= \mathrm{Hom}_{V^\flat}(b_0^* \mathcal{K}_0, b_r^! \mathcal{K}_r) \\ &= \mathrm{Hom}_{V^\flat}\big(\widetilde{b}_0^*(h_0^V)^* \mathcal{K}_0, \widetilde{b}_r^!(h_r^V)^! \mathcal{K}_r\big). \end{aligned} \tag{9.4.15}$$

The equality $\beta \circ b_0 = \beta \circ b_r$, viewed as an equality of morphisms $V^\flat \to \mathbf{A}^1_{\mathbf{F}_q}$, induces an isomorphism

$$u_\beta \colon b_0^* \mathscr{G}_\beta \xrightarrow{\sim} b_r^* \mathscr{G}_\beta, \tag{9.4.16}$$

and similarly $(-\beta) \circ b_0 = (-\beta) \circ b_r$ induces $u_{-\beta}$.

We define $\mathfrak{c}_V \in \mathrm{Corr}_{V^\flat}(\overline{\mathbf{Q}}_{V_0}, \overline{\mathbf{Q}}_{V_r}\langle -d(h_r)\rangle)$ to be the pullback of the identity correspondence on pt. Concretely, it is the composition

$$\mathfrak{c}_V \colon b_0^* \overline{\mathbf{Q}}_{V_0} \cong b_r^* \overline{\mathbf{Q}}_{V_r} \xrightarrow{[b_r]} b_r^! \overline{\mathbf{Q}}_{V_r}\langle -d(h_r)\rangle \tag{9.4.17}$$

where $[b_r]$ is the relative fundamental class attached to the quasi-smooth map b_r. We compare it with $\mathfrak{c}_{\widetilde{V}}$ through (9.4.15). Since h_r^V is the base change of the smooth map h_r, relative purity for h_r^V gives

$$[h_r^V] \colon (h_r^V)^* \overline{\mathbf{Q}}_{V_r} \xrightarrow{\sim} (h_r^V)^! \overline{\mathbf{Q}}_{V_r}\langle -d(h_r)\rangle. \tag{9.4.18}$$

Compositional compatibility of relative derived fundamental classes (cf. [Kha19, §3.2]) identifies $[b_r]$ with the composite of $\widetilde{b}_r^*([h_r^V])$ followed by $[\widetilde{b}_r]$; note that for the smooth map h_r^V the fundamental class coincides with the purity isomorphism (9.4.18). Tensoring (9.4.17) with (9.4.16) gives a cohomological correspondence

$$\mathfrak{c}_V \otimes u_\beta \in \mathrm{Corr}_{V^\flat}\big(\mathscr{G}_\beta, \mathscr{G}_\beta\langle -d(h_r)\rangle\big). \tag{9.4.19}$$

Similarly, tensoring with $u_{-\beta}$ gives

$$\mathfrak{c}_V \otimes u_{-\beta} \in \mathrm{Corr}_{V^\flat}\big(\mathscr{G}_{-\beta}, \mathscr{G}_{-\beta}\langle -d(h_r)\rangle\big). \tag{9.4.20}$$

Let $\mathrm{FT}_{V^\flat}$ denote, in this varying-base formulation, the Fourier transform of (7.4.6). After applying the Gaussian isomorphism of Lemma 6.4.2, the normalized varying-base Fourier transform is

$$\begin{aligned} [2]_V^* \mathbb{T}_{[-d(V/S)](-d(V/S))}\, \mathrm{FT}_{V^\flat} \colon \quad & \mathrm{Corr}_{V^\flat}\big(\mathscr{G}_\beta, \mathscr{G}_\beta\langle -d(h_r)\rangle\big) \\ & \longrightarrow \mathrm{Corr}_{V^\flat}\Big((\mathscr{G}_{-\beta} \otimes \pi^* \mathfrak{G}_Q)_{[-d(V/S)](-d(V/S))}, \\ & \qquad\qquad (\mathscr{G}_{-\beta} \otimes \pi^* \mathfrak{G}_Q)_{[-d(V/S)](-d(V/S))\langle -d(h_r)\rangle}\Big). \end{aligned} \tag{9.4.21}$$

Under (9.4.15), the target group in (9.4.21) becomes the target group of (9.4.10); in the notation of (9.4.3), this is the pair of identifications

$$(h_0^V)^*\big((\mathscr{G}_{-\beta}\otimes\pi^*\mathfrak{G}_Q)[-d(V/S)](-d(V/S))\big)=\big(\mathscr{G}_{-\beta,0}\otimes\widetilde{\pi}_0^*\mathfrak{G}_{Q,0}\big)[-d(V/S)](-d(V/S)),$$
$$(h_r^V)^!\big((\mathscr{G}_{-\beta}\otimes\pi^*\mathfrak{G}_Q)[-d(V/S)](-d(V/S))\langle -d(h_r)\rangle\big)\cong\big(\mathscr{G}_{-\beta,r}\otimes\widetilde{\pi}_r^*\mathfrak{G}_{Q,r}\big)[-d(V/S)](-d(V/S)).$$

The second identification is obtained by tensoring (9.4.18) with $(h_r^V)^*(\mathscr{G}_{-\beta}\otimes\pi^*\mathfrak{G}_Q)[-d(V/S)](-d(V/S))$; the twist $\langle -d(h_r)\rangle$ cancels the purity shift and Tate twist in $(h_r^V)^!$.

Proposition 9.4.6 (Varying-base Gaussian correspondence)**.** *With respect to the normalized varying-base map* (9.4.21)*, one has*

$$[2]_V^*\big(\mathbb{T}_{[-d(V/S)](-d(V/S))}\,\mathrm{FT}_{V^\flat}(\mathfrak{c}_V\otimes u_\beta)\big)=(\mathfrak{c}_V\otimes u_{-\beta})\otimes\xi_r[-d(V/S)](-d(V/S)) \tag{9.4.22}$$
$$\in\mathrm{Corr}_{V^\flat}\Big((\mathscr{G}_{-\beta}\otimes\pi^*\mathfrak{G}_Q)[-d(V/S)](-d(V/S)),(\mathscr{G}_{-\beta}\otimes\pi^*\mathfrak{G}_Q)[-d(V/S)](-d(V/S))\langle -d(h_r)\rangle\Big), \tag{9.4.23}$$

where ξ_r is the isomorphism of Proposition 9.4.5, viewed through (9.4.15) *(note that $b_i^*\pi^*\mathfrak{G}_Q=\widetilde{b}_i^*\widetilde{\pi}_i^*\mathfrak{G}_{Q,i}$).*

Proof. By definition of the varying-base Fourier transform (7.4.6), it can be related to the fixed-base Fourier transform for different sheaves. In this case, the translation goes through (9.4.15). After applying that identification, $\mathfrak{c}_V\otimes u_\beta$ becomes $\mathfrak{c}_{\widetilde{V}}\otimes u_\beta$: the class $[b_r]$ becomes $[\widetilde{b}_r]$ by the composition and purity discussion preceding (9.4.19), and $(h_r^V)^!(\mathscr{G}_{\beta}\langle -d(h_r)\rangle)\cong\mathscr{G}_{\beta,r}$. The same argument identifies $\mathfrak{c}_V\otimes u_{-\beta}$ with $\mathfrak{c}_{\widetilde{V}}\otimes u_{-\beta}$, and the Gauss factors correspond via $\mathfrak{G}_{Q,i}=h_i^*\mathfrak{G}_Q$, matching the shifted target groups.

By the definition of the varying-base transform (7.4.6), the varying-base Fourier transform of cohomological correspondences (9.4.21) corresponds, under the identification (9.4.15), to the fixed Fourier transform (9.4.10). Therefore, after translated between fixed-base and varying-based groups of cohomological correspondences using (9.4.15), the image of (9.4.22) becomes precisely the identity of Proposition 9.4.5, with the same ξ_r. □

10. Proof of the Modularity Conjecture

In this section we complete the proof of the Modularity Conjecture of [FYZ25]. We first recall its formulation. A point of $\mathrm{Bun}_{GU^-(2m)}$ is a triple $(\mathcal{G},\mathfrak{M},h)$, where $\mathcal{G}$ is a rank $2m$ vector bundle on X', $\mathfrak{M}$ is a line bundle on X, and

$$h\colon\mathcal{G}\xrightarrow{\sim}\sigma^*\mathcal{G}^\vee\otimes\nu^*\mathfrak{M}=\sigma^*\mathcal{G}^*\otimes\nu^*(\omega_X\otimes\mathfrak{M})$$

is skew-Hermitian. Put $\mathfrak{L}:=\omega_X\otimes\mathfrak{M}$. Let $\mathrm{Bun}_{\widetilde{P}_m}$ be the stack of quadruples $(\mathcal{G},\mathfrak{M},h,\mathcal{E})$, with $(\mathcal{G},\mathfrak{M},h)\in\mathrm{Bun}_{GU^-(2m)}$ and $\mathcal{E}\subset\mathcal{G}$ a Lagrangian sub-bundle. Forgetting $\mathcal{E}$ defines a surjective map

$$\mathrm{Bun}_{\widetilde{P}_m}(\mathbf{F}_q)\twoheadrightarrow\mathrm{Bun}_{GU^-(2m)}(\mathbf{F}_q). \tag{10.0.1}$$

With the normalization of [FYZ25, §4.6], the higher theta series is the function

$$\widetilde{Z}_m^{n,r}\colon\ \mathrm{Bun}_{\widetilde{P}_m}(\mathbf{F}_q)\longrightarrow\mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{GU(n)}), \tag{10.0.2}$$

whose value at $(\mathcal{G},\mathfrak{M},h,\mathcal{E})$ is

$$\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E})=\chi(\det\mathcal{E})\,q^{n(\deg\mathcal{E}-\deg\mathfrak{L}-\deg\omega_X)/2}\sum_{a\in\mathcal{A}_{\mathcal{E},\mathfrak{L}}(\mathbf{F}_q)}\psi(\langle e_{\mathcal{G},\mathcal{E}},a\rangle)\,[\mathcal{Z}^r_{\mathcal{E},\mathfrak{L}}(a)] \tag{10.0.3}$$

with the notation explained as follows:

- $\chi:\mathrm{Pic}_{X'}(k)\to\overline{\mathbf{Q}}^\times$ is a character satisfying $\chi\circ\nu^*=\eta^n$ on $\mathrm{Pic}_X(k)$, where $\eta:\mathrm{Pic}_X(k)\to\{\pm1\}$ is the character corresponding to the double cover X'/X. We also write $\eta_{F'/F}=\eta$.
- $\psi:\mathbf{F}_q\to\overline{\mathbf{Q}}^\times$ is a nontrivial character.
- $\mathcal{A}_{\mathcal{E},\mathfrak{L}}(\mathbf{F}_q)$ is the set of Hermitian maps $a:\mathcal{E}\to\sigma^*\mathcal{E}^\vee\otimes\nu^*\mathfrak{L}$.
- $e_{\mathcal{G},\mathcal{E}}\in\mathrm{Ext}^1(\sigma^*\mathcal{E}^*\otimes\nu^*\mathfrak{L},\mathcal{E})$ is the extension class for the short exact sequence

$$0\longrightarrow\mathcal{E}\longrightarrow\mathcal{G}\longrightarrow\sigma^*\mathcal{E}^*\otimes\nu^*\mathfrak{L}\longrightarrow0, \tag{10.0.4}$$

  and $\langle e_{\mathcal{G},\mathcal{E}},a\rangle$ is the Serre duality pairing.

Conjecture 10.0.1 (The Modularity Conjecture)**.** *The function $\widetilde{Z}_m^{n,r}$ descends along* (10.0.1)*; in other words, $\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E})$ is independent of the choice of Lagrangian $\mathcal{E}\subset\mathcal{G}$.*

To keep the notation manageable, we write out the proof below only in the case of the trivial-similitude fiber $\mathfrak{L}=\mathcal{O}_X$. The same ideas apply straightforwardly to the general case. On this fiber $\mathfrak{M}=\omega_X^{-1}$, the skew-Hermitian form is an isomorphism $\mathcal{G}\cong\sigma^*\mathcal{G}^*$. We write $\mathrm{Bun}_{GU^-(2m)}^{\mathrm{triv}}$ and Bun_{P_m} for the corresponding fibers of $\mathrm{Bun}_{GU^-(2m)}$ and $\mathrm{Bun}_{\widetilde{P}_m}$, and we write $\mathrm{Sht}_{U(n)}^r=\mathrm{Sht}_{U(n),\mathcal{O}_X}^r$. We likewise abbreviate $\mathcal{A}_{\mathcal{E}}:=\mathcal{A}_{\mathcal{E},\mathcal{O}_X}$ and $\mathcal{Z}_{\mathcal{E}}^r(a):=\mathcal{Z}_{\mathcal{E},\mathcal{O}_X}^r(a)$.

10.1. **Modularity in low corank.** We first prove the Modularity Conjecture in the range $m\le n/3$, where we have established the Trace Conjecture in Theorem 5.1.3.

Theorem 10.1.1 (Modularity in low corank)**.** *Assume $m\le n/3$. Then for every $\mathcal{G}\in\mathrm{Bun}_{GU^-(2m)}^{\mathrm{triv}}(\mathbf{F}_q)$ and every pair of Lagrangian sub-bundles $\mathcal{E}_1,\mathcal{E}_2\subset\mathcal{G}$,*

$$\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E}_1)=\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E}_2)\in\mathrm{CH}_{r(n-m)}(\mathrm{Sht}_{U(n)}^r). \tag{10.1.1}$$

Proof of Theorem 10.1.1.

10.1.1. *Reductions.* Let $S=\mathrm{Bun}_{U(n)}^{\le\mu}$ be a Harder–Narasimhan truncation and let $\mathrm{Sht}_S^r\subset\mathrm{Sht}_{U(n)}^r$ be the induced open substack. We have

$$\mathrm{CH}_{r(n-m)}(\mathrm{Sht}_{U(n)}^r)=\varprojlim_{\mu}\mathrm{CH}_{r(n-m)}(\mathrm{Sht}_S^r). \tag{10.1.2}$$

Thus it is enough to prove (10.1.1) after restriction to every $\mathrm{CH}_{r(n-m)}(\mathrm{Sht}_S^r)$. The transverse reduction of [FYZ23, §2.2] says that, in a fixed fiber of p_m, the equivalence relation generated by transverse pairs of Lagrangians is the full relation. Thus we may assume that $\mathcal{E}_1,\mathcal{E}_2\subset\mathcal{G}$ are transverse. We will prove that

$$\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E}_1)\big|_S=\widetilde{Z}_m^{n,r}(\mathcal{G},\mathcal{E}_2)\big|_S. \tag{10.1.3}$$

Now the formulations and results of §9 apply, and we maintain the notation from there.

10.1.2. *Preparations.* To summarize the setup, we have short exact sequences

$$0\longrightarrow\sigma^*\mathcal{E}_2\longrightarrow\mathcal{E}_1^*\longrightarrow Q_1\longrightarrow 0,\qquad 0\longrightarrow\sigma^*\mathcal{E}_1\longrightarrow\mathcal{E}_2^*\longrightarrow Q_2\longrightarrow 0, \tag{10.1.4}$$

where $Q_2\cong\sigma^*Q_1^*$ and $Q_1^*=\mathcal{E}xt^1_{X'}(Q_1,\mathcal{O}_{X'})$. Let D_Q be the effective divisor on X, counted with lengths, whose pullback to X' is the scheme-theoretic support divisor of Q_1, equivalently of Q_2. Thus $\det(Q_1)\cong\nu^*\mathcal{O}_X(D_Q)$. In what follows $U,V,W,U^\perp,\widehat{V},W^\perp$, their Hecke correspondences, and the maps $f,\pi,f^\perp,\pi^\perp$ are as in (9.1.2), (9.1.3), (9.1.5), and (9.1.8). Abbreviate $d:=d(V/S)$.

The skew-Hermitian form on $\mathcal{G}$ induces two opposite Hermitian forms on the common torsion quotient, as in [FYZ23, §2.3]:

$$h_{12},h_{21}\colon Q_1\xrightarrow{\sim}\sigma^*Q_1^*,\qquad h_{21}=-h_{12}. \tag{10.1.5}$$

(In particular $\mathrm{div}(Q_1)$ is σ-invariant, which justifies the existence of the divisor D_Q on X introduced above.) Let $\beta=\beta_{12}\colon V\to\mathbf{A}_S^1$ be the quadratic function attached to h_{12}, and let $-\beta=\beta_{21}$ be the quadratic function attached to h_{21}. We write $\mathscr{G}_\beta=\beta^*\,\mathrm{AS}_{\psi,S}$ and $\mathscr{G}_{-\beta}=(-\beta)^*\,\mathrm{AS}_{\psi,S}$.

Lemma 10.1.2. *The relative Gauss cohomology of* (9.4.8)*, equivalently*

$$\mathfrak{G}_Q:=\pi_{V!}\mathscr{G}_{\beta[d]}\in\mathcal{D}_{\mathrm{mot}}(S;\overline{\mathbf{Q}}) \tag{10.1.6}$$

has constant Frobenius trace function on $S(\mathbf{F}_q)$, which for every $s\in S(\mathbf{F}_q)$ is equal to

$$\mathrm{Tr}(\mathrm{Frob},\mathfrak{G}_Q)(s)=(-1)^d q^{d/2}\eta_{F'/F}(D_Q)^n. \tag{10.1.7}$$

More generally, for every finite extension $\mathbf{F}_{q^k}$ and every $s\in S(\mathbf{F}_{q^k})$, one has

$$\mathrm{Tr}(\mathrm{Frob}_q^k,\mathfrak{G}_Q)(s)=(-1)^d q^{kd/2}\eta_{F'_k/F_k}(D_{Q,k})^n=\big((-1)^d q^{d/2}\eta_{F'/F}(D_Q)^n\big)^k$$

where the subscript k indicates base change to $\mathbf{F}_{q^k}$.

Proof. The statement is pointwise, so fix $s \in S(\mathbf{F}_{q^k})$; to lighten notation take $k = 1$. The formation of $\pi_{V!}\mathscr{G}_\beta$ commutes with the base change along $s \to S$, and since Q_1 is torsion, the fiber $V_s = \mathrm{Hom}(\mathcal{F}_s^*, Q_1)$ is a finite-dimensional $\mathbf{F}_q$-quadratic space with quadratic form $\mathfrak{q}_{12}$. By the Grothendieck–Lefschetz trace formula,

$$\mathrm{Tr}(\mathrm{Frob}, \pi_{V!}\mathscr{G}_\beta)(s) = \sum_{v \in V_s(\mathbf{F}_q)} \psi(\beta_s(v)), \tag{10.1.8}$$

the Gauss sum of $(V_s, \mathfrak{q}_{12})$. The computation of [FYZ23, Lemma 2.3.8] gives

$$\sum_{v \in V_s(\mathbf{F}_q)} \psi(\beta_s(v)) = q^{d/2} \eta_{F'/F}(D_Q)^n,$$

independently of the Hermitian bundle $\mathcal{F}_s$. For general k with the additive character $\psi_k := \psi \circ \mathrm{Tr}_{\mathbf{F}_{q^k}/\mathbf{F}_q}$, the same placewise calculation over X_k gives

$$\sum_{v \in V_s(\mathbf{F}_{q^k})} \psi_k(\beta_s(v)) = q^{kd/2} \eta_{F'_k/F_k}(D_{Q,k})^n. \tag{10.1.9}$$

Moreover, $\eta_{F'_k/F_k}(D_{Q,k}) = \eta_{F'/F}(D_Q)^k$. Finally, the shift $[d]$ in (10.1.6) contributes the sign $(-1)^d$. Here $d = n \deg_{X'}(\nu^* D_Q) = 2n \deg D_Q$ is even, so this sign is 1. Together with (10.1.9), this proves the result. □

Let $\pi^\flat \colon V^\flat \to \mathrm{Hk}_S^1$ be the projection. This is a vector bundle in the classical sense by the proof of Lemma 9.3.4. Therefore $(\pi^\flat)^*$ is fully faithful (by homotopy invariance), and via the identifications

$$(\pi^\flat)^* h_i^* \mathfrak{G}_Q = b_i^* \pi^* \mathfrak{G}_Q = \widetilde{b}_i^* \widetilde{\pi}_i^* \mathfrak{G}_{Q,i},$$

the isomorphism ξ_1 of Proposition 9.4.5 descends uniquely to an isomorphism

$$\bar{\xi}_1 \colon h_0^* \mathfrak{G}_Q \xrightarrow{\sim} h_1^* \mathfrak{G}_Q.$$

For $0 \le i \le r$, let $h_i \colon \mathrm{Hk}_S^r \to S_i$ record the i-th intermediate bundle, where each S_i is a copy of S. For $1 \le i \le r$, let $\rho_i \colon \mathrm{Hk}_S^r \to \mathrm{Hk}_S^1$ record the i-th modification. Then $\rho_i^* \bar{\xi}_1$ is an isomorphism from $h_{i-1}^* \mathfrak{G}_Q$ to $h_i^* \mathfrak{G}_Q$. For general r, let

$$\bar{\xi}_r := \rho_r^* \bar{\xi}_1 \circ \cdots \circ \rho_1^* \bar{\xi}_1 \colon h_0^* \mathfrak{G}_Q \xrightarrow{\sim} h_r^* \mathfrak{G}_Q \tag{10.1.10}$$

and set $\bar{\xi}_0 = \mathrm{Id}$. The one-leg factorization is displayed schematically as

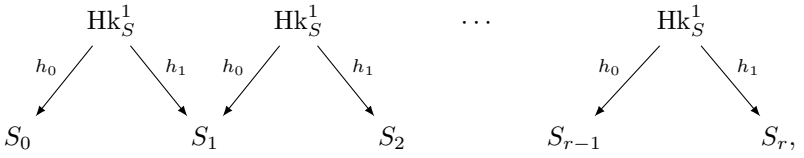

(10.1.11)

where the i-th copy of Hk_S^1 is the correspondence from S_{i-1} to S_i, pulled back to Hk_S^r along ρ_i. Thus the pullback of $\bar{\xi}_r$ along $\mathrm{Hk}_V^\flat \to \mathrm{Hk}_S^r$ is the isomorphism ξ_r of Proposition 9.4.5.

Corollary 10.1.3. *The trace of $(\mathfrak{G}_Q, \bar{\xi}_r)$ twisted by Frobenius is the constant class*

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{G}_Q, \bar{\xi}_r) = (-1)^d q^{d/2} \eta_{F'/F}(D_Q)^n \cdot 1_{\mathrm{Sht}_S^r}$$

in $\mathrm{H}^0(\mathrm{Sht}_S^r, \overline{\mathbf{Q}})$*, where* $1_{\mathrm{Sht}_S^r}$ *denotes the unit section.*

Proof. By Lemma 6.4.3, applied as in the paragraph following (9.4.8), the object $\mathfrak{G}_Q$ is tensor-invertible. Hence the trace is given by an element of $\mathrm{H}^0(\mathrm{Sht}_S^r, \overline{\mathbf{Q}})$, i.e., a locally constant function. It is enough to compute this locally constant function on a nonempty open substack of each connected component.

Let $\mathrm{Sht}_S^{r,\circ} \subset \mathrm{Sht}_S^r$ be the locus where all legs avoid $\nu^{-1}(D_Q)$. If $r = 0$, this is all of Sht_S^0. If $r > 0$, let C be a connected component of Sht_S^r. By [FYZ24, §6], the leg map $\mathrm{Sht}_{U(n)}^r \to (X')^r$ is smooth, hence open; therefore the image of C is a nonempty open subset of a connected component of $(X')^r$. Since $(X' \setminus \nu^{-1}(D_Q))^r$ is dense in every connected component of $(X')^r$, the component C meets $\mathrm{Sht}_S^{r,\circ}$.

Over $\mathrm{Sht}_S^{r,\circ}$, all bundle modifications in the shtuka restrict to isomorphisms on a neighborhood of $\nu^{-1}(D_Q)$. Thus the restriction of $\bar{\xi}_r$ is the r-fold composition of the tautological isomorphism (9.4.12). Fix

a geometric point y of $\mathrm{Sht}_S^{r,\circ}$. Let τ_y be the isometry from $\mathcal{F}_0|_{\nu^{-1}(D_Q)}$ to $\mathcal{F}_r|_{\nu^{-1}(D_Q)}$ obtained by composing the Hecke modifications, and let $\varphi_y\colon \mathcal{F}_r \xrightarrow{\sim} \mathrm{Frob}_q^*\,\mathcal{F}_0$ be the identification in the definition of shtuka. The composite $(\varphi_y|_{\nu^{-1}(D_Q)})\circ\tau_y$ is the Frobenius-semilinear isometry that supplies descent on the finite quadratic space $V_y = \mathrm{Hom}(\mathcal{F}_0^*, Q_1)$. By the definition of ξ° in (9.4.12) and of $\bar{\xi}_r$, the Frobenius-twisted endomorphism of the Gauss cohomology at y is precisely the transport induced by this semilinear isometry. The twisted Grothendieck–Lefschetz formula identifies its trace with the Gauss sum on the fixed subspace, which is the descended finite Hermitian quadratic space over $\mathbf{F}_q$. The computation of [FYZ23, Lemma 2.3.8], applied to that descended space, gives $q^{d/2}\eta_{F'/F}(D_Q)^n$, and the shift $[d]$ contributes $(-1)^d$. □

10.1.3. *Sheaf-cycle correspondence.* Recall the cohomological correspondence $\mathfrak{c}_V$ from (9.4.17). On the Fourier-dual row, write

$$\mathfrak{c}_{\widehat{V}} \in \mathrm{Corr}_{\mathrm{Hk}^\flat_{\widehat{V}}}\left(\overline{\mathbf{Q}}_{\widehat{V}_0}, \overline{\mathbf{Q}}_{\widehat{V}_r}\langle -d(h_r)\rangle\right)$$

for the analogous pullback of the identity correspondence, defined using the quasi-smooth map $\beta_r : \mathrm{Hk}^\flat_{\widehat{V}} \to \widehat{V}_r$. Under the self-duality identifications used in Proposition 9.4.6, $\mathfrak{c}_V$ identifies with $\mathfrak{c}_{\widehat{V}}$. The pullback cohomological correspondences $\mathfrak{c}_U = f^*\,\mathfrak{c}_V$ and $\mathfrak{c}_{U^\perp} = (f^\perp)^*\,\mathfrak{c}_{\widehat{V}}$ exist by Lemma 9.3.1. We also use the Gaussian pullback isomorphisms u_β and $u_{-\beta}$ defined in and immediately after (9.4.16): u_β is induced by the equality $\beta\circ b_0 = \beta\circ b_r$, and $u_{-\beta}$ is induced by the equality $(-\beta)\circ b_0 = (-\beta)\circ b_r$.

Lemma 10.1.4. *In* $\mathrm{CH}_{r(n-m)}(\mathrm{Sht}_S^r)$*, one has*

$$\begin{aligned}&\mathrm{Sht}(\pi)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U\otimes f^*u_\beta)\\ &\quad= q^{-d+d(U/S)+d/2}\eta_{F'/F}(D_Q)^n\ \mathrm{Sht}(\pi^\perp)_!\,\mathrm{Tr}^{\mathrm{Sht}}\left(\mathfrak{c}_{U^\perp}\otimes(f^\perp)^*u_{-\beta}\right).\end{aligned} \tag{10.1.12}$$

Proof. Under the self-duality identifications used in Proposition 9.4.6, the right hand side of (9.4.22) is the cohomology correspondence $\mathfrak{c}_{\widehat{V}}\otimes u_{-\beta}$ on the row with $\mathrm{Hk}^\flat_{\widehat{V}}$ in (9.1.8). By definition of $\mathfrak{c}_U$ and $\mathfrak{c}_{U^\perp}$, and by compatibility of pullback with tensoring (the pullable analogue of Lemma 7.2.2, proved by the same concatenation-of-base-change argument), we have

$$f^*(\mathfrak{c}_V\otimes u_\beta) = \mathfrak{c}_U\otimes f^*u_\beta, \qquad (f^\perp)^*(\mathfrak{c}_{\widehat{V}}\otimes u_{-\beta}) = \mathfrak{c}_{U^\perp}\otimes(f^\perp)^*u_{-\beta}. \tag{10.1.13}$$

We next compare the pushed correspondences on the base correspondence $S \xleftarrow{h_0} \mathrm{Hk}_S^r \xrightarrow{h_r} S$. Apply the Base Change Theorem for cohomological correspondences, [FK24, Theorem 4.4.2], to the square of maps of correspondences induced by the first square in (9.1.5). The corresponding arrows for its top, right, left, and bottom maps are, respectively, f, g, π, and $z_{\mathrm{Hk}^\flat_W}$. Parts (1)–(4) of Lemma 9.3.1 say respectively that f is pullable, π is pushable, g is pushable, and $z_{\mathrm{Hk}^\flat_W}$ is pullable. The corresponding commutative squares are the base changes of the derived Cartesian square (9.1.3), hence are themselves Cartesian. Thus the hypotheses of the Base Change Theorem hold, and it gives the canonical identification

$$\pi_! f^*(\mathfrak{c}_V\otimes u_\beta) \cong z^*_{\mathrm{Hk}^\flat_W}\, g_!(\mathfrak{c}_V\otimes u_\beta), \tag{10.1.14}$$

where $z_{\mathrm{Hk}^\flat_W}\colon \mathrm{Hk}_S^r \to \mathrm{Hk}^\flat_W$ denotes the zero section in (9.1.5).

We now apply the motivic derived Fourier transform to (10.1.14). On the S-row the Fourier transform is the identity, since S is the zero derived vector bundle over itself. The required hypotheses are precisely the global-presentation hypotheses asserted for both diagrams after (9.1.8) in §9.1. First, the left-pushable part of Proposition 7.4.1, applied to the map with vertical arrow g, gives

$$\mathbb{T}_{[d(g_0)]}\,\mathrm{FT}_{\mathrm{Hk}^\flat_W}(g_!(\mathfrak{c}_V\otimes u_\beta)) \cong (f^\perp)^*\,\mathrm{FT}_{\mathrm{Hk}^\flat_V}(\mathfrak{c}_V\otimes u_\beta). \tag{10.1.15}$$

Here the Fourier-dual map has vertical arrow $f^\perp = \widehat{g^\sharp}$, denoted $\widehat{g}$ in the second diagram of (9.1.5). Second, the right-pullable part of Proposition 7.4.1, applied to the zero-section map $z_{\mathrm{Hk}^\flat_W}\colon \mathrm{Hk}_S^r \to \mathrm{Hk}^\flat_W$, gives

$$\mathrm{FT}_{\mathrm{Hk}^r_S}\left(z^*_{\mathrm{Hk}^\flat_W}\, g_!(\mathfrak{c}_V\otimes u_\beta)\right) \cong \pi_!^\perp\mathbb{T}_{[-d(z_{W,0})](-d(z_{W,0}))}\,\mathrm{FT}_{\mathrm{Hk}^\flat_W}(g_!(\mathfrak{c}_V\otimes u_\beta)). \tag{10.1.16}$$

The Fourier-dual of the zero section $z_W\colon S \to W$ is the projection $U^\perp = \widehat{W} \to S$, hence the pushforward in (10.1.16) is $\pi_!^\perp$. Here $z_{W,0}\colon S \to W_0$ denotes the zero section over the first copy of S. Combining (10.1.14), (10.1.15), and (10.1.16), and then using Proposition 9.4.6, gives the following identity of cohomological correspondences on Hk_S^r:

$$\pi_! f^*(\mathfrak{c}_V \otimes u_\beta) \cong \big(\pi_!^\perp (f^\perp)^*(\mathfrak{c}_{\widehat{V}} \otimes u_{-\beta})\big) \otimes (\mathfrak{G}_Q, \bar{\xi}_r)_{[d-2d(U/S)](d-d(U/S))}. \tag{10.1.17}$$

Here $\bar{\xi}_r$ is as in (10.1.10), and $\otimes(\mathfrak{G}_Q, \bar{\xi}_r)$ is the tensor operation of Definition 7.1.1 with this datum. The shift and Tate twist in (10.1.17) are obtained by combining the following contributions:

operation	shift and twist
normalized FT (9.4.21)	$[d](d)$
Fourier commutation with $g_!$ (10.1.15)	$\mathbb{T}_{[-d(g_0)]}$
zero-section functoriality (10.1.16)	$\mathbb{T}_{[-d(z_{W,0})](-d(z_{W,0}))}$
common normalization on the Gauss factor	$[-d](-d)$
total after rank additivity	$[d-2d(U/S)](d-d(U/S))$.

The extra row records that the normalized Gaussian Fourier identity carries the factor $\mathfrak{G}_{Q[-d](-d)}$, while (10.1.17) is written as a twist by $\mathfrak{G}_Q$ together with the displayed common shift and Tate twist. Finally, the scalar multiplication $[2]_V^*$ appearing in (9.4.21) commutes with all linear maps in (9.1.3) and (9.1.5), and becomes the identity on S.

Taking Frobenius traces in (10.1.17), and using Corollary 10.1.3 for the Gauss twist, gives

$$\mathrm{Tr}^{\mathrm{Sht}}(\pi_! f^*(\mathfrak{c}_V \otimes u_\beta)) = (-1)^d q^{-d+d(U/S)}\,\mathrm{Tr}(\mathrm{Frob}, \mathfrak{G}_Q)\ \mathrm{Tr}^{\mathrm{Sht}}\big(\pi_!^\perp (f^\perp)^*(\mathfrak{c}_{\widehat{V}} \otimes u_{-\beta})\big)\,. \tag{10.1.18}$$

It remains to commute the pushforwards past the formation of the trace, using Proposition 8.1.1. We check that the maps of correspondences in question satisfy its hypotheses. Indeed, $\mathcal{F}_\bullet^\flat$ has tor-amplitude $[0,1]$, so its dual has amplitude $[-1,0]$; applying relative RHom into the fixed vector bundles and then relative cohomology on the curve produces coconnective perfect complexes. The same calculation applies to the Fourier-dual row. The comparison maps $\widetilde{a}_i$ and $\widetilde{a}_i^\perp$ are closed embeddings of derived vector bundles by [FYZ23, Lemma 9.1.3(1)]; equivalently, the fibers of the defining maps of perfect complexes have tor-amplitude $[1,\infty)$ by [FYZ23, Lemma 6.1.5]. Finally, AS_ψ is a direct summand of the pushforward of the unit along the finite étale Artin–Schreier cover (6.1.1); hence it is geometric, and so are its Gaussian pullbacks and their tensor products with the unit coefficient objects. Thus Proposition 8.1.1 applies and gives

$$\begin{aligned} \mathrm{Tr}^{\mathrm{Sht}}(\pi_! f^*(\mathfrak{c}_V \otimes u_\beta)) &= \mathrm{Sht}(\pi)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U \otimes f^* u_\beta)\,, \\ \mathrm{Tr}^{\mathrm{Sht}}\big(\pi_!^\perp (f^\perp)^*(\mathfrak{c}_{\widehat{V}} \otimes u_{-\beta})\big) &= \mathrm{Sht}(\pi^\perp)_!\,\mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}_{U^\perp} \otimes (f^\perp)^* u_{-\beta}\big)\,. \end{aligned} \tag{10.1.19}$$

Thus (10.1.18) and (10.1.19) imply

$$\begin{aligned} &\mathrm{Sht}(\pi)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U \otimes f^* u_\beta) \\ &\quad = (-1)^d q^{-d+d(U/S)}\,\mathrm{Tr}(\mathrm{Frob}, \mathfrak{G}_Q)\ \mathrm{Sht}(\pi^\perp)_!\,\mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}_{U^\perp} \otimes (f^\perp)^* u_{-\beta}\big)\,. \end{aligned} \tag{10.1.20}$$

Finally, substituting the relative Gauss-sum calculation (10.1.7), namely

$$\mathrm{Tr}(\mathrm{Frob}, \mathfrak{G}_Q) = (-1)^d q^{d/2} \eta_{F'/F}(D_Q)^n,$$

into (10.1.20) gives (10.1.12), with scalar

$$(-1)^d q^{-d+d(U/S)} \cdot (-1)^d q^{d/2} \eta_{F'/F}(D_Q)^n = q^{-d+d(U/S)+d/2} \eta_{F'/F}(D_Q)^n.$$

□

We now express higher theta functions as a trace.

Lemma 10.1.5. *The two pushed trace classes satisfy*

$$\begin{aligned} \widetilde{Z}_m^{n,r}(\mathcal{G}, \mathcal{E}_1)\big|_S &= \chi(\det \mathcal{E}_1) q^{n(\deg \mathcal{E}_1 - \deg \omega_X)/2}\,\mathrm{Sht}(\pi)_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U \otimes f^* u_\beta)\,, \\ \widetilde{Z}_m^{n,r}(\mathcal{G}, \mathcal{E}_2)\big|_S &= \chi(\det \mathcal{E}_2) q^{n(\deg \mathcal{E}_2 - \deg \omega_X)/2}\,\mathrm{Sht}(\pi^\perp)_!\,\mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}_{U^\perp} \otimes (f^\perp)^* u_{-\beta}\big)\,. \end{aligned} \tag{10.1.21}$$

Proof. Since Artin–Schreier sheaves are tensor-invertible, Example 5.4.1 and Lemma 5.4.2 yield

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U \otimes f^* u_\beta) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U) \cap \mathrm{Tr}^{\mathrm{Sht}}(f^* u_\beta),$$
$$\mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}_{U^\perp} \otimes (f^\perp)^* u_{-\beta}\big) = \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{U^\perp}) \cap \mathrm{Tr}^{\mathrm{Sht}}((f^\perp)^* u_{-\beta}).$$

We identify the two pushed trace classes with the two unnormalized theta sums. Let $\mathcal{M}_{\mathcal{E}_i,S} \to \mathcal{A}_{\mathcal{E}_i}$ be the Hitchin stack over S for the Lagrangian $\mathcal{E}_i$. The moduli descriptions of the two U-rows give canonical identifications

$$\mathrm{Sht}(\mathrm{Hk}^\flat_U) \cong \mathcal{Z}^r_{\mathcal{E}_1}|_S, \qquad \mathrm{Sht}(\mathrm{Hk}^\flat_{U^\perp}) \cong \mathcal{Z}^r_{\mathcal{E}_2}|_S.$$

Under these identifications, $\mathfrak{c}_U$ and $\mathfrak{c}_{U^\perp}$ are, respectively, the restrictions over S of the canonical Hitchin cohomological correspondences $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}_1}}$ and $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}_2}}$. The fixed-point Hitchin base is (using [FYZ25, Lemma 6.14])

$$\mathcal{A}_{\mathcal{E}_i} \times_{(\mathrm{Id},\mathrm{Frob}),\, \mathcal{A}_{\mathcal{E}_i} \times \mathcal{A}_{\mathcal{E}_i},\, \Delta} \mathcal{A}_{\mathcal{E}_i} \cong \coprod_{a \in \mathcal{A}_{\mathcal{E}_i}(\mathbf{F}_q)} \mathrm{Spec}\ \mathbf{F}_q. \tag{10.1.22}$$

The low-corank Trace Conjecture, Theorem 5.2.1, applies because $m \le n/3$. Restricted to the open-and-closed summand indexed by a in (10.1.22), it identifies the unweighted trace class with the derived fundamental class of $\mathcal{Z}^r_{\mathcal{E}_i}(a)|_S$. (Here we use that $\mathrm{Tr}^{\mathrm{Sht}}$ commutes with restriction to open substacks, by [FK24, Proposition 6.2.2].) The extension-class calculation of [FYZ23, §2.3 and Lemma 10.2.7] identifies the trace of $f^* u_\beta$ on the a-summand for $\mathcal{E}_1$ with $\psi(\langle e_{\mathcal{G},\mathcal{E}_1}, a\rangle)$, and identifies the trace of $(f^\perp)^* u_{-\beta}$ on the a-summand for $\mathcal{E}_2$ with $\psi(\langle e_{\mathcal{G},\mathcal{E}_2}, a\rangle)$. Thus

$$\mathrm{Sht}(\pi)_!\, \mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}_U \otimes f^* u_\beta) = \sum_{a \in \mathcal{A}_{\mathcal{E}_1}(\mathbf{F}_q)} \psi(\langle e_{\mathcal{G},\mathcal{E}_1}, a\rangle)[\mathcal{Z}^r_{\mathcal{E}_1}(a)]\big|_S, \tag{10.1.23}$$

and

$$\mathrm{Sht}(\pi^\perp)_!\, \mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}_{U^\perp} \otimes (f^\perp)^* u_{-\beta}\big) = \sum_{a \in \mathcal{A}_{\mathcal{E}_2}(\mathbf{F}_q)} \psi(\langle e_{\mathcal{G},\mathcal{E}_2}, a\rangle)[\mathcal{Z}^r_{\mathcal{E}_2}(a)]\big|_S. \tag{10.1.24}$$

Comparing (10.1.23) and (10.1.24) with (the trivial similitude specialization of) (10.0.3) gives (10.1.21). □

Combining Lemma 10.1.5 with Lemma 10.1.4 gives

$$\begin{aligned}\widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_1)\big|_S = q^{-d+d(U/S)+d/2+\frac{n(\deg \mathcal{E}_1 - \deg \omega_X)}{2} - \frac{n(\deg \mathcal{E}_2 - \deg \omega_X)}{2}} \\ \cdot \eta_{F'/F}(D_Q)^n \chi(\det \mathcal{E}_1)\chi(\det \mathcal{E}_2)^{-1}\, \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_2)\big|_S.\end{aligned} \tag{10.1.25}$$

10.1.4. *Completion of the proof of Theorem 10.1.1.* It remains to show that the scalar factor in (10.1.25), namely

$$q^{-d+d(U/S)+d/2+\frac{n(\deg \mathcal{E}_1 - \deg \omega_X)}{2} - \frac{n(\deg \mathcal{E}_2 - \deg \omega_X)}{2}} \eta_{F'/F}(D_Q)^n \chi(\det \mathcal{E}_1)\chi(\det \mathcal{E}_2)^{-1}$$

is equal to 1.

Applying $\underline{\mathrm{RHom}(\mathcal{F}^*_{\mathrm{univ}}, -)}$ to the first exact sequence in (10.1.4) gives an exact triangle

$$\underline{\mathrm{RHom}(\mathcal{F}^*_{\mathrm{univ}}, \sigma^*\mathcal{E}_2)} \longrightarrow \underline{\mathrm{RHom}(\mathcal{F}^*_{\mathrm{univ}}, \mathcal{E}_1^*)} \longrightarrow \underline{\mathrm{RHom}(\mathcal{F}^*_{\mathrm{univ}}, Q_1)}.$$

The total spaces of the middle and right terms are U and V, respectively, and $d = d(V/S)$. Hence additivity of virtual ranks gives, on each connected component of S,

$$-d + d(U/S) + \frac{d}{2} = \chi_{\mathrm{Eul}}(X', \mathcal{F}_{\mathrm{univ}} \otimes \sigma^*\mathcal{E}_2) + \frac{1}{2}\chi_{\mathrm{Eul}}(X', \mathcal{F}_{\mathrm{univ}} \otimes Q_1). \tag{10.1.26}$$

For a geometric point $\mathcal{F}$ of S, the Hermitian condition $\mathcal{F} \cong \sigma^*\mathcal{F}^\vee$, with $\mathcal{F}^\vee = \mathcal{H}om(\mathcal{F}, \omega_{X'})$, implies

$$\deg \mathcal{F} = \frac{n}{2} \deg \omega_{X'} = n \deg \omega_X.$$

Riemann–Roch on X' therefore gives

$$\chi_{\mathrm{Eul}}(X', \mathcal{F} \otimes \sigma^*\mathcal{E}_2) = n \deg \mathcal{E}_2.$$

Since Q_1 is torsion,

$$\chi_{\mathrm{Eul}}(X', \mathcal{F} \otimes Q_1) = n \deg Q_1.$$

The first exact sequence in (10.1.4) gives $\deg Q_1 = -\deg \mathcal{E}_1 - \deg \mathcal{E}_2$. Substituting these two Euler characteristic computations into (10.1.26) gives

$$-d + d(U/S) + \frac{d}{2} = \frac{n(\deg \mathcal{E}_2 - \deg \mathcal{E}_1)}{2}. \tag{10.1.27}$$

Consequently the total exponent of q in (10.1.25) is

$$\frac{n(\deg \mathcal{E}_2 - \deg \mathcal{E}_1)}{2} + \frac{n(\deg \mathcal{E}_1 - \deg \omega_X)}{2} - \frac{n(\deg \mathcal{E}_2 - \deg \omega_X)}{2} = 0.$$

It remains to identify the character factor. Taking determinants in the first exact sequence of (10.1.4) gives

$$\det(\mathcal{E}_1^*) \cong \det(\sigma^* \mathcal{E}_2) \otimes \det(Q_1).$$

By the definition of D_Q, the determinant of the torsion sheaf Q_1 is $\nu^* \mathcal{O}_X(D_Q)$. Equivalently,

$$\det \mathcal{E}_1 \otimes \nu^* \mathcal{O}_X(D_Q) \cong \sigma^*(\det \mathcal{E}_2^*). \tag{10.1.28}$$

The character χ restricts to $\eta_{F'/F}^n$ on $\nu^* \operatorname{Pic}_X(\mathbf{F}_q)$. Moreover $\chi(\sigma^* L) = \chi(L)^{-1}$ for every $L \in \operatorname{Pic}_{X'}(\mathbf{F}_q)$, because $\sigma^* L \otimes L = \nu^* \operatorname{Nm}(L)$ and $\eta_{F'/F}$ is trivial on norms. Applying χ to (10.1.28) gives

$$\chi(\det \mathcal{E}_1) \eta_{F'/F}(D_Q)^n = \chi(\det \mathcal{E}_2). \tag{10.1.29}$$

Equation (10.1.27) makes the q-power in (10.1.25) equal to 1, and (10.1.29) does the same for the character and Gauss factor. Hence the scalar is 1, proving Theorem 10.1.1. □

10.2. **The embedding trick.** We now pass from Theorem 10.1.1 to the full Modularity Conjecture using a version of the "embedding trick" for theta series. A similar idea has appeared before, for instance, in the work of Yuan–Zhang–Zhang [YZZ09], Kudla [Kud21], and Howard–Madapusi [HM25] on modularity of theta series on orthogonal Shimura varieties. This will involve considering Hermitian shtukas of varying ranks, so we introduce some temporary notation that makes the rank explicit. We write

$$\widetilde{Z}_m^{N,r}(\mathcal{G}, \mathcal{E}) \in \mathrm{CH}_{r(N-m)}(\mathrm{Sht}^r_{U(N)})$$

for the higher theta series of Hermitian rank N and corank m. For each N, the theta series is normalized using an admissible character

$$\chi_N \colon \operatorname{Pic}_{X'}(k) \longrightarrow \overline{\mathbf{Q}}^\times$$

whose restriction along $\nu^* \colon \operatorname{Pic}_X(k) \to \operatorname{Pic}_{X'}(k)$ is $\eta_{F'/F}^N$. Note that this condition implies $\chi_N(\sigma^* L) = \chi_N(L)^{-1}$, since $\sigma^* L \otimes L = \nu^* \operatorname{Nm}(L)$ and $\eta_{F'/F}^N(\operatorname{Nm}(L)) = 1$.

Lemma 10.2.1. *The direct-sum morphism*

$$\mathrm{add}_{N_1,N_2} \colon \mathrm{Sht}^0_{U(N_1)} \times \mathrm{Sht}^r_{U(N_2)} \longrightarrow \mathrm{Sht}^r_{U(N_1+N_2)} \tag{10.2.1}$$

is LCI, so that the pullback $\mathrm{add}^*_{N_1,N_2}$ *exists on Chow groups, via the Gysin pullback* (3.5.2)*. Let* $N_1, N_2 \geq m$*, and suppose*

$$\chi_{N_1+N_2} = \chi_{N_1} \chi_{N_2}. \tag{10.2.2}$$

Then for every $(\mathcal{G}, \mathcal{E}) \in \mathrm{Bun}_{P_m}(k)$ *one has*

$$\mathrm{add}^*_{N_1,N_2}\left(\widetilde{Z}_m^{N_1+N_2,r}(\mathcal{G}, \mathcal{E})\right) = \widetilde{Z}_m^{N_1,0}(\mathcal{G}, \mathcal{E}) \boxtimes \widetilde{Z}_m^{N_2,r}(\mathcal{G}, \mathcal{E}). \tag{10.2.3}$$

Proof. The source and target of add_{N_1,N_2} are smooth: $\mathrm{Sht}^r_{U(N)}$ is smooth by [FYZ24, §6], while $\mathrm{Sht}^0_{U(N_1)}$ is the discrete groupoid $\mathrm{Bun}_{U(N_1)}(\mathbf{F}_q)$. Hence add_{N_1,N_2} is LCI. By [FYZ24, Lemma 6.9(2)], $\mathrm{Sht}^r_{U(N)}$ has pure dimension rN, including the r leg parameters, whereas $\mathrm{Sht}^0_{U(N_1)}$ has dimension 0. Thus

$$d(\mathrm{add}_{N_1,N_2}) = rN_2 - r(N_1 + N_2) = -rN_1.$$

Consequently its Gysin pullback carries $\mathrm{CH}_{r(N_1+N_2-m)}$ to $\mathrm{CH}_{r(N_1+N_2-m)-rN_1} = \mathrm{CH}_{r(N_2-m)}$, which is precisely the degree of the product class on the right side of (10.2.3). For target rank N, let $\mathcal{Z}_{\mathcal{E}}^{N,r}(a) \subset$

$\mathrm{Sht}^r_{U(N)}$ denote the special cycle of Hermitian parameter $a \in \mathcal{A}_{\mathcal{E}}(k)$. Functoriality in the unitary variable [FYZ25, Proposition 7.5] gives, for every a,

$$\mathrm{add}^*_{N_1,N_2}[\mathcal{Z}^{N_1+N_2,r}_{\mathcal{E}}(a)] = \sum_{a_1+a_2=a} [\mathcal{Z}^{N_1,0}_{\mathcal{E}}(a_1)] \boxtimes [\mathcal{Z}^{N_2,r}_{\mathcal{E}}(a_2)] \tag{10.2.4}$$

in $\mathrm{CH}_*(\mathrm{Sht}^0_{U(N_1)} \times \mathrm{Sht}^r_{U(N_2)})$. Additivity of the Artin–Schreier character and (10.2.4) now give the unnormalized factorization

$$\begin{aligned}
&\mathrm{add}^*_{N_1,N_2}\left(\sum_a \psi(\langle e_{\mathcal{G},\mathcal{E}}, a\rangle)[\mathcal{Z}^{N_1+N_2,r}_{\mathcal{E}}(a)]\right) \\
&= \sum_{a_1,a_2} \psi(\langle e_{\mathcal{G},\mathcal{E}}, a_1+a_2\rangle)[\mathcal{Z}^{N_1,0}_{\mathcal{E}}(a_1)] \boxtimes [\mathcal{Z}^{N_2,r}_{\mathcal{E}}(a_2)] \\
&= \left(\sum_{a_1} \psi(\langle e_{\mathcal{G},\mathcal{E}}, a_1\rangle)[\mathcal{Z}^{N_1,0}_{\mathcal{E}}(a_1)]\right) \boxtimes \left(\sum_{a_2} \psi(\langle e_{\mathcal{G},\mathcal{E}}, a_2\rangle)[\mathcal{Z}^{N_2,r}_{\mathcal{E}}(a_2)]\right).
\end{aligned}$$

The character identity (10.2.2) likewise gives

$$\chi_{N_1+N_2}(\det \mathcal{E})\, q^{(N_1+N_2)(\deg \mathcal{E} - \deg \omega_X)/2} = \prod_{i=1}^{2} \chi_{N_i}(\det \mathcal{E})\, q^{N_i(\deg \mathcal{E} - \deg \omega_X)/2}.$$

Together, these two factorizations prove (10.2.3). □

Throughout the rest of this subsection, we use the identification of $\mathrm{CH}_0(\mathrm{Sht}^0_{U(N)})$ with the space of $\overline{\mathbf{Q}}$-valued functions on the groupoid $\mathrm{Bun}_{U(N)}(\mathbf{F}_q)$, as in [FYZ25, §4] and [FYZ23, §2.3]; under this identification, $[\mathcal{Z}^0_{\mathcal{E}}(a)](\mathcal{F}) = \#\{t\colon \mathcal{E} \to \mathcal{F} \mid a(t) = a\}$, with no automorphism weighting.

Lemma 10.2.2. *Fix $m \le n$ and $(\mathcal{G}, \mathcal{E}) \in \mathrm{Bun}_{P_m}(k)$. There is a point $\mathcal{F}_0 \in \mathrm{Sht}^0_{U(2n)}(k)$ such that*

$$\widetilde{Z}^{2n,0}_m(\mathcal{G}, \mathcal{E})(\mathcal{F}_0) \neq 0.$$

Proof. Choose a rank n vector bundle $\mathcal{H}$ on X' defined over $\mathbf{F}_q$, for instance $\mathcal{O}^{\oplus n}_{X'}$. For

$$\mathcal{F}_0 = \mathcal{H} \oplus \sigma^*\mathcal{H}^\vee$$

we have $\sigma^*\mathcal{F}_0^\vee \simeq \sigma^*\mathcal{H}^\vee \oplus \mathcal{H}$, and the isomorphism $\mathcal{F}_0 \xrightarrow{\sim} \sigma^*\mathcal{F}_0^\vee$ exchanging the two summands is Hermitian. Thus $\mathcal{F}_0$ defines a k-point of $\mathrm{Bun}_{U(2n)}$, hence a point $\mathcal{F}_0 \in \mathrm{Sht}^0_{U(2n)}(k)$. Then

$$\widetilde{Z}^{2n,0}_m(\mathcal{G}, \mathcal{E})(\mathcal{F}_0) = \chi_{2n}(\det \mathcal{E}) q^{n(\deg \mathcal{E} - \deg \omega_X)} \sum_{t \in \mathrm{Hom}_{X'}(\mathcal{E}, \mathcal{F}_0)} \psi\big(\langle e_{\mathcal{G},\mathcal{E}}, a(t)\rangle\big), \tag{10.2.5}$$

where $a(t) = \sigma^* t^\vee \circ h_{\mathcal{F}_0} \circ t$ is the Hermitian parameter induced by t. The prefactor in (10.2.5) is nonzero, and the function $t \mapsto \langle e_{\mathcal{G},\mathcal{E}}, a(t)\rangle$ is a homogeneous quadratic form on the finite k-vector space $\mathrm{Hom}_{X'}(\mathcal{E}, \mathcal{F}_0)$. Hence (10.2.5) is nonzero. □

Theorem 10.2.3 (The Modularity Conjecture)**.** *For all $m \le n$ and all $r \ge 0$, the map*

$$\widetilde{Z}^{n,r}_m\colon \mathrm{Bun}_{P_m}(k) \longrightarrow \mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{U(n)})$$

descends to $\mathrm{Bun}^{\mathrm{triv}}_{GU^-(2m)}(k)$. Equivalently, if $\mathcal{G}$ is a trivial-similitude skew-Hermitian bundle of rank $2m$ and $\mathcal{E}_1, \mathcal{E}_2 \subset \mathcal{G}$ are Lagrangian subbundles, then

$$\widetilde{Z}^{n,r}_m(\mathcal{G}, \mathcal{E}_1) = \widetilde{Z}^{n,r}_m(\mathcal{G}, \mathcal{E}_2)$$

in $\mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{U(n)})$.

Proof. Fix the rank n character χ_n. For the auxiliary ranks set $\chi_{2n} = 1$ and $\chi_{3n} = \chi_n$.

Let $\mathcal{G}$ be a trivial-similitude skew-Hermitian bundle of rank $2m$, and let $\mathcal{E}_1, \mathcal{E}_2 \subset \mathcal{G}$ be Lagrangian subbundles. Since $m \le n = (3n)/3$, Theorem 10.1.1 applies to give

$$\widetilde{Z}^{3n,r}_m(\mathcal{G}, \mathcal{E}_1) = \widetilde{Z}^{3n,r}_m(\mathcal{G}, \mathcal{E}_2) \in \mathrm{CH}_{r(3n-m)}(\mathrm{Sht}^r_{U(3n)}). \tag{10.2.6}$$

Pulling (10.2.6) back along

$$\mathrm{add}_{2n,n}\colon \mathrm{Sht}^0_{U(2n)} \times \mathrm{Sht}^r_{U(n)} \longrightarrow \mathrm{Sht}^r_{U(3n)}$$

gives, by Lemma 10.2.1,

$$\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1) \boxtimes \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_1) = \widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_2) \boxtimes \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_2). \tag{10.2.7}$$

The $r = 0$ case of [FYZ25, Conjecture 4.15] is the classical automorphy of the theta series for the dual pair $(GU^-(2m), GU(2n))$, and is proved in [FYZ23, §§2.2–2.3]. Therefore we have $\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1) = \widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_2)$. Putting this into (10.2.7) yields

$$\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1) \boxtimes \left(\widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_1) - \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_2)\right) = 0. \tag{10.2.8}$$

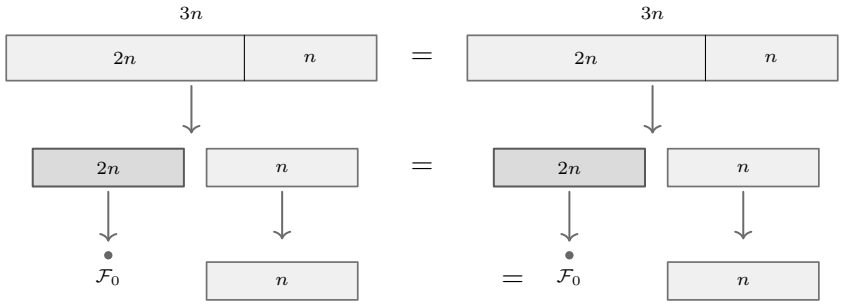

FIGURE 6. Cartoon for the embedding-and-cancellation step following (10.2.8). Pullback from rank $3n$ splits into a rank $2n$ factor and a rank n factor; evaluation at $\mathcal{F}_0$ makes the common rank $2n$ factor nonzero, so the rank n factors agree.

By Lemma 10.2.2, there is a point $\mathcal{F}_0 \in \mathrm{Sht}^0_{U(2n)}(k)$ with $\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1)(\mathcal{F}_0) \neq 0$. Pulling (10.2.8) back along $\{\mathcal{F}_0\} \times \mathrm{Sht}^r_{U(n)} \to \mathrm{Sht}^0_{U(2n)} \times \mathrm{Sht}^r_{U(n)}$ gives

$$\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1)(\mathcal{F}_0) \left(\widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_1) - \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_2)\right) = 0 \quad \text{in} \quad \mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{U(n)}).$$

Since the scalar $\widetilde{Z}^{2n,0}_m(\mathcal{G},\mathcal{E}_1)(\mathcal{F}_0)$ is nonzero, we conclude that

$$\widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_1) = \widetilde{Z}^{n,r}_m(\mathcal{G},\mathcal{E}_2) \quad \text{in} \quad \mathrm{CH}_{r(n-m)}(\mathrm{Sht}^r_{U(n)}).$$

This proves the theorem. □

Part 4. Supermodularity for general linear groups

11. Supermodularity

11.1. **Formulation of supermodularity.** We begin by formulating supermodularity. Fix integers $m \leq n$ and a decomposition $m = m_1 + m_2$ with $m_1, m_2 \geq 0$. Let $P_{(m_1,m_2)} \subset \mathrm{GL}(m)$ be the standard parabolic subgroup with Levi quotient $\mathrm{GL}(m_1) \times \mathrm{GL}(m_2)$, with the convention $\mathrm{GL}(0) = \{e\}$. Thus a point of $\mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$ is a rank m vector bundle $\mathcal{G}$ on X together with a rank m_1 subbundle $\mathcal{E}_1 \subset \mathcal{G}$ whose quotient has the form $\mathcal{E}_2^*$ for a rank m_2 vector bundle $\mathcal{E}_2$. Equivalently, $(\mathcal{E}_1, \mathcal{G})$ determines a short exact sequence

$$0 \longrightarrow \mathcal{E}_1 \longrightarrow \mathcal{G} \longrightarrow \mathcal{E}_2^* \longrightarrow 0, \tag{11.1.1}$$

and we denote its extension class by $e_{\mathcal{G}} \in \mathrm{Ext}^1(\mathcal{E}_2^*, \mathcal{E}_1)$.

Let $\mu = (\mu_1, \ldots, \mu_r)$ be a sequence with $\mu_i \in \{\pm 1\}$ and $\sum_i \mu_i = 0$. For a pair $(\mathcal{E}_1, \mathcal{E}_2)$ as in (11.1.1), the special cycle of [FYZ25, §4] gives a class

$$[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}] \in \mathrm{CH}_{\frac{r}{2}(2n-m)}(\mathrm{Sht}^{\mu}_n).$$

The cycle is naturally decomposed over the Hitchin base $\mathrm{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee)$, where $\mathcal{E}_2^\vee = \mathcal{E}_2^* \otimes \omega_X$:

$$[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}] = \sum_{a \in \mathrm{Hom}(\mathcal{E}_1,\mathcal{E}_2^\vee)} [\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a)]. \tag{11.1.2}$$

For a non-trivial additive character $\psi\colon \mathbf{F}_q \to \overline{\mathbf{Q}}^\times$, the associated higher theta series is

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2}\colon \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q) \longrightarrow \mathrm{CH}_{\frac{r}{2}(2n-m)}(\mathrm{Sht}^\mu_n), \tag{11.1.3}$$

whose value at $(\mathcal{E}_1, \mathcal{G})$ is

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G}) = q^{n\deg \mathcal{E}_2} \sum_{a \in \mathrm{Hom}(\mathcal{E}_1,\mathcal{E}_2^\vee)} \psi(\langle e_{\mathcal{G}}, a\rangle)[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a)]. \tag{11.1.4}$$

Here $\langle e_{\mathcal{G}}, a\rangle \in \mathbf{F}_q$ is the Serre duality pairing between $e_{\mathcal{G}}$ and a. When $m_2 = 0$, the set $\mathrm{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee)$ has a single element and (11.1.4) reduces to $[\mathcal{Z}^\mu_{\mathcal{G},0}]$. The notation in Theorem 11.1.1 suppresses ψ; Lemma 13.1.1 below proves that the value in (11.1.4) is independent of the choice of ψ.

The Modularity Conjecture for general linear groups [FYZ25, Conjecture 4.24] asserts that (11.1.3) descends from $\mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$ to $\mathrm{Bun}_m(\mathbf{F}_q)$, i.e., that (11.1.4) depends on the subbundle $\mathcal{E}_1 \subset \mathcal{G}$ only through $\mathcal{G}$. Supermodularity is the stronger assertion that the higher theta series is also essentially independent of the parabolic.

Theorem 11.1.1 (Supermodularity)**.** *For a cycle $[Z] \in \mathrm{CH}_*(\mathrm{Sht}^\mu_n)$ we write*

$$[Z]^{/\!/} := \sum_{d \in \mathbf{Z}} q^{-m_2 d + m_2 n(g-1)} [Z]_d,$$

where $[Z]_d$ is the projection to the component of Sht^μ_n on which the initial vector bundle has degree d. The sum is understood componentwise with respect to this open-and-closed degree decomposition: equivalently, $[Z]^{/\!/}$ is the class whose restriction to the degree-d component is $q^{-m_2 d + m_2 n(g-1)}[Z]_d$. Then we have

$$\widetilde{Z}^\mu_{m_1,m_2}(\mathcal{E}_1,\mathcal{G}) = [\mathcal{Z}^\mu_{\mathcal{G},0}]^{/\!/} = \sum_{d \in \mathbf{Z}} q^{-m_2 d + m_2 n(g-1)} [\mathcal{Z}^\mu_{\mathcal{G},0}]_d. \tag{11.1.5}$$

11.2. **Proof of Theorem 11.1.1 for $r = 0$.** In this subsection we prove Theorem 11.1.1 in the special case $r = 0$, as a toy model for the general case. In this case, all the geometry becomes trivial and the result is classical. However, it is delicate to formulate a proof that generalizes to $r > 0$, and we have not found our particular argument elsewhere in the literature.

For a rank n vector bundle $\mathcal{F}$ on X, the $r = 0$ theta function is

$$\widetilde{Z}^0_m(\mathcal{E}_1,\mathcal{G})_{\mathcal{F}} := q^{n(\deg \mathcal{E}_2)} \sum_{\substack{t_1 \in \mathrm{Hom}(\mathcal{E}_1,\mathcal{F}) \\ t_2 \in \mathrm{Hom}(\mathcal{E}_2,\mathcal{F}^\vee)}} \psi(\langle e_{\mathcal{G}}, a(t_1,t_2)\rangle). \tag{11.2.1}$$

We write the theta function (11.2.1) as follows. Serre duality gives a pairing

$$\mathrm{ev}\colon \mathrm{Hom}(\mathcal{E}_1,\mathcal{F}) \times \mathrm{Hom}(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X[1]) \to \mathbf{F}_q.$$

Let $f\colon \mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee) \to \mathrm{Ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X)$ be the boundary map in the long exact sequence

$$\ldots \to \mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee) \xrightarrow{f} \mathrm{Ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X) \xrightarrow{g} \mathrm{Ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X) \to \ldots \tag{11.2.2}$$

For $u \in \mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee)$, corresponding to $t_2 \in \mathrm{Hom}(\mathcal{E}_2, \mathcal{F}^\vee)$, set

$$a(t_1,t_2) := u \circ t_1 \in \mathrm{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee).$$

With the standard connecting-homomorphism convention, $f(u)$ is the pullback of the extension

$$0 \to \mathcal{E}_1 \otimes \omega_X \to \mathcal{G} \otimes \omega_X \to \mathcal{E}_2^\vee \to 0$$

along u. Hence

$$\mathrm{ev}(t_1, f(u)) = \mathrm{Tr}_{\mathcal{E}_1}\big(f(u) \circ t_1\big) = \mathrm{Tr}_{\mathcal{E}_2^*}\big((u \circ t_1) \circ e_{\mathcal{G}}\big) = \langle e_{\mathcal{G}}, a(t_1,t_2)\rangle.$$

Noting that $\mathrm{Hom}(\mathcal{E}_2, \mathcal{F}^\vee) \cong \mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee)$, the theta function (11.2.1) can be reformulated as

$$\widetilde{Z}_m^0(\mathcal{E}_1, \mathcal{G})_{\mathcal{F}} = q^{n(\deg \mathcal{E}_2)} \sum_{\substack{t_1 \in \mathrm{Hom}(\mathcal{E}_1, \mathcal{F}) \\ s_1 \in \mathrm{Ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X)}} \psi(\mathrm{ev}(t_1, s_1))(f_! \mathbb{1}_{\mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee)})(s_1), \tag{11.2.3}$$

where $\mathbb{1}$ is the indicator function. Set

$$\phi := f_! \mathbb{1}_{\mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee)}.$$

Consider the commutative diagram below:

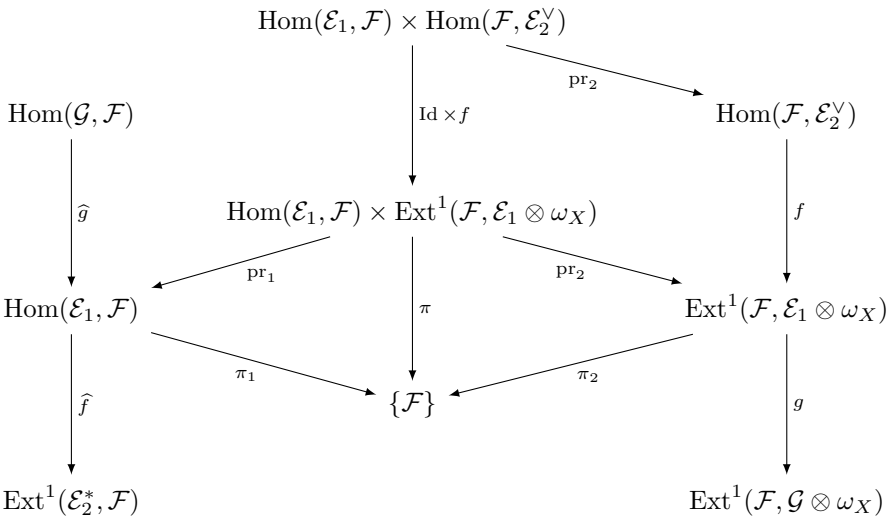

(11.2.4)

Here $\widehat{g}$ and $\widehat{f}$ are the Serre-dual maps of g and f, respectively, and the left column is the linear dual of the right column under Serre duality. In these terms, we have

$$\widetilde{Z}_m^0(\mathcal{E}_1, \mathcal{G})_{\mathcal{F}} = q^{n(\deg \mathcal{E}_2)} \pi_! (\mathrm{ev}^*\, \psi \cdot \mathrm{pr}_2^*\, \phi). \tag{11.2.5}$$

With the finite Fourier transform FT normalized as in [FYZ23, §2.3.7], we have[5]

$$\mathrm{pr}_{1!}(\mathrm{ev}^*\, \psi \cdot \mathrm{pr}_2^*\, \phi) = (-1)^{\hom(\mathcal{E}_1, \mathcal{F})}\, \mathrm{FT}(\phi).$$

We may thus rewrite (11.2.5) as

$$\widetilde{Z}_m^0(\mathcal{E}_1, \mathcal{G})_{\mathcal{F}} = q^{n(\deg \mathcal{E}_2)} (-1)^{\hom(\mathcal{E}_1, \mathcal{F})} \pi_{1!}\, \mathrm{FT}(\phi). \tag{11.2.6}$$

Now, by exactness of (11.2.2), we have

$$\phi = q^{\dim \ker(f)} g^* \delta_{\mathrm{Ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)}, \tag{11.2.7}$$

where δ is the delta function with value 1 at $0 \in \mathrm{Ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)$. Therefore, using the functoriality of FT as in [FYZ23, (2.3.58) and Example 2.3.7], we have

$$\begin{aligned} \mathrm{FT}(\phi) &= q^{\dim \ker(f)}\, \mathrm{FT}(g^* \delta_{\mathrm{Ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)}) \\ &= q^{\dim \ker(f) + \mathrm{ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X) - \mathrm{ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)} (-1)^{\mathrm{ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X)} \widehat{g}_! \mathbb{1}_{\mathrm{Hom}(\mathcal{G}, \mathcal{F})}. \end{aligned}$$

Combining this with (11.2.6), and using $\hom(\mathcal{E}_1, \mathcal{F}) = \mathrm{ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X)$ to simplify signs, we obtain

$$\widetilde{Z}_m^0(\mathcal{E}_1, \mathcal{G})_{\mathcal{F}} = q^{n \deg \mathcal{E}_2 + \dim \ker(f) + \mathrm{ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X) - \mathrm{ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)} \pi_{1!} \widehat{g}_! \mathbb{1}_{\mathrm{Hom}(\mathcal{G}, \mathcal{F})}. \tag{11.2.8}$$

By the long exact sequence (11.2.2), we have

$$\dim \ker(f) = \hom(\mathcal{F}, \mathcal{G} \otimes \omega_X) - \hom(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X).$$

Therefore, the exponent of q in (11.2.8) can be rewritten as

$$n \deg \mathcal{E}_2 + \chi_{\mathrm{Eul}}(X, \mathrm{RHom}(\mathcal{F}, \mathcal{G} \otimes \omega_X)) - \chi_{\mathrm{Eul}}(X, \mathrm{RHom}(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X)) = n \deg \mathcal{E}_2 + \chi_{\mathrm{Eul}}(X, \mathrm{RHom}(\mathcal{F}, \mathcal{E}_2^\vee)).$$

[5]We write $\hom(\dots) := \dim \mathrm{Hom}(\dots)$ and $\mathrm{ext}(\dots) := \dim \mathrm{Ext}(\dots)$, etc.

By Riemann–Roch,

$$\chi_{\mathrm{Eul}}(X, \mathrm{RHom}(\mathcal{F}, \mathcal{E}_2^\vee)) = -m_2 \deg \mathcal{F} - n \deg \mathcal{E}_2 + m_2 n(g-1),$$

so the exponent of q in (11.2.8) may be rewritten as $-m_2 \deg \mathcal{F} + m_2 n(g-1)$. We find that

$$\widetilde{Z}^0_m(\mathcal{E}_1, \mathcal{G})_{\mathcal{F}} = q^{-m_2 \deg \mathcal{F} + m_2 n(g-1)} \pi_{1!} \widehat{g}_! \mathbb{1}_{\mathrm{Hom}(\mathcal{G}, \mathcal{F})}. \tag{11.2.9}$$

Since $\pi_{1!}\widehat{g}_! \mathbb{1}_{\mathrm{Hom}(\mathcal{G},\mathcal{F})} = \#\,\mathrm{Hom}(\mathcal{G}, \mathcal{F}) = [\mathcal{Z}^0_{\mathcal{G},0}]_{\mathcal{F}}$, and the prefactor is the shear factor of Theorem 11.1.1 on the component $d = \deg \mathcal{F}$, this confirms Theorem 11.1.1 in the case $r = 0$. In particular, this visibly depends on the choice of sub-bundle $\mathcal{E}_1 \subset \mathcal{G}$ only through its rank. □

11.3. **Outline for general r.** For $m \le n/3$, the proof of Theorem 11.1.1 follows the same pattern as the case $r = 0$. Its four ingredients are these.

(1) Derived vector bundles $U \to V \to W$ over Bun_n geometrize the sequence

$$\mathrm{Hom}(\mathcal{F}, \mathcal{E}_2^\vee) \to \mathrm{Ext}^1(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X) \to \mathrm{Ext}^1(\mathcal{F}, \mathcal{G} \otimes \omega_X)$$

in the right column of (11.2.4).

The bundles $U' = \widehat{W}$, $V' = \widehat{V}$, and $W' = \widehat{U}$ form a second sequence $U' \to V' \to W'$ geometrizing

$$\mathrm{Hom}(\mathcal{G}, \mathcal{F}) \to \mathrm{Hom}(\mathcal{E}_1, \mathcal{F}) \to \mathrm{Ext}^1(\mathcal{E}_2^*, \mathcal{F})$$

in the left column of (11.2.4).

(2) The six Hecke correspondences $\mathrm{Hk}^\mu_U, \mathrm{Hk}^\mu_V, \mathrm{Hk}^\mu_W$ and $\mathrm{Hk}^\mu_{U'}, \mathrm{Hk}^\mu_{V'}, \mathrm{Hk}^\mu_{W'}$ are derived vector bundles over Hk^μ_n.

(3) The cycle $[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}]$, whose components over the Hitchin base are the Fourier coefficients of $\widetilde{Z}^\mu_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})$, is the trace of a cohomological correspondence $\mathfrak{c}_U$ on Hk^μ_U; the cycle $[\mathcal{Z}^\mu_{\mathcal{G},0}]$ is the trace of a correspondence $\mathfrak{c}_{U'}$ on $\mathrm{Hk}^\mu_{U'}$. The first calculation uses Theorem 13.2.1 under the separate bounds $m_1, m_2 \le n/3$ and does not require $m_1 + m_2 \le n/3$. The second applies it to $(m, 0)$ and uses $m = m_1 + m_2 \le n/3$. This lifts supermodularity to the correspondence level. The dotted arrows in the diagram are Fourier dualities.

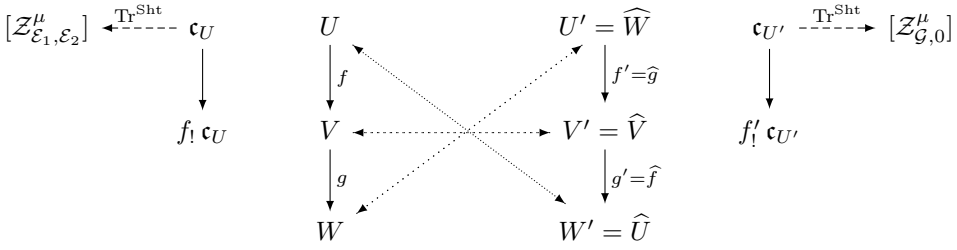

(4) The derived Fourier transform identifies $f_!\,\mathfrak{c}_U$ with $f'_!\,\mathfrak{c}_{U'}$ up to the explicit shift and twist. Pushforward to Hk^μ_S, followed by the trace formula of §8, gives Theorem 11.1.1 for $m \le n/3$; here $S \subset \mathrm{Bun}_n$ is the Harder–Narasimhan truncation of §12.1.1.

For arbitrary m, apply the low-corank identity at target rank $3n$, pull it back along $\mathrm{Sht}^0_{2n} \times \mathrm{Sht}^\mu_n \to \mathrm{Sht}^\mu_{3n}$, and cancel the nonzero zero-leg factor.

12. The genetic pattern of supermodularity

In this section we prove the sheaf-theoretic results that can be regarded as the "genetics of supermodularity" in the sense of Figure 1. Again, this section is essentially agnostic of the ground field.

12.1. **The bestiary of derived vector bundles.** We suppose a short exact sequence of vector bundles on X,

$$0 \to \mathcal{E}_1 \to \mathcal{G} \to \mathcal{E}_2^* \to 0 \tag{12.1.1}$$

where $\operatorname{rank}\mathcal{G} = m$, $\operatorname{rank}\mathcal{E}_1 = m_1$, and $\operatorname{rank}\mathcal{E}_2 = m_2$. The constructions and arguments below apply uniformly for all $m_1, m_2 \geq 0$; for instance, when $m_1 = 0$ one has $V = S$, and the Fourier transform along V is the identity. (The case $m_2 = 0$ is in any case tautological, cf. the proof of Theorem 13.3.3.)

Below we will construct the collection of spaces promised in §11.3. For reasons of familiarity we will use the same notation for the setup as used in §9 (and also in the prior works [FYZ23], and [FK24]), but the actual meaning of the spaces is very different from taking those earlier definitions and specializing them to the situation at hand.

12.1.1. *Definition of U, V and W.* Below we write S for a Harder–Narasimhan truncation $\mathrm{Bun}_n^{\leq \nu}$. (The reason for this truncation is to guarantee *global presentability* in the sense of [FYZ23, §6.3.1].) We define several derived vector bundles over S. Let $\mathcal{F}_{\mathrm{univ}}$ be the universal vector bundle over $X \times S$. Below R denotes any animated $\mathbf{F}_q$-algebra.

- Define $\mathcal{U} := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}, \mathcal{E}_2^\vee)}$ to be the perfect complex on S whose pullback to an R-point $\mathcal{F} \in S(R)$ is naturally in R isomorphic to $\mathrm{RHom}_{X_R}(\mathcal{F}, \mathcal{E}_2^\vee \otimes R)$ regarded as a derived R-module. Let $U := \mathrm{Tot}_S(\mathcal{U})$ be the associated derived vector bundle over S.
- Define $\mathcal{V} := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}, \mathcal{E}_1 \otimes \omega_X[1])}$ to be the perfect complex on S whose pullback to an R-point $\mathcal{F} \in S(R)$ is naturally in R isomorphic to $\mathrm{RHom}_{X_R}(\mathcal{F}, \mathcal{E}_1 \otimes \omega_X \otimes R[1])$ regarded as a derived R-module. Let $V := \mathrm{Tot}_S(\mathcal{V})$ be the associated derived vector bundle over S.
- Define $\mathcal{W} := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ}}, \mathcal{G} \otimes \omega_X[1])}$ to be the perfect complex on S whose pullback to an R-point $\mathcal{F} \in S(R)$ is naturally in R isomorphic to $\mathrm{RHom}_{X_R}(\mathcal{F}, \mathcal{G} \otimes \omega_X \otimes R[1])$ regarded as a derived R-module. Let $W := \mathrm{Tot}_S(\mathcal{W})$ be the associated derived vector bundle over S.

The short exact sequence (12.1.1) induces an exact triangle $\mathcal{U} \to \mathcal{V} \to \mathcal{W}$ in $\mathrm{Perf}(S)$, which geometrically is equivalent to the derived Cartesian square of total spaces over S:

$$\begin{array}{ccc} U & \xrightarrow{f} & V \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle g} \\ S & \xrightarrow{z_W} & W \end{array} \tag{12.1.2}$$

where z_W is the zero section.

12.1.2. *Hecke stacks.* Let $\mu = (\mu_1, \mu_2, \ldots, \mu_r)$ be a sequence such that $\mu_i \in \{\pm 1\}$ for each i and $\sum_i \mu_i = 0$. Then one has the Hecke stack Hk_n^μ, whose R-points are diagrams

$$\mathcal{F}_0 \overset{x_1}{\dashrightarrow} \mathcal{F}_1 \overset{x_2}{\dashrightarrow} \ldots \dashrightarrow \mathcal{F}_{r-1} \overset{x_r}{\dashrightarrow} \mathcal{F}_r \tag{12.1.3}$$

where each $x_i \in X(R)$ and each $\mathcal{F}_i \in \mathrm{Bun}_n(R)$, and the dashed arrow $\mathcal{F}_{i-1} \dashrightarrow \mathcal{F}_i$ is an elementary modification along the graph of x_i, which is upper of length 1 if $\mu_i = +1$ and lower of length 1 if $\mu_i = -1$, in the sense of [YZ17, Definition 5.1]. We shall abbreviate such diagrams as $(\mathcal{F}_\star) \in \mathrm{Hk}_n^\mu(R)$.

Given $(\mathcal{F}_\star)$, we define for each $i = 1, \ldots, r$,

$$\mathcal{F}^\flat_{i-1/2} := \mathcal{F}_{i-1} \cap \mathcal{F}_i \quad \text{and} \quad \mathcal{F}^\sharp_{i-1/2} := \mathcal{F}_{i-1} + \mathcal{F}_i$$

inside the common meromorphic bundle determined by the modification. This fits into an obvious pullback square

$$\begin{array}{ccccc} & & \mathcal{F}^\flat_{i-1/2} & & \\ & \swarrow & & \searrow & \\ \mathcal{F}_{i-1} & & & & \mathcal{F}_i \\ & \searrow & & \swarrow & \\ & & \mathcal{F}^\sharp_{i-1/2} & & \end{array} \tag{12.1.4}$$

There are maps

$$h_i\colon \mathrm{Hk}_n^\mu \to \mathrm{Bun}_n, \quad i = 0, \dots, r \tag{12.1.5}$$

projecting to the datum of $\mathcal{F}_i$, as well as $\mathrm{pr}\colon \mathrm{Hk}_n^\mu \to X^r$ projecting to the data of $(x_1, \dots, x_r)$.

We define the open substack $\mathrm{Hk}_S^\mu \subset \mathrm{Hk}_n^\mu$ as

$$\mathrm{Hk}_S^\mu := h_0^{-1}(S) \cap h_1^{-1}(S) \cap \ldots \cap h_r^{-1}(S).$$

Therefore the maps (12.1.5) restrict to give

$$h_i\colon \mathrm{Hk}_S^\mu \to S, \quad i = 0, \dots, r. \tag{12.1.6}$$

Given a diagram $(\mathcal{F}_\star) \in \mathrm{Hk}_S^\mu(R)$, define $\mathcal{F}_\bullet^\flat$ to be the object of $\mathrm{Perf}(X_R)$ represented by the two-term complex in cohomological degrees 0 and 1

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

where the map sends $(s_{1/2}, \cdots, s_{r-1/2})$ to $(s_{1/2} - s_{3/2}, \cdots, s_{r-3/2} - s_{r-1/2})$, with reference to the commutative diagram (12.1.4). This object need not be quasi-isomorphic to a coherent sheaf.

Define $\mathcal{F}_\bullet^\sharp$ to be the object represented by the two-term complex in cohomological degrees -1 and 0

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

where the map sends $(s_1, \dots, s_{r-1})$ to $(s_1, s_1 - s_2, \dots, s_{r-2} - s_{r-1}, -s_{r-1})$. As in the bestiary construction of [FYZ23, §9.1], this object is quasi-isomorphic to a coherent *sheaf* on X_R, which may not be locally free; we view it as a perfect complex concentrated in degree 0.

We have a natural map of perfect complexes on X_R

$$\mathfrak{p}_i^\flat : \mathcal{F}_\bullet^\flat \to \mathcal{F}_i, \quad i = 0, 1, \cdots, r \tag{12.1.7}$$

that is the composition of the projection to $\mathcal{F}_{i-1/2}^\flat$ and the natural inclusion $\mathcal{F}_{i-1/2}^\flat \hookrightarrow \mathcal{F}_i$ when $i > 0$, and the composition of the projection to $\mathcal{F}_{i+1/2}^\flat$ and the natural inclusion $\mathcal{F}_{i+1/2}^\flat \hookrightarrow \mathcal{F}_i$ when $i < r$. (Both constructions give the same map up to explicit chain homotopy when $0 < i < r$.) Similarly, we have a natural map

$$\mathfrak{p}_i^\sharp\colon \mathcal{F}_i \to \mathcal{F}_\bullet^\sharp, \quad i = 0, \dots, r \tag{12.1.8}$$

As $\mathcal{F}_\star$ varies in Hk_S^μ, the construction $\mathcal{F}_\star \mapsto \mathcal{F}_\bullet^\flat$ gives a universal perfect complex $\mathcal{F}_{\mathrm{univ},\bullet}^\flat$ over $X \times \mathrm{Hk}_S^\mu$, and similarly $\mathcal{F}_\star \mapsto \mathcal{F}_\bullet^\sharp$ gives a universal coherent sheaf $\mathcal{F}_{\mathrm{univ},\bullet}^\sharp$ over $X \times \mathrm{Hk}_S^\mu$, viewed as a perfect complex concentrated in degree 0. These fit into an exact triangle in $\mathrm{Perf}(X \times \mathrm{Hk}_S^\mu)$,

$$\mathcal{F}_{\mathrm{univ},\bullet}^\flat \to h_0^* \mathcal{F}_{\mathrm{univ}} \oplus h_r^* \mathcal{F}_{\mathrm{univ}} \to \mathcal{F}_{\mathrm{univ},\bullet}^\sharp \tag{12.1.9}$$

in which the first map has components $(\mathfrak{p}_0^\flat, \mathfrak{p}_r^\flat)$ and the second is the difference $\mathfrak{p}_0^\sharp - \mathfrak{p}_r^\sharp$.

When $r = 0$, the previous paragraph should be understood as follows: $\mathrm{Hk}_S^\mu = S$, the Hecke correspondence is the identity correspondence, and $\mathcal{F}_\bullet^\flat := \mathcal{F}_0 =: \mathcal{F}_\bullet^\sharp$. Both correspondence maps are the identity, and the triangle (12.1.9) is defined to be the split triangle

$$\mathcal{F}_0 \xrightarrow{(1,1)} \mathcal{F}_0 \oplus \mathcal{F}_0 \xrightarrow{(1,-1)} \mathcal{F}_0.$$

12.1.3. *Hecke stacks for U, V and W.* We now define several perfect complexes on Hk_S^μ, using similar notation as in §12.1.1, so the explanations will be abbreviated.

- Define $\mathcal{U}_{\mathrm{Hk}}^\mu := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ},\bullet}^\sharp, \mathcal{E}_2^\vee)}$. Let $\mathrm{Hk}_U^\mu := \mathrm{Tot}_{\mathrm{Hk}_S^\mu}(\mathcal{U}_{\mathrm{Hk}}^\mu)$ be the associated derived vector bundle over Hk_S^μ.
- Define $\mathcal{V}_{\mathrm{Hk}}^\mu := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ},\bullet}^\sharp, \mathcal{E}_1 \otimes \omega_X[1])}$. Let $\mathrm{Hk}_V^\mu := \mathrm{Tot}_{\mathrm{Hk}_S^\mu}(\mathcal{V}_{\mathrm{Hk}}^\mu)$ be the associated derived vector bundle over Hk_S^μ.
- Define $\mathcal{W}_{\mathrm{Hk}}^\mu := \underline{\mathrm{RHom}(\mathcal{F}_{\mathrm{univ},\bullet}^\sharp, \mathcal{G} \otimes \omega_X[1])}$. Let $\mathrm{Hk}_W^\mu := \mathrm{Tot}_{\mathrm{Hk}_S^\mu}(\mathcal{W}_{\mathrm{Hk}}^\mu)$ be the associated derived vector bundle over Hk_S^μ.

From (12.1.1), we get an exact triangle $\mathcal{U}^\mu_{\mathrm{Hk}} \to \mathcal{V}^\mu_{\mathrm{Hk}} \to \mathcal{W}^\mu_{\mathrm{Hk}}$ of perfect complexes on Hk^μ_S, which geometrically is equivalent to the derived Cartesian square of derived vector bundles over Hk^μ_S:

$$\begin{array}{ccc} \mathrm{Hk}^\mu_U & \xrightarrow{f} & \mathrm{Hk}^\mu_V \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle g} \\ \mathrm{Hk}^\mu_S & \xrightarrow{z} & \mathrm{Hk}^\mu_W \end{array} \tag{12.1.10}$$

where z is the zero section.

12.1.4. *Definition of U', V', W'.* We define further perfect complexes on S, using the same notational pattern as in §12.1.1.

- Define $\mathcal{U}' := \underline{\mathrm{RHom}(\mathcal{G}, \mathcal{F}_{\mathrm{univ}})}$. Let $U' := \mathrm{Tot}_S(\mathcal{U}')$ be its associated derived vector bundle over S.
- Define $\mathcal{V}' := \underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{F}_{\mathrm{univ}})}$. Let $V' := \mathrm{Tot}_S(\mathcal{V}')$ be its associated derived vector bundle over S.
- Define $\mathcal{W}' := \underline{\mathrm{RHom}(\mathcal{E}_2^*, \mathcal{F}_{\mathrm{univ}[1]})}$. Let $W' := \mathrm{Tot}_S(\mathcal{W}')$ be its associated derived vector bundle over S.

From (12.1.1), we have a derived Cartesian square

$$\begin{array}{ccc} U' & \xrightarrow{f'} & V' \\ {\scriptstyle \pi'}\downarrow & & \downarrow{\scriptstyle g'} \\ S & \xrightarrow{z_{W'}} & W' \end{array} \tag{12.1.11}$$

where $z_{W'}$ is the zero section.

12.1.5. *Hecke stacks for U', V', W'.* We now define several perfect complexes on Hk^μ_S which will be used to construct the Hecke stacks for U', V', W'.

- Define $(\mathcal{U}')^\mu_{\mathrm{Hk}} := \underline{\mathrm{RHom}(\mathcal{G}, \mathcal{F}^\flat_{\mathrm{univ},\bullet})}$. Let $\mathrm{Hk}^\mu_{U'} := \mathrm{Tot}_{\mathrm{Hk}^\mu_S}((\mathcal{U}')^\mu_{\mathrm{Hk}})$ be the associated derived vector bundle over Hk^μ_S.
- Define $(\mathcal{V}')^\mu_{\mathrm{Hk}} := \underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{F}^\flat_{\mathrm{univ},\bullet})}$. Let $\mathrm{Hk}^\mu_{V'} := \mathrm{Tot}_{\mathrm{Hk}^\mu_S}((\mathcal{V}')^\mu_{\mathrm{Hk}})$ be the associated derived vector bundle over Hk^μ_S.
- Define $(\mathcal{W}')^\mu_{\mathrm{Hk}} := \underline{\mathrm{RHom}(\mathcal{E}_2^*, \mathcal{F}^\flat_{\mathrm{univ},\bullet[1]})}$. Let $\mathrm{Hk}^\mu_{W'} := \mathrm{Tot}_{\mathrm{Hk}^\mu_S}((\mathcal{W}')^\mu_{\mathrm{Hk}})$ be the associated derived vector bundle over Hk^μ_S.

From (12.1.1), we get an exact triangle $(\mathcal{U}')^\mu_{\mathrm{Hk}} \to (\mathcal{V}')^\mu_{\mathrm{Hk}} \to (\mathcal{W}')^\mu_{\mathrm{Hk}}$ of perfect complexes on Hk^μ_S, which geometrically is equivalent to the derived Cartesian square of derived vector bundles over Hk^μ_S:

$$\begin{array}{ccc} \mathrm{Hk}^\mu_{U'} & \xrightarrow{f'} & \mathrm{Hk}^\mu_{V'} \\ \downarrow{\scriptstyle \pi'} & & \downarrow{\scriptstyle g'} \\ \mathrm{Hk}^\mu_S & \xrightarrow{z'} & \mathrm{Hk}^\mu_{W'} \end{array} \tag{12.1.12}$$

where z' is the zero section.

12.2. **Dualities.** We record the dualities that relate U, V, W and their Hecke stacks with U', V', W' and their Hecke stacks.

12.2.1. *Dual bundles.*

Lemma 12.2.1. *As derived vector bundles over S, Serre duality identifies V' with the dual $\widehat{V}$ of V, U' with the dual $\widehat{W}$ of W, and W' with the dual $\widehat{U}$ of U. Under this identification, the derived fiber square*

(12.1.11) *is identified with the dual fiber square to* (12.1.2)

$$\begin{array}{ccc} \widehat{W} & \xrightarrow{\widehat{g}} & \widehat{V} \\ \downarrow & & \downarrow{\scriptstyle \widehat{f}} \\ S & \xrightarrow{z_{\widehat{U}}} & \widehat{U} \end{array}$$

Proof. Immediate from the definitions and Serre duality (applied on X_R for all animated $\mathbf{F}_q$-algebras R). □

12.2.2. *Hecke stacks.* Let $\widehat{\mathrm{Hk}}^{\mu}_{E}$ be the dual derived bundle to Hk^{μ}_{E}.

For each $i \in \{0, r\}$, let S_i be the corresponding copy of S. For $E \in \{U, V, W, U', V', W'\}$, write E_i for the corresponding copy over S_i, and let $\widetilde{E}_i \to \mathrm{Hk}^{\mu}_{S}$ be the pullback of $E_i \to S_i$ along $h_i \colon \mathrm{Hk}^{\mu}_{S} \to S_i$. The base change maps are denoted

$$h_i^E : \widetilde{E}_i \to E_i. \tag{12.2.1}$$

For $i \in \{0, r\}$, the map $\mathrm{Hk}^{\mu}_{E} \to \widetilde{E}_i$ dualizes to $\widehat{\widetilde{E}}_i \to \widehat{\mathrm{Hk}}^{\mu}_{E}$.

Lemma 12.2.2. *Under the identification of* $U'_i \xrightarrow{f'_i} V'_i \xrightarrow{g'_i} W'_i$ *with* $\widehat{W}_i \xrightarrow{\widehat{g}_i} \widehat{V}_i \xrightarrow{\widehat{f}_i} \widehat{U}_i$ *from Lemma 12.2.1, the six diagrams in* (12.2.2) *and* (12.2.3) *are derived Cartesian.*

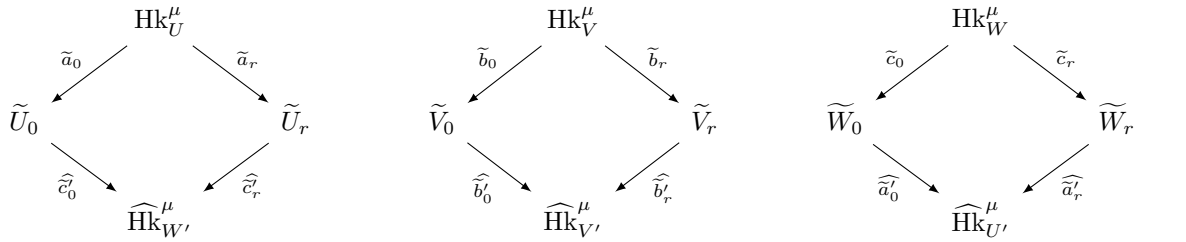

(12.2.2)

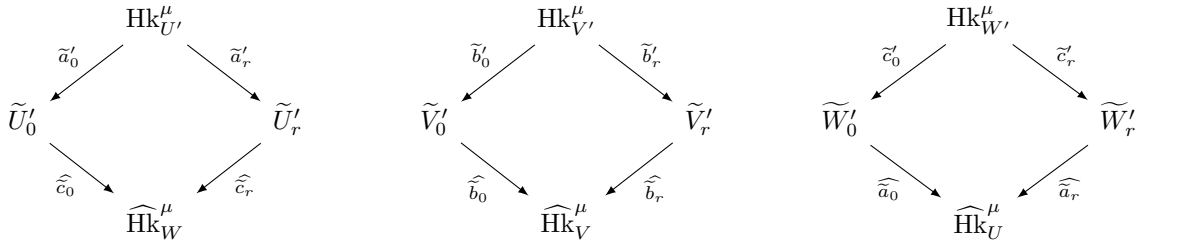

(12.2.3)

Proof. This follows from Serre duality (applied on X_R for all animated $\mathbf{F}_q$-algebras R) and the exact triangle (12.1.9). □

12.2.3. *Summary.* For $i \in \{0, r\}$, the maps (12.1.8) induce the unprimed maps and the maps (12.1.7) induce the primed maps in

$$\begin{aligned} &\widetilde{a}_i : \mathrm{Hk}^{\mu}_{U} \to \widetilde{U}_i, \quad \widetilde{a}'_i : \mathrm{Hk}^{\mu}_{U'} \to \widetilde{U}'_i, \\ &\widetilde{b}_i : \mathrm{Hk}^{\mu}_{V} \to \widetilde{V}_i, \quad \widetilde{b}'_i : \mathrm{Hk}^{\mu}_{V'} \to \widetilde{V}'_i, \\ &\widetilde{c}_i : \mathrm{Hk}^{\mu}_{W} \to \widetilde{W}_i, \quad \widetilde{c}'_i : \mathrm{Hk}^{\mu}_{W'} \to \widetilde{W}'_i. \end{aligned} \tag{12.2.4}$$

Composing the maps in (12.2.4) with the corresponding maps h_i^E of (12.2.1), we obtain

$$\begin{aligned} a_i &:= h_i^U \circ \widetilde{a}_i : \mathrm{Hk}^{\mu}_{U} \to U_i, \quad a'_i := h_i^{U'} \circ \widetilde{a}'_i : \mathrm{Hk}^{\mu}_{U'} \to U'_i, \\ b_i &:= h_i^V \circ \widetilde{b}_i : \mathrm{Hk}^{\mu}_{V} \to V_i, \quad b'_i := h_i^{V'} \circ \widetilde{b}'_i : \mathrm{Hk}^{\mu}_{V'} \to V'_i, \\ c_i &:= h_i^W \circ \widetilde{c}_i : \mathrm{Hk}^{\mu}_{W} \to W_i, \quad c'_i := h_i^{W'} \circ \widetilde{c}'_i : \mathrm{Hk}^{\mu}_{W'} \to W'_i. \end{aligned} \tag{12.2.5}$$

We have a pair of commutative diagrams

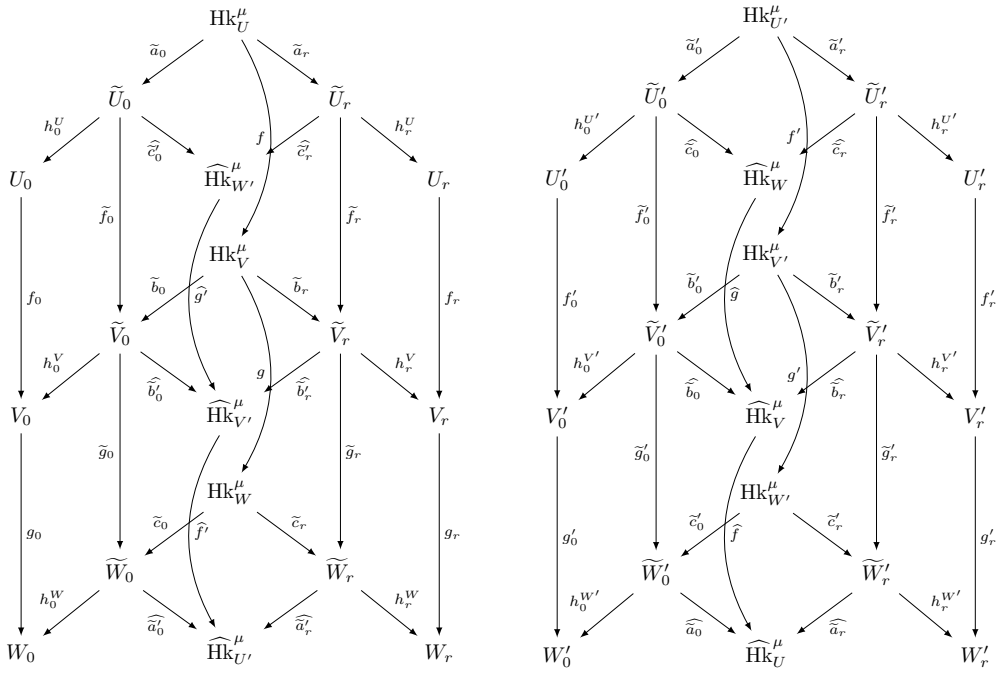

(12.2.6)

In each diagram:

- The maps in the columns come from exact triangles of perfect complexes.
- The three diamonds in the middle are derived Cartesian.
- The four parallelograms on the left and right sides are derived Cartesian.

The primed diagram in (12.2.6) is dual to the unprimed diagram. The duality exchanges U with W', V with V', and W with U'.

Remark 12.2.3 (Global presentability)**.** By the same considerations as in [FYZ23, §9.3], the diagram (12.2.6) is globally presented.

12.3. **Geometric properties.** We establish some needed geometric properties of the maps in (12.2.6).

Lemma 12.3.1. *For $i \in \{0, r\}$, we have the following properties of the morphisms in* (12.2.4)*.*

(1) Each of $\widetilde{a}_i$, $\widetilde{a}'_i$, and $\widetilde{b}'_i$ is a quasi-smooth closed embedding.

(2) Each of $\widetilde{b}_i, \widetilde{c}_i$, and $\widetilde{c}'_i$ is a smooth vector bundle.

Proof. We will first discuss the maps $\widetilde{a}_i$, $\widetilde{b}_i, \widetilde{c}_i$. For an R-point over $(\mathcal{F}_\star)$, let

$$\mathcal{T}_i := \mathrm{cofib}(\mathcal{F}_i \xrightarrow{\mathfrak{p}_i^\sharp} \mathcal{F}_\bullet^\sharp).$$

This glues to a universal complex $\mathcal{T}_{i,\mathrm{univ}}$ on $X \times \mathrm{Hk}_S^\mu$. Note that on classical truncations, $\mathcal{T}_i$ is a torsion coherent sheaf on $X \times \mathrm{Hk}_S^\mu$, supported on the graph of the legs. The relative tangent complex of $\widetilde{a}_i$ to the R-point $(\mathcal{F}_\star)$ is $\mathrm{RHom}_{X_R}(\mathcal{T}_i, \mathcal{E}_2^\vee \otimes R)$, i.e., the pullback of $\underline{\mathrm{RHom}(\mathcal{T}_{i,\mathrm{univ}}, \mathcal{E}_2^\vee)}$ from $X \times \mathrm{Hk}_S^\mu$. As in [FYZ23, Lemma 9.1.3], this complex has tor-amplitude $[0, 1]$, so [FYZ23, Lemma 6.1.5] makes $\widetilde{a}_i$ quasi-smooth. Closedness may be checked on classical truncations, where since $\mathcal{T}_i$ is torsion, this complex is represented by a vector bundle in degree 1. Thus $\widetilde{a}_i$ is a closed embedding, by [FYZ23, Lemma 6.1.5(3)].

The analysis of $\widetilde{b}_i$ and $\widetilde{c}_i$ is similar. The relative tangent complexes of $\widetilde{b}_i$ and $\widetilde{c}_i$ are the pullbacks of $\underline{\mathrm{RHom}(\mathcal{T}_{i,\mathrm{univ}}, \mathcal{E}_1 \otimes \omega_X[1])}$ and $\underline{\mathrm{RHom}(\mathcal{T}_{i,\mathrm{univ}}, \mathcal{G} \otimes \omega_X[1])}$. Both are locally vector bundles in degree 0; hence $\widetilde{b}_i$ and $\widetilde{c}_i$ are smooth vector bundles.

Next we turn to $\widetilde{a}'_i, \widetilde{b}'_i$, and $\widetilde{c}'_i$. Let $\mathcal{K}_i := \mathrm{fib}(\mathcal{F}_\bullet^\flat \xrightarrow{\mathfrak{p}_i^\flat} \mathcal{F}_i)$ be the derived fiber in $\mathrm{Perf}(X_R)$, and let $\mathcal{Q}_{i,\mathrm{univ}}$ be the universal cofiber of $\mathcal{F}_{\mathrm{univ},\bullet}^\flat \to h_i^* \mathcal{F}_{\mathrm{univ}}$. Then $\mathcal{K}_i$ is the pullback of $\mathcal{Q}_{i,\mathrm{univ}[-1]}$. At an R-point,

the relative tangent complexes of $\widetilde{a}'_i$, $\widetilde{b}'_i$, and $\widetilde{c}'_i$ are computed by $\mathrm{RHom}_{X_R}(\mathcal{G}, \mathcal{K}_i)$, $\mathrm{RHom}_{X_R}(\mathcal{E}_1, \mathcal{K}_i)$, and $\mathrm{RHom}_{X_R}(\mathcal{E}_2^*, \mathcal{K}_{i[1]})$, respectively. Globally, they are the pullbacks of $\underline{\mathrm{RHom}(\mathcal{G}, \mathcal{Q}_{i,\mathrm{univ}})}_{[-1]}$, $\underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{Q}_{i,\mathrm{univ}})}_{[-1]}$, and $\underline{\mathrm{RHom}(\mathcal{E}_2^*, \mathcal{Q}_{i,\mathrm{univ}})}$. The first two are locally vector bundles in degree 1, and the last in degree 0. Thus $\widetilde{a}'_i$ and $\widetilde{b}'_i$ are quasi-smooth closed embeddings, while $\widetilde{c}'_i$ is a smooth vector bundle. □

By the argument of [FYZ24, Lemma 6.9], each map $h_i\colon \mathrm{Hk}^\mu_S \to S$ is smooth, schematic, and separated. Since, for every $E \in \{U, V, W, U', V', W'\}$, the map h_i^E in (12.2.1) is a base change of h_i, we get the following corollary.

Corollary 12.3.2. *For $i \in \{0, r\}$, we have the following properties of the morphisms in* (12.2.5).

(1) Each of a_i, a'_i, b'_i is quasi-smooth, schematic, and separated.
(2) Each of b_i, c_i, c'_i is smooth, schematic, and separated.

Lemma 12.3.3. *The commutative diagram*

$$\begin{array}{ccccccc}
U_0 & \xleftarrow{a_0} & \mathrm{Hk}^\mu_U & \xrightarrow{a_r} & U_r & & \\
\downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_r} & & \\
V_0 & \xleftarrow{b_0} & \mathrm{Hk}^\mu_V & \xrightarrow{b_r} & V_r & & \\
\downarrow{\scriptstyle \pi_0} & & \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi_r} & & \\
S_0 & \xleftarrow{g_0} & \mathrm{Hk}^\mu_S & \xrightarrow{g} & S_r & \xleftarrow{g_r} & \\
\downarrow{\scriptstyle z_0} & & \downarrow{\scriptstyle z} & & \downarrow{\scriptstyle z_r} & & \\
W_0 & \xleftarrow{c_0} & \mathrm{Hk}^\mu_W & \xrightarrow{c_r} & W_r & &
\end{array} \tag{12.3.1}$$

satisfies the conditions in [FK24, §4.4.1]. *(Here, z_i and z are the inclusions of zero sections, and π_i and π are the natural projections.) The analogous diagram, obtained by replacing every object and map in the list $U, V, W, f, g, \pi, z, a_i, b_i, c_i$ by its counterpart with $'$, satisfies the same conditions.*

Proof. The schematic and separated properties needed for the horizontal maps in (12.3.1) follow from Corollary 12.3.2.

The vertical squares (U_0, V_0, S_0, W_0), $(\mathrm{Hk}^\mu_U, \mathrm{Hk}^\mu_V, \mathrm{Hk}^\mu_S, \mathrm{Hk}^\mu_W)$ and (U_r, V_r, S_r, W_r) are derived Cartesian by (12.1.2) and (12.1.10). In particular, the rightmost square is pullable with defect zero. It remains to check the pushability or pullability of the remaining squares.

(a) The square $(\mathrm{Hk}^\mu_U, U_0, \mathrm{Hk}^\mu_V, V_0)$ is pushable. For this it suffices to base change all relevant spaces to Hk^μ_S and show instead that

$$\begin{array}{ccc}
\widetilde{U}_0 & \xleftarrow{\widetilde{a}_0} & \mathrm{Hk}^\mu_U \\
\downarrow{\scriptstyle \widetilde{f}_0} & & \downarrow{\scriptstyle f} \\
\widetilde{V}_0 & \xleftarrow{\widetilde{b}_0} & \mathrm{Hk}^\mu_V
\end{array}$$

is pushable. As $\widetilde{a}_0$ is a closed embedding and $\widetilde{b}_0$ is separated by Lemma 12.3.1, the pushability follows from [FYZ23, Example 3.1.2].

(b) The square $(\mathrm{Hk}^\mu_S, S_0, \mathrm{Hk}^\mu_W, W_0)$ is pushable. After base change to Hk^μ_S, it suffices to show that

$$\begin{array}{ccc}
\mathrm{Hk}^\mu_S & = & \mathrm{Hk}^\mu_S \\
\downarrow{\scriptstyle \widetilde{z}_0} & & \downarrow{\scriptstyle z} \\
\widetilde{W}_0 & \xleftarrow{\widetilde{c}_0} & \mathrm{Hk}^\mu_W
\end{array}$$

is pushable. As $\mathrm{Id}_{\mathrm{Hk}^\mu_S}$ is proper and $\widetilde{c}_0$ is separated by Lemma 12.3.1, the pushability follows from [FYZ23, Example 3.1.2].

(c) The square $(\mathrm{Hk}^\mu_U, U_r, \mathrm{Hk}^\mu_S, S_r)$ is pullable. It suffices to check that

$$\begin{array}{ccc} \mathrm{Hk}^\mu_U & \xrightarrow{\widetilde{a}_r} & \widetilde{U}_r \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \widetilde{\pi}_r} \\ \mathrm{Hk}^\mu_S & = & \mathrm{Hk}^\mu_S \end{array}$$

is pullable. As $\widetilde{a}_r$ is quasi-smooth (by Lemma 12.3.1) and Id is smooth, this follows from [FYZ23, Example 3.1.3].

(d) The square $(\mathrm{Hk}^\mu_V, V_r, \mathrm{Hk}^\mu_W, W_r)$ is pullable. It suffices to check that

$$\begin{array}{ccc} \mathrm{Hk}^\mu_V & \xrightarrow{\widetilde{b}_r} & \widetilde{V}_r \\ \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle \widetilde{g}_r} \\ \mathrm{Hk}^\mu_W & \xrightarrow{\widetilde{c}_r} & \widetilde{W}_r \end{array}$$

is pullable. As $\widetilde{c}_r$ and $\widetilde{b}_r$ are smooth by Lemma 12.3.1, this follows from [FYZ23, Example 3.1.3].

The analogous diagram with $'$ is treated identically, using (12.1.11) and (12.1.12) for the vertical derived Cartesian squares. The four side squares are checked by the same base-change reductions as in (a)–(d), with Lemma 12.3.1 applied to the $'$ maps; in the analogue of (d), one uses that $\widetilde{b}'_r$ is quasi-smooth and $\widetilde{c}'_r$ is smooth. □

13. The proof of supermodularity

Finally, in this section we complete the proof of Theorem 11.1.1 (supermodularity).

13.1. **Recollections on higher theta series.** The higher theta series

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2}\colon \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q) \to \mathrm{CH}_{\frac{r}{2}(2n-m)}(\mathrm{Sht}^\mu_n)$$

is defined in [FYZ25, §4.8][6]. Unlike in [FYZ25], we emphasize here the a priori dependence on the non-trivial additive character ψ, although we will shortly see that (in this general linear case) it is independent of the choice of ψ. The value of $\widetilde{Z}^{\psi,\mu}_{m_1,m_2}$ on a pair $(\mathcal{E}_1, \mathcal{G}) \in \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$ is

$$q^{n\deg \mathcal{E}_2} \sum_{a\in \mathrm{Hom}(\mathcal{E}_1,\mathcal{E}_2^\vee)} \psi(\langle e_{\mathcal{G}}, a\rangle)[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a)]$$

where $e_{\mathcal{G}} \in \mathrm{Ext}^1(\mathcal{E}_2^*, \mathcal{E}_1)$ is the extension class of $\mathcal{G}$, paired with a using Serre duality. When $m_2 = 0$, the Hom set in the displayed formula has the single element 0, and the formula reduces to $[\mathcal{Z}^\mu_{\mathcal{G},0}]$.

Lemma 13.1.1. *For any two non-trivial additive characters $\psi, \psi'\colon \mathbf{F}_q \to \overline{\mathbf{Q}}^\times$, we have*

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2} = \widetilde{Z}^{\psi',\mu}_{m_1,m_2}. \tag{13.1.1}$$

Proof. There are two scaling $\mathbf{F}_q^\times$-actions on $\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}$: one scaling the map $t_1\colon \mathcal{E}_1 \to \mathcal{F}_i$, and the other scaling the map $t_2\colon \mathcal{F}_i \to \mathcal{E}_2^\vee$. The derived fundamental class $[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}] \in \mathrm{CH}(\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2})$ is equivariant for each of these two actions; equivalently, it is pulled back from $\mathrm{CH}(\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}/(\mathbf{F}_q^\times \times \mathbf{F}_q^\times))$.

Fix the action scaling t_1, and set $\psi_s(x) := \psi(sx)$ for $s \in \mathbf{F}_q^\times$. If $p_a\colon \mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a) \to \mathrm{Sht}^\mu_n$ is the projection, scaling t_1 induces an isomorphism $\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a) \xrightarrow{\sim} \mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(sa)$. Therefore, their pushforwards to $\mathrm{CH}_{\frac{r}{2}(2n-m)}(\mathrm{Sht}^\mu_n)$ satisfy $[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(sa)] = [\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(a)]$. Using the notation of (11.1.4) and reindexing by $b = sa$, we obtain

$$\widetilde{Z}^{\psi_s,\mu}_{m_1,m_2} = q^{n\deg \mathcal{E}_2} \sum_b \psi(\langle e_{\mathcal{G}}, b\rangle)[\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}(s^{-1}b)] = \widetilde{Z}^{\psi,\mu}_{m_1,m_2}.$$

[6]Our E_1 and E_2 here correspond to $E^{(1)}$ and $E^{(2)}$ of [FYZ25, §4.8].

Every two non-trivial additive characters of $\mathbf{F}_q$ are related by $\psi' = \psi_s$ for a unique $s \in \mathbf{F}_q^\times$, proving (13.1.1). □

The stack Sht^μ_S, defined as the fibered product

$$\begin{array}{ccc} \mathrm{Sht}^\mu_S & \longrightarrow & \mathrm{Hk}^\mu_S \\ \downarrow & & \downarrow \\ S & \xrightarrow{\mathrm{Id}\times\mathrm{Frob}} & S \times S \end{array}$$

is the open Harder–Narasimhan truncation of Sht^μ_n on which every intermediate bundle $\mathcal{F}_i$ lies in $S = \mathrm{Bun}_n^{\leq \nu}$. As ν ranges over all HN polygons, the union of the Sht^μ_S exhausts Sht^μ_n. Therefore, we have

$$\mathrm{CH}_*(\mathrm{Sht}^\mu_n) = \varprojlim_\nu \mathrm{CH}_*(\mathrm{Sht}^\mu_S). \tag{13.1.2}$$

Thus, to prove Theorem 11.1.1, it suffices to prove the restriction of the statement to $\mathrm{CH}(\mathrm{Sht}^\mu_S)$ for all S. The pullback of $\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}$ to Sht^μ_S is $\mathrm{Sht}^\mu_{V'\times_S U}$, the derived fibered product

$$\begin{array}{ccc} \mathrm{Sht}^\mu_{V'\times_S U} & \longrightarrow & \mathrm{Hk}^\mu_{V'} \times_{\mathrm{Hk}^\mu_S} \mathrm{Hk}^\mu_U \\ \downarrow & & \downarrow \\ V' \times_S U & \xrightarrow{\mathrm{Id}\times\mathrm{Frob}} & (V' \times_S U) \times (V' \times_S U) \end{array}$$

13.2. **Supermodularity in low corank.** We first prove supermodularity in the low-corank range using the following refinement of the Trace Conjecture in the split case.

Let $X' = X^{(1)} \coprod X^{(2)}$, with the involution exchanging the two components, and identify a Hermitian bundle with its restriction to $X^{(1)}$. Under this convention, a leg on $X^{(2)}$ gives an upper modification and has sign $\mu_i = +1$, whereas a leg on $X^{(1)}$ gives a lower modification and has sign $\mu_i = -1$.

Theorem 13.2.1 (Trace Conjecture in low corank, split case)**.** *Let $\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2)$ be a pair of vector bundles on X of ranks m_1, m_2, such that $0 \leq m_1 \leq n/3$ and $0 \leq m_2 \leq n/3$. Let $\mathcal{M}_{\mathcal{E}_1,\mathcal{E}_2}$ be the corresponding localized split Hitchin stack. For a sequence $\mu \in \{\pm 1\}^r$ with $\sum_i \mu_i = 0$, put*

$$r_\pm := \#\{i : \mu_i = \pm 1\}, \qquad d_\mu := r_+(n-m_1) + r_-(n-m_2) = \frac{r}{2}(2n - m_1 - m_2).$$

Let $\mathfrak{c}^\mu_{\mathcal{M}_{\mathcal{E}_1,\mathcal{E}_2}}$ denote the restriction of the cohomological correspondence $\mathfrak{c}_{\mathcal{M}_{\mathcal{E}_1,\mathcal{E}_2}}$, defined as in §5.2, to the open-and-closed component of the leg space $(X')^r$ indexed by μ. Then

$$\mathrm{Tr}^{\mathrm{Sht}}\big(\mathfrak{c}^\mu_{\mathcal{M}_{\mathcal{E}_1,\mathcal{E}_2}}\big) = [\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}] \in \mathrm{CH}_{d_\mu}\big(\mathcal{Z}^\mu_{\mathcal{E}_1,\mathcal{E}_2}\big). \tag{13.2.1}$$

Proof. Fix μ and work on the corresponding open-and-closed component of the leg space $(X')^r$. The split Hitchin stack, its cohomological correspondence, and its kernel decomposition are the split-cover specializations of the Hermitian constructions used in Theorem 5.2.1. The proof of Theorem 5.2.1 respects this component and goes through with no essential changes. □

In the application to Theorem 13.2.2, $m = m_1 + m_2 \leq n/3$. Therefore the split theorem applies both to $(\mathcal{E}_1, \mathcal{E}_2)$ and to $(\mathcal{G}, 0)$, whose rank pairs are (m_1, m_2) and $(m, 0)$, respectively.

Theorem 13.2.2 (Supermodularity in low corank)**.** *If $m \leq n/3$, we have*

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = [\mathcal{Z}^\mu_{\mathcal{G},0}]^{/\!/} = \sum_{d\in\mathbf{Z}} q^{-m_2 d + m_2 n(g-1)} [\mathcal{Z}^\mu_{\mathcal{G},0}]_d \in \mathrm{CH}_{\frac{r}{2}(2n-m)}(\mathrm{Sht}^\mu_n). \tag{13.2.2}$$

Proof of Theorem 13.2.2. We write $d(U/S)$, $d(V/S)$, $d(W/S)$, etc. for the virtual relative dimensions of the corresponding derived vector bundles over S. These integers are locally constant on connected components of S. In formulas involving the degree of the universal bundle, such as (13.2.10), all shifts, Tate twists, and Chow degrees are read componentwise on the fixed-degree connected component under consideration.

13.2.1. *Theta sheaves.* Consider the commutative diagram

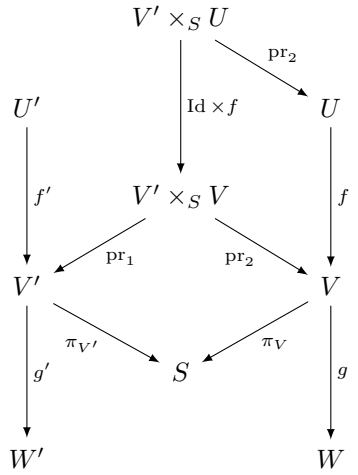

(13.2.3)

Let

$$\mathcal{L}_V := \langle\, ,\, \rangle_V^* \, \mathrm{AS}_{\psi,S} \in \mathcal{D}_{\mathrm{mot}}(V' \times_S V; \overline{\mathbf{Q}})$$

be the Artin–Schreier kernel appearing in (6.2.2). We define

$$\Theta_{\mathcal{E}_1,\mathcal{E}_2} = \pi_{V'!}\, \mathrm{pr}_{1!}(\mathrm{pr}_2^*\, f_! \overline{\mathbf{Q}}_U \otimes \mathcal{L}_V) \in \mathcal{D}_{\mathrm{mot,gm}}(S; \overline{\mathbf{Q}}),$$

and

$$\Theta_{\mathcal{G},0} = \pi_{V'!} f'_!(\overline{\mathbf{Q}}_{U'}) \in \mathcal{D}_{\mathrm{mot,gm}}(S; \overline{\mathbf{Q}}).$$

Unwinding the definition of the motivic Fourier transform (6.2.2), we may rewrite

$$\Theta_{\mathcal{E}_1,\mathcal{E}_2} = \pi_{V'!}\, \mathrm{FT}_V(f_! \overline{\mathbf{Q}}_U)_{[-d(V/S)]} \tag{13.2.4}$$

where FT_V is the motivic Fourier transform of (6.2.2).

Construction 13.2.3. We will construct an isomorphism

$$\Theta_{\mathcal{E}_1,\mathcal{E}_2} \cong \Theta_{\mathcal{G},0}\langle -d(U/S)\rangle \in \mathcal{D}_{\mathrm{mot,gm}}(S; \overline{\mathbf{Q}}). \tag{13.2.5}$$

Note that $f_! \overline{\mathbf{Q}}_U \cong g^* \delta_W$ by proper base change applied to the (derived) Cartesian square (12.1.2). Then, using the Fourier Transform functorialities obtained from Theorem 6.3.1 and Lemma 12.2.1, we have

$$\begin{aligned}
\Theta_{\mathcal{E}_1,\mathcal{E}_2} &= \pi_{V'!}\, \mathrm{FT}_V(f_! \overline{\mathbf{Q}}_U)_{[-d(V/S)]} \\
&\cong \pi_{V'!}\, \mathrm{FT}_V(g^* \delta_W)_{[-d(V/S)]} \\
&\cong \pi_{V'!} \widehat{g}_!\, \mathrm{FT}_W(\delta_W)\langle -d(V/S)\rangle [d(W/S)](d(W/S)) \\
&\cong \pi_{V'!} f'_! \overline{\mathbf{Q}}_{U'}\langle -d(V/S)+d(W/S)\rangle \\
&= \Theta_{\mathcal{G},0}\langle -d(U/S)\rangle.
\end{aligned}$$

13.2.2. *Cohomological correspondences.* We define the following cohomological correspondences.

(i) Since $h_r \colon \mathrm{Hk}_S^\mu \to S$ is smooth, the map of correspondences

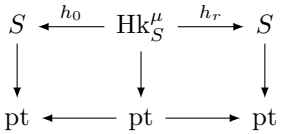

is right pullable. We define $\mathfrak{c}_S \in \mathrm{Corr}_{\mathrm{Hk}_S^\mu}(\overline{\mathbf{Q}}_S, \overline{\mathbf{Q}}_S\langle -d(h_r)\rangle)$ to be the pullback of the identity correspondence in $\mathrm{Corr}_{\mathrm{pt}}(\overline{\mathbf{Q}}_{\mathrm{pt}}, \overline{\mathbf{Q}}_{\mathrm{pt}})$. More explicitly, it is the composition

$$h_0^* \overline{\mathbf{Q}}_S \cong \overline{\mathbf{Q}}_{\mathrm{Hk}_S^\mu} \cong h_r^* \overline{\mathbf{Q}}_S \xrightarrow{\mathrm{gys}_{h_r}} h_r^! \overline{\mathbf{Q}}_S\langle -d(h_r)\rangle.$$

(ii) As shown in part (c) of the proof of Lemma 12.3.3, the map of correspondences

$$\begin{array}{ccccc} U & \xleftarrow{a_0} & \mathrm{Hk}_U^\mu & \xrightarrow{a_r} & U \\ \downarrow{\scriptstyle \pi_0} & & \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi_r} \\ S & \xleftarrow{h_0} & \mathrm{Hk}_S^\mu & \xrightarrow{h_r} & S \end{array}$$

is right pullable. We define $\mathfrak{c}_U = \pi^* \mathfrak{c}_S \in \mathrm{Corr}_{\mathrm{Hk}_U^\mu}(\overline{\mathbf{Q}}_U, \overline{\mathbf{Q}}_U\langle -d(a_r)\rangle)$. Note that we could also have defined $\mathfrak{c}_U$ as the pullback from the identity correspondence in $\mathrm{Corr}_{\mathrm{pt}}(\overline{\mathbf{Q}}_{\mathrm{pt}}, \overline{\mathbf{Q}}_{\mathrm{pt}})$. Explicitly, $\mathfrak{c}_U$ is the composition

$$a_0^* \overline{\mathbf{Q}}_U \cong \overline{\mathbf{Q}}_{\mathrm{Hk}_U^\mu} \cong a_r^* \overline{\mathbf{Q}}_U \xrightarrow{\mathrm{gys}_{a_r}} a_r^! \overline{\mathbf{Q}}_U\langle -d(a_r)\rangle.$$

(iii) The analogous construction to (ii) defines $\mathfrak{c}_{U'} \in \mathrm{Corr}_{\mathrm{Hk}_{U'}^\mu}(\overline{\mathbf{Q}}_{U'}, \overline{\mathbf{Q}}_{U'}\langle -d(a_r')\rangle)$.

(iv) As shown in part (b) of the proof of Lemma 12.3.3, the map of correspondences

$$\begin{array}{ccccc} S & \xleftarrow{h_0} & \mathrm{Hk}_S^\mu & \xrightarrow{h_r} & S \\ \downarrow{\scriptstyle z_0} & & \downarrow{\scriptstyle z} & & \downarrow{\scriptstyle z_r} \\ W & \xleftarrow{c_0} & \mathrm{Hk}_W^\mu & \xrightarrow{c_r} & W \end{array}$$

is left pushable. We define $\mathfrak{c}_W = z_! \mathfrak{c}_S \in \mathrm{Corr}_{\mathrm{Hk}_W^\mu}(\delta_W, \delta_W\langle -d(h_r)\rangle)$.

Definition 13.2.4 (Twist and shift of cohomological correspondences)**.** For a cohomological correspondence $\mathfrak{c} \in \mathrm{Corr}_C(\mathcal{K}_0, \mathcal{K}_1)$ and $c \in \mathbf{Z}$, we abbreviate

$$\mathbb{T}_{\langle c\rangle}\, \mathfrak{c} := \mathbb{T}_{[2c](c)}\, \mathfrak{c} \in \mathrm{Corr}_C(\mathcal{K}_0\langle c\rangle, \mathcal{K}_1\langle c\rangle)$$

where the notation is as in [FK24, §6.5].

Proposition 13.2.5. *(1) With respect to the isomorphism $f_!\overline{\mathbf{Q}}_U \cong g^*\delta_W$ coming from proper base change for the Cartesian square* (12.1.2)*, we have*

$$f_!\, \mathfrak{c}_U = g^*\, \mathfrak{c}_W \in \mathrm{Corr}_{\mathrm{Hk}_V^\mu}(f_!\overline{\mathbf{Q}}_U, f_!\overline{\mathbf{Q}}_U\langle -d(a_r)\rangle). \tag{13.2.6}$$

(2) We have

$$\begin{aligned} \mathrm{FT}(f_!\, \mathfrak{c}_U) &\cong \mathbb{T}_{[2d(W/S)-d(V/S)](d(W/S)-d(V/S))} f_!'(\mathfrak{c}_{U'}) \\ &\in \mathrm{Corr}_{\mathrm{Hk}_{V'}^\mu}\big(\mathrm{FT}_V(f_!\overline{\mathbf{Q}}_U), \mathrm{FT}_V(f_!\overline{\mathbf{Q}}_U)\langle -d(a_r')\rangle\big). \end{aligned} \tag{13.2.7}$$

Proof. (1) Thanks to Lemma 12.3.3, we may apply [FK24, Theorem 4.4.2], in the middle isomorphism below:

$$f_!\, \mathfrak{c}_U = f_!\pi^*\, \mathfrak{c}_S \cong g^* z_!\, \mathfrak{c}_S = g^*\, \mathfrak{c}_W\,.$$

(2) By the global-presentation statement preceding Lemma 12.3.1, Proposition 7.4.1 applies to these diagrams. Lemma 12.3.3(b), together with Lemmas 12.2.1 and 12.2.2, identifies the dual map with $\pi' = \widehat{z}$ and gives

$$\mathrm{FT}(\mathfrak{c}_W) = \mathrm{FT}(z_!\, \mathfrak{c}_S) \cong \mathbb{T}_{[d(W/S)]}(\pi')^*\, \mathrm{FT}_0(\mathfrak{c}_S) = \mathbb{T}_{[d(W/S)]}(\pi')^*\, \mathfrak{c}_S = \mathbb{T}_{[d(W/S)]}\, \mathfrak{c}_{U'}, \tag{13.2.8}$$

where the isomorphism is the left-pushable case of Proposition 7.4.1. We used here that $\mathrm{FT}_0(\mathfrak{c}_S) = \mathfrak{c}_S$ because the zero derived vector bundle is self-dual.

Lemma 12.3.3(d) gives right pullability for g. The right-pullable case of Proposition 7.4.1, together with (1), gives

$$\begin{aligned} \mathrm{FT}(f_!\, \mathfrak{c}_U) &\overset{(13.2.6)}{\cong} \mathrm{FT}(g^*\, \mathfrak{c}_W) = \mathbb{T}_{[d(W/S)-d(V/S)](d(W/S)-d(V/S))} \widehat{g}_!\, \mathrm{FT}(\mathfrak{c}_W) \\ &\overset{(13.2.8)}{=} \mathbb{T}_{[2d(W/S)-d(V/S)](d(W/S)-d(V/S))} f_!'(\mathfrak{c}_{U'}), \end{aligned}$$

as desired. □

13.2.3. *Supermodularity.* We will prove (13.2.2) after restriction to $\mathrm{CH}_*(\mathrm{Sht}^\mu_S)$; as explained in §13.1, this suffices to prove Theorem 13.2.2. We write $(\dots)|_S$ for base change of all projections along $S \hookrightarrow \mathrm{Bun}_n$: for Hecke correspondences it means restriction to $\mathrm{Hk}^\mu_S = \bigcap_{i=0}^r h_i^{-1}(S)$, and for shtuka stacks and Chow classes it means restriction to Sht^μ_S.

By Lemma 12.3.1, the map of correspondences

$$\begin{array}{ccccc} V_0' & \xleftarrow{b_0'} & \mathrm{Hk}^\mu_{V'} & \xrightarrow{b_r'} & V_r' \\ \downarrow & & \downarrow{\scriptstyle \pi_{V'}} & & \downarrow \\ S_0 & \xleftarrow{h_0} & \mathrm{Hk}^\mu_S & \xrightarrow{h_r} & S_r \end{array}$$

is pushable. It also satisfies the remaining hypotheses of Proposition 8.1.1. Indeed, the complexes $\mathcal{V}' = \underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{F}_{\mathrm{univ}})}$ and $(\mathcal{V}')^\mu_{\mathrm{Hk}} = \underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{F}^\flat_{\mathrm{univ},\bullet})}$ are coconnective. The correspondence maps have fiber $\underline{\mathrm{RHom}(\mathcal{E}_1, \mathcal{Q}_{i,\mathrm{univ}})}[-1]$, which has tor-amplitude $[1,\infty)$, and Lemma 12.3.1 identifies $\widetilde{b}_i'$ as the resulting closed embedding of derived vector bundles. The relevant coefficient objects are obtained from unit objects by pullback, tensor product, and pushforward along schematic maps locally of finite type, so they are geometric by [FK24, Remark 3.3.2]. Proposition 8.1.1 therefore implies that pushforward along $\pi_{V'}$ is compatible with formation of $\mathrm{Tr}^{\mathrm{Sht}}$.

With respect to the identification $\Theta_{\mathcal{E}_1,\mathcal{E}_2} \cong \Theta_{\mathcal{G},0}\langle -d(U/S)\rangle$ from (13.2.5), Proposition 13.2.5(2), together with the shift in (13.2.4), gives an identification

$$\mathbb{T}_{[-d(V/S)]}\pi_{V'!}\,\mathrm{FT}(f_!\,\mathfrak{c}_U) \cong \mathbb{T}_{\langle -d(U/S)\rangle}\pi_{V'!}f'_!(\mathfrak{c}_{U'}) \in \mathrm{Corr}_{\mathrm{Hk}^\mu_S}(\Theta_{\mathcal{E}_1,\mathcal{E}_2}, \Theta_{\mathcal{E}_1,\mathcal{E}_2}\langle -d(a_r')\rangle). \tag{13.2.9}$$

Indeed, Proposition 13.2.5(2) gives

$$\pi_{V'!}\,\mathrm{FT}(f_!\,\mathfrak{c}_U) \cong \mathbb{T}_{[d(V/S)]}\mathbb{T}_{\langle -d(U/S)\rangle}\pi_{V'!}f'_!(\mathfrak{c}_{U'}).$$

Recall from (13.2.4) that $\Theta_{\mathcal{E}_1,\mathcal{E}_2} = \pi_{V'!}\,\mathrm{FT}_V(f_!\overline{\mathbf{Q}}_U)[-d(V/S)]$.

Now we invoke the *motivic sheaf-cycle correspondence* of [FK24, §6]. It is clear from [FK24, Remark 3.3.2] that $\Theta_{\mathcal{E}_1,\mathcal{E}_2} \cong \Theta_{\mathcal{G},0}\langle -d(U/S)\rangle$ lies in $\mathcal{D}_{\mathrm{mot,gm}}(S;\overline{\mathbf{Q}}) \subset \mathcal{D}_{\mathrm{mot}}(S;\overline{\mathbf{Q}})$, so we may form the (shtuka-twisted) *trace* of (13.2.9) in the sense of [FK24, §5.4]. On the RHS, Theorem 13.2.1 gives the final equality in

$$\begin{aligned} q^{n\deg\mathcal{E}_2}\,\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{\langle -d(U/S)\rangle}\pi_{V'!}f'_!(\mathfrak{c}_{U'})) &\overset{\text{Prop 8.1.1}}{=} q^{n\deg\mathcal{E}_2}\,\mathrm{Sht}(\pi_{V'}f')_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{\langle -d(U/S)\rangle}\,\mathfrak{c}_{U'}) \\ &= \sum_{d\in\mathbf{Z}} q^{-m_2 d + m_2 n(g-1)}[\mathcal{Z}^\mu_{\mathcal{G},0}]_d|_S \end{aligned} \tag{13.2.10}$$

Here $[\mathcal{Z}^\mu_{\mathcal{G},0}]_d$ is the projection of $[\mathcal{Z}^\mu_{\mathcal{G},0}]$ to the component of Sht^μ_n consisting of shtukas with degree $d \in \mathbf{Z}$. On that component, the trace identity of [FK24, §6.5] gives

$$\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{\langle -d(U/S)\rangle}\,\mathfrak{c}) = q^{d(U/S)}\,\mathrm{Tr}^{\mathrm{Sht}}(\mathfrak{c}),$$

because $\langle -d(U/S)\rangle = [-2d(U/S)](-d(U/S))$. Using Theorem 13.2.1 to compute the shtuka-twisted trace of $\mathfrak{c}_{U'}$ and the calculation from §11.2, we obtain

$$d(U/S) = \chi_{\mathrm{Eul}}(X, \mathrm{RHom}(\mathcal{F}, \mathcal{E}_2^\vee)) = -n\deg\mathcal{E}_2 - m_2\deg\mathcal{F} + m_2 n(g-1),$$

so that $q^{n\deg\mathcal{E}_2}q^{d(U/S)} = q^{-m_2 d + m_2 n(g-1)}$. Finally, the virtual dimension calculation gives $d(a_r') = \frac{r}{2}(2n-m)$, so every term lies in the Chow degree asserted in (13.2.2).

On the left side of (13.2.9), referring to the notation of (13.2.3), we will show that

$$q^{n\deg\mathcal{E}_2}\,\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{[-d(V/S)]}\pi_{V'!}\,\mathrm{FT}(f_!\,\mathfrak{c}_U)) = \widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G})|_S. \tag{13.2.11}$$

To begin, we have

$$\begin{aligned} \mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{[-d(V/S)]}\pi_{V'!}\,\mathrm{FT}(f_!\,\mathfrak{c}_U)) &= \mathrm{Tr}^{\mathrm{Sht}}(\pi_{V'!}\,\mathrm{pr}_{1!}(\mathrm{pr}_2^*\,f_!\,\mathfrak{c}_U \otimes \mathcal{L}_V)) \\ &\overset{(1)}{=} \mathrm{Tr}^{\mathrm{Sht}}(\pi_{V'!}\,\mathrm{pr}_{1!}(\mathrm{Id}\times f)_!(\mathrm{pr}_2^*\,\mathfrak{c}_U \otimes (\mathrm{Id}\times f)^*\mathcal{L}_V)) \\ &\overset{(2)}{=} \mathrm{Sht}(\pi_{V'}\,\mathrm{pr}_1(\mathrm{Id}\times f))_!\,\mathrm{Tr}^{\mathrm{Sht}}(\mathrm{pr}_2^*\,\mathfrak{c}_U \otimes (\mathrm{Id}\times f)^*\mathcal{L}_V) \end{aligned}$$

where the steps are justified as follows.

(1) This follows from the Base Change Theorem for cohomological correspondences, [FK24, Theorem 4.4.2], whose hypotheses are justified by Lemma 12.3.3.
(2) This follows from Proposition 8.1.1 applied to the composite projection from $V' \times_S U$, whose hypotheses are satisfied by Lemma 12.3.1, and (6.1.3).

We claim that

$$q^{n \deg \mathcal{E}_2} \operatorname{Sht}(\pi_{V'} \operatorname{pr}_1(\mathrm{Id} \times f))_! \operatorname{Tr}^{\mathrm{Sht}}(\operatorname{pr}_2^* \mathfrak{c}_U \otimes (\mathrm{Id} \times f)^* \mathcal{L}_V) = \widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})|_S, \tag{13.2.12}$$

which would complete the proof of (13.2.11). First, write $\mathcal{M} := \mathcal{M}_{\mathcal{E}_1,\mathcal{E}_2}$ for the localized Hitchin stack of Theorem 13.2.1, so that $V' \times_S U = \mathcal{M}|_S$. Both $\operatorname{pr}_2^* \mathfrak{c}_U$ and $\mathfrak{c}_{\mathcal{M}}$ are obtained by pulling back the identity correspondence on the point, so $\operatorname{pr}_2^* \mathfrak{c}_U = \mathfrak{c}_{\mathcal{M}}$ by transitivity of pullback of cohomological correspondences. Since $m_1, m_2 \le m \le n/3$, (13.2.1) gives

$$\operatorname{Tr}^{\mathrm{Sht}}(\operatorname{pr}_2^* \mathfrak{c}_U) = \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}}) = [\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]|_S \in \mathrm{CH}_*(\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}|_S). \tag{13.2.13}$$

The map of correspondences, in which the vertical maps are the composite of the projection to the Hitchin base $\operatorname{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee)$ with the linear functional $\langle e_{\mathcal{G}}, - \rangle$,

$$\begin{array}{ccccc} \mathcal{M} & \longleftarrow & \mathrm{Hk}^{\mu}_{\mathcal{M}} & \longrightarrow & \mathcal{M} \\ \downarrow & & \downarrow & & \downarrow \\ \mathbf{A}^1 & = & \mathbf{A}^1 & = & \mathbf{A}^1 \end{array} \tag{13.2.14}$$

induces after applying the Sht construction a map $\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2} \xrightarrow{\mathrm{ev}} \mathbf{F}_q$, which in turn induces a decomposition

$$\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2} = \coprod_{\lambda \in \mathbf{F}_q} \underbrace{(\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2})_\lambda}_{\mathrm{ev}^{-1}(\lambda)}.$$

By inspection of the definitions, the map ev coincides with the pairing of $e_{\mathcal{G}} \in \operatorname{Ext}^1(\mathcal{E}_2^*, \mathcal{E}_1)$ with the map from $\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}$ to the Hitchin base $\operatorname{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee)$. Therefore, we may rewrite the higher theta series as

$$\widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = q^{n \deg \mathcal{E}_2} \sum_{a \in \operatorname{Hom}(\mathcal{E}_1, \mathcal{E}_2^\vee)} \psi(\langle e_{\mathcal{G}}, a \rangle)[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}(a)] = q^{n \deg \mathcal{E}_2} \sum_{\lambda \in \mathbf{F}_q} \psi(\lambda)[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]_\lambda \tag{13.2.15}$$

where $[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]_\lambda$ is the projection of $[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]$ to $\mathrm{CH}_*((\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2})_\lambda)$[7] followed by pushforward to $\operatorname{Sht}^{\mu}_n$. After restricting to S, this pushforward is what was denoted $\operatorname{Sht}(\pi_{V'} \operatorname{pr}_1(\mathrm{Id} \times f))_!$ in (13.2.12), but we suppress it below for ease of notation.

For $\lambda \in \mathbf{F}_q \subset \mathbf{A}^1_{\mathbf{F}_q}$, write $\mathcal{M}_\lambda$, etc. for the derived fiber over λ, and $\iota_\lambda \colon \mathcal{M}_\lambda \hookrightarrow \mathcal{M}$ for the tautological closed embedding. By the commutativity of (13.2.14), the correspondence $\mathcal{M} \leftarrow \mathrm{Hk}^{\mu}_{\mathcal{M}} \to \mathcal{M}$ *stabilizes* $\mathcal{M}_\lambda$ in the sense of [FK24, Definition 7.2.1], hence its Frobenius twist is *contracting* near $\mathcal{M}_\lambda$ [FK24, Example 7.2.2]. Hence [FK24, Construction 7.2.3] applies, to give a restriction of (cohomological) correspondences

$$\iota_\lambda^* \colon \operatorname{Corr}_{\mathrm{Hk}^{\mu}_{\mathcal{M}}}(\mathcal{K}_0, \mathcal{K}_1) \to \operatorname{Corr}_{\mathrm{Hk}^{\mu}_{\mathcal{M}_\lambda}}(\iota_\lambda^* \mathcal{K}_0, \iota_\lambda^* \mathcal{K}_1).$$

By [FK24, Theorem 7.5.1], this restriction is compatible with the formation of $\operatorname{Tr}^{\mathrm{Sht}}$, so that we have

$$\operatorname{Tr}^{\mathrm{Sht}}(\iota_\lambda^* \mathfrak{c}_{\mathcal{M}} \otimes \iota_\lambda^* (\mathrm{Id} \times f)^* \mathcal{L}_V) = \operatorname{Sht}(\iota_\lambda)^* \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}} \otimes (\mathrm{Id} \times f)^* \mathcal{L}_V) = \psi(\lambda)[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]_\lambda|_S \tag{13.2.16}$$

Here, the first equality is the compatibility [FK24, Theorem 7.5.1]. For the second equality, note that $\iota_\lambda^*(\mathrm{Id} \times f)^* \mathcal{L}_V$ is the constant invertible object with Frobenius trace $\psi(\lambda)$ (the restriction of the Artin–Schreier kernel along the constant map λ), so the left side equals $\psi(\lambda) \cdot \operatorname{Sht}(\iota_\lambda)^* \operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}})$, which is $\psi(\lambda)[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]_\lambda|_S$ by (13.2.13) and [FK24, Theorem 7.5.1] applied to $\mathfrak{c}_{\mathcal{M}}$. Summing (13.2.16) over all $\lambda \in \mathbf{F}_q$, we find that

$$\operatorname{Tr}^{\mathrm{Sht}}(\mathfrak{c}_{\mathcal{M}} \otimes (\mathrm{Id} \times f)^* \mathcal{L}_V) = \sum_{\lambda \in \mathbf{F}_q} \psi(\lambda)[\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2}]_\lambda|_S,$$

[7] This projection is the same as the derived fundamental class of $(\mathcal{Z}^{\mu}_{\mathcal{E}_1,\mathcal{E}_2})_\lambda$.

which, in combination with (13.2.15) and the normalizing factor $q^{n\deg\mathcal{E}_2}$ from (11.1.4), finally establishes (13.2.12), hence also (13.2.11).

Putting things together, we have found

$$\begin{aligned}\widetilde{Z}^{\psi,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G})|_S &\overset{(13.2.11)}{=} q^{n\deg\mathcal{E}_2}\,\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{[-d(V/S)]}\pi_{V'!}\,\mathrm{FT}(f_!\,\mathfrak{c}_U))\\ &\overset{(13.2.9)}{=} q^{n\deg\mathcal{E}_2}\,\mathrm{Tr}^{\mathrm{Sht}}(\mathbb{T}_{\langle -d(U/S)\rangle}\pi_{V'!}f'_!(\mathfrak{c}_{U'}))\\ &\overset{(13.2.10)}{=} \sum_{d\in\mathbf{Z}} q^{-m_2d+m_2n(g-1)}[\mathcal{Z}^{\mu}_{\mathcal{G},0}]_d|_S\end{aligned}$$

as desired. □

13.3. **Embedding and cancellation in all coranks.** We now pass from the low-corank supermodularity identity to the full statement by the same embedding-and-cancellation mechanism used in §10.2. Since the argument compares shtukas of different ranks, we temporarily write the target rank as a superscript:

$$\widetilde{Z}^{N,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G})\in \mathrm{CH}_*(\mathrm{Sht}^{\mu}_N),\qquad [\mathcal{Z}^{N,\mu}_{\mathcal{G},0}]\in \mathrm{CH}_*(\mathrm{Sht}^{\mu}_N).$$

For a cycle $Z\in \mathrm{CH}_*(\mathrm{Sht}^{\mu}_N)$, define

$$Z^{/\!\!/_N} := \sum_{d\in\mathbf{Z}} q^{-m_2d+m_2N(g-1)}Z_d, \tag{13.3.1}$$

where Z_d is the component on which the initial rank N bundle has degree d, and the sum is interpreted componentwise as in Theorem 11.1.1. Thus $/\!\!/_n$ is the operator denoted $/\!\!/$ in (11.1.5).

Lemma 13.3.1 (Direct-sum factorization)**.** *Let $N_1,N_2\geq m$. The direct-sum morphism*

$$\mathrm{add}^{\mu}_{N_1,N_2}\colon \mathrm{Sht}^{0}_{N_1}\times\mathrm{Sht}^{\mu}_{N_2}\longrightarrow \mathrm{Sht}^{\mu}_{N_1+N_2} \tag{13.3.2}$$

is LCI (both source and target being smooth Deligne–Mumford stacks, cf. [YZ17, §5]*), so that the Gysin pullback $(\mathrm{add}^{\mu}_{N_1,N_2})^*$ of* (3.5.2) *exists on Chow groups. For every $(\mathcal{E}_1,\mathcal{G})\in \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$, one has*

$$(\mathrm{add}^{\mu}_{N_1,N_2})^*\widetilde{Z}^{N_1+N_2,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G}) = \widetilde{Z}^{N_1,0}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G})\boxtimes\widetilde{Z}^{N_2,\mu}_{m_1,m_2}(\mathcal{E}_1,\mathcal{G}). \tag{13.3.3}$$

Moreover, the sheared special cycles attached to $(\mathcal{G},0)$ satisfy

$$(\mathrm{add}^{\mu}_{N_1,N_2})^*\big[\mathcal{Z}^{N_1+N_2,\mu}_{\mathcal{G},0}\big]^{/\!\!/_{N_1+N_2}} = \big[\mathcal{Z}^{N_1,0}_{\mathcal{G},0}\big]^{/\!\!/_{N_1}}\boxtimes\big[\mathcal{Z}^{N_2,\mu}_{\mathcal{G},0}\big]^{/\!\!/_{N_2}}. \tag{13.3.4}$$

Proof. The dimension formula for the smooth stacks in [YZ17, §5] gives $\dim\mathrm{Sht}^{\mu}_N = rN$ and $\dim\mathrm{Sht}^{0}_{N_1} = 0$. Hence

$$d(\mathrm{add}^{\mu}_{N_1,N_2}) = -rN_1.$$

In particular, the Gysin pullback sends the target degree $\frac{r}{2}(2(N_1+N_2)-m)$ to

$$\frac{r}{2}(2(N_1+N_2)-m)-rN_1 = \frac{r}{2}(2N_2-m),$$

the degree of either product class in (13.3.3) and (13.3.4). Write $\mathcal{Z}^{N,\mu}_{\mathcal{E}_1,\mathcal{E}_2}(a)$ when the target rank is N. Functoriality of derived special cycles in the target variable, [FYZ25, Proposition 7.5], and the identity $a(t^{(1)}\oplus t^{(2)}) = a(t^{(1)})+a(t^{(2)})$ give

$$(\mathrm{add}^{\mu}_{N_1,N_2})^*[\mathcal{Z}^{N_1+N_2,\mu}_{\mathcal{E}_1,\mathcal{E}_2}(a)] = \sum_{a_1+a_2=a}[\mathcal{Z}^{N_1,0}_{\mathcal{E}_1,\mathcal{E}_2}(a_1)]\boxtimes[\mathcal{Z}^{N_2,\mu}_{\mathcal{E}_1,\mathcal{E}_2}(a_2)].$$

Multiplying by $\psi(\langle e_{\mathcal{G}},a\rangle)$, summing over a, and using $q^{(N_1+N_2)\deg\mathcal{E}_2} = q^{N_1\deg\mathcal{E}_2}q^{N_2\deg\mathcal{E}_2}$ proves (13.3.3).

The same derived pullback for $(\mathcal{G},0)$, decomposed by the degrees of the two summands, gives, for every d,

$$(\mathrm{add}^{\mu}_{N_1,N_2})^*[\mathcal{Z}^{N_1+N_2,\mu}_{\mathcal{G},0}]_d = \sum_{d_1+d_2=d}[\mathcal{Z}^{N_1,0}_{\mathcal{G},0}]_{d_1}\boxtimes[\mathcal{Z}^{N_2,\mu}_{\mathcal{G},0}]_{d_2}.$$

Equation (13.3.4) follows after summing over d_1,d_2, since the exponent in (13.3.1) is additive under $(N,d) = (N_1+N_2,d_1+d_2)$. □

Lemma 13.3.2. *Assume $m \leq n$ and fix $(\mathcal{E}_1, \mathcal{G}) \in \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$. Then there is a point $\mathcal{F}_0 \in \mathrm{Sht}^0_{2n}(k)$ such that*

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})(\mathcal{F}_0) \neq 0.$$

Proof. Note that $\mathrm{Sht}^0_{2n}(k) \simeq \mathrm{Bun}_{2n}(k)$ is always non-empty. For any $\mathcal{F}_0 \in \mathrm{Bun}_{2n}(k)$, the calculation of §11.2, applied with target rank $2n$, gives

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})(\mathcal{F}_0) = q^{-m_2 \deg \mathcal{F}_0 + 2m_2 n(g-1)} \, \# \mathrm{Hom}(\mathcal{G}, \mathcal{F}_0) \neq 0,$$

as in (11.2.9): the value is a positive integer times a power of q. □

Theorem 13.3.3 (Supermodularity)**.** *For all $m \leq n$, all decompositions $m = m_1 + m_2$, all μ with entries in $\{\pm 1\}$ and $\sum_i \mu_i = 0$, and all $(\mathcal{E}_1, \mathcal{G}) \in \mathrm{Bun}_{P_{(m_1,m_2)}}(\mathbf{F}_q)$, one has*

$$\widetilde{Z}^{\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = \left[\mathcal{Z}^{\mu}_{\mathcal{G},0}\right]^{/\!\!/} = \sum_{d \in \mathbf{Z}} q^{-m_2 d + m_2 n(g-1)} [\mathcal{Z}^{\mu}_{\mathcal{G},0}]_d. \tag{13.3.5}$$

Proof. If $m_2 = 0$, the identity is tautological from (11.1.4) and (11.1.5). We therefore assume $m_2 > 0$. Since $m \leq n = (3n)/3$, the low-corank supermodularity identity proved in Theorem 13.2.2 gives

$$\widetilde{Z}^{3n,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = \left[\mathcal{Z}^{3n,\mu}_{\mathcal{G},0}\right]^{/\!\!/_{3n}} \quad \text{in} \quad \mathrm{CH}_*(\mathrm{Sht}^{\mu}_{3n}). \tag{13.3.6}$$

Pulling (13.3.6) back along $\mathrm{add}^{\mu}_{2n,n} \colon \mathrm{Sht}^0_{2n} \times \mathrm{Sht}^{\mu}_n \longrightarrow \mathrm{Sht}^{\mu}_{3n}$ gives, using Lemma 13.3.1,

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) \boxtimes \widetilde{Z}^{n,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = \left[\mathcal{Z}^{2n,0}_{\mathcal{G},0}\right]^{/\!\!/_{2n}} \boxtimes \left[\mathcal{Z}^{n,\mu}_{\mathcal{G},0}\right]^{/\!\!/_n}. \tag{13.3.7}$$

The $r = 0$ calculation in §11.2, applied with target rank $2n$, gives

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = \left[\mathcal{Z}^{2n,0}_{\mathcal{G},0}\right]^{/\!\!/_{2n}}.$$

Substituting this identity into (13.3.7) yields

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) \boxtimes \left(\widetilde{Z}^{n,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) - \left[\mathcal{Z}^{n,\mu}_{\mathcal{G},0}\right]^{/\!\!/_n}\right) = 0 \tag{13.3.8}$$

in $\mathrm{CH}_*(\mathrm{Sht}^0_{2n} \times \mathrm{Sht}^{\mu}_n)$.

By Lemma 13.3.2, there is a point $\mathcal{F}_0 \in \mathrm{Sht}^0_{2n}(k)$ such that

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})(\mathcal{F}_0) \neq 0. \tag{13.3.9}$$

Pulling back along $\{\mathcal{F}_0\} \times \mathrm{Sht}^{\mu}_n \to \mathrm{Sht}^0_{2n} \times \mathrm{Sht}^{\mu}_n$ gives

$$\widetilde{Z}^{2n,0}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G})(\mathcal{F}_0) \left(\widetilde{Z}^{n,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) - \left[\mathcal{Z}^{n,\mu}_{\mathcal{G},0}\right]^{/\!\!/_n}\right) = 0.$$

The scalar in (13.3.9) is nonzero, so

$$\widetilde{Z}^{n,\mu}_{m_1,m_2}(\mathcal{E}_1, \mathcal{G}) = \left[\mathcal{Z}^{n,\mu}_{\mathcal{G},0}\right]^{/\!\!/_n},$$

completing the proof. □

References

[BHK+20] Jan H. Bruinier, Benjamin Howard, Stephen S. Kudla, Michael Rapoport, and Tonghai Yang, *Modularity of generating series of divisors on unitary Shimura varieties*, Astérisque (2020), no. 421, Diviseurs arithmétiques sur les variétés orthogonales et unitaires de Shimura, 7–125.

[Bor99] Richard E. Borcherds, *The Gross-Kohnen-Zagier theorem in higher dimensions*, Duke Math. J. **97** (1999), no. 2, 219–233.

[BR25] Jan Hendrik Bruinier and Martin Raum, *Correction to "Kudla's Modularity Conjecture and Formal Fourier-Jacobi Series"*, 2025.

[BRZ24] Jan Hendrik Bruinier, Eugenia Rosu, and Shaul Zemel, *Modularity of special* 0*-cycles on toroidal compactifications of Shimura varieties*, 2024.

[BWR15] Jan Hendrik Bruinier and Martin Westerholt-Raum, *Kudla's modularity conjecture and formal Fourier-Jacobi series*, Forum Math. Pi **3** (2015), e7, 30.

[CD19] Denis-Charles Cisinski and Frédéric Déglise, *Triangulated categories of mixed motives*, Springer Monographs in Mathematics, Springer, Cham, [2019] ©2019.

[Che24a] Ryan Chen, *Co-rank 1 arithmetic siegel–weil i: Local non-archimedean*, 2024.

[Che24b] ———, *Co-rank 1 arithmetic siegel–weil iii: Geometric local-to-global*, 2024.
[FH25] Tony Feng and Michael Harris, *Derived structures in the Langlands correspondence*, The Langlands Program, Proceedings of Symposia in Pure Mathematics, vol. 112.2, American Mathematical Society, Providence, RI, 2025, pp. 471–568.
[FHM25] Tony Feng, Benjamin Howard, and Mikayel Mkrtchyan, *Higher Siegel–Weil formula for unitary groups II: corank one terms*, 2025.
[FK24] Tony Feng and Adeel Khan, *Modularity of higher theta series II: Chow group of the generic fiber*, 2024.
[FYZ23] Tony Feng, Zhiwei Yun, and Wei Zhang, *Modularity of higher theta series I: cohomology of the generic fiber*, 2023, To appear in Publ. Math. Inst. Hautes Études Sci.; arXiv:2308.10979.
[FYZ24] ———, *Higher Siegel–Weil formula for unitary groups: the non-singular terms*, Invent. Math. **235** (2024), no. 2, 569–668.
[FYZ25] ———, *Higher theta series for unitary groups over function fields*, Ann. Sci. Éc. Norm. Supér. (4) **58** (2025), no. 2, 275–388.
[HM25] Benjamin Howard and Keerthi Madapusi, *Kudla's modularity conjecture on integral models of orthogonal Shimura varieties*, Compos. Math. **161** (2025), no. 12, 3380–3454.
[Jin24] Fangzhou Jin, *Trace maps in motivic homotopy and local terms*, Trans. Am. Math. Soc., Ser. B **11** (2024), 215–247 (English).
[Kha19] Adeel A. Khan, *Virtual fundamental classes of derived stacks I*, 2019.
[KM90] Stephen S. Kudla and John 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. (1990), no. 71, 121–172.
[KR11] Stephen Kudla and Michael Rapoport, *Special cycles on unitary Shimura varieties I. Unramified local theory*, Invent. Math. **184** (2011), no. 3, 629–682.
[KR14] ———, *Special cycles on unitary Shimura varieties II: Global theory*, J. Reine Angew. Math. **697** (2014), 91–157.
[KRY06] Stephen S. Kudla, Michael Rapoport, and Tonghai Yang, *Modular forms and special cycles on Shimura curves*, Annals of Mathematics Studies, vol. 161, Princeton University Press, Princeton, NJ, 2006.
[Kud97] Stephen S. Kudla, *Algebraic cycles on Shimura varieties of orthogonal type*, Duke Math. J. **86** (1997), no. 1, 39–78.
[Kud04] ———, *Special cycles and derivatives of Eisenstein series*, Heegner points and Rankin L-series, Math. Sci. Res. Inst. Publ., vol. 49, Cambridge Univ. Press, Cambridge, 2004, pp. 243–270.
[Kud21] ———, *Remarks on generating series for special cycles on orthogonal Shimura varieties*, Algebra Number Theory **15** (2021), no. 10, 2403–2447.
[Luo25] Yu Luo, *Kudla-rapoport conjecture for unramified maximal parahoric level*, 2025.
[Mad22] Keerthi Madapusi, *Derived special cycles on Shimura varieties*, 2022.
[Mae21] Yota Maeda, *The modularity of special cycles on orthogonal Shimura varieties over totally real fields under the Beilinson-Bloch conjecture*, Canad. Math. Bull. **64** (2021), no. 1, 39–53.
[Pol24] Aaron Pollack, *Automatic convergence for siegel modular forms*, 2024.
[Rau26] Martin Raum, *The geometric unitary kudla conjecture*, 2026.
[SSTT22] Ananth N. Shankar, Arul Shankar, Yunqing Tang, and Salim Tayou, *Exceptional jumps of Picard ranks of reductions of K3 surfaces over number fields*, Forum Math. Pi **10** (2022), Paper No. e21, 49.
[Sta20] The Stacks Project Authors, *Stacks Project*, `https://stacks.math.columbia.edu`, 2020.
[YZ17] Zhiwei Yun and Wei Zhang, *Shtukas and the Taylor expansion of L-functions*, Ann. of Math. (2) **186** (2017), no. 3, 767–911.
[YZZ09] Xinyi Yuan, Shou-Wu Zhang, and Wei Zhang, *The Gross-Kohnen-Zagier theorem over totally real fields*, Compos. Math. **145** (2009), no. 5, 1147–1162. MR 2551992
[Zef25] Chaim Avram Zeff, *Toward generic modularity of higher theta series over global function fields and p-adic local fields*, Ph.D. thesis, Columbia University, 2025, ProQuest document ID 3201333388.
[Zha09] Wei Zhang, *Modularity of generating functions of special cycles on Shimura varieties*, ProQuest LLC, Ann Arbor, MI, 2009, Thesis (Ph.D.)–Columbia University.
[Zha21] W. Zhang, *Weil representation and arithmetic fundamental lemma*, Ann. of Math. (2) **193** (2021), no. 3, 863–978.
[Zho26] Tong Zhou, *A motivic derived Fourier transform*, 2026.

University of California Berkeley, Department of Mathematics, Berkeley, CA 94720, USA
Email address: `fengt@berkeley.edu`

Massachusetts Institute of Technology, Department of Mathematics, 77 Massachusetts Avenue, Cambridge, MA 02139, USA
Email address: `zyun@mit.edu`

Massachusetts Institute of Technology, Department of Mathematics, 77 Massachusetts Avenue, Cambridge, MA 02139, USA
Email address: `weizhang@mit.edu`