

SKEIN TRACES FROM CURVE COUNTING

TOBIAS EKHOLM, PIETRO LONGHI, SUNGHYUK PARK, AND VIVEK SHENDE

ABSTRACT. Given a 3-manifold M , and a branched cover arising from the projection of a Lagrangian 3-manifold L in the cotangent bundle of M to the zero-section, we define a map from the skein of M to the skein of L , via the skein-valued counting of holomorphic curves. When M and L are products of surfaces and intervals, we show that wall crossings in the space of the branched covers obey a skein-valued lift of the Kontsevich-Soibelman wall-crossing formula.

Holomorphic curves in cotangent bundles correspond to Morse flow graphs; in the case of branched double covers, this allows us to give an explicit formula for the skein trace. After specializing to the case where M is a surface times an interval, and additionally specializing the HOMFLYPT skein to the $\mathfrak{gl}(2)$ skein on M and the $\mathfrak{gl}(1)$ skein on L , we recover an existing prescription of Neitzke and Yan.

1.	INTRODUCTION	3
2.	THE SKEIN DILOGARITHM	6
2.1.	Skein of the solid torus	6
2.2.	Skein dilogarithm	6
3.	BRANCHED DOUBLE COVERS FROM TRIANGULATIONS	7
3.1.	Smooth covers from marked triangulations	7
3.2.	Singular branched double covers	8
3.3.	Angle structures	10
3.4.	An effectivity condition	11
3.5.	HOMFLYPT skein modules of singular branched double covers	15
3.6.	Specializing the singular skein	16
4.	A SKEIN TRACE FOR DOUBLE COVERS OF 3-MANIFOLDS	19
4.1.	Turning numbers	19
4.2.	Proof of Theorem 4.1	22
4.3.	The trivial double cover	25

TE is supported by the Knut and Alice Wallenberg Foundation, KAW2020.0307 Wallenberg Scholar and by the Swedish Research Council, VR 2022-06593, Centre of Excellence in Geometry and Physics at Uppsala University and VR 2024-04417, project grant.

The work of P.L. is supported by the Knut and Alice Wallenberg Foundation, KAW2020.0307 Wallenberg Scholar and by the Swedish Research Council, VR 2022-06593, Centre of Excellence in Geometry and Physics at Uppsala University, and he records the preprint number UUITP-30/25.

S.P. gratefully acknowledges support from Simons Foundation through Simons Collaboration on Global Categorical Symmetries.

V.S. is supported by Villum Fonden Villum Investigator grant 37814, Novo Nordisk Foundation grant NNF20OC0066298, and Danish National Research Foundation grant DNRF157.

5.	CONORMALS	28
6.	FLOW GRAPHS AND THE FIRST SKEIN RELATION	32
6.1.	Flow graph review	32
6.2.	Description of the result	33
6.3.	A choice of perturbation	33
6.4.	Flow loops	35
6.5.	Flow graphs with corners for $\tilde{U}(0, \rho)$	36
6.6.	Flow graphs for $\tilde{U}(\sigma, \rho)$, $\sigma \neq 0$	38
7.	REVIEW OF SKEIN-VALUED CURVE COUNTING	41
7.1.	Spin structures – remarks and extensions	41
7.2.	4-chains – remarks and extensions	42
7.3.	Basic classes	42
7.4.	Normalization of the basic cylinder	43
8.	SKEIN TRACE FROM CURVE COUNTING	43
8.1.	4-chains for branched covers	44
8.2.	Proof of Theorem 1.3	45
8.2.1.	First skein relation	45
8.2.2.	Second skein relation	47
8.2.3.	Third skein relation	48
8.3.	A Symplectic Field Theory perspective	49
8.3.1.	Sketch of another proof	49
8.3.2.	The skein trace map and Lagrangian cabling	50
9.	BRANCHED DOUBLE COVERS, SKEIN COUNTS, AND FLOW GRAPHS	51
9.1.	Lagrangian double covers away from the cone points	51
9.2.	Smoothing the cones	52
9.3.	4-chains for branched double covers	53
9.4.	Intersections of holomorphic curves with the 4-chain	55
9.5.	Spin structures	59
9.6.	Flow graphs on a branched double cover	59
9.7.	3d spectral network as a moduli space of holomorphic disks	62
10.	WALL CROSSINGS	63
10.1.	Lagrangian isotopies and proof of Theorem 1.6	64
10.2.	Transverse (-1) -disk and the proof of Theorem 1.7	65
11.	EXAMPLES OF SKEIN-VALUED WALL-CROSSING FORMULAS	66
11.1.	The pentagon relation	67
11.2.	The Seiberg-Witten wall-crossing	69
	APPENDIX A. FLOW GRAPHS AND HOLOMORPHIC CURVES	72
A.1.	The flow graph limit	72
A.2.	Flow graphs attached to curves	73
	REFERENCES	75

1. INTRODUCTION

The use of what would later be called ‘cluster coordinates’ on the moduli space of $\mathrm{SL}(2, \mathbb{C})$ local systems on a Riemann surface S dates back at least to the work of Penner [55]; it is by now a part of the ‘higher Teichmüller theory’ of Fock and Goncharov [23], which treats local systems of arbitrary rank n . It was subsequently discovered by Gaiotto, Moore, and Neitzke [30, 29] that the resulting coordinate charts can be identified with the moduli of abelian local systems on a spectral curve $\Sigma \subset T^*S$, whose projection to the zero section is a degree n branched cover $\pi: \Sigma \rightarrow S$. The coordinate systems have at least two interpretations, indicated already in [30], either as governing Stokes phenomena in the exact WKB analysis of the Schrödinger equation on a Riemann surface [38, 3, 40], or in terms of certain moduli of objects in the Fukaya category [63], especially as calculated in the flow-tree limit [10] by [51, 7].

Fock and Goncharov have shown that cluster varieties in general can be q -deformed by q -deforming their charts [24]. On the other hand, Turaev had previously established that moduli of rank n local systems on S admit q -deformations to the $\mathfrak{gl}(n)$ skein modules [67]. Here, we recall that given a 3-manifold M , its HOMFLYPT skein $\mathrm{Sk}_{a,z}(M)$ is the $\mathbb{Z}[a^\pm, z^\pm]$ -module generated by framed links in M , modulo the following relations:

$$(1.1) \quad \begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} - \begin{array}{c} \diagdown \diagup \\ \diagup \diagdown \end{array} = z \begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array},$$

$$(1.2) \quad a \begin{array}{c} \circlearrowleft \end{array} - a^{-1} \begin{array}{c} \circlearrowright \end{array} = z \begin{array}{c} \circlearrowright \end{array},$$

$$(1.3) \quad \begin{array}{c} \uparrow \\ \curvearrowright \end{array} = a \begin{array}{c} \uparrow \end{array}.$$

We will always set $z = q^{1/2} - q^{-1/2}$ and extend scalars to $\mathbb{Z}[q^{\pm 1/2}, (1-q)^{-1}, (1-q^2)^{-1}, \dots]$. The $\mathfrak{gl}(n)$ skein modules are the corresponding $\mathbb{Z}[q^{\pm 1/2}]$ -modules arising from the specialization of the above relations at $a = q^{n/2}$.¹

For the Fock-Goncharov deformation to agree with that of Turaev, there must be corresponding maps from the $\mathfrak{gl}(n)$ skein of S to the $\mathfrak{gl}(1)$ skein of Σ , as indeed Fock conjectured explicitly [22]. Such a map was constructed by Bonahon and Wong [6] for $n = 2$ and termed a *quantum trace*; their methods have been pursued further by various authors [42, 53, 32]. Neitzke and Yan [49, 50], building on earlier work of Gabella [27], showed that a q -deformation of the Gaiotto-Moore-Neitzke formulas would also provide such maps; this construction was later shown to recover the Bonahon-Wong map [54]. We will give in Theorem 4.2 a generalization of the Neitzke-Yan formulas, both to more general 3-manifolds (rather than surface times interval) and to the HOMFLYPT skein on both base and cover (rather than $\mathfrak{gl}(2)$ and $\mathfrak{gl}(1)$ skeins, respectively).

Our main purpose in the present article is to give a geometric construction of these and more general skein traces, by counting holomorphic curves.

¹To be precise, the $\mathfrak{gl}(n)$ skein modules are quotients of such specializations that kill, e.g., all the λ -colored strands, for any Young diagram λ with more than n rows.

Let us recall from [19, 17, 16] that counting holomorphic curves with (Maslov zero) Lagrangian boundary conditions in Calabi-Yau 3-folds by the class of their boundary in the HOMFLYPT skein modules of the Lagrangians lead to invariant quantities. This is because it is possible to match the wall-crossings in parameter space corresponding to boundaries of moduli space with the HOMFLYPT skein relations.

Consider a 3-manifold M , possibly noncompact with cylindrical ends $\partial_\infty M \times \mathbb{R}_{>0}$. We fix $L \subset T^*M$; if M is noncompact, we ask that L is cylindrical in the cylindrical end, i.e., there is a Lagrangian $\partial_\infty L \subset T^*\partial_\infty M$ and, in the cylindrical end, L takes the form $\partial_\infty L \times \mathbb{R}_{>0} \subset T^*\partial_\infty M \times T^*\mathbb{R}_{>0}$.

Consider a knot or link $K \subset M$ and its conormal $N^*K \subset T^*M$. We identify M with the 0-section in T^*M , $M \subset T^*$; then $N^*K \cap M = K$. We may shift N^*K along the 1-form dual to the tangent of K . This is a (non-Hamiltonian) symplectic isotopy which takes N^*K to a non-exact Lagrangian $N^\circ K$ off the 0-section, $N^\circ K \cap M = \emptyset$. In the geometry $(T^*M; M \sqcup N^\circ K)$, there is a single bounded holomorphic curve, namely the trace of K along the isotopy. (When $M = S^3$, this geometry motivated the Ooguri-Vafa conjecture expressing the HOMFLYPT invariant of a knot in as a count of curves in the holomorphic curves in the resolved conifold [52], and its proof [19, 15].) We may shift further so that $N^\circ K$ sits *outside* the unit cotangent bundle; in particular, so that it is disjoint from L .

After such further shift, we count curves in the geometry $(T^*M; L \sqcup N^\circ K)$. In particular, if we restrict attention to those which go once positively around the longitude $K \subset N^\circ K$ (or count all curves and take the appropriate coefficient with respect to a basis of $\text{Sk}(N^\circ K)$, see Section 7.3) we obtain an element of $\text{Sk}(L)$.

Definition 1.1. We write $[K]_L \in \text{Sk}(L)$ for the count of curves in $(T^*M; L \sqcup N^\circ K)$ which have boundary going once positively around the longitude k of the solid torus $N^\circ K$.

Remark 1.2. This formulation is similar to the proposal of [49, Sec. 1.6.2-3]. One difference is that we push the conormal off the zero section. Another is that we work in the HOMFLYPT skein on both L and M . Yet another is that the skein-valued curve counting is most immediately expressed in the variable $z = e^{g_s/2} - e^{-g_s/2}$, see [16, Theorem 1.4], rather than g_s , and in the variable a rather than $e^{Ng_s/2}$, see [15]. All of these may be related to the assertion in loc. cit. that those authors work “‘before’ the conifold transition”.

We understand the Lagrangians appearing in Definition 1.1 as coming equipped with certain additional “brane data”. This data includes – as always for studying holomorphic curves with boundary – a choice of spin structure. Related issues lead us to introduce certain *sign lines* into our skeins; these being some fixed 1-cycles in the 3-manifold which our links should avoid, but may cross at the cost of changing sign.

$$(1.4) \quad \begin{array}{c} \circ \\ \text{---} \uparrow \text{---} \\ \circ \end{array} = (-1) \cdot \begin{array}{c} \circ \\ \uparrow \text{---} \\ \circ \end{array},$$

The effect on the isomorphism class of the skein module only depends on the class of the sign lines in $H^1(M, \mathbb{Z}/2\mathbb{Z})$. These lines correspond to the appearance of twisted local systems in [30], and appeared previously in a skein-valued curve-counting setting in [59]. Brane data also includes, uniquely to the skein-valued curve counting, a choice of a 4-chain $W \subset T^*M$ where ∂W is twice the Lagrangian of interest [19]. We discuss appropriate choices for these data in Section 8.1 below.

We show that $[\cdot]_L$ defines a skein trace:

Theorem 1.3. *The map $K \mapsto [K]_L$ factors through $\text{Sk}(M)|_{a_M=a_L^n}$, hence defining a map*

$$\begin{aligned} \text{Sk}(M)|_{a_M=a_L^n} &\rightarrow \text{Sk}(L) \\ K &\mapsto [K]_L \end{aligned}$$

This result does not follow immediately from the well-definedness of the skein-valued curve counting, and requires us to compare the counts of curves ending on conormals to knots which are related by a skein relation. We do so by studying holomorphic curves in the Morse flow graph limit of [10]; the proof of Theorem 1.3 is given in Section 8.2.

In general, finding higher genus Morse flow graphs is a difficult problem, but when the cover has degree two, it simplifies enough that we can solve it explicitly, showing:

Theorem 1.4 (Theorem 9.13). *When $L \rightarrow M$ has degree two, the skein trace is computed by an explicit finite sum of graphs, recovering the prescription of Neitzke and Yan [49], and, more generally, our Theorem 4.2.*

For a Riemann surface S , we write $\text{Sk}(S) := \text{Sk}(S \times \mathbb{R})$. Recall that concatenation of the \mathbb{R} factor makes $\text{Sk}(S)$ an algebra, and, similarly, if M is a 3-manifold and S is a component of its boundary, then $\text{Sk}(S)$ acts on $\text{Sk}(M)$.

Fix $\Sigma \subset T^*S$ a Lagrangian which bounds no (non-constant) holomorphic curves, which can be arranged by choosing Σ either to be an exact Lagrangian with Legendrian ends as in [63]², or, as in the original [30], considering holomorphic S and Σ , and, if Ω is the holomorphic symplectic form on T^*S , taking the real symplectic form to be $\text{Re}(e^{i\theta}\Omega)$ where θ is generic. The following structure is immediate:

Theorem 1.5. *Suppose $\Sigma \subset T^*S$ bounds no nonconstant holomorphic curves and $\Sigma \rightarrow S$ is a degree n cover. Then*

$$[\cdot]_\Sigma : \text{Sk}_{a^n}(S) \rightarrow \text{Sk}_a(\Sigma)$$

is an algebra homomorphism.

*More generally, if M is a 3-manifold, S is a component of ∂M , and $L \subset T^*M$ is a Lagrangian eventually cylindrical on $\Sigma \subset T^*S$ then the map $\text{Sk}_{a^n}(M) \rightarrow \text{Sk}_a(L)$ is a map of $\text{Sk}_{a^n}(S)$ modules, where $\text{Sk}_{a^n}(S)$ acts on $\text{Sk}_a(L)$ through the algebra map $[\cdot]_\Sigma$.*

Proof. The algebra structure concerns knots in $S \times \mathbb{R}$ whose projections to the \mathbb{R} direction are disjoint. The holomorphic curves involved in determining $[\cdot]_\Sigma$ then necessarily have correspondingly disjoint projections to $T^*\mathbb{R}$. \square

In fact, letting (totally real) traces of sufficiently slow Lagrangian isotopies play the role of the Lagrangian L in Theorem 1.5, we also establish the following ‘wall crossing’ formula in Section 10.1:

Theorem 1.6. *Let $\phi_t : \Sigma \rightarrow T^*S$ be a path of Lagrangians interpolating between some Σ_- and Σ_+ . Then there is an invertible element $\Omega(\phi) \in \text{Sk}(\Sigma)$ such that, for every $[K] \in \text{Sk}(S)$,*

$$(1.5) \quad [K]_{\Sigma_-} \cdot \Omega(\phi) = \Omega(\phi) \cdot [K]_{\Sigma_+}$$

Moreover, $\Omega(\phi)$ is invariant under homotopy of the path ϕ with fixed endpoints.

²In [63], the Legendrian ends were drawn in the cosphere bundle of a surface S . However, those ends projected to a neighborhood of certain marked points on S ; if we puncture S at these points to get S° and view T^*S° as having convex contact boundary, then we may isotope the ends to lie in the ‘vertical’ part of the boundary, compatibly with our current conventions.

Furthermore, for the simplest such paths, we determine $\Omega(\phi)$ explicitly in Section 10.2. The result is formulated in terms of the exponentiated skein dilogarithm, which is a certain lift of the q -dilogarithm of Fock and Goncharov from the $\mathfrak{gl}(1)$ skein to the HOMFLYPT skein, see Definition 2.1. The skein dilogarithm first appeared as the skein-valued count of curves ending on an Aganagic-Vafa brane in \mathbb{C}^3 [18].

Theorem 1.7. *If $\phi_t(\Sigma)$ bounds no holomorphic disks away from $t = 0$, and $\phi_0(\Sigma)$ bounds a single embedded index -1 holomorphic disk D , which is a transversely cut out holomorphic disk instance in the 1-parameter family, see Section 10.2 for details, then $\Omega(\phi)$ is the exponentiated skein dilogarithm on ∂D .*

Theorem 1.7 is the skein-quantized version of the fact that the simplest wall crossings of the charts $\text{Loc}_1(\Sigma) \rightarrow \text{Loc}_n(S)$ in [30] or [63] are cluster transformations. A related result, showing that a certain algebra of boundary operators $A(\Sigma)$ transforms similarly, has previously appeared in [60, 36].

2. THE SKEIN DILOGARITHM

2.1. Skein of the solid torus. Consider a solid torus, $D^2 \times S^1$. Choose an orientation of the longitude, hence distinguishing an oriented meridian $P_{1,0}$ of the boundary which links the longitude positively. According to [34], $P_{1,0}$ acts diagonally with distinct eigenvalues, one for each pair of partitions (λ, μ) equal to

$$\frac{a - a^{-1}}{q^{1/2} - q^{-1/2}} + (q^{1/2} - q^{-1/2}) (aC_\lambda(q) - a^{-1}C_\mu(q^{-1})),$$

where C_λ denotes the content polynomial of λ , on the $\text{Sk}(D^2 \times S^1)$ and correspondingly determines, up to scalar multiples, a basis $W_{\lambda,\mu}$ indexed by pairs of partitions λ, μ . A choice of framing of the longitude determines the scalars by requiring that, when completing the torus to the three-sphere, $W_{\lambda,\mu}$ goes to the appropriate colored HOMFLYPT invariant of the longitude-turned-unknot with the given framing. Note that the framing of the longitude also determines a choice of longitude in the boundary of the solid torus. Thus for $\Psi \in \text{Sk}(D^2 \times S^1)$, we may expand:

$$(2.1) \quad \Psi = \sum \psi_{\lambda,\mu} W_{\lambda,\mu}$$

The $W_\lambda := W_{\lambda,\emptyset}$ are known to span the part of the skein generated by links whose tangent in the S^1 direction is everywhere positive; we refer to this as the *positive skein* and denote it $\text{Sk}_+(D^2 \times S^1)$. We write $\widehat{\text{Sk}}(D^2 \times S^1)$ for the completion in which we allow infinite sums of links, so long as for any N , there are only finitely many whose winding around S^1 is bounded above by N .

We will later want to consider the $a = 1$ specialization of $\text{Sk}_+(D^2 \times S^1)$. As the eigenvalues of the meridian operator among the W_λ remain distinct, these remain linearly independent.

2.2. Skein dilogarithm. Let $P_{1,0}$ and $P_{0,1}$ be the meridian and longitude of the boundary of the solid torus.

Definition 2.1. The *exponentiated skein dilogarithm* $\Psi[\xi] \in \mathbb{Q}[\xi] \otimes \widehat{\text{Sk}}_+(D^2 \times S^1)$ is the unique solution to the 3-term recurrence relation

$$(2.2) \quad (\bigcirc - P_{1,0} - a\xi P_{0,1})\Psi[\xi] = 0$$

of the form $\Psi[\xi] = 1 + \dots$, i.e. the coefficient of W_\emptyset is 1.

Explicitly,

$$(2.3) \quad \Psi[\xi] := \sum_{\lambda} \prod_{\square \in \lambda} \frac{-q^{-c(\square)/2} \xi}{q^{h(\square)/2} - q^{-h(\square)/2}} W_{\lambda} = \exp \left(- \sum_d \frac{1}{d} \frac{\xi^d P_d}{q^{d/2} - q^{-d/2}} \right).$$

For the existence and uniqueness of the solution, and the first expression for it, see [18]. In the second expression, the P_d are the images of the power sum symmetric functions under the identification of symmetric functions and $\text{Sk}_+(D^2 \times S^1)$ determined by identifying W_{λ} with the corresponding Schur function [2, 44]. The equality of the two formulas can be derived by manipulation of symmetric functions, see e.g. [48, Theorem 1.1 (d)]. As is customary in the literature, we will often omit the word ‘exponentiated’; we always mean $\Psi[\xi]$, rather than its logarithm. Also, we will simply write Ψ for $\Psi[1]$.

The skein dilogarithm satisfies a relative version of the 3-term recurrence relation (2.2) in a thickened annulus with 2 marked points on the boundary, see [18] or [48, Equation (104)]:

$$(2.4) \quad \begin{array}{c} \text{Diagram 1} \end{array} \Psi = \begin{array}{c} \text{Diagram 2} \end{array} \Psi + \begin{array}{c} \text{Diagram 3} \end{array} \Psi .$$

The skein dilogarithm has an inverse:

$$(2.5) \quad \Psi^{-1} = \sum_{\lambda} \prod_{\square \in \lambda} \frac{q^{c(\square)/2}}{q^{h(\square)/2} - q^{-h(\square)/2}} W_{\lambda}$$

which satisfies an analogous 3-term recurrence relation

$$(2.6) \quad (\bigcirc - P_{-1,0} - a^{-1} P_{0,1}) \Psi^{-1} = 0$$

where $P_{-1,0}$ is the meridian with opposite orientation, and the corresponding relative version.

3. BRANCHED DOUBLE COVERS FROM TRIANGULATIONS

In this section we recall how to associate smooth double covers to certain decorated ideal triangulations of 3-manifolds, and how these can be understood as smoothings of certain singular double covers depending on the triangulation alone. Then we define a skein module associated to the singular double cover, and identify conditions under which it maps to the skein module of a smoothing.

3.1. Smooth covers from marked triangulations. Given any manifold M containing a cooriented codimension two submanifold $\sigma \subset M$, there is a canonical double cover $\widetilde{M} \rightarrow M$ branched over σ . In this context we will write $\widetilde{\sigma}$ for the preimage of σ .

For triangulated 3-manifolds M , we will define branch loci $\sigma \subset M$ by a patching local pieces in each tetrahedron of the triangulation. Recall that an ideal tetrahedron is a tetrahedron minus its vertices, and an ideal triangulation Δ of a noncompact 3-manifold M is a stratification of M by ideal tetrahedra. We will use the notation Δ^k to denote the set of k -cells in the ideal triangulation.

Definition 3.1. A *marking* of an ideal triangulation Δ of a 3-manifold M is a choice of a pair of opposed edges on each tetrahedron in Δ .




FIGURE 1. Red tangle associated to the marked blue edges.




FIGURE 2. Boundaries of slicing quadrilaterals.

A marking of a tetrahedron determines a tangle by the local rules in Figure 1, and hence an associated smooth branched double cover with topology $S^1 \times D^2 \xrightarrow{2:1} D^3$. We may alternatively specify a marking by choosing quadrilaterals which slice the tetrahedron in a way separating the tangle associated to the marking, as in Figure 2.

3.2. Singular branched double covers. In fact, given an ideal triangulation of a 3-manifold M without any additional structure, there is a singular branched double cover $\widetilde{M}^{\text{sing}}$ of M canonically associated to it; see e.g. [8, 25]. In this section, we briefly review this construction.

Let M be a 3-manifold with an ideal triangulation Δ . We write $M^\circ \subset M$ for the complement of the barycenters of the tetrahedra in Δ . We write $\tau \subset M$ for the singular tangle that consists of the edges (1-cells) of the polygonal decomposition dual to the triangulation, i.e. the edges which connect barycenters in adjacent tetrahedra of Δ , passing through the barycenters of faces, see Figure 3.

We write \widetilde{M}° for branched double cover of M° corresponding to the smooth tangle $\tau \cap M^\circ$. The topology of \widetilde{M}° near the barycenters is $T^2 \times \mathbb{R}_{\geq 0} \subset \widetilde{M}^\circ$, where the torus T^2 double covers the sphere S^2 in the punctured neighborhood $\mathbb{R}_{\geq 0} \times S^2 \subset M^\circ$ of the barycenter. We compactify \widetilde{M}° by gluing in a cone on the torus T^2 near each barycenter. We call the resulting singular space $\widetilde{M}^{\text{sing}}$. Then $\widetilde{M}^{\text{sing}}$ is smooth save for the isolated points (corresponding to the barycenters) where it is the cone on a torus and it comes with a branched double cover $\widetilde{M}^{\text{sing}} \rightarrow M$. We write $\widetilde{\tau} \subset \widetilde{M}^{\text{sing}}$ for the preimage of τ under this branched double cover.

Consider a tetrahedron in Δ and consider inside this tetrahedron the triangles with the following vertices: one vertex of the tetrahedron, the barycenter of a face adjacent to this tetrahedron vertex, and the barycenter of the tetrahedron. For each vertex of the tetrahedron there are three such triangles and therefore there are $4 \cdot 3 = 12$ such triangles in each tetrahedron, see Figure 4. In [25], the complex formed by all such triangles in M is called




FIGURE 3. The branch locus for the singular branched double cover of a tetrahedron




FIGURE 4. 3d spectral network on a tetrahedron (twelve green triangles).

a 3d spectral network in M , as an analogue in the 3-dimensional context of the spectral networks for surfaces [29].

Consider next the complement of the 3d spectral network triangles in a tetrahedron of Δ . It has six connected components, each one contains exactly one edge of the tetrahedron and exactly two of its ideal vertices in its boundary. We next consider the faces (2-cells) of the polygonal decomposition Δ^\vee dual to Δ . In each tetrahedron these are the quadrilaterals depicted in Figure 5. (We will use the 2-skeleton $\Delta_{(2)}^\vee$ as a branch cut for the double cover $\widetilde{M}^{\text{sing}} \rightarrow M$, see Definition 3.2.) Each component of the complement of the 3d spectral network triangles retracts to one of the quadrilaterals in $\Delta_{(2)}^\vee$. In fact, it retracts along the lines of a foliation in which each line connects the two ideal vertices in the boundary of the complement component and intersects the quadrilateral exactly once. We extend the foliation across the 3d spectral network by allowing critical leaves: the critical leaves connect points on the branch locus to ideal vertices as follows: each edge in the branch locus is a boundary component of three of the spectral network triangles and the critical leaves through a point in the branch locus connects this point inside the network triangles to the ideal vertex which is also a vertex of the triangle, see Figure 6 for an illustration of this 1-dimensional foliation.

Consider the foliation inside a tetrahedron $\delta \in \Delta$ in a connected component $U_\delta \subset \delta \setminus \Delta_{(2)}^\vee$ of the complement of the branch cut. It admits two orientations, either all leaves are oriented away from the branch cut toward the ideal vertex or vice versa. We will lift these oriented foliations to the branched double cover. Consider two copies of U_δ that are the sheets of the double cover over U_δ .

Definition 3.2 (Sheet labeling convention). In “sheet 1” over U_δ , the leaves are oriented away from the branch cut. In “sheet 2” over U_δ , they are oriented towards the branch cut.




FIGURE 5. Leaf space / branch cuts for the branched cover (six orange quadrilaterals).




FIGURE 6. 1-dimensional foliation on a tetrahedron




FIGURE 7. Orientation convention for the foliation, as seen on a boundary triangle of the tetrahedron

Note that as we go across the branch cut, the two sheets (hence the labels for these orientations) flip. See Figure 7 for an illustration of this orientation convention. Consider the lift of the foliation in M to the branched double cover $\widetilde{M}^{\text{sing}}$. Note that the two orientations of the foliation in M together give a genuine orientation of the lifted foliation in $\widetilde{M}^{\text{sing}}$.

3.3. Angle structures. In this section we introduce the notions of angle groups and structures used to encode natural geometric structures on the leaf spaces of the foliations discussed in Section 3.2.

Definition 3.3. For a 3-manifold M with ideal triangulation Δ , the *angle group* $\mathbf{a}(\Delta)$ is the abelian group with generators a set of *angle symbols*, one for each opposing pair of dihedral angles in each tetrahedron of Δ (see Figure 8) and one auxiliary generator that we denote π , subject to the following relations:

- (1) The three angle symbols θ , θ' , and θ'' of any tetrahedron in Δ satisfies

$$\theta + \theta' + \theta'' = \pi.$$




FIGURE 8. Dihedral angles $\theta, \theta', \theta''$




FIGURE 9. Distinguished cycle (green) in the branched double cover associated to a marked ideal tetrahedron; the numbers indicate the sheet labels away from the standard branch cut (Figure 5). These cycles remain the same if we reverse the orientation and flip the sheet labels at the same time.

- (2) The sum of angle symbols around any internal edge is 2π .
- (2') In case M has boundary, the sum of angle symbols around a boundary edge is π .

In case Δ was an ideal hyperbolic triangulation, then specializing the angle symbols to the actual values of the angles gives a map $\mathbf{a}(\Delta) \rightarrow \mathbb{R}$ that takes the generator π to the number π and each angle symbol to some number in $(0, \pi)$. More generally, any map of the angle group with these properties is called an *angle structure*. Even more generally, any map $\mathbf{a}(\Delta) \rightarrow \mathbb{R}$ sending π to π is called a *generalized angle structure*. If a generalized angle structure maps each angle symbol to 0 or π , it is called a *taut angle structure*.

While the angle group may become trivial if some of the imposed relations are inconsistent, it is well-known, see e.g. [43, Theorem 1], that any ideal triangulation of M admits a generalized angle structure – and hence has a non-trivial angle group – if all cusps of M are tori.

We will write $\mathbb{Z}[\mathbf{a}(\Delta)]$ for the group ring of the angle group and denote the generators of this group ring as $a := [\pi]$ and correspondingly $a^{\theta/\pi} := [\theta]$; a -powers then have the expected behavior of exponentials.

3.4. An effectivity condition. For each marked tetrahedron, there is a *distinguished cycle* in the first homology of the corresponding branched double cover, see Figure 9. Throughout this article, we will consider only covers satisfying the following:

Definition 3.4 (Effectivity condition). We say that a marking on an ideal triangulation is *effective* if the distinguished cycles span a strictly convex cone in $H^1(\widetilde{M}, \mathbb{R})$, or equivalently, if there exists a closed 1-form ζ on \widetilde{M} that is positive on all distinguished cycles.




FIGURE 10. The 1-chain $\gamma_{\delta}^{\tilde{e}} \in H_1(\tilde{\delta}, \partial\tilde{\delta})$, drawn on the boundary, viewed from outside of $\tilde{\delta}$.

We will need this condition to define the following completion of the skein:

Definition 3.5. For a cover \tilde{M} associated to an effective marking, and some (any) choice of corresponding closed 1-form ζ , we write $\widehat{\text{Sk}}^{\zeta}(\tilde{M}; \tilde{\sigma})$ for the completion of $\text{Sk}(\tilde{M}; \tilde{\sigma})$ that consists of all infinite sums $\sum_{\alpha} \kappa_{\alpha}$, $\kappa_{\alpha} \in \text{Sk}(\tilde{M}, \tilde{\sigma})$, such that if $[\kappa_{\alpha}] \in H_1(\tilde{M})$ denotes the homology class represented by κ_{α} then for any $N > 0$, $\int_{[\kappa_{\alpha}]} \zeta < N$ for only finitely many κ_{α} .

We split the homology and cohomology of \tilde{M} by its characters for the involution p :

$$\begin{aligned} H^1(\tilde{M}, \mathbb{R}) &= H^1(\tilde{M}, \mathbb{R})^{\text{even}} \oplus H^1(\tilde{M}, \mathbb{R})^{\text{odd}} \\ H_1(\tilde{M}, \mathbb{R}) &= H_1(\tilde{M}, \mathbb{R})^{\text{even}} \oplus H_1(\tilde{M}, \mathbb{R})^{\text{odd}} \end{aligned}$$

and we use similar notation for the corresponding splitting of 1-forms or 1-chains; i.e.

$$\zeta^{\text{odd}} := \frac{1}{2}(\zeta - p^*\zeta) \quad \text{and} \quad \zeta^{\text{even}} := \frac{1}{2}(\zeta + p^*\zeta).$$

As always for Galois covers, the invariant co/homology is identified with the co/homology of M by pushforward or pullback. A distinguished cycle has image contained in a tetrahedron of M , so has vanishing even class. Thus if ζ is positive on all distinguished cycles, so is ζ^{odd} .

To an edge \tilde{e} of the double cover $\tilde{\delta}$ of a marked tetrahedron $\delta \in \Delta^3$, we associate a 1-chain $\gamma_{\delta}^{\tilde{e}}$ in $\tilde{\delta}$ connecting the two barycenters of the adjacent faces, oriented so that the puncture on sheet 1 (resp., 2) is to the right (resp., left); see Figure 10. We will be interested in the periods

$$x_{\delta}^{\tilde{e}} := \int_{\gamma_{\delta}^{\tilde{e}}} \zeta^{\text{odd}}.$$

If \tilde{e}_1 and \tilde{e}_2 are two edges of $\tilde{\delta}$ that project down to the same edge e of δ , then $\gamma_{\delta}^{\tilde{e}_1} + \gamma_{\delta}^{\tilde{e}_2}$ is a cycle, see Figure 11. This cycle is odd: $p(\gamma_{\delta}^{\tilde{e}_1}) = -\gamma_{\delta}^{\tilde{e}_2}$. It follows that

$$\int_{\gamma_{\delta}^{\tilde{e}_1} + \gamma_{\delta}^{\tilde{e}_2}} \zeta = \int_{\gamma_{\delta}^{\tilde{e}_1}} (\zeta - p^*\zeta) = 2x_{\delta}^{\tilde{e}_1} = 2x_{\delta}^{\tilde{e}_2}.$$

Since $x_{\delta}^{\tilde{e}_1} = x_{\delta}^{\tilde{e}_2}$, we simply write x_{δ}^e for this common value.

Lemma 3.6. *If three edges $\tilde{e}, \tilde{e}', \tilde{e}''$ of $\tilde{\delta}$ are adjacent to the same vertex, then $\gamma_{\delta}^{\tilde{e}} + \gamma_{\delta}^{\tilde{e}'} + \gamma_{\delta}^{\tilde{e}''}$ is homologous to zero. Moreover, fixing a lift \tilde{e} of an edge $e \in \Delta^1$, the sum (with multiplicity)*

$$(3.1) \quad \sum_{\delta \text{ abutting } e} \gamma_{\delta}^{\tilde{e}}$$




FIGURE 11. $\gamma_\delta^{\tilde{e}_1} + \gamma_\delta^{\tilde{e}_2}$ is the distinguished cycle if e is an edge clockwise from the marked edges when viewed from a vertex.




FIGURE 12. Branch locus for the smoothing of a marked tetrahedron. The quantity $z > 0$ records the half of the value of a 1-form on the distinguished 1-cycle (Figure 9) on the boundary.

is also homologous to zero.

Proof. For the first claim, it is enough to observe that $\tilde{\gamma}_\delta + \tilde{\gamma}_\delta^{e'} + \tilde{\gamma}_\delta^{e''}$ is homologous to a small contractible loop close to the vertex shared by the three edges. For the second claim, the sum is homologous to a link of \tilde{e} , which is contractible. \square

Corollary 3.7.

- (1) If e, e', e'' are three edges of δ sharing a vertex, $x_\delta^e + x_\delta^{e'} + x_\delta^{e''} = 0$.
- (2) If $\{e, f\}$ is a pair of opposite edges of δ , then $x_\delta^e = x_\delta^f$.
- (3) If $e \in \Delta^1$ is an edge of the triangulation, then

$$(3.2) \quad \sum_{\delta \text{ abutting } e} x_\delta^e = 0.$$

(If δ has multiple edges abutting e , then x_δ^e should be counted with multiplicity accordingly in (3.2).)

Proof. (1) and (3) are immediate corollaries of Lemma 3.6, and (2) follows from (1). \square

Corollary 3.8. For a marked tetrahedron δ and a closed 1-form ζ on the branched double cover \tilde{M} , there is a unique $z_\delta \in \mathbb{R}$ such that

$$(3.3) \quad x_\delta^e = \begin{cases} 0 & \text{if } e \text{ is a marked edge,} \\ \pm z_\delta & \text{otherwise,} \end{cases}$$

where the sign is + (resp., -) if the marked edges are counterclockwise (resp., clockwise) from e when viewed from vertices; see the right-hand side of Figure 12.

Proof. The period of ζ along one of the three cycles $(1, 0)$, $(0, 1)$, and $(-1, -1)$ – corresponding to quadrilateral slices in Figure 2 – is equal to $2x_\delta^e$, where e is an edge avoiding the quadrilateral slice; see Figure 11. Since the cycle corresponding to the marked edges is collapsed, the result follows. \square

We call z_δ the *angular period* of ζ over δ . The period of ζ along the distinguished cycle of $\tilde{\delta}$ is $2z_\delta$, so ζ is positive on all distinguished cycles if and only if the z_δ are all positive. By rewriting the gluing equations (3.2) in terms of $\{z_\delta\}_{\delta \in \Delta}$ using (3.3), we obtain a system of homogeneous equations

$$(3.4) \quad \left\{ \sum_{\delta \text{ abutting } e} \epsilon_\delta^e z_\delta^e = 0 \right\}_{e \in \Delta^1},$$

where Δ^1 denotes the edges in the triangulation Δ , and $\epsilon_\delta^e \in \{0, \pm 1\}$ is determined as in (3.3). Since the sum of angles in the generalized angle structure around any edge $e \in \Delta^1$ equals zero, each boundary component of M is a torus. It follows that the Euler characteristic of M is zero. Then, if there are t tetrahedra in Δ^3 , there must be $2t$ faces in Δ^2 , and hence t edges in Δ^1 . The t equations (3.4) are well-studied in the context of normal surface theory [43], and it is well-known that if M has b boundary components then at most $t - b$ of the t equations are linearly independent. In particular, it admits an at least b -dimensional space of solutions.

Lemma 3.9. *Any solution (z_1, \dots, z_t) to (3.4) induces a closed and bounded 1-form ζ on \tilde{M} , odd with respect to p , with given periods.*

Proof. Each z_δ determines a closed 1-form on $\tilde{\delta}$, odd with respect to p , with the given period z_δ . Since we are specifying the periods along the 1-chains ending on the barycenters, these 1-forms glue together to give a closed 1-form on the 3-manifold obtained by gluing $\tilde{\delta}$'s along small neighborhoods of barycenters. The resulting 3-manifold is \tilde{M} minus the preimages of the edges in Δ^1 . We obtain \tilde{M} by attaching 2-handles along the links of those preimages. The gluing equations (3.4) are exactly the conditions that the closed 1-form extends over these 2-handles as a closed 1-form. After extension we get a closed 1-form near the boundary of \tilde{M} , $\sqcup_{j=1}^{2b} T^2 \times [0, 1)$, and we extend it to a bounded closed 1-form over $\sqcup_{j=1}^{2b} T^2 \times \mathbb{R}_+$. \square

To this end, the effectivity condition simply requires all z_δ 's to be positive. We summarize our discussion so far with the following:

Proposition 3.10. *Let Δ be a marked ideal triangulation of M with t tetrahedra and b boundary components. Then \tilde{M} admits a closed 1-form ζ positive on all distinguished cycles if and only if there is a tuple of angular periods (z_1, \dots, z_t) that satisfies (3.4) and has all $z_j > 0$. In particular, such solutions exist provided the d -dimensional solution space to (3.4), $d \geq b \geq 1$, intersects the positive orthant $\bigcap_{1 \leq j \leq t} \{z_j > 0\}$. \square*

Remark 3.11. By Stiemke's alternative [64], the criterion in Proposition 3.10 is equivalent to the following condition: there is no $\{p_e\}_{e \in \Delta^1} \in \mathbb{R}^t$ such that $\{\sum_e p_e \epsilon_\delta^e\}_{\delta \in \Delta^3}$ is nonnegative but not identically zero.

Example 3.12. Consider the figure-eight knot complement ideally triangulated as in Figure 13. Since there are only two tetrahedra (hence 2 edges), besides $\theta + \theta' + \theta'' = \pi$ and




FIGURE 13. Figure-eight knot complement triangulated into two ideal tetrahedra

$\eta + \eta' + \eta'' = \pi$, there is only one equation that defines a generalized angle structure, which is

$$2\theta + \theta'' + 2\eta + \eta'' = 2\pi.$$

As a result, this triangulation admits 3 taut angle structures:

$$\begin{aligned} (\theta, \theta', \theta'', \eta, \eta', \eta'') &= (\pi, 0, 0, 0, \pi, 0), \\ \text{or } (\theta, \theta', \theta'', \eta, \eta', \eta'') &= (0, 0, \pi, 0, 0, \pi), \\ \text{or } (\theta, \theta', \theta'', \eta, \eta', \eta'') &= (0, \pi, 0, \pi, 0, 0). \end{aligned}$$

Note, for any pair $\theta^\square \in \{\theta, \theta', \theta''\}$ and $\eta^\square \in \{\eta, \eta', \eta''\}$, there is a taut angle structure for which both θ^\square and η^\square are zeros, meaning, for any choice of tangle (corresponding to a choice of marked quad type) in each tetrahedron, there is a taut angle structure compatible with it.

In this example, the equation (3.4) is given by

$$(3.5) \quad 2x_1^\theta + x_1^{\theta''} + 2x_2^\eta + x_2^{\eta''} = 0,$$

where

$$x_1^{\theta^\bullet} = \begin{cases} 0 & \text{if } \theta^\bullet \text{ is the marked quad type } \theta^\square, \\ z_1 & \text{if } \theta^\square \text{ is counterclockwise from } \theta^\bullet, \\ -z_1 & \text{if } \theta^\square \text{ is clockwise from } \theta^\bullet, \end{cases}$$

and likewise for $x_2^{\eta^\bullet}$. There are solutions (z_1, z_2) such that $z_1, z_2 > 0$ if and only if the coefficient of z_1 and that of z_2 in (3.5) have opposite signs, which is the case when the marking $(\theta^\square, \eta^\square)$ is either (θ, η'') , (θ', η'') , (θ'', η) , or (θ'', η') .

3.5. HOMFLYPT skein modules of singular branched double covers.

Definition 3.13. Let M be a 3-manifold with ideal triangulation Δ , associated (singular) branched double cover $\widetilde{M}^{\text{sing}}$ and branch locus $\widetilde{\tau} \subset \widetilde{M}^{\text{sing}}$. We write $\text{Sk}_{a,z}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})$ for the $\mathbb{Z}[\mathbf{a}(\Delta)][z^\pm]$ module generated by the isotopy classes of framed oriented links in $\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau}$, subject to the following relations

- the usual skein relations (1.1), (1.2), (1.3),
- relation (1.4) holds as links cross $\widetilde{\tau}$; i.e., the branch locus $\widetilde{\tau}$ is a sign line,




FIGURE 14. The 3-term relation (3.6) near the cone point




FIGURE 15. The 3-term relation (3.6) drawn on the projection to the torus

- the following relation near the singular cone point in any tetrahedron of Δ :

$$(3.6) \quad \begin{array}{c} \theta' \quad 1 \quad 2 \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \theta'' \end{array} = a^{\frac{\theta}{\pi}} \begin{array}{c} \theta' \quad 1 \quad 2 \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \theta'' \end{array} + a^{-\frac{\theta'}{\pi}} \begin{array}{c} \theta' \quad 1 \quad 2 \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \text{---} \quad \text{---} \quad \text{---} \\ \theta'' \end{array} .$$

Unlike the first four relations in $\text{Sk}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})$, which are drawn in $\widetilde{M}^{\text{sing}}$, the relation (3.6) is drawn in the projection to M – the numbers 1 and 2 denote the sheet labels in accordance with the convention in Definition 3.2. It is a local relation near the singular cone point at the barycenter of some tetrahedron in Δ , here projected from a vertex of tetrahedron to the face opposing it. For the corresponding 3-dimensional picture in M , see Figure 14, and for the relation in $\widetilde{M}^{\text{sing}}$, projected to the torus T^2 which is the link of the cone point in $\widetilde{M}^{\text{sing}}$, see Figure 15.

3.6. Specializing the singular skein. Following [41], we say that a generalized angle structure $\Theta : \mathbf{a}(\Delta) \rightarrow \mathbb{R}$ is *taut* if it maps each angle symbol to 0 or π .³

Given a taut structure, we produce a marking (and hence a tangle in M , and associated smooth branched double cover) specifying a sign $\{+, -\}$ on each tetrahedron, and following the local rules of Figure 17. We will say a taut triangulation equipped with such a choice of signs is a *signed taut triangulation*. In this case, as usual, we denote the smooth branched cover $\widetilde{M} \rightarrow M$ and its branch locus $\widetilde{\sigma} \subset \widetilde{M}$ over $\sigma \subset M$.

We write $\text{Sk}(\widetilde{M}^{\text{sing}})_{\Theta}$ for specialization of the singular skein by using Θ to evaluate the angle variables.

³To be precise, our notion of “taut” is weaker than that of Lackenby, in that Lackenby’s notion also requires a choice of transverse orientation. If we are also given a transverse orientation, then half-integer powers of a would not appear in our skein trace map. See also Remark 9.5.




FIGURE 16. A taut ideal tetrahedron




FIGURE 17. Positive (left) and negative (right) compatible tangles drawn in red. The edges they link are drawn in blue. The sign is the sign of the angle at a vertex from a π edge to the linked (blue) edge, viewed from outside the tetrahedron.




FIGURE 18. Solid torus double-covering the 3-ball. The top and the bottom faces of the right-hand side each correspond to a pair of faces of the tetrahedron sharing a π -edge.

Theorem 3.14. *Fix a taut structure Θ , and a signing so that the corresponding marking is effective. Then there is a unique morphism $\text{Sk}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})_{\Theta} \rightarrow \widehat{\text{Sk}}^{\zeta}(\widetilde{M}; \widetilde{\sigma})$ commuting with the map of skeins induced by the inclusion $\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau} \hookrightarrow \widetilde{M} \setminus \widetilde{\sigma}$.*

Proof. Such a morphism is unique if it exists, since any skein in $\text{Sk}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})_{\Theta}$ can be represented as a linear combination of links in $\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau}$.

Consider \widetilde{M}° , the complement in $\widetilde{M}^{\text{sing}}$ of small neighborhoods of the singular points. This is a manifold with one boundary torus near each singular point. Then $(\widetilde{M}, \widetilde{\sigma})$ is obtained from $(\widetilde{M}^{\circ}, \widetilde{\tau})$ by gluing in solid tori with appropriate sign lines. This is illustrated in Figure 18.




FIGURE 19. Skein operators on the boundary of the solid torus.



FIGURE 20. W_λ in the solid torus with sign lines

FIGURE 21. Resolving the singular point by inserting a skein dilogarithm

Now, by definition, $\text{Sk}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})$ is the quotient of $\text{Sk}(\widetilde{M}^\circ; \widetilde{\tau})$ by the 3-term relation (3.6) (imposed near each singularity). Thus a map from $\text{Sk}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})$ is the same as a map from $\text{Sk}(\widetilde{M}^\circ; \widetilde{\tau})$ which annihilates this relation.

But with the specialization of angles, the 3-term skein relation (3.6) becomes the 3-term recurrence relation (2.4) satisfied by the skein dilogarithm in case of the positive resolution, or the corresponding recurrence relation for its inverse in case of the negative resolution, inserted along the distinguished cycle.⁴ Thus, we construct our desired map by gluing a solid torus

⁴The sign lines don't change the recursion since we can pull the sign lines away symmetrically. See Figure 19 for the meridian and longitude operators drawn on the boundary torus, and Figure 20 for the elements W_λ in the presence of sign lines in the solid torus. The orientation of the longitude $P_{0,1}$ is chosen to match that of the distinguished cycle.

$D^2 \times S^1$ containing a skein dilogarithm $\Psi \in \text{Sk}(D^2 \times S^1)$ (or its inverse $\Psi^{-1} \in \text{Sk}(D^2 \times S^1)$) cabling the distinguished cycle to each boundary of \widetilde{M}° ; see Figure 21. \square

4. A SKEIN TRACE FOR DOUBLE COVERS OF 3-MANIFOLDS

Building on ideas from [30, 29, 25] and especially [49, 50], we will prove:

Theorem 4.1. *Let M be a 3-manifold with ideal triangulation Δ and associated (singular) branched double cover $\widetilde{M}^{\text{sing}}$ with branch locus $\widetilde{\tau}$. Then there is a well-defined $\mathbb{Z}[a^\pm, z^\pm]$ -module homomorphism*

$$[\cdot]_{\widetilde{M}^{\text{sing}}}^{\text{net}} : \text{Sk}_{a^2, z}(M) \rightarrow \text{Sk}_{a, z}(\widetilde{M}^{\text{sing}}; \widetilde{\tau}).$$

Here, “net” is to remind us of spectral networks. The morphism is an explicit formula for lifting links, described in (4.2) below. We prove it is well defined, in Section 4.2, by checking by hand that the explicit formula respects the skein relations. The fact that the morphism is nontrivial can be seen already by calculations for the trivial double cover, which we give in Section 4.3. By the results of the previous section, we derive the following consequence for smooth double covers:

Corollary 4.2. *Let M be a 3-manifold equipped with effective signed taut ideal triangulation, and let $\widetilde{M} \rightarrow M$ be the branched double cover associated to the tangle σ it specifies. Then there is a well-defined map*

$$[\cdot]_{\widetilde{M}}^{\text{net}} : \text{Sk}_{a^2, z}(M) \rightarrow \widehat{\text{Sk}}_{a, z}(\widetilde{M}; \widetilde{\sigma}).$$

Proof. Compose the map of Theorem 4.1 with that of Theorem 3.14. \square

4.1. Turning numbers. We describe additional geometric data on a 3-manifold M that allows us to compute framing data numerically. Assume that the tangent bundle of M is decomposed as $TM = \mathbb{R} \oplus E$, where \mathbb{R} is the trivial bundle trivialized by some vector field ξ and E is an $SO(2)$ -bundle, then links in M that are nowhere tangent to ξ gets an induced framing as follows. If the tangent to K is t , let e_1 denote the projection of t to E and frame K by the unit normal vector field e_2 in E that is perpendicular to e_1 in E and such that (e_1, e_2) is a positively oriented basis. If furthermore E is equipped with a flat connection ∇ then the section e_2 along K can be compared to a parallel section in E , if $s \in [0, L]$ parametrizes K then

$$(4.1) \quad \int_0^L \langle \nabla_s e_2, -e_1 \rangle ds = \int_0^L \langle \nabla_s e_1, e_2 \rangle ds \in 2\pi\mathbb{Z},$$

measures to total turning of e_2 , or equivalently e_1 , with respect to the parallel section along K . We thus have an integer that measures the total turning of the induced framing of a link nowhere parallel to ξ . Consider now a general framing of K , given by a vector field n generic with respect to ξ . Then n is parallel to ξ only at a finite number of points along K , we call such a point a *framing tangency*, and if n_E denotes the projection of n to E and p is a point where n_E vanishes then $\nabla_s n_E(p)$ is linearly independent with $e_1(p)$. We take the sign $\sigma(p) = \pm 1$ of the framing tangency p to be the orientation sign of $(e_1(p), \nabla_s n_E(p))$. It is easy to see that the framing given by n is homotopic to the framing given by $\hat{n}_E = \pm e_2$, the unit vector in direction n_E outside small intervals of framing tangencies and a $\sigma(p)\pi$ -rotation

$$\hat{n}_E(\text{before}) \rightarrow \sigma(p)\xi \rightarrow \hat{n}_E(\text{after})$$




FIGURE 22. Euclidean structure on each sheet of the leaf space

inside the small intervals.

Lemma 4.3. *Two generic framings of a knot K are homotopic if and only if their sums of the turning numbers and discrete π -turns at framing tangencies:*

$$\int_0^L \langle \nabla_s e_1, e_2 \rangle ds + \sum_p \sigma(p)\pi \in 2\pi\mathbb{Z}$$

agree.

Proof. The sum of the turning number and the π -turns is clearly invariant under generic homotopy of framings. Further one can homotope any framing to a framing given by e_2 outside a small arc where all π rotations take place and it is clear that any two such framings with the same algebraic number of π -turns are homotopic. \square

We now consider an analogue of this for branched double covers. Consider an ideal triangulation Δ on M equipped with a generalized angle structure. The generalized angle structure gives a geometric polygonal decomposition of the leaf space of Δ , or equivalently the 2-skeleton $\Delta_{(2)}^\vee$ of the dual polygonal decomposition, as follows. Consider a 2-dimensional polygon η of $\Delta_{(2)}^\vee$ inside a tetrahedron of Δ . It has four corners, two at barycenters of the boundary triangles, one at the barycenter of an edge and one at the barycenter of the tetrahedron itself. We represent η as a planar polygon with angles at the corners corresponding to barycenters of the boundary triangles equal to $\frac{\pi}{2}$, the angle at the corner of the barycenter of the edge equal to θ as determined by the generalized angle structure and the angle at the corner of the barycenter equal to $\pi - \theta$, see Figure 22.

Consider the tangent bundles $T\eta$ of such polygonal representatives of the 2-dimensional polygons η of Δ^\vee and equip them with the $SO(2)$ -structure induced by the flat metric on $\eta \subset \mathbb{R}^2$ as above. Note that these bundles glue naturally along the edges of the polygons in η and thus gives a flat $SO(2)$ -bundle over $\Delta_{(2)}^\vee \setminus \Delta_{(0)}^\vee$. Pulling back along the foliation gives a flat $SO(2)$ -bundle E over $\widetilde{M}^{\text{sing}} \setminus (\widetilde{\tau} \cup \widetilde{\Delta}_{(1)})$.

Lemma 4.4. *The $SO(2)$ -bundle E extends over the 1-skeleton as a flat bundle over $\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau}$ and gives a decomposition $T(\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau}) = E \oplus \mathbb{R}$, where \mathbb{R} is trivial bundle with section the tangent ξ to the foliation of $\widetilde{M}^{\text{sing}}$ induced by Δ . Moreover, the monodromy in E around the a loop in the three faces of Δ^\vee adjacent to a vertex of a tetrahedron in Δ equals 2π .*

Proof. The first statement follows from the condition that angles sum to 2π around edges of the triangulation. To see the second statement, note that E can be viewed as the quotient

of the tangent bundle by the foliation. The final statement follows from adding the angles at the barycenter of the three faces of the leaf space that are adjacent to a vertex of the tetrahedron:

$$(\pi - \theta) + (\pi - \theta') + (\pi - \theta'') = 3\pi - (\theta + \theta' + \theta'') = 2\pi.$$

□

The orientation on the foliation together with the orientation of $\widetilde{M}^{\text{sing}}$ induces an orientation on E . We use this flat bundle E to define the turning number of a link diagram on the leaf space:

Definition 4.5. Let \widetilde{K} be an oriented link in $\widetilde{M}^{\text{sing}} \setminus \widetilde{\tau}$ with tangent vectors everywhere transverse to the foliation. Equip \widetilde{K} with the framing given by the tangent ξ to the foliation. Then projecting to the leaf space, we get a link diagram framed by a vector in the projection direction. Consider the $SO(2)$ -bundle E along \widetilde{K} with the section e_1 which is the normalized projection of the tangent vector to \widetilde{K} to E . Define the *turning number* of \widetilde{K} , $\text{turn}(\widetilde{K}) \in \mathbb{R} = T_1SO(2)$, to be the monodromy of this global section with respect to the flat connection ∇ determined by the generalized angle structure, with notation as in (4.1):

$$\text{turn}(\widetilde{K}) = \frac{1}{2\pi} \int_0^L \langle \nabla_s e_1, e_2 \rangle ds \in \mathbb{R}.$$

We next show how to compute turning numbers diagrammatically. Let \widetilde{K} be a knot that is generic with respect to the triangulation Δ . Project \widetilde{K} first to $K \subset M$ and then to the leaf space $\kappa \subset \Delta_{(2)}^\vee$. Let η be a face in $\Delta_{(2)}^\vee$. Then $\kappa \cap \eta$ is a collection of curves $\kappa[\eta]$ in η the lift of the curve $\kappa[\eta]$ to \widetilde{K} then subdivides $\kappa[\eta]$ into $\kappa[\eta]_1 \cup \kappa[\eta]_2$ where $\kappa[\eta]_j$ has lift that lies in sheet $j \in \{1, 2\}$ (with respect to the branch cut). We write

$$\text{turn}_{ij}(\widetilde{K}_i) = \sum_{\eta} \text{turn}_{ij}(\kappa[\eta]_i), \quad i \neq j \in \{1, 2\},$$

where $\text{turn}_{ij}(\kappa[\eta]_i)$ denotes the turning number, in the usual elementary sense, of $\kappa[\eta]_i$ projected to that face of the ij -leaf space, oriented in such a way that the i - (resp., j -) direction of the foliation is pointing away (resp., toward) the point of projection.

Lemma 4.6. *We have*

$$\text{turn}(\widetilde{K}) = \text{turn}_{12}(\widetilde{K}_1) + \text{turn}_{21}(\widetilde{K}_2).$$

Furthermore, $\text{turn}(\widetilde{K})$ is invariant under isotopies of \widetilde{K} that are nowhere tangent to ξ and changes by 2π (-2π) at positive (negative) generic tangencies with ξ .

Proof. By definition the contribution to the turning number from a subarc of the link diagram that lies in a polygonal face equals the usual turning number of the tangent of the arc considered as plane curve. The first formula follows. The invariance statement is immediate from the definition. To see the last statement note that positive and negative tangencies correspond to the appearance or disappearance of positive or negative kinks in the diagram. □

Remark 4.7. The turning number is a \mathbb{Z} -linear combination of ratios $\frac{\theta}{2\pi}$, where θ comes from the generalized angle structure, and $\frac{1}{2}$, corresponding to half-turns near the seams of the leaf space, so the turning number is in general *not* an integer.

Remark 4.8. A generic framing of a knot \tilde{K} is tangent to ξ at a finite number of points. Adding π -turns at such tangencies to the turning number exactly as in Lemma 4.3 gives a numerical invariant of framed knots in \tilde{M} that remains invariant under framed isotopies by Lemma 4.6.

4.2. Proof of Theorem 4.1. Let M be a 3-manifold with ideal triangulation Δ , associated branched double cover \tilde{M}^{sing} , with branch locus $\tau \subset M$ and correspondingly $\tilde{\tau} \subset \tilde{M}^{\text{sing}}$. In this section, we construct a $\mathbb{Z}[a^\pm, z^\pm]$ -module homomorphism:

$$[\cdot]_{\tilde{M}^{\text{sing}}}^{\text{net}} : \text{Sk}_{a^2, z}(M) \rightarrow \text{Sk}_{a, z}(\tilde{M}^{\text{sing}}, \tilde{\tau}).$$

The construction is almost identical to that of [49, 54], except that we are working with HOMFLYPT skeins, instead of \mathfrak{gl}_2 -skeins.


Let $K \subset M \setminus \tau$ be a framed link in general position with respect to the foliation. Projecting K to the leaf space of the foliation gives a link diagram on the leaf space. Thinking of K as a ribbon link, the link diagram will be flat on the leaf space, except at finitely many points where it is half-twisted (i.e., where the framing is parallel to the foliation).

We will define $[K]_{\tilde{M}^{\text{sing}}}^{\text{net}}$ as a finite linear combination $\sum_{\beta} c(\tilde{K}_{\beta})[\tilde{K}_{\beta}]$, where the summands $\tilde{K}_{\beta} \subset \tilde{M}^{\text{sing}} \setminus \tilde{\tau}$ ranges over all possible lifts of K that can be constructed by applying combinations of three local lifts: *direct lifts*, *detours*, and *exchanges* defined as follows:


- *Direct lift*: an arc of K can be lifted to either sheet of the branched double cover, we label the arc by the number of the sheet to where it lifts:



- *Detour*: arcs in K that intersect a critical leaf (i.e. that cross the binding of the leaf space), can be lifted as follows, where the numbering indicates to which sheet we lift:



- *Exchange*: two arcs in K that lie on the same line in the foliation (i.e. crossings in the link diagram of K on the leaf space) can be lifted in the following way, where the numbering indicates to which sheet we lift:



It is easy to see that there are only finitely many possible lifts.

We will associate a weight $c(\tilde{K}_{\beta}) \in R^{\Theta}$ to each lift $\tilde{K}_{\beta} \subset \tilde{M}^{\text{sing}} \setminus \tilde{\tau}$ of $K \subset M \setminus \tau$ that is constructed by the these local lifts. The weight is the product of *turning factors*, *exchange factors*, and *framing factors* defined as follows.

- *Turning factor*:

$$a^{\text{turn}(\tilde{K}_{\beta})} = \prod_{\substack{1 \leq i, j \leq 2 \\ i \neq j}} a^{\text{turn}_{ij} \tilde{K}_{\beta; i}},$$

where $\tilde{K}_{\beta;i}$ is the part of the link diagram of \tilde{K}_β that lies in sheet i over a polygon in the leaf space and where the product ranges over all polygons in the leaf space.

- *Exchange factor:*

$$\text{exch}(\tilde{K}_\beta) = \prod_{\text{exchanges}} \pm z$$

where the product ranges over all exchanges of \tilde{K}_β and where the sign is the sign of the crossing on the leaf space.

- *Framing factor:*

$$a^{\text{frame}(\tilde{K}_\beta)} = \prod_{\text{framing tangencies}} a^{\pm \frac{1}{2}},$$

where the product ranges over all framing tangencies and the sign in the exponent is the sign of the corresponding half-twist in the ribbon link diagram of K .

We then define the weight of \tilde{K}_β as the monomial in a^\pm, z^\pm given by

$$c(\tilde{K}_\beta) = a^{\text{turn}(\tilde{K}_\beta)} a^{\text{frame}(\tilde{K}_\beta)} \text{exch}(\tilde{K}_\beta).$$


Finally, we define

$$(4.2) \quad [K]_{\tilde{M}^{\text{sing}}}^{\text{net}} = \sum_{\beta} c(\tilde{K}_\beta) [\tilde{K}_\beta] \in \text{Sk}_{a,z}(\tilde{M}^{\text{sing}}; \tilde{\tau}),$$


where the sum is over all possible lifts \tilde{K}_β .

We show that $[K]_{\tilde{M}^{\text{sing}}}^{\text{net}} \in \text{Sk}_{a,z}(\tilde{M}^{\text{sing}}; \tilde{\tau})$ depends only on the class $[K] \in \text{Sk}_{a^2,z}(M)$; i.e., that it is invariant under framed isotopy of K and that it respects the skein relations. Any framed isotopy of K can be decomposed into isotopies that induce isotopies of the diagram of K on the leaf space, which clearly does not affect $[K]_{\tilde{M}^{\text{sing}}}^{\text{net}}$, and a finite composition of diagram moves that are isotopies of the following types:


- (I) 3 types of framed Reidemeister moves
- (II) 4 types of moves near the bindings:
 - (a) moving an exchange across a critical leaf




- (b) height exchange for detours




(c) moving a strand across a critical leaf



(d) moving a strand across the branch locus⁵




(III) Moving a strand across a critical leaf ending on the 4-valent singular point




The proof of invariance of $[K]_{\widetilde{M}^{\text{sing}}}^{\text{net}}$ under isotopies of type (I) and (II), as well as the proof that $[\cdot]_{\widetilde{M}^{\text{sing}}}^{\text{net}}$ respects the skein relations of $\text{Sk}_{a,z}(M)$, follows almost verbatim from the proof in [49] after the substitutions $(q - q^{-1}) \mapsto z$ and $q \mapsto a$, see [49, Section 7] for the details.

It then only remains to check invariance under the move (III), where the only non-trivial case is when there are detours. The image of the left hand side of (III) under F that uses a detour is



whereas the corresponding image of the right hand side of (III) under F is a sum of two terms, each using one detour:



where the coefficients come from alteration of the turning factors. Thus, the difference between $[\cdot]_{\widetilde{M}^{\text{sing}}}^{\text{net}}$ acting on the left- and right-hand sides is exactly the 3-term relation (3.6) near the four-valent point of the branch locus at the barycenter of a tetrahedron of Δ in the definition of $\text{Sk}_{a,z}(\widetilde{M}^{\text{sing}}; \widetilde{\tau})$ (Definition 3.13). The theorem follows. \square

⁵The strands are drawn thicker here to emphasize the half-twist on the right-hand side.

4.3. The trivial double cover. In this section, we consider trivial double covers. If M is a 3-manifold and α is an everywhere non-zero 1-form on M , the trivial double cover of M is the two component Lagrangian in T^*M with one component the zero section and the other the graph of α .

We first calculate the skein trace map

$$\text{Sk}_{a_1 a_2, z}(\mathbb{R}^3) \rightarrow \text{Sk}_{a_1, z}(\mathbb{R}^3) \otimes \text{Sk}_{a_2, z}(\mathbb{R}^3)$$

for the trivial double cover of \mathbb{R}^3 . For triangulated 3-manifolds, this situation applies to links which lie in a ball in the interior of a tetrahedron that is disjoint from the branch locus. Since the double cover has two \mathbb{R}^3 components it is possible to use separate framing variables a_1 and a_2 for the two sheets. The framing variable of the underlying \mathbb{R}^3 is then $a_1 a_2$, and the case with non-separated framing variables is obtained simply by setting $a_1 = a_2 = a$. We express the results below with separate framing variables a_1 and a_2 . In this case, the turning factor should be replaced by

$$a_1^{\text{turn}_1(\tilde{K}_\beta)} a_2^{\text{turn}_2(\tilde{K}_\beta)} = \prod_{\substack{1 \leq i, j \leq 2 \\ i \neq j}} a_j^{\text{turn}_{ij} \tilde{K}_{\beta; i}}.$$

While this example may appear trivial – since $\text{Sk}(\mathbb{R}^3)$ is freely spanned by the empty skein $[\emptyset]$, and the skein trace in this case sends $[\emptyset] \mapsto [\emptyset] \otimes [\emptyset]$ – it nonetheless reveals, as we illustrate below, intricate algebraic structures underlying the Hecke algebra.

Example 4.9. We discuss the skein lift of the unknot to the trivial double cover of \mathbb{R}^3 . The skein trace map applied to the unknot is as follows:

$$\begin{array}{ccc} \otimes_2^1 \text{ (circle)} & \mapsto & a_2 \text{ (circle)}^1 + a_1^{-1} \text{ (circle)}^2 \\ \parallel & & \parallel \\ \frac{a_1 a_2 - a_1^{-1} a_2^{-1}}{z} \emptyset & \mapsto & \left(a_2 \frac{a_1 - a_1^{-1}}{z} + a_1^{-1} \frac{a_2 - a_2^{-1}}{z} \right) \emptyset \end{array},$$

where \otimes_1 and \otimes_2 indicates the natural oriented lift to the two sheets of the vector field dual of α .⁶

Example 4.10. We compute the skein lift of a positive kink:

⁶That is, the foliation is orthogonal to the page, oriented in such a way that the orientation of the leaves in sheet 1 (resp., sheet 2) is pointed away from (resp., toward) our eyes.

$$\begin{array}{ccc}
\otimes_2^1 \uparrow \circlearrowleft & \mapsto & a_2^{-1} \uparrow^1 + a_1 \uparrow^2 + a_1 z \uparrow^1 \circlearrowleft^2 \\
& & \parallel \\
& & a_1 a_2^{-1} \uparrow^1 + a_1 a_2 \uparrow^2 + a_1 z \frac{a_2 - a_2^{-1}}{z} \uparrow^1 \\
& & \parallel \\
a_1 a_2 \uparrow & \mapsto & a_1 a_2 (\uparrow^1 + \uparrow^2)
\end{array}$$

This illustrates that skein lifting respects change of framing.

We next consider lifting skein elements γ which are of the form: a framed link K with components colored by representations (i.e., Young diagrams) other than the fundamental representation. To use our definition of the skein lifting map F above, we would first express γ as a linear combination of links in a tubular neighborhood of K all colored by the fundamental representation and then apply F from (4.2). This is a rather complicated procedure, and one could ask for some way to carry out the calculation of the skein trace map acting on γ that bypasses going to the fundamental representation. In the case of trivial double cover and when K has no crossings (i.e., when K is an unlink), there is such a way:

Theorem 4.11. *Under the skein trace map applied to the trivial double cover of the thickened annulus $S^1 \times [0, 1] \times \mathbb{R}$, with turning numbers corresponding to the product metric on $S^1 \times [0, 1]$ where the turning number of the central circle equals zero,*

$$(4.3) \quad W_\lambda \mapsto \sum_{\mu, \nu} c_{\mu\nu}^\lambda W_\mu \otimes W_\nu,$$

where $c_{\mu\nu}^\lambda$ denotes the Littlewood-Richardson coefficients. (Note that this is just the standard coproduct of the self-dual Hopf algebra of symmetric functions.)

Proof. Gluing the inner boundary of an annulus with the outer boundary of another gives the usual product on $\text{Sk}_{a,z}^+(\text{Ann}) \cong \Lambda$, where Λ is the algebra of symmetric functions. Since the leaves of the foliation are along the I -direction of $\text{Ann} \times I$, it follows that the skein trace map

$$\Lambda \cong \text{Sk}_{a_1 a_2, z}^+(\text{Ann}) \rightarrow \text{Sk}_{a_1, z}^+(\text{Ann}) \otimes \text{Sk}_{a_2, z}^+(\text{Ann}) \cong \Lambda \otimes \Lambda$$

is an algebra map, and hence Λ is a bialgebra. We need to show that this coproduct is the standard coproduct (4.3). Since the power sum elements p_n generate Λ , it suffices to show that the corresponding elements in $\text{Sk}^+(\text{Ann})$,

$$P_n := \frac{1}{[n]} (A_{0, n-1} + A_{1, n-2} + \cdots + A_{n-1, 0}),$$

where $A_{i,j}$ is the closure of the braid $\sigma_1 \cdots \sigma_i \sigma_{i+1}^{-1} \cdots \sigma_{i+j}^{-1}$ read from bottom to top [2, 44] and $[n] = (q^{n/2} - q^{-n/2}) / (q^{1/2} - q^{-1/2})$ is the quantum integer n , are primitive with respect

to our coproduct, i.e.,

$$P_n \mapsto P_n \otimes 1 + 1 \otimes P_n.$$

In other words, we need to show that all the non-trivial lifts of P_n cancel out so that only the two direct lifts survive.

By direct calculation, observe that

$$\begin{aligned} A_{i,j} &\mapsto A_{i,j} \otimes 1 + 1 \otimes A_{i,j} \\ &+ z \sum_{0 \leq k \leq i-1} A_{k,0} \otimes A_{i-1-k,j} - z \sum_{0 \leq l \leq j-1} A_{0,l} \otimes A_{i,j-1-l} \\ &- z^2 \sum_{\substack{0 \leq k \leq i-1 \\ 0 \leq l \leq j-1}} A_{k,0} A_{0,l} \otimes A_{i-1-k,j-1-l}, \end{aligned}$$

so that

$$(4.4) \quad \begin{aligned} P_n &\mapsto P_n \otimes 1 + 1 \otimes P_n \\ &+ \frac{z}{[n]} \left(\sum_{\substack{i_1, i_2, j_1, j_2 \geq 0 \\ i_1 + i_2 + j_1 + j_2 = n-2}} A_{i_1,0} \otimes A_{i_2, j_2} - \sum_{\substack{i_2, j_1, j_2 \geq 0 \\ i_2 + j_1 + j_2 = n-2}} A_{0, j_1} \otimes A_{i_2, j_2} \right. \\ &\quad \left. - z \sum_{\substack{i_1, i_2, j_1, j_2 \geq 0 \\ i_1 + i_2 + j_1 + j_2 = n-3}} A_{i_1,0} A_{0, j_1} \otimes A_{i_2, j_2} \right). \end{aligned}$$

Since

$$\begin{aligned} z \sum_{\substack{i_1, i_2, j_1, j_2 \geq 0 \\ i_1 + i_2 + j_1 + j_2 = n-3}} A_{i_1,0} A_{0, j_1} \otimes A_{i_2, j_2} &= \sum_{\substack{i_1, i_2, j_1, j_2 \geq 0 \\ i_1 + i_2 + j_1 + j_2 = n-3}} (A_{i_1+1, j_1} - A_{i_1, j_1+1}) \otimes A_{i_2, j_2} \\ &= \sum_{\substack{k, i_2, j_2 \geq 0 \\ k + i_2 + j_2 = n-2}} (A_{k,0} - A_{0,k}) \otimes A_{i_2, j_2}, \end{aligned}$$

all the terms inside the brackets in (4.4) cancel out, and therefore

$$P_n \mapsto P_n \otimes 1 + 1 \otimes P_n.$$

□

We next take the knot K as the framed unknot in \mathbb{R}^3 represented by a circle in a plane framed by the outward planar normal. Identify the thickened annulus in Theorem 4.11 with a neighborhood of the unknot. Since the turning number of the unknot in the plane equals ± 1 (depending on the orientation) and the turning number of the central circle in Theorem 4.11 equals 0, we must substitute $W_\mu \otimes W_\nu \mapsto a_2^{|\mu|} W_\mu \otimes a_1^{-|\nu|} W_\nu$ in the right hand side of (4.3) when the links are included in the plane. We get:




FIGURE 23. The sharp front in $J^0(M) \approx M \times \mathbb{R}$ of the conormal of a knot.

Corollary 4.12. *Consider the unknot colored by λ . Then, under the skein trace map $\text{Sk}_{a_1 a_2, z}(\mathbb{R}^3) \rightarrow \text{Sk}_{a_1, z}(\mathbb{R}^3) \otimes \text{Sk}_{a_2, z}(\mathbb{R}^3)$,*

$$\otimes_{\frac{1}{2}}^{\lambda} \text{circle} \mapsto \sum_{\mu, \nu} c_{\mu\nu}^{\lambda} a_2^{|\mu|} \text{circle}_{\mu}^1 \otimes a_1^{-|\nu|} \text{circle}_{\nu}^2,$$

and taking HOMFLYPT polynomials gives

$$\prod_{\square \in \lambda} \frac{a_1 a_2 q^{c(\square)/2} - a_1^{-1} a_2^{-1} q^{-c(\square)/2}}{q^{h(\square)/2} - q^{-h(\square)/2}} = \sum_{\mu, \nu} c_{\mu, \nu}^{\lambda} a_2^{|\mu|} a_1^{-|\nu|} \prod_{\square \in \mu} \frac{a_1 q^{c(\square)/2} - a_1^{-1} q^{-c(\square)/2}}{q^{h(\square)/2} - q^{-h(\square)/2}} \prod_{\square \in \nu} \frac{a_2 q^{c(\square)/2} - a_2^{-1} q^{-c(\square)/2}}{q^{h(\square)/2} - q^{-h(\square)/2}}. \quad \square$$

5. CONORMALS

The basic construction of Definition 1.1 requires a non-exact push-off of a link conormal. Our proof of Theorem 1.3 will involve considering a family of such conormals along the crossing change for the link. Here we give explicit constructions of such push-offs and discuss their basic properties.

We first give a front description of the conormal of the central curve $I \times 0 \subset I \times D^2$, where I is a 1-manifold.

Let I denote either an interval $I \approx [0, 1]$ or a circle $I \approx S^1$. Let D^2 denote the unit disk in \mathbb{R}^2 , $D^2 = \{v = (v_1, v_2) : v_1^2 + v_2^2 < 1\}$. Consider the product $I \times D^2$. Take coordinates $(t, v) \in I \times D^2$. Let $c > 0$ and consider a proper function $h : [0, 1] \rightarrow \mathbb{R}$ with the following properties

- $h(r) = cr^2$ for $0 \leq r \leq \frac{1}{2}$,
- $r = 0$ is the only critical point of the convex function h , and
- $h(r) \rightarrow \infty$ as $|r| \rightarrow 1$.

Let $f: I \times D^2 \rightarrow \mathbb{R}$ be given by

$$(5.1) \quad f(t, v) = f(v_1, v_2) = h \left(\sqrt{v_1^2 + v_2^2} \right)$$

and consider a closed 1-form τ along $I \times D^2$ with everywhere non-zero dt -component. Then the graph $\Gamma \subset T^*(I \times D^2)$ of $df + \tau$ is a graphical Lagrangian push-off of the conormal to $I \times 0 \subset T^*(I \times D^2)$. We denote it \mathbf{U}_τ .

Given any embedding $K: S^1 \rightarrow M$ with an (embedded) tubular neighborhood $N: D^2 \times S^1 \rightarrow M$, $N(0 \times S^1) = K$, we may transport the above perturbed shifted conormal to T^*M , giving a graphical Lagrangian push-off of the conormal of K . We denote it by $\mathbf{U}_{N,\tau}$. When the choice of N is implicit or irrelevant and $\tau = dt$, we sometimes write \mathbf{U}_K instead. The construction is local along K , so makes sense when K is only immersed, or when (even if K is embedded) its tubular neighborhood is only immersed; though in these cases the resulting conormal push-off need not be graphical or even embedded.

We will now fix specific choices for tubular neighborhoods and shifting 1-forms for a model family of links participating in the first skein relation, and make some geometric observations regarding the resulting conormals.

Consider coordinate cubes

$$B_\rho = \{(x_1, x_2, x_3) : |x_j| \leq \rho; j = 1, 2, 3\},$$

and a 1-parameter family of immersed intervals, $k(\sigma)$, $-\epsilon < \sigma < \epsilon$, in B_4 :

$$(5.2) \quad k(\sigma) := k_1(\sigma) \cup k_2(\sigma) := \{x_2 = x_3 = 0\} \cup \{x_1 = 0, x_3 = \sigma\phi(x_2)\},$$

where $\phi: [-4, 4] \rightarrow [0, 1]$ is a smooth cut-off function equal to 0 outside $[-\frac{3}{4}, \frac{3}{4}]$ and equal to 1 in $[-\frac{1}{2}, \frac{1}{2}]$. We think of the intervals as oriented in the positive x_1 - and x_2 -directions, respectively. Then the sign of σ determines the sign of the crossing, see Figure 24 (leftmost and middle links), and when $\sigma = 0$, $k(0)$ is immersed with a double point at the origin and when $\sigma \neq 0$, $k(\sigma)$ is embedded. Note that $k(\sigma) \cap \partial B_1$ is the four points

$$p_{1,\pm} = (\pm 1, 0, 0), \quad p_{2,\pm} = (0, \pm 1, 0),$$

for all σ , where the subscript sign is $-$ where $k(\sigma)$ enters the ball and $+$ where it exits.

Fix tubular neighborhoods of the components $k_1(\sigma), k_2(\sigma)$. Let $\delta > 0$, $0 < \eta \ll \delta$, and let $\phi: [-1, 1] \rightarrow [0, 1]$ be as above. We choose closed 1-forms in these tubular neighborhoods:

$$\begin{aligned} \tau_1 &= (\delta + \eta x_2 \phi'(x_1)) dx_1 + \eta \phi(x_1) dx_2, \\ \tau_2 &= (\delta - \eta x_1 \phi'(x_2)) dx_2 - \eta \phi(x_2) dx_1. \end{aligned}$$

Let $\mathbf{U}_{k(\sigma)}$ denote the union of the corresponding two graphical Lagrangian conormals shifted by τ_1 and τ_2 ,

$$(5.3) \quad \mathbf{U}_{k(\sigma)} = \mathbf{U}_{k_1(\sigma), \tau_1} \cup \mathbf{U}_{k_2(\sigma), \tau_2}.$$

We use coordinates (x, y) on $T^*\mathbb{R}^3$, $x = (x_1, x_2, x_3)$ and $y = (y_1, y_2, y_3)$.

Lemma 5.1. *In the above situation, $\mathbf{U}_{k(\sigma)}$, $\sigma \neq 0$ is embedded and $\mathbf{U}_{k(0)}$ has a clean self-intersection along the line ($\approx \mathbb{R}$):*

$$(x_1, x_2, x_3, y_1, y_2, y_3) = \left(\xi, \xi, x_3, \delta, \delta, \frac{\partial f}{\partial v_2}(\xi, x_3) \right), \quad |x_3| < \sqrt{1 - \xi^2},$$

where $\xi = \frac{\delta - \eta}{c}$ and f is as in (5.1).




FIGURE 24. Three terms of the skein relation

Proof. The Lagrangians $\mathbf{U}_{k_j(\sigma), \tau_j}$ are graphs of 1-forms λ_j , $j = 1, 2$, where λ_1 and λ_2 are supported in

$$\{(x, y) : |x_1| \leq 1, |(x_2, x_3)| < r\} \quad \text{and} \quad \{(x, y) : |x_2| \leq 1, |(x_1, x_3 - \sigma\phi(x_2))| < r\},$$

respectively. By definition $\lambda_j = df_j + \tau_j$, and for sufficiently small $|x_1|$ and $|x_2|$ we have

$$(5.4) \quad \tau_1 = \delta dx_1 + \eta dx_2, \quad \tau_2 = \delta dx_2 - \eta dx_1,$$

$$df_1 = \left(\frac{\partial f}{\partial v_1}(x_2, x_3) dx_2 + \frac{\partial f}{\partial v_2}(x_2, x_3) dx_3 \right) = cx_2 dx_2 + \frac{\partial f}{\partial v_2}(x_2, x_3) dx_3,$$

$$df_2 = \left(\frac{\partial f}{\partial v_1}(x_2, x_3 - \sigma) dx_2 + \frac{\partial f}{\partial v_2}(x_1, x_3 - \sigma) dx_3 \right) = cx_1 dx_1 + \frac{\partial f}{\partial v_2}(x_1, x_3 - \sigma) dx_3,$$

see (5.1), for small $\delta > 0$, $0 < \eta \ll \delta$, and large $c > 0$.

Intersections correspond to zeros of the difference of the defining forms. Now, $\lambda_1 - \lambda_2 = 0$ only if

$$((\delta - \eta) - cx_1) dx_1 + (cx_2 - (\delta - \eta)) dx_2 = 0,$$

which holds exactly when $x_1 = x_2 = \frac{\delta - \eta}{c} = \xi$ and then $y_1 = y_2 = \delta$. It remains to consider the (x_3, y_3) -coordinates where the Lagrangians are given by

$$\mathbf{U}_{k_1(\sigma), \tau_1} : (x_3, y_3) = \left(x_3, \frac{\partial f}{\partial v_2}(\xi, x_3) \right), \quad |x_3| < \sqrt{1 - \xi^2},$$

$$\mathbf{U}_{k_2(\sigma), \tau_2} : (x_3, y_3) = \left(x_3, \frac{\partial f}{\partial v_2}(\xi, x_3 - \sigma) \right), \quad |x_3 - \sigma| < \sqrt{1 - \xi^2}.$$

By convexity of f , the partial derivative $\frac{\partial f}{\partial v_2}(\xi, \cdot)$ is a strictly increasing function. Consequently, the curves are distinct for $\sigma \neq 0$ and agree for $\sigma = 0$. The lemma follows. \square

The curves $k(\sigma)$ and the corresponding Lagrangians $\mathbf{U}_{k(\sigma)}$ give the local model for links in a family with a crossing.

The link corresponding to the spliced crossing meets our box in an embedded curve $k_{\sphericalangle} \subset B_1 \cap \{x_3 = 0\}$ connecting $p_{1,-}$ to $p_{2,+}$ and $p_{2,-}$ to $p_{1,+}$, as in Figure 24. Let N be an embedded tubular neighborhood of k_{\sphericalangle} , and dt a pullback to N of a 1-form on k_{\sphericalangle} dual to the tangent vector. Then we take the (embedded, graphical) conormal

$$(5.5) \quad \mathbf{U}_{k_{\sphericalangle}} := \mathbf{U}_{N, dt}.$$

It will later be necessary to localize the difference between the model knots close to the self-intersection of the original link. We achieve this by scaling the above models as follows.




FIGURE 25. The link $k(0)$ and the front conormal fibers at the intersection point, with the clean intersection line.

We first consider the links $k(\sigma) \subset B_4$, $-\epsilon < \sigma < \epsilon$ and $k_{\succ} \subset B_4$. For $0 < \rho \leq 1$ consider the scaling map

$$s_\rho: B_4 \rightarrow B_{4\rho}, \quad s_\rho(x_1, x_2, x_3) = (\rho x_1, \rho x_2, \rho x_3).$$

Let $k(\sigma; \rho)$ and $k_{\succ}(\rho)$ denote the links in B_4 that consists of $s_\rho(k(\sigma))$ and $s_\rho(k_{\succ})$, respectively, with the segments of the x_1 and x_2 coordinate axes in $B_4 \setminus B_{4\rho}$ added.

We next consider the Lagrangians. Consider a tubular neighborhood N_ρ of $k_{\succ}(\rho)$ that agrees with N near ∂B_4 and agrees with $s_\rho(N)$ in $B_{4\rho}$. Note that the symplectomorphism $S_\rho: T^*B_1 \rightarrow T^*B_\rho$,

$$S_\rho(x, y) = (\rho \cdot x, \rho^{-1} \cdot y)$$

takes $\mathbf{U}_{k(\sigma)} \cap T^*(B_4 \setminus B_{4(1-\epsilon)})$ to $\mathbf{U}_{s_\rho(k(\sigma))} \cap T^*(B_{4\rho} \setminus B_{4\rho(1-\epsilon)})$ and similarly for $\mathbf{U}_{k_{\succ}}$ and $\mathbf{U}_{s_\rho(k_{\succ})}$. We take

$$(5.6) \quad \mathbf{U}_{k(\sigma; \rho)} = (\mathbf{U}_{N_\rho} \cap T^*(B_4 \setminus B_{4\rho})) \cup S_\rho(\mathbf{U}_{k(\sigma)}),$$

$$(5.7) \quad \mathbf{U}_{k_{\succ}(\rho)} = (\mathbf{U}_{N_\rho} \cap T^*(B_4 \setminus B_{4\rho})) \cup S_\rho(\mathbf{U}_{k_{\succ}}).$$

Finally, consider an immersed link $K_\times \subset M$ with a single double point where the tangent vectors of the branches are linearly independent. We will choose conormals of the corresponding three links participating in the first skein relation by choosing models which agree with the above standard constructions inside some box

$$(5.8) \quad b: B_4 = [-4, 4]^3 \rightarrow M \quad b^{-1}(K_\times) = k(0).$$

Let $K_{\text{out}} = K_\times \setminus b(B_4)$. Define links

$$K(\sigma) = K_{\text{out}} \cup b(k(\sigma)), \quad -\epsilon < \sigma < \epsilon,$$

$$K_{\succ} = K_{\text{out}} \cup b(k_{\succ}).$$

Fix a tubular neighborhood N_{out} of $K_\times \cap (M \setminus b(B_3))$ that agrees with the tubular neighborhood $b(N)$ of $b(k(0))$ in $b(B_4 \setminus B_3)$. Define

$$(5.9) \quad \mathbf{U}(\sigma; \rho) = \mathbf{U}_{N_{\text{out}}} \cup b_*(\mathbf{U}_{k(\sigma; \rho)}), \quad -\epsilon < \sigma < \epsilon, \quad 0 < \rho \leq 1,$$

$$(5.10) \quad \mathbf{U}_{\succ}(\rho) = \mathbf{U}_{N_{\text{out}}} \cup b_*(\mathbf{U}_{k_{\succ}(\rho)}), \quad 0 < \rho \leq 1.$$

Then $U(\sigma; \rho)$ are Lagrangian conormals of $K(\sigma)$ and $U_{\prec}(\rho)$ are Lagrangian conormals of $K_{\prec}(\rho)$.

If K_{\times} is connected, then $U(\sigma, \rho)$, $\sigma \neq 0$, and $U_{\prec}(\sigma)$ have either one or two connected components and that $U(\sigma, \rho)$ has one component exactly when $U_{\prec}(\rho)$ has two and vice versa. We call the homology classes in $U(\sigma; \rho)$ and $U_{\prec}\rho$ represented by the natural lifts of the oriented links $K(\sigma; \rho)$ and $K_{\prec}(\rho)$, respectively the *basic homology class*.

6. FLOW GRAPHS AND THE FIRST SKEIN RELATION

The purpose of this section is to describe the behavior of flow graph counts associated to curves stretching between shifted knot conormals and the zero section, as the knot undergoes a crossing change. That is, we study the flow graphs for the family $U(\sigma; \rho)$ for $|\rho|$ small.

6.1. Flow graph review. Consider the cotangent bundle T^*M of a manifold M and let $S \subset T^*M$ be a Lagrangian. We take S to be generic with respect to the projection $\pi: T^*M \rightarrow M$, meaning that $\pi|_S$ is an immersion outside a codimension one submanifold and outside the image of this locus S is locally given by the graph of a finite number of 1-forms, $\Gamma_{\beta_1} \cup \dots \cup \Gamma_{\beta_m}$. (The numbering of the 1-forms can only be locally defined.) Having fixed a metric on M , the duals of the differences $(\beta_i - \beta_j)^*$ give local vector fields along M . Consider a flow line of $(\beta_i - \beta_j)^*$ and its cotangent lift, the flow line lifted to the graphs of the 1-forms that defines it. A flow graph is a union of flow lines such that the closure of a corresponding union of cotangent lifts of its flow lines gives a closed oriented curve in $S \subset T^*M$, see [10, Definition 2.10] for details.

Given a flow graph we associate a continuous map of a surface Σ into $S \subset T^*M$ with boundary in S as follows. Consider its *flow strips*, i.e., the union of the lines in the fibers of T^*M over the flow lines of the flow graph. The union of all its flow strips is a surface with two types of boundary, the cotangent lift of the flow graph and residual boundary components in the cotangent fiber over vertices of valency 3 or higher (flow graphs in general position have no vertices of valency > 3). The residual boundary components are closed polygonal curves in a cotangent fiber, we fill them by polygonal disk. The filled surface is the surface Σ associated to the flow graph. Note that Σ is naturally oriented by the complex orientation of its flow strips. We will refer to Σ as the *flow surface*.

The following is a trivial but important example of a flow graph. Let $K \subset M$ be a knot, $N: S^1 \times D^2 \rightarrow M$ a fixed tubular neighborhood, and $U_N \subset T^*M$ the corresponding shifted conormal.

Lemma 6.1. *There is a unique embedded flow graph for $(T^*M, U_N \sqcup M)$. It is a loop, and the boundary of its flow surface is $S^1 \times 0$ in U_N and K in M .*

Proof. By definition of U_N , the difference vector field in $(S^1 \times D^2)$ -coordinates is $\partial_t - \nabla f$ and $f: D^2 \rightarrow [0, \infty)$ has only one critical point, a non-degenerate minimum at the origin. \square

Flow graphs are related to holomorphic curves [21, 26, 37]. This relation was studied in [10] for Legendrians in 1-jet spaces that project to immersed exact Lagrangians in the corresponding cotangent bundles. Under fiber scaling, holomorphic curves ending on these Lagrangians (Gromov-) limit to holomorphic maps into the zero-section (necessarily constant), but may be recovered by rescaling leading to a one-to-one correspondence between rigid holomorphic disks and rigid flow trees. A more general situation was considered in [11] where the limit Lagrangian does carry non-constant holomorphic curves; in this case,

the one-to-one correspondence relates the curves in the original geometry with holomorphic curves on the limit Lagrangian with flow trees attached. Here we consider somewhat more general situations but the results we need can be extracted from the analysis in the papers mentioned. We review and extract the relevant results relating holomorphic curves and flow graphs in Appendix A.

6.2. Description of the result. Let us outline the result of the calculations in the remainder of this section. We study the family $\mathbf{U}(\sigma; \rho)$. For $\sigma \gg 0$, these conormals are graphical, hence their flow objects are described entirely by Lemma 6.1. However, this is not the case in a neighborhood of $\sigma = 0$.

For these conormals, we will call a flow graph *basic* if the boundary of the flow surface traces out the basic class on the conormal.

We will further perturb the conormal to the singular knot to a $\tilde{\mathbf{U}}(0; \rho)$, which we will show carries exactly two basic flow graphs. One is a flow loop tracing out a knot isotopic to K_- , and this flow loop persists for similarly perturbed $\tilde{\mathbf{U}}(\sigma; \rho)$ with $|\sigma|$ sufficiently small (Lemma 6.6).

The other is a flow graph ‘with corners’ corresponding to a holomorphic curve with two corners along the self-intersection of $\tilde{\mathbf{U}}(\sigma; 0)$ (Lemma 6.9). This graph ‘glues up’ (as would the curve) to a flow graph on $\tilde{\mathbf{U}}(\sigma; \rho)$ for $0 < \sigma < \epsilon$. The boundary of the holomorphic curve corresponding to this flow graph traces out K_{\prec} along the zero section.

This behavior illustrates the well-definedness (and need for) skein-valued curve counting. Indeed, we should have $\psi(\tilde{\mathbf{U}}(\sigma; \rho))$ independent of $\sigma > 0$. For $\sigma \gg 0$, the conormal is graphical, and so the only flow graph is the loop tracing out K_- itself. For $\sigma < \epsilon$, we have flow graphs tracing out K_+ and K_{\prec} ; we will see later (confer the proof of Theorem 9.13 for the sign effect of attaching a flow line) that both come with positive sign. The holomorphic curve corresponding to the nontrivial graph has Euler characteristic -1 . Then with $\sigma_+ \gg 0$ and $0 < \sigma < \epsilon$, invariance demands:

$$(6.1) \quad [K_+] = \psi(\tilde{\mathbf{U}}(\sigma_+; \rho)) = \psi(\tilde{\mathbf{U}}(\sigma; \rho)) = [K_-] + z[K_{\prec}].$$

This is precisely the first skein relation.

More relevant to work in the present article is a *different* appearance of the first skein relation. The above description of flow graphs (and the correspondence between flow graphs and holomorphic curves) amounts to the assertion, for $0 < \sigma < \epsilon$,

$$\begin{aligned} \psi(\tilde{\mathbf{U}}(\sigma; \rho)) &= [K_-] + z[K_{\prec}] \\ \psi(\tilde{\mathbf{U}}(-\sigma; \rho)) &= [K_-]. \end{aligned}$$

We could have deduced these formulas by comparing to the corresponding ψ for $|\sigma|$ large, but our analysis here has shown something sharper, which will be crucial for our later work: the terms in the above formulas are a direct description of the holomorphic curves which appear. But in any case, the promised appearance of the skein relation, following from the above, is:

$$\psi(\tilde{\mathbf{U}}(\sigma; \rho)) - \psi(\tilde{\mathbf{U}}(-\sigma; \rho)) = [K_{\prec}] = \psi(\mathbf{U}_{\prec}).$$

6.3. A choice of perturbation. Recall we have fixed in (5.8) a choice of coordinates $B_4 = [-4, 4]^3 \rightarrow M$; all the knots of interest agree in the complement of the image of $B_1 = [-1, 1]^3$. We keep our notation for rescalings: $s_\rho(x_1, x_2, x_3) = (\rho x_1, \rho x_2, \rho x_3)$, and $S_\rho := T^*s_\rho$.

Let

$$(6.2) \quad \mathbf{U}^{\text{out}}(\rho) = \mathbf{U}(\sigma; \rho) \cap T^*(M \setminus B_\rho) = \mathbf{U}_{\prec}(\rho) \cap T^*(M \setminus B_\rho).$$

Note that the first intersection is independent of σ . Also, let

$$(6.3) \quad \mathbf{U}^{\text{in}}(\sigma, \rho) = \mathbf{U}(\sigma; \rho) \cap S_\rho(T^*B_1),$$

$$(6.4) \quad \mathbf{U}_{\prec}^{\text{in}}(\rho) = \mathbf{U}_{\prec}(\rho) \cap S_\rho(T^*B_1).$$

Then S_ρ identifies $\mathbf{U}^{\text{in}}(\sigma, \rho)$ and $\mathbf{U}_{\prec}^{\text{in}}(\rho)$ with $\mathbf{U}_{k(\sigma)} \subset T^*B_1$, see (5.3) and $\mathbf{U}_{k_{\prec}} \subset T^*B_1$, see (5.5), respectively. Similarly, S_ρ identifies $\mathbf{U}^{\text{out}}(\rho) \cap T^*B_{3\rho} \setminus B_{2\rho}$ with $\mathbf{U}_{k(\rho)} \cap T^*(B_3 \setminus B_2) = \mathbf{U}_{k_{\prec}} \cap T^*(B_3 \setminus B_2)$.

Recall the intersection points

$$p_{1,\pm} = (\pm 1, 0, 0), \quad p_{2,\pm} = (0, \pm 1, 0).$$

of $k(\sigma)$ and k_{\prec} with ∂B_1 . Identify ∂B_1 with $\partial B_\rho \subset M$.

The Lagrangian $\mathbf{U}^{\text{out}}(\rho)$ is the graph of a closed 1-form in a tubular neighborhood of K in $M \setminus B_\rho$. Let $V^{\text{out}}(\rho)$ denote the vector field dual to this 1-form. Then the flow $\Phi^{\text{out}}(\rho)$ of $V^{\text{out}}(\rho)$ has flow lines that start at $p_{1,+}$ and $p_{2,+}$ and end at $p_{1,-}$ and $p_{2,-}$. The flow correspondingly takes a neighborhood $P_+ \subset \partial B_1$ of $\{p_{1,+}, p_{2,+}\}$ to a neighborhood $P_- \subset \partial B_1$ of $\{p_{1,-}, p_{2,-}\}$.

We write $P_{\pm}(\rho) = P_{1,\pm}(\rho) \cup P_{2,\pm}(\rho)$ for the components of these neighborhoods around the corresponding $\{p_{1,\pm}, p_{2,\pm}\}$. Note that when K_{\prec} is connected, $\Phi^{\text{out}}(\rho)$ takes $P_{1,+}(\rho)$ to $P_{1,-}(\rho)$ and $P_{2,+}(\rho)$ to $P_{2,-}(\rho)$; when K_{\prec} has two components, $\Phi^{\text{out}}(\rho)$ takes $P_{1,+}(\rho)$ to $P_{2,-}(\rho)$ and $P_{2,+}(\rho)$ to $P_{1,-}(\rho)$.

The vector field $V^{\text{out}}(\rho)$ has the form

$$V^{\text{out}} = \partial_t - \nabla f,$$

in coordinates $(t, v) \in I \times D^2$ on a tubular neighborhood, where f is as in (5.1). This implies that its flow contracts the fibers. More precisely, if we pick coordinates $\xi \in D^2$ on $P_{1,\pm}$ and $P_{2,\pm}$, respectively, centered at the flow line where $\nabla f = 0$ we find that the flow map acts as multiplication by e^{-kT_j} , $k > 0$, where T_j is the flow time (the flow time T_j is proportional to δ^{-1} , where δ is the size of the shift of the conormal),

$$(6.5) \quad \Phi^{\text{out}}(\xi) = e^{-kT_j} \xi, \quad j = 1, 2$$

where ξ in the left-hand lies in $P_{j,+}$ and the right-hand side lies in $P_{i,-}$, where $i = j$, if K_{\prec} is connected and $i \neq j$ is if it has two components.

We will consider small perturbations of the flow, corresponding to a fixed model perturbation of $\mathbf{U}_{k(\sigma)}|_{T^*(B_3 \setminus B_2)} = \mathbf{U}_{k_{\prec}}|_{T^*(B_3 \setminus B_2)}$, supported in a neighborhood of the negative x_1 -axis. We will denote perturbed objects as e.g., $\tilde{\mathbf{U}}, \tilde{V}$. Note that $\tilde{V}^{\text{out}}(\rho)$ then differs from $V^{\text{out}}(\rho)$ only on $B_{3\rho} \setminus B_{2\rho}$. The exact form of the perturbation will not be important but its effect on the x_3 -coordinate of the return map will.

Definition 6.2. We say that the perturbation is positive (negative) if:

- K_{\prec} is connected and $x_3(\tilde{\Phi}^{\text{out}}(0 \in P_{1,+}))$ is positive (negative).
- K_{\prec} is disconnected and $x_3(\tilde{\Phi}^{\text{out}}(0 \in P_{2,+}))$ is positive (negative).

We will write $\kappa > 0$ for the size of the perturbation with respect to the standard metric on B_4 .

Lemma 6.3. *For given perturbation of size $\kappa > 0$, if the shift δ is sufficiently small, then for small enough choice of the original disk $P_{2,+}$ (or $P_{1,+}$ as the case may be), its image under $\tilde{\Phi}^{\text{out}}(\rho)$ in $P_{1,-}$ lies in the region $x_3 > 0$ ($x_3 < 0$) if the perturbation is positive (negative).*

Proof. Note that the flow time $T_j \rightarrow \infty$ as $\delta \rightarrow 0$. Thus, for sufficiently small shift δ , if $x_3(\tilde{\Phi}^{\text{out}}(0)) = \eta > 0$ then $\tilde{\Phi}^{\text{out}}$ takes the whole disk $P_{j,+}$ to the region $\frac{1}{2}\eta < x_3 < \frac{3}{2}\eta$, see (6.5). \square

6.4. Flow loops. As the Lagrangian $U_{\succ}(\rho)$ is graphical, Lemma 6.1 shows that it supports exactly one flow graph, the basic flow loop(s), in the basic homology class. The same is true for the perturbed $\tilde{U}_{\succ}(\rho)$, so we have:

Lemma 6.4. *The flow surface of the basic flow loop(s) of $\tilde{U}_{\succ}(\rho)$ has boundary which is at C^1 -distance $\mathcal{O}(\kappa\rho)$ to $S^1 \times 0$ in all components of $\tilde{U}_{\succ}(\rho)$ and of C^1 -distance $\mathcal{O}(\kappa\rho)$ to $K_{\succ}(\rho)$ in M . In particular, the boundary component of the flow loop(s) in M converges to K_{\times} as $\rho \rightarrow 0$.*

Proof. Uniqueness of flow loops under small perturbation follows from the transversality of the original flow loop (which we discussed in Lemma 6.1). To see the convergence, note that $K_{\succ}(\rho)$ agrees with K_{\times} in $M \setminus s_{\rho}(B_1)$ by construction, and that $s_{\rho}(B_1)$ tends to a point as $\rho \rightarrow 0$. \square

We will re-describe the determination of the above loop as a fixed point calculation. Inside B_1 , the conormal $U_{k_{\succ}} \subset T^*B_1$ is the graph of a 1-form; we denote its dual vector field as V_{\succ}^{in} and the corresponding flow map

$$\Phi_{\succ}^{\text{in}}: P_{-} \rightarrow P_{+}.$$

It is clear from the definition of $U_{k_{\succ}}$ that the map Φ_{\succ}^{in} is a contraction of the 2-disks in P_{\pm} .

We denote by $\Phi_{\succ}^{\text{in}}(\rho)$ the corresponding map when rescaled and transplanted to the image of B_{ρ} in M , which gives the corresponding flow map for $U^{\text{in}}(\rho)$.

Lemma 6.5. *If $K_{\succ}(\rho)$ has one component then its flow loop appears as the unique fixed point of*

$$\tilde{\Phi}^{\text{out}}(\rho) \circ \Phi_{\succ}^{\text{in}}(\rho) \circ \tilde{\Phi}^{\text{out}}(\rho) \circ \Phi_{\succ}^{\text{in}}(\rho)$$

and if it has two, then the two components of the basic flow loop correspond to the unique fixed points of

$$\tilde{\Phi}^{\text{out}}(\rho) \circ \Phi_{\succ}^{\text{in}}.$$

Proof. Certainly any flow loop is such a fixed point. Uniqueness holds because both $\tilde{\Phi}^{\text{out}}$ and Φ_{\succ}^{in} are contractions. \square

We turn to study the flow loops for $\tilde{U}(\sigma, \rho)$. As above we will consider these as fixed points of flow maps. Outside B_{ρ} , $\tilde{U}(\sigma, \rho) = \tilde{U}(\rho)$ and we have the corresponding flow map $\tilde{\Phi}^{\text{out}}(\rho): P_{+} \rightarrow P_{-}$ exactly as before. Inside $B_{\rho} \approx B_1$ we have instead the flow $\Phi^{\text{in}}(\sigma, \rho): P_{j,-} \rightarrow P_{j,+}$, $j = 1, 2$, defined by $U_{k(\sigma; \rho)}$, see (5.9) and (5.3). As for the flows discussed above also this flow gives a contraction $P_{j,-} \rightarrow P_{j,+}$, $j = 1, 2$.

Lemma 6.6. *For all sufficiently small $|\sigma|$, there is a unique basic flow loop (or union of flow loops) for $U(\sigma, \rho)$. This loop is isotopic to K_{+} for positive perturbations and to K_{-} for negative perturbations.*

Proof. The flow loops correspond to fixed points of flow maps $\tilde{\Phi}^{\text{out}}(\rho) \circ \Phi^{\text{in}}(\sigma, \rho)$ in case K_{\pm} has two components and $\tilde{\Phi}^{\text{out}}(\rho) \circ \Phi^{\text{in}}(\sigma, \rho) \circ \tilde{\Phi}^{\text{out}}(\rho) \circ \Phi^{\text{in}}(\sigma, \rho)$ in case it is connected.

Consider first the two component case. Then the flow loop of the conormal of the component that contains the x_2 -axis lies in the $x_3 = 0$ plane near the crossing: the map $\tilde{\Phi}^{\text{out}}(\rho): P_{2,+} \rightarrow P_{2,-}$ was not perturbed, $\tilde{\Phi}^{\text{out}}(\rho) = \Phi^{\text{out}}(\rho): P_{2,+} \rightarrow P_{2,-}$, the knot agrees with the x_2 -axis near the origin, and the flow loop is just the original standard flow loop.

Consider now the other flow loop as a fixed point of the map $\tilde{\Phi}^{\text{out}}(\rho) \circ \Phi^{\text{in}}(\sigma; \rho): P_{1,-} \rightarrow P_{1,-}$. This fixed point has positive or negative x_3 -coordinate depending on the sign of the perturbation. Since the flow of the fiber disk that defines $\Phi^{\text{in}}(\sigma)$ has the form $(x_2, x_3) = e^{-kt}(x_2^0, x_3^0)$, the sign of the x_3 -coordinate does not change along the flow line and the result follows in this case.

The case of a single knot is similar: The sign of the x_3 -coordinate does not change along the flow lines near the crossing (the flow lines giving $\Phi^{\text{in}}(\sigma; \rho)$). Along the flow lines corresponding to $\tilde{\Phi}^{\text{out}}(\rho) = \Phi^{\text{out}}(\rho): P_{1,+} \rightarrow P_{2,-}$, the x_3 -coordinate is scaled along the flow by e^{-kT} , where T is the flow time. It follows as before that the sign of the crossing agrees with that of the perturbation. \square

6.5. Flow graphs with corners for $\tilde{U}(0, \rho)$. In this subsection, we will study flow graphs for the immersed Lagrangian $\tilde{U}(0; \rho)$ which is our perturbed conormal of the singular knot K_{\times} .

Recall that the flow graphs appropriate to immersed Lagrangians may have edges asymptotic to the critical locus of differences of the local defining functions of the Lagrangians i.e., zeroes of differences of the local 1-forms of the Lagrangian. Such flow graphs correspond to holomorphic curves with corners at the immersed points of the Lagrangian over the critical locus [10]. We will refer to these as flow graphs with corners.

The local model near the immersed point of K_{\times} was described in (5.3); it has two components $U_{k(0)} = U_{k_1(0); \tau_1} \cup U_{k_2(0); \tau_2}$. We have the following.

Lemma 6.7. *For $U_{k(0)}$ and the zero section, there is a unique flow graph which starts at $q_{-,1} \in P_{-,1}$ ($q_{-,2} \in P_{-,2}$) and has a corner corresponding to the intersection $U_{k_1(0); \tau_1} \cap U_{k_2(0); \tau_2}$. It is a rigid flow tree with exactly one corner, and ends at some point $q_{+,2} \in P_{+,2}$ ($q_{+,1} \in P_{+,1}$).*

We denote the resulting map as

$$\Phi_0^{\text{in}}: P_- \rightarrow P_+, \quad \Phi^{\text{in}}(q_{-,k}) = \Phi^{\text{in}}(q_{+,k'}),$$

where $k, k' \in \{1, 2\}$ and $k \neq k'$.

Proof. Consider flow graphs with one corner that starts at $q_{-,1} \in P_{-,1}$. The flow of the difference vector field dual to the difference of 1-forms defining L_1 and L_2 oriented as the cotangent lift are shown in blue in Figure 26. In order for the one corner flow tree to have a corner at the intersection of L_1 and L_2 , the initial flow line between L_1 and the zero-section must intersect the stable/un-stable manifold of the difference vector field that is oriented toward the intersection in L_1 (vertical in Figure 26). (In the language of [10], the flow tree has a Y_0 -vertex at this point.) Following the flow line to the intersection between L_1 and L_2 , then jumping to L_2 and follow the difference flow line back to the Y_1 -vertex we switch flows and then follow the flow line between L_2 and the zero section to $P_{2,+}$.

Figure 26 shows $U_{k_j(0); \tau_j}$, $j = 1, 2$, with the flows of their difference together with the flows of their respective differences with the zero section. The first statement follows from




FIGURE 26. Flows along L_1 and L_2 ; flow lines in blue refer to the difference between L_1 and L_2 ; flow lines in red refer to flows between L_j and the zero section. The small parameter $\eta > 0$ in (5.4) makes the central flow lines of L_j and the 0-section non-parallel to the stable/unstable flow lines for the flow of the difference between L_1 and L_2 . The critical point in the flow between L_1 and L_2 corresponds to the intersection line in $L_1 \cap L_2$.

the fact that for each flow line between L_1 (L_2) and the zero section there is a unique intersection point (which then gives the Y_0 -vertex) with the stable/unstable manifolds of the (the projection to B_1) of the intersection line in $\mathbf{U}_{k_1(0); \tau_1} \cap \mathbf{U}_{k_2(0); \tau_2}$. (See Figure 27 for an illustration of a flow tree.)

The case of flow graphs that start at $q_{-,2} \in P_{-,2}$ is similar using instead the stable/unstable manifold oriented toward the critical locus in L_2 (the horizontal line in Figure 26). \square

Lemma 6.8. *The map $\Phi_0^{\text{in}}: P_- \rightarrow P_+$ is a contraction.*

Proof. This follows from the properties of the flows of the vector fields dual to $d\tau_j + df$, which are the difference forms of $\mathbf{U}_{k_j(0); \tau_j}$ and the zero section B_1 . The flows are depicted in Figure 26. Consider the flow map with initial condition in $P_{-,1}$ near $(x_1, x_2, x_3) = (-1, 0, 0)$. The map follows the flow line of $d\tau_1 + df_1$ until it hits the stable/unstable manifold of $(d\tau_j + df_1) - (d\tau_2 + df_2)$, then it follows the flow line of $d\tau_2 + df_2$. We note that both flows contract the x_3 -direction in the standard way. We check that the total flow map contracts the remaining direction in $P_{-,1} \rightarrow P_{+,2}$. The first flow is the solution to the differential equation

$$\begin{cases} \dot{x}_1 &= \delta, \\ \dot{x}_2 &= cx_2 + \eta \end{cases}$$

with general solution

$$\begin{cases} \dot{x}_1(t) = x_1(0) + \delta t, \\ \dot{x}_2(t) = (x_2(0) + \frac{\eta}{c}) e^{ct} - \frac{\eta}{c} \end{cases} .$$

The second is similar with the roles of x_1 and x_2 switched. This means the flow lines of the first flow focuses around the central flow line starting at $x_2 = \frac{\eta}{c}$, and the flow lines of the second focuses around the flow line ending at $x_1 = \frac{\eta}{c}$. It is then straightforward to see that the flow map is a contraction. \square




FIGURE 27. The unique flow tree with initial condition $q_{1,-} \in P_{1,-}$. To the left the tree itself, to the right, the boundary of its flow surface in the various Lagrangians.

Recall that $\tilde{U}_{\succ}(\rho)$ also had a unique flow loop; we write $\gamma_{\succ}(\rho)$ for the boundary of the corresponding flow surface in the 0-section M (which is just the flow loop itself). Fix a tubular neighborhood N of the central immersed curve K_{\times} close to $\gamma_{\succ}(\rho)$ lies in an $\mathcal{O}(\rho)$ disk sub-bundle of N and $\gamma_{\succ}(\rho) \cap N \setminus s_{\rho}(B_1)$ is given by a smooth section over $K_{\times} \setminus s_{\rho}(B_1)$ since it is a flow of a smooth differential equation. .

Lemma 6.9. *There is a unique basic flow graph with corners for $\tilde{U}(0, \rho)$. It has two corners. We write $\gamma(0; \rho)$ for the boundary in the zero section M of the corresponding flow surface. The curve $\gamma(0, \rho)$ lies in an $\mathcal{O}(\rho)$ disk subbundle of N and is given by a smooth section over $K_{\times} \setminus s_{\rho}(B_1)$. Furthermore, the C^1 -distance between the section of $\gamma_{\succ}(\rho)$ and $\gamma(0; \rho)$ in $N \setminus s_{\rho}(B_1)$ is of the order of magnitude ρ .*

Proof. The flow graphs can be computed exactly as in Lemma 6.5, save with $\Phi_{\succ}^{\text{in}}(\rho)$ replaced with the appropriate rescaling $\Phi_0^{\text{in}}(\rho)$ of Φ_0^{in} .

Both the actual inside maps are defined by applying s_{ρ} to fixed maps on the unit box. This implies that their distance in N is $\mathcal{O}(\rho)$. The fact that the outside map is a standard contraction given by flows of the same differential equation then gives the 1-1 correspondence with corresponding sections at $\mathcal{O}(\rho)$ C^1 -distance as claimed. \square

6.6. Flow graphs for $\tilde{U}(\sigma, \rho)$, $\sigma \neq 0$. The family of Lagrangians $\tilde{U}(\sigma, \rho)$, $-\sigma_0 < \sigma < \sigma_0$ constitutes a Lagrangian regular homotopy. For $\sigma \neq 0$, $U(\sigma, \rho)$ are embedded and $\tilde{U}(0, \rho)$ self-intersects cleanly along a curve, i.e., has a Bott degenerate intersection along \mathbb{R} , see Lemma 5.1. Take σ_0 small, then far from the intersection curve the Lagrangian regular homotopy is simply a small Lagrangian isotopy, which corresponds to a shift along a closed 1-form in the cotangent bundle of $U(0, \rho)$.

From the point of view of the Bott intersection, Lagrangians nearby the Bott degenerate Lagrangian correspond to closed 1-forms on the Bott manifold. We will next describe $U(\sigma, \rho)$ from this perspective and then show that it allows us to determine flow graphs of $U(\sigma; \rho)$ from those of $U(0, \rho)$.

We use T^*B_4 -coordinates near the self-intersection of $U(0, \rho)$. Here, the components of the Lagrangian $U(\sigma, \rho)$ have two components $U_{k_1(\sigma), \tau_1}$ and $U_{k_2(\sigma), \tau_2}$ given by graphs of the 1-forms in (5.4). Lemma 5.1 shows that the intersection curve $U_{k_1(0), \tau_1} \cap U_{k_2(0), \tau_2}$ lies over a line segment at $(x_1, x_2) = (\xi, \xi)$ parallel to the x_3 -axis where the difference 1-form of $U_{k_2(\sigma), \tau_2}$ and $U_{k_1(\sigma), \tau_1}$ along this intersection is given by

$$(6.6) \quad (d\tau_2 + df_2) - (d\tau_1 + df_1) = \frac{\partial f}{\partial v_2}(\xi, x_3 - \sigma) - \frac{\partial f}{\partial v_2}(\xi, x_3) = \sigma g(x_3) dx_3,$$

where $g(x_3) > 0$ for all x_3 . We use this line segment parallel to the x_3 -axis to parameterize the intersection, the shifting 1-form is then $\sigma g(x_3)dx_3$.

Describing flow graphs of $U(\sigma, \rho)$ from data on $U(0, \rho)$ is a simple case of Morse theory for functions with Bott-degenerate critical loci, see e.g., [21, 56, 12] for applications to Lagrangians with clean intersections. Note first that any flow graph of $U(\sigma, \rho) \cup M$ converges to a flow graph of $U(0, \rho)$, possibly with corners at the self-intersection of $U(0, \rho)$. (This convergence should be understood in the Gromov compactness sense: the convergence of flow lines is uniform on compact subsets of the domain of the flow graph.) Further, if we rescale $U(\sigma, \rho)$ in a neighborhood of its intersection we see that if the limit is a flow graph with corners, then the rescaled flow converges to a flow line of $\pm g(x_3)\partial_{x_3}$ connecting corners. We call such a flow line of $\pm g(x_3)\partial_{x_3}$ a *flow line inside the intersection line*. Using standard finite dimensional Morse theory, it is easy to glue a flow graph with corners of $U(0, \rho)$, connected by a flow line in the intersection line to a flow graph of $U(\sigma, \rho)$ provided the direction of the flow line (given by $\pm g(x_3)\partial_{x_3}$) agrees with the shift given by $\sigma g(x_3)dx_3$, $\sigma \neq 0$. As in the case of flow graphs, we consider the cotangent lift of the flow line inside the intersection line, which is the flow line itself lifted to the components and oriented as cotangent lifts of edges of flow graphs.

We will use the following notation for flow graphs with corners connected by flow lines in the intersection line. Given a flow graph of $U(0, \rho)$, possibly with corners, a *positive (resp. negative) pregluing* is the data of flow lines of $g(x_3)dx_3$ (resp. of $-g(x_3)dx_3$) inside the intersection line that connect the corners in pairs such that the flow graph together with the flow line in the intersection admits a continuous cotangent lift to $U(\sigma_0; \rho)$ (resp. to $U(-\sigma_0; \rho)$) for $\sigma_0 > 0$. Note that flow graphs without corners tautologically admit both positive and negative gluing.

The correspondence between flow graphs with flow lines in the intersection line of the immersed Lagrangian and actual flow graph on nearby Lagrangians is then as follows.

Lemma 6.10. *As $\sigma \rightarrow 0$, flow graphs determined by $U(\sigma; \rho)$, $\sigma \neq 0$, converge (in the sense of Gromov compactness to flow graphs on $U(0; \rho)$ with corners at the intersection.*

Furthermore, for all sufficiently small $|\sigma|$ there is a 1 – 1 correspondence between flow graphs on $U(\sigma; \rho)$, $\sigma < 0$ ($\sigma > 0$) and flow graphs of $U(0, \rho)$ with positive (negative) pregluing. The cotangent lifts of corresponding flow graphs are arbitrarily C^0 -close and C^1 -close outside $\mathcal{O}(\sigma)$ -neighborhoods of the corners.

Proof. The first statement is standard as discussed above. The second statement of C^0 - and C^1 -convergence follows easily from properties of the flow of the ordinary differential equation

$$\begin{cases} \dot{x}_1 &= cx_1, \\ \dot{x}_2 &= cx_2, \\ \dot{x}_3 &= \sigma \end{cases},$$

for small $|\sigma| \geq 0$. See Figure 28 for illustration. □

Corollary 6.11. *For $|\rho| \ll 0$, there is at most one basic flow graph with (a non-zero number of) vertices for $\tilde{U}(\sigma; \rho)$. For positive (negative) perturbation, this graph appears only for $\sigma > 0$ ($\sigma < 0$). It has two trivalent vertices and the components of its cotangent lift in M is C^0 -close to $K_{\prec}(\rho)$ and the corresponding holomorphic curve has boundary smoothly isotopic to it.*

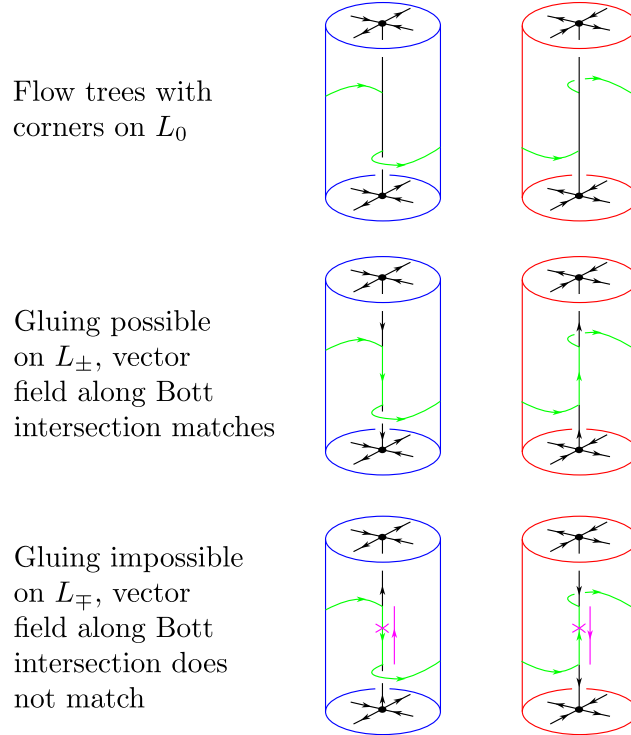


FIGURE 28. Flow tree gluing

Proof. By Lemma 6.9, there is a unique flow graph with corners for $\tilde{U}(\sigma; \rho)$. It has exactly two corners, hence admits either positive or negative pregluing (there is a flow line between them going one way or the other). To determine which, the key point is the following fact: if the perturbation of the outside flow is positive (negative) then the corner of the flow tree emanating from $P_{1,-}$ lies above (below) the corner in the flow tree emanating from $P_{2,-}$.

To see this consider the positive case, the negative case is similar. Note first that the x_3 -component of along flow line of the difference of the two components of $L_{k(0)}$ (the flow line that asymptotes to the intersection) is constant, set $\sigma = 0$ in (5.4). It follows that the corner in the flow tree of the flow map Φ_0^{in} emanating from the fixed point in $P_{1,-}$ has x_3 -coordinate given by scaling the x_3 -coordinate of the fixed point by e^{-kT} of the fixed point, where T is the flow time from the fixed point to the splitting Y_0 -vertex. Similarly the following corner of the flow tree of the flow tree map Φ_0^{in} emanating from the fixed point in $P_{2,-}$ lies at further scaling by $e^{-kT'}$, where T' is the flow time along the unperturbed outer arc and the short inside flow time to the splitting. It follows that the x_3 -coordinates of the corners are as claimed.

Then Lemma 6.10 asserts the existence of the desired flow graph for $\tilde{U}(\sigma; \rho)$, and the non-existence of any others. (Note we already studied the flow graphs without vertices and/or corners in Lemma 6.6.)

Finally, the boundary of the cotangent lift in M are the flow lines through the fixed point that intersect the $\tilde{U}(0, \rho)$ with the stable/unstable manifolds of the line of intersections. For small perturbation this is C^0 -close to $K_{\prec}(\rho)$. The isotopy statement follows for the local gluing model for Y_0 -vertices, see [10, Section 6.1.5]. \square

7. REVIEW OF SKEIN-VALUED CURVE COUNTING

Here we briefly recall the setup and main results of [19, 17, 16], with particular attention to the various topological ‘brane data’ involved (as these brane data choices will play an important role in our subsequent calculations). Let X be a symplectic 6-manifold with trivial(ized) first Chern class, and $L \subset X$ a (not necessarily connected) smooth Lagrangian submanifold with trivial(ized) Maslov class. Fix an almost complex structure J , standard in a neighborhood of L . In case X and L are non-compact, we require that (X, L, J) at infinity is such as to ensure that Gromov compactness arguments remain valid. Fix (as usual for holomorphic curve counting) a spin structure on L , or more generally, a background class in $H^2(X, \mathbb{Z}/2\mathbb{Z})$ and a relative spin structure.

The skein valued curve counting requires additional ‘brane data’ to control various framings. We fix a vector field ξ on L with transverse zeroes, and a 4-chain Ξ with $\partial\Xi = 2L$, such that, in a Weinstein neighborhood of L , the neighborhood of $\partial\Xi$ in Ξ is $\overline{\mathbb{R}_{\geq 0} \cdot \pm J\xi}$. Fix a bounding chain b for $\gamma := (\Xi \setminus \partial\Xi) \cap L$. Given a map $u: (C, \partial C) \rightarrow (X, L)$ which is holomorphic and immersed near ∂L , the *corrected framing* of ∂C is

$$(7.1) \quad \xi + u(C) \cap \Xi + u(\partial C) \cap b.$$

Here, the first term is a vector field; for generic u it will be nowhere parallel to ∂C and hence define a framing; the next two terms are integers, and the sum is understood via the fact that topological types of framings of a knot in a 3-manifold form a \mathbb{Z} -torsor. As explained in [19], this corrected framing is invariant under deformations of ξ, Ξ, b . The space of vector fields is connected, and the dependence on global aspects of the choices (Ξ, b) is as follows: the difference between two such choices gives an element of $H_4^{BM}(X) \oplus H_2^{BM}(L)$, and the corresponding difference in corrected framings is the intersection of this class with $[u] \in H_2(X, L)$.

Fix a class $d \in H_2(X, L)$. Recall that we call a map $u: (C, \partial C) \rightarrow (X, L)$ *bare* if $u^*(\omega)$ gives positive symplectic area to every irreducible component of C . We write $\chi(u)$ for Euler characteristic of a smoothing of the domain of u . Given some system of (weighted, multivalued) perturbations for the holomorphic curve equation, if u is a transverse solution to the perturbed equation, we write $\text{wt}(u)$ for the local contribution of this solution. The main result of [19, 17, 16] is that there is a system of perturbations to the holomorphic curve equation so that the following sum over bare solutions is well defined and deformation invariant:

$$(7.2) \quad Z_{X,L,d} = \sum_{[u]=d} \text{wt}(u) \cdot z^{-\chi(u)} \cdot [u(\partial\Sigma)] \in \text{Sk}(L) \otimes_{\mathbb{Z}[z]} \mathbb{Q}[[z]]$$

To sum over d , one generally must extend scalars by the Novikov ring $\mathbb{Q}[[H_2(X, L)]]$. However, when (as will be the case throughout this article) the symplectic form is exact, $\omega = d\lambda$, we detect the symplectic area by integrating λ over the image of $[u]$ in $H_1(L)$, which in turn can be read from the class $[u(\partial\Sigma)] \in \text{Sk}(L)$. Thus we may instead simply complete $\text{Sk}(L)$ along the cone of λ positive classes. Going forward, we will omit this completion from the notation, and write $Z_{X,L} = \sum_d Z_{X,L,d} \in \text{Sk}(L)$.

7.1. Spin structures – remarks and extensions.

- (1) Recall that the choice of spin structure on L is used to orient the moduli space of maps, so ultimately to determine signs. If we change the spin structure on L , we change the signs by a factor determined by the degree $d \in H_2(X, L)$. In our exact

setting, this factors through $H_1(L)$ and hence through the skein. That is, changing the spin structure is intertwined with a certain automorphism of the skein, multiplying each curve by a sign depending on its homology class.

- (2) Rather than choose a spin structure on L , we may fix a link $\eta \subset L$ and choose a spin structure on $L \setminus \eta$ which does not extend across η . Generically, the boundaries of holomorphic curves will not meet η , so this causes no problem for defining the count. As far as invariance is concerned, we pick up a sign when the boundary of a curve crosses η . So, the skein-valued curve count is well defined in the skein of L with a sign line, see (1.4), along η .

7.2. 4-chains – remarks and extensions.

- (1) The theory depends on the vector field ξ on L only through the corresponding section of $\underline{\xi}$ of the sphere bundle SL . The condition that ξ has transverse zeroes translates to asking $\underline{\xi}$ is defined away from finitely many points p_i , and near each such point, $\underline{\xi}$ defines a degree ± 1 map from boundary of some ϵ ball around p_i to the tangent sphere. (Note that the extra boundary fiber disks of the 4-chains of ξ and of $-\xi$ at zeros of ξ cancel in the 4-chain of $\pm\xi$ with boundary $2L$.) We could work throughout specifying only $\underline{\xi}$ rather than ξ .
- (2) Rather than choose a class b with $\gamma = (\Xi \setminus L) \cap L$, we may work with a skein with a framing line along γ . That is, we take framed links in $L \setminus \gamma$, and impose the usual skein relations, along with the relation:

$$(7.3) \quad \begin{array}{c} \uparrow \\ \leftarrow \text{---} \gamma \end{array} = a \cdot \begin{array}{c} \uparrow \\ \leftarrow \text{---} \gamma \end{array} .$$

The sign ± 1 in the exponent of a in (7.3) reflects the change in local linking number with γ .

- (3) If $L = \bigsqcup L_i$ and compatibly $\Xi = \sum \Xi_i$ (i.e. $\partial \Xi_i = 2L_i$), then we may use a different framing variables a_i on each L_i .
- (4) Similarly, we may take several 4-chains ξ_i, Ξ_i for one given Lagrangian L . Then we may consider the 4-chain $\Xi_L := \frac{1}{n} \sum \Xi_i$. Of course we should count 4-chain intersections with each Ξ_i by $a_L^{1/n}$.
- (5) Given a local system of sets of n vector fields on L , the collection of which however may undergo monodromy around L , we may consider a 4-chain Ξ for L which is locally near L of the form $\Xi_L := \frac{1}{n} \sum \Xi_i$. Counting holomorphic curves in the skein using such a 4-chain gives a deformation invariant result since the relationship between the 4-chain/vector field pair and the skein relations in [19] are all local over L .
- (6) We may specialize at $a = 1$, or, if there are many components as above, specialize some of the a variables to 1. Evidently, for a component L_i for which we set $a_i = 1$, the invariant is independent of the choice of ξ_i, Ξ_i ; in fact, an inspection of [19] shows that we do not need to make such choices at all to define the $a_i = 1$ specialization of the invariant. In the $a = 1$ skein of S^3 , the class of every nontrivial knot or link is zero, but in the solid torus that is not the case, see Section 7.3.

7.3. Basic classes. Suppose X is an exact symplectic manifold, and $L \subset X$ is a Maslov zero Lagrangian with the topology of the solid torus; assume L is equipped with vector field and 4-chain as appropriate to define skein-valued curve counts per [19]. Fix a generic longitude

of L . The longitude acquires a framing from the vector field, and an orientation by asking that the fixed symplectic primitive be positive on the longitude.

A choice of framed oriented longitude determines expansions as in (2.1) for elements in the skein of L . We will write $\psi(L)$ for the coefficient of $\psi_{\square, \emptyset}$ in the skein-valued curve count for (X, L) . More generally, if M is another Maslov zero Lagrangian, again equipped with appropriate data and disjoint from L , then we write $\psi(L) \in \text{Sk}(M)$ for the element obtained by taking the original curve count in $\text{Sk}(M \sqcup L) = \text{Sk}(M) \otimes \text{Sk}(L)$ and extracting the coefficient of $W_{\square, \emptyset}$ in the second factor.

Similarly, when L is a disjoint union of n solid tori, we similarly write $\psi(L)$ for the coefficient of $W_{\square, \emptyset}^{\otimes n}$ in the skein-valued curve count; again $\psi(L)$ is valued in the skein of the other Lagrangians.

For Ψ which are known a priori to admit an expansion in $W_{\lambda, \emptyset}$, the foregoing discussion and definition of $\psi(L) \in \text{Sk}(M)$ remain valid in the $a = 1$ specialization.

7.4. Normalization of the basic cylinder. Let M be a 3-manifold. Let $K \subset M$ be a link. Then $\psi(\mathbf{U}_K)$ is the count of one embedded cylinder for each component of K ; their geometric boundary in M is K . Here we explain how to normalize the 4-chain on M so that, if we put the $a = 1$ skein on \mathbf{U}_K , then

$$\psi(\mathbf{U}_K) = [K] \in \text{Sk}(M).$$

Fix a vector field ξ on M and take the 4-chain for the zero section $M \subset T^*M$ to be $\Xi := \mathbb{R}_{\geq 0} \cdot \pm J\xi$.

Fix a metric on M such that ξ is perpendicular to K only at isolated points and that $\langle \xi, \tau \rangle$ has transverse zeros at these points. Then:

Lemma 7.1. *Let $K_{\xi > 0} := \langle \xi, TK \rangle^{-1}(\mathbb{R}_{> 0})$ and similarly $K_{\xi < 0} := \langle \xi, TK \rangle^{-1}(\mathbb{R}_{< 0})$. Then, for small push-offs \mathbf{U}_K , the locus $\Xi \cap \mathbf{U}_K$ is obtained by smoothing the loci*

$$K_{\xi > 0} \cup \mathbb{R}_{\leq 0} \xi|_{\partial K_{\xi > 0}} \quad \text{and} \quad K_{\xi < 0} \cup \mathbb{R}_{\geq 0} \xi|_{\partial K_{\xi < 0}}.$$

In particular, $\Xi \cap \mathbf{U}_K = 0 \in H_1^{BM}(\mathbf{U}_K, \mathbb{Z})$. \square

Note that the transverse disk generating $H_2^{BM}(\mathbf{U}_K, \mathbb{Z})$ pairs as 1 with the longitude ℓ of the solid torus; thus, we may choose a Borel-Moore chain b with $db = \Xi \cap \mathbf{U}_K$ so that $\ell \cap b \in \mathbb{Z}$ takes any value. We choose b so that $\ell \cap b = 0$.

Remark 7.2. A version of the above discussion, identifying $\psi(\mathbf{U}_K) = K$ was the first step of the proof of the Ooguri-Vafa conjecture in [19]. In that article, we were free to use the full skein (not specialized at $a = 1$) on \mathbf{U}_K because we were working with $M = S^3$, so K necessarily bounds and \mathbf{U}_K admits a 4-chain. Indeed, since $T^*M \cong M \times \mathbb{R}^3$, pulling back gives an isomorphism $H_{\star}^{BM}(M) \cong H_{\star+3}^{BM}$ carrying $K \mapsto \mathbf{U}_K$. That is, \mathbf{U}_K admits a 4-chain if and only if $[K]$ is in the image of $H_1(\partial M) \rightarrow H_1(M)$. This is why in the present more general setting, we will use the $a = 1$ skein on \mathbf{U}_K to avoid restricting attention to links which are null-homologous rel. ∂M .

8. SKEIN TRACE FROM CURVE COUNTING

In this section, we will prove Theorem 1.3 on the well-definedness of the skein trace. First we must clarify the statement, or more precisely Definition 1.1, by giving the choice of 4-chain on $L \subset T^*M$ used to define the skein-valued curve counting. Then we check that

the prescription of Definition 1.1 for the map $K \mapsto [K]_L \in \text{Sk}(L)$ satisfies the three skein relations in M , hence factors through $\text{Sk}(M)$.

8.1. 4-chains for branched covers. Let M be a 3-manifold; fix a vector field with transverse zeros ξ_M on M and 4-chain $\Xi_M := \overline{\mathbb{R}_{\geq 0}} \cdot \pm J\xi$. We write $\Xi = \Xi[J]$ when we want to emphasize the dependence of the almost complex structure.

Consider a Lagrangian $L \subset T^*M$ so that $\pi: L \rightarrow M$ is a degree n cover. We will say that a 4-chain Ξ_L for L is *compatible* with Ξ_M if, in the complement of a fiberwise compact subset of T^*M , we have $\Xi_L = n\Xi_M$.

If $B \subset M$ is an open ball where π is not branched then, over B , L is given by the graphs of the differentials df_1, \dots, df_n of n functions. We parameterize the graph of df_j by a map $\gamma_j: B \rightarrow T^*M$: for $q \in B$, $\gamma_j(q) = (q, df_j(q)) \in T^*M$ and extend γ_j to a symplectomorphism $\Gamma_j: T^*B \rightarrow T^*B$ fiberwise, $\Gamma_j(q, p) = (q, df_j(q) + p)$.

Definition 8.1. Choose local coordinates $M \supset B \rightarrow \mathbb{R}^3$, on a region where ξ has no zeroes and $L \rightarrow M$ is a (not branched) cover. We say that L is *standard over B* if, inside some fiberwise compact region of T^*B containing $L \cap T^*B$, one has, in local coordinates:

- (1) The 4-chain for L is $\Xi_L = \sum \Xi_{L_i}$ where Ξ_{L_i} corresponds to the push-forward vector field on the graph of df_i :

$$(8.1) \quad \begin{aligned} \xi_i(\gamma_i(q)) &:= (d\gamma_i(q)) \xi(q), \\ \Xi_{L_i} &= \Gamma_i(\Xi[J_\Gamma]) = \bigcup_{q \in B} (\gamma_i(q) \pm J(\gamma_i(q)) \cdot \mathbb{R}_{>0} \xi_i(\gamma_i(q))), \end{aligned}$$

where $J_\Gamma = d\Gamma^{-1} \circ J \circ d\Gamma$.

- (2) Ξ_{L_i} is disjoint from the graph of df_j for all $j \neq i$.

We will show that Theorem 1.3 holds for any 4-chain which is standard over some ball B . In the remainder of this subsection, we establish the existence of such chains.

Let ξ be a vector fields on M and let $B \subset M$ be a ball with compact closure where ξ has no zeros. Let $\Xi = \sum \Xi_{L_i}$, where Ξ_{L_i} is as in (8.1). Let $S_\lambda: T^*M \rightarrow T^*M$ denote the fiber scaling by λ .

Lemma 8.2. *If $df_i - df_j$, $1 \leq i < j \leq n$, is not a multiple of the co-vector ξ^* dual to ξ then condition (2) of Definition 8.1 holds for $S_\lambda(L)$ for all sufficiently small $\lambda > 0$.*

Proof. Since $d\Gamma = 1 + \mathcal{O}(\lambda)$, $J_\Gamma \xi_i = J\xi + \mathcal{O}(\lambda)$. Then, if, for some $s \in \mathbb{R}$,

$$\lambda df_j(q) = \lambda df_i(q) + sJ\xi_i(q),$$

we have

$$df_j(q) - df_i(q) = \lambda^{-1}s(\xi^* + \mathcal{O}(\lambda))$$

and in the limit $\lambda \rightarrow 0$ we find that $df_j - df_i$ is a multiple of ξ^* . The lemma follows. \square

We next consider how condition (2) fails generically. Consider a vector field ξ in general position with respect to $df_i - df_j$. Then the locus where $df_i - df_j$ is parallel to ξ is a transversely cut out 1-submanifold $K_{ij} \subset B$.

Lemma 8.3. *For all sufficiently small $\lambda > 0$, the intersection $\Xi_{L_i} \cap L_j$ is a 1-submanifold isotopic to $\Gamma_j(K_{ij})$; in fact the intersection C^1 -converges to $\Gamma_j(K_{ij})$ as $\lambda \rightarrow 0$.*

Proof. As above, the intersection $\Xi_{L_i} \cap L_j$ corresponds to the solutions $q \in B$ to the equation

$$df_j(q) - df_i(q) = \lambda^{-1} s (\xi^* + \mathcal{O}(\lambda))$$

and K_{ij} corresponds to solutions to the equation

$$df_j(q) - df_i(q) = \sigma \xi^*.$$

The lemma follows by transversality of the solutions to the last equation. \square

We next consider compatibility of the 4-chain Ξ .

Lemma 8.4. *Let $\Xi = \sum_i \Xi_{L_i}$ be a 4-chain given by the formula (8.1) in $D_1^*M|_B$. Then, for all sufficiently small $\lambda > 0$, Ξ can be continued to a 4-chain in $T^*M|_B$ that agrees with $n\Xi_M$ outside $D_2^*M|_B$.*

Proof. Since $\Xi \rightarrow n\Xi_M$ as $\lambda \rightarrow 0$, there is a homotopy of chains that take $\Xi \cap \partial D_1^*M|_B$ to $n\Xi_M \cap \partial D_1^*M|_B$. We use this homotopy to interpolate in $D_2^*M|_B \setminus D_1^*M|_B$. The lemma follows. \square

In fact, 4-chains which are standard in the complement of some neighborhood of the branch locus are generic for all sufficiently small scalings of a given branched cover.

Consider a degree n branched cover $\pi: L \rightarrow M$ where $L \subset T^*M$ is a Lagrangian. Let $\tau \subset M$ be the link which is the branch locus of π . Then L defines a system of n closed differential forms $(\alpha_1, \dots, \alpha_n)$ over $M \setminus \tau$. Let $K \subset M$ denote the subset where the dual of ξ^* is proportional to some difference $\alpha_i - \alpha_j$. For ξ in general position K is a transversely cut out 1-submanifold. Let $\tilde{K} \subset K$ denote the natural lift of K that lifts each arc of K where ξ^* is parallel to $\alpha_i - \alpha_j$ to the corresponding sheet of L . We assume that ξ is generic and that the closure of K is disjoint from τ . Then $K \subset M$ is a closed submanifold.

Corollary 8.5. *Let $N(\tau)$ and $N(K)$ be a tubular neighborhoods of τ and K respectively. Then for all sufficiently small $\lambda > 0$ there is a 4-chain Ξ_L on $S_\lambda(L)$ that is compatible, and standard over any ball $B \subset M \setminus (N(\tau) \cup N(K))$. Furthermore, the intersection $\text{int}(\Xi_L) \cap L$ is isotopic to $\tilde{K} \subset L$ and C^1 -converges to \tilde{K} as $\lambda \rightarrow 0$.*

Proof. For a fixed ball in $B \subset M \setminus (N(\tau) \cup N(K))$ this follows from Lemmas 8.2, 8.3, and 8.4. To see that it holds for any ball first observe that we can get compatibility near any vertex. The closure of the complement of the vertices in $M \setminus (N(\tau) \cup N(K))$ is compact and can then be covered by a finite number of balls. The argument for a fixed ball then implies the lemma. \square

Remark 8.6. The analogous proposition holds with the identical proof if in place of single vector fields ξ , one has local systems of vector fields as in (5) in Section 7.2.

8.2. Proof of Theorem 1.3.

8.2.1. *First skein relation.* Consider a Lagrangian contained in a unit cotangent ball, $L \subset B^*M \subset T^*M$. In particular, under the fiber scaling, $\lambda L \rightarrow M$ as $\lambda \rightarrow 0$. We have the following general fact:

Lemma 8.7. *Consider $L \subset T^*M$ as above, and a Lagrangian $U \subset T^*M \setminus B^*M$.*

*Then in the flow tree limit, holomorphic curves for $(T^*M, U \sqcup \lambda L)$ converge to holomorphic curves for $(T^*M, U \sqcup M)$ with attached flow graphs for λL in T^*M .*

Proof. This is a special case of Lemma A.5. \square

We also have:

Lemma 8.8. *Consider $L \subset T^*M$. Let K be an immersed link in M with conormal U_\times with clean intersection and let $U_\pm(\rho)$ and $U_\succ(\rho)$ be the nearby conormals of the links K_\pm and K_\succ as above. Let p be the intersection point of K . Then for generic L (open dense set of L) there is a neighborhood of V around p in M such that no rigid flow graph in the basic class on U_\times has any flow line of L inside V . It follows that the same holds for generic resolutions $U_\pm(\rho)$ and $U_\succ(\rho)$ for all sufficiently small ρ .*

Proof. The genericity statement is a consequence of general position: the formal dimension of a rigid flow graph of L on K_\times with the additional constraint that some flow line in the graph is attached at a particular point along K_\times equals -1 . The statement about U_\times follows.

For $U_\pm(\rho)$ and $U_\succ(\rho)$ rigid flow graphs on $K_\pm(\rho)$ and $K_\succ(\rho)$ converge to rigid flow graphs on K_\times as $\rho \rightarrow 0$. Thus, if K_\times is generic there is a neighborhood V of p such that for large enough ρ , rigid flow graphs of $K_\pm(\rho)$ ($K_\succ(\rho)$) do not have any flow graphs attached to $K_\pm(\rho)$ ($K_\succ(\rho)$) in V that leave V .

It remains to show that they also do not have flow lines entirely contained in V . Such a flow line would converge to the constant flow line at the intersection point, the limit of its tangent vector is the limit of a gradient difference of sheets of L at p . Thus, as long as all gradient differences are not parallel to the shifts of the two resolution directions of $K_\pm(\rho)$ and $K_\succ(\rho)$ there are no such flow graphs. The lemma follows. \square

We now prove:

Proposition 8.9. *Let $K_\times \subset M$ be a link with an immersed point, nowhere parallel to (and hence framed by) the fixed vector field v_M . Let $K_+, K_-, K_\succ \subset M$ be the three corresponding links for the first skein relation (so $K_+ - K_- = zK_\succ$ in $\text{Sk}(M)$). Then*

$$[K_+]_L - [K_-]_L = z[K_\succ]_L \in \text{Sk}(L)$$

Proof. We should compare holomorphic curves between L and (any choice of, see Section 7.3) perturbed conormals to K_+, K_-, K_\succ . Isotope K_\times so that its singular point is in the ball in M where the 4-chain of L satisfies the ‘standardness-condition’, and moreover is in the open set U of Lemma 8.8. We will choose ρ sufficiently small that the B_ρ neighborhood around the singular point remains in U .

We use the corresponding $\tilde{U}_\succ(\rho)$, $\tilde{U}(\pm\sigma; \rho)$ for sufficiently small $\sigma, \rho > 0$, and negative perturbation in the sense of Definition 6.2 and choose them so that Lemma 8.8 holds.

Then, as described in Lemma 6.10 and Corollary 6.11, there is one holomorphic curve each in the basic class for $(T^*M, M \sqcup \tilde{U}(-\sigma, \rho))$ and $(T^*M, M \sqcup \tilde{U}_\succ(\rho))$. Denote these curves as C_- and C_\succ , respectively. There are two holomorphic curves for $(T^*M, M \sqcup \tilde{U}(\sigma, \rho))$, which we will denote C'_-, C'_\succ . These curves have the following properties:

- C'_- is C^1 close to C_- near M
- C'_\succ is C^1 close to C_\succ at least near $M \setminus B_\rho$
- there is an isotopy $\partial C'_\succ \sim \partial C_\succ$ which is C^1 small in $M \setminus B_\rho$ and standard in B_ρ , see Lemma 6.10.
- The domains of C_-, C'_-, C_\succ are (possibly unions of) cylinders, while C' is a three-holed sphere, $\chi(C'_\succ) = -1$.

By Lemma 8.7, the holomorphic curves ending on $L \sqcup \tilde{U}_\succ(\rho)$ limit to flow trees for $M \cup L$ attached to C_\succ . The limit is not necessarily transverse, but since no tree enters the region

U where C_{\succ} and C'_{\succ} differs and since perturbations can be taken supported in a small neighborhood of these flow trees it follows, using the same perturbation for the gluing in the two cases that the solutions for C_{\succ} are in canonical bijection with the solutions to the corresponding problem for C'_{\succ} . Let us write respectively $\text{Count}(C_{\succ})$ and $\text{Count}(C'_{\succ})$ for the skein count of the resulting glued curves.

The same argument applied to C_{\prec} and C'_{\prec} again gives canonical identifications of moduli spaces and we use analogous notation. We conclude:

$$\begin{aligned} \psi(\tilde{U}(\sigma, \rho)) &= \text{Count}(C'_{\prec}) + \text{Count}(C'_{\succ}) = \text{Count}(C_{\prec}) + z \cdot \text{Count}(C_{\succ}) \\ &= \psi(\tilde{U}(-\sigma, \rho)) + z\psi(\tilde{U}_{\prec}(\rho)). \end{aligned}$$

Here the z power appears because, as noted above, $\chi(C'_{\prec}) = \chi(C_{\prec}) - 1$, and we count curves by $z^{-\chi}$. No a powers appear because the geometries match away from B_{ρ} , and the 4-chain in B_{ρ} is standard. This completes the proof. \square

8.2.2. *Second skein relation.* We first study a model geometry.

Definition 8.10. Consider $M = \mathbb{R}^3$. Consider a trivial cover $L = L_1 \sqcup \dots \sqcup L_n \subset T^*\mathbb{R}^3$ where each L_i is a constant Lagrangian of the form $\mathbb{R}^3 \times (\ell_i^1, \ell_i^2, \ell_i^3) \subset \mathbb{R}^3 \times (\mathbb{R}^3)^* = T^*\mathbb{R}^3$. We assume all the ℓ_i^3 are distinct and order the L_i such that $\ell_1^3 > \ell_2^3 > \dots > \ell_n^3$.

We fix the vector field $v_M = \partial/\partial x_3$, and the corresponding linear 4-chain $W_M = \pm\mathbb{R}_{>0}J \cdot \partial/\partial x_3$. (Here J is determined by the choice of a metric on \mathbb{R}^3 ; we take the standard Euclidean metric.)

We choose 4-chains for the L_i by taking the linear 4-chain whose corresponding vector field v_{L_i} is $(\pi_i : L_i \rightarrow M)^*v_i = v_i + (\ell_i^1, \ell_i^2, \ell_i^3)$. We write a_i for the ‘a’ variable associated to the component L_i .

Lemma 8.11. *Let U be a standard unknot given by $(\cos \theta, \sin \theta, 0) \subset \mathbb{R}^3$ in the model geometry of Def. 8.10. Then*

$$[U]_L = \frac{a_1 \dots a_n - a_1^{-1} \dots a_n^{-1}}{z} \in \text{Sk}(L)$$

Proof. Let L_U be the shifted conormal of the unknot. It is straightforward to see that there is a unique holomorphic cylinder between L_U and each L_i , meeting L_i in the lift of U to L_i . By construction, U is the standardly framed unknot with respect to v_{L_i} (i.e. projects to an embedded unknot in the leaf space of the vector field), hence the boundary contribution of this holomorphic curve is $(a_i - a_i^{-1})/z$.

It remains to compute the 4-chain contributions. Because the Lagrangians are constant, we may compute the intersection after projecting to the cotangent fiber coordinates. The image of the holomorphic cylinder is $(\ell_i^1, \ell_i^2, \ell_i^3) + t(\cos(\theta), \sin(\theta), 0)$ with $t > 0$. The image of the 4-chain for L_j is the line $(\ell_j^1, \ell_j^2, \ell_j^3 + s)$. There is a unique $(t \cos(\theta), t \sin(\theta))$ so that the first two coordinates match. Considering the last coordinate, we see there is a unique intersection, which is in the positive part of the 4-chain if $s = \ell_i^3 - \ell_j^3 > 0$, and the negative part of the 4-chain otherwise. That is, the contribution of the 4-chain of L_j to the cylinder ending on L_i is a_j if $j > i$ and a_j^{-1} if $j < i$.

All in all, we find

$$[U]_L = \sum_{j=1}^n a_1^{-1} a_2^{-1} \dots a_{j-1}^{-1} \left(\frac{a_j - a_j^{-1}}{z} \right) a_{j+1} \dots a_n = \frac{a_1 \dots a_n - a_1^{-1} \dots a_n^{-1}}{z}$$

as desired. \square

Corollary 8.12. *Fix a Lagrangian $L \subset T^*M$ such that $L \rightarrow M$ is a degree n cover, and L is equipped with a somewhere standard 4-chain. Then we have, for a standardly framed unknot U in M ,*

$$[U]_L = \left(\frac{a_L^n - a_L^{-n}}{z} \right) [\emptyset]_L \in \text{Sk}(L)$$

Proof. We may put our unknot U anywhere on M , and take it arbitrarily small.

We fix attention at an open ball of the standard neighborhood through which no flow trees pass (which exists per Lemma 8.8), and there use v_M to choose coordinates so that $v_M = \partial/\partial x_3$. Taking an arbitrarily small neighborhood of a point, we may arbitrarily closely approximate the model geometry of Definition 8.10.

When we compute $[U]_L$, each contribution is a disjoint union of holomorphic curves corresponding to L flow graphs (which are by hypothesis disjoint from the neighborhood) and the model cylinders. The contribution of the external flow graphs are identical for each model cylinder, and total, by definition, $[\emptyset]_L$.

The result then follows from Lemma 8.11. \square

8.2.3. Third skein relation.

Proposition 8.13. *Fix a Lagrangian $L \subset T^*M$ such that $L \rightarrow M$ is a degree n cover, and L is equipped with a somewhere standard 4-chain. Let $\beta \subset M$ be a knot framed by v_M , and β^+ an isotopic knot whose v_M framing is one greater, so in particular $[\beta^+] = a_M[\beta] \in \text{Sk}(M)$. Then*

$$[\beta^+]_L = a_L^n [\beta]_L.$$

Proof. We isotope β, β^+ so that they differ only in the standard neighborhood and in a region Ω in the complement of all flow graphs. We take this region small enough to be arbitrarily close to the model geometry 8.10 and choose coordinates so that β and β^+ are as follows. Let $\psi: \mathbb{R} \rightarrow [0, 1]$ be a smooth function that equals 0 on $[-2, 2]$ and equals 1 on $\mathbb{R} \setminus [-3, 3]$ and let $\epsilon, \delta > 0$ be small, then

- β is parameterized by

$$t \mapsto (1, t, 0),$$

- and β^+ is parametrized by

$$t \mapsto ((1 - \psi(t))t^2 + \psi(t), t((1 - \psi(t))(t^2 - \delta^2) + \psi(t)), (1 - \psi(t))\epsilon t)$$

We take graphical perturbed conormals for β, β^+ . As in the proof of Proposition 8.9, the total curve count is given by gluing the basic cylinders for β (resp. β^+) to flow graphs ending on β (resp. β^+). By general position, compare Lemma 8.8, all flow graphs either attach outside our neighborhood, or lie entirely contained in our neighborhood. The former kind are exactly the same for β, β^+ , and, while they may not necessarily be transverse, their gluing problems are similarly identical so they will contribute identically. Of the latter kind, there is only one: the interval $\{(\delta^2, 0, \epsilon s) \mid -\delta \leq s \leq \delta\}$.

Let us fix some choice of external contributions (which again, will be the same in both cases), and analyze how the internal geometry contributes. We now pass to the model geometry. Let us write β_1, \dots, β_n for the lifts of β to L_i ; similarly β_i^+ for the corresponding lifts of β^+ . In addition to these trivial lifts, we have the contributions corresponding to β^+

with the attached flow line; there is one such contribution for each $i > j$. The corresponding holomorphic curve has boundary approximately the union of β_i and a standard-framed unknot U_j on L_j ; the curve itself has Euler characteristic one less (corresponding to the flow line) than the trivial contributions.

The a power computation is (for small enough ϵ) similar to that of Lemma 8.11. (See the end of Section 9.4 for a detailed discussion of flow line gluing and framings). We compute:

$$\begin{aligned}
[\beta^+|_\Omega]_L &= \beta_1^+ a_2 \dots a_n \\
&+ (a_1^{-1} \beta_2^+ a_3 \dots a_n + z \beta_2 U_1 a_2 \dots a_n) \\
&+ (a_1^{-1} a_2^{-1} \beta_3^+ a_4 \dots a_n + z \beta_3 (a_1^{-1} U_2 a_3 \dots a_n + U_1 a_2 \dots a_n)) \\
&+ \dots \\
&= (a_1 \dots a_n) (\beta_1 + \dots + \beta_n) \\
&= (a_1 \dots a_n) [\beta|_\Omega]_L
\end{aligned}$$

The above equality would make sense in a skein-with-boundary-marked-points for $(L|_\Omega, \partial L_\Omega)$, but what we really mean by it here is that it holds after attaching the remainder of β and any fixed external flow graph contribution. The result follows. \square

8.3. A Symplectic Field Theory perspective. Theorem 1.3 and the results it is built on can be used to express related skein-valued curve counts in many situations. We discuss some of them in this section but start out with the symplectic field theory perspective that underlies all such applications.

8.3.1. Sketch of another proof. Let us note another approach to the skein trace result; the idea is to stretch along a cosphere bundle $S^*M \subset T^*M$ which separates L from \mathbf{U}_K .

Recall the Abbondandolo-Schwarz isomorphism comparing the symplectic homology of a cotangent bundle with the chains on the loop space of the base [1]; recall it has also an S^1 -equivariant version. We will give elsewhere a higher genus analogue for three dimensional cotangent bundles, which compares a ‘bare curve’ version of the degree zero (all genus) SFT of S^*M with the skein module $\text{Sk}(M)$. The isomorphism will be constructed by counting curves in $(T^*M, 0_M)$ that are asymptotic to Reeb orbits at infinity, by their boundary in the skein of M . Accept for now (or regard as conjectural) this isomorphism; we denote it:

$$\Omega : \text{SFT}_{\text{bare}}^0(S^*M) \xrightarrow{\sim} \text{Sk}(M).$$

Given this isomorphism, Theorem 1.3 can be deduced by SFT stretching from the well-definedness of skein-valued curve counting. Indeed, consider the geometry $(T^*M, \mathbf{U}_K \sqcup L)$. By definition, we set $[K]_L := \psi_{T^*M, L}(\mathbf{U}_K) \in \text{Sk}(L)$. However, let us now consider the geometry $(S^*M \times \mathbb{R}, \mathbf{U}_K)$. Counting holomorphic curves here which may end at \mathbf{U}_K and at the negative end of $S^*M \times \mathbb{R}$, we may define correspondingly $\psi_{S^*M \times \mathbb{R}}(\mathbf{U}_K) \in \text{SFT}_{\text{bare}}^0(S^*M)$. Finally, let us define a similar map (now not an isomorphism)

$$\Omega_L : \text{SFT}_{\text{bare}}^0(S^*M) \rightarrow \text{Sk}(L)$$

by counting holomorphic curves in T^*M which are asymptotic to Reeb orbits and end on L . Now, by SFT stretching,

$$\begin{aligned}
\psi_{T^*M, M} &= \Omega \circ \psi_{S^*M} \\
\psi_{T^*M, L} &= \Omega_L \circ \psi_{S^*M}
\end{aligned}$$

Now,

$$\Omega \circ \psi_{S^*M}(\mathbf{U}_K) = \psi_{T^*M,M}(\mathbf{U}_K) = [K] \in \text{Sk}(M)$$

and, by definition,

$$\Omega_L \circ \psi_{S^*M}(\mathbf{U}_K) = \psi_{T^*M,L}(\mathbf{U}_K) = [K]_L \in \text{Sk}(L).$$

Thus we find that the skein trace is given by the formula

$$[K]_L = \Omega_L \Omega^{-1}[K]$$

and, in particular, $[K]_L$ depends only on the class of K in the skein of M , as desired.

8.3.2. The skein trace map and Lagrangian cabling. Now, let X be any Calabi-Yau 3-fold, and $M \subset X$ a (Maslov zero) Lagrangian submanifold. Choose $L \subset T^*M$ a Lagrangian which branch covers M ; after fiber rescaling, of the cotangent bundle, we may implant L into X . Fix, if desired, also some other Lagrangian $\Lambda \subset X$ away from a neighborhood of M containing L . Fix brane data for all involved objects, and assume that the 4-chain for L is compatible with that of M , and somewhere standard.

Corollary 8.14. *The skein-valued curve counts $Z(X, \Lambda \sqcup M) \in \text{Sk}(\Lambda) \otimes \text{Sk}(M)$ and $Z(X, \Lambda \sqcup L) \in \text{Sk}(\Lambda) \otimes \text{Sk}(L)$ are related by the skein trace:*

$$Z(X, \Lambda \sqcup L) = [(Z(X, \Lambda \sqcup M))]_L.$$

Proof. Any bare holomorphic curve with boundary on M is arbitrarily C^1 -close to the basic cylinder on its boundary in a small neighborhood of the zero section. The result then follows from Lemma A.5. \square

Remark 8.15. There is also an argument along the lines of the sketch above: we may count curves in $X \setminus M$ ending on Λ and asymptotic to Reeb orbits in the negative S^*M end; this will take value in $Z(X \setminus M; \Lambda) \in \text{Sk}(\Lambda) \otimes \text{SFT}_{bare}^0(S^*M)$. Now, by SFT stretching,

$$\begin{aligned} Z(X, \Lambda \sqcup M) &= \Omega(Z(X \setminus M; \Lambda)) \\ Z(X, \Lambda \sqcup L) &= \Omega_L(Z(X \setminus M; \Lambda)) \end{aligned}$$

so the result follows by inverting Ω .

Example 8.16. For a framed knot $K \subset M$ and any skein σ in the standard solid torus, we write $\sigma(K) \in \text{Sk}(M)$ for the cabling of K by σ , i.e. we implant the solid torus in a neighborhood of K and take the image of σ . It is well known that $W_{\mu,\emptyset}(K)$ is the skein counterpart of the coloring of a knot K by the representation corresponding to the partition μ .

We recall from [15] the formula:

$$Z(T^*M, \mathbf{U}_K \sqcup M) = \sum_{\mu} W_{\mu,\emptyset} \otimes W_{\mu,\emptyset}(K) \in \text{Sk}(\mathbf{U}_K) \otimes \text{Sk}(M),$$

Applying Corollary 8.14,

$$Z(T^*M, \mathbf{U}_K \sqcup L) = \sum_{\mu} W_{\mu,\emptyset} \otimes [W_{\mu,\emptyset}(K)]_L \in \text{Sk}(\mathbf{U}_K) \otimes \text{Sk}(L),$$

This is a geometric interpretation of the skein traces of all colorings $W_{\mu,\emptyset}(K)$: they are obtained by extracting the coefficients of the curve count $Z(T^*M, \mathbf{U}_K \sqcup L)$ in the $W_{\lambda,\mu}$ basis of $\text{Sk}(\mathbf{U}_K)$.

This argument leads to an alternative proof of Theorem 4.11 after we observe that the only holomorphic curves there are the two annuli with boundary in the two sheets and their multiple covers:

$$\sum_{\mu, \nu} W_{\mu, \emptyset} W_{\nu, \emptyset} \otimes W_{\mu, \emptyset}(K) \otimes W_{\nu, \emptyset}(K) = \sum_{\lambda} W_{\lambda, \emptyset} \otimes \left(\sum_{\mu, \nu} c_{\mu, \nu}^{\lambda} W_{\mu, \emptyset}(K) \otimes W_{\nu, \emptyset}(K) \right)$$

in $\text{Sk}(\mathbf{U}_K) \otimes \text{Sk}(\text{Ann}) \otimes \text{Sk}(\text{Ann})$, where K is the core of the thickened annulus.

9. BRANCHED DOUBLE COVERS, SKEIN COUNTS, AND FLOW GRAPHS

This section considers holomorphic curves on Lagrangians associated to a branched double cover. In Sections 9.1 and 9.2 we construct Lagrangians from branched double covers. In Section 9.3 we equip them with 4-chains. The Lagrangian give flow graphs which we determine in Section 9.6. The flow graphs correspond to holomorphic curves and we give combinatorial formulas for the intersections of such curves with the 4-chain in Section 9.4. Finally in Section 9.7, we show that the foliation induced by an ideal triangulation has a direct interpretation in terms of moduli spaces of holomorphic disks with two positive punctures on the Lagrangian cobordism corresponding to a double cover.

9.1. Lagrangian double covers away from the cone points. We recall from [66, 61] certain constructions of Lagrangian double covers in T^*M from fronts in M .

Let us first consider the case of a surface M and cubic graph $\Gamma \subset M$. Then there is a singular hypersurface $F_{\Gamma} \subset M \times \mathbb{R}$ such that (1) it is a trivial two sheet cover over $M \setminus \Gamma$, (2) it has transverse self-crossings over Γ , and (3) over the vertices of Γ , it has the D_4^- singularity of Arnol'd. It is the front of a smooth Legendrian $\Lambda_{\Gamma} \subset J^1M$, see Figure 29. The figure also indicates a choice of perturbation of the Lagrangian to one whose front has generic singularities (only double points, cusps, and swallowtails), which we will later use in flow tree calculations.

Assuming the connected components of $M \setminus \Gamma$ (the faces of the graph) are contractible, we may ensure Λ_{Γ} has exactly one Reeb chord over each such component, corresponding to a maximum of the difference between the local functions whose jets give the two branches of the Legendrian. So, the corresponding Lagrangian immersion similarly has one double point over each face. This construction was described and explored in [66]. An analogous discussion goes through for any M and any singular hypersurface $\Gamma \subset M$ which is locally a product (cubic graph in surface) times (smooth manifold).

A motivating instance for this construction is the case when Γ is the 1-skeleton \mathbb{T} of a tetrahedron, viewed as a graph on a sphere. In this case, the corresponding

$$\Lambda_{\mathbb{T}} \subset J^1(S^2 = \partial_{\infty}\mathbb{R}^3) \subset S^5 = \partial_{\infty}T^*\mathbb{R}^3$$

is isotopic to the ideal contact boundary of the Harvey-Lawson cone L_{HL} . In coordinates, $L_{\text{HL}} \approx \mathbb{T} \times \mathbb{R}$ is the cone on the torus

$$\Lambda_{\text{HL}; \theta} = \{(e^{i\theta_1}, e^{i\theta_2}, e^{i\theta_3}) \mid \theta_1 + \theta_2 + \theta_3 = \theta\} \subset S^5 \subset \mathbb{C}^3;$$

we sometimes write $\Lambda_{\text{HL}} := \Lambda_{\text{HL}; \frac{\pi}{2}}$.

Let us now turn to to a 3-manifold M equipped with an ideal triangulation Δ ; we write M° for a complement of the neighborhood of the barycenters of the tetrahedra of Δ . We will write Δ^{\vee} for the dual ‘polyhedral’ decomposition; note however that the maximal cells will be topologically of the form $S \times \mathbb{R}$ for some component S of the ideal boundary $\partial_{\infty}M$. Consider

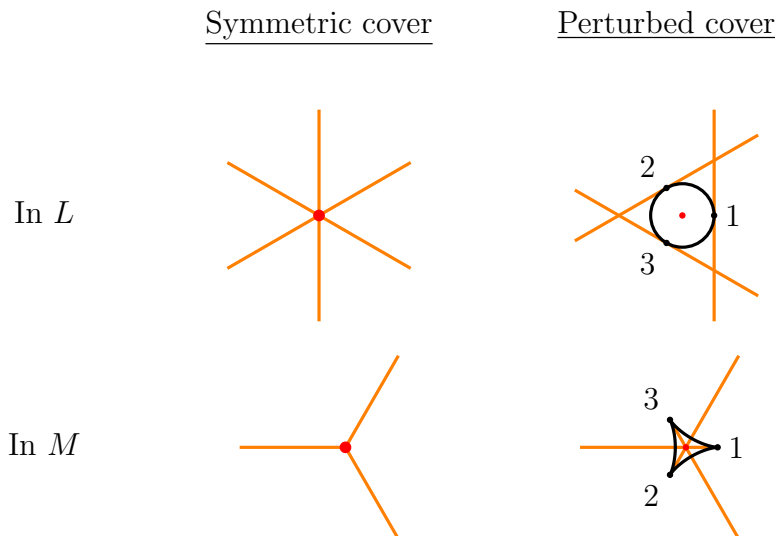


FIGURE 29. Left: the front of the Lagrangian near the vertex of a cubic vertex projected to the base. Right: the front determining our choice of perturbed Lagrangian. The projection of the perturbed Lagrangian to the zero section is modeled by the map $z \mapsto w = z^2 + \epsilon \bar{z}$. The circle $|z| = |\epsilon|/2$ is the critical set, and the deltoid $w = \frac{\epsilon^2}{4}(e^{i\theta} + 2e^{-i\theta})$ is the caustic.

the hypersurface $\Delta_{\leq 2}^\vee$ given by the 2-skeleton of Δ^\vee . Then $(M^\circ, \Delta_{\leq 2}^\vee)$ is a manifold equipped with a hypersurface with cubic planar singularities, and so we may form the corresponding Lagrangian $L_\Delta^\circ := L_{\Delta_{\leq 2}^\vee} \subset T^*(M^\circ)$. We arrange that in each cell of Δ^\vee , the difference of local defining functions has no maximum and grows as one approaches $\partial_\infty M$ and decreases as one approaches the vertices of $\Delta_{\leq 2}^\vee$.

We will also write $\Lambda_{\partial\Delta}$ for the Legendrian in $J^1\partial_\infty M$ corresponding to the cubic graph given by $\Delta^\vee|_{\partial_\infty M}$. We glue in Harvey-Lawson cones at the negative end and obtain a singular Lagrangian $L_\Delta \subset T^*M$.

We collect some properties of this construction:

Proposition 9.1. *Let M be a 3-manifold with ideal triangulation Δ . Let $(T^*M)^\circ$ be the complement in T^*M of the barycenters of the tetrahedra of Δ . We regard this as a Weinstein cobordism with positive end $\partial_\infty T^*M$ and negative end $\bigsqcup S^5 \times [0 - \infty)$.*

*Then L_Δ° is an embedded exact Lagrangian cobordism with positive end $\Lambda_{\partial\Delta} \subset J^1(\partial_\infty M) \subset \partial_\infty T^*M$ and negative end identified with $\bigsqcup \Lambda_{\mathbb{T}} \subset \bigsqcup S^5$.*

Furthermore, the positive end $\Lambda_{\partial\Delta}$ is a two-component Legendrian link with components obtained by shifting $\partial_\infty M$ in the positive and negative Reeb direction. There are thus natural $\partial_\infty M$ Bott families of Reeb chords connecting the two shifted components. At the negative end, each component $\Lambda_{\mathbb{T}} \approx \Lambda_{\text{HL}}$ has \mathbb{T} family of minimal action Reeb chords of action $\frac{2\pi}{3}$, see [18].

Proof. By inspection of the above construction. □

9.2. Smoothing the cones. The Harvey-Lawson cone L_{HL} has three distinguished smoothings to non-exact Lagrangian solid tori L_{HL}^* , known also as the ‘toric Aganagic-Vafa branes’ for \mathbb{C}^3 . They are distinguished from each other by the 1-cycle in Λ_{HL} which is collapsed; with respect to some choice of basis of $H_1(\Lambda_{\text{HL}}, \mathbb{Z})$, these cycles are $(0, 1)$, $(1, 0)$, $(-1, -1)$. In

the branched double cover setting, these three smoothings are topologically the smoothings we associated to a marking of the tetrahedron in Section 3.1. Furthermore, the pull-back of the Liouville form on \mathbb{C}^3 to the smoothing L_{HL}^* gives a 1-form that is positive on the distinguished cycle of the tetrahedron. i.e., it gives an effective marking in the sense of Definition 3.4.

Consider a 3-manifold M with a marked ideal triangulation Δ with corresponding smooth branch locus τ and corresponding distinguished cycle in each tetrahedron of Δ , see Section 3.4. Proposition 9.1 gives an exact Lagrangian cobordism $L_\Delta^\circ \subset (T^*M)^\circ$ with negative conical ends on the Legendrian tori $\Lambda_{\mathbb{T}} \subset S^5$ near the barycenter of each tetrahedron, identified with $\Lambda_{\text{HL}} \subset S^5$. We think of T^*M as obtained from $(T^*M)^\circ$ by gluing symplectic \mathbb{C}^3 's at the negative ends of $(T^*M)^\circ$.

We say that a Lagrangian $L_\tau \subset T^*M$ is a *smoothing of L_Δ° respecting τ* if it agrees with smoothings L_{HL}^* in balls $B_j(\delta)$ around the barycenter of each tetrahedron $\delta \in \Delta$ that are positive on the distinguished cycle over δ determined by the smoothing τ , and if $L_\tau \setminus \bigcup_j B_j$ is graphical over $L_\Delta^\circ \setminus \bigcup_j B_j$. We then have the following result connecting smoothings and effective 1-forms.

Proposition 9.2. *The following are equivalent:*

- (1) *There is a smoothing L_τ of L_Δ° respecting τ .*
- (2) *The marking of Δ that gives the smoothing τ is effective in the sense of Definition 3.4.*

Proof. If (1) holds then along $L_\Delta^\circ \setminus \bigcup_j B_j$, the desired 1-form in (2) is the 1-form of the graphical submanifold $L_\tau \subset T^*L_\Delta^\circ$. Since L_τ is a smoothing of L_Δ° , the 1-form is a pull-back of a closed 1-form on $\Lambda_{\mathbb{T}}$ which then extends and gives a 1-form positive on all distinguished cycles.

Conversely, given a 1-form ζ positive on distinguished cycles we find, for all sufficiently small $\epsilon > 0$, (graphical) exact Lagrangian isotopies near each negative end such that the 1-form $\epsilon\zeta$ agrees with the form of the smoothing L_{HL}^* of the Harvey Lawson cone. Let L_τ be the Lagrangian that agrees with the smoothing of the Harvey Lawson cones near the negative ends and that outside the negative end is the graph of $\epsilon\zeta$ over L_Δ° . \square

Note, the only condition on the shift form ζ in each torus $\Lambda_{\mathbb{T}}$ is that, it vanishes along the collapsing cycle and is positive along the distinguished cycle. Therefore, (2) is equivalent to the existence of a closed 1-form ζ used in Definition 3.5 to define a completion of the skein module.

9.3. 4-chains for branched double covers. We will use a 4-chain associated to a local system ξ_L of pairs of vector fields over the double cover L , as in (5) of Section 7.2. To construct such a chain, it suffices to exhibit the local system of vector fields satisfying the hypotheses of Proposition 8.5.

Recall that over a ball in M where $\pi: L \rightarrow M$ is not branched, there are $\frac{1}{2}n(n-1)$ vector fields on M given by $(\alpha_i - \alpha_j)^*$, $i < j$ where α_i is a local 1-form on M whose graph is the i^{th} local sheet of L . In the present case of interest, $n = 2$, so there is only one such local vector field. In fact, outside the branch locus difference 1-form is defined up to sign and hence gives a field of involutive line segments. Note that this field of line segments has a natural lift to a vector field v on $L \setminus \tilde{\tau}$: if for $p \in M$ if $\pi^{-1}(p) = \{p_1, p_2\} \subset L$ then $v(p_1) = d\Gamma_{\alpha_1}(\alpha_1(p_1) - \alpha_2(p_2))^*$ and $v(p_2) = d\Gamma_{\alpha_2}(\alpha_2(p_2) - \alpha_1(p_1))^*$. In order to extend v

over the branch locus, fix a tubular neighborhood of the branch locus $\tau \subset M$. Let r denote the radius of the tubular neighborhood and note that the difference vector fields decay as \sqrt{r} as we approach τ . In order to get a C^k -extension we scale the difference vector field by r^{k+1} in the fiber direction of the tubular neighborhood. We write ξ_{Mor} for the corresponding vector field on L ; extended by zero along the branch locus.

We write $\overset{\leftrightarrow}{\xi}_{\text{Mor}}$ for the field of involutive intervals on M with positive and negative halves that pullback to ξ_{Mor} on L . We will use a small perturbation of $\overset{\leftrightarrow}{\xi}_{\text{Mor}}$ to construct a 4-chain for L using Proposition 8.5. Consider a 3-manifold M with a taut triangulation Δ and corresponding branched double cover $L \subset T^*M$. Then the line field $\overset{\leftrightarrow}{\xi}_{\text{Mor}}$ is tangent to the foliation with leaf space given by the 2-cells of the polygonal decomposition dual to Δ , see Section 3.2, or equivalently, the vector field ξ_{Mor} is tangent to the foliation of L by boundaries of holomorphic strips in $\mathcal{M}(L_\Delta^\circ)$, see Section 9.7.

Lemma 9.3. *With M and Δ as above, there exists a line field $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ on M which has the following properties:*

- $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is invariant under the flow of the foliation of Δ and induces a line field on the leaf space,
- $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is perpendicular to the faces of all tetrahedra in Δ ,
- $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is everywhere non-vanishing (see Remark 9.4),
- $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is parallel to the branch locus τ ,
- away from τ and the edges of Δ , $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is nowhere parallel to $\overset{\leftrightarrow}{\xi}_{\text{Mor}}$,

Proof. We may define the line field in the complement of the edges using only the condition that the dihedral angles of each individual tetrahedron are $0, 0, \pi$. We do so as follows: Choose one of the two π -edges and call the two adjacent faces “incoming” and call the other two faces “outgoing”. Then, upon choosing an orientation, the line field $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ looks like a vector field coming in through the incoming faces and going out through the outgoing faces; see Figure 30. It is clear from the picture that we may choose the line field parallel to the branch locus. This line field glues along the faces to give a line field on M minus the edges. Finally, the line field extends over the edges as a nonzero line field if the sum of angles around edges is 2π , i.e., when the ideal triangulation is taut. \square

Remark 9.4. If the condition on the triangulation Δ that the angle sum around edges is relaxed from being 2π to being just a multiple of π , then the line field vanishes along edges with angle sum $\neq 2\pi$, and there is a weighted framing line (compare (7.3)) there.

Because $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ is parallel to the branch locus τ , it lifts to a line field on L which is tangent to $\tilde{\tau}$.

Consider an ideal tetrahedron $\delta \in \Delta$ and its preimage in $\tilde{\delta} \subset L$. Let $\overset{\rightarrow}{\xi}_{\text{pert}}$ and $\overset{\leftarrow}{\xi}_{\text{pert}}$ denote the two vector fields with opposite orientations on $\tilde{\delta}$ that lift the line field $\overset{\leftrightarrow}{\xi}_{\text{pert}}$ on δ . Let

$$\xi_{\text{pert}} := \frac{1}{2}(\overset{\rightarrow}{\xi}_{\text{pert}} \oplus \overset{\leftarrow}{\xi}_{\text{pert}})$$

be the *formal* average of these two vector fields. Then ξ_{pert} glues smoothly over the faces of tetrahedra in Δ to give a local system of two vector fields along L .




FIGURE 30. The line field $\overset{\leftrightarrow}{\xi}_{\text{pert}}$, and the branch locus of the double cover (drawn in red).

Define $C_L = C_{\xi_L}$ to be the compatible 4-chain of L constructed as in Proposition 8.5 from the vector field $\pm(\xi_{\text{Mor}} \pm \epsilon\xi_{\text{pert}})$, i.e., so that the pushforward vector fields of the 4-chain along L is the system of vector fields

$$\xi_L := \xi_{\text{Mor}} + \epsilon\xi_{\text{pert}}.$$

Here we require that $\epsilon > 0$ is sufficiently small in the following sense. Fix a tubular neighborhood of $\tilde{\tau}$ such that the tangent vector of the lift of any flow tree near the branch locus is linearly independent from ξ_{pert} , then take $\epsilon > 0$ sufficiently small compared to the separation of the two branches of the double cover outside the tubular neighborhood, so that here $|\epsilon\xi_{\text{pert}}| \ll |\xi_{\text{Mor}}|$.

Remark 9.5. If we are given a taut ideal triangulation in the sense of Lackenby, i.e., if we are also given a compatible choice of transverse orientation on each tetrahedron, then the local system of vector fields constructed above has a global section, namely the vector field $\vec{\xi}_{\text{pert}}$ along the direction of the transverse orientation. In this case, we can simply use the vector field $\vec{\xi}_{\text{pert}}$ to get a genuine 4-chain (instead of a local system of it).

9.4. Intersections of holomorphic curves with the 4-chain. In this section, we give a diagrammatic calculation of the a -powers in the skein lifting map. To this end, we consider a holomorphic curve u in T^*M with boundary on M . Assume that M has an ideal triangulation Δ and let L be the corresponding branched double cover. Then by Lemma A.5, u gives rise to holomorphic curves \tilde{u} with boundary on L which are obtained from u by attaching flow graphs to u . Here we will give a combinatorial formula for the intersection numbers between \tilde{u} and the 4-chain C_L constructed in Section 9.3 provided L is sufficiently close to M .

Note that there are three different local models for points in \tilde{u} :

- points where \tilde{u} is simply a lift of u ,
- points where \tilde{u} agrees with the holomorphic strip of a flow line,
- and junction points where a flow line is joined to u .

We consider first simple lifts. Consider an arc γ in the boundary ∂u that maps to a region in a tetrahedron of Δ in the complement of the 3d spectral network, see Section 3.2. We pick local coordinates $x = (x_1, x_2, x_3) \in \mathbb{R}^2 \times \mathbb{R}$ in this region of the tetrahedron in such a way that the foliation corresponds to the foliation of lines in the x_3 -direction (with constant flow speed around γ) and the leaf space corresponds to the x_1x_2 -plane. We write $\pi_{\mathbb{R}^2}$ for the projection to the x_1x_2 -plane and $\pi_{\mathbb{R}}$ for the projection to the x_3 -axis. The coordinates x give corresponding coordinates (x, y) on the cotangent bundle T^*M , where $y = (y_1, y_2, y_3)$ are momenta dual to x .

With this choice of coordinates we then have the two sheets of the λ -fiber scaled Lagrangian $\lambda \cdot L$ given by

$$\lambda \cdot L_{\pm} = \Gamma_{\pm\lambda\alpha}, \quad \alpha = \lambda \left(\pm \frac{1}{2} \delta dy_3 + \mathcal{O}(\lambda) \right),$$

for some $\delta > 0$, see Section 3.2. On $\lambda \cdot L_{\pm}$, the associated vector field, the dual of the difference covector pushed forward, is

$$\xi_{\text{Mor}} = \lambda (\pm \delta \partial_{x_3} + \mathcal{O}(\lambda)).$$

Consider next the perturbation vector field ξ_{pert} . Since it is invariant under the flow of the foliation and is small compared to ξ_{Mor} , we write

$$\xi_{\text{pert}} = \pm \lambda \epsilon v(x_1, x_2),$$

where $\epsilon \ll \delta$ and where $v(x_1, x_2)$ is a unit length vector field on \mathbb{R}^2 pulled back to \mathbb{R}^3 .

The framing vector field (i.e. the 4-chain vector field) on L_{\pm} is then given by:

$$\xi_{L_{\pm}} = \xi_{\text{Mor}} + \xi_{\text{pert}} = \lambda (\pm \delta \partial_{x_3} + \epsilon v + \mathcal{O}(\lambda)).$$

In fact, we have a local system of two such vector fields, the second component is obtained by changing the sign of v , see Section 9.3. Let $C_L = C_{L_+} + C_{L_-}$ be the compatible 4-chain with $\partial C_{L_{\pm}} = 2L_{\pm}$ constructed as in Section 8.5 using this data. Consider a holomorphic curve u with boundary arc $\gamma \subset \partial u$ on M and the corresponding lifts \tilde{u}_{\pm} with boundary in L_{\pm} .

Lemma 9.6. *For all $\lambda > 0$ sufficiently small, there is exactly one intersection point in $\tilde{u}_{\pm} \cap C_{L_{\mp}}$ in a neighborhood of any point on γ where the projection of γ to the leaf space \mathbb{R}^2 has tangent vector $\pi_{\mathbb{R}^2} \dot{\gamma}$ equal to a negative multiple of v , and no other intersections. Furthermore, the orientation sign of the intersection point in $\tilde{u}_{\pm} \cap C_{L_{\mp}}$ agrees with the sign of the tangency, i.e. the orientation sign of $\pi_{\mathbb{R}^2} \dot{\gamma} \wedge \pi_{\mathbb{R}^2} \dot{\gamma}$ in the orientation of \mathbb{R}^2 induced by $\xi_{\text{Mor}} = \pm \delta \partial_{x_3}$ and the orientation of L_{\pm} .*

Proof. Parameterize u by $\sigma + i\tau$ in a neighborhood of the origin in the upper half plane, so that its boundary curve γ has projected tangent vector field $\pi_{\mathbb{R}^2}(\dot{\gamma}(\sigma))$ of length 1.

Near $\lambda \cdot L_{\pm}$, the holomorphic curve is then as follows in $\mathbb{R}_x^3 \times \mathbb{R}_y^3$ -coordinates on T^*M

$$\tilde{u}_{\pm}(\sigma, \tau) = (\gamma(\sigma), (0, 0, \pm \lambda \delta / 2) + (\tau \dot{\gamma}(\sigma) + \mathcal{O}(\tau^2))) + \mathcal{O}(\lambda^2).$$

Similarly, the 4-chain $C_{L_{\pm}}$ is

$$\begin{aligned} C_{L_{\pm}} = & \{((x_1, x_2, x_3), \pm(0, 0, \lambda \delta / 2) + t\lambda((\epsilon v, \pm \delta) + \mathcal{O}(\lambda))); t \geq 0\} \\ & \cup \{((x_1, x_2, x_3), \pm(0, 0, \lambda \delta / 2) + t\lambda(-(\epsilon v, \pm \delta) + \mathcal{O}(\lambda))); t \geq 0\}, \end{aligned}$$

oriented so that

$$\partial C_{L_{\pm}} = 2L_{\pm}.$$

The intersections of \tilde{u}_{\pm} and $C_{L_{\mp}}$ then correspond to solutions of the following equation

$$(0, 0, \pm \lambda \delta / 2) + \tau \dot{\gamma}(\sigma) + \mathcal{O}(\tau^2) = (0, 0, \mp \lambda \delta / 2) + t\lambda(\beta(\epsilon v, \mp \delta) + \mathcal{O}(\lambda)), \quad t \geq 0,$$

where $\beta = \pm 1$. For any solution $\tau = \mathcal{O}(\lambda)$, and hence for all sufficiently small $\lambda > 0$ solutions correspond to solutions to

$$\begin{aligned} \tau \pi_{\mathbb{R}^2} \dot{\gamma}(\sigma) &= t\lambda \beta \epsilon v(\pi(\gamma(\sigma))), \\ \tau \pi_{\mathbb{R}} \dot{\gamma}(\sigma) &= -\lambda \beta \delta t - \lambda \delta. \end{aligned}$$




FIGURE 31. Projected to the leaf space of the Morse flow, intersection points with the 4-chain are in 1-1 correspondence with the points of γ where the tangent vector is negatively parallel to the perturbation vector field.




FIGURE 32. Sign of the intersection points

From the first equation, we see that $\tau = \lambda\epsilon|t|$. The second equation then gives $\beta = -1$: if $\beta = 1$ the right hand side is of size $|\lambda\delta t|$, whereas the left hand side is of the smaller size $|\lambda\epsilon t|$. For $\beta = -1$ we get

$$\tau\pi_{\mathbb{R}^2}\dot{\gamma} = -t\lambda\epsilon v, \quad \tau\pi_{\mathbb{R}^2}\dot{\gamma} = \lambda\delta(t-1).$$

It follows that any point where $\pi_{\mathbb{R}^2}\dot{\gamma}$ is negatively parallel to v is an intersection point, take

$$\tau = t\lambda\epsilon > 0, \quad t = \frac{\lambda\delta}{\lambda\delta - \lambda\epsilon\pi_{\mathbb{R}^2}\dot{\gamma}} > 0.$$

Note, that the intersection lies very close to the boundary of the holomorphic curve, close to the boundary point where $\pi_{\mathbb{R}^2}\dot{\gamma}$ is negatively parallel to v ; see Figure 31. The sign of the intersection point is given by the sign of the orientation of

$$\partial_{\sigma}u \wedge \partial_{\tau}u \wedge J\xi_{\text{Mor}}$$

at that point, which is the same as the sign of the orientation of

$$\pi_{\mathbb{R}^2}\ddot{\gamma} \wedge \pi_{\mathbb{R}^2}\dot{\gamma} \wedge \xi_{\text{Mor}},$$

see Figure 32. The lemma follows. \square

Remark 9.7. In Lemma 9.6, we computed intersections from one local 4-chain. Our actual 4-chain is built from a local system of vector fields and the intersection is the weighted sum of two contributions.

We next consider contributions from Morse flow line parts of the lift \tilde{u} . Let u_{strip} denote the holomorphic strip close to a flow line in \tilde{u} .

Lemma 9.8. *For all $\lambda > 0$ sufficiently small, $C_L \cap u_{\text{strip}} = \emptyset$.*

Proof. The holomorphic strip is within distance $\mathcal{O}(\lambda^2)$ of the fiber strip over the flow line, with fibers in direction $\xi_{\text{Mor}}^* + \mathcal{O}(\lambda)$. Since C_L has inner normal $\xi_{\text{Mor}} + \epsilon v$, where v is not parallel to ξ_{Mor} , the lemma follows. \square




FIGURE 33. A local model for \tilde{u} with a holomorphic strip attached. Here, $\xi_{\text{Mor}} = \pm \partial_{x_3}$ in L_{\pm} and $\xi_{\text{pert}} = -\epsilon \partial_{x_1}$.

We finally consider the contribution at junction points, where a flow line strip u_{strip} is attached to u . Let γ denote the arc in ∂u where u_{strip} is attached.

Lemma 9.9. *If the tangent 2-plane containing $\dot{\gamma}$ and ξ_{Mor} at the junction point does not contain ξ_{pert} then there are no 4-chain intersections of \tilde{u} near the junction point.*

Proof. The holomorphic curve \tilde{u} with boundary on L close to this configuration is constructed by Floer gluing as in [10]. It follows in particular that \tilde{u} is arbitrarily close to the pregluing of the explicit local models of the curves given in [10, Section 6]. Thus, outside a small neighborhood of the point where the flow line is joined to the curve, the tangent vectors of the curve on L are proportional to ξ_{Mor} and $\dot{\gamma}$, where γ is the arc in u where u_{strip} is attached. Inside a small region, the tangent vector is arbitrarily close to the linear interpolation of these vectors and in particular lies arbitrarily close to the plane they span. It follows that, as long as ξ_{pert} at the junction point does not lie in this plane, there are no 4-chain intersections near the boundary. \square

The results above lead to the following combinatorial description of the a -power of the skein lift.

Corollary 9.10. *The a -power of the skein lift is given by the turning number of the lifted curve projected to the leaf space.*

Proof. Contributions from points where \tilde{u} is simply a lift of u , and points where \tilde{u} agrees with the holomorphic strip of a flow line are given by the turning number by Lemmas 9.6 and 9.8. Consider next a junction point, by Lemma 9.9 there are no local 4-chain intersections near the junction point. However, as the boundary of the glued curve is arbitrarily close to the pregluing we find that there are kinks in the projection to the leaf space. It is straightforward to check that these give the same contribution as turning numbers of smoothings, see Remark 9.11. \square

Remark 9.11. As in the proofs above, to check the relation between actual gluing and smoothing near junctions for all small $\lambda > 0$, it suffices to carry out the calculation in a model situation. We take \mathbb{R}^3 coordinates with leaf space \mathbb{R}^2 of the first two coordinates and use the model described in Figure 33. The same model projected along the framing vector field $\xi_{\text{Mor}} + \xi_{\text{pert}}$ is shown in Figure 34. Note that while the framing of the link $\partial \tilde{u}_{\pm}$ changes depending on the sign of ϵ , the extra a -powers are exactly compensated by the 4-chain intersections of \tilde{u}_{\mp} : the extra a -powers appearing in the case $\epsilon > 0$ can be interpreted in two skein-equivalent ways, either as the framing factor for $\partial \tilde{u}_{\pm}$ (left-hand sides) or as the 4-chain




FIGURE 34. The same model projected along the framing vector field $\xi_{\text{Mor}} + \xi_{\text{pert}}$

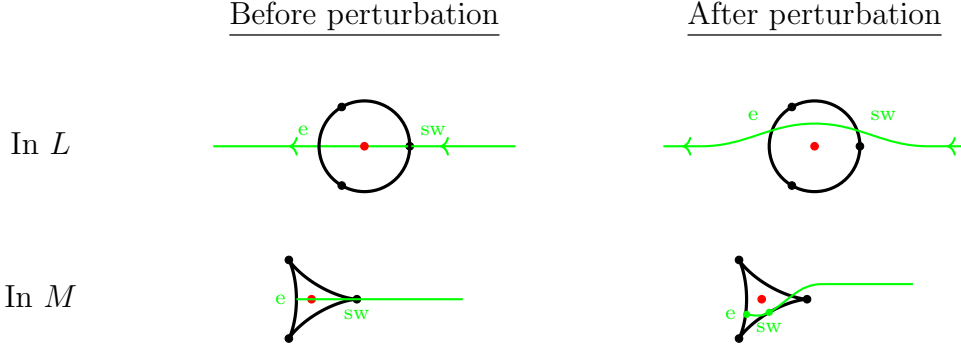


FIGURE 35. Flow tree near the branch locus with a 2-valent switch vertex and a 1-valent end vertex. (e denotes an end and sw denotes a switch.)

intersection of \tilde{u}_{\mp} (right-hand sides). This shows that, up to skein equivalence, the lifted link obtained by attaching a flow line can be taken to be the resolution of the crossing, in the projection along the Morse flow vector field ξ_{Mor} .

9.5. Spin structures. We will make the following spin structure choices. We fix a spin structure on the base 3-manifold M .⁷ On L , we pull back the spin structure along $L \rightarrow M$. This makes sense away from the branch locus, and defines a spin structure on the complement of the branch locus which does not extend across the branch locus. As noted in Section 7.1 above, we may still do skein-valued curve counting using this broken spin structure, but must correspondingly introduce a sign line into the skein of L along the branch locus.

The significance of this particular spin structure choice is that for flow graphs corresponding holomorphic curves along L , because both edges of the holomorphic curve corresponding to a given flow line are along pieces of Lagrangian with canonically identified spin structures, the signs introduced by the spin structure cancel, save possibly where the flow graph ends along image of the branch locus.

9.6. Flow graphs on a branched double cover. Let M be a 3-manifold with a signed taut triangulation Δ with t tetrahedra. Let $L^\circ \subset T^*M^\circ$ be the singular double cover as in Proposition 9.1. This is an exact Lagrangian cobordism connecting t Legendrian tori $\Lambda_{\mathbb{T}}$ at

⁷A 3-manifold is spin if orientable. We can generalize to the case of non-orientable M as follows. On T^*M , there is a choice of symplectic background class coming from the Lagrangian polarization by cotangent fibers, and there is a canonical choice of relative spin structure on the base manifold with respect to the background class. A spin structure on the base manifold gives an identification of this choice of background class with the standard symplectic background class.

the barycenter of each tetrahedron in Δ in the negative end $\approx \sqcup S^5$ of T^*M° to b Legendrian tori at the vertices of Δ in the positive end of T^*M° .

Let $K \subset M$ be a link and let $\mathbf{U}_K \subset T^*M$ be its pushed-off conormal. By Lemma A.5, holomorphic curves ending on $\mathbf{U}_K \sqcup L^\circ$ with boundary in \mathbf{U}_K that go once around the longitude are in 1-1 correspondence with configurations Γ of the standard holomorphic cylinder between \mathbf{U}_K and M with flow graphs for $L^\circ \sqcup M$ on it.

Lemma 9.12. *In any flow configuration Γ of the basic cylinder of \mathbf{U}_K with flow graphs of L° the flow graph does not have 3-valent vertices. Furthermore, the flow graphs attached are of two types*

- (i) *Flow line segments connecting two points on K .*
- (ii) *Flow line segments with an end and a switch connecting a point of K to the branch locus, see Figure 35.*

Proof. By exactness L° has no flow loops, and since the multiplicity of $L^\circ \rightarrow M^\circ$ is two outside a small neighborhood of the branch locus, there can be no trivalent vertices. Therefore, in this region, all flow graphs are simply flow lines. It remains to understand flow graphs near the branch locus. It is straightforward to check that any such flow graph of dimension zero has one switch and one end and that there are three such flow graphs along the branch locus, see Figure 35. It follows that rigid flow graph components are either flow lines connecting distinct points on K or intersections of K with the flow manifolds of flow lines with a switch and an end emanating from the branch locus. \square

Assume next that the ideal triangulation Δ of M is equipped with signed taut structure that satisfies the criterion of Proposition 3.10. Let $L_\tau \subset T^*M$ denote the smoothing of L° with these markings and periods.

Theorem 9.13. *Fix the brane data on L_τ constructed as in Sections 9.3 and 9.5. Then the maps from Theorems 1.3 and 4.2 agree, i.e., for all knots $K \subset M$, we have*

$$[K]_L = [K]_L^{\text{net}} \in \text{Sk}(L)$$

Proof. We will compute $[K]_L$ directly from its definition in the flow tree limit. Let \mathbf{U}_K be the pushed-off conormal; we count holomorphic curves ending on $\mathbf{U}_K \sqcup L$ whose boundary in \mathbf{U}_K goes once around the longitude.

We first consider the cobordism L° with negative ends. By Lemmas A.5 and A.6, there is 1-1 correspondence between flow graphs on the holomorphic cylinder of \mathbf{U}_K and holomorphic curves with boundary on L° . By Lemma 9.12 the flow graphs are exactly the configurations that contribute to $[K]_L^{\text{net}}$. We need to check signs. The sign of a glued curve can be calculated from a gluing sequence as in [11], as there this calculation reduces to a comparison of evaluations of kernel/cokernel elements at the junction points and the resulting sign is the crossing sign for exchanges and the crossing sign with the detour manifold at detours.

For exchanges, the kernel of the marked points gives the tangent vectors of the curve branches at the crossing and the kernels of the strip of the flow line gives the vertical vector.

For detours, we use the following model. We think of the critical leaves near the branch locus as a moduli space of disks with Reeb chord asymptotics, see Figure 37, where we stretch around the perturbed branch locus.

We orient this moduli space as follows. Choosing an orientation of the branch locus orients the normal bundle of the branch locus which in turn gives oriented zero dimensional disks, see Figure 37. We then orient the 1-dimensional moduli space using the branch locus

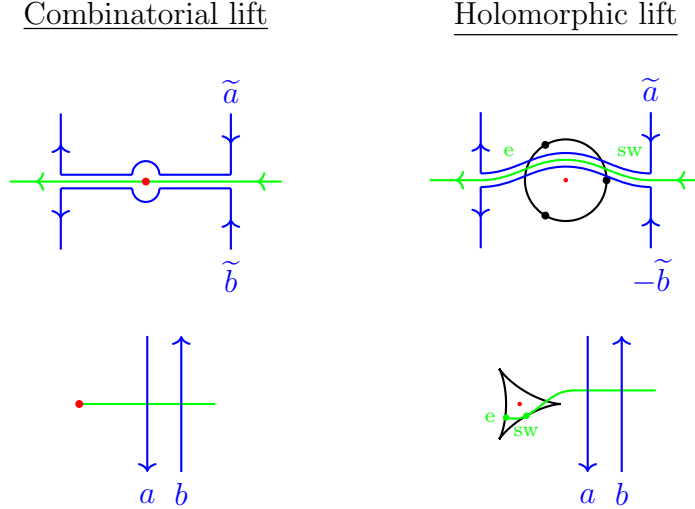


FIGURE 36. The combinatorial and holomorphic lifts agree: isotope \tilde{b} in the combinatorial lift across the sign line.

orientation again. Note that this uses the branch locus orientation twice and hence the resulting orientation is well defined. This then gives a coorientation to the two-dimensional surface in L that is the evaluation along the boundary of the disks in the moduli space.

We next use the gluing sequence to orient the disks that results from gluing flow graphs as discussed to holomorphic curves with boundary on L , compare [14]. Here we should compare the oriented kernel vector fields of the gluing point evaluation maps along the moduli spaces and the orientation of L . As in [14], signs of relevant L^2 inner products can be computed using evaluation maps and here the gluing sign is given by the intersection sign of the boundary orientation of the big curve and the oriented moduli space. This matches the combinatorial calculation up to sign, see Figure 36.

To see that the holomorphic curve sign matches the combinatorial sign exactly, we use deformation invariance: Consider the boundary of a holomorphic curve that intersects the projection of the evaluation map of the moduli space of detour disks, i.e., a critical leaf transversely at one point. Consider its lifts, with the boundary condition that specifies the sheet to which we lift at the incoming and out going ends – this corresponds to specifying a Reeb chord of the detour attached. We take these lifting conditions so that the only possible choice is the lift with one detour. On the other hand, if we isotope the boundary of the curve across the branch locus, then it intersects two critical leaves, but with the given boundary conditions, the only possible lift is the direct lift without using any detour and for the direct lift there is no extra sign. The result follows for lifting to the Lagrangian L_Δ° .

Consider now the smoothing L_τ . Note the angular periods can be scaled $(z_1, \dots, z_t) \mapsto \epsilon(z_1, \dots, z_t)$ and still satisfy the condition in Proposition 3.10. This means we can consider the SFT-stretched picture, where L_τ is obtained by gluing Harvey Lawson cones $\sqcup L_{\mathbb{T}}$ to L° in $(T^*M)^\circ$ to $\sqcup \Lambda_{\mathbb{T}}$ at the negative end, according to $\epsilon(z_1, \dots, z_t)$ for all sufficiently small $\epsilon > 0$. Because $\Lambda_{\mathbb{T}} \subset S^5$ has no index ≤ 0 Reeb chords or orbits, the curve counting problem separates. We have already determined the curves for L° , and the curves in the smoothed Harvey-Lawson cone $L_{\mathbb{T}}$ were determined in [18]. The theorem follows. \square

Remark 9.14. It is interesting to see how the sign rules above are affected by an orientation reversal of M . For the exchanges, all crossings change sign, and hence there is a -1 for each strip attached. Since the strips also change the number of boundary components, we find that the change of sign matches the overall sign. For the detours, they do not change the number of boundary components, so the sign changes of boundary components of curves before and after lifting commute. Note also that the curves in the relation at negative ends are unchanged but that the a -powers changes sign which takes recursion relations of disks to those of anti-disks.

9.7. 3d spectral network as a moduli space of holomorphic disks. In this section, we show that the 2-skeleton $\Delta_{\leq 2}^{\vee}$ of the polygonal decomposition dual to an ideal triangulation Δ on a 3-manifold M has a natural interpretation as the compactified moduli space of certain holomorphic curves with boundary on L_{Δ}° . This gives a holomorphic curve interpretation of the foliation used to define the skein lifting map that will be useful if one considers similar questions for ‘large’ deformations of the Lagrangian to which we lift in the cotangent bundle.

Let M and Δ be as above and construct the Lagrangian cobordism L_{Δ}° . Consider the moduli space $\mathcal{M}(L_{\Delta}^{\circ})$ of holomorphic disks in T^*M with two positive punctures asymptotic to Reeb chords in $\Lambda_{\partial\Delta}$.

Lemma 9.15. *The interior of $\mathcal{M}(L_{\Delta}^{\circ})$ is naturally diffeomorphic to the union of the interiors of the 2-cells of Δ^{\vee} .*

Proof. We use the correspondence between flow lines and holomorphic disks. There is exactly one smooth flow line in the foliation determined by difference of the sheets of L_{Δ}° through each interior point of a 2-cell. This flow line connects vertices of Δ and has unique asymptotics in $\partial_{\infty}M$. By the correspondence, we then find associated holomorphic strips. The diffeomorphism is obtained by intersection with $\Delta_{\frac{1}{2}}^{\vee}$. It is clearly a diffeomorphism for the Morse flows, and since this is an open condition, also for the associated holomorphic curves. \square

We next consider the boundary of $\mathcal{M}(L_{\Delta}^{\circ})$. There is one boundary configuration that corresponds to boundary breaking, where the strip with two positive punctures breaks into two disks with one positive puncture each. We denote this boundary component $\partial_{\times}\mathcal{M}$.

Lemma 9.16. *The boundary $\partial_{\times}\mathcal{M}(L_{\Delta}^{\circ})$ is naturally diffeomorphic to the union of the 1-cells of Δ^{\vee} .*

Proof. Consider first the flow lines. As a flow line goes to the boundary of a 2-cell it breaks into two flow lines from the vertices at its end to the branch locus which is the 1-cell at its boundary. Note that these flow lines correspond to flow trees with a switch and an end, see Figure 35, and that their lifts intersect transversely in a single point and that these are the only such flow trees. We thus find that there are pairs of intersecting disks corresponding to the boundary configurations, and by transversality of the evaluation maps and uniqueness of gluing we find that the pairs of disks give the boundary $\partial_{\times}\mathcal{M}$ as claimed, see Figure 37. \square

Finally, we consider the boundary of the moduli space where parts of the curve fall into the negative end. We call this boundary $\partial_{-}\mathcal{M}$.

Lemma 9.17. *The boundary $\partial_{-}\mathcal{M}$ corresponds to the barycenters in Δ . Curves in $\partial_{-}\mathcal{M}$ consist of two strips with a positive and a negative puncture each and two disks in the negative*




FIGURE 37. The disks in the boundary corresponding to flow lines ending in the branch locus.

end, one at each negative asymptotic. The curves in the negative end form a 2-dimensional space: each disk can be translated in the \mathbb{R} -direction.

Proof. Remaining flow lines connect vertices to barycenters. These correspond to strips with a positive and a negative Reeb chord asymptote. Remaining claims follow from the description of the moduli space of curve on $\Lambda_{\mathbb{T}}$, see [13]. \square

We summarize the results in Lemmas 9.15, 9.16, and 9.17.

Corollary 9.18. *The moduli space $\mathcal{M}(L_{\Delta}^{\circ})$ of holomorphic disks in T^*M° with two positive punctures and boundary on L_{Δ}° is naturally isomorphic to $\Delta_{\leq 2}^{\vee}$, where 2-dimensional stratum correspond to one level unbroken curves, the 1-dimensional stratum to nodal disks with one boundary node, and the 0-dimensional stratum to SFT broken curves with fixed positive end two disks with one positive and one negative puncture each and two negative end disks at the negative end Reeb chords, each parameterized by \mathbb{R} .*

In particular, evaluation at a boundary marked point $\text{ev}: \mathcal{M}(L_{\Delta}^{\circ}) \rightarrow L_{\Delta}^{\circ}$ is an injective degree one map. \square

10. WALL CROSSINGS

In this section, we will prove Theorems 1.6 and 1.7. The basic idea of the proofs of these results is the following. Suppose there was a Lagrangian $L \subset T^*(S \times \mathbb{R})$ whose symplectic reductions over $t \in \mathbb{R}$ gave $\Sigma_t \subset T^*S$. For a knot $K \subset S$, we could then consider $[K]_L$ and compute it by putting K at either very negative or very positive t . Then by Theorem 1.5, we would deduce as desired:

$$[K]_{\Sigma_-}[\emptyset]_L = [K]_L = [\emptyset]_L[K]_{\Sigma_+}.$$

Such a Lagrangian L exists when the Lagrangian isotopy is exact, but the isotopies of interest here (for which $[\emptyset]_L \neq 0$) are not exact and the Lagrangian L does not exist. However, as we will see below, it is possible to construct a totally real \tilde{L} which, despite not being

Lagrangian, can play a role analogous to that of L and using \tilde{L} we prove Theorem 1.6. To then establish Theorem 1.7 we show that under its hypotheses, \tilde{L} bounds a single disk, whose multiple cover count was previously determined in [18].

10.1. Lagrangian isotopies and proof of Theorem 1.6. We consider first a general construction of totally real submanifolds associated to ‘slowed down’ Lagrangian isotopies and show that there exists almost complex structures for which these totally real submanifolds are good boundary conditions for holomorphic curves.

Let $L \subset X$ be a Lagrangian, possibly non-compact in which case we assume it is asymptotic to a Legendrian. Let $\iota_t: L \rightarrow X$, $0 \leq t \leq 1$ be a Lagrangian isotopy which we take to be constant near the ends. We will slow down the isotopy so that its trace is suitable for holomorphic curve counting.

Fix an integer $N > 0$, which we will ultimately take sufficiently large. First, rescale the isotopy to $\phi'_t = \iota_{\frac{1}{N}t}$, $0 \leq t \leq N$. Then $|\frac{\partial \phi'_t}{\partial t}| = \mathcal{O}(1/N)$. Second, transport the isotopy ϕ'_t to $[0, 2N]$ and make it constant in intervals $[2j-1, 2j]$. We call the resulting isotopy ϕ_t , $t \in \mathbb{R}$, where we extend it to be constant in $(-\infty, 0]$ and $[2N, \infty)$. More formally, let $\kappa: [0, 1] \rightarrow [0, 1]$ be an approximation of the identity map that is constant in small neighborhoods of the end points. Then

$$\phi_t = \begin{cases} \iota_0, & t \in (-\infty, 0], \\ \iota_{\frac{\kappa(t-2j)+j}{N}}, & t \in [2j, 2j+1], \quad 0 \leq j \leq N-1, \\ \iota_{\frac{j+1}{N}}, & t \in [2j+1, 2j+2], \quad 0 \leq j \leq N-1, \\ \iota_1, & t \in [2N, \infty). \end{cases}$$

Consider now the trace $\Phi: L \times \mathbb{R} \rightarrow X \times T^*\mathbb{R}$ of ϕ_t :

$$\Phi(q, t) = (\phi_t(q), t, 0).$$

Lemma 10.1. *The trace is Lagrangian over $(-\infty, 0]$, over $[2N, \infty]$, and over each interval $[2j+1, 2j+2]$, $0 \leq j \leq N-1$. Furthermore over each interval $[2j, 2j+1]$, $0 \leq j \leq N-1$, there is a diffeomorphism $\psi_{2j}: X \times T^*[2j, 2j+1] \rightarrow X \times T^*[2j, 2j+1]$ of C^1 -distance $\mathcal{O}(1/N)$ from the identity such that $\psi_{2j}(\Phi(q, t)) = (\phi_{2j+\frac{1}{2}}(q), t, 0)$.*

Proof. Direct consequence of the construction of ϕ_t . □

It follows from Lemma 10.1 that the trace $\Phi(L \times \mathbb{R})$ is totally real for N sufficiently large. Furthermore, if J_X is an almost complex structure on X compatible with its symplectic form ω , then the complex structures $J_{2j} = d\psi_{2j}^{-1} \circ J_X \circ d\psi_{2j}$, $0 \leq j \leq N$, are compatible with form ω provided N is sufficiently large. Define an almost complex structure J compatible with $\omega_0 = \omega + dt \wedge d\tau$ on $X \times T^*\mathbb{R}$ by taking $J = J_{2j} \oplus J_{\mathbb{R}}$ over $[2j, 2j+1]$, where $J_{\mathbb{R}}$ is the standard complex structure on $T^*\mathbb{R}$, and $J_t \oplus J_{\mathbb{R}}$, $t \in [2j+1, 2j+2]$, where J_t interpolates between J_{2j} and J_{2j+2} over $[2j+1, 2j+2]$. More formally,

$$\begin{cases} J|_{X \times T^*(-\infty, 0]} & = J_X \oplus J_{\mathbb{R}}, \\ J|_{X \times T^*[2j, 2j+1]} & = J_{2j} \oplus J_{\mathbb{R}}, \quad 0 \leq j \leq N-1, \\ J|_{X \times T^*[2j+1, 2j+2]_t} & = J_t \oplus J_{\mathbb{R}}, \quad 0 \leq j \leq N-1, \\ J|_{X \times T^*[2N, \infty)} & = J_X \oplus J_{\mathbb{R}}. \end{cases}$$

We next show that J -holomorphic curves in $X \times T^*\mathbb{R}$ with boundary on $\Phi(L \times \mathbb{R})$ are confined to intervals of length 1.

Lemma 10.2. *Any connected J -holomorphic curve with boundary on $\Phi(L \times \mathbb{R})$ that takes some point into $X \times T^*[2j, 2j + 1]$, into $X \times T^*(-\infty, 0]$ or into $X \times T^*[2N, \infty)$ lies entirely in a slice $(t, \tau) = (c, 0)$.*

Proof. By our choice of complex structure and the definition of Φ , the J -holomorphic curve is holomorphic for a split complex structure $J_t \oplus J_{\mathbb{R}}$ and has boundary on a Lagrangian of the form $L \times \mathbb{R} \subset T^*\mathbb{R}$. The projection of the curve to $T^*\mathbb{R}$ is then holomorphic and hence constant by the maximum principle. \square

Proof of Theorem 1.6. We use notation as above for $L = \Sigma$ and $X = T^*S$. Assume that the moduli space of $\phi_0(\Sigma)$ and $\phi_{2N}(\Sigma)$ in T^*S are empty.

We claim then that there exists a perturbation for bare curves with boundary on $\Phi(\Sigma \times \mathbb{R}) \subset T^*S \times T^*\mathbb{R}$ and $\epsilon > 0$ such that no bare curve intersects the regions $T^*S \times T^*[-\infty, 0]$ and $T^*S \times T^*[2N, \infty)$.

To see this, note that by Lemma 10.2, any connected J -holomorphic curve which intersects these regions in fact lies in the fiber over $(t, 0) \in T^*\mathbb{R}$ and is J_0 -holomorphic respectively J_{2N} -holomorphic. Since the corresponding moduli spaces are empty, moduli spaces are empty in 4ϵ -intervals around $t = 0$ and $t = 2N$ by Gromov compactness. Since bare curve perturbations are supported in a small neighborhood of the original unperturbed moduli spaces, the claim holds.

Equation (1.5) then follows: take $\Omega(\phi)$ to be the skein-valued bare curve count of

$$\Phi(L) \cap T^*S \times T^*[\epsilon, 2N - \epsilon]$$

in $\text{Sk}(L)$ and observe that the left and right hand sides correspond to including the conormal of K on the far left and on the far right, respectively; they are equal by deformation isotopy invariance of the skein count where the isotopy translates the conormal in the \mathbb{R} -direction.

We next consider homotopy invariance. Let $\phi_{s,t}$, $0 \leq s \leq 1$, be a 1-parameter families of isotopies. By continuity in s we find N so that after rescaling, the size of the derivative $\frac{\partial \phi_{s,t}}{\partial t}$ is sufficiently small for all s . As for $\Phi(L \times \mathbb{R})$, the 1-parameter family of totally real submanifolds $\Phi_s(L \times \mathbb{R})$ provides an isotopy of good boundary conditions for bare curve counts, and hence the desired homotopy invariance of $\Omega(\phi)$ follows from deformation invariance of the skein-valued curve count. \square

10.2. Transverse (-1) -disk and the proof of Theorem 1.7. We next consider specific 1-parameter families $\phi_t: \Sigma \rightarrow T^*S$ such that there are no holomorphic curves for $t \neq 0$ and such that there is a single embedded disk D of index -1 and boundary on $\phi_0(\Sigma)$. We assume that this disk is transversely cut out in the sense of 1-parameter families. In terms of Cauchy-Riemann operators, this means the following.

Consider a cotangent neighborhood $T^*\Sigma$ of $\phi_0(\Sigma)$. The Lagrangian isotopy ϕ_t moves $\phi_0(\Sigma)$ inside $T^*\Sigma$ by shifting along closed 1-forms. We extend it to an ambient isotopy by cutting off and denote it $\tilde{\phi}_t$. The transversality condition then means that the linearized $\bar{\partial}_J$ at D has cokernel of dimension 1 and that the linearized variation of D given by $\frac{\partial \tilde{\phi}_t}{\partial t}|_{t=0}$ spans the cokernel. We take the sign of the disk to be the sign induced by the standard orientation of the cokernel.

Proof of Theorem 1.7. Consider the trace $\Phi: L \times \mathbb{R} \rightarrow T^*S \times T^*\mathbb{R}$ around 0 and note that under the assumptions above, the disk D is an embedded disk with boundary on $\Phi(L \times \mathbb{R})$.

The tangent bundle $T(T^*S)|_D$ splits as two complex line bundles $TD \oplus ND$, the tangent and normal bundles to D . Consider first the multiplication of $\phi_0(\Sigma)$ with $T^*\mathbb{R}$. In this case, the boundary condition splits as the sum of two real line bundles in TD and ND , respectively. Trivializing these line bundles, the boundary conditions give rotating real lines along the boundary. Since the index of the disk equals -1 and the cokernel is 1-dimensional it follows that the Maslov index of the line in the tangent direction equals 2 and the line in the normal direction equals -2 . The additional direction is not rotating at all, and the kernel consists of constant sections in this direction. This means that the linearized Cauchy-Riemann equation in the normal direction for this trivial product is conjugate to the standard $\bar{\partial}$ -operator on the unit disk D with values in \mathbb{C}^2 and Lagrangian boundary condition given by the columns in the matrix

$$\begin{pmatrix} e^{-i\theta} & 0 \\ 0 & 1 \end{pmatrix}, \quad e^{i\theta} \in \partial D.$$

Consider next the boundary condition of the trace of a deformation that spans the cokernel. Since it is a small deformation of the constant condition, Fourier expanding the boundary condition, it is straightforward to check that it is conjugate to

$$\frac{1}{1+r^2} \begin{pmatrix} e^{-i\theta} & -re^{-i\theta/2} \\ re^{-i\theta/2} & 1 \end{pmatrix}, \quad e^{i\theta} \in \partial D,$$

where $r > 0$ is small. This in turn agrees with the boundary condition of the toric brane, up to $\mathcal{O}(r^2)$, studied in [18], where the multiple cover formula is shown to be the exponentiated skein dilogarithm. The theorem follows. \square

Remark 10.3. Using Morse flow graphs one can give another proof of the recursion relation for multiple covers of a disk corresponding to an isolated flow line as follows. Consider a 3-manifold M with a branched double cover L with branch locus $\tau \subset M$. Consider a rigid basic disk w corresponding to a flow line of $L \rightarrow M$ that connects two points on τ . We study what happens when a moving big holomorphic curve u_s , $-\epsilon < s < \epsilon$ with boundary on M crosses the boundary of w at $s = 0$. (The map u_s could for example be the basic cylinder of a conormal of a knot K_s that moves in such a way that K_0 intersects the boundary of w and such that the tangent vector of K_0 together with the deformation vector and the tangent vector of ∂w spans the tangent space to M .)

The basic disk w itself has multiple covers that correspond to flow lines with multiplicity. Thus, at the moment $s = 0$ when u_0 hits the disk, there are broken flow graph and big curve configurations $u_0 \cup w$, but at moments before or after there are not. As u_s is embedded, Lemma A.6 shows that there can be no curves with multiply covered edges attached in the limit. Looking at possible limits before (one flow line to the branch locus attached to u) and after the crossing (two flow lines to the branch locus attached) then reveals the recursion relation for the disk, where the operators correspond to the curves with simple flow lines attached. This gives another illustration of why there are simple recursion relations that control the many skein instances involved in a wall crossing: there are no very complicated big curves nearby.

11. EXAMPLES OF SKEIN-VALUED WALL-CROSSING FORMULAS

In this section, we apply Theorems 1.6 and 1.7 to derive skein-valued versions of well-known wall crossing formulas.

Let C be a Riemann surface. A quadratic differential $\phi \in H^0(C, K_C^{\otimes 2})$ on C defines a two-fold ramified covering $\Sigma \rightarrow C$:

$$(11.1) \quad \lambda^2 = \phi$$

with λ the complex Liouville 1-form in T^*C . The cover Σ is smooth so long as ϕ has nondegenerate zeroes.

We will consider such Σ as real Lagrangians with respect to the real symplectic form $\omega = \operatorname{Re} d\lambda$. We will count holomorphic curves in a complex structure J tamed by ω , *not* the complex structure on T^*C coming from the complex structure of C .

Let us recall the basic consequences of Theorem 1.6. Suppose given 1-parameter family of such quadratic differentials, ϕ_t for t varying in some interval $I = [0, 1]$, such that the corresponding Σ_0 and Σ_1 bound no J -holomorphic curves. Then, by Theorem 1.6, this family determines an element $Z(\phi) \in \widehat{\operatorname{Sk}}(\Sigma)$, where here we write Σ to indicate the total space of the family Σ_t . Moreover, suppose given a 2-parameter family $\phi_t[u]$ over some $(t, u) \in [0, 1] \times I$ where the corresponding $\Sigma_0[\cdot]$ and $\Sigma_1[\cdot]$ never bound holomorphic curves. Then, again by Theorem 1.6, the element $Z(\phi[u])$ is independent of u . (This independence should be regarded as a general ‘‘wall crossing formula’’.)

To actually evaluate $Z(\phi)$, we will use Theorem 1.7. To do so we must show that our family Σ_t bound holomorphic curves at some isolated collection of t , and each such instance is an embedded disk, transverse in an appropriate sense. Choosing a metric on the base and passing to flow graphs, we must ask that the only flow graphs which appear are flow lines (and that they appear transversely in t). Such flow lines necessarily connect the branch points of the cover.

Lemma 11.1. *For an appropriate choice of metric, Morse flow lines agree with the horizontal foliation on C defined by the quadratic differential, i.e.*

$$(11.2) \quad \operatorname{Im}(e^{i\theta} \sqrt{\phi}) = 0.$$

Proof. Indeed, this correspondence of flow lines and horizontal foliation is well known for the singular conformal metric associated to the quadratic differential, see e.g., [65] and [7, Example 2.23]. This metric is flat except at a finite number of isolated points that are modeled on Euclidean cones. Rounding the tip of these cones we get a smooth metric that agrees with the flat metric outside arbitrarily small disk neighborhoods of the cone points. In particular, flow lines in the smooth metric converge to flow lines for the singular metric. \square

Flow lines connecting branch points are the ‘saddle connections’ of such foliations, which are precisely what enter into the Gaiotto-Moore-Neitzke formulas [30]; correspondingly, these lines have already been determined in the examples we study below.

11.1. The pentagon relation. Let $C = \mathbb{C}$ with quadratic differential

$$(11.3) \quad \phi_t[u](z) = e^{2\pi it}(z^3 - 3z + 2u) dz^2.$$

This has no double zeros so long as $u \in \mathbb{C} \setminus \{\pm 1\}$. The variety in T^*C given by (11.1) is the Seiberg-Witten curve for Argyres-Douglas theory introduced in [4] and presented in this form in [30]. It is well-known that the count of holomorphic disks depends on u piecewise continuously, and there are two regions corresponding to $|u| \lesssim 1$ and $|u| \gtrsim 1$.

For $u = 0$ there are two holomorphic disks, corresponding to the trajectories shown in Figure 38. Disk boundaries correspond to generators of $H_1(\Sigma, \mathbb{Z})$, respectively denoted $(0, 1)$ for the first disk (with smaller value of t) and $(1, 0)$ for the second one.




FIGURE 38. Two holomorphic disks at $u = 0$ for $0 \leq \theta \leq \pi$.FIGURE 39. Three holomorphic disks at $u = 2i$ for $0 \leq \theta \leq \pi$.

FIGURE 40. The skein-valued pentagon relation on the punctured torus.

Moving across the ‘line of marginal stability’, approximately described by $|u| \approx 1$, causes the two disks to swap ordering in t . For $u = 2i$ there are three holomorphic disks, shown in Figure 39. In addition to disks with boundaries $(1, 0)$ (now with largest value of t) and $(0, 1)$ (with lowest value of t) there is now another disk whose boundary has class $(1, 1)$.

From Theorems 1.6 and 1.7, we deduce the skein-valued pentagon relation:

$$(11.4) \quad \Psi_{(1,0)} \Psi_{(0,1)} = \Psi_{(0,1)} \Psi_{(1,1)} [-1] \Psi_{(1,0)}.$$

Here $\Psi_{(i,j)}$ is the skein-valued dilogarithm (2.3) with P_d replaced by $P_{di,dj}$, corresponding to insertions on the boundary torus in homology class (di, dj) , see [46, Definition 2.4]. Relation (11.4) is illustrated in Figure 40, and was originally derived in [48, 35], and generalized in [58]. As noted in those references, the skein-valued pentagon identity in the punctured torus implies, via [47, 5], the pentagon identity in the elliptic Hall algebra, originally proven in [33, 68].

Specializing to the \mathfrak{gl}_1 skein algebra of T^2 with generators $\hat{y}\hat{x} = q\hat{x}\hat{y}$ is given by the replacements

$$(11.5) \quad P_{i,j} \rightarrow q^{-ij/2} \hat{y}^i \hat{x}^j, \quad \Psi_{i,j}[\xi] \rightarrow (\xi q^{1/2-ij/2} \hat{y}^i \hat{x}^j; q)_\infty^{-1},$$




FIGURE 41. Top row: direct lift of disk boundaries from Figure 39 to the torus double covering of C . Sign defects are marked by crosses. An additional sign defect is located in correspondence of the puncture, which is marked by a thick dot. Bottom row: a perturbation of disk boundaries away from sign defects.

and recovers the pentagon identity for the quantum dilogarithm, originally established in [20]. The pentagon identity has a celebrated derivation via a wall crossing in the space of stability conditions for the A_2 quiver [57], later reinterpreted in terms of the Donaldson-Thomas invariants of an appropriate CY3 category [39]. We do not know a similar interpretation of the skein-valued pentagon identity.


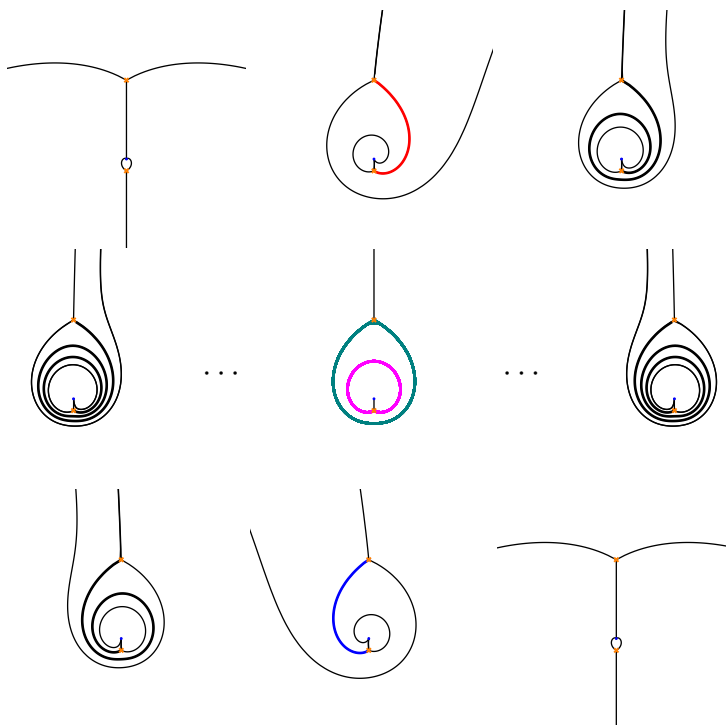
Remark 11.2. The sign in $\Psi_{(1,1)}[-1]$ can be understood geometrically as follows. In Figures 38 and 39 disk boundaries projected to C attach directly to zeroes of ϕ_2 , which correspond to projections of sign defects of the double covering. The direct lift of disk boundaries to the punctured torus is shown in the top row of Figure 41, where sign defects are also marked. In the same figure, bottom row, a perturbation of disk boundaries is shown. Note that the perturbed boundstate disk (in purple) differs from the curve obtained by smoothing the intersection of generating curves (purple line in Figure 40), by crossing the sign defect in the center. This is the geometric origin of the sign in $\Psi_{(1,1)}[-1]$.

Algebraically, the sign can be absorbed by twisting the skein algebra of the torus by a quadratic refinement, such as the one adopted in [30]. Let $\sigma : \mathbb{Z}^2 \rightarrow \mathbb{C}$ such that $\sigma(i, j)\sigma(k, l) = \sigma(i + k, j + l)(-1)^{jk-il}$. Note that this implies $\sigma(di, dj) = \sigma(i, j)^d$. We twist the skein algebra of the torus by introducing $\tilde{P}_{i,j} := \sigma(i, j)P_{i,j}$. Replacing $P_{i,j}$ into (2.3) gives the twisted skein-valued dilogarithm $\tilde{\Psi}_{(i,j)} \equiv \Psi_{(i,j)}[\sigma(i, j)]$. This obeys the pentagon identity without additional signs

$$(11.6) \quad \tilde{\Psi}_{(1,0)}\tilde{\Psi}_{(0,1)} = \tilde{\Psi}_{(0,1)}\tilde{\Psi}_{(1,1)}\tilde{\Psi}_{(1,0)}.$$

11.2. The Seiberg-Witten wall-crossing. Let $C = \mathbb{C}^*$ with a family of quadratic differentials

$$(11.7) \quad \phi_t[u](z) = e^{2\pi it} \left(\frac{1}{z^3} + \frac{2u}{z^2} + \frac{1}{z} \right) dz^2$$

FIGURE 42. Two holomorphic disks at $u = 0$ for $0 \leq \theta \leq \pi$.FIGURE 43. Infinitely many holomorphic disks at $u = -2i$ for $0 \leq \theta \leq \pi$.

parameterized by $u \in \mathbb{C} \setminus \{\pm 1\}$. The variety in T^*C given by (11.1) is the Seiberg-Witten curve for $\mathcal{N} = 2$ supersymmetric $SU(2)$ Yang-Mills theory introduced in [62] and presented in this form in [30].

Again there are two regions in parameter space, corresponding (roughly) to $|u| < 1$ and $|u| > 1$, with different curve counts.

Choosing $u = 0$ gives two holomorphic disks. The corresponding flow lines on C are shown in Figure 42. In a suitable choice of basis for $H_1(\Sigma, \mathbb{Z})$, the boundaries of the two curves are in class $(-1, 1)$ for the first disk (with smaller value of t) and $(1, 1)$ for the second one.

Deforming parameters to $|u| > 1$, the two disks swap order in t , but there are (in contrast to the previous example) infinitely many more embedded holomorphic disks, see Figure 43.

From Theorems 1.6 and 1.7, we deduce:

$$(11.8) \quad \Psi_{(1,1)} \Psi_{(-1,1)} = \Psi_{(-1,1)} \cdots \Psi_{(-1,2n+1)} \cdots \Psi_{A_{10}}^{-1} \Psi_{A_{01}}^{-1} \cdots \Psi_{(1,2n+1)} \cdots \Psi_{(1,1)}$$




FIGURE 44. Curves $(1, 1)$ (red), $(-1, 1)$ (blue), A_{10} (magenta) and A_{01} (green) after moving $(1, 1)$ across $(-1, 1)$.




FIGURE 45. Curves $(1, 1)$, $(-1, 1)$, A_{10} and A_{01} .

where ellipses contain contributions from all values of $n \in \mathbb{N}$. Here $A_{01}, A_{10} \in \text{Sk}(T^2 \times I)$ are connected curves in class $(0, 2)$ with a negative (respectively, positive) crossing shown in Figure 44. See Remark 11.4 for details about the associated factors $\Psi_{A_{10}}^{-1} \Psi_{A_{01}}^{-1}$.

Remark 11.3. Formula (11.8) can also be derived by repeated application of the pentagon relation of Figure 40. The first steps of this procedure are shown in Figure 44: we first push the curves $(1, 1)$ and $(-1, 1)$ through each other, creating curves A_{10} and A_{01} . After this step, further pushing curve $(1, 1)$ down will cause it to cross A_{10} generating further curves in class $(1, 2n + 1)$. Similarly, further pushing curve $(-1, 1)$ upward will also generate curves in class $(2n + 1)$ by unlinking with A_{10} . Note that A_{01} does not participate in further creation of curves, as shown in Figure 45.

Remark 11.4. A_{10} and A_{01} can be expressed in terms certain skein elements $P_{i,j}$ introduced by Morton and Samuelson [46]:

$$\begin{aligned}
 (11.9) \quad A_{10} &= \frac{1}{2}(q^{1/2} + q^{-1/2})P_{0,2} + \frac{1}{2}(q^{1/2} - q^{-1/2})P_{0,1}^2 \\
 A_{01} &= \frac{1}{2}(q^{1/2} + q^{-1/2})P_{0,2} - \frac{1}{2}(q^{1/2} - q^{-1/2})P_{0,1}^2
 \end{aligned}$$

The factors appearing in (11.8) are defined as follows

$$\begin{aligned}
 (11.10) \quad \Psi_{A_{10}}^{-1} &= \exp \left(\sum_{d \geq 1} \frac{1}{2d} \left(\frac{q^{1/2} + q^{-1/2}}{q^{1/2} - q^{-1/2}} P_{0,2d} + P_{0,d}^2 \right) \right) \\
 \Psi_{A_{01}}^{-1} &= \exp \left(\sum_{d \geq 1} \frac{1}{2d} \left(\frac{q^{1/2} + q^{-1/2}}{q^{1/2} - q^{-1/2}} P_{0,2d} - P_{0,d}^2 \right) \right).
 \end{aligned}$$

Each term can be given a geometric interpretation by an application of [45, Theorem 13], which allows to compute insertions of any $Q \in \text{Sk}(S^1 \times D^2)$ along a curve in $\text{Sk}(T^2 \times I)$. We have checked numerically that $\Psi_{A_{10}}^{-1}$ coincides with the insertion of a disk, i.e. a skein-valued dilogarithm (2.3), along curve A_{10} . Similarly, we have checked numerically that $\Psi_{A_{01}}$ (i.e. the second line in (11.10) with opposite overall sign in the exponent) coincides with the insertion of an anti-disk, i.e. the inverse of the skein-valued dilogarithm (2.5), along A_{01} . To obtain $\Psi_{A_{01}}^{-1}$ one needs to compute the multiplicative inverse. As an operation in the skein, this is complicated by the fact that annihilation of curves in the product $\Psi_{A_{01}}^{-1} \Psi_{A_{01}}$ involves moving (multiple coverings of) the disk and anti-disk onto each other, which induces repeated wall-crossings among their boundaries.

Remark 11.5. Formula (11.8) is a skein-valued lift of the wall-crossing formula for 4d $\mathcal{N} = 2$ $SU(2)$ Yang-Mills theory which was previously identified with the Kontsevich-Soibelman wall-crossing formula for the CY3 category associated to the Kronecker quiver [39, discussion below Conjecture 1] and [28, Equation (2.26)]. The ‘motivic’ counterpart of the Kontsevich-Soibelman wall-crossing formula is known to give, in this case, the spin refinement of the Seiberg-Witten wall crossing [9, Equation (1.3) and Section 4.1]. Our wall crossing formula (arrived at by entirely different considerations, having nothing to do with the Donaldson-Thomas invariants of a Calabi-Yau category) in fact specializes to the refined/motivic formula by specializing (11.8) to the $\mathfrak{gl}(1)$ skein via (11.5). In particular, we note that the $\mathfrak{gl}(1)$ -specialization of the middle factors $\Psi_{A_{10}}^{-1} \Psi_{A_{01}}^{-1}$ is $\Phi(q^{1/2}\hat{x}^2)^{-1} \Phi(q^{-1/2}\hat{x}^2)^{-1}$ where $\Phi(\xi) = (q^{1/2}\xi; q)_\infty^{-1}$, in agreement with [9, Equation (4.8)], [31, Section 4.1], and [49, Section 8.4].

Question 11.6. What (if any) question in the Donaldson-Thomas theory of the CY3 category associated to the Kronecker quiver – or what (if any) question about $\mathcal{N} = 2$ supersymmetric Yang-Mills theory – corresponds to the skein valued wall crossing (11.8)?

APPENDIX A. FLOW GRAPHS AND HOLOMORPHIC CURVES

A.1. The flow graph limit. Consider a Lagrangian $Q \subset T^*M$. We will take Q to lie very close to the zero section. We obtain this by fiber scaling $\sigma_\lambda: T^*M \rightarrow T^*M$, $\sigma_\lambda(q, p) = (q, \lambda p)$. Since $\sigma_\lambda^* \omega = \lambda \omega$, σ_λ preserves Lagrangians and in the non-exact case it scales the flux. Let $Q_\lambda = \sigma_\lambda(Q)$.

We first consider holomorphic curves with boundary on Q_λ only. If we consider the standard Gromov limit of such curves we have the following straightforward result.

Lemma A.1. *Let $u_\lambda: (\Sigma_\lambda, \partial\Sigma_\lambda) \rightarrow (T^*M, Q_\lambda)$ be a family of holomorphic curves, then some subsequence u_λ Gromov converges to a constant curve $u_0: (\Sigma_0, \partial\Sigma_0)$ mapping into M as $\lambda \rightarrow 0$.*

Proof. This follows from the standard argument for Gromov compactness. (Note that there cannot be any derivative blow up: such blow up would result in a non-constant holomorphic curve in T^*M with boundary in M). \square

We write (q, p) for standard coordinates on T^*M and use the almost complex structure J on T^*M (in some neighborhood of the zero section) induced by a Riemannian metric on M . The following result follows from the maximum principle for harmonic functions.

Lemma A.2. [10, Lemma 5.4] *If $u = (q, p): (\Sigma, \partial\Sigma) \rightarrow (T^*M, Q_\lambda)$ is a J -holomorphic map then $\sup_\Sigma |p| = \mathcal{O}(\lambda)$.* \square

Lemma A.2 shows that we can rescale holomorphic curves u_λ locally by λ^{-1} . Such rescaled map also satisfies a Cauchy-Riemann equation and, provided the derivative of the rescaled map is bounded, it converges to the fiber strip over a local Morse flow line. As explained in [10, Section 5, Corollary 5.24], it is possible to add, given that S meets a finite number of general position conditions, finitely many boundary punctures on the domain of u_λ so that for the punctured domain the rescaled derivative is uniformly bounded. Hence after adding a finite number of punctures the holomorphic map converges (in this rescaled sense as is standard for flow trees, see [10, Section 5.4.4]) to a flow graph. We refer to this rescaled convergence as *flow graph convergence*, which in particular gives C^0 -convergence everywhere (and C^1 -convergence outside neighborhoods of $\mathcal{O}(\lambda \log \lambda)$ -neighborhoods of the graph vertices, see [10, Lemma 5.13]). We state this result on flow graph convergence as a lemma.

Lemma A.3. *Any sequence of holomorphic curves u_λ with boundary on Q_λ has a subsequence that flow graph converges to a flow graph Γ of Q .* \square

Dimension formulas for flow graphs and holomorphic curves agree, and after small perturbation, there is therefore a discrete set of flow graphs that are finite below any fixed Euler characteristic, compare [10, Section 3].

A.2. Flow graphs attached to curves. Suppose now that $M \subset X$ is a Lagrangian, $Q \subset D^*M \subset X$ is contained in a Weinstein neighborhood of M , and we study curves with boundary on Q and possibly also on other Lagrangians $L \subset X \setminus D^*M$. We take $Q_\lambda \subset D^*M$, $\lambda > 0$, as the corresponding family of fiber-scaled Lagrangians.

In this case, as $\lambda \rightarrow 0$, holomorphic curves with boundary on $Q_\lambda \sqcup L$ concentrate near holomorphic curves with boundary on $M \sqcup L$ with flow graphs of $Q \subset T^*M$ attached along their boundaries in M . Such configurations were called quantum flow trees in [11] and were the key to the calculation of the knot contact homology differential for any link. Technical results needed to establish the results stated here are found in [11, Section 5].

We first consider the standard Gromov limit of curves in (X, Q_λ) . We have the following.

Lemma A.4. *Let $u_\lambda: (\Sigma_\lambda, \partial\Sigma_\lambda) \rightarrow (X, Q_\lambda)$ be a family of holomorphic curves. Then some subsequence u_λ Gromov converges to a curve $u_0: (\Sigma_0, \partial\Sigma_0) \rightarrow (X, M)$ as $\lambda \rightarrow 0$.*

Proof. This follows from the standard argument for Gromov compactness. The number of non-zero area bubbles in the limit is controlled by the initial area bound of the curve. For sequences with bounded derivative, the usual Arzela-Ascoli argument gives a sequence that converges to a holomorphic curve with boundary on $L \cup M$. Since the action and flux of S_λ is $\mathcal{O}(\lambda)$ this accounts for all area and we get Gromov convergence. \square

We point out that unlike Gromov compactness in the case when the Lagrangian does not degenerate, in Lemma A.4, constant components of the limit may be located far away from the positive area components of the lifting curve. As before, to get a more precise description, we need to rescale.

We consider this more refined picture of the limit as $\lambda \rightarrow 0$. As in [11, Section 5], we consider the domains $(\Sigma_\lambda, \partial\Sigma_\lambda)$ in a sequence of curves u_λ . Arguing as in [11, Lemma 5.8 and 5.13] we find subsets $(\Sigma_0, \partial\Sigma_0)$ of the domain bounded by extremal short segments where

the map satisfies $\mathcal{O}(\lambda)$ derivative bounds and where $\text{area}(u_\lambda|_{\Sigma_0}) = \mathcal{O}(\lambda)$. As in Lemma A.3 we then find that the restriction $u_\lambda|_{\Sigma_0}$ can be rescaled and converges to flow graphs, and by the area argument in [11, Lemma 5.13] these flow graphs begin and end on the boundary of the positive area part of the limit. We say that such a configuration is a *holomorphic curve with a flow graph attached*. This then leads to the following result.

Lemma A.5. *Any sequence of holomorphic curves $u_\lambda: (\Sigma_\lambda, \partial\Sigma_\lambda) \rightarrow (X, Q_\lambda)$ has a subsequence that flow graph converges to a holomorphic curve $u_M: (\Sigma, \partial\Sigma) \rightarrow (X, M)$ with a flow graph Γ of $Q_\lambda \subset T^*M$ attached along the boundary components in M . If C is embedded and transversely cut out and the flow graph configuration is transversely cut out as well, then for sufficiently small $\lambda > 0$ there is a unique transversely cut out holomorphic curve in N with boundary (piecewise) C^1 -close to the cotangent lift of the flow graph configuration.*

Proof. The convergence statement was discussed above. The last statement follows from results in [11, Section 5.4]. We give a brief description of the argument given there.

The existence and uniqueness of holomorphic curves near flow graph configurations is a consequence of a gluing theorem that contains the two usual ingredients in Floer gluing: establish conditions for Newton iteration and show that the resulting gluing map is surjective. In the case at hand, the domain of the curve constructed is a fat graph attached to the domain of the big curve. For generic configurations there are two local domain models in the graph: strip regions (along flow edges, switches, and ends) and strips with a slit (near trivalent vertices). In all these local models both solutions and their linearizations near the limit can be explicitly determined, see [10, Section 6.1]. There is also a local model for junction points where the graph pieces are attached to the big curve, see [11, Section 5.4.2].

Along these local model domains there are Sobolev spaces of maps with positive exponential weights, pieces are glued and interpolated in bounded regions where the weight function is bounded, see [10, Sections 6.2–3]. This leads to a uniform Fredholm problem and a pregluing that satisfies the necessary quadratic estimate for Floer iteration, see [10, Section 6.4] and [11, Section 5.4.3]. The surjectivity is then established using the fact that along the standard pieces, control of the function norm in the bounded patching regions gives control in the weighted norm over the whole standard piece, see [10, Section 6.4.4] and [11, Section 5.4.3]. \square

Lemma A.5 gives a one-to-one correspondence between holomorphic curves with boundary on Q_λ and holomorphic curves with boundary on M with flow graphs attached provided the holomorphic curves in the flow graphs are transversely cut out and simple (i.e., nowhere multiple covered). We next show that if the big holomorphic curve C in the limit is embedded near its boundary, then flow graph edges of limiting curves that are attached to C cannot be multiply covered.

Lemma A.6. *Let $u_\lambda: (\Sigma_\lambda, \partial\Sigma_\lambda) \rightarrow (X, Q_\lambda)$ be a sequence that flow graph converges to a holomorphic curve $u_M: (\Sigma, \partial\Sigma) \rightarrow (X, M)$ with a flow graph Γ of $Q_\lambda \subset T^*M$ attached along the boundary components in M with all attaching points distinct and assume that u_M is an embedding in a neighborhood of $\partial\Sigma$. Then any edge of Γ is non-multiply covered.*

Proof. We first consider an edge in Γ attached to u_M . Consider a neighborhood $(\mathbb{C}^3, \mathbb{R}^3)$, where \mathbb{R}^3 corresponds to L of the point where the flow line is attached. We take coordinates so that the big curve lies in the first coordinate line. We take one of the sheets of Q to be \mathbb{R}^3 and the other to be the graph of λdx_2 . (Note that the flow tree rescaling, multiplication by

λ^{-1} , takes the holomorphic curve near the flow graph configuration to a standard holomorphic curve, with boundary on \mathbb{R}^3 and the graph of dx_2 , up to an error term of size $\mathcal{O}(\lambda)$, see [10, Lemma 5.12].)

Projecting the curve with flow line attached to the complex line spanned by $\partial_{x_1} + \partial_{x_2}$ we find that the curve projects to the upper half-plane with boundary on the real line and a slit along the line with imaginary part $\lambda > 0$. The assumption that u_M is an embedding near the boundary implies that the degree of the half plane equals 1 above the line with imaginary part λ . If the edge attached was a multiple cover, then the degree below the slit would be > 1 . This is not possible and hence the edge cannot be multiply covered.

We next show that no other edge of Γ can be multiply covered. Note first that it is clear that the multiplicity of a curve is constant along an edge in Γ . We check that it remains constant over vertices. As above, near a vertex there is a standard model for the projected boundary condition after rescaling, sheets without cusp edges correspond to lines in \mathbb{C} parallel to the real axis, a those with cusp edge correspond to a half circles with lines attached. In the limit, the projected curve is holomorphic, possibly with interior and boundary branch points, and it is straightforward to check that edge multiplicity does not change. Since edges of Γ attached to the big curve u_M have multiplicity one, it follows that all edges of Γ have multiplicity one. \square

REFERENCES

- [1] Alberto Abbondandolo and Matthias Schwarz. On the Floer homology of cotangent bundles. *Communications on Pure and Applied Mathematics*, 59(2):254–316, 2006.
- [2] Anna Aiston. Adams operators and knot decorations. *arXiv:q-alg/9711015*.
- [3] Dylan Allegretti. Voros symbols as cluster coordinates. *Journal of Topology*, 12(4):1031–1068, 2019.
- [4] Philip Argyres and Michael Douglas. New phenomena in $SU(3)$ supersymmetric gauge theory. *Nucl. Phys. B*, 448:93–126, 1995.
- [5] Léa Bittmann, Alex Chandler, Anton Mellit, and Chiara Novarini. Type A DAHA and doubly periodic tableaux. *Advances in Mathematics*, 416:108919, 2023.
- [6] Francis Bonahon and Helen Wong. Quantum traces for representations of surface groups in $SL_2(\mathbb{C})$. *Geometry & Topology*, 15(3):1569–1615, 2011.
- [7] Roger Casals and Yoon Jae Nho. Spectral Networks and Betti Lagrangians. *arXiv:2504.08144*.
- [8] Sergio Cecotti, Clay Cordova, and Cumrun Vafa. Braids, walls, and mirrors. *arXiv:1110.2115*.
- [9] Tudor Dimofte, Sergei Gukov, and Yan Soibelman. Quantum Wall Crossing in $N=2$ Gauge Theories. *Lett. Math. Phys.*, 95:1–25, 2011.
- [10] Tobias Ekholm. Morse flow trees and Legendrian contact homology in 1-jet spaces. *Geom. Topol.*, 11:1083–1224, 2007.
- [11] Tobias Ekholm, John Etnyre, Lenhard Ng, and Michael Sullivan. Knot contact homology. *Geometry & Topology*, 17(2):975–1112, 2013.
- [12] Tobias Ekholm, John Etnyre, and Joshua Sabloff. A duality exact sequence for Legendrian contact homology. *Duke Mathematical Journal*, 150(1):1–75, 2009.
- [13] Tobias Ekholm, Pietro Longhi, and Vivek Shende. The skein valued mirror of the topological vertex. *arXiv:2412.15454*.
- [14] Tobias Ekholm, Lenhard Ng, and Vivek Shende. A complete knot invariant from contact homology. *Inventiones mathematicae*, pages 1–52, 2018.
- [15] Tobias Ekholm and Vivek Shende. Colored HOMFLYPT counts holomorphic curves. *arXiv:2101.00619*.
- [16] Tobias Ekholm and Vivek Shende. Counting bare curves. *arXiv preprint arXiv:2406.00890*.
- [17] Tobias Ekholm and Vivek Shende. Ghost bubble censorship. *arXiv:2212.05835*. To appear in *Communications in Analysis and Geometry*.
- [18] Tobias Ekholm and Vivek Shende. Skein recursion for holomorphic curves and invariants of the unknot. *arXiv:2012.15366*.

- [19] Tobias Ekholm and Vivek Shende. Skeins on branes. *arXiv preprint arXiv:1901.08027*.
- [20] Ludvig Faddeev and Rinat Kashaev. Quantum Dilogarithm. *Mod. Phys. Lett. A*, 9:427–434, 1994.
- [21] Andreas Floer. Witten’s complex and infinite-dimensional Morse theory. *J. Differential Geom.*, 30(1):207–221, 1989.
- [22] Vladimir Fock. Dual Teichmüller spaces. *dg-ga/9702018*.
- [23] Vladimir Fock and Alexander Goncharov. Moduli spaces of local systems and higher Teichmüller theory. *Publications Mathématiques de l’IHÉS*, 103:1–211, 2006.
- [24] Vladimir Fock and Alexander Goncharov. Cluster ensembles, quantization and the dilogarithm. *Annales scientifiques de l’École normale supérieure*, 42(6):865–930, 2009.
- [25] Daniel Freed and Andrew Neitzke. 3d spectral networks and classical Chern-Simons theory. In *Surveys in differential geometry 2021. Chern: a great geometer of the 20th century*, volume 26 of *Surv. Differ. Geom.*, pages 51–155. Int. Press, 2024.
- [26] Kenji Fukaya and Yong-Geun Oh. Zero-loop open strings in the cotangent bundle and Morse homotopy. *Asian J. Math.*, 1(1):96–180, 1997.
- [27] Maxime Gabella. Quantum Holonomies from Spectral Networks and Framed BPS States. *Commun. Math. Phys.*, 351(2):563–598, 2017.
- [28] Davide Gaiotto, Gregory Moore, and Andrew Neitzke. Four-dimensional wall-crossing via three-dimensional field theory. *Communications in Mathematical Physics*, 299(1):163–224, 2010.
- [29] Davide Gaiotto, Gregory Moore, and Andrew Neitzke. Spectral networks. *Annales Henri Poincaré*, 14(7):1643–1731, 2013.
- [30] Davide Gaiotto, Gregory Moore, and Andrew Neitzke. Wall-crossing, Hitchin systems, and the WKB approximation. *Advances in Mathematics*, 234:239–403, 2013.
- [31] Dmitry Galakhov, Pietro Longhi, and Gregory Moore. Spectral Networks with Spin. *Commun. Math. Phys.*, 340(1):171–232, 2015.
- [32] Stavros Garoufalidis and Tao Yu. A quantum trace map for 3-manifolds. *arXiv:2403.12424*.
- [33] Adriano Garsia and Anton Mellit. Five-term relation and Macdonald polynomials. *Journal of Combinatorial Theory, Series A*, 163:182–194, 2019.
- [34] Richard Hadji and Hugh Morton. A basis for the full Homfly skein of the annulus. *Math. Proc. Cambridge Philos. Soc.*, 141(1):81–100, 2006.
- [35] Mingyuan Hu. A proof of the pentagon relation for skeins. *2401.10817*.
- [36] Mingyuan Hu, Gus Schrader, and Eric Zaslow. Skeins, clusters and wavefunctions. *arXiv:2312.10186*.
- [37] Michael Hutchings and Yi-Jen Lee. Circle-valued Morse theory and Reidemeister torsion. *Geometry and Topology*, 3:369–396, 1999.
- [38] Kohei Iwaki and Tomoki Nakanishi. Exact WKB analysis and cluster algebras. *Journal of Physics A: Mathematical and Theoretical*, 47(47):474009, 2014.
- [39] Maxim Kontsevich and Yan Soibelman. Stability structures, motivic Donaldson-Thomas invariants and cluster transformations. *arXiv:0811.2435*.
- [40] Tatsuki Kuwagaki. Sheaf quantization from exact WKB analysis. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2024(814):165–203, 2024.
- [41] Marc Lackenby. Taut ideal triangulations of 3-manifolds. *Geometry & Topology*, 4(1):369–395, 2000.
- [42] Thang Lê. Quantum Teichmüller spaces and quantum trace map. *Journal of the Institute of Mathematics of Jussieu*, 18(2):249–291, 2019.
- [43] Feng Luo and Stephan Tillmann. Angle structures and normal surfaces. *Trans. Amer. Math. Soc.*, 360(6):2849–2866, 2008.
- [44] Hugh Morton. Power sums and Homfly skein theory. In *Invariants of knots and 3-manifolds (Kyoto, 2001)*, volume 4 of *Geom. Topol. Monogr.*, pages 235–244. Geom. Topol. Publ., 2002.
- [45] Hugh Morton and Pedro María González Manchón. Geometrical relations and plethysms in the Homfly skein of the annulus. *Journal of the London Mathematical Society*, 78(2):305–328, 2008.
- [46] Hugh Morton and Peter Samuelson. The HOMFLYPT skein algebra of the torus and the elliptic Hall algebra. *Duke Mathematical Journal*, 166(5):801–854, 2017.
- [47] Hugh Morton and Peter Samuelson. DAHAs and skein theory. *Communications in Mathematical Physics*, 385(3):1655–1693, 2021.
- [48] Lukas Nakamura. Recursion relations and BPS-expansions in the HOMFLY-PT skein of the solid torus. *arXiv:2401.10730*.

- [49] Andrew Neitzke and Fei Yan. q-nonabelianization for line defects. *Journal of High Energy Physics*, 2020(9):1–67, 2020.
- [50] Andrew Neitzke and Fei Yan. The quantum UV-IR map for line defects in $\mathfrak{gl}(3)$ -type class S theories. *Journal of High Energy Physics (Online)*, 2022(9), 2022.
- [51] Yoon Jae Nho. Family Floer theory, non-abelianization, and spectral networks. *arXiv:2307.04213*.
- [52] Hiroshi Ooguri and Cumrun Vafa. Knot invariants and topological strings. *Nucl.Phys.*, B577:419–438, 2000.
- [53] Samuel Panitch and Sunghyuk Park. 3d quantum trace map. *arXiv:2403.12850*. *To appear in Algebraic & Geometric Topology*.
- [54] Samuel Panitch and Sunghyuk Park. Compatibility of quantum trace and UV-IR maps. *arXiv:2509.09100*.
- [55] Robert Penner. The decorated Teichmüller space of punctured surfaces. *Communications in Mathematical Physics*, 113:299–339, 1987.
- [56] Marcin Poźniak. Floer homology, Novikov rings and clean intersections. In *Northern California Symplectic Geometry Seminar*, volume 196 of *Amer. Math. Soc. Transl. Ser. 2*, pages 119–181. Amer. Math. Soc., Providence, RI, 1999.
- [57] Markus Reineke. The Harder-Narasimhan system in quantum groups and cohomology of quiver moduli. *Inventiones mathematicae*, 152(2):349–368, 2003.
- [58] Matthias Scharitzer. A skein-valued lift of the unidirectional An-quiver q-dilogarithm relations. *2410.17773*.
- [59] Matthias Scharitzer and Vivek Shende. Quantum mirrors of cubic planar graph Legendrians. *arXiv:2304.01872*.
- [60] Matthias Scharitzer and Vivek Shende. Skein valued cluster transformation in enumerative geometry of Legendrian mutation. *arXiv:2312.10625*.
- [61] Gus Schrader, Linhui Shen, and Eric Zaslow. The chromatic Lagrangian: wavefunctions and open Gromov-Witten conjectures. *Adv. Theor. Math. Phys.*, 28(6):1781–1879, 2024.
- [62] Nathan Seiberg and Edward Witten. Electric - magnetic duality, monopole condensation, and confinement in $N=2$ supersymmetric Yang-Mills theory. *Nucl. Phys. B*, 426:19–52, 1994. [Erratum: *Nucl.Phys.B* 430, 485–486 (1994)].
- [63] Vivek Shende, David Treumann, Harold Williams, and Eric Zaslow. Cluster varieties from Legendrian knots. *Duke Mathematical Journal*, 168(15), 2019.
- [64] Erich Stiemke. Über positive Lösungen homogener linearer Gleichungen. *Math. Ann.*, 76(2-3):340–342, 1915.
- [65] Kurt Strebel. *Quadratic differentials*. Springer, 1984.
- [66] David Treumann and Eric Zaslow. Cubic planar graphs and Legendrian surface theory. *Adv. Theor. Math. Phys.*, 22(5):1289–1345, 2018.
- [67] Vladimir Turaev. Skein quantization of Poisson algebras of loops on surfaces. *Ann. Sci. École Norm. Sup. (4)*, 24(6):635–704, 1991.
- [68] Yegor Zenkevich. On pentagon identity in Ding-Iohara-Miki algebra. *Journal of High Energy Physics*, 2023(3):1–14, 2023.

DEPARTMENT OF MATHEMATICS, UPPSALA UNIVERSITY, BOX 480, 751 06 UPPSALA, SWEDEN AND
INSTITUT MITTAG-LEFFLER, AURAV 17, 182 60 DJURSHOLM, SWEDEN

Email address: tobias.ekholm@math.uu.se

DEPARTMENT OF PHYSICS AND ASTRONOMY AND DEPARTMENT OF MATHEMATICS, UPPSALA UNIVERSITY,
BOX 516, 751 20 UPPSALA, SWEDEN

Email address: pietro.longhi@physics.uu.se

DEPARTMENT OF MATHEMATICS AND CENTER OF MATHEMATICAL SCIENCES AND APPLICATIONS, HARVARD
UNIVERSITY, CAMBRIDGE, MA 02138, USA

Email address: sunghyukpark@math.harvard.edu

CENTER FOR QUANTUM MATHEMATICS, SYDDANSK UNIV., CAMPUSVEJ 55 5230 ODENSE DENMARK
AND DEPARTMENT OF MATHEMATICS, UC BERKELEY, 970 EVANS HALL, BERKELEY CA 94720 USA

Email address: vivek.vijay.shende@gmail.com