Coherent sheaves and quantum Coulomb branches II: quiver gauge theories and knot homology

BEN WEBSTER[1]
Department of Pure Mathematics, University of Waterloo &
Perimeter Institute for Theoretical Physics
Waterloo, ON
Email: ben.webster@uwaterloo.ca

Abstract. We continue our study of noncommutative resolutions of Coulomb branches in the case of quiver gauge theories. These resolutions are based on vortex line defects in quantum field theory, but have a precise mathematical description, which in the quiver case is a modification of the formalism of weighted KLR algebras. While best understood in a context which depends on the geometry of the affine Grassmannian and representation theory in characteristic p, we give a description of the Coulomb branches and their commutative and non-commutative resolutions which can be understood purely in terms of algebra.

This allows us to construct a purely algebraic version of the knot homology theory defined by Aganagić from a physical perspective, categorifying the Reshetikhin-Turaev invariants for minuscule representations of type ADE Lie algebras. We show that this homological invariant agrees with the categorification of these invariants previously defined by the author, and thus with Khovanov-Rozansky homology in type A.

Slodowy slices in type A and symmetric powers in $\mathbb{C}^2$ are special cases of these Coulomb branches, and in this case, we recover the noncommutative Springer resolution of Bezrukavnikov and those constructed using the Cherednik algebra by Bezrukavnikov, Finkelberg and Ginzburg.

Author's note: As this is a continuation of the first part of this paper [Weba], we will use notation and constructions from that paper without additional reference or comment. You can spot links to part I as they will have blue outline instead of red. Note that in some PDF viewers, these links will not open correctly, due to a `#` getting coverted to `%23`. Using Adobe Acrobat seems to solve this issue, as does manually changing the `%23` in the URL back to `#`.

5. (Re)introduction

In [Weba, Webb, Webc], we developed a general theory of Coulomb branches from an algebraic perspective. We showed that the Coulomb branch algebra itself, its extended category (of line operators) and various related algebras, such as the non-commutative

[1]Supported by the NSF under Grant DMS-1151473 and the Alfred P. Sloan Foundation. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities.

resolution constructed through quantization by Bezrukavnikov and Kaledin [BK08, Kal08, Bez06] and the category controlling their Gelfand-Tsetlin modules all have explicit combinatorial descriptions. In this sequel to these papers, we focus in on understanding this construction in the quiver gauge case, especially on the non-commutative resolution and corresponding geometric constructions with coherent sheaves. Applications to the representation theory of quantum Coulomb branches in characteristic 0 have already been covered extensively in [KTW$^+$, Webb, Web19a], so we will only discuss these in passing.

At the root of this perspective is a description of the Coulomb branch algebra as paths in the space $T_{\mathbb{R}}/W$, the quotient of the compact torus of the gauge group G modulo the Weyl group W, modulo certain relations (see [Webc, (2.5a–c)]). In the case of GL_n, this space can be identified with the configuration of n points on the circle $\mathbb{R}/\mathbb{Z}$ (allowing collisions), and thus a path in this space with a diagram drawn on the cylinder.

Fix a quiver Γ with vertex set $\mathcal{V}(\Gamma)$, and dimension vectors $\mathbf{v}, \mathbf{w} \colon \mathcal{V}(\Gamma) \to \mathbb{Z}_{\geq 0}$ for this quiver. We should emphasize that we do allow edge loops. By a **quiver gauge theory** we mean the one attached to the gauge group and matter (G, V) given by:

$$(5.1) \qquad G = \prod GL(\mathbb{C}^{v_i}) \qquad V = \Big(\bigoplus_{i \to j} \operatorname{Hom}(\mathbb{C}^{v_i}, \mathbb{C}^{v_j})\Big) \bigoplus \Big(\bigoplus_{i \in \mathcal{V}(\Gamma)} \operatorname{Hom}(\mathbb{C}^{w_i}, \mathbb{C}^{v_i})\Big),$$

As described above, we can think of a path in $T_{\mathbb{R}}/W$ as a path in a labeled configuration space where v_i points have label i, that is, as a string diagram on the cylinder where strands are labeled by points in the Dynkin diagram. When we translate the relations [Webc, (2.5a–c)] into this framework, they suddenly become very familiar to any one used to categorification: those of KLRW (weighted KLR) algebra, as presented in [Web19b]. The author and his collaborators exploited this in [KTW$^+$] to study the representation theory of shfited Yangians, but here we apply the same idea with a more geometric perspective. Based on these ideas, we attach to the choice of $\Gamma, \mathbf{v}, \mathbf{w}$ and some additional auxilliary data a **cylindrical KLRW algebra**; these are defined by diagrams of red and black strands drawn on a cylinder subject to local relations like those of those of the **weighted KLR algebra** [Web19b].

Recall that a **non-commutative crepant resolution of singularities** for a commutative algebra A (or the variety $\operatorname{Spec} A$) is an associative algebra R such that R-mod behaves like the category of coherent sheaves on a crepant resolution. For our purposes, we only need consider the case where there is an idempotent e such that $A = eRe$: in this case, R is a NCCR if eR is Cohen-Macaulay as an A-module, and R has global dimension equal to the Krull dimension of A. Since a (usual) crepant resolution of a symplectic singularity is symplectic, we may as well call a NCCR of A a **non-commutative symplectic resolution** in the case where $\operatorname{Spec} A$ has symplectic singularities.

Our main result is that:

Theorem D *If $\mathfrak{M}$ is the Coulomb branch of an quiver gauge theory that admits a BFN resolution, then the cylindrical KLRW algebra $\mathring{R}$ with the same underlying combinatorial data defines a non-commutative symplectic resolution of singularities of $\mathfrak{M}$.*

Since we'll be interested in considering versions of these varieties in positive characteristic, let us note that this result is only proven for these varieties over a field of characteristic 0 or sufficiently large (to the worried physicist: in particular, this result holds over $\mathbb{C}$), though we believe it actually is true in arbitrary characteristic.

By a **BFN resolution**, we mean a symplectic resolution of singularities arising from the constructions of Braverman-Finkelberg-Nakajima in [BFN]. In physics terms, this means that the Coulomb branch becomes smooth at a generic choice of mass parameters (and trivial FI parameters).

The quiver gauge theories which admit BFN resolutions include those with cyclic or linear quivers (affine type A), and those for type D and E quivers with w_i only non-zero on nodes with minuscule fundamental representations, in both cases chosen so they correspond to dominant weight spaces (these theories are often called "good" in the physics literature, to contrast them with "bad" and "ugly" theories). The most familiar examples of these are resolved Slodowy slices (or more generally, S3 varieties) for a linear quiver and the Hilbert scheme of points on $\mathbb{C}^2$ (or more generally, the resolved Kleinian singularity $\mathbb{C}^2/\mathbb{Z}/\ell\mathbb{Z}$). We'll discuss these examples in more detail in Appendix A.

These non-commutative resolutions arise as the endomorphisms of tilting generators on symplectic resolutions of the same varieties; in fact, this is how we prove that they give non-commutative resolutions.

Theorem E *If $\mathfrak{M}$ is the Coulomb branch of a quiver gauge theory, and $\tilde{\mathfrak{M}}$ a BFN resolution, then:*

1. *the homogeneous coordinate ring of $\tilde{\mathfrak{M}}$ is an algebra of twisted cylindrical KLR diagrams, modulo local relations.*
2. *For each NCSR R of Theorem D, the variety $\tilde{\mathfrak{M}}$ admits a tilting generator $\mathcal{T}$, described as an explicit module over the homogeneous coordinate ring, such that* $\operatorname{End}(\mathcal{T}) = R$*; in particular,* $\operatorname{End}(\mathcal{T})$ *is a cylindrical KLRW algebra $\mathring{R}$.*
3. *The wall-crossing functors relating different tilting generators are given by tensor product with explicit bimodules, modeled on the braiding functors of* [Web17a, §6]*; more generally, the Schober connected to these functors can constructed using the representation theory of related algebras.*

As with Theorem D, this is only currently proven in characteristic 0, though we believe it holds in all characteristics.

Finally, we turn to applying these algebras in topology for the construction of homological knot invariants. Recent work [Aga] of Aganagić defines such a knot invariant based on the coherent sheaves on $\tilde{\mathfrak{M}}$ (which she denotes $\mathcal{X}$; see [Aga, §3.1.1]) in the case where Γ is of ADE type. In this case, our quiver gauge theory must correspond to a tensor product of minuscule representations in order to construct a BFN resolution. Aganagić's construction depends on an action of affine tangles of the category $D^b(\mathrm{Coh}\,\tilde{\mathfrak{M}})$, which arises naturally from a family of central charge functions on the Grothendieck group of this category: the action of affine braids is given by wall-crossing functors associated to certain singular loci for this function, and the action of cups and caps from a filtration that arises on the category as we approach certain walls. We can capture this precisely in a **real variation of stability** in the sense of Anno, Bezrukavnikov and Mirković [ABM15].

By Theorem E, we can translate Aganagić's action into complexes of bimodules over a cylindrical KRLW algebra. In particular, we give a combinatorial description of the central charge function, which you can think of this as "integration over a non-commutative resolution." The affine braid group action is covered in Theorem E, and we add cup and cap functors that extend this to an annular tangle action on the categories $D^b(\mathrm{Coh}\,\tilde{\mathfrak{M}})$ for different theories. This gives us an annular knot invariant defined using only KLRW algebras, which agrees with that of Aganagić; in particular, this removes any questions about mathematical rigor from [Aga, Th. 5*]. It's a tricky philosophical question how rigorously you can check that a construction which does not claim to be mathematically rigorous matches another one, but at each point, our construction matches the description in [Aga] under the equivalences of E.

More precisely, let Γ be of ADE type, and let $\mathring{R}^{\mathbf{j}}$ be the corresponding cylindrical KLR algebra, summed over all possible numbers of black strands. Consider an oriented affine ribbon tangle T labeled with minuscule fundamental representations. We can read off the labels at the bottom of this tangle (taking the dual representation if the strand is oriented downward) to get a sequence $\mathbf{j}$ of minuscule fundamental representations, and similarly read off $\mathbf{j}'$ from the top.

Theorem F *We have an induced functor* $\Phi(T)\colon D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod}) \to D^b(\mathring{R}^{\mathbf{j}'}\text{-fdmod})$ *which is compatible with composition of tangles up to isomorphism of functors. This defines a link invariant which agrees with those of* [Aga, Th. 5*] *and* [Web17a, §8]*, and in type A with Khovanov-Rozansky homology.*

It's worth nothing that this defines an invariant of annular links (we think of a usual link as an annular link by embedding B^3 into the annulus times an interval). We expect that in type A, this agrees with the annular Khovanov-Rozansky homology of Queffelec and Rose [QR18]; we lay out some preliminary steps to checking this fact, but verifying it carefully is beyond the scope of this paper.

Acknowledgements. Many thanks to Mina Aganagić, Roman Bezrukavnikov, Alexander Braverman, Kevin Costello, Tudor Dimofte, Michael Finkelberg, Justin Hilburn, Joel Kamnitzer, Gus Lonergan, Ivan Losev, Alex Weekes and Philsang Yoo for useful discussions on these topics and Youssef Mousaaid and Alistair Savage for not complaining about the theft of their diagrams.

6. Cylindrical KLRW algebras

6.1. **The definition of cylindrical KLRW algebras.** As in the introduction, let Γ be a quiver, and $\mathbf{v}, \mathbf{w} \colon \mathcal{V}(\Gamma) \to \mathbb{Z}_{\geq 0}$ dimension vectors. First, we note some basic facts about quiver gauge theories. In this case, the group $H = N_{GL(V)}(G)$ is generated by G, the product $GL(\mathbb{C}^{w_i})$ acting by precomposition in the obvious way, and by $GL(\mathbb{C}^{\chi_{i,j}})$ where $\chi_{i,j}$ is the number of edges $i \to j$, acting by taking linear combinations of the maps along these edges, i.e. via the isomorphism

$$\bigoplus_{i\to j} \mathrm{Hom}(\mathbb{C}^{v_i}, \mathbb{C}^{v_j}) \cong \bigoplus_{(i,j)\in\mathcal{V}(\Gamma)} \mathrm{Hom}(\mathbb{C}^{v_i}, \mathbb{C}^{v_j}) \otimes \mathbb{C}^{\chi_{i,j}}.$$

Thus, for any unitary element $g \in H$, after multiplying by an element of G and conjugating, we can find classes $\beta_e, \gamma_{i,k} \in \mathbb{R}/\mathbb{Z}$ for $e \in \mathcal{E}$ and $i \in \mathcal{V}(\Gamma), k = 1, \dots, w_i$ such that g acts on the homomorphism along e by the scalar $\exp(2\pi i\beta_e)$ and on $\mathbb{C}^{w_i}$ with eigenvalues $\exp(2\pi i\gamma_{i,k})$.

6.1.1. *Cylindrical KLRW diagrams.* In this section, we will discuss how to utilize these classes to define a quite different object: a **cylindrical KLRW algebra**. For simplicity, we will assume that $\beta_e, \gamma_{i,k}$ satisfy the following genericity assumptions:

- For all edges e, $\beta_e \neq 0$.
- We have $\gamma_{i,k} \neq \gamma_{j,\ell}$ for all $(i,k) \neq (j,\ell)$.
- If two edges e, e' have the same tail and head, then $\beta_e \neq \beta_{e'}$. If the tail of e is the head of e' and *vice versa*, then $\beta_e \neq \beta_{e'}$.

These restrictions are primarily to simplify the relations we draw below; the non-generic relations can be derived from these in a method we will describe.

Definition 6.1 *A* **cylindrical KLRW** *diagram is a collection of finitely many oriented curves in $\mathbb{R}/\mathbb{Z} \times [0,1]$ of the form $\{(\bar{\pi}(t), t) \mid t \in [0,1]\}$ for some path $\bar{\pi} \colon [0,1] \to \mathbb{R}/\mathbb{Z}$. Each curve is labeled with an element $i \in \mathcal{V}(\Gamma)$ and decorated with finitely many dots. For each curve with label i, and each edge $e : i \to j$, we draw a "ghost" curve at $\{(\bar{\pi}(t) - \beta_e, t) \mid t \in [0,1]\}$, which we label with j, and draw as dashed. We also draw red lines at $x = -\gamma_{i,k}$, labeled with i. These should satisfy the same genericity conditions as weighted KLR diagrams from* [Web19b, Def. 2.3]*: we have no tangencies or triple points involving any combination of strands and ghosts (except the case where $\beta_e = 0$, so a strand and ghost coincide), and no dots on crossings. We identify any diagrams* *that differ by isotopies preserving these genericity conditions. The only local structures we allow are those shown below:*

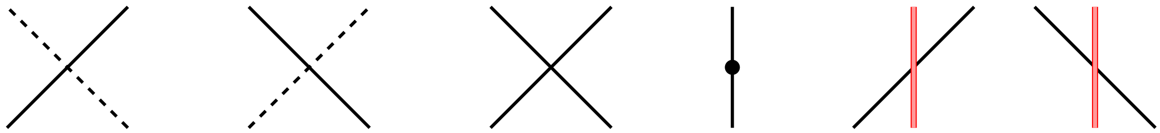

We'll draw these on the page in the rectangle $[0,1] \times [0,1]$ *with seams on the left and right side of the diagram where we should glue to obtain the cylindrical diagram. If our graph* Γ *is given by* $i \leftarrow j \leftarrow k$*, then an example of such a diagram is*

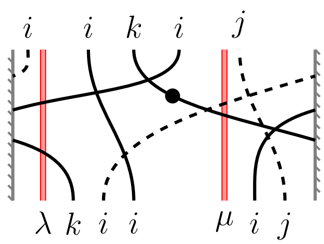

If there were any strands with label j*, they would have ghosts with label* k*, but we have left these out for simplicity of drawing.*

As in [Web19b]*, we let a* **cylindrical loading** *be a map to* $\mathcal{V}(\Gamma)$ *from a finite subset of* $\mathbb{R}/\mathbb{Z}$ *which avoids* $x = -\gamma_{i,k}$ *and such that if there is an edge* $e\colon i \to j$*, then there is no pair of elements* x *and* y *mapping to* i *and* j *with* $x - y = \beta_e$*. Note slicing a cylindrical KLRW diagram at a fixed value of* y *gives a loading. We will record a cylindrical loading by the list* $\mathbf{i} = (i_1, \dots, i_n) \in \mathcal{V}(\Gamma)^n$ *of the labels on the subset as we read from* $x = 0$ *to* $x = 1$ *in increasing order (so this is their cyclic order around the circle), and by the images* $\mathbf{a} = (a_1, \dots, a_n) \in (\mathbb{R}/\mathbb{Z})^n$*.*

Note that the cylindrical loadings divide naturally into equivalence classes that are bounded by the subtori where $a_j - a_k = \beta_e$ when $e\colon i_j \to i_k$ and $a_j = -\gamma_{i,k}$ with $i = i_j$, where the a_j's are the positions of the strands and i_j the label on the jth strand; compare with [Web19b, Prop. 2.12] for the linear case. If we deform β_e and $\gamma_{i,k}$, usually the set of these chambers will not change, but it will when we hit a point where there is a redundancy between these equations:

Lemma 6.2 *If the set of possible equivalence classes of cylindrical loadings is not locally constant near* $(\beta_e, \gamma_{i,k})$ *then for some* $i, j \in \mathcal{V}(\Gamma)$*, we have a path*

$$i = i_0 \xrightarrow{e_1} i_1 \xrightarrow{e_2} i_2 \xrightarrow{e_3} \cdots \xrightarrow{e_n} i_n = j$$

such that in $\mathbb{R}/\mathbb{Z}$*:*

$$\gamma_{j,\ell} - \gamma_{i,k} + \sum_{p=1}^{n} \beta_{e_i} = 0 \tag{6.1}$$

The set of cylindrical loadings inherits a measure from the usual measure with volume 1 on $(\mathbb{R}/\mathbb{Z})^n$ (that is, its Haar measure as a Lie group).

6.1.2. *Relations.*

Definition 6.3 *The* **cylindrical KLRW algebra** $\mathring{R}$ *attached to the data* $\Gamma, \mathbf{v}, \mathbf{w}, \beta_*, \gamma_{*,*}$ *is the quotient of the formal span over a ring* $\Bbbk$ *of cylindrical wKLR diagrams for these data by the local relations*

(6.2a) 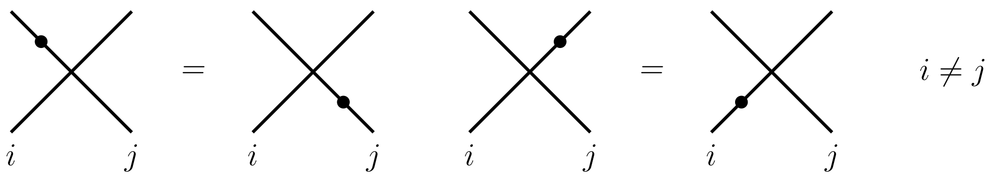

(6.2b) 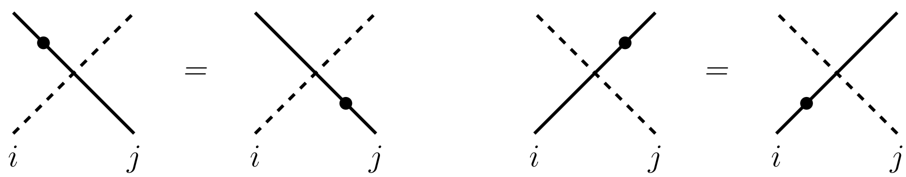

(6.2c) 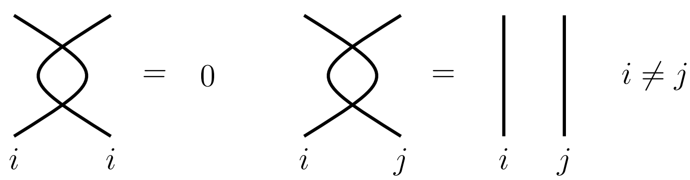

(6.2d) 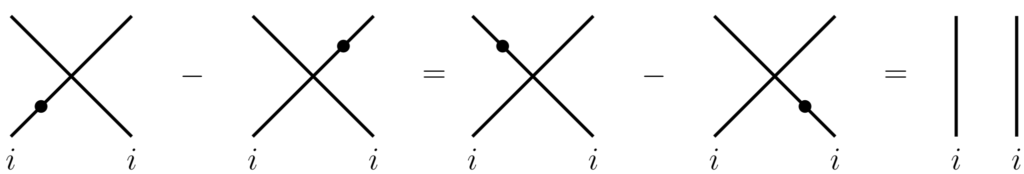

(6.2e) 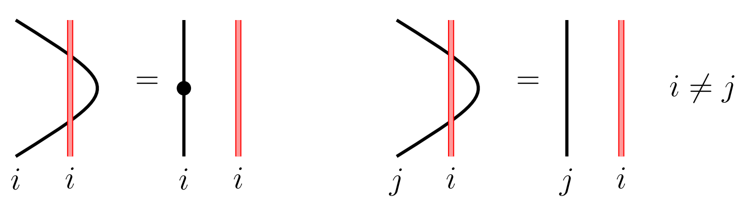

(6.2f) 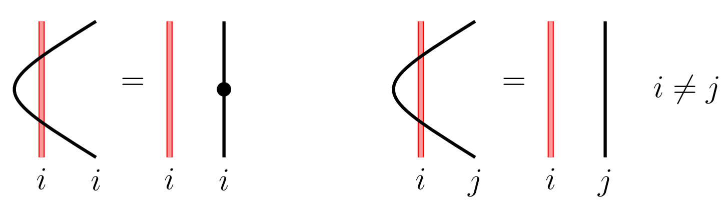

Given an edge $e\colon j \to i$, we have that:

(6.2g)

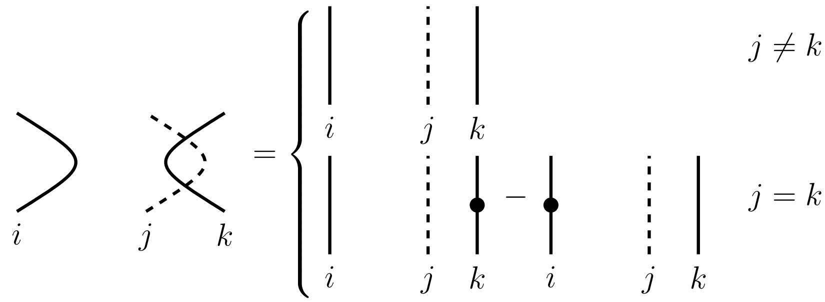

(6.2h)

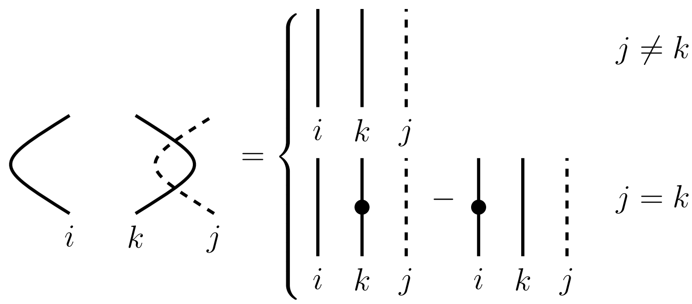

(6.2i)

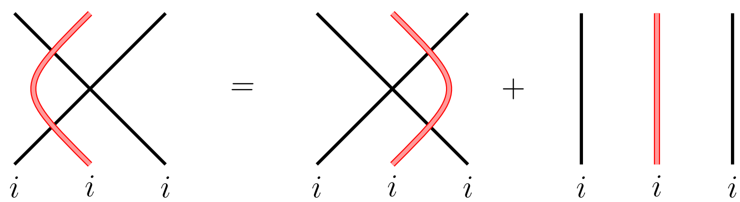

(6.2j)

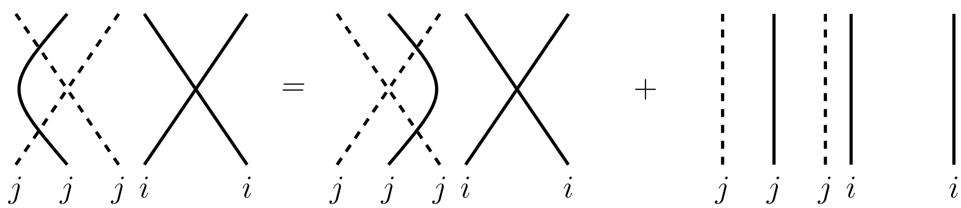

(6.2k)

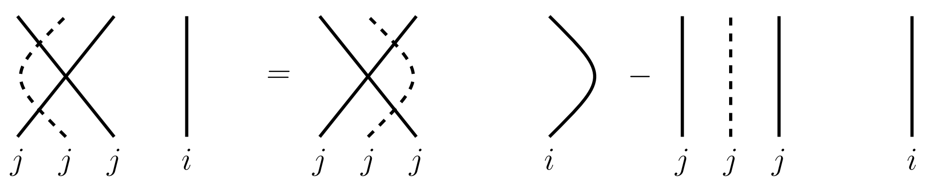

For all other triple points, we set the two sides of the isotopy through it equal.

Remark 6.4. As mentioned before, to simplify our relations, we have assumed that our parameters are generic. If you wish to work with non-generic parameters, they will have the same form as above, relating diagrams isotopic across a non-generic diagram (tangency, dot on crossing or triple point). You can find the appropriate relations in that case by taking a very small deformation of the parameters $\beta_*, \gamma_{*,*}$, applying the relations above (possibly multiple times), and then deforming the parameters back to the non-generic ones of interest.

Remark 6.5. Note that in many earlier works, such as [Web17a, Web17b], we had an additional non-local relation setting a diagram to 0 if it had a black strand at far left of the diagram; it not clear how one could interpret this relation on the circle. Ultimately this corresponds to the fact that coherent sheaves have no obvious notion of category $\mathcal{O}$.

One useful observation about relations: as discussed in [Web19b, Prop. 2.15], if $\eta\colon \mathcal{V}(\Gamma) \to \mathbb{R}/\mathbb{Z}$ is any function, we can replace $\beta_*, \gamma_{*,*}$ with the cohomologous value:

$$\beta'_e = \beta_e - \eta_{h(e)} + \eta_{t(e)} \qquad \gamma'_{i,k} = \gamma_{i,k} - \eta_i \tag{6.3}$$

We can define a map S_η between the cylindrical KLRW algebras for these different choices of parameters by shifting the position of each solid strand with label i by η_i; by (6.3), this has the effect of shifting each ghost with the same label by the same amount, so all ghost/solid crossings remain intact. This shows:

Lemma 6.6 *The map S_η is an isomorphism of algebras.*

The case that will be of most interest to use is when Γ is a tree. In this case, we can always use a choice of η to reduce to the case where $\beta_e = 0$ for all $e \in \mathcal{E}$, since thinking of β_* as a 1-cocycle, it must be a coboundary. In this case, there are no visible ghosts, since the coincide with the original strand so our diagrams look like.

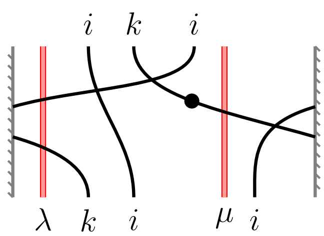

The relations can be deduced in this case from Remark 6.4, but for completeness, let us note that if we have a single edge $j \to i$, then applying this approach to the relations (6.2d,6.2g,6.2h), we have that

(6.4)

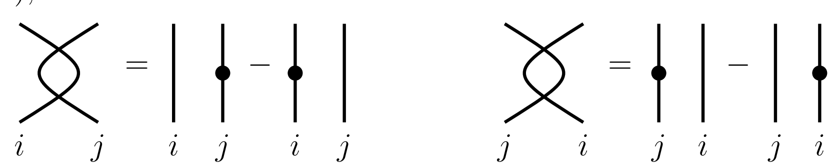

Similarly in (6.2j,6.2k), we just move the solid strand with label i on top of the ghost with label j. Note that while the relations (6.2a–6.2k) did not depend at all on the number of edges between vertices, because different edges will have separate ghosts, the relations above need to be modified if there are multiple edges $i \to j$ (following the standard recipe of [KL11]).

6.1.3. *Relation to planar KLRW algebras.* As suggested by the name, cylindrical KLRW algebras are related to KLRW (or weighted KLR) algebras. We will distinguish these by calling them "planar." Let us just start for the record that this does not refer to the structure of the corresponding quiver (which is arbitrary) but the surface on which the diagrams are drawn. These are spanned by diagrams with the same local properties as cylindrical KLRW diagrams, but in $\mathbb{R} \times [0,1]$ instead of $S^1 \times [0,1]$, and subject to the same local relations (8.4c–6.2k). We discuss these algebras largely with the aim of transferring certain calculations done for planar KLRW algebras to cylindrical ones. We'll only discuss this connection in the case where $\beta_e = 0$.

In this case, we consider the algebra $\tilde{T}^{\mathbf{j}}$ where we fix the labels on red strands to be $\mathbf{j} = (j_1, \dots, j_\ell)$, but allow any number of black strands. Note that in this case $\tilde{T}^{\emptyset}$ is the usual KLR algebra defined in [KL11]. The category of $\tilde{T}^{\emptyset}$-modules has a monoidal structure induced by the induction functor defined in [KL09, §2.6]: put simply, horizontal composition induces an algebra map $\tilde{T}^{\emptyset} \otimes \tilde{T}^{\emptyset} \to \tilde{T}^{\emptyset}$, and induction is pushforward by this map.

Similarly, horizontal composition also gives a map $\tilde{T}^{\emptyset} \otimes \tilde{T}^{\mathbf{j}} \otimes \tilde{T}^{\emptyset} \to \tilde{T}^{\mathbf{j}}$, and pushforward gives a functor sending a triple (K, M, N) with of $\tilde{T}^{\emptyset}$-modules K, N, and a $\tilde{T}^{\mathbf{j}}$-module M to a $\tilde{T}^{\mathbf{j}}$-module $K \circ M \circ N$.

We can think of the cylindrical KLRW algebra as an *affinization* of the $\tilde{T}^{\emptyset}$-mod bimodule structure on $\tilde{T}^{\mathbf{j}}$-mod. We can say this a little more precisely when we think about the **planar KLRW category** $\tilde{\mathcal{T}}$, where the objects are words in black and red copies of $\mathcal{V}(\Gamma)$, and morphisms are KLRW diagrams joining these words, modulo the relations (6.2a–6.2k); for notational purposes, let $\mathcal{U}^-$ be the planar KLRW category with no red lines. Similarly, there is a **cylindrical KLRW category** with objects given by cyclic words (or if you prefer, periodic words) and morphisms by cylindrical KLRW diagrams modulo the same local relations.

Consider the category $\mathcal{Q}$ obtained by adjoining to $\tilde{\mathcal{T}}$ an isomorphism $\xi_{\mathbf{i},\mathbf{i}'} \colon \mathbf{i} \circ \mathbf{i}' \to \mathbf{i}' \circ \mathbf{i}$ for $\mathbf{i}'$ an object in $\mathcal{U}^-$ (a word only in the black copy of $\mathcal{V}(\Gamma)$), and $\mathbf{i}'$ an object in $\tilde{\mathcal{T}}$ (a word in the red and black copies of $\mathcal{V}(\Gamma)$), and impose the additional relations:

$$\xi_{\mathbf{i},\mathbf{i}'\circ\mathbf{i}''} = \xi_{\mathbf{i}''\circ\mathbf{i},\mathbf{i}'}\xi_{\mathbf{i}\circ\mathbf{i}',\mathbf{i}''} \tag{6.5}$$

$$\xi_{\mathbf{m},\mathbf{m}'} \circ (f \otimes g) = (g \otimes f) \circ \xi_{\mathbf{i},\mathbf{i}'} \tag{6.6}$$

for $\mathbf{i}, \mathbf{m}$ words with red strands labeled by $\underline{\boldsymbol{\lambda}}$, $\mathbf{i}', \mathbf{i}'', \mathbf{m}'$ words only in the black strands and $f \colon \mathbf{i} \to \mathbf{m}$ and $g \colon \mathbf{i}' \to \mathbf{m}'$ arbitrary morphisms. Note the similarity to the work of Mousaaid and Savage on affinization of monoidal categories; this not quite a special case of their work, since we are using a bimodule category, and they only consider the action of a monoidal category on itself on the left and right.

Proposition 6.7 *The category $\mathcal{Q}$ defined above is equivalent to the cylindrical KLRW category via the functor sending* $\mathbf{i}$ *to the same word considered cyclically, sending any morphism in the planar KLRW category to the morphism drawn in in $S^1 \times [0,1]$ by* *embedding $\mathbb{R}$ as $S^1\backslash\{*\}$, and $\xi_{\mathbf{i},\mathbf{i}'}$ to the diagram with bottom given by the concatenation $\mathbf{i}\circ\mathbf{i}'$, which moves the strands in $\mathbf{i}'$ around the back of the cylinder in the positive direction. That is:*

$\xi_{\mathbf{i},\mathbf{i}'} \mapsto$ 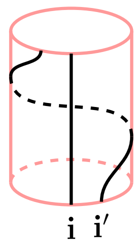

Proof. To show this functor is well-defined, we need only check that the relations (6.5–6.6) hold, which is an easy geometric verification by the relations shown below:

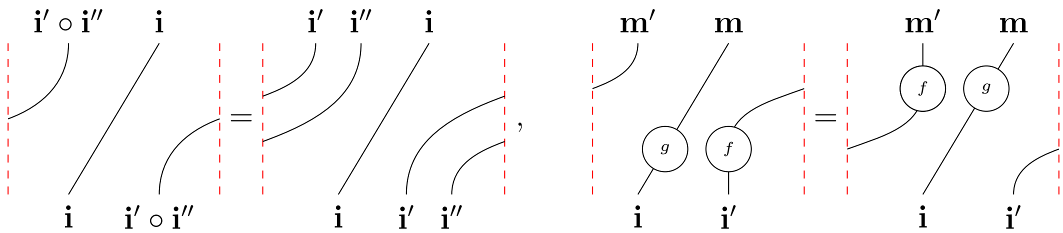

Now, we wish to show this functor is fully faithful. Fullness is clear from the fact that we can isotope crossings and dots off of the seam and we can use ξ and its inverse to get any crossings over the seam where we glue $\mathbb{R}$ to form S^1. Thus, only faithfulness remains. If a planar combination of diagrams is 0, this is derived by a finite chain of local relations and isotopies. All the local relations are applied when away from the seam, and so they hold in the planar category as well. Thus, we need only consider isotopies. We can any isotopy in the cylinder can be written using the cancelation of $\xi_{*,*}$ and its inverse, the relations ((6.5)–(6.6)) and isotopies avoiding the seam. This shows faithfulness. $\square$

This equivalence has a manifestation on the level of modules: given a $\tilde{T}^{\mathbf{j}}$-module M, we can consider the tensor product $\mathring{R}^{\mathbf{j}} \otimes_{\tilde{T}^{\underline{\boldsymbol{\lambda}}}} -$. This is the pushforward by the inclusion of the planar KLRW category in the cylindrical.

Lemma 6.8 *The functor $\mathring{R}^{\mathbf{j}} \otimes_{\tilde{T}^{\underline{\boldsymbol{\lambda}}}} -$ is exact.*

This is proven on page 37. This perspective is useful in that it shows that a bimodule $\mathfrak{B}$ between $\tilde{T}^{\mathbf{j}}$ and $\tilde{T}^{\mathbf{j}'}$ compatible with the bimodule structure on these categories induces a bimodule between cylindrical KLRW categories. For our purposes, it will be easier to say this in terms of functors between derived categories. We say that the functor $\mathbb{B} = \mathfrak{B} \overset{L}{\otimes} -$ is **strongly equivariant** if it commutes with the action of $\mathcal{U}^-$ on the left and the right, i.e the functors $\mathcal{U}^- \times \tilde{\mathcal{T}} \times \mathcal{U}^- \to \tilde{\mathcal{T}}$ defined by

$$K \circ \mathbb{B}M \circ N \leftarrow\!\shortmid (K, M, N) \mapsto \mathbb{B}(K \circ M \circ N)$$

are isomorphic.

Lemma 6.9 *If $\mathbb{B}$ is strongly equivariant for the left and right actions of $\mathcal{U}^-$, then there is an induced functor*

$$\mathring{\mathbb{B}}\colon D^b(\mathring{R}^{\mathbf{j}}\text{-mod}) \to D^b(\mathring{R}^{\mathbf{j}'}\text{-mod})$$

which is compatible with composition: if $\mathbb{B} = \mathbb{B}_1\mathbb{B}_2$, then $\mathring{\mathbb{B}} = \mathring{\mathbb{B}}_1\mathring{\mathbb{B}}_2$.

Of course, we have

$$\mathring{\mathbb{B}}(M) = \mathring{\mathfrak{B}} \otimes_{\mathring{R}^{\mathbf{j}}} M \qquad \mathring{\mathfrak{B}} = \mathring{\mathbb{B}}(\mathring{R}^{\mathbf{j}}).$$

Proof. It's enough to define the functor $\mathring{\mathbb{B}}$ on modules of the form $\mathring{R}^{\mathbf{j}}e(\mathbf{i})$ for $\mathbf{i}$ a cyclic word. This is just $\mathring{R}^{\mathbf{j}'} \otimes_{\tilde{T}^{\mathbf{j}'}} \mathfrak{B}e(\tilde{\mathbf{i}})$ for $\tilde{\mathbf{i}}$ any lift of the cyclic word to a usual planar word. This is well-defined and functorial by the strong equivariance of $\mathbb{B}$.

Note in particular that if M is a $\tilde{T}^{\mathbf{j}}$-module, then

$$\mathring{\mathbb{B}}(\mathring{R}^{\mathbf{j}} \otimes_{\tilde{T}^{\mathbf{j}}} M) \cong \mathring{R}^{\mathbf{j}'} \otimes_{\tilde{T}^{\mathbf{j}'}} \mathbb{B}(M). \tag{6.7}$$

This is what we need to prove the composition, since we have a functorial isomorphism

$$\mathring{\mathbb{B}}(\mathring{R}^{\mathbf{j}}e(\mathbf{i})) \cong \mathring{R}^{\mathbf{j}'} \otimes_{\tilde{T}^{\mathbf{j}'}} \mathbb{B}(\tilde{T}^{\mathbf{j}}e(\tilde{\mathbf{i}})) = \mathring{R}^{\mathbf{j}'} \otimes_{\tilde{T}^{\mathbf{j}'}} \mathbb{B}_1\mathbb{B}_2(\tilde{T}^{\mathbf{j}}e(\tilde{\mathbf{i}})) = \mathring{\mathbb{B}}_1\mathring{\mathbb{B}}_2(\mathring{R}^{\mathbf{j}}e(\mathbf{i}))$$

with the last step applying (6.7) twice, to $\mathbb{B}_1$ first and then $\mathbb{B}_2$. □

6.1.4. *Comparison with Coulomb branches.* For each cylindrical loading $(\mathbf{i}, \mathbf{a})$, we have a corresponding idempotent $e(\mathbf{i}, \mathbf{a})$. As before, we assume that our order is chosen that $0 < a_1 < a_2 < \cdots < a_n < 1$.

We will now refine this definition to also assign an idempotent to $e(\mathbf{i}, \mathbf{a})$ where $\mathbf{a}$ is defined to be only weakly increasing and may have repeated values: $0 \leq a_1 \leq a_2 \leq \cdots \leq a_n < 1$.

We must apply an operation to each group $i_k, \dots, i_{k+p}$ with $a_{k-1} \neq a = a_k = a_{k+1} = \cdots = a_{k+p} \neq a_{k+p+1}$.

(1) We fix any order on the set $\mathcal{V}(\Gamma)$, and order the multi-set of indices $\{i_k, \dots, i_{k+p}\}$ in increasing order to obtain $i'_k, \dots, i'_{k+p}$. This gives us a sequence $(\mathbf{i}', \mathbf{a})$, and we will have $e(\mathbf{i}, \mathbf{a}) = e(\mathbf{i}', \mathbf{a})$.
(2) Now, we deform to $\mathbf{a}$ to be strictly increasing by defining $a'_m = a_m + m\epsilon$ for $\epsilon > 0$ very small. For ϵ sufficiently small, the isotopy class of this loading is independent of ϵ. Now $e(\mathbf{i}', \mathbf{a}')$ has a well-defined idempotent, but this is not what we wish to call $e(\mathbf{i}', \mathbf{a})$.
(3) Since the indices $i'_k, \dots, i'_{k+p}$ are ordered, all copies of an given vertex i occur together. Let $\mu_{i,a}$ be the multiplicity with which i appears here.
(4) Crossing the corresponding strands in $i'_k, \dots, i'_{k+p}$ with label i gives a homomorphism of the nilHecke algebra of rank $\mu_{i,a}$ to $e(\mathbf{i}', \mathbf{a}')\mathring{R}e(\mathbf{i}', \mathbf{a}')$. This nilHecke algebra contains a primitive idempotent projecting to the $S_{\mu_{i,a}}$ invariants in the usual polynomial representation.

Definition 6.10 *We let $e(\mathbf{i}, \mathbf{a}) \in \mathring{R}$ be the product of the images of these idempotents for all vertices i and values a.*

Note that the idempotent is independent of the choice of order of $\mathcal{V}(\Gamma)$ up to isomorphism because of we can reorder strands with different labels by (6.2b).

Definition 6.11 *Let $e(a)$ be the idempotent where we set all a_i's are equal to a single value $a \neq -\gamma_{i,j}$. In this case, the ordering of $\mathbf{i}$ is irrelevant; when constructing $\mathbf{i}'$, we will put the indices in our fixed order, and as discussed above, the order won't change the isomorphism type of the corresponding idempotent.*

One primary reason for our interest in this algebra is the following result:

Theorem 6.12 *The algebra $A_0 = e(a)\mathring{R}e(a)$ is isomorphic to the (undeformed) Coulomb branch algebra $\Bbbk[\mathfrak{M}]$ of the quiver gauge theory associated to Γ with the dimension vectors $\mathbf{v}, \mathbf{w}$.*

This is proven on page 33. Note that this theorem is independent of the choice of a, β_* and $\gamma_{*,*}$; we obtain the same algebra A_0 here regardless of these parameters. This is a generalization of [Weea, Cor. 4.13]. More generally, for each $e(\mathbf{i}, \mathbf{a})$, we have a corresponding vortex line operator, that is, an object in the category $\mathscr{B}$, and $\mathring{R}$ is the sum of the morphism spaces between these objects; see Physics Motivation 2.5.

6.2. **Change of flavor.** In this section, we discuss a modification of the structures introduced in the previous subsection by allowing β_e and $\gamma_{i,k}$ to vary with respect to the height t. To emphasize the difference with the previous section, we will make these symbols bold when they vary with respect to t.

6.2.1. *Twisted diagrams.* We now fix smooth functions $\boldsymbol{\beta}_e \colon [0,1] \to \mathbb{R}/\mathbb{Z}$ and $\boldsymbol{\gamma}_{i,k} \colon [0,1] \to \mathbb{R}/\mathbb{Z}$, and let

$$\beta_e = \boldsymbol{\beta}_e(0) \qquad \beta'_e = \boldsymbol{\beta}_e(1) \qquad \gamma_{i,k} = \boldsymbol{\gamma}_{i,k}(0) \qquad \gamma'_{i,k} = \boldsymbol{\gamma}_{i,k}(1).$$

We'll most often want $\boldsymbol{\beta}_e$ and $\boldsymbol{\gamma}_{i,k}$ to be linear, but at the moment, we can allow them to be any smooth functions. To save a great deal of notation in the section below, we interpret any equality involving the functions $\boldsymbol{\beta}_e$ to implicitly include the $\boldsymbol{\gamma}_{i,k}$'s as well.

Definition 6.13 *A* **twisted cylindrical KLRW diagram** *is a diagram satisfying all the conditions of a cylindrical KLRW diagram, but with the position of red lines and distance of ghosts strands now varying with respect to time. That is, the ghost strands are given by $\{(\bar{\pi}(t) - \boldsymbol{\beta}_e(t), t) \mid t \in [0,1]\}$ and the red lines by $\{(-\boldsymbol{\gamma}_{i,k}(t), t) \mid t \in [0,1]\}$.*

Let $B_{\boldsymbol{\beta},\boldsymbol{\gamma}}$ be the formal span of the twisted cylindrical KLRW diagrams modulo the local relations (6.2a–6.2k)

Below, we show an example of a twisted cylindrical KLRW diagram, once in our usual convention where the cylinder is cut open, and the same example in perspective on a cylinder:

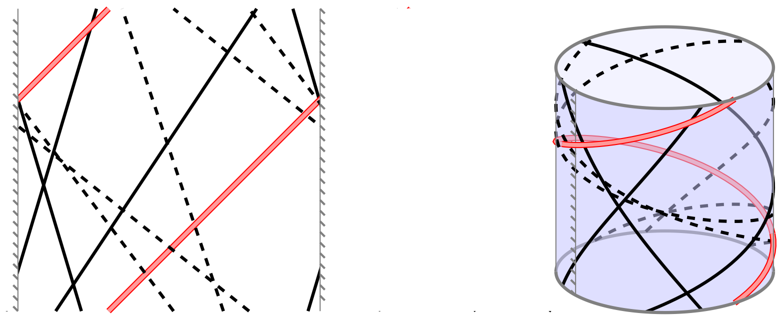

Note that we can compose twisted diagrams for functions $\boldsymbol{\beta}_e, \boldsymbol{\beta}'_e$ with $\boldsymbol{\beta}_e(1) = \boldsymbol{\beta}'_e(0)$ (and implicitly $\boldsymbol{\gamma}_{i,k}(1) = \boldsymbol{\gamma}'_{i,k}(0)$; we'll stop pointing this out now). Important special cases of this are:

(1) The module $B_{\boldsymbol{\beta},\boldsymbol{\gamma}}$ is a right module over the algebra $\mathring{R}$ for the parameters $\beta_e = \boldsymbol{\beta}_e(0)$, and a left module over the algebra for parameters $\beta'_e = \boldsymbol{\beta}_e(1)$. Note that after naively composing these diagrams, we need to apply an isotopy to get back to the original functions $\boldsymbol{\beta}_e$.

(2) If $\boldsymbol{\beta}_e(1) = \boldsymbol{\beta}_e(0)$, then we can define a function $\boldsymbol{\beta}_e^{(n)}(t) = \boldsymbol{\beta}_e(nt)$ which runs through this function n times over $[0, 1]$. Using appropriate isotopies, we have multiplications

$$B_{\boldsymbol{\beta}^{(k)},\boldsymbol{\gamma}^{(k)}} \otimes B_{\boldsymbol{\beta}^{(m)},\boldsymbol{\gamma}^{(m)}} \to B_{\boldsymbol{\beta}^{(k+m)},\boldsymbol{\gamma}^{(k+m)}}; \tag{6.8}$$

the left and right module structures discussed above are the case where k or m is 0.

6.2.2. *Relation to resolved Coulomb branches.* For fixed integers $b_e, c_{i,k}$ and parameters $\beta_e, \gamma_{i,k}$, we let

$$\boldsymbol{\beta}_e(t) = \beta_e + b_e t \qquad \boldsymbol{\gamma}_{i,k} = \gamma_{i,k} + c_{i,k} t$$

We can consider the direct sum $\mathbf{R}^{\mathbf{b},\mathbf{c}} = \bigoplus_{k\geq 0} B_{\boldsymbol{\beta}^{(k)},\boldsymbol{\gamma}^{(k)}}$; this is a $\mathbb{Z}_{\geq 0}$-graded ring with the multiplication (6.8).

Theorem 6.14 *The ring* $\mathbf{A}^{\mathbf{b},\mathbf{c}} = e(a)\mathbf{R}^{\mathbf{b},\mathbf{c}}e(a)$ *is commutative, and* $\operatorname{Proj}(\mathbf{A}^{\mathbf{b},\mathbf{c}})$ *is the partial resolution* $\tilde{\mathfrak{M}}^{\mathbf{b},\mathbf{c}}$ *defined in* [BFN].

This is a more precise statement of Theorem E(1). This is proven on page 36. If the space $\tilde{\mathfrak{M}}^{\mathbf{b},\mathbf{c}}$ is a resolution of singularities for some $\mathbf{b}, \mathbf{c}$, we call it a **BFN resolution**. By [Weeb, Thm. 5], a BFN resolution is necessarily symplectic. By [Nam08, Prop. 19], the space $\tilde{\mathfrak{M}}^{\mathbf{b}}$ is a symplectic resolution if and only if the twistor deformation for the corresponding ample line bundle is generically smooth (i.e. the Coulomb branch becomes smooth when we consider $\mathbf{b}, \mathbf{c}$ as FI parameters).

This allows us to state one of the main results of our paper:

Theorem 6.15 *If the space $\tilde{\mathfrak{M}}^{\mathbf{b},\mathbf{c}}$ is a symplectic resolution, then $\mathring{R}$ is a non-commutative symplectic resolution of singularities for any generic choice of $\beta_e, \gamma_{i,k}$ and $\mathcal{D}^b(\mathrm{Coh}(\tilde{\mathfrak{M}}^{\mathbf{b},\mathbf{c}})) \cong D^b(\mathring{R}\text{-mod})$.*

This follows from Lemma 8.10. One point we should emphasize here is that $\beta_e, \gamma_{i,k}$ is in no way related to $\mathbf{b}, \mathbf{c}$. We prove this result by showing that $\tilde{\mathfrak{M}}^{\mathbf{b}}$ possesses a tilting generator with endomorphisms given by $\mathring{R}$, which proves the desired properties.

6.2.3. *Wall crossing functors.* Let $V = \mathbb{R}^{\mathcal{E}} \times \prod_{i\in\mathcal{V}(\Gamma)} \mathbb{R}^{w_i}$.

Given a point $\beta = (\beta_e, \gamma_{i,k})$ in this space, we have a corresponding algebra $\mathring{R}_\beta$, and for any pair β, β', we have have a bimodule $B_{\boldsymbol{\beta}}$, where $\boldsymbol{\beta} = (1-t)\beta + t\beta$. Let $\bar{\boldsymbol{\beta}}(t) = \boldsymbol{\beta}(1-t)$.

We can define an equivalence relation on V by $\beta \sim \beta'$ if $B_{\boldsymbol{\beta}}$ is a Morita equivalence. Considering the bimodules $B_{\boldsymbol{\beta}}$ and $B_{\bar{\boldsymbol{\beta}}}$ as a Morita context (or "pre-equivalence data" in the terminology of [Bas68]), by [Bas68, II.3.4], this defines a Morita equivalence if and only the multiplication maps

$$B_{\boldsymbol{\beta}} \otimes B_{\bar{\boldsymbol{\beta}}} \to \mathring{R}_\beta \qquad B_{\bar{\boldsymbol{\beta}}} \otimes B_{\boldsymbol{\beta}} \to \mathring{R}_{\beta'}$$

are surjective.

Proposition 6.16 *This equivalence relation is refined by the alcoves of the hyperplane arrangement*

$$\gamma_{j,\ell} - \gamma_{i,k} + \sum_{p=1}^{n} \beta_{e_i} = n \tag{6.9}$$

for all paths

$$i = i_0 \xrightarrow{e_1} i_1 \xrightarrow{e_2} i_2 \xrightarrow{e_3} \dots \xrightarrow{e_n} i_n = j$$

in the graph Γ.

Note that we allow the "lazy path" of length 0 at some i; this is the only one which we need to consider if $\beta_e = 0$.

Proof. If there is a path from β to β' that never crosses one of these hyperplanes, then by Lemma 6.2 shows that the set of equivalence classes of loadings never changes, and so we can draw diagrams $B_{\boldsymbol{\beta}}$ and $B_{\bar{\boldsymbol{\beta}}}$ with no strands crossing, joining any loading in $\mathring{R}_\beta$ to the same loading in $\mathring{R}_{\beta'}$. Multiplying these in either order gives the desired idempotent, and shows the Morita equivalence. □

If the bimodules $B_{\boldsymbol{\beta}}$ and $B_{\bar{\boldsymbol{\beta}}}$ do not induce Morita equivalences, then the situation is more complicated.

Conjecture 6.17 *Given a planar path $\boldsymbol{\beta} = (1-t)\beta + t\beta'$, the bimodule $B_{\boldsymbol{\beta}} \overset{L}{\otimes} -$ induces an equivalence of derived categories for any ring $\Bbbk$.*

Establishing this in full generality requires more algebraic machinery than we want to develop at this moment, but we can use already known results to show it holds in the most important cases for us:

Proposition 6.18 *For each $\Gamma, \mathbf{v}, \mathbf{w}, \beta_e, \gamma_{i,k}$, Conjecture 6.17 holds for $\Bbbk = \mathbb{Q}$ (or more generally any characteristic 0 field) if the space $\tilde{\mathfrak{M}}^{\mathbf{b},\mathbf{c}}$ is a symplectic resolution for generic $\mathbf{b}, \mathbf{c}$.*

This is proven on page 36.

6.3. **Real variation of stability conditions.** Recall the definition of a real variation of stability conditions: let $\mathcal{D}$ be a finite type triangulated category and V a real vector space. Suppose that a discrete collection Σ of affine hyperplanes in V is fixed, let V^0 denote their complement. For each hyperplane in Σ consider the parallel hyperplane passing through zero, let Σ_{lin} be the set of those planar hyperplanes. Fix a component V^+ of the complement to the union of hyperplanes in Σ_{lin}. The choice of V^+ determines for each $H \in \Sigma$ the choice of the positive half-space $(V \setminus H)^+ \subset V \setminus H$, where $(V \setminus H)^+ = H + V^+$. By an *alcove* we mean a connected component of the complement to hyperplanes in Σ and we let ∇ denote the set of alcoves. For two alcoves A, $A' \in \nabla$ sharing a codimension one face which is contained in a hyperplane $H \in \Sigma$ we will say that A' is *above* A and A is below A' if $A' \in (V \setminus H)^+$.

Definition 6.19 *A* **real variation of stability conditions** *on $\mathcal{D}$ parameterized by V^0 and directed to V^+ is the data (Z, τ), where Z (the central charge) is a polynomial map $Z : V \to (K^0(\mathcal{D}) \otimes \mathbb{R})^*$, and τ is a map from ∇ to the set of bounded t-structures on $\mathcal{D}$, subject to the following conditions.*

1. *If M is a nonzero object in the heart of $\tau(A)$, $A \in \nabla$, then $\langle Z(x), [M] \rangle > 0$ for $x \in A$.*
2. *Suppose A, $A' \in \nabla$ share a codimension one face H and A' is above A. Let $\mathcal{C}$ be the heart of $\tau(A)$; for $n \in \mathbb{N}$ let $\mathcal{C}_n \subset \mathcal{C}$ be the full subcategory in $\mathcal{C}$ given by: $M \in \mathcal{C}_n$ if the polynomial function on V, $x \mapsto \langle Z(x), [M] \rangle$ has zero of order at least n on H. One can check that $\mathcal{C}_n$ is a Serre subcategory in $\mathcal{C}$, thus $\mathcal{D}_n = \{C \in \mathcal{C} \mid H^i_{\tau(A)}(C) \in \mathcal{C}_n\}$ is a thick subcategory in $\mathcal{D}$. We require that*
   - (a) *The t-structure $\tau(A')$ is compatible with the filtration by thick subcategories $\mathcal{D}_n$.*
   - (b) *The functor of shift by n sends the t-structure on $gr_n(\mathcal{D}) = \mathcal{D}_n/\mathcal{D}_{n+1}$ induced by $\tau(A)$ to that induced by $\tau(A')$. In other words,*

$$gr_n(\mathcal{C}') = gr_n(\mathcal{C})[n]$$

where $\mathcal{C}'$ is the heart of $\tau(A')$, $gr_n = \mathcal{C}'_n/\mathcal{C}'_{n+1}$, $\mathcal{C}'_n = \mathcal{C}' \cap \mathcal{C}_n$.

In our context, $V = \mathbb{R}^{\mathcal{E}} \times \prod_{i\in\mathcal{V}(\Gamma)} \mathbb{R}^{w_i}$, and Σ is the set of hyperplanes defined by the conditions (6.9), and as defined above, ∇ is the set of alcoves for this arrangement. We fix a generic integral element χ in $\mathbb{Z}^{\mathcal{E}} \times \prod_{i\in\mathcal{V}(\Gamma)} \mathbb{Z}^{w_i}$, which we can also interpret as a cocharacter into Q; we define the positive side each hyperplane as the side any ray parallel to ν points toward.

Now, let us define a real variation of stability conditions for each alcove C_0 in ∇. First, for each alcove C, we let $\mathring{R}_C$ be the cylindrical KLRW algebra corresponding to parameters in that chamber.

Given C, we define an equivalence $\mathbb{B}_{C,C_0}\colon D^b(\mathring{R}_C\text{-fdmod}) \to D^b(\mathring{R}_{C_0}\text{-fdmod})$ by the rule that for any pair of chambers $C^{\pm}$ on the positive/negative sides of a hyperplane (i.e. C^+ is above C^-), the equivalences are related by

$$\mathbb{B}_{C^-,C_0}(M) = \mathbb{B}_{C^+,C_0}(B_{\boldsymbol{\beta}} \overset{L}{\otimes} M),$$

where $\boldsymbol{\beta}$ is a path from the interior of C^- to the interior of C^+.

Definition 6.20 *Let τ be the map sending the alcove C to the image in $D^b(\mathring{R}_{C_0}\text{-fdmod})$ of the standard t-structure of $D^b(\mathring{R}_C\text{-fdmod})$ under the derived equivalence $\mathbb{B}_{C,C_0}$. In particular, it sends C_0 to the standard t-structure on $\mathring{R}$-modules.*

Now, we need to define the central charge Z. Since this depends polynomially on V, we need only define it on C_0. Let $\mathrm{cyc}_n \subset (\mathbb{R}/\mathbb{Z})^n$ be the set of cyclically ordered elements. For a generic element $\mathbf{a}$ of this set, and a choice of $(i_1, \dots, i_n)$, we have idempotent. Note that the associated idempotent depends on the choice of β, γ, since these define the positions of red lines and distance of ghosts. We then define the central charge by the formula:

$$Z_{\beta,\gamma}(M) = \sum_{\mathbf{i}\in I^n_{\mathbf{v}}} \int_{\mathbf{a}\in\mathrm{cyc}_n} \dim e(\mathbf{a},\mathbf{i})M. \tag{6.10}$$

This is a weighted sum of the dimensions of the images of idempotents, weighted by the volume of the corresponding chamber.

This is a combinatorial version of the central charge function $\mathcal{Z}^0$ defined in [Aga, (4.7)], as we will explain in more detail in Remark 9.3.

Note that the alcove C_0 is arbitrary, so we can define this central charge function on $K^0(\mathring{R}_C\text{-fdmod})$ for any alcove using the formula (6.10) on the alcove C; note that for any integral ν, the translation $C_0+\nu$ is another alcove with $\mathring{R}_{C_0} \cong \mathring{R}_{C_0+\nu}$, but the corresponding functions of the K-group differ by translation in V. Of course, we can reasonably ask how to match these functions for different alcoves:

Lemma 6.21 *Using the isomorphism $[\mathbb{B}_{C,C_0}]\colon K^0(\mathring{R}_C\text{-fdmod}) \cong K^0(\mathring{R}_{C_0}\text{-fdmod})$ to identify K-groups, the functions $Z_{\beta,\gamma}(M)$ match.*

This is proven on page 39.

Theorem 6.22 *The data (Z, τ) define a real variation of t-structures.*

This is proven on page 41.

One structure we need to understand when studying this real variation of t-structures is the structure of the categories $\mathcal{C}_n$ corresponding to a given wall. Consider a wall of C_0, defined by an equality of the form (6.9) for some path (possibly lazy). Let β, γ denote a choice of parameters in the interior of C_0, and β', γ' a generic point on the wall. In the case where $\beta_e = 0$, this wall must be of the form $\gamma'_{i,k} = \gamma'_{j,\ell}$.

When we follow a path from the interior of C_0 to this wall, each equivalence class of loadings is defined by a set of inequalities depending on β, γ. The closure of this set of loadings becomes a face of some codimension d as we reach the wall.

This codimension d is the number of separate groupings of strands that trace out the circuit path from (6.9). In particular, if $\beta_e = 0$, it is just the number of black strands between the red strands on corresponding to $\gamma'_{i,k}$ and $\gamma'_{j,\ell}$. We call this the **codimension** of the equivalence class (with respect to the wall H).

Thus, the objects in the subcategory $\mathcal{C}_n$ defined with respect to C_0 and H are those killed by all idempotents of codimension $< n$.

7. The tangle action

7.1. **Affine braids.** Throughout this section, we assume that Γ is an ADE quiver and that $\beta_e = 0$; by Lemma 6.6, we lose no generality by assuming this. Recall that the operation of dualizing fundamental representations induces an involution of Γ, which coincides with the action of $-w_0$ on fundamental representations; this is the unique nontrivial diagram automorphism for A_n, D_{2n+1} and E_6, and trivial for D_{2n} and E_7, E_8. We denote this involution by $j \mapsto j^*$.

We fix dimension vectors $\mathbf{v}, \mathbf{w}$ as usual, and let $\ell = \sum w_i$. We let $\mathring{R}^{\mathbf{j}}\mathring{R}^{\mathbf{j}}$ be the cylindrical KLRW algebra where the labels on the red strands, reading from 0 to 1, are given by the ℓ-tuple $\mathbf{j} = (j_1, \dots, j_\ell)$

Let $\widehat{B}_\ell$ be the extended braid group of affine type A acting as usual on the set of ℓ-tuples with w_i appearances of i. This is the group generated by elements $s_0, \dots, s_{\ell-1}, \sigma$ with relations:

$$s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1} \qquad s_i s_j = s_j s_i \quad |i-j| > 1$$
$$s_{i+1}\sigma = \sigma s_i$$

Let $\Sigma_{\mathbf{w}}$ be the set of sequences in $\mathcal{V}(\Gamma)$ where i appears in w_i many times. The group $\widehat{W}$ acts on $\Sigma_{\mathbf{w}}$ with s_i acting by transposition of the i and $i+1$st entries and σ by by the cyclic permutation $\sigma \cdot (j_1, \dots, j_\ell) = (j_\ell, j_1, \dots, j_{\ell-1})$.

Definition 7.1 *The* **affine braid groupoid** *is the action groupoid for the action of $\widehat{B}_\ell$ on $\Sigma_{\mathbf{w}}$.*

The notation $\mathring{R}^{\mathbf{j}}$ is useful, since the algebra does not depend up to isomorphism on the position of the red lines, but it will be useful for us to fix parameters $\gamma_{i,k}$, so they are evenly spaced around the circle; that is, we define k_m so that j_m is the k_m appearance of j_m reading from the start of the word $\mathbf{j}$, and let $\gamma_{j_m,k_m} = \frac{m}{\ell} - \frac{1}{2\ell}$. For each i, we define a linear path for $t \in [0,1]$ that swaps the ith and ist parameters, that is,

$$\boldsymbol{\gamma}_{j_m,k_m} = \frac{m + \delta_{i,m}t - \delta_{i+1,m}t}{\ell} - \frac{1}{2\ell}.$$

Similarly, to σ we associate the path rotating one "click" around the cylinder:

$$\boldsymbol{\gamma}_{j_m,k_m} = \frac{m+t}{\ell} - \frac{1}{2\ell}.$$

These are maybe easier to visualize in terms of the path traced by the red strands:

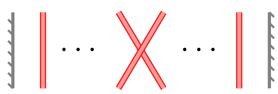

s_i

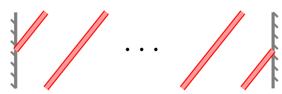

σ

Consider the functors of tensor product with the bimodule $B_{\boldsymbol{\gamma}}$ corresponding to these paths:

$$\mathbb{B}_i \colon D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod}) \to D^b(\mathring{R}^{s_i\mathbf{j}}\text{-fdmod}) \qquad \mathbb{B}_\sigma \colon D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod}) \to D^b(\mathring{R}^{\sigma\mathbf{j}}\text{-fdmod}).$$

Theorem 7.2 *The functors $\mathbb{B}_i$ and $\mathbb{B}_\sigma$ define an affine braid groupoid action on the categories $D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod})$.*

This is proven on page 42.

7.2. **Cups and caps.** The functor $\mathbb{B}_i$ corresponds to a path through a wall where the i-th and $(i+1)$-st red strands cross. Associated to the passage to this wall, we have a filtration of the derived category $D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod})$ by subcategories $\mathcal{D}_m$; let e_m be the sum of all idempotents corresponding to loadings with $< m$ black strands between these red strands. As is discussed earlier, we have:

Lemma 7.3 *The subcategory $\mathcal{D}_m$ is the subcategory of complexes $M^\bullet$ such that the complex of vector spaces $e_m M^\bullet$ is exact.*

Proof. For a chamber representing an equivalence class of cylindrical loadings, its volume depends polynomially on the choice of $\gamma_{i,k}$'s, and as two of the strands come together, the order of vanishing of this polynomial is the number of black strands between the two reds which are coming together. Thus, the order of vanishing of $Z_{\beta,\gamma}(M)$ for a module M less than m if and only if $e_m M \neq 0$, and taking contrapositive, $M \in \mathcal{C}_m$ if and only if $e_m M = 0$. The result follows immediately. $\square$

It follows immediately that:

Lemma 7.4 *The quotient $\mathcal{C}_m/\mathcal{C}_{m+1}$ is isomorphic to the representations of the ring $R^{(m)} = e_{m+1}\mathring{R}e_{m+1}/e_{m+1}\mathring{R}e_m\mathring{R}e_{m+1}$.*

We'll be particularly interested in the case where the consecutive strands have labels j, j^*. Without loss of generality, we can assume these are the labels $j_{\ell-1}, j_\ell$. Let $\mathbf{j}' = (j_1, \dots, j_{\ell-2})$. In this case, the deepest level of the filtration is $\mathcal{D}_z$ where

$$z = \rho^\vee(\varpi_j + \varpi_{j^*}) = \rho^\vee(\varpi_j - w_0\varpi_j) = 2\rho^\vee(\varpi_j).$$

Let $\varpi_j + \varpi_{j^*} = \sum_{i\in\mathcal{V}(\Gamma)} z_i\alpha_i$. Note that $z = \sum z_i$.

Lemma 7.5 *The ring $R^{(z)}$ in this case is Morita equivalent to the algebra $\mathring{R}^{\mathbf{j}'}$ with dimension vector $\mathbf{v}' = \mathbf{v} - \mathbf{z}$.*

This is proven on page 42. This defines a functor $\cup_\ell \colon \mathring{R}^{\mathbf{j}'}\text{-mod} \to \mathring{R}^{\mathbf{j}}\text{-mod}$, which is exact; we can think of this as $\mathring{R}^{\mathbf{j}}e_{m+1}/\mathring{R}^{\mathbf{j}}e_m\mathring{R}^{\mathbf{j}}e_{m+1}\otimes_{\mathring{R}^{\mathbf{j}'}}$. We can try to define a left adjoint $\cup_\ell^*$ to this functor by derived tensor product with $e_{m+1}\mathring{R}^{\mathbf{j}}/e_{m+1}\mathring{R}^{\mathbf{j}}e_m\mathring{R}^{\mathbf{j}}$. We say "try" above because there's no guarantee that this functor will preserve being a bounded complex.

Recall that we call a fundamental weight ϖ_j **minuscule** if all of the non-zero weight spaces of its corresponding representation are extremal. This holds for all fundamental weights in type A, for the vector and both spin representations in type D, two representations for E_6 and one for E_7. For simplicity, we call $i \in \mathcal{V}(\Gamma)$ minuscule if the corresponding fundamental representation is minuscule, and let $\mathcal{V}(\Gamma)_{\min}$ be the subset of minuscule elements of $\mathcal{V}(\Gamma)$.

Lemma 7.6 *We have a well-defined functor $\cup_\ell^* \colon D^b(\mathring{R}^{\mathbf{j}}\text{-fdmod}) \to D^b(\mathring{R}^{\mathbf{j}'}\text{-fdmod})$ if $j \in \mathcal{V}(\Gamma)_{\min}$.*

Proof. This is functor is well-defined if and only if $\mathring{R}^{\mathbf{j}}/\mathring{R}^{\mathbf{j}}e_z\mathring{R}^{\mathbf{j}}$ has a finite resolution as a module over $\mathring{R}^{\mathbf{j}}$. This follows from the fact that $\mathring{R}^{\mathbf{j}}$ has finite global dimension, since it is an NCCR, by Theorem D. □

Remark 7.7. This result holds in effectively no non-minuscule cases. As in [Web17a], we could consider a category of complexes which are bounded above or below with weaker finiteness properties and use non-minuscule weights; for simplicity we don't work through the details of this. The interested reader can infer them from [Web17a].

Definition 7.8 *Let $\cap_\ell = \cup_\ell^*[-m](-m)$; we remind the reader that $[-m](-m)$ means that we shift degree by increasing the internal and homological degree of any element by m. Note that this is the opposite of the shift in* [Web17a, Def. 7.4], *because we are defining our functor using a tensor product. We also have $\cap_\ell = \cup_\ell^![m](m)$ where $\cup_\ell^!$ is the right adjoint to $\cup_\ell$ by* [Web17a, Thm. 8.11].

The cup and cap functors $\cup_k$ and $\cap_k$ are obtained from rotation in the cylinder so that there are k red strands left of the two red strands which meet (i.e. they are the $(k+1)$st and $(k+2)$nd when reading around the circle from 0). We can write this in terms of the equivalence $\mathbb{B}_\sigma$, which precisely acts by twisting the cylinder.

Draw the image of 1 in $\mathring{R}^{\mathbf{j}}e_{m+1}/\mathring{R}^{\mathbf{j}}e_m\mathring{R}^{\mathbf{j}}e_{m+1}$ as a cup where the two red strands meet a point carrying an element of L_0, the unique simple over $\tilde{T}^{j,j^*}$ killed by e_m, constructed in [Web17a, Lem. 7.3] (denoted L_{ϖ_j} there). We'll fix a generator of L_0 in each case, and just draw this generator as two red strands meeting at the bottom of a cup, with the black strands coming out of the minimum. These are annular versions of the cup and cap bimodules in [Web17a, §7.3]. In type A, these can be seen as a special case of the the Y-ladders with $c = n$ introduced in [MW18, 3.1].

7.3. **Annular tangles.** Let $\mathbb{A}$ be the annulus $\mathbb{R}^2 \setminus \{(0,0)\}$. Let $\pi\colon \mathbb{A} \to S^1$ be the obvious projection along rays; we'll use the same symbol to denote the induced map $\mathbb{A} \times [0,1] \to S^1 \times [0,1]$.

Definition 7.9 *A* **oriented ribbon annular tangle** *T is a framed tangle in $\mathbb{A} \times [0,1]$, that is a 1-dimensional oriented submanifold with boundary lying in $S^1 \times \{0,1\}$. As usual, we number the boundary points of the tangle by their cyclic order around $S^1 \cong \mathbb{R}/\mathbb{Z}$, starting at 0; as usual, we consider these up to isotopy avoiding 0. A* **projection** *of an annular tangle T is the image of T under π in $S^1 \times [0,1]$ when it is isotoped so that the projection map is an immersion on T, and any point in $S^1 \times [0,1]$ has at most two pre-images in T, whose images cross transversely (i.e. we avoid triple points and tangencies). We account for the framing on tangles by only using projections with the "blackboard framing" i.e. where the derivative of the projection always induces an isomorphism between the framing space and the tangent space of $S^1 \times [0,1]$.*

As always when considering tangles in a thickened surface, we have that any two projections for isotopic ribbon tangles are related by a finite chain of isotopies and Reidemeister moves II and III, as well as cancelling pairs of Reidemeister I moves which preserve the blackboard framing; a single Reidemeister I move will not.

We consider labelings of the components of an oriented ribbon tangle with elements $j \in \mathcal{V}(\Gamma)_{\min}$ (which we think of as the corresponding fundamental representation). As usual, these induce a labeling of the boundary of T, where we use the same element of $\mathcal{V}(\Gamma)$ if the orientation on T matches the upward orientation of $[0,1]$ under projection, and the "dual" j^* if the orientations are opposite.

Definition 7.10

Let Tang be the category such that:

- objects are finite subsets of S^1 labeled with elements of $\mathcal{V}(\Gamma)_{\min}$,

- morphisms $S \to S'$ are annular ribbon tangles with boundary in $\mathbb{A} \times \{0\}$ given by S and in $\mathbb{A} \times \{1\}$ by S'.
- composition is just stacking of tangles (followed by appropriate isotopy).

We let the list $(j_1, \dots, j_\ell)$ denote any fixed set where this is the list of labels.

We have an obvious functor from the affine braid groupoid to Tang giving the tangles with no minima or maxima. In order to generate all tangles, we need only add cup and cap functors joining two adjacent points, which we also denote $\cup_k, \cap_k$ when the cup and cap use the $(k+1)$st and $(k+2)$nd when reading around the circle from 0; of course, there are many different versions of these functors with different labels $j \in \mathcal{V}(\Gamma)_{\mathrm{min}}$, but we leave these implicit. Consider two lists $\mathbf{j} \in \mathcal{V}(\Gamma)^s_{\mathrm{min}}$ and $\mathbf{j}' \in \mathcal{V}(\Gamma)^{s'}_{\mathrm{min}}$; these give corresponding dimension vectors $\mathbf{w}$ and $\mathbf{w}'$ where w_i is the number of k with $j_k = i$, and similarly with w'_i. Fix a vector $\mathbf{v}$, and let $\mathbf{v}' = \mathbf{v} + C^{-1}(\mathbf{w}' - \mathbf{w})$ where C is the Cartan matrix of Γ; of course, this is not necessarily integral, but it will be if there is a morphism $T \colon \mathbf{j} \to \mathbf{j}'$.

Theorem 7.11 *For each morphism $T \colon \mathbf{j} \to \mathbf{j}'$ in Tang, there is an associated functor $\Phi(T) \colon D^b(\mathring{R}^{\mathbf{j}}_{\mathbf{v}}\text{-fdmod}) \to D^b(\mathring{R}^{\mathbf{j}'}_{\mathbf{v}'}\text{-fdmod})$ satisfying $\Phi(T_1 \circ T_2) \cong \Phi(T_1) \circ \Phi(T_2)$ such that $\Phi(\tau)$ for an affine braid τ is the wall-crossing functor $\mathbb{B}_\tau$, and for a cups and caps, we have $\Phi(\cup_k) = \cup_k$ and $\Phi(\cap_k) = \cap_k$.*

This is proven on page 43. You can think of this as defining a functor from Tang to the category whose objects are triangulated categories, and whose morphisms are exact functors up to isomorphism.

Of course, this assigns a finite dimensional bigraded vector space to any annular knot or link with components labeled by minuscule representations by considering this link as a tangle and applying it to $\mathring{R}^{\emptyset}_{\mathbf{0}}$.

Theorem 7.12 *If $K \subset B^3 \subset \mathbb{R}^3$ is any link, and we embed K as an annular link via any embedding $B^3 \hookrightarrow \mathbb{A} \times [0,1]$ (all such embeddings are isotopic), the invariant $\Phi(K)$ is the same as the invariant $\Phi_{\mathsf{L}}(K)$ defined in* [Web17a, §8.1] *for the same labeling. In particular,*

(1) If Γ is of type A_1, then $\Phi(K)$ coincides with Khovanov homology.
(2) If Γ is of type A_n, then $\Phi(K)$ coincides with Khovanov-Rozansky $\mathfrak{sl}_{n+1}$-homology.

This is proven on page 43. A more comprehensive description of all the different manifestations of this invariant in type A is given in [MW18, Th. A].

However, constructing an invariant of knots in $\mathbb{R}^3$ is only using a small portion of the power of this construction: we obtain an invariant for each annular link, which will depend on how the link wraps around the origin. Of course, in type A, invariants of this type are well-known: annular Khovanov-Rozansky homology. This is defined by Queffelec and Rose in [QR18].

It also seems natural to compare with the following construction: given an annular link K, we can cut along the plane over the positive x-axis in $\mathbb{A}$ to obtain a usual tangle K' in $\mathbb{R}^2 \times [0,1]$. This has an associated complex of $T^{\underline{\boldsymbol{\lambda}}}$-$T^{\underline{\boldsymbol{\lambda}}}$-bimodules $\Phi_{\mathsf{L}}(K')$ over the (planar) KLRW algebras defined by the construction of [Web17a, §8.1]. The Hochschild homology $HH_{T^{\underline{\boldsymbol{\lambda}}}}(\Phi_{\mathsf{L}}(K'))$ is easily seen to be an annular link invariant:

Conjecture 7.13 *The following invariants of annular knots labeled with minuscule representations coincide:*

(1) the invariant $\Phi(K)$ constructed above;
(2) the Hochschild homology $HH_{T^{\underline{\boldsymbol{\lambda}}}}(\Phi_{\mathsf{L}}(K'))$;
(3) if $\Gamma = A_n$, the annular Khovanov-Rozansky homology of [QR18].

We don't expect this conjecture to be exceptionally difficult; the equivalence of (1-3) in type A should be approachable by rephrasing constructions (1) and (2) using actions of foams as in [MW18], and showing that (1) and (2) are both defined by annular evaluation of foams. The main difficulty here is showing that a single essential circle on the annulus evaluates to the corresponding representation over the Lie algebra $\mathfrak{g}_\Gamma$. This is effectively just the observation that $T^{\underline{\boldsymbol{\lambda}}}$ has finite global dimension in situation (2), since this means that higher Hochschild homology of $T^{\underline{\boldsymbol{\lambda}}}$ vanishes, and in degree 0 it matches the Grothendieck group. For (1), this not obvious for $n > 2$, and requires a rather complex calculation. It seems promising to think of the equivalence of (1) and (2) as a generalization of Queffelec and Rose's comparison of the horizontal and vertical traces in [QR18, §3], but we have not made much progress on making this precise.

8. Diagrams for quiver gauge theories

In order to establish the results of Sections 6 and 7, we need to describe how the results of [Webc, Weba] can be interpreted in the quiver case. In those papers, we explain how the quantum Coulomb branch $\mathcal{A}^{\mathrm{sph}}$ can be written as an endomorphism ring in a larger category $\mathscr{B}$, which we call the *extended BFN category.* This larger category is more easily presented and more amenable to algebraic methods; this will allow us to make the connection between Coulomb branches and cylindrical KLR algebras.

The definition of $\mathscr{B}$ depends on a parameter δ. For simplicity, we'll assume throughout this paper that $\delta = \frac{1}{2}$ (and so in pth root conventions, we have $\delta = \frac{1}{2p}$).

8.1. **Unrolled diagrams.**

In the case of a quiver gauge theory, the extended BFN category $\mathscr{B}$ has a more graphical description.

Recall that this is a category whose objects are elements of $\mathfrak{t}_{1,\tilde{H}} \cong \mathfrak{t}_{\tilde{H}}$, the (real) Lie algebra of the maximal torus of the normalizer $H = N^\circ_{GL(V)}(G)$. As discussed earlier, this is generated by the diagonal matrices in $\mathfrak{gl}(\mathbb{R}^{v_i}), \mathfrak{gl}(\mathbb{R}^{w_i}), \mathfrak{gl}(\mathbb{R}^{\epsilon_{i,j}})$; thus, each object in this category can be represented by choosing diagonal entries. We let $\{z_{i,k}\}_{i\in\mathcal{V}(\Gamma),k=1,\dots,v_i}$ be the diagonal entries in $\mathfrak{gl}(v_i)$; for slightly complicated reasons, we let $\{\gamma_{i,k}\}_{i\in\mathcal{V}(\Gamma),k=1,\dots,w_i}$ and $\{\beta_e\}_{e\colon i\to j}$ be the appropriate diagonal entries in $\mathfrak{gl}(w_i)$ and $\mathfrak{gl}(\epsilon_{i,j})$ plus $\frac{1}{2}$. This is to cancel the shift in the definition of φ_i^{mid}; in particular, the element τ corresponds to $z_{i,k}=0$ and $\gamma_{i,k}=\beta_e=\frac{1}{2}$. Readers familiar with this field are presumably used to odd shifts of $\frac{1}{2}$ appearing and disappearing. This representation isn't unique; you get the same element of the Lie algebra $\mathfrak{t}_{\tilde{H}}$ if you shift by

$$z'_{i,k}=z_{i,k}+\eta_i \qquad \gamma'_{i,k}=\gamma_{i,k}-\eta_i \qquad \beta'_e=\beta_e-\eta_{h(e)}+\eta_{t(e)}.$$

Note the similarity to (6.3).

In these terms, the unrolled arrangements defined in Section 2.2 are given by the unrolled root hyperplanes $\{\alpha(\eta)=n\mid n\in\mathbb{Z}\}$ of the form:

$$z_{i,k}-z_{i,m}=n \qquad \text{for all } k\neq m\in[1,v_i], n\in\mathbb{Z}, \tag{8.1a}$$

and the unrolled matter hyperplanes $\{\varphi_i^{\mathrm{mid}}(\eta)=n\mid n\in\mathbb{Z}\}$ of the form:

(8.1b)
$$z_{j,k}-z_{i,m}+\beta_e=n \qquad \text{for all edges } i\to j,\ \text{for all } k\in[1,v_j], m\in[1,v_i], n\in\mathbb{Z}$$

$$z_{i,m}+\gamma_{i,k}=n \qquad \text{for all } i\in\mathcal{V}(\Gamma), m\in[1,v_i], n\in\mathbb{Z} \tag{8.1c}$$

Note that unlike in Section 6, our parameters $\gamma_{*,*}$ and β_* are real numbers, but if we change them by an integer amount, we will not change the set of hyperplanes (8.1b–8.1c).

Fix a choice of $\gamma_{*,*}$ and β_*; let $\mathscr{B}_{\beta,\gamma}$ be the subcategory $\mathscr{B}$ where we only consider objects in the coset $\mathfrak{t}_{\beta,\gamma}$ of this choice under $\mathfrak{t}_{\mathbb{R}}$, the Cartan of $\mathfrak{gl}(\mathbb{R}^{v_i})$. For simplicity, we assume that $\gamma_{*,*}$ and β_* are generic in the same sense discussed in Section 6. We can describe the objects in this category just by the values $z_{*,*}\in\mathbb{R}$.

The category $\mathscr{B}_{\beta,\gamma}$ also depends on a flavor ϕ, which we can take to be an element of the Lie algebra of $\mathrm{Aut}_G(V)$, and thus is induced by a regular element of $\prod_i \mathfrak{gl}(\mathbb{k}^{w_i})\times\prod_{(i,j)\in\mathcal{V}(\Gamma)^2}\mathfrak{gl}(\mathbb{k}^{\chi_{i,j}})$. We can assume that this cocharacter lands in the usual torus of diagonal matrices; let $c_{i,1},\dots,c_{i,w_i}$ be the diagonal entries of its components into $\mathfrak{gl}(\mathbb{k}^{w_i})$ and b_e for each edge e the weights of its components into $\mathfrak{gl}(\mathbb{k}^{\chi_{i,j}})$. Let us just emphasize that while b_* and $c_{*,*}$ live in a related space to β_* and $\gamma_{*,*}$, the latter parameters are always chosen over $\mathbb{R}$, whereas the former are defined over the base field (which we will often want to think of as $\mathbb{C}$ or $\mathbb{F}_p$) and they are chosen independently from β_* and $\gamma_{*,*}$. We can think of b_* and $c_{*,*}$ either as numerical values or as formal variables.

Thus, we can decompose a path $\pi\colon[0,1]\to\mathfrak{t}_{\beta,\gamma}$ into a v_i-tuple of paths $\pi_{i,k}$ for each $i\in\mathcal{V}(\Gamma)$. We can visualize this by superimposing the graphs of these paths (though we will use the opposite convention from calculus class, using the y-axis for the independent variable and the x-axis for the dependent). That is, we consider the path $t\mapsto(\pi_{i,k}(t),t)$ landing in $\mathbb{R}\times[0,1]$.

We cross root hyperplanes when two of these paths for the same element of $\mathcal{V}(\Gamma)$ are integer distance from each other, and matter hyperplanes when the x-value of a strand with label i is an integer minus $\gamma_{i,j}$ or the difference between two such strands is an integer plus β_e. It's thus convenient to label the points $(\pi_{i,k}(t)+n,t)$ for $n\in\mathbb{Z}$ with "partner" strands and those for $n\in\mathbb{Z}-\beta_e$ for each edge $e\colon j\to i$ with "ghost" strands so that we can see when these crossings occur. We'll label the strand corresponding to n with $i;n$. We'll draw partners as solid lines, and ghosts as dashed lines. We'll also draw in dashed lines at $x=n$ for $n\in\mathbb{Z}-\gamma_{i,j}$, which we'll label with $\infty;n$.

There's an obvious action of the affine Weyl group $\widehat{W}$ on the set of such paths with the finite Weyl group $\prod_{i\in\mathcal{V}(\Gamma)} S_{v_i}$ acting on paths by permutation of the second indices in $\pi_{i,k}$, and the coweight lattice acting by translations $\pi_{i.k}(t)\mapsto\pi_{i,k}(t)+n_{i,k}$ for integers $n_{i,k}$. This leaves the collections of the original curves and their partners unchanged, just changing the indices and which curves are partners, and which are originals. Thus, we can visualize the action of an affine Weyl group element by changing the labels accordingly at some fixed value of $y=a$. The equations (2.6e) and (2.6i) assure that the result does not depend on the value of a.

We can also visualize the action of S_h by identifying the weights $\epsilon_{i,k}$ with a dot on the corresponding path, at the top for multiplying at the left and at the bottom for multiplication on the right. It will be useful to also draw dots on partner strands with label $i;n$, which we will associate to $\epsilon_{i,k}+nh$. This likewise assures that inserting an element of the affine Weyl group above or below a dot will give the same answer, by equation (2.6a).

Thus, the morphisms in $\mathscr{B}_{\beta,\gamma}$ can all be expressed as one of these diagrams. More precisely:

Definition 8.1 *An* **unrolled diagram** *is a collection of paths in $\mathbb{R}\times[0,1]$ of the form $\{(\pi_{i,k}(t)+n,t)\mid t\in[0,1]\}$ for $n\in\mathbb{Z}$ for some piecewise smooth map $\pi_{i,k}\colon[0,1]\to\mathbb{R}$. At a finite number of values $t_1,\dots,t_k$, we apply elements $w_1,\dots,w_k\in\widehat{W}$ to the labels on the curves and their partners, which can create a discontinuity in the function $\pi_{i,k}$, but we assume that the resulting curves are still smooth and simply change labeling. We'll draw one of these changes as a squiggly green line.*

We also add ghost strands at $\{(\pi_{i,k}(t)+n-\beta_e,t)\mid t\in[0,1]\}$ for each edge $e\colon j\to i$ and at $\{(n-\gamma_{i,k},t)\mid t\in[0,1]\}$ for all $n\in\mathbb{Z}$, which we draw as dashed. Each curve is decorated with finitely many dots.

The diagram must be locally of the form

(8.2) 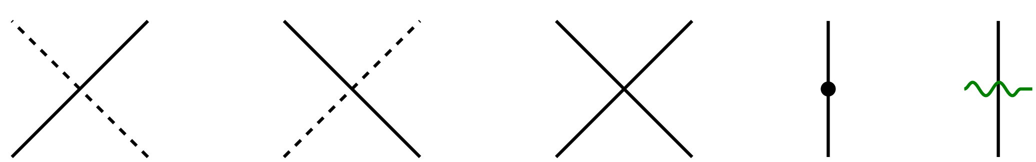

That is, there are no tangencies, triple crossings or dots on crossings. The curves (including ghosts) must meet the lines at $y=0$ and $y=1$ at distinct points. We consider these diagrams up to isotopy preserving the conditions above.

We'll draw an example below; lacking infinitely wide paper, we can only draw part of the diagram. In the interest of comprehensibility, we draw an example with $\mathbf{w} = 0$, so there are no vertical strands.

(8.3)

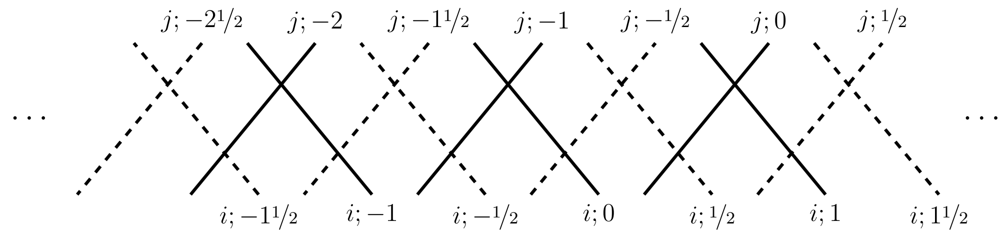

Let the **bottom** $D(0)$ of a diagram defined by the path π be the object in $\mathscr{B}_{\beta,\gamma}$ defined $z_{i,k} = \pi_{i,k}(0)$, and the **top** $D(1)$ the object defined by $z_{i,k} = \pi_{i,k}(1)$.

To an unrolled diagram D, we can associate a morphism $\mathbb{r}_D \colon D(0) \to D(1)$ from the bottom of D to its top as follows:

Definition 8.2

(1) *Given an unrolled diagram D with no dots and no elements of $\widehat{W}$, let $\mathbb{r}_D$ be the morphism $\mathbb{r}_\pi$ from* [Webc, Def. 3.12] *for the corresponding path.*

(2) *Given an unrolled diagram with all $\pi_{i,k}$'s constant, and one dot on the strand at $(\pi_{i,k} + n, t)$, we let $\mathbb{r}_D$ be multiplication by $\epsilon_{i,k} + nh$.*

(3) *Given an unrolled diagram with no dots, all strands vertical, and a single relabeling by $w \in \widehat{W}$, we let $\mathbb{r}_D = y_w$.*

Any other diagram can be written as a composition of these, and we define $\mathbb{r}_D$ to be the composition of the corresponding morphisms.

This associates a morphism to a diagram, and since we hit all the generators of the category $\mathscr{B}_{\beta,\gamma}$, every morphism is a sum of $\mathbb{r}_D$'s; this is clear from [Webc, Cor. 3.13]. However, this description is redundant, since there are relations between these generators.

All the relations of the category $\mathscr{B}_{\beta,\gamma}$ can be described locally in terms of unrolled diagrams. We can visualize Definition 8.2(2) as a convention of "dot migration" for interpreting dots on partner strands:

(8.4a)

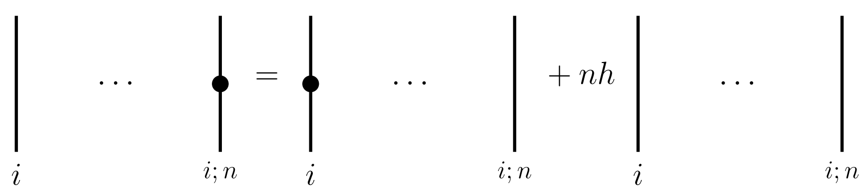

The fact that isotopy leaves $\mathbb{r}_D$ invariant is a combination of relations (2.6a) and the "boring" cases of (2.6c,2.6h,2.6j,2.6k) where the portions of diagrams commuting past each other are distant in $\mathbb{R}$. As discussed before, (2.6b,2.6e,2.6f,2.6i) imply that green lines can isotope past all crossings and dots, and (2.6d) that green lines can merge by multiplying their labels.

(8.4b)
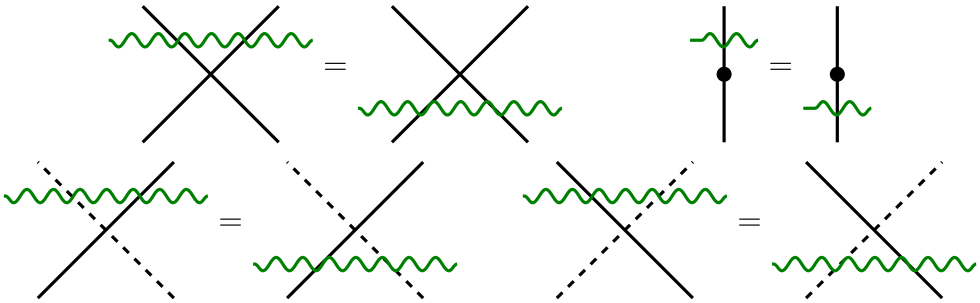

The relation (2.6a) implies that dots can commute past strands with a different label or their partners, or past all ghosts:

(8.4c)
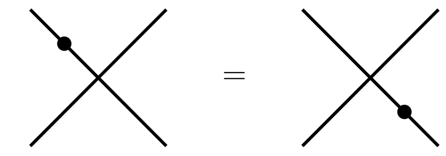

unless $i = j$

(8.4d)
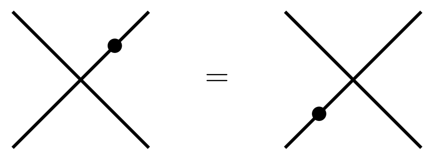

unless $i = j$

(8.4e)


The previous relations have not depended on the weights $c_{i,1}, \dots, c_{i,w_i}$ and b_e defined before. Let

(8.4f)

$$p_{i,m}(u) = \prod_{\substack{k \text{ s. t.} \\ \gamma_{i,k} - m \in \mathbb{Z}}} (u + (c_{i,k} - m + \gamma_{i,k})h) \in \Bbbk[u, c_{i,*}, h]$$

(8.4g)

$$q_{ij,m}(u) = \prod_{\substack{e \colon j \to i \text{ s. t.} \\ \beta_e - m \in \mathbb{Z}}} (u + (b_e - \beta_e + m)h) \cdot \prod_{\substack{e \colon i \to j \text{ s. t.} \\ \beta_e + m \in \mathbb{Z}}} (-u + (b_e - \beta_e - m)h) \in \Bbbk[u, \beta_*, h].$$

Note that $p_{i,m}(u) = p_{i,m+n}(u + nh)$ for all $n \in \mathbb{Z}$, and $q_{ij,m}(u) = q_{ij,m+n}(u - nh)$.

If $\gamma_{*,*}$ and β_* satisfy the genericity assumptions of Section 6, then these functions always have degree ≤ 1 in u, but we allow somewhat less generic parameters here, just assuming that $\beta_e \neq 0$ for all e. It might seem strange to see $\gamma_{*,*}$ and β_* since these are in $\mathbb{R}$, but if the difference or sum of one of these with m is an integer, we can consider that integer as an element of 𝕜. We can now use these to write out our version of the relation (2.6c):

(8.4h)

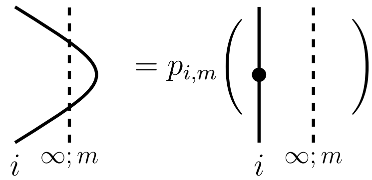

(8.4i)

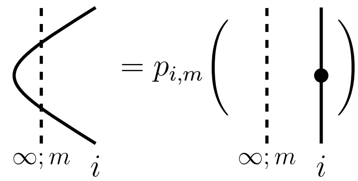

(8.4j)

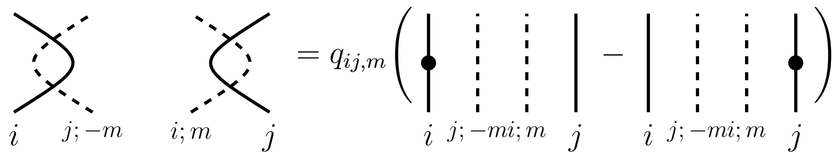

(8.4k)

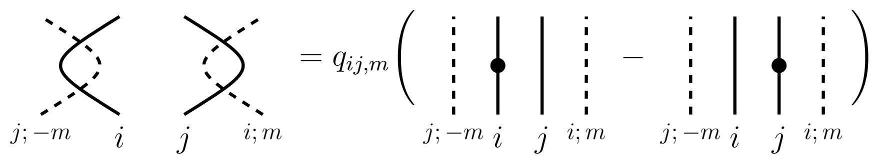

Note that the equations above are written assuming that $m > 0$, but they are equally valid if $m < 0$, with the requisite reordering of strands. The relation (2.6g) implies that:

(8.4l)

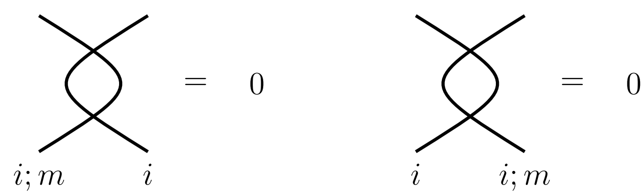

The relation (2.6j) is equivalent to isotopy and

(8.4m)

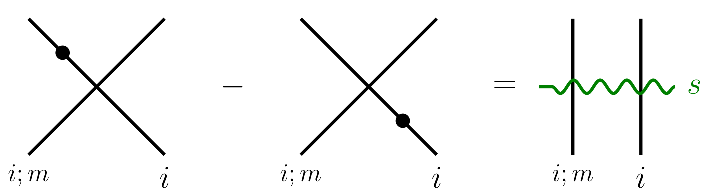

(8.4n)

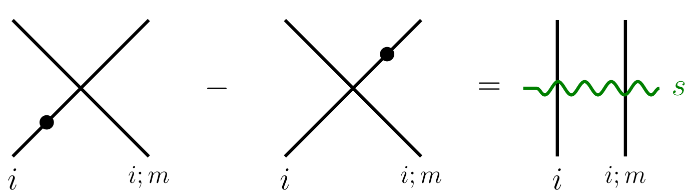

with s denoting the unique reflection in the affine Weyl group switching the top and bottom labels of the diagram.

Finally, the codimension 2 relations show how to relate the two resolutions of a triple point. The correct relation depends on the number of strands with the same label going through the triple point: if all three have the same label, then we use (2.6h), if two have the same label, we use (2.6k) and if there is no such pair, then (2.6c). These imply we can isotope through any triple point unless it involves exactly two partners with the label $i \in \mathcal{V}(\Gamma)$. In order to cover this last case, let $\partial f(u_1, u_2) = \frac{f(u_1)-f(u_2)}{u_1-u_2}$ for any polynomial $f(u)$.

(8.4o)

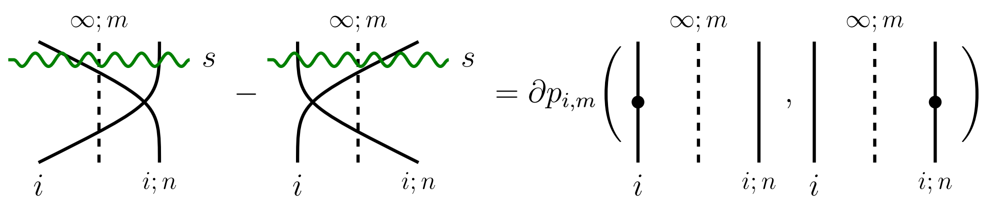

(8.4p)

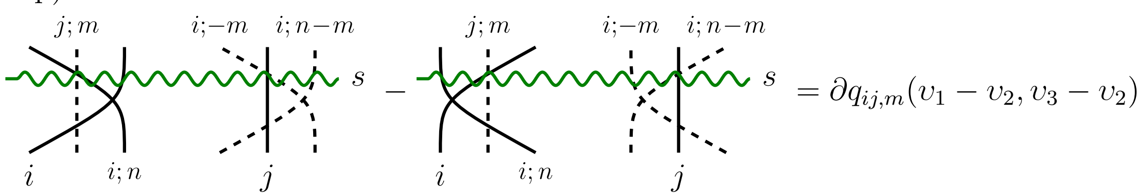

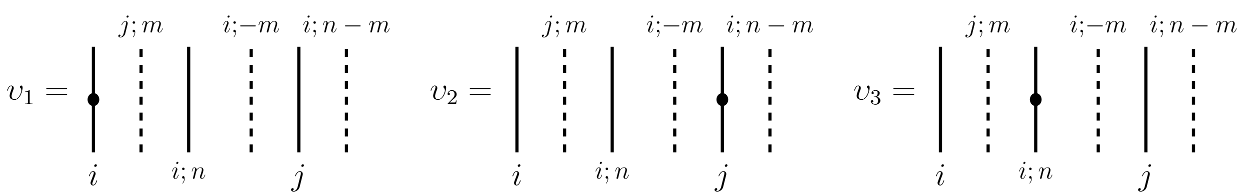

with s denoting the unique reflection in the affine Weyl group making the top and bottom match.

Lemma 8.3 *Given $\eta, \eta' \in \mathfrak{t}_{\beta,\gamma}$, the Hom space* $\mathrm{Hom}_{\mathscr{B}}(\eta, \eta')$ *is spanned by the morphisms $\mathbb{r}_D$ for unrolled diagrams D with top η' and bottom η, modulo the local relations (8.4a–8.4p).*

Proof. We have justified in each individual case why the relations (8.4a–8.4p) hold. Thus, we have a map from the formal span of unrolled diagrams modulo these relations to $\mathrm{Hom}_{\mathscr{B}}(\eta, \eta')$. This is surjective because the generating morphisms of the category $\mathscr{B}$ are given by the basic unrolled diagrams in (8.2). On the other hand, the relations (8.4a–8.4p) suffice to write any $\mathbb{r}_D$ as a sum of diagrams corresponding to a reduced word in $\widehat{W}$ with all dots and green lines at the bottom, and to relate any two reduced words for $w \in \widehat{W}$ modulo the diagrams for shorter elements of $\widehat{W}$. Thus, we find that the unrolled diagrams corresponding to the basis of [Webc, Cor. 3.13] are a spanning set of this quotient. This is only possible if the map is injective as well. □

8.2. **Quantum cylindrical KLRW algebras.** The reader may have noticed that these diagrams are actually quite difficult to draw and interpret, but there is a symmetry that we have not exploited, the action of the extended affine Weyl group. The quotient of $\prod_i \mathbb{R}^{v_i}$ by the extended Weyl group $\widehat{W}$ is given by the space $\prod_i (\mathbb{R}/\mathbb{Z})^{v_i}/S_{v_i}$, which we can interpret as the moduli space of multisubsets of the circle $\mathbb{R}/\mathbb{Z}$ labeled with elements of $\mathcal{V}(\Gamma)$, such that v_i elements have label $i \in \mathcal{V}(\Gamma)$. Thus the path $[0,1] \to \prod_i \mathbb{R}^{v_i}$ composed with the projection $\prod_i \mathbb{R}^{v_i} \to \prod_i (\mathbb{R}/\mathbb{Z})^{v_i}/S_{v_i}$ can be thought of as a path in this moduli space.

We draw this by considering our diagrams in $\mathbb{R} \times [0,1]$, and considering the quotient of this plane by $\mathbb{Z}$ acting by addition to the x-coordinate. Note that this sends all the partners to a single curve in $\mathbb{R}/\mathbb{Z} \times [0,1]$, and all ghosts to a single curve, whose position depends only on the coset of β_e modulo $\mathbb{Z}$. Of course, these are exactly the cylindrical KLRW diagrams.

Every cylindrical KLRW diagram has a unique lift to a path $[0,1] \to \mathfrak{t}_{\beta,\gamma}$ which starts in the fundamental region of $\widehat{W}$ where the coordinates $z_{i,k}$ satisfy

$$-\frac{1}{2} < z_{i,1} < z_{i,2} < \cdots < z_{i,v_i} < \frac{1}{2}. \tag{8.5}$$

That is, by the path lifting property of the universal cover, each of the curves $\bar{\pi}\colon [0,1] \to \mathbb{R}/\mathbb{Z}$ has a unique lift π with $-\frac{1}{2} < \pi(0) < \frac{1}{2}$, and we can number these so that

$$-\frac{1}{2} < \pi_{i,1}(0) < \cdots < \pi_{i,v_i}(0) < \frac{1}{2}. \tag{8.6}$$

Let $\tilde{D}$ be the unrolled diagram defined by the paths $\pi_{i,k}$, followed by the unique element of $\widehat{W}$ sending the top of this diagram back to the fundamental region.

Definition 8.4 *Given a cylindrical KLRW diagram D with no dots, let $\mathbb{r}_D$ denote the morphism $\mathbb{r}_{\tilde{D}}$ associated to the lifted unrolled diagram $\tilde{D}$.*

If the diagram contains dots, then place these in the lifted diagram on the unique partner preimage which has x-value in $(-\frac{1}{2}, \frac{1}{2})$.

We have to be careful about lifting cylindrical KLRW diagrams with dots, because if we do so in the most naive way, the result will not be compatible with composition, which the definition above is. We could accomplish the same effect if instead of applying a Weyl group element at the end, we applied one immediately whenever we left the fundamental region to move back into it.

We can also interpret the relations (8.4a–8.4p) as relations on cylindrical KLRW diagrams, as before following the rule that $\sum a_i D_i = 0$ if we have that $\sum a_i \mathbb{r}_{D_i} = 0$. Some of these can interpreted locally exactly as they appear above: (8.4c–8.4d) and (8.4l–8.4n) are of this type. On the other, if a dot on an cylindrical KLRW diagram is slid over the half-integer ghost with label ∞, then it goes between lifting to the ghost just right of $x = -\frac{1}{2}$ to that just left of $x = \frac{1}{2}$. Thus, if we draw $x = \frac{1}{2}$ as a fringed grey line, the effect of (8.4a) is thus the following relation on cylindrical KLRW diagrams:

(8.7) 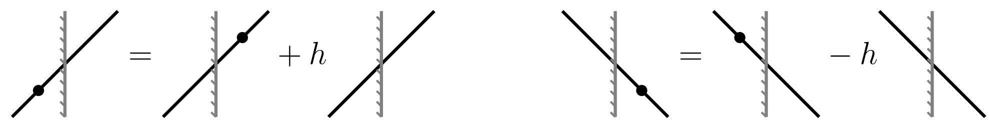

The other relations need to interpreted carefully to be compatible with lifting. For example, the relations (8.4h–8.4i) and the relations (8.4j–8.4j) need to be applied with in the version where all original strands have x-values in $(-\frac{1}{2}, \frac{1}{2})$ (we can isotope to avoid any values in the coset $\mathbb{Z} + \frac{1}{2}$). For (8.4h–8.4i), this means that $m \in (-\frac{1}{2}, \frac{1}{2})$, and for (8.4j–8.4j) that $m \in (-1, 1)$, with the sign determined by how the strands are cyclically ordered compared with $x = 1/2$.

Definition 8.5 *The* **quantum cylindrical KLRW (qcKLRW) algebra** *for the dimension vectors* $\mathbf{v}, \mathbf{w}$ *and parameters* $\beta_*, \gamma_*, b_*, c_{*,*}$ *is the quotient of the formal span of cylindrical KLRW diagrams for these parameters over* $\Bbbk[c_{*,*}, b_*, h]$ *modulo the relations induced by (8.4a–8.4p), in particular by (8.7), with the usual rule of multiplication by stacking.*

Some might prefer to think about the quantum cylindrical KLRW category, the category whose objects are cylindrical loadings, and morphisms are diagrams with fixed loadings at top and bottoms. Lemma 8.3 can thus be rephrased as:

Proposition 8.6 *The qcKLRW category is equivalent to the category* $\mathscr{B}_{\beta,\gamma}$.

Since the relations of the nilHecke algebra are not deformed in the qcKLRw algebra, it contains idempotents $e(a)$. These correspond to elements $\mathfrak{t}_{\beta,\gamma}$ which are central, since they correspond to scalar matrices. Thus the corresponding object η_a lies on the root hyperplanes corresponding to a copy of the finite root system for G; from this perspective, it is as singular as possible. In fact, the corresponding parahoric Iwa_{η_a} is exactly $G[[t]]$.

Lemma 8.7 *The endomorphisms of* η_a *are isomorphic to the quantum Coulomb branch* $\mathcal{A}^{\mathrm{sph}}$ *for a choice of flavor depending on* a.

Proof. Consider the cocharacter $\mu\colon \mathbb{C}^* \to \prod_{i\in\mathcal{V}(\Gamma)} GL(\mathbb{C}^{w_i})$ which acts on the kth basis vector in $\mathbb{C}^{w_i}$ with weight $\lfloor a+\gamma_{i,k}\rfloor$; note that this is locally constant in terms of a, only changing when a passes one of the red lines. The subspace U_{η_a} is given by $t^{-\mu}V[[t]]$, and so multiplication by t^{μ} induces a $G[[t]]$-equivariant isomorphism between the spaces ${}_\tau\mathcal{X}_\tau \cong {}_{\eta_a}\mathcal{X}_{\eta_a}$. However, this action does not commute with the action of the cocharacter φ (since it doesn't commute with the loop action); it intertwines the action of φ with its product with μ, so this gives the necessary shift of flavor. We have an isomorphism $\mathrm{End}_{\mathscr{B}}(\tau,\tau) \cong \mathcal{A}^{\mathrm{sph}}$ by [Webc, (3.1)]. □

Note, this construction is closely related to the flag Yangian introduced in [KTW$^+$, Def. 4.12]. In that paper, we assumed that Γ was bipartite (with the sets of nodes called **even** and **odd**), that if $i \to j$ then i is even and j odd. Furthermore, the definition depended on a polynomial p_i. If $h = 2$, $b_e = 0$ and the scalars $c_{i,k}$ are the roots (with multiplicity) of $p_i(2u-1)$, then the cylindrical KLR category is closely related to the flag Yangian category, via the transformation of diagrams sending all odd strands to their ghosts. Since there are some other minor differences of convention between these categories, we will not make a precise statement about the relationship between them.

Let us give a simple example. Consider $G = GL(2)$ and $V \cong \mathbb{C}^2 \oplus \mathfrak{gl}_2$. We have a natural isomorphism $\mathfrak{t}_{\mathbb{R}} \cong \mathbb{R}^2$ with the coordinates given by z_1, z_2. The unrolled matter hyperplanes are $z_1, z_2, z_1 - z_2 \in \mathbb{Z} - \frac{1}{2}$ and the unrolled root hyperplanes are $\alpha = z_1 - z_2 \in \mathbb{Z}$.

With these conventions, we match morphisms of the extended category with cylindrical KLRW diagrams. We'll draw these on a cylinder sliced open at $x = \frac{1}{2}$.

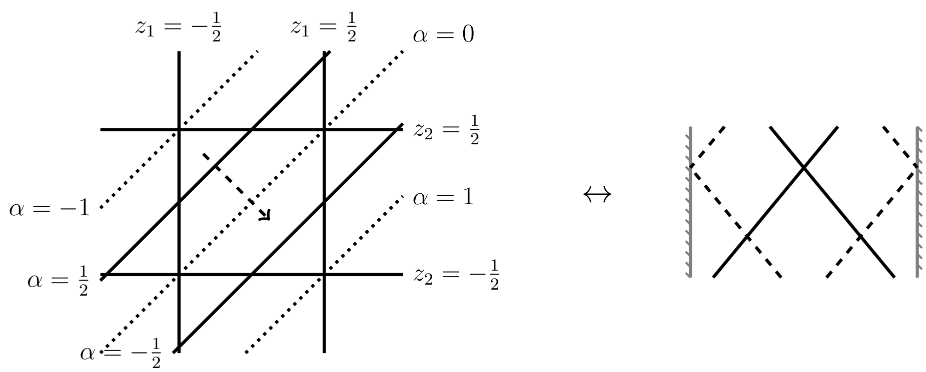

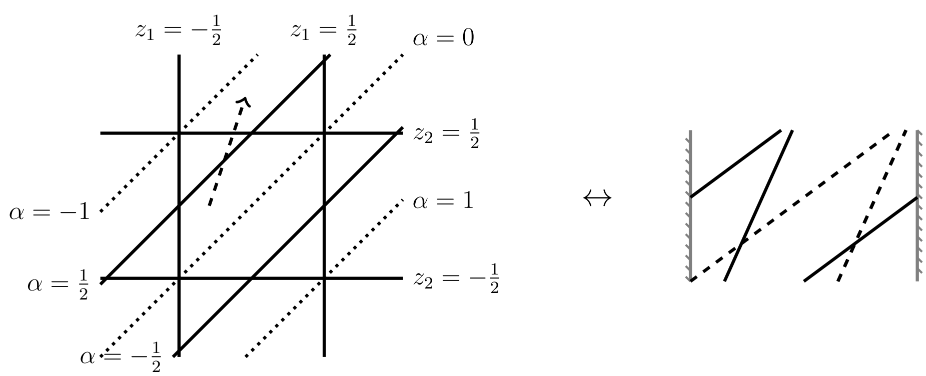

8.3. **The classical limit.** Now, let use consider the classical limit where we set $h = 0$. In this case, the relations (8.4a–8.4p) become exactly the relations (6.2a–6.2k). As the name suggests, we thus have:

Lemma 8.8 *The cylindrical KLRW algebra of Definition 6.3 is the specialization of the qcKLRW algebra at $h = 0$.*

This observation allows us to finally begin proving results from Section 6:

Proof of Thm. 6.12. Specialize Lemma 8.7 at $h = 0$. □

Note that after this specialization, the shift of flavor that was needed in the quantum version this theorem disappears; since the classical Coulomb branch can be written as $G[[t]]$-equivariant homology of ${}_\tau\mathfrak{X}_\tau$, the map in the proof of Lemma 8.7 induces an algebra isomorphism on the nose.

Note that this implies a basis theorem for the cylindrical KLRW algebras. Given a cylindrical loading with v_i elements mapping to $i \in \mathcal{V}(\Gamma)$, we have a unique way of lifting to real numbers $z_{i,1}, \dots, z_{i,v_i}$ in the fundamental region (that is, satisfying $-\frac{1}{2} < z_{i,1} < \cdots < z_{i,v_i} < \frac{1}{2}$), and the extended affine Weyl group $\widehat{W}$ acts freely transitively on the set of possible lifts. Having fixed two cylindrical loadings S and T, there is an unrolled diagram with a minimal number of crossings with the bottom given by this lift of S and the image of this lift of T under $w \in \widehat{W}$; this diagram is not unique, but as usual, any two choices differ by the diagram for a shorter permutation by the relations (6.2i–6.2k). The image of this diagram D_w on the cylinder $\mathbb{R}/\mathbb{Z} \times [0, 1]$ gives a weighted antipodal diagram.

From [Webc, Cor. 3.13], we find that:

Lemma 8.9 *The Hom space between two objects in the cylindrical wKLR category is a free module for the left action of polynomials in the dots, with basis D_w for $w \in \widehat{W}$.*

Studying the $h = 0$ case is particularly important because of its connection to coherent sheaves, as shown in [Weba]. In Theorem 2.21, we show that the category of representations of the quantum Coulomb branch over $\mathbb{F}_p$ is equivalent to the representations of a subcategory $\widehat{\mathsf{A}}_p(\mathbb{F}_p)$ in the extended BFN category with a different set of parameters, which we call "pth root" (Definition 2.12). Using the structure of quantizations in characteristic p, we use this isomorphism to give a derived equivalence between the representations of B_τ, the extended BFN category with these parameters, and the coherent sheaves $\mathrm{Coh}(\tilde{\mathfrak{M}}_{\mathbb{F}_p})$ as long as p is sufficiently large, and the parameters are generic (Theorem 3.28). Note that the category B_τ and the induced equivalence depend on parameters $b_e, c_{i,k} \in \mathbb{F}_p$, but not on a choice of $\beta_e, \gamma_{i,k}$; these are fixed to all be $-\frac{1}{2}$ to match τ).

By Lemma 8.8, representations of the category B_τ is equivalent to the representations of a cylindrical KLRW algebra with $\Bbbk = \mathbb{F}_p$ for the same dimension vectors $\mathbf{v}, \mathbf{w}$, but we need to explain how its parameters depend on $b_e, c_{i,k} \in \mathbb{F}_p$. Under pth root conventions, since we take a pth root of the flavor, we have that

$$\beta_e = \frac{b_e}{p} + \frac{1}{2p} \qquad \gamma_{i,k} = \frac{c_{i,k}}{p} + \frac{1}{2p}. \tag{8.8}$$

This gives us a result (Corollary 3.22) relating representations of $\mathring{R}$ over $\mathbb{F}_p$, and coherent sheaves on $\mathrm{Coh}(\tilde{\mathfrak{M}}_{\mathbb{F}_p})$, again, as long as p is sufficiently large, and the parameters are generic. Those familiar with the work of Kaledin [Kal08] will expect this to be phrased in terms of a tilting generator, which we define in Definition 3.16; this vector bundle is defined over $\mathbb{Z}$, and is a sum of terms, one for each loading such that all x-values are in $\frac{1}{p}\mathbb{Z}/\mathbb{Z} \subset \mathbb{R}/\mathbb{Z}$.

Note that while our possible choices of parameter are constrained by p, this is not the case up to equivalence. For any alcove C defined by the hyperplanes (6.9), we can choose that $p, b_e, c_{i,k}$ so that

(1) the parameters (8.8) lie in C
(2) every equivalence class of cylindrical loadings has a representative with values in $\frac{1}{p}\mathbb{Z}/\mathbb{Z} \subset \mathbb{R}/\mathbb{Z}$

Theorem 3.28 shows that:

Lemma 8.10 *If $\Bbbk$ be characteristic 0 (or just sufficiently large) then the variety $\tilde{\mathfrak{M}}$ carries a tilting generator $\mathcal{Q}_{\mathbb{Q}}$ whose endomorphisms A (defined in Section 3.6) are the cylindrical KLRW algebra $\mathring{R}$ for the parameters* (8.8).

This establishes Theorem 6.15 and allows to us to prove two of the claimed results in the introduction: Theorem E(2) now follows from Theorem 3.28 and Theorem D from Corollary 3.30. .

Example 8.i. Consider the case where $G = \mathbb{C}^*$ acting on $\mathbb{C}^2$ by scalars. In this case, the Coulomb branch is $T^*\mathbb{P}^1$. The corresponding cylindrical wKLR algebra has two red strands, and one black strand, all with the same label. There are two idempotents in this algebra, corresponding to the two cyclic orders of the 3 strands. Since the corresponding quiver has no edges, the black strand has no ghosts.

x $\qquad x^*$ $\qquad y$ $\qquad y^*$

These satisfy the quadratic relations

$$xx^* = yy^* \qquad x^*x = y^*y, \tag{8.9}$$

and it's easy to check that these are a complete set of relations. This algebra is Koszul and its Koszul/quadratic dual is easily seen to be defined by

$$xx^* = -yy^* \qquad x^*x = -y^*y \qquad y^*x = x^*y = yx^* = xy^* = 0. \tag{8.10}$$

This latter set of relations defines an 8-dimensional algebra studied by Nandakumar and Zhao in [NZ], which appears as the endomorphisms of a projective generator for exotic sheaves on $T^*\mathbb{P}^1$. ♣

Example 8.ii. The most important example of a case where weighting is useful is case of $\mathrm{Sym}^n(\mathbb{C}^2)$; in this case we have only a single node in our quiver, which carries a loop, equipped with the weight $\beta_e = \vartheta$, and a single red strand (of course, labeled with this node). This describes B when $\frac{b_e}{p} + \frac{1}{2p} \approx \theta$.

The objects in the cylindrical wKLR category are thus n-tuples of distinct points in S^1, where each point has a ghost ϑ units to its right, which the other points avoid. This information can be recorded by listing the order in which one encounters dots and ghosts; the set of possible configurations for a given ϑ corresponds to the set $\bar{\Lambda}$ discussed earlier.

Note that the set of possible configurations is locally constant, and will only change at values of θ where one has a non-simple hyperplane arrangment this can only be the case if there is a loop of equations

$$\begin{aligned} z_1 - z_2 &\equiv \vartheta \pmod{\mathbb{Z}} \\ z_2 - z_3 &\equiv \vartheta \pmod{\mathbb{Z}} \\ &\vdots \\ z_k - z_1 &\equiv \vartheta \pmod{\mathbb{Z}} \end{aligned}$$

for $k \leq n$. This implies that $k\vartheta \in \mathbb{Z}$, i.e. that ϑ is rational with denominator $\leq n$. Of course, this same set of values has shown up in the structure of Hilbert schemes and Cherednik algebras in other contexts. ♣

8.4. **Change of flavor.** Let us now discuss how we can interpret twisted diagrams in this context.

We can choose flavors ϕ, ϕ' corresponding to choices of parameters $b_e, c_{i,k} \in \mathbb{F}_p$ for ϕ, and $b'_e, c'_{i,k} \in \mathbb{F}_p$ for ϕ'. Choose integers

$$\nu_e \equiv b'_e - b_e \pmod{p} \qquad \nu_{i,k} \equiv c'_{i,k} - c_{i,k} \pmod{p}.$$

We can interpret ν here as a real cocharacter of the flavor torus T_F (with the usual caveats about redundancy). Associated to ν, we have a twisting bimodule ${}_{\phi+\nu}\mathscr{T}_{\phi}$ over the categories $\mathscr{B}_{\phi+\nu}$ and $\mathscr{B}_{\phi}$, defined in (2.8). The elements of these bimodules are morphisms in the BFN category $\mathscr{B}^Q$ attached to the larger group Q acting on V. Thus, applying Proposition 2.18 to $\mathscr{B}_{\phi+\nu}, \mathscr{B}_{\phi}$ and $\mathscr{B}^Q$, we obtain that the twisting bimodule ${}_{\phi+\nu}\mathscr{T}_{\phi}$ is intertwined with the corresponding a similar bimodule ${}_{\phi'_{1/p}}\mathsf{T}_{\phi_{1/p}}$ with pth root conventions. Since $\frac{1}{p}\nu$ might not be intergral, we cannot apply the definition of ${}_{\phi+\nu}\mathscr{T}_{\phi}$ directly, and we take the description above to be the definition, but let us say a few words about why the fact that $h = 0$ allows us to extend this definition to arbitrary cocharacters of $\mathfrak{t}_F$.

We will spare the reader the blizzard of notation required to say this carefully, but in brief $\mathsf{B}_{\phi_{1/p}}$ and $\mathsf{B}_{\phi'_{1/p}}$ can be realized as subcategories of B^Q modulo the action of polynomial morphisms $\mathfrak{t}_F^*$; note that this quotient is only well-defined because $h = 0$. We identify their object sets with orbits of $\mathfrak{t}$ in $\mathfrak{t}_{1;Q}$ which differ by $\frac{1}{p}\nu$, and note that morphisms in $\mathsf{B}_{\phi_{1/p}}$ and $\mathsf{B}_{\phi'_{1/p}}$ are precisely those in the larger category generated by paths, polynomials, and u_α and the extended affine Weyl group of G (as opposed to the affine Weyl group of Q, which has more translations). We can define ${}_{\phi'_{1/p}}\mathsf{T}_{\phi_{1/p}}$ as the space of morphisms in $\mathsf{B}^Q/(\mathfrak{t}_F^*)$ generated by paths, polynomials, and u_α and the extended affine Weyl group of G which begin in one coset and end in the other.

In the quiver case, we can also recover the twisting bimodules ${}_{\phi+\nu}\mathscr{T}_\phi$ and ${}_{\phi'_{1/p}}\mathsf{T}_{\phi_{1/p}}$ using the appropriate modification of Proposition 8.6 and Lemma 8.8. Let

$$\boldsymbol{\beta}_e = \beta_e + \nu_e t \qquad \boldsymbol{\gamma}_{i,k} = \beta_e + \nu_{i,k} t$$

We can define a bimodule over qcKLRW algebras given by the span of twisted cylindrical KLRW diagrams for these functions modulo the local relations (8.4a–8.4p); note that applying these relations requires some real care: the real number labeling ghost strands will vary as t changes, and that must be incorporated into the relations (8.4j,8.4k,8.4o,8.4p).

Proposition 8.11 *The isomorphism of Proposition 8.6 extends to a bimodule isomorphism of ${}_{\phi+\nu}\mathscr{T}_\phi$ with the span of twisted cylindrical KLRW diagrams for the functions $\boldsymbol{\beta}_e$ and $\boldsymbol{\gamma}_{i,k}$ modulo the local relations (8.4a–8.4p);*

With pth root conventions, this means that ${}_{\phi'_{1/p}}\mathsf{T}_{\phi_{1/p}}$ is isomorphic to the span of twisted cylindrical KLRW diagrams with $h = 0$ for the functions

$$\boldsymbol{\beta}_e = \frac{b_e + \nu_e t}{p} + \frac{1}{2p} \qquad \boldsymbol{\gamma}_{i,k} = \frac{c_{i,k} + \nu_{i,k} t}{p} + \frac{1}{2p}. \tag{8.11}$$

This now establishes Theorem E(3) by Theorem C.

Proof of Thm. 6.14. By the $h = 0$ special case of Proposition 8.11 above, we can rewrite $e(a)B_{\boldsymbol{\beta}^{(k)},\boldsymbol{\gamma}^{(k)}}e(a)$ as the bimodule ${}_{\phi+n\nu}\mathscr{T}_\phi(\eta_a, \eta_a)$. By the definition (2.8), we thus have

$$e(a)B_{\boldsymbol{\beta}^{(k)},\boldsymbol{\gamma}^{(k)}}e(a) \cong H_*^{BM}({}_{\eta_a}\mathfrak{X}_{\eta_a}^{(k\nu)}).$$

Since the space ${}_{\eta_a}\mathfrak{X}_{\eta_a}^{(k\nu)}$ is precisely the same as the quotient by $G[[t]]$ of that denoted $\tilde{\mathcal{R}}^{(k\nu)}$ in [BFN], this shows that our definition matches exactly the projective coordinate ring of the partial resolution attached to the cocharacter ν in [BFN]. □

Proof of Prop. 6.18. For some large prime p, we can approximate β and β' by choices of $b_e, c_{i,k}$ as in (8.11), such that derived localization holds at these parameters by Theorem 3.25. Thus, the corresponding wall-crossing functor is a derived equivalence by Lemma 4.4, and so Proposition 8.11 shows the same is true for tensor product with $B_{\boldsymbol{\beta}}$ with coefficients in $\mathbb{F}_p$.

For any base ring $\Bbbk$, derived tensor product with $B_{\boldsymbol{\beta}}$ is an equivalence if and only if the natural map $\mathring{R}_\beta \to \mathrm{RHom}_{\mathring{R}_{\beta'}}(B_{\boldsymbol{\beta}}, B_{\boldsymbol{\beta}})$ is an isomorphism, and similarly with β, β' reversed. First consider the cone $K^\bullet$ of this map with $\Bbbk = \mathbb{Z}$. The result is a complex whose cohomology is a finitely generated graded bimodule over $\mathring{R}_\beta$ in each degree. Since each graded degree of $\mathring{R}_\beta$ is a finitely generated abelian group, the same is true of $H^k(K^\bullet)$ for each homological degree k. In particular, we can use the universal coefficient theorem to compute the cohomology of $K^\bullet \otimes_{\mathbb{Z}} \Bbbk$ for any ring $\Bbbk$. We have shown above that $K^\bullet$ is exact after base change to $\mathbb{F}_p$ for some p; this shows that the cohomology of $H^k(K^\bullet)$ is torsion of order coprime to p, and so becomes trivial for $\Bbbk = \mathbb{Q}$. □

9. Proofs from Sections 6 and 7

We deferred the proofs of a number of results in Sections 6 and 7 which required the results of Section 8; in this section we will cover these and any preliminary lemmata needed for them.

Proof of Lem. 6.8. By Lemma 8.9, the module $e(\mathbf{i})\mathring{R}^{\mathbf{j}}$ has a basis as a free right module over the action of the dots indexed by affine permutations, with labels on strands fixed by the labeling on the top given by $\mathbf{i}$. This basis is given by diagrams that trace out this affine permutation on the cylinder with a minimal number of crossings and no dots. We can easily check that the set of these diagrams in a fixed left coset[2] of the finite permutation group span a projective right $\tilde{T}^{\boldsymbol{\lambda}}$ module, freely generated by the unique shortest element of this coset. This is isomorphic to $e(\mathbf{i}')\tilde{T}^{\boldsymbol{\lambda}}$, where $\mathbf{i}'$ is determined by the bottom of this shortest coset diagram (the one with a minimal number of crossings). □

In order to prove Lemma 6.21, we need to give a lemma comparing the central charge Z with a representation theoretic central charge similar to the ones considered in [ABM15].

Fix a choice of parameters $b_e, c_{i,k} \in \mathbb{Z}$ so that the associated parameters of the form (8.11) lie in our preferred alcove C_0. Consider the associated sheaf of algebras $\hat{\mathscr{W}}_\phi$ for the corresponding flavor defined by Definition 3.5. For any other $b'_e, c'_{i,k} \in \mathbb{Z}$, we can consider the quantized line bimodule ${}_{\phi+\nu}\mathscr{T}_\phi$ with $\nu_e = b'_e - b_e, \nu_{i,k} = c'_{i,k} - c_{i,k}$. Tensor product ${}_{\phi+\nu}\mathscr{T}_\phi \otimes -$ gives an equivalence of abelian categories $\hat{\mathscr{W}}_\phi$-mod $\cong \hat{\mathscr{W}}_{\phi+\nu}$-mod, so we can implicitly identify all these categories. Let $\mathscr{W}_\phi$-mod$_0$ be the subcategory of sheaves of modules which are set-theoretically supported on the fiber over the cone point in $\mathfrak{M}$; note that we have left out the completion here, since the action on any such module automatically extends to the completion.

[2]As always confuses the author, this means an orbit for right multiplication of a subgroup.

Definition 9.1 *The representation theoretic central charge on* $K^0(\hat{\mathscr{W}}_\phi\text{-mod}_0)$ *is defined by*

$$\mathcal{Z}_{\mathbf{b}',\mathbf{c}'}([\mathcal{M}]) = \frac{1}{p^V}\chi(\mathbb{R}\Gamma(\mathcal{M}))$$

Since the parameters (8.11) lie in the alcove C_0, then we have an equivalence

$$\mathcal{G}\colon D^b(\mathscr{W}_\phi\text{-mod}_0) \cong D^b(\mathring{R}_{C_0}\text{-fdmod}).$$

This is effectively Lemma 3.10 but also follows from applying Lemma 8.10, together with equivalence of $\hat{\mathscr{W}}_\phi$-mod$_0$ with coherent sheaves with the same support; the characteristic assumption on Lemma 8.10 is unnecessary by Theorem 3.25. Using the equivalence, we can interpret $\mathcal{Z}_{\mathbf{b}',\mathbf{c}'}$ as a function on $K^0(\mathring{R}_{C_0}\text{-fdmod})$.

Lemma 9.2 *The RT central charge* $\mathcal{Z}_{\mathbf{b}',\mathbf{c}'}$ *is:*

(1) polynomial in the variables $\frac{b'_e}{p}$ *and* $\frac{c'_{i,k}}{p}$.

(2) for $(\beta_*, \gamma_{*,*}) \in C_0$*, this function limits to* $Z_{\beta,\gamma}$ *in the sense that if* $p^{(m)}$ *is a series of primes with* $\lim_{m\to\infty} p^{(m)} = \infty$*, and* $b_e^{(m)}, c_{i,k}^{(m)}$ *a series of parameters such that* $\frac{b_e^{(m)}}{p^{(m)}}$ *and* $\frac{c_{i,k}^{(m)}}{p^{(m)}}$ *converge to* β_e *and* $\gamma_{i,k}$*, then* $\mathcal{Z}_{\mathbf{b}^{(m)},\mathbf{c}^{(m)}}$ *converges to* $Z_{\beta,\gamma}$ *as* $m \to \infty$.

Note any point in $\mathbb{R}^\ell$ is a limit of the desired form, so $Z_{\beta,\gamma}$ is determined by this property.

Proof. Given an $\hat{\mathscr{W}}_\phi$-module $\mathcal{M}$, we can consider its Rees module as a module over the usual microlocalization W of $\mathcal{A}^{\mathrm{sph}}$ to $\tilde{\mathfrak{M}}$ (essentially just forgetting the coherent sheaf structure); we can then reduce this sheaf modulo h, and obtain a different coherent sheaf $\mathcal{N}$.

This will perhaps be easier to picture in terms of $\mathbb{Z}$-algebras as discussed in Section 3.2. Let ${}_{\phi+k\chi}\mathscr{L}_\phi$ be the quantizations of the line bundles $\mathcal{O}(k)$ as left W-modules. The sheaf $\mathcal{M}$ is associated to the module over the projective coordinate ring given by $\mathbb{R}\Gamma_{\mathbb{S}}({}_{\phi+kp\chi}\mathscr{L}_\phi \otimes \mathcal{M})$, with sections of the line bundles acting by the quantum Frobenius map. On the other hand, $\mathcal{N}$ corresponds to taking the associated graded for $\mathbb{R}\Gamma_{\mathbb{S}}({}_{\phi+k\chi}\mathscr{L}_\phi \otimes \mathcal{M})$, with induced action of the associated graded of the $\mathbb{Z}$-algebra.

Note that $\mathbb{R}\Gamma(\mathcal{N})$ has the same Euler characteristic of $\mathbb{R}\Gamma(\mathcal{M})$, since the former is the limit of a spectral sequence with E_1 page given by the latter.

Of course, if we consider the tensor product $\mathcal{M}' = {}_{\phi+\nu}\mathscr{T}_\phi \otimes \mathcal{M}$, then the resulting sheaf $\mathcal{N}'$ is simply $\mathcal{N}' \cong \mathcal{N} \otimes \mathcal{O}(\nu)$. This shows that the polynomiality in b'_e/p and $c'_{i.k}/p$, since changing these parameters is just tensoring $\mathcal{N}$ with a line bundle.

By Theorem 2.24, the central charge $p^V \prod v_i! \cdot \mathcal{Z}_{\mathbf{b}',\mathbf{c}'}$ is equal to a sum over the set Λ of the dimension of the image of the corresponding idempotent eM times the number of p-torsion points in the corresponding polytope.

On the other hand, the $Z_{\beta,\gamma}$ is the sum over the same set, weighted by volume, and divided by $\prod v_i!$ to account for the fact that in our set of longitudes $a_1 < \cdots < a_n$ is increasing.

Thus, the agreement of these in the limit is just the fact that as $p \to \infty$, the number of p-torsion points in a polytope inside of $(\mathbb{R}/\mathbb{Z})^n$ times $\frac{1}{p^n}$ limits to the volume of this set (the leading coefficient of the Erhart polynomial). □

Remark 9.3. Applying Grothendieck-Hirzebruch-Riemann-Roch, we can also compute this representation theoretic central charge as an integral over the zero fiber of the resolution $\tilde{\mathfrak{M}} \to \mathfrak{M}$. Up to normalization, this should match [Aga, (5.18)].

Proof of Lem. 6.21. Consider two alcoves C^+ and C^- with C^+ above C^- across a single hyperplane, and the algebras $\mathring{R}_{C^\pm}$.

This implies that for any sufficiently large prime p there are parameters $b_e, c_{i,k}$ and a integer $m > 0$ such that taking $\nu = m\chi$, the path (8.11) begins in C^-, ends in C^+ and passes through no other alcoves. By Lemma 4.4, we have a commutative diagram:

$$\begin{array}{ccc} D^b(\mathscr{W}_\phi\text{-mod}_0) & \xrightarrow{\ {}_{\phi+\nu}\mathscr{T}_\phi \otimes -\ } & D^b(\mathscr{W}_{\phi+\nu}\text{-mod}_0) \\ \Big\downarrow \mathcal{G} & & \Big\downarrow \mathcal{G} \\ D^b(\mathring{R}_{C^-}\text{-fdmod}) & \xrightarrow[\ B_{\boldsymbol{\beta}} \otimes -\ ]{} & D^b(\mathring{R}_{C^+}\text{-fdmod}) \end{array} \tag{9.1}$$

Thus, we have that the RT central charge defined with respect to C^- for an $\mathring{R}_{C^-}$-module M is the same as that defined with respect to C^+ for $B_{\boldsymbol{\beta}} \overset{L}{\otimes} M$. Since the equivalence $\mathbb{B}_{C,C_0}$ is uniquely defined by the relation that $\mathbb{B}_{C^-,C_0}(M) = \mathbb{B}_{C^+,C_0}(B_{\boldsymbol{\beta}} \overset{L}{\otimes} M)$, for any such pair, we have that the RT central charge with respect to C_0 for $\mathbb{B}_{C,C_0}(M)$ is the same as this charge for M with respect to C. This establishes the analogue of Lemma 6.21 for the RT central charge, and thus taking limit as in Lemma 9.2 gives the desired result. □

Before proving Theorem 6.22, we need to establish some facts about the perverse structure on wall-crossing/translation functors given by Losev [Los, Prop. 7.3]; we are interested in this filtration in the characteristic p case, but since it is defined by starting with ideals in characteristic 0, we will need to consider the characteristic 0 case as well.

We will be focused on the special case relevant for us, where we have fixed $b_e = 0$ for all edges and we are crossing over a wall defined by $c_{i,k} = c_{j,m}$. This perverse structure is based on a chain of ideals $\mathcal{I}^0 \subset \cdots \subset \mathcal{I}^V$, defined as follows for parameters $c_{i,k} - c_{j,m} \in \mathbb{Z}_{\neq 0}$: consider $\mathcal{A}_\phi^{\mathrm{sph},\mathbb{C}}$, with all other $c_{*,*}$'s generic, and let $\mathcal{I}^k$ be the minimal ideal such that $\mathcal{A}_\phi^{\mathrm{sph},\mathbb{C}}/\mathcal{I}^k$ has Gelfand-Kirillov dimension $2(V-k)$. This is also well-defined for arbitrary $c_{*,*}$ by continuity. This defines an ideal in $\mathcal{A}_\phi^{\mathbb{Z}}$ by intersection, and then over any other base field by base change.

We want to describe these ideals using only Gelfand-Tsetlin modules, that is modules locally finite under the ring S_1. Let us just remind the reader of a few facts: the ring S_1 is a polynomial ring generated by variables $z_{i,k}$ for $k = 1, \dots, v_k$. Thus, a maximal ideal of this ring is described by choosing scalars $a_{i,k} \in \mathbb{C}$. There is a natural notion of equivalence of maximal ideals: the important property is that maximal ideals $a_{i,k}$ and $a'_{i,k}$ are equivalent if:

(1) $a_{i,k} - a'_{i,k} \in \mathbb{Z}$ for all i and $k = 1, \dots, v_i$,
(2) if $a_{i,k} - a_{j,m} \in \mathbb{Z}$, then $a'_{i,k} - a'_{j,m} \in \mathbb{Z}$ is an integer of the same sign, and *vice versa*,
(3) if $c_{i,k} - a_{j,m} \in \mathbb{Z}$, then $c_{i,k} - a'_{j,m} \in \mathbb{Z}$ is an integer of the same sign, and *vice versa*.

In order to give the sharpest version of this theorem, we need to carefully specify when 0 counts as positive or negative, which depends on various choices of convention; for our purposes, this is irrelevant. By [Webb, Lem. 4.15], we have that:

Lemma 9.4 *For any Gelfand-Tsetlin module M over $\mathcal{A}_\phi^{\mathrm{sph},\mathbb{C}}$, the weight spaces of two equivalent maximal ideals (the subspaces on which they act nilpotently) are naturally isomorphic.*

While only stated in the case of a type A quiver, this follows by the same logic as [Webb, Prop. 5.4].

Lemma 9.5 *For any quotient $\mathcal{A}_\phi^{\mathrm{sph},\mathbb{C}}/I$ with GK dimension m, there is a faithful GT module over this quotient with GK dimension $\frac{m}{2}$.*

Proof. Let S' denote the image of S in $A = \mathcal{A}_\phi^{\mathrm{sph},\mathbb{C}}/I$. Let $\mathfrak{p} \subset S'$ be the ideal of a component of $\operatorname{Spec} S'$ of maximal dimension $d = \dim V(\mathfrak{p}) = \dim \operatorname{Spec} S'$, and let $\mathfrak{m}$ be a generic maximal ideal containing $\mathfrak{p}$. This ideal is defined by scalars $a_{i,k} \in \mathbb{C}$ as above, and in particular, generic means that we minimize the number of pairs of $(a_{i,k}, a_{j,m})$ or $(c_{i,j}, a_{j,m})$ with integral difference. We can thus have at most $V - d$ independent equations of the form $a_{i,k} = a_{j,m} + p$ or $c_{i,k} = a_{j,m} + p$ which are satisfied. That is, we can find d disjoint subsets $\Omega_1, \dots, \Omega_d$ of Ω such that if (i,k) in Ω_q, then $a_{i,k}$ does not have integral difference with any $c_{j,m}$ or with $a_{j,m}$ outside of Ω_q. Thus, for any $(x_1, \dots, x_d) \in \mathbb{Z}^d$, we can add x_q to all elements of Ω_q and obtain an equivalent maximal ideal. Let X be the Zariski closure of these points, which is a d-dimensional affine subspace.

Now consider the module $A/A\mathfrak{m}^N$, which is a GT module. Note that this has non-zero multiplicity at all points in $\operatorname{Spec} S$ that lie in the equivalence class of $\mathfrak{m}$; in particular, at the points obtained by translation as above. This shows that $X \subset \operatorname{Spec} S'$, and for dimension reasons, it must be a component of this variety. Since X contains $V(\mathfrak{m})$, this is only possible if $X = V(\mathfrak{p})$.

Thus, the annihilator of $A/A\mathfrak{m}^N$ is a 2-sided ideal I' of A/I; for N sufficiently large, the ideal $I' \cap S'$ has trivial $\mathfrak{p}$-primary component. Thus, $I' \cap S'$ has strictly fewer associated primes of dimension d than S'.

Now, apply the same logic to I' as a left A-module: let S'' be the quotient of S_1 by the elements annihilating $A/A\mathfrak{m}^N$. Note that S'' is a quotient of S', and as noted above, it either has dimension $< d$ or fewer components of dimension d. Choose $\mathfrak{m}_2$ generic in a component of $\operatorname{Spec} S''$, etc.

We can inductively define I_k as the annihilator of $\bigoplus_{q=1}^{k-1} A/A\mathfrak{m}_q^{N_q}$, the ring S_k as S_1 modulo the annihilator of this module, and $\mathfrak{m}_k$ as a generic maximal ideal in a component of $\operatorname{Spec} S_k$. Since at each step, the number of components or dimension of $\operatorname{Spec} S_k$ drops, eventually, this process will terminate at a module M which is faithful, since its annihilator is killed by all elements of S'.

Finally, we need to show that $2d$ is the Gelfand-Kirillov dimension of A/I. As usual, if we let $J_1, \dots, J_n$ be the primitive ideals killing the composition factors of M, then $J_1 \cdots J_n \subset I \subset J_i$ for all i, so the GK dimension of A/I is the same as the max of the GK dimensions A/J_i, and we can assume without loss of generality that I is the annihilator of a simple GT module. The support in $\operatorname{Spec} S_1$ of this simple module M is a finite union of equivalence classes, so its Zariski closure is a finite union of affine subspaces of dimension $< d$. As before, let X be one of these spaces of maximal dimension, and let $\mathfrak{p}$ be its maximal ideal. Then the tensor product of A/I on the right with the localization $S'_{\mathfrak{p}}$ gives a finite dimensional module over the fraction field K of $S'_{\mathfrak{p}}$.

The action of A/I on this tensor product gives an embedding of A/I into a finite rank matrix algebra over the twisted group algebra of the lattice of differences between integral points on X with coefficients in the fraction field K. Furthermore, this map becomes an isomorphism after tensor product with K. This shows that A/I has GK dimension $2d$. □

Proof of Thm. 6.22. By Lemma 6.21, it's enough to check these properties for the chamber C_0. In this case $\langle Z(x), [M]\rangle$ is the integral of a positive function, and thus is positive; this shows (1).

Now we turn to (2a). The standard structure is compatible with the filtration by $\mathcal{D}_n$ by definition, so this is clear.

Finally, consider (2b); this requires us to show the equivalence $\mathbb{B}_{C,C_0}$ for C below C_0 is perverse with respect to this filtration. By Proposition 8.11, it is equivalent to the same question as whether the functor ${}_{\phi+\nu}\mathscr{T}_\phi \overset{L}{\otimes} -$ is perverse. This is proven in [Los, Prop. 7.3], and so we need only prove that our filtration of the category agrees with Losev's.

We have already described Losev's filtration above. It is based on the ideals $\mathcal{I}^q$, which by Lemma 9.5 is the annihilator of all GT modules with GK dimension $\leq V - q$.

Thus, $\mathcal{I}^q$ kills a GT module if and only if all equivalence classes with non-zero weights have Zariski closure of dimension $\leq V - q$. The only way this can happen in the case where $c_{*,*}$ are generic, except that $c_{i,k} - c_{j,m} \in \mathbb{Z}_{\geq 0}$, is for there there to be q values of $a_{\ell,n}$ between $c_{i,k}$ and $c_{j,m}$. In the corresponding idempotent over the KLRW algebra, this means that there are at least q black strands between this particular pair of reds. That is, it is killed by the planar equivalent of the idempotent e_q defined in Section 7.2.

This description of the ideal is unchanged when we specialize to non-generic values of $c_{*,*}$ and reduce mod p. Thus, we have that the modules killed by Losev's $\mathcal{I}^q$ are those whose corresponding cylindrical KLRW module is killed by the idempotent e_q exactly as defined in Section 7.2. This completes the proof. □

Proof of Thm. 7.2. This follows immediately from Proposition 4.7: the space $\mathring{T}_{1,F}$ is the torus minus the toric braid arrangment, so we can write the affine braid groupoid inside the fundamental groupoid of this space as usual, and the action of the functors $\Phi_w^{\phi',\phi''}$ match the bimodules $\mathbb{B}_\tau$ by Proposition 8.11. □

Proof of Lem. 7.5. Recall that in [Web17a, §7.3], we defined a $\tilde{T}^{\mathbf{j}'}$-$\tilde{T}^{\mathbf{j}}$ bimodule $\mathfrak{k}_{\mathbf{j}}^{\mathbf{j}'}$ where the red strands trace out a cap with an element of L_0.

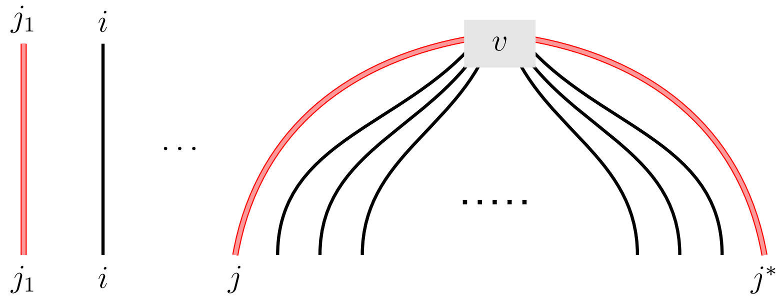

This bimodule is strongly equivariant, and we let $\mathring{\mathfrak{k}}$ be the corresponding cylindrical bimodule as in Lemma 6.9. We let $\mathring{\mathfrak{k}}'$ be the reflection of this bimodule through a horizontal line.

The Morita equivalence we require is the given by a cylindrical version of the cup bimodules, which is a $\mathring{R}^{\mathbf{j}}$-$\mathring{R}^{\mathbf{j}'}$ bimodule. We consider the rolled version of this bimodule over the planar KLRW algebra. By the basis theorem [Web17a, Lemma 7.17], we see that this module is killed by $e_z \in \mathring{R}^{\mathbf{j}}$, so the left $\mathring{R}^{\mathbf{j}}$-module structure on the cup bimodule factors through the quotient $R^{(z)} = \mathring{R}^{\mathbf{j}}/\mathring{R}^{\mathbf{j}} e_z \mathring{R}^{\mathbf{j}}$.

We now show that this bimodule induces a Morita equivalence. The difficult step in this is to show that it forms a Morita correspondence with the bimodule $\mathring{\mathfrak{k}}'$ given by reflecting this diagrams in the horizontal axis (and keeping all relations the same).

That, is, we need to define maps

$$\alpha\colon \mathring{\mathfrak{k}}' \otimes_{\mathring{R}^{\mathbf{j}}} \mathring{\mathfrak{k}} \to \mathring{R}^{\mathbf{j}'} \qquad \omega\colon \mathring{\mathfrak{k}} \otimes_{\mathring{R}^{\mathbf{j}'}} \mathring{\mathfrak{k}}' \to \mathring{R}^{\mathbf{j}}$$

that make the matrix space $\begin{bmatrix} \mathring{R}^{\mathbf{j}} & \mathring{\mathfrak{k}} \\ \mathring{\mathfrak{k}}' & \mathring{R}^{\mathbf{j}'} \end{bmatrix}$ into an associative algebra.

Elements of the simple module L_0 whose elements sit at the vertex of the cup and cap; we are thinking of this as a left module, but in $\mathring{\mathfrak{k}}'$, we've reflected, and thus have elements of the right module $\dot{L}_0$, using the notation of [Web17a, §5.2]. This is isomorphic to the dual of L_0 by [Web17a, Prop. 5.11], and so we can choose a non-degenerate pairing $\alpha\colon L_0 \otimes L_0 \to \Bbbk$ which realizes this duality, satisfying $\alpha(av, v') = \alpha(v, \dot{a}v')$. The map α on a pair of diagrams is defined by stacking the diagrams, which creates a closed circle, which we can simplify to avoid any strands outside it. We can also simplify so the circle has an element v at the top, an element v' and only straight strands between; in this case, we delete the circle, and multiply the diagram by the scalar $\alpha(v, v')$.

In order to define ω, we have to note that if we consider the planar KLRW algebra T with $j, j*$ as labels, z black strands of any label, and consider the quotient T/Te_zT by every idempotent with black strands outside the two red, then L_0 is the unique simple module that factors through this quotient, since a simple only factors through this quotient if it is highest weight for the categorical action, and L_0 is the only such module with this number of black strands. Furthermore, [Web17a, Lemma 7.3] shows that L_0 is projective over this quotient: the standard module S appearing that result is the quotient by all idempotent where a black strand is left of all reds, and that result show that the kernel of the map $S \to L_0$ is generated by the image of e_z. Thus, T/Te_zT is a matrix algebra so we have an isomorphism

$$T/Te_zT \cong L_0 \otimes_{\Bbbk} \dot{L}_0 \cong \mathrm{End}_{\Bbbk}(L_0).$$

Thus, we can define the map ω by using the relations [Web17a, §7.5-6] to remove any strands from between the cup and cap, and then replacing the pair of elements v at the bottom of the cup and v' at the top of the cap with the linear map $w \mapsto \alpha(v', w)v$, thought of as an element T/Te_zT. This shows that we have a Morita correspondence. It show that this gives a Morita equivalence, we need only show that the maps α and ω are surjective, and it's enough to show that 1 is in the image of both maps. This is easy to see from the non-degeneracy of the pairing α. □

Proof of Thm. 7.11. To prove that we have an action of affine ribbon tangles, we use the annular version of the formalism of [Oht02, Ch. 3]: we define a sliced affine tangle diagram to be such a diagram where no pair where the elements are crossings, minima or maxima of the tangle are at the same height, and where we have cut the tangle into the corresponding pieces. The assignment in the theorem gives a well-defined functor for each sliced affine tangle, and so we need only show that any two ways of doing the slicing will result in the same functor. By [Oht02, Th. 3.1], this requires showing the ribbon Turaev moves: the ribbon Reidemeister moves, the S-move, the pitchfork move, and the commutation of distant tangles. Each one of these is proven in the planar case in [Web17a, Th. 8.6], and thus follows in the annular case by Lemma 6.9. □

Proof of Thm. 7.12. Note that when we have no red or black strands, the algebra $\mathring{R}^{\emptyset}_{\mathbf{0}}$ is just $\mathbb{C}$; in this case, the planar and cylindrical KLRW algebras coincide. If a link arises from an inclusion of the 3-ball, then it can be presented as a sliced tangle in $\mathbb{R}^2 \times [0, 1]$, and then hit with the map compactifying one of the $\mathbb{R}$-directions to S^1.

Its value $\Phi(K)$ on a link is the affinization of the functor $\Phi_{\mathsf{L}}(K)$, since this is true for each individual slice, and Lemma 6.9 implies the compatibility of affinization with composition. Since the planar and cylindrical KLRW algebras are the same in the source and target categories, the functor is tensor product with the same vector spaces in either case, and the theories coincide. The correspondence with other knot homologies follows from [MW18, Th. A]. □

Appendix A. Slodowy slices in type A

One particularly interesting special case of the constructions we have discussed are the **S3 varieties for $\mathfrak{sl}_n$**. These are resolutions of the intersections of Slodowy slices and nilpotent orbits in $\mathfrak{sl}_n$. Every one of these varieties can be written as a Nakajima quiver variety and as an affine Grassmannian slice (both in type A). That is, they have a realization both as Higgs and as Coulomb branches of quiver gauge theories.

Let us remind the reader of the combinatorics underlying this realization. Given a partition $\lambda = (\lambda_1 \geq \lambda_2 \geq \cdots)$ of N with n parts, we can consider λ as a (co)weight of $\mathfrak{sl}_n$, in the usual way. Given μ, another partition of N, we let

$$w_i = \lambda_i - \lambda_{i+1} \qquad v_i = \sum_{k=1}^{i} \lambda_k - \mu_k.$$

The significance of these are more easily seen from the familiar formulae

$$\lambda = \sum_{i=1}^{n} w_i \omega_i \qquad \mu = \lambda - \sum_{i=1}^{n} v_i \alpha_i.$$

Consider the S3 variety $\mathfrak{X}^{\lambda}_{\mu}$ given by the slice to nilpotent matrices of Jordan type μ in the closure of those of Jordan type λ. Note that all symplectic resolutions of the nilpotent orbit closures with generic Jordan type λ are of the form $T^*G/P \to \bar{\mathscr{O}}$ for P the parabolic of block upper-triangular matrices with block sizes given by the transpose partition to λ. Different orders of block sizes can potentially give non-isomorphic resolutions; one can check w_i as defined above is the number of blocks of size i in along the diagonal. All resolutions of S3 varieties are given by the fiber product of T^*G/P with a slice to the orbit with Jordan type μ.

Theorem A.1 ([MV07, Th. 1.2], [BFN, Th. 5.6]) *The S3 variety $\mathfrak{X}^{\lambda}_{\mu}$ is isomorphic to the affine Grassmannian slice to Gr^{μ} inside $\mathrm{Gr}^{\bar{\lambda}}$, that is, to the Coulomb branch $\mathfrak{M}$ of the quiver gauge theory with dimension vectors $\mathbf{w}$ and $\mathbf{v}$. The resolution $\tilde{\mathfrak{M}}$ attached to a cocharacter $\xi\colon \mathbb{C}^* \to G_{\mathbf{w}}$ is isomorphic to the convolution resolution of $\mathrm{Gr}^{\bar{\lambda}}_{\mu}$ with order on fundamental coweights induced by the cocharacter ξ, and to the resolution $T^*G/P \times_{\mathfrak{sl}_n^*} \mathfrak{X}^{\lambda}_{\mu}$ where ξ determines the order on blocks in P.*

Thus, Theorems D and E give us a non-commutative resolution of the S3 variety which is D-equivalent to any symplectic resolution of this variety. This is given by a cylindrical KLRW algebra with w_i red strands and v_i black strands of label i. Below we'll discuss in a bit more detail what we see in different cases of interest.

A.1. **Kleinian singularities.** The simplest special case is the Kleinian singularity $\mathbb{C}^2/(\mathbb{Z}/\ell\mathbb{Z})$. This is isomorphic to the slice to the subregular orbit of $\mathfrak{sl}_\ell$ in the full nilcone, i.e. Jordan types $\lambda = (\ell, 0)$ and $\mu = (\ell - 1, 1)$. Thus, this corresponds to the case where $w_1 = \ell$ and $v_1 = 1$. That is, we have ℓ red strands with the same label and a single black strand.

We have ℓ different idempotents depending on the position of the black strand, which we think of as positioned cyclically on a circle. The algebra of endomorphisms is generated by these idempotents, and by the degree 1 maps joining adjacent chambers by crossing the red strand:

Thus, this algebra can be written as a quotient of the path algebra of the quiver with ℓ cyclically ordered nodes and edges joining adjacent pairs of edges in both directions. The only relations needed are that the two length two paths starting and ending at a given node are equal: they are both multiplication by a single dot on the single black strand by (6.2e). Example 8.i covers the $\ell = 2$ case. In the $\ell = 3$ case, we have the quiver shown below, with the diagrams above corresponding to a single pair of edges (with the others coming from rotations of these diagrams).

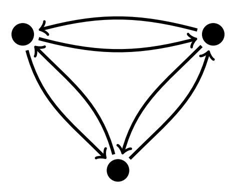

A.2. **2-row Slodowy slices.** Another case which has attracted considerable attention is that of 2-row Slodowy slices. That is, for $k \leq \ell/2$, we consider the case $\lambda = (\ell, 0)$ and $\mu = (\ell - k, k)$. Thus, we have $w_1 = \ell, v_1 = k$. The result is that all red and black strands are labeled by the same simple root. This is thus a cylindrical version of the algebras $\tilde{T}^\ell_k$ defined in [Web16, Def. 2.3].

Anno and Nandakumar [AN] show that there is an action of the category of affine tangles on the category of coherent sheaves in this case, which obviously we expect to match ours. In fact, it is virtually certain that this is the case. The same affine braid group action appears (up to reflection of braids), so we need only check that the cup and cap functors match. From the fact that the composition of a cup and then a crossing is cup shifted by ± 1 (depending on the sign of the crossing) both in our tangle action by Theorem 6.22 and in Anno-Nandakumar's action, as shown in [AN, Prop. 4.7], we see that the cup functors just have the same subcategory as image (all the eigenobjects for the crossing functor). However, this does not show that the functors are isomorphic. We leave carefully matching these functors to another time, or perhaps an industrious reader.

A.3. **Cotangent bundles to projective spaces and Grassmannians.** Dual to the examples of Kleinian singularities and 2-row Slodowy slices respectively are the cotangent bundles to projective spaces and Grassmannians.

The example of $T^*\mathbb{P}^n$ corresponds to thinking of this as the S3 variety for the minimal orbit in type A, that is, for the Jordan types $\lambda = (2, 1, \dots, 1, 0)$ and $\mu = (1, \dots, 1)$. This corresponds to the quiver gauge theory attached to a linear quiver with $n-1$ nodes, and vectors $\mathbf{w} = (1, 0, \cdots, 0, 1)$ and $\mathbf{v} = (1, \dots, 1)$. One can easily check that the associated representation is that of $G = D \cap SL_n$, the diagonal matrices of determinant 1 acting on $V = \mathbb{C}^n$, so indeed the associated Higgs branch is the Kleinian singularity $\mathbb{C}^2/(\mathbb{Z}/n\mathbb{Z})$.

To obtain the cotangent bundle to a the Grassmannian $\mathrm{Gr}(n,k)$ of k planes in $\mathbb{C}^n$ with $n \geq 2k$, we consider $\lambda = (2^k, 1^{n-2k}, 0^k)$, and $\mu = (1, \dots, 1)$. This gives $w_p = \delta_{k,p} + \delta_{n-k,p}$ and $\mathbf{v} = (1, 2, 3, \dots, k-1, k, \dots, k, k-1, \dots, 2, 1)$. For example, $\mathrm{Gr}(4, 2)$ corresponds to $\mathbf{w} = (0, 2, 0), \mathbf{v} = (1, 2, 1)$ and $\mathrm{Gr}(5, 2)$ to $\mathbf{w} = (0, 1, 1, 0), \mathbf{v} = (1, 2, 2, 1)$. We'll study these examples in more detail in forthcoming work with Zhou [WZ].

A.4. **The noncommutative Springer resolution.** One final variety of considerable interest that appears here is the cotangent bundle to the type A flag variety T^*GL_n/B. This arises from the dimension vectors $\mathbf{v} = (1, 2, 3, \dots, n-1)$ and $\mathbf{w} = (0, \cdots, 0, n)$. In this case, $\mu = \lambda^t = (1, \dots, 1)$ and $\lambda = \mu^t = (n)$. It is a well-known theorem of Nakajima that the Higgs branch of this theory is T^*G/B. The equality $\mu = \lambda^t$ shows that this example is self-dual, and the Coulomb branch arises the same way.

It's also well-known that the quantum Coulomb branch that arises this way is essentially the universal enveloping algebra of $\mathfrak{gl}_n$; if we fix the flavors to numerical values, then this is the quotient of this ring by a maximal ideal of its center, but keeping the flavors as variables, it's easy to construct $U(\mathfrak{gl}_n)$ on the nose.

The construction of a tilting generator in Section 3 is thus just a rephrasing of the noncommutative Springer resolution as constructed by Bezrukavnikov, Mirković, Rumynin and Riche [BMRR08, Bez06]. Recall that this construction operates by turning differential operators on the flag variety X in characteristic p into an Azumaya algebra $\mathcal{D}$ on T^*X, completing in a formal neighborhood of the zero-section and finding a splitting of the resulting Azumaya algebra $\hat{\mathcal{D}}$ on this formal neighborhood.

Thus, the constructions of Section 3 can be recast in this case in purely Lie theoretic terms. The ring homomorphism σ of Theorem 3.1 is just the map $\sigma(X) = X^p - X^{(p)}$ for $X \in \mathfrak{gl}_n$. The integrable system given by equivariant parameters on the Coulomb branch is precisely the Gelfand-Tsetlin system as discussed in [WWY]. Identifying the Harish-Chandra center of $U(\mathfrak{gl}_n)$ with $\mathbb{F}_p[\mathfrak{t}]^W$, the homomorphism σ is pullback by the Artin-Schreier map AS: $\mathfrak{t}^* \to \mathfrak{t}^*$ (which is W-equivariant). This describes the induced map on the full Gelfand-Tsetlin subalgebra Γ, which is the tensor product of the Harish-Chandra centers of $\mathfrak{gl}_k$ for $k = 1, \dots, n$.

That is, we can think of a point in $\mathrm{MaxSpec}(\Gamma \otimes \overline{\mathbb{F}}_p)$ as a choice $\mathbf{a}_1, \dots, \mathbf{a}_n$ with $\mathbf{a}_k$ an unordered k-tuple in $\overline{\mathbb{F}}_p$, and σ the map sending $a_{i,j} \mapsto a_{i,j}^p - a_{i,j}$. In particular, when we complete in a formal neighborhood of the zero-section, the ideal in $\sigma(\Gamma)$ generated by $a_{i,j}^p - a_{i,j} = 0$ acts topologically nilpotently.

The sections of $\mathcal{D}$ are identified with a completion of the cylindrical KLRW algebra by Lemma 3.10 and Proposition 8.6. By the Chinese remainder theorem applied to the image of Γ, this completion also contains idempotents attached to each maximal ideals lying over that in $\sigma(\Gamma)$, that is, those with $a_{i,j} \in \mathbb{F}_p$. These correspond to the idempotents in the KLR algebra, given by $e(\mathbf{a}')$ where the labels i are on the point $a_{i,j}/p \in \mathbb{R}/\mathbb{Z}$ (following Definition 6.10); as before we let $e(a/p)$ denote when we have fixed $a_{i,j} = a$ for all i, j. The image $\mathcal{D}e(a/p)$ is a splitting bundle for this Azumaya algebra by Lemma 3.11.

Thus, while we obtain a familiar object, we obtain a new perspective on it, since this KLR presentation is not at all obvious from the Lie theoretic perspective. Developing its consequences will have to wait for future work.

Glossary

T_*	A maximal torus of the group $*$.	2, 35, 49
G	The gauge group.	2, 5, 32, 34, 36, 48, 49
W	The Weyl group of G.	2, 49
Γ	The quiver used to define the quiver gauge theory as in (1.1).	2, 5, 32, 48
$\mathcal{V}(\Gamma)$	The vertex set of Γ.	2, 5, 6, 29, 30
$\mathbf{v}$	A dimension vector with components v_i for Γ that gives the ranks of the factors of the gauge group G.	2, 5, 24, 30, 33, 44, 46
$\mathbf{w}$	A dimension vector with components w_i for Γ that gives the ranks of the factors of the flavor group F.	2, 5, 24, 27, 44, 46
V	The matter representation.	2, 32, 49
$\mathfrak{M}$	The Coulomb branch—the quotient of the convolution algebra of a modified affine Grassmannian as defined in [BFN18].	3, 13, 37, 44, 48, 49
$\mathring{R}$	The cylindrical KLRW algebra; see Definition 6.3.	3, 7, 11, 13, 15, 34
$\tilde{\mathfrak{M}}$	The resolution of the Coulomb branch $\mathfrak{M}$ defined by taking symplectic redution with a GIT quotient of $\mathfrak{M}_Q$.	3, 4, 14, 34, 44, 49
$\mathring{R}^{\mathbf{j}}$	The cylindrical KLRW algebra with $\beta_e = 0$ for all e and where the word $\mathbf{j}$ defines the labels on the red strands.	4, 18, 19, 37, 42
H	The normalizer $N^\circ_{GL(V)}(G)$.	5, 23, 49
β_*	A $\mathbb{R}/\mathbb{Z}$ valued function on the edge set of Γ; these control the distance between solid and ghost strands in cylindrical KLRW diagrams.	5, 10, 13–16, 18, 24, 25, 30, 31, 34–36, 48

$\gamma_{*,*}$	A w_i-tuple of elements in $\mathbb{R}/\mathbb{Z}$ for each node in Γ. These control the position of red strands in cylindrical KLRW diagrams.	5, 13–16, 19, 24, 25, 31, 34, 48
$\mathscr{B}$	The extended BFN category, defined in Definition 2.4.	13, 23–25, 29, 31, 35, 49
$B_{\boldsymbol{\beta},\boldsymbol{\gamma}}$	The bimodule defined by twisted cylindrical KLRW diagrams for some path $\boldsymbol{\beta},\boldsymbol{\gamma}$ in the space of parameters β_*, $\gamma_{*,*}$.	13–15, 17, 36, 39
Q	The preimage in the normalizer H of a fixed torus T_F in the flavor group F.	17, 35, 49
$\mathring{R}_C$	The cylindrical KLRW algebra with parameters in the alcove C.	17, 38, 39
$Z_{\beta,\gamma}(M)$	The central charge function, depending on β_* and $\gamma_{*,*}$, defined by polynomially extending the formula (6.10) on $K^0(\mathring{R}_{C_0}\text{-fdmod})$.	17, 19, 38, 41
e_m	As we pass to a wall where two red strands collide, e_m is the sum of idempotents corresponding to loadings with $< m$ black strands between these red strands.	19, 20, 42, 43
$\mathcal{A}^{\mathrm{sph}}$	The quantum Coulomb branch $\mathcal{A}^{\mathrm{sph}} = H_*^{BM,\widetilde{G((t))}}(\mathcal{Y}_{V[[t]]} \times_{V((t))} \mathcal{Y}_{V[[t]]})$.	23, 31, 32, 38, 39, 49
δ	A parameter between the open interval $(0,1) \subset \mathbb{R}$ used in the definition of $\mathscr{B}$.	23
pth root	The conventions for the extended category adopted in Definition 2.12.	23, 35, 49
φ_i^{mid}	The average of φ_i^+ and $-\varphi_i^-$.	24
τ	The cocharacter $\tau\colon \mathbb{C}^* \to GL(T^*V)$ that acts trivially on V and weight -1 on V^*.	24, 33
$c_{*,*}$	The diagonal entries $c_{i,1},\dots,c_{i,w_i}$ of the flavor ϕ into $\mathfrak{gl}(\Bbbk^{w_i})$.	24, 27, 31, 34–37, 39
b_e	The diagonal entries b_e of the flavor ϕ into $\mathfrak{gl}(\Bbbk^{\chi_{i,j}})$.	24, 27, 31, 34–37, 39
$\widehat{W}$	The affine Weyl group of G. The semi-direct product of W and the coweight lattice of T.	25, 29, 30
S_*	The symmetric algebra on $\tilde{\mathfrak{t}}^*$, that is, the ring of functions on the affine variety $\tilde{\mathfrak{t}}$, with the parameter h specialized at $h = *$.	25, 40
$\mathtt{r}_D$	The morphism $\mathtt{r}_D\colon D(0) \to D(1)$ in the category $\mathscr{B}$ associated to an unrolled diagram described in Definition 8.2. In Definition 8.4, we extend this to cylindrical KLRW diagrams.	26, 29, 30
$\widehat{\mathsf{A}}_p$	The completion with respect to grading of A_p.	33
B	The extended BFN category with pth root conventions.	33, 36, 49
F	The flavor group H/G.	35, 48
${}_{\phi+\nu}\mathscr{T}_{\phi}$	The $\mathscr{B}_{\phi+\nu}$-$\mathscr{B}_{\phi}$ bimodule formed by the appropriate quotient of $\mathscr{T}(\nu)$, the morphisms of weight ν in $\mathscr{B}^Q$.	35, 36
$\mathscr{B}^Q$	The extended BFN category attached to the pair (Q,V), defined in Definition 2.4.	35

$\mathfrak{t}_{1,\mathbb{R}}$	Let $\mathfrak{t}_{1,\mathbb{R}} \subset \tilde{\mathfrak{t}}_{\mathbb{R}} = d\nu^{-1}(1)$ be the real lifts of the cocharacter ϕ.	36
$\hat{\mathscr{W}}_\phi$	The restriction of $\mathscr{W}_\phi$ to the formal neighborhood $\hat{\mathfrak{M}}$.	37
$\mathfrak{M}_Q$	The Coulomb branch $\mathfrak{M}$ attached to the group Q acting on V with its usual action.	48
A_p	The subcategory of B defined in Definition 2.20.	49
$\mathscr{W}_\phi$	The coherent sheaf of generically Azumaya algebras on $\tilde{\mathfrak{M}}$ or its pushforward to $\mathfrak{M}$, such that $\Gamma(\tilde{\mathfrak{M}}; \mathscr{W}_\phi) = \mathcal{A}_1^{\mathrm{sph}}$ with the quantization parameter ϕ; see Definition 3.5.	49
$\hat{\mathfrak{M}}$	The formal neighborhood in $\mathfrak{M}$ of the fiber over the origin in $\mathfrak{t}/W$.	49

References

[ABM15] Rina Anno, Roman Bezrukavnikov, and Ivan Mirković, *Stability conditions for Slodowy slices and real variations of stability*, Mosc. Math. J. **15** (2015), no. 2, 187–203, 403. MR 3427420

[Aga] Mina Aganagić, *Knot categorification from mirror symmetry, part i: Coherent sheaves*, arXiv:2004.14518.

[AN] Rina Anno and Vinoth Nandakumar, *Exotic t-structures for two-block springer fibers*, arXiv:1602.00768.

[Bas68] Hyman Bass, *Algebraic K-theory*, W. A. Benjamin, Inc., New York-Amsterdam, 1968. MR 0249491

[Bez06] R. Bezrukavnikov, *Noncommutative counterparts of the Springer resolution*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1119–1144.

[BFN] Alexander Braverman, Michael Finkelberg, and Hiraku Nakajima, *Line bundles over coulomb branches*, arXiv:1805.11826.

[BFN18] Alexander Braverman, Michael Finkelberg, and Hiraku Nakajima, *Towards a mathematical definition of Coulomb branches of 3-dimensional $\mathcal{N} = 4$ gauge theories, II*, Adv. Theor. Math. Phys. **22** (2018), no. 5, 1071–1147. MR 3952347

[BK08] R. Bezrukavnikov and D. Kaledin, *Fedosov quantization in positive characteristic*, J. Amer. Math. Soc. **21** (2008), no. 2, 409–438. MR 2373355

[BMRR08] Roman Bezrukavnikov, Ivan Mirković, Dmitriy Rumynin, and Simon Riche, *Localization of modules for a semisimple lie algebra in prime characteristic*, Annals of Mathematics (2008), 945–991.

[Kal08] Dmitry Kaledin, *Derived equivalences by quantization*, Geom. Funct. Anal. **17** (2008), no. 6, 1968–2004.

[KL09] Mikhail Khovanov and Aaron D. Lauda, *A diagrammatic approach to categorification of quantum groups. I*, Represent. Theory **13** (2009), 309–347.

[KL11] ———, *A diagrammatic approach to categorification of quantum groups II*, Trans. Amer. Math. Soc. **363** (2011), no. 5, 2685–2700. MR 2763732 (2012a:17021)

[KTW+] Joel Kamnitzer, Peter Tingley, Ben Webster, Alex Weekes, and Oded Yacobi, *On category $\mathcal{O}$ for affine Grassmannian slices and categorified tensor products*, arXiv:1806.07519.

[Los] Ivan Losev, *On modular categories $\mathcal{O}$ for quantized symplectic resolutions*, arXiv:1712.07726.

[MV07] I. Mirković and M. Vybornov, *Quiver varieties and Beilinson-Drinfeld Grassmannians of type A*, 2007, arXiv:0712.4160.

[MW18] Marco Mackaay and Ben Webster, *Categorified skew Howe duality and comparison of knot homologies*, 2018, pp. 876–945.

[Nam08] Yoshinori Namikawa, *Flops and Poisson deformations of symplectic varieties*, Publ. Res. Inst. Math. Sci. **44** (2008), no. 2, 259–314.

[NZ] Vinoth Nandakumar and Gufang Zhao, *Categorification via blocks of modular representations for $\mathfrak{sl}(n)$*, `arXiv:1612.06941`.

[Oht02] Tomotada Ohtsuki, *Quantum invariants*, Series on Knots and Everything, vol. 29, World Scientific Publishing Co. Inc., River Edge, NJ, 2002, A study of knots, 3-manifolds, and their sets. MR MR1881401 (2003f:57027)

[QR18] Hoel Queffelec and David E. V. Rose, *Sutured annular Khovanov-Rozansky homology*, Trans. Amer. Math. Soc. **370** (2018), no. 2, 1285–1319. MR 3729501

[Weba] Ben Webster, *Coherent sheaves and quantum Coulomb branches I: tilting bundles from integrable systems*, `arXiv:1905.04623`.

[Webb] ———, *Gelfand-Tsetlin modules in the Coulomb context*, `arXiv:1904.05415`.

[Webc] ———, *Koszul duality between Higgs and Coulomb categories $\mathcal{O}$*, `arXiv:1611.06541`.

[Web16] ———, *Tensor product algebras, Grassmannians and Khovanov homology*, Physics and mathematics of link homology, Contemp. Math., vol. 680, Amer. Math. Soc., Providence, RI, 2016, pp. 23–58. MR 3591642

[Web17a] ———, *Knot invariants and higher representation theory*, Mem. Amer. Math. Soc. **250** (2017), no. 1191, 141.

[Web17b] ———, *Rouquier's conjecture and diagrammatic algebra*, Forum Math. Sigma **5** (2017), e27, 71. MR 3732238

[Web19a] ———, *Representation theory of the cyclotomic Cherednik algebra via the Dunkl-Opdam subalgebra*, New York J. Math. **25** (2019), 1017–1047. MR 4017214

[Web19b] ———, *Weighted Khovanov-Lauda-Rouquier algebras*, Doc. Math. **24** (2019), 209–250. MR 3946709

[Weea] Alex Weekes, *Generators for Coulomb branches of quiver gauge theories*, `arXiv:1903.07734`.

[Weeb] ———, *Quiver gauge theories and symplectic singularities.*

[WWY] Ben Webster, Alex Weekes, and Oded Yacobi, *A quantum Mirković–Vybornov isomorphism*, `arXiv:1706.03841`.

[WZ] Ben Webster and Baorui Zhou, forthcoming.