HORI-MOLOGICAL PROJECTIVE DUALITY

JØRGEN VOLD RENNEMO AND ED SEGAL

Abstract. Kuznetsov has conjectured that Pfaffian varieties should admit non-commutative crepant resolutions which satisfy his Homological Projective Duality. We prove half the cases of this conjecture, by interpreting and proving a duality of non-abelian gauged linear sigma models proposed by Hori.

Contents

1. Introduction

Let V be a vector space of odd dimension v. For any even number $0 \leq 2q < v$, we have a Pfaffian variety

$$\mathrm{Pf}_q \subset \mathbb{P}(\wedge^2 V^\vee)$$

consisting of all 2-forms on V whose rank is at most $2q$. This variety is not a complete intersection, and is usually highly singular – the singularities occur where the rank drops below $2q$. We only get smooth varieties in the cases $q = 1$, which gives the Grassmannian $\mathrm{Gr}(V, 2)$, and $q = \frac{1}{2}(v-1)$, which gives the whole of $\mathbb{P}(\wedge^2 V^\vee)$.

The projective dual of Pf_s is another Pfaffian variety; it's the locus

$$\mathrm{Pf}_s \subset \mathbb{P}(\wedge^2 V)$$

consisting of bivectors of rank at most $2s$, where $2s = v - 1 - 2q$.

This paper achieves two closely-connected goals. The first is to establish that *Homological Projective Duality* (HPD) holds for this pair of varieties. This is a conceptual framework due to Kuznetsov [Kuz07], for understanding how we should compare the derived categories $D^b(X)$ and $D^b(Y)$ for a pair of projectively-dual varieties X and Y. The idea is that we should pick a generic linear subspace L and then look at the derived category of the slice $X \cap \mathbb{P}L$ and of the dual slice $Y \cap \mathbb{P}L^\perp$; then the "interesting part" of these categories will be equivalent. Often there is a critical value of $\dim L$ such that both slices are Calabi-Yau, and they are derived equivalent.

A complicating factor here is that Pfaffian varieties are singular, and it seems that it is not sensible to try to apply HPD to singular varieties. Instead, we replace both of them with non-commutative crepant resolutions. A non-commutative resolution of a variety X is a sheaf of non-commutative algebras A on X which has an appropriate smoothness property, and is Morita-equivalent to $\mathcal{O}_X$ over some open subset. Then instead of working with $\mathcal{O}_X$-modules we work with A-modules, and obtain a category that behaves a lot like the derived category of a geometric resolution. Špenko and Van den Bergh [ŠVdB15] have constructed non-commutative crepant resolutions for Pfaffian varieties, we denote the resolution of Pf_s by A and the resolution of Pf_q by B.

We prove that the non-commutative varieties (Pf_s, A) and (Pf_q, B) are HP dual to each other. For example, in the Calabi-Yau case we have the following result (a special case of Theorem 4.23 and Proposition 4.4):

Theorem 1.1. *Let $L \subset \wedge^2 V^\vee$ be a generic subspace of dimension sv. Then the sheaves A and B restrict to give non-commutative crepant resolutions of the varieties $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$ and $\mathrm{Pf}_q \cap \mathbb{P}L$, the categories*

$$D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp,\ A|_{\mathbb{P}L^\perp}) \qquad \text{and} \qquad D^b(\mathrm{Pf}_q \cap \mathbb{P}L,\ B|_{\mathbb{P}L})$$

are both Calabi-Yau of dimension $2qs + 1$, and they are equivalent.

Our results build on Kuznetsov's own pioneering work [Kuz06], where he proved that for $v \leq 7$ the Grassmannian $\mathrm{Gr}(V, 2)$ is HP dual to a non-commutative resolution of Pf_{v-3}. In the same work, he conjectures that this could be made to work for all Pfaffians, if one could find the correct non-commutative resolutions. Hence we have confirmed Kuznetsov's conjecture (in the case where $\dim V$ is odd; the even case is discussed more below).

Our second goal is to interpret and prove a physical duality proposed by Hori [Hor13]. This duality relates certain gauged linear sigma models (GLSMs). These are gauge theories in two dimensions with $N = (2, 2)$ supersymmetry, and the models in question have symplectic gauge groups. On the basis of various physical arguments, Hori proposes that each such model has a dual model producing an equivalent theory. These dual GLSMs are closely related to projectively-dual Pfaffian varieties, and it is clear that there is a connection to HPD; in fact Hori states that this connection was one of the motivations that lead to his proposal.

In this paper we give a rigorous formulation of this duality at the level of B-branes. Each GLSM should have an associated category of B-branes, with a purely algebro-geometric construction, and then the duality predicts that dual models will produce equivalent categories. We propose a definition of these categories, and prove that with our definition the predicted equivalence does indeed hold. We then use this result to deduce our HPD statement, using some ideas from the physical arguments, along with other mathematical ingredients.

Our proposal for the category of B-branes appears to be new – it is not present in Hori's paper – so we hope that this aspect is a useful contribution to the physics literature.

As mentioned above, another important input for this paper is the work of Špenko and Van den Bergh [ŠVdB15]. They describe a very general procedure for constructing non-commutative (crepant) resolutions for quotient singularities, but our proposed category of B-branes turns out to be an example of their procedure. Hence their results imply that we obtain non-commutative resolutions of the Pfaffian varieties.

As well as these external sources, this paper is very much a continuation of previous work of the authors and their coauthors [ADS15, ST14, Ren15]. Indeed, once we have defined the correct non-commutative resolution and the functor relating the two sides, (both of which were only available for $s, q \leq 4$ before), the proof that HPD holds over the smooth locus is a simple extension of these previous works. However, extending this to a full HPD statement valid over the singular loci requires completely new arguments.

In our first sentence we assumed that $v = \dim V$ was odd, but one can also consider Pfaffians for even-dimensional vector spaces. These too come in projectively-dual pairs Pf_s and Pf_q, where now $2s = v - 2q$, and Kuznetsov conjectures that HPD holds in this case too. However it seems that there is no physical duality in this case, because the field theories defined by the GLSMs are not regular if v is even. Despite this, we can prove some partial results. For ease of exposition we keep v odd for most of the paper, and in Section 5 we explain what is different about the even case.

For a more significant variation, one can swap two-forms $\wedge^2 V^\vee$ for quadratic forms $\mathrm{Sym}^2 V^\vee$. In the GLSMs this corresponds to replacing the symplectic gauge groups with orthogonal groups, and Hori predicts a similar duality.

We have a proof of HPD in this situation which parallels this paper, and which we will present in forthcoming work. This generalises work of Kuznetsov [Kuz08] corresponding to quadratic forms of rank 1, and of Hosono–Takagi [HT13] and the first author [Ren15] for forms of rank ≤ 2.

The remainder of the introduction discusses the constructions used in this paper and sketches the main ideas of the proofs.

Conventions. For an algebra A, we write $D^b(A)$ for the derived category whose objects are complexes of left A-modules with finitely generated cohomology. If A is graded, then $D^b(A)$ means the category of complexes of graded modules.

If (X, A) is a variety with a coherent sheaf of algebras on it, we write $D^b(X, A)$ for the derived category whose objects are complexes of left A-modules with coherent cohomology.

If $\mathcal{E}$ is a chain-complex in some abelian category, we write $h_\bullet(\mathcal{E})$ for its homology object.

Acknowledgements. It's a pleasure to thank Michel Van den Bergh for his patient explanations of [ŠVdB15], and Kentaro Hori, Špela Špenko and Michael Wemyss for illuminating discussions. E.S. would also like to thank Nick Addington, Will Donovan, and Richard Thomas – this paper draws heavily on his previous joint work with them – and Dan Halpern-Leistner for some helpful conversations.

1.1. **The non-commutative resolutions.** Let V be an odd-dimensional vector space as before, and let Q be a symplectic vector space of dimension $2q$. Let $\widetilde{\mathcal{Y}}$ be the stack:

$$\widetilde{\mathcal{Y}} = \big[\operatorname{Hom}(V, Q) \,/\, \operatorname{Sp}(Q)\big]$$

This stack maps to $\wedge^2 V^\vee$ by pulling-back the symplectic form, and its image is the cone

$$\widetilde{\operatorname{Pf}_q} \subset \wedge^2 V^\vee$$

over the Pfaffian variety Pf_q. Thus $\widetilde{\mathcal{Y}}$ is, in some sense, a resolution of $\widetilde{\operatorname{Pf}_q}$.

Since we're ultimately interested in projective varieties, we need to also quotient by rescaling. Let $\operatorname{GSp}(Q)$ denote the 'symplectic similitude' group of Q, the subgroup of $\operatorname{GL}(Q)$ that preserves the symplectic form up to scale. It's a semi-direct product:

$$\operatorname{GSp}(Q) = \operatorname{Sp}(Q) \rtimes \mathbb{C}^*$$

Then we let:

$$\mathcal{Y} = \big[\operatorname{Hom}(V, Q) \,/\, \operatorname{GSp}(Q)\big]$$

This stack is a quotient of $\widetilde{\mathcal{Y}}$ by an additional $\mathbb{C}^*$, and it maps to the stack $[\wedge^2 V^\vee \,/\, \mathbb{C}^*]$ with image $[\widetilde{\operatorname{Pf}_q} \,/\, \mathbb{C}^*]$. The pre-image of the origin is the locus where V maps to an isotropic subspace in Q, if we delete this locus then we get an open substack

$$\mathcal{Y}^{ss} \subset \mathcal{Y}$$

whose underlying scheme is the projective variety Pf_q. Our notation here reflects the fact that $\mathcal{Y}^{ss}$ is the semi-stable locus for the obvious GIT stability condition on $\mathcal{Y}$, but note that if $q > 1$ then it is still an Artin stack.

This stack $\mathcal{Y}^{ss}$ is in some sense a resolution of Pf_q, but it is not yet what we need because $D^b(\mathcal{Y}^{ss})$ is very large, for example its Hochschild homology will be infinite-dimensional. To get a more finite category, more akin to the derived category of an honest geometric resolution, we pick out a subcategory $\mathsf{Br}(\mathcal{Y}^{ss}) \subset D^b(\mathcal{Y}^{ss})$. The notation here stands for 'B-branes', as we'll explain in the next section. To define this subcategory, we start by defining a subcategory of $D^b(\widetilde{\mathcal{Y}})$, as follows.

Recall that irreps of $\operatorname{Sp}(Q)$ are indexed by Young diagrams of height at most $q = \frac{1}{2}\dim(Q)$. For any such Young diagram δ there is a corresponding vector bundle on $\widetilde{\mathcal{Y}}$, associated to that irrep, and we'll denote this vector bundle by:

$$\mathbb{S}^{\langle \delta \rangle} Q$$

Here the operation $\mathbb{S}^{\langle \delta \rangle}$ is a 'symplectic Schur functor', we'll use the notation $\mathbb{S}^{\delta}$ for ordinary (GL) Schur functors.

For any $a, b \in \mathbb{N}$, let's write $Y_{a,b}$ for the set of Young diagrams of height at most a and width at most b. To define the category $\mathsf{Br}(\widetilde{\mathcal{Y}})$, we consider the set $Y_{q,s}$, where

$$s = \tfrac{1}{2}(v-1) - q$$

as in the classical projective duality discussed above. There is a corresponding set of vector bundles on $\widetilde{\mathcal{Y}}$, and we define $\mathsf{Br}(\widetilde{\mathcal{Y}})$ to be the subcategory generated by these vector bundles, that is:

$$\mathsf{Br}(\widetilde{\mathcal{Y}}) = \langle\, \mathbb{S}^{\langle \delta \rangle} Q,\ \gamma \in Y_{q,s} \,\rangle \quad \subset D^b(\widetilde{\mathcal{Y}})$$

Our original motivations for this definition were the analogy with Kapranov's exceptional collections on Grassmannians [Kap84], and Hori's calculation of the Witten index of the associated GLSM (see the next section). However, a much more compelling reason to consider it is found in the recent work of Špenko–Van den Bergh [ŠVdB15]. The category $\mathsf{Br}(\widetilde{\mathcal{Y}})$ has a tilting bundle

$$\widetilde{T} = \bigoplus_{\delta \in Y_{q,s}} \mathbb{S}^{\langle \delta \rangle} Q$$

essentially by definition, and then $\widetilde{B} = \mathrm{End}_{\widetilde{\mathcal{Y}}}(\widetilde{T})$ is a non-commutative algebra defined over the ring of functions on the quotient singularity $\mathrm{Hom}(V,Q)/\mathrm{Sp}(Q) = \widetilde{\mathrm{Pf}_q}$. Špenko and Van den Bergh prove that $\widetilde{B}$ is a non-commutative crepant resolution of this singularity. In particular, we have an equivalence

$$\mathrm{Hom}(\widetilde{T}, -) : \mathsf{Br}(\widetilde{\mathcal{Y}}) \xrightarrow{\sim} D^b(\widetilde{B}\text{-mod})$$

and $\mathsf{Br}(\widetilde{\mathcal{Y}})$ has all the properties of the derived category of a geometric crepant resolution. For example, it is a 'non-compact Calabi–Yau' in the following sense: given objects $E, F \in \mathsf{Br}(\widetilde{\mathcal{Y}})$ with F supported over $0 \in \widetilde{\mathrm{Pf}_q}$, we have

$$\mathrm{Hom}^\bullet(E,F) \cong \mathrm{Hom}^{\dim \widetilde{\mathrm{Pf}_q} - \bullet}(F,E)^\vee,$$

by [VdB04, Lemma 6.4.1]. In fact Špenko–Van den Bergh give a very general construction that produces a non-commutative resolution for any quotient of a smooth affine variety by a reductive group, but they cover $\widetilde{\mathrm{Pf}_q}$ as an explicit example, and show that in this case the resolution can be chosen to be crepant [ŠVdB15, Section 6].

Now we add in our additional $\mathbb{C}^*$-action, and define a subcategory

$$\mathsf{Br}(\mathcal{Y}) \subset D^b(\mathcal{Y})$$

to be the pre-image of $\mathsf{Br}(\widetilde{\mathcal{Y}})$ under pull-back along the map $\widetilde{\mathcal{Y}} \to \mathcal{Y}$. This category can also be defined by a 'grade-restriction-rule' at the origin, see Section 3.1.

We treat $\mathcal{Y}$ as being defined relative to the base $[\wedge^2 V^\vee / \mathbb{C}^*]$, so the morphisms in $D^b(\mathcal{Y})$ are (chain-complexes of) graded modules over $\mathcal{O}_{\wedge^2 V^\vee}$, supported on $\widetilde{\mathrm{Pf}_q}$. We can lift $\widetilde{T}$ to a vector bundle T on $\mathcal{Y}$, then we have a graded algebra $B = \mathrm{End}_{\mathcal{Y}}(T)$ whose underlying ungraded algebra is $\widetilde{B}$, and $\mathsf{Br}(\mathcal{Y})$ is equivalent to $D^b(B\text{-mod})$. Deleting the origin, we obtain a sheaf of non-commutative algebras on the projective variety Pf_q. This is our non-commutative resolution.

As mentioned above, for HPD we are interested in slicing Pf_q, *i.e.* considering $\mathrm{Pf}_q \cap \mathbb{P}L$ for a subspace $L \subset \wedge^2 V^\vee$. We can restrict B to $\mathrm{Pf}_q \cap \mathbb{P}L$ and get a sheaf of algebras $B|_{\mathbb{P}L}$; for generic L this will yield a non-commutative crepant resolution of $\mathrm{Pf}_q \cap \mathbb{P}L$, by the 'non-commutative Bertini theorem' [RSVdB17].

We can of course do everything in a precisely analogous way for the projectively-dual Pfaffian $\mathrm{Pf}_s \subset \wedge^2 V$. We fix a symplectic vector space S of dimension $2s$, and define the stacks:

$$\widetilde{\mathcal{X}} = \big[\mathrm{Hom}(S,V)\,/\,\mathrm{Sp}(S)\big] \qquad \text{and} \qquad \mathcal{X} = \big[\mathrm{Hom}(S,V)\,/\,\mathrm{GSp}(s)\big]$$

To define a subcategory $\mathsf{Br}(\widetilde{\mathcal{X}})$ we just 'rotate our rectangle', and consider the set of Young diagrams $Y_{s,q}$. These all correspond to irreps of $\mathrm{Sp}(S)$, and we define:

$$\mathsf{Br}(\widetilde{\mathcal{X}}) = \langle\, \mathbb{S}^{\langle\gamma\rangle} Q,\ \gamma \in Y_{s,q} \,\rangle \quad \subset D^b(\widetilde{\mathcal{X}})$$

Now we proceed through the same steps: we have an algebra $\widetilde{A}$ which is a non-commutative crepant resolution of $\widetilde{\mathrm{Pf}_s}$, and a graded algebra A which restricts to give a sheaf of algebras over Pf_s.

It should already be evident that at the most crude level, our duality comes down to a bijection between the sets $Y_{q,s}$ and $Y_{s,q}$. However even at this level one must take care to choose the right bijection! For any $\gamma \in Y_{s,q}$, let us denote by $\gamma^c \in Y_{s,q}$ the complement of γ in a rectangle of height s and width q. The relevant bijection for us is the function:

$$\begin{aligned} Y_{s,q} &\xrightarrow{\sim} Y_{q,s} \\ \gamma &\mapsto (\gamma^c)^\top \end{aligned} \tag{1.2}$$

Unsurprisingly, it will quite a lot of work to lift this to an actual comparison of any categories.

1.2. **A sketch of the physics.** For a string-theorist, a stack like $\mathcal{Y}$ or $\widetilde{\mathcal{Y}}$ can be thought of as the input data for a non-abelian GLSM. This kind of GLSM was analysed in some detail in [Hor13] (see also [HK13]) and a duality was proposed, as discussed above. In this section we'll give a very rough summary of this proposed duality, and its connection to HPD for Pfaffians.

Let's begin with the stack $\mathcal{Y}$, which corresponds to a theory with gauge group $\mathrm{GSp}(Q)$. This theory has a Fayet–Iliopoulos parameter r, which roughly corresponds to the value of the moment map or GIT stability condition, and for $r \gg 0$ the classical space of vacua is the GIT quotient $Y^{ss}/\mathrm{GSp}(Q) = \mathrm{Pf}_q$. One might expect that the theory reduces, in some limit, to the σ-model with target Pf_q. This can only be approximately correct because Pf_q is singular, but one might hope that quantum corrections somehow resolve the singularities.

To get a σ-model on a slice $\mathrm{Pf}_q \cap \mathbb{P}L$, we perform a standard trick that goes back to Witten. Write $L^\perp \subset \wedge^2 V$ for the annihilator of L, and define an action of $\mathrm{GSp}(Q)$ on $L^\perp$ by setting the subgroup $\mathrm{Sp}(Q)$ to act trivially, and the residual $\mathbb{C}^*$ to act diagonally with weight -1. Then we add this into our stack/GLSM data, forming:

$$\mathcal{Y} \times_{\mathbb{C}^*} L^\perp = \left[\mathrm{Hom}(V, Q) \times L^\perp \,/\, \mathrm{GSp}(Q) \right]$$

This stack has a canonical invariant function on it, namely

$$W(y, a) = \omega_Q(\wedge^2 y(a)) \tag{1.3}$$

where $y \in \mathrm{Hom}(V, Q)$ and $a \in L^\perp$ and ω_Q is the symplectic form. We add W to our GLSM as a superpotential.

If we didn't add W, then in the $r \gg 0$ phase we might expect to get the σ-model on the corresponding GIT quotient, which is the total space of the vector bundle $L^\perp(-1)$ over Pf_q. The presence of W localizes the theory onto $\mathrm{Crit}(W)$. After quotienting by $\mathrm{GSp}(Q)$ the superpotential becomes quadratic, and $\mathrm{Crit}(W)$ is the subvariety:

$$\{a = 0,\ y \circ a = 0\} = \mathrm{Pf}_q \cap \mathbb{P}L$$

In fact this is only true over the smooth locus in Pf_q, really $\mathrm{Crit}(W)$ has some non-compact branches over the singular locus. But again we might dream that quantum corrections will somehow solve this problem:

The addition of $L^\perp$ has another important consequence. In the previous model if we set $r \ll 0$ then the (classical) vacuum space is empty, but in the new model this is not true and we have a second interesting phase; however, this second phase is more difficult to analyze. In GIT terms, only the locus $\{a = 0\}$ is unstable, and we can think of this phase as family of models living over $\mathbb{P}L^\perp$. Each fibre is the stack/GLSM $\widetilde{\mathcal{Y}}$, equipped with a superpotential W_a which varies with the point $[a] \in \mathbb{P}L^\perp$. To connect to (homological) projective duality, we need to show that these fibre-wise models are very simple: they give no contribution at all unless $[a]$ lies in the dual Pfaffian Pf_s, and when $[a]$ does lie in Pf_s they look like a σ-model on a point.

In earlier work [HT07], Hori and Tong directly analyse these fibre-wise models in the case $q = 1$, and give some arguments that this desired conclusion holds. However in [Hor13], Hori gives a much cleaner approach. Based on various physical arguments (that we do not understand well enough to summarize), he proposes that $\widetilde{\mathcal{Y}}$ has a dual description as the model:

$$\widetilde{\mathcal{X}}\times\wedge^2V^\vee = \big[\operatorname{Hom}(S,V)\,/\,\operatorname{Sp}(S)\big]\times\wedge^2V^\vee$$

Here, as in the previous section, S is a symplectic vector space of dimension $2s = v-2q-1$. This dual model comes equipped with a tautological superpotential

$$W'(x,b) = b(\wedge^2 x(\beta_S)) \tag{1.4}$$

where $x\in\operatorname{Hom}(S,V)$ and $b\in\wedge^2V$ and β_S is the Poisson bivector. We will explain shortly how we interpret this duality as predicting an equivalence of categories, but let us first fill in the final steps connecting it to projective duality.

As just stated this duality looks very asymmetric, since the dual side has 'extra directions' $\wedge^2V^\vee$ and a superpotential. To correct this asymmetry we choose a subspace $L\subset\wedge^2V^\vee$, and cross both sides with $L^\perp$. On the original side this gives the model $\widetilde{\mathcal{Y}}\times L^\perp$, and to this we can add a superpotential W as in (1.3). On the dual side, we get:

$$\widetilde{\mathcal{X}}\times\wedge^2V^\vee\times L^\perp$$

Under the duality, the variable $\omega_Q(\wedge^2 y)\in\wedge^2V^\vee$ corresponds to b. Hence adding W on the original side corresponds, on the dual side, to adding the term $b(a)$ to the existing superpotential W'. This term is quadratic, so we may integrate out its non-degenerate part, which means deleting the directions $(L^\perp)^\vee\times L^\perp$. What remains is the model $\mathcal{X}\times L$. So a more symmetric way to state the duality is that it exchanges

$$\widetilde{\mathcal{Y}}\times L^\perp \quad\leftrightarrow\quad \widetilde{\mathcal{X}}\times L$$

with their tautological superpotentials. Now we simply add the additional $\mathbb{C}^*$ action, promoting our gauge groups to $\operatorname{GSp}(Q)$ and $\operatorname{GSp}(S)$. This introduces an FI parameter, and the duality exchanges the 'easy' ($r\gg 0$) and 'difficult' ($r\ll 0$) phases of the two sides. We have argued that the original model reduces in the easy phase to a σ-model on the target $\operatorname{Pf}_q\cap\mathbb{P}L$, or some kind of resolution thereof, so when we pass to the difficult phase and apply the duality we must reduce to a σ-model with the target $\operatorname{Pf}_s\cap\mathbb{P}L^\perp$.

Now we discuss the implications of this story for B-branes. A σ-model on a smooth variety X should have an associated category of B-branes, and everyone knows that this is the derived category $D^b(X)$. A GLSM should also have an associated category of B-branes, but this is much less well understood, particularly if the model has FI parameters. The GLSM $\widetilde{\mathcal{Y}}$ has no FI parameters because the symplectic group is simple, so one can predict with some confidence that there should be a single associated category. Our proposal is that the category of B-branes in this GLSM is the subcategory

$$\mathsf{Br}(\widetilde{\mathcal{Y}})\subset D^b(\widetilde{\mathcal{Y}})$$

defined in the previous section. Our evidence for this is:

(1) $\mathsf{Br}(\widetilde{\mathcal{Y}})$ is a non-commutative crepant resolution, so it behaves like a non-compact Calabi–Yau variety. This would seem to be a desirable property for the B-brane category. Furthermore, Van den Bergh conjectures that all (non-commutative or commutative) crepant resolutions are derived equivalent [VdB04, Conj. 4.6], if this is correct then this property uniquely determines the B-brane category.

(2) Hori calculates [Hor13, Section 5.3] that the Witten index of this model is $\binom{q+s}{q}$, and this is the size of the set $Y_{q,s}$ indexing the generators of $\mathsf{Br}(\widetilde{\mathcal{Y}})$. One might conjecture that it is also the Euler characteristic of the cyclic homology of the category.

Of course we also propose that the category of B-branes for the GLSM $\widetilde{\mathcal{X}}$ should be the subcategory $\mathsf{Br}(\widetilde{\mathcal{X}}) \subset D^b(\widetilde{\mathcal{X}})$. To interpret Hori's duality we need a little bit more, we need to identify the category of B-branes in the GLSMs $\widetilde{\mathcal{Y}}\times L^\perp$ and $\widetilde{\mathcal{X}}\times L$ with their tautological superpotentials W and W'. Fortunately there is an obvious way to generalize our proposal – we consider analogous subcategories inside the categories of matrix factorizations:

$$D^b(\widetilde{\mathcal{Y}}\times L^\perp, W) \qquad \text{and} \qquad D^b(\widetilde{\mathcal{X}}\times L, W')$$

A matrix factorization can be represented as a vector bundle, equipped with a 'twisted differential', and we define full subcategories

$$\mathsf{Br}(\widetilde{\mathcal{Y}}\times L^\perp, W) \qquad \text{and} \qquad \mathsf{Br}(\widetilde{\mathcal{X}}\times L, W')$$

where we insist that this vector bundle is a direct sum of the bundles coming from the sets $Y_{q,s}$ or $Y_{s,q}$ (a slightly more elegant definition of these categories is given in Section 3.1).

So for us, Hori's duality becomes the predicted equivalence of categories

$$\mathsf{Br}(\widetilde{\mathcal{Y}}\times L^\perp, W) \cong \mathsf{Br}(\widetilde{\mathcal{X}}\times L, W')$$

for any L. In fact we really want this with the additional $\mathbb{C}^*$ in place, so it's an equivalence between categories defined on $\mathcal{Y}\times_{\mathbb{C}^*} L^\perp$ and $\mathcal{X}\times_{\mathbb{C}^*} L$.

1.3. **A sketch of our proof.** There are two things we need to prove: firstly that our interpretation of Hori's duality holds, and secondly that this implies HPD for Pfaffians. The first point will be proved in Section 3, and the second in Section 4. Here we give a sketch of both proofs.

We begin with Hori's duality in the extreme case $L = 0$, so we want to prove the equivalence:

$$\mathsf{Br}(\mathcal{X}) \;\cong\; \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} \wedge^2 V,\, W)$$

Recall that the stack $\mathcal{X}$ maps to $[\wedge^2 V/\mathbb{C}^*]$, hitting the locus $[\widetilde{\mathrm{Pf}_s}/\mathbb{C}^*]$. The other stack $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$ also maps to $[\wedge^2 V/\mathbb{C}^*]$, just by projection onto the second factor. An important aspect of our argument is that our equivalence will be constructed relative to this common base, *i.e.* it comes from a Fourier-Mukai kernel defined on their relative product. The definition of this kernel is fairly straight-forward (see Section 3.3).

Having defined our functor, we can then base-change to open substacks of $[\wedge^2 V/\mathbb{C}^*]$ and examine it there. In particular we can delete the rank $< 2s$ locus, this removes all the singularities of $\widetilde{\mathrm{Pf}_s}$, and $\mathcal{X}$ becomes equivalent to the smooth locus in Pf_s. Here the methods of our previous papers apply, and we use them to prove that over this locus our functor gives an equivalence between the Br subcategories on each side (see Section 3.4).

Next we need to extend this over the singularities (Section 3.5). As we discussed in Section 1.1, $\mathsf{Br}(\mathcal{X})$ is generated by a finite set of objects, and is equivalent to the derived category of their endomorphism algebra A. On the other side, we identify a 'dual' set of generating objects in the category $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V, W)$, and prove that the endomorphisms of these dual objects also form a Cohen–Macaulay algebra. Because our functor is generically an equivalence these two algebras are generically isomorphic, and then we can use the Cohen–Macaulay property to deduce that they are isomorphic everywhere. It follows immediately that our functor is an equivalence.

This proves the case $L = 0$, and then it's easy to prove it for general L using the physical sketch of the previous section - just replace 'integrating out the quadratic term' with Knörrer periodicity. Then we have an equivalence

$$\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \;\cong\; \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$$

which is our version of Hori's duality (Theorem 3.2).

Now we move on to deducing HPD. Our equivalence is relative to $\wedge^2 V$, so restricting to the complement of the origin gives an equivalence:

$$\mathsf{Br}(\mathcal{X}^{ss} \times_{\mathbb{C}^*} L,\, W') \;\cong\; \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0),\, W)$$

In the terminology of the previous section, this relates the 'easy' phase on the left-hand side with the 'difficult' phase on the right-hand side. Let's discuss the left-hand side first. The non-stacky locus in $\mathcal{X}^{ss}$ is the smooth locus in Pf_s, and here a standard application of Knörrer periodicity implies that (assuming L is generic)

$$D^b(\mathrm{Pf}_s^{sm} \times_{\mathbb{C}^*} L,\, W') \cong D^b(\mathrm{Pf}_s^{sm} \cap \mathbb{P}L^\perp)$$

as in the physical sketch. It's fairly straight-forward to extend this over the singular locus and we prove in Section 4.1 that we have an equivalence

$$\mathsf{Br}(\mathcal{X}^{ss} \times_{\mathbb{C}^*} L,\, W') \cong D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp})$$

where A is the sheaf of non-commutative algebras discussed in Section 1.1. This proves that, for generic L, our non-commutative resolution of $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$ is equivalent to the 'difficult' phase of the dual model.

A more challenging step is to compare $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0),\, W)$ with the category for the 'easy' phase of the same model, *i.e.* $\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W)$. This is a kind of variation-of-GIT process, and we use the idea of 'windows' [Seg11, HL15, BFK12] (which was also used in [ADS15, ST14, Ren15]). What we do is to lift both categories to the ambient stack $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp$, by finding subcategories of $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$ to which they are equivalent. The window for the 'difficult' phase is essentially standard, but the window that we need for the 'easy' phase is not the one provided by general theory, and although it's easy to describe it takes us quite a lot of new calculations to prove that it works (see Section 4.2).

One of our windows is obviously contained in the other, with the direction of containment depending on the dimension of L, and equality in the case $\dim L = sv$. This means that we have an embedding

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0),\, W) \;\hookrightarrow\; \mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W) \tag{1.5}$$

or vice-versa. Putting this together with our previous results, we get an embedding

$$D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp}) \;\hookrightarrow\; D^b(\mathrm{Pf}_q \cap \mathbb{P}L,\, B|_{\mathbb{P}L})$$

or vice-versa. In the critical case $\dim L = sv$ the two categories are equivalent, and in this case both are Calabi-Yau (as stated in Theorem 1.1).

The claims of the preceding paragraph are the most important consequences of HPD, but are some way from being the full statement as Kuznetsov wrote it. No doubt it is possible to prove the full statement directly in our situation, but instead we take a slight digression so that we can bootstrap off some other work of the first author. If we look at (1.5) in the case $\dim L^\perp = 1$, we get an embedding:

$$\mathsf{Br}(\widetilde{Y}) \;\hookrightarrow\; \mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} \mathbb{C},\, W)$$

Notice that $\mathcal{Y}^{ss} \times_{\mathbb{C}^*} \mathbb{C}$ is the total space of a line bundle on $\mathcal{Y}^{ss}$, and that W is essentially a choice of section of the dual line bundle. We can do this in families where we let $L^\perp$ vary in $\wedge^2 V$ – which means letting W vary – and the universal such family gives an embedding:

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(\wedge^2 V\setminus 0),\, W) \;\hookrightarrow\; D^b\big((\mathcal{Y}^{ss}\times_{\mathbb{C}^*}\mathbb{C})\times_{\mathbb{C}^*}(\wedge^2 V\setminus 0),\, W\big)$$

The target space here is the total space of a line bundle over $\mathcal{Y}^{ss}\times\mathbb{P}(\wedge^2 V)$.

It's shown in [Ren17] that given a stack Z with a line bundle $\mathcal{L}$, and an appropriate subcategory of $D^b(Z)$, one can construct a 'tautological' HP dual to that subcategory. It lives in $D^b(\operatorname{Tot}\mathcal{L}^\vee(-1), W)$, where $\mathcal{L}^\vee(-1)$ is a line bundle on $Z\times\mathbb{P}(H^0(\mathcal{L}))$ and W is the tautological superpotential on this space (see Section 2.1). We check that the image of our embedding is exactly the HP dual to $\mathsf{Br}(\mathcal{Y}^{ss})$. But our previous results show that the category being embedded is equivalent to $\mathsf{Br}(\mathcal{X}^{ss})$, so $\mathsf{Br}(\mathcal{Y}^{ss})$ is HP dual to $\mathsf{Br}(\mathcal{X}^{ss})$. Or in other words, $D^b(\mathrm{Pf}_q, B)$ is HP dual to $D^b(\mathrm{Pf}_s, A)$.

2. Technical background

2.1. **Homological projective duality via LG models.** We recall the basic definitions and theorems of HP duality, phrased in terms of LG models as in [Ren17] or [BDF+13]. The original source is [Kuz07], see also [BDF+13, Tho15, Ren15] for further background.

Let Y be an algebraic stack with a line bundle $\mathcal{L}$ such that $V = H^0(Y,\mathcal{L}) \neq 0$. We assume that Y satisfies the hypotheses used in [BFK12], *i.e.* we assume that $Y = \big[X\,/\,G\big]$ for a smooth quasi-projective X and a reductive G.

Let $\mathcal{W}\subseteq D^b(Y)$ be a full subcategory, which is assumed to be admissible and saturated, i.e. every functor $\mathcal{W}\to D^b(\mathbb{C})^{\mathrm{op}}$ and $\mathcal{W}^{\mathrm{op}}\to D^b(\mathbb{C})^{\mathrm{op}}$ is representable. We assume $\mathcal{W}\otimes\mathcal{L} = \mathcal{W}$, and that $\mathcal{W}$ is equipped with a Lefschetz decomposition with respect to $\mathcal{L}$, which means that there is a semiorthogonal decomposition

$$\mathcal{W} = \big\langle \mathcal{A}_0,\ \mathcal{A}_1(1),\ \ldots,\ \mathcal{A}_N(N)\big\rangle,$$

where $\mathcal{A}_0\supseteq\mathcal{A}_1\supseteq\cdots\supseteq\mathcal{A}_N$ are full admissible subcategories, and we're writing $\mathcal{A}_i(i)$ as shorthand for $\mathcal{A}_i\otimes\mathcal{L}^i$.

We can then describe the HP dual of $\mathcal{W}$ in the following way. Let $\widetilde{Y} = \operatorname{Tot}(\mathcal{L}^\vee)$, and let $Z = \big[\widetilde{Y}\times V\,/\,\mathbb{C}^*\big]$, where $\mathbb{C}^*$ acts with weight 1 on the fibres of $\widetilde{Y}\to Y$, and with weight -1 on V. We equip Z with an action of $\mathbb{C}^*_R$ which fixes $\widetilde{Y}$ and acts with weight 2 on V. Any element $v\in V$ describes a section s_v of $\mathcal{L}$, dualising this gives a function $s_v^\vee:\widetilde{Y}\to\mathbb{C}$. Thus there is a canonical $W: Z\to\mathbb{C}$ given by $W(x,v) = s_v^\vee(x)$, and (Z,W) is an LG model.

Let $\pi: Z\to\big[V\,/\,\mathbb{C}^*\big]$ be the projection, and let $Z^{\mathrm{ss}} = \pi^{-1}(\mathbb{P}(V))$. Let $\mathcal{W}^\vee\subset D^b(Z^{\mathrm{ss}}, W)$ be the subcategory consisting of those objects $\mathcal{E}$ such that for all $[v]\in\mathbb{P}V$, the restriction $\mathcal{E}|_{Y\times[v]}$ is contained in $\mathcal{A}_0\subseteq\mathcal{W}$. We say the category $\mathcal{W}^\vee$ is the HP dual of $\mathcal{W}$.

This does not agree with Kuznetsov's terminology in [Kuz07], where he reserves the term 'HP dual' for varieties derived equivalent to a certain subcategory $\mathcal{C}\subseteq D^b(\mathcal{H})$, where $\mathcal{H}\subseteq Y\times\mathbb{P}(V)$ is the universal hyperplane section. Using the techniques of [BDF+13], it is shown in [Ren17] that $\mathcal{W}^\vee\cong\mathcal{C}$. Furthermore, the main theorem of HP duality holds for the appropriate linear sections of $\mathcal{W}^\vee$, thus justifying the name HP dual for $\mathcal{W}$.

Let N' be the minimal integer such that $\mathcal{A}_{N'}\neq\mathcal{A}_0$, and let $M = \dim V - N' - 1$.

Proposition 2.1 ([Kuz07, Thm. 6.3], [Ren17])**.** *The category $\mathcal{W}^\vee$ admits a Lefschetz decomposition*

$$\mathcal{W}^\vee = \big\langle \mathcal{B}_{-M}(-M),\ \ldots,\ \mathcal{B}_1(1),\ \mathcal{B}_0\big\rangle.$$

For any linear subspace $L \subset V^\vee$, with orthogonal $L^\perp \subset V$, we define the base-changed categories $\mathcal{W}_L$ and $\mathcal{W}^\vee_{L^\perp}$ as follows.

Definition 2.2. *Let $\mathcal{W}_L \subset D^b(\widetilde{Y}^{\rm ss} \times_{\mathbb{C}^*} L^\perp, W)$ be the inverse image of $\mathcal{W}$ under the restriction functor*

$$D^b(\widetilde{Y}^{\rm ss} \times_{\mathbb{C}^*} L^\perp, W) \longrightarrow D^b(\widetilde{Y}^{\rm ss} \times_{\mathbb{C}^*} 0) = D^b(Y).$$

Let $\mathcal{W}^\vee_{L^\perp} \subset D^b(\widetilde{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$ be the full subcategory of objects $\mathcal{E}$ such that for each $p \in \mathbb{P}(L^\perp)$, the restriction of $\mathcal{E}$ to $D^b(Y \times p, 0)$ is contained in $\mathcal{A}_0 \subseteq \mathcal{W}$.

Let $l = \dim L$, and $l' = \dim L^\perp = \dim V - l$. It is not hard to show that each of the categories $\mathcal{A}_{l'}(l'), \ldots, \mathcal{A}_N(N)$ are mapped fully faithfully to $\mathcal{W}_L$, and that the semiorthogonality relations between them are preserved, *i.e.* that there are no maps from objects in $\mathcal{A}_i(i)|_{\mathcal{W}_L}$ to objects $\mathcal{A}_j(j)|_{\mathcal{W}_L}$ if $i > j$. Similarly the $\mathcal{B}_{-M}(-M), \ldots, \mathcal{B}_{-l}(-l)$ are preserved upon restriction to $\mathcal{W}^\vee_{L^\perp}$, and their semiorthogonality preserved. Far more non-trivial is the claim that the orthogonal complements of these sequences of categories agree:

Theorem 2.3 ([Kuz07, Ren17])**.** *There exist semiorthogonal decompositions*

$$\mathcal{W}_L = \big\langle \mathcal{C}_L,\ \mathcal{A}_{l'}(l'),\ \ldots,\ \mathcal{A}_N(N) \big\rangle$$

and

$$\mathcal{W}^\vee_{L^\perp} = \big\langle \mathcal{B}_{-M}(-M),\ \ldots,\ \mathcal{B}_{-l}(-l),\ \mathcal{C}_L \big\rangle.$$

If Y_L has the expected dimension, then $D^b(Y_L) \cong \mathcal{W}_L$, by Knörrer periodicity. Likewise, if X is a variety over $\mathbb{P}(V)$ such that $D^b(X) \cong \mathcal{W}^\vee$, where the equivalence is induced by an FM kernel in $D^b(X \times_{\mathbb{P}(V)} Z, W)$, and if $X_{L^\perp}$ has the expected dimension, then Knörrer periodicity implies that $D^b(X_{L^\perp}) \cong \mathcal{W}^\vee_{L^\perp}$. Using these observations, we obtain Kuznetsov's original [Kuz07, Thm. 6.3].

2.2. **Matrix factorisation categories.** As is clear from the introduction, our proofs and results involve (graded) matrix factorisation categories. We will here restrict ourselves to stating precisely what we mean by these categories – for further background on definitions and tools, see [ADS15, Ren15, BDF+13].

By a *Landau–Ginzburg model* we mean the data of

- a stack $\mathcal{X}$ of the form $[X/(G \times \mathbb{C}^*)]$, where X is a smooth, quasi-projective variety, G is a reductive group.
- a function $W : X \to \mathbb{C}$.

We denote the distinguished $\mathbb{C}^*$-factor of the group by $\mathbb{C}^*_R$, and refer to it as the R-charge.

This set of data is subject to some restrictions:

- The function W is G-invariant and has degree 2 with respect to the $\mathbb{C}^*_R$-action.
- The element $-1 \in \mathbb{C}^*_R$ acts trivially on $[X/G]$, i.e. there exists a $g \in G$ such that $(g, -1)$ acts trivially on X.

Given this data, by work of Positselski and Orlov [Pos11, EP15, Orl12] one can define a category of matrix factorisations $D^b(\mathcal{X}, W)$. Let's say briefly what this category looks like.

There is a natural map $\mathcal{X} \to [\mathrm{pt}/\mathbb{C}^*_R]$, we denote the pullback of the standard line bundle on $[\mathrm{pt}/\mathbb{C}^*_R]$ via this map by $\mathcal{O}_\mathcal{X}[1]$. More generally, for any sheaf $\mathcal{E}$ on $\mathcal{X}$, we write $\mathcal{E}[1]$ for $\mathcal{E} \otimes \mathcal{O}_\mathcal{X}[1]$. The objects of $D^b(\mathcal{X}, W)$ can then be described as curved dg sheaves $(\mathcal{E}, d_\mathcal{E})$ on $\mathcal{X}$, meaning the data of

- A coherent sheaf $\mathcal{E}$ on $\mathcal{X}$
- A 'twisted differential' $d_\mathcal{E} : \mathcal{E} \to \mathcal{E}[1]$ such that $d^2_\mathcal{E} = W \otimes 1_\mathcal{E} : \mathcal{E} \to \mathcal{E}[2]$.

Given two curved dg sheaves $(\mathcal{E}, d_{\mathcal{E}})$ and $(\mathcal{F}, d_{\mathcal{F}})$, the sheaf $\mathcal{H}om(\mathcal{E}, \mathcal{F})$ on $\mathcal{X}$ inherits a differential, so becomes a dg sheaf, and one can take its cohomology to turn the set of curved dg sheaves into a category. Just as for the ordinary derived category, this definition is too naive, and defining the morphism spaces and the triangulated category structure on $D^b(\mathcal{X}, W)$ requires that we take the Verdier quotient by some subcategory of 'acyclic' curved dg sheaves.

Without going into details of this, let's just mention that given two objects $(\mathcal{E}, d_{\mathcal{E}})$ and $(\mathcal{F}, d_{\mathcal{F}})$, the morphism spaces $\mathrm{Hom}(\mathcal{E}, \mathcal{F})$ can be computed as follows. We can find a curved dg sheaf $(\mathcal{E}', d_{\mathcal{E}'})$ such that $\mathcal{E}'$ is locally free, together with a quasi-isomorphism $\mathcal{E}' \to \mathcal{E}$. As mentioned above, $\mathcal{H}om(\mathcal{E}', \mathcal{F})$ is a complex, and we can compute

$$\mathrm{RHom}(\mathcal{E}, \mathcal{F}) \cong R\Gamma(\mathcal{H}om(\mathcal{E}', \mathcal{F})).$$

A morphism between LG models is a map $f : \mathcal{X} \to \mathcal{Y}$ such that $W_{\mathcal{X}} = W_{\mathcal{Y}} \circ f$ and such that $f^*(\mathcal{O}_{\mathcal{Y}}[1]) \cong \mathcal{O}_{\mathcal{X}}[1]$. The usual ensemble of (derived) functors exists, e.g. if f is proper, we have a functor $f_* : D^b(\mathcal{X}, W_{\mathcal{X}}) \to D^b(\mathcal{Y}, W_{\mathcal{Y}})$, and in general we have a pull-back functor $f^* : D^b(\mathcal{Y}, W_{\mathcal{Y}}) \to D^b(\mathcal{X}, W_{\mathcal{X}})$. Given an object $\mathcal{E} \in D^b(\mathcal{X}, W)$, we get a functor $- \otimes \mathcal{E} : D^b(\mathcal{X}, W') \to D^b(\mathcal{X}, W' + W)$.

Composing these functors, the formalism of Fourier–Mukai kernels can be used: Given LG models $(\mathcal{X}, W_{\mathcal{X}})$ and $(\mathcal{Y}, W_{\mathcal{Y}})$ and a kernel object $\mathcal{E} \in D^b(\mathcal{X} \times \mathcal{Y}, W_{\mathcal{Y}} - W_{\mathcal{X}})$, we get a functor $(\pi_{\mathcal{Y}})_*(\mathcal{E} \otimes \pi_{\mathcal{X}}^*(-))$, where $\pi_{\mathcal{X}}$ and $\pi_{\mathcal{Y}}$ are the projections from $\mathcal{X} \times \mathcal{Y}$, assuming the pushforward along $\pi_{\mathcal{Y}}$ exists.

Remark 2.4. From this point on, we will abuse notation and write $D^b([X/G], W)$ for what is here denoted $D^b([X/G \times \mathbb{C}^*_R], W)$ – i.e. we will leave the $\mathbb{C}^*_R$-action implicit. This notational choice reflects the idea that it is best to think of $[X/G]$ as the geometric object underlying these categories, as illustrated by the fact [BDF+13, Prop. 2.1.6] that if $W = 0$ and $\mathbb{C}^*_R$ acts trivially, then $D^b([X/G \times \mathbb{C}^*_R], 0)$ is equivalent to the usual derived category $D^b([X/G])$.

Remark 2.5. We should highlight one piece of our terminology which may differ from other authors: we reserve the term *matrix factorization* for a curved dg-sheaf $(\mathcal{E}, d_{\mathcal{E}})$ where $\mathcal{E}$ is actually a finite-rank vector bundle. Every object in $D^b(\mathcal{X}, W)$ is equivalent to a matrix factorization; for us this statement is part of the definition of the category (though really one should define $D^b(\mathcal{X}, W)$ as some category of compact objects and then prove the statement).

2.3. **Minimal models.** In the introduction, we defined the subcategories $\mathsf{Br}(\mathcal{X}) \subseteq D^b(\mathcal{X})$ and the matrix factorisation versions $\mathsf{Br}(\mathcal{X}, W) \subseteq D^b(\mathcal{X}, W)$ by choosing some set Y of $\mathrm{GSp}(S)$-irreps, and declaring $\mathcal{E}$ to lie in the brane subcategory if it can be represented by a complex (or matrix factorisation), whose underlying sheaf is a shifted sum of vector bundles corresponding to $\mathrm{GSp}(S)$-irreps in Y. This definition is simple, but it can be hard to check if a given $\mathcal{E}$ is of this form, since it requires replacing $\mathcal{E}$ with a locally free model.

An alternative definition can be given by restricting $\mathcal{E}$ to $0 \in \mathrm{Hom}(S, V)$, noting that $\mathcal{E}|_0$ decomposes as a shifted sum of $\mathrm{GSp}(S)$-representations, and then declare $\mathcal{E}$ to lie in the brane subcategory if all $\mathrm{GSp}(S)$-representations appearing in $\mathcal{E}|_0$ lie in Y (see Section 3.1 below). It is easy to see that a brane object in the former sense is also a brane object in the latter; Lemma 2.6 below implies that in fact the two definitions are equivalent.

Let G be a reductive group, and V a G-representation. Add a linear R-charge $\mathbb{C}^*_R$ action on V and a G-invariant superpotential W.

Lemma 2.6. *Assume that the combined group $G \times \mathbb{C}^*_R$ on V includes a central 1-parameter-subgroup σ which acts as some non-zero power of the usual dilation* *action on V. Then any object $\mathcal{E} \in D^b_G(V, W)$ is equivalent to a matrix factorization (E, d_E) such that $d_E|_0 = 0$, and E is the vector bundle associated to the G-representation $h_\bullet(\mathcal{E}|_0)$.*

If $W = 0$ this is the standard theory of minimal resolutions, e.g. [Wey03, Prop. 5.4.2]. Without the G-action, the statement is essentially [KR08, Prop. 7]. The following proof works whether $W = 0$ or not.

Proof. Since $\mathcal{O}_V$ is graded by σ (either non-positively or non-negatively) with $(\mathcal{O}_V)_0 = \mathbb{C}$, it's elementary that any $(G \times \mathbb{C}^*_R)$-equivariant vector bundle on V must be the bundle associated to some representation; see e.g. [Wey03]. By definition, any object $\mathcal{E} \in D^b_G(V, W)$ is equivalent to a matrix factorization (E, d_E), and then E is the vector bundle associated to the $(G \times \mathbb{C}^*_R)$-representation $E|_0$. If $d_E|_0 = 0$, then it is immediate that E is the vector bundle associated to the G-representation $h_\bullet(\mathcal{E}|_0)$.

Suppose that $d_E|_0 \neq 0$. We claim that then E contains a contractible sub-object whose underlying sheaf is a subbundle of E. Quotienting by this subbundle produces an equivalent matrix factorization of smaller rank, so applying this recursively results in a model of the required form.

To prove the claim, we consider $(E|_0, d_E|_0)$. This is a bounded complex of G-representations, which we can decompose into graded irreps. If $d_E|_0 \neq 0$, there must be some component of $d_E|_0$ of the form

$$U \xrightarrow{1_U} U[-1]$$

where U is a (shift of an) irrep. Then E contains two associated subbundles, and there is a component of d_E

$$\delta : U \longrightarrow U[-1]$$

mapping between them. This δ is a $(G \times \mathbb{C}^*_R)$-invariant element of $\mathrm{Hom}_V(U, U)$, reducing to 1_U at the origin. In fact δ must be constant – the 1-PS σ is central so it acts on U as scalar multiple of the identity, and then the only σ-invariant elements of $\mathrm{Hom}_V(U, U)$ are constant. So $\delta = 1_U$.

Now let $\iota : U \hookrightarrow E$ denote the inclusion of the first subbundle. The map $d_E \circ \iota : U[-1] \to E$ is $\mathbb{C}^*_R$-invariant, and one of its components is $\delta \circ \iota$ which is the inclusion of the second subbundle. Hence the map $(\iota, d_E \circ \iota) : U \oplus U[-1] \to E$ has full rank at all points so it's the inclusion of a subbundle. This map embeds the contractible matrix factorization

$$U \underset{W 1_U}{\overset{1_U}{\rightleftarrows}} U[-1]$$

as a subobject of (E, d_E). □

Remark 2.7. The proof still works if V is just a cone instead of a vector space, *i.e.* $\mathcal{O}_V$ is a non-negatively graded (by σ) ring with $(\mathcal{O}_V)_0 = \mathbb{C}$. In fact it still works if $(\mathcal{O}_V)_0$ is a local ring, and we replace 'restrict to the origin' with 'restrict to the unique closed point of $\mathrm{Spec}\,(\mathcal{O}_V)_0$'. The only difference in the proof is that now the map δ need not be exactly 1_U, but it must be of the form $t1_U$ for some unit $t \in (\mathcal{O}_V)_0$ (the rest of the argument is identical). We will need this more general case in the proof of Lemma 3.21.

Corollary 2.8. *In the situation of Lemma 2.6, suppose that $\mathcal{E} \in D^b_G(V, W)$ is such that $\mathcal{E}|_0 \cong 0$. Then $\mathcal{E} = 0$.*

3. The affine duality

3.1. **The statement.** Let's recall our notation from the introduction. We fix a vector space V of dimension v, and two symplectic vector spaces S and Q of dimensions $2s$ and $2q$ respectively, where $v = 2s + 2q + 1$. We also fix isomorphisms:

$$\mathrm{GSp}(S)/\mathrm{Sp}(S) \cong \mathbb{C}^* \cong \mathrm{GSp}(Q)/\mathrm{Sp}(Q)$$

We let $\mathcal{X}$ and $\mathcal{Y}$ denote the following stacks:

$$\mathcal{X} = \big[\mathrm{Hom}(S, V) \,/\, \mathrm{GSp}(S)\big] \qquad \text{and} \qquad \mathcal{Y} = \big[\mathrm{Hom}(V, Q) \,/\, \mathrm{GSp}(Q)\big]$$

Now we pick a subspace $L \subset \wedge^2 V^\vee$ and its annihilator $L^\perp \subset \wedge^2 V$. We let $\mathbb{C}^*$ act with weight -1 on $\wedge^2 V$ and, dually, with weight 1 on $\wedge^2 V^\vee$. Then we can form the products

$$\mathcal{X} \times_{\mathbb{C}^*} L := \big[\mathrm{Hom}(S, V) \times L \,/\, \mathrm{GSp}(S)\big]$$

and:

$$\mathcal{Y} \times_{\mathbb{C}^*} L^\perp := \big[\mathrm{Hom}(V, Q) \times L^\perp \,/\, \mathrm{GSp}(Q)\big]$$

We let W' and W denote the tautological invariant functions ('superpotentials') on these two stacks, see (1.4) and (1.3). We define an R-charge on both stacks by letting $\mathbb{C}^*_R$ act with weight zero on $\mathrm{Hom}(S, V)$ and $\wedge^2 V$, with weight 1 on $\mathrm{Hom}(V, Q)$, and with weight 2 on $\wedge^2 V^\vee$. Then both W' and W have R-charge 2, and both pairs

$$(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \qquad \text{and} \qquad (\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$$

are Landau-Ginzburg B-models.

Both stacks can be mapped ($\mathbb{C}^*_R$-equivariantly) to the stack $[\wedge^2 V/\mathbb{C}^*]$ - the first by projection onto $\mathcal{X}$ and then applying the evident map from $\mathrm{Hom}(S, V)$ to $\wedge^2 V$, and the second simply by projection onto $L^\perp$ and then inclusion. In this section (Section 3) we will treat both of them as being defined over this base. In particular, when we discuss the categories of matrix factorizations

$$D^b(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \qquad \text{and} \qquad D^b(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$$

we mean that both are enriched over graded $\mathcal{O}_{\wedge^2 V}$-modules – that is, the morphisms in these categories are obtained by taking the morphism sheaves and pushing them down to $[\wedge^2 V/\mathbb{C}^*]$ instead of taking their absolute global sections.

Remark 3.1. It would be more symmetric to define the R-charge on $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp$ to have weight zero on $\mathrm{Hom}(V, Q)$ and weight 2 on $L^\perp$. If we work with the 'absolute' category $D^b(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$ then it does not matter which choice we make, because the two actions differ by the diagonal 1-parameter subgroup $\Delta : \mathbb{C}^* \hookrightarrow \mathrm{GSp}(Q)$ (see the proof of Proposition 3.10 below). However it is simplest if we keep the map to $[\wedge^2 V/\mathbb{C}^*]$ explicitly $\mathbb{C}^*_R$-equivarant.

In Section 4 we will stop working relative to this base, and make use of this other choice of R-charge.

Now we define/recall our 'B-brane' subcategories. Recall that $Y_{s,q}$ denotes the set of Young diagrams of height at most s and width at most q, and we'll use the same notation for the corresponding set of irreps of $\mathrm{Sp}(S)$. Similarly $Y_{q,s}$ denotes the transposed set of Young diagrams, or the corresponding set of irreps of $\mathrm{Sp}(Q)$. Given any object

$$\mathcal{E} \in D^b(\mathcal{X} \times_{\mathbb{C}^*} L,\, W')$$

we can consider its restriction $\mathcal{E}|_0$ to the origin, which is a chain-complex of representations of $\mathrm{GSp}(S)$. The homology $h_\bullet(\mathcal{E}|_0)$ of this is a graded $\mathrm{GSp}(S)$-representation, which we may view just as a representation of $\mathrm{Sp}(S)$. We define

$$\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \subset D^b(\mathcal{X} \times_{\mathbb{C}^*} L,\, W')$$

to be the full subcategory of all objects $\mathcal{E}$ satisying the following 'grade-restriction-rule':

$$\text{all irreps of } \mathrm{Sp}(S) \text{ occuring in } h_\bullet(\mathcal{E}|_0) \text{ lie in the set } Y_{s,q}$$

Similarly

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W) \;\subset\; D^b(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)$$

is the full subcategory of all objects $\mathcal{F}$ satisying the grade-restriction-rule:

$$\text{all irreps of } \mathrm{Sp}(Q) \text{ occuring in } h_\bullet(\mathcal{F}|_0) \text{ lie in the set } Y_{q,s}$$

Note that both subcategories are obviously triangulated because the grade-restriction-rule is preserved under taking mapping cones.

Let's connect these definitions to the ones used in the introduction. Any irrep of $\mathrm{Sp}(S)$ lifts to $\mathbb{Z}$-many irreps of $\mathrm{GSp}(S)$, differing by characters, so the set $Y_{s,q}$ determines an infinite set of irreps of $\mathrm{GSp}(S)$. To any such irrep there is an associated vector bundle on $\mathcal{X}\times_{\mathbb{C}^*}L$. By Lemma 2.6, an object $\mathcal{E}\in D^b(\mathcal{X}\times_{\mathbb{C}^*}L,\, W')$ lies in the subcategory $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W')$ if and only if it can be represented as a matrix factorization whose underlying vector bundle is a direct sum of the bundles coming from $Y_{s,q}$. In particular in the case $L=0$, the statement is that $\mathsf{Br}(\mathcal{X})\subset D^b(\mathcal{X})$ is the subcategory generated by this infinite set of vector bundles. Both versions of the definition will be useful in the course of our proofs.

Now we can state our interpretation of Hori's duality for B-branes.

Theorem 3.2. *For any $L\subset \wedge^2 V^\vee$, we have an equivalence*

$$\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W') \;\xrightarrow{\sim}\; \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)$$

of categories over $[\wedge^2 V/\mathbb{C}^*]$.

When we say that our equivalence is 'over $[\wedge^2 V/\mathbb{C}^*]$' we mean simply that it is given by a Fourier–Mukai kernel on the relative product over this base. It follows that we may restrict our categories and our kernel to open substacks in the base, and still get an equivalence.

Of course we could have chosen to work over $[\wedge^2 V^\vee/\mathbb{C}^*]$ instead, in which case we'd get an equivalence relative to that base. This might suggest that we should really be able to define our kernel over $[\wedge^2 V\times\wedge^2 V^\vee/\mathbb{C}^*]$, but this appears not to be the case.[1]

If we forget the $\mathbb{C}^*$-action, we obtain an equivalence

$$\mathsf{Br}(\widetilde{\mathcal{X}}\times L,\, W') \;\cong\; \mathsf{Br}(\widetilde{\mathcal{Y}}\times L^\perp,\, W)$$

which expresses the duality between symplectic GLSMs. Here $\mathsf{Br}(\widetilde{\mathcal{X}}\times L,\, W')$ and $\mathsf{Br}(\widetilde{\mathcal{Y}}\times L^\perp,\, W)$ are defined to be the images of $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W')$ and $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)$ under the pull-back functor; they can also be defined to be the subcategories consisting of matrix factorizations built from the vector bundles associated to the sets $Y_{s,q}$ and $Y_{q,s}$, as in Section 1.1. However they cannot be defined by a grade-restriction-rule at the origin, because now the group action has closed orbits other than the origin and Lemma 2.6 doesn't apply.

The remainder of Section 3 will be devoted to the proof of Theorem 3.2. As sketched in Section 1.3, most of the work will be in proving the extreme case $L=0$, from there we will deduce the general case quite easily.

[1] This same phenomenon can be seen in the simplest examples of Knörrer periodicity.

3.2. **Relating the brane subcategories to (curved) algebras.** On the stack $\widetilde{\mathcal{X}}$ we have the subcategory $\mathsf{Br}(\widetilde{\mathcal{X}})$, generated by the vector bundles associated to irreps in $Y_{s,q}$. We discussed in Section 1.1 how we may also view this as the derived category of an algebra: we take the tilting bundle $\widetilde{T} = \bigoplus_{\gamma\in Y_{s,q}} \mathbb{S}^{\langle\gamma\rangle}S$ and consider its endomorphism algebra $\widetilde{A} = \mathrm{End}_{\widetilde{\mathcal{X}}}(\widetilde{T})$, then $\mathsf{Br}(\widetilde{\mathcal{X}})$ is equivalent to the derived category of $\widetilde{A}$. We'll now do a similar thing for the category $\mathsf{Br}(\mathcal{X})$.

Every $\mathrm{Sp}(S)$-irrep $\mathbb{S}^{\langle\gamma\rangle}S$ has a standard lift to a $\mathrm{GSp}(S)$-irrep, determined by the property that the diagonal subgroup $\Delta : \mathbb{C}^* \to \mathrm{GSp}(S)$ act with weight $\sum \gamma_i$. The group $\mathrm{GSp}(S)$ has a character $\langle\omega_S\rangle$ given by its action on the line spanned by ω_S in $\wedge^2 S$, and $\langle\omega_S\rangle$ generates the character lattice. The higher-dimensional irreps of $\mathrm{GSp}(S)$ are then given by $\mathbb{S}^{\langle\gamma\rangle}S\otimes\langle\omega_S\rangle^p$, for γ a dominant weight of $\mathrm{Sp}(S)$ and $p\in\mathbb{Z}$, so each $\mathrm{Sp}(S)$-irrep can be lifted to a $\mathrm{GSp}(S)$-irrep in $\mathbb{Z}$-many ways. Correspondingly, the vector bundle $\mathbb{S}^{\langle\gamma\rangle}S$ on $\widetilde{\mathcal{X}}$ can be lifted to a vector bundle on $\mathcal{X}$ in $\mathbb{Z}$-many ways, differing by powers of the line-bundles $\langle\omega_S\rangle^p$. If we choose an integer p_γ for each $\gamma\in Y_{s,q}$, we get a lift

$$T = \bigoplus_{\gamma\in Y_{s,q}} \mathbb{S}^{\langle\gamma\rangle}S\otimes\langle\omega_S\rangle^{p_\gamma} \;\in D^b(\mathcal{X}). \tag{3.3}$$

of the vector bundle $\widetilde{T}$. Let A be the algebra:

$$A = \mathrm{End}_{\mathcal{X}}(T)$$

By our definitions we only take $\mathrm{Sp}(S)$-invariants when forming A instead of $\mathrm{GSp}(S)$-invariants, so A is a graded algebra over the ring of functions on $\wedge^2 V$, and it's supported on $\widetilde{\mathrm{Pf}_s}$. We can also think of it as a quiver algebra (with relations), where the underlying quiver has vertices indexed by the set $Y_{s,q}$.

Remark 3.4. This definition of A is slightly ambiguous because it depends on our choices p_γ for the lift of each summand. However the different choices only affect the grading on A, and all the resulting algebras are trivially Morita equivalent.

Lemma 3.5. *We have an equivalence:*

$$Hom(T,-) : \mathsf{Br}(\mathcal{X}) \xrightarrow{\sim} D^b(A)$$

For each $\gamma\in Y_{s,q}$, this equivalence takes the vector bundle $\mathbb{S}^{\langle\gamma\rangle}S$ to the projective module at the corresponding vertex.

Proof. The functors $F := \mathrm{Hom}(T,-)$ and $F^* = T\otimes_A -T$ are an adjoint pair of functors between $D(\mathcal{X})$ and $D(A)$ (the unbounded derived categories), and F^* is fully faithful.

Since A has finite global dimension as an ungraded algebra by [ŠVdB15, Thm. 1.4.2], it has finite global dimension as a graded algebra by [NvO82, Cor. I.2.7].

Hence $D^b(A)$ equals the smallest thick subcategory of $D(A)$ containing $\{A(i)\}_{i\in\mathbb{Z}}$. Now $F^*(A(i)) = T(i)\in\mathsf{Br}(\mathcal{X})$, and so since $\mathsf{Br}(\mathcal{X})$ is a thick subcategory of $D^b(\mathcal{X})$ and F^* preserves direct sums, it follows that F^* takes $D^b(A)$ into $\mathsf{Br}(\mathcal{X})$.

Let $0\neq\mathcal{E}\in\mathsf{Br}(\mathcal{X})$. We may represent $\mathcal{E}$ by a minimal complex

$$0\to\mathcal{E}_i\to\cdots\to\mathcal{E}_j\to 0$$

as in Lemma 2.6. Let $\mathbb{S}^{\langle\delta,k\rangle}S$ be an irreducible summand of $\mathcal{E}_j$. The inclusion map $\mathbb{S}^{\langle\delta,k\rangle}S\to\mathcal{E}_j$ does not factor through $\mathcal{E}_{j-1}\to\mathcal{E}_j$, by minimality of the resolution, hence we get a non-zero map $\mathbb{S}^{\langle\delta,k\rangle}S\to\mathcal{E}[j]$. It follows that $\mathrm{Hom}(T,\mathcal{E}[j])\neq 0$, and so $F(\mathcal{E})\neq 0$. Thus $\ker F\cap\mathsf{Br}(\mathcal{X}) = \{0\}$, and by [Kuz07, Thm. 3.3], $F^* : D^b(A)\to\mathsf{Br}(\mathcal{X})$ is then essentially surjective. □

The above equivalence is linear over our base $[\wedge^2 V/\mathbb{C}^*]$, so it makes sense to base-change it to open subsets. For example, suppose we delete the origin in $\wedge^2 V$. The restriction of $\mathcal{X}$ to this locus is the open substack $\mathcal{X}^{ss} \subset \mathcal{X}$, and we define $\mathsf{Br}(\mathcal{X}^{ss})$ to be the full subcategory of $D^b(\mathcal{X}^{ss})$ generated by the image of $\mathsf{Br}(\mathcal{X})$ under restriction. Equivalently, this the subcategory generated by the (infinite) set of vector bundles corresponding to the set $Y_{s,q}$.

If we delete the origin from the stack $[\widetilde{\mathrm{Pf}_s}/\mathbb{C}^*]$ we get the projective variety Pf_s. The algebra A restricts to give a sheaf of algebras on Pf_s, and have an associated abelian category of coherent modules over it, together with the various flavours of derived category.

Corollary 3.6. *We have an equivalence:*

$$\mathrm{Hom}(T|_{\mathcal{X}^{ss}}) : \mathsf{Br}(\mathcal{X}^{ss}) \xrightarrow{\sim} D^b(\mathrm{Pf}_s, A)$$

Proof. Let $F_{ss} = \mathrm{Hom}(T|_{\mathcal{X}^{ss}}, -) : D(\mathcal{X}^{ss}) \to D(\mathrm{Pf}_s, A)$, and let F^*_{ss} be the left adjoint. The functor F^*_{ss} is automatically fully faithful.

If $\mathcal{E} \in \mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$, then there exists an $\widetilde{\mathcal{E}} \in \mathsf{Br}(\mathcal{X})$ restricting to $\mathcal{E}$. Since F is linear over $[\wedge^2 V/\mathbb{C}^*]$ we find

$$\mathcal{E} = F^*_{\mathrm{ss}}(F(\widetilde{\mathcal{E}})|_{\mathrm{Pf}_s}),$$

and so $F^*_{\mathrm{ss}} : D^b(\mathrm{Pf}_s, A) \to \mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ is essentially surjective. □

Remark 3.7. This same proof works if we restrict to other ($\mathbb{C}^*$-invariant) open subsets in $\wedge^2 V$.

Now fix $L \subset \wedge^2 V^\vee$ and consider our Landau–Ginzburg B-model $(\mathcal{X} \times_{\mathbb{C}^*} L, W')$. We can adapt the equivalence of $\mathsf{Br}(\mathcal{X})$ with $D^b(A)$ to give another description of the category $\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L, W')$.

Let's start by forgetting W', and just tensoring the previous lemma by $\mathcal{O}_L$. The vector bundle T on $\mathcal{X}$ can be pulled up to give a vector bundle on $\mathcal{X} \times_{\mathbb{C}^*} L$ which we'll continue to denote by T. and we have a functor:

$$\mathrm{Hom}(T, -) : D^b(\mathcal{X} \times_{\mathbb{C}^*} L) \longrightarrow D^b(A \otimes \mathcal{O}_L)$$

The algebra $A \otimes \mathcal{O}_L$ has an obvious set of projective modules, obtained by taking the 'vertex projective' A-modules and tensoring with $\mathcal{O}_L$. The adjoint functor

$$T \otimes - : D^b(A \otimes \mathcal{O}_L) \longrightarrow D^b(\mathcal{X} \times_{\mathbb{C}^*} L)$$

sends each projective module to the corresponding vector bundle, *i.e.* the corresponding summand of T). It is an embedding with image $\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L)$.

Now we introduce the superpotential W'. We can view W' as a central element of the algebra $A \otimes \mathcal{O}_L$, since this an algebra over the ring of functions on $\wedge^2 V \times L$. So the pair $(A \otimes \mathcal{O}_L, W')$ is a 'curved algebra', and there is an associated category

$$D^b(A \otimes \mathcal{O}_L, W')$$

of curved dg-modules. By definition every object in this category is equivalent to a curved dg-module whose underlying module is projective.

Lemma 3.8. *We have a functor*

$$\mathrm{Hom}(T, -) : D^b(\mathcal{X} \times_{\mathbb{C}^*} L, W') \longrightarrow D^b(A \otimes \mathcal{O}_L, W')$$

with a left adjoint $T \otimes -$. The adjoint is an embedding and gives an equivalence:

$$T \otimes - : D^b(A \otimes \mathcal{O}_L, W') \xrightarrow{\sim} \mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L, W')$$

In particular the subcategory $\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L, W')$ is right-admissible.

Proof. If we have an object $(\mathcal{E}, d_\mathcal{E}) \in D^b(\mathcal{X} \times_{\mathbb{C}^*} L,\, W')$ then applying the functor $\mathrm{Hom}(T, -)$ to $\mathcal{E}$ gives an (R-charge equivariant) $A \otimes \mathcal{O}_L$-module. Under this functor the endomorphism $d_\mathcal{E}$ maps to a endomorphism of the module $\mathrm{Hom}(T, \mathcal{E})$, which squares to W' – note that the functor $\mathrm{Hom}(T, -)$ is exact so there are no 'up-to-homotopy' complications here. We claim that this defines a functor

$$\mathrm{Hom}(T, -) : D^b(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \longrightarrow D^b(A \otimes \mathcal{O}_L, W')$$

The technical issue in this claim is to check that the functor does indeed land in $D^b(A \otimes \mathcal{O}_L, W')$ and not some larger category of curved dg-modules. However, $\mathrm{Hom}(R, \mathcal{E})$ will be a finitely-generated $A \otimes \mathcal{O}_L$ module so it has a finite projective resolution (this algebra has finite global dimension, since A does). Then the perturbation technique of [Seg11, Lemma 3.6] implies that the curved dg-module $\mathrm{Hom}(T, \mathcal{E})$ does lie in $D^b(A \otimes \mathcal{O}_L, W')$.

Now we claim that the only projective $A \otimes \mathcal{O}_L$-modules are the obvious ones corresponding to the vertex projective A-modules. To see this observe that $A \otimes \mathcal{O}_L$ is bi-graded; one grading is by the diagonal $\Delta \subset \mathrm{GSp}(S)$ and the other is the R-charge. If we collapse to a single grading appropriately then it becomes non-negatively graded, with its degree-zero part the semi-simple algebra $\mathbb{C}^{Y_{s,q}}$. Then the claim follows by the graded Nakayama lemma.

Therefore any object in $D^b(A \otimes \mathcal{O}_L, W')$ is equivalent to a curved dg-module built from these projective modules. The adjoint functor $T \otimes -$ identifies this category of curved dg-modules with the category of matrix factorizations on $\mathcal{X} \times_{\mathbb{C}^*} L$ built from the summands of T (this is an equivalence even at the chain level). So $T \otimes -$ gives an equivalence between $D^b(A \otimes \mathcal{O}_L, W')$ and $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L,\, W')$.

The final statement of the lemma follows formally. □

We shall see later in Section 4.1 that if L is generic and its dimension is not too big, then there is a third way to describe the category $\mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L,\, W')$, by using Knörrer periodicity to remove the L directions entirely.

The fact that A (or $\mathsf{Br}(\mathcal{X})$) is a non-commutative *crepant* resolution of $\widetilde{\mathrm{Pf}}_s$ is reflected in the following partial Serre duality statement.

Proposition 3.9. *Let $\mathcal{E}, \mathcal{F} \in \mathsf{Br}(\mathcal{X})$ with $\mathcal{F}|_{\mathcal{X}^{ss}} = 0$. Then*

$$\mathrm{Hom}(\mathcal{E}, \mathcal{F}) = \mathrm{Hom}(\mathcal{F}, \mathcal{E} \otimes (\det S)^v [\dim \widetilde{\mathrm{Pf}}_s])^\vee.$$

Proof. The singular variety $\widetilde{\mathrm{Pf}}_s$ is Gorenstein, and the pull-back of $\omega_{\widetilde{\mathrm{Pf}}_s}$ to $\mathcal{X}$ is $(\det S)^v$. Applying the equivalence $\mathsf{Br}(\mathcal{X}) \cong D^b(A)$, the statement is a graded version of [VdB04, Lemma 6.4.1]. □

We have precisely analogous results on the $\mathcal{Y}$ side. We have a graded algebra B, defined as the endomorphisms of a vector bundle on $\mathcal{Y}$, and an equivalence:

$$\mathsf{Br}(\mathcal{Y}) \xrightarrow{\sim} D^b(B)$$

Then for any $L^\perp$ we have a curved algebra $(B \otimes \mathcal{O}_{L^\perp}, W)$, and an equivalence:

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W) \xrightarrow{\sim} D^b(B \otimes \mathcal{O}_{L^\perp}, W)$$

In particular the subcategory $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)$ is right-admissible.

There is one numerical difference on the $\mathcal{Y}$ side which we need to highlight: the non-trivial R-charge affects the Serre functor.

Proposition 3.10. *Let $\mathcal{E}, \mathcal{F} \in \mathsf{Br}(\mathcal{Y})$ with $\mathcal{F}|_{\mathcal{Y}^{ss}} = 0$. Then:*

$$\mathrm{Hom}(\mathcal{E}, \mathcal{F}) = \mathrm{Hom}(\mathcal{F}, \mathcal{E} \otimes (\det Q)^{-v} [\dim \widetilde{\mathrm{Pf}}_q - 2qv])^\vee$$

We observed in Remark 3.1 that the non-trivial R-charge on $\mathcal{Y}$ is optional, since it can be absorbed into the $\mathrm{GSp}(Q)$ action. However it does affect how we label our line bundles, which is why the Serre functor has this different form. Also note that since $\widetilde{\mathrm{Pf}}_q = \mathrm{Hom}(V,Q)/\mathrm{Sp}(Q)$ the shift here is actually negative; it's:

$$\dim \widetilde{\mathrm{Pf}}_q - 2qv = -\dim \mathrm{Sp}(Q) = -\binom{2q+1}{2}$$

Proof. Objects in $D^b(\mathcal{Y})$ are by definition sheaves on the stack

$$\mathcal{Y}_1 = \big[\,\mathrm{Hom}(V,Q)\,/\,\mathrm{GSp}(Q)\times\mathbb{C}^*_R\big]$$

where the $\mathbb{C}^*_R$ acts by scaling with weight 1. Let's write $\mathcal{Y}_2$ for the same quotient stack, but with the $\mathbb{C}^*_R$ acting trivially. Recall that $\Delta\subset\mathrm{GSp}(Q)$ denotes the diagonal 1-parameter subgroup. There is an isomorphism $f:\mathcal{Y}_2\to\mathcal{Y}_1$ induced by the map of groups

$$\begin{pmatrix}1 & \Delta\\ 0 & 1\end{pmatrix} : \mathrm{GSp}(Q)\times\mathbb{C}^*_R \to \mathrm{GSp}(Q)\times\mathbb{C}^*_R,$$

and f_* induces an equivalence $D^b(\mathcal{Y}_2)\to D^b(\mathcal{Y}_1)$. Under this equivalence the line bundle $(\det Q)^{-v}[\dim\widetilde{\mathrm{Pf}}_q]$ on $\mathcal{Y}_2$ becomes the line bundle $(\det Q)^{-v}[\dim\widetilde{\mathrm{Pf}}_q - 2qv]$. on $\mathcal{Y}_1$, and now the result follows by the analogue of Proposition 3.9. □

3.3. **The kernel.** In this section we construct the Fourier–Mukai kernel for a functor:

$$D^b(\mathcal{X}) \longrightarrow D^b(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$$

This functor will induce the equivalence of Theorem 3.2 in the case $L=0$.

3.3.1. *The definition of the kernel.* As stated above, our kernel will be a matrix factorization living on the relative product of $\mathcal{X}$ and $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$ over the base $[\wedge^2 V/\mathbb{C}^*]$.

Recall that $\langle\omega_S\rangle$ is the character of $\mathrm{GSp}(S)$ contained in $\wedge^2 S$. We let $\mathrm{GSp}(S,Q)\subset \mathrm{GSp}(S)\times\mathrm{GSp}(Q)$ denote the kernel of the character $\langle\omega_S\rangle^{-1}\langle\omega_Q\rangle$. The relative product can then be described as

$$\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y} = \big[\,\mathrm{Hom}(S,V)\times\mathrm{Hom}(V,Q)\,/\,\mathrm{GSp}(S,Q)\big]$$

Notice that $\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}$ admits a map to the stack

$$\mathcal{Z} = \big[\,\mathrm{Hom}(S,Q)\,/\,\mathrm{GSp}(S,Q)\big]$$

by composing the two factors. We denote this map by:

$$\psi : \mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}\longrightarrow\mathcal{Z}$$

There's an obvious superpotential $W\in\Gamma(\mathcal{O}_{\mathcal{Z}})$ which sends $z\in\mathrm{Hom}(S,Q)$ to

$$W(z) = \omega_Q\big(\wedge^2 z(\beta_S)\big)$$

where as before ω_Q is the symplectic form on Q and β_S is the Poisson bivector on S. If we pull this up via ψ we obtain the superpotential we already have, *i.e.* it agrees with the pull-up of the function W on $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$ (1.3), so it seems reasonable to denote all of them by W.

We also give $\mathcal{Z}$ an R-charge by letting $\mathbb{C}^*_R$ act with weight 1 on the underlying vector space, this makes $(\mathcal{Z},W)$ a Landau–Ginzburg B-model, and the map ψ R-charge equivariant.

Our kernel in fact comes from the stack $\mathcal{Z}$, that is it is the pull-up of an object $\mathcal{K}\in D^b(\mathcal{Z},W)$. The Landau–Ginzburg model $(\mathcal{Z},W)$ is very simple, it is just a vector space with a non-degenerate quadratic superpotential, plus a group action. The following form of Knörrer periodicity applies:

Lemma 3.11. *There is an object $\mathcal{K} \in D^b(\mathcal{Z}, W)$ such that the functor*

$$\mathcal{H}om(\mathcal{K}, -) : D^b(\mathcal{Z}, W) \longrightarrow D^b\left(\left[\,\mathrm{pt}\,/\,\mathrm{GSp}(S,Q)\right]\right)$$

is an equivalence, sending $\mathcal{K}$ to $\mathcal{O}_{pt}$.

Proof. If we forget the group action then this statement is basic Knörrer periodicity. *A priori* this might not descend to give the equivariant equivalence, because there could be a Brauer class obstruction. However, in [ST14, Prop. 4.6] it was shown that there is an object $\mathcal{K} \in D^b(\mathcal{Z}, W)$ such that when we forget the group action $\mathcal{K}$ becomes equivalent to the sky-scraper sheaf along a maximal-isotropic subspace in $\mathrm{Hom}(S, Q)$. Since this sky-scraper sheaf induces the non-equivariant equivalence, it follows easily that $\mathcal{K}$ induces the equivariant equivalence. □

In other words, all objects of the category $D^b(\mathcal{Z}, W)$ can be obtained by tensoring $\mathcal{K}$ with (a shifted sum of) $\mathrm{GSp}(S,Q)$-representations. So picking

$$\psi^*\mathcal{K} \quad \in D^b(\mathcal{X} \times_{\mathbb{C}^*} \mathcal{Y},\, W)$$

as our Fourier–Mukai kernel seems like a natural choice. It defines a functor:

$$\widehat{\Phi} : D^b(\mathcal{X}) \to D^b(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$$

We shall see shortly that it in fact defines a functor:

$$\Phi : \mathsf{Br}(\mathcal{X}) \to \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$$

3.3.2. *An important property of $\mathcal{K}$.* We will need the following fact about $\mathcal{K}$, which begins to make the duality manifest. Recall that at the most basic combinatorial level, our duality is the bijection $\gamma \mapsto \gamma^{c\top}$ (1.2) between Young diagrams in $Y_{s,q}$ and $Y_{q,s}$.

If we restrict $\mathcal{K}$ to the origin in $\mathrm{Hom}(S, Q)$ and take its homology we obtain a representation of $\mathrm{GSp}(S, Q)$, which has an underlying $\mathrm{Sp}(S)\times\mathrm{Sp}(Q)$-representation.

Proposition 3.12. *The vector space $h_\bullet(\mathcal{K}|_0)$ is concentrated in degree zero, and as an $\mathrm{Sp}(S) \times \mathrm{Sp}(Q)$ representation we have:*

$$h_\bullet(\mathcal{K}|_0) = \bigoplus_{\gamma \in Y_{s,q}} \mathbb{S}^{\langle\gamma\rangle} S \otimes \mathbb{S}^{\langle\gamma^{c\top}\rangle} Q$$

Recall that it in our notation $\mathbb{S}^{\langle\gamma\rangle}$ denotes a symplectic Schur functor, and $\mathbb{S}^{\gamma}$ denotes an ordinary (GL) Schur functor.

Proof. The object $\mathcal{K}$ was constructed to give an equivalence

$$\mathrm{Hom}(\mathcal{K}, -) : D^b(\mathcal{Z}, W) \xrightarrow{\sim} D^b(\left[\,\mathrm{pt}\,/\,\mathrm{GSp}(S,Q)\right])$$

under which $\mathcal{K}$ maps to the 1-dimensional trivial representation (in degree zero), see Lemma 3.11. Applying this functor to the sky-scraper sheaf $\mathcal{O}_0$ of the origin gives some graded representation R, and then we must have $\mathcal{O}_0 \cong \mathcal{K} \otimes R$. Hence $h_\bullet(\mathcal{K}|_0) = R$, and also $\mathrm{End}(\mathcal{O}_0) = R \otimes R^\vee$. However, $\mathcal{O}_0$ is also equivalent to a matrix factorization given by taking the Koszul complex and perturbing it in the standard way. It follows that as graded representations we must have:

$$R \otimes R^\vee = \mathrm{End}(\mathcal{O}_0) = \wedge^\bullet\big(\mathrm{Hom}(S, Q)[1]\big)$$

Since the vector space $\mathrm{Hom}(S, Q)$ already has R-charge 1 this vector space is actually concentrated in degree zero, so the same is true of R.

As an $\mathrm{Sp}(S) \times \mathrm{Sp}(Q)$-representation, R is determined uniquely by the equality $R^{\otimes 2} = \wedge^\bullet(S \otimes Q)$, so the remainder of the argument is a computation in the representation ring.

We start by computing the character of R. Let $x_1^{\pm 1}, \ldots, x_s^{\pm 1}$ and $y_1^{\pm 1}, \ldots, y_q^{\pm 1}$ denote the standard characters of the maximal tori in $\mathrm{Sp}(S)$ and $\mathrm{Sp}(Q)$ respectively, so that the characters of the standard representations of S and Q are $\sum_{i=1}^{s} x_i + x_i^{-1}$ and $\sum_{j=1}^{q} y_j + y_j^{-1}$. The character of $S \otimes Q$ is then:

$$\sum_{i,j}(x_i + x_i^{-1})(y_j + y_j^{-1})$$

We claim that the character of $\wedge^\bullet(S \otimes Q)$ can be expressed as:

$$\prod_{i,j}(x_i + x_i^{-1} + y_j + y_j^{-1})^2$$

This is an easy computation if $s = q = 1$, and then the general case follows from the fact that $\wedge^\bullet$ converts sums to products. Consequently, the character of R is:

$$\prod_{i,j}(x_i + x_i^{-1} + y_j + y_j^{-1})$$

All monomials appearing in the the above expression have total degree $\leq sq$. To get a monomial of degree exactly sq, we choose a subset Γ of the rectangle $[1, s] \times [1, q]$, then there is a corresponding monomial :

$$\left(\prod_{(i,j)\in\Gamma} x_i\right)\left(\prod_{(i,j)\notin\Gamma} y_j\right)$$

Now choose a partition $\gamma \in Y_{s,q}$, *i.e.* a non-increasing sequence $(\gamma_1, \gamma_2, \ldots, \gamma_s)$ with $\gamma_1 \leq q$. We can define an associated subset

$$\Gamma = \{(i, j),\ q - a_i \leq j \leq q\}\ \subset [1, s] \times [1, q]$$

and the corresponding monomial is

$$m = \left(x_1^{\gamma_1} x_2^{\gamma_2} \cdots x_s^{\gamma_s}\right)\left(y_1^{\beta_1} y_2^{\beta_2} \cdots y_q^{\beta_q}\right)$$

where $\beta = (\beta_1, \ldots, \beta_q)$ is the partition $\beta = \gamma^{c\top}$. This is a dominant weight of $\mathrm{Sp}(S) \times \mathrm{Sp}(Q)$, which we claim is the highest weight of a subrepresentation of R.

To see this, we show that m is a maximal element among the weights of R, in the standard partial ordering of the weight lattice $X(\mathrm{Sp}(S) \times \mathrm{Sp}(Q))$. Recall that the partial ordering on $X(\mathrm{Sp}(S))$ (and similarly on $X(\mathrm{Sp}(Q))$) is such that $\prod x_i^{\gamma_i} \leq \prod x_i^{\gamma'_i}$ if and only if $\sum_{i=1}^{k} \gamma_i \leq \sum_{i=1}^{k} \gamma'_i$ for all $k \in [1, s]$.

We assume for a contradiction that there exists a monomial $\prod x_i^{\gamma'_i} \prod y_i^{\beta'_i}$ among the weights of R such that

$$\prod x_i^{\gamma_i} \prod y_i^{\beta_i} < \prod x_i^{\gamma'_i} \prod y_i^{\beta'_i} \tag{3.13}$$

We must then have

$$sq = \sum \gamma_i + \sum \beta_i \leq \sum \gamma'_i + \sum \beta'_i \leq sq,$$

and so $\sum \gamma'_i + \sum \beta'_i = sq$. Then $\prod x_i^{\gamma'_i} \prod y_i^{\beta'_i}$ must arise from a set $\Gamma' \subset [1, s] \times [1, q]$ as above.

Using this description, it is easy to see that for a fixed γ'_i the monomial $\prod x_i^{\gamma'_i} \prod y_i^{\beta'_i}$ is maximal (i.e. any other choice of β' is strictly smaller) when $\beta' = (\beta'_1, \ldots, \beta'_q) = (\gamma')^{c\top}$, so we may assume that this is the case.

Let now $k \in [1, s]$ be the smallest integer such that $\gamma_k < \gamma'_k$. Then we must have $\beta_i = \beta'_i$ for $i = 1, \ldots, q - \gamma'_k$, and $\beta'_{q-\gamma'_k+1} < \beta_{q-\gamma'_k+1}$, which contradicts (3.13). This proves that R contains the irrep $\mathbb{S}^{\langle\gamma\rangle} S \otimes \mathbb{S}^{\langle\gamma^{c\top}\rangle} Q$.

It remains to show that no other irreps occur in R. Let N be the number of irreps appearing in R (with multiplicities). We know that $N \geq |Y_{s,q}|$, and we want to show that this is an equality. Since $\mathrm{Hom}(R, R) = R^{\otimes 2} = \wedge^\bullet(S \otimes Q)$, we have

$N = \dim \bigwedge^\bullet(S \otimes Q)^{\mathrm{Sp}(S)\times \mathrm{Sp}(Q)}$. A standard computation of Littlewood–Richardson coefficients [FH91, p. 80] shows that as a $\mathrm{GL}(S)\times\mathrm{GL}(Q)$-representation we have:

$$\bigwedge^\bullet(S\otimes Q) = \bigoplus_\alpha \mathbb{S}^\alpha S \otimes \mathbb{S}^{\alpha^\top} Q$$

By Lemma 3.14 below, and its analogue for $\mathrm{Sp}(Q)$, the dimension N of the space of $\mathrm{Sp}(S)\times\mathrm{Sp}(Q)$-invariants equals the number of partitions $\alpha = (\alpha_1, \ldots, \alpha_s)$ such that:

- Each α_i is even (so $(\mathbb{S}^{\alpha^\top}Q)^{\mathrm{Sp}(Q)} = \mathbb{C}$).
- Each number appears an even number of times in α (so $(\mathbb{S}^\alpha S)^{\mathrm{Sp}(S)} = \mathbb{C}$).
- $\alpha_1 \leq 2q$.

Mapping each such partition (α_i) to the partition $(\gamma_i) = (\frac{1}{2}\alpha_{2i})$ gives a bijection onto the set $Y_{s,q}$, so $N = |Y_{s,q}|$. □

Lemma 3.14. *Let α be a partition of length $\leq 2s$. The space of* $\mathrm{Sp}(S)$ *invariants in $\mathbb{S}^\alpha S$ is 1-dimensional if each entry in α occurs an even number of times – equivalently, if each entry in $\alpha^\top$ is even – and zero-dimensional otherwise.*

Proof. By [Sun86, Thms 12.1 and 9.4], an irrep $\mathbb{S}^{\langle\beta\rangle}S$ of $\mathrm{Sp}(S)$ occurs in $\mathbb{S}^\alpha S$ if and only if $\beta \leq \alpha$ and the skew-shape α/β can be labelled by a semi-standard Young tableau of shape δ, such that

(1) every entry in $\delta^\top$ is even, and
(2) for each i, if $2i+1$ appears in row k of the tableau, then $k \leq s+i$.

The multiplicity is given by the number of such tableaux (if the length of α is $\leq s$ then the second condition is vacuous, this case is a classical result due to Littlewood). We are interested in the case $\beta = \varnothing$, in which case there is only one way to label $\alpha/\beta = \alpha$ by a semi-standard Young tableau, we must label all boxes by their row number. Hence $\delta = \alpha$ and the second condition is always satisfied. □

Corollary 3.15. *$\mathcal{K}$ is equivalent to a matrix factorization whose underlying vector bundle is associated to the representation*

$$\bigoplus_{\gamma\in Y_{s,q}} \mathbb{S}^{\langle\gamma\rangle} S \otimes \mathbb{S}^{\langle\gamma^{c\top}\rangle} Q$$

up to twisting each summand by $\mathrm{GSp}(S,Q)$*-characters (but not shifts).*

Proof. The previous proposition computed that $h_\bullet(\mathcal{K}|_0)$ is this $\mathrm{GSp}(S,Q)$-representation, up to characters, but not R-charge shifts. Now apply Lemma 2.6. □

It would be nice to have an explicit construction of the differentials in this matrix factorization, but we don't know how to do this.

Now we take the object $\phi^*\mathcal{K} \in D^b(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}, W)$, and consider the associated functor

$$\widehat{\Phi} = (\pi_2)_*(\pi_1^*(-)\otimes\psi^*\mathcal{K}) :\ D^b(\mathcal{X}) \longrightarrow D^b(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$$

where π_1 and π_2 denote the projection maps from $\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}$ to the two factors $\mathcal{X}$ and $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$.

Corollary 3.16. *The image of $\widehat{\Phi}$ lies inside* $\mathsf{Br}(\mathcal{Y}\times\wedge^2 V,\, W)$.

Proof. From Corollary 3.15, $\psi^*\mathcal{K}$ is equivalent to a matrix factorization built only from the vector bundles $\mathbb{S}^{\langle\gamma\rangle}S\otimes\mathbb{S}^{\langle\gamma^{c\top}\rangle}Q$ with $\gamma \in Y_{s,q}$, up to characters of $\mathrm{GSp}(S,Q)$. Hence the image under $\widehat{\Phi}$ of any object is equivalent to an infinite-rank matrix factorization built from the vector bundles $\mathbb{S}^{\langle\delta\rangle}Q$ with $\delta \in Y_{q,s}$, up to characters of $\mathrm{GSp}(Q)$. Such a matrix factorization obviously satisfies the necessary grade-restriction-rule. □

If we let Φ denote the restriction of $\widehat{\Phi}$ to the subcategory $\mathsf{Br}(\mathcal{X})$, it follows that Φ defines a functor:

$$\Phi : \mathsf{Br}(\mathcal{X}) \to \mathsf{Br}(\mathcal{Y}\times\wedge^2 V,\, W)$$

3.3.3. *Adjoints.* We'll also need to understand the adjoint to this functor Φ. To get a functor from $D^b(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$ to $D^b(\mathcal{X})$ we need a kernel which is an object in $D^b(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, -W)$. The obvious guess for the adjoint to Φ is to use the kernel $\psi^*\mathcal{K}^\vee$, up to some line bundle and shift. This is correct, but the proof is a little involved.

Recall that π_2 is the projection from $\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}$ to $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$. Define a subcategory $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W) \subset D^b(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W)$ by enforcing both the $\mathrm{Sp}(S)$ and $\mathrm{Sp}(Q)$ grade-restriction rules. We have a functor $(\pi_2)_* : \mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W) \to \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$ and it's left adjoint is π_2^*. We define:

$$\pi_2^! = \pi_2^*\left(sv - \binom{v}{2}\right)\left[-\binom{v-2s}{2}\right]$$

Lemma 3.17. *The functor*

$$\pi_2^! : \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W) \;\longrightarrow\; \mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W)$$

is the right adjoint to $(\pi_2)_*$.

Proof. To see this, we pass to the description of the brane categories in terms of algebras as in Section 3.2. The category $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W)$ is equivalent to the derived category of the curved algebra $(A\otimes B, W)$. The functor $(\pi_2)_*$ becomes

$$\mathrm{Hom}(P,-) : D^b(A\otimes B,\, W) \to D^b(\mathcal{O}_{\wedge^2 V}\otimes B,\, W),$$

where $P \in D^b(A\otimes B)$ is the object corresponding to:

$$\bigoplus_{\delta\in Y_{q,s}} \mathbb{S}^{\langle\delta\rangle}Q \in D^b(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y})$$

The functor $\pi_2^!$ becomes $P\otimes_{\mathcal{O}_{\wedge^2 V}\otimes B} -$, up to twists.

In order to apply results from the literature, let us instead consider the functors

$$\phi^* := (A\otimes B)\otimes_{\mathcal{O}_{\wedge^2 V}\otimes B} - :\; D^b(\mathcal{O}_{\wedge^2 V}\otimes B,\, W) \to D^b(A\otimes B,\, W)$$

and:

$$\phi_* := \mathrm{Hom}(A\otimes B, -) :\; D^b(A\otimes B,\, W) \;\longrightarrow\; D^b(\mathcal{O}_{\wedge^2 V}\otimes B,\, W)$$

We claim that some twist of ϕ^* is the right adjoint to ϕ_*. Since P is a summand of $A\otimes B$, the claim $\pi_2^! \vdash (\pi_2)_*$ follows from this.

By [YZ06, Ex. 6.4], $A\otimes B$ admits a rigid dualising complex $\omega_{A\otimes B}$, and since $A\otimes B$ is a self-dual maximal Cohen–Macaulay module over a Gorenstein ring, it is easy to check that $\omega_{A\otimes B}\cong A\otimes B$ up to twist and shift. We define the dualising functor $\mathbb{D}_{A\otimes B} = \mathcal{H}om(-, \omega_{A\otimes B})$.

Similarly, we get a dualising complex $\omega_{\mathcal{O}_{\wedge^2 V}\otimes B}$ for $\mathcal{O}_{\wedge^2 V}\otimes B$, and an induced duality functor $\mathbb{D}_{\mathcal{O}_{\wedge^2 V}\otimes B}$. We now define

$$\phi^! = \mathbb{D}_{\mathbb{A}\otimes\mathbb{B}} \circ \phi^* \circ \mathbb{D}_{\mathcal{O}_{\wedge^2 V}\otimes B},$$

and see that up to twist $\phi^! \cong \phi^*$.

Let us first ignore the presence of the superpotential and prove $\phi^! \vdash \phi_*$ when $W = 0$. In this case by [YZ04, Thm. 4.13] we have

$$\phi^! \cong \phi^\flat := \mathcal{H}om(A\otimes B, -) : D^b(\mathcal{O}_{\wedge^2 V}\otimes B) \to D^b(A\otimes B).$$

It is shown in [YZ04] that a nondegenerate trace map $\phi_*\phi^\flat \to 1$ exists, and arguing as in [Har66, Sec. III.6] the adjoint property $\phi^! \vdash \phi_*$ follows.

We now turn to the case $W \neq 0$. Note first that the counit map $\phi_*\phi^! \to 1$ is induced by a map of Fourier–Mukai kernels supported along the diagonal in (the centre of) $(\mathcal{O}_{\wedge^2 V}\otimes B)\otimes(\mathcal{O}_{\wedge^2 V}\otimes B)^{op}$. Hence we get a counit map in the case with superpotential as well.

Given $\mathcal{E}\in D^b(A\otimes B, W)$ and $\mathcal{F}\in D^b(\mathcal{O}_{\wedge^2 V}\otimes B, W)$, we have now shown that the induced map

$$\mathrm{Hom}(\mathcal{E},\pi^!\mathcal{F})\to\mathrm{Hom}(\pi_*\mathcal{E},\pi_*\pi^!\mathcal{F})\to\mathrm{Hom}(\pi_*\mathcal{E},\mathcal{F})$$

is an isomorphism if $W=0$. For general W, we may degenerate both $\mathcal{E}$ and $\mathcal{F}$ to honest complexes, and then the upper semicontinuity of cohomology implies that the map is an isomorphism in the general case as well. □

Now consider the projection $\pi_1:\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}\to\mathcal{X}$. There is a problem here: this map is not proper (not even equivariantly so), so $(\pi_1)_*$ really lands in the derived category $D^b(QCoh(\mathcal{X}))$ of quasi-coherent sheaves. Nevertheless, we claim the following:

Proposition 3.18. *The composition*

$$(\pi_1)_*\mathcal{H}om\big(\psi^*\mathcal{K},\pi_2^!(-)\big)$$

defines a functor

$$\Phi^\dagger:\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2V,\,W)\;\longrightarrow\;\mathsf{Br}(\mathcal{X})$$

and this is the right adjoint to Φ.

Proof. Using Corollary 3.15 it's clear that we have a functor

$$\mathcal{H}om\big(\psi^*\mathcal{K},\pi_2^!(-)\big):\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2V,\,W)\;\longrightarrow\;\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y})\tag{3.19}$$

and Lemma 3.17 shows that this is the right adjoint to the functor:

$$(\pi_2)_*(\mathcal{K}\otimes -)$$

What we have to prove is that if we apply $(\pi_1)_*$ to something in the image of (3.19) then we get a coherent sheaf; then since $(\pi_1)_*$ is right-adjoint to π_1^* it follows immediately that $\Phi^\dagger$ is the right adjoint to Φ.

The functor $(\pi_1)_*$ respects the $\mathrm{Sp}(S)$ grade-restriction-rule, or in terms of algebras, it maps $A\otimes B$-modules to (perhaps infinitely-generated) A-modules. Now take $\mathcal{F}\in\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V, W)$, and consider the A-module $M=(\pi_1)_*\mathcal{H}om(\psi^*\mathcal{K},\pi_2^!\mathcal{F})$. We want to know that M is in fact finitely-generated. By adjunction, we have

$$\begin{aligned}M=\mathrm{Hom}_A(A,M)&=\mathrm{Hom}_{\mathcal{X}}\big(T,\,(\pi_1)_*\mathcal{H}om(\psi^*\mathcal{K},\pi_2^!\mathcal{F})\big)\\&=\mathrm{Hom}\big((\pi_2)_*(\mathcal{K}\otimes\pi_1^*T),\,\mathcal{F}\big)\end{aligned}$$

This last one is a Hom space in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2V,W)$, so it is a finitely generated graded module over $\mathcal{O}_{\wedge^2V}$. Since the functors are linear over $\wedge^2V$, it follows that M is finitely-generated module over $\mathcal{O}_{\wedge^2V}$ and hence also over A. □

3.4. **The equivalence over the smooth locus.** Since we've constructed our kernel $\psi^*\mathcal{K}$ relative to the base $[\wedge^2V/\mathbb{C}^*]$, we can restrict ourselves to open sets in this base and examine the functor there.

Specifically, we're going to look at the open subset

$$U\subset\wedge^2V$$

consisting of bivectors having rank at least $2s$. This gives a substack $[U/\mathbb{C}^*]\subset[\wedge^2V/\mathbb{C}^*]$, which is just a quasi-projective variety $\mathbb{P}U\subset\mathbb{P}(\wedge^2V)$. The intersection of $\mathbb{P}U$ with the Pfaffian Pf_s is exactly the smooth locus Pf_s^{sm}, so $\mathcal{X}|_{\mathbb{P}U}$ is equivalent to Pf_s^{sm}. This brings us very close to the situation considered in [ADS15, ST14], so we'll now spend a little time making the connection to the point-of-view of those earlier papers.

On the dual side, restricting $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$ to U simply gives $\mathcal{Y}\times_{\mathbb{C}^*}U$. We can think of this a bundle of stacks over the variety $\mathbb{P}U$, with fibres $\widetilde{\mathcal{Y}} = \big[\operatorname{Hom}(V,Q)\,/\,\operatorname{Sp}(Q)\big]$. Locally in $\mathbb{P}U$ it's just a vector bundle, quotiented by a fibre-wise $\operatorname{Sp}(Q)$ action. This is not true globally, but it is a vector bundle over the stack $[U/\operatorname{GSp}(Q)]$, which is a bundle of stacks over $\mathbb{P}U$ with fibres $B\operatorname{Sp}(Q)$.

Our kernel lives on the product of $\mathcal{X}|_{\mathbb{P}U}$ and $\mathcal{Y}\times_{\mathbb{C}^*}U$ relative to $\mathbb{P}U$, which is simply the restriction of the bundle of stacks $\mathcal{Y}\times_{\mathbb{C}^*}\mathcal{U}$ to the subvariety $\operatorname{Pf}_s^{sm}\subset\mathbb{P}U$. So we have a functor

$$D^b(\operatorname{Pf}_s^{sm}) \to D^b(\mathcal{Y}\times_{\mathbb{C}^*}U,\, W)$$

given by pulling up to $\mathcal{Y}\times_{\mathbb{C}^*}U|_{\operatorname{Pf}_s^{sm}}$, tensoring with the object $\psi^*\mathcal{K}|_{\mathbb{P}U}$, and then pushing-forward along the inclusion map.

Now, to any point $a\in\operatorname{Pf}_s^{sm}$ we can associate the image of a, this gives a rank $2s$ subbundle

$$\Sigma\subset V(1)$$

which carries a symplectic form, determined up to scale. There's an associated bundle of stacks over Pf_s^{sm} whose fibres are $\big[\operatorname{Hom}(\Sigma,Q)\,/\,\operatorname{Sp}(Q)\big]$. To be precise: take the smooth locus $(\widetilde{\operatorname{Pf}_s})^{sm}$ of the affine cone over Pf_s, this is a $\mathbb{C}^*$-bundle over Pf_s^{sm}. We have a vector bundle $\widetilde{\Sigma}$ over $(\widetilde{\operatorname{Pf}_s})^{sm}$, which is a subbundle of the trivial bundle $V\times(\widetilde{\operatorname{Pf}_s})^{sm}$. Then we form the stack:

$$\overline{\mathcal{Z}} = \big[\operatorname{Hom}(\widetilde{\Sigma},Q)\,/\,\operatorname{GSp}(Q)\big]$$

This is a bundle of stacks over Pf_s^{sm}, or a vector bundle over $[(\widetilde{\operatorname{Pf}_s})^{sm}/\operatorname{GSp}(Q)]$.

We have a quotient map

$$\overline{\psi}:\mathcal{Y}\times_{\mathbb{C}^*}U|_{\operatorname{Pf}_s^{sm}}\longrightarrow\overline{\mathcal{Z}}$$

and, by construction, our kernel is pulled-back from $\overline{\mathcal{Z}}$. If this is not immediately obvious, observe that we can factor the map

$$\psi:\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}\longrightarrow\mathcal{Z}=\big[\operatorname{Hom}(S,Q)\,/\,\operatorname{GSp}(S,Q)\big]$$

through an intermediate step:

$$\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}\longrightarrow\mathcal{X}\times_{\operatorname{GSp}(S)}\mathcal{Z}=\big[\operatorname{Hom}(S,V)\times\operatorname{Hom}(S,Q)\,/\,\operatorname{GSp}(S,Q)\big]$$

Restricting this to $\mathbb{P}U$, we obtain the stack $\overline{\mathcal{Z}}$ and the map $\overline{\psi}$. We can pull-back the superpotential W to $\overline{\mathcal{Z}}$, and we have an object

$$\overline{\mathcal{K}}\in D^b(\overline{\mathcal{Z}},W)$$

given by pulling back $\mathcal{K}$. So over $\mathbb{P}U$, our kernel is equal to $\overline{\psi}^*\overline{\mathcal{K}}$.

The superpotential on $\overline{\mathcal{Z}}$ gives a non-degenerate quadratic form on each fibre, so Knörrer periodicity should apply unless there's a Brauer class obstruction. However the whole point of the object $\mathcal{K}$ was that it gives a universal way to implement Knörrer periodicity for this kind of bundle, so indeed there is no Brauer class obstruction, and the object $\overline{\mathcal{K}}$ induces an equivalence between $D^b(\overline{\mathcal{Z}},W)$ and $D^b([(\widetilde{\operatorname{Pf}_s})^{sm}/\operatorname{GSp}(Q)])$. Another point to note is that if we restrict to open neighbourhoods in Pf_s^{sm} then there are other ways to implement this equivalence. For example, we may pick a Lagrangian subbundle $\Lambda\subset\Sigma$ (this is possible in a small-enough neighbourhood), which induces an $\operatorname{Sp}(Q)$-invariant maximal isotropic subbundle $\operatorname{Hom}(S/\Lambda,Q)$ in $\overline{\mathcal{Z}}$. The sky-scraper sheaf on this subbundle gives a second object in $D^b(\overline{\mathcal{Z}},W)$ which is equivalent to $\overline{\mathcal{K}}$.

Now let's take our subcategories $\mathsf{Br}(\mathcal{X})$ and $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\,W)$ and base-change them to the open set U, meaning (as in Section 3.2) that we consider the full triangulated subcategories generated by their images under restriction.

When we do this to $\mathsf{Br}(\mathcal{X})$ we get the whole of $D^b(\mathrm{Pf}_s^{sm})$, because $\mathsf{Br}(\mathcal{X}) = D^b(A)$ is the derived category of a non-commutative resolution of Pf_s^{sm}. Indeed by Corollary 3.6 (and Remark 3.7) $\mathsf{Br}(\mathrm{Pf}_s^{sm})$ is the derived category of the sheaf of algebras $A = \mathrm{End}_{\mathcal{X}}(T)$, but the restriction of T to $\mathcal{X}|_U = \mathrm{Pf}_s^{sm}$ is still a vector bundle so this is a trivial Azumaya algebra. However, on the dual side we get a proper subcategory:

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U, W) \subset D^b(\mathcal{Y} \times_{\mathbb{C}^*} U, W)$$

Everything in this subcategory satisfies the grade-restriction-rule along the zero section $\mathbb{P}U$ – this follows from Lemma 2.6 and the fact that the grade-restriction rule is preserved under taking mapping cones. Presumably this rule exactly characterizes $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U, W)$ but we won't need this fact.

We've seen that over the whole of $[\wedge^2 V/\mathbb{C}^*]$, our kernel induces adjoint functors Φ and $\Phi^\dagger$ between our Br subcategories (Corollary 3.16). If we restrict to $\mathbb{P}U$, it follows immediately that we get a pair of adjoint functors between $D^b(\mathrm{Pf}_s^{sm})$ and the subcategory $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U, W)$. Let's denote these functors by Φ_U and $\Phi_U^\dagger$.

Theorem 3.20. *These functors give an equivalence*

$$\Phi_U : D^b(\mathrm{Pf}_s^{sm}) \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U, W)$$

and its inverse.

Most of this theorem was proven in [ADS15, ST14]. We'll give a complete proof here, partly for convenience, and also because those previous papers inexplicably failed to prove essential-surjectivity. We'll get the result as a corollary of the Lemma 3.21 below.

Fix a point $a \in \mathbb{P}U$, and let $\widehat{\mathbb{P}U}_a$ denote the formal neighbourhood of a; this is isomorphic to a formal neighbourhood of 0 in $\wedge^2 V/\langle a \rangle$. In the case that a lies in Pf_s, we write $\widehat{(\mathrm{Pf}_s)}_a$ for the formal neighbourhood of Pf_s at a.

If we restrict $\mathcal{Y} \times_{\mathbb{C}^*} U$ to the formal neighbourhood $\widehat{\mathbb{P}U}_a$ it becomes:

$$\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a = \big[\mathrm{Hom}(V, Q) \, / \, \mathrm{Sp}(Q)\big] \times \widehat{\mathbb{P}U}_a$$

We define a subcategory

$$\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a, W) \subset D^b(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a, W)$$

as the full subcategory of objects $\mathcal{E}$ such that all $\mathrm{Sp}(Q)$-irreps occuring in $h_\bullet(\mathcal{E}|_{(0,a)})$ come from the set $Y_{q,s}$. The restriction of an object in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U, W)$ lands in this subcategory, and presumably these objects generate the category. If that is true then it follows formally that Φ and $\Phi^\dagger$ induce adjoint functors between $D^b(\widehat{(\mathrm{Pf}_s)}_a)$ and $\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a, W)$; however this can be proven directly using the arguments of Corollary 3.16.

Lemma 3.21.

(1) *If $a \in \mathrm{Pf}_s^{\mathrm{sm}}$ then Φ induces an equivalence:*

$$D^b(\widehat{(\mathrm{Pf}_s)}_a) \xrightarrow{\sim} \mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a, W)$$

(2) *If $a \notin \mathrm{Pf}_s$ then $\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a, W)$ is zero.*

Proof. Let the rank of a be $2t$, so a lies in the smooth locus of the Pfaffian variety Pf_t. Along Pf_t there is a rank $2t$ bundle $\Sigma \subset V(1)$ given by the images of the 2-forms at each point, and a rank $v - 2t$ bundle $K \subset V^\vee$ given by the kernels. In the formal neighbourhood $\widehat{\mathbb{P}U}_a$ we can do a gauge transformation to trivialize the family of 2-forms on the bundle $V \times \mathrm{Pf}_t$; then both Σ and K become trivial bundles with fibre the image and kernel of a. Furthermore, the normal bundle to Pf_t at a is $\wedge^2 K^\vee$, so after a formal change of co-ordinates:

$$\widehat{\mathbb{P}U}_a \cong \widehat{\wedge^2 K^\vee}_0 \times \widehat{(\mathrm{Pf}_t)}_a$$

If we choose a splitting $V^\vee = K \oplus \Sigma$, we can write:

$$\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a = \big[\mathrm{Hom}(\Sigma, Q) \,/\, \mathrm{Sp}(Q)\big] \times \big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\wedge^2 K^\vee}_0 \times \widehat{(\mathrm{Pf}_t)}_a$$

The superpotential is now a a sum of a quadratic W_q and a cubic term W_c. The quadratic term is the tautological superpotential on the first factor, and the cubic term is the tautological superpotential on $\big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\wedge^2 K^\vee}_0$. We can use Knörrer periodicity to remove the quadratic term, and get an equivalence:

$$\Psi : D^b\big(\big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\mathbb{P}U}_a,\, W_c\big) \;\xrightarrow{\sim}\; D^b\big(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a,\, W\big)$$

We can do this using Lemma 3.11, which provides (forgetting the $\mathrm{GSp}(S)$ action) an object $\mathcal{L} \in D^b([\mathrm{Hom}(\Sigma, Q)/\mathrm{Sp}(Q)],\, W_q)$; then the equivalence is given by pulling-up and tensoring with $\mathcal{L}$. Alternatively we may pick a Lagrangian $\Lambda \subset \Sigma$ and use the sky-scraper sheaf along $\mathrm{Hom}(\Sigma/\Lambda, Q)$.

Let's examine this construction in the case that $a \in \mathrm{Pf}_s$. Within the formal neighbourhood, the functor Φ is given by pulling-up to $[\mathrm{Hom}(V, Q)/\mathrm{Sp}(Q)] \times \widehat{(\mathrm{Pf}_s)}_a$, tensoring with the object $\overline{\psi}^* \mathcal{K}$, and then pushing-forward into $[\mathrm{Hom}(V, Q)/\mathrm{Sp}(Q)] \times \widehat{\mathbb{P}U}_a$. We can form a commutative diagram:

$$\begin{array}{ccc}
D^b\big(\big[\mathrm{Hom}(V, Q) \,/\, \mathrm{Sp}(Q)\big] \times \widehat{(\mathrm{Pf}_s)}_a,\, W\big) & \longrightarrow & D^b\big(\big[\mathrm{Hom}(V, Q) \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\mathbb{P}U}_a,\, W\big) \\
\Psi|_{\widehat{(\mathrm{Pf}_s)}_a} \big\uparrow & & \Psi \big\uparrow \\
D^b\big(\big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{(\mathrm{Pf}_s)}_a\big) & \longrightarrow & D^b\big(\big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\mathbb{P}U}_a,\, W_c\big) \\
\big\uparrow & & \\
D^b\big(\widehat{(\mathrm{Pf}_s)}_a\big) & &
\end{array}$$

Here the horizontal arrows are just push-foward along the inclusion maps. Note that the top row is natural, as is the composition of the two vertical arrows on the left; however the middle row depends on our choices of co-ordinates and splittings. Moving from the bottom left corner to the top right corner gives the functor Φ, because the object $\overline{\psi}^* \mathcal{K}$ is exactly the restriction of the object $\mathcal{L}$.

Next we need to examine what happens to the subcategory $\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a,\, W)$ under the equivalence Ψ, in both cases $a \in \mathrm{Pf}_s$ and $a \notin \mathrm{Pf}_s$. Given an object $\mathcal{E} \in D^b(\big[K \otimes Q \,/\, \mathrm{Sp}(Q)\big] \times \widehat{\mathbb{P}U}_a,\, W_c)$, if we want to find the representation $h_\bullet\big(\Psi\mathcal{E}|_{(0,a)}\big)$, we take $h_\bullet(\mathcal{E}|_{(0,a)})$ and tensor it with $h_\bullet(\mathcal{L}|_0)$. By Proposition 3.12, the $\mathrm{Sp}(Q)$-irreps that occur in $h_\bullet(\mathcal{L}|_0)$ are exactly those of the form $\mathbb{S}^{\langle\delta\rangle} Q$ where the width of δ is at most t (one can also reach this conclusion by taking the Koszul resolution of $\mathcal{O}_{\mathrm{Hom}(\Sigma/\Lambda, Q)}$).

Now suppose that $a \notin \mathrm{Pf}_s$, so $t > s$. If $\mathcal{E}$ is such that $h_\bullet(\mathcal{E}|_{(0,a)}) \neq 0$ then $\Psi\mathcal{E}$ does not lie in the subcategory $\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a,\, W)$, but $h_\bullet(\mathcal{E}|_{(0,a)}) = 0$ implies that $\mathcal{E} \simeq 0$ by Lemma 2.6. This proves (2).

Suppose instead that $a \in \mathrm{Pf}_s$, so $t = s$. If $\Psi\mathcal{E} \in \mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathcal{U}}_a,\, W)$ then we must have that $h_\bullet(\mathcal{E}|_{(0,a)})$ is a trivial $\mathrm{Sp}(Q)$-representation, which means (by Lemma 2.6 again) that $\mathcal{E}$ is equivalent to a matrix factorization whose underlying vector bundle is equivariantly trivial. The category of such matrix factorizations is exactly the same as the category of matrix factorizations on the underlying scheme, *i.e.* the scheme-theoretic quotient of $K \otimes Q$ by $\mathrm{Sp}(Q)$. This quotient is itself a Pfaffian, it's the rank $2q$ locus in $\wedge^2 K$. But since $a \in \mathrm{Pf}_s$ we have $\dim K = 2q+1$, so in fact the quotient is the whole of $\wedge^2 K$. So in this case, Ψ induces an equivalence:

$$\mathsf{Br}(\widetilde{\mathcal{Y}} \times \widehat{\mathbb{P}U}_a,\, W) \;\xrightarrow{\sim}\; D^b(\wedge^2 K \times \widehat{\mathcal{U}}_a,\, W_c)$$

Note that once we've quotiented, the superpotential W_c becomes quadratic, it's just the pairing between $\wedge^2 K$ and $\wedge^2 V/\langle a\rangle$.

Comparing this with our diagram above, we see that within our formal neighbourhood we can identify the functor Φ with the functor

$$D^b(\widehat{(\mathrm{Pf}_s)}_a) \longrightarrow D^b(\wedge^2 K \times \widehat{\mathcal{U}}_a,\, W_c)$$

given by 'pull-up to $\wedge^2 K \times \widehat{(\mathrm{Pf}_s)}_a$, then push-forward'. By Knörrer periodicity this functor is an equivalence. □

Proof of Theorem 3.20. Let's first prove that Φ_U is fully faithful. This question is local in $\mathbb{P}U$, and locally $D^b(\mathrm{Pf}_s^{sm})$ is generated by the structure sheaf, so it's sufficent to check that

$$\Phi_U : \mathcal{O}_{\mathrm{Pf}_s} \longrightarrow \mathrm{Hom}\big(\Phi_U(\mathcal{O}_{\mathrm{Pf}_s}), \Phi_U(\mathcal{O}_{\mathrm{Pf}_s})\big)$$

is an isomorphism. Part (1) of the previous lemma says that this map is an isomorphism in the formal neighbourhood of any point $a \in \mathrm{Pf}_s^{sm}$, so it's an isomorphism.

Now we show that Φ_U is essentially surjective, which is is equivalent to the adjoint $\Phi_U^\dagger$ having no kernel. Suppose that $\mathcal{E} \in \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} U,\, W)$. By part (2) of the previous lemma, $\mathcal{E} \equiv 0$ in the formal neighbourhood of any point $a \notin \mathrm{Pf}_s^{sm}$. If $\Phi_U^\dagger \mathcal{E} = 0$, then (by part (1)) this is also true when $a \in \mathrm{Pf}_s^{sm}$. Hence $\mathcal{H}om(\mathcal{E}, \mathcal{E})$ is acyclic in the formal neighbourhood of any point $a \in \mathbb{P}U$, so it is acyclic, and $\mathcal{E}$ is contractible. □

3.5. **Extending over the singular locus.** In Section 3.2 we saw that $\mathsf{Br}(\mathcal{X})$ is equivalent to $D^b(A)$, and it contains a finite generating set of objects $\{\mathbb{S}^{\langle\gamma\rangle} S,\, \gamma \in Y_{s,q}\}$ corresponding to the obvious projective A-modules. To prove our equivalence, we need to show that the category $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$ has exactly the same structure.

3.5.1. *Some objects on the dual side.* First we must identify the corresponding generating objects in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$.

For each Young diagram $\delta \in Y_{q,s}$, we define an object $\widetilde{\mathcal{P}}_\delta \in D^b(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$ as:

$$\widetilde{\mathcal{P}}_\delta \;=\; \mathcal{O}_{0 \times \wedge^2 V} \otimes \mathbb{S}^{\langle\delta\rangle} Q$$

This is an object in $D^b(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$ because W vanishes along the locus $0 \times \wedge^2 V$, but it does not lie in the subcategory $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$. However, this subcategory is right-admissible (Lemma 3.8) so for each $\delta \in Y_{q,s}$, we can define

$$\mathcal{P}_\delta \in \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V,\, W)$$

be the image of $\widetilde{\mathcal{P}}_\delta$ under the right-adjoint to the inclusion.

Lemma 3.22. *The functor $\Phi^\dagger$ sends the object $\mathcal{P}_\delta$ to the vector bundle $\mathbb{S}^{\langle \delta^{\top c}\rangle} S$, up to a shift and line bundle.*

The shift is constant (independent of δ), the line bundle might not be.

Proof. Recall (from Section 3.3.3) that $\Phi^\dagger$ is the functor

$$\Phi^\dagger : \mathcal{E} \mapsto (\pi_1)_* \mathcal{H}om(\psi^* \mathcal{K}, \pi_2^* \mathcal{E})$$

up to some shift and line-bundle. We claim that:

$$\Phi^\dagger \mathcal{P}_\delta = \Phi^\dagger \widetilde{\mathcal{P}}_\delta$$

This is because the kernel $\psi^*\mathcal{K}$ lives in the admissible subcategory $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y},\, W)$, so $\Phi^\dagger\widetilde{\mathcal{P}}_\delta$ only depends on the projection of $\pi_2^*\widetilde{\mathcal{P}}_\delta$ into this subcategory, and this is the same as $\pi_2^*\mathcal{P}_\delta$.

Now we compute $\Phi^\dagger(\mathcal{P}_\delta)$. The pull-up $\pi_2^*\mathcal{O}_{0\times\wedge^2 V}$ is just the sky-scraper sheaf along the subspace $\mathrm{Hom}(S,V)\times 0$ in $\mathcal{X}\times_{\mathbb{C}^*}\mathcal{Y}$, and now we can use the model for $\psi^*\mathcal{K}$ provided by Corollary 3.15 to see the result. □

Remark 3.23. Another way to characterize these objects $\mathcal{P}_\delta$ is to consider their images in the equivalent category $D^b(B\otimes\mathcal{O}_{\wedge^2 V}, W)$. Suppose we work on $\mathcal{Y}$ instead, and consider a twist of the sky-scraper sheaf at the origin $\mathcal{O}_0\otimes\mathbb{S}^{\langle\delta\rangle}Q$. To project this into $\mathsf{Br}(\mathcal{Y})$ we first apply the functor $\mathrm{Hom}_{\mathcal{Y}}(T,-): D^b(\mathcal{Y})\to D^b(B)$, where T is the direct sum of the vector bundles from $Y_{q,s}$. The result is a one-dimensional B-module M_δ, it's the 'vertex simple' at the vertex δ. Now simply cross everything with $\wedge^2 V$: we see that $\mathcal{P}_\delta$ corresponds to the module $M_\delta\otimes\mathcal{O}_{\wedge^2 V}$. Note that this is indeed a module over the curved algebra $(B\otimes\mathcal{O}_{\wedge^2 V}, W)$ because W acts as zero on it.

We will see in due course that this set of objects $\{\mathcal{P}_\delta\}$ generates $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$. The essential point is the following:

Lemma 3.24. *The set of objects $\{\mathcal{P}_\delta, \delta\in Y_{q,s}\}$ in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$ has no left orthogonal.*

Proof. Take a non-zero object $\mathcal{E}\in\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$, we want to prove that there is some $\delta\in Y_{q,s}$ such that:

$$\mathrm{Hom}(\mathcal{E},\mathcal{P}_\delta)=\mathrm{Hom}(\mathcal{E},\widetilde{\mathcal{P}}_\delta)\neq 0$$

By Lemma 2.6 we can assume the sheaf underlying $\mathcal{E}$ is a vector bundle and the differential vanishes at the origin, so $\mathcal{E}$ is the bundle associated to the representation $\mathcal{E}|_0$. In particular, $\mathcal{E}^\vee|_0$ is a non-zero representation, and by the grade-restriction rule it only contains $\mathrm{Sp}(Q)$-irreps from the set $Y_{q,s}$.

Decompose $\mathcal{E}^\vee|_0$ under R-charge, let t be the highest weight occuring, and let $(\mathcal{E}^\vee|_0)_t$ denote the highest weight space. Let $(\mathcal{E}^\vee)_t$ denote the corresponding sub-bundle of $\mathcal{E}^\vee$. Now consider:

$$\mathcal{H}om(\mathcal{E},\mathcal{O}_{0\times\wedge^2 V})=\mathcal{E}^\vee|_{0\times\wedge^2 V}$$

The R-charge action on $\wedge^2 V$ is trivial, so this is a bounded complex of $\mathrm{GSp}(Q)$-equivariant vector bundles on $\wedge^2 V$, whose final term is the restriction of $(\mathcal{E}^\vee)_t$. The differential vanishes at the origin, so the final differential cannot be surjective, and there is a non-zero homology sheaf $h_t(\mathcal{E}^\vee|_{0\times\wedge^2 V})$ in the top degree. We have surjections:

$$(\mathcal{E}^\vee)_t|_{0\times\wedge^2 V}\longrightarrow h_t(\mathcal{E}^\vee|_{0\times\wedge^2 V})\longrightarrow h_t(\mathcal{E}^\vee|_{0\times\wedge^2 V})|_0=(\mathcal{E}^\vee|_0)_t$$

By the grade-restriction-rule, there is some $\delta\in Y_{q,s}$ such that $(\mathcal{E}^\vee|_0)_t\otimes\mathbb{S}^{\langle\delta\rangle}Q$ contains non-zero $\mathrm{Sp}(Q)$-invariants, which implies that $h_t(\mathcal{E}^\vee|_{0\times\wedge^2 V})\otimes\mathbb{S}^{\langle\delta\rangle}Q$ doesn't vanish after taking $\mathrm{Sp}(Q)$-invariants. Since taking $\mathrm{Sp}(Q)$ invariants is exact, it follows that

$$\mathrm{Hom}(\mathcal{E},\mathcal{P}_\delta)=\left(\mathcal{E}^\vee|_{0\times\wedge^2 V}\otimes\mathbb{S}^{\langle\delta\rangle}Q\right)^{\mathrm{Sp}(Q)}$$

has non-zero homology in degree t. □

3.5.2. *The algebra on the dual side.* Let A' denote the dg-algebra:

$$A'=\mathrm{End}_{D^b(\mathcal{Y}\times\wedge^2 V,\, W)}\left(\bigoplus_{\delta\in Y_{q,s}}\mathcal{P}_\delta\right)$$

By definition, this is a $\mathbb{C}^*$-equivariant dga over the ring of functions on $\wedge^2 V$.

In Section 3.2 we saw that $\mathsf{Br}(\mathcal{X})$ is equivalent to the derived category of an algebra A. This algebra A is also defined over $[\wedge^2 V/\mathbb{C}^*]$, and it is Cohen–Macaulay. In this section we will prove that A' is in fact quasi-isomorphic to A. The results of the previous section show that A' and A are quasi-isomorphic over the open set $U \subset \wedge^2 V$. We will show that A' has homology only in degree 0 – so it's really an algebra – and that it's Cohen–Macaulay. Together these facts will imply that $A' \simeq A$ globally.

Recall (Remark 3.4) that our definition of A was slightly ambiguous, because we didn't specify which lift of $\mathbb{S}^{\langle\gamma\rangle}S$ from $\widetilde{\mathcal{X}}$ to $\mathcal{X}$ we wanted to take. Let's now fix this ambiguity. For each $\delta \in Y_{q,s}$ the object $\Phi^\dagger(\mathcal{P}_\delta)$ is a lift of the vector bundle $\mathbb{S}^{\langle\delta^{\top c}\rangle}S$, up to a constant shift; this is the lift we'll take.

Lemma 3.25. *Over the open set $U \subset \wedge^2 V$, we have a quasi-isomorphism:*

$$\Phi_U^\dagger : A'|_U \xrightarrow{\sim} A|_U$$

Proof. This follows immediately from Theorem 3.20. ☐

Lemma 3.26. *The dga A' has homology concentrated in degrees $[-\binom{2q+1}{2}, 0]$.*

Proof. By Remark 3.23, the object $\mathcal{P}_\delta$ corresponds to the module $M_\delta \otimes \mathcal{O}_{\wedge^2 V}$ over the curved algebra $(B \otimes \mathcal{O}_{\wedge^2 V}, W)$, where M_δ is the one-dimensional B-module corresponding to the vertex δ. Consider the dga $E = \mathrm{End}_B(\oplus_{\delta\in Y_{q,s}} M_\delta)$; we first claim that the homology of E is concentrated in degrees $[-\binom{2q+1}{2}, 0]$.

To see this, observe that the B-module $\oplus_{\delta\in Y_{q,s}} M_\delta$ is the quotient $B/\overline{B}$ where $\overline{B}$ is the ideal in B generated by the arrows. Hence it has a bar resolution whose underlying module is:

$$\bigoplus_{p\geq 0} B \otimes \overline{B}^{\otimes p}[p]$$

Since the R-charge acts with weight 1 on $\mathrm{Hom}(V, Q)$, the ideal $\overline{B}$ is the subspace of B where the R-charge is ≥ 1, so every term in this bar resolution is concentrated in non-negative degrees. It follows that E is concentrated in non-positive degrees. For the lower bound, we use the Serre functor on $D^b(B) = \mathsf{Br}(\mathcal{Y})$ from Proposition 3.10.

To compute A', we can take a projective resolutions of each M_δ, tensor them with $\mathcal{O}_{\wedge^2 V}$, then perturb the differential to get a module over $(B \otimes \mathcal{O}_{\wedge^2 V}, W)$ which is equivalent to M_δ [Seg11, Lemma 3.6]. This means that A' can be computed from a spectral sequence that starts with E. The result follows. ☐

To show that A' in fact has homology only in degree zero, and is Cohen–Macaulay, we develop the following general criterion:

Lemma 3.27. *Let $\mathcal{F}$ be a $\mathbb{C}^*$-equivariant complex on $\mathbb{A}^n$ which has support on a set $Z \subset \mathbb{A}^n$ of codimension c. If $\mathcal{F}|_0$ is a complex with homology concentrated in degrees $[-c, 0]$, then the homology of $\mathcal{F}$ is concentrated in degree zero and is a Cohen–Macaulay sheaf.*

Proof. Let k denote the skyscraper sheaf at the origin. Note that the condition on $\mathcal{F}|_0$ is equivalent to the requirement that $\mathrm{Ext}^i(\mathcal{F}, k) = 0$ if $i \notin [0, c]$.

We first prove the claim when $n = c$, so that $Z = 0$. There is a spectral sequence where $E_2^{p,q} = \mathrm{Ext}^p(\mathcal{H}^{-q}(\mathcal{F}), k)$, which converges to $\mathrm{Ext}^{p+q}(\mathcal{F}, k)$. Note that as all $\mathcal{H}^j(\mathcal{F})$ are supported in 0 and hence are CM sheaves, we have $\mathrm{Ext}^i(\mathcal{H}^j(\mathcal{F}), k) \neq 0$ if and only if $i \in [0, c]$ and $\mathcal{H}^j(\mathcal{F}) \neq 0$. Let now $j \leq 0$ be minimal such that $\mathcal{H}^j(\mathcal{F}) \neq 0$. Then $\mathrm{Ext}^c(\mathcal{H}^j(\mathcal{F}), k) \neq 0$, and since this group lies in the upper right-hand corner of non-vanishing terms in the E_2 page, we find $\mathrm{Ext}^{c-j}(\mathcal{F}, k) =$

$\mathrm{Ext}^c(\mathcal{H}^j(\mathcal{F}), k) \neq 0$. Hence we must have $j = 0$. Similarly let $m \geq 0$ be maximal such that $\mathcal{H}^m(\mathcal{F}) \neq 0$, then $\mathrm{Ext}^{-m}(\mathcal{F}, k) \neq 0$ so $m = 0$. Hence $\mathcal{F}$ is a sheaf, and, having 0-dimensional support, must be CM.

Let now $n > c$, and let $\mathbb{A}^{n-1} \subset \mathbb{A}^n$ be a generic hyperplane through the origin. Then $\mathcal{F}|_{\mathbb{A}^{n-1}}$, as a complex on $\mathbb{A}^{n-1}$, satisfies the assumptions of the lemma, hence must be a CM sheaf.

Using the long exact sequence associated with $\mathcal{F}(-1) \to \mathcal{F} \to \mathcal{F}|_{\mathbb{A}^{n-1}}$ now shows that the maps $\mathcal{H}^i(\mathcal{F}(-1)) \to \mathcal{H}^i(\mathcal{F})$ must be surjections for $i \neq 0$. By the $\mathbb{C}^*$-equivariance of these maps and Nakayama's lemma, we must then have $\mathcal{H}^i(\mathcal{F}) = 0$ for all $i \neq 0$.

Since

$$n - c = \mathrm{depth}_0(\mathcal{F}|_{\mathbb{A}^{n-1}}) + 1 \leq \mathrm{depth}_0(\mathcal{F}) \leq \dim Z = n - c,$$

the sheaf $\mathcal{F}$ is CM at the point 0. Since the locus where $\mathcal{F}$ is Cohen–Macaulay is open and $\mathbb{C}^*$-invariant, we see that $\mathcal{F}$ is Cohen–Macaulay at every point of $\mathbb{A}^n$. $\square$

Proposition 3.28. *The dga A' has homology only in degree 0, and its homology is a Cohen–Macaulay sheaf on $\wedge^2 V$.*

Proof. The algebra A is, by construction, supported on the locus $\widetilde{\mathrm{Pf}_s} \subset \wedge^2 V$. The codimension of this locus is $\binom{v-2s}{2} = \binom{2q+1}{2}$. By Lemma 3.25 the dga A' is supported on this same locus, since the complement of U is contained in $\widetilde{\mathrm{Pf}_s}$. Lemma 3.26 implies that $\mathcal{A}'|_0$ has homology concentrated in $[-\binom{2q+1}{2}, 0]$, and so Lemma 3.27 applies. $\square$

So we may replace A' with its homology, and declare that A' is a Cohen–Macaulay algebra.

Lemma 3.29. *If $\mathcal{E}$ and $\mathcal{F}$ are Cohen–Macaulay modules on a regular variety X with $\mathrm{supp}(\mathcal{E}) = \mathrm{supp}(\mathcal{F}) = Z$ for an irreducible $Z \subset X$, and $\phi : \mathcal{E} \to \mathcal{F}$ is a homomorphism which is an isomorphism away from a locus $Z' \subset Z$ of codimension 2, then ϕ is an isomorphism.*

Proof. Let $\mathcal{K} = \ker(\phi)$ and $\mathcal{C} = \mathrm{cok}(\phi)$. If $\mathcal{K} \neq 0$, then for any $x \in \mathrm{supp}(\mathcal{K})$, we would have $\mathrm{depth}\, \mathcal{E}_x \leq \dim \mathrm{supp}\, \mathcal{K}_x \leq \dim Z'$. This contradicts the CM-ness of $\mathcal{E}$, hence $\mathcal{K} = 0$. It follows that $\mathrm{pd}(\mathcal{C}) \leq \max(\mathrm{pd}(\mathcal{F}) + 1, \mathrm{pd}(\mathcal{E})) = c + 1$, where c is the codimension of Z in X. On the other hand, if $\mathcal{C} \neq 0$, then

$$\mathrm{pd}(\mathcal{C}) \geq \dim X - \max_{x \in X} \mathrm{depth}(\mathcal{C}_x) \geq \mathrm{codim}_X \mathrm{supp}(\mathcal{C}) \geq c + 2,$$

using the Auslander–Buchsbaum formula. Hence we must have $\mathcal{C} = 0$, and ϕ must be an isomorphism. $\square$

Setting $\mathcal{E} = A'$ and $\mathcal{F} = A$ and $\phi = \Phi^\dagger$, and noting that the complement of U is $\widetilde{\mathrm{Pf}}_{s-1}$ which has codimension $2(v - 2s) + 1$ in $\widetilde{\mathrm{Pf}}_s$, we immediately obtain:

Proposition 3.30. *We have an isomorphism:*

$$\Phi^\dagger : A' \xrightarrow{\sim} A$$

This could perhaps be viewed as a form of Koszul duality between the algebra A and the curved algebra $(B \otimes \mathcal{O}_{\wedge^2 V}, W)$.

3.5.3. *Completing the proof.* We can now complete the proof our 'Hori duality' statement, Theorem 3.2. The preceding proposition (Proposition 3.30) essentially proves the special case $L = 0$.

Proposition 3.31. *We have an equivalence*

$$\Phi : \mathsf{Br}(\mathcal{X}) \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$$

of categories over $[\wedge^2 V/\mathbb{C}^*]$.

Proof. The algebra A was defined as the endomorphism algebra $\mathrm{End}_{\mathcal{X}}(T)$ of a vector bundle. The bundle T is self-dual (up to line bundles), so A is isomorphic to A^{op} (up to grading), and the functor $\mathrm{Hom}_{\mathcal{X}}(-,T)$ gives an equivalence between $\mathsf{Br}(\mathcal{X})^{op}$ and the derived category $D^b(\text{mod-}A)$ of right A-modules. If we let $T' = \oplus_{\delta\in Y^{q,s}}\mathcal{P}_\delta$, then the functor $\Phi^\dagger$ maps T' to T and induces an isomorphism between $A' = \mathrm{End}(T')$ and A (Proposition 3.30). It follows that the category $D^b(\text{mod-}A')$ has finite global dimension, so we have a well-defined functor

$$\mathrm{Hom}(-,T') : \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)^{op} \longrightarrow D^b(\text{mod-}A')$$

with adjoint $\otimes_{A'}T'$, and this is an equivalence by Lemma 3.24. The result follows. □

From here we prove the case of general L, with another application of Knörrer periodicity.

Choose any $L \subset \wedge^2 V^\vee$, and equip it with a weight 1 action of the group $\mathbb{C}^* = \mathrm{GSp}(S)/\mathrm{Sp}(S) = \mathrm{GSp}(Q)/\mathrm{Sp}(Q)$ and a weight 2 R-charge (as in Section 3.1). Form the stacks $\mathcal{X}\times_{\mathbb{C}^*}L$ and:

$$\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L = \left[\mathrm{Hom}(V,Q)\times\wedge^2 V\times L \,/\, \mathrm{GSp}(Q)\right]$$

Equip the former with the zero superpotential, and the latter with the superpotential W (pulled-up from $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V$), then they are both Landau–Ginzburg B-models. We can define a category $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L)$ by our usual grade-restriction-rule at the origin, this is generated by vector bundles and agrees with the triangulated closure of the pull-up of $\mathsf{Br}(\mathcal{X})$. On $\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L$ it doesn't make sense to apply a grade-restriction rule at the origin because the hypothesis of Lemma 2.6 doesn't hold, so instead we define

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L,\, W)$$

as the triangulated closure of the pull-up of $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V,\, W)$ (or as the matrix factorizations on our usual set of vector bundles). Then Proposition 3.31 immediately implies that we have an equivalence

$$\Phi : \mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L) \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L,\, W)$$

of categories relative to the base $[\wedge^2 V\times L \;/\; \mathbb{C}^*]$. This base carries a canonical quadratic superpotential which we'll call W', if we pull this up to $\mathcal{X}\times_{\mathbb{C}^*}L$ it becomes the W' already defined. If we pull it up to the other side we can add it on to the existing W, getting a 'perturbed' superpotential $W + W'$. It's immediate that Φ induces a functor between the categories with this extra W' added in, we claim that in fact:

Lemma 3.32. *The functor*

$$\Phi : \mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W') \longrightarrow \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L,\, W+W')$$

is an equivalence.

Morally the reason this is true is that W' defines an (unobstructed) class in Hochschild cohomology for both sides, and we are deforming both categories along this class. Since the original categories are equivalent, the deformed categories must also be equivalent. Since we are ignorant of the necessary foundations to state this argument precisely, we present an ad-hoc proof instead.

Proof. Suppose we have objects

$$(E, d_E) \in \mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L,\, W') \qquad \text{and} \qquad (F, d_F) \in \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W + W')$$

where E and F are (possibly infinite) direct sums of the usual vector bundles. Let's introduce an additional $\mathbb{C}^*$ action coming from rescaling just the L directions. Since E and F are non-equivariantly trivial we can canonically lift this extra action to them, using the trivial action. Then we can decompose $d_E \in \mathrm{End}(E)$ and $d_F \in \mathrm{End}(F)$ with respect to this extra grading as

$$d_E = (d_E)_0 + (d_E)_{>0} \qquad \text{and} \qquad d_F = (d_F)_0 + (d_F)_{>0}$$

(they cannot have negative terms), and it follows that $((d_E)_0)^2 = 0$ and $((d_F)_0)^2 = W1_F$. So the pair $(E, (d_E)_0)$ defines an object in $\mathsf{Br}(\mathcal{X})$, and the pair $(F, (d_F)_0)$ defines an object in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W)$; we'll denote these objects by $\widehat{E}$ and $\widehat{F}$. Moreover, it's evident that

$$\Phi \widehat{E} = \widehat{\Phi E} \qquad \text{and} \qquad \Phi^\dagger \widehat{F} = \widehat{\Phi^\dagger F}$$

since the kernel is trivial in the L directions.

Now take two objects $E_1, E_2 \in \mathsf{Br}(\mathcal{X} \times_{\mathbb{C}^*} L,\, W')$. We have a chain map

$$\Phi : \mathrm{Hom}(E_1, E_2) \longrightarrow \mathrm{Hom}(\Phi E_1, \Phi E_2) \tag{3.33}$$

These are complexes of modules over the ring of functions on $\wedge^2 V \times L$, but we will instead view them as double-complexes of vector spaces, using our extra grading. If we look at the first pages of the associated spectral sequences we see the map

$$\Phi : \mathrm{Hom}(\widehat{E_1}, \widehat{E_2}) \longrightarrow \mathrm{Hom}(\Phi \widehat{E_1}, \Phi \widehat{E_2})$$

and this is a quasi-isomorphism. Therefore (3.33) is also a quasi-isomorphism, so Φ is fully-faithful. A similar argument shows that the adjoint $\Phi^\dagger$ is also fully-faithful, hence Φ is an equivalence.

□

Since here W' is just the pairing between L and $\wedge^2 V / L^\perp$, Knörrer periodicity gives an equivalence:

$$D^b(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W + W') \xrightarrow{\sim} D^b(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)$$

This can be chosen to be linear over $[\wedge^2 V / \mathbb{C}^*]$; indeed the relative product over this base is $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp \times_{\mathbb{C}^*} L$ which is a maximal isotropic subbundle for W'.

Lemma 3.34. *Knörrer periodicity induces an equivalence*

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W + W') \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)$$

of categories over $[\wedge^2 V / \mathbb{C}^*]$.

Proof. An object $\mathcal{E} \in \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W + W')$ can be represented as a matrix factorization built from the infinite set of vector bundles corresponding to $Y_{q,s}$. If we restrict to $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp \times_{\mathbb{C}^*} L$ and push down we get an (infinite-rank) matrix factorization which satisfies the grade-restriction rule at the origin, so it lies in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)$.

Going in the other direction, choose a splitting $\wedge^2 V = L^\perp \oplus L^\vee$, and correspondingly write $W = W_1 + W_2$. Then we have a pull-up functor:

$$D^b(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W) \longrightarrow D^b(\mathcal{Y} \times_{\mathbb{C}^*} \wedge^2 V \times_{\mathbb{C}^*} L,\, W_1)$$

The sky-scraper sheaf along $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp \times_{\mathbb{C}^*} L$ can be viewed as a curved dg-sheaf for the superpotential $W_2 + W'$, and it's equivalent to a Koszul-type matrix factorization whose underlying vector bundle is the exterior algebra on L. The inverse to our Knörrer periodicity functor can be described as 'pull up, then tensor with this matrix factorization'. Since L is trivial as an $\mathrm{Sp}(Q)$-representation this functor sends objects in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp, W)$ to objects in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}\wedge^2 V\times_{\mathbb{C}^*}L,\, W+W')$. □

This completes the proof of Theorem 3.2.

4. The projective duality

Recall (from Section 3.2) that we have subcategories

$$\mathsf{Br}(\mathcal{X}^{ss})\subset D^b(\mathcal{X}^{ss})\qquad\text{and}\qquad \mathsf{Br}(\mathcal{Y}^{ss})\subset D^b(\mathcal{Y}^{ss})$$

defined as the triangulated subcategories generated by the images of $\mathsf{Br}(\mathcal{X})$ and $\mathsf{Br}(\mathcal{Y})$; or equivalently as the subcategories generated by the vector bundles corresponding to the sets $Y_{s,q}$ and $Y_{q,s}$. In this section we examine these categories in more detail.

In Section 4.1 we show that various natural notions of base change for these categories from $\mathcal{X}^{\mathrm{ss}}$ to a slice $\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp}$ agree, under a genericity assumption.

In Section 4.2 we find explicit descriptions of the categories $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(L^\perp\setminus 0), W)$ and $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}\times_{\mathbb{C}^*}L^\perp, W)$ as 'window' subcategories of $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp, W)$. The first case is handled by the general results of [HL15, BFK12], while the second case requires more work.

Finally, Section 4.3 shows how these window results together with Hori duality imply that $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ is HP dual to $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$.

4.1. Slicing the non-commutative resolution.

We saw in Section 3.2 that $\mathsf{Br}(\mathcal{X})$ is equivalent to the derived category of the graded algebra A, which is a ($\mathbb{C}^*$-equivariant) non-commutative resolution of the cone $\widetilde{\mathrm{Pf}}_s$. If we restrict to the complement of the origin then A becomes a sheaf of algebras on the projective variety Pf_s, and by Cor. 3.6 we have an equivalence

$$\mathsf{Br}(\mathcal{X}^{ss})\xrightarrow{\sim} D^b(\mathrm{Pf}_s, A).$$

So the category $\mathsf{Br}(\mathcal{X}^{ss})$ is a non-commutative resolution of Pf_s.

If we pick a subspace $L\subset\wedge^2 V$, then we have our category $\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W')$, which after deleting the fibre over the origin in $\wedge^2 V$ gives a category $\mathsf{Br}(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$. In this section we'll show how this relates to a non-commutative resolution of the slice $\mathrm{Pf}_s\cap\mathbb{P}L^\perp$.

Let's delete all the singularities in $\widetilde{\mathrm{Pf}}_s$ (*i.e.* restrict to the open set $U\cap\widetilde{\mathrm{Pf}}_s$), so that $\mathcal{X}$ becomes equivalent to the quasi-projective variety $\mathrm{Pf}_s^{\mathrm{sm}}$. Here the subcategory $\mathsf{Br}(\mathcal{X})$ becomes the whole of the derived category $D^b(\mathrm{Pf}_s^{\mathrm{sm}})$, since A is a trivial Azumaya algebra on this subset. The stack $\mathcal{X}\times_{\mathbb{C}^*}L$ becomes the total space of the vector bundle $L(-1)$ over $\mathrm{Pf}_s^{\mathrm{sm}}$. The embedding $\mathrm{Pf}_s^{\mathrm{sm}}\hookrightarrow\mathbb{P}(\wedge^2 V)$ gives a canonical section of the dual bundle $L^\vee(1)$, and the superpotential W' is just the pairing of this section with the fibre co-ordinate. If we assume that L is generic, then the section is transverse, and the critical locus of W' is the zero locus of the section, which is the slice $\mathrm{Pf}_s\cap\mathbb{P}L^\perp$. This is a standard setup for 'global Knörrer periodicity', and we have an equivalence

$$D^b(\mathrm{Pf}_s^{\mathrm{sm}}\cap\mathbb{P}L^\perp)\xrightarrow{\sim} D^b(\mathrm{Tot}(L(-1)),\, W')$$

(see for example [Shi12, Hir16] – note that the R-charge is acting fibre-wise on this vector bundle as required). We want to extend this fact over the singular locus in Pf_s.

The fibre product $\mathcal{X}|_{L^\perp}$ is a quotient of a cone inside $\mathrm{Hom}(S,V)$, and if we intersect this with $\mathcal{X}^{ss}$ we get a stack $\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}$ mapping to the singular variety $\mathrm{Pf}_s\cap\mathbb{P}L^\perp$.

We define $\mathsf{Br}(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp})$ to be the subcategory of $D^b(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp})$ consisting of those objects which land in $\mathsf{Br}(\mathcal{X}^{ss})$ under the pushforward functor.

Lemma 4.1. *Assume that $\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}$ has the expected dimension. Then we have an equivalence*
$$D^b(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}) \xrightarrow{\sim} D^b(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$$
inducing an equivalence:
$$\mathsf{Br}(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}) \xrightarrow{\sim} \mathsf{Br}(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$$

Proof. This is very similar to Lemma 3.34, but we can be a bit slicker in this case. Under Knörrer periodicity, the push-forward functor from $D^b(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp})$ to $D^b(\mathcal{X}^{ss})$ corresponds to the restriction functor from $D^b(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$ to $D^b(\mathcal{X}^{ss})$. It follows immediately that the brane subcategories are mapped to each other. □

We can also restrict the algebra A to the subspace $L^\perp$, giving an algebra defined over $\widetilde{\mathrm{Pf}_s}\cap L^\perp$. If we further delete the origin, we get a sheaf of algebras $A|_{\mathbb{P}L^\perp}$ on the projective variety $\mathrm{Pf}_s\cap\mathbb{P}L^\perp$. The non-commutative analogue of the Bertini theorem [RSVdB17] ensures that for generic $L^\perp$ this sheaf of algebras is a non-commutative resolution.

Lemma 4.2. *Assume that $\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}$ has the expected dimension. Then we have an equivalence:*
$$\mathsf{Br}(\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}) \xrightarrow{\sim} D^b\big(\mathrm{Pf}_s\cap\mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp}\big)$$

Proof. Let $\pi : \mathcal{X}^{\mathrm{ss}}\times_{\mathbb{C}^*}L \to \mathbb{P}(\wedge^2 V)$ be the projection. Let $T_L \in \mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$ be the restriction of T from (3.3). Using the assumption that $\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp}$ has the expected dimension, we find that $\pi_*\mathcal{H}om(T_L,T_L)=A|_{\mathbb{P}L^\perp}$ (where homomorphisms are derived).

We now have an adjoint pair of functors $F_L=\pi_*\mathcal{H}om(T_L,-)$ and $G_L=-\otimes_A T_L$ going between $D(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$ and $D(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$. A standard computation shows that G_L is fully faithful, and we must show that it sends $D^b(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$ surjectively onto $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$.

Consider first the case $L=0$. Then Corollary 3.6 shows that G_0 gives an equivalence $D^b(\mathrm{Pf}_s, A)\xrightarrow{\sim}\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$.

Let now L be any subspace, let $i:\mathrm{Pf}_s\cap\mathbb{P}L^\perp\to\mathrm{Pf}_s$ and $j:\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp}\to\mathcal{X}^{\mathrm{ss}}$ denote the inclusions, and let $i_*: D(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})\to D(\mathrm{Pf}_s,A)$ be the pushforward functor. There is an equality of functors
$$j_*\circ G_L = G_0\circ i_* : D(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})\to D(\mathcal{X}^{ss}).$$
Now as $G_0\circ i_*$ sends $D^b(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$ to $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$, and $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{PL^\perp})=(j_*)^{-1}(\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}))$, it follows that G_L sends $D^b(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{PL^\perp})$ to $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$.

If $\mathcal{E}\in\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$, then $F_L(\mathcal{E})=0$ implies that
$$i_*\pi_*\mathcal{H}om(T_L,\mathcal{E})=\pi_*\mathcal{H}om(T_0,j_*\mathcal{E})=F_0(j_*(\mathcal{E}))=0,$$
hence $j_*(\mathcal{E})=0$, and so $\mathcal{E}=0$. Thus we find that $\ker F_L\cap\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})=0$, which by [Kuz07, Thm. 3.3] means that $G_L(D^b(\mathrm{Pf}_s\cap\mathbb{P}L^\perp, A|_{PL^\perp}))=\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp})$. □

Corollary 4.3. *Assume that $\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}$ has the expected dimension. Then*
$$\mathsf{Br}(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')\cong D^b\big(\mathrm{Pf}_s\cap\mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp}\big).$$

If $\mathcal{X}^{ss}|_{\mathbb{P}L^\perp}$ fails to have the expected dimension then it seems appropriate to view $\mathsf{Br}(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$ as the 'correct' base-change of the category $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ to $\mathbb{P}L^\perp$.

We remark that these results also hold 'in the affine case', *i.e.* if we don't restrict to $\mathcal{X}^{ss}$ then it's still true that
$$\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W)\cong D^b\big(\big[\widetilde{\mathrm{Pf}_s}\cap L^\perp\,/\,\mathbb{C}^*\big],\, A|_{L^\perp}\big),$$
under the stronger assumption that $\mathcal{X}|_{L^\perp}$ has the expected dimension.

Now we examine the Serre functor on these categories.

Proposition 4.4. *Suppose that* $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$ *has the expected dimension, and that* $A|_{\mathbb{P}L^\perp}$ *is a non-commutative resolution of it. Then there exists a Serre functor on* $D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$ *given by:*

$$\mathcal{E} \mapsto \mathcal{E} \otimes \mathcal{O}_{\mathbb{P}(\wedge^2 V)}(sv - \dim L)[\dim \mathrm{Pf}_s \cap \mathbb{P}L^\perp]$$

Note that for generic choice of L both conditions hold. In particular, for generic L with $\dim L = sv$, the category $D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$ is Calabi-Yau.

Proof. Note that $\mathcal{O}_{\mathbb{P}(\wedge^2 V)}(sv - \dim L)$ is the canonical bundle on $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$. By [YZ06, Example 6.4], and using the fact that A is a maximal CM sheaf, we find that $A(sv - \dim L^\perp)|_{\mathbb{P}L^\perp}$ is a rigid dualising complex for $(\mathrm{Pf}_s \cap \mathbb{P}L^\perp, A|_{\mathbb{P}L^\perp})$. The claim then follows from [YZ06, Prop. 6.14]. □

There are of course exactly analogous results for $\mathcal{Y}$: we have an equivalence

$$\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp, W) \xrightarrow{\sim} D^b(\mathrm{Pf}_q \cap \mathbb{P}L,\, B|_{\mathbb{P}L})$$

under the assumption that $\mathcal{Y}^{ss}|_L$ has the expected dimension. There are two subtleties to point out:

- Until this point there has been no harm in continuing to regard everything as being relative to the base $[\wedge^2 V/\mathbb{C}^*]$, we can view the category $D^b\big([\widetilde{\mathrm{Pf}_s} \cap L^\perp / \mathbb{C}^*],\, A|_{L^\perp}\big)$ as enriched over $D^b(\mathrm{Pf}_s)$ if we wish, and the equivalence to $\mathsf{Br}(\mathcal{X}^{ss} \times_{\mathbb{C}^*} L,\, W')$ respects this. However, at this point we are forced to abandon this point of view, because the category $D^b(\mathrm{Pf}_q \cap \mathbb{P}L,\, B|_{\mathbb{P}L})$ is not defined relative to this base. We could view this category (and the above equivalence) as being relative to $[\wedge^2 V^\vee/\mathbb{C}^*]$ instead, but in the next section we will have to abandon that too.
- Recall that we can choose the R-charge on $\mathcal{Y}^{\mathrm{ss}}$ to be explicitly trivial (see Remark 3.1), then the algebra B will be concentrated in homological degree zero.

4.2. **Windows.** Since the equivalence of Theorem 3.2 is defined relative to the base $[\wedge^2 V/\mathbb{C}^*]$, we can restrict it to the complement of the origin and get:

Corollary 4.5. *We have an equivalence:*

$$\mathsf{Br}(\mathcal{X}^{\mathrm{ss}} \times_{\mathbb{C}^*} L,\, W') \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0),\, W)$$

Proof. This follows formally from Theorem 3.2, because these categories are by definition the subcategories generated by the images of the corresponding brane subcategories on $\mathcal{X} \times_{\mathbb{C}^*} L$ and $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp$. □

If L is sufficiently generic then by Corollary 4.3 the first category is equivalent to $D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp})$, our non-commutative resolution of the variety $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$. However, our non-commutative resolution for the dual slice $\mathrm{Pf}_q \cap \mathbb{P}L$ is equivalent to the category $\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W)$. So we need to understand how the categories $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0),\, W)$ and $\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W)$ are related.

The stacks $\mathcal{Y} \times_{\mathbb{C}^*} (L^\perp \setminus 0)$ and $\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp$ are related by variation of GIT stability; they are the two possible semi-stable loci in the ambient stack $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp$. We can use the technique of 'windows' [Seg11, HL15, BFK12] to compare them, by lifting their associated categories to subcategories defined on the ambient stack.

The existence of such lifts has been worked out in large generality in the mentioned literature, but our situation is complicated by two factors: we really care about the B-brane subcategories of each stack rather than the full derived category, and we want the window categories to be of a specific form, in order that they be comparable as subcategories of $D^b(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)$. It will therefore take us some work to construct these lifts.

Let us now introduce some notation for GSp(Q)-irreps. To describe the highest weights of these irreps, let's first define $\Delta : \mathbb{C}^* \to \mathrm{GSp}(Q)$ to be the diagonal 1-parameter subgroup $\Delta(t) = t \cdot 1_Q$, then a weight of GSp(Q) can be described by a pair (δ, k), where δ is a weight of Sp(Q) and k is a weight of Δ. Under this description, all GSp(Q)-weights satisfy $\sum_i \delta_i + k \equiv 0 \bmod 2$. We'll write $\mathbb{S}^{\langle \delta, k\rangle} Q$ for the irrep with highest weight (δ, k); in the notation used previously, we have

$$\mathbb{S}^{\langle \delta,k\rangle} Q = \mathbb{S}^{\langle \delta\rangle} Q \otimes \langle \omega_Q \rangle^{(k-\sum \delta_i)/2}.$$

Our subcategory $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W) \subset D^b(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)$ is defined by a rule which restricts the allowed Sp(Q)-representations occuring at the origin $0 \in \mathrm{Hom}(V, Q) \times L^\perp$. We now define some smaller subcategories: for an interval $I \subset \mathbb{Z}$, we define

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_I \;\subset\; \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)$$

to be the full subcategory of objects $\mathcal{E}$ such that the Δ-weights of $h_\bullet(\mathcal{E}|_0)$ all lie in I. By Lemma 2.6, this is equivalent to requiring that $\mathcal{E}$ can be represented by a matrix factorization built only from the vector bundles $\mathbb{S}^{\langle\delta,k\rangle}Q$, where $\delta \in Y_{q,s}$ and $k \in I$.

Note that these subcategories are not defined relative to any base and are not preserved by tensoring with line bundles.

General techniques give the first window result:

Proposition 4.6. *Set $l' = \dim L^\perp$. For any $n \in \mathbb{Z}$, the restriction functor*

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')} \longrightarrow \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$$

is an equivalence.

Proof. Define a subcategory

$$D^b(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')} \;\subset\; D^b(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)$$

by taking only objects whose Δ-weights at the origin lie in the interval $[n, n+2l')$. By general theory [Seg11, HL15, BFK12], the restriction functor

$$D^b(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')} \longrightarrow D^b(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$$

is an equivalence – note that the width of this interval here is $2l'$ (instead of l') because Δ acts with weight 2 on $L^\perp$. So we just need to argue that this equivalence matches up the brane subcategories.

If an object lies in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')}$ then (by definition) it restricts to give an object in $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$. This shows that restriction gives an embedding:

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')} \hookrightarrow \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$$

Given an object $\mathcal{E} \in \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$, it is the restriction of some object $\widetilde{\mathcal{E}} \in \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)$. The general recipe of [Seg11, HL15, BFK12] gives a way to modify $\widetilde{\mathcal{E}}$ to a new object $\widetilde{\mathcal{E}}'$, which still restricts to $\mathcal{E}$, but which lies in the subcategory $D^b(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')}$. The object $\widetilde{\mathcal{E}}'$ is constructed by taking cones over objects supported at $\mathcal{Y} \times 0$. It is easy to see that this process of taking cones will not introduce new Sp(Q)-weights, and so since all the Sp(Q)-weights of $\widetilde{\mathcal{E}}$ lie in $Y_{q,s}$, the same will be true of $\widetilde{\mathcal{E}}'$. Hence $\widetilde{\mathcal{E}}' \in \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')}$, which proves that the functor $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[n,n+2l')} \to \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} (L^\perp \setminus 0), W)$ is essentially surjective. □

For the other GIT quotient, we have:

Theorem 4.7. *The restriction functor induces an equivalence:*

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[-qv,qv)} \xrightarrow{\sim} \mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*} L^\perp, W)$$

This is the main technical result of Section 4, we will prove it as Propositions 4.13 and 4.18 in the next two sections. Note that this does not follow from the general theory of [HL15, BFK12], which would tell us to consider a Kempf–Ness stratification of $\mathcal{Y} \setminus \mathcal{Y}^{ss}$.

Set $n = -qv$ in Proposition 4.6. Our two subcategories of $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)$ are then contained one inside the other; we have

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)_{[-qv, 2l'-qv)} \;\subset\; \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)_{[-qv, qv)}$$

if $l' \leq qv$, and vice-versa if $l' \geq qv$. Combining this with our other results, we get the following 'HPD lite' statement:

Corollary 4.8. *If $l' \leq qv$ we have a fully faithful functor*

$$\mathsf{Br}(\mathcal{X}^{ss} \times_{\mathbb{C}^*} L,\, W') \hookrightarrow \mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W),$$

and if $l' \geq qv$ we have a fully faithful functor

$$\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W) \hookrightarrow \mathsf{Br}(\mathcal{X}^{ss} \times_{\mathbb{C}^*} L,\, W')$$

If $l' = qv$ the two categories are equivalent.

If L and $L^\perp$ are sufficiently generic then (by Corollary 4.3) we may replace the two categories by $D^b(\mathrm{Pf}_s \cap \mathbb{P}L^\perp,\, A|_{\mathbb{P}L^\perp})$ and $D^b(\mathrm{Pf}_q \cap \mathbb{P}L,\, B|_{\mathbb{P}L})$.

4.2.1. *Fully faithfulness.* In this section we prove half of Theorem 4.7: that the restriction functor from $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)_{[-qv,qv)}$ to $\mathsf{Br}(\mathcal{Y}^{ss} \times_{\mathbb{C}^*} L^\perp,\, W)$ is fully faithful.

Lemma 4.9. *Let R be a graded local Gorenstein ring of dimension n with $\omega_R \cong R(-d)$, and write $k = R/\mathfrak{m}$ for the unique 1-dimensional R-module. If M is a maximal Cohen–Macaulay module on R then*

$$\mathrm{Ext}^i_R(k, M) = 0$$

for $i \neq n$ and:

$$\mathrm{Ext}^n_R(k, M)^\vee = M^\vee(-d) \otimes k$$

as graded k-vector spaces.

Proof. Let $i : \operatorname{Spec} k \to \operatorname{Spec} R$ be the inclusion. We have

$$i^!(M) = (i^*(\mathrm{RHom}(M, \omega_R[n])))^\vee,$$

e.g. by [Sta16, Tag 0AU2]. Since M is maximal Cohen–Macaulay, we have

$$\mathrm{RHom}(M, \omega_R[n]) \cong M^\vee \otimes \omega_R[n],$$

e.g. by [Eis95, Prop. 21.12]. The claim follows. □

Let π denote the map:

$$\pi : \mathcal{Y} \to \left[\widetilde{\mathrm{Pf}_q} \,/\, \mathbb{C}^*\right]$$

Lemma 4.10. *If $\mathcal{E}$ is a locally free sheaf on $\mathcal{Y}$, then $\pi_*(\mathcal{E}^\vee) \cong \pi_*(\mathcal{E})^\vee$.*

Proof. Since π does not send any divisor to a codimension ≥ 2 subset, the push-forward functor preserves reflexive sheaves by [Bri93, Prop. 1.3], and so $\pi_*(\mathcal{E}^\vee)$ is reflexive.

Over the smooth locus $\widetilde{\mathrm{Pf}}_q^{\mathrm{sm}}$ the map π is an equivalence, so within this locus $\pi_*(\mathcal{E}^\vee) \cong \pi_*(\mathcal{E})^\vee$. Let $j : \widetilde{\mathrm{Pf}}_q^{\mathrm{sm}} \hookrightarrow \widetilde{\mathrm{Pf}}_q$ be the inclusion. Since the singular locus has codimension ≥ 2, and both $\pi_*(\mathcal{E}^\vee)$ and $\pi_*(\mathcal{E})^\vee$ are reflexive, we have

$$\pi_*(\mathcal{E}^\vee) \cong j_*(\pi_*(\mathcal{E}^\vee)|_{\widetilde{\mathrm{Pf}}_q^{\mathrm{sm}}}) \cong j_*(\pi_*(\mathcal{E})^\vee)|_{\widetilde{\mathrm{Pf}}_q^{\mathrm{sm}}}) \cong \pi_*(\mathcal{E})^\vee,$$

using [Har80, Prop. 1.6]. □

Let $Z = \{0\} \subset \widetilde{\mathrm{Pf}}_q$ be the vertex of the cone, and let $H^\bullet_Z(-) : D^b([\widetilde{\mathrm{Pf}}_q \,/\, \mathbb{C}^*]) \to D^b(\mathrm{Vect})$ denote the functor of taking local cohomology at Z (and $\mathbb{C}^*$-invariants).

Lemma 4.11. *Take two irreps of* $\mathrm{GSp}(Q)$ *with highest weights* (α_1, k_1) *and* (α_2, k_2) *lying in the set* $Y_{q,s} \times [-qv, qv)$. *We have associated vector bundles* $\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q$ *and* $\mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q$ *on* $\mathcal{Y}$. *Let:*

$$N = \pi_*\big(\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q^\vee \otimes \mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q\big) \quad \in \mathrm{Coh}\big([\widetilde{\mathrm{Pf}}_q \,/\, \mathbb{C}^*]\big)$$

Then $H^*_Z(N(i)) = 0$ *for all* $i \geq 0$.

Proof. Let us restrict to the case $i = 0$, the general case is shown the same way. Decomposing the representation $\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q^\vee \otimes \mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q$ into irreps, we obtain a decomposition of N. We may write $H^\bullet_Z(N) = \lim_{\to} \mathrm{Hom}(\mathcal{O}/I^p_Z, N)$, where $\lim_{\to}$ denotes a homotopy colimit [Lip02, Prop. 1.5.3]. Using the short exact sequence $I^p_Z/I^{p+1}_Z \hookrightarrow \mathcal{O}/I^{p+1}_Z \twoheadrightarrow \mathcal{O}/I^p_Z$, the required vanishing then follows if $\mathrm{Hom}(I^p_Z/I^{p+1}_Z, N) = 0$ for all $p \geq 0$. As I^p_Z/I^{p+1}_Z is isomorphic to a direct sum of copies of $\mathcal{O}_Z(-p)$, it finally suffices to show that $\mathrm{RHom}(\mathcal{O}_Z, N(p)) = 0$ vanishes when $p \geq 0$.

Write n for the dimension of $\widetilde{\mathrm{Pf}}_q$, and recall that its canonical bundle is $\mathcal{O}(-2qv)$. By [ŠVdB15, Thm. 1.6.4] the module $N(p)$ is Cohen–Macaulay. By Lemma 4.9 we have that $\mathrm{Ext}^j(\mathcal{O}_Z, N(p)) = 0$ if $j \neq n$, and $\mathrm{Ext}^n(\mathcal{O}_Z, N(p))^{\mathbb{C}^*}$ is the dual of the invariants in $N^\vee(-p-2qv)|_Z$. Now by Lemma 4.10, we have $N^\vee = \pi_*(\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q \otimes \mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q^\vee)$, and this module is generated in degrees $\geq k_1 - k_2 > -2qv$. So for $p \geq 0$ the module $N^\vee(-p-2qv)$ is generated in positive degree, hence its restriction to Z has no $\mathbb{C}^*$-invariants. □

Lemma 4.12. *Let* $[\mathbb{A}^n \,/\, \mathbb{C}^*]$ *be a vector space with a diagonal* $\mathbb{C}^*$*-action of weight* -1. *Suppose that* $\mathcal{E} \in D^b([\mathbb{A}^n \,/\, \mathbb{C}^*])$ *has the property that* $(\mathcal{E}(p)|_0)^{\mathbb{C}^*} = 0$ *for all* $p \geq 0$. *Then* $\mathcal{E}$ *has no global sections.*

Proof. This follows easily from Lemma 2.6. □

Let's abuse notation and continue to use π for the map:

$$\pi : \mathcal{Y} \times_{\mathbb{C}^*} L^\perp \longrightarrow [\widetilde{\mathrm{Pf}}_q \times L^\perp \,/\, \mathbb{C}^*]$$

Proposition 4.13. *The restriction functor*

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)_{[-qv,qv)} \to \mathsf{Br}(\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^\perp,\, W)$$

is fully faithful.

Proof. Let $\mathcal{E}, \mathcal{F} \in \mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, W)_{[-qv,qv)}$. We can assume $\mathcal{E}$ and $\mathcal{F}$ are matrix factorizations whose underlying vector bundles are direct sums of the bundles $\{\mathbb{S}^{\langle \alpha, k \rangle}Q,\ (\alpha, k) \in Y_{q,s} \times [-qv, qv)\}$. The morphisms between $\mathcal{E}$ and $\mathcal{F}$ are just the homology of the chain-complex $\Gamma(\mathcal{E}^\vee \otimes \mathcal{F}) = (\mathcal{E}^\vee \otimes \mathcal{F})^{\mathrm{GSp}(Q)}$. If we restrict to $\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^\perp$ then the morphism are *a priori* more complicated, because we must take the derived global sections of $\mathcal{E}^\vee \otimes \mathcal{F}$. However, we claim that in fact the vector bundle $\mathcal{E}^\vee \otimes \mathcal{F}$ has the property that

$$R\Gamma(\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^\perp,\, \mathcal{E}^\vee \otimes \mathcal{F}) = \Gamma(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp, \mathcal{E}^\vee \otimes \mathcal{F})$$

so the morphisms between $\mathcal{E}$ and $\mathcal{F}$ do not change upon restriction. This is the statement of the proposition.

Now we prove the claim. Obviously it's sufficient to prove it for the summands of $\mathcal{E}$ and $\mathcal{F}$, so pick two vector bundles $\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q$ and $\mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q$ with (α_1, k_1) and (α_2, k_2) both in $Y_{q,s} \times [-qv, qv)$. Let $M = \pi_*(\mathbb{S}^{\langle \alpha_1, k_1 \rangle}Q^\vee \otimes \mathbb{S}^{\langle \alpha_2, k_2 \rangle}Q)$. We are claiming that the map

$$\Gamma(M) \to R\Gamma((\widetilde{\mathrm{Pf}}_q \setminus Z) \times L^\perp, M)$$

is an isomorphism, *i.e.* that the local cohomology $H^\bullet_{Z\times L^\perp}(M)$ vanishes. Arguing as in the proof of Lemma 4.11, this reduces to checking that $\mathrm{RHom}(\mathcal{O}_{Z\times L^\perp}, M(p)) = 0$ for all $p \geq 0$. In other words, if we let $i : Z\times L \hookrightarrow \widetilde{\mathrm{Pf}}_q \times L$ be the inclusion, we want to see that $i^! M(p)$ has no global sections when $p \geq 0$.

Everything here is flat over $L^\perp$, so we can use Lemma 4.11 to understand the restriction of $i^! M(p)$ to the origin in $L^\perp$. In particular, $i^! M(p)|_0$ has no invariants for $p \geq 0$. Now apply Lemma 4.12. □

4.2.2. *Essential surjectivity.* In this section we prove that the restriction functor from $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp, W)_{[-qv,qv)}$ to $\mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*} L^\perp, W)$ is essentially surjective.

Pick an irrep $\mathbb{S}^{\langle\alpha,k\rangle}Q$ with highest weight (α, k) lying in $Y_{q,s}\times\mathbb{Z}$. On $\mathcal{Y}$ we have a corresponding twist $\mathcal{O}_0 \otimes \mathbb{S}^{\langle\alpha,k\rangle}Q$ of the skyscraper sheaf at the origin, and by Lemma 3.8 we can project this into $\mathsf{Br}(\mathcal{Y})$ to get an object:

$$P_{\alpha,k} \in \mathsf{Br}(\mathcal{Y})$$

Up to a grading shift this object corresponds to the 'vertex simple' B-module at the vertex α. It also agrees with the restriction of the object $\mathcal{P}_\alpha$ from Section 3.5.1 to the slice $\mathcal{Y}\times 0$, again up to grading; c.f. Remark 3.23.

For any object $\mathcal{E} \in \mathsf{Br}(\mathcal{Y})$, if we restrict to the origin we get a $\mathrm{GSp}(Q)$-representation $h_\bullet(\mathcal{E}|_0)$, and by definition of $\mathsf{Br}(\mathcal{Y})$ this is a direct sum of irreps with highest weight in $Y_{q,s}\times\mathbb{Z}$. We also know (by Lemma 2.6) that $\mathcal{E}$ is equivalent to a matrix factorization on the associated vector bundle. For brevity, let's refer to the set of irreps that occur in $h_\bullet(\mathcal{E}|_0)$ as the *weights* of $\mathcal{E}$.

Our first task is to determine the weights of the object $P_{\alpha,k}$.

Lemma 4.14. *For any $\mathcal{E} \in \mathsf{Br}(\mathcal{Y})$, we have that (α, k) is a weight of $\mathcal{E}$ if and only if* $\mathrm{Hom}_{\mathcal{Y}}(\mathcal{E}, P_{\alpha,k}) \neq 0$.

Proof. By adjunction:

$$\mathrm{Hom}\big(\mathcal{E}|_0, \mathbb{S}^{\langle\alpha,k\rangle}Q\big) = \mathrm{Hom}_{\mathcal{Y}}\big(\mathcal{E}, \mathcal{O}_0 \otimes \mathbb{S}^{\langle\alpha,k\rangle}Q\big) = \mathrm{Hom}_{\mathcal{Y}}\big(\mathcal{E}, \mathcal{P}_{\alpha,k}\big)$$

□

Lemma 4.15. *The set of weights of $P_{\alpha,k}$ contains one copy of (α, k), and all the remaining weights are contained in the set $Y_{q,s} \times [k - 2qv, k)$.*

Proof. To compute the weights of $P_{\alpha,k}$ we need to resolve it by vector bundles associated to the set $Y_{q,s}\times\mathbb{Z}$. As a first step, replace $\mathbb{S}^{\langle\alpha,k\rangle}Q \otimes \mathcal{O}_0$ by a twist of the usual Koszul resolution on $\mathrm{Hom}(V, Q)$. This complex has a single copy of $\mathbb{S}^{\langle\alpha,k\rangle}Q$, and the remaining summands are of the form $\mathbb{S}^{\langle\alpha',k'\rangle}Q$ with $k' < k$.

Now we need to project into the subcategory $\mathsf{Br}(\mathcal{Y})$. If $\alpha' \in Y_{q,s}$ then $\mathbb{S}^{\langle\alpha',k'\rangle}Q$ is unaffected by this projection, however most summands will not be of this form and will become something more complicated. To understand what they become, we use the results of Špenko–Van den Bergh. They construct a large set of objects in $D^b(\mathcal{Y})$ which are orthogonal to $\mathsf{Br}(\mathcal{Y})$, and also find locally-free resolutions of them. Projecting the resolutions into $\mathsf{Br}(\mathcal{Y})$ gives exact sequences, which allows us to express the projections of vector bundles in terms of projections of other vector bundles. Using these exact sequences repeatedly, we can eventually express the projection of any vector bundle in terms of vector bundles associated to the set $Y_{q,s}\times\mathbb{Z}$ (this is how they prove that $\mathsf{Br}(\mathcal{Y})$ has finite global dimension).

To be precise, we consider the complexes which are denoted $C_{\lambda,\chi}$ in [ŠVdB15, p. 34] – taking χ as a $\mathrm{GSp}(Q)$-weight and λ a 1-parameter subgroup of $\mathrm{Sp}(Q)$ we obtain a complex on $\mathcal{Y}$. We only need to know one fact: the resulting exact sequences replace a bundle $\mathbb{S}^{\langle\alpha',k'\rangle}Q$ with bundles $\mathbb{S}^{\langle\alpha'',k''\rangle}Q$ where $k'' \leq k'$. This follows from [ŠVdB15, Lemma 11.2.1] and the fact that Δ acts with positive weights on $\mathrm{Hom}(V, Q)$. This proves that the set of weights of $P_{\alpha,k}$ contains one copy of (α, k) and the other weights satisfy $k' < k$.

Now we apply Serre duality for the category $\mathsf{Br}(\mathcal{Y})$ (Proposition 3.10), recalling that the Serre functor is given (up to a shift) by tensoring by $(\det Q)^{-v} = \mathbb{S}^{\langle \mathbf{0}, -2qv\rangle} Q$. By Lemma 4.14, if (α', k') is a weight of $P_{\alpha,k}$ then

$$\mathrm{Hom}_{\mathcal{Y}}(P_{\alpha,k}, P_{\alpha',k'}) \neq 0 \implies \mathrm{Hom}_{\mathcal{Y}}(P_{\alpha',k'}, P_{\alpha,k-2qv}) \neq 0$$

so $(\alpha, k - 2qv)$ is a weight of $P_{\alpha',k'}$. Hence (by the first part of the proof) we must have $k - 2qv \leq k'$. □

In fact the Serre duality argument proves slightly more: if (α', k') is a weight of $P_{\alpha,k}$ with k' minimal, then k' is exactly $k - 2qv$, and $\alpha' = \alpha$.

Lemma 4.16. *The objects $P_{\alpha,k} \in \mathsf{Br}(\mathcal{Y})$ all restrict to 0 in $\mathsf{Br}(\mathcal{Y}^{ss})$.*

Proof. Obviously the object $\mathcal{O}_0 \otimes \mathbb{S}^{\langle \alpha,k\rangle} Q$ is supported at the origin in $\widetilde{\mathrm{Pf}_q}$. Now observe that the projection $D^b(\mathcal{Y}) \to \mathsf{Br}(\mathcal{Y})$ is linear over $\widetilde{\mathrm{Pf}_q}$. □

Lemma 4.17. *For any $(\alpha, k) \in Y_{q,s} \times \mathbb{Z}$, the vector bundle $\mathbb{S}^{\langle \alpha,k\rangle} Q$ on $\mathcal{Y}^{ss}$ has a finite resolution by vector bundles associated to the set $Y_{q,s} \times [-qv, qv)$.*

Proof. Using Lemma 2.6, the claim is equivalent to $\mathbb{S}^{\langle \alpha,k\rangle} Q$ being in the image of the restriction functor $F : \mathsf{Br}(\mathcal{Y})_{[-qv,qv)} \to D^b(\mathcal{Y}^{\mathrm{ss}})$.

We argue by induction on $|k|$. The case when $|k| < qv$ and $k = -qv$ are obvious. If $k > qv$, we consider the following exact triangle in $D^b(\mathcal{Y})$:

$$\mathcal{E} \to \mathbb{S}^{\langle \alpha,k\rangle} Q \to P_{\alpha,k}.$$

The weights of $\mathcal{E}$ are precisely the weights of $P_{\alpha,k}$, minus (k, α). Hence any (α', k') appearing as a weight of $\mathcal{E}$ are such that $(\alpha', k') \in Y_{q,s} \times [k - 2qv, k)$, and hence $\mathcal{E}$ admits a resolution in terms of vector bundles $\mathbb{S}^{\langle \alpha',k'\rangle} Q$. By induction these vector bundles are in the in the image of F, and since $F(P_{k,\alpha}) = 0$, it follows that $\mathbb{S}^{\langle \alpha,k\rangle} Q$ is in the image of F.

If $k < -qv$, we argue similarly starting from the exact triangle

$$P^{\vee}_{\alpha,-k} \to \mathbb{S}^{\langle \alpha,k\rangle} Q \to \mathcal{F}$$

where the weights of $\mathcal{F}$ are contained in $Y_{q,s} \times (k, k + 2qv]$. □

Proposition 4.18. *The restriction functor*

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^{\perp}, W)_{[-qv,qv)} \to \mathsf{Br}(\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^{\perp}, W)$$

is essentially surjective.

Proof. Take $\mathcal{E} \in \mathsf{Br}(\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^{\perp}, W)$. By definition $\mathcal{E}$ is equivalent to the image of an object in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^{\perp}, W)$, so we can assume that $\mathcal{E}$ is a vector bundle built from the set $Y_{q,s} \times \mathbb{Z}$. This vector bundle (although not the twisted differential) is pulled up from $\mathcal{Y}^{ss}$, so it has a resolution by bundles from the set $Y_{q,s} \times [-qv, qv)$. By Proposition 4.13 this set of bundles has no higher Ext groups between them, so we can use the perturbation process of [ADS15, Lemma 4.10], [Seg11, Lemma 3.6] to conclude that $\mathcal{E}$ is equivalent to a matrix factorization built from this set. Then such a matrix factorization is automatically the restriction of a matrix factorization in $\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^{\perp}, W)_{[-qv,qv)}$ (because the twisted differential must extend). □

4.3. **The homological projective duality statement.** In this section, we produce a Lefschetz decomposition of $\mathsf{Br}(\mathcal{Y}^{ss})$, and show that $\mathsf{Br}(\mathcal{X}^{ss})$ is equivalent to the tautological HP dual defined in Section 2.1.

Lemma 4.19. *Let $(\alpha, k) \in Y_{q,s} \times \mathbb{Z}$. The vector bundle $\mathbb{S}^{\langle \alpha,k\rangle}Q$ on $\mathcal{Y}^{ss}$ is exceptional,* i.e. $\mathrm{REnd}(\mathbb{S}^{\langle\alpha,k\rangle}Q) = \mathbb{C}$.

Let $(\alpha_1, k_1), (\alpha_2, k_2) \in Y_{q,s} \times [-qv, qv)$, with $k_1 > k_2$. Then

$$\mathrm{RHom}_{\mathcal{Y}^{ss}}(\mathbb{S}^{\langle\alpha_1,k_1\rangle}Q, \mathbb{S}^{\langle\alpha_2,k_2\rangle}Q) = 0$$

Proof. By the fully faithfulness result of Proposition 4.13, we may replace the bundles on $\mathcal{Y}^{ss}$ with the corresponding bundles on $\mathcal{Y}$ and compute Hom spaces for these. The claims are then easy to verify. □

Let $\mathcal{A}_0 = \mathsf{Br}(\mathcal{Y})_{[-qv,-qv+1]}$, and let $\mathcal{A}_0(i) = \mathcal{A}_0 \otimes \mathcal{O}_{\mathbb{P}(\wedge^2 V^\vee)}(i)$.

Lemma 4.20. *There is a Lefschetz decomposition*

$$\mathsf{Br}(\mathcal{Y}^{ss}) = \langle \mathcal{A}_0, \mathcal{A}_0(1), \ldots, \mathcal{A}_0(qv-1)\rangle.$$

Proof. Recall that by Proposition 4.7 we have $\mathsf{Br}(\mathcal{Y}^{ss}) \cong \mathsf{Br}(\mathcal{Y})_{[-qv,qv)}$. Using this equivalence, the claim follows from Lemmas 4.17 and 4.19. □

Recall from Section 2.1 that for any $L \subset \wedge^2 V^\vee$, there is a tautological HP dual category $\mathcal{W}^\vee_{L^\perp}$ living over $\mathbb{P}L^\perp$ such that the categories $D^b(\mathcal{Y}^{ss} \times L, W)$ are related by Theorem 2.3.

Proposition 4.21. *For any linear subspace $L \subseteq \wedge^2 V^\vee$, there is an equivalence* $\mathsf{Br}((L^\perp \setminus 0) \times_{\mathbb{C}^*} \mathcal{Y}, W) \cong \mathcal{W}^\vee_{L^\perp}$.

This is really a family version of Proposition 4.6, over the base $\mathbb{P}L^\perp$.

Proof. Let

$$Z = \left[(L^\perp \setminus 0) \times \mathbb{C} \times \mathrm{Hom}(V, Q) \,/\, (\mathrm{GSp}(Q) \times \mathbb{C}^*)\right],$$

where $\mathrm{GSp}(Q)$ acts standardly on $\mathrm{Hom}(V, Q)$, with Δ-weight -2 on $\mathbb{C}$, and trivially on $L^\perp$. The factor $\mathbb{C}^*$ acts with weights -1 on $\mathbb{C}$ and 1 on $L^\perp$ Letting x be the coordinate of the $\mathbb{C}$-factor, we define a potential W on Z by $W = x \cdot W_{std}$, where W_{std} is the usual potential on $L^\perp \times \mathrm{Hom}(V, Q)$. We let the R-charge be scaling $(L^\perp \setminus 0)$ with weight 2 and fixing the other two factors.

We consider the open substacks

$$(L^\perp \setminus 0) \times_{\mathbb{C}^*} \mathcal{Y} = \left[(L^\perp \setminus 0) \times (\mathbb{C} \setminus 0) \times \mathrm{Hom}(V, Q) \,/\, (\mathrm{GSp}(Q) \times \mathbb{C}^*)\right]$$

and

$$X := \left[(L^\perp \setminus 0) \times \mathbb{C} \times \mathrm{Hom}(V, Q)^{ss} \,/\, (\mathrm{GSp}(Q) \times \mathbb{C}^*)\right].$$

By definition, $\mathcal{W}^\vee_{L^\perp} \subseteq D^b(X, W)$ is the subcategory of those objects $\mathcal{E}$ such that for each point $p \in \mathbb{P}L^\perp$, upon restriction to $\mathcal{X}^{ss} = p \times 0 \times \mathcal{X}^{ss} \subset X_L$, we have $\mathcal{E}|_{\mathcal{X}^{ss}} \in \mathcal{A}_0$.

The techniques of [BFK12, HL15] show that there is an inclusion

$$F : D^b((L^\perp \setminus 0) \times_{\mathbb{C}^*} \mathcal{Y}, W) \hookrightarrow D^b(Z, W)$$

whose image is the subcategory of those objects such that for every point $p \in \mathbb{P}L^\perp$, upon restriction to $\left[p \times 0 \times 0 \,/\, \mathrm{GSp}(Q)\right]$ we get an object with Δ-weights in $[-qv, -qv+1]$. The image of $\mathsf{Br}((L^\perp \setminus 0) \times_{\mathbb{C}^*} \mathcal{Y}, W)$ under F is the set of objects which after restriction to $\left[p \times 0 \times 0 \,/\, \mathrm{GSp}(Q)\right]$ have weights in $Y_{q,s}$.

We similarly get an inclusion $D^b(X, W) \to D^b(Z, W)$, where the image is the window subcategory of those objects which upon restriction to $\left[p \times 0 \times 0 \,/\, \mathrm{GSp}(Q)\right]$ have Δ-weights in $[-qv, qv)$. Restricting to a single point, we find that under this equivalence, $\mathcal{A}_0$ consists of the objects which restrict to the origin to lie in the window $Y_{q,s}\times[-qv, -qv+1]$.

Summing up, we see that the inclusion of $\mathsf{Br}((L^\perp\setminus 0)\times_{\mathbb{C}^*}\mathcal{Y})$ into $D^b(Z,W)$ has image precisely $\mathcal{W}^\vee_{L^\perp}$, which is what we wanted. ☐

Without any further hypotheses on L, we thus arrive at an HP duality statement:

Theorem 4.22. *Let $L\subset\wedge^2V^\vee$ be a linear subspace. We get semiorthogonal decompositions*

$$\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}\times_{\mathbb{C}^*}L^\perp,\,W)=\Big\langle\mathcal{C}_L,\;\mathcal{A}_{l'}(l'),\;\ldots,\;\mathcal{A}_{qv-1}(qv-1)\Big\rangle$$

and

$$\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}\times_{\mathbb{C}^*}L,\,W)=\Big\langle\mathcal{B}_{1-sv}(1-sv),\;\ldots,\;\mathcal{B}_{-l}(-l),\;\mathcal{C}_L\Big\rangle$$

where $l=\dim L$ and $l'=\dim L^\perp$.

Proof. Since $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ admits a full exceptional collection, it has a strong generator given by the direct sum of the objects in the exceptional collection. Hence by the main result of [BVdB03], every functor $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})\to D^b(\mathbb{C})^{op}$ is representable. Since $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ has a Serre functor (Prop. 4.4), this means that every functor $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})^{\mathrm{op}}\to D^b(\mathbb{C})^{\mathrm{op}}$ is also representable, hence $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ is saturated. As $D^b(\mathcal{Y}^{\mathrm{ss}})$ is ext finite, $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})\subseteq D^b(\mathcal{Y}^{\mathrm{ss}})$ is admissible.

Now combine Theorem 2.3, Proposition 4.21 and Corollary 4.5, observing that $\dim\wedge^2V^\vee=sv+qv$. ☐

This statement is an upgrade of Corollary 4.8. If $l\geq sv$ (so $l'\leq qv$) then the second semi-orthogonal decomposition contains only a single piece, so the category $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}}\times_{\mathbb{C}^*}L,\,W)=\mathcal{C}_L$ embeds as an admissible subcategory of $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}\times_{\mathbb{C}^*}L^\perp,\,W)$. If the inequalities are reversed then the embedding goes the other way, and if $l=sv$ (so $l'=qv$) then the categories are equivalent.

Combining this theorem with Lemma 4.2, we arrive at the claim that (Pf_q,B) is HP dual to (Pf_s,A):

Theorem 4.23. *If $L\subset\wedge^2V^\vee$ is a linear subspace such that $\mathcal{X}^{\mathrm{ss}}|_{\mathbb{P}L^\perp}$ and $\mathcal{Y}^{\mathrm{ss}}|_{\mathbb{P}L}$ have the expected dimensions, then we get semiorthogonal decompositions*

$$D^b(\mathrm{Pf}_q\cap\mathbb{P}L,\;B|_{\mathbb{P}L})=\Big\langle\mathcal{C}_L,\;\mathcal{A}_{l'}(l'),\;\ldots,\;\mathcal{A}_{qv-1}(qv-1)\Big\rangle$$

and

$$D^b(\mathrm{Pf}_s\cap\mathbb{P}L^\perp,\;A|_{\mathbb{P}L^\perp})=\Big\langle\mathcal{B}_{1-sv}(1-sv),\;\ldots,\;\mathcal{B}_l(l),\;\mathcal{C}_L\Big\rangle.$$

Remark 4.24. The category $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ is equipped with two natural Lefschetz decompositions: The first via the natural variant of Lemma 4.20, and the second obtained from the fact that it is the HP dual of $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$. We don't know if these two decompositions are the same, but this matching up is not needed for the relation of HP duality.

5. The case when $\dim V$ is even

5.1. **What is still true.**

In this section – as mentioned briefly in the introduction – we switch to the case when $v=\dim V$ is even, instead of odd. Many of our previous results continue to hold, but a few crucial ones fail. We will first summarize which parts work or do not work, and then go on to discuss one part (the 'windows' result) which holds in a modified form, but requires some extra effort to prove.

We keep all of our notation from before: the vector spaces V,S,Q, the stacks $\mathcal{X}$ and $\mathcal{Y}$, etc. We now set the dimension v of V to be:

$$v=2s+2q$$

Then Pf_s is still the classical projective dual to Pf_q. We continue to use $Y_{s,q}$ for the set of Young diagrams of height at most s and width at most q, and define the subcategory

$$\mathsf{Br}(\mathcal{X}\times L,\, W') \;\subset\; D^b(\mathcal{X}\times L,\, W')$$

in exactly the same way as before (and the same is true on the $\mathcal{Y}$ side). The crucial thing that changes is that, although $\mathsf{Br}(\mathcal{X})$ is still a non-commutative resolution of $\widetilde{\mathrm{Pf}_s}$ by [ŠVdB15], *it is no longer crepant* in general. This means that we still have an equivalence between $D^b(\mathcal{X})$ and the derived category of an algebra A (defined exactly as before), but A is no longer Cohen–Macaulay. All other statements of Section 3.2 continue to hold.

Section 3.3 in which we define and study the kernel is completely unchanged, since V plays essentially no role here. Section 3.4, in which we prove that Φ is generically an equivalence, also continues to work without modification (indeed the v even case was considered in our earlier papers on which this section is based).

However, our proof that Φ is an equivalence everywhere fails. We can still define our objects $\mathcal{P}_\delta$ (Section 3.5.1) and their endomorphism dga A', but since A is not Cohen–Macaulay we cannot prove Proposition 3.30. In fact the first failure is at Lemma 3.26, which means we cannot even prove that A' is an algebra.

The final steps of Section 3.5.3 continue to work, and the end result is that we have a pair of adjoint functors

$$\mathsf{Br}(\mathcal{X}\times_{\mathbb{C}^*}L,\, W') \rightleftarrows \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)$$

but we only know that they form an equivalence once we restrict to the open set $U \subset \wedge^2 V$ of bivectors having rank $\geq 2s$.

Now we move on to Section 4. Base-changing the above adjunction to the complement of the origin in $\wedge^2 V$, we get an adjunction between $\mathsf{Br}(\mathcal{X}^{ss}\times_{\mathbb{C}^*}L,\, W')$ and $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(L^\perp\setminus 0),\, W)$, which we know to be an equivalence over the open set $\mathbb{P}U \subset \mathbb{P}(\wedge^2 V)$. The results of Section 4.1 are unaffected, so if L is generic then the former category is equivalent to the derived category of the sheaf of algebras $A|_{\mathbb{P}L^\perp}$ on the slice $\mathrm{Pf}_s \cap \mathbb{P}L^\perp$. Since the intersection of Pf_s with $\mathbb{P}U$ is exactly the smooth locus Pf_s^{sm}, this shows that over the smooth locus we have

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(U\cap L^\perp),\, W) \;\cong\; D^b(\mathrm{Pf}_s^{sm}, A|_{\mathbb{P}L^\perp})$$

which is simply $D^b(\mathrm{Pf}_s^{sm})$. But once we include the singular locus we don't know whether these two categories are the same.

The real extra work – which will take up most of the remainder of Section 5 – concerns the 'windows' of Section 4.2, which allow us to compare $\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(L^\perp\setminus 0),\, W)$ with $\mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*}L^\perp,\, W)$. The window for the 'difficult phase' (Proposition 4.6) needs no adjusting, but we must modify the window for the 'easy phase' (Theorem 4.7).

Recall that the problem is essentially to lift the category $\mathsf{Br}(\mathcal{Y}^{ss})$ to an equivalent subcategory in $\mathsf{Br}(\mathcal{Y})$. When v was odd, we did this using 'prism' in the set of weights of $\mathrm{GSp}(Q)$, restricting the weights of $\mathrm{Sp}(Q)$ and the diagonal 1-parameter subgroup Δ separately. However, one sees already in the case $q=1$ that this does not work for v even, because in this case Kuznetsov's Lefschetz decomposition of $D^b(\mathrm{Gr}(V,2))$ is not rectangular. So our set of allowed weights must have a slightly more complicated shape.

Recall that irreps of $\mathrm{GSp}(Q)$ have highest weights (δ, k) for δ a dominant weight of $\mathrm{Sp}(Q)$ and k a weight of Δ such that $\sum_i \delta_i + k \cong 0 \bmod 2$. We define a subset

$$Y_{\mathrm{proj}} \subset Y_{q,s}\times\mathbb{Z}$$

as the set of weights (δ, k) such that either

- $k \in \big[-qv, (q-1)v\big)$, or
- $k \in \big[(q-1)v, qv\big)$ and $\delta \in Y_{q,s-1}$.

Then we define a corresponding full subcategory

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{\mathrm{proj}} \;\subset\; D^b(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)$$

of objects $\mathcal{E}$ such that $h_\bullet(\mathcal{E}|_0)$ contains only irreps from the set Y_{proj}. The analogue of Theorem 4.7 for the v even case is:

Theorem 5.1. *The restriction functor*

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{\mathrm{proj}} \;\longrightarrow\; \mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*}L^\perp,\, W)$$

is an equivalence.

We will prove this theorem in the next two sections, but first we finish discussing to what extent our main results continue to hold for the even case.

Let $S = Y_{q,s}\times\{-qv, -qv+1\} \subseteq Y_{\mathrm{proj}}$, and let $S' = Y_{q,s-1}\times\{-qv,-qv+1\} \subseteq Y_{\mathrm{proj}}$. Applying Theorem 5.1 with $L^\perp = 0$ and arguing as in Lemma 4.20, we see that $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ has a Lefschetz decomposition

$$\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}) = \Big\langle \mathcal{A},\, \mathcal{A}(1),\, \dots,\, \mathcal{A}(qv - \tfrac{1}{2}v),\, \mathcal{A}'(qv+\tfrac{1}{2}v),\, \dots,\, \mathcal{A}'(qv) \Big\rangle \tag{5.2}$$

where

$$\mathcal{A} = \big\langle \mathbb{S}^{\langle\delta,k\rangle}Q \;\mid\; (\delta,k)\in S\big\rangle \qquad \text{and} \qquad \mathcal{A}' = \big\langle \mathbb{S}^{\langle\delta,k\rangle}Q \;\mid\; (\delta,k)\in S'\big\rangle.$$

Proposition 4.21 holds with the same proof, *i.e.* $\mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*}(\wedge^2 V\setminus 0),\, W)$ is the HP dual of $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$.

Thus we get relations between $\mathsf{Br}(\mathcal{Y}^{ss}\times_{\mathbb{C}^*}(L^\perp\setminus 0),\, W)$ and $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}|_{\mathbb{P}L})$, which can also be seen concretely in terms of the window categories. Recall that the window for the 'difficult phase' is obtained by restricting the Δ weights to the interval $[-qv, 2l'-qv)$ where $l' = \dim L^\perp$. For the right ranges of l' one window is contained in the other, we have

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{\mathrm{proj}} \;\subset\; \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{[-qv,2l'-qv)} \qquad \text{if } l'\geq qv$$

and:

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{[-qv,2l'-qv)} \;\subset\; \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}L^\perp,\, W)_{\mathrm{proj}} \qquad \text{if } l' \leq (q-\tfrac{1}{2})v$$

Therefore, if $l'\geq qv$ and L is generic then we get an embedding

$$D^b(\mathrm{Pf}_q\cap\mathbb{P}L, B|_{\mathbb{P}L}) \;\hookrightarrow \mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(L^\perp\setminus 0),\, W)$$

and we know that the latter is a categorical resolution of Pf_s. In the other direction, if $l' \geq (q-\frac{1}{2})v$ and L is generic then we get an embedding:

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*}(L^\perp\setminus 0),\, W) \;\hookrightarrow\; D^b(\mathrm{Pf}_q\cap\mathbb{P}L, B|_{\mathbb{P}L}).$$

A similar decomposition to (5.2) exists for $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$, of course. The Lefschetz pieces of $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ and $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ all have full exceptional collections, and it is easy to check that the sizes of these Lefschetz pieces agree with what they would be if $\mathsf{Br}(\mathcal{X}^{\mathrm{ss}})$ and $\mathsf{Br}(\mathcal{Y}^{\mathrm{ss}})$ were HP dual to each other. This lends some support to the possibility that Theorem 3.2 and our other results in fact hold when $\dim V$ is even as well.

5.2. **Squeezing Van den Bergh.** Our proof of Theorem 5.1 will depend on computations of certain local cohomology groups along the unstable locus $\mathcal{Y} \setminus \mathcal{Y}^{ss}$ in $\mathcal{Y}$. In the paper [VdB91], a partial computation of these groups was carried out, in order to prove Cohen–Macaulayness of certain modules of covariants. As pointed out in that paper, the results obtained are not the strongest possible, and we will here refine some statements from the paper to prove what we need.

Van den Bergh's computations are rather complicated and involve several spectral sequences. We will import his techniques and results without explaining them; we apologise that this makes this section of the paper rather opaque unless the reader is familiar with [VdB91].

5.2.1. *General notation.* We first recall some general notation from [VdB91].

Let G be a reductive group, and let M be a linear representation. Choose a maximal torus T in G, and let $X(T)$ and $Y(T)$ denote the lattices of characters and cocharacters respectively. We write $\Phi_G \subset X(T)$ for the set of roots, and Φ_G^+ for the set of positive roots. We let B be the Borel subgroup defined such that $\Phi_B = -\Phi_G^+$.

Choose a diagonalisation of the action of T on M, and let $\mathrm{Chars}(M)$ be the associated sequence of weights $\alpha_1, \ldots, \alpha_{\dim M} \in X(T)$. Given a cocharacter $\lambda \in Y(T)$, we let $\mathrm{Chars}_\lambda^{\leq 0}(M) \subseteq \{1, \ldots, \dim M\}$ denote the set of i such that $(\alpha_i, \lambda) \leq 0$. We denote the complement of this subset by $\mathrm{Chars}_\lambda^{>0}(M)$.

Let $\mathcal{Q}$ denote the set of parabolic subgroups containing B, and let:

$$\mathcal{R} = \{(R, P) \in \mathcal{Q} \times \mathcal{Q} \mid R \subseteq P\}$$

Given $(R, P) \in \mathcal{R}$, the notation π_R^P will be used to denote the maps $G/R \to G/P$ and $G \times^R M \to G \times^P M$.

For any cocharacter $\lambda \in Y(T)$, let $M_\lambda = \{x \in M \mid \lim_{t \to 0} \lambda(t) \cdot x = 0\}$. Given a set of cocharacters $U \subset Y(T)$, let $M_U = \bigcup_{\lambda \in U} M_\lambda$. For a parabolic subgroup $P \subseteq G$, let $A_P = \{\lambda \in Y(T) \mid (\lambda, \rho) \geq 0 \text{ for all } \rho \in \Phi_P\}$. We use the abbreviation $M_P = M_{A_P}$.

5.2.2. *Notation in our case.* For the remainder of Section 5.2, we let $M = \mathrm{Hom}(V, Q)$ and $G = \mathrm{Sp}(Q)$. We choose the standard maximal torus T and the standard positive roots. Concretely, a character $\chi \in X(T)$ is specified by an q-tuple $(\chi_1, \ldots, \chi_q)$. Letting E_i denote the character with i-th coordinate 1 and all other coordinates 0, the simple positive roots of G are $E_1 - E_2, \ldots, E_{q-1} - E_q$, and $2E_q$. We denote this set by Φ_{sim}^+.

We use the convention that $\mathbf{i}_j$ denotes the sequence where the integer i is repeated j times, so that e.g. $(\mathbf{2}_2, \mathbf{1}_3) = (2, 2, 1, 1, 1)$. For each $i \in [0, q]$ we define a cocharacter $\mu_i = (-\mathbf{1}_i, \mathbf{0}_{q-i}) \in Y(T)$. These will be the only cocharacters that are relevant for our calculations, because of the following:

Lemma 5.3. *For $P \in \mathcal{Q}$, we have $M_P = M_{\mu_i}$, where i is the largest integer such that $\mu_i \in A_P$.*

Proof. Suppose $\lambda \in A_P$. Then in particular $\lambda \in A_B$, which means that $(\lambda, \rho) \leq 0$ for all $\rho \in \Phi_{\mathrm{sim}}^+$. Writing $\lambda = (a_1, \ldots, a_n)$, this means $a_1 \leq \cdots \leq a_n \leq 0$. Letting j be the largest integer such that $a_j < 0$, one checks that $M_\lambda = M_{\mu_j}$. Finally, since $M_{\mu_0} \subseteq M_{\mu_1} \subseteq \cdots \subseteq M_{\mu_q}$, and the claim then follows. □

We will also need to consider the group $\overline{G} = \mathrm{GSp}(Q)$, and we write $\overline{T}$ for its maximal torus. A character $\chi \in X(\overline{T})$ is specified by a pair $(\chi_T, \chi_\Delta) = ((\chi_1, \ldots, \chi_q), \chi_\Delta)$, where these are subject to the condition that $\sum_i \chi_i + \chi_\Delta \equiv 0 \pmod 2$.[2] The roots of $\overline{G}$ are $\{(\rho, 0),\ \rho \in \Phi_G\}$.

[2] In previous sections we have denoted this data by (δ, k).

5.2.3. *Local cohomology computations.* Let $M^{us} \subset M$ denote the locus where the image of V in Q is isotropic; so in the notation of the rest of the paper, we have:

$$\mathcal{Y}^{ss} = \big[(M \setminus M^{us}) \,/\, \mathrm{GSp}(Q)\big]$$

M^{us} is the instability locus for one of the two possible stability conditions for the action of $\overline{G} = \mathrm{GSp}(Q)$, it is also the null-cone for the action of $G = \mathrm{Sp}(Q)$. This means it must be the orbit of the linear subspace M_B, which equals M_{μ_q} by Lemma 5.3. One can also see this in an elementary way: the cocharacter μ_q has the effect of scaling a specific Lagrangian subspace in Q, and scaling a complementary Lagrangian in the opposite direction; therefore M_{μ_q} is the locus where V lands inside this specific Lagrangian, and so $M^{us} = GM_{\mu_q}$.

For a character (χ_T, χ_Δ), with χ_T a dominant weight of $\mathrm{Sp}(Q)$, we are interested in computing the $\overline{G}$-invariants in the local cohomology group:

$$H^\bullet_{M^{us}}\Big(M,\, \mathbb{S}^{\langle \chi_T, \chi_\Delta\rangle} Q\Big)$$

Equivalently, we wish to understand the local cohomology $H^\bullet_{M^{us}}(M, \mathcal{O}_M)$ as a representation of $\overline{G}$.

Remark 5.4. The computations in [VdB91] only concern a single group G, and they compute local cohomology along the null-cone for G, as a G-representation. However, the difference between G-representations and $\overline{G}$-representations is just an additional grading, and all the computations remain valid if we keep track of this too. Note that we can't just replace G by $\overline{G}$ in Van den Bergh's method, because the null-cone for the $\overline{G}$-action is the whole of M.

Lemma 5.5. *Let $\lambda \in Y(T)$. A $\overline{T}$-character χ appears in $H^*_{M_\lambda}(M, \mathcal{O}_M)$ with multiplicity equal to the number of sequences $(a_i), (b_i)$ such that*

$$\chi = \sum_{i \in \mathrm{Chars}_\lambda^{\leq 0}(M)} (a_i + 1)\alpha_i - \sum_{j \in \mathrm{Chars}_\lambda^{>0}(M)} b_j \alpha_j \qquad a_i, b_i \in \mathbb{Z}_{\geq 0}.$$

Proof. This follows from [VdB93, Prop. 3.3.1]. Note that λ is a cocharacter for $T \subset G$, and χ is a character for $\overline{T} \subset \overline{G}$; see Remark 5.4. ☐

Lemma 5.6.

(1) *The $\overline{T}$-characters appearing in $H^*_{M_{\mu_0}}(M, \mathcal{O}_M) = H^*_0(M, \mathcal{O}_M)$ are of the form*
$$((d_1, \ldots, d_q), 2qv + k) \qquad d_i, k \in \mathbb{Z}$$
with $k \geq \sum |d_i|$.

(2) *For $i = 1, \ldots, s$, the $\overline{T}$-characters appearing in $H^*_{M_{\mu_i}}(M, \mathcal{O}_M)$ are of the form*
$$((v + c_1, \ldots, v + c_i, d_{i+1}, \ldots, d_q), 2qv - iv + k) \qquad c_i, d_i, k \in \mathbb{Z}$$
with $c_j \geq 0$, and with $\sum c_j + |k| \geq \sum |d_j|$ and $k + \sum c_i \geq 0$.
The character $((\mathbf{v}_i, \mathbf{0}_{q-i}), 2qv - iv)$ appears exactly once.

Proof. Straightforward from Lemma 5.5. ☐

We now have enough information to prove the vanishing of certain local cohomology groups.

Proposition 5.7. *Let $\chi = (\chi_T, \chi_\Delta) \in X(\overline{T})^+$. Assume that one of the following holds:*

(1) *$\chi_T \in Y_{q,2s-1}$ and $\chi_\Delta < 2qv$.*
(2) *$\chi_T \in Y_{q,2s}$ and $\chi_\Delta < 2qv - v$.*

Then $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ *does not appear as a summand of* $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$*, or equivalently:*

$$H^*_{M^{\mathrm{us}}}\left(M,\, \mathbb{S}^{\langle\chi_T,-\chi_\Delta\rangle}Q\right)^{\overline{G}} = 0$$

Proof. Corollary 6.8 of [VdB91] says that (bearing in mind Remark 5.4) if $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ appears in $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$ we must have

$$\chi = \chi' + \sum_{\rho\in S}\rho \tag{5.8}$$

where χ' is a $\overline{T}$-character appearing in some $H^*_{M_\lambda}(M,\mathcal{O}_M)$, and S is some subset of $\Phi_{\overline{G}}$.

By the proof, we see that in fact we must have χ' appearing in some $H^*_{M_P}(M,\mathcal{O}_M)$ for some $P\in\mathcal{Q}$, hence χ' appears in some $H^*_{M_{\mu_i}}(M,\mathcal{O}_M)$, by Lemma 5.3.

We define the following two norms on the space $X(T)_{\mathbb{Q}}$:

$$|\psi|_1 = \max_i |\psi_i| \qquad \text{and} \qquad |\psi|_2 = \max_{i\neq j}\frac{|\psi_i|+|\psi_j|}{2}$$

We'll use the same notation for the pull-back of these functions to $X(\overline{T})_{\mathbb{Q}}$; there they become only semi-norms, but still satisfy the triangle inequality.

Suppose that condition (1) holds. Then in (5.8) we must have $|\sum_{\rho\in S}\rho|_1 \leq 2q$, and by assumption $|\chi|_1\leq 2s-1$, hence:

$$|\chi'|_1 \leq 2s+2q-1 = v-1$$

Using Lemma 5.6 and the fact that $\chi_\Delta = \chi'_\Delta < 2qv$, we see that χ cannot appear in $H^*_{M_{\mu_i}}(M,\mathcal{O}_M)$ for any i, and hence $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ does not appear in $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$.

Suppose next that condition (2) holds. In (5.8) we must have $|\sum_{\rho\in S}\rho|_2\leq 2q-1$ and by assumption $|\chi|_2\leq 2s$, hence $|\chi|_2\leq v-1$. By Lemma 5.6 and the fact that $\chi_\Delta=\chi'_\Delta<2qv-v$, we see that if χ appears in $H^*_{M_{\mu_i}}(M,\mathcal{O}_M)$, we must have $i=1$. But in the terminology of that lemma, we must have $k<0$, and hence $c_1>0$, which implies $|\chi'|_1>v$, and hence $|\chi|_1>2s$, which contradicts our assumed condition on χ. Hence $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ does not appear in $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$. □

We next want to prove that certain local cohomology groups do *not* vanish. This is harder, and we have to look deeper into Van den Bergh's method.

We present the essential part of the argument first; we then follow this with several computational lemmas that it requires.

Proposition 5.9. *Let* $\delta_T = (2s,\mathbf{0}_{q-1})$ *and* $\delta_\Delta = 2qv-v$*. The irrep* $\mathbb{S}^{\langle\delta_T,\delta_\Delta\rangle}Q$ *appears in* $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$*, and so:*

$$H^*_{M^{\mathrm{us}}}\left(M,\, \mathbb{S}^{\langle\delta_T,-\delta_\Delta\rangle}Q\right)^{\overline{G}} \neq 0$$

Proof. We will show that $\mathbb{S}^{\langle\delta_T,\delta_\Delta\rangle}Q$ appears in $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$ with odd multiplicity. The strategy is to compute $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$ via several spectral sequences from [VdB91], and apply the observation that working modulo 2, the multiplicity of an irrep in any page of a spectral sequence equals its multiplicity in the E^∞-page. The bigradings on the spectral sequences involved are irrelevant and will not be specified.

By [VdB91, Thm. 5.2.1], there is a spectral sequence converging to $H^*_{M^{\mathrm{us}}}(M,\mathcal{O}_M)$ whose E^1-page is

$$\bigoplus_{(R,P)\in\mathcal{R}} H^{DR}_*(G\times^R M_P/M).$$

Here H^{DR}_* denotes relative algebraic de Rham homology, see [VdB91].

Each summand $H_*^{DR}(G \times^R M_P/M)$ is computed by a spectral sequence whose E^1-page is

$$H^*_{G\times^B M_P}\big(G \times^B M,\, (\pi_R^B)^*(\wedge^\bullet \Omega_{G\times^R M/M})\big), \tag{5.10}$$

by the proof of [VdB91, Lemmas 6.2 & 6.3]. There is a spectral sequence converging to (5.10), whose E^2-page is

$$H^*\big(G/B,\, (\pi_R^B)^*(\wedge^\bullet \Omega_{G/R}) \otimes \widetilde{H^*_{M_P}(M, \mathcal{O}_M)}\big), \tag{5.11}$$

by [VdB91, Lemma 6.4]. Here the notation $\widetilde{H^*_{M_P}(M, \mathcal{O}_M)}$ denotes the vector bundle on G/B induced by treating $H^*_{M_P}(M, \mathcal{O}_M)$ as a B-representation.

Let $n(R,P) \in \mathbb{Z}/2\mathbb{Z}$ be the multiplicity of $\mathbb{S}^{\langle \delta_T, \delta_\Delta\rangle}Q$ in $H_*^{DR}(G \times^R M_P/M)$ (taken modulo 2), this is the same as its multiplicity in (5.11). The (mod 2)-multiplicity of $\mathbb{S}^{\langle \delta_T, \delta_\Delta\rangle}Q$ in $H^*_{M^{\mathrm{us}}}(M, \mathcal{O}_M)$ is $\sum_{(R,P)\in\mathcal{R}} n(R,P)$, and by Lemma 5.16 below we find that this equals 1. □

Let $P_1 \in \mathcal{Q}$ be the following parabolic in G:

$$P_1 = \{g \in G \mid \lim_{t\to 0} \mu_1(t) g \mu_1(t)^{-1} \text{ exists}\}$$

This is the parabolic subgroup whose roots are $\Phi_B \cup \{\alpha \in \Phi_G \mid (\alpha, \mu_1) \geq 0\}$.

Lemma 5.12. *Choose a parabolic $P \in \mathcal{Q}$. Then the character $((v, \mathbf{0}_{q-1}), 2qv - v)$ appears in $H^*_{M_P}(M, \mathcal{O}_M)$ with multiplicity 1 if $P = P_1$, and does not appear if $P \neq P_1$.*

Proof. Let $S = \Phi^+_{\mathrm{sim}} \cap \Phi_P$. If $(1, -1, \mathbf{0}_{q-2}) \in S$, then $\mu_1 \notin A_P$. If there is an $i \in [2, q-1]$ such that $(\mathbf{0}_{i-1}, 1, -1, \mathbf{0}_{q-i-1}) \notin S$, then $\mu_i \in A_P$. If $(\mathbf{0}_{q-1}, 2) \notin S$, then $\mu_q \in A_P$.

Thus we have "$\mu_1 \in A_P$ and $\mu_i \notin A_P$ for all $i > 1$" if and only if $S = \Phi^+_{\mathrm{sim}} \setminus (1, -1, \mathbf{0}_{q-2})$, which is if and only if $P = P_1$. Combining Lemmas 5.6 and 5.3 we get the claim we want. □

For $\chi \in \overline{T}$, let $BWB(\chi) = w(\chi + \overline{\rho}) - \overline{\rho}$ if there exists a (necessarily unique) Weyl group element $w \in W_G$ such that $w(\chi + \overline{\rho}) - \overline{\rho}$ is a dominant weight; if no such w exists, let $BWB(\chi)$ be undefined. (The notation BWB here is short for the Borel–Weil–Bott theorem, which is how this operation on characters enters the computation.)

Lemma 5.13. *Let C be the set of pairs (R, S), with $R \in \mathcal{Q}$, and $S \subset -\Phi_{G/P}$ such that $BWB(\sum_{\rho\in S} \rho) = 0$. Then $|C|$ is odd.*

Proof. If (R, S) is such a pair, then there exists a $w \in W_G$ such that $w(\sum_{\rho \in S} \rho + \overline{\rho}) - \overline{\rho} = 0$. Since $S \subseteq -\Phi^+$, this tranforms into

$$w\Big(\sum_{\rho\in S} \rho + \sum_{\rho \in \Phi^+ \setminus -S} \rho\Big) = \sum_{\rho\in\Phi^+} \rho.$$

Since the left hand side is a sum of distinct roots, and the right hand side admits only one expression as a sum of distinct roots, we have $S \cup (\Phi^+ \setminus -S) = w^{-1}(\Phi^+)$. From this it follows that

$$S = w^{-1}(\Phi^+) \cap \Phi^-. \tag{5.14}$$

We next note that

$$S \subseteq -\Phi_{G/P} \;\Leftrightarrow\; -\Phi_P \cap S = \varnothing \;\Leftrightarrow\; -\Phi_P \cap -\Phi^+_{\mathrm{sim}} \cap S = \varnothing, \tag{5.15}$$

where $\Phi^+_{\mathrm{sim}} \subseteq \Phi^+$ are the simple roots. To see the last equivalence, note that $-\Phi_P \cap S \subseteq -\Phi_P \cap -\Phi^+$, and any element in $-\Phi_P \cap -\Phi^+$ can be written as a sum of elements in $-\Phi_P \cap -\Phi^+_{\mathrm{sim}}$ [Hum75, p. 184]. If

$$\varnothing = -\Phi_P \cap -\Phi^+_{\mathrm{sim}} \cap S = -\Phi_P \cap -\Phi^+_{\mathrm{sim}} \cap w^{-1}(\Phi^+) \cap \Phi^- = -\Phi_P \cap w^{-1}(\Phi^+),$$

then

$$-\Phi_P \cap -\Phi^+_{\mathrm{sim}} \subseteq \Phi \setminus w^{-1}(\Phi^+) = w^{-1}(-\Phi^+).$$

Since $w^{-1}(-\Phi^+)$ is closed under addition, it follows that $-\Phi_P \subseteq w^{-1}(-\Phi^+)$ and then since $w^{-1}(-\Phi^+) \cap S = \varnothing$, we get $-\Phi_P \cap S = \varnothing$.

The map $P \mapsto -\Phi_P \cap -\Phi^+_{\mathrm{sim}}$ gives a bijection between $\mathcal{Q}$ and the power set of $\{-\Phi^+_{\mathrm{sim}}\}$. Thus by (5.15), we find that for a fixed S, the number of P such that $S \subseteq -\Phi_{G/P}$ equals $2^{|-\Phi^+_{\mathrm{sim}} \setminus S|}$. Hence the parity of $|C|$ equals the number of S such that $-\Phi^+_{\mathrm{sim}} \subseteq S$. But $-\Phi^+_{\mathrm{sim}} \subseteq S$ implies that in (5.14) we must take $w \in W_G$ to be the element which acts via multiplication by -1 on $X(T)$. In particular there is only one such S, hence we are done. □

Let $\delta = (\delta_T, \delta_\Delta) = ((2s, \mathbf{0}_{q-1}), 2qv - v)$ as above.

Lemma 5.16. *Let $P \in \mathcal{Q}$. The representation $\mathbb{S}^{\langle \delta_T, \delta_\Delta \rangle} Q$ appears in*

$$\bigoplus_{(R,P)\in\mathcal{R}} H^*\big(G/B,\, (\pi^B_R)^*(\wedge^\bullet \Omega_{G/R}) \otimes_{\mathcal{O}_{G/B}} \widetilde{H^*_{M_P}(M, \mathcal{O}_M)}\big) \tag{5.17}$$

with odd multiplicity if $P = P_1$, and does not appear if $P \neq P_1$.

Proof. In the following, we identify $\Phi_{\overline{G}}$ with Φ_G in the obvious way, so that a $\rho \in \Phi_G$ is thought of as an element of $X(\overline{T})$.

Let us first consider one summand of (5.17), i.e. one corresponding to a fixed $(R, P) \in \mathcal{R}$. By the proof of [VdB91, Lemma 6.5], the multiplicity of the representation $\mathbb{S}^{\langle \delta_T, \delta_\Delta \rangle} Q$ in this summand equals the cardinality of the set

$$\Big\{(\chi, S) \ \mid\ \chi \text{ appears in } H^*_{M_P}(M, \mathcal{O}_M),\ S \subseteq -\Phi_{G/R},\ BWB(\chi + \sum_{\rho\in S} \rho) = \delta\Big\} \tag{5.18}$$

where (χ, S) is counted with the multiplicity of χ in $H^*_{M_P}(M, \mathcal{O}_M)$. Since the diagonal cocharacter $\Delta \in Y(\overline{T})$ is fixed by $W_{\overline{G}}$, and $\rho_\Delta = 0$ for all $\rho \in \Phi(G/R)$, we see that any χ appearing in (5.18) must satisfy $\chi_\Delta = 2qv - v$.

Let $\psi = ((v, \mathbf{0}_{q-1}), 2qv - v) \in X(\overline{T})$. We now claim that if (χ, S) appears in (5.18), then $\chi = \psi$ and $P = P_1$. So suppose that (χ, S) appears, and let $\chi' = \chi + \sum_{\rho\in S} \rho$. If $BWB(\chi') = \delta$, then we must have $\chi'_1 \leq 2s$ and $\chi'_2 \leq 2s + 1$. In general, we have $(\sum_{\rho\in S} \rho)_1 \geq -2q$ and $(\sum_{\rho\in S} \rho)_2 \geq -2q + 2$. It follows that $\chi_1 \leq v$ and $\chi_2 \leq v - 1$.

We know that χ appears in $H^*_{M_P}(M, \mathcal{O}_M) = H^*_{M_{\mu_i}}(M, \mathcal{O}_M)$ for some $i \in [0, q]$. The inequality $\chi_2 \leq v - 1$ implies that $i \leq 1$, while $\chi_\Delta = 2qv - v$ implies $i \neq 0$, hence $i = 1$, and so we must have $P = P_1$, by Lemma 5.12.

Finally, since $\chi_1 \leq v$, Lemma 5.6 shows that $\chi = \psi$, and that χ_1 appears precisely once in $H^*_{M_P}(M, \mathcal{O}_M)$.

Hence the cardinality of (5.18) equals the cardinality of

$$\Big\{S \subseteq -\Phi_{G/R} \ \mid\ BWB(\sum_{\rho\in S} \rho + \chi) = \delta\Big\} \tag{5.19}$$

if $P = P_1$, and equals 0 otherwise.

Let $A = \{-E_1 \pm E_i\}_{i\in[2,q]} \cup \{-2E_1\} \subseteq -\Phi^+$. Since $(\sum_{\rho\in S} \rho + \psi)_1 \leq 2s$, we must have $(\sum_{\rho\in S} \rho)_1 \leq -2q$, which implies that $(\sum_{\rho\in S} \rho)_1 = -2q$ and hence $A \subseteq S$. Let now $S' = S \setminus A$. The set (5.19) is then in bijection with the set

$$\Big\{S' \subseteq -\Phi_{\overline{G}/R} \setminus A \ \mid\ BWB\Big(\sum_{\rho\in S'} \rho + ((2q, \mathbf{0}_{q-1}), 0)\Big) = \delta\Big\}.$$

Let now Ψ be the root system for $\mathrm{GSp}(2q-2)$, and let E_i' be the standard generators of the character lattice of $\mathrm{Sp}(2q-2)$. There is an inclusion of the character lattice of $\mathrm{GSp}(2q-2)$ into that of $\mathrm{GSp}(2q)$ by

$$(E_i', k) \mapsto (E_{i+1}, k).$$

This gives a bijection $F : -\Phi^+ \setminus A \to -\Psi^+$. The map F then also induces a bijection between the $R \in \mathcal{Q}$ such that $R \subseteq P_1$ and the set of all parabolic R in $\mathrm{GSp}(2q-2)$ containing the standard Borel subgroup.

Furthermore, it is easy to see that $BWB\big(\sum_{\rho\in S'}\rho + ((2q, \mathbf{0}_{q-1}), 0)\big) = \delta$ holds if and only if $BWB(\sum_{\rho\in F(S')}\rho) = 0$, where the latter equation is in the character lattice of $\mathrm{GSp}(2q-2)$. Thus Lemma 5.13 (applied to $G = \mathrm{Sp}(2q-2)$) concludes the proof. □

This completes the proof of Proposition 5.9. We need to prove one more non-vanishing result, as follows:

Proposition 5.20. *Let $\chi = (\chi_T, \chi_\Delta) = ((\mathbf{0}_q), 2qv)$. Then $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ appears as a summand of $H^*_{M^{\mathrm{us}}}(M, \mathcal{O}_M)$, so:*

$$H^*_{M^{\mathrm{us}}}\Big(M, \mathbb{S}^{\langle\chi_T,-\chi_\Delta\rangle}Q\Big)^{\overline{G}} \neq 0$$

Proof. An argument similar to that in the proof of Lemma 5.16 shows that if $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ appears in

$$H^*\big(G/B,\, (\pi_R^B)^*(\wedge^\bullet\Omega_{G/R}) \otimes_{\mathcal{O}_{G/B}} \widetilde{H^*_{M_P}(M, \mathcal{O}_M)}\big)$$

then we must have $M_P = M_{\mu_0} = 0$, and arguing similarly to Lemma 5.12, this means that $P = G$. The proof of Lemma 5.16 shows that the multiplicity of $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ inside

$$\bigoplus_{R\in\mathcal{Q}} H^*\big(G/B,\, (\pi_R^B)^*(\wedge^\bullet\Omega_{G/R}) \otimes_{\mathcal{O}_{G/B}} \widetilde{H^*_0(M, \mathcal{O}_M)}\big)$$

reduces to the cardinality of the pairs (R, S) such that $R \in \mathcal{Q}$, $S \in -\Phi_{G/R}$, and $BWB(\sum_{\rho\in S}\rho) = 0$. But by Lemma 5.13 there is an odd number of such. Arguing as in the proof of Proposition 5.9 completes the proof. □

5.3. **Proof of the window equivalence in the even case.** We will now use the results from the previous section to prove Theorem 5.1, essentially in the same way as we proved Theorem 4.7 when v was odd. We will prove the essential surjectivity and fully faithfulness of the functor separately.

5.3.1. *Fully faithfulness.*

Lemma 5.21. *Given $\chi, \psi \in X(T)^+$, every summand $\mathbb{S}^{\langle\alpha\rangle}Q$ of $\mathbb{S}^{\langle\chi\rangle}Q \otimes \mathbb{S}^{\langle\psi\rangle}Q$ satisfies $\alpha_1 \leq \chi_1 + \psi_1$.*

Proof. Given any character $\phi \in X(T)$, let $|\phi|_1 = \max|\phi_i|$. Let now ϕ be a T-weight occurring in the $\mathrm{Sp}(Q)$-irrep $\mathbb{S}^{\langle\delta\rangle}Q$. We claim

$$|\phi|_1 \leq |\delta|_1. \tag{5.22}$$

Suppose not, then there would be a k such that $|\phi_k| > |\delta|_1 = \delta_1$. We can find an element w of the Weyl group of $\mathrm{Sp}(Q)$ such that $(w\phi)_1 = |\phi_k|$. But since $w\phi$ is a T-weight of $\mathbb{S}^{\langle\delta\rangle}Q$, it is dominated by δ, hence $(w\phi)_1 \leq \delta_1$, and we have a contradiction, so (5.22) holds.

Now $\alpha = \phi + \phi'$ for ϕ some T-weight of $\mathbb{S}^{\langle\chi\rangle}Q$ and ϕ' some T-weight of $\mathbb{S}^{\langle\psi\rangle}Q$, so we find:

$$\alpha_1 = |\alpha|_1 \;\leq\; |\phi|_1 + |\phi'|_1 \;\leq\; |\chi|_1 + |\psi|_1 = \chi_1 + \phi_1$$

□

Proposition 5.23. *Let $\chi, \psi \in Y_{\mathrm{proj}}$. Then restriction induces an isomorphism:*

$$\mathrm{Hom}_{\mathcal{Y}}\big(\mathbb{S}^{\langle \chi_T, \chi_\Delta\rangle}Q,\, \mathbb{S}^{\langle \psi_T, \psi_\Delta\rangle}Q\big) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{Y}^{ss}}\big(\mathbb{S}^{\langle \chi_T, \chi_\Delta\rangle}Q,\, \mathbb{S}^{\langle \psi_T, \psi_\Delta\rangle}Q\big)$$

Proof. We need to prove the vanishing of the local cohomology of the bundle $\mathbb{S}^{\langle \chi_T, \chi_\Delta\rangle}Q^\vee \otimes \mathbb{S}^{\langle \psi_T, \psi_\Delta\rangle}Q$ along the locus $\mathcal{Y} \setminus \mathcal{Y}^{ss}$. Examining the definition of Y_{proj}, it's sufficient to check the following two cases:

a) $\chi_T \in Y_{q,s-1},\ \psi_T \in Y_{q,s},\ \chi_\Delta - \psi_\Delta < 2qv$.
b) $\chi_T \in Y_{q,s},\ \psi_T \in Y_{q,s},\ \chi_\Delta - \psi_\Delta < 2qv - v$.

Let $\mathbb{S}^{\langle \alpha_T, \alpha_\Delta\rangle}Q$ be a summand of $\mathbb{S}^{\langle \chi_T, \chi_\Delta\rangle}Q^\vee \otimes \mathbb{S}^{\langle \psi_T, \psi_\Delta\rangle}Q$. In case (a) we must have $\alpha_T \in Y_{q,2s-1}$ (by Lemma 5.21) and $\alpha_\Delta = \psi_\Delta - \chi_\Delta < 2qv$; this is the first hypothesis of Lemma 5.7 so the result follows. Similarly in case (b) we have $\alpha_T \in Y_{q,2s}$ and the second hypothesis of that lemma is satisfied. □

Obviously this result still holds if we work on $\mathcal{Y} \times_{\mathbb{C}^*} L^\perp$ instead, and then the first paragraph of the proof of Proposition 4.13 shows immediately that:

Corollary 5.24. *The restriction functor*

$$\mathsf{Br}(\mathcal{Y} \times_{\mathbb{C}^*} L^\perp,\, W)_{\mathrm{proj}} \to \mathsf{Br}(\mathcal{Y}^{\mathrm{ss}} \times_{\mathbb{C}^*} L^\perp,\, W)$$

is fully faithful.

5.3.2. *Essential surjectivity.* As in the v odd case, we define an object

$$P_{\delta,k} \in \mathsf{Br}(\mathcal{Y})$$

for any dominant weight (δ, k) of $\mathrm{GSp}(Q)$, by projecting the twist $\mathcal{O}_0 \otimes \mathbb{S}^{\langle \delta, k\rangle}Q$ of the skyscraper sheaf at the origin into the admissible subcategory $\mathsf{Br}(\mathcal{Y}) \subset D^b(\mathcal{Y})$. As before these objects restrict to zero in $\mathsf{Br}(\mathcal{Y}^{ss})$ (Lemma 4.16), and each one has a finite resolution by vector bundles which hence becomes an exact sequence on $\mathcal{Y}^{ss}$. The bundles appearing in this sequence correspond to the irreps in $h_\bullet(P_{\delta,k}|_0)$, and as before we refer to these as the *weights* of $P_{\delta,k}$.

Proposition 5.25. *Let (δ, k) be a dominant weight of $\mathrm{GSp}(Q)$. Then the set of weights of $P_{\delta,k}$ contains one copy of (δ, k). Furthermore:*

- *If $\delta \in Y_{q,s-1}$, then the remaining weights are contained in*
$$Y_{s,q} \times [k - 2qv,\, k).$$
- *If $\delta \in Y_{q,s} \setminus Y_{q,s-1}$ then the remaining weights are contained in*
$$Y_{s,q} \times [k - 2(q-1)v,\, k).$$

Proof. The first claim, and the fact that all remaining weights lie in $Y_{s,q} \times (-\infty, k)$, can be argued exactly as in the v odd case (see the proof of Lemma 4.15); we just have to prove the lower bounds.

So, let m be the minimum of the set $\{m' \mid (\psi, m') \text{ is a weight of } P_{\delta,k}\}$. If we consider the locally-free resolution of $P_{\delta,k}$, we get a short exact sequence:

$$(P_{\delta,k})_{>min} \to P_{\delta,k} \to (P_{\delta,k})_{min} \tag{5.26}$$

Here $(P_{\delta,k})_{min}$ is the final term in the resolution, it is a direct sum of vector bundles of the form $\mathbb{S}^{\langle \psi, m\rangle}Q$. The object $(P_{\delta,k})_{>min}$ consists of all the remaining terms, it is a complex of vector bundles whose summands are all of the form $\mathbb{S}^{\langle \psi, m'\rangle}Q$ with $m < m' \leq k$.

Now choose one of the minimal weights (ψ, m), *i.e.* choose one of the summands of $(P_{\delta,k})_{min}$. We make the following claim, which immediately implies the proposition:

- If $\psi \in Y_{q,s-1}$ then $\psi = \delta$ and $m = k - 2qv$.
- If $\psi \in Y_{q,s} \setminus Y_{q,s-1}$ then $\psi = \delta$ and $m = k - 2(q-1)v$.

(So in particular $(P_{\delta,k})_{min}$ is actually a direct sum of copies of $\mathbb{S}^{\langle\delta,m\rangle}Q$.)

For any vector bundle E on $\mathcal{Y}$, let's write $\mathrm{Hom}_{\mathcal{Y}^{us}}(E,-)$ for the functor that sends an object $\mathcal{F}$ to the local cohomology of $\mathrm{Hom}(E,\mathcal{F})$ along the substack $\mathcal{Y}^{us} = [M^{us}/\mathrm{GSp}(Q)]$. To prove our claim we're going to apply this functor to the short-exact-sequence (5.26); we get the two statements by using different choices of vector bundle E.

Let's start with the case $\psi \in Y_{q,s-1}$. Set E to be the bundle $\mathbb{S}^{\langle\psi,m+2qv\rangle}Q$, and apply the functor. To evaluate the first term, we need to know the local cohomology of each bundle

$$\mathrm{Hom}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, \mathbb{S}^{\langle\psi',m'\rangle}Q\big)$$

with $m' > m$. But this always vanishes, by part (1) of Lemma 5.7 (c.f. the proof of Proposition 5.23). So

$$\mathrm{Hom}_{\mathcal{Y}^{us}}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, (P_{\delta,k})_{>min}\big) = 0$$

and we conclude that the map

$$\mathrm{Hom}_{\mathcal{Y}^{us}}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, P_{\delta,k}\big) \longrightarrow \mathrm{Hom}_{\mathcal{Y}^{us}}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, (P_{\delta,k})_{min}\big) \tag{5.27}$$

is an isomorphism.

Let's examine the third term, the target of this isomorphism (5.27). The bundle $(P_{\delta,k})_{min}$ contains $\mathbb{S}^{\langle\psi,m\rangle}Q$ as a summand, and the bundle

$$\mathrm{Hom}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, \mathbb{S}^{\langle\psi,m\rangle}Q\big)$$

contains $\mathbb{S}^{\langle(\mathbf{0}_q),-2qv\rangle}Q$ as a summand. So by Proposition 5.20, we have:

$$\mathrm{Hom}_{\mathcal{Y}^{us}}\big(\mathbb{S}^{\langle\psi,m+2qv\rangle}Q,\, (P_{\delta,k})_{min}\big) \neq 0$$

This means that the second term (the source in (5.27)) is also non-zero. However, the object $P_{\delta,k}$ is supported along $\mathcal{Y}^{us}$, so for this term the local cohomology is just the global sections of the Hom bundle. But by Lemma 4.14 this can only be non-zero if we have $\delta = \psi$ and $k = m + 2qv$. This proves the first case of the claim.

Now we consider the second case, when $\psi \in Y_{q,s} \setminus Y_{q,s-1}$. We set E to be the bundle $\mathbb{S}^{\langle\psi,m+2(q-1)v\rangle}Q$, and proceed in a very similar way. The vanishing of the first term now holds by part (2) of Lemma 5.7. To understand the third term we need the observation that since the first entry of ψ is s, the bundle

$$\mathrm{Hom}\big(\mathbb{S}^{\langle\psi,m+2(q-1)v\rangle}Q,\, \mathbb{S}^{\langle\psi,m\rangle}Q\big)$$

contains $\mathbb{S}^{\langle(2s,\mathbf{0}_{q-1}),-2(q-1)v\rangle}Q$ as a summand; see Lemma 5.30 below. Then Proposition 5.9 shows that the third term does not vanish. The remainder of the argument is identical to the first case. □

A trivial modification of the proof of Lemma 4.17 yields:

Corollary 5.28. *For any $(\chi_T, \chi_\Delta) \in Y_{q,s} \times \mathbb{Z}$, the vector bundle $\mathbb{S}^{\langle\chi_T,\chi_\Delta\rangle}Q$ on $\mathcal{Y}^{ss}$ has a finite resolution by vector bundles associated to the set Y_{proj}.*

Now we can use the proof of Proposition 4.18 verbatim to deduce that:

Corollary 5.29. *The restriction functor*

$$\mathsf{Br}(\mathcal{Y}\times_{\mathbb{C}^*} L^\perp,\, W)_{\mathrm{proj}} \to \mathsf{Br}(\mathcal{Y}^{\mathrm{ss}}\times_{\mathbb{C}^*} L^\perp,\, W)$$

is essentially surjective.

This completes the proof of Theorem 5.1, except for the following representation-theoretic calculation needed in the proof of Proposition 5.25.

Lemma 5.30. *Let χ be a dominant weight of $\mathrm{Sp}(Q)$. Then $(\mathbb{S}^{\langle\chi\rangle}Q)^{\otimes 2}$ contains $\mathbb{S}^{\langle(2\chi_1,\mathbf{0}_{q-1})\rangle}Q$ as a summand.*

Proof. The claim is equivalent to the claim that $\mathbb{S}^{\langle(2\chi_1,\mathbf{0}_{q-1})\rangle}Q \otimes \mathbb{S}^{\langle\chi\rangle}Q$ contains $\mathbb{S}^{\langle\chi\rangle}Q$ as a summand. We apply the Littlewood–Newell formula [BKW83, 5.7]. The formula tells us that if β is a partition and n_β denotes the multiplicity of $\mathbb{S}^{\langle\beta\rangle}Q$ in $\mathbb{S}^{\langle\alpha\rangle}Q \otimes \mathbb{S}^{\langle\psi\rangle}Q$, we have the equality

$$\sum n_\beta x^\beta = \sum_{\lambda,\gamma,\theta} c(\beta) N^\psi_{\lambda,\gamma} N^\alpha_{\lambda,\theta} N^\beta_{\gamma,\theta} x^{m(\beta)}.$$

Here the $N^*_{*,*}$ are Littlewood–Richardson coefficients, and we sum over all partitions. We have $c(\beta) \in \{-1, 0, 1\}$ and $m(\beta)$ is a partition; both are determined by 'modification rules' [BKW83, Sec. 3]. We only need to know that if the length of β is at most q then $c(\beta) = 1$ and $m(\beta) = \beta$, and if the length of β is $q + 1$ then $c(\beta)$ is 0 (and $m(\beta)$ is undefined).

Let us write (i) for the partition $(i, \mathbf{0}_\infty)$. For the product $\mathbb{S}^{\langle\chi\rangle}Q \otimes \mathbb{S}^{\langle(2\chi_1)\rangle}Q$, the Littlewood–Newell formula gives:

$$\sum_\beta n_\beta x^\beta = \sum_{\theta,\beta} \sum_{k=0}^{2\chi_1} c(\beta) N^\chi_{(k),\theta} N^\beta_{(2\chi_1-k),\theta} x^{m(\beta)},$$

since $N^{(2\chi_1)}_{\lambda,\gamma} = 1$ if $\lambda = (k)$ and $\gamma = (2\chi_1 - k)$ and $N^{(2\chi_1)}_{\lambda,\gamma} = 0$ otherwise. To find the coefficient n_χ here we must find all β which appear with non-zero coefficient in this sum and which satisfy $m(\beta) = \chi$. However, if $N^\chi_{(k),\theta} \neq 0$, then $l(\theta) \leq l(\chi)$, which if $N^\beta_{(2\chi_1-k),\theta} \neq 0$ means:

$$l(\beta) \leq l(\theta) + 1 \leq l(\chi) + 1 \leq q + 1$$

So by the modification rules, the only contribution to n_χ comes from $\beta = \chi$. For both $N^\chi_{(k),\theta}$ and $N^\chi_{(2\chi_1-k),\theta}$ to be non-zero we must have $k = \chi_1$, so:

$$n_\chi = \sum_\theta (N^\chi_{(\chi_1),\theta})^2$$

Since $N^\chi_{(\chi_1),\theta} = 1$ if $\theta = (\chi_2, \chi_3, \ldots, \chi_q)$ and is zero otherwise, this shows that $n_\chi = 1$. □

References

[ADS15] Nicolas Addington, Will Donovan, and Ed Segal. The Pfaffian-Grassmannian equivalence revisited. *Alg. Geom.*, 2(3):332–364, 2015.

[BDF+13] Matthew Ballard, Dragos Deliu, David Favero, M. Umut Isik, and Ludmil Katzarkov. Homological projective duality via variation of geometric invariant theory quotients. *JEMS (to appear)*, 2013. arXiv:1306.3957.

[BFK12] Matthew Ballard, David Favero, and Ludmil Katzarkov. Variation of geometric invariant theory quotients and derived categories. *J. Reine Angew. Math. (to appear)*, 2012. arxiv:1203.6643.

[BKW83] G. R. E. Black, R. C. King, and B. G. Wybourne. Kronecker products for compact semisimple Lie groups. *J. Phys. A*, 16(8):1555–1589, 1983.

[Bri93] Michel Brion. Sur les modules de covariants. *Ann. Sci. École Norm. Sup. (4)*, 26(1):1–21, 1993.

[BVdB03] A. Bondal and M. Van den Bergh. Generators and representability of functors in commutative and noncommutative geometry. *Mosc. Math. J.*, 3(1):1–36, 258, 2003.

[Eis95] David Eisenbud. *Commutative algebra*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

[EP15] Alexander I. Efimov and Leonid Positselski. Coherent analogues of matrix factorizations and relative singularity categories. *Algebra Number Theory*, 9(5):1159–1292, 2015.

[FH91] William Fulton and Joe Harris. *Representation theory*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.

[Har66] Robin Hartshorne. *Residues and duality*. Lecture notes of a seminar on the work of A. Grothendieck, given at Harvard 1963/64. With an appendix by P. Deligne. Lecture Notes in Mathematics, No. 20. Springer-Verlag, Berlin-New York, 1966.

[Har80] Robin Hartshorne. Stable reflexive sheaves. *Math. Ann.*, 254(2):121–176, 1980.

[Hir16] Yuki Hirano. Derived Knörrer periodicity and Orlov's theorem for gauged Landau-Ginzburg models. *arXiv:1602.04769*, 2016.

[HK13] Kentaro Hori and Johanna Knapp. Linear sigma models with strongly coupled phases – one parameter models. *J. High Energy Phys.*, (10), 2013.

[HL15] Daniel Halpern-Leistner. The derived category of a GIT quotient. *J. Amer. Math. Soc.*, 28(3):871–912, 2015.

[Hor13] Kentaro Hori. Duality in two-dimensional $(2,2)$ supersymmetric non-Abelian gauge theories. *J. High Energy Phys.*, (10), 2013.

[HT07] Kentaro Hori and David Tong. Aspects of non-abelian gauge dynamics in two-dimensional $N = (2,2)$ theories. *J. High Energy Phys.*, (5), 2007.

[HT13] Shinobu Hosono and Hiromichi Takagi. Duality between $S^2\mathbb{P}^4$ and the double quintic symmetroid. *arXiv:1302.5881*, 2013.

[Hum75] James E. Humphreys. *Linear algebraic groups*. Springer-Verlag, New York-Heidelberg, 1975. Graduate Texts in Mathematics, No. 21.

[Kap84] M. Kapranov. Derived category of coherent sheaves on Grassmann manifolds. *Izv. Akad. Nauk SSSR Ser. Mat.*, 48(1):192–202, 1984.

[KR08] Mikhail Khovanov and Lev Rozansky. Matrix factorizations and link homology. *Fund. Math.*, 199(1):1–91, 2008.

[Kuz06] Alexander Kuznetsov. Homological projective duality for Grassmannians of lines. *arXiv:math/0610957*, 2006.

[Kuz07] Alexander Kuznetsov. Homological projective duality. *Publications mathématiques de l'IHÉS*, 105(1):157–220, June 2007.

[Kuz08] Alexander Kuznetsov. Derived categories of quadric fibrations and intersections of quadrics. *Advances in Mathematics*, 218(5):1340–1369, 2008.

[Lip02] Joseph Lipman. Lectures on local cohomology and duality. In *Local cohomology and its applications (Guanajuato, 1999)*, volume 226 of *Lecture Notes in Pure and Appl. Math.*, pages 39–89. Dekker, New York, 2002.

[NvO82] C. Năstăsescu and F. van Oystaeyen. *Graded ring theory*, volume 28 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam-New York, 1982.

[Orl12] Dmitri Orlov. Matrix factorizations for nonaffine LG-models. *Math. Ann.*, 353(1):95–108, 2012.

[Pos11] Leonid Positselski. Two kinds of derived categories, Koszul duality, and comodule-contramodule correspondence. *Mem. Amer. Math. Soc.*, 212(996):vi+133, 2011.

[Ren15] Jørgen Rennemo. The homological projective dual of $\mathrm{Sym}^2\,\mathbb{P}(V)$. *arXiv:1509.04107*, 2015.

[Ren17] Jørgen Rennemo. The fundamental theorem of homological projective duality via variation of GIT stability. *arxiv:1705.01437*, 2017.

[RSVdB17] Jørgen Rennemo, Ed Segal, and Michel Van den Bergh. A non-commutative Bertini theorem. *arxiv:1705.01366*, 2017.

[Seg11] Ed Segal. Equivalences between GIT quotients of Landau-Ginzburg B-models. *Comm. Math. Phys.*, 304(2):411–432, 2011.

[Shi12] Ian Shipman. A geometric approach to Orlov's theorem. *Compos. Math.*, 148(5):1365–1389, 2012.

[ST14] Ed Segal and Richard Thomas. Quintic threefolds and fano elevenfolds. *J. Reine Angew. Math. (to appear)*, 2014. 1410.6829.

[Sta16] The Stacks Project Authors. Stacks Project. http://stacks.math.columbia.edu, 2016.

[Sun86] Sheila Sundarem. *On the combinatorics of representations of* $\mathrm{Sp}(2n,\mathbb{C})$. PhD thesis, MIT, 1986.

[ŠVdB15] Špela Špenko and Michel Van den Bergh. Non-commutative resolutions of quotient singularities. *Invent. Math. (to appear)*, 2015.

[Tho15] R. P. Thomas. Notes on HPD. *arXiv:1512.08985 [math]*, 2015. arXiv: 1512.08985.

[VdB91] Michel Van den Bergh. Cohen-Macaulayness of modules of covariants. *Invent. Math.*, 106(2):389–409, 1991.

[VdB93] Michel Van den Bergh. Cohen-Macaulayness of semi-invariants for tori. *Trans. Amer. Math. Soc.*, 336(2):557–580, 1993.

[VdB04] Michel Van den Bergh. Non-commutative crepant resolutions. In *The legacy of Niels Henrik Abel*, pages 749–770. Springer, Berlin, 2004.

[Wey03] Jerzy Weyman. *Cohomology of vector bundles and syzygies*, volume 149 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2003.

[YZ04] Amnon Yekutieli and James J. Zhang. Rigid Dualizing Complexes on Schemes. *arXiv:math/0405570*, 2004. arXiv: math/0405570.

[YZ06] Amnon Yekutieli and James J. Zhang. Dualizing complexes and perverse sheaves on noncommutative ringed schemes. *Selecta Math. (N.S.)*, 12(1):137–177, 2006.