Discriminants and Semi-orthogonal Decompositions

Alex Kite, Ed Segal

University College London, London, UK. E-mail: e.segal@ucl.ac.uk

Received: 3 March 2021 / Accepted: 8 December 2021
 Published online: 1 February 2022 – © The Author(s) 2022

Abstract: The derived categories of toric varieties admit semi-orthogonal decompositions coming from wall-crossing in GIT. We prove that these decompositions satisfy a Jordan–Hölder property: the subcategories that appear, and their multiplicities, are independent of the choices made. For Calabi–Yau toric varieties wall-crossing instead gives derived equivalences and autoequivalences, and mirror symmetry relates these to monodromy around the GKZ discriminant locus. We formulate a conjecture equating intersection multiplicities in the discriminant with the multiplicities appearing in certain semi-orthogonal decompositions. We then prove this conjecture in some cases.

Contents

1.	Introduction	908
1.1	Spherical functors from wall-crossing	908
1.2	FI parameter spaces	909
1.3	The rank 1 case	910
1.4	Components of the discriminant	910
1.5	Factorizations and multiplicities	911
2.	Toric Background	912
2.1	Notation and assumptions	912
2.2	The Calabi–Yau case	913
2.3	Higgs and Coloumb GIT problems	913
3.	Semi-orthogonal Decompositions for Toric Varieties	914
3.1	Crossing a single wall	914
3.2	The algorithm	915
3.3	Relevant subspaces	917
3.4	Proof of the main theorem	918
4.	FI Parameter Spaces and Discriminants	921
4.1	Spherical functors	921

4.2 Discriminants	922
4.2.1 Horn uniformization	923
4.2.2 Components of the discriminant	924
4.3 Faces and subspaces	927
4.4 The conjecture	928

1. Introduction

Let X be a toric variety, constructed as a GIT quotient of a vector space V by a torus T . There is a well-established theory [Kaw, Seg, BFK, HL] that tells us how to produce semi-orthogonal decompositions of the derived category $D^b(X)$. We do it by considering other birational models of X , *i.e.* crossing walls in the GIT problem $T \curvearrowright V$. If we cross to a quotient X' , and $K_{X'}$ is ‘more negative’ than K_X , then $D^b(X)$ decomposes as

$$D^b(X) = \langle D^b(X'), D^b(Z), \dots, D^b(Z) \rangle \quad (1.1)$$

where Z is another toric variety of smaller dimension. We do this repeatedly until we arrive at a ‘minimal’ chamber. Since the extra pieces are always equivalent to the derived category of a toric variety they themselves can be decomposed by the same procedure, and we get a recursive algorithm which terminates after a finite number of steps.

If X is projective then the result of this algorithm is a full exceptional collection for X , *i.e.* every piece of the final decomposition is equivalent to $D^b(\mathbb{C})$. But for quasi-projective varieties there will usually be many different categories occurring, each one with some multiplicity. Moreover the decomposition is not unique; at each step of the algorithm one may have a choice about which wall to cross through and these choices result in different decompositions. The main technical result of this paper is the following Jordan–Hölder type theorem:

Theorem A (Theorem 3.12). *Let X be a toric variety. If we decompose $D^b(X)$ using the wall-crossing algorithm then the subcategories occurring in the final decomposition, and their multiplicities, are independent of all choices.*

This result is not particularly hard to prove and neither is it an abstract result; we prove it by analysing the algorithm. But it is notable that the Jordan–Hölder property does not hold for semi-orthogonal decompositions in general [BBS, Kuz].

Our real motivation for proving the theorem above was to be able to understand a conjecture appearing in a physics paper by Aspinwall–Plesser–Wang [APW]. Part of what they state is already understood in the mathematical literature but there remains a significant unsolved problem which we are able to formulate precisely using our theorem (Conjecture 4.17). This generalizes a conjecture made by Halpern-Leistner–Shipman [HLSh].

We will use the remainder of this introduction to explain the motivation and context for this conjecture

1.1. Spherical functors from wall-crossing. Our conjecture concerns the special case when the torus action $T \curvearrowright V$ is through the subgroup $SL(V)$. In this case all the GIT quotients X will be Calabi–Yau, meaning $K_X \cong \mathcal{O}_X$, and not projective. In this situation the wall-crossing theory does not provide any decompositions of $D^b(X)$, instead it proves that all the GIT quotients are derived equivalent since the decomposition (1.1)

just becomes $D^b(X) = D^b(X')$. However the category $D^b(Z)$ still has an important role.

The derived equivalence between X and X' is not unique, the theory gives us multiple equivalences for every wall-crossing, and by composing them we get autoequivalences of $D^b(X)$. From work of Halpern-Leistner–Shipman [HLS_{Sh}] it is known that each of these autoequivalences can be described as a twist T_F around a spherical functor

$$F : D^b(Z) \rightarrow D^b(X)$$

where Z is the same toric variety that appears in (1.1).

By combining these, and the Picard groups of each GIT quotient, we can get many autoequivalences of $D^b(X)$. So the interesting problem becomes to understand this large group of autoequivalences.

1.2. FI parameter spaces. Now we explain some heuristics from physics and mirror symmetry. In string theory the data of T acting on V determines an abelian gauged linear sigma model, a widely studied class of $N = (2, 2)$ superconformal field theories. In this theory there are certain important parameters called complexified Fayet–Iliopoulos parameters, they take values in a complex manifold which we call the FI Parameter Space (FIPS). They are related to stability conditions in the GIT problem and in certain limiting regions of the FIPS the theory reduces to a sigma model whose target is one of the quotients X . In physical terminology X is a *phase* of the model. In this region we can identify the FI parameters with the complexified Kähler moduli of X so the FIPS is closely related to the *extended* or *stringy* Kähler moduli space of X .¹

Under mirror symmetry the FI parameters become complex parameters, so the FIPS is the base of the mirror family. Since toric mirror symmetry has a mathematically precise formulation this gives us a rigorous definition of the FIPS: it's the complement of the GKZ discriminant locus ∇ inside the dual torus $T^\vee$ (Sect. 4.2). It is helpful to think of $T^\vee$ as an open subset of the secondary toric variety $\mathfrak{F}$ and to take the closure $\overline{\nabla} \subset \mathfrak{F}$, because then the phases correspond to the toric fixed points in $\mathfrak{F}$. From this point-of-view the FIPS is obtained by deleting $\overline{\nabla}$ and the toric boundary from $\mathfrak{F}$.

The mirror family is a locally-trivial family of symplectic manifolds over the FIPS with fibre $\check{X}$. The monodromy of this family gives an action of $\pi_1(\text{FIPS})$ on $\check{X}$ as symplectomorphisms, and hence as autoequivalences of the Fukaya category $\text{Fuk}(\check{X})$. On the mirror side this predicts an action:

$$\pi_1(\text{FIPS}) \curvearrowright D^b(X)$$

This is the ‘B-brane monodromy’. Examples and physical calculations suggest that this is essentially the group of autoequivalences that arise via wall-crossing as described in Sect. 1.1. This prediction appears in many places in the maths and physics literature (e.g. [HHP, HW, HLS_{Sam}, HLS_{Sh}]) and has been verified for some examples [DS, Kite]. It seems to be a difficult problem to verify it in general, mainly because it is hard to understand $\pi_1(\text{FIPS})$.

¹ The FIPS is not quite the same as the SKMS, the latter should be intrinsic to X whereas the former depends on its presentation as $V//T$. Also note that the SKMS is expected to be a complex submanifold of the space of Bridgeland stability conditions; on the mirror side this is the difference between small and big quantum cohomology. The FIPS is easier to compute than either the SKMS or the space of stability conditions.

1.3. The rank 1 case. The case where $T = \mathbb{C}^*$ is quite well-known and easy to understand directly. In this case there are two possible phases which we denote by $X_{\pm}$. If we split V by weights as $V_+ \oplus V_0 \oplus V_-$ then it's easy to see that $X_{\pm}$ is a vector bundle over $\mathbb{P}V_{\pm} \times V_0$, where $\mathbb{P}V_{\pm}$ is a weighted projective space.

In this rank 1 case the discriminant locus is always a single point δ so the FIPS is $\mathbb{C}^* \setminus \delta$ (see Example 4.4). Or we can say that the secondary toric variety $\mathfrak{F}$ is a $\mathbb{P}^1$ and that the FIPS is obtained from it by deleting the two toric fixed points and one more non-fixed point. The phase X_+ corresponds to the region near one of the toric fixed points, and the loop around that fixed point simply acts as $\otimes \mathcal{O}(1)$ on $D^b(X_+)$.

More interesting is the loop around the non-fixed point δ —often called the *conifold point*—which corresponds to wall-crossing to X_- and back again. If there are no zero weights then the resulting autoequivalence is the twist T_S around a spherical object

$$S = \mathcal{O}_{PV_+}$$

given by the sky-scraper sheaf along the zero section in X_+ . If there are zero weights we upgrade this to a twist around the spherical functor

$$F : D^b(V_0) \rightarrow D^b(X_+)$$

given by pulling-up to $\mathbb{P}V_+ \times V_0$ and then pushing-forward along the inclusion into X_+ . In the notation of Sect. 1.1 the variety Z is V_0 .

Remark 1.2. If there is only one positive weight then X_+ is an affine orbifold and $\text{Pic}(X_+)$ is a finite cyclic group $\mathbb{Z}/k$. In this case it's sensible to allow that toric fixed point as part of the FIPS. The reason is that $\mathfrak{F}$ is (if we're careful) an orbifold $\mathbb{P}^1$ and this fixed point has isotropy group $\mathbb{Z}/k$, so we get an action of the orbifold fundamental group.

This subtlety is interesting in the rank 1 case since it occurs in the well-known ‘Calabi–Yau/Landau–Ginsburg correspondence’. In higher rank it happens very rarely and is of no significance for this paper. For us the FIPS will contain none of the toric boundary and hence we can ignore any orbifold structure on $\mathfrak{F}$.

1.4. Components of the discriminant. Suppose we have a higher rank torus $T \cong (\mathbb{C}^*)^r$. The discriminant locus ∇ is now some hypersurface in $(\mathbb{C}^*)^r$ and it is usually the union of several irreducible components:

$$\nabla = \nabla_0 \cup \dots \cup \nabla_k$$

Aspinwall–Plesser–Wang [APW] observed that there is a correspondence between these components ∇_i and certain toric varieties Z_i , built from subsets of the original toric data. They conjecture that for each phase X there should be a spherical functor

$$F_i : D^b(Z_i) \rightarrow D^b(X) \tag{1.3}$$

and that T_{F_i} corresponds to the monodromy around the component ∇_i . There is some deliberate ambiguity here; there is no canonical loop around ∇_i (even up to homotopy), so the functors F_i are at best defined up to composition by autoequivalences.

Fig. 1. (L) A real picture of C_W as the straight line connecting the two points marked by $X_{\pm}$. (R) A complex picture of a 2-sphere near to the rational curve C_W , where the point δ has split into three. A loop from X_+ to X_- and back again will factor into two loops around Δ_0 and one loop around Δ_1

1.5. Factorizations and multiplicities. To understand this conjecture of [APW] more clearly we pick two adjacent chambers of the secondary fan, separated by a wall W . This is the situation we discussed in Sect. 1.1. The two chambers give two phases $X_{\pm}$ which are derived equivalent, and we get an autoequivalence of $D^b(X_+)$ which is the twist around a spherical functor

$$F : D^b(Z) \rightarrow D^b(X_+) \quad (1.4)$$

for some smaller toric variety Z .

In the secondary toric variety $\mathfrak{F}$ our wall W corresponds to a rational curve C_W connecting the toric fixed points corresponding to our two phases. It turns out that the discriminant locus $\bar{\nabla}$ always intersects C_W in a single point δ (Corollary 4.13). This is the same picture that we saw in Sect. 1.3, and the reason for this is that by focusing on a single wall-crossing we are essentially reducing to a rank 1 GIT problem. There is a 1-parameter subgroup $\lambda_W \subset T$ normal to the wall and it is only stability with respect to λ_W that is changing. So, just as in the rank 1 case, a loop in C_W that goes around the point δ should correspond to the autoequivalence T_F .

However, C_W is not part of the FIPS since it lies in the toric boundary of $\mathfrak{F}$. To get an actual element of $\pi_1(FIPS)$ we have to perturb C_W (or an open subset of it) off the toric boundary, and take a loop in the perturbed curve.

When we do this perturbation the point δ may split into several points because $\bar{\nabla}$ typically meets C_W with some multiplicity. This means that our element of $\pi_1(FIPS)$ is naturally a composite of several loops, one around each of our new missing points. In fact each component $\bar{\nabla}_i$ might meet C_W with multiplicity, and we can group the new missing points according to these components (see Fig. 1).

So the loop around δ naturally factors into several loops around the different components of $\bar{\nabla}$, with each component possibly appearing multiple times. This suggests that we should look for a corresponding factorization of the autoequivalence T_F .

This factorization does indeed exist. The toric variety Z is not usually a Calabi–Yau, which means that $D^b(Z)$ (unlike $D^b(X)$) can be decomposed using the wall-crossing algorithm. Moreover, the subcategories that appear in this decomposition are always

equivalent to $D^b(Z_i)$ where Z_i is one of the varieties considered by Aspinwall–Plesser–Wang (Sect. 1.4). So we get a semi-orthogonal decomposition

$$D^b(Z) = \langle D^b(Z_0), D^b(Z_0), \dots, D^b(Z_k), D^b(Z_k) \rangle \quad (1.5)$$

where each $D^b(Z_i)$ occurs some number of times (possibly zero). The order of the factors here depends on the choices made in the algorithm, but by our Theorem 3.12 the multiplicities do not.

Halpern-Leistner–Shipman [HLSh] observed that this decomposition gives us a factorization of the autoequivalence T_F . If we restrict the spherical functor F (1.4) to each piece of $D^b(Z)$ then we again get a spherical functor, and T_F is the composition of all the corresponding twists. This provides the spherical functors F_i required by Aspinwall–Plesser–Wang and matches with our discussion of loops in the FIPS.

However, for this story to make sense there is one essential numerical condition:

Conjecture 1.6 (Conjecture 4.17). *The multiplicity of $D^b(Z_i)$ in the decomposition (1.5) agrees with the intersection multiplicity of $\bar{\nabla}_i$ with C_W .*

We finish by proving our conjecture in some special cases, the strongest of which is:

Theorem 2 (Theorem 4.23). *If the torus T has rank 2 then Conjecture 1.6 holds.*

Remark 1.7. A significant part of this story was already understood by Halpern-Leistner–Shipman. They only consider the case when Z is projective, meaning that the decomposition of $D^b(Z)$ is actually a full exceptional collection, and they conjecture that the number of exceptional objects agrees with the intersection multiplicity of $\bar{\nabla}$ with C_W [HLSh, Remark 4.7]. Our conjecture is a synthesis of theirs with the work of [APW].

2. Toric Background

2.1. Notation and assumptions. We are interested in toric varieties constructed as GIT quotients of a vector space V by a torus T . We specify the data of the torus action as a complex of lattices

$$0 \longrightarrow L \xrightarrow{Q^\vee} \mathbb{Z}^n \xrightarrow{A} N \longrightarrow 0 \quad (2.1)$$

or its dual:

$$0 \longrightarrow M \xrightarrow{A^\vee} \mathbb{Z}^n \xrightarrow{Q} L^\vee \longrightarrow 0 \quad (2.2)$$

here:

- L is the lattice of 1-parameter subgroups of the torus T , so $T = L_{\mathbb{C}^*}$.
- $\mathbb{Z}^n$ is the lattice of Laurent monomials on V , i.e. $V = \text{Spec}[\mathbb{N}^n]$ for the submonoid $\mathbb{N}^n \subset \mathbb{Z}^n$.
- Q is the *weight map*. The images $q_i = Q(e_i)$ of the standard basis vectors are the *weights* of the action.
- N is the cokernel of $Q^\vee$ modulo torsion.
- M is the kernel of Q and the dual of N .
- A is the *ray map*. The images $a_i = A(e_i)$ are the *rays*.

By definition A is surjective and $A^\vee$ is injective. We will always assume that $Q^\vee$ is injective, so Q is surjective modulo torsion—this is the assumption that generic points of our GIT quotient stacks do not have infinite isotropy groups. It follows that (2.1) and (2.2) are exact apart from a possible torsion group $L^\vee / \text{Im } Q \cong \text{Ker } A / \text{Im } Q^\vee$.

A *stability condition* is an element of $L_{\mathbb{R}}^\vee$. A choice of stability condition θ defines a semi-stable locus in V and hence a GIT quotient, which for us means the quotient stack:

$$X_\theta = [V_\theta^{ss} / T]$$

We'll generally only be interested in quotients with respect to generic θ , in which case X_θ is at worst a DM stack. We'll also refer to these generic GIT quotients as the *phases* of the GIT problem. Each phase is a toric orbifold and has a corresponding fan in N . The rays of this fan are always (some subset of) the a_i 's, hence the name. The higher dimensional cones change depending on the phase.

Remark 2.3. If the weight map Q has some finite cokernel then the representation $T \rightarrow GL(V)$ has a finite kernel, so the GIT quotients X_θ have finite isotropy groups at all points. We need to allow this possibility, since even if it doesn't apply to our initial toric variety X it can happen for the smaller-dimensional varieties Z that appear in wall-crossing.

Note that in this situation A does not determine Q . There is a theory of *stacky fans* but which solves this issue but we won't need it because for us Q is the fundamental piece of data.

The space of stability conditions has a wall-and-chamber structure whose chambers correspond to phases. If we consider all (non-empty) GIT quotients we get a fan in $L^\vee$ called the *secondary fan*—the top-dimensional cones correspond to non-empty phases and the lower-dimensional cones correspond to non-generic GIT quotients. The rays of the secondary fan include those generated by the weights q_i , but in general there more rays than this. Corresponding to the secondary fan is a toric variety, the *secondary toric variety* $\mathfrak{F}$.

2.2. The Calabi–Yau case. An important special case is when the torus T acts through $SL(V)$, which implies that each phase is Calabi–Yau.

In terms of the toric data, the Calabi–Yau case is when the sum of the weights q_i is zero. Equivalently, the rays a_i are all contained in (and hence affinely span) an affine hyperplane of height 1. In this case is helpful to consider the polytope

$$\Pi \subset N_{\mathbb{R}}$$

given by the convex hull of the rays. Each phase corresponds to a fan in N , which when intersected with the affine hyperplane determines a decomposition of Π . These decompositions are exactly the *coherent triangulations*, i.e. triangulations induced by a piece-wise linear function.

2.3. Higgs and Coloumb GIT problems. From our original GIT problem $T \curvearrowright V$ we will often extract a smaller GIT problem involving some subset of the toric data, either by picking a subset of the weights, or a subset of the rays. The two main ways this will happen are:

- (1) Suppose $W \subset L_{\mathbb{R}}^{\vee}$ is a wall in the secondary fan, normal to some 1-parameter subgroup $\lambda \in L$. Then we can consider the subset of weights which are orthogonal to λ , *i.e.* which lie in the subspace $\langle W \rangle$.
- (2) In the Calabi–Yau case we can choose a face $\Gamma \subset \Pi$ of the toric polytope, and consider the set of rays lying in this face.

Formally, suppose we pick a subset $S \subset \{1, \dots, n\}$. We can view S as a subset of the standard basis vectors $\{e_1, \dots, e_n\}$ in $\mathbb{Z}^n$ so there is a corresponding set of rays $A(S) \subset N$. We set $N_S \subset N$ to be the sublattice spanned by $A(S)$, write $A_S : \mathbb{Z}^S \rightarrow N_S$ for the restriction of A , and set $L_S = \text{Ker } A_S$. Then we get a GIT problem:

$$L_S \xrightarrow{Q_S^{\vee}} \mathbb{Z}^S \xrightarrow{A_S} N_S$$

We'll refer to this as the *Coloumb GIT problem* associated to the subset S .

Alternatively we pick a subset $T \subset \{1, \dots, n\}$ and consider the corresponding set of weights $Q(T) \subset L^{\vee}$. We define $L_T^{\vee}$ as the primitive sublattice generated by these weights

$$L_T^{\vee} = L^{\vee} \cap \langle Q(T) \rangle_{\mathbb{R}} \subset L_{\mathbb{R}}^{\vee}$$

and we get a GIT problem:

$$M_T \xrightarrow{A_T^{\vee}} \mathbb{Z}^T \xrightarrow{Q_T} L_T^{\vee}$$

We'll call this the *Higgs GIT problem* associated to T . Note that Q_S is by definition surjective but Q_T might not be (*c.f.* Remark 2.3).

Our ‘Higgs’ and ‘Coloumb’ terminology is based on the ‘Higgs GLSM’ and ‘Coloumb GLSM’ from [APW], which are related to the Higgs and Coloumb branches of the vacuum moduli space at singular values of the FI parameters.

Remark 2.4. If our original GIT problem is Calabi–Yau then the Coloumb GIT problem is also Calabi–Yau for any subset S . But the Higgs GIT problems may not be.

3. Semi-orthogonal Decompositions for Toric Varieties

3.1. Crossing a single wall. Fix a toric GIT problem $T \curvearrowright V$. Let C_+ and C_- be two adjacent chambers of the secondary fan separated by a wall W , and labelled such that C_+ lies on the same side of W as the character $\det(V)$. Let $\lambda_W \in L$ be the primitive 1-parameter subgroup normal to this wall, oriented such that

$$\kappa = (\det V)(\lambda_W) \geq 0$$

i.e. C_+ lies on the $\lambda_W > 0$ side. Write $X_{\pm}$ for the phases corresponding to these two chambers.

For this wall we have a Higgs GIT problem as described in Sect. 2.3. Let $\mathcal{T}$ be the indexing set for the weights orthogonal to λ_W , so $Q(\mathcal{T})$ are all the weights lying in the subspace $\langle W \rangle$. The vector space corresponding to $\mathbb{Z}^{\mathcal{T}}$ is the fixed subspace $V^{\lambda_W} \subset V$. Also $Q(\mathcal{T})$ necessarily span $\langle W \rangle$, so $L_{\mathcal{T}}^{\vee}$ is exactly the orthogonal to λ_W , *i.e.* it's the character lattice of T/λ_W . Hence this Higgs GIT problem is just describing the action of T/λ_W on V^{λ_W} .

The secondary fan for this Higgs GIT problem lives in the vector space $\langle W \rangle$ and the cone W lies in some chamber of it. We write Z for the corresponding phase.

Theorem 3.1 [BFK, Theorem 5.2.1]. *We have a semi-orthogonal decomposition*

$$D^b(X_+) = \langle D^b(X_-), D^b(Z), \dots, D^b(Z) \rangle$$

where κ copies of $D^b(Z)$ occur.

Remark 3.2. This theorem is an application of the general theory of ‘windows’ relating GIT and derived categories [BFK, HL, Seg], which applies to a general GIT quotient of a variety by a reductive group. However, in the current state-of-the-art you cannot use this theory to compare two *different* GIT quotients unless you assume that the wall-crossing is of a particularly simple form. Which these ones are.

Remark 3.3. If $\det(V)$ lies on the wall then $\kappa = 0$ and the theorem states that $D^b(X_+)$ and $D^b(X_-)$ are equivalent. This is a toric flop.

Example 3.4. If we consider the standard action of $\mathbb{C}^*$ on $\mathbb{C}^{n+1}$ then $X_- = \emptyset$ and we get

$$D^b(\mathbb{P}^n) = \langle D^b(pt), \dots, D^b(pt) \rangle$$

which recovers Beilinson’s result that $\mathbb{P}^n$ has full exceptional collection of length $n + 1$.

Remark 3.5. If X_+ happens to be a blow-up of X_- then Theorem 3.1 recovers Orlov’s blow-up formula for this toric situation. It’s possible to formulate the theorem more generally in such a way that it directly generalizes Orlov’s result.

3.2. The algorithm. Theorem 3.1 immediately suggests the following recursive algorithm for decomposing the derived category of a phase X :

- (1) Starting at the chamber for X we cross through a sequence of walls, always moving away from $\det(V)$. At each wall we refine our decomposition.
- (2) We stop when we reach a *minimal phase* where no further such wall-crossings are possible.
- (3) Every factor occurring in this decomposition is the derived category of a phase of a smaller GIT problem, so we can apply this algorithm to each factor.

Note that a phase is minimal if $-(\det V)$ lies in the closure of that chamber, or equivalently if the canonical bundle of that phase is nef.

Remark 3.6. If X is projective then you can use this algorithm to recover Kawamata’s result [Kaw] that a projective toric variety has a full exceptional collection [BFK, Thm 5.2.3]. This is because the minimal phase will be empty (as in Example 3.4), and moreover the minimal phase is empty in every Higgs GIT problem that occurs in the algorithm.

In this paper we are more interested in quasi-projective examples.

Example 3.7. Take $V = \mathbb{C}^6$ and quotient by $(\mathbb{C}^*)^2$ using the following matrix of weights:

$$\begin{pmatrix} 1 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & -1 \end{pmatrix}$$

Observe that $\det(V) = (1, 2)^\top$. This GIT problem has four phases and the secondary fan is drawn in Fig. 2. The phases are:

Fig. 2. The phases of Example 3.7

- (1) $X_1 = \mathbb{A}^4$. This is the unique minimal phase.
- (2) $X_2 = \mathcal{O}(-1)_{\mathbb{P}^3}$, the total space of the tautological line bundle on $\mathbb{P}^3$.
- (3) $X_3 = \mathcal{O}(-1)_{\mathbb{P}^1} \times \mathbb{A}^2$.
- (4) $X_4 = \mathcal{O}(-1)_P$, the total space of the relative $\mathcal{O}(-1)$ line bundle over the projective bundle $P = \mathbb{P}(\mathcal{O}^{\oplus 2} \oplus \mathcal{O}(-1)) \rightarrow \mathbb{P}^1$.

Firstly we decompose $D^b(X_2)$ by crossing the wall into chamber (1). The 1-parameter subgroup for this wall is $(1, 1)$ so $\kappa = 3$. The Higgs GIT is $\mathbb{C}^* \curvearrowright \mathbb{C}$ with weight 1, and Z is the non-empty phase $Z = pt$. Hence Theorem 3.1 in this case says

$$D^b(X_2) = \langle D^b(X_1), D^b(pt), D^b(pt), D^b(pt) \rangle$$

which is an instance of Orlov's blow-up formula (see Remark 3.5).

To make the rest of this example more readable we'll write this SOD and all following ones in the compressed form:

$$D^b(X_2) = \langle X_1, pt, pt, pt \rangle$$

For this phase no further refinements are possible, and the algorithm is finished.

Next we apply the algorithm to phase 4. Let us choose to cross to phase (2) and then to phase (1). The wall-crossing between (2) and (4) is again a blow-up, it blows up the codimension 2 subvariety $\mathcal{O}(-1)_{\mathbb{P}^1}$. So crossing both walls gives:

$$D^b(X_4) = \langle X_2, \mathcal{O}(-1)_{\mathbb{P}^1} \rangle = \langle X_1, pt, pt, pt, \mathcal{O}(-1)_{\mathbb{P}^1} \rangle$$

We are not yet finished, because we can still apply the algorithm to the factor $D^b(\mathcal{O}(-1)_{\mathbb{P}^1})$. But this variety is just the blow-up of $\mathbb{A}^2$ at the origin, so the next refinement is:

$$D^b(X_4) = \langle X_1, pt, pt, pt, \mathbb{A}^2, pt \rangle \quad (3.8)$$

No further refinements are possible.

What happens if we make a different choice? We could instead have crossed to phase (3) before crossing to phase (1). The crossing $(1) \rightsquigarrow (3)$ blows up a plane, and the crossing $(3) \rightsquigarrow (4)$ blows up a $\mathbb{P}^1$, so crossing these walls gives the decomposition:

$$D^b(X_4) = \langle X_3, \mathbb{P}^1, \mathbb{P}^1 \rangle = \langle X_1, \mathbb{A}^2, \mathbb{P}^1, \mathbb{P}^1 \rangle$$

The factor $D^b(\mathbb{P}^1)$ can be split into two exceptional objects (as in Example 3.4) so the final step is:

$$D^b(X_4) = \langle X_1, \mathbb{A}^2, pt, pt, pt, pt \rangle \quad (3.9)$$

Note that in this example the quotienting torus is $(\mathbb{C}^*)^2$, and for each phase we needed to apply the recursive algorithm (at most) two times. For rank r it would need r applications.

In the preceding example we noticed that when decomposing $D^b(X_4)$ we had two choices, since there were two possible paths from chamber (4) to chamber (1).² In the second step of the algorithm there was no such choice, since the Higgs GIT problems were all rank 1 and had only two chambers. In a higher rank example there will be many more choices because we need to choose a path at every step except the last one.

However, examining the decompositions (3.8) and (3.9) that resulted from our two paths we can see evidence of our Jördan-Holder property—the decompositions are different, but the multiplicities of the ‘irreducible factors’ agree. To state this precisely we need to think about what these ‘irreducible factors’ really are.

3.3. Relevant subspaces. Recall that our initial GIT problem is given by a weight matrix $Q : \mathbb{Z}^n \rightarrow L^\vee$ specifying a torus action $T \curvearrowright V$. At any step in the algorithm the Higgs GIT problem arises as the fixed subspace $V^{T'}$ for some sub-torus $T' \subset T$, with a corresponding sublattice $L' \subset L$. The weights $q_1, \dots, q_h$ of this Higgs GIT problem are those weights which are orthogonal to L' , and they always span the subspace $(L')^\perp_{\mathbb{R}} \subset L_{\mathbb{R}}^\vee$.

The ‘irreducible factors’ of our decompositions are the derived categories of the minimal phases of each such Higgs GIT problem. However, some of these minimal phases will be empty. Since the stability condition that produces the minimal phase is

$$-\det(V^{T'}) = -\sum_{i=1}^h q_i \in (L')^\perp$$

we get a non-empty minimal phase iff the vector $-\sum q_i$ lies in the cone spanned by $q_1, \dots, q_h$.

Definition/Lemma 3.10. Let $H \subset L_{\mathbb{R}}^\vee$ be a subspace, let $q_1, \dots, q_h$ be the weights lying in H , and let $\sigma_H \subset H$ be the cone spanned by these weights. We call H *relevant* if one of the following two equivalent conditions hold:

- (i) σ_H is the whole of H .
- (ii) H is spanned by $q_1, \dots, q_h$ and also $-\sum q_i \in \sigma_H$.

Proof. Obviously (i) implies (ii). Conversely (ii) implies that $-q_i \in \sigma_H$ for all i , so if the q_i ’s span H then any vector in H can be written as a positive linear combination of them.

Clearly there can only be finitely-many relevant subspaces. We allow $H = 0$ (which is always relevant) and $H = L_{\mathbb{R}}^\vee$ (which might not be). A 1-dimensional relevant subspace is a line which has weights on both its rays.

The relevant subspaces index the ‘irreducible factors’ in our semi-orthogonal decompositions. Each one defines a Higgs GIT problem with a non-empty minimal phase Z_H , and the corresponding factor is $D^b(Z_H)$.

² By ‘path’ we really mean a sequence of adjacent chambers.

Example 3.11. In Example 3.7 there are three relevant subspaces: the whole of $\mathbb{R}^2$, the vertical axis, and the origin. They contribute the factors $D^b(\mathbb{A}^4)$, $D^b(\mathbb{A}^2)$ and $D^b(pt)$ respectively.

Theorem 3.12. *Let X be a phase of a toric GIT problem and let H be a relevant subspace. The multiplicity of $D^b(Z_H)$ in the semi-orthogonal decomposition of $D^b(X)$ is independent of all choices of paths.*

Presumably the different decompositions resulting from different choices of paths are always related by mutations, but we haven't checked this.

Remark 3.13. The actual value of the multiplicity of $D^b(Z_H)$ in $D^b(X)$ is determined algorithmically from the toric data. It would be interesting—and probably helpful for Conjecture 4.17—to find something like a closed-form expression for it.

We don't know how to do this, except in the case when H has codimension 1 when it follows easily from the discussion in Sect. 3.1. Let λ_H be a primitive 1-parameter subgroup normal to H , oriented so that it pairs positively with the chamber for X , and set $\kappa = (\det V)(\lambda_H)$. Then the multiplicity of $D^b(Z_H)$ in $D^b(X)$ is

$$\max\{\kappa, 0\}$$

since the algorithm only tells us to cross H if $\kappa > 0$.

3.4. Proof of the main theorem. We'll prove Theorem 3.12 using the recursive structure of the algorithm to reduce to the rank 2 case, *i.e.* when the GIT problem consists of $(\mathbb{C}^*)^2 \curvearrowright V = \mathbb{C}^n$. In the rank 1 case the theorem is vacuous since there are no choices.

Lemma 3.14. *Theorem 3.12 holds in the rank 2 case.*

Proof. If $\det(V)$ is the trivial character then all phases are derived equivalent and the theorem is vacuously true, so we can assume $\det(V) \neq 0$. For simplicity we assume that neither $\det(V)$ nor $-\det(V)$ lie on a wall, so there is a unique minimal phase and a unique 'maximal' phase $X_{\max}$, whose chamber contains $\det(V)$. In fact there could be up to two minimal or maximal phases, but crossing the walls between them is a derived equivalence and we can ignore it. If we start at any non-maximal phase then there are no choices to be made in the algorithm, but if we start at $X_{\max}$ then we have exactly two choices of paths to reach $X_{\min}$. So the only thing to check is that these two choices produce the same multiplicities.

There are three classes of relevant subspace:

- (1) $H = \mathbb{C}^2$. This is relevant iff $X_{\min}$ is non-empty, in which case $D^b(X_{\min})$ occurs in $D^b(X_{\max})$ with multiplicity one for either choice of path.
- (2) H a line, both rays of which are walls. The Higgs GIT for H has two non-empty phases, let Z_H be a minimal one and Z'_H be the other one.

By assumption $\det(V)$ doesn't lie on H , so if λ_H is a primitive normal 1-parameter subgroup to H then $\kappa = |\lambda_H(\det V)|$ is strictly positive. The minimal and maximal chambers lie on opposite sides of H so either choice of path crosses it; one choice contributes κ copies of $D^b(Z_H)$ and the other contributes κ copies of $D^b(Z'_H)$. But the decomposition of $D^b(Z'_H)$ includes exactly one copy of $D^b(Z_H)$ so either way the multiplicity of $D^b(Z_H)$ in $D^b(X_{\max})$ is κ .

- (3) $H = \{0\}$. This contributes the factor $D^b(V^T)$, the subspace of V fixed by the whole torus.

Consider a line $l \subset L_{\mathbb{R}}^{\vee}$ containing at least one weight, let q_l be the sum of the weights on this line, and let $\mu_l = |q_l|$ be the lattice length of q_l . There are two possibilities:

- (a) There are weights on both rays of l . Then l is a relevant subspace as in case (2), both rays are walls and the Higgs GIT has a non-empty minimal phase Z_l . The derived category of the other phase Z'_l decomposes into one copy of $D^b(Z_l)$ and μ_l copies of $D^b(V^T)$.
- (b) There are only weights on one ray so only that ray is a wall. The Higgs GIT has an empty phase and the other phase decomposes into μ_l copies of $D^b(V^T)$.

In either case only one of our two paths will pick up any factors of $D^b(V^T)$ from this line l ; it's the path that crosses l on the same side as q_l , and the number of such factors it picks up is

$$\mu_l \kappa_l = \mu_l |\lambda_l(\det(V))|$$

where λ_l is a primitive 1-parameter subgroup normal to l . So we may as well assume that each such line contains only a single weight $q_i = q_l$, and hence only that ray of the line is a wall.

Now fix an orientation on our lattice $L^{\vee}$. Since the lattice is rank 2 this is the same as a unit symplectic form ω . This means that for the wall through q_i we can produce a primitive normal subgroup by setting $\lambda = \omega(\hat{q}_i, -)$ where $\hat{q}_i$ is a primitive vector in the direction of q_i . With this choice one of our paths always crosses walls in the direction of increasing λ and the other path always crosses walls in the direction of decreasing λ . So if the first path crosses the rays through $q_1, \dots, q_s$ and the second path crosses the rays through $q_{s+1}, \dots, q_n$ then the equality we want to show is:

$$\sum_{i=1}^s \mu_i \lambda_i(\det(V)) = - \sum_{i=s+1}^n \mu_i \lambda_i(\det(V))$$

But this is true since

$$\sum_{i=1}^n \mu_i \lambda_i = \sum_{i=1}^n \omega(q_i, -) = \omega(\det(V), -)$$

and $\omega(\det(V), \det(V)) = 0$.

Now suppose we have a higher rank problem, and we choose a phase X corresponding to a chamber C_X . To run the algorithm we first pick a path from C_X to the chamber for a minimal phase, C_{min} , always moving away from $\det(V)$. It doesn't matter which minimal phase we pick since moving between them is a derived equivalence. But there might be many possible paths from C_X to C_{min} .

To visualize this clearly pass from the secondary fan to the dual 'secondary polytope' in $L_{\mathbb{R}}$. This has a vertex for each chamber, an edge for each wall, and higher-dimensional faces for fans of higher codimension. The element $\det(V) \in L^{\vee}$ defines a linear function on the polytope and induces a direction on (most of) the edges, since we never allow this function to increase when we traverse an edge.

Choose a path, *i.e.* a directed sequence of edges, between the vertices c_X and c_{min} corresponding to the chambers C_X and C_{min} . Now pick a polygon P (a two-dimensional

face) in the secondary polytope which meets our path; let's say the path meets P at some vertex c_1 , traverses some edges of the polygon, then leaves it again at c_2 . If the remaining edges in P also happen to form a directed path then we can produce a new path from c_X to $c_{\min}$ by choosing to go the other way around P . This is possible iff c_1 maximizes $\det(V)$ among vertices of P and c_2 minimizes $\det(V)$.

Let's call this kind of operation on paths a *simple modification*.

Lemma 3.15. *Any two paths from c_X to $c_{\min}$ are connected by a sequence of simple modifications.*

Proof. The secondary polytope is a cell decomposition of an $(r - 1)$ -sphere, and the subset where $\det(V) \leq \det(V)(c_X)$ is a decomposition of a disc. Any two paths in this disc from c_X to $c_{\min}$ are homotopic, and given such a homotopy we can move it orthogonally to $\det(V)$ until it lies in the 2-skeleton of the polytope. We then have a collection of polygons whose boundary is the union of our two paths, with $\det(V) \leq \det(V)(c_X)$ everywhere. We just need to show that we can perform a simple modification to one of our paths at one of these polygons; then the result follows by induction.

Choose one of our paths. Let P_1 and P_2 be the first two polygons this path meets and let c_1 be the vertex where the path switches between them. By considering the edge between P_1 and P_2 , and remembering that $\det(V)$ is a linear function on these affine polygons, we can see that either c_1 minimizes $\det(V)$ in P_1 , or else it maximizes $\det(V)$ in P_2 . In the first case we can modify the path at P_1 . In the second case we move to the next pair of polygons along the path and repeat the argument. Note that $c_{\min}$ certainly minimizes $\det(V)$ in the final polygon so the algorithm terminates there if not before.

Proof of Theorem 3.12. Pick two paths from the chamber for X to the chamber for a minimal phase, always moving away from $\det(V)$. By Lemma 3.15 it's enough to deal with the case when our two paths are related by a simple modification; this means they agree except at a single codimension-two cone Λ in the secondary fan where they travel opposite ways around. For every Higgs GIT problem that our paths encounter we also need to make choices, but those GIT problems have lower rank so by induction we can assume that those choices do not matter.

Let $U \subset V$ be the semi-stable locus for a character lying on our codimension-two cone. Then we have a GIT problem $T \curvearrowright U$ whose phases are exactly those phases of $T \curvearrowright V$ whose chambers are adjacent to the cone. This new GIT problem is 'non-linear' in that U is not a vector space, and there is an important sense in which it is rank two. If we let $L' \subset L$ be the rank 2 sublattice normal to our codimension-two cone, and $T' \subset T$ be the corresponding subtorus, then only subgroups lying in T' can have fixed points in U . It follows that the GIT fan for $T \curvearrowright U$ is just the GIT fan for $T' \curvearrowright U$, pulled-back via the projection $L_{\mathbb{R}}^{\vee} \rightarrow (L')_{\mathbb{R}}^{\vee}$.

So in the region where our two paths differ we can think of them as paths in the GIT fan for $T' \curvearrowright U$. And since they are different they both must start in a maximal chamber and end in a minimal chamber.

Now consider the linear GIT problem $T' \curvearrowright V$. The GIT fan for $T' \curvearrowright U$ is a coarsening of the one for $T' \curvearrowright V$; every wall of the former is a wall of the latter, but not necessarily vice-versa since a subgroup $\lambda \subset T'$ could have fixed points in V but none in U . However, from the point-of-view of our algorithm there is no harm in regarding every wall for $T' \curvearrowright V$ as corresponding to a wall for $T \curvearrowright U$ —it just happens that some of them will be 'fake walls' where the semi-stable locus does not change. In the semi-orthogonal decomposition crossing a fake wall adds some number of copies of the zero category $D^b(U^{\lambda} // T) = D^b(\emptyset)$. Note that if both rays of a line are a wall for

$T' \curvearrowright V$ then either both give genuine walls for $T \curvearrowright U$ or both are fake. Also if there are any fake walls then $U^{T'}$ is empty, which means that the codimension-two cone itself also contributes the zero category.

If we include these zero categories then we have a bijection between the factors in the decomposition algorithms for $T' \curvearrowright V$ and for $T \curvearrowright U$, and their multiplicities agree since these depend only on the restriction of the character $\det(V)$ to the subtorus T' . Hence the result follows from Lemma 3.14. $\square$

4. FI Parameter Spaces and Discriminants

In this section we consider a *Calabi–Yau* GIT problem $T \curvearrowright V$ where T acts through the subgroup $SL(V)$. This has a different flavour to the previous section, since all phases are Calabi–Yau and every wall-crossing is a derived equivalence, so no semi-orthogonal decompositions occur. Instead (as discussed in the introduction) we focus on autoequivalences of the phases and relate these to the fundamental group of the FI parameter space.

4.1. Spherical functors. Let $T \curvearrowright V$ be a Calabi–Yau toric GIT problem. Let X_+ and X_- be two phases coming from two adjacent chambers C_+ and C_- , separated by a wall W . Let Z be the phase of the associated Higgs GIT problem for a character lying on W .

Since $\det(V) = 0$, Theorem 3.1 tells us that $D^b(X_+)$ and $D^b(X_-)$ are equivalent. However, what the theory actually gives us is a countable set of equivalences

$$\Phi_i : D^b(X_+) \xrightarrow{\sim} D^b(X_-)$$

indexed by the integers. They are related by the Picard groups of X_+ and X_- .

Theorem 4.1 [HLS_h, Prop. 3.4]. *There is a spherical functor*

$$F : D^b(Z) \rightarrow D^b(X_+)$$

such that $\Phi_1^{-1}\Phi_0$ is the twist around F .

Recall that the *twist* around F is the endofunctor of $D^b(X_+)$ defined by the cone on the counit

$$T_F = [FR \rightarrow 1]$$

where R is the right adjoint to F , and that the key property of a *spherical* functor is that T_F is an autoequivalence. See [AL] for more detail on spherical functors. Note that this cone of functors makes sense since we can interpret it as a cone of Fourier–Mukai kernels (or insert the prefix ‘dg’ where needed).

The variety Z is toric—it’s a phase of the Higgs GIT problem—but it will not usually be Calabi–Yau. So using the algorithm of Sect. 3 we can produce a semi-orthogonal decomposition:

$$D^b(Z) = \langle \mathcal{C}_1, \dots, \mathcal{C}_r \rangle \tag{4.2}$$

Halpern-Leistner and Shipman observed that this implies:

- (1) The restriction of F to each piece gives a spherical functor $F_i : \mathcal{C}_i \rightarrow D^b(X_+)$.

(2) The twist T_F factors as:

$$T_F = T_{F_1} \circ \cdots \circ T_{F_r} \quad (4.3)$$

The formal result is [HLSh, Theorem 4.14] and it applies in this situation since the cotwist around F is (up to a shift) the Serre functor on $D^b(Z)$.

The factors in the semi-orthogonal decomposition (4.2) are indexed by the relevant subspaces in the Higgs GIT problem for W , but these are simply the relevant subspaces $H \subset L_{\mathbb{R}}^{\vee}$ which are contained in the hyperplane $\langle W \rangle$.

4.2. Discriminants. We now recall some of the theory of discriminant loci developed by Gelfand–Kapranov–Zelevinsky [GKZ].

Recall that our GIT problem is specified by a sequence of lattices, exact modulo torsion, or its dual:

$$\begin{array}{ccccc} L & \xrightarrow{Q^{\vee}} & \mathbb{Z}^n & \xrightarrow{A} & N \\ M & \xrightarrow{A^{\vee}} & \mathbb{Z}^n & \xrightarrow{Q} & L^{\vee} \end{array}$$

From this point on we need to make two mild additional assumptions:

- (1) We assume that the rays $A(e_i)$ are all distinct. We need this because for [GKZ] A is a subset of N . This excludes 1-parameter subgroups acting with weights $(0, \dots, 0, 1, -1, 0, \dots, 0)$ but these are very uninteresting from a wall-crossing perspective.
- (2) We assume the weights $Q(e_i)$ are all non-zero. This is just for simplicity. A zero weight just contributes a factor of $\mathbb{A}^1$ to each phase.

Tensoring our lattices by $\mathbb{C}^*$ gives two exact sequences of tori:

$$\begin{array}{ccccc} L_{\mathbb{C}^*} & \xrightarrow{Q_{\times}^{\vee}} & (\mathbb{C}^*)^n & \xrightarrow{A_{\times}} & N_{\mathbb{C}^*} \\ M_{\mathbb{C}^*} & \xrightarrow{A_{\times}^{\vee}} & (\mathbb{C}^*)^n & \xrightarrow{Q_{\times}} & L_{\mathbb{C}^*}^{\vee} \end{array}$$

The map $A_{\times}^{\vee}$ provides us with n characters of the torus $M_{\mathbb{C}^*}$. If we pick a vector of coefficients $a \in \mathbb{C}^n$ we can take a linear combination of these characters, this gives us a Laurent monomial:

$$\begin{aligned} W_a : M_{\mathbb{C}^*} &\rightarrow \mathbb{C} \\ x &\mapsto \langle a, A_{\times}^{\vee}(x) \rangle \end{aligned}$$

In explicit coordinates this means

$$W_a = \sum_{i=1}^n a_i \prod_{t=1}^m X_t^{A_{it}}$$

where $X_1, \dots, X_m$ are coordinates on $M_{\mathbb{C}^*}$. This is the *Hori-Vafa mirror* to our toric GIT problem (or abelian GLSM); it's a family of Landau-Ginzburg models parametrized by a .

Since our GIT problem is Calabi–Yau we can choose co-ordinates such that the first column of A is entirely 1's, hence

$$W_a = X_1 \tilde{W}_a$$

where X_1 doesn't appear in $\tilde{W}_a$.

For a generic a the zero locus W_a will be a smooth hypersurface in $M_{\mathbb{C}^*}$. Consider the subset of non-generic a , *i.e.*

$$D_A = \{a \in \mathbb{C}^n, \exists x \in M_{\mathbb{C}^*} \text{ such that } W_a(x) = 0 \text{ and } dW_a(x) = 0\}$$

This, or perhaps its closure, is the *discriminant locus* of the family W_a . This definition is the correct one for general A ; since we're in the Calabi–Yau case the first condition is redundant as $\partial_{X_1} W_a = 0$ implies $W_a = 0$.

The closure of D_A is an affine variety, which is always irreducible and usually a hypersurface [GKZ, Ch. 9]. To understand why this is true observe that D_A is a cone so there is an associated projective variety in $\mathbb{P}^{n-1}$. It's not hard to compute that its projective dual is the closure of the image of $M_{\mathbb{C}^*}$ in $\mathbb{P}^{n-1}$, which is evidently irreducible. But the projective dual to an irreducible variety is always irreducible, and usually a hypersurface [GKZ, Ch. 1]. If D_A is a hypersurface then we denote its defining polynomial by Δ_A .

As well as being a cone D_A is invariant under rescaling the X_i variables, *i.e.* it is invariant under the action of the torus $M_{\mathbb{C}^*}$ on $\mathbb{C}^n$. We can replace D_A with the open subset $D_A \cap (\mathbb{C}^*)^n$ —if D_A is a hypersurface this loses no information—and then the quotient by $M_{\mathbb{C}^*}$ is a subvariety:

$$\nabla_A \subset L_{\mathbb{C}^*}^\vee$$

D_A is a hypersurface iff ∇_A is, and in this case Δ_A is really a function on $L_{\mathbb{C}^*}^\vee$.

4.2.1. Horn uniformization In the Calabi–Yau case there is a useful dominant rational map

$$\mathbb{P}L_{\mathbb{C}} \dashrightarrow \nabla_A$$

called Horn uniformization, given by:

$$[\lambda] \mapsto Q_\times \circ Q^\vee(\lambda)$$

In explicit co-ordinates this says:

$$\lambda_1 : \dots : \lambda_r \mapsto \left(\prod_{i=1}^n \left(\sum_{k=1}^r Q_{ik} \lambda_k \right)^{Q_{i1}}, \dots, \prod_{i=1}^n \left(\sum_{k=1}^r Q_{ik} \lambda_k \right)^{Q_{ir}} \right)$$

Example 4.4. Suppose $L = \mathbb{Z}$ has rank one, and write $(q_1, \dots, q_n)$ for the vector of weights. Then by the above ∇_A consists of the single point

$$q_1^{q_1} \dots q_n^{q_n} \in \mathbb{C}^*$$

(recall we are assuming that no weights are zero). In particular ∇_A is a hypersurface and non-empty.

Let's explain why this works. We have:

$$\partial_{X_s} W_a = \frac{1}{X_s} \sum_{i=1}^n a_i A_{is} \prod_{t=1}^m X_t^{A_{it}}$$

Invariantly, for a fixed $x \in M_{\mathbb{C}^*}$ this says that $dW_a(x)$ is the linear map

$$dW_a(x) : M_{\mathbb{C}} \longrightarrow \mathbb{C}$$

given by composing:

$$M_{\mathbb{C}} \xrightarrow{x^{-1}} M_{\mathbb{C}} \xrightarrow{A^\vee} \mathbb{C}^n \xrightarrow{A_\times^\vee(x)} \mathbb{C}^n \xrightarrow{a} \mathbb{C}$$

Here the first map is the action of the element $x^{-1} \in M_{\mathbb{C}^*}$ on $M_{\mathbb{C}}$, and similarly for the third map. So dW_a has a critical point at x iff $a \circ A_\times^\vee(x)$ annihilates $M_{\mathbb{C}}$, *i.e.* iff

$$a \circ A_\times^\vee(x) = Q^\vee(\lambda)$$

for some $\lambda \in L_{\mathbb{C}}$. So the image of the map

$$\begin{aligned} M_{\mathbb{C}^*} \times L_{\mathbb{C}} &\longrightarrow \mathbb{C}^n \\ (x, \lambda) &\mapsto a = (A_\times^\vee(x))^{-1} Q^\vee(\lambda) \end{aligned}$$

is the subset where W_a has a critical point, and in the Calabi–Yau case this is exactly D_A .

Next we compose this with the quotient map $Q_\times : D_A \dashrightarrow \nabla_A$ and observe that

$$Q_\times(a) = (Q_\times A_\times^\vee(x))^{-1} Q_\times Q^\vee(\lambda) = Q_\times Q^\vee(\lambda)$$

is independent of x , since $Q_\times A_\times^\vee(x) = 1$. Hence $Q_\times \circ Q^\vee$ is a dominant rational map from $L_{\mathbb{C}}$ to ∇_A . Finally, the Calabi–Yau condition implies that this map descends to $\mathbb{P}L_{\mathbb{C}}$.

If ∇_A is a hypersurface it has the same dimension as $\mathbb{P}L_{\mathbb{C}}$, and in this case Horn uniformization is a birational equivalence [GKZ, Ch. 9, Thm 3.3]. The inverse is the logarithmic Gauss map.

4.2.2. Components of the discriminant Recall that the convex hull of the rays $A(e_i)$ is a polytope $\Pi \subset N_{\mathbb{R}}$, which lies in an affine hyperplane of height 1.

Choose a face Γ of Π . Associated to this face there is a Coulomb GIT problem, as described in Sect. 2.3. We consider all the rays that lie in this face, and (abusing notation) write $\Gamma \subset \{1, \dots, n\}$ for the subset that indexes these rays. Then the Coulomb GIT problem is specified by an exact sequence of lattices

$$L_\Gamma \xrightarrow{Q^\vee} \mathbb{Z}^\Gamma \xrightarrow{A_\Gamma} N_\Gamma \tag{4.5}$$

where N_Γ is the sublattice spanned by the face.

We can define a discriminant locus associated to this face in the same way as we did for the whole polytope. For any vector of coefficients $a' \in \mathbb{C}^\Gamma$ there is a Laurent monomial $W'_{a'}$ on the torus $M_{\mathbb{C}^*}^\Gamma$, where M^Γ is the dual lattice to N_Γ . To obtain $W'_{a'}$ from W_a you just delete all the terms that don't correspond to rays on Γ , then since only some variables remain this function descends from $M_{\mathbb{C}^*}$ to the quotient $M_{\mathbb{C}^*}^\Gamma$. Proceeding as before, we obtain a discriminant subset $D_\Gamma \subset \mathbb{C}^\Gamma$, a subvariety

$$\nabla'_\Gamma \subset (L_\Gamma^\vee)_{\mathbb{C}^*}$$

and its preimage:

$$\nabla_\Gamma \subset (L^\vee)_{\mathbb{C}^*}$$

Remark 4.6. What we’ve just done works for any subset of the rays, not just the subsets corresponding to faces of Π . But the faces are the most important. Also note that the Coulomb GIT problems are all Calabi–Yau (Remark 2.4) so we still have Horn uniformization.

Roughly, we are interested in the union of these subvarieties over all faces of Π . However, some faces don’t contribute anything. For example if Γ is a simplex then $L_\Gamma = 0$ so ∇_Γ must be empty; indeed it’s easy to see that D_Γ is just the origin in this case.

More generally suppose Γ contains a ray $A(e_i)$ which is linearly independent of the other rays in Γ . Then D_Γ will be contained in the hyperplane $a'_i = 0$ and hence ∇_Γ is empty. If we want to access D_Γ then we should try deleting this ray $A(e_i)$; this will give us a subface $\Sigma \subset \Gamma$ with one less ray but with $L_\Sigma = L_\Gamma$. Then $D_\Sigma = D_\Gamma$ under the inclusion $\mathbb{C}^\Sigma \hookrightarrow \mathbb{C}^\Gamma$, but ∇'_Σ might be a non-empty subvariety of the torus $(L^\vee_\Sigma)_{\mathbb{C}^*} = (L^\vee_\Gamma)_{\mathbb{C}^*}$. This observation leads to us to the following:

Definition 4.7. A subset $\mathcal{S} \subset \{1, \dots, n\}$ is *minimal* if, for all $i \in \mathcal{S}$, the ray $A(e_i)$ is linearly dependent on the remaining rays $\{A(e_j), j \in \mathcal{S} \setminus i\}$.

A face $\Gamma \subset \Pi$ is *minimal* if the set of all rays lying in Γ is indexed by a minimal subset.

So a face Γ is minimal iff we can remove any ray from Γ without making the linear span smaller. Then we define:

Definition 4.8. The *discriminant locus* $\nabla \subset (L^\vee)_{\mathbb{C}^*}$ is the union of the subvarieties ∇_Γ , for each minimal face $\Gamma \subset \Pi$ such that ∇_Γ is a hypersurface.

The whole polytope Π is minimal since we’re assuming that there are no zero weights. If $\nabla_\Pi = \nabla_A$ is a hypersurface then we call it the *principal component* of ∇ .

Remark 4.9. This definition comes from [GKZ]. It is not entirely clear to us why one disregards the subvarieties ∇_Γ which are not hypersurfaces. In the examples we’ve calculated it makes no difference, *i.e.* each discriminant subvariety of higher codimension is contained in one which is a hypersurface. But we don’t know if this is always true.

If ∇_Γ is a hypersurface we write Δ_Γ for its defining polynomial, then the product of these cuts out the hypersurface ∇ . Gelfand–Kapranov–Zelevinsky modify this by introducing some multiplicities μ_Γ and then taking the product

$$E_A = \prod_{\Gamma} (\Delta_\Gamma)^{\mu_\Gamma}$$

which they call the *principal A-determinant* [GKZ, Ch. 10, 1.B]. The μ_Γ ’s are not relevant for us but there are two important theorems that they prove that are stated in terms of E_A .

Theorem 4.10 [GKZ, Ch. 10, Thm 1.4]. *The Newton polytope of E_A is dual to the secondary fan.*

In fact they give a more precise definition of the *secondary polytope* $\check{\Pi}$ —which is in particular dual to the secondary fan—and their theorem is that the Newton polytope of E_A is $\check{\Pi}$. There is a potential sign confusion here: the theorem is that the cones of the secondary fan are the same as the cones spanned by the inward normal vectors at each vertex of $\check{\Pi}$.

Recall that the secondary fan is the fan of the secondary toric variety $\mathfrak{F}$. This is a compactification of $L_{\mathbb{C}^*}^\vee$ so we can consider the closure:

$$\overline{\nabla} \subset \mathfrak{F}$$

The theorem above suggests that this is a natural choice of compactification for ∇ . In particular it implies:

Corollary 4.11. $\overline{\nabla}$ avoids all the toric fixed points in $\mathfrak{F}$.

Proof. A fixed point is the origin in one of the toric charts. Each chart corresponds to a vertex of the Newton polytope of E_A , and when we write E_A in that chart we get a non-zero constant term.

Recall also that phases of our GIT problem correspond to coherent triangulations of the polytope Π , meaning triangulations induced by a piece-wise linear function [GKZ]. More generally a non-generic stability condition induces a *coherent subdivision* of Π where not all the pieces are simplices. Such a stability condition corresponds to a face of the secondary polytope $\check{\Pi}$ whose vertices are the phases refining this subdivision to a triangulation.

Suppose we fix a coherent subdivision of Π , corresponding to a face $\check{\Gamma} \subset \check{\Pi}$. Now choose one of the pieces of the subdivision, it is some polytope $\Sigma_i \subset \Pi$. As usual we abuse notation and also write $\Sigma_i \subset \{1, \dots, n\}$ for the indexing set of the rays appearing in this polytope. Associated to this subset Σ_i we have a Coloumb GIT problem and a corresponding discriminant locus $\nabla_{\Sigma_i} \subset L_{\mathbb{C}^*}^\vee$ (see Remark 4.6). If Σ_i is a simplex this discriminant locus is empty, so it's only worth considering the non-simplicial pieces of our subdivision.

Going further we can consider the principal determinant E_{Σ_i} , which we may view as a function on $L_{\mathbb{C}^*}^\vee$ by pulling-back under the projection $L^\vee \rightarrow L_{\Sigma_i}^\vee$. The zero locus of E_{Σ_i} consists of the discriminant locus associated to Σ_i as well as the discriminant loci coming from all the faces of Σ_i .

On the other hand, Theorem 4.10 tells us that the face $\check{\Gamma}$ corresponds to some subset of the monomials appearing in E_A . Let us write $(E_A)_{\check{\Gamma}}$ for the sum of this set of monomials.

Theorem 4.12 [GKZ, Ch. 10, Thm 1.12]. *For some positive integer multiplicities μ_i and some non-zero constant v we have*

$$(E_A)_{\check{\Gamma}} = v \prod_i (E_{\Sigma_i})^{\mu_i}$$

where the product runs over the non-simplicial pieces of the subdivision.

In fact we only care about one special case of this theorem: the case when $\check{\Gamma}$ is an edge of $\check{\Pi}$. Such an edge connects two phases, and corresponds to a wall W in the secondary fan. In the secondary toric variety $\mathfrak{F}$ the phases correspond to toric fixed points, and the wall W (or edge $\check{\Gamma}$) corresponds to a toric rational curve

$$C_W \subset \mathfrak{F}$$

connecting the two fixed points. We discussed this in Sect. 1.5.

Corollary 4.13. *The discriminant locus $\overline{\nabla}$ intersects C_W in exactly one point.*

Proof. The intersection of $\overline{\nabla}$ with C_W is the zero locus of the restriction $E_A|_{C_W}$ and this restriction is the sum $(E_A)_{\check{\Gamma}}$ of the monomials appearing in the edge $\check{\Gamma}$. This edge corresponds to a coherent subdivision of Π which has exactly one non-simplicial piece Σ , having two possible triangulations. By Theorem 4.12 the zero locus of $(E_A)_{\check{\Gamma}}$ agrees with the zero locus of E_{Σ} .

But the zero locus of E_{Σ} is the discriminant locus for Coloumb GIT problem associated to Σ . This GIT problem has rank $L_{\Sigma} = 1$ so by Example 4.4 its discriminant locus is a single point. $\square$

In Lemma 4.16 below we will refine this result by identifying which components of $\overline{\nabla}$ can intersect with C_W .

4.3. Faces and subspaces. In Sect. 3.3 we discussed *relevant subspaces* in $L_{\mathbb{R}}^{\vee}$, these index the factors appearing in our SODs. In this section we show that relevant subspaces biject with minimal faces of the polytope Π ; this is an elementary observation but crucial for formulating our conjecture.

Consider a subset $\mathcal{S} \subset \{1, \dots, n\}$ and its complement $\mathcal{S}^c$. Let's consider the Coloumb GIT problem associated to $\mathcal{S}$ and the Higgs GIT problem associated to $\mathcal{S}^c$ (Sect. 2.3). These are related by the following diagram:

$$\begin{array}{ccccc}
 M_{\mathcal{S}^c} & \xrightarrow{A_{\mathcal{S}^c}^{\vee}} & \mathbb{Z}^{\mathcal{S}^c} & \xrightarrow{Q_{\mathcal{S}^c}} & L_{\mathcal{S}^c}^{\vee} \\
 \downarrow & & \downarrow & & \downarrow \\
 M & \xrightarrow{A^{\vee}} & \mathbb{Z}^n & \xrightarrow{Q} & L^{\vee} \\
 \downarrow & & \downarrow & & \downarrow \\
 M_{\mathcal{S}} & \xrightarrow{A_{\mathcal{S}}^{\vee}} & \mathbb{Z}^{\mathcal{S}} & \xrightarrow{Q_{\mathcal{S}}} & L_{\mathcal{S}}^{\vee}
 \end{array} \tag{4.14}$$

The middle column is obviously exact, the other columns are exact modulo torsion. Let us also write

$$H_{\mathcal{S}^c} \subset L_{\mathbb{R}}^{\vee}$$

for the subspace spanned by $L_{\mathcal{S}^c}^{\vee}$.

As a special case we could consider a face of the polytope Π and let Γ be the indexing set for the rays on that face. Then we get an associated subspace $H_{\Gamma^c} \subset L_{\mathbb{R}}^{\vee}$.

Proposition 4.15. *The map $\Gamma \mapsto H_{\Gamma^c}$ is a bijection between the minimal faces of Π and the relevant subspaces of $L_{\mathbb{R}}^{\vee}$.*

Proof. Recall from Definition 4.7 that a subset $\mathcal{S} \subset \{1, \dots, n\}$ is called minimal if the set of rays $A(\mathcal{S}) \subset N$ has the property that every ray in $A(\mathcal{S})$ is linearly dependent on the remaining rays in $A(\mathcal{S})$. This is the statement that no basis vectors map to zero under the map $\mathbb{Z}^{\mathcal{S}} \rightarrow L_{\mathcal{S}}^{\vee}$, or equivalently that the only weights lying in $H_{\mathcal{S}^c}$ are $Q(\mathcal{S}^c)$.

Conversely, pick a subspace $H \subset L_{\mathbb{R}}^{\vee}$ which is spanned by the weights it contains, and let $\mathcal{S}$ be the set of weights which do not lie in H . Then $H = H_{\mathcal{S}^c}$ and $\mathcal{S}$ is minimal. Hence the assignment $\mathcal{S} \mapsto H_{\mathcal{S}^c}$ is a bijection between the minimal subsets of $\{1, \dots, n\}$ and the subspaces of $L_{\mathbb{R}}^{\vee}$ which are spanned by the weights they contain.

It follows immediately from part (ii) of Definition/Lemma 3.10 that the subspace $H_{\mathcal{S}^c}$ is relevant iff there is a vector $k \in \mathbb{Z}^{\mathcal{S}^c}$ with strictly positive entries which maps

to zero under Q . Such a vector is exactly an element $k \in N^\vee$ such that $k(A(e_i)) = 0$ if $i \in S$ and $k(A(e_i)) > 0$ if $i \in S^c$. Since the polytope Π lives in an affine hyperplane of height 1, the existence of such a k is the statement that S is all the rays on a face of Π .

The zero subspace $H = 0$ is always relevant, and since we assume there are no zero weights it corresponds to the whole polytope Π . The empty set is a face of Π (in the sense of the above proof), it corresponds to the subspace $H = L_{\mathbb{R}}^\vee$ which is therefore relevant.

4.4. The conjecture. Let W be a wall separating two chambers in the secondary fan. Recall that we have the following two objects associated to W :

- (1) A toric variety Z_W . The wall W has an associated Higgs GIT problem, and Z_W is the phase of this problem coming from a character on the relative interior of W .
- (2) A toric rational curve C_W in the secondary stack. W is a codimension 1 cone in the secondary fan and C_W is the associated curve.

We can decompose $D^b(Z_W)$ using the algorithm of Sect. 3, and the factors that appear are indexed by the relevant subspaces $H \subset L_{\mathbb{R}}^\vee$ contained in the hyperplane $\langle W \rangle$. Each such subspace defines a Higgs GIT problem with a non-empty minimal phase Z_H , and by Theorem 3.12 the multiplicity of $D^b(Z_H)$ in $D^b(Z_W)$ is well-defined.

Relevant subspaces correspond (by Proposition 4.15) to minimal faces Γ of the polytope Π , and these in turn index the components of the discriminant locus. As discussed in Sect. 1.5 we are interested in the intersection of $\overline{\nabla}_\Gamma$ with the curve C_W .

Lemma 4.16. *Let Γ be a minimal face. If H_{Γ^c} is not contained in W then $\overline{\nabla}_\Gamma$ doesn't meet the curve C_W .*

Proof. Consider the projection map

$$\pi : L^\vee \rightarrow L_\Gamma^\vee$$

or its real version $L_{\mathbb{R}}^\vee \rightarrow (L_\Gamma^\vee)_{\mathbb{R}}$, whose kernel is H_{Γ^c} . This map takes a stability condition for the original GIT problem and restricts it to give one for the Coloumb GIT problem associated to Γ . If we take a chamber of stability conditions and restrict them then they will all lie in a single chamber for the Coloumb GIT problem (if two stability conditions induce the same triangulation of Π then they evidently induce the same triangulation of the face Γ). This says that π is a map of fans, between the secondary fan for the original problem and the secondary fan for the Coloumb problem, hence it induces a toric morphism

$$\pi : \mathfrak{F} \rightarrow \mathfrak{F}_\Gamma$$

between the two secondary toric varieties.

Recall that ∇_Γ is defined as the preimage of the discriminant locus $\nabla'_\Gamma \subset (L_\Gamma)_{\mathbb{C}^*}^\vee$ under the projection $\pi : (L^\vee)_{\mathbb{C}^*} \rightarrow (L_\Gamma)_{\mathbb{C}^*}^\vee$. Since π extends to the toric boundary we can also say that $\overline{\nabla}_\Gamma \subset \mathfrak{F}$ is contained in the preimage of $\overline{\nabla}'_\Gamma \subset \mathfrak{F}_\Gamma$.

The wall W is a codimension 1 cone in $L^\vee$. If it doesn't contain H_{Γ^c} then $\pi(W)$ is a top-dimensional cone in $L_\Gamma^\vee$ so $\pi(C_W)$ is one of the toric fixed points in $\mathfrak{F}_\Gamma$. Corollary 4.11 says that $\overline{\nabla}'_\Gamma$ avoids all the toric fixed points hence $\overline{\nabla}_\Gamma$ misses C_W .

Conjecture 4.17. *Let $W \subset L_{\mathbb{R}}^\vee$ be a wall, let $\Gamma \subset \Pi$ be a minimal face and let $H = H_{\Gamma^c}$ be the corresponding relevant subspace. Assume that $H \subseteq \langle W \rangle$. Write*

$n_{\Gamma, W}$ for the multiplicity of $D^b(Z_H)$ in $D^b(Z_W)$, and
 $m_{\Gamma, W}$ for the intersection multiplicity of $\overline{\nabla}_\Gamma$ with C_W .

Then $n_{\Gamma, W} = m_{\Gamma, W}$.

Remark 4.18. We could allow the case when H doesn't lie in $\langle W \rangle$: then $D^b(Z_H)$ is not a factor in $D^b(Z_W)$ so we should set $n_{\Gamma, W} = 0$, and by Lemma 4.16 $m_{\Gamma, W} = 0$ also.

We will now prove various special cases of this conjecture. The most straight-forward case is when $\text{rank } L_\Gamma = 1$ so H is a hyperplane, hence $H = \langle W \rangle$.

Proposition 4.19. *If $\text{rank } L_\Gamma = 1$ then $n_{\Gamma, W} = m_{\Gamma, W} = 1$.*

Proof. In this case Z_H is the minimal phase for the Higgs GIT problem that produces Z_W , so $n_{\Gamma, W} = 1$.

As in Lemma 4.16 we consider the map $\pi : \mathfrak{F} \rightarrow \mathfrak{F}_\Gamma$. This map induces an isomorphism from C_W to $\mathfrak{F}_\Gamma$. The discriminant locus $\nabla'_\Gamma \subset \mathfrak{F}_\Gamma$ is a single non-fixed point (Example 4.4) and $\overline{\nabla}_\Gamma$ is its pre-image, as a divisor on $\mathfrak{F}$. So intersecting this divisor with C_W gives $m_{\Gamma, W} = 1$.

Remark 4.20. This proposition includes the case when L itself has rank 1 and hence $H = W$ is the origin. This is a vacuous case of our conjecture: C_W is the whole of $\mathfrak{F}$, there is only the principal component of $\overline{\nabla}$ which is a single point, Z_W is a point, and the decomposition of $D^b(Z_W)$ is trivial.

We can get a less trivial special case by increasing the rank by one.

Proposition 4.21. *If $\text{rank } L_\Gamma = 2$ then $n_{\Gamma, W} = m_{\Gamma, W}$.*

Proof. In this case H is a hyperplane in $\langle W \rangle$. Since the projection $\pi : L^\vee \rightarrow L_\Gamma^\vee$ is a map of fans $\pi(W)$ must be a ray, so W must lie completely on one side of H . Pick a primitive one-parameter subgroup λ normal to W , then the GIT problem producing Z_W consists of the vector space V^λ —these are weights that lie in $\langle W \rangle$ —acted on by the torus T/λ . Then H is normal to some primitive one-parameter subgroup $\mu \in L/\langle \lambda \rangle$ and we orient μ so that it pairs positively with W . Recall that Z_W is defined to be the phase associated to a generic character in W , so such a character pairs positively with μ . For the decomposition of $D^b(Z_W)$ the important quantity is

$$\kappa = (\det V^\lambda)(\mu) = \sum_{\text{weights } q_i \in \langle W \rangle} q_i(\tilde{\mu})$$

for any $\tilde{\mu} \in L$ lifting μ , and

$$n_{\Gamma, W} = \max\{\kappa, 0\}$$

(see Remark 3.13).

Now we compute the intersection multiplicity $m_{\Gamma, W}$. To start with, let's assume that $\Gamma = \Pi$ so L itself has rank 2 and H is the origin. Then we wish to compute the intersection multiplicity of the principal component $\overline{\nabla}_A$ with the boundary curve C_W . To do this we use the Horn uniformization map

$$\mathbb{P}(L_\Gamma)_\mathbb{C} \dashrightarrow \overline{\nabla}_A$$

from Sect. 4.2.1, which in this case is actually a morphism since $\mathbb{P}(L_{\mathbb{C}}) \cong \mathbb{P}^1$. In explicit co-ordinates, as a rational map to $(\mathbb{C}^*)^2$, this is given by:

$$\lambda_1 : \lambda_2 \mapsto \left(\prod_{i=1}^n (Q_{i1}\lambda_1 + Q_{i2}\lambda_2)^{Q_{i1}}, \prod_{i=1}^n (Q_{i1}\lambda_1 + Q_{i2}\lambda_2)^{Q_{i2}} \right)$$

Without loss of generality we may assume that W is the ray through $(1, 0)$. This ray in the secondary stack corresponds to a partial compactification of the torus $(\mathbb{C}^*)^2$, the subset:

$$\mathbb{C} \times \mathbb{C}^* \subset \mathfrak{F}$$

The subset where the first co-ordinate is zero is C_W with its fixed points deleted. Since $\overline{\nabla}_A$ avoids the fixed points the only way that $\lambda_1 : \lambda_2$ can map to C_W is if there exists an i such that $Q_{i1}\lambda_1 + Q_{i2}\lambda_2 = 0$ and $Q_{i2} = 0$, hence $\lambda_1 : \lambda_2 = 0 : 1$. Then the intersection multiplicity is given by

$$\sum_{i \mid Q_{i2}=0} Q_{i1}$$

if this sum is strictly positive, and zero otherwise. But these rows of Q are precisely the weights q_i that lie on $\langle W \rangle$, and we may set $\tilde{\mu} = (1, 0)^\top$, so this sum is κ and hence $m_{\Gamma, W} = n_{\Gamma, W}$ in this case.

To finish we must compute $m_{\Gamma, W}$ for rank $L_\Gamma = 2$ but $\Gamma \subsetneq \Pi$. Once again we use the projection $\pi : \mathfrak{F} \rightarrow \mathfrak{F}_\Gamma$. It maps W to a wall W' for the rank 2 GIT problem, and hence it maps C_W isomorphically (at least away from its fixed points) onto the boundary curve $C_{W'} \subset \mathfrak{F}_\Gamma$. Since $\overline{\nabla}_\Gamma$ is the preimage of $\overline{\nabla}'_\Gamma \subset \mathfrak{F}_\Gamma$ it is enough to compute the intersection multiplicity of $\overline{\nabla}_\Gamma$ with $C_{W'}$ inside the two-dimensional space $\mathfrak{F}_\Gamma$. But this was the calculation we just performed, and the result is still $\max\{\kappa, 0\}$ since $q_i \in \langle W \rangle$ iff $\pi(q_i) \in \langle W' \rangle$. $\square$

Remark 4.22. The above result and its proof are quite close to [HLS_h, Prop. 4.4.].

Theorem 4.23. *If rank $L = 2$ then Conjecture 4.17 holds.*

Proof. The wall W is a ray and there are only two possibilities for H : either $H = \langle W \rangle$ (if this is a relevant subspace) or $H = 0$. The first case is covered by Proposition 4.19 and the second by Proposition 4.21.

The main obstacle to extending our proofs to higher rank is the fact that Horn uniformization may no longer be a morphism so it becomes harder to compute the intersection multiplicity $m_{\Gamma, W}$. However in special cases it is still possible to verify the conjecture—see [Kite, Sect. 10.2] for some more examples.

Acknowledgements. E.S. would like to thank Paul Aspinwall and Lars Louder for helpful conversations. This Project has received funding from the European Research Council (ERC) under the European Union Horizon 2020 research and innovation programme (Grant Agreement No. 725010). A.K. was supported by the EPSRC [EP/L015234/1] via the LSGNT Centre for Doctoral Training.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is

not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

- [AL] Anno, R., Logvinenko, T.: Spherical DG-functors. *J. EMS* **19**(9), 2577–2656. [arXiv:1309.5035](#)
- [APW] Aspinwall, P., Plesser, M., Wang, K.: Mirror symmetry and discriminants. [arXiv:1702.04661](#)
- [BFK] Ballard, M., Favero, D., Katzarkov, L.: Variation of geometric invariant theory quotients and derived categories. *J. Reine Angew. Math.* **746**, 235–303 (2019). [arXiv:1203.6643](#)
- [BBS] Böhning, C., Graf von Bothmer, H.-C., Sosna, P.: On the Jordan–Hölder property for geometric derived categories. *Adv. Math.* **256**, 479–492. [arXiv:1211.1229](#)
- [DS] Donovan, W., Segal, E.: Mixed braid group actions from deformations of surface singularities. *Commun. Math. Phys.* **335**(1), 497–543 (2015). [arXiv:1310.7877](#)
- [GKZ] Gelfand, I., Kapranov, M., Zelevinsky, A.: *Discriminants, Resultants, and Multidimensional Determinants*. Birkhäuser, Boston (1994)
- [HL] Halpern-Leistner, D.: The derived category of a GIT quotient. *J. AMS* **28**(3), 871–912. [arXiv:1203.0276](#)
- [HLSam] Halpern-Leistner, D., Sam, S.: Combinatorial constructions of derived equivalences. *J. AMS* **33**, 735–773 (2020). [arXiv:1601.02030](#)
- [HLSh] Halpern-Leistner, D., Shipman, I.: Autoequivalences of derived categories via geometric invariant theory. *Adv. Math.* **303**, 1264–1299. [arXiv:1303.5531](#)
- [HHP] Herbst, M., Hori, K., Page, D.: Phases of $N = 2$ Theories in 1+1 Dimensions with Boundary, DESY-07-154, CERN-PH-TH/2008-048 (2008). [arXiv:0803.2045](#)
- [HW] Herbst, M., Walcher, J.: On the unipotence of autoequivalences of toric complete intersection Calabi–Yau categories. *Math. Ann.* **353**(3), 783–802. [arXiv:0911.4595](#)
- [Kaw] Kawamata, Y.: Derived categories of toric varieties. *Mich. Math. J.* **54**(3), 517–535 (2006). [arXiv:0503.102](#)
- [Kite] Kite, A.: Fundamental group actions on derived categories. PhD thesis, (2018). www.homepages.ucl.ac.uk/~ucaheps/
- [Kuz] Kuznetsov, A.: A simple counterexample to the Jordan–Hölder property for derived categories. [arXiv:1304.0903](#)
- [Seg] Segal, E.: Equivalences between GIT quotients of Landau–Ginzburg B-models. *Commun. Math. Phys.* **304**, 411–432 (2011). [arXiv:0910.5534](#)

Communicated by C. Schweigert