

MIXED BRAID GROUP ACTIONS FROM DEFORMATIONS OF SURFACE SINGULARITIES

WILL DONOVAN AND ED SEGAL

ABSTRACT. We consider a set of toric Calabi–Yau varieties which arise as deformations of the small resolutions of type A surface singularities. By careful analysis of the heuristics of B-brane transport in the associated GLSMs, we predict the existence of a mixed braid group action on the derived category of each variety, and then prove that this action does indeed exist. This generalizes the braid group action found by Seidel and Thomas for the undeformed resolutions. We also show that the actions for different deformations are related, in a way that is predicted by the physical heuristics.

CONTENTS

1. Introduction	1
1.1. Main result	3
1.2. Outline	3
2. Examples and results	4
2.1. The A_1 examples	4
2.2. Statement of results	7
2.3. Future directions	10
3. Toric calculations	10
3.1. Representations of the free quiver	10
3.2. Subvarieties associated to partitions	12
3.3. Families of rational curves	17
3.4. Derived equivalences from wall-crossing	19
4. Fayet–Iliopoulos parameter spaces	21
4.1. Toric Mirror Symmetry heuristics	21
4.2. Results for our examples	22
4.3. Large-radius limits	24
4.4. Fundamental groups	24
5. Mixed braid group actions	25
5.1. Heuristics for the generators	25
5.2. Braid relations on the ambient space	27
5.3. The poset of groupoid actions	33
Appendix A. List of notations	36
References	37

1. INTRODUCTION

Let V be a vector space, and let T be a torus acting on V . If we pick a character θ of T , we have a stability condition for the associated GIT problem, and can form a GIT quotient $X_\theta = V //_\theta T$. We'll borrow some physics terminology and refer to these different possible GIT quotients as *phases*.

2010 *Mathematics Subject Classification*. Primary 14F05, 18E30; Secondary 14J33, 20F36.

W.D. is grateful for the support of EPSRC grant EP/G007632/1, and E.S. for the support of an Imperial College Junior Research Fellowship.

There is now a well-developed general theory [HL12, BFK12], inspired by the physics paper [HHP08] but going back to ideas of Kawamata, that allows us to compare the derived categories of different phases. In particular, if T acts through $\mathrm{SL}(V)$, so that all the phases X_θ are Calabi–Yau, then all their derived categories $D^b(X_\theta)$ will be equivalent. These equivalences are not canonical: whenever we cross a wall in the space of stability conditions, we have a countably infinite set of equivalences between the two phases that lie on either side. Consequently we can produce autoequivalences of the derived category of a single phase X_θ : we pass through an equivalence to the derived category of a different phase, then we pass back again, using a different equivalence. In this way we can produce a whole group of autoequivalences acting on $D^b(X_\theta)$. Then we can ask the question: what group is it?

To get a prediction, we can turn to quantum field theory. The data of T acting on V determines a Gauged Linear Sigma Model, a particular kind of supersymmetric 2-dimensional QFT. The model has a parameter, called the (complexified) Fayet–Iliopoulos parameter, which takes values in a space which we’ll call \mathcal{F} . This parameter space \mathcal{F} is a complex manifold (or orbifold) locally modelled on the dual Lie algebra \mathfrak{t}^* of T , but it can have non-trivial global structure. In particular limiting regions in \mathcal{F} , the theory is expected to reduce to a sigma model with target one of the phases X_θ , with different phases occurring at different limits.

There is also believed to be a triangulated category associated to the GLSM, called the category of B-branes. If we assume the Calabi–Yau condition then this category is independent of the FI parameter, but only up to isomorphism. We should think of it as a ‘local system’ of categories over the space \mathcal{F} . When we approach one of the limiting regions, the category of B-branes becomes identified (not quite canonically¹) with the derived category $D^b(X_\theta)$ of the corresponding phase. If we travel along a path in \mathcal{F} from one limit to another then we can perform ‘parallel transport’ of the B-branes, and produce a derived equivalence between the corresponding phases. Similarly if we travel around a large loop in \mathcal{F} then we get a monodromy autoequivalence, acting on the derived category of a single phase. The derived equivalences and autoequivalences that arise from this ‘brane transport’ are believed to be exactly the ones that arise from variation of GIT quotient in the mathematical constructions mentioned above. So this picture tells us which group we should expect to see acting on $D^b(X_\theta)$: it’s the fundamental group of \mathcal{F} .²

Providing a rigorous definition of this QFT is an immensely difficult problem which will probably not be resolved for many years, and consequently an intrinsic definition of the FI parameter space \mathcal{F} is not known. Fortunately, what is known is a completely precise heuristic recipe that tells us how to construct \mathcal{F} . The justification for this recipe (at least mathematically) is toric mirror symmetry, since \mathcal{F} is related to the complex moduli space of the mirror theory.³

If we apply this recipe, and can compute $\pi_1(\mathcal{F})$, then our heuristics will have given us a precise prediction about derived autoequivalences which we can then attempt to prove. In this paper we’re going to carry out this program, and prove the prediction, for a particular set of examples.

Our examples arise from a well-known piece of geometry, namely resolutions of the A_k surface singularities, and their deformations. The resulting set of GLSMs has another important feature: certain phases of each model embed, in a natural way, into particular phases of some of the other models. In other words, we’re considering some toric varieties, and also some torically-embedded subvarieties.

When two GLSMs are related in this way then a rough physical argument suggests that there should be a map from the FI parameter space for the ambient variety to the FI parameter space for the subvariety. We don’t know a mathematical justification for this argument, but in our examples we see that these maps do indeed exist, and in fact they’re covering maps. Consequently the fundamental groups of the two FI parameter spaces are related, and this suggests that the brane transport autoequivalences

¹The ambiguity is tensoring by line bundles.

²This picture doesn’t tell us whether or not the action is faithful, however.

³The fact that T is a torus is crucial here, if we replace it with a non-abelian group then as far as we know no recipe for \mathcal{F} exists.

for the two models should also be related. We show that these predicted relationships between the examples do indeed hold.

1.1. Main result. We now briefly state our main theorem, leaving some of the details to Section 2. Firstly, let Y denote a minimal resolution of a local model for an A_k surface singularity, as described explicitly in Section 2.2. In a seminal paper [ST01], Seidel and Thomas construct a faithful action $B_{k+1} \curvearrowright \mathrm{D}^b(Y)$ of the braid group on $k+1$ strands on the derived category of Y , by spherical twists. Now given a partition Γ of the set of strands, we may define:

- (1) A *mixed braid group* B_Γ , namely that subgroup of the braid group B_{k+1} consisting of braids which respect the partition Γ (see (2.10)).
- (2) A certain deformation X_Γ of the surface Y (see Section 2.2.1). This deformation has s parameters, where $s+1$ is the number of pieces in the partition.

Theorem A (Corollary 5.4). *For each partition Γ of $\{0, \dots, k\}$, there is a faithful action of the mixed braid group $B_\Gamma \curvearrowright \mathrm{D}^b(X_\Gamma)$ on the derived category of the deformation X_Γ .*

When Γ has only one part, the theorem recovers the original Seidel–Thomas braid group action on Y : in this case, the usual braid generators t_i act by spherical twists. When Γ is the finest partition, with $k+1$ parts, X_Γ is a versal deformation of Y , and B_Γ is the *pure braid group*, consisting of braids which return each strand to its original position. In this case, the braids t_i^2 act by family spherical twists: these fibre over the deformation base (Section 3.4.3), deforming the original spherical twists on Y .

The actions given in Theorem A also appear in the context of geometric representation theory. They may be obtained from work of Bezrukavnikov–Riche [BR12, Section 4], who construct such actions by representation-theoretic means on slices of the Grothendieck–Springer simultaneous resolution. Similar constructions appear in Cautis–Kamnitzer [CK12, Section 2.5], in the study of categorified quantum group representations. However our focus is somewhat different, as we see these actions as arising from the toric geometry of the varieties X_Γ , and our main aim is to relate this to the physical heuristics.

We remark very briefly on our method of proof (which is quite different to [BR12, CK12] or [ST01]). In our construction, the autoequivalences which generate the group action turn out to correspond to ‘windows’ in the derived category of the stack associated to a certain GIT problem. After a careful analysis of these windows, a conceptually simple proof of the braid relations emerges, in Section 5.2. We hope that this approach will be useful in more general situations, where the machinery of [HL12, BFK12] can be applied. See Remark 5.14 for more discussion on this point.

1.2. Outline. The structure of this paper is as follows.

- Section 2 explains some simple examples, and then gives a detailed statement of our main result in Section 2.2.
- Section 3 details the toric geometry of our examples X_Γ , and of their flops. This yields derived equivalences between phases in Section 3.4.
- Section 4 recalls the heuristics which allow us to describe the FI parameter spaces for our examples in Section 4.1, and explains how to apply them in Section 4.2. We go on to identify the large-radius limits in Section 4.3, and the fundamental groups which we expect to act on the $\mathrm{D}^b(X_\Gamma)$ in Section 4.4.
- Section 5 proves the physical prediction of Section 4. Specifically, we show that a certain groupoid T_Γ acts (faithfully) via derived equivalences and autoequivalences on the phases of X_Γ . The isotropy group of T_Γ is B_Γ , and so Theorem A follows as an immediate corollary. Section 5.3 proves the expected relationships between the actions for different partitions Γ .

Appendix A lists our main notations, along with cross-references to their explanations.

Acknowledgements. This project grew out of enlightening discussions between W.D. and Iain Gordon. W.D. is also grateful to Adrien Brochier, Jim Howie and Michael Wemyss for valuable conversations and helpful suggestions. E.S. would like to thank Tom Coates, Kentaro Hori and especially Hiroshi Iritani for their patient explanations and many extremely useful comments and suggestions.

2. EXAMPLES AND RESULTS

2.1. The A_1 examples.

2.1.1. *The 3-folds.* To set the stage, we recall some very well-known constructions. Consider the stack

$$\mathcal{X} = [\mathbb{C}_{b_0, a_0, b_1, a_1}^4 / \mathbb{C}^*]$$

where the \mathbb{C}^* acts with weights $(1, -1, -1, 1)$. The coarse moduli space of this stack is the 3-fold ODP singularity $\{uv = z_0z_1\}$, which is the versal deformation of the A_1 surface singularity.⁴ If we pick a character of \mathbb{C}^* then we can form the associated GIT quotient. There are two possible quotients which we'll denote by X_+ and X_- : each one is isomorphic to the total space of the bundle $\mathcal{O}(-1)_{\mathbb{P}^1}^{\oplus 2}$. This is the standard Atiyah flop.

These two resolutions X_+ and X_- have a common roof

$$\begin{array}{ccc} & \mathcal{O}(-1, -1)_{\mathbb{P}^1 \times \mathbb{P}^1} & \\ & \parallel & \\ & \tilde{X} & \\ \swarrow \pi_+ & & \searrow \pi_- \\ X_+ & & X_- \end{array} \quad (2.1)$$

Bondal and Orlov [BO95] showed that this correspondence induces a derived equivalence

$$F = (\pi_-)_*(\pi_+)^* : D^b(X_+) \xrightarrow{\sim} D^b(X_-).$$

There is another approach to this derived equivalence, introduced by the second-named author in [Seg11] and based on the physics arguments of [HHP08]. We view the GIT quotients as open substacks

$$\iota_{\pm} : X_{\pm} \hookrightarrow \mathcal{X},$$

and we define certain subcategories of $D^b(\mathcal{X})$, nicknamed ‘windows’, by

$$\mathcal{W}(k) = \langle \mathcal{O}(k), \mathcal{O}(k+1) \rangle \subset D^b(\mathcal{X})$$

for $k \in \mathbb{Z}$, i.e. $\mathcal{W}(k)$ is the full triangulated subcategory generated by this pair of line bundles. Then it is easy to show that for any k both functors

$$\iota_{\pm}^* : \mathcal{W}(k) \rightarrow D^b(X_{\pm})$$

are equivalences, so we can define a set of derived equivalences between X_+ and X_- by

$$\psi_k : D^b(X_+) \xrightarrow{(\iota_+^*)^{-1}} \mathcal{W}(k) \xrightarrow{\iota_-^*} D^b(X_-). \quad (2.2)$$

The relationship between these two approaches is the following statement:

Proposition 2.3. *The window equivalence ψ_{-1} and the Bondal–Orlov equivalence F coincide:*

$$\psi_{-1} = F$$

Proof. Consider the autoequivalence

$$\hat{F} = (\iota_-^*)^{-1} F \iota_+^* : \mathcal{W}(-1) \rightarrow \mathcal{W}(-1).$$

The proposition is the statement that \hat{F} is the identity. Since $\mathcal{W}(-1)$ is generated by $\mathcal{O}(-1)$ and \mathcal{O} , it's enough to check that \hat{F} acts as the identity on these two objects, and on all morphisms between them. For the objects, one can check that F sends \mathcal{O} to \mathcal{O} and $\mathcal{O}(-1)$ to $\mathcal{O}(1)$, which is the required statement because $\iota_-^* \mathcal{O}(-1) = \mathcal{O}(1)$. For the morphisms, we note that all morphisms between these two line bundles come from functions on \mathbb{C}^4 (there are no higher Exts), and F is the identity away from a subvariety of codimension 2. \square

⁴The variables here are the invariant functions $u = a_0a_1$, $v = b_0b_1$, $z_0 = a_0b_0$, and $z_1 = a_1b_1$.

The other ψ_k can be obtained by modifying F by the appropriate line bundle on \tilde{X} . Also see [HLS, Section 3.1] for a more general version of this statement.

Now we connect the above with the GLSM heuristics. For this GIT problem, the FI parameter space turns out to be a \mathbb{P}^1 with three punctures:

$$\mathcal{F}_{\mathcal{X}} = \mathbb{P}^1 - \{1:0, 0:1, 1:1\}$$

The first two punctures (or more accurately, neighbourhoods of these punctures) are ‘large-radius’ (LR) limits corresponding to the two phases X_{\pm} . The third puncture is the ‘conifold point’ where the theory becomes singular (we’ll explain this picture further in Section 4). Traversing a loop around one of the large-radius limits is supposed to produce an autoequivalence on the derived category of the corresponding phase, and the correct autoequivalence is just tensoring by the $\mathcal{O}(1)$ line bundle.

More interestingly, travelling along a path which starts near one LR limit and ends near the other should produce a derived equivalence between the two phases. Identify $\mathcal{F}_{\mathcal{X}}$ with $\mathbb{C}_{\zeta} - \{0, 1\}$, then take $\log(\zeta)$ to get the infinite-sheeted cover $\mathbb{C}_{\log(\zeta)} - 2\pi i\mathbb{Z}$. The LR limits lie at $\operatorname{Re}(\log(\zeta)) \ll 0$ and $\operatorname{Re}(\log(\zeta)) \gg 0$, so consider a path that travels from one to the other with $\operatorname{Re}(\log \zeta)$ increasing monotonically. The homotopy class of such a path is determined by which interval it lies in when it crosses the imaginary axis. The interpretation of the physical arguments in [HHP08] is that the path that goes through the interval between $2\pi ik$ and $2\pi i(k+1)$ produces the derived equivalence ψ_k described above.

Now consider a loop that begins in $\operatorname{Re}(\log(\zeta)) \ll 0$ and circles the origin once clockwise, without encircling any other punctures. It follows that this loop produces the autoequivalence

$$(\psi_{-1})^{-1} \circ \psi_0 : D^b(X_+) \rightarrow D^b(X_+).$$

In [DS12] autoequivalences of this form were called ‘window shifts’. It was shown there, for a larger class of examples, that these autoequivalences can be described as twists of certain spherical functors. This particular window shift is equal to a Seidel–Thomas spherical twist [ST01] around the spherical object $\mathcal{S} = \mathcal{O}_{\mathbb{P}^1}(-1)$ in $D^b(X_+)$.⁵ That is, it’s equal to a functor

$$T_{\mathcal{S}} : \mathcal{E} \mapsto \operatorname{Cone}(\mathcal{S} \otimes \operatorname{Hom}(\mathcal{S}, \mathcal{E}) \rightarrow \mathcal{E}) \quad (2.4)$$

Given Proposition 2.3, we can also describe this autoequivalence as a composition of ‘flop’ and ‘flop back again’ (with the appropriate line bundles inserted).

2.1.2. *The surfaces.* Now consider the stack

$$\mathcal{Y} = [\mathbb{C}_{s_0, s_1, p}^3 / \mathbb{C}^*]$$

where the \mathbb{C}^* acts with weights $(1, 1, -2)$. The coarse moduli space of this stack is the A_1 surface singularity $\{uv = z^2\}$. There are again two possible GIT quotients: Y_+ is the total space of $\mathcal{O}(-2)_{\mathbb{P}^1}$, and Y_- is the orbifold $[\mathbb{C}^2 / \mathbb{Z}_2]$. Most of the analysis of the previous example can be repeated verbatim, so there is an equivalence

$$\tilde{F} : D^b(Y_+) \rightarrow D^b(Y_-)$$

coming from the common birational roof, and there are windows $\tilde{\mathcal{W}}(k) \subset D^b(\mathcal{Y})$ defined in the same way and giving rise to equivalences

$$\tilde{\psi}_k : D^b(Y_+) \rightarrow D^b(Y_-).$$

Also, $\tilde{F} = \tilde{\psi}_{-1}$ by the same argument.

The FI parameter space for this example turns out to be an orbifold: it’s a \mathbb{P}^1 with two punctures

$$\mathcal{F}_{\mathcal{Y}} = \mathbb{P}^1 - \{1:0, 1:1\}$$

and an orbifold point at $0:1$ with \mathbb{Z}_2 isotropy group. The phases Y_+ and Y_- correspond to the regions near $1:0$ and $0:1$ respectively, and $1:1$ is again a conifold point. Using the functors $\tilde{\psi}_k$, and tensoring by line bundles, we can again produce an action of $\pi_1(\mathcal{F}_{\mathcal{Y}})$ on the derived category of either of the two

⁵This particular relation was well-known long before [DS12]: see [Kaw02, Example 5.10] for a much earlier discussion.

phases. The existence of the orbifold point in \mathcal{F}_Y corresponds to the fact that on Y_- , tensoring with $\mathcal{O}(1)$ has order 2. Also, one can show that looping around the conifold point produces a spherical twist autoequivalence

$$T_{\tilde{S}} : D^b(Y_+) \rightarrow D^b(Y_+),$$

where \tilde{S} is the spherical object $\mathcal{O}_{\mathbb{P}^1}(-1)$ in $D^b(Y_+)$.

2.1.3. Relating the examples. Now we can ask about the relationship between these two examples. It is a well-known fact that we can include Y_+ as a subvariety into either of X_+ or X_- as the zero locus of the invariant function $b_0a_0 - b_1a_1$ on \mathbb{C}^4 , so that we have

$$j_{\pm} : Y_+ \hookrightarrow X_{\pm}.$$

Indeed if we put the obvious symplectic form on \mathbb{C}^4 then we can view Y_+ as a hyperkähler quotient, and the invariant function as the complex moment map.

We claim that this fact is reflected in the FI parameter spaces for the two examples. Specifically, we make the observation that there is a 2-to-1 covering map

$$\mathcal{F}_X \rightarrow \mathcal{F}_Y$$

sending the conifold point to the conifold point, and identifying the two LR limits in \mathcal{F}_X with the limit corresponding to Y_+ in \mathcal{F}_Y . To see this, just notice that we can identify \mathcal{F}_Y with

$$[\mathbb{C}_{\zeta} - \{0, 1\} / \mathbb{Z}_2]$$

where the involution is $\zeta \mapsto 1/\zeta$. We'll explain this more systematically in Section 4.

Now this covering map suggests that the derived equivalences that arise from the two examples are related, in a way that corresponds to the map between the fundamental groupoids of \mathcal{F}_X and \mathcal{F}_Y . For example, a loop around one of the LR limits in \mathcal{F}_X projects to a loop around the LR limit in \mathcal{F}_Y , which corresponds to the obvious fact that

$$j_{\pm}^*((-) \otimes \mathcal{O}(1)) = j_{\pm}^*(-) \otimes \mathcal{O}(1).$$

More interestingly, a path between the two LR limits in \mathcal{F}_X projects to a loop around the conifold point in \mathcal{F}_Y , which suggests the following proposition.

Proposition 2.5. *The square below commutes.*

$$\begin{array}{ccc} D^b(X_+) & \xrightarrow{\psi_0} & D^b(X_-) \\ j_+^* \downarrow & & \downarrow j_-^* \\ D^b(Y_+) & \xrightarrow{T_{\tilde{S}}} & D^b(Y_+) \end{array}$$

Proof. This is very similar to the proof of Proposition 2.3. It's enough to check the statement on the line bundles \mathcal{O} and $\mathcal{O}(1)$, and all morphisms between them. It's easy to check that $T_{\tilde{S}}$ sends \mathcal{O} to \mathcal{O} , and $\mathcal{O}(1)$ to the cone on the non-trivial extension

$$\text{Cone}(\mathcal{O}_{\mathbb{P}^1}(-1)[-1] \rightarrow \mathcal{O}(1))$$

which is $\mathcal{O}(-1)$, as required. The argument for the morphisms is similar to our previous one. (The $\mathbb{P}^1 \subset Y_+$ is codimension 1, but all morphisms between these line bundles do extend uniquely.) \square

Remark 2.6. One can also prove the above proposition by considering the Fourier–Mukai kernel for ψ_0 given to us from its geometric description, calculating the derived restriction of this kernel to $Y_+ \times Y_+$, and checking that the result agrees with the kernel for the spherical twist.

We may then deduce the following corollary.

Corollary 2.7. *The square below commutes.*

$$\begin{array}{ccc} \mathrm{D}^b(X_+) & \xrightarrow{T_S} & \mathrm{D}^b(X_+) \\ j_+^* \downarrow & & \downarrow j_+^* \\ \mathrm{D}^b(Y_+) & \xrightarrow{(T_S)^2} & \mathrm{D}^b(Y_+) \end{array}$$

This fact can also be deduced from the results in [HT06]. It corresponds to the fact that a loop around the conifold point in $\mathcal{F}_{\mathcal{X}}$ projects to a *double* loop around the conifold point in $\mathcal{F}_{\mathcal{Y}}$.

Remark 2.8. It's important not to confuse the FI parameter space with the ‘stringy Kähler moduli space’ (SKMS). Under renormalization, a GLSM is believed to flow to a (super)conformal field theory.⁶ The SKMS is a particular slice through the moduli space of this conformal field theory, obtained by only varying the Kähler parameters of the theory. If we assume the Calabi–Yau condition, varying the FI parameter corresponds, after renormalization, to a variation of the Kähler parameters, and so there should be a map

$$\mathcal{F} \longrightarrow \text{SKMS}.$$

For the example \mathcal{Y} , it is believed that this map is an isomorphism, so $\mathcal{F}_{\mathcal{Y}}$ is exactly the Kähler moduli space of the associated CFT.

However, this cannot be true in the 3-fold example \mathcal{X} . The two phases X_+ and X_- are isomorphic and so produce the same CFT, therefore $\mathcal{F}_{\mathcal{X}}$ certainly does not inject into the moduli space of CFTs. For this example, it is believed that the true SKMS is in fact $\mathcal{F}_{\mathcal{Y}}$, so the FI parameter space is actually a double cover of the true SKMS.⁷

It should be possible to produce derived autoequivalences of the phases X_{\pm} by performing parallel transport of B-branes over the SKMS, not just over the FI parameter space, so the autoequivalence group that we are seeing is actually an index 2 subgroup of a larger possible group. Concretely, this means that the autoequivalence T_S of $\mathrm{D}^b(X_+)$ should have a square root. This is indeed true, and the required autoequivalence can be produced by composing F with pullback over the obvious isomorphism between X_+ and X_- . There is also a beautiful construction of Hori [Hor11, Section 2.4] that produces it using a different GLSM (equipped with a superpotential). Note however that although T_S is compactly supported (i.e. it is trivial away from a compact subvariety), the square root of T_S is not.

2.2. Statement of results. In this section we'll explain the class of examples that we're going to consider, and the results that we obtain.

2.2.1. Construction. Fix an integer $k \geq 1$. Consider the quiver obtained by taking the affine Dynkin diagram of type A_k , and replacing each edge with a pair of arrows, one in each direction (see Figure 1). We'll label the clockwise arrows by a_0, \dots, a_k , and the anti-clockwise arrows by b_0, \dots, b_k . Now consider the Artin stack $[V/T]$ parametrizing representations of this quiver which have dimension 1 at each vertex. More explicitly, we let

- $V = \mathbb{C}^{2(k+1)}$ be the vector space (whose dual is) spanned by the arrows, so that it has coordinates $a_0, \dots, a_k, b_0, \dots, b_k$, and
- $T = (\mathbb{C}^*)^{k+1}$ be the torus with one \mathbb{C}^* factor for each vertex.

There's an obvious action of T on V , by letting the \mathbb{C}^* associated to the i^{th} vertex act with

- weight +1 on the two incoming arrows a_{i-1}, b_i , and
- weight -1 on the two outgoing arrows a_i, b_{i-1} ,

reading indices modulo $k+1$. Note that the diagonal \mathbb{C}^* acts trivially, and also that the whole torus acts with trivial determinant on V .

⁶If the associated GIT quotients are non-compact then this statement is problematic, and this is true in the examples that we care about. But we'll skip over this point.

⁷In these two examples the map from \mathcal{F} to the Kähler moduli space is at least a local isomorphism, but if we add superpotentials to our GLSMs then it's easy to find examples where this fails.




FIGURE 1. Quivers associated to affine Dynkin diagrams.

The T -invariant functions on V are generated by the monomials

$$u = a_0 a_1 \dots a_k, \quad v = b_0 b_1 \dots b_k$$

and

$$z_i = a_i b_i, \quad \text{for } i \in [0, k].$$

So the affine quotient V/T is the singularity

$$uv = z_0 z_1 \dots z_k$$

which is the universal unfolding of the A_k surface singularity.

If we choose a character θ of T , we can take a GIT quotient

$$X = V //_{\theta} T.$$

For X to be non-empty, we need to choose a θ that annihilates the diagonal \mathbb{C}^* . Then, for a generic such θ , the quotient X is a smooth toric variety resolving the singularity V/T . It's also Calabi–Yau, since T acts through $\mathrm{SL}(V)$, and as we shall see later it's independent of θ , i.e. all the phases are isomorphic.

This construction is well-known in the context of Nakajima quiver varieties. In that approach, one equips V with a T -invariant symplectic form by making a_i and b_i symplectically dual, then takes a hyperkähler quotient Y . As a complex variety, Y is a subvariety of X formed by taking level sets of all the ‘complex moment maps’, which are the invariant functions

$$\mu_i = z_i - z_{i-1} \tag{2.9}$$

for $1 \leq i \leq k$. Using these functions, we can view X as a family

$$X \xrightarrow{\mu} \mathbb{C}^k.$$

The fibres are all smooth surfaces: these are the (underlying varieties of the) possible hyperkähler quotients Y . The fibre over zero Y_0 is the very famous small resolution of the A_k surface singularity

$$uv = z^{k+1},$$

which appears in the McKay correspondence and Kronheimer’s ALE classification. The larger space X is a versal deformation of Y_0 . Note that it may also be obtained as the inverse image of a Slodowy slice under the Grothendieck–Springer resolution [Slo].

In the case $k = 1$, we get the 3-fold X and the surface $(Y_0 =) Y_+ \subset X$ that we discussed in Section 2.1. For higher k , we are also interested in some intermediate subvarieties lying between Y_0 and X . These intermediate subvarieties are indexed by partitions, as we will now describe.

Firstly, note that for any $i, j \in [0, k]$ the invariant function $z_i - z_j$ lies in the space \mathbb{C}^k spanned by the complex moment maps. Now let Γ be a partition of the set

$$\{z_0, \dots, z_k\}.$$

There is a corresponding subspace of \mathbb{C}^k , where we include the function $z_i - z_j$ if and only if the variables z_i and z_j lie in the same part of Γ . We let

$$X_\Gamma \subset X$$

be the subvariety defined as the vanishing locus of the subspace of the complex moment maps corresponding to a partition Γ .

Partitions form a poset, ordered by refinement, and obviously we have

$$X_\Gamma \subset X_{\Gamma'}$$

if Γ' is a refinement of Γ . We'll let Γ_{fin} denote the finest possible partition, which has $(k + 1)$ parts, then $X_{\Gamma_{\text{fin}}}$ is the ambient space X . At the other extreme, the coarsest possible partition Γ_{crs} , which has only one part, corresponds to the surface $X_{\Gamma_{\text{crs}}} = Y_0$.

2.2.2. Physical heuristics. As we will justify in Section 3, each of these varieties X_Γ is a smooth toric Calabi–Yau, which arises as the GIT quotient of a vector space by a torus. As such, we can view each one as a phase of an abelian GLSM, and so we can run the physicists' recipe and compute the FI parameter space \mathcal{F}_Γ for each one. We will do this in Section 4, and the results are as follows.

For the ambient space $X = X_{\Gamma_{\text{fin}}}$, the FI parameter space is the set

$$\mathcal{F}_{\text{fin}} = \left\{ (\zeta_0 \cdots \zeta_k) \mid \begin{array}{l} \zeta_i \neq 0 \ \forall i \\ \zeta_i \neq \zeta_j \ \forall i \neq j \end{array} \right\} \subset \mathbb{P}^k$$

of $(k + 1)$ -tuples of distinct non-zero complex numbers, up to overall scale. Now take a partition Γ , and let

$$S_\Gamma \subset S_{k+1}$$

be the Young subgroup of permutations that preserve Γ . As we shall show, the FI parameter space associated to the subvariety $X_\Gamma \subset X$ is the orbifold

$$\mathcal{F}_\Gamma = [\mathcal{F}_{\text{fin}} / S_\Gamma]$$

using the obvious action of S_{k+1} on \mathcal{F}_{fin} . So the FI parameter spaces also show the poset structure, since we have a covering map

$$\mathcal{F}_{\Gamma'} \rightarrow \mathcal{F}_\Gamma$$

whenever Γ' is a refinement of Γ .

The FI parameter space associated to the surface Y_0 is $\mathcal{F}_{\text{crs}} = [\mathcal{F}_{\text{fin}} / S_{k+1}]$, which is a \mathbb{C}^* quotient of the configuration space of $k + 1$ points in \mathbb{C}^* . The fundamental group of this space is generated by two subgroups: a lattice, generated by letting each ζ_i loop around zero, and a copy of the braid group B_{k+1} . For our purposes the lattice is not very interesting, and we'll focus on the braid group. The heuristic picture of brane transport over the FI parameter space predicts that there should be an action

$$B_{k+1} \curvearrowright D^b(Y_0).$$

This braid group action was famously constructed by Seidel–Thomas [ST01].

Now choose a general partition Γ . It's clear that (the interesting part of) the fundamental group of \mathcal{F}_Γ is the group B_Γ , defined as the fibre product

$$\begin{array}{ccc} B_\Gamma & \hookrightarrow & B_{k+1} \\ \downarrow & & \downarrow \\ S_\Gamma & \hookrightarrow & S_{k+1} \end{array} \tag{2.10}$$

where S_Γ is the Young subgroup of elements of S_{k+1} preserving the partition Γ . This group B_Γ is sometimes called a ‘mixed braid group’: it consists of braids that permute their endpoints in a way that preserves the partition Γ . The special case $\Gamma = \Gamma_{\text{fin}}$, and so $S_\Gamma = \{1\}$, produces the ‘pure braid group’ P_{k+1} .




FIGURE 2. Elements of the mixed braid group B_Γ for the partition $\Gamma = (01)(2)$.

The FI parameter space heuristics suggest the following result, which we shall prove:

Theorem 2.11 (Corollary 5.4). *There is an action of the mixed braid group B_Γ on the derived category $D^b(X_\Gamma)$.*

In the surface case, Seidel–Thomas proved the rather deep result that the action is faithful. We can leverage their result to prove that all our actions are also faithful.

Note that in the $k = 1$ case (Section 2.1), we did indeed see the above structure, but there wasn't much to prove since both B_2 and P_2 are isomorphic to \mathbb{Z} . The action of \mathbb{Z} on both $D^b(X_+)$ and $D^b(Y_+)$ was generated in each case by a spherical twist. The non-trivial fact was how these two actions were related to each other: we saw (Corollary 2.7) that the action of P_2 on $D^b(X_+)$ intertwined, via restriction, with the action of the subgroup $P_2 \subset B_2$ on $D^b(Y_+)$. Heuristically, this was a reflection of the fact that the FI parameter space for X_+ was a double cover of the FI parameter space for Y_+ .

For general k , similar FI parameter space heuristics predict that our mixed braid group actions should be related to each other, by the same poset structure. Precisely, we should expect that if Γ' is a refinement of Γ , then the action of $B_{\Gamma'}$ on $D^b(X_{\Gamma'})$ will intertwine via the restriction functor

$$D^b(X_{\Gamma'}) \rightarrow D^b(X_\Gamma)$$

with the action of the subgroup $B_{\Gamma'} \subset B_\Gamma$ on $D^b(X_\Gamma)$. In the course of our proof we will show that this prediction does indeed hold (Proposition 5.17).

Remark 2.12. As we explained in Remark 2.8, the FI parameter spaces are not the true Kähler moduli spaces for these theories, indeed the true SKMS appears to be always given by \mathcal{F}_{crs} (for a fixed value of k). On the level of derived categories, this means that our mixed braid group actions could be extended to an action of the full braid group, by including square roots of the relevant spherical twists.

2.3. Future directions. In principle, one can carry out this program for any example of a torus acting on a vector space. The main obstacle appears to be finding a meaningful description of the group $\pi_1(\mathcal{F})$.

A first guess for a good generalization of our examples is to follow the standard technique for Nakajima quiver varieties, and replace the affine A_k Dynkin diagram by some other graph. Unfortunately, as soon as the graph has vertices with valency greater than 2, the obvious subvarieties of the representation space (the analogues of our X_Γ 's) will not be toric, and so we won't get the whole poset of GLSMs and covering maps that we see in our examples.

In fact this covering map phenomenon, and its relationship with derived categories, is probably the most interesting feature of our construction. It would be worthwhile to investigate the general conditions under which it arises.

3. TORIC CALCULATIONS

In this section we'll apply some completely standard toric techniques to understand our varieties X_Γ .

3.1. Representations of the free quiver. Our first task is to find the toric fan for the ambient variety X . As an initial step, we can package our construction as an exact sequence of lattices as follows:

$$\begin{array}{ccccccc}
 \mathbb{Z} & \xrightarrow{(1 \ \cdots \ 1)} & \mathbb{Z}^{k+1} & \xrightarrow{Q} & \mathbb{Z}^{2(k+1)} & \xrightarrow{P} & \mathbb{Z}^{2+(k+1)} \xrightarrow{\begin{pmatrix} +1 \\ +1 \\ -1 \\ \vdots \\ -1 \end{pmatrix}} \mathbb{Z} \\
 & & \parallel & & \updownarrow & & \updownarrow \\
 & & \text{Hom}(\mathbb{C}^*, T) & & \text{basis vectors } b_i, a_i & & \text{co-ordinates } u, v, z_i
 \end{array}$$

$$Q = \begin{pmatrix}
 & b_0 & a_0 & b_1 & a_1 & b_2 & a_2 & \cdots & b_{k-1} & a_{k-1} & b_k & a_k & \\
 -1 & +1 & +1 & -1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 & 1 \\
 0 & 0 & -1 & +1 & +1 & -1 & & 0 & 0 & 0 & 0 & 0 & 2 \\
 \vdots & & & & & & & & & & & \vdots & \vdots \\
 0 & 0 & 0 & 0 & 0 & 0 & & -1 & +1 & +1 & -1 & -1 & k \\
 +1 & -1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & -1 & +1 & +1 & 0
 \end{pmatrix}$$

$$P = \begin{pmatrix}
 & u & v & z_0 & z_1 & z_2 & \cdots & \\
 0 & 1 & 1 & 0 & 0 & \cdots & & b_0 \\
 1 & 0 & 1 & 0 & 0 & & & a_0 \\
 0 & 1 & 0 & 1 & 0 & & & b_1 \\
 1 & 0 & 0 & 1 & 0 & & & a_1 \\
 0 & 1 & 0 & 0 & 1 & & & b_2 \\
 1 & 0 & 0 & 0 & 1 & & & a_2 \\
 \vdots & & & & & & & \vdots
 \end{pmatrix}$$

 FIGURE 3. Toric data for GIT quotient X .

The matrices for the maps Q and P are given in Figure 3. The \mathbb{Z} at the left of the sequence corresponds to the diagonal $\mathbb{C}^* \subset T$, which acts trivially on V . The toric fan for X lies in the rank $k + 2$ lattice $\text{Im } P \subset \mathbb{Z}^{2+(k+1)}$. However, we will often view it as a fan in $\mathbb{Z}^{2+(k+1)}$, since the larger lattice has a convenient system of co-ordinates.

For any choice of θ we know the rays in the toric fan immediately: they're generated by the images of the standard generators for $\mathbb{Z}^{2+(k+1)}$ (i.e. the rows of P). Consider the generators corresponding to the arrows a_0, \dots, a_k , and look at their images under P . This gives a set of vectors $\alpha_0, \dots, \alpha_k$ (i.e. the a -indexed rows of P) which form the vertices of a standard k -simplex in the affine subspace $u = 1, v = 0$. The other half of the generators give a set of vectors β_0, \dots, β_k (i.e. the b -indexed rows of P) which span a standard k -simplex in $u = 0, v = 1$. So all the generators together span a polytope

$$\Pi \cong \Delta^k \times I.$$

Remark 3.1. Notice that Π lies in an affine hyperplane of height 1, namely $\{z_0 + \dots + z_k = 1\}$: this is equivalent to the Calabi–Yau condition.

Now choose a character $\theta = (\theta_0, \dots, \theta_k) \in (\mathbb{Z}^{k+1})^\vee$ of the torus T which annihilates the diagonal \mathbb{C}^* , and lift it to an element $\vartheta \in (\mathbb{Z}^{2+(k+1)})^\vee$. We can choose ϑ to be of the form

$$\vartheta = (0, \vartheta_0, 0, \vartheta_1, 0, \vartheta_2, \dots, 0, \vartheta_k). \quad (3.2)$$

Then

$$\theta_0 = \vartheta_k - \vartheta_0, \quad \theta_1 = \vartheta_0 - \vartheta_1, \quad \dots, \quad \theta_k = \vartheta_{k-1} - \vartheta_k,$$

and so these values ϑ_i are unique up to adding an overall constant. Choose a character such that

$$\vartheta_0 > \vartheta_1 > \dots > \vartheta_k \quad (3.3)$$

i.e. all the θ_i are positive for $i \geq 1$. We'll refer to the corresponding quotient X as the *standard phase*.

The toric fan for X is the set of cones on some subdivision of the polytope Π . To find out which subdivision it is, we use a standard shortcut from toric geometry [Ful93, Section 3.4]. View ϑ as an integer-valued function on the vertices of Π . After subdividing Π according to the toric fan for X , this function ϑ will extend to a strictly concave piecewise linear (PL) function over Π .

This shortcut lets us guess the answer. There is a standard way to triangulate $\Pi = \Delta^k \times I$ into $(k+1)$ -simplices, by cutting it into the pieces

$$\Delta_i = [\alpha_0, \dots, \alpha_i, \beta_i, \dots, \beta_k], \quad (3.4)$$

where $0 \leq i \leq k$. Our standard ϑ gives a strictly concave PL function on this subdivision, and this is the only subdivision for which this is true, therefore this subdivision gives the correct fan for the standard phase. In summary:

Lemma 3.5. *The toric fan for the standard phase X is given by all the cones on the simplices Δ_i .*

Remark 3.6. For future reference in Section 3.2, note that the cone on the polytope Π (i.e. the union of all the cones in the fan) is just the intersection of the positive orthant

$$\{u, v, z_0, \dots, z_k \geq 0\} \subset \mathbb{Z}^{2+(k+1)}$$

with the hyperplane

$$u + v = z_0 + \dots + z_k \quad (3.7)$$

spanned by the lattice $\text{Im } P$. To get the cone on the simplex $\Delta_i \subset \Pi$, we additionally impose the inequalities

$$u \geq z_0 + \dots + z_{i-1}, \quad v \geq z_{i+1} + \dots + z_k. \quad (3.8)$$

If we vary the character ϑ , then the above argument shows that the GIT quotient does not change as long as we remain in the region (3.3), and conversely as soon as two of the ϑ_i become equal then we hit a wall. Consequently the chambers for the GIT problem are the regions

$$\vartheta_{\sigma(0)} > \vartheta_{\sigma(1)} > \dots > \vartheta_{\sigma(k)}$$

for some permutation $\sigma \in S_{k+1}$. We'll denote the corresponding quotients by X^σ , but we'll continue to let $X = X^{(1)}$ denote the standard phase.

To get the fan for a non-standard phase X^σ , we just re-order the vertices of Π by σ before performing the standard triangulation. Notice that all the phases are in fact isomorphic, because of this S_{k+1} symmetry (i.e. the Weyl group). This is not immediately obvious from the original quiver description.

Example 3.9. (1) Let $k = 1$. The polytope Π lies in the affine subspace

$$\{u + v = z_0 + z_1, \quad z_0 + z_1 = 1\} \subset \mathbb{C}^4.$$

We can use (z_1, u) as co-ordinates on this subspace, then we can draw the triangulation of Π (see Figure 4a).

(2) For $k = 2$, the polytope Π lies in the affine subspace

$$\{u + v = z_0 + z_1 + z_2, \quad z_0 + z_1 + z_2 = 1\} \subset \mathbb{C}^5.$$

See Figure 4b for the triangulation of Π .

3.2. Subvarieties associated to partitions. In this section, we'll investigate the subvarieties $X_\Gamma \subset X$ cut out by subsets of the complex moment map equations, as defined in Section 2.2.1.

To start with, suppose X_Γ is the divisor $\{\mu_i = 0\}$. The function μ_i is not a character of the torus which acts on X , so X_Γ is not a torus-invariant divisor. Nevertheless, X_Γ is a toric variety. To see this, we first restrict our attention to the Zariski open torus $(\mathbb{C}^*)^{k+2} \subset X$. Within this locus, the set $\{\mu_i = 0\}$ is just a corank 1 subtorus. So X_Γ is the closure of this subtorus, and is invariant under the action of the subtorus on X , so X_Γ is a toric variety. Furthermore, to get the toric fan for X_Γ we take the toric fan for X and slice it along the corresponding corank 1 sublattice, namely $\{z_i = z_{i-1}\} \subset \mathbb{Z}^{2+(k+1)}$.




FIGURE 4. Triangulations of the polytope Π whose cones give the fan for the standard phase X .

Similarly, if $X_\Gamma \subset X$ is a subvariety of higher codimension (associated to a partition with fewer parts), then X_Γ is also a toric variety, and we can obtain its toric fan by slicing along the appropriate sublattice. We'll now perform this procedure explicitly, and work out the toric fans for all the X_Γ .

3.2.1. *Toric data.* Fix a partition Γ , and encode it as a surjective function

$$\gamma : [0, k] \rightarrow [0, s]$$

whose level sets are the pieces of the partition. The fan for X_Γ lies in the sublattice

$$\mathbb{Z}^{2+(s+1)} \subset \mathbb{Z}^{2+(k+1)} \quad (3.10)$$

cut out by the equations $z_i = z_j$, for all i, j such that $\gamma(i) = \gamma(j)$.

Notation 3.11. We choose co-ordinates


$$(\tilde{u}, \tilde{v}, \tilde{z}_0, \dots, \tilde{z}_s)$$

on the sublattice $\mathbb{Z}^{2+(s+1)}$ of (3.10), induced from co-ordinates on the bigger lattice $\mathbb{Z}^{2+(k+1)}$ in the natural way: we set $\tilde{u} = u$, $\tilde{v} = v$, and for each $t \in [0, s]$ we have a co-ordinate \tilde{z}_t , which is the restriction of z_i for all $i \in \gamma^{-1}(t)$. We write

$$N_t := \#(\gamma^{-1}(t))$$

for the sizes of the pieces of the partition.

Example 3.12. Let $k = 4$, and let $\Gamma = (01)(234)$. Then $s = 1$, and the corresponding function γ and the associated co-ordinates \tilde{z}_i are as follows:



Recall from Remark 3.6 that the union of all the cones in the fan for X is the intersection of the positive orthant in $\mathbb{Z}^{2+(k+1)}$ with the hyperplane (3.7). Consequently, the union of all the cones in the fan for X_Γ must be the intersection of the positive orthant

$$\{\tilde{u}, \tilde{v}, \tilde{z}_0, \dots, \tilde{z}_s \geq 0\} \subset \mathbb{Z}^{2+(s+1)}$$

with the hyperplane

$$\tilde{u} + \tilde{v} = \sum_{t=0}^s N_t \tilde{z}_t. \quad (3.13)$$

The lattice points that lie in this locus are easy to determine, they form a cone generated by the set of points

$$\tilde{\alpha}_t(\delta) := (\delta, N_t - \delta, 0, \dots, 0, \underset{\uparrow t}{1}, 0, \dots, 0) \quad (3.14)$$

where $t \in [0, s]$, and $\delta \in [0, N_t]$. The convex hull of all these generating points $\tilde{\alpha}_t(\delta)$ is a polytope

$$\tilde{\Pi} \cong \Delta^s \times I \quad (3.15)$$

which again is isomorphic to a prism on a simplex. The points at the vertices of $\tilde{\Pi}$ are the two s -simplices

$$[\tilde{\alpha}_0(0), \tilde{\alpha}_1(0), \dots, \tilde{\alpha}_s(0)] \quad \text{and} \quad [\tilde{\alpha}_0(N_1), \tilde{\alpha}_1(N_2), \dots, \tilde{\alpha}_s(N_s)]$$

which are embedded in a non-standard way. All the other points $\tilde{\alpha}_t(\delta)$ lie along edges of $\tilde{\Pi}$.

We've determined that the union of all the cones in the fan for X_Γ is the cone on this polytope $\tilde{\Pi}$. To determine the individual cones, we recall that the polytope Π for X is subdivided into $(k+1)$ -simplices $\Delta_0, \dots, \Delta_k$, and the cones on these simplices define the toric fan for X . To get the cones in the fan for X_Γ , we take the cone on each simplex Δ_i and slice it along the hyperplane (3.13).

The cone on Δ_i is cut out of the positive orthant in $\mathbb{Z}^{2+(k+1)}$ by the inequalities (3.8), so after slicing we obtain the subset of the positive orthant in $\mathbb{Z}^{2+(s+1)}$ cut out by the inequalities

$$\tilde{u} \geq \sum_{t=0}^s n_t^i \tilde{z}_t, \quad \tilde{v} \geq \sum_{t=0}^s m_t^i \tilde{z}_t,$$

where the integers n_t^i and m_t^i are defined as follows:

$$\begin{aligned} n_t^i &= \#(\gamma^{-1}(t) \cap [0, i-1]), & m_t^i &= \#(\gamma^{-1}(t) \cap [i+1, k]) \\ & & &= \begin{cases} N_t - n_t^i, & \text{for } t \neq \gamma(i) \\ N_{\gamma(i)} - n_{\gamma(i)}^i - 1, & \text{for } t = \gamma(i) \end{cases} \end{aligned}$$

The lattice points that lie in this region necessarily form a cone generated by some subset of the lattice points in $\tilde{\Pi}$. This subset consists of

$$\tilde{\alpha}_0(n_0^i), \tilde{\alpha}_1(n_1^i), \dots, \tilde{\alpha}_s(n_s^i), \text{ and } \tilde{\alpha}_{\gamma(i)}(n_{\gamma(i)}^i + 1). \quad (3.16)$$

These points span an $(s+1)$ -simplex, which we'll denote by $\tilde{\Delta}_i$. Together these simplices form a triangulation of the polytope $\tilde{\Pi}$. In summary:

Lemma 3.17. *The toric fan for X_Γ consists of the cones on the simplices $\tilde{\Delta}_i$. In particular, the complete set of points $\tilde{\alpha}_t(\delta)$, as defined by (3.14), generate all the rays.*

Remark 3.18. The whole of $\tilde{\Pi}$ lies in the affine hypersurface

$$\{\tilde{z}_0 + \dots + \tilde{z}_s = 1\},$$

and hence X_Γ is Calabi–Yau.

Remark 3.19. This triangulation of $\tilde{\Pi} \cong \Delta^s \times I$ is easy to visualize: start with the base s -simplex

$$[\tilde{\alpha}_0(0), \dots, \tilde{\alpha}_s(0)],$$

and turn it into an $(s+1)$ -simplex by adjoining the point $\tilde{\alpha}_{\gamma(0)}(1)$. Now take the ‘top’ face of this $(s+1)$ -simplex, namely

$$[\tilde{\alpha}_0(0), \dots, \tilde{\alpha}_{\gamma(0)}(1), \dots, \tilde{\alpha}_s(0)],$$


and extend it to a new $(s+1)$ -simplex by adjoining the point $\tilde{\alpha}_{\gamma(1)}(n_{\gamma(1)}^2)$. This new point will either be $\tilde{\alpha}_{\gamma(1)}(2)$, if $\gamma(0) = \gamma(1)$, or $\tilde{\alpha}_{\gamma(1)}(1)$, if $\gamma(0) \neq \gamma(1)$. We continue in this way, increasing the ‘heights’ of the vertices according to the values of γ , until we reach all the maximum heights N_0, \dots, N_s and have triangulated the whole of $\tilde{\Pi}$.

Example 3.20. We return to the very simple example of 3.9, the conifold. The subvariety corresponding to the partition $\Gamma = (01)$ is the A_1 surface of Section 2.1.2, and the fan is the cone on the polytope in Figure 5a.

Example 3.21. Let $k = 4$, so that X is a 6-fold, and let Γ be the partition as in Example 3.12. Then $s = 1$, so X_Γ is a 3-fold. The polytope $\tilde{\Pi}$ lies in the affine subspace

$$\{\tilde{u} + \tilde{v} = 2\tilde{z}_0 + 3\tilde{z}_1, \quad \tilde{z}_0 + \tilde{z}_1 = 1\} \subset \mathbb{C}^4.$$

We can use (\tilde{z}_1, \tilde{u}) as co-ordinates on this subspace, then we can draw the triangulation of $\tilde{\Pi}$ (see Figure 5b).



(a) Triangulation for $k = 1$, $\Gamma = (01)$.
The associated X_Γ is the A_1 surface.

(b) Triangulation for $k = 4$, $\Gamma = (01)(234)$.
In this case, X_Γ is a 3-fold.

FIGURE 5. Triangulations of polytope $\tilde{\Pi}$ associated with subvarieties X_Γ .

3.2.2. GIT description. We've constructed the toric data for X_Γ , however we claim that X_Γ is in fact a GIT quotient of a vector space by a torus, which for us is a more useful description. (Of course this is not surprising, as it's true of 'most' toric varieties.) We'll first argue this abstractly, then we'll explicitly identify the GIT data.

If we forget the individual cones in the fan (i.e. the triangulation of $\tilde{\Pi}$), then we have only the data of the set of $(k + s + 2)$ vectors $\tilde{\alpha}_t(\delta)$ that generate the rays. Let $\{e_t(\delta)\}$ be an abstract set bijecting with the set of vectors $\{\tilde{\alpha}_t(\delta)\}$. Then we have a map between lattices

$$\begin{aligned} \tilde{P} : \mathbb{Z}^{k+s+2} := \langle e_t(\delta) \rangle_{\mathbb{Z}} &\rightarrow \mathbb{Z}^{2+(s+1)}, \\ e_t(\delta) &\mapsto \tilde{\alpha}_t(\delta). \end{aligned}$$

The kernel of this map is a rank k lattice, and the associated torus $\tilde{T} = (\mathbb{C}^*)^k$ acts on the vector space \tilde{V} generated by the $e_t(\delta)$. This is a GIT problem, and the possible phases $\tilde{V} // \tilde{T}$ correspond to particular subdivisions of the polytope $\tilde{\Pi}$.

Proposition 3.22. X_Γ is one of the phases of the GIT problem $\tilde{V} // \tilde{T}$.

Proof. Our choice of character θ specifies a line bundle L on X , making X a projective-over-affine variety. By definition, X_Γ is also projective-over-affine: it is Proj of the graded ring of sections of powers of the line bundle $\tilde{L} = L|_{X_\Gamma}$. The line bundle \tilde{L} is necessarily toric, so it corresponds to some character $\tilde{\theta}$ of the torus $(\mathbb{C}^*)^k$, i.e. a stability condition for our new GIT problem. The corresponding GIT quotient is by definition X_Γ . \square

We'll refer to X_Γ as the 'standard phase' for this GIT problem. There are some other obvious phases: if we start with a non-standard phase X^σ for the ambient space (associated to some $\sigma \in S_{k+1}$), and then impose the moment map equations corresponding to Γ , we get another toric variety X_Γ^σ which is birational to X_Γ . It's clear how to get the toric fan for X_Γ^σ : we just apply the permutation σ before running the recipe from the previous section. Consequently X_Γ^σ is constructed by the same GIT problem that constructs X_Γ , so it's another phase.

From the recipe for the toric fans, it's clear that if σ fixes the partition γ then X_Γ and X_Γ^σ are identical (i.e. isomorphic over the affine base). Therefore these phases depend only on the coset σS_Γ , where $S_\Gamma \subset S_{k+1}$ is the symmetry group of Γ .

As we will see later, these are not the only phases of this GIT problem: there are phases which are orbifolds, and so cannot arise in this way.

Next we'll explicitly identify the kernel of \tilde{P} , i.e. the action of $\tilde{T} = (\mathbb{C}^*)^k$ on \tilde{V} . Note that for any t , and any $\delta \in [0, N_t - 1]$, we have

$$e_t(\delta + 1) - e_t(\delta) \xrightarrow{\tilde{P}} (-1, 1, 0, 0, \dots, 0),$$

so if we pick any two of these elements then their difference

$$e_t(\delta + 1) - e_t(\delta) - e_{t'}(\delta' + 1) + e_{t'}(\delta')$$

lies in the kernel of \tilde{P} , and it's evident that the whole kernel is spanned by vectors of this form.

We can write down a convenient basis for $\ker \tilde{P}$ by considering pairs of adjacent cones in the toric data. For example, consider the two simplices $\tilde{\Delta}_0$ and $\tilde{\Delta}_1$. There are two situations to consider:

- If $\gamma(0) \neq \gamma(1)$, then the four vectors

$$\tilde{\alpha}_{\gamma(0)}(1), \tilde{\alpha}_{\gamma(0)}(0), \tilde{\alpha}_{\gamma(1)}(1), \tilde{\alpha}_{\gamma(1)}(0) \in \tilde{\Delta}_0 \cup \tilde{\Delta}_1$$

are coplanar, and we let

$$\tau_1 = e_{\gamma(0)}(1) - e_{\gamma(0)}(0) - e_{\gamma(1)}(1) + e_{\gamma(1)}(0) \in \mathbb{Z}^{k+s+2}$$

which lies in the kernel of \tilde{P} .

- Alternatively, if $\gamma(0) = \gamma(1)$ then the three vectors

$$\tilde{\alpha}_{\gamma(0)}(2), \tilde{\alpha}_{\gamma(0)}(1), \tilde{\alpha}_{\gamma(0)}(0) \in \tilde{\Delta}_0 \cup \tilde{\Delta}_1$$

are collinear, and we let

$$\tau_1 = 2e_{\gamma(0)}(1) - e_{\gamma(0)}(2) - e_{\gamma(0)}(0) \in \mathbb{Z}^{k+s+2},$$

which also lies in the kernel of \tilde{P} .

In general, we have the following proposition.

Proposition 3.23. *For $i \in [1, k]$, two adjacent simplices $\tilde{\Delta}_{i-1}$ and $\tilde{\Delta}_i$ give a vector*

$$\tau_i = e_{\gamma(i-1)}(n_{\gamma(i-1)}^{i-1} + 1) - e_{\gamma(i-1)}(n_{\gamma(i-1)}^{i-1}) - e_{\gamma(i)}(n_{\gamma(i)}^i + 1) + e_{\gamma(i)}(n_{\gamma(i)}^i)$$

in the kernel of \tilde{P} , and this set of vectors forms a basis.

Example 3.24. For example 3.20, the A_1 surface, the above Proposition 3.23 gives a single generator

$$\tau_1 = 2e_0(1) - e_0(0) - e_0(2)$$

for the kernel. This tells us that X_Γ is a GIT quotient of \mathbb{C}^3 by \mathbb{C}^* acting with weights $(-1, 2, -1)$, which of course we knew.

Example 3.25. Returning to the setting of Example 3.21, we find that the two simplices $\tilde{\Delta}_1$ and $\tilde{\Delta}_2$ yield coplanar vectors as in case (1) above, and the other pairs of adjacent simplices give collinear vectors as in case (2). So this 3-fold X_Γ is a GIT quotient of \mathbb{C}^7 by the torus $(\mathbb{C}^*)^4$ acting with weights as follows:

$$\begin{pmatrix} -1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 2 & -1 \end{pmatrix}$$

As we explained in the proof of Proposition 3.22, our choice of character θ when constructing the ambient space X induces a character $\tilde{\theta}$ of \tilde{T} . We'll now identify this character explicitly. This will tell

us what stability conditions we can choose on our GIT problem $\tilde{V} // \tilde{T}$ to produce the standard phase X_Γ .

Recall that our standard ϑ extends to a concave PL function on the toric fan for X . We just restrict this function to the fan for X_Γ , then its values on the vertices $\tilde{\alpha}_t(\delta)$ give a lift $\tilde{\vartheta}$ of the character $\tilde{\theta}$. The vertex $\tilde{\alpha}_t(\delta)$ can be written as a sum

$$\tilde{\alpha}_t(\delta) = \sum_{i \in D} \alpha_i + \sum_{j \in \gamma^{-1}(t) - D} \beta_j$$

for any subset $D \subset \gamma^{-1}(t)$ of size δ (this is immediate from the definition (3.14)). Because of our choice (3.2) of ϑ we have that $\tilde{\vartheta}(\beta_i) = 0$, so the fact that the function ϑ is concave implies that

$$\tilde{\vartheta} : \tilde{\alpha}_t(\delta) \mapsto \max_{\substack{D \subset \gamma^{-1}(t) \\ |D| = \delta}} \sum_{d \in D} \vartheta_d, \quad (3.26)$$

i.e. the sum of the δ biggest values lying in the t^{th} part of the partition. If we restrict this function $\tilde{\vartheta}$ to the kernel of \tilde{P} , this will give us our character $\tilde{\theta}$ of the torus $(\mathbb{C}^*)^k$. To see what the result is, note first that

$$\tilde{\vartheta}(\tilde{\alpha}_t(\delta + 1) - \tilde{\alpha}_t(\delta)) = \vartheta_d$$

where ϑ_d is the $(\delta + 1)^{\text{th}}$ biggest value lying in the t^{th} part of the partition. Consequently

$$\tilde{\vartheta}(\tilde{\alpha}_{\gamma(i)}(n_{\gamma(i)}^i + 1) - \tilde{\alpha}_{\gamma(i)}(n_{\gamma(i)}^i)) = \vartheta_i$$

and therefore

$$\tilde{\theta}(\tau_i) = \vartheta_{i-1} - \vartheta_i = \theta_i. \quad (3.27)$$

So in this basis, setting all coefficients of the character $\tilde{\theta}$ to be positive produces the standard phase X_Γ . In fact we will see in the next section that this is precisely the chamber of characters that produce this phase, so we'll refer to it as the *standard chamber* for this GIT problem.

3.3. Families of rational curves. In this section we'll make some very straightforward observations on the geometry of X and its subvarieties X_Γ .

Recall that the fan for X is the union of the cones on the simplices Δ_i (3.4). Each such cone corresponds to a Zariski open set $U_i \subset X$, which is isomorphic to \mathbb{C}^{k+2} with co-ordinates $a_0, \dots, a_i, b_i, \dots, b_k$. We can also think of U_i as (the image of) the affine subspace in V where we set all the remaining co-ordinates to 1.

If we take two neighbouring simplices Δ_{i-1} and Δ_i , and consider the corresponding open sets U_{i-1} and U_i , we find that they glue together to give an open set

$$\mathcal{N}_i := U_{i-1} \cup U_i \cong \mathcal{O}(-1)_{\mathbb{P}^1}^{\oplus 2} \times \mathbb{C}^{k-1}.$$

The ratio $a_i : b_{i-1}$ gives co-ordinates on the \mathbb{P}^1 , the co-ordinates a_{i-1} and b_i are the fibre directions on the bundle, and the remaining co-ordinates $a_0, \dots, a_{i-2}, b_{i+1}, \dots, b_k$ parametrize the \mathbb{C}^{k-1} . We will be interested in the codimension 2 subvarieties

$$S_i = \{a_{i-1} = b_i = 0\} \subset X \quad (3.28)$$


corresponding to the zero section of the bundle, where here $1 \leq i \leq k$. Each S_i lies entirely within the open set \mathcal{N}_i , so it's a trivial family of rational curves, and it has this very simple Zariski neighbourhood.

Remark 3.29. If we vary our GIT quotient from the standard chamber (3.3) to the neighbouring chamber

$$\vartheta_0 > \dots > \vartheta_{i-2} > \vartheta_i > \vartheta_{i-1} > \vartheta_{i+1} > \dots > \vartheta_k \quad (3.30)$$

by passing through the wall $\theta_i = 0$, then S_i is exactly the locus that becomes unstable. To see this, observe that a cone disappears from the fan if and only if it contains the cone on the interval $[\alpha_{i-1}, \beta_i]$, which is the cone corresponding to the toric subvariety S_i . See Figure 6a for an example.

3.3.1. Subvarieties X_Γ . Now we pick a partition Γ with s pieces. We want to find the intersection of these S_i with the subvariety X_Γ , which is cut out by some subset of the equations $z_j - z_l = 0$. Within

FIGURE 6. Flop geometry for $k = 2$, for 4-fold X .

\mathcal{N}_i , the global functions z_j become

$$z_j = \begin{cases} a_j & j \leq i - 2, \\ a_j b_j & i \leq j \leq i + 1, \\ b_j & j \geq i + 1. \end{cases}$$

There are two cases to distinguish:

- (a) If i and $i - 1$ are not in the same piece of Γ , so we don't impose $z_{i-1} = z_i$, then we have

$$\tilde{\mathcal{N}}_i := \mathcal{N}_i \cap X_\Gamma \cong \mathcal{O}(-1)_{\mathbb{P}^1}^{\oplus 2} \times \mathbb{C}^{s-1}.$$

The subvarieties S_i and X_Γ intersect transversely, and the subvariety

$$\tilde{S}_i := S_i \cap X_\Gamma \tag{3.31}$$

is an $(s - 1)$ -parameter trivial family of rational $(-1, -1)$ -curves.

- (b) If i and $i - 1$ are in the same piece of Γ then

$$\tilde{\mathcal{N}}_i = \mathcal{N}_i \cap X_\Gamma \cong \mathcal{O}(-2)_{\mathbb{P}^1} \times \mathbb{C}^s.$$

In this case the intersection of S_i and X_Γ is not transverse, and \tilde{S}_i is an s -parameter trivial family of rational (-2) -curves.

In each case, \tilde{S}_i has this very simple Zariski neighbourhood $\tilde{\mathcal{N}}_i$.




FIGURE 7. Flop loci for $k = 2$, partition $\Gamma = (01)(2)$. In this case $s = 1$, and X_Γ is a 3-fold. The locus S_1 is a 1-parameter families of (-2) -curves, and S_2 is a single $(-1, -1)$ -curve. Compare Figure 6b.

We can also find these \tilde{S}_i in the toric data for X_Γ , or equivalently in terms of the GIT problem $\tilde{V} // \tilde{T}$. Recall that the toric fan for X_Γ is the union of the cones on the simplices $\tilde{\Delta}_i$ (3.16). Each such cone corresponds to a ‘toric chart’ $\tilde{U}_i \subset X_\Gamma$, isomorphic to affine space, and $\tilde{\mathcal{N}}_i = \tilde{U}_{i-1} \cup \tilde{U}_i$, so $\tilde{\mathcal{N}}_i$ is a

toric variety whose fan has only two cones. If we inspect this toric data, we see that it's constructing $\tilde{\mathcal{N}}_i$ as a GIT quotient

$$\tilde{\mathcal{N}}_i = \tilde{V}_i //_{\tilde{\theta}} \tau_i$$

where $\tilde{V}_i \subset \tilde{V}$ is the subspace spanned by the vertices that appear in $\tilde{\Delta}_{i-1} \cup \tilde{\Delta}_i$, and $\tau_i \subset \tilde{T}$ is the 1-parameter subgroup identified in Proposition 3.23. The torus action is trivial in some directions (these contribute the trivial directions in $\tilde{\mathcal{N}}_i$), and in the remaining directions it's the usual GIT problem from Section 2.1 constructing either (a) $\mathcal{O}(-1)_{\mathbb{P}^1}^{\oplus 2}$ or (b) $\mathcal{O}(-2)_{\mathbb{P}^1}$.

By (3.27), the value of a standard stability condition $\tilde{\theta}$ on the 1-parameter subgroup τ_i is $(\vartheta_{i-1} - \vartheta_i)$. Consequently, when we leave the standard chamber across the wall corresponding to $\vartheta_i = \vartheta_{i-1}$ the family of curves \tilde{S}_i becomes unstable. What happens when we cross the wall differs in the two cases:

- (a) Across the wall, there is a neighbouring phase for X_Γ where the locus \tilde{S}_i has been flopped: locally around \tilde{S}_i , this is just a trivial family of standard 3-fold flops. This neighbouring phase is a subvariety inside a non-standard phase of the ambient space X : we move to the chamber (3.30) and then impose moment maps corresponding to the same Γ . In other words, this flop of the locus \tilde{S}_i is just induced by a flop of the locus S_i inside X .
- (b) Across the wall there is a neighbouring phase for X_Γ where the locus \tilde{S}_i has been removed, and replaced with a trivial family of orbifold points with \mathbb{Z}_2 isotropy groups. This neighbouring phase is an orbifold, and so cannot be a subvariety inside any phase of X . The characters $\tilde{\theta}$ in this neighbouring chamber do not arise from any θ .

If we travel further away from the standard chamber then we can reach other kinds of phases, which can have larger isotropy groups. Note also that the above description still works if we start in any phase X_Γ^σ induced from a non-standard phase X^σ of the ambient space: we just have to permute the variables.

3.4. Derived equivalences from wall-crossing. In this section we'll discuss the derived equivalences that correspond to the wall-crossing described in the previous section, following the general theory of [HL12, BFK12] and especially [HLS].⁸

Start with a standard phase X_Γ of one our subvarieties, pick a value of i , and let X'_Γ be the phase obtained by crossing the wall corresponding to $\vartheta_i = \vartheta_{i-1}$. As we saw in Section 3.3, the birational transformation between X_Γ and X'_Γ is locally just a trivial family of the two kinds of A_1 flops discussed in Section 2.1. Consequently, there are \mathbb{Z} -many derived equivalences

$$\psi_k : D^b(X_\Gamma) \xrightarrow{\sim} D^b(X'_\Gamma)$$

which roughly-speaking come from family versions of the derived equivalences we saw in the A_1 case. Let's explain this more precisely.

3.4.1. General theory. Let Y be a smooth projective-over-affine variety equipped with a \mathbb{C}^* action, and let

$$\sigma : Z \hookrightarrow Y$$

be the inclusion of the fixed locus. For simplicity, assume that Z is connected. To form a GIT quotient, we need to choose an equivariant line bundle on Y , so let's choose $\mathcal{O}_Y(1)$ (i.e. the trivial line bundle equipped with a weight 1 action). Then the unstable locus is the subvariety $S_+ \subset Y$ of points that flow to Z under the \mathbb{C}^* action (as $t \in \mathbb{C}^*$ goes to 0), so the semistable locus is $Y^{ss} := Y - S_+$, and the GIT quotient is $[Y^{ss}/\mathbb{C}^*]$.⁹

Now let η_+ be the \mathbb{C}^* -weight of the line bundle $\det N_{S_+}^\vee Y$ along Z , which is necessarily a positive integer. For any choice of integer k , we define a 'window'

$$\mathcal{W}(k) \subset D^b([Y/\mathbb{C}^*])$$

⁸The general theory deals with GIT quotients under arbitrary reductive groups, since we only care about the abelian case we'll only be using a small part of it.

⁹Purists may insist that the GIT quotient is really the coarse moduli space of this stack, but this abuse-of-language is becoming quite common.

to be the full subcategory of objects \mathcal{E} such that all homology sheaves of $L\sigma^*\mathcal{E}$ have \mathbb{C}^* -weights lying in the interval

$$[k, k + \eta_+).$$

So for example, the equivariant line bundle $\mathcal{O}_Y(i)$ lies in $\mathcal{W}(k)$ if and only if i lies in the above interval. In fact if Y is a vector space then we may equivalently define $\mathcal{W}(k)$ to be the subcategory generated by this set of line bundles, as we did in Section 2.1. A basic result of the theory cited above is that the restriction map

$$\mathcal{W}(k) \rightarrow \mathrm{D}^b([Y^{ss}/\mathbb{C}^*])$$

is an equivalence (for any k). If we change our stability condition to $\mathcal{O}_Y(-1)$ then the unstable locus changes to the subvariety S_- of points that flow away from Z (i.e. they flow to Z as $t \in \mathbb{C}^*$ goes to infinity), and we have a similar description of the derived category of the other GIT quotient $[(Y - S_-)/\mathbb{C}^*]$ based on the weight η_- of $(\det N_{S_-}Y)|_Z$. In particular if we happen to have $\eta_- = \eta_+$ then the two GIT quotients are derived equivalent, and each choice of window gives a specific derived equivalence ψ_k .

3.4.2. Our examples. Now let's apply this theory to the wall-crossing between the phases X_Γ and X'_Γ . We saw in Section 3.2.2 that both are GIT quotients of a vector space \tilde{V} by a torus \tilde{T} . The wall between the phases is part of the annihilator of the 1-parameter subgroup $\tau_i \subset \tilde{T}$. Choose a stability condition that lies exactly on the wall (but not also on any other walls), and let \tilde{V}^{ss} be the semistable locus for this stability condition. If we pick a splitting $\tilde{T} = \tau_i \oplus \tau_i^\perp$, then the torus τ_i^\perp acts freely on \tilde{V}^{ss} , the quotient $Y = \tilde{V}^{ss}/\tau_i^\perp$ is a smooth projective-over-affine variety, and we have an isomorphism of stacks

$$\left[\tilde{V}^{ss} / \tilde{T} \right] = \left[Y / \tau_i \right].$$

Our phases X_Γ and X'_Γ are the two possible GIT quotients of Y by τ_i . By definition the subvariety $S_- \subset Y$ is the closure of the locus that becomes unstable when we cross the wall, so we know from Remark 3.29 that S_- is (the closure in Y of) \tilde{S}_i , our trivial family of rational curves (3.28). Similarly S_+ is (the closure of) the subset that replaces \tilde{S}_i after the flop: by Section 3.3 it's either (a) another family of rational curves, or (b) a family of orbifold points.

Now we explain the calculation of the numerical data η_\pm in our examples. We can calculate in the Zariski open neighbourhood $[\tilde{V}_i/\tau_i]$, i.e. the GIT problem that constructs the neighbourhood \tilde{N}_i . The action of τ_i is trivial in most directions, so we can reduce to one of our two basic A_1 examples from Section 2.1. Then it's easy to calculate¹⁰ that

$$\eta_- = \eta_+ = 2$$

and so the general theory tells us that we have derived equivalences ψ_k between the two phases, defined using windows

$$\mathcal{W}(k) \subset \mathrm{D}^b\left(\left[\tilde{V}^{ss}/\tilde{T}\right]\right).$$

3.4.3. Geometric description of equivalences. The family version of Proposition 2.3 holds true [HLS, Section 3.1], so the equivalence

$$\psi_{-1} : \mathrm{D}^b(X_\Gamma) \rightarrow \mathrm{D}^b(X'_\Gamma)$$

can also be described using the common birational roof.

Similarly, the autoequivalence

$$(\psi_{-1})^{-1} \circ \psi_0 : \mathrm{D}^b(X_\Gamma) \rightarrow \mathrm{D}^b(X_\Gamma)$$

¹⁰In fact the equality of η_- and η_+ , though not their value, follows from the fact that both sides are Calabi–Yau.

is a family version of the spherical twist around $\mathcal{O}_{\mathbb{P}^1}(-1)$ that we discussed in Section 2.1. More precisely, let

$$\begin{array}{ccc} \tilde{S}_i & \xrightarrow{\iota} & X_\Gamma \\ \mu \downarrow & & \\ \mathbb{C}^r & & \end{array}$$

be the trivial family of rational curves (here r is either $(s-2)$ or $(s-1)$ depending on which case we're in), and let $\mathcal{O}(-1)$ be the relative tautological line bundle for μ . Then we define a functor \mathcal{S} with right adjoint \mathcal{R} , as follows:

$$\begin{array}{ccc} & \mathcal{S} & \\ \text{D}^b(\mathbb{C}^r) & \xrightarrow{\quad} & \text{D}^b(X_\Gamma^\sigma) \\ & \mathcal{R} & \end{array}$$

$$\begin{aligned} \mathcal{S}(\mathcal{E}) &= \iota_*(\mu^*\mathcal{E} \otimes \mathcal{O}(-1)) \\ \mathcal{R}(\mathcal{F}) &= \mu_*(\iota^!\mathcal{F} \otimes \mathcal{O}(1)) \end{aligned}$$

The functor \mathcal{S} is a spherical functor in the sense of [Ann07, AL13], and we can define an associated spherical twist functor

$$T_{\mathcal{S}} : \mathcal{E} \mapsto \text{Cone}(\mathcal{S}\mathcal{R}\mathcal{E} \rightarrow \mathcal{E}). \quad (3.32)$$

Recall that \tilde{S}_i has a nice Zariski neighbourhood $\tilde{\mathcal{N}}_i$ which is a trivial family of copies of either (a) $\mathcal{O}(-1, -1)$ or (b) $\mathcal{O}(-2)$ over \mathbb{P}^1 . Within this neighbourhood $T_{\mathcal{S}}$ is just a trivial family of ordinary spherical twists around $\mathcal{O}_{\mathbb{P}^1}(-1)$, and outside $\tilde{\mathcal{N}}_i$ it's just the identity functor. By [HLS, Section 3.2], we have that

$$(\psi_{-1})^{-1} \circ \psi_0 = T_{\mathcal{S}}.$$

Again this discussion holds equally well if we start our wall-crossing from any phase X_Γ^σ .

4. FAYET–ILIOPOULOS PARAMETER SPACES

4.1. Toric Mirror Symmetry heuristics. Choose an action of a rank r torus T on an n -dimensional vector space V . After picking diagonal co-ordinates, the action is given by the matrix of weights

$$Q : \mathbb{Z}^r \rightarrow \mathbb{Z}^n.$$

We let

$$P : \mathbb{Z}^n \rightarrow \mathbb{Z}^{n-r}$$

be the cokernel of Q (modulo torsion). Notice that \mathbb{Z}^n has a canonical positive orthant (since it comes from a vector space), so it has a canonical set of generators, up to permutation. If we were to pick a character θ of T and construct the corresponding toric variety $X = V //_{\theta} T$, then the toric fan for X lies in the lattice \mathbb{Z}^{n-r} , and the rays in the fan are (some subset of) the images of the canonical generators under P .

As we discussed in the introduction, we can also view this as the input data for a gauged linear sigma model, with gauge group T . Associated to the model there is a complex orbifold \mathcal{F} , which is supposed to be the space of possible values for the complexified Fayet–Iliopoulos parameter that occurs in the Lagrangian of the model. Fortunately we don't need to understand what this means, since toric mirror symmetry provides a precise heuristic recipe for constructing \mathcal{F} (see for example [CCIT09]). We'll now describe this recipe.

Let Π be the polytope in \mathbb{Z}^{n-r} spanned by the images of the canonical generators for \mathbb{Z}^n . For simplicity, we assume that all these images are distinct, and that Π contains only these n lattice points and no others. Now pick $(n-r)$ formal variables x_1, \dots, x_{n-r} , so each point $p = (p_1, \dots, p_{n-r}) \in \mathbb{Z}^{n-r}$ corresponds to a Laurent monomial $x_1^{p_1} \dots x_{n-r}^{p_{n-r}}$, which we denote by x^p . We consider the set of Laurent

polynomials

$$L = \left\{ \sum_{p \in \Pi} c_p x^p \mid c_p \in \mathbb{C} \right\}$$

spanned by Laurent monomials corresponding to the vertices of Π . Then L is naturally the dual vector space to V , and it carries an action of $(\mathbb{C}^*)^{n-r}$ by rescaling each variable x_1, \dots, x_{n-r} . The weight matrix for this $(\mathbb{C}^*)^{n-r}$ -action is exactly P^T , so if we try and quotient L by this action, then we are exactly looking at the dual GIT problem.

The FI parameter space \mathcal{F} is defined to be the quotient of a certain open set in L by the torus $(\mathbb{C}^*)^{n-r}$. To obtain this open set, we remove two kinds of points from L , as follows:

- (1) The hyperplanes $\{c_p = 0\}$, for every point $p \in \mathbb{Z}^{n-r}$ that corresponds to a vertex of Π . In other words, we insist that our Laurent polynomials have Newton polytope exactly Π .

We'll call the complement of this set of hyperplanes L' . Every point in L' has finite stabilizer, so the space

$$\overline{\mathcal{F}} := [L' / (\mathbb{C}^*)^{n-r}]$$

is an orbifold. It's actually a (non-compact) toric orbifold, via the residual action of the torus T^\vee . The space \mathcal{F} is an open set in $\overline{\mathcal{F}}$, which we obtain by removing a second kind of point:

- (2) The 'discriminant' locus of 'non-generic' polynomials. Genericity here means the following: for every face F of Π (of any dimension) that doesn't contain the origin, consider the 'restricted' Laurent polynomial obtained by setting to zero each coefficient that doesn't correspond to a vertex of F . We require that every such restricted Laurent polynomial has no critical points in the torus $(\mathbb{C}^*)^{n-r}$.

The discriminant locus is the degeneracy locus of the associated GKZ system of differential equations [Ado94]. The discriminant locus is not usually invariant under the action of T^\vee , so there is no torus action on \mathcal{F} .

The 1-parameter subgroups in T^\vee correspond to characters θ for our original GIT problem. They're divided into chambers corresponding to the different possible GIT quotients X , these are the chambers of the secondary fan. For each possible GIT quotient we have a corresponding 'large-radius limit' in \mathcal{F} , which is roughly the limit of a generic point in $\overline{\mathcal{F}}$ under the action of any of the 1-parameter subgroups in the corresponding chamber. If this limit exists then it will be a torus fixed point in $\overline{\mathcal{F}}$, if not then we think of it as a region lying at infinity.

Remark 4.1. Although this recipe appears to correctly produce the FI parameter space, it is not enough to accurately produce the mirror family. By construction, \mathcal{F} is a moduli space of Laurent polynomials, and Hori–Vafa [HV00] argued that the mirror family is the associated family of Landau–Ginzburg models. This works in some compact examples, but is known to be only an approximation when the phases are non-compact, as happens in our examples. The actual construction of the mirrors to our examples is rather subtle, see [CPU13].

4.2. Results for our examples. Now we run the recipe of the previous section on the toric varieties that we're interested in. Fix a partition Γ encoded by

$$\gamma : [0, k] \rightarrow [0, s],$$

and let X_Γ be the corresponding toric variety. In Section 3.2 we computed the full toric fan for X_Γ , in particular we described (Lemma 3.17) the set of generators for all the rays in the fan. This is the same as giving the matrix P , so it's enough information to run the recipe and compute the corresponding FI parameter space \mathcal{F}_Γ .

Recall that the rays in the fan for X_Γ are generated by the vectors

$$\tilde{\alpha}_t(\delta) \in \mathbb{Z}^{2+(s+1)}$$

defined in (3.14). Here $t \in [0, s]$ and $\delta \in [0, N_t]$, where $N_t = \#\gamma^{-1}(t)$. The corresponding Laurent monomials are

$$\tilde{u}^\delta \tilde{v}^{N_t - \delta} \tilde{z}_t,$$

so the space L consists of Laurent polynomials of the form

$$f_0(\tilde{u}, \tilde{v})\tilde{z}_0 + \dots + f_s(\tilde{u}, \tilde{v})\tilde{z}_s$$

where each f_t is a homogeneous polynomial of degree N_t . We have a $(\mathbb{C}^*)^{2+(s+1)}$ action on L , but there is a global \mathbb{C}^* stabilizer which is an artifact of our decision to draw the toric fan in a lattice whose rank was 1 larger than necessary. We need to remember to neglect this global stabilizer when forming the quotient.

Next we have to identify which loci we should delete from L . Step 1 is easy, we just require that the first and last coefficients of each f_t cannot go to zero, i.e. each polynomial $f_t(\tilde{u}, 1)$ cuts out a length N_t subscheme in the punctured line $\mathbb{C}_{\tilde{u}}^*$.

For step 2, recall (3.15) that the vectors $\tilde{\alpha}_t(\delta)$ span a polytope $\tilde{\Pi}$, which is abstractly isomorphic to a prism $\Delta^s \times I$. Hence there are two kinds of faces of $\tilde{\Pi}$. For any subset $J \subset [0, s]$, there are subsimplices $\Delta_J \subset \Delta^s$ at either end of the prism, and there is also the prism $\Delta_J \times I \subset \Delta^s \times I$.

- For the first kind of face, the corresponding restricted Laurent polynomial is of the form

$$\sum_{t \in J} \lambda_t \tilde{u}^{N_t} \tilde{z}_t \quad \text{or} \quad \sum_{t \in J} \mu_t \tilde{v}^{N_t} \tilde{z}_t,$$

and these never have critical points on the torus.

- For the second kind of face, the corresponding restricted Laurent polynomial is

$$\sum_{t \in J} f_t(\tilde{u}, \tilde{v}) \tilde{z}_t.$$

Now we have a non-trivial condition, because this Laurent polynomial has a critical point in the torus if and only if the corresponding complete intersection

$$\{f_t(\tilde{u}, 1) = 0, \forall t \in J\} \subset \mathbb{C}_{\tilde{u}}^*$$

is non-generic. When $\#J = 1$, this gives us the condition that no single f_t can have a repeated root. When $\#J \geq 2$, we get the condition that no two of the f_t can have a shared root.

In summary, we have the following proposition.

Proposition 4.2.

$$\mathcal{F}_\Gamma = \{(f_0, \dots, f_s)\} / (\mathbb{C}^*)^{s+2}$$

where each f_t is a homogeneous polynomial of degree N_t in \tilde{u} and \tilde{v} such that

- no f_t has roots at $\tilde{u} = 0$ or $\tilde{v} = 0$,
- no f_t has repeated roots, and
- no two of the f_t share a root.

The torus acts by rescaling \tilde{u} and \tilde{v} , and each f_t .

Now consider the special case when we choose the finest possible partition Γ_{fin} , so that each $N_t = 1$. In this case we can replace each linear function f_t with its root $\zeta_t \in \mathbb{C}_{\tilde{u}}^*$, so we have (using the abbreviation $\mathcal{F}_{\text{fin}} := \mathcal{F}_{\Gamma_{\text{fin}}}$)

$$\mathcal{F}_{\text{fin}} = \{(\zeta_0 : \dots : \zeta_k)\} \subset \mathbb{P}^k$$

where the ζ_t are all distinct and non-zero.

At the opposite extreme, suppose we choose the coarsest possible partition Γ_{crs} . This has only one piece, so $s = 0$ and $N_0 = k + 1$. In this case, we can replace the single f_0 with its set of roots $\{\zeta_0, \dots, \zeta_k\} \subset \mathbb{C}_{\tilde{u}}^*$. Writing $\mathcal{F}_{\text{crs}} := \mathcal{F}_{\Gamma_{\text{crs}}}$, we then have

$$\mathcal{F}_{\text{crs}} = \left[\{(\zeta_0 : \dots : \zeta_k)\} / S_{k+1} \right] \subset \left[\mathbb{P}^k / S_{k+1} \right]$$

where again the ζ_t are all distinct and non-zero. So \mathcal{F}_{crs} is a quotient of \mathcal{F}_{fin} by the action of the symmetric group.

Finally, suppose that Γ is some intermediate partition. We can identify each f_t with its set of roots, and so get a set of distinct non-zero roots $(\zeta_0, \dots, \zeta_k)$. However, this set of roots is partitioned by Γ ,

and we allow relabelings that preserve the partition. In other words, we have

$$\mathcal{F}_\Gamma = [\mathcal{F}_{\text{fin}} / S_\Gamma]$$

where S_Γ is the Young subgroup $S_\Gamma = S_{N_0} \times \dots \times S_{N_s} \subset S_{k+1}$ that fixes the partition Γ .

4.3. Large-radius limits. We now identify some of the large-radius limits in the spaces \mathcal{F}_Γ . Pick a 1-parameter subgroup of T^\vee corresponding to a character $\theta \in (\mathbb{Z}^{k+1})^\vee$ with a lift $\vartheta \in (\mathbb{Z}^{2(k+1)})^\vee$ lying in the standard chamber (3.3). This acts on the space \mathcal{F}_{fin} by

$$\theta(\lambda) : \zeta_i \mapsto \lambda^{\vartheta_i} \zeta_i,$$

so the corresponding ‘region at infinity’ in \mathcal{F}_{fin} is the locus where

$$\log |\zeta_0| \gg \log |\zeta_1| \gg \dots \gg \log |\zeta_k|. \quad (4.3)$$

The other large-radius limits in \mathcal{F}_{fin} are of the same form, but with the order of the roots permuted by some $\sigma \in S_{k+1}$.

Next we look at \mathcal{F}_{crs} , where we have a single polynomial f_0 . A 1-parameter subgroup $\tilde{\theta}$ in the standard chamber is induced from a standard θ as explained in Section 3.2.2, and it acts on the polynomial f_0 by rescaling the coefficient of $\tilde{u}^{k+1-\delta} \tilde{v}^\delta \tilde{z}_0$ with weight

$$\tilde{\vartheta}(\tilde{\alpha}_0(\delta)) = \vartheta_0 + \vartheta_1 + \dots + \vartheta_{\delta-1}$$

(see (3.26)). Therefore it acts on the roots of f_0 by rescaling them with weights $\vartheta_0, \dots, \vartheta_k$, and so the standard LR limit in \mathcal{F}_{crs} is again the region where

$$\log |\zeta_0| \gg \log |\zeta_1| \gg \dots \gg \log |\zeta_k|$$

(here the labelling of the roots is arbitrary, so we can choose to label them according to size). So the standard LR limit in \mathcal{F}_{crs} is the common image of all the LR limits in \mathcal{F}_{fin} , exactly as our heuristic picture suggests.

Now leave the standard chamber by crossing the wall $\tilde{\theta}(\tau_i) = \vartheta_{i-1} - \vartheta_i = 0$, i.e. violate the inequality

$$2\tilde{\vartheta}(\tilde{\alpha}_0(i)) - \tilde{\vartheta}(\tilde{\alpha}_0(i+1)) - \tilde{\vartheta}(\tilde{\alpha}_0(i-1)) > 0$$

while preserving all the other such inequalities. The corresponding LR limit in \mathcal{F}_{crs} is the region where

$$\log |\zeta_0| \gg \dots \gg \log |\zeta_{i-1}| \approx \log |\zeta_i| \gg \dots \gg \log |\zeta_k|$$

and

$$\log |\zeta_{i-1} + \zeta_i| \ll \log |\zeta_i|.$$

These are the LR limits which are ‘adjacent’ to the standard one, if we go further into the parameter space (i.e. cross further walls in the secondary fan) we reach other limits where more roots become commensurable.

Now let Γ be an arbitrary partition, and \mathcal{F}_Γ the corresponding parameter space. From the discussion above, the standard LR limit in \mathcal{F}_Γ is the image of the standard LR limit in \mathcal{F}_{fin} (and all LR limits obtained from the standard one under the action of S_Γ), as the heuristic picture suggests. When we leave the standard chamber by crossing the wall $\tilde{\theta}(\tau_i) = 0$, there are two possibilities:

- (a) If $i-1$ and i are in distinct pieces of Γ , then we move to a LR limit in \mathcal{F}_Γ which is the image of a LR limit in \mathcal{F}_{fin} . It’s obtained from the standard limit by transposing $i-1$ and i .
- (b) If $i-1$ and i are in the same piece of Γ , then we move to a LR limit in \mathcal{F}_Γ where the two roots ζ_{i-1} and ζ_i of the polynomial $f_{\gamma(i)}$ have comparable sizes.

The phase corresponding to this new LR limit was discussed in Section 3.3.

4.4. Fundamental groups. In this section we make some simple observations on the fundamental groups of the FI parameter spaces computed in the previous section. These spaces are orbifolds, and the symbol π_1 will always denote the orbifold fundamental group.

Consider the FI parameter space \mathcal{F}_{crs} associated to the surface $X_{\Gamma_{\text{crs}}}$. As was noted in [CCIT09], it has fundamental group

$$\pi_1(\mathcal{F}_{\text{crs}}) = \tilde{B}_{k+1} \rtimes C_{k+1},$$

where \tilde{B}_{k+1} is the affine braid group associated to the affine Dynkin diagram \tilde{A}_k , and C_{k+1} is a cyclic group. We should explain this briefly: by rescaling we can insist that $\zeta_0 \zeta_1 \dots \zeta_k = 1$, and this leaves a residual action of the cyclic group C_{k+1} . So if

$$\mathcal{G} = \{ \zeta_0, \dots, \zeta_k \mid \prod_i \zeta_i = 1 \}$$

with all the ζ_i distinct, then $\mathcal{F}_{\text{crs}} = [\mathcal{G} / S_{k+1} \times C_{k+1}]$. By taking logs of all the ζ_i we get a principle \mathbb{Z}^k -bundle over \mathcal{G} given by

$$\bar{\mathcal{G}} = \{ w_0, \dots, w_k \mid \sum_i w_i = 0, w_i - w_j \notin \mathbb{Z} \text{ for } i \neq j \},$$

and this is the complement of the affine complex hyperplane arrangement associated to \tilde{A}_k . When we quotient by the affine Coxeter group $\mathbb{Z}^k \rtimes S_{k+1}$, we get

$$\pi_1(\mathcal{G} / S_{k+1}) = \tilde{B}_{k+1},$$

and so $\pi_1(\mathcal{F}_{\text{crs}})$ is $\tilde{B}_{k+1} \rtimes C_{k+1}$, as claimed.

We want to take a slightly different point of view on this group. Choose a point in \mathcal{F}_{crs} which is the orbit of a tuple $(\zeta_0, \zeta_1, \dots, \zeta_k)$ with

$$\log |\zeta_0| \gg \dots \gg \log |\zeta_k|.$$

We can think of such a point as lying ‘near the large-radius limit’ for the standard phase. If we stay in this region, we can only see a subgroup of $\pi_1(\mathcal{F}_{\text{crs}})$ which is a lattice $\mathbb{Z}^k = \mathbb{Z}^{k+1} / \mathbb{Z}$ coming from rotating the phases of the ζ_i .

Alternatively, we could leave this large-radius limit, but insist that we stay in (the orbit of) the region where $\text{Re}(\zeta_i) < 0$ for all i . The subgroup of $\pi_1(\mathcal{F}_{\text{crs}})$ contributed by this region is the ordinary braid group B_{k+1} . These two subgroups generate $\pi_1(\mathcal{F}_{\text{crs}})$, indeed we can view it as a semidirect product

$$\pi_1(\mathcal{F}_{\text{crs}}) = \langle \mathbb{Z}^k \rangle \rtimes B_{k+1},$$

where $\langle \mathbb{Z}^k \rangle$ is the normal closure of \mathbb{Z}^k .

From this, it’s easy to deduce $\pi_1(\mathcal{F}_\Gamma)$ for any other partition Γ . The inclusion of \mathcal{F}_{crs} into $[\mathbb{P}^k / S_{k+1}]$ gives a map from $\pi_1(\mathcal{F}_{\text{crs}})$ to S_{k+1} , and $\pi_1(\mathcal{F}_\Gamma)$ is the fibre product of $\pi_1(\mathcal{F}_{\text{crs}})$ with the Young subgroup S_Γ . So it’s generated by a lattice \mathbb{Z}^k , and the mixed braid group B_Γ , which by definition (2.10) is the fibre product

$$B_\Gamma := B_{k+1} \times_{S_{k+1}} S_\Gamma.$$

5. MIXED BRAID GROUP ACTIONS

For each Γ , we wish to produce an action of $\pi_1(\mathcal{F}_\Gamma)$ on the derived category $\text{D}^b(X_\Gamma)$. The FI parameter space heuristics give a precise prediction for what this action should be, as we now explain.

5.1. Heuristics for the generators. Pick a partition Γ , and pick a base point in \mathcal{F}_Γ which is close to the large-radius limit for the standard phase, i.e. it is the orbit of a tuple $(\zeta_0, \dots, \zeta_k)$ where

$$\log |\zeta_0| \gg \dots \gg \log |\zeta_k|.$$

Recall from Section 4.4 that $\pi_1(\mathcal{F}_\Gamma)$ is generated by two subgroups: a lattice \mathbb{Z}^k arising from rotating the phases of the ζ_i , and the mixed braid group B_Γ . The action of the lattice on $\text{D}^b(X_\Gamma)$ is obvious, we can canonically identify this lattice with the set of toric line bundles on X_Γ , and these act on the derived category by the tensor product. Consequently we will ignore the lattice in all of the subsequent discussion, and focus on the mixed braid group B_Γ . This means we only consider the region in \mathcal{F}_Γ which is the orbit of the set

$$\{\text{Re}(\zeta_i) < 0, \forall i\}.$$

Let’s denote this region by $\mathcal{F}_\Gamma^{<0}$.

If we let Γ vary over all partitions (of size $k + 1$), then the spaces $\mathcal{F}_\Gamma^{<0}$ form a poset of covering spaces, with $\mathcal{F}_{\text{fin}}^{<0}$ at the top and $\mathcal{F}_{\text{crs}}^{<0}$ at the bottom. Let's apply some completely elementary algebraic topology to this situation. Pick a base point $b \in \mathcal{F}_{\text{crs}}$ which is close to the standard LR limit. For any Γ , the preimage of b in $\mathcal{F}_\Gamma^{<0}$ is some finite set of points, which we can identify with the set of cosets

$$\pi_1(\mathcal{F}_{\text{crs}}^{<0}) / \pi_1(\mathcal{F}_\Gamma^{<0}) = B_{k+1} / B_\Gamma = S_{k+1} / S_\Gamma$$

by matching the identity coset with the point near the standard LR limit in $\mathcal{F}_\Gamma^{<0}$.

Let \mathcal{G}_Γ be the groupoid of homotopy classes of paths in $\mathcal{F}_\Gamma^{<0}$ that start and end somewhere in this set of points. This is the same thing as the action groupoid for the action of B_{k+1} on this set of cosets. The isotropy group in \mathcal{G}_Γ is of course B_Γ . If we let Γ vary we get a set of groupoids, and the covering maps induce functors between them which are faithful, and also surjective on the total sets of arrows.

Every point σS_Γ in \mathcal{G}_Γ lies near a LR limit in \mathcal{F}_Γ , and there is a corresponding phase X_Γ^σ of the GIT problem. If we move along a path in the groupoid to another point $\sigma' S_\Gamma$, the FI parameter space heuristics predict that we get an associated derived equivalence

$$D^b(X_\Gamma^\sigma) \xrightarrow{\sim} D^b(X_\Gamma^{\sigma'}).$$

Putting these together, we expect to get a functor

$$T : \mathcal{G}_\Gamma \rightarrow \mathbf{Cat}_1$$

sending σS_Γ to $D^b(X_\Gamma^\sigma)$. Here \mathbf{Cat}_1 denotes the category of categories with morphisms being functors up to natural isomorphism.¹¹

The FI parameter space heuristics can also tell us exactly what T should do to arrows. To explain this, we need to identify some generators for \mathcal{G}_Γ .

Firstly set $\Gamma = \Gamma_{\text{fin}}$. The resulting groupoid \mathcal{G}_{fin} is very easy to describe: its arrows are ordinary braid diagrams but with the strands labelled from 0 to k in some order. Pick a point $\sigma \in S_{k+1}$ in this groupoid, and suppose that i and j are adjacent labels at this point (i.e. i and j are adjacent after applying σ to $[0, k]$), with i to the left of j . Let $t_{i,j}^\sigma$ be the braid which crosses these two adjacent strands (left-over-right, as in Figure 2) and leaves all the other strands alone. There is a corresponding path in \mathcal{F}_{fin} , starting in the LR limit

$$\log |\zeta_{\sigma(0)}| \gg \dots \gg \log |\zeta_i| \gg \log |\zeta_j| \gg \dots \gg \log |\zeta_{\sigma(k)}| \quad (5.1)$$

and ending in the LR limit

$$\log |\zeta_{\sigma(0)}| \gg \dots \gg \log |\zeta_j| \gg \log |\zeta_i| \gg \dots \gg \log |\zeta_{\sigma(k)}|$$

where only ζ_i and ζ_j change in value along the path, their norms change monotonically, and we have $\arg(\zeta_i) > \arg(\zeta_j)$ at the point where their norms are equal.

The set of all these braids (or paths) $t_{i,j}^\sigma$, for all σ , evidently generates the groupoid \mathcal{G}_{fin} . Therefore for any partition Γ , their images under the map

$$\mathcal{G}_{\text{fin}} \rightarrow \mathcal{G}_\Gamma$$

generate \mathcal{G}_Γ , since the map is surjective on the total set of arrows. We'll denote the image of $t_{i,j}^\sigma$ by $t_{i,j}^{\sigma S_\Gamma}$, since it only depends on the coset σS_Γ . When $\Gamma = \Gamma_{\text{crs}}$ we just have the standard generators for $\mathcal{G}_{\text{crs}} = B_{k+1}$, indeed the set $t_{i,j}^{\sigma S_\Gamma}$ are just the preimages of these standard generators.

Now fix an arbitrary Γ again, and let

$$t_{i,j}^{\sigma S_\Gamma} : \sigma S_\Gamma \rightarrow \sigma' S_\Gamma$$

be one of these generating arrows, where σ' is the composition $\sigma \cdot (ij)$. Let's describe the corresponding path in $\mathcal{F}_\Gamma^{<0}$, and deduce the derived equivalence that it should correspond to. There are two cases to consider:

- (a) If i and j don't lie in the same part of Γ then this path moves us from one LR limit to a neighbouring LR limit, and the two phases X_Γ^σ and $X_\Gamma^{\sigma'}$ are related by a flop of a trivial family

¹¹This is the natural 1-category associated to the 2-category of categories.

of rational $(-1, -1)$ -curves, as described in Section 3.3. The path $t_{i,j}^{\sigma S_\Gamma}$ should correspond to the derived equivalence

$$T_{i,j}^{\sigma S_\Gamma} := \psi_0 : D^b(X_\Gamma^\sigma) \xrightarrow{\sim} D^b(X_\Gamma^{\sigma'})$$

discussed in Section 3.4, which as we saw can be defined either using VGIT and ‘windows’, or geometrically using the birational roof.

- (b) If i and j lie in the same part of Γ then σS_Γ and $\sigma' S_\Gamma$ are the same coset, and the path $t_{i,j}^{\sigma S_\Gamma}$ is actually a loop. We start at the LR limit (5.1), move into the neighbouring LR limit where $|\zeta_i| \approx |\zeta_j|$, and return, having looped the discriminant locus $\zeta_i = \zeta_j$. The phase X_Γ^σ contains a trivial family of rational (-2) -curves. When we pass to the neighbouring LR limit this family gets ‘flopped’ and becomes a family of \mathbb{Z}_2 -orbifold points. Consequently this loop in $\mathcal{F}_\Gamma^{<0}$ should correspond to the autoequivalence

$$T_{i,j}^{\sigma S_\Gamma} := (\psi_{-1})^{-1} \circ \psi_0 : D^b(X_\Gamma^\sigma) \xrightarrow{\sim} D^b(X_\Gamma^\sigma)$$

discussed in Section 3.4, which is equal to a family spherical twist around the family of rational curves.

If we let $\mathcal{G}_\Gamma^{\text{free}}$ be the free groupoid generated by the arrows $t_{i,j}^{\sigma S_\Gamma}$, then it’s tautological that this assignment

$$T_\Gamma : t_{i,j}^{\sigma S_\Gamma} \rightarrow T_{i,j}^{\sigma S_\Gamma}$$

defines a functor from $\mathcal{G}_\Gamma^{\text{free}}$ to \mathbf{Cat}_1 . In agreement with the FI parameter space heuristics, we will show that in fact we have:

Theorem 5.2. *T_Γ gives a well-defined functor*

$$T_\Gamma : \mathcal{G}_\Gamma \rightarrow \mathbf{Cat}_1.$$

In the process we will show the following stronger result, for which we don’t have a heuristic justification:

Theorem 5.3. *T_Γ is faithful.*

A priori this second theorem is much more difficult. However for the special case when $\Gamma = \Gamma_{\text{cfs}}$, and so X_Γ is a surface, the faithfulness was proved by Seidel and Thomas [ST01], and we can deduce the result for general Γ fairly easily.

Restricting T_Γ to the isotropy group of the standard phase we deduce:

Corollary 5.4. *T_Γ defines a faithful action of the mixed braid group B_Γ on the derived category $D^b(X_\Gamma)$.*

The proofs of these two theorems will be presented in the next two sections.

5.2. Braid relations on the ambient space. In this section we’ll prove the following special case of Theorem 5.2:

Proposition 5.5. *When $\Gamma = \Gamma_{\text{fin}}$, so each phase X_Γ^σ is the versal deformation space of the resolution of the A_k singularity, then $T_{\text{fin}} := T_{\Gamma_{\text{fin}}}$ gives a well-defined functor*

$$T_{\text{fin}} : \mathcal{G}_{\text{fin}} \rightarrow \mathbf{Cat}_1.$$

The key case to consider is when $k = 2$. (When $k = 1$, there’s nothing to prove.) In this case there are six phases X^σ , indexed by elements σ of S_3 . We can denote the functors between them by $T_{i,j}^\sigma$, dropping the trivial group $S_{\Gamma_{\text{fin}}}$ from the notation. Also, we’ll write elements of S_3 as orderings of the set $\{0, 1, 2\}$ rather than as products of cycles, so for example the standard phase will be denoted X^{012} .

A prototypical braid relation in \mathcal{G}_{fin} is drawn in Figure 8, and expressed in the following diagram of functors:

Proposition 5.6. *The hexagon below commutes.*

$$\begin{array}{ccccc}
 & & \mathrm{D}^b(X^{102}) & \xrightarrow{T_{0,2}^{102}} & \mathrm{D}^b(X^{120}) \\
 & \nearrow^{T_{0,1}^{012}} & & & \searrow^{T_{1,2}^{120}} \\
 \mathrm{D}^b(X^{012}) & & & & \mathrm{D}^b(X^{210}) \\
 & \searrow_{T_{1,2}^{012}} & & & \nearrow_{T_{0,1}^{201}} \\
 & & \mathrm{D}^b(X^{021}) & \xrightarrow{T_{0,2}^{021}} & \mathrm{D}^b(X^{201})
 \end{array} \tag{5.7}$$




FIGURE 8. Prototypical relation in $\mathcal{G}_{\mathrm{fin}}$ for $k = 2$.

This is a special case of Proposition 5.5, but it quickly implies the whole proposition:

Proof of Proposition 5.5. Let $k > 2$. We can draw a similar hexagon moving between the standard phase $X^{0123\dots k}$ on the left, and the phase $X^{2103\dots k}$ on the right, and the commutativity of this hexagon is one of the braid relations. If we start with the fan for the standard phase, as described in Section 3, then the changes to the fan that take place when we move around the hexagon will only affect the first three cones: the remaining $\Delta_3, \dots, \Delta_k$ stay unchanged, and so the birational modifications are concentrated inside the open set $\{a_3 \neq 0, \dots, a_k \neq 0\}$. Inside this open set, the geometry we're studying is just a trivial family (over $\mathbb{C}_{b_3, \dots, b_k}^{k-2}$) of copies of the geometry appearing in the $k = 2$ case. Therefore Proposition 5.6 implies that this hexagon also commutes.

We've shown that one particular braid relation holds for each $k \geq 2$, so all the others hold by the S_{k+1} symmetry. \square

The rest of this section will be devoted to the proof of Proposition 5.6, our prototype braid relation.

5.2.1. *Formal structure of the proof.* To see why the hexagon (5.7) commutes, we need to go back to the fundamental definition of our functors $T_{i,j}^\sigma$ from Section 3.4.

Let $V \cong \mathbb{C}^6$ and $T \cong (\mathbb{C}^*)^3$ be the data of the GIT problem that we're considering, and let $T' = T/\mathbb{C}^*$ be the quotient of T by the trivially-acting diagonal subgroup. Let \mathcal{X} denote the Artin stack

$$\mathcal{X} = [V/T'].$$

Each of the six phases X^σ of the GIT problem is an open substack in \mathcal{X} . Let's identify these open substacks completely explicitly.

Let τ_0, τ_1 and τ_2 be the three 1-parameter subgroups of T corresponding to the nodes of the quiver, so in the quotient torus T' we have $\tau_0 + \tau_1 + \tau_2 = 0$. These three are the only 1-parameter subgroups which fix more than the origin in V , and their fixed loci are as follows:

$$\begin{aligned}
 \tau_0 = (1, 0, 0) & \text{ fixes } \bar{Z}_0 = \{a_0 = b_0 = a_2 = b_2 = 0\}, \\
 \tau_1 = (0, 1, 0) & \text{ fixes } \bar{Z}_1 = \{a_0 = b_0 = a_1 = b_1 = 0\}, \\
 \tau_2 = (0, 0, 1) & \text{ fixes } \bar{Z}_2 = \{a_1 = b_1 = a_2 = b_2 = 0\}.
 \end{aligned}$$

If we pick a stability condition, then each τ_i destabilizes either the attracting or the repelling subspace of the corresponding \bar{Z}_i , and the union of these three subspaces is precisely the unstable locus for that stability condition. For example, in the standard phase X^{012} the unstable locus is the union of the three subspaces

$$\bar{S}_0 = \{b_0 = a_2 = 0\}, \quad \bar{S}_1 = \{b_0 = a_1 = 0\}, \quad \bar{S}_2 = \{b_1 = a_2 = 0\}.$$

(see also Figure 9). As another example, if we move to the adjacent phase X^{021} then only \bar{S}_2 changes: it gets replaced by the subspace $\bar{S}'_2 = \{a_1 = b_2 = 0\}$.




FIGURE 9. For the standard phase X^{012} , unstable subspaces \bar{S}_i flowing under 1-parameter subgroups τ_i to fixed loci \bar{Z}_i : each locus is defined by setting the dotted arrows to zero.

Now pick a (non-zero) stability condition that lies on the wall between the two phases X^{012} and X^{021} , so the corresponding unstable locus is $\bar{S}_0 \cup \bar{S}_1$. The partial quotient

$$Y_{1,2}^{012} := (V - \bar{S}_0 - \bar{S}_1) / \tau_0$$

is smooth, and carries a residual action of τ_2 . The stack

$$\mathcal{Y}_{1,2}^{012} = [(V - \bar{S}_0 - \bar{S}_1) / T'] = [Y_{1,2}^{012} / \tau_2]$$

is an open substack of \mathcal{X} , and it contains both X^{012} and X^{021} as open substacks. We can view this as a GIT problem, and it describes this particular wall-crossing.

Recall that we defined windows

$$\mathcal{W}_{1,2}^{012}(k) \subset \mathrm{D}^b(\mathcal{Y}_{1,2}^{012})$$

as the full subcategories

$$\mathcal{W}_{1,2}^{012}(k) = \left\langle \mathcal{E} \mid \begin{array}{l} \text{all homology sheaves of } L\sigma^*\mathcal{E} \text{ have } \tau_2\text{-weights} \\ \text{lying in the interval } [k, k+2] \end{array} \right\rangle$$

where $\sigma : \bar{Z}_2 \hookrightarrow Y_{1,2}^{012}$ is the inclusion of the fixed locus (recall also that we've calculated that the numerical invariant η is 2). Each of these windows is equivalent, under the restriction functor, to the derived categories of both X^{012} and X^{021} , and the wall-crossing functor $T_{1,2}^{012}$ is defined by the following commuting triangle:

$$\begin{array}{ccc} & \mathcal{W}_{1,2}^{012}(0) & \\ & \swarrow \quad \searrow & \\ \mathrm{D}^b(X^{012}) & \xrightarrow{T_{1,2}^{012}} & \mathrm{D}^b(X^{021}) \end{array}$$

Very similar constructions apply for the other five wall crossings in the hexagon (5.7).

Remark 5.8. For the horizontal arrows in (5.7), the 1-parameter subgroup controlling the wall-crossing is τ_0 . However, the value of the stability condition $\theta(\tau_0)$ is changing from negative to positive when we cross these walls, not vice versa. This means that when we define the two corresponding windows, $\mathcal{W}_{0,2}^{102}(0)$ and $\mathcal{W}_{0,2}^{021}(0)$, we must measure weights with respect to the subgroup $-\tau_0$, not τ_0 .¹²

Lemma 5.9. *There is a subcategory $\mathcal{V} \subset \mathrm{D}^b(\mathcal{X})$ such that, for any wall-crossing in the hexagon (5.7), the restriction functor*

$$\mathcal{V} \rightarrow \mathrm{D}^b(\mathcal{Y}_{i,j}^\sigma)$$

is an embedding with image $\mathcal{W}_{i,j}^\sigma(0)$.

We'll discuss the definition of this category \mathcal{V} shortly. Before we do, let's show that Proposition 5.6 follows from it as an immediate formality.

Proof of Proposition 5.6. For each wall-crossing we have a commuting diagram

$$\begin{array}{ccc} & \mathcal{V} & \\ & \downarrow & \\ & \mathcal{W}_{i,j}^\sigma(0) & \\ & \swarrow \quad \searrow & \\ \mathrm{D}^b(X^\sigma) & \xrightarrow{T_{i,j}^\sigma} & \mathrm{D}^b(X^{\sigma \cdot (ij)}) \end{array}$$

in which all arrows are equivalences. So we can complete the hexagon (5.7) to the following diagram

$$\begin{array}{ccccc} & & \mathrm{D}^b(X^{102}) & \xrightarrow{T_{0,2}^{102}} & \mathrm{D}^b(X^{120}) & & \\ & & \swarrow & & \searrow & & \\ & & \mathcal{V} & & \mathcal{V} & & \\ \mathrm{D}^b(X^{012}) & \xleftarrow{T_{0,1}^{012}} & & & & \xrightarrow{T_{1,2}^{120}} & \mathrm{D}^b(X^{210}) \\ & & \swarrow & & \searrow & & \\ & & \mathrm{D}^b(X^{021}) & \xrightarrow{T_{0,2}^{021}} & \mathrm{D}^b(X^{201}) & & \\ & & \swarrow & & \searrow & & \\ & & & & & & \mathrm{D}^b(X^{201}) \end{array}$$

in which every arrow is an equivalence and every triangle commutes. \square

5.2.2. *The category \mathcal{V} .* If we consider a GIT problem where the group acting is just \mathbb{C}^* , then as we've explained in Section 3.4.1, the machinery of [BFK12, HL12] tells us how to produce a section of the restriction functor

$$\mathrm{D}^b([Y_{i,j}^\sigma / \mathbb{C}^*]) \rightarrow \mathrm{D}^b(Y_{i,j}^\sigma // \mathbb{C}^*),$$

i.e. a lift of the derived category of the GIT quotient into the equivariant derived category of $Y_{i,j}^\sigma$. However, this general machinery doesn't just apply when the group is \mathbb{C}^* , it works for any reductive group. In particular, we can apply it to the GIT problem considered in the previous section, and for any phase X^σ it will tell us how to produce a section of the restriction functor

$$\mathrm{D}^b(\mathcal{X}) \rightarrow \mathrm{D}^b(X^\sigma).$$

We'll now explain how to run this machinery for our example.¹³ We shall see that it constructs a subcategory of $\mathrm{D}^b(\mathcal{X})$ which (after a little more work) we can show has the right properties to be the category \mathcal{V} required by Lemma 5.9.

¹²The functors $T_{0,2}^{102}$ and $T_{0,2}^{021}$ correspond to left-over-right crossings. If we were to use τ_0 -weights then we'd get the functors corresponding to right-over-left crossings.

¹³Again, the fact that our group is abelian makes our application considerably simpler than the general case.

We begin with the standard phase X^{012} . We need to fix an ordering of the three unstable subspaces, and for simplicity we'll choose the obvious one. Then we define

$$\begin{aligned} Z_0 &= \bar{Z}_0 & S_0 &= \bar{S}_0 \\ Z_1 &= \bar{Z}_1 - S_0 & S_1 &= \bar{S}_1 - S_0 \\ Z_2 &= \bar{Z}_2 - S_1 - S_0 & S_2 &= \bar{S}_2 - S_1 - S_0 \end{aligned} \quad (5.10)$$

so the unstable locus is the disjoint union of the three locally-closed subvarieties S_0, S_1 and S_2 . If we define $\tilde{\tau}_i = \tau_i$ for $i \in [1, 2]$, and $\tilde{\tau}_0 = -\tau_0$, then the subvariety S_i is precisely the attracting subvariety of Z_i under the subgroup $\tilde{\tau}_i$. This data defines a *KN stratification* of the unstable locus, in the sense of [HL12].¹⁴

Now we proceed as we did in the rank 1 case (Section 3.4.1), but for each of the strata simultaneously. So for each stratum S_i , we define the numerical invariant

$$\eta_i = \tilde{\tau}_i\text{-weight of } (\det N_{S_i}^\vee V)|_{Z_i}$$

and in this example we have $\eta_i = 2$, for all i , as we've already calculated (Section 3.4.2). Make a choice of three integers $k_0, k_1, k_2 \in \mathbb{Z}$, one for each stratum. We define the corresponding 'window'

$$\mathcal{W}(k_0, k_1, k_2) \subset \mathrm{D}^b(\mathcal{X})$$

to be the full subcategory

$$\mathcal{W}(k_0, k_1, k_2) = \left\langle \mathcal{E} \mid \begin{array}{l} \forall i, \text{ all homology sheaves of } L\sigma_i^* \mathcal{E} \text{ have } \tilde{\tau}_i\text{-weights} \\ \text{lying in the interval } [k_i, k_i + \eta_i) \end{array} \right\rangle$$

where $\sigma_i : Z_i \hookrightarrow V$ are the inclusion maps. The theorem is that the restriction functor

$$\mathcal{W}(k_0, k_1, k_2) \rightarrow \mathrm{D}^b(X^{012})$$

is an equivalence, for any values of k_0, k_1, k_2 .

Now let's move to the adjacent phase X^{021} , and construct a KN stratification by the same process (again using the obvious ordering on the unstable subspaces). Only the stratum S_2 changes. S_0 and S_1 stay the same, as do the fixed loci Z_0, Z_1, Z_2 . Also the destabilizing 1-parameter subgroups are now $\tilde{\tau}_0 = -\tau_0$, $\tilde{\tau}_1 = \tau_1$, and $\tilde{\tau}_2 = -\tau_2$. If we pick an integer l_i for each stratum then again have a window

$$\mathcal{W}^{021}(l_0, l_1, l_2) \subset \mathrm{D}^b(\mathcal{X})$$

where the weight-restriction condition is now defined with respect to this new KN stratification. The general machinery tells us that this window is equivalent, under restriction, to $\mathrm{D}^b(X^{021})$. However, because the Z_i have not changed, it's immediate that

$$\mathcal{W}^{021}(k_0, k_1, -k_2 - 1) = \mathcal{W}(k_0, k_1, k_2)$$

(the shift in the parameters here is just due to the slight asymmetry in the definition of a window). Consequently, we get some derived equivalences between X^{012} and X^{021} , by lifting into one of these windows and then restricting down again. These equivalences are not new, because this is just another way to describe the derived equivalences that we discussed in Section 3.4.2.

To see this, suppose that we don't restrict all the way down to X^{012} , but only to the larger open substack $\mathcal{Y}_{1,2}^{012}$. It's immediate from these definitions that any object in the window $\mathcal{W}(k_0, k_1, k_2)$ restricts to give an object in the window $\mathcal{W}_{1,2}^{012}(k_2)$, and it follows that the restriction functor

$$\mathcal{W}(k_0, k_1, k_2) \rightarrow \mathcal{W}_{1,2}^{012}(k_2) \subset \mathrm{D}^b(\mathcal{Y}_{1,2}^{012})$$

is an equivalence. So our equivalence $T_{1,2}^{012}$, which we defined via the window $\mathcal{W}_{1,2}^{012}(0)$, can also be defined via any of the windows $\mathcal{W}(k_0, k_1, 0)$.

From this discussion, it's clear that these windows in $\mathrm{D}^b(\mathcal{X})$ are candidates for the category \mathcal{V} required by Lemma 5.9. Unfortunately there is a subtlety which means we cannot apply this machinery

¹⁴This is a mild abstraction of a Kirwan–Kempf–Ness stratification. In the latter construction, the ordering on the strata is (partly) determined by a choice of inner product on the Lie algebra of T' , but that is irrelevant for our purposes.

immediately. If we move to a more distant phase, more unstable loci will change, and because of the precise definition (5.10) of the KN strata this means that the Z_i will change (their closures, the \bar{Z}_i , remain constant). Consequently, the definitions of the windows for different phases are different, and it is not immediately obvious that there is a single window lifting all of the $\mathcal{W}_{i,j}^\sigma(0)$.¹⁵ To get around this, we use an alternative characterization of these windows.

Proposition 5.11. *Let $\mathcal{O}(a, b, c)$ (with $a+b+c=0$) denote the equivariant line bundle on V associated to the corresponding character of the torus T' , and let*

$$\mathcal{V} = \langle \mathcal{O}(0, 0, 0), \mathcal{O}(-1, 1, 0), \mathcal{O}(-1, 0, 1) \rangle \subset \mathrm{D}^b(\mathcal{X})$$

denote the subcategory generated by these three line bundles. Let $\mathcal{W}(k_0, k_1, k_2) \subset \mathrm{D}^b(\mathcal{X})$ be a window defined with respect to the KN stratification for the standard phase as described above. Then

$$\mathcal{W}(0, 0, 0) = \mathcal{V}.$$

Proof. If we evaluate any of these three characters on any of the three subgroups τ_1, τ_2 or $-\tau_0$ then the result is either 0 or 1, so these three line bundles do lie in $\mathcal{W}(0, 0, 0)$ (and they are the only line bundles on \mathcal{X} which do). Furthermore, the definition of windows implies that any element of $\mathrm{D}^b(\mathcal{X})$ which has a finite resolution by copies of these three line bundles also lies in $\mathcal{W}(0, 0, 0)$. This proves that $\mathcal{V} \subset \mathcal{W}(0, 0, 0)$, and hence the restriction functor

$$\mathcal{V} \rightarrow \mathrm{D}^b(X^{012}) \tag{5.12}$$

is fully faithful. We claim that this functor is also essentially surjective. The proposition then follows since the inclusion of \mathcal{V} into $\mathcal{W}(0, 0, 0)$ must be an equivalence.

Since V is a vector space, any element of $\mathrm{D}^b(\mathcal{X})$, and hence any element of $\mathrm{D}^b(X^{012})$, has a finite resolution by line bundles from the set $\{\mathcal{O}(a, b, c)\}$. Also, on X^{012} we have three short exact sequences

$$\begin{aligned} 0 &\longrightarrow \mathcal{O}(-2, 1, 1) \xrightarrow{\begin{pmatrix} -a_2 & b_0 \end{pmatrix}} \mathcal{O}(-1, 1, 0) \oplus \mathcal{O}(-1, 0, 1) \xrightarrow{\begin{pmatrix} b_0 \\ a_2 \end{pmatrix}} \mathcal{O} \longrightarrow 0 \\ 0 &\longrightarrow \mathcal{O}(-1, 2, -1) \xrightarrow{\begin{pmatrix} -a_1 & b_0 \end{pmatrix}} \mathcal{O}(-1, 1, 0) \oplus \mathcal{O}(0, 1, -1) \xrightarrow{\begin{pmatrix} b_0 \\ a_1 \end{pmatrix}} \mathcal{O} \longrightarrow 0 \\ 0 &\longrightarrow \mathcal{O}(-1, -1, 2) \xrightarrow{\begin{pmatrix} -a_2 & b_1 \end{pmatrix}} \mathcal{O}(0, -1, 1) \oplus \mathcal{O}(-1, 0, 1) \xrightarrow{\begin{pmatrix} b_1 \\ a_2 \end{pmatrix}} \mathcal{O} \longrightarrow 0 \end{aligned}$$

associated to the three strata of the unstable locus. By repeatedly using twists of these sequences we can resolve any of the line bundles $\mathcal{O}(a, b, c)$ on X^{012} in terms of the three line bundles that generate \mathcal{V} . Consequently any element of $\mathrm{D}^b(X^{012})$ has a finite resolution in terms of these three line bundles, and so (5.12) is essentially surjective as claimed. \square

Remark 5.13. There are some delicate features of this argument that are worth highlighting. The fact that the values of the three parameters $k_0 = k_1 = k_2 = 0$ are correctly aligned is crucial: for most choices of these integers the window $\mathcal{W}(k_0, k_1, k_2)$ contains no line bundles at all. Also the precise action of \tilde{T} on \tilde{V} is important: for a general example of a torus action on a vector space there won't be *any* choice of integers on the strata such that the associated window is generated by line bundles.

Now we can show that this definition of \mathcal{V} suffices to produce all six of our wall-crossing functors.

Proof of Lemma 5.9. For each phase X^σ , we define $k_i^\sigma = 0$ or $k_i^\sigma = -1$ according to whether τ_i or $-\tau_i$ is the destabilizing subgroup in this phase, i.e. according to whether $\theta(\tau_i)$ is positive or negative.

Now let σ and $\sigma' = \sigma \cdot (ij)$ be two consecutive phases in the hexagon (5.7). For the phase X^σ , pick any ordering on the unstable subspaces, and then define a window $\mathcal{W}^\sigma(k_0^\sigma, k_1^\sigma, k_2^\sigma) \subset \mathrm{D}^b(\mathcal{X})$ using the corresponding KN stratification. Arguing as above, the restriction functor

$$\mathcal{W}^\sigma(k_0^\sigma, k_1^\sigma, k_2^\sigma) \rightarrow \mathrm{D}^b(\mathcal{Y}_{i,j}^\sigma)$$

¹⁵We are free to change our ordering on the strata for different phases, but this doesn't solve this problem.

is an embedding with image $\mathcal{W}_{i,j}^\sigma(0)$. On the other hand, the proof of Proposition 5.11 adapts immediately to show that

$$\mathcal{W}^\sigma(k_0^\sigma, k_1^\sigma, k_2^\sigma) = \mathcal{V},$$

and so our claim is proved. \square

Remark 5.14. As we move around between phases, the definitions of our windows in $D^b(\mathcal{X})$ change, but this result says that the actual windows do not change (for appropriate choices of the parameters). We've proved this using the fact that the relevant windows are generated by line bundles, but this is a very special condition and we believe that this result should actually hold in more generality. If true, this would imply analogous relations between autoequivalences for other GLSM examples.

In fact this is the main reason that we've included some discussion of the general machinery of windows on \mathcal{X} . If one wants to avoid it then it's reasonably straightforward to prove Lemma 5.9 directly from our definition of \mathcal{V} .

5.3. The poset of groupoid actions. When discussing functors between derived categories of schemes, one doesn't normally work in \mathbf{Cat}_1 but rather in the subcategory \mathbf{FM} where the morphisms between $D^b(Y)$ and $D^b(Z)$ are (isomorphism classes of) Fourier–Mukai kernels, that is objects of $D^b(Y \times Z)$. This subcategory is better-behaved in a number of ways. In our situation, we're actually working in a slightly more specialized subcategory.

Recall that each of our varieties X_Γ^σ is equipped with a fibration

$$\mu : X_\Gamma^\sigma \rightarrow \mathbb{C}^s,$$

where $(s+1)$ is the number of parts of Γ . Define $\mathbf{FM}_{\mathbb{C}^s}$ to be the category whose objects are schemes flat over \mathbb{C}^s , and whose morphisms are relative Fourier–Mukai kernels, i.e. objects of $D^b(Y \times_{\mathbb{C}^s} Z)$. Now let $\tilde{\Gamma}$ be a coarsening of Γ , with $\tilde{\Gamma}$ having $(\tilde{s}+1)$ parts. Associated to this data is a linear inclusion

$$j : \mathbb{C}^{\tilde{s}} \hookrightarrow \mathbb{C}^s,$$

and each subvariety $X_{\tilde{\Gamma}}^\sigma$ is the fibre product of the corresponding X_Γ^σ over this subspace.

Proposition 5.15. *There is a functor*

$$\begin{aligned} \text{Res} : \mathbf{FM}_{\mathbb{C}^s} &\rightarrow \mathbf{FM}_{\mathbb{C}^{\tilde{s}}} \\ Y &\mapsto Y \times_{\mathbb{C}^s} \mathbb{C}^{\tilde{s}}, \end{aligned}$$

acting on morphisms by the derived restriction functor

$$\begin{aligned} D^b(Y \times_{\mathbb{C}^s} Z) &\xrightarrow{j^*} D^b(\text{Res}(Y \times_{\mathbb{C}^s} Z)) \\ &= D^b(\text{Res}(Y) \times_{\mathbb{C}^{\tilde{s}}} \text{Res}(Z)). \end{aligned}$$

Proof. Indeed we may replace the linear inclusion $j : \mathbb{C}^{\tilde{s}} \hookrightarrow \mathbb{C}^s$ with any morphism of schemes $f : T \rightarrow S$, so that we define $\text{Res}(Y) := Y_T = Y \times_S T$, and use derived pullback f^* for the morphisms. Then taking schemes X, Y , and Z , all flat over S , and Fourier–Mukai kernels

$$P \in D^b(X \times_S Y), \quad Q \in D^b(Y \times_S Z),$$

we verify that Res is indeed a functor: composition of morphisms in \mathbf{FM}_S is by convolution of kernels, so we require

$$\text{Res}(P * Q) \cong \text{Res}(P) * \text{Res}(Q).$$

We have a diagram

$$\begin{array}{ccccc}
X_T \times_T Y_T \times_T Z_T & & & & \\
\downarrow p'_{13} & \searrow f' & & & \\
X_T \times_T Z_T & & X \times_S Y \times_S Z & & \\
\downarrow f & \swarrow p_{12} & \downarrow p_{13} & \searrow p_{23} & \\
X \times_S Y & & X \times_S Z & & Y \times_S Z
\end{array}$$

and by definition

$$\begin{aligned}
\text{Res}(P * Q) &:= \text{Res}(p_{13*}(p_{12}^*P \otimes p_{23}^*Q)) \\
&:= f^*p_{13*}(p_{12}^*P \otimes p_{23}^*Q) \\
&\cong p'_{13*}f'^*(p_{12}^*P \otimes p_{23}^*Q) \\
&\cong p'_{13*}(f'^*p_{12}^*P \otimes f'^*p_{23}^*Q)
\end{aligned}$$

where the first isomorphism is base change using flatness of p_{13} [L09, Theorem 3.10.3]. The proof is finished using the commutativity of the following two squares

$$\begin{array}{ccccc}
X_T \times_T Y_T & \xleftarrow{p'_{12}} & X_T \times_T Y_T \times_T Z_T & \xrightarrow{p'_{23}} & Y_T \times_T Z_T \\
\downarrow f & & \downarrow f' & & \downarrow f \\
X \times_S Y & \xleftarrow{p_{12}} & X \times_S Y \times_S Z & \xrightarrow{p_{23}} & Y \times_S Z
\end{array}$$

and the definition of $\text{Res}(P) * \text{Res}(Q)$. \square

Recall that in Section 5.1 we defined derived equivalences

$$T_{i,j}^{\sigma S_\Gamma} : D^b(X_\Gamma^\sigma) \rightarrow D^b(X_\Gamma^{\sigma \cdot (ij)})$$

corresponding to generators of the groupoid \mathcal{G}_Γ .

Proposition 5.16. *Each of our functors $T_{i,j}^{\sigma S_\Gamma}$ is actually a morphism in $\mathbf{FM}_{\mathbb{C}^s}$, i.e. their kernels are well-defined relative to the fibration μ .*

Proof. When $T_{i,j}^{\sigma S_\Gamma}$ corresponds to a flop between two distinct phases (case (a)) then this is clear from the geometric construction of the functor, because the common birational roof of the two phases is evidently defined relative to μ . When $T_{i,j}^{\sigma S_\Gamma}$ is a family spherical twist acting on a single phase (case (b)) then one can either verify the claim directly by inspecting the kernel for a family spherical twist, or just deduce it from case (a) using the next proposition. \square

Proposition 5.17. *Choose a partition Γ , and let $\tilde{\Gamma}$ be a coarsening of Γ . Pick σ, i, j such that the equivalence $T_{i,j}^\sigma$ is defined (and hence so are $T_{i,j}^{\sigma S_\Gamma}$ and $T_{i,j}^{\sigma S_{\tilde{\Gamma}}}$). Then*

$$\text{Res}\left(T_{i,j}^{\sigma S_\Gamma}\right) = T_{i,j}^{\sigma S_{\tilde{\Gamma}}}, \quad (5.18)$$

and in particular $T_{i,j}^{\sigma S_\Gamma}$ and $T_{i,j}^{\sigma S_{\tilde{\Gamma}}}$ are intertwined by pullback along the inclusions.

Proof. Let $\sigma' = \sigma \cdot (ij)$. We have four varieties as follows:

$$\begin{array}{ccc}
X_\Gamma^\sigma & \dashrightarrow & X_\Gamma^{\sigma'} \\
\uparrow j & & \uparrow j \\
X_\Gamma^\sigma & \dashrightarrow & X_\Gamma^{\sigma'}
\end{array}$$

The last part is the statement that

$$j^* \circ T_{i,j}^{\sigma S_\Gamma} \cong T_{i,j}^{\sigma S_{\tilde{\Gamma}}} \circ j^*,$$

and this follows from (5.18) using standard base change results [BBH, Proposition 6.1].

To prove (5.18), suppose firstly that i and j are in different parts of $\tilde{\Gamma}$, and so also in different parts of Γ . Then both birational maps are flops of a trivial family of rational $(-1, -1)$ -curves (as described in Section 3.3), and the functors $T_{i,j}^{\sigma\Gamma}$ and $T_{i,j}^{\sigma\tilde{\Gamma}}$ are the associated birational equivalences. Away from the flopping families both functors are trivial, so we only need to check the equality inside our very simple Zariski neighbourhoods of the flopping families. In these neighbourhoods, we're considering a trivial family of standard flops over \mathbb{C}^{s-1} , and then just restricting to a subspace $\mathbb{C}^{\tilde{s}-1} \subset \mathbb{C}^{s-1}$, so the equality of the functors is clear.

Now suppose that i and j are in the same part of Γ , and so also in the same part of $\tilde{\Gamma}$. A similar argument applies, except that now the birational maps are identities and the functors are family spherical twists.

The final case is that i and j are in different parts of Γ , but in the same part of $\tilde{\Gamma}$. By the same argument we can reduce to our simple Zariski neighbourhoods, and then we're considering a trivial family of copies of the geometry considered in Proposition 2.5, so the argument from there suffices. \square

Combining the above three propositions with Proposition 5.5, we can formally deduce the whole of Theorem 5.2.

Proof of Theorem 5.2. As before we let $\mathcal{G}_{\Gamma}^{\text{free}}$ be the free groupoid generated by all the arrows $t_{i,j}^{\sigma S_{\Gamma}}$, so \mathcal{G}_{Γ} is the quotient of $\mathcal{G}_{\Gamma}^{\text{free}}$ by the braid relations. If $\tilde{\Gamma}$ is a coarsening of Γ , then the above three propositions give us a commutative square as follows:

$$\begin{array}{ccc} \mathcal{G}_{\Gamma}^{\text{free}} & \xrightarrow{T_{\Gamma}} & \mathbf{FM}_{\mathbb{C}^s} \\ \downarrow & & \downarrow \text{Res} \\ \mathcal{G}_{\tilde{\Gamma}}^{\text{free}} & \xrightarrow{T_{\tilde{\Gamma}}} & \mathbf{FM}_{\mathbb{C}^{\tilde{s}}} \end{array}$$

If we set $\Gamma = \Gamma_{\text{fin}}$, then Proposition 5.5 says that T_{Γ} factors through \mathcal{G}_{Γ} . Therefore $T_{\tilde{\Gamma}}$ must factor through $\mathcal{G}_{\tilde{\Gamma}}$, and this holds for all partitions $\tilde{\Gamma}$. \square

Informally, the above proof just says that since the braid relations hold between the equivalences on the ambient space, they must continue to hold when we restrict these equivalences to subvarieties.

Finally, we use the faithfulness result of Seidel–Thomas to deduce Theorem 5.3.

Proof of Theorem 5.3. Let Γ be any partition, and consider coarsening it to Γ_{crs} . We have a commutative square as follows:

$$\begin{array}{ccc} \mathcal{G}_{\Gamma} & \xrightarrow{T_{\Gamma}} & \mathbf{FM}_{\mathbb{C}^s} \\ \downarrow & & \downarrow \text{Res} \\ B_{k+1} \simeq \mathcal{G}_{\text{crs}} & \xrightarrow{T_{\text{crs}}} & \mathbf{FM}_{\text{pt}} \end{array}$$

The functor T_{crs} is precisely the braid group action on the derived category of the surface $Y_0 = X_{\text{crs}}$ considered in [ST01], and it is proved there that the action is faithful. Since the functor $\mathcal{G}_{\Gamma} \rightarrow B_{k+1}$ is also faithful, we must have that T_{Γ} is faithful. \square

APPENDIX A. LIST OF NOTATIONS

$T \curvearrowright V$	torus T acting on vector space V	
θ	character of T	
X_θ	GIT quotient $V //_\theta T$	
Y_0	small resolution of A_k surface singularity	
\mathcal{X}	Artin quotient stack $[V/T]$	
\mathcal{F}	Fayet–Iliopoulos (FI) parameter space	
D^b	bounded derived category of coherent sheaves	
B_{k+1}	braid group on $k+1$ strands, labelled $[0, k]$	
\tilde{B}_{k+1}	affine braid group associated to affine Dynkin diagram \tilde{A}_k	
\mathcal{W}	window, in the bounded derived category of a stack	§2.1, §3.4
ψ	window equivalence, corresponding to wall-crossing	(2.2), §3.4
$T_{\mathcal{S}}$	spherical twist, around spherical object (or functor) \mathcal{S}	(2.4), (3.32)
a_i, b_i	co-ordinates on V for clockwise, and anti-clockwise, arrows	§2.2
μ_i	complex moment maps	(2.9)
Γ	partition of $[0, k]$, with pieces indexed by $[0, s]$	§2.2.1
S_Γ	Young subgroup of S_{k+1} , preserving partition Γ	(2.10)
B_Γ	mixed braid group: subgroup of B_{k+1} , preserving partition Γ	(2.10)
Q	toric data: matrix of weights	§3.1
Δ^k	k -simplex	§3.1
Π	polytope	§3.1
α_i, β_i	vectors in polytope, generating rays in the fan for X	§3.1
ϑ	lift of character θ	(3.2)
X_Γ	moment map subvariety of X	§3.2
γ	function $[0, k] \rightarrow [0, s]$ encoding partition Γ	§3.2
N_t	size of t^{th} piece of partition Γ	§3.11
$\tilde{\alpha}_t(\delta)$	vectors in polytope, generating rays in the fan for X_Γ	(3.14)
$\tilde{V} // \tilde{T}$	GIT quotient giving subvariety X_Γ	§3.2.2
X_Γ^σ	phase corresponding to partition Γ , permutation σ	§3.2.2
S_i, \tilde{S}_i	flopping subvarieties of X, X_Γ	(3.28), (3.31)
\mathcal{F}_Γ	FI parameter space for variety X_Γ	§4.2
ζ_t	FI parameter	§4.2
$\Gamma_{\text{fin}}, \Gamma_{\text{crs}}$	finest, and coarsest, partitions Γ	§4.2
\mathcal{G}_Γ	a path groupoid, associated to \mathcal{F}_Γ	§5.1
$t_{i,j}^{\sigma S_\Gamma}$	a morphism in \mathcal{G}_Γ , crossing adjacent strands i and j	§5.1
$T_{i,j}^{\sigma S_\Gamma}$	an equivalence, with source $D^b(X_\Gamma^\sigma)$	§5.1
\mathbf{Cat}_1	category of categories, with functors up to isomorphism	§5.1
τ_i	1-parameter subgroup of T	§5.2.1
S_i	unstable stratum	§5.2.1
Z_i	fixed loci of unstable stratum	§5.2.1
η_i	window width	§5.2.1
$\mathbf{FM}, \mathbf{FM}_B$	category of Fourier–Mukai kernels between varieties (relative to B)	§5.3

REFERENCES

- [Ado94] Alan Adolphson. Hypergeometric functions and rings generated by monomials. *Duke Math. Journal*, 73 (1994), no. 2, pp. 269–290.
- [Ann07] Rina Anno. Spherical functors. [arXiv:0711.4409](#).
- [AL13] Rina Anno and Timothy Logvinenko. Spherical DG-functors. [arXiv:1309.5035](#).
- [BFK12] Matthew Ballard, David Favero, and Ludmil Katzarkov. Variation of geometric invariant theory quotients and derived categories. [arXiv:1203.6643](#).
- [BBH] Claudio Bartocci, Ugo Bruzzo, and Daniel Hernández Ruipérez. Fourier–Mukai and Nahm Transforms in Geometry and Mathematical Physics. *Progress in Mathematics*, no. 276, Birkhäuser, 2009.
- [BR12] Roman Bezrukavnikov and Simon Riche. Affine braid group actions on derived categories of Springer resolutions. *Annales scientifiques de l'École normale supérieure*, 45, fascicule 4, pp. 535–599, 2012. [arXiv:1101.3702](#).
- [BO95] Alexei Bondal and Dmitri Orlov. Semiorthogonal decomposition for algebraic varieties. [arXiv:alg-geom/9506012](#).
- [CK12] Sabin Cautis and Joel Kamnitzer. Braiding via geometric Lie algebra actions. *Compositio Math.*, 148:2, pp. 464–506, 2012. [arXiv:1001.0619](#).
- [CPU13] Kwokwai Chan, Daniel Pomerleano, Kazushi Ueda. Lagrangian torus fibrations and homological mirror symmetry for the conifold. [arXiv:1305.0968](#).
- [CCIT09] Tom Coates, Alessio Corti, Hiroshi Iritani, Hsian-Hua Tseng. Computing genus-zero twisted Gromov-Witten invariants. *Duke Math. Journal*, 147 (2009), no. 3, 377–438. [arXiv:math/0702234](#).
- [DS12] Will Donovan and Ed Segal. Window shifts, flop equivalences and Grassmannian twists. *Compositio Math.* (to appear). [arXiv:1206.0219](#).
- [Ful93] William Fulton. Introduction to toric varieties. *Annals of Mathematics Studies*, no. 131, Princeton University Press, 1993.
- [HHP08] Manfred Herbst, Kentaro Hori, and David Page. Phases Of $\mathcal{N} = 2$ Theories In $1 + 1$ Dimensions With Boundary. [arXiv:0803.2045](#).
- [HL12] Daniel Halpern-Leistner. The derived category of a GIT quotient. [arXiv:1203.0276](#).
- [HLS] Daniel Halpern-Leistner and Ian Shipman. Autoequivalences of derived categories via geometric invariant theory. [arXiv:1303.5531](#).
- [Hor11] Kentaro Hori. Duality In Two-Dimensional (2,2) Supersymmetric Non-Abelian Gauge Theories. [arXiv:1104.2853](#).
- [HV00] Kentaro Hori and Cumrun Vafa. Mirror symmetry. [arXiv:hep-th/0002222](#).
- [HT06] Daniel Huybrechts and Richard P Thomas. \mathbb{P} -objects and autoequivalences of derived categories. *Math. Research Letters*, 13 (2006), pp. 87–98. [arXiv:math/0507040](#).
- [Kaw02] Yujiro Kawamata. D -equivalence and K -equivalence. *Journal of Differential Geometry*, 61(1):147–171, 2002. [arXiv:math/0205287](#).
- [L09] Joseph Lipman. Notes on derived functors and Grothendieck duality. In Foundations of Grothendieck Duality for Diagrams of Schemes, *Lecture Notes in Math.*, no. 1960, Springer–Verlag, New York, 2009, pp. 1–259. <http://www.math.purdue.edu/~lipman/Duality.pdf>.
- [Seg11] Ed Segal. Equivalences between GIT quotients of Landau–Ginzburg B-models. *Communications in Math. Physics*, 304(2):411–432, 2011. [arXiv:0910.5534](#).
- [ST01] Paul Seidel and Richard P Thomas. Braid group actions on derived categories of coherent sheaves. *Duke Math. Journal*, 108:37–108, 2001. [arXiv:math/0001043](#).
- [Slo] Peter Slodowy. Simple singularities and simple algebraic groups. *Lecture Notes in Math.*, no. 815, Springer–Verlag, New York, 1980.

WILL DONOVAN, THE MAXWELL INSTITUTE, SCHOOL OF MATHEMATICS, JAMES CLERK MAXWELL BUILDING, THE KING'S BUILDINGS, MAYFIELD ROAD, EDINBURGH, EH9 3JZ, U.K.

E-mail address: Will.Donovan@ed.ac.uk

ED SEGAL, DEPARTMENT OF MATHEMATICS, IMPERIAL COLLEGE LONDON, LONDON, SW7 2AZ, U.K.

E-mail address: edward.segal04@imperial.ac.uk