

FUKAYA CATEGORIES OF SURFACES, SPHERICAL OBJECTS, AND MAPPING CLASS GROUPS

DENIS AUROUX AND IVAN SMITH

ABSTRACT. We prove that every spherical object in the derived Fukaya category of a closed surface of genus at least two whose Chern character represents a primitive Hochschild homology class is quasi-isomorphic to a simple closed curve equipped with a rank one local system. (The homological hypothesis is necessary.) This largely answers a question of Haiden, Katzarkov and Kontsevich. It follows that there is a natural surjection from the autoequivalence group of the Fukaya category to the mapping class group. The proofs appeal to and illustrate numerous recent developments: quiver algebra models for wrapped categories, sheafifying the Fukaya category, equivariant Floer theory for finite and continuous group actions, and homological mirror symmetry. An application to high-dimensional symplectic mapping class groups is included.

1. INTRODUCTION

The mapping class group Γ_g of a closed surface Σ_g arises naturally in different contexts: in algebra as the outer automorphism group $\text{Out}(\pi_1 \Sigma_g)$, in topology as the component group $\pi_0 \text{Diff}^+(\Sigma_g)$, in algebraic geometry as the orbifold fundamental group $\pi_1^{\text{orb}}(\mathcal{M}_g)$ of the moduli space of curves. In Floer theory, and mirror symmetry, a symplectic manifold Z appears through the Fukaya category $\mathcal{F}(Z)$ and its derived category $D^\pi \mathcal{F}(Z)$, and the natural symmetry group of a surface in that context is the group of autoequivalences $\text{Auteq}(D^\pi \mathcal{F}(\Sigma_g))$. This comes with a map $\Gamma_g \rightarrow \text{Auteq}(D^\pi \mathcal{F}(\Sigma))$ (which depends on additional choices), which has no obvious intrinsic categorical or Floer-theoretic description; in many cases, we know of autoequivalences of Fukaya categories which are not geometric, cf. [7]. This paper shows that the mapping class group arises naturally from the Fukaya category.

Let (Σ_g, ω) denote a closed surface of genus $g \geq 2$, equipped with an area form of area 1. Let $\mathcal{F}(\Sigma_g)$ denote the Fukaya category of Σ_g , which is a $\mathbb{Z}/2$ -graded A_∞ -category, linear over the one-variable Novikov field $\Lambda = \mathbb{C}((q^{\mathbb{R}}))$. Objects of the category are immersed unobstructed closed curves, equipped with a finite rank local system and auxiliary brane data (including a choice of spin structure on the underlying curve). We will denote by (ξ, γ) the object associated to an immersed closed curve $\gamma : S^1 \rightarrow \Sigma_g$ and local system $\xi \rightarrow S^1$ on the domain of γ .

We denote by $D^\pi \mathcal{F}(\Sigma_g) = \mathcal{F}(\Sigma)^{\text{per}}$ the category of perfect modules over $\mathcal{F}(\Sigma_g)$, equivalently the split-closure of twisted complexes, which is triangulated in the classical sense; write \simeq for quasi-isomorphism in this category. The composite of the Chern character and the open-closed map defines a class

$$ch(A) \in HH_0(D^\pi \mathcal{F}(\Sigma), D^\pi \mathcal{F}(\Sigma)) \cong H_1(\Sigma; \Lambda)$$

for any object $A \in D^\pi \mathcal{F}(\Sigma)$. Recall that an object $A \in D^\pi \mathcal{F}(\Sigma_g)$ is *spherical* if

$$\mathrm{Hom}_{D^\pi \mathcal{F}(\Sigma_g)}(A, A) \cong H^*(S^1; \Lambda).$$

Lemma 2.19 shows that $ch(X)$ is an integral class when X is spherical. Our main result is the following geometricity theorem for spherical objects.

Theorem 1.1. *If $X \in D^\pi \mathcal{F}(\Sigma_g)$ is spherical and $ch(X)$ is primitive, then there is a simple closed curve $\gamma \subset \Sigma_g$ and a rank one local system $\xi \rightarrow \gamma$ with $X \simeq (\xi, \gamma)$.*

En route, we prove the corresponding result for surfaces with non-empty boundary (Corollary 4.14). When $g = 1$, Theorem 1.1 is a consequence of homological mirror symmetry for elliptic curves and Atiyah's classification of bundles on such curves. When $g > 1$, there are spherical objects with vanishing Chern character which are *not* quasi-isomorphic to any simple closed curve with local system, so the result is in some sense sharp; see Lemma 2.23. Theorem 1.1 largely answers [19, Problem 2].

Let $\Gamma(\Sigma_g)$ denote the symplectic mapping class group of Σ_g . There is a homomorphism $\Gamma(\Sigma_g) \rightarrow \mathrm{Auteq}(D^\pi \mathcal{F}(\Sigma_g))$ (this depends on choices, cf. Section 2.6).

Corollary 1.2. *There is a natural surjective homomorphism $\mathrm{Auteq}(D^\pi \mathcal{F}(\Sigma_g)) \rightarrow \Gamma(\Sigma_g)$.*

The same conclusion also holds for surfaces of genus ≥ 1 with boundary, by combining Corollary 4.14 with the argument of Proposition 7.7.

It follows that the map $\Gamma(\Sigma_g) \rightarrow \mathrm{Auteq}(D^\pi \mathcal{F}(\Sigma_g))$ splits. The existence of such a split homomorphism has cohomological implications; for instance, the autoequivalence group has infinite-dimensional second bounded cohomology, and admits families of unbounded quasimorphisms.

We conjecture that, for $g \geq 2$ (so the flux group is trivial), the kernel of the natural map from autoequivalences to the mapping class group is generated by tensoring by flat unitary line bundles and the actions of symplectomorphisms of non-trivial flux, i.e. that

$$\mathrm{Auteq}(D^\pi \mathcal{F}(\Sigma)) \stackrel{?}{=} H^1(\Sigma; \Lambda^*) \rtimes \Gamma(\Sigma_g),$$

where the map $\Gamma(\Sigma_g) \rightarrow \mathrm{Auteq}(D^\pi \mathcal{F}(\Sigma_g))$ is only well-defined up to the action of $H^1(\Sigma; \mathbb{R})$. In Section 7.2 we outline an argument suggesting that the subgroup $H^1(\Sigma; \Lambda^*)$ is normal, which is consistent with the speculation.

The proof of Theorem 1.1 is surprisingly involved, and breaks into the following steps (the main text treats these in somewhat different order). Let $X \in D^\pi \mathcal{F}(\Sigma)$ be a spherical object.

- (1) The open-closed image $ch(X) \in HH_0(\mathcal{F}(\Sigma)) = H_1(\Sigma; \Lambda)$ defines an integral class, i.e. lies in the image of $H_1(\Sigma; \mathbb{Z})$. *Assume henceforth this class is primitive, and in particular non-zero.*
- (2) A non-zero integral class $a \in H^1(\Sigma; \mathbb{Z})$ defines a \mathbb{G}_m -action on $\mathcal{F}(\Sigma)$.
- (3) If $\langle a, ch(X) \rangle = 0$, then X defines a \mathbb{G}_m -equivariant object, hence a \mathbb{Z}/N -equivariant object for any finite $\mathbb{Z}/N \leq \mathbb{G}_m$. The \mathbb{Z}/N -equivariant Fukaya category of Σ is the Fukaya category of an N -fold cover $\hat{\Sigma}$ of Σ .

- (4) A choice of equivariant structure on X defines a lift \hat{X} of X to $\hat{\Sigma}$, and for large enough N , there are disjoint homologically independent simple closed curves $\gamma_1, \gamma_2 \subset \hat{\Sigma}$ with $H^*(\text{hom}_{\mathcal{F}(\hat{\Sigma})}(\hat{X}, \gamma_i)) = 0$.
- (5) There are annular neighbourhoods $\gamma_i \subset A_i \subset \hat{\Sigma}$ for which \hat{X} lifts to define a spherical object (which we still call) \hat{X} in the wrapped category $\mathcal{W}(C)^{per}$ of $C = \hat{\Sigma} \setminus (A_1 \cup A_2)$.
- (6) \hat{X} is represented by a strictly unobstructed immersed closed curve $\sigma \subset C$ equipped with a unitary local system.
- (7) If it is not embedded, the immersed curve σ supporting \hat{X} bounds an embedded bigon.
- (8) Bigons on σ may be “emptied” and then “cancelled”, so \hat{X} is quasi-isomorphic to a simple closed curve with rank one local system in $\mathcal{F}(C)$.
- (9) X is quasi-isomorphic to a simple closed curve with rank one local system in $\mathcal{F}(\Sigma)$.

Steps (2)-(4) rely on work of Seidel on equivariant Floer theory [30, 32], and ideas of family Floer theory à la Abouzaid and Fukaya. Step (5) relies on H. Lee’s restriction technology [24] for sheafifying wrapped categories of surfaces. Step (6) appeals to the work of Haiden, Katzarkov and Kontsevich [19] on quiver algebra models for wrapped categories of punctured surfaces, and to a split-closure result for such wrapped categories which we infer from homological mirror symmetry and a K-theoretic characterisation of split-closure for derived categories of singularities due to Abouzaid, Auroux, Efimov, Katzarkov and Orlov [4]. Step (7) invokes delicate classical results of Steinitz [39, 14] and Hass and Scott [21] in surface topology. The “cancelling” move for bigons in Step (8) uses the homological hypothesis from Step (1). Corollary 1.2 follows by considering the action of autoequivalences of $\mathcal{F}(\Sigma_g)$ on a Floer-theoretic “Schmutz graph” of non-separating curves [27].

As an application of Corollary 1.2, in Section 8 we prove:

Theorem 1.3. *There is a smooth manifold Z with symplectic forms ω_δ , $\delta \in (0, 1]$, for which $\pi_0 \text{Symp}(Z, \omega_\delta)$ surjects to a free group of rank $N(\delta)$ where $N(\delta) \rightarrow \infty$ as $\delta \rightarrow 0$.*

Explicitly, Z is the product of Σ_2 with the blow-up of the four-torus at a point, but equipped with an irrational perturbation (of size δ) of the standard Kähler form. The free group quotients arise from subgroups of the genus two Torelli group. Gromov [18], Abreu-McDuff [8] and others showed that the topology of the symplectomorphism group can vary rather wildly as one continuously deforms the symplectic form, but this seems to be the first example in which a fixed degree homotopy group is known to have unbounded rank.

Notation: Throughout the paper, surfaces are connected, and (immersed) curves are assumed to be homotopically non-trivial unless explicitly stated otherwise.

Acknowledgements. Thanks to Paul Seidel for suggesting we use equivariance and for several related explanations; and to Mohammed Abouzaid, Fabian Haiden, Yankı Lekili, Dhruv Ranganathan and Henry Wilton for helpful conversations. DA is partially supported by NSF grant DMS-1937869 and by Simons Foundation grant #385573 (Simons Collaboration on Homological Mirror Symmetry). IS is partially supported by Fellowship EP/N01815X/1 from the Engineering and Physical Sciences Research Council, U.K.

2. THE FUKAYA CATEGORY OF A SURFACE

2.1. Background. Fix a coefficient field k . Let $\mathbb{K} = \Lambda_k$ denote the single-variable Novikov field, with formal variable q , of formal series $\sum_i a_i q^{t_i}$ with $a_i \in k$, $t_i \in \mathbb{R}$ and $\lim t_i = +\infty$. The valuation map

$$val : \Lambda \rightarrow \mathbb{R} \cup \{\infty\}, \quad val(0) = +\infty$$

associates to a non-zero element of Λ its smallest q -power. The subring $\Lambda_{\geq 0} = val^{-1}[0, \infty]$ comprises series with $t_i \geq 0$ for every i , and there is a homomorphism

$$\eta : \Lambda_{\geq 0} \rightarrow k$$

which extracts the constant coefficient. The kernel of this homomorphism is denoted $\Lambda_{>0}$. The unitary subgroup $U_\Lambda = val^{-1}(0)$ is the subgroup of elements $a + \sum_{t_i > 0} a_i q^{t_i}$ where $a \in k^*$ is non-zero. The field $\Lambda_{\mathbb{C}}$ is algebraically closed of characteristic zero.

A non-unital A_∞ -category \mathcal{A} over the field \mathbb{K} comprises: a set of objects $\text{Ob } \mathcal{A}$; for each $X_0, X_1 \in \text{Ob } \mathcal{A}$ a $\mathbb{Z}/2$ -graded \mathbb{K} -vector space $\text{hom}_{\mathcal{A}}(X_0, X_1)$; and \mathbb{K} -linear composition maps, for $k \geq 1$,

$$\mu_{\mathcal{A}}^k : \text{hom}_{\mathcal{A}}(X_{k-1}, X_k) \otimes \cdots \otimes \text{hom}_{\mathcal{A}}(X_0, X_1) \longrightarrow \text{hom}_{\mathcal{A}}(X_0, X_k)[2 - k]$$

of degree $2 - k$ (here $[j]$ denotes downward shift by $j \in \mathbb{Z}/2$, and all degrees are mod 2 degrees; we write $2 - k$ rather than $-k$ since \mathbb{Z} -graded categories may be more familiar). The maps $\{\mu^k\}$ satisfy a hierarchy of quadratic equations

$$\sum_{m,n} (-1)^{\mathfrak{X}_n} \mu_{\mathcal{A}}^{k-m+1}(a_k, \dots, a_{n+m+1}, \mu_{\mathcal{A}}(a_{n+m}, \dots, a_{n+1}), a_n, \dots, a_1) = 0$$

with $\mathfrak{X}_n = \sum_{j=1}^n |a_j| - n$ and where the sum runs over all possible compositions: $1 \leq m \leq k$, $0 \leq n \leq k - m$. If the characteristic of \mathbb{K} is not equal to 2, the signs in the A_∞ -associativity equations depend on the mod 2 degrees of generators, and the existence of a $\mathbb{Z}/2$ -grading on \mathcal{A} is essential.

We denote by $\text{Tw}^\pi(\mathcal{A}) = \mathcal{A}^{per}$ the idempotent completion of the category of twisted complexes $\text{Tw}(\mathcal{A})$ [28]. If the smallest split-closed triangulated A_∞ -category containing a subcategory $\mathcal{A}' \subset \mathcal{A}$ is \mathcal{A}^{per} , then we will say that \mathcal{A}' split-generates \mathcal{A} .

In a curved A_∞ -category each object A comes with a curvature $\mu^0 \in \text{hom}_{\mathcal{A}}^{ev}(A, A)$, and the quadratic associativity relations are modified to take account of μ^0 . In general the resulting sums could *a priori* be infinite, so some geometric conditions are required to ensure convergence. The following result controls the curvature of mapping cones on closed morphisms.

Lemma 2.1. *If \mathcal{A} is a curved A_∞ -category and $X_1 \xrightarrow{f} X_2$ is a mapping cone in $\text{Tw}(\mathcal{A})$ between objects X_i with $\mu^0(X_i) = 0$, then*

$$\mu_{\mathcal{A}}^0(X_1 \xrightarrow{f} X_2) = \mu_{\mathcal{A}}^1(f).$$

Proof. Standard. □

The Hochschild cohomology $HH^*(\mathcal{A}, \mathcal{A}) = \text{Hom}_{(\mathcal{A}\text{-mod-}\mathcal{A})}(\text{id}_{\mathcal{A}}, \text{id}_{\mathcal{A}})$ of an (uncurved) A_∞ -category is the endomorphisms of the diagonal bimodule. Hochschild cohomology of an A_∞ -category is invariant under passing to the derived category, so we will not distinguish notationally between $HH^*(\mathcal{A}, \mathcal{A})$ and $HH^*(D^\pi \mathcal{A}, D^\pi \mathcal{A})$.

2.2. Immersed curves. Let Σ be a closed oriented surface of genus $g(\Sigma) = g \geq 2$ equipped with an area form ω of unit total area, $\int_\Sigma \omega = 1$.

The Fukaya category $\mathcal{F}(\Sigma)$ has objects which are called Lagrangian branes, and which comprise an immersed curve $\iota : S^1 \rightarrow \Sigma$ with the following properties and additional data. First note that if $\gamma = \iota(S^1)$ is immersed, it may bound “teardrop discs”, i.e. images of holomorphic discs with a unique boundary puncture where the puncture is mapped to a self-intersection point of γ . Let $\mu^0(\gamma) \in CF^0(\gamma, \gamma)$ be the algebraic sum of all such discs, which is the obstruction term in the self-Floer complex. We insist:

- $\iota(S^1) = \gamma$ is unobstructed, meaning that $\mu^0(\gamma) \in CF^0(\gamma, \gamma)$ vanishes;
- all self-intersections of γ are transverse.
- the domain S^1 is equipped with a flat unitary local system of Λ -vector spaces;
- γ is equipped with a *spin* structure.

Note that the first condition is that $\mu^0(\gamma)$ vanishes identically (i.e. at each self-intersection point). We could equally well insist that objects of $\mathcal{F}(\Sigma)$ are tautologically unobstructed in the sense of bounding no teardrop discs; in fact the two conditions are equivalent:

Lemma 2.2. *If γ is a homotopically non-trivial immersed curve, then the following are equivalent: (1) $\mu^0(\gamma) = 0$; (2) γ does not bound any teardrop discs; (3) γ lifts to a properly embedded arc in the universal cover of Σ .*

Proof. Since γ is homotopically non-trivial, it lifts to a properly immersed arc $\tilde{\gamma}$ in the universal cover $\tilde{\Sigma}$. Teardrop discs bounded by γ lift to teardrop discs in $\tilde{\Sigma}$ with boundary on $\tilde{\gamma}$; thus (3) implies (2), and (2) implies (1).

Conversely, assume that $\tilde{\gamma}$ is not embedded, and let $a \in \tilde{\Sigma}$ be any self-intersection of $\tilde{\gamma}$. If the arc $\tilde{\gamma}_a$ consisting of the portion of $\tilde{\gamma}$ which connects a to itself is not embedded, then there exists another “nested” self-intersection b such that the arc $\tilde{\gamma}_b$ connecting b to itself is a strict subset of $\tilde{\gamma}_a$. Considering b instead of a , and repeating the process if needed, we can assume that $\tilde{\gamma}_a$ is embedded in $\tilde{\Sigma}$; the portion D_a of $\tilde{\Sigma}$ enclosed by $\tilde{\gamma}_a$ is then an embedded teardrop disc with a corner at a .

The teardrop D_a may be either locally convex or locally concave near its corner, meaning it occupies either 1 or 3 of the quadrants delimited by the two branches of $\tilde{\gamma}$ intersecting at a . Recall that only locally convex teardrops contribute to $\mu^0(\gamma)$. We claim that, if D_a has a locally concave corner at a , then there are smaller embedded teardrop discs contained inside it. Indeed, in the locally concave case, the portions of $\tilde{\gamma}$ just before and after the arc $\tilde{\gamma}_a$ lie inside D_a . Continuing along $\tilde{\gamma}$ until it exits D_a (which must eventually happen since $\tilde{\gamma}$ is properly immersed), we find another self-intersection b lying on the boundary of D_a , such that the arc $\tilde{\gamma}_b$ connecting b to itself is entirely contained in D_a (and is not the entire boundary of D_a). If $\tilde{\gamma}_b$ is not embedded, then we replace b by a nested self-intersection as above; this allows us to assume that $\tilde{\gamma}_b$ is embedded, while still ensuring that $\tilde{\gamma}_b$ is entirely

contained in D_a . The region of $\tilde{\Sigma}$ bounded by $\tilde{\gamma}_b$ then gives a new embedded teardrop disc D_b which is strictly contained in D_a .

Repeating the process, we conclude that if $\tilde{\gamma}$ is not embedded then it must bound an embedded teardrop disc with a locally convex corner. The generator of $CF^0(\gamma, \gamma)$ corresponding to this self-intersection must then have a non-zero coefficient in $\mu^0(\gamma)$, since there are no other teardrops with the same corner. Thus, if $\mu^0(\gamma) = 0$ then $\tilde{\gamma}$ must be embedded. \square

Remark 2.3. *Regular homotopies of immersed curves do not preserve the absence of teardrops or the vanishing of μ^0 . For instance, pushing an embedded arc sideways then back through itself to create a pair of self-intersections gives rise to a teardrop (and a bigon). Thus, in the arguments below we will take care to only consider regular homotopies which preserve unobstructedness.*

Remark 2.4. *On a two-dimensional surface, rigid J -holomorphic discs with pairwise distinct boundary conditions are immersed polygons with convex corners, which are purely combinatorial. One can set up the Fukaya category either via moving Lagrangian boundary conditions and honest J -holomorphic curves or via (the more usual approach with) Hamiltonian perturbation terms in the Floer equation. See [29] for an implementation.*

Remark 2.5. *The results of this section apply, mutatis mutandis, to the case of a surface S with non-empty boundary. The main difference is that it is sometimes necessary to consider immersed curves which live in the (infinite area) completion of S , rather than in S itself. We will pay careful attention to this issue when it arises in the sequel.*

2.3. Isotopies and twists. For each $a \in H^1(\Sigma; \mathbb{R})$ there is a symplectomorphism ϕ_a with flux a , and one can move Lagrangian submanifolds by such symplectomorphisms. This obviously preserves unobstructedness, since it preserves all teardrop discs and their areas. There is a more general statement for isotopies not induced by global symplectomorphisms.

Lemma 2.6. *Let $\gamma \subset \Sigma$ be an immersed curve and σ a simple closed curve. One can isotope γ by a regular homotopy so as to meet σ minimally. Moreover, the regular homotopy can be chosen to preserve the unobstructedness of γ .*

Proof. Suppose the intersection is not minimal, so there is a not necessarily embedded bigon bound by $\sigma \cup \gamma$. Pull back σ to the domain of such a bigon; changing the choice of bigon to one that is “innermost” if necessary, we can assume that there is a bigon H which does not meet σ in its interior. The boundary of this bigon therefore contains a proper subset of γ , since the boundary arc of the bigon lying along γ connects two intersections in $\gamma \cap \sigma$ which are consecutive as read along γ . (Note that the σ -edge of this bigon might loop fully around σ even with the bigon being innermost.) Now push the arc of γ lying along H across H , to decrease the total intersection number of γ and σ . This can be iterated to reduce to a minimal situation.

By Lemma 2.2, γ is unobstructed if and only if its lift $\tilde{\gamma}$ to the universal cover of Σ is embedded. Lifting to $\tilde{\Sigma}$ the bigon H along which we slide γ in the above argument, we obtain a bigon \tilde{H} with boundary on $\tilde{\gamma}$ and on a lift $\tilde{\sigma}$ of σ , whose interior is moreover disjoint from all lifts of σ . If the interior of \tilde{H} is disjoint from $\tilde{\gamma}$ then sliding γ across H preserves unobstructedness. Otherwise, by pulling back $\tilde{\gamma}$ to the domain of the bigon \tilde{H}

(which yields a disjoint collection of embedded arcs, since $\tilde{\gamma}$ is embedded) and changing the choice of bigon to one that is innermost, we obtain a bigon with boundary on $\tilde{\gamma} \cup \tilde{\sigma}$ whose interior is disjoint from $\tilde{\gamma}$ and from all lifts of σ . Projecting back to Σ and sliding across this bigon decreases the intersection number without affecting unobstructedness. \square

Remark 2.7. *This argument also gives the following generalization of Lemma 2.6: let S be a compact surface with boundary, σ a simple closed curve, and γ either an immersed curve or an immersed arc with ends in ∂S . Then one can isotope γ by a regular homotopy so as to meet σ minimally, in a way that preserves the absence of teardrops with boundary on γ .*

Lemma 2.8. *For any immersed curve $\gamma \subset \Sigma$ with $\mu^0(\gamma) = 0$, and $a \in \mathbb{R}$, there is a regular isotopy γ_t of γ through (generically) immersed unobstructed curves which sweeps area a .*

Proof. Pick a non-separating simple closed curve σ which meets γ minimally and with non-zero geometric intersection number. Such a curve can be obtained from Lemma 2.6. (The construction of σ in that Lemma involves an isotopy of γ which might now have non-trivial flux, but the conclusion of the current Lemma for the modified curve would then yield the same result for the initial one, just with a different base-point to the one-parameter family). One obtains γ_t by sliding a portion of γ along σ , introducing a cancelling pair of self-intersections each time γ is being pushed across itself. The resulting curves bound no teardrops other than those bound by γ . Indeed, in the universal cover, we can choose lifts of γ and σ which meet exactly once, by the minimal intersection assumption; and then the family γ_t which slides along σ bounds no new teardrops since its lift to the universal cover acquires no new self-intersections. \square

Remark 2.9. *Let γ be an unobstructed immersed curve in a compact surface with boundary S . If γ is not boundary-parallel (i.e., it cannot be homotoped to a curve contained in a collar neighborhood of ∂S), then there exists a simple closed curve which has non-zero geometric intersection number with γ , and the above argument shows that there are unobstructed regular isotopies of γ in S which sweep arbitrary area $a \in \mathbb{R}$. On the other hand, if γ is boundary-parallel then the total area between γ and ∂S is finite, and any regular homotopy sweeping more than this amount of area must introduce a teardrop.*

Lemma 2.10. *Let $\gamma \subset \Sigma$ be an embedded simple closed curve. If γ is non-separating, then there are embedded simple closed curves γ_t smoothly isotopic to γ and obtained by an isotopy of flux t , for every $t \in \mathbb{R}$.*

Proof. Fix a simple closed curve σ with non-zero algebraic intersection number with γ , and fix a small embedded cylinder centred on σ . Let σ^\pm be disjoint embedded curves in this cylinder which bound a subcylinder of area $\varepsilon > 0$. The curve

$$(\tau_{\sigma^+} \circ \tau_{\sigma^-}^{-1})(\gamma)$$

is then smoothly isotopic to γ but differs from it by a flux of area ε times the algebraic intersection number. Since it is the image of γ under a diffeomorphism, it is obviously embedded. By varying ε and iterating the map $\tau_{\sigma^+} \circ \tau_{\sigma^-}^{-1}$ or its inverse, one obtains embedded curves differing from γ by arbitrary real values of flux. \square

The analogue of Lemma 2.10 does not hold for separating simple closed curves, cf. Figure 3 below.

For embedded curves, isotopies that sweep zero area are induced by Hamiltonian diffeomorphisms, and the invariance of Floer theory is classical. We record the following consequence (and note that the result also holds for surfaces with boundary):

Corollary 2.11. *Let γ_1, γ_2 be two embedded simple closed curves with non-zero algebraic intersection number, and let X_1, X_2 be the objects of $\mathcal{F}(\Sigma)$ obtained by equipping γ_1, γ_2 with rank one local systems. Then the rank of the Floer cohomology group $HF^*(X_1, X_2)$ is equal to the geometric intersection number of γ_1 and γ_2 .*

Proof. Move γ_2 by an isotopy in order to obtain a simple closed curve γ'_2 which intersects γ_1 minimally, so that there are no bigons bound by $\gamma_1 \cup \gamma'_2$. The isotopy from γ_2 to γ'_2 may sweep a non-zero amount of flux, but this can be remedied by sliding γ'_2 along γ_1 : applying the construction in the proof of Lemma 2.10 to γ'_2 , taking $\sigma = \gamma_1$, yields γ''_2 which is Hamiltonian isotopic to γ_2 and intersects γ_1 minimally. Replacing X_2 by the quasi-isomorphic object X''_2 given by γ''_2 with the appropriate local system, we find that the Floer complex $CF^*(X_1, X''_2)$ has rank equal to the geometric intersection number and vanishing Floer differential. \square

Invariance of Floer cohomology under isotopies that sweep zero area holds more generally for immersed curves, even when we allow the areas of the regions bounded by the curve to vary, or if we allow isotopies through a curve that does not self-intersect transversely, creating or cancelling self-intersections. We note that such regular homotopies are still induced by a (time-dependent) Hamiltonian on the *domain* of the immersion, and thus we will abusively refer to them as Hamiltonian isotopies.

Lemma 2.12. *If unobstructed immersed curves γ_0 and γ_1 are regular homotopic through (generically self-transverse) unobstructed immersed curves by an isotopy that sweeps zero area (and which identifies their Spin structures and local systems), then γ_0 and γ_1 define quasi-isomorphic objects of $\mathcal{F}(\Sigma)$.*

Proof. By concatenation, this follows from the result for C^1 -small isotopies sweeping zero area. When γ_0 and γ_1 are C^1 -close to each other, there exists a small time-independent Morse function H on the domain of the immersion whose Hamiltonian vector field generates the isotopy. Each critical point of H gives rise to an intersection of γ_0 and γ_1 ; considering separately the maxima and minima, we let

$$p = \sum_{p_i \in \min(H)} q^{H(p_i)} p_i \in CF^0(\gamma_0, \gamma_1) \quad \text{and} \quad p' = \sum_{p_j \in \max(H)} q^{-H(p_j)} p_j \in CF^0(\gamma_1, \gamma_0).$$

When γ_0 and γ_1 are embedded, a classical argument considering the bigons that connect the consecutive minima and maxima of H shows that $\mu^1(p) = 0$, $\mu^1(p') = 0$, i.e. p and p' are Floer cocycles, and $\mu^2(p', p) = \text{id}_{\gamma_0}$, $\mu^2(p, p') = \text{id}_{\gamma_1}$, i.e. p and p' provide an explicit isomorphism between γ_0 and γ_1 . When γ_0 and γ_1 are immersed, there are additional generators of $CF(\gamma_0, \gamma_1)$ and $CF(\gamma_1, \gamma_0)$ near the self-intersections. However, by considering lifts of γ_0 and γ_1 to the universal cover (which are embedded by our assumption of unobstructedness), we find that the lifts of holomorphic polygons which contribute to $\mu^1(p)$, $\mu^1(p')$, $\mu^2(p', p)$ and $\mu^2(p, p')$ cannot involve the generators coming from the self-intersections, and the outcome of the calculation is exactly the same as in the embedded case. \square

Any simple closed curve $\sigma \subset \Sigma$ has an associated Dehn twist τ_σ . The identity id and τ_σ both define A_∞ -equivalences of $\mathcal{F}(\Sigma)$, and viewed as functors, there is a distinguished morphism $\Phi_\sigma : \text{id} \rightarrow \tau_\sigma$ in $\text{Hom}_{nu-fun}(\text{id}, \tau_\sigma)$ obtained from the count of sections of a Lefschetz fibration over a disc with one interior critical point and vanishing cycle σ .

Lemma 2.13. *If σ is non-separating, then τ_σ determines σ up to Hamiltonian isotopy. If σ is separating, then the pair $(\tau_\sigma, \Phi_\sigma)$ determines σ up to Hamiltonian isotopy.*

Proof. It is straightforward to see that τ_σ determines the smooth isotopy class of σ . Deforming a simple closed curve σ by an isotopy that sweeps area α , the Dehn twist τ_σ changes by a non-Hamiltonian isotopy whose flux is $\alpha \cdot \text{PD}([\sigma])$. If $[\sigma] \neq 0$ the result follows.

If $[\sigma] = 0$, then σ separates Σ into two subsurfaces. Pick a simple closed curve γ whose geometric intersection number with σ is two, and let γ' be a simple closed curve which is Hamiltonian isotopic to $\tau_\sigma(\gamma)$ and intersects γ and σ minimally, as in Figure 1.

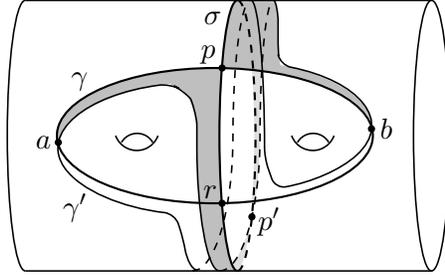


FIGURE 1. Dehn twisting about a separating curve

By the work of Seidel [28], there is an exact triangle in $\mathcal{F}(\Sigma)$ taking the form

$$CF(\sigma, \gamma) \otimes \sigma \xrightarrow{e} \gamma \xrightarrow{f} \gamma' \xrightarrow{\quad} CF(\sigma, \gamma) \otimes \sigma$$

where e is a tautological evaluation map, which can be written in terms of the generators p, r of $CF(\sigma, \gamma)$ as $e = (p, r) : \sigma \oplus \sigma[1] \rightarrow \gamma$, and $f \in CF^0(\gamma, \tau_\sigma(\gamma))$ is given by the natural transformation Φ_σ . Since $\sigma \oplus \sigma[1]$ is not isomorphic to $\gamma \oplus \gamma'[1]$, the morphism f is a non-zero linear combination of the generators a, b of $CF^0(\gamma, \gamma')$, and determined uniquely up to scaling by the property that $\mu^2(f, p) = \mu^2(f, r) = 0$.

Denoting by A and B the symplectic areas of the two shaded triangles on Figure 1, we find that $\mu^2(a, p) = \pm q^A p'$ and $\mu^2(b, p) = \mp q^B p'$, so the vanishing of $\mu^2(f, p)$ implies that f is proportional to $q^B a + q^A b$. (As expected, $\mu^2(f, r)$ then vanishes as well: since γ' is Hamiltonian isotopic to $\tau_\sigma(\gamma)$, the areas of the two triangles contributing to $\mu^2(f, r)$ differ by the same amount $B - A$).

Keeping γ and γ' fixed, when σ moves by an isotopy of flux α (without creating new intersections) the quantity $B - A$ changes by α , so that the class $[f] = \Phi_\sigma(\gamma) \in HF^0(\gamma, \tau_\sigma(\gamma))$ must change as well. This implies that Φ_σ detects the Hamiltonian isotopy class of σ . \square

An object $Y \in \mathcal{F}(\Sigma)^{per}$ is *spherical* if $H^*(\text{hom}_{\mathcal{F}^{per}}(Y, Y)) = H^*(S^1; \Lambda)$. A spherical object Y has an associated twist functor T_Y .

Lemma 2.14. *Let Y, Y' be spherical objects on a surface S . Suppose the twist functors T_Y and $T_{Y'}$ are quasi-isomorphic. If Y is a homologically non-trivial simple closed curve, then Y and Y' are quasi-isomorphic in $\text{Tw}^\pi \mathcal{F}(S)$.*

Proof. Write T for $T_Y \simeq T_{Y'}$. Let δ be a simple closed curve with geometric intersection number 1 with Y . The group $HF(\delta, T(\delta))$ has rank one, so up to quasi-isomorphism there are only two distinct mapping cones

$$Y \simeq \{\delta \xrightarrow{x} T(\delta)\} \quad \text{and} \quad \delta \oplus T(\delta)[1] \simeq \{\delta \xrightarrow{0} T(\delta)\}.$$

Since the functor $T = T_{Y'}$ is a cone over an evaluation functor which has image in the subcategory with objects $V \otimes Y'$ for graded vector spaces V , one of these mapping cones is isomorphic to direct sums of copies of Y' . The sum $\delta \oplus T(\delta)[1]$ is a sum of two non-isomorphic indecomposables, so by the Krull-Schmidt property cannot be a sum of copies of a single indecomposable Y' . Therefore Y is isomorphic to a direct sum of copies of Y' , and indecomposability of Y implies $Y \simeq Y'$ as required. \square

2.4. Generation. Consider an A_{2g} -chain of $2g$ curves $\{\zeta_i\}_{1 \leq i \leq 2g}$ on Σ , as depicted in Figure 2 when $g = 2$.

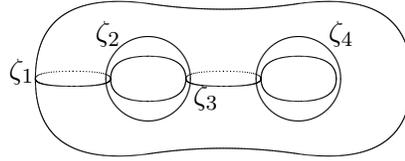


FIGURE 2. Split-generating curves for $\mathcal{F}(\Sigma_2)$

Proposition 2.15. *The curves ζ_i split-generate $\mathcal{F}(\Sigma)$.*

Proof. This is a consequence of [37, Proposition 3.8], a variant of [29, Lemma 6.4]. Briefly, if $u = \prod_{i=1}^{2g} \tau_i$, the Dehn twists τ_i in the curves ζ_i satisfy the positive relation $u^{4g+2} = 1 \in \Gamma_g$. The square of this relation defines a Lefschetz fibration $X \rightarrow \mathbb{P}^1$ with fibre Σ , with $2g(8g+4)$ critical fibres, and for which every section of the fibration has square ≤ -2 .

Let $\delta \in \mathcal{F}(\Sigma)$ be an arbitrary curve (equipped with a local system, which we suppress from the notation). There are exact triangles associated to the Dehn twists τ_i , on the cohomological category $H(\mathcal{F}(\Sigma))$ taking the form

$$HF(\delta, \phi(\delta)) \xrightarrow{P} HF(\delta, \tau_i \circ \phi(\delta)) \longrightarrow HF(\delta, \zeta_i) \otimes HF(\zeta_i, \phi(\delta))$$

for a subword $\phi = \prod_{j < i} \tau_j$ of the monodromy. This triangle is induced from an exact triangle in $\mathcal{F}(\Sigma)$ of the form

$$\phi(\delta) \xrightarrow{P} \tau_i \circ \phi(\delta) \longrightarrow V \otimes \zeta_i$$

where $V = HF(\zeta_i, \phi(\delta))$ is a $\mathbb{Z}/2$ -graded vector space and the arrow p is multiplication by the section count Φ_{τ_i} . Concatenating such triangles for all the twists in u^{8g+4} yields

$$(2.1) \quad \delta \xrightarrow{\hat{p}} u^{8g+4}(\delta) \cong \delta$$

The morphism \hat{p} counts sections of $X \rightarrow \mathbb{P}^1$, and there are no holomorphic such for generic almost complex structure, since all sections have square ≤ -2 and live in moduli spaces of virtual dimension < 0 . Therefore the arrow \hat{p} vanishes, and δ is exhibited as a summand in a triangle whose third entry is a twisted complex on the vanishing cycles ζ_i . \square

Corollary 2.16. *The closed-open map $H^*(\Sigma; \Lambda) \rightarrow HH^*(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ is an isomorphism; similarly for the open-closed map $HH_*(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma)) \rightarrow H_*(\Sigma; \Lambda)$.*

Proof. This is a special case of [37, Corollary 3.11], where the required hypotheses are obtained from Proposition 2.15. \square

Two curves on a surface meet minimally if they meet transversely in their geometric intersection number of points.

Corollary 2.17. *Given a simple closed curve $\sigma \subset \Sigma$, there are curves $\{\xi_1, \dots, \xi_{2g}\}$ which are split-generators for $\mathcal{F}(\Sigma)$ and which meet σ minimally.*

Proof. Up to automorphism, there are only finitely many possibilities for σ , so the result follows by inspection of the pattern of curves in Figure 2. \square

The same methods that underlie Corollary 2.16 also show:

Lemma 2.18. *If Σ is a closed surface, then $\mathcal{F}(\Sigma)$ is homologically smooth and proper.*

Sketch. Properness is immediate since the objects are closed Lagrangians. Smoothness follows from the fact that one can resolve the diagonal on $\Sigma \times \Sigma$ by product Lagrangians, which follows for instance from the argument of [37, Section 3.4]. (An alternative is to use that $\mathcal{F}(\Sigma)$ is equivalent to a category of matrix factorizations, and such categories are always smooth and proper.) \square

2.5. Mukai pairing. Let \mathcal{A} be a proper (cohomologically finite) A_∞ -category, linear over \mathbb{K} . The *Chern character* is a map

$$ch : K_0(\mathcal{A}) \rightarrow HH_0(\mathcal{A}),$$

whilst the *Mukai pairing* is a graded bilinear pairing

$$\langle \bullet, \bullet \rangle : HH_*(\mathcal{A}) \otimes HH_*(\mathcal{A}) \rightarrow \mathbb{K}.$$

These were introduced by Shklyarov in the case of dg -categories [36]; since any A_∞ category is quasi-equivalent to a dg -category, their definitions and basic properties extend to the A_∞ -setting. Shklyarov proved that

$$(2.2) \quad \langle ch(X), ch(Y) \rangle = -\chi(X, Y).$$

In the case of $\mathcal{F}(\Sigma)$, these notions are quite explicit, using the open-closed map to identify $HH_0(\mathcal{F}(\Sigma))$ with $H_1(\Sigma; \Lambda)$. If γ is an unobstructed immersed curve (equipped with a rank

one local system), the absence of non-constant holomorphic discs with boundary on γ implies that the image of $ch(\gamma) = [\text{id}_\gamma] \in HH_0(\mathcal{F}(\Sigma))$ under the open-closed map is exactly $[\gamma] \in H_1(\Sigma; \mathbb{Z}) \subset H_1(\Sigma; \Lambda)$. Comparing (2.2) with the classical identity $\chi HF(\gamma, \gamma') = -[\gamma] \cdot [\gamma']$, we find that the Mukai pairing is simply the intersection pairing on $H_1(\Sigma; \Lambda)$.

Lemma 2.19. *Let $X \in \mathcal{F}(\Sigma)^{\text{per}}$ be a spherical object. Then $ch(X) \in H_1(\Sigma; \Lambda)$ represents an integral class, i.e. $ch(X) \in \text{image}\{H_1(\Sigma; \mathbb{Z}) \rightarrow H_1(\Sigma; \Lambda)\}$.*

Proof. By (2.2), $\langle ch(X), ch(\gamma) \rangle = -\chi HF(X, \gamma) \in \mathbb{Z}$ for any simple closed curve $\gamma \subset \Sigma$. As noted above, under the open-closed map $ch(\gamma)$ maps to $[\gamma] \in H_1(\Sigma; \mathbb{Z})$. It follows that $ch(X)$ has integral pairing with all of $H_1(\Sigma; \mathbb{Z})$, which is only possible if the class is integral. \square

Corollary 2.20. *Let $X \in \mathcal{F}(\Sigma)^{\text{per}}$ be spherical. There is a non-zero class $a \in H^1(\Sigma; \mathbb{Z})$ with $\langle a, ch(X) \rangle = 0$.*

Proof. Evident from Lemma 2.19. \square

2.6. Balancing. There is a natural map

$$\text{Symp}(\Sigma) \rightarrow \text{nu-fun}(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$$

which takes any Hamiltonian symplectomorphism to an equivalence which is quasi-isomorphic to the identity. This yields a map $\text{Symp}(\Sigma)/\text{Ham}(\Sigma) \rightarrow \text{Auteq}(D^\pi \mathcal{F}(\Sigma))$ (here the domain is viewed as a discrete group). For a surface of genus ≥ 2 ,

$$\text{Symp}(\Sigma)/\text{Ham}(\Sigma) = H^1(\Sigma; \mathbb{R}) \rtimes \Gamma(\Sigma)$$

by Moser's theorem and the vanishing of the flux group. To build a homomorphism $\Gamma_g \rightarrow \text{Auteq}(D^\pi \mathcal{F}(\Sigma_g))$ requires some additional choice. Suppose $g \geq 2$, and fix a primitive θ for the pullback of ω_Σ to the unit tangent bundle $S(T\Sigma)$. For any simple closed curve $\sigma \subset \Sigma$, a choice of orientation of σ defines a canonical lift $\sigma \subset S(T\Sigma)$, and we then have a real number $t_\sigma = \int_\sigma \theta$. Say σ is balanced if $t_\sigma = 0$, and define a *balanced* symplectomorphism $f : \Sigma \rightarrow \Sigma$ to be one which takes balanced curves to balanced curves, i.e. for which $t_{f(\sigma)} = t_\sigma$ for every oriented simple closed curve $\sigma \subset \Sigma$.

Lemma 2.21. $\{\text{Balanced symplectomorphisms}\}/\text{Ham}(\Sigma) \simeq \Gamma_g$.

Proof. See [29]. \square

It follows that the choice of θ defines a map

$$(2.3) \quad \Gamma_g \rightarrow \text{Auteq}(D^\pi \mathcal{F}(\Sigma_g)).$$

Given two primitives θ, θ' for $p^*\omega_\Sigma$, with $p : S(T\Sigma) \rightarrow \Sigma$, one obtains a class $[\theta - \theta'] \in H^1(S(T\Sigma); \mathbb{R})$. Note that, since $\chi(\Sigma) \neq 0$, the pullback p^* induces an isomorphism on first cohomology, so we can think of $[\theta - \theta']$ as an element of $H^1(\Sigma; \mathbb{R})$. Changing θ to θ' conjugates the image of (2.3) by the action of this element of $H^1(\Sigma; \mathbb{R})$.

Remark 2.22. *The analogous construction for punctured surfaces may be more familiar to the reader: the choice of an exact symplectic structure on a punctured surface S (i.e. a primitive θ for the symplectic form itself, rather than its lift to the unit tangent bundle)*

determines a homomorphism from the mapping class group of S to $\text{Auteq}(D^\pi \mathcal{F}(S))$ by considering exact symplectomorphisms of S , i.e. those which take exact curves to exact curves, up to Hamiltonian isotopy.

2.7. Cautionary examples. Despite its generally elementary character, there are some surprises in Floer theory for curves on surfaces.

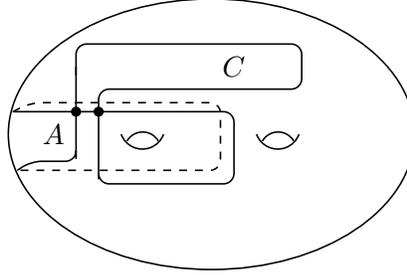


FIGURE 3. An exotic spherical object when $C > A$

Lemma 2.23. *There are spherical objects in $\mathcal{F}(\Sigma)$ which are not quasi-isomorphic to any simple closed curve with local system.*

Sketch. See Figure 3, which shows an immersed curve by pushing a separating simple closed through itself to create a single bigon. The region labelled C contributes to a non-trivial Floer differential, so this immersed curve is spherical. However, the hypothesis that the area $C > A$ implies that one cannot deform the curve to be embedded through an isotopy which has trivial flux: the end result would have to separate Σ into two regions, one of area $A - C < 0$. It is not hard to see that, if this immersed curve was quasi-isomorphic to a simple closed curve, that curve would have to be in the same homotopy class, and then consideration of Lemma 2.13 would show that the natural transformation from the identity to the associated twist functor would involve a different linear combination of Floer generators. \square

Lemma 2.24. *There is an immersed curve $S^1 \rightarrow \Sigma$ with rank one local system which admits non-trivial idempotents.*

Proof. Take an immersed curve γ which is homotopic to the double cover of a simple closed curve σ . Then a rank one local system on γ defines an object quasi-isomorphic to a rank two local system ξ on σ , and rank one sub-local-systems of ξ define idempotents. \square

Lemma 2.25. *Let p be a transverse intersection point of curves γ, γ' such that p is a Floer cocycle in $CF^*(\gamma, \gamma')$. The immersed curve resulting from surgery at p need not be quasi-isomorphic to the mapping cone $\gamma \xrightarrow{p} \gamma'$.*

Proof. Consider Figure 4, where in the first image the Floer complex $CF^*(\gamma, \gamma')$ has

$$dr_1 = q^{\alpha_1} \cdot p, \quad dr_2 = -q^{\alpha_2} \cdot p$$

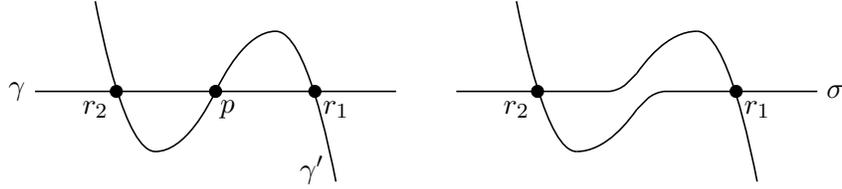


FIGURE 4. The mapping cone is not the surgery

with α_1, α_2 the areas of the bigons; we assume there are no holomorphic strips other than those in the picture. Then p is an *exact* Floer cocycle, so

$$\text{Cone}(p) \simeq \gamma[1] \oplus \gamma'.$$

In this case, the Lagrange surgery σ at p is quasi-isomorphic to the cone on $q^{\alpha_1}r_1$ viewed as a (closed) morphism in the reverse direction, from γ' to γ (or equivalently, $q^{\alpha_2}r_2$, which is cohomologous). However, there are also examples where the Lagrange surgery at an exact Floer cocycle between a pair of simple closed curves γ, γ' yields an immersed curve σ which lies outside of their triangulated envelope (Figure 5). In all these examples, the surgered curve σ remains cobordant to $\gamma \cup \gamma'$, but the Lagrangian cobordism between them is obstructed, and the cobordism only yields an exact triangle after deforming σ by a suitable bounding cochain (which amounts geometrically to smoothing a self-intersection of σ , to obtain a curve in a different homotopy class – in fact, homotopic to $\gamma \cup \gamma'$ rather than σ). This is a purely one-dimensional phenomenon – for instance, in higher dimensions, there would be no rigid strip passing through the neck region after surgery, in contrast to the visible strip in the right image of Figure 4. \square

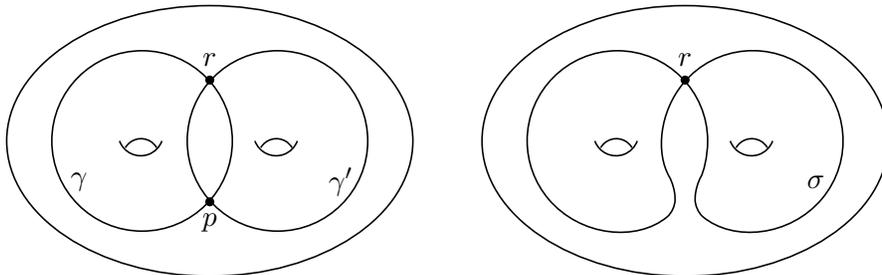


FIGURE 5. The mapping cone differs from the surgery by a bounding cochain

3. IMMERSED CURVES AND BIGONS

3.1. Bigons. In this section, S denotes a compact surface which may have empty or non-empty boundary. Let $\gamma \subset S$ be an unobstructed immersed curve, which we always assume has only transversal self-intersections. An *embedded bigon* with boundary on γ is a map $u : D \rightarrow S$ from the closed disc to S which takes $\pm 1 \in \partial D$ to self-intersection points of γ , which takes the boundary $\partial D \setminus \{\pm 1\}$ to γ , and which is an embedding $D \rightarrow S$. Note that

there will in general be arcs of γ which meet the interior of the bigon. If there are no such arcs, so $\gamma \cap u(D) = u(\partial D)$, then we say the bigon is *empty*.

Lemma 3.1. *An immersed curve $\gamma \subset S$ bounds at most finitely many embedded bigons.*

Proof. For a given pair of intersection points, there are only finitely many possible boundary arcs in γ between them. A pair of arcs which cuts out a disc in S defines a unique bigon. \square

Lemma 3.2 (Hass, Scott). *Let $\gamma \subset \Sigma$ be an immersed closed curve which does not bound any teardrops. If γ is homotopic to a simple closed curve, but is not embedded, then γ bounds an embedded bigon.*

Proof. This is proved in [21]. \square

Hass and Scott show by examples that one can not in general assume that γ bounds an empty bigon. We introduce a combinatorial move on immersed curves:

- The *triple-point move* changes a configuration of three intersecting arcs which cut out a small downwards-pointing triangle to one defining a small upwards-pointing triangle, cf. Figure 6.

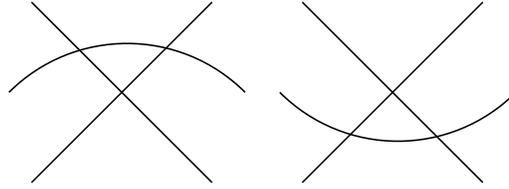


FIGURE 6. The triple-point move

Lemma 3.3. *The triple-point move preserves the total number of self-intersections of γ , and does not increase the total number of embedded bigons on γ .*

Proof. Evident. \square

Lemma 3.4 (Steinitz). *If γ bounds an embedded bigon, then after a finite sequence of triple-point moves, γ bounds an empty bigon.*

Proof. This is proved in [38, 39], see also [14]. \square

3.2. Removing bigons. Throughout this section, suppose S has non-empty boundary. Let $v : \gamma \rightarrow S$ be an immersion with $v(\gamma) \in H_1(S; \mathbb{Z})$ a primitive class.

Lemma 3.5. *Suppose there is an embedded polygon $u(D) \subset S$ bound by $v(\gamma)$. Fix a point $p \in \text{int}(u(D))$. Then there is an embedded arc $\sigma : [0, 1] \rightarrow S$ with (i) $\sigma(0) = p$ and $\sigma(1) \in \partial S$, (ii) with $\sigma \cap u(D)$ connected, and (iii) with algebraic intersection number $\langle \sigma, v(\gamma) \rangle = 0$.*

Proof. Since $v(\gamma)$ is primitive, there is a simple closed curve $\kappa \subset S$ with $\langle v(\gamma), \kappa \rangle = 1$ (for appropriate orientation of κ). Sliding κ off $u(D)$ if necessary, we may suppose $\kappa \cap u(D) = \emptyset$. Since κ is nonseparating, there exists an embedded arc in $S \setminus u(D)$ connecting the boundary of $u(D)$ to the boundary of S and intersecting κ just once. Extending this arc into $u(D)$, we obtain an embedded arc σ' from p to ∂S , with $\sigma' \cap u(D)$ connected, and with $|\sigma' \cap \kappa| = 1$. If the algebraic intersection number of the arc σ' with $v(\gamma)$ is non-zero, then by Dehn twisting σ' repeatedly around κ (which amounts to connect summing copies of κ while preserving embeddedness), we obtain a new arc σ with the desired properties. \square

Example 3.6. *Let S be an annulus and γ an immersed curve with a single self-intersection representing twice the generator of $H_1(S; \mathbb{Z})$. There is no arc from a point p inside the bigon bound by γ to ∂S having zero algebraic intersection with γ .*

Lemma 3.7. *If $v(\gamma)$ bounds an empty embedded polygon, there is a Hamiltonian isotopy of $v(\gamma)$ which decreases the area of this polygon to be arbitrarily small without creating any self-intersections.*

Proof. Pick a point p in the empty polygon, and an arc σ as in the previous Lemma. Bend $v(\gamma)$ along σ towards the point p , i.e. deform $v(\gamma)$ near each point of $v(\gamma) \cap \sigma$ in the direction of p , for instance by the flow of a vector field supported near σ and constant and non-zero along σ . One can choose the local bending to be by an equal small amount at each intersection point. The vanishing of the algebraic intersection number with σ implies that this isotopy sweeps zero area (i.e. it is induced by a Hamiltonian on the domain of the immersion). Inside the empty polygon containing p , this pushes one side towards the others, decreasing the area of the polygon. We can iterate this process until the area is arbitrarily small. \square

Lemma 3.8. *If γ bounds an empty bigon, then γ is quasi-isomorphic to a curve γ' , with two fewer self-intersections, obtained by cancelling the bigon.*

Proof. The previous Lemma shows that we can deform γ by a Hamiltonian isotopy to decrease the area of the bigon to be arbitrarily small. Once the bigon is sufficiently small, we can cancel the pair of self-intersections by a regular homotopy that sweeps zero area, without creating any other intersections; the result then follows from Lemma 2.12. \square

Lemma 3.9. *If γ bounds an empty triangle, then γ is quasi-isomorphic to a curve γ' obtained by performing a triple point move.*

Proof. The argument is the same as for bigons: Lemma 3.7 shows we can shrink the triangle to have arbitrarily small area by an isotopy that sweeps zero area, and once the triangle is sufficiently small we can perform the triple point move by a regular homotopy that sweeps zero area. The result then follows from Lemma 2.12. \square

3.3. Analyticity of Floer cohomology. Let (X, ω) be a symplectic manifold and $\{F_b\}_{b \in B}$ a family of unobstructed Lagrangian submanifolds parametrized by a smooth manifold B . (The prototypical situation in the literature would be that B is a subset of a tropical SYZ base, and we will sometimes refer to the F_b as fibres.)

Over a small disk $b_0 \in P \subset B$, the fibres F_b are graphs of closed one-forms α_b over F_{b_0} ; the Hamiltonian isotopy class of F_b depends only on the de Rham cohomology class of the one-form α_b . The space of choices $(b, \xi \rightarrow F_b)$ is therefore naturally a domain centred on the base-point $b_0 = (0, 1)$ inside $(\Lambda^*)^k = H^1(F_{b_0}; \Lambda^*) = H^1(F_{b_0}; \mathbb{R}) \times H^1(F_{b_0}; U_\Lambda)$, where $k = \text{rk}_{\mathbb{Z}} H^1(F_{b_0}; \mathbb{Z})$. To be more explicit, fixing a basis a_1, \dots, a_k for $H_1(F_{b_0}; \mathbb{Z})$, we have Λ^* -valued co-ordinates z_1, \dots, z_k given by $z_i(b, \xi) = q^{[\alpha_b] \cdot a_i} \text{hol}_\xi(a_i)$. More intrinsically, to every element $\gamma = \sum \gamma_i a_i \in H_1(F_b; \mathbb{Z})$ corresponds a monomial

$$z_{b, \xi}^\gamma = z_1^{\gamma_1} \dots z_k^{\gamma_k} = q^{[\alpha_b] \cdot \gamma} \text{hol}_\xi(\gamma).$$

Consider a Lagrangian submanifold $L \subset X$ transverse to the fibres F_b over P . Over P , L defines an unbranched cover of P , so the intersection points $\{L \pitchfork F_b\}$ may be identified with $L \pitchfork F_{b_0}$. Fix a local section of the family $\{F_b\}$, so that for each fibre F_b we can fix a smoothly varying base-point $\star \in F_b$. In F_{b_0} , fix a (homotopy class of) path γ_x from x to \star , for each $x \in L \pitchfork F_{b_0}$. We fix a rank one local system $\xi \rightarrow F_{b_0}$ and an arbitrary identification $\xi_\star \cong \Lambda$; by parallel transport along γ_x , this identifies $\xi_x \cong \Lambda$ for each intersection $x \in L \cap F_{b_0}$.

Now consider another point $b \in P$. Assuming P is convex, over a path joining b and b_0 the path γ_x sweeps a two-chain Γ_x , yielding an area $a_x(b) = \int_{\Gamma_x} \omega$. By a trick due to Fukaya (essentially the observation that taming is an open condition on almost complex structures), a rigid holomorphic strip u with boundaries on L and F_{b_0} will deform, for P sufficiently small, to a rigid holomorphic strip u' with boundaries on L and F_b . By concatenating the boundary arcs of u' with the reference paths, we get a well-defined element $[\partial u'] \in H_1(F_b; \mathbb{Z})$.

Lemma 3.10. *Suppose the arc of the boundary of u connects intersection points x, y . We have*

$$q^{E(u')} \text{hol}_\xi(\partial u') = q^{a_y(b) - a_x(b)} \cdot q^{E(u)} z_{b, \xi}^{[\partial u']}.$$

Proof. See [2]. □

The key point is that this expression involves the monomial $z_{b, \xi}^{[\partial u']}$, whereas the quantity $q^{E(u)}$ is *constant*, depending only on the reference point b_0 . In particular, under the rescaling

$$CF^*(L, (\xi \rightarrow F_{b_0})) \longrightarrow CF^*(L, (\xi \rightarrow F_b)), \quad x \mapsto q^{a_x(b)} \cdot x = x'$$

one finds that the Floer differential becomes analytic as a function of the co-ordinates $z_{b, \xi}$; schematically,

$$\langle \mu^1(x'), y' \rangle = \sum_{[u]} \#\mathcal{M}([u]) \cdot q^{E([u])} z_{b, \xi}^{[\partial u]}.$$

The dependence on the choice of base-point b_0 and of the homotopy classes of the paths γ_x are also analytic; the former changes the values $a_x(b)$ by some fixed constant, rescaling x' by a value which does not depend on $z_{b, \xi}$, whilst the latter rescales x' by a monomial.

Corollary 3.11. *There is an affinoid neighbourhood P of $b_0 \in (\Lambda^*)^k$ such that the Floer cohomology groups $HF^*(L, (\xi \rightarrow F_b))$ are the fibres of an analytic sheaf over P .*

The same argument would apply to a family of immersed Lagrangians equipped with analytically varying bounding cochains.

Up to this point, we have assumed that the Lagrangian L is transverse to all the fibres $\{F_b\}_{b \in P}$. To obtain a more global statement, given a family of Lagrangians $\{F_b\}_{b \in B}$ one chooses a finite set of Hamiltonian perturbations L_i of L for which the corresponding caustics of the projections $L_i \rightarrow B$ have empty intersection. This finite set of Hamiltonian perturbations can be spanned by a simplex of Hamiltonian perturbations, and there are (higher) continuation maps on Floer cochains associated to the edges and higher-dimensional facets of this simplex. These yield a module for a Čech complex of the corresponding covering of B , and (for a sufficiently fine cover) gluing maps of the local analytic sheaves over affinoids in B constructed previously. Summing up:

Theorem 3.12 (Abouzaid). *Let $\{F_b\}_{b \in B}$ be a family of unobstructed Lagrangians parametrized by an open $B \subset H^1(F_b; \mathbb{R})$. There is an analytic dg-sheaf over the annulus $B \times U_\Lambda \subset H^1(F_b; \Lambda^*)$ with stalk quasi-isomorphic to $CF^*(L, (\xi \rightarrow F_b))$. Furthermore, this association yields an A_∞ -functor from the tautologically unobstructed Fukaya category of (X, ω) into the dg-category of complexes of sheaves over $B \times U_\Lambda$.*

Proof. See [3]. □

Remark 3.13. *Suppose we fix a class $a \in H^1(F_b; \mathbb{R})$ and consider the corresponding real one-parameter family of Lagrangians F_b^t given by moving F_b by symplectomorphisms of flux $t \cdot a$. Then for a test Lagrangian L , there is a discrete set of values of $t \in \mathbb{R}$ where the transversality condition $L \pitchfork F_b^t$ fails. Picking Hamiltonian perturbations, the Čech complex above amounts to a zig-zag diagram of quasi-isomorphisms associated to neighbouring such intervals; generically, these quasi-isomorphisms furthermore correspond to the simplest birth-death bifurcations of Floer complexes, adding or subtracting an acyclic subcomplex with two generators and a unique minimal area bigon.*

3.4. Spherical objects. Let S be a surface which may be closed or have non-empty boundary. Let $\gamma \subset S$ be an unobstructed immersed curve and $\xi \rightarrow \text{domain}(\gamma)$ a local system. Equip S with a symplectic form ω and compatible complex structure j .

Lemma 3.14. *If (ξ, γ) defines a spherical object of $\mathcal{F}(S)$, then γ is regular homotopic to a simple closed curve.*

Proof. Equipping S with a hyperbolic metric, it is a classical fact that there is a unique geodesic η in the homotopy class of γ . In fact, performing the homotopy by a suitable curve-shortening flow (see e.g. [22]), one finds that γ is regular homotopic to η among generically self-transverse unobstructed immersed curves. We claim that (ξ, γ) being spherical implies that the geodesic η is a simple closed curve.

Recall that the Floer complex $CF^*((\xi, \gamma), (\xi, \gamma))$ splits into a direct sum of complexes corresponding to the various lifts of γ to the universal cover of S which have non-trivial intersection with a fixed lift $\tilde{\gamma}$. (This is because Floer generators which correspond to intersections between different pairs of lifts of γ cannot be connected by bigons.) The summand which corresponds to the trivial homotopy class (i.e., intersections of $\tilde{\gamma}$ with a small Hamiltonian perturbation of itself) contributes $H^*(S^1; \text{End}(\xi))$, which has rank at least two (considering the identity endomorphism of ξ). Thus, all other summands must be acyclic.

If the geodesic η is multiply covered, then by homotoping γ to η , then translating along the underlying simple geodesic, and homotoping back, we find that there is a regular homotopy of γ to itself which sweeps zero area and turns the chosen lift $\tilde{\gamma}$ into a different lift. Hamiltonian isotopy invariance (cf. the proof of Lemma 2.12) then implies that this pair of lifts also contributes non-trivially to the Floer cohomology of (ξ, γ) to itself, which contradicts the assumption that (ξ, γ) is spherical. Hence η cannot be multiply covered.

Finally, observe that each lift of γ to the universal cover of S lies within bounded distance of a lift of the geodesic η . If two lifts of η intersect each other (necessarily at a unique point, since they are hyperbolic geodesics), then the corresponding lifts of γ also have algebraic intersection number equal to 1, and hence they must contribute non-trivially to Floer cohomology, which contradicts the assumption. Thus, the lifts of η to the universal cover of S are pairwise disjoint, which implies that η is embedded. \square

Lemma 3.15. *Let $\xi \rightarrow \gamma$ be an indecomposable local system over a simple closed curve γ . The endomorphism ring $H^*(\text{hom}_{\mathcal{F}}((\xi, \gamma), (\xi, \gamma))) = H^*(\gamma; \text{End}(\xi))$ has rank 2 if and only if ξ has rank 1.*

Proof. An indecomposable rank r local system ξ is determined by its monodromy A . The fibre Λ^r is cyclic as a $\Lambda[t^{\pm 1}]$ -module (where t acts by A), generated by any vector, and can thus be identified with $\Lambda[t^{\pm 1}]/\langle \chi_A(t) \rangle$, where χ_A is the characteristic polynomial of the monodromy. Then $H^0(\text{End}(\xi))$ contains the maps $\Lambda^r \rightarrow \Lambda^r$ which commute with the monodromy, i.e. $\Lambda[t^{\pm 1}]$ -module maps

$$\Lambda[t^{\pm 1}]/\langle \chi_A(t) \rangle \longrightarrow \Lambda[t^{\pm 1}]/\langle \chi_A(t) \rangle$$

are determined by the image of 1, which can be any element of $\Lambda[t^{\pm 1}]/\langle \chi_A(t) \rangle \simeq \Lambda^r$, we conclude that $H^0(\text{End}(\xi))$ has rank r over Λ ; the same is then true for H^1 by considering Euler characteristic. \square

Corollary 3.16. *If $(\xi, \gamma) \subset S$ is an immersed curve with local system which defines a spherical object $X \in \mathcal{F}(S)$, and if $[\gamma] \in H_1(S; \mathbb{Z})$ is primitive, then γ is quasi-isomorphic to an embedded simple closed curve and ξ has rank one.*

Proof. Being spherical implies that γ is homotopic to a simple closed curve, hence bounds an embedded bigon. By a sequence of triple point moves, we may find an empty bigon bound by γ , which we may then shrink by Hamiltonian isotopy and cancel to obtain a new immersed curve γ' in the same quasi-isomorphism class. By repeated applying Lemmas 3.8 and 3.9, we eventually arrive at a simple closed curve. Finally, Lemma 3.15 implies that ξ must have rank one. \square

4. GEOMETRIZATION ON PUNCTURED SURFACES

4.1. The wrapped Fukaya category of a non-compact surface. Let $(S, \partial S)$ be a symplectic surface with non-empty boundary; fix a finite subset $\Lambda \subset \partial S$ of boundary marked points (the ‘‘stops’’) and a homotopy class of line field $\eta \subset TS$. Associated to this data is a \mathbb{Z} -graded partially wrapped Fukaya category $\mathcal{W}(S, \Lambda, \eta)$. If one does not make a choice of line field, there is also a $\mathbb{Z}/2$ -graded category $\mathcal{W}(S, \Lambda)$; when the set $\Lambda \subset \partial S$ of stops is empty,

we will simply write $\mathcal{W}(S)$ or $\mathcal{W}(S, \eta)$. There are forgetful functors $\mathcal{W}(S, \Lambda, \eta) \rightarrow \mathcal{W}(S, \Lambda)$ which forget the grading structure, and localization functors $\mathcal{W}(S, \Lambda) \rightarrow \mathcal{W}(S, \Lambda')$ which forget some of the stops whenever $\Lambda' \subset \Lambda$; in particular there are “acceleration” functors $\mathcal{W}(S, \Lambda) \rightarrow \mathcal{W}(S)$ which are (by definition) surjective on objects.

The objects of $\mathcal{W}(S, \Lambda)$ (resp. $\mathcal{W}(S, \Lambda, \eta)$) are (graded) unobstructed properly immersed curves or arcs with boundary in $\partial S \setminus \Lambda$, equipped with local systems. While the construction is usually carried out in the exact setting, we work over the Novikov field and allow non-exact objects into our category. Recall that generators of the wrapped Floer complex arise not only from intersection points but also from (positively oriented) boundary chords in $\partial S \setminus \Lambda$ connecting the end points of a pair of arcs. Fixing a Liouville structure on S , the structure maps of the wrapped category count isolated solutions of Floer’s equation with a suitable Hamiltonian perturbation in the Liouville completion of S . These counts are weighted by the topological energy of the solutions and by holonomy terms. In the case of arcs, the weights can be cancelled out by trivializing the local systems and rescaling generators by their Floer action; it is only for non-exact objects that Novikov coefficients are necessary. Note that the wrapped category only depends on the Liouville completion of (S, Λ) ; in particular it is independent of the choice of Liouville structure.

Remark 4.1. *As in the case of Fukaya categories of closed Riemann surfaces, the structure maps of $\mathcal{W}(S, \Lambda)$ can be determined combinatorially in terms of immersed polygons with convex corners. Indeed, solutions to Floer’s equation which do not lie entirely in the cylindrical ends of the completion can be reinterpreted as immersed polygons in S whose boundary lies partially on the given Lagrangians and partially along chords in $\partial S \setminus \Lambda$; and solutions which lie entirely in the cylindrical ends only contribute a “classical” term to μ^2 which concatenates two boundary chords with a common end point.*

4.2. Geometricity of twisted complexes. We will say that an object Y of $\mathrm{Tw}^\pi \mathcal{W}(S, \Lambda)$ is *geometric* if it is quasi-isomorphic to a union of immersed arcs or curves with local systems in S or its Liouville completion. The category $\mathrm{Tw} \mathcal{W}(S, \Lambda, \eta)$ has a combinatorial model, due to [19], using which they prove the remarkable:

Theorem 4.2 (Haiden, Katzarkov, Kontsevich). *Let $(\Sigma, \partial\Sigma, \Lambda)$ be a surface with non-empty boundary and a (possibly empty) collection of boundary marked points $\Lambda \subset \partial\Sigma$. Let $Y \in \mathrm{Tw} \mathcal{W}(\Sigma, \Lambda)$ be a $\mathbb{Z}/2$ -graded twisted complex. Then Y is geometric.*

We include a brief discussion of the proof, to illustrate why its ingredients do not readily generalise to the case of closed surfaces treated in this paper, and to clarify its applicability to the $\mathbb{Z}/2$ -graded case and to the non-exact setting.

Sketch. Suppose that $\Lambda \cap C \neq \emptyset$ for each component $C \subset \partial\Sigma$; the general case will follow from this by localization. A “full formal arc system” is a collection $\{a_i\}$ of disjoint embedded arcs with boundary in $\partial\Sigma \setminus \Lambda$ which decompose Σ into a union of discs each containing exactly one point of Λ . Any such system of arcs generates the category $\mathcal{W}(\Sigma, \Lambda)$. The A_∞ -algebra of endomorphisms of the collection of objects $\{a_i\}$ has a particularly simple description: it is formal, and the only non-trivial products correspond to the concatenation of boundary chords on each component of $\partial\Sigma \setminus \Lambda$. This can be described by a nilpotent quiver algebra,

with vertices the arcs and arrows the boundary chords connecting successive end points along each component of $\partial\Sigma \setminus \Lambda$.

Since the arcs a_i generate the wrapped category, any object of $\text{Tw } \mathcal{W}(\Sigma, \Lambda)$ can be expressed as a twisted complex $A = (\bigoplus V_i \otimes a_i, \delta)$ for some $(\mathbb{Z}$ or $\mathbb{Z}/2)$ graded vector spaces V_i . This twisted complex can be viewed as a representation of a “net”, i.e. a collection of vector spaces which carry two filtrations, together with prescribed isomorphisms between certain pieces of the associated graded.

To a component c of $\partial\Sigma \setminus \Lambda$, containing end points of the arcs a_{i_1}, \dots, a_{i_k} in that order, we associate the $(\mathbb{Z}$ or $\mathbb{Z}/2)$ graded vector space

$$V_c := V_{i_1}[d_1] \oplus \dots \oplus V_{i_k}[d_k],$$

where the shifts d_1, \dots, d_k reflect the gradings of the arcs near the relevant end points. This carries two filtrations. One comes from the ordering of the arcs along the boundary of Σ :

$$V_{i_1}[d_1] \subset V_{i_1}[d_1] \oplus V_{i_2}[d_2] \subset \dots \subset V_c.$$

The other comes from the part of the differential of the twisted complex which involves boundary chords lying along c , viewed as an endomorphism $\delta_c \in \text{End}(V_c)$ which squares to zero, giving the filtration

$$\text{im}(\delta_c) \subset \ker(\delta_c) \subset V_c.$$

In the latter filtration, δ_c induces isomorphisms $V_c/\ker(\delta_c) \simeq \text{im}(\delta_c)$, whereas in the former, each V_i appears twice in the associated graded (once for each end point of the arc a_i).

Because the language of nets is formulated for ungraded vector spaces, the argument of [19, Section 4.4] actually splits the vector spaces V_c according to cohomological degree: in the \mathbb{Z} -graded case, the indexing set for the net is $\pi_0(\partial\Sigma \setminus \Lambda) \times \mathbb{Z}$, and one considers the collection of vector spaces V_c^d for all $c \in \pi_0(\partial\Sigma \setminus \Lambda)$ and $d \in \mathbb{Z}$ (each equipped with the two filtrations described above).

With this understood, a classification theorem generalising results of [25] implies that any indecomposable representation of a net is pushed forward from an indecomposable representation of a net of “height 1”, i.e. one in which all the filtrations have length 1. These correspond to twisted complexes “locally” built from pieces that involve a single arc among the a_i , and connecting differentials that are isomorphisms between multiplicity vector spaces that are concentrated in a single degree and correspond to a single boundary chord (not a linear combination). Such twisted complexes look like either

$$\begin{array}{ccccccc} V_{j_1} \otimes a_{j_1} & \longrightarrow & V_{j_2} \otimes a_{j_2} & \longleftarrow & \dots & \longrightarrow & V_{j_\ell} \otimes a_{j_\ell} & \text{or} \\ V_{j_1} \otimes a_{j_1} & \longrightarrow & V_{j_2} \otimes a_{j_2} & \longleftarrow & \dots & \longrightarrow & V_{j_\ell} \otimes a_{j_\ell} & \\ & & & & & \searrow & & \\ & & & & & & & \nearrow \end{array}$$

where the vector spaces V_{j_i} are all isomorphic up to grading shift, and the arrows between $V_{j_i} \otimes a_{j_i}$ and $V_{j_{i+1}} \otimes a_{j_{i+1}}$ can point in either direction and each correspond to a single boundary chord. Interpreting the mapping cone of a boundary chord geometrically as a boundary connected sum surgery, these two kinds of twisted complexes correspond respectively to immersed arcs and curves equipped with local systems. (See [19]; see also [13] for an earlier classification of objects of the derived category of a cycle of rational curves based on the same algebraic formalism.)

Since we are working over the Novikov field, there is one more subtlety that arises: when using boundary connected sum surgeries to build a geometric object out of the second kind of indecomposable twisted complex, one arrives at an immersed closed curve carrying a local system which is not necessarily unitary. However, irreducibility implies that the holonomy has a single eigenvalue, whose valuation can be adjusted by modifying the boundary connected sum construction by an isotopy that sweeps a suitable amount of flux. (The isotopy may however require replacing Σ with a larger domain inside the Liouville completion; as noted in Remark 2.9, for boundary-parallel curves this is unavoidable, whereas for all other curves one can find enough space within Σ by applying the trick of Lemma 2.8). We then arrive at an immersed curve with a unitary local system.

In [19] the classification theorem is stated for objects of the \mathbb{Z} -graded category (where gradings are defined with respect to any choice of line field η). However, the argument above uses filtrations coming from the boundary structure of the full formal arc system and from the differential of the twisted complex, and *not* from the \mathbb{Z} indexing degrees. The argument of [19, Section 4.4] carries over without modification to the $\mathbb{Z}/2$ -graded case simply by using $\pi_0(\partial\Sigma \setminus \Lambda) \times \mathbb{Z}/2$ instead of $\pi_0(\partial\Sigma \setminus \Lambda) \times \mathbb{Z}$ as indexing set for the net and reducing the second factor mod 2 in all the statements. This yields geometricity for $\mathbb{Z}/2$ -graded twisted complexes. (In the special case of a once-punctured torus, an explicit algorithm for producing the geometric replacement of a $\mathbb{Z}/2$ -graded twisted complex guaranteed by Theorem 4.2 is given in [20].) \square

Remark 4.3. *Fix a grading structure on S and a full formal arc system \mathcal{A} . The category $\mathrm{Tw}_{gr}(\mathcal{A})$ of graded twisted complexes over \mathcal{A} is split-closed, since it admits a stability condition by [19]; hence $\mathrm{Tw}_{gr}(\mathcal{A}) = \mathrm{Tw}_{gr}^\pi(\mathcal{A})$. There is a commuting diagram*

$$\begin{array}{ccc} \mathrm{Tw}_{gr}(\mathcal{A}) & \longrightarrow & \mathrm{Tw}^\pi(\mathcal{A}) \\ \downarrow & & \downarrow \\ SH^*(S) & \xlongequal{\quad} & SH^*(S) \end{array}$$

where the top arrow collapses the \mathbb{Z} -graded structure to its underlying $\mathbb{Z}/2$ -grading and the vertical arrows are open-closed maps (these factor through the localisation functors from partially to fully wrapped categories). Both vertical maps hit the unit, by [17], so the image of $\mathrm{Tw}_{gr}(\mathcal{A})$ in $\mathrm{Tw}^\pi(\mathcal{A})$ is a split-generating subcategory. Nonetheless, the only idempotents which are admitted in $\mathrm{Tw}_{gr}(\mathcal{A})$ are those of degree zero. Concretely, a simple closed curve which separates a genus one subsurface of S containing no punctures cannot be graded for any choice of line field, since the winding number of the line field is necessarily non-zero by the Poincaré-Hopf theorem. Such a curve defines an object of $\mathrm{Tw}^\pi(\mathcal{A})$ which does not lift to $\mathrm{Tw}_{gr}^\pi(\mathcal{A}) = \mathrm{Tw}_{gr}(\mathcal{A})$ for any choice of grading structure.

Theorem 4.2 shows that objects of $\mathrm{Tw} \mathcal{W}(S)$ are geometric, but not that objects of the split-closure $\mathrm{Tw}^\pi \mathcal{W}(S)$ are geometric. For a line field η on S we have a \mathbb{Z} -graded full subcategory $\mathrm{Tw} \mathcal{W}(S, \eta) \subset \mathrm{Tw} \mathcal{W}(S)$. The existence of stability conditions on $\mathrm{Tw} \mathcal{W}(S, \eta)$ (as constructed in [19]) implies that it is split-closed. However, Remark 4.3 implies that one cannot reduce geometricity of idempotent summands of $\mathrm{Tw} \mathcal{W}(S)$ to geometricity of

objects of $\mathrm{Tw} \mathcal{W}(S, \eta)$, since there are objects in the former which don't lift to the latter for any choice of η .

Section 4.4, following a strategy from [4, Appendix B], uses homological mirror symmetry to prove that the $\mathbb{Z}/2$ -graded category $\mathrm{Tw} \mathcal{W}(S)$ is split-closed whenever $S \subset (\mathbb{C}^*)^2$ is a very affine curve in a maximally degenerating family. The next section reviews the relevant mirror symmetry input, due to Heather Lee [24].

4.3. Homological mirror symmetry for punctured surfaces. We consider a finite subset $A \subset \mathbb{Z}^2$ and a function $\rho : A \rightarrow \mathbb{R}$ which is the restriction of a convex piecewise-linear function $\bar{\rho} : \mathrm{Conv}(A) \rightarrow \mathbb{R}$. We assume that the maximal domains of linearity of $\bar{\rho}$ are the cells of a “regular” polyhedral decomposition \mathcal{P} of $\mathrm{Conv}(A)$, i.e. one with vertex set A and for which every maximal cell is congruent to a $GL(2, \mathbb{Z})$ -image of the standard simplex. We consider a punctured surface

$$S_t = \left\{ \sum_{a \in A} c_a t^{-\rho(a)} z^a = 0 \right\} \subset \mathbb{C}^* \times \mathbb{C}^*, \quad t \gg 0$$

with its natural exact convex symplectic structure. Explicitly, we can take

$$\omega_t = \frac{i}{2|\log(t)|^2} \sum_{j=1}^2 d \log z_j \wedge d \log \bar{z}_j.$$

The regularity hypothesis ensures that the genus and number of punctures of S_t are independent of the choice of ρ , and the wrapped Fukaya category $\mathcal{W}(S_t, \omega_t)$ is independent of $t \gg 0$ up to quasi-isomorphism. Let

$$\mathrm{Log}_t : (\mathbb{C}^*)^2 \rightarrow \mathbb{R}^2, \quad (z_1, z_2) \mapsto \frac{1}{|\log(t)|} (\log |z_1|, \log |z_2|).$$

As $t \rightarrow \infty$ the images $\mathrm{Log}_t(S_t)$ Gromov-Hausdorff converge to the “tropical amoeba”, the 1-dimensional polyhedral complex Π which is the singular locus of the Legendre transform of the convex function ρ , defined by

$$(4.1) \quad L_\rho : \mathbb{R}^2 \rightarrow \mathbb{R}, \quad \xi \mapsto \max \{ \langle a, \xi \rangle - \rho(a) \mid a \in A \}.$$

This is combinatorially the 1-skeleton of the dual cell complex of \mathcal{P} . The regions $R^a \subset \mathbb{R}^2 \setminus \Pi$ in the complement of the tropical curve Π are labelled by elements of A , according to which term in (4.1) achieves the maximum. Let

$$(4.2) \quad \Delta_{Z, \rho} = \{ (\xi, \eta) \in \mathbb{R}^2 \times \mathbb{R} \mid \eta \geq L_\rho(\xi) \},$$

let Z be the corresponding 3-dimensional toric variety, and $W : Z \rightarrow \mathbb{C}$ the function defined by the toric monomial $(0, 0, 1)$, which vanishes to order 1 along each component of $W^{-1}(0)$. The mirror to (S_t, ω_t) is the Landau-Ginzburg model (Z, W) , i.e. the symplectic geometry of S_t is reflected in the singularities of the toric divisor $W^{-1}(0) = Z_0$. For further discussion and context, see [5, Section 3]. In particular, we point out that the topology of Z depends on the choice of polyhedral decomposition \mathcal{P} , with different choices differing by flops. The irreducible toric divisors of Z are labelled by the components $R^a \subset \mathbb{R}^2 \setminus \Pi$ (whose closures are their moment polytopes), and their intersections are determined by the combinatorics of the tropical curve.

Remark 4.4. *By considering general elements of linear systems of curves in \mathbb{P}^2 of degree d or in $\mathbb{P}^1 \times \mathbb{P}^1$ of bidegree (a, b) , one obtains punctured surfaces of genus g with ℓ punctures for pairs (g, ℓ) of the form $((d-1)(d-2)/2, 3d)$ and $((a-1)(b-1), 2(a+b))$. More generally, one can obtain punctured surfaces for any (g, ℓ) with $3 \leq \ell \leq 2g+4$ by considering the family of tropical plane curves depicted in Figure 7.*

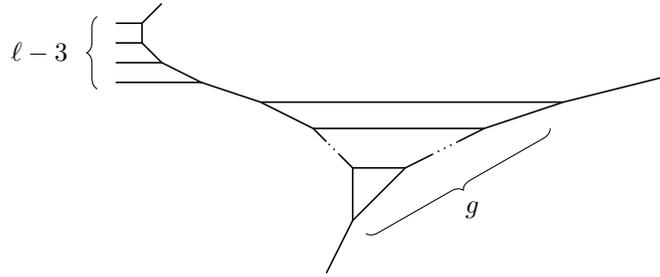


FIGURE 7. Tropical plane curves of genus g with ℓ punctures

For any scheme Z , we will write $\text{Perf}(Z)$ for the dg-category of perfect complexes over Z , which is the full subcategory of $D^b \text{Coh}(Z)$ of objects admitting locally free resolutions. The dg-quotient of the latter by the former is the derived category of singularities $D_{sg}(Z)$. This is not in general split-closed and its split-closure is denoted $D_{sg}^\pi(Z)$. The main result of [24] asserts:

Theorem 4.5 (Heather Lee). *There is an equivalence of $\mathbb{Z}/2$ -graded split-closed \mathbb{C} -linear triangulated categories*

$$D^\pi \mathcal{W}(S_t) \simeq D_{sg}^\pi(Z_0).$$

Lee's proof of this theorem involves writing both sides as limits over restrictions to certain simple pieces, and matching the two restriction diagrams in order to conclude that their limits are equivalent. On one hand, the tropical curve Π induces a decomposition of S_t into pairs of pants (indexed by the vertices of Π) glued together along cylinders (corresponding to the bounded edges of Π); on the other hand, the Landau-Ginzburg model (Z, W) admits a matching decomposition into toric affine charts $(\mathbb{C}^3, z_1 z_2 z_3)$. In Section 5.3, we will revisit the argument in order to apply analogous technology to a closed symplectic surface. For now, we note the following consequence:

Proposition 4.6. *Let $S'_t \subset S_t$ be the union of the pairs of pants corresponding to a given subset of the vertices of Π . Let Z' be the union of the corresponding affine charts of Z , i.e the toric 3-fold obtained from Z by removing all the toric strata whose closure does not contain any of the selected vertices, and $Z'_0 = Z' \cap Z_0$. Then $D^\pi \mathcal{W}(S'_t) \simeq D_{sg}^\pi(Z'_0)$.*

Proof. This follows immediately from Lee's proof of Theorem 4.5 by considering only the parts of the restriction diagrams that correspond to the pairs of pants and cylinders in S'_t on one hand, and to the affine charts of Z' on the other hand. The limits of these diagrams compute $D^\pi \mathcal{W}(S'_t)$ and $D_{sg}^\pi(Z'_0)$ respectively, which yields the result. \square

Remark 4.7. *Proposition 4.6 allows us to apply Lee’s result to punctured surfaces of arbitrary genus and with any number $\ell \geq 3$ punctures. For $\ell > 2g + 4$ the graphs of Figure 7 are not entire tropical plane curves because two of the legs would need to intersect each other outside of the depicted region of the plane; however they describe subsurfaces S'_t of higher genus curves $S_t \subset (\mathbb{C}^*)^2$, to which we can apply the Proposition. The mirror configuration Z'_0 is again a union of smooth toric divisors of Z' , as depicted in Figure 7; the only difference with the previous setting is that one of the components of Z'_0 (corresponding to the upper-right region of the figure) now arises as the complement of a toric divisor (e.g. a \mathbb{P}^1 with negative normal bundle) inside a compact component of Z_0 ; this does not have any incidence on the properties of the derived category of singularities, and the results below apply without modification to these examples.*

Lee’s proof furthermore matches certain specific objects on the two sides of the mirror.

Given any irreducible component $Z_0^a \subset Z_0$, and a line bundle $\mathcal{L}_a \rightarrow Z_0^a$, the push-forward of \mathcal{L}_a from Z_0^a to Z_0 defines an object of $D_{sg}(Z_0)$, which by abuse of notation we also denote by \mathcal{L}_a .

Proposition 4.8. *The equivalence of Theorem 4.5 matches the sheaves \mathcal{L}_a with properly embedded arcs or simple closed curves in S_t , hence with objects in the image of $\mathrm{Tw} \mathcal{W}(S_t) \rightarrow \mathrm{Tw}^\pi \mathcal{W}(S_t)$.*

Sketch. This follows from the proof of Theorem 4.5 as given in [24]. Specifically, \mathcal{L}_a is matched with an arc or curve which lies on the portion of S_t whose projection under Log_t collapses, as $t \rightarrow \infty$, to the boundary of the corresponding region $R^a \subset \mathbb{R}^2 \setminus \Pi$. The specific arc or curve is determined up to Hamiltonian isotopy, and hence quasi-isomorphism in $\mathcal{W}(S_t)$, by its winding over each of the cylindrical regions of S_t which collapse to finite edges of ∂R^a , and by a normalization condition; the winding numbers are determined explicitly by the degrees of the restriction of \mathcal{L}_a to the corresponding projective lines in Z_0^a (see [24, Section 3.1]). The details of the correspondence will not matter in the sequel. (The discussion in [24] concerns specifically those objects \mathcal{L}_a which arise from powers of the polarization determined by the polytope $\Delta_{Z,\rho}$, but the construction easily extends to general line bundles.) \square

4.4. Split-closure. Following a strategy from [4], this section will prove that the categories $\mathrm{Tw} \mathcal{W}(S)$ and $D_{sg}(Z_0)$ appearing in Theorem 4.5 are in fact already idempotent complete, i.e. split-closed. For a scheme Z , we write $K_j(Z)$ for $K_j(\mathrm{Perf}(Z))$.

Proposition 4.9. *$D_{sg}(Z)$ is idempotent complete if and only if $K_{-1}(Z) = 0$.*

Proof. This is [4, Proposition B.1]. \square

Lemma 4.10. *Let $W : Z \rightarrow \mathbb{C}$ be as after (4.2). Then $D_{sg}(Z_0)$ is split-closed.*

Proof. Recall $W : Z \rightarrow \mathbb{C}$ is a toric monomial morphism on a toric 3-fold Z with central fibre Z_0 a union of (not necessarily compact) toric surfaces. Let $\Gamma \subset Z_0$ denote the one-dimensional subscheme of singular points of Z_0 , and let $Z'_0 \rightarrow Z_0$ be the normalisation and

$\Gamma' = \Gamma \times_{Z_0} Z'_0 \subset Z'_0$. Concretely, Z'_0 is the disjoint union of the toric surfaces appearing in Z_0 , and Γ' is the union of their toric boundaries. There is an exact sequence

$$K_0(Z'_0) \oplus K_0(\Gamma) \rightarrow K_0(\Gamma') \rightarrow K_{-1}(Z_0) \rightarrow K_{-1}(Z'_0) \oplus K_{-1}(\Gamma) \rightarrow K_{-1}(\Gamma')$$

(cf. the proof of [4, Proposition B.2], which is itself inspired by [40]). In our case, Z'_0 is a union of smooth surfaces, so $K_{-1}(Z'_0) = 0$. We claim that

- (1) the map $K_{-1}(\Gamma) \rightarrow K_{-1}(\Gamma')$ is an isomorphism;
- (2) the map $K_0(Z'_0) \oplus K_0(\Gamma) \rightarrow K_0(\Gamma')$ is surjective.

By [40, Lemma 2.3], for the curves Γ and Γ' (whose irreducible components are all \mathbb{P}^1 's and \mathbb{A}^1 's) the K_{-1} -group is $\mathbb{Z}^{b_1(\bullet)}$ where $b_1(\bullet)$ denotes the first Betti number of the curve. Suppose S_t has genus g . Then $b_1(\Gamma) = g = b_1(\Gamma')$, and the natural map $\Gamma' \rightarrow \Gamma$ identifies the corresponding cycles of \mathbb{P}^1 's, which implies the first statement.

For the second statement, recall that Γ' is the disjoint union of the toric boundaries $\Gamma^a = \Gamma' \cap Z_0^a$ of the components of Z_0 . For any non-compact component $Z_0^a \subset Z_0$, the map $K_0(Z_0^a) \rightarrow K_0(\Gamma^a)$ is surjective. There are g compact components of Z_0 , on each of which the corresponding map has rank two cokernel. This is essentially a cohomological computation, since the relevant K_0 -groups for rational curves and toric surfaces are given by ranks of cohomology. By classical toric geometry, there is an exact sequence

$$H^2(Z_0^a) \rightarrow H^2(\Gamma^a) \rightarrow \mathbb{Z}^2 \rightarrow 0,$$

where the second map sends each component of Γ^a to the primitive normal vector of the corresponding facet of the moment polytope. The cokernel of the first map is therefore generated by any two irreducible toric divisors whose corresponding normal vectors form a basis of \mathbb{Z}^2 : for instance, by the Delzant condition, any two irreducible toric divisors which meet in one point.

Given this, an easy inductive argument shows that the total map from $K_0(Z'_0) \oplus K_0(\Gamma)$ to $K_0(\Gamma') = \bigoplus_a K_0(\Gamma^a)$ is surjective. Namely, pick an ordering of the components of Z_0 such that, for each compact component $Z_0^a \subset Z_0$, there exist two components of Γ^a whose normal vectors generate \mathbb{Z}^2 and which arise as intersections of Z_0^a with two other components of Z_0 which appear before it in the chosen ordering. (For example, order the components by scanning $\mathbf{R}^2 \setminus \Pi$ from top to bottom: then the edges meeting at a top-most vertex of a compact component have the requisite property.) We then show that the map is surjective onto each summand $K_0(\Gamma^a)$ by induction on a : for non-compact components the map $K_0(Z_0^a) \rightarrow K_0(\Gamma^a)$ is surjective, and for compact components our assumption yields two generators of $K_0(\Gamma)$ whose images, after quotienting by the previously encountered summands of $K_0(\Gamma')$, span the cokernel of $K_0(Z_0^a) \rightarrow K_0(\Gamma^a)$. This implies surjectivity. \square

Lemma 4.11. *The sheaves \mathcal{L}_a of Proposition 4.8 generate $D_{sg}(Z_0)$.*

Proof. For each irreducible component Z_0^a of Z_0 , line bundles over the toric surface Z_0^a generate its derived category $D^b \text{Coh}(Z_0^a)$ [11, Corollary 4.8], so their images under inclusion generate the full subcategory $D_{Z_0^a}^b \text{Coh}(Z_0)$ of complexes whose cohomology is supported on the component Z_0^a . Considering all components, these sheaves taken together generate $D^b \text{Coh}(Z_0)$; the result follows. \square

Corollary 4.12. *Let S be a surface with $\ell \geq 3$ punctures. Then the category of $\mathbb{Z}/2$ -graded twisted complexes $\mathrm{Tw} \mathcal{W}(S)$ is split-closed.*

Proof. The hypotheses imply that S can be realised as a hypersurface in $(\mathbb{C}^*)^2$ defined by a Laurent polynomial as in the setting of Theorem 4.5 and Remark 4.4, or as a subsurface as in Proposition 4.6 and Remark 4.7. Lemma 4.10 (which applies equally well to the examples of Remark 4.7) shows that Theorem 4.5 in fact gives an equivalence $\mathrm{Tw}^\pi \mathcal{W}(S) \simeq D_{sg}(Z_0)$. Furthermore, the right-hand side is generated (and not just split-generated) by objects which lie in $\mathrm{Tw} \mathcal{W}(S)$, by Proposition 4.8 and Lemma 4.11. It follows that $\mathrm{Tw} \mathcal{W}(S)$ is split-closed. \square

This incidentally shows that, for such a surface S , the category $D^\pi \mathcal{W}(S)$ has finite rank Grothendieck group; the corresponding result is false for a closed elliptic curve.

Corollary 4.13. *Let S be a surface of genus g with $\ell \geq 3$ punctures. Any irreducible object $X \in \mathcal{W}(S)^{per}$ with finite-dimensional endomorphism ring is quasi-isomorphic to a union of immersed closed curves with finite rank local system.*

Proof. By Corollary 4.12, X is quasi-isomorphic to a twisted complex. The geometricity result, Theorem 4.2, then says that X is quasi-isomorphic to the direct sum of some immersed arcs and immersed closed curves with finite rank local systems. However, non-compact arcs have infinite-dimensional endomorphisms in the wrapped category, so cannot appear. \square

Corollary 4.14. *Let S be a surface with at least one puncture. If $X \in D^\pi \mathcal{F}(S)$ is a spherical object with primitive Chern character, then X is quasi-isomorphic to a simple closed curve with rank one local system.*

Proof. For surfaces with $\ell \geq 3$ punctures, this follows directly from Corollary 4.13 and Corollary 3.16. If S has fewer than three punctures, we reduce to the previous case by considering the surface S^+ obtained by attaching a 4-punctured sphere P to S along an annular neighborhood A of a puncture, $S^+ = S \cup_A P$. There is a fully faithful inclusion functor $\mathcal{F}(S) \rightarrow \mathcal{F}(S^+)$ which comes from viewing S as a subsurface of S^+ and observing that none of the holomorphic polygons contributing to the A_∞ -operations can escape into $S^+ \setminus S$ (due to the open mapping principle). This extends to a fully faithful functor $\mathcal{F}(S)^{per} \rightarrow \mathcal{F}(S^+)^{per}$, and we can view X as a spherical object of the latter category. The inclusion into S^+ also preserves the property that $ch(X)$ is primitive. Since S^+ has at least 3 punctures, we conclude that X is quasi-isomorphic to a simple closed curve γ in S^+ with a rank one local system ξ . Given any properly embedded arc η contained in $S^+ \setminus S$ (e.g. connecting two punctures of P), the vanishing of the wrapped Floer cohomology $HW^*(X, \eta)$ implies that the geometric intersection number of γ with η is zero; this in turn implies that γ can be isotoped away from $S^+ \setminus S$. Alternatively, Lemma 5.10 below (applied to the object $(\gamma, \xi) \in \mathcal{F}(S^+)$ and the waist curve of the annulus A) implies that γ is isotopic to a simple closed curve in the completion of S . Either way, we conclude that X is quasi-isomorphic to (γ, ξ) in $D^\pi \mathcal{F}(S)$. \square

To extend this result to closed surfaces, we will use equivariant Floer theory and restriction functors to subsurfaces to prove that a spherical object on a closed surface in fact comes from some open subsurface.

5. RESTRICTION TO OPEN SUBSURFACES

The first three subsections below review material from [24], which is subsequently applied in our setting.

5.1. Dipping Hamiltonians. Let Σ be a surface (closed or with punctures) and $\sigma \subset \Sigma$ a simple closed curve. (Given a finite union of disjoint curves σ_j one can consider the corresponding Hamiltonians H_k which dip near each; in an abuse of notation we will continue to write A for the union of annular neighbourhoods of the σ_j , and refer to A as an annulus.) Let $\sigma \subset A \subset \Sigma$ be an annular neighbourhood of σ . Following [24], we consider a sequence of functions

$$H_k : \Sigma \rightarrow \mathbb{R}$$

which are small perturbations of the constant (say zero) function on $\Sigma \setminus A$ but “dip” inside the annulus. We work in co-ordinates $(r, \theta) \in (-2, 2) \times S^1 = A$ such that the symplectic form is $\omega = c dr \wedge d\theta$ for some constant $c > 0$, and define $H_k(r, \theta) = c f_k(r)$ where

$$f_k(r) = \begin{cases} -k.\pi(r+2)^2 & -2 < r < -1 \\ k.\pi r^2 - 2k.\pi & -1 \leq r \leq 1 \\ -k.\pi(r-2)^2 & 1 < r < 2 \end{cases}$$

The time-1 Hamiltonian flow of H_k lifted to the universal cover $(-2, 2) \times \mathbb{R}$ of the annulus A is then given by

$$\phi_{H_k}(r, \theta) = \begin{cases} (r, \theta - 2k\pi.(r+2)) & -2 < r < -1 \\ (r, \theta + 2k\pi.r) & -1 \leq r \leq 1 \\ (r, \theta - 2k\pi.(r-2)) & 1 < r < 2 \end{cases}$$

Since $f_k = k.f_1$, the time-1 flow $\phi_{H_k}^1$ of H_k is exactly the time- k flow of H_1 , and indeed there is a well-defined time- t flow $\Phi_{H_1}^t$ for non-integer times t which interpolates between the time-1 flows of the H_k . If $\gamma \subset A$ is an arc $\{\theta = \text{constant}\}$ crossing the annulus, the time-1 Hamiltonian flow of H_k applied to γ yields an arc which wraps k times clockwise around A , then $2k$ times anticlockwise, and then k times clockwise again. (Note that clockwise corresponds to negative Reeb flow and anticlockwise to positive Reeb flow.) It will be important later to divide

$$A = A_{in} \cup A_{out}$$

into the inner region $A_{in} = (-1, 1) \times S^1$ in which the wrapping is by positive Reeb flow, and the outer region $A_{out} = ((-2, -1) \sqcup (1, 2)) \times S^1$ in which the wrapping is negative.

For a pair of Lagrangians $L, L' \subset (X, \omega)$, and a Hamiltonian $H : X \rightarrow \mathbb{R}$, denote by

$$CF^*(L, L'; H) := CF^*(\phi_H^1(L), L')$$

given by flowing L by the time-1 flow of H ; the group is generated by time-1 chords of H from L to L' , or equivalently by intersections of $\phi_H^1(L)$ and L' . For a given finite collection of

Lagrangians, a generic small perturbation of H (which we shall suppress from the notation) will make all such chords non-degenerate. In the setting at hand, given a pair of distinct arcs $\gamma_0, \gamma_1 \subset \Sigma$ which both cross the annulus A , the set of intersections $\phi_{H_k}^1(\gamma_0) \cap \gamma_1$ will grow in size with k , as more and more intersections appear in the “wrapping” regions.

Lemma 5.1. *There is $n(\gamma_0, \gamma_1) > 0$ with the following property: for any integer $w > n(\gamma_0, \gamma_1)$, any point $p \in \phi_{H_w}^1(\gamma_0) \cap \gamma_1$ belongs to a unique smooth arc*

$$(5.1) \quad [w, \infty) \rightarrow \Sigma, \quad t \mapsto p(t) \in \phi_{H_t}^1(\gamma_0) \cap \gamma_1$$

of transverse intersections between the time $t \geq w$ flow by H_1 of γ_0 and γ_1 .

Proof. See [24, Section 3.5, properties (P1,2)]. □

Lemma 5.1 means that, once any pair of arcs has been sufficiently wrapped, their intersection points persist (and remain transverse) for all further time, even though new intersections keep being created (at non-integer times). By using a cascade model for continuation maps of Floer complexes, as in [28, Section 10e], in which one counts exceptional holomorphic discs and flow-trees for isolated times (J_t, H_t) in a one-parameter family of almost complex structures and Hamiltonians, Lee infers:

Lemma 5.2. *For $N > n > n(\gamma_0, \gamma_1)$, the continuation map*

$$CF^*(\gamma_0, \gamma_1; H_n) \rightarrow CF^*(\gamma_0, \gamma_1; H_N)$$

maps each generator p to the summand generated by the intersection point which lies on the smooth arc $p(t)$ of (5.1).

Similarly, because of the non-existence of exceptional holomorphic discs on a Riemann surface, a cascade model for higher continuation maps shows that continuation-type products

$$CF^*(\gamma_{k-1}, \gamma_k; H_n) \otimes \cdots \otimes CF^*(\gamma_0, \gamma_1; H_n) \rightarrow CF^*(\gamma_0, \gamma_k; H_N)[1 - k]$$

with $k \geq 2$ inputs vanish whenever $N > kn > n(\gamma_0, \dots, \gamma_k)$ is sufficiently large. This leads to a well-defined A_∞ -inclusion functor $CF^*(\gamma_0, \gamma_1; H_n) \hookrightarrow CF^*(\gamma_0, \gamma_1; H_N)$, for $N > n$ sufficiently large, which has no higher-order (nonlinear) terms. It is simplest to formalise this by passing to telescope models for wrapped Floer complexes.

5.2. Telescope models and A_∞ -ideals. The telescope complex for exact manifolds comes from [6], and a detailed exposition in the monotone case (under geometric hypotheses which also apply in the case of a closed surface) is given in [26]. We will incorporate an action-rescaling of generators of Floer complexes, so we briefly review the set-up; for simplicity we suppress local systems, which are discussed in [26, Section 3.17].

Let X be closed or convex at infinity, and fix a pair of Lagrangian branes $L_i, L_j \in \mathcal{W}(X)$, which might be compact or cylindrical at infinity. If outside a compact set $(X, \omega) \cong (\partial X \times [1, \infty), d(r \cdot \alpha))$ for a contact form $\alpha \in \Omega^1(\partial X)$ and co-ordinate $r \in [1, \infty)$, we will fix a Hamiltonian H with $H(y, r) = r$, and which has no integer-length chords from L_i to L_j . The telescope model for the wrapped Floer complex is then

$$(5.2) \quad CW^*(L_i, L_j) = \bigoplus_{w=1}^{\infty} CF^*(L_i, L_j; wH)[\mathbf{q}]$$

where \mathbf{q} is a formal variable of degree -1 satisfying $\mathbf{q}^2 = 0$, equipped with the differential

$$(5.3) \quad \mu^1(x + \mathbf{q}y) = (-1)^{|x|} \mathfrak{d}x + (-1)^{|y|} (\mathbf{q}\mathfrak{d}y + \mathfrak{K}y - y)$$

where \mathfrak{d} denotes the usual Floer boundary operator, and where \mathfrak{K} denotes the Floer continuation map

$$\mathfrak{K} : CF^*(L_i, L_j; wH) \rightarrow CF^*(L_i, L_j; (w+1)H).$$

The part of this complex which does not involve \mathbf{q} is the direct sum of Floer complexes $CF^*(L_i, L_j; wH)$, with the usual Floer differential. For a \mathfrak{d} -cocycle y , $\mathbf{q}y$ serves to identify y and $\mathfrak{K}y$ in cohomology, as expected in the cohomology direct limit

$$HW^*(L_i, L_j) = \varinjlim HF^*(L_i, L_j; wH)$$

Let $u : \mathbb{R} \times [0, 1] \rightarrow X$ be a non-constant isolated (modulo \mathbb{R} -translation) solution of Floer's equation $\partial_s u + J(\partial_t u - wX_H) = 0$, with Lagrangian boundary conditions $u(\cdot, 0) \in L_i$ and $u(\cdot, 1) \in L_j$, and asymptotic conditions $u \rightarrow x, y$ as $s \rightarrow -\infty, +\infty$ respectively, where x, y are chords of X_H . Usually the differential \mathfrak{d} counts isolated such solutions $u : S \rightarrow X$ weighted by their "topological energy" $E_{\text{top}}(u) \geq 0 \in \mathbb{R}$, defined by:

$$E_{\text{top}}(u) = \int_S u^* \omega - d(u^*(wH) dt) = \int_S u^* \omega + wH(x) - wH(y),$$

which is also equal to the geometric energy

$$E_{\text{geo}}(u) = \frac{1}{2} \int_S \|du - wX_H \otimes dt\|^2 ds \wedge dt.$$

(There is also an orientation sign, which we will not discuss: for comprehensive treatments see [6, 1].) Similarly the continuation map \mathfrak{K} counts isolated solutions to the corresponding equation where J and wH (in fact, in our case just w) depend on the parameter s (so there is no \mathbb{R} -reparametrization), again weighted by topological energy (which now provides an upper bound on the geometric energy, up to an additive constant given by the minimum of the Hamiltonian). Similar remarks apply to the counts of isolated curves contributing to the rest of the A_∞ -structure, cf. [26, Section 3.11–15].

Remark 5.3. *Floer's equation $\partial_s u + J(\partial_t u - wX_H) = 0$ for maps $u : \mathbb{R} \times [0, 1] \rightarrow X$ with boundary on L_i and L_j can be recast in terms of pseudo-holomorphic curves with boundary on $\phi_{wH}^1(L_i)$ and L_j , by considering $\tilde{u}(s, t) = \phi_{wH}^{1-t}(u(s, t))$. Continuation maps are then naturally defined via pseudo-holomorphic curves with moving boundary conditions, or (under suitable conditions), as mentioned in the previous section, cascades of honest pseudo-holomorphic curves. The appropriate notion of topological energy is then obtained by transcribing the above definition through the dictionary between the two viewpoints.*

Specialise now to the case where $X = \Sigma$ is a surface and we have functions $H_w = w \cdot H$ for a Hamiltonian which dips in some collection of annuli, as in the previous section. For a given collection of Lagrangians γ_j , Lemma 5.2 asserts that the continuation maps in the telescope complex (5.2), (5.3) are eventually diagonal inclusions of based vector spaces. To make them inclusions of subcomplexes, for a Hamiltonian H_w -chord $x \subset A$ lying in the cylinder at radial parameter $r \in (-2, 2)$, introduce the action values

$$\mathcal{A}_w(x) = H_w(r) - r \cdot H'_w(r)$$

(i.e. the intercept of the tangent line to the graph of the dipping function at r with the vertical axis).

Lemma 5.4. *Consider the rescaled bases of Floer complexes $CF^*(\gamma_j, \gamma_k; H_w)$ in which the chords x inside the cylindrical regions A are rescaled by their action $\mathcal{A}_w(x)$. Then for N sufficiently large, the continuation maps \mathfrak{K} are inclusions for $w > N$.*

Proof. Given a perturbed holomorphic polygon $u : S \rightarrow \Sigma$ with input chords $x_1^{in}, \dots, x_j^{in}$ with weights $w_1^{in}, \dots, w_j^{in}$ and output chord x^{out} with weight w^{out} , and supposing the image $u(S)$ is contained in a single connected component of A , then

$$E_{top}(u) = \mathcal{A}_{w^{out}}(x^{out}) - \sum_{i=1}^j \mathcal{A}_{w_i^{in}}(x_i^{in}).$$

The Floer solution u contributes to the usual continuation map with weight $q^{E_{top}(u)}$, which means that it takes rescaled generators to rescaled generators. \square

The key property of dipping Hamiltonians, explained in [24, Section 3.8], is that they lead to a “localisation” result for holomorphic polygons, which precludes non-trivial Floer products from crossing the inner region $A_{in} \subset A$ of the annulus in which one wraps. To state this precisely, denote by S the surface obtained from $\Sigma \setminus A_{in}$ by reattaching a copy of the corresponding annular region A_{in} to each end (see Definition 5.6 below): S is a (possibly disconnected) surface with boundary, and there is a natural surjective map $S \rightarrow \Sigma$ which is two-to-one over A_{in} and one-to-one outside of A_{in} .

Fix $\gamma_0, \dots, \gamma_k \subset \Sigma$ which intersect the boundary of A_{in} minimally; equivalently, we require that the preimages of γ_i in S do not contain any trivial arcs connecting ∂S to itself. Assume moreover that, if γ_i is parallel to the annulus A_{in} (i.e., isotopic to a curve contained in A_{in}), then it is actually contained in A_{in} . Also fix a collection of Floer generators p_0, \dots, p_{k-1} with $p_i \in CF^*(\gamma_i, \gamma_{i+1}; H_n)$. By Lemma 5.1, there are unique corresponding generators $\kappa(p_i) \in CF^*(\gamma_i, \gamma_{i+1}; H_N)$ for any $N > n$, once n is sufficiently large.

Proposition 5.5. *Fix a collection of Floer generators $\{p_0, \dots, p_{k-1}\}$ as above.*

- (1) *For $N > n$ sufficiently large, every rigid perturbed holomorphic polygon in Σ with inputs $\kappa(p_i) \in CF^*(\gamma_i, \gamma_{i+1}; H_N)$ lifts to a perturbed holomorphic polygon inside (a single component of) S .*
- (2) *For $N > n$ sufficiently large, if a rigid perturbed holomorphic polygon in Σ with inputs $\kappa(p_i) \in CF^*(\gamma_i, \gamma_{i+1}; H_N)$ is not entirely contained in A_{in} (for example if at least one p_i lies outside of A_{in}) then its output also lies in the complement of A_{in} .*

Proof. These follow respectively from [24, Lemma 3.7], which states that rigid perturbed holomorphic polygons cannot cross the inner annular region A_{in} , and [24, Lemma 3.5], which states that if part of the disc lies outside of A_{in} then so does the output. Lee states and proves these results for a specific collection of objects of $\mathcal{F}(\Sigma)$ (and [24, Lemma 3.5] is stated under the stronger assumption that one input lies outside of A_{in}), but the same arguments apply verbatim to our setting, as we now explain.

The proof of [24, Lemma 3.5] considers a rigid holomorphic polygon whose output lies in A_{in} , and shows by contradiction that it must be entirely contained in A_{in} using a two-step argument. Lee first shows (“Case 1” of the proof in [24]) that the boundary of the polygon cannot backtrack in $A \setminus A_{in}$; thus, if the polygon is not entirely contained in A_{in} , some part of it must lie outside of A . Our assumptions on the γ_i ensure that, inside $A \setminus A_{in}$, they are locally graphs $r = f(\theta)$, just like the curves considered in [24]; this (and the ordering on the slopes of these graphs imposed by the negative wrapping) is what prevents the backtracking. Next, Lee argues (“Case 2” of the proof) that the portion of the polygon which lies outside of A determines the width of the interval(s) along which it intersects ∂A , and as the amount of negative wrapping in $A \setminus A_{in}$ increases, the width of the corresponding interval(s) along which the polygon intersects ∂A_{in} decreases and eventually becomes negative, preventing it from entering A_{in} altogether. Our assumption that all γ_i intersect ∂A_{in} minimally ensures that no portion of disc crossing into A can look like a strip with both of its boundaries on the same γ_i ; this in turn ensures that negative wrapping does indeed decrease the available width at the boundary of A_{in} . Moreover, excluding curves which are parallel to the annulus but lie outside of A_{in} ensures that the width of the intersection of the polygon with ∂A is completely determined by the collection of input generators in $\Sigma \setminus A$. (By contrast, polygons with corners at fixed generators on an annulus-parallel curve outside A_{in} can enter the annulus with an arbitrarily large width, as the boundary of the polygon could wrap around the annulus-parallel curve arbitrarily many times.)

The proof of [24, Lemma 3.7] uses similar considerations to show that a rigid holomorphic polygon cannot cross completely through A_{in} . First, Lee observes that, when the γ_i are locally given by graphs $r = f(\theta)$, due the ordering of the slopes imposed by the positive wrapping inside A_{in} , any input contained in A_{in} forces the boundary of the polygon to backtrack, which prevents it from crossing from one end of A_{in} to the other. This remains true if we allow curves $\{r = r_0\}$ contained in A_{in} , as the local convexity property of rigid holomorphic polygons ensures that if part of the boundary of the polygon lies on such a curve then it must backtrack. The rest of the argument is then similar to the proof of [24, Lemma 3.5]: if a polygon crosses A_{in} , then regardless of whether A_{in} contains the output, its width at the center of A_{in} is determined by the inputs on one side of A_{in} or on the other, and is not affected by the wrapping, while the available width at one of the two boundaries of A_{in} decreases with the amount of wrapping and eventually becomes negative. \square

5.3. Lee restriction functors. Fix a finite set $\gamma_1, \dots, \gamma_\ell$ of split-generators of $\mathcal{F}(\Sigma)$ each of which meets σ minimally, and none of which is parallel to σ (unless it is σ itself). The category $\mathcal{F}(\Sigma)$ embeds into the category of modules over the A_∞ -algebra $\oplus_{i,j} CF^*(\gamma_i, \gamma_j; H_k)$ for any fixed k , and hence into modules over the telescope algebra $\oplus_{i,j} CW^*(\gamma_i, \gamma_j)$.

Fix an annular neighbourhood $\sigma \subset A \subset \Sigma$ of a simple closed curve; as in the previous section we regard $\sigma = \{0\} \times S^1 \subset (-2, 2) \times S^1 = A$ as divided into three sub-annuli, the outer of which comprise A_{out} and the inner of which is labelled A_{in} . As above, we construct a surface S by reattaching a copy of A_{in} to each end of $\Sigma \setminus A_{in}$:

Definition 5.6.

- If σ is separating, we view Σ as a union of two subsurfaces with non-empty boundary which overlap in the central “positive wrapping” region $A_{in} \subset A$, i.e. the subannulus

$(-1, 2) \times S^1$ lies inside one subsurface and $(-2, 1) \times S^1$ lies inside the other. Write $S = S_{left} \sqcup S_{right}$ for the disjoint union of these subsurfaces.

- If σ is non-separating, we define S to be a surface with two boundary components, neighborhoods of which are respectively modelled on $(-2, 1) \times S^1 \subset A$ and $(-1, 2) \times S^1 \subset A$. Thus, S is not strictly a subsurface of Σ , but there is an obvious map $S \rightarrow \Sigma$ which is two-to-one over the central subannulus $A_{in} \subset A$.

Given any finite collection of disjoint simple closed curves in Σ and annular neighbourhoods of those curves, there are associated open surfaces given by cutting along the annuli in a similar manner. For our purposes, it is sufficient to focus on the case of a single annulus.

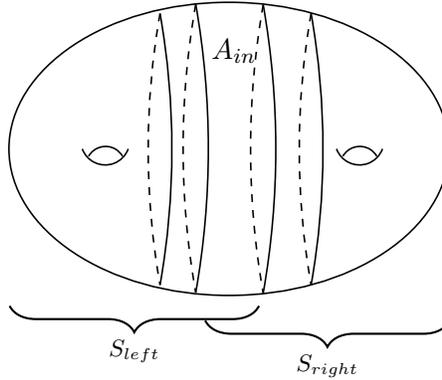


FIGURE 8. Restriction to subsurfaces

Note that the Hamiltonians H_w considered previously induce the “standard” positive wrapping at the ends of S and A_{in} , hence are appropriate for defining their wrapped categories.

Lemma 5.7. *Let $A \subset \Sigma$ be a single annulus. There is an A_∞ -restriction functor $\Phi_{\Sigma, S} : \mathcal{F}(\Sigma) \rightarrow \mathcal{W}(S)$, and a pair of restriction functors $\Phi_{S, A_{in}}^{left/right} : \mathcal{W}(S) \rightarrow \mathcal{W}(A_{in})$ coming from the inclusions of A_{in} into the two ends of S .*

Proof. It suffices to construct the functors on a finite set of split-generators γ_i ; recall that we do not allow any of these to be parallel to the annulus (but see Example 5.9 below for a description of the functor on annulus-parallel curves). We take telescope models for the respective morphism groups (even for the compact surface Σ) and note that, for any given set of inputs, replacing the telescope complex by a quasi-isomorphic truncation allows us to only consider arbitrarily large wrapping Hamiltonians.

Proposition 5.5 implies that the generators which lie outside of A_{in} form an ideal for the A_∞ -operations in either Σ or S , and that quotienting the Floer complex by this ideal recovers the wrapped Floer complex inside A_{in} and its A_∞ -operations. Thus the quotient maps

$$\Phi_{\Sigma, A_{in}} : CW_{\Sigma}^*(\gamma_i, \gamma_j) \rightarrow CW_{A_{in}}^*(\gamma_i, \gamma_j) \quad \text{and} \quad \Phi_{S, A_{in}}^{left/right} : CW_S^*(\gamma_i, \gamma_j) \rightarrow CW_{A_{in}}^*(\gamma_i, \gamma_j),$$

where the subscripts denote the surfaces in which we compute Floer cohomology, define restriction A_∞ -functors with no higher order terms.

Next, we consider the pullback map

$$\Phi_{\Sigma, S} : CW_{\Sigma}^*(\gamma_i, \gamma_j) \rightarrow CW_S^*(\gamma_i, \gamma_j)$$

mapping each Floer generator to its preimage under the map $S \rightarrow \Sigma$ if it lies outside of A_{in} , or to the sum of its two preimages if it lies in A_{in} . Proposition 5.5 implies that this map is compatible with the A_{∞} -operations (again after restricting to sufficiently large Hamiltonians), and defines a A_{∞} -functor with no higher order terms, identifying the wrapped Floer complex in Σ with the sub-algebra of the wrapped Floer complex in S consisting of elements in which pairs of generators in the two copies of $A_{in} \subset S$ appear with the same coefficients. \square

Corollary 5.8. $\mathcal{F}(\Sigma)^{per}$ is the pullback in the diagram of restriction functors

$$\mathrm{Tw}^{\pi} \mathcal{F}(\Sigma) \longrightarrow \mathrm{Tw}^{\pi} \mathcal{W}(S) \rightrightarrows \mathrm{Tw}^{\pi} \mathcal{W}(A_{in}).$$

Proof. This follows directly from Lemma 5.7. After truncating the telescope complexes to only consider Floer complexes $CF^*(\gamma_i, \gamma_j; H_n)$ where the wrapping n is sufficiently large, $CW_{\Sigma}^*(\gamma_i, \gamma_j)$ is exactly the subcomplex of $CW_S^*(\gamma_i, \gamma_j)$ of elements which restrict compatibly under the two restrictions to $A_{in} \subset S$, i.e. the equalizer of the diagram $\mathcal{W}(S) \rightrightarrows \mathcal{W}(A_{in})$. This description is compatible with the A_{∞} -operations since the arrows are given by quotienting by an ideal. \square

Example 5.9. Let a and b be two simple closed curves which cross an annulus A , meeting once just outside the annulus, and differ from each other by a Dehn twist parallel to the annulus (cf. the first part of Figure 9). We consider the twisted complex $X = a \xrightarrow{p} b$, which is isomorphic to a simple closed curve γ' running parallel to the annulus, as can be seen by considering the Lagrange surgery of a and b at p (the red curve in Figure 9). (This is the most natural way in which an annulus-parallel curve can be replaced by a twisted complex to which the machinery of Lee restriction functors can be applied). The Lee deformation effects the change in the second part of Figure 9, which gives a quasi-isomorphic model

$$X \simeq a \xrightarrow{p_{main} + p_{cyl} + p'_{cyl}} b.$$

Indeed, the deformation creates two holomorphic strips contributing to $\partial(p_{main})$, and the image of p under the continuation map is $\kappa(p) = p_{main} + p_{cyl} + p'_{cyl}$ (up to a rescaling of the generators). Restricting a and b to the inner annulus A_{in} yields two isotopic arcs, and $\kappa(p)$ restricts to $p_{cyl} + p'_{cyl}$, so $\Phi_{\Sigma, A_{in}}(X) = \mathrm{Cone}(p_{cyl} + p'_{cyl}) \in \mathrm{Tw} \mathcal{W}(A_{in})$ is isomorphic to a simple closed curve γ'_{in} in the Liouville completion of A_{in} . Just like the simple closed curve γ' that represents X in Σ , this curve may or may not actually lie inside A_{in} ; in the latter case, the difference with the naive restriction of γ' illustrates the need to exclude such annulus-parallel curves in the construction of Lee restriction functors.

The curves a and b restrict (or rather lift) to arcs in the open surface S , representing isomorphic objects of $\mathcal{W}(S)$. Denoting by p_{left} and p_{right} , resp. p'_{left} and p'_{right} the two preimages of p_{cyl} , resp. p'_{cyl} in the ends of the open surface S , the twisted complex $\Phi_{\Sigma, S}(X)$ is the mapping cone of $p_{main} + p_{left} + p'_{left} + p_{right} + p'_{right}$, which is isomorphic to the direct sum of two boundary-parallel simple closed curves in the completion of S . (If γ' lies inside A_{in} then these are the lifts of γ' to S ; otherwise one curve is the preimage of γ' and the other one

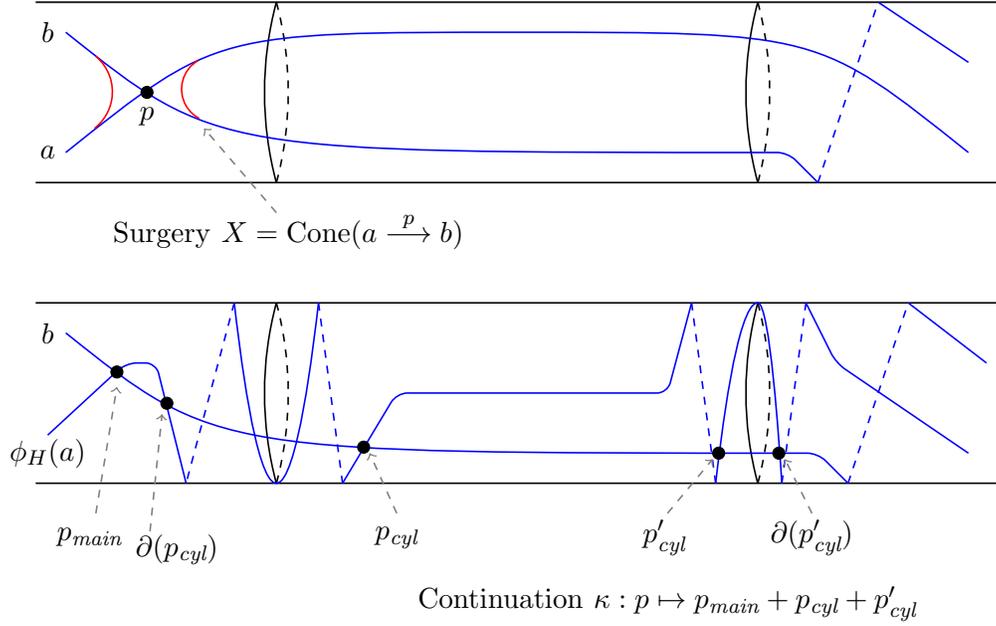


FIGURE 9. Continuation maps associated to dipping Hamiltonians

lies past the boundary of S inside the other end of the completion). Further restriction to either copy of A_{in} yields the mapping cone of $p_{\text{cyl}} + p'_{\text{cyl}}$ in $\text{Tw } \mathcal{W}(A_{\text{in}})$, as expected.

By contrast, if a and b did not run parallel to each other in $\Sigma \setminus A$, so that no holomorphic strip connected p_{main} to $\partial(p'_{\text{cyl}})$ in the second part of Figure 9, then $\kappa(p)$ would be $p_{\text{main}} + p_{\text{cyl}}$, whose restriction to A_{in} is p_{cyl} , a quasi-isomorphism between the restrictions of a and b ; $\Phi_{\Sigma, A_{\text{in}}}(X)$ would then be acyclic. This is as expected, since in that case the Lagrange surgery of a and b at p can be pulled through A_{in} and away from it altogether. Meanwhile, the twisted complex $\Phi_{\Sigma, S}(X)$ becomes the mapping cone of $p_{\text{main}} + p_{\text{left}} + p_{\text{right}}$, where p_{left} and p_{right} are the preimages of p_{cyl} in the ends of S ; it can be checked that this is isomorphic to a simple closed curve in $\Sigma \setminus A_{\text{in}} \subset S$, obtained from a and b by performing the surgery in Σ and sliding the result through A_{in} and away from it.

5.4. Applications of restriction. We continue with the notation of the previous subsection. Given a simple closed curve $\gamma \subset \Sigma$, we consider an annulus A with waist curve γ , and the Lee restriction functors $\Phi_{\Sigma, S} : \mathcal{F}(\Sigma) \rightarrow \mathcal{W}(S)$, and similarly $\Phi_{S, A_{\text{in}}}^{\text{left/right}}$ and $\Phi_{\Sigma, A_{\text{in}}}$.

Lemma 5.10. *Let $Y \in \text{Tw } \mathcal{F}(\Sigma)$. If $HF^*(Y, \gamma) = 0$ then $\Phi_{\Sigma, S}(Y)$ and $\Phi_{\Sigma, A_{\text{in}}}(Y)$ are isomorphic to direct sums of immersed closed curves with local systems in the completions of S and A_{in} .*

Proof. As in Section 5.3, we consider the equalizer diagram

$$CW_{\Sigma}^*(Y, \gamma) \longrightarrow CW_S^*(\Phi_{\Sigma, S}(Y), \Phi_{\Sigma, S}(\gamma)) \rightrightarrows CW_{A_{\text{in}}}^*(\Phi_{S, A_{\text{in}}}^{\text{left/right}}(Y), \Phi_{S, A_{\text{in}}}^{\text{left/right}}(\gamma)),$$

and recall that the restriction to A_{in} is the quotient of $CW_{\Sigma}^*(Y, \gamma)$ by the A_{∞} -ideal spanned by the generators which lie outside of A_{in} . Since $\gamma \subset A_{in}$, all the generators of $CW_{\Sigma}^*(Y, \gamma)$ lie in A_{in} , and Proposition 5.5 (2) ensures that the bigons contributing to the Floer differential are also entirely contained in A_{in} . Therefore, $CW_{A_{in}}^*(\Phi_{\Sigma, A_{in}}(Y), \Phi_{\Sigma, A_{in}}(\gamma))$ is isomorphic to $CW_{\Sigma}^*(Y, \gamma)$, hence acyclic. It then follows from the equalizer diagram that $CW_S^*(\Phi_{\Sigma, S}(Y), \Phi_{\Sigma, S}(\gamma))$ is also acyclic.

Next, we observe that $\Phi_{\Sigma, S}(\gamma) \simeq \gamma_{\text{left}} \cup \gamma_{\text{right}}$ is the disjoint union of the two boundary-parallel simple closed curves obtained by lifting γ to the open surface S . On the other hand, the restriction $\Phi_{\Sigma, S}(Y) \simeq \Gamma_Y$ has a geometric replacement Γ_Y in the Liouville completion of S , which is a union of immersed arcs and curves with local systems, by Theorem 4.2. The vanishing of $HF^*(Y, \gamma)$ implies that $CW^*(\Gamma_Y, \gamma_{\text{left}})$ and $CW^*(\Gamma_Y, \gamma_{\text{right}})$ are acyclic.

We claim that Γ_Y contains no arcs. Indeed, Assume that Γ_Y contains an immersed arc η which reaches the boundary component of S adjacent to γ_{left} . As noted in Remark 2.7, the argument of Lemma 2.6 constructs an immersed arc η' which is unobstructedly regular homotopic and hence quasi-isomorphic to η and intersects γ_{left} minimally, so that there are no bigons contributing to the Floer differential on $CF^*(\eta', \gamma_{\text{left}})$. However, the acyclicity of $CW^*(\Gamma_Y, \gamma_{\text{left}})$ implies that of $CF^*(\eta, \gamma_{\text{left}})$, since η is a direct summand in Γ_Y . Therefore $\eta' \cap \gamma_{\text{left}} = \emptyset$, and the arc η' is entirely contained in the cylindrical region that lies beyond γ_{left} in the completion of S , hence trivial as an object of $\mathcal{W}(S)$. Arcs which reach the other boundary component of S are similarly excluded by considering $CW^*(\Gamma_Y, \gamma_{\text{right}})$.

It follows that the geometric replacement $\Gamma_Y \in \mathcal{W}(S)$ is a union of immersed closed curves with local systems in the completion of S . Moreover, as noted at the end of the proof of Theorem 4.2, it can be assumed that all non boundary-parallel curves in Γ_Y are strictly contained in the interior of S , or more precisely, in the subsurface $\Sigma \setminus A_{in} \subset S$ (see Remark 2.9); whereas the boundary-parallel curves can be “straightened” to run parallel to γ .

Next we apply the restriction functors $\Phi_{S, A_{in}}^{\text{left/right}}$ to Γ_Y . The non boundary-parallel summands are represented by curves in $\Sigma \setminus A_{in}$, and hence mapped to zero. Meanwhile, the boundary-parallel curves which lie outside of A_{in} need to be resolved by mapping cones as in Example 5.9 in order to apply the machinery of Section 5.3; the upshot is that curves which are parallel to γ_{left} (resp. γ_{right}) are mapped by $\Phi_{S, A_{in}}^{\text{left}}$ (resp. $\Phi_{S, A_{in}}^{\text{right}}$) to closed curves in the completion of A_{in} . \square

Corollary 5.11. *Let $Y \in \text{Tw}\mathcal{F}(\Sigma)$. If $HF^*(Y, \gamma) = 0$ then Y is quasi-isomorphic to a direct sum of immersed closed curves with local systems in Σ .*

Proof. We start again from the geometric replacement $\Gamma_Y \in \mathcal{W}(S)$ of $\Phi_{\Sigma, S}(Y)$ constructed in the proof of the previous Lemma, which is a direct sum of immersed closed curves with local systems in the completion of S . Since $\Phi_{S, A_{in}}^{\text{left}}(\Gamma_Y) \simeq \Phi_{S, A_{in}}^{\text{right}}(\Gamma_Y) \in \mathcal{W}(A_{in})$, the boundary-parallel summands of Γ_Y are “the same” curves near the two boundaries of S .

Define $\hat{Y} \in \mathcal{F}(\Sigma)$ to be the *reduced* projection of Γ_Y from S to Σ , i.e. the direct sum of the non-boundary-parallel summands of Γ_Y (which by Remark 2.9 can be assumed to lie in $\Sigma \setminus A_{in}$), and for each pair of boundary-parallel summands of Γ_Y , a curve in Σ which runs parallel to the annulus A and differs from the waist curve γ by the same amount of symplectic

area. (Recall that “parallel” means “homotopic to a curve in A ”: if γ is separating, achieving the desired amount of symplectic area might require a non-embedded regular homotopy, cf. Lemmas 2.8 and 2.10). By construction, the restriction $\Phi_{\Sigma,S}(\hat{Y})$ (defined by first resolving the annulus-parallel summands of \hat{Y} as in Example 5.9) is isomorphic to Γ_Y . Thus, the two restriction diagrams

$$\Phi_{\Sigma,S}(Y) \rightrightarrows \Phi_{\Sigma,A_{in}}(Y) \quad \text{and} \quad \Phi_{\Sigma,S}(\hat{Y}) \rightrightarrows \Phi_{\Sigma,A_{in}}(\hat{Y})$$

involve quasi-isomorphic objects of $\mathcal{W}(S)$ and $\mathcal{W}(A_{in})$ (namely, Γ_Y and its restriction to A_{in}). In order to conclude from Corollary 5.8 that Y and \hat{Y} are quasi-isomorphic in $\text{Tw } \mathcal{F}(\Sigma)$, we need to verify that the restriction maps in the two diagrams are also the same.

More precisely, $\Phi_{\Sigma,S}(Y) \simeq \Gamma_Y$ determines a Yoneda module over $\mathcal{W}(S)$, whose pullback along $\Phi_{\Sigma,S}$ is a $\mathcal{F}(\Sigma)$ -module which we denote by \mathcal{Y}_S . Similarly, $\Phi_{\Sigma,A_{in}}(Y) \in \mathcal{W}(A_{in})$ determines a Yoneda module over $\mathcal{W}(A_{in})$, and we denote by $\mathcal{Y}_{A_{in}}$ its pullback along $\Phi_{\Sigma,A_{in}}$. The functors $\Phi_{S,A_{in}}^{\text{left/right}}$ induce two $\mathcal{F}(\Sigma)$ -module homomorphisms

$$(5.4) \quad \phi_Y^{\text{left/right}} : \mathcal{Y}_S \rightrightarrows \mathcal{Y}_{A_{in}},$$

and Corollary 5.8 states that the Yoneda module associated to Y is isomorphic to the equalizer of this diagram in the category of $\mathcal{F}(\Sigma)$ -modules.

Assume for now that all annulus-parallel curves in Γ_Y actually lie inside A_{in} (near both ends of S). Then, using the correspondence between holomorphic polygons in Σ and in S given by Proposition 5.5, the module $\mathcal{Y}_S \in \text{mod-}\mathcal{F}(\Sigma)$ can be represented by a direct sum \tilde{Y}_S of curves with local systems in Σ , namely the (total) projection of Γ_Y under the map $S \rightarrow \Sigma$. This differs from \hat{Y} in that \tilde{Y}_S contains two copies of each of the curves which lie in A_{in} . Meanwhile, $\mathcal{Y}_{A_{in}}$ is represented by the direct sum \tilde{Y}_{in} of the curves contained in a single copy of A_{in} . Since the non-boundary-parallel summands of Γ_Y are orthogonal to the curves in A_{in} (i.e., their Floer cohomology vanishes), the morphisms from \tilde{Y}_S to \tilde{Y}_{in} which represent the module homomorphisms $\phi_Y^{\text{left/right}}$ must be supported on the summands which are parallel to A_{in} . For each annulus-parallel curve $\sigma \subset A_{in}$ which appears as a summand of \tilde{Y} , the restriction $\Phi_{\Sigma,S}(\sigma)$ is a disjoint union $\sigma_{\text{left}} \cup \sigma_{\text{right}}$ of the two preimages of σ in the ends of S (each of which appears as a summand in Γ_Y , and hence in \tilde{Y}_S). By considering the diagram

$$(5.5) \quad CW_{\Sigma}^*(Y, \sigma) \longrightarrow CW_S^*(\Gamma_Y, \sigma_{\text{left}} \cup \sigma_{\text{right}}) \rightrightarrows CW_{A_{in}}^*(\Phi_{\Sigma,A_{in}}(Y), \sigma),$$

we find that ϕ_Y^{left} maps the summand σ_{left} of \tilde{Y}_S isomorphically to the summand σ of \tilde{Y}_{in} , and vanishes on σ_{right} ; and vice-versa for ϕ_Y^{right} . This implies in turn that the equalizer of the diagram (5.4) is isomorphic to \hat{Y} .

When Γ_Y contains annulus-parallel curves which lie outside of A_{in} , the argument is essentially the same, but now requires a detour via the construction of Example 5.9 to show that the modules \mathcal{Y}_S and \mathcal{Y}_{in} are represented by the total projection \tilde{Y}_S of Γ_Y (containing two copies of each annulus-parallel curve) and by \tilde{Y}_{in} as in the previous case; and once again to analyze the restriction diagrams (5.5) for each annulus-parallel summand σ of \tilde{Y} . \square

Remark 5.12. *Instead of dealing with annulus-parallel curves which lie outside of A_{in} via Example 5.9, one might try to simply enlarge the annulus A_{in} inside Σ in order to ensure that it contains all of the annulus-parallel curves whose Floer cohomology with Y is nonzero. However, this is not always possible, as the area of an immersed cylinder bound by a pair of annulus-parallel curves can be greater than that of the whole surface Σ .*

When Y is an idempotent summand rather than a twisted complex, a similar result holds provided that Y pairs trivially with *two* simple closed curves:

Lemma 5.13. *Let $Y \in \text{Tw}^\pi \mathcal{F}(\Sigma)$. Let $\gamma_1, \gamma_2 \subset \Sigma$ be two disjoint simple closed curves whose homology classes are linearly independent, and denote by A_1, A_2 (disjoint) annular neighborhoods of γ_1, γ_2 , and by S the surface obtained by cutting Σ open along both annuli. If $HF^*(Y, \gamma_1) = HF^*(Y, \gamma_2) = 0$ then $\Phi_{\Sigma, S}(Y)$ and $\Phi_{\Sigma, A_{i, in}}(Y)$ are isomorphic to direct sums of immersed closed curves with local systems in the completions of S and $A_{i, in}$.*

Proof. The argument proceeds as in the proof of Lemma 5.10, considering two annuli parallel to γ_1 and γ_2 . The assumption that γ_1 and γ_2 are homologically independent ensures that cutting Σ open along both annuli yields a connected surface with four punctures, so that the existence of a geometric replacement Γ_Y for $\Phi_{\Sigma, S}(Y)$ is guaranteed by Corollary 4.13. \square

Corollary 5.14. *Let $Y \in \text{Tw}^\pi \mathcal{F}(\Sigma)$. Suppose there are disjoint, homologically independent simple closed curves $\gamma_1, \gamma_2 \subset \Sigma$ with $HF(Y, \gamma_i) = 0$. Then Y is quasi-isomorphic to a direct sum of immersed closed curves with local systems in Σ .*

Proof. This follows from Lemma 5.13 by the same argument as Corollary 5.11. \square

Corollary 5.15. *Let $X \in D^\pi \mathcal{F}(\Sigma)$ be spherical with $ch(X)$ primitive. If there are two disjoint homologically independent simple closed curves γ_j , $j = 1, 2$, with $HF^*(X, \gamma_j) = 0$, then X is quasi-isomorphic to an embedded simple closed curve with a rank one local system.*

Proof. The previous corollary gives a geometric replacement for X ; since a spherical object is indecomposable, the replacement consists of a single immersed curve σ with local system. The assumption that $ch(X) \in H_1(\Sigma; \mathbb{Z})$ is primitive implies that $[\sigma] \in H_1(\Sigma; \mathbb{Z})$ is primitive. Corollary 3.16 then implies that σ is embedded and the local system has rank one. \square

6. GROUP ACTIONS

6.1. Finite group actions on categories. Let G be a finite group. Following [30, 32], we say G acts strictly on a strictly unital A_∞ -category \mathcal{A} if there is an action of G on $\text{Ob } \mathcal{A}$, and corresponding maps between morphism spaces which strictly satisfy the relations of G and for which the A_∞ -operations are strictly equivariant. An action of G on \mathcal{A} induces one on the category \mathcal{A}^{mod} of A_∞ -modules over \mathcal{A} , which is a strictly unital A_∞ -category. Necessarily G will then preserve the strict units.

A *strictly equivariant* object Y of \mathcal{A} is one for which we have closed morphisms

$$\rho_Y^1(g) \in \text{hom}_{\mathcal{A}}^0(g(Y), Y), \quad \rho_Y^1(e_G) = e_Y$$

satisfying

$$\mu_{\mathcal{A}}^2(\rho_Y^1(g), g \cdot \rho_Y^1(h)) = \rho_Y^1(gh)$$

for all $g, h \in G$. Any object $Y \in \mathcal{A}$ which is fixed by G , meaning $g(Y) = Y$ for every $g \in G$, defines a strictly equivariant object for each character $\chi : G \rightarrow \mathbb{K}^*$ via $\rho_Y^1(g) = \chi(g)e_Y$.

In the case of A_∞ -modules, a strictly equivariant structure on $M \in \mathcal{A}^{mod}$ is given by A_∞ -homomorphisms $\rho_M^1(g) \in \text{hom}_{\mathcal{A}^{mod}}^0(g(M), M)$ which generally include higher order terms. The special case where the module homomorphisms $\rho_M^1(g)$ are ordinary linear maps, with all higher order terms identically zero, corresponds to the situation where M is equivariant in the *naive* sense, i.e., there is a linear action of G on the vector spaces underlying M , with respect to which the structure maps of M are equivariant.

A *weakly equivariant* object is one which satisfies the cohomological analogue of strict equivariance, i.e. there are classes $[\rho_Y^1(g)] \in \text{Hom}_{H(\mathcal{A})}^0(g(Y), Y)$ which satisfy

$$[\rho_Y^1(gh)] = [\rho_Y^1(g)] \cdot (g \cdot [\rho_Y^1(h)]), \quad [\rho_Y^1(e_G)] = e_Y.$$

If Y and Y' are both weakly equivariant, then $H^*(\text{hom}_{\mathcal{A}}(Y, Y'))$ becomes a G -representation via $a \mapsto [\rho_{Y'}^1(g)] \cdot (g \cdot a) \cdot [\rho_Y^1(g)]^{-1}$.

Lemma 6.1 (Seidel). *Assume \mathbb{K} has characteristic zero. If Y is a weakly equivariant object of \mathcal{A} or \mathcal{A}^{per} , then the Yoneda module of Y is quasi-isomorphic to a naively equivariant module.*

Proof. This follows from [32, Propositions 14.5 & 14.7] (which in turn are direct analogues of [30, Lemmas 8.2 & 8.3]; in particular they do not rely on the problematic Lemma 13.7 in *op. cit.*). Proposition 14.5 is essentially an obstruction theory calculation (see also [30, Section 8c]) showing that, if $H^r(G; \text{Hom}_{H(\mathcal{A})}^{1-r}(Y, Y)) = 0$ for $r \geq 2$, then a weakly equivariant structure can be upgraded to a *homotopy equivariant* (also known as *coherently equivariant*) structure. In our case, the cohomology $H^r(G; M)$ vanishes for $r > 0$ whenever G is finite and M is a G -module in which $|G|$ is invertible in M , in particular for all modules in characteristic zero. Once this is done, Proposition 14.7 of [32] constructs a naively equivariant A_∞ -module (see also [30, Section 8b]; the rationality requirement which imposes an extra assumption on [30, Lemma 8.2] is not relevant here). \square

6.2. Finite actions and coverings. Let G be a finite abelian group with dual group $G^\vee = \text{Hom}(G, \mathbb{C}^*)$. Given a surface Σ , recall that $H^1(\Sigma; \mathbb{C}^*)$ acts on the Fukaya category $\mathcal{F}(\Sigma)$ by tensoring by flat (Novikov-unitary) line bundles. Thus, a homomorphism from G to $H^1(\Sigma; \mathbb{C}^*)$ determines an action of G on the Fukaya category $\mathcal{F}(\Sigma)$, where each $g \in G$ acts by twisting by a local system ζ^g (with $\zeta^{g_1 g_2} \simeq \zeta^{g_1} \otimes \zeta^{g_2}$). The main example that we consider is the following:

Lemma 6.2. *A non-zero class $a \in H^1(\Sigma; \mathbb{Z})$ defines a \mathbb{Z}/p -action on $\mathcal{F}(\Sigma)$, for any $p \geq 2$.*

Proof. The action on objects is given by tensoring by flat unitary line bundles with holonomy given by the class defined by a in $\text{Hom}(\pi_1(\Sigma); \mathbb{Z}/p)$. Choosing the Floer perturbation data to be independent of local systems in the construction of $\mathcal{F}(\Sigma)$ ensures that this action is strict. (Note that, in general, there is a standard trick to make a finite action strict on a quasi-isomorphic model of $\mathcal{F}(\Sigma)$, by passing to a category whose objects are pairs of an object of $\mathcal{F}(\Sigma)$ and an element of the given finite group, and choosing perturbation data independently for such pairs. See e.g. [35, Appendix A] for a closely related case.) \square

A homomorphism from G to $H^1(\Sigma; \mathbb{C}^*)$ is the same thing as an element of $H^1(\Sigma; G^\vee)$, or a homomorphism $\mu : \pi_1(\Sigma) \rightarrow G^\vee$. Thus, it determines a finite Galois covering $\tilde{\Sigma} \rightarrow \Sigma$ with deck group G^\vee .

Given an immersed curve with a local system $(\xi, \gamma) \in \mathcal{F}(\Sigma)$, the action of $g \in G$ twists ξ by the rank one local system $\zeta_{|\gamma}^g$, whose holonomy is $\langle \mu([\gamma]), g \rangle$. Thus, if ξ has rank one then the G -action preserves the isomorphism class of (ξ, γ) if and only if $[\gamma] \in \text{Ker } \mu$, i.e. if and only if γ lifts to the covering $\tilde{\Sigma}$. The set of G -equivariant structures on (ξ, γ) is then a G^\vee -torsor, and so is the set of lifts of (ξ, γ) to $\tilde{\Sigma}$. To be explicit, fix a base point in Σ , and a trivialization of the G -family of local systems $\{\zeta^g\}$ at the base point. The choice of an arc connecting γ to the base point then determines on one hand, a trivialization of the collection of local systems $\{\zeta^g\}$ over γ , and hence a G -equivariant structure on the object $(\xi, \gamma) \in \mathcal{F}(\Sigma)$ induced by the isomorphisms $\xi \otimes \zeta_{|\gamma}^g \xrightarrow{\sim} \xi$; and on the other hand, a lift $\tilde{\gamma}$ of γ to $\tilde{\Sigma}$. Moving the base arc by a loop along which μ takes the value $\chi \in G^\vee$, we obtain another G -equivariant structure on (ξ, γ) , which we denote by $(\xi, \gamma)^\chi$, where the isomorphism $\xi \otimes \zeta_{|\gamma}^g \xrightarrow{\sim} \xi$ is modified by $\chi(g) \in \mathbb{C}^*$; and another lift $\tilde{\gamma}^\chi$ of γ to $\tilde{\Sigma}$, which differs from $\tilde{\gamma}$ by the deck transformation χ .

Let $\{\gamma_i\}$ be a finite collection of split-generators of $\mathcal{F}(\Sigma)$, whose homotopy classes all lie in $\text{Ker } \mu$. (One way to construct such γ_i is to choose simple closed curves in $\tilde{\Sigma}$ which satisfy Abouzaid's split-generation criterion for $\mathcal{F}(\tilde{\Sigma})$, and project them down to Σ . Compatibility of the open-closed map with the projection $\tilde{\Sigma} \rightarrow \Sigma$ implies that the corresponding immersed curves in Σ split-generate $\mathcal{F}(\Sigma)$.) Fixing base arcs as above, we equip each γ_i with a preferred (strict) G -equivariant structure, and a preferred lift $\tilde{\gamma}_i$ to $\tilde{\Sigma}$.

The chosen G -equivariant structures on γ_i equip the A_∞ -algebra $\mathcal{A} = \bigoplus_{i,j} CF(\gamma_i, \gamma_j)$ with an action of G . Explicitly, $g \in G$ acts on each Floer generator $p \in \gamma_i \cap \gamma_j$ by multiplication by $\chi_p(g)$, where $\chi_p \in G^\vee$ is the image under μ of the loop formed by connecting the intersection point p to the base point along the base arcs for γ_i and γ_j . The G -equivariant part of the Floer complex, $CF^G(\gamma_i, \gamma_j)$, is then generated by those intersections for which $\chi_p = 1$. Those correspond exactly to the intersections between the chosen lifts of γ_i and γ_j to $\tilde{\Sigma}$; therefore $CF^G(\gamma_i, \gamma_j) \simeq CF_{\tilde{\Sigma}}(\tilde{\gamma}_i, \tilde{\gamma}_j)$. Varying the G -equivariant structures by $\chi_i, \chi_j \in G^\vee$, $CF^G(\gamma_i^{\chi_i}, \gamma_j^{\chi_j})$ is generated by those intersections for which $\chi_p = \chi_i^{-1} \chi_j$, and isomorphic to $CF_{\tilde{\Sigma}}(\tilde{\gamma}_i^{\chi_i}, \tilde{\gamma}_j^{\chi_j})$. These isomorphisms are compatible with the A_∞ -operations (if the Floer perturbation data are chosen consistently in Σ and $\tilde{\Sigma}$), and give an isomorphism of A_∞ -algebras

$$(6.1) \quad \tilde{\mathcal{A}} = \bigoplus_{i,j,\chi_i,\chi_j} CF_{\tilde{\Sigma}}(\tilde{\gamma}_i^{\chi_i}, \tilde{\gamma}_j^{\chi_j}) \simeq \bigoplus_{i,j,\chi_i,\chi_j} CF^G(\gamma_i^{\chi_i}, \gamma_j^{\chi_j}).$$

Recall from [33, Ch. 4c] that the semidirect product A_∞ -algebra $\mathcal{A} \rtimes G$ is defined by considering the tensor product $\mathcal{A} \otimes_{\mathbb{K}} \mathbb{K}[G]$ of \mathcal{A} with the group ring of G , with the operations

$$\mu_{\mathcal{A} \rtimes G}^d(a_d \otimes g_d, \dots, a_1 \otimes g_1) = \mu_{\mathcal{A}}^d(a_d, g_d \cdot a_{d-1}, (g_d g_{d-1}) \cdot a_{d-2}, \dots) \otimes (g_d \dots g_1).$$

Recalling that each $\chi \in G^\vee$ determines an idempotent $e_\chi = |G|^{-1} \sum \chi(g) g \in \mathbb{K}[G]$, an easy calculation shows that the linear map $\varphi : \tilde{\mathcal{A}} \rightarrow \mathcal{A} \rtimes G$ which takes $a \in CF^G(\gamma_i^{\chi_i}, \gamma_j^{\chi_j}) \subset \tilde{\mathcal{A}}$

to $\varphi(a) = a \otimes e_{\chi_i^{-1}} \in CF(\gamma_i, \gamma_j) \otimes e_{\chi_i^{-1}} \subset \mathcal{A} \rtimes G$ defines an isomorphism of A_∞ -algebras. Summarizing, we have:

Proposition 6.3 (Seidel [33, Ch. 4 & 8b]). $\mathcal{A} \rtimes G \simeq \tilde{\mathcal{A}}$.

An immersed curve with local system $Y = (\xi, \gamma) \in \mathcal{F}(\Sigma)$, with a G -equivariant structure (i.e., a choice of isomorphisms $\xi \otimes \zeta_{|\gamma}^g \xrightarrow{\sim} \xi$), determines a module $\mathcal{Y}_{(\xi, \gamma)} = \bigoplus_i CF(\gamma_i, Y)$ over the A_∞ -algebra \mathcal{A} , which is G -equivariant in the naive sense (i.e. G acts linearly on $\mathcal{Y}_{(\xi, \gamma)}$ and all the module maps are equivariant). A naive G -equivariant structure on an \mathcal{A} -module M equips M with the structure of a module over $\mathcal{A} \rtimes G$, with structure maps given by

$$\mu^{d+1}(m, a_d \otimes g_d, \dots, a_1 \otimes g_1) = (g_1^{-1} \dots g_d^{-1}) \cdot \mu^{d+1}(m, a_d, g_d \cdot a_{d-1}, (g_d g_{d-1}) \cdot a_{d-2}, \dots).$$

In the case of $\mathcal{Y}_{(\xi, \gamma)}$, this has a geometric interpretation in terms of the lift \tilde{Y} of Y to $\tilde{\Sigma}$:

$$\mathcal{Y}_{(\xi, \gamma)} \simeq \bigoplus_{i, \chi_i} CF^G(\gamma_i^{\chi_i}, Y) \simeq \bigoplus_{i, \chi_i} CF_{\tilde{\Sigma}}(\tilde{\gamma}_i^{\chi_i}, \tilde{Y})$$

as modules over $\mathcal{A} \rtimes G \simeq \tilde{\mathcal{A}}$.

Transcribing these statements in the language of modules over A_∞ -categories, rather than the endomorphism algebra of a given set of generators, a G -equivariant structure (in the naive sense) on a module over $\mathcal{F}(\Sigma)$ determines a lift to a module over $\mathcal{F}(\tilde{\Sigma})$, and for Yoneda modules of equivariant objects of $\mathcal{F}(\Sigma)$ this coincides with the geometric lifting under the covering map $\tilde{\Sigma} \rightarrow \Sigma$.

Example 6.4. Consider the action of $G = \mathbb{Z}/p$ on $\mathcal{F}(\Sigma)$ determined by a class $a \in H^1(\Sigma; \mathbb{Z})$ as in Lemma 6.2, and let $(\xi, \gamma) \in \mathcal{F}(\Sigma)$ be an immersed curve with a rank one local system.

- (1) If $\langle a, [\gamma] \rangle = 0$, then (ξ, γ) is strictly equivariant for G , and its Yoneda module $\mathcal{Y}_{(\xi, \gamma)}$ is equivariant in the naive sense. A choice of equivariant structure corresponds to a choice of lift of γ to $\tilde{\Sigma}$, and the corresponding lift of $\mathcal{Y}_{(\xi, \gamma)}$ to $\mathcal{F}(\tilde{\Sigma})^{mod}$ is then the Yoneda module for the chosen lift of (ξ, γ) .
- (2) If $\langle a, [\gamma] \rangle = 1$ (or is coprime with p), then the the Yoneda module of (ξ, γ) is not equivariant, but its full orbit $\mathcal{Y}_{(\xi, \gamma)} \rtimes G = \bigoplus_{g \in G} \mathcal{Y}_{g(\xi, \gamma)}$ admits a unique equivariant structure. The lift of this equivariant module to $\mathcal{F}(\tilde{\Sigma})^{mod}$ is the Yoneda module for the lift of (ξ, γ) to $\tilde{\Sigma}$ (which consists of a single curve covering γ p -fold).

Generalising Example 6.4, we wish to prove that an object $X \in D^\pi \mathcal{F}(\Sigma)$ with $\langle a, ch(X) \rangle = 0$ admits a G -equivariant structure. However, it is not obvious why the purely homological hypothesis on X should force it to be even weakly equivariant. To prove this we will embed the finite group action into a \mathbb{G}_m -action, where the homological condition will yield infinitesimal equivariance. General machinery due to Seidel then implies weak \mathbb{G}_m -equivariance, and *a posteriori* G -equivariance.

6.3. Categorical \mathbb{G}_m -actions. Let \mathcal{A} be a proper $\mathbb{Z}/2$ -graded A_∞ category with a strict \mathbb{G}_m -action. There are several notions of a \mathbb{G}_m -action on a module $M \in \mathcal{A}^{mod}$. Briefly, following [30] one says:

- (1) \mathbb{G}_m acts *naively* if it acts linearly on the underlying vector spaces of M , with all the structure maps being equivariant.
- (2) \mathbb{G}_m acts *strictly* if there are higher order (in \mathcal{A}) terms to the action, i.e. for each $g \in \mathbb{G}_m$ we have

$$\rho_g^1 = \{\rho_g^{1,d+1} : M \otimes \mathcal{A}^d \rightarrow g^*M\}_{d \geq 0} \in \text{hom}_{\mathcal{A}^{mod}}^0(M, g^*M)$$

satisfying unitality and cocycle conditions, and a rationality condition that the maps $\rho_g^{1,d+1}$ are coherently induced from a single map $\rho^{1,d+1} : M \otimes \mathcal{A}^d \rightarrow \mathbb{K}[\mathbb{G}_m] \otimes M[-d]$.

- (3) \mathbb{G}_m acts *weakly* if there are module homomorphisms ρ_g^1 as above, for which the unitality and cocycle conditions hold at the cohomological level (in $H^0(\mathcal{A}^{mod})$).
- (4) \mathbb{G}_m acts *up to homotopy* if there are higher order (in \mathbb{G}_m) terms to the action, i.e. we now have maps $\rho_{(g_i, \dots, g_1)}^i \in \text{hom}_{\mathcal{A}^{mod}}^{1-r}(M, g_1^*g_2^* \cdots g_i^*M)$, which should again satisfy appropriate unitality, cocycle and rationality conditions.

A naive action yields a strict action with no higher order terms (so $\rho_g^{1,d+1} = 0$ for $d > 0$), whilst a strict action on a module gives a naive action on the quasi-isomorphic module $M \otimes_{\mathcal{A}} \mathcal{A}$, so these are essentially equivalent notions.

Unless \mathcal{A} admits a set of generators which are strictly fixed by \mathbb{G}_m , the notion of rationality is best formulated in the algebro-geometric language of [34, Appendix A]. We will consider situations where \mathbb{G}_m acts freely on the objects of \mathcal{A} and strictly on morphisms. Concretely, in the case of Fukaya categories, this is achieved by setting the objects to be pairs (L, g) where L is a Lagrangian brane and g is an element of the group; the object (L, g) , also denoted $g(L)$, is obtained from L by the action of g and equipped with Floer perturbation data which is pulled back from that for L .

A $\mathbb{K}[\mathbb{G}_m]$ -module is just a quasicoherent sheaf over \mathbb{K}^* , hence has stalks at points $g \in \mathbb{K}^*$. We will identify $\mathbb{K}[\mathbb{G}_m]^{\otimes d}$ with $\mathbb{K}[\mathbb{G}_m^{\times d}]$, and note there is a natural morphism $\mathbb{K}[\mathbb{G}_m] \rightarrow \mathbb{K}[\mathbb{G}_m]^{\otimes d}$ dual to the total multiplication map $\mathbb{G}_m^{\times d} \rightarrow \mathbb{G}_m$. A rational \mathbb{G}_m -action on \mathcal{A} (with \mathbb{G}_m acting freely on objects and strictly on morphisms) is then, by definition, an A_∞ -category $\underline{\mathcal{A}}$, with the same objects as \mathcal{A} , in which

- (1) the morphism groups $\text{hom}_{\underline{\mathcal{A}}}(X_0, X_1)$ are $\mathbb{Z}/2$ -graded projective $\mathbb{K}[\mathbb{G}_m]$ -modules, whose fibre at $g \in \mathbb{G}_m$ is $\text{hom}_{\mathcal{A}}(X_0, X_1)_g := \text{hom}_{\mathcal{A}}(g(X_0), X_1)$; using strictness, $\text{hom}_{\mathcal{A}}(g_0(X_0), g_1(X_1)) = \text{hom}_{\underline{\mathcal{A}}}(X_0, X_1)_{g_1^{-1}g_0}$ for all $X_0, X_1 \in \text{ob } \mathcal{A}$ and $g_0, g_1 \in \mathbb{G}_m$;

- (2) the A_∞ -operations define $\mathbb{K}[\mathbb{G}_m^{\times d}]$ -module maps

$$\mu_{\underline{\mathcal{A}}}^d : \text{hom}_{\underline{\mathcal{A}}}(X_{d-1}, X_d) \otimes_{\mathbb{K}} \cdots \otimes_{\mathbb{K}} \text{hom}_{\underline{\mathcal{A}}}(X_0, X_1) \rightarrow \mathbb{K}[\mathbb{G}_m^{\times d}] \otimes_{\mathbb{K}[\mathbb{G}_m]} \text{hom}_{\underline{\mathcal{A}}}(X_0, X_d)[2-d]$$

which satisfy appropriate associativity equations, and a unitality condition (fiberwise, these are just the usual axioms for \mathcal{A} to be an A_∞ -category);

- (3) for all $X_0, X_1 \in \text{ob } \mathcal{A}$ and $g_0, g_1 \in \mathbb{G}_m$, the $\mathbb{K}[\mathbb{G}_m]$ -module $\text{hom}_{\underline{\mathcal{A}}}(g_0(X_0), g_1(X_1))$ is the pullback of $\text{hom}_{\underline{\mathcal{A}}}(X_0, X_1)$ under multiplication by $g_1^{-1}g_0$, and the A_∞ -operations on these modules strictly coincide under these identifications.

(A small difference between our exposition and [34, Appendix A] is that Seidel takes objects of \mathcal{A} to be pairs consisting of an object of $\underline{\mathcal{A}}$ and an element of \mathbb{G}_m , i.e. the objects of $\underline{\mathcal{A}}$ are \mathbb{G}_m -orbits of objects of \mathcal{A} ; the additional objects in our version of $\underline{\mathcal{A}}$ contain no additional

information, since their morphism spaces and the A_∞ -operations on those are completely determined by condition (3)).

We will consider the case where $\underline{\mathcal{A}}$ is *proper*, i.e. the cohomology groups $H^*(\mathrm{hom}_{\underline{\mathcal{A}}}(X_0, X_1))$ are bounded $\mathbb{K}[\mathbb{G}_m]$ -modules and finitely generated in each degree; this implies properness of \mathcal{A} .

Informally, an object $X \in \mathcal{A}$ is homotopy \mathbb{G}_m -equivariant if it is isomorphic to all of its images $g(X)$ for $g \in \mathbb{G}_m$, in a manner which is coherent up to higher homotopy data. Formally, a (rational) *homotopy equivariant* structure on X is a sequence

$$\rho_X^i \in \mathbb{K}[\mathbb{G}_m^{\times i}] \otimes_{\mathbb{K}[\mathbb{G}_m]} \mathrm{hom}_{\underline{\mathcal{A}}}^{1-i}(X, X), \quad i \geq 1$$

which, stalkwise, give elements

$$\rho_{X, g_i, \dots, g_1}^i \in \mathrm{hom}_{\underline{\mathcal{A}}}^{1-i}(g_i \cdots g_1(X), X)$$

which satisfy the following associativity equations [34, Appendix]:

$$\sum_{\substack{d \geq 1 \\ i_1 + \cdots + i_d = i}} \mu_{\mathcal{A}}^d(\rho_{X, g_i, \dots, g_{i_1 + \cdots + i_{d-1} + 1}}^{i_d}, \dots, \rho_{X, g_{i_1}, \dots, g_1}^{i_1}) + \sum_{1 \leq k < i} (-1)^k \rho_{X, g_i, \dots, g_{k+1} g_k, \dots, g_1}^{i-1} = 0$$

and the unitality condition $[\rho_{X, e}^1] = [e_X] \in H^0(\mathrm{hom}_{\mathcal{A}}(X, X))$.

In this language, a weakly \mathbb{G}_m -equivariant structure on X is the data of ρ_X^1 which satisfies unitality, and for which there exists ρ_X^2 such that the first two associativity equations hold, namely

$$\mu_{\mathcal{A}}^1(\rho_{X, g}^1) = 0 \quad \text{and} \quad \mu_{\mathcal{A}}^2(\rho_{X, g_2}^1, \rho_{X, g_1}^1) - \rho_{X, g_2 g_1}^1 + \mu_{\mathcal{A}}^1(\rho_{X, g_2, g_1}^2) = 0.$$

The machinery in [30, Section 8] upgrades weak equivariant structures to homotopy equivariant structures using an obstruction theory argument, and turns homotopy equivariant objects into naively equivariant modules. For applications, the essential point is therefore to find criteria for the existence of weak actions.

6.4. Infinitesimal equivariance. Denoting by \mathfrak{g} the Lie algebra of \mathbb{G}_m , we have a short exact sequence of $\mathbb{K}[\mathbb{G}_m]$ -modules

$$0 \rightarrow \mathfrak{g}^\vee \rightarrow \mathbb{K}[\mathbb{G}_m]/\mathcal{I}_e^2 \rightarrow \mathbb{K} \rightarrow 0$$

where \mathcal{I}_e is the maximal ideal at the point $e \in \mathbb{G}_m$, and $\mathbb{K} = \mathbb{K}[\mathbb{G}_m]/\mathcal{I}_e$. Tensoring this with the (projective, hence flat) module $\mathrm{hom}_{\underline{\mathcal{A}}}(X, X)$ and considering the resulting long exact sequence in cohomology, the image of $[e_X] \in H^0(\mathrm{hom}_{\mathcal{A}}(X, X))$ under the connecting morphism yields an element

$$(6.2) \quad \mathrm{Def}_X^0 \in \mathfrak{g}^\vee \otimes H^1(\mathrm{hom}_{\mathcal{A}}(X, X)).$$

If this vanishes then we say X is “infinitesimally equivariant”.

Proposition 6.5. *Suppose \mathcal{A} is proper, has a rational \mathbb{G}_m -action (in the sense of the preceding section), and \mathbb{K} has characteristic zero. If X is infinitesimally equivariant and simple, i.e. $H^0(\mathrm{hom}_{\mathcal{A}}(X, X)) = \mathbb{K}$, then X admits a weakly equivariant structure.*

Proof. The proof follows along the same lines as Seidel's argument for the case $\mathbb{K} = \mathbb{C}$, which is stated as [34, Theorem 2.7 and Corollary A.12] for A_∞ -categories with naive and rational \mathbb{C}^* -actions, respectively. Rationality of the action and properness imply that the cohomology $H^0(\mathrm{hom}_{\underline{\mathcal{A}}}(X, X))$ defines a coherent sheaf over \mathbb{G}_m , and infinitesimal equivariance equips that sheaf with an algebraic connection. More precisely, a choice of primitive

$$\mathrm{def}_X^0 \in \mathfrak{g}^\vee \otimes \mathrm{hom}_{\underline{\mathcal{A}}}^0(X, X) \quad \text{for} \quad \mathrm{Def}_X^0 = 0 \in \mathfrak{g}^\vee \otimes H^1(\mathrm{hom}_{\underline{\mathcal{A}}}(X, X))$$

equips the sheaf with a distinguished connection (cf. [34, Lemma A.3] and [30, Section 7a]). In characteristic zero, any coherent sheaf admitting a connection is locally free [9, Corollary 2.5.2.2], so the sheaf $\mathcal{F} = H^0(\mathrm{hom}_{\underline{\mathcal{A}}}(X, X))$ is locally free of rank one, i.e. a line bundle.

Over \mathbb{C} , the construction in [30, Lemma 7.12] uses surjectivity of the exponential map to modify the algebraic connection and trivialise the monodromy. A covariant constant section taking the value $[e_X]$ at $e \in \mathbb{G}_m$ then provides a \mathbb{G}_m -family of cohomological isomorphisms which obey the group law, and hence defines a weak \mathbb{G}_m -action (which can be lifted to a homotopy \mathbb{G}_m -action using a general obstruction theory argument).

Following [32, Remark 14.22], we instead argue as follows. The total space of the line bundle \mathcal{F} carries an action of an algebraic group \mathcal{G} which fits into an extension

$$1 \rightarrow \mathrm{Aut}(X) \rightarrow \mathcal{G} \rightarrow \mathbb{G}_m \rightarrow 1$$

where the subgroup $\mathrm{Aut}(X) = H^0 \mathrm{hom}_{\underline{\mathcal{A}}}(X, X)^*$ acts by multiplication on the fibres $\mathcal{F}_g = H^0 \mathrm{hom}_{\underline{\mathcal{A}}}(gX, X)$, and the quotient \mathbb{G}_m acts on the base by multiplication.

Since X is simple, $\mathrm{Aut}(X) = \mathbb{G}_m$. Over any perfect (e.g. characteristic zero) field, the group of self-extensions $\mathrm{Ext}^1(\mathbb{G}_m, \mathbb{G}_m) = 0$, cf. for instance [12, Chapter 5]. Therefore $\mathcal{G} \cong \mathbb{G}_m \times \mathbb{G}_m$, the extension splits, and the action of \mathbb{G}_m on itself by multiplication admits a lift to an action of \mathbb{G}_m on the total space of \mathcal{F} . The orbit of $[e_X] \in \mathcal{F}_e$ then provides the desired \mathbb{G}_m -family of cohomological isomorphisms. \square

6.5. \mathbb{G}_m -actions from the Picard group. We now apply the machinery of rational \mathbb{G}_m -actions to the setting of Fukaya categories of surfaces. We first state the result for a closed surface Σ of genus $g \geq 2$, which is our main focus.

Proposition 6.6. *A choice of class $a \in H^1(\Sigma; \mathbb{Z})$ defines a rational \mathbb{G}_m -action on $\mathcal{F}(\Sigma)$, for which $\underline{\mathcal{F}}(\Sigma)$ is proper in the sense of Section 6.3.*

Proof. Choose a closed differential form $\alpha \in \Omega^1(\Sigma; \mathbb{R})$ representing a , and let $V \in \Gamma(T\Sigma)$ be the symplectic vector field ω -dual to α . Consider the actions of $s \in \mathbb{R}$ and of $\lambda \in U_\Lambda$ on Lagrangian branes in Σ given by

- the symplectomorphisms ϕ^s associated to the time s flow of V , with flux $s \cdot a$, and
- twisting by the line bundle $\xi_\lambda \rightarrow \Sigma$ which is topologically trivial and has connection $d + \lambda \cdot \alpha$ and hence holonomy $\exp(\lambda \cdot \alpha)$.

Using the fact that $\mathcal{L}_V(\alpha) = 0$, one can check that the actions of \mathbb{R} and U_Λ *strictly* commute, and define an action of \mathbb{G}_m on the set of Lagrangian submanifolds with local systems, where $q^s \lambda \in \mathbb{G}_m$ acts by $\phi^s \circ (\otimes \xi_\lambda)$.

We now enlarge the Fukaya category, following [28, Section 10b], to allow pairs (L, g) with L a brane and $g \in \mathbb{G}_m$, where the perturbation data for (L, g) is the g -pullback of that for L . Then \mathbb{G}_m acts strictly on Lagrangian branes, via A_∞ -functors with no higher order terms. The ideas of Section 3.3, cf. [31] and [23], and in particular Theorem 3.12 imply that the morphism groups $CF^*(g(L), L')$ form analytic sheaves of complexes over \mathbb{G}_m . We will explain below, using ideas from Section 5, that these morphism groups can in fact be arranged to be the stalks of (infinite rank but cohomologically proper) projective $\mathbb{K}[\mathbb{G}_m]$ -modules. The fact that the A_∞ -operations are compatible with the group operation follows from strictness of the action.

The reason why the projective $\mathbb{K}[\mathbb{G}_m]$ -modules $\text{hom}_{\mathcal{F}(\Sigma)}(L, L')$ typically have infinite rank, is that the Floer complexes $CF^*(g(L), L')$ have unbounded rank, since deforming L by an increasingly large amount of flux may create unboundedly many new intersections with L' . These new intersections come in cancelling pairs, and because Σ has genus at least two, only finitely many such pairs of intersections ever contribute to the Floer cohomology $HF^*(g(L), L')$, with the rest belonging to summands which remain acyclic for all g . Still, in order to define the chain-level A_∞ -structure it is simpler to work with the whole Floer complex. In our setting, there is a direct geometric approach to turning the Floer complexes $CF^*(g(L), L')$ into projective (in fact, free) $\mathbb{K}[\mathbb{G}_m]$ -modules.

Namely, represent the Poincaré dual to a (or some multiple of it, if a is not a primitive element of $H^1(\Sigma; \mathbb{Z})$) by a simple closed curve σ , and pick the 1-form α and vector field V to be supported in an annular neighbourhood $\sigma \subset A \subset \Sigma$, which we identify with $(-2, 2) \times S^1$ as in Section 5.3, with α pulled back from $(-2, 2)$ so that V is everywhere parallel to the S^1 -factor. As in Section 5.3, only consider objects of $\mathcal{F}(\Sigma)$ which are not parallel to the annulus A and intersect its boundary minimally (annulus-parallel curves get replaced by twisted complexes as in Example 5.9). Then the Lee perturbations commute with the \mathbb{G}_m -action, and the various properties of the Floer complexes $CF^*(L, L'; wH)$ for large w carry over to $CF^*(g(L), L'; wH)$ as long as w is sufficiently large compared to the valuation of $g \in \mathbb{G}_m$. Because the continuation maps $CF^*(g(L), L'; wH) \rightarrow CF^*(g(L), L'; (w+1)H)$ for sufficiently large w are chain-level inclusions which define strict A_∞ -homomorphisms with no higher order terms, the telescope models for $CW^*(g(L), L') = \varinjlim CF^*(g(L), L'; wH)$ can be replaced with naive limits: define $\text{hom}(g(L), L')$ to be the union of the increasing sequence of complexes $CF^*(g(L), L'; wH)$ for $w = N, N+1, \dots$ starting from sufficiently large $N \gg \text{val}(g)$, identifying each with a subcomplex of the next by inclusion. Equivalently, declare the generators of $\text{hom}(g(L), L')$ to be the arcs of intersections given by Lemma 5.1, with A_∞ -operations given by counts of perturbed holomorphic discs for any value of w which is sufficiently large relative to the given inputs and the valuation of g , corrected by action rescalings as in Lemma 5.4. There is an obvious identification between generators for varying values of g , and the fact that the A_∞ -operations vary algebraically with $g \in \mathbb{G}_m$ is a direct consequence of Lemma 3.10; thus the Floer complexes constructed in this manner naturally define projective $\mathbb{K}[\mathbb{G}_m]$ -modules underlying a rational \mathbb{G}_m -action.

Even though the $\mathbb{K}[\mathbb{G}_m]$ -modules $\text{hom}_{\mathcal{F}(\Sigma)}(L, L')$ constructed in this manner are not finitely generated, the cohomology $H^* \text{hom}_{\mathcal{F}(\Sigma)}(L, L')$ is a finitely generated $\mathbb{K}[\mathbb{G}_m]$ -module, because all but finitely many of the generators created upon wrapping come in pairs which are connected by a single bigon contributing to the Floer differential, hence lie in acyclic

summands of the chain-level $\mathbb{K}[\mathbb{G}_m]$ -module. Here we use crucially the fact that Σ is not a (closed) torus. We conclude that $\underline{\mathcal{F}}(\Sigma)$ is proper. \square

Remark 6.7. *Proposition 6.6 also holds for a surface S of genus ≥ 1 with non-empty boundary, under the additional assumption that the restriction of $a \in H^1(S; \mathbb{Z})$ to ∂S vanishes. This ensures that the class a can be represented by a one-form supported in an annular neighborhood of a simple closed curve, and the argument then proceeds exactly as in the case of closed surfaces. Moreover, the construction applies equally well to the wrapped Fukaya category, and gives a rational \mathbb{G}_m -action on $\mathcal{W}(S)$ (however $\underline{\mathcal{W}}(S)$ is not proper).*

Example 6.8. *Suppose $\gamma \subset \Sigma$ is a simple closed curve, and $a \in H^1(\Sigma; \mathbb{Z})$. Then it is clear from the definition of the \mathbb{G}_m -action that the class $\text{Def}_\gamma^0 \in \mathfrak{g}^\vee \otimes HF^1(\gamma, \gamma) \simeq H^1(\gamma; \mathbb{K})$ is $a|_\gamma$, or more precisely, the homomorphism from $\mathfrak{g} \simeq \mathbb{K}$ to $HF^1(\gamma, \gamma)$ given by multiplication by $a|_\gamma$. Thus, γ is infinitesimally equivariant if and only if $\langle a, \gamma \rangle = 0$, as expected.*

The rational \mathbb{G}_m -action on $\mathcal{F}(\Sigma)$ immediately extends to twisted complexes and to the idempotent completion $\mathcal{F}(\Sigma)^{\text{per}}$; the properness of $\underline{\mathcal{F}}(\Sigma)$ implies that of $\underline{\mathcal{F}}(\Sigma)^{\text{per}}$.

Lemma 6.9. *Let $X \in \mathcal{F}(\Sigma)^{\text{per}}$ be a spherical object with $ch(X) \neq 0$, and $a \in H^1(\Sigma; \mathbb{Z})$. Then X is infinitesimally equivariant with respect to the action of Proposition 6.6 if and only if $\langle a, ch(X) \rangle = 0$.*

Proof. The deformation classes Def_X^0 , $X \in \text{ob } \mathcal{F}(\Sigma)$ are induced by a Hochschild cohomology class $\mathbf{Def} \in \mathfrak{g}^\vee \otimes HH^1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$, which measures the infinitesimal action of \mathbb{G}_m on morphisms [34]. (Since $\mathfrak{g}^\vee \simeq \mathbb{K}$, we can think of this as a single Hochschild cohomology class, and henceforth we drop \mathfrak{g}^\vee from the notation). Recall that the closed-open map $CO : H^1(\Sigma; \Lambda) \rightarrow HH^1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ is an isomorphism. Thus, the fact that $\text{Def}_\gamma^0 = a|_\gamma$ for all simple closed curves implies that $\mathbf{Def} = CO(a)$.

The Yoneda product makes Hochschild cohomology a unital algebra, over which the *cap product* endows the Hochschild homology $HH_*(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ with a module structure. It is well-known to experts that the closed-open map is a ring homomorphism, and that the open-closed map $OC : HH_*(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma)) \rightarrow H^{*+1}(\Sigma; \Lambda) \simeq H_{1-*}(\Sigma; \Lambda)$ is a homomorphism of $H^*(\Sigma; \Lambda)$ -modules (see e.g. [16], [26]). With this understood, we can view $\text{Def}_X^0 \in H^1 \text{hom}(X, X)$ as a Hochschild cycle, and the corresponding class $[\text{Def}_X^0] \in HH_1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ can be expressed as $\mathbf{Def} \cap [e_X] = CO(a) \cap [e_X]$, which implies that

$$OC([\text{Def}_X^0]) = OC(CO(a) \cap [e_X]) = a \cap OC([e_X]) = a \cap ch(X) \in H_0(\Sigma; \Lambda).$$

Since the open-closed map is an isomorphism, we conclude that $[\text{Def}_X^0] \in HH_1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ vanishes if and only if $\langle a, ch(X) \rangle = 0$. In order to reach the same conclusion for $\text{Def}_X^0 \in H^1 \text{hom}(X, X)$, it remains to verify that the map $H^1 \text{hom}(X, X) \rightarrow HH_1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$ is an isomorphism. Since X is spherical, both sides have rank one, and it is enough to show that the map does not vanish identically; since $ch(X) \neq 0$ this follows from the existence of classes a for which $\langle a, ch(X) \rangle \neq 0$. \square

We note that the assumption that $ch(X) \neq 0$ can be removed by using the Calabi-Yau structure on $\mathcal{F}(\Sigma)$; more generally the statement is expected to hold for all indecomposable objects of $\mathcal{F}(\Sigma)^{\text{per}}$.

We also remark that a similar criterion for infinitesimal equivariance can be formulated for spherical objects on a surface with boundary, using inclusion into a larger closed surface to reduce to Lemma 6.9.

7. CONCLUSIONS

7.1. Spherical objects revisited. We can now conclude the proof of Theorem 1.1.

Corollary 7.1. *Let $X \in \mathcal{F}(\Sigma)^{per}$ be spherical, with $ch(X) \neq 0$. Suppose there is a class $a \in H^1(\Sigma; \mathbb{Z})$ with the property that $\langle a, ch(X) \rangle = 0$. Then for each p , X is quasi-isomorphic to an object of $\mathcal{F}(\Sigma)^{per}$ which admits a naive equivariant structure for the \mathbb{Z}/p -action of Lemma 6.2 and lifts to the degree p covering $\tilde{\Sigma} \rightarrow \Sigma$ as an object $\hat{X} \in \mathcal{F}(\tilde{\Sigma})^{per}$.*

Proof. By Lemma 6.9, the hypothesis $\langle a, ch(X) \rangle = 0$ implies that X is infinitesimally equivariant for the \mathbb{G}_m -action associated to a by Proposition 6.6. Proposition 6.5 then shows that X is weakly \mathbb{G}_m -equivariant in the sense of Section 6.3. Restricting to the finite subgroup of p -th roots of unity, we find that X is weakly \mathbb{Z}/p -equivariant with respect to the action of Lemma 6.2, and so by Lemma 6.1 its Yoneda module is quasi-isomorphic to a naively equivariant module. One can then appeal to the results of Section 6.2 and reinterpret this naively \mathbb{Z}/p -equivariant $\mathcal{F}(\Sigma)$ -module as a module over $\mathcal{F}(\tilde{\Sigma})$.

Since $\mathcal{F}(\Sigma)$ is proper, the Yoneda module of X is proper, and so are the equivariant module and its lift; the homological smoothness of $\mathcal{F}(\tilde{\Sigma})$ then implies that the lifted module is perfect, hence can be represented by an object $\hat{X} \in \mathcal{F}(\tilde{\Sigma})^{per}$. Finally, projecting \hat{X} back to Σ yields a naively equivariant object of $\mathcal{F}(\Sigma)^{per}$ which is quasi-isomorphic to X . \square

We now show that certain objects are supported on subsurfaces of $\tilde{\Sigma}$, in the following sense.

Corollary 7.2. *If $X \in \mathcal{F}(\Sigma)^{per}$ is spherical, with $ch(X) \neq 0$, we can find a cover $\tilde{\Sigma} \rightarrow \Sigma$ and a lift \hat{X} of X to $\tilde{\Sigma}$ and a simple closed curve $\gamma \subset \tilde{\Sigma}$ for which $HF^*(\hat{X}, \gamma) = 0$.*

Proof. Pick a class $a \in H^1(\Sigma; \mathbb{Z})$ with $\langle a, ch(X) \rangle = 0$. By the previous Corollary, X is $G = \mathbb{Z}/p$ -equivariant for the action of Lemma 6.2, and can be lifted to the degree p covering $\tilde{\Sigma} \rightarrow \Sigma$, giving an object $\hat{X} \in \mathcal{F}(\tilde{\Sigma})^{per}$.

Let $\delta \subset \Sigma$ be any simple closed curve such that $\langle a, [\delta] \rangle = 0$: then δ lifts to a simple closed curve $\hat{\delta}$ in $\tilde{\Sigma}$, and by the results of Section 6.2 we have

$$(7.1) \quad HF^*(X, \delta) = \bigoplus_{\chi \in G^\vee} HF^*(\hat{X}, \hat{\delta}^\chi),$$

where we recall that the objects $\hat{\delta}^\chi$, $\chi \in G^\vee$ are the p lifts of δ to $\tilde{\Sigma}$. Then for sufficiently large p ($p > \text{rank } HF^*(X, \delta)$) some summand on the RHS vanishes, giving a simple closed curve $\gamma = \hat{\delta}^\chi \subset \tilde{\Sigma}$ with $HF^*(\hat{X}, \gamma) = 0$. \square

Lemma 7.3. *In the situation of Corollary 7.2, one can choose $\tilde{\Sigma}$ so that there are two disjoint homologically independent simple closed curves $\gamma_1, \gamma_2 \subset \tilde{\Sigma}$ with $HF^*(\hat{X}, \gamma_i) = 0$.*

Proof. We proceed as in the proof of Corollary 7.2. Since the genus of Σ is at least two, we can find two simple closed curves δ, δ' in Σ which have geometric intersection number one and both pair trivially with a . The p lifts of δ to $\tilde{\Sigma}$ are disjoint simple closed curves, and each of them intersects precisely one of the p lifts of δ' , so their homology classes are linearly independent in $H_1(\tilde{\Sigma}; \mathbb{Z})$. The result now follows by arguing as in the previous Corollary and taking p sufficiently large to force the vanishing of at least two of the Floer cohomology groups $HF^*(\hat{X}, \hat{\delta}^\chi)$. \square

Corollary 7.4 (=Theorem 1.1). *A spherical object $X \in D^x\mathcal{F}(\Sigma)$ with primitive Chern character is quasi-isomorphic to a simple closed curve equipped with a rank one local system.*

Proof. By Lemma 7.3, we can find an action of $G = \mathbb{Z}/p$ and a finite covering $\tilde{\Sigma} \rightarrow \Sigma$ for which X is G -equivariant and lifts to an object $\hat{X} \in \mathcal{F}(\tilde{\Sigma})^{per}$, and two disjoint, homologically independent simple closed curves $\gamma_1, \gamma_2 \subset \tilde{\Sigma}$ with $HF^*(\hat{X}, \gamma_i) = 0$. We now check that \hat{X} is spherical and $ch(\hat{X})$ is primitive.

Recalling that the Fukaya category of $\tilde{\Sigma}$ gives a model for the G -equivariant Fukaya category of Σ (cf. Section 6.2), $HF^*(\hat{X}, \hat{X})$ is isomorphic to the G -invariant summand of $HF^*(X, X) \simeq H^*(S^1; \mathbb{K})$. General properties of equivariant objects imply that the unit e_X is G -invariant, so $HF^0(\hat{X}, \hat{X})$ has rank one, and Poincaré duality (or vanishing of the skew-symmetric pairing $\chi\text{Hom}(\hat{X}, \hat{X})$) implies that $HF^1(\hat{X}, \hat{X})$ also has rank one. Hence \hat{X} is spherical. Meanwhile, the assumption that $ch(X) \in H_1(\Sigma; \mathbb{Z})$ is primitive means that there exists a simple closed curve η such that $\chi\text{Hom}(X, \eta) = 1$. Denoting by $\hat{\eta}$ the total preimage of η in $\tilde{\Sigma}$ (which may consist of one or more simple closed curves), we find that $\chi\text{Hom}(\hat{X}, \hat{\eta}) = \chi\text{Hom}(X, \eta) = 1$, so $ch(\hat{X})$ is a primitive element of $H_1(\tilde{\Sigma}; \mathbb{Z})$.

Corollary 5.15 now implies that \hat{X} is quasi-isomorphic to a simple closed curve $(\hat{\xi}, \hat{\sigma})$ with rank one local system in $\mathcal{F}(\tilde{\Sigma})$. The last step is to descend back from $\tilde{\Sigma}$ to Σ . The projection of $\hat{\sigma}$ to Σ is a closed (*a priori* immersed) curve σ in Σ , whose homology class satisfies $\langle a, [\sigma] \rangle = 0$, and the local system $\hat{\xi}$ descends to a rank one local system ξ on σ . As noted in Section 6.2, the object $(\xi, \sigma) \in \mathcal{F}(\Sigma)$ admits p distinct G -equivariant structures, and one of these corresponds to the lift $(\hat{\xi}, \hat{\sigma}) \in \mathcal{F}(\tilde{\Sigma})$. A quasi-isomorphism between \hat{X} and $(\hat{\xi}, \hat{\sigma})$ then descends to a G -equivariant quasi-isomorphism between X and (ξ, σ) (i.e., a quasi-isomorphism which lies in the G -invariant part of the Floer complex); from which it follows that X and (ξ, σ) are also quasi-isomorphic (non-equivariantly) in $\mathcal{F}(\Sigma)^{per}$. Corollary 3.16 then implies that σ is quasi-isomorphic to an embedded simple closed curve. \square

Remark 7.5. *The exotic spherical object of Figure 3 lifts to an embedded simple closed curve on a double cover of the surface, so the ability to descend from $\tilde{\Sigma}$ back to Σ above is again making essential use of the homological hypothesis on X .*

7.2. A Floer-theoretic Schmutz graph. The ‘‘Schmutz graph’’, introduced in [27], has vertices non-separating simple closed curves up to isotopy, and two such are joined by an edge exactly when they have geometric intersection number one. The main theorem of *op. cit.* asserts that the group of simplicial automorphisms of the Schmutz graph is the extended mapping class group (of both orientation-preserving and reversing automorphisms) for a

surface of genus $g \geq 3$, and the quotient of the mapping class group by the hyperelliptic involution ι when $g = 2$.

We now introduce a Floer-theoretic analogue. By Theorem 1.1, spherical objects of $D^\pi \mathcal{F}(\Sigma)$ with primitive Chern character correspond to homologically essential simple closed curves with rank one local systems. Moreover, the simple closed curves underlying two spherical objects X_1, X_2 have geometric intersection number one if and only if the rank of $HF^*(X_1, X_2)$ is equal to 1, by Corollary 2.11. We introduce an equivalence relation on the set of spherical objects by declaring that $X \sim X'$ if and only if they have Floer cohomology rank one with the same set of spherical objects:

$$(7.2) \quad X \sim X' \iff (\text{for all spherical } Y, \text{rk } HF^*(X, Y) = 1 \iff \text{rk } HF^*(X', Y) = 1).$$

Lemma 7.6. *Two spherical objects $X, X' \in D^\pi \mathcal{F}(\Sigma)$ with primitive Chern characters are equivalent if and only if the simple closed curves γ, γ' underlying X, X' are isotopic.*

Proof. Using Corollary 2.11, $X \simeq X'$ if and only if γ, γ' have geometric intersection number one with the exact same set of simple closed curves. It is a general fact of surface topology that this can only happen when γ, γ' are isotopic. For example, if γ and γ' were not isotopic, then the Schmutz graph would admit an automorphism which exchanges the vertices corresponding to γ and γ' while preserving every other vertex; this cannot be induced by an element of the extended mapping class group. \square

Since our equivalence relation on spherical objects is clearly preserved by auto-equivalences of the Fukaya category, Lemma 7.6 implies that auto-equivalences act on the set of isotopy classes of non-separating simple closed curves. This remains true for surfaces with boundary, by the same argument; cf. Remark 7.8. (However, spherical objects supported on boundary-parallel curves can be permuted arbitrarily by auto-equivalences of the compact Fukaya category; auto-equivalences of the wrapped Fukaya category are better-behaved.)

On closed surfaces, the machinery of [31] actually allows us to obtain a stronger result: recalling that every homologically essential spherical object of $\mathcal{F}(\Sigma)$ comes in a \mathbb{G}_m -family, any auto-equivalence must act on such objects in a way that maps \mathbb{G}_m -families to \mathbb{G}_m -families. Since we will not need this result, we only sketch the argument.

Sketch. Given a spherical object $X \in \mathcal{F}(\Sigma)$ with $ch(X)$ primitive, and a class $a \in H^1(\Sigma; \mathbb{Z})$ with $\langle a, ch(X) \rangle = 1$, the rational \mathbb{G}_m -action of Proposition 6.6 yields an object of $\underline{\mathcal{F}}(\Sigma)$ representing the \mathbb{G}_m -orbit of X , which in turn determines an object \mathcal{X} of the category $\mathcal{F}_{\mathbb{K}^*}^{per}$ of perfect families of $\mathcal{F}(\Sigma)$ -modules over the base \mathbb{K}^* [31, Section 1f]. Example 6.8 and the proof of Lemma 6.9 show that this family of modules follows a deformation field $[\gamma] \in HH^1(\mathcal{F}_{\mathbb{K}^*}^{per}, \Omega_{\mathbb{K}^*}^1 \otimes \mathcal{F}_{\mathbb{K}^*}^{per})$ which is constant over the base \mathbb{K}^* and given fibrewise by $CO(a) \in HH^1(\mathcal{F}(\Sigma), \mathcal{F}(\Sigma))$. Now, an autoequivalence G of $\mathcal{F}(\Sigma)$ induces a functor \mathcal{G}^{per} on $\mathcal{F}_{\mathbb{K}^*}^{per}$, and $\mathcal{G}^{per}(\mathcal{X})$ is a perfect family of modules which follows the deformation field $[\gamma'] = G_*([\gamma])$, which is constant over the base and given fibrewise by $G_*(CO(a))$ [31, Section 1i]. Since the closed-open map is an isomorphism, there exists $a' \in H^1(\Sigma; \Lambda)$ such that $G_*(CO(a)) = CO(a')$; in fact $a' \in H^1(\Sigma; \mathbb{Z})$, because $\langle a', ch(G(Y)) \rangle = \langle a, ch(Y) \rangle \in \mathbb{Z}$ for all spherical objects Y . Repeat the above construction for the spherical object $X' = G(X)$ and the \mathbb{G}_m -action determined by the cohomology class a' , to build a perfect family of

modules $\mathcal{X}' \in \mathcal{F}_{\mathbb{K}^*}^{per}$, which follows the deformation field $[\gamma']$. The two families of modules $\mathcal{G}^{per}(\mathcal{X})$ and \mathcal{X}' agree over the origin and both follow the deformation field $[\gamma']$, so by [31, Proposition 1.21] their fibres are quasi-isomorphic at every point of \mathbb{K}^* . \square

The group of auto-equivalences $\text{Auteq}(\mathcal{F}(\Sigma))$ contains a subgroup isomorphic to $H^1(\Sigma, \Lambda^*)$, given by symplectic isotopies of arbitrary flux and tensoring by flat line bundles. We expect that a further elaboration on the above argument implies that this is a normal subgroup.

We now return to the proof of Corollary 1.2:

Proposition 7.7. *There is a natural homomorphism $\text{Auteq}(D^\pi \mathcal{F}(\Sigma)) \rightarrow \Gamma_g$, which is split by the (non-canonical) homomorphism $\Gamma_g \rightarrow \text{Auteq}(D^\pi \mathcal{F}(\Sigma))$ constructed in Section 2.6.*

Proof. Define a graph $\Upsilon(\Sigma)$ as follows:

- vertices are spherical objects $X \in \mathcal{F}(\Sigma)^{per}$ with primitive Chern character, modulo the equivalence relation (7.2);
- two distinct vertices X_1 and X_2 are joined by an edge whenever $HF(X_1, X_2)$ has rank 1 (this is clearly invariant under the equivalence relation).

Theorem 1.1 and Lemma 7.6 imply that the vertices of the graph are in bijection with isotopy classes of homologically essential simple closed curves on Σ , while Corollary 2.11 shows that the edges correspond to pairs of curves with geometric intersection number one. Thus, $\Upsilon(\Sigma)$ agrees with the Schmutz graph from [27]. On the other hand, it is manifest that $\text{Auteq}(\mathcal{F}(\Sigma)^{per})$ acts on $\Upsilon(\Sigma)$ by simplicial automorphisms. Thus, we obtain a homomorphism from $\text{Auteq}(\mathcal{F}(\Sigma)^{per})$ to the extended mapping class group Γ_g^\pm for $g \geq 3$, or to the quotient $\Gamma_2^\pm / \langle \iota \rangle$ when $g = 2$.

Since any autoequivalence preserves the pairing $\chi HF(\cdot, \cdot)$, i.e. algebraic intersection numbers of simple closed curves, its action on the set of isotopy classes cannot be that of an orientation-reversing diffeomorphism. Therefore, the homomorphism actually takes values in the ordinary (oriented) mapping class group Γ_g for $g \geq 3$, or in $\Gamma_2 / \langle \iota \rangle$ when $g = 2$. When $g = 2$, one can consider the action of $\text{Auteq}(\mathcal{F}(\Sigma_2)^{per})$ on the graph $\Upsilon^+(\Sigma)$ whose vertices are spherical objects modulo a graded version of (7.2) which requires $HF^*(X, Y)$ and $HF^*(X', Y)$ to be in the same degree when they both have rank 1; i.e., we now consider isotopy classes of *oriented* simple closed curves. Since the hyperelliptic involution reverses the orientation on all simple closed curves, this allows one to lift the homomorphism from $\Gamma_2 / \langle \iota \rangle$ to Γ_2 . The fact that the action on the Floer-theoretic Schmutz graph is split by the construction of Section 2.6 is straightforward. \square

Remark 7.8. *The Schmutz graph can be defined analogously for isotopy classes of non-separating simple closed curves on a punctured surface of genus $g \geq 1$ with n punctures. It is known [27] that its simplicial isometry group reproduces the extended mapping class group, modulo the hyperelliptic involution for $(g, n) \in \{(1, 1), (1, 2)\}$. Starting from here, and using the geometricity of spherical objects obtained in Corollary 4.14, one sees that the analogue of Proposition 7.7 holds for punctured surfaces, i.e. the autoequivalence group of the compact Fukaya category determines (and surjects onto) the mapping class group.*

8. AN APPLICATION TO SYMPLECTIC MAPPING CLASS GROUPS

In this final section, we prove Theorem 1.3 as an application of Corollary 1.2. The argument is a fairly straightforward adaption of ideas from [31, 37], but involves somewhat different technology from that in the main body of the paper, so we will be relatively brief. We also leave the realm of strictly unobstructed Lagrangians; our main examples satisfy a weak monotonicity property, but we will not labour foundational aspects of the Fukaya category.

8.1. Fukaya category summands. Let (M, ω) be a closed symplectic manifold. Define the preliminary category $\mathcal{F}_{pr}(M)$ to be a curved \mathbb{Z}_2 -graded linear A_∞ -category over $\Lambda_{\geq 0}$ which has objects oriented spin Lagrangian submanifolds $L \subset M$ equipped with finite-dimensional $\Lambda_{\geq 0}$ -local systems $\xi \rightarrow L$. The morphism groups, in the two most important cases, are given by

$$(8.1) \quad \text{hom}_{\mathcal{F}_{pr}(M)}((L, \xi), (L', \xi')) = \begin{cases} C^*(L; \text{Hom}(\xi, \xi')) & L = L' \\ \bigoplus_{x \in L \cap L'} \xi_x^\vee \otimes \xi'_x & L \pitchfork L' \end{cases}$$

where in the first case we take any fixed finite-dimensional chain-level model for the classical cohomology of L with coefficients in the bundle $\text{Hom}(\xi, \xi')$. When $\xi = \xi'$ we take this chain-level model to be strictly unital, and denote the unit by $1_{(L, \xi)}$ (or 1_L if we suppress local systems from the notation). Floer theory defines a curved A_∞ -structure $\{\mu^d\}_{d \geq 0}$ on $\mathcal{F}_{pr}(M)$.

Let $\lambda \in \Lambda_{> 0}$. The Fukaya category $\mathcal{F}(M; \lambda)$ has objects pairs (L, α) where $L \in \text{Ob}(\mathcal{F}_{pr}(M))$, where $\alpha \in \text{hom}_{\mathcal{F}_{pr}(M)}^1(L, L)$ vanishes in $\text{hom}_{\mathcal{F}_{pr}(M)}^1(L, L) \otimes_{\Lambda_0} \mathbb{C}$, and where α is a solution of the *weakly unobstructed* Maurer-Cartan equation

$$\mu^0 + \mu^1(\alpha) + \mu^2(\alpha, \alpha) + \cdots = \lambda \cdot 1_L \in \text{hom}_{\mathcal{F}_{pr}(M)}^0(L, L).$$

The morphism spaces in $\mathcal{F}(M; \lambda)$ are given by Floer cochains

$$\text{hom}_{\mathcal{F}(M; \lambda)}(L, L') = \text{hom}_{\mathcal{F}_{pr}(M)}(L, L') \otimes_{\Lambda_0} \Lambda$$

and inherit a non-curved A_∞ -structure obtained by all possible insertions of Maurer-Cartan elements; thus, the differential in the Floer complex for (L_0, α_0) and (L_1, α_1) is given by

$$(8.2) \quad \begin{aligned} \mu_{\mathcal{F}(M; \lambda)}^1(x) &= \mu_{\mathcal{F}_{pr}(M)}^1(x) + \mu_{\mathcal{F}_{pr}(M)}^2(\alpha_1, x) + \mu_{\mathcal{F}_{pr}(M)}^2(x, \alpha_0) + \mu_{\mathcal{F}_{pr}(M)}^3(\alpha_1, \alpha_1, x) \\ &\quad + \mu_{\mathcal{F}_{pr}(M)}^3(\alpha_1, x, \alpha_0) + \mu_{\mathcal{F}_{pr}(M)}^3(x, \alpha_0, \alpha_0) + \cdots \end{aligned}$$

The Floer differential squares to zero, i.e. $\mathcal{F}(M; \lambda)$ has vanishing curvature, since $\lambda \cdot 1_L$ is central, because $\mathcal{F}_{pr}(M)$ was assumed to be strictly unital. The mapping class group $\pi_0 \text{Symp}(M)$ acts on $\mathcal{F}(M; \lambda)$ for each $\lambda \in \Lambda_{> 0}$ separately.

The open-closed map $OC : HH_*(\mathcal{F}(M; \lambda), \mathcal{F}(M; \lambda)) \rightarrow QH^*(M; \Lambda)$ takes values in the generalized λ -eigenspace of quantum multiplication by $c_1(TM)$, which is a subring

$$QH^*(M)_\lambda = \{a \in QH^*(M) : (c_1(TM) - \lambda)^N * a = 0 \text{ for some } N \in \mathbb{N}\} \subset QH^*(M).$$

The analogue of Abouzaid's generation criterion in this setting states that, if the restriction of the open-closed map to a full subcategory of $\mathcal{F}(M; \lambda)^{per}$ hits an invertible element of $QH^*(M)_\lambda$, then the full subcategory split-generates $\mathcal{F}(M; \lambda)^{per}$ [26, Theorem 11.3].

8.2. Relative parallel transport. Let (X, ω) be a symplectic manifold and $(Y_t)_{t \in [0,1]}$ a smooth family of symplectic submanifolds of X .

Lemma 8.1. *There is a time-dependent symplectomorphism $\Phi_t : X \rightarrow X$ with $\Phi_0 = \text{id}$ and $\Phi_t(Y_0) = Y_t$, well-defined up to isotopy through symplectomorphisms with the same property.*

Proof. The existence is [10, Proposition 4]. For uniqueness up to isotopy, observe that two choices $\Phi_{0,t}$ and $\Phi_{1,t}$ differ by a time-dependent symplectomorphism $\rho_{1,t} = \Phi_{0,t}^{-1} \circ \Phi_{1,t}$ which preserves Y_0 setwise, i.e. a path based at the origin in the group $\text{Symp}(X, Y_0)$ of symplectomorphisms preserving Y_0 setwise. The path $\rho_{1,t}$ can be deformed continuously to the constant path $\rho_{0,t} \equiv \text{id}$ (e.g. setting $\rho_{s,t} = \rho_{1,st}$), and $\Phi_{s,t} = \Phi_{0,t} \circ \rho_{s,t}$ gives an isotopy between $\Phi_{0,t}$ and $\Phi_{1,t}$. \square

Recall that a symplectic fibration $\mathcal{X} \rightarrow B$ with fibre (X, ω) is a smooth fibre bundle with a globally closed 2-form $\Omega_{\mathcal{X}}$ such that $(X_b, \Omega_{\mathcal{X}|_{X_b}}) \cong (X, \omega)$ for each $b \in B$.

Corollary 8.2. *Given a symplectic fibration $\mathcal{X} \rightarrow B$ with fibre X and a locally trivial symplectic subfibration $\mathcal{Y} \hookrightarrow \mathcal{X}$ with fibre $Y \subset X$ at a base-point $b \in B$, there is a relative monodromy representation*

$$\pi_1(B; b) \longrightarrow \pi_0 \text{Symp}(X, Y)$$

into the mapping class group of the group of symplectomorphisms preserving Y setwise.

Proof. Fix a 1-dimensional submanifold $\gamma \subset B$ and trivialise the bundle $\mathcal{X}|_{\gamma}$ over γ by symplectic parallel transport. This brings us into the situation of Lemma 8.1, meaning that we have a one-parameter family of symplectic embeddings (parametrized by a co-ordinate $t \in \gamma$) of Y into a fixed (X, ω) . By differentiating the relative Moser maps Φ_t of Lemma 8.1, we obtain closed 1-forms $a_t \in \Omega^1(X)$ for which the ω -dual vector fields Z_t flow the submanifolds Y_t into one another. If we subtract $dt \wedge da_t$ from $\Omega_{\mathcal{X}}$, we obtain a new globally closed 2-form on $\mathcal{X}|_{\gamma}$ with the correct fibrewise restriction and for which parallel transport preserves the subfibration $\mathcal{Y}|_{\gamma}$. We can apply the preceding construction to the 1-skeleton of B to obtain relative parallel transport maps for loops generating $\pi_1(B)$. The uniqueness up to isotopy in Lemma 8.1 shows the construction descends to a representation of $\pi_1(B)$. \square

Let ω_{st} denote the standard constant coefficient Kähler form on the four-torus T^4 . Consider $(\Sigma_2 \times T^4, \omega \oplus \omega_{st})$. We fix a sufficiently small $\varepsilon > 0$ and let $p : Z \rightarrow \Sigma_2 \times T^4$ denote the ε -symplectic blow-up of $\Sigma_2 \times T^4$ along the symplectic submanifold $C = \Sigma_2 \times \{0\}$. The exceptional divisor $E = C \times \mathbb{P}^1$ is canonically a product; indeed Z is just the product $\Sigma_2 \times Bl_{pt}(T^4)$, and carries a symplectic form Ω with cohomology class $p^*[\omega \oplus \omega_{st}] - \varepsilon \cdot E$.

The cohomology of Z admits a splitting

$$H^*(Z; \mathbb{Z}) = H^*(\Sigma_2 \times T^4; \mathbb{Z}) \oplus H^*(\Sigma_2; \mathbb{Z}) \cdot u$$

where $u = -PD(E)$ has degree 2. Let $\{\eta_j\}$ denote a basis for $H^1(Z; \mathbb{Z})$, and set

$$\Omega_{irr}^{\delta} = \Omega + \delta \cdot \sum_{i,j} c_{ij} \eta_i \wedge \eta_j$$

for coefficients $c_{ij} \in (0, 1)$. If $\delta > 0$ is sufficiently small then Ω_{irr}^{δ} is a symplectic form on Z , because the symplectic condition is open and the c_{ij} are bounded.

Lemma 8.3. *If the coefficients c_{ij} are linearly independent over \mathbb{Q} , then every symplectomorphism of (Z, Ω_{irr}^δ) acts trivially on $H^*(Z)$.*

Proof. Assume that the c_{ij} are rationally linearly independent. We will show that any diffeomorphism of Z preserving the cohomology class $[\Omega_{irr}]$ acts trivially on cohomology. As a ring, $H^*(Z)$ is generated by $H^1(Z) \simeq H^1(\Sigma_2) \oplus H^1(T^4)$ and by the class $u \in H^2(Z)$. Note that $\pi_2(Z)$ is generated by a fibre $F \subset E$, so any diffeomorphism preserves the class $[F] \in H_2(Z; \mathbb{Z})$ and its intersection Poincaré dual $E \in H_4(Z; \mathbb{Z})$. The action on $H^2(Z; \mathbb{R})$ of a diffeomorphism of Z which fixes $[\Omega_{irr}^\delta]$ has the eigenvalue 1 appearing with multiplicity at least two (since u and $[\Omega_{irr}^\delta]$ are both preserved). However, the action on $H^2(Z)/\langle u \rangle$ is the action on $\Lambda^2(H^1(Z))$. Since the action on H^1 is through an integral matrix, the coefficients of any eigenvector for the eigenvalue 1 must be linearly dependent over \mathbb{Q} . Therefore, preservation of $[\Omega_{irr}^\delta]$ implies that the diffeomorphism acts trivially on $H^1(Z)$ and hence on $H^*(Z)$. \square

The Torelli group $I_2 \leq \Gamma_2$ is an infinitely generated free group, generated by the Dehn twists on separating simple closed curves. The construction of relative parallel transport applied to a family of blow-ups yields a representation

$$\Gamma_2 \rightarrow \pi_0 \text{Symp}(Z, \Omega)$$

which depends on the same kinds of choice as in Section 2.6. Any element of Γ_2 acting non-trivially on cohomology cannot deform to a symplectomorphism with respect to the perturbed symplectic structure Ω_{irr}^δ , by Lemma 8.3. A given element of the Torelli group, however, will deform for δ sufficiently small.

Corollary 8.4. *Given $N > 0$, there is $\delta(N) > 0$ such that $I_2 \rightarrow \pi_0 \text{Symp}(Z, \Omega)$ deforms on a rank N free subgroup $\mathbb{F}_N \leq I_2$ to $\mathbb{F}_N \rightarrow \pi_0 \text{Symp}(Z, \Omega_{irr}^\delta)$ for all $\delta \in (0, \delta(N))$.*

Proof. The graph of $f \in I_2$ defines a Lagrangian submanifold $\Gamma(f) \subset (Z \times Z, \Omega \oplus -\Omega)$. Since the cohomology classes of the perturbing forms $\eta_i \wedge \eta_j$ restrict trivially to $\Gamma(f)$, using that $f^*([\eta]) = [\eta]$ for all η , if δ is sufficiently small there is a Lagrangian isotopy of $\Gamma(f)$ to a submanifold Lagrangian with respect to Ω_{irr}^δ . Since being graphical is an open condition, if δ is sufficiently small this is again the graph of a symplectomorphism f^{irr} . Pick N separating simple closed curves σ_j on Σ_2 . Then the corresponding elements f_{σ_j} admit common deformations $f_{\sigma_j}^{irr}$ to symplectomorphisms of (Z, Ω_{irr}^δ) if $\delta > 0$ is sufficiently small. (*A priori* the size of δ depends on geometric bounds on the Dehn twists about σ_j , hence cannot be made uniform as $N \rightarrow \infty$.) \square

Subsequently we will show that the homomorphism $\mathbb{F}_N \rightarrow \pi_0 \text{Symp}(Z, \Omega_{irr}^\delta)$ is faithful.

8.3. Unbounded rank. Let $\mathcal{L} = \mathcal{O}_{\mathbb{P}^1}(-1)$ denote the ε -blow-up of \mathbb{C}^2 at the origin, equipped with its toric Kähler form in cohomology class $\varepsilon \cdot u$, where $u = -PD(E)$ is the negative of the Poincaré dual to the exceptional divisor (zero-section). The Gromov invariant of E is non-trivial, and

$$QH^*(\mathcal{L}; \Lambda) = \Lambda[u]/\langle u(u + q^\varepsilon) \rangle$$

Note that $c_1(T\mathcal{L}) = u$. Implanting the local model into the four-torus, one finds that if $Y = Bl_{pt}(T^4)$ with the natural Kähler form $p^*\omega_{st} + \varepsilon \cdot u$, then

$$QH^*(Y; \Lambda) \cong H^*(T^4; \Lambda) \oplus \Lambda \cdot u$$

and the first Chern class $c_1(Y) = u = -PD(E)$ acts, under quantum multiplication, nilpotently on all cohomology classes of positive degree in $H^*(T^4; \Lambda)$ and invertibly on the second summand.

Lemma 8.5. $\mathcal{F}(Y; -q^\varepsilon)^{per}$ is semisimple and generated by an idempotent summand T^+ of a Lagrangian torus $T \subset Y$.

Proof. \mathcal{L} contains a (monotone) Lagrangian torus $T \subset \mathcal{O}_{\mathbb{P}^1}(-1)$, which is the orbit of the torus action corresponding to the unique critical point of the toric potential function $W(x, y) = x + y + q^{-\varepsilon}xy$. The torus T is weakly unobstructed, defines an object of $\mathcal{F}(\mathcal{L}; -q^\varepsilon)$ whose Floer cohomology is semisimple in characteristic zero, and splits into the direct sum of two idempotent summands T^\pm , which are isomorphic up to shift [37, Section 4.4].

A neighborhood of the zero section in \mathcal{L} (large enough to contain T) embeds into a neighborhood of the exceptional divisor in $Y = Bl_{pt}(T^4)$; since all holomorphic discs bounded by T in Y must be contained inside the neighborhood of the exceptional divisor, the Floer cohomology of the torus T in Y is exactly as in \mathcal{L} . Hence T also defines an object of $\mathcal{F}(Y; -q^\varepsilon)$ with semisimple Floer cohomology, which splits into two idempotent summands T^\pm . An explicit calculation shows that the images of the two idempotents of $HF(T, T)$ under the open-closed map are $\pm u = \mp PD(E)$; the generation criterion [26, Theorem 11.3] then implies that T split-generates $\mathcal{F}(Y; -q^\varepsilon)^{per}$. \square

Proposition 8.6. *There is a fully faithful functor $\mathcal{F}(\Sigma_2) \rightarrow \mathcal{F}(Z, \Omega_{irr}^\delta; \lambda)^{per}$ whose image split-generates $\mathcal{F}(Z, \Omega_{irr}^\delta; \lambda)^{per}$ for the eigenvalue $\lambda = -q^\varepsilon$.*

Sketch. For the product form Ω on Z , there is a Künneth functor associated to the A_∞ -tensor product $\mathcal{F}(Z; \lambda) \simeq \mathcal{F}(\Sigma_2) \otimes \mathcal{F}(Y; \lambda)$ and the semisimple piece of the second factor afforded by Lemma 8.5. Note that $T \subset Y$ survives arbitrary small perturbations of the given Kähler form $p^*\omega_{st} + \varepsilon \cdot u$ on Y , in the sense that it deforms as a Lagrangian to any sufficiently nearby symplectic form, since the restriction map $H^2(Y) \rightarrow H^2(T)$ vanishes. Since its Floer cohomology with respect to the initial symplectic form is semisimple, it must remain semisimple after small deformation.

Choose the 1-forms η_i on $Z = \Sigma_2 \times Y$ so that that η_1, \dots, η_4 are the pullbacks of closed 1-forms $\alpha_1, \dots, \alpha_4$ on Σ_2 representing a basis of $H^1(\Sigma_2; \mathbb{Z})$, and η_5, \dots, η_8 are the pullbacks of closed 1-forms on Y which vanish everywhere in a neighborhood of the exceptional divisor (and in particular on the torus T).

Then, for a fixed simple closed curve $\gamma \subset \Sigma_2$, the submanifold $\gamma \times T$ is Lagrangian not just for the product form Ω , but also for $\Omega_{irr}^\delta = \Omega + \delta \sum_{ij} c_{ij} \eta_i \wedge \eta_j$.

The association

$$\gamma \mapsto \gamma \times T^+$$

is globally realised by an A_∞ -functor associated to a Lagrangian correspondence

$$G \subset \Sigma_2^- \times Z = \Sigma_2^- \times (\Sigma_2 \times Bl_{pt}(T^4))$$

which fibres over the diagonal of $\Sigma_2^- \times \Sigma_2$ with fibre $T \subset \mathcal{L} \subset Y$. Note that $\Delta_{\Sigma_2} \times T$ remains Lagrangian after perturbing the symplectic form on Z by $\delta \sum_{ij} c_{ij} \eta_i \wedge \eta_j$, and that on Σ_2 by $\delta \sum_{ij} c_{ij} \alpha_i \wedge \alpha_j$. Since the correspondence is globally a product, the local-to-global spectral sequence $H^*(\Sigma_2; HF^*(T, T)) \rightarrow HF^*(G, G)$ degenerates; therefore the correspondence G itself has an idempotent summand G^+ associated to a choice of idempotent for T , compare to [31, Section 5.4]. This yields the desired functor \mathcal{G}^+ from $\mathcal{F}(\Sigma_2)$ (where we suppress from the notation the fact that the symplectic form depends on δ) to $\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per}$. This functor maps every object of $\mathcal{F}(\Sigma_2)$ to its product with T^+ ; since we are in a product situation and $\text{End}_{\mathcal{F}(Y; -q^\varepsilon)^{per}}(T^+) = \Lambda$, the functor is fully faithful.

The fact that the image of the functor \mathcal{G}^+ split-generates $\mathcal{F}(Z; -q^\varepsilon)^{per}$ follows from Ganatra's automatic generation result [15], since $HH^*(\mathcal{F}(\Sigma_2), \mathcal{F}(\Sigma_2)) \simeq H^*(\Sigma_2; \Lambda)$ has the same rank as $QH^*(Z)_{-q^\varepsilon} \simeq u \cdot H^*(\Sigma_2; \Lambda)$, and $\mathcal{F}(\Sigma_2)$ is homologically smooth. One can also proceed more directly: consider a collection of curves $\gamma_i \in \mathcal{F}(\Sigma_2)$ which satisfy Abouzaid's split-generation criterion, i.e. the full subcategory with this set of objects has a Hochschild cycle α which maps to the unit $1 \in H^*(\Sigma_2; \Lambda)$ under the open-closed map. Then the Hochschild cycle $\alpha^+ = \mathcal{G}_*(\alpha)$, formed by replacing every morphism which appears in α with its tensor product with 1_{T^+} , maps to $1 \otimes u \in QH^*(Z; \Lambda)$ (using the fact that we are locally in a product situation and $OC(1_{T^+}) = u$). Since u is invertible in $QH^*(Z)_{-q^\varepsilon}$, [26, Theorem 11.3] implies that the objects $\gamma_i \times T^+ = \mathcal{G}^+(\gamma_i)$ split-generate $\mathcal{F}(Z; -q^\varepsilon)^{per}$. \square

Together with Corollary 1.2, this proposition yields a natural map

$$\text{Auteq}((\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per}) \simeq \text{Auteq}(\mathcal{F}(\Sigma_2)^{per}) \rightarrow \Gamma_2,$$

which then induces a map

$$\text{Auteq}_{HH}(\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per}) \rightarrow I_2$$

where the domain denotes those autoequivalences which act trivially on Hochschild cohomology. We know that the closed-open map

$$H^*(\Sigma_2) \cdot u = QH^*(Z)_{-q^\varepsilon} \rightarrow HH^*(\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per})$$

is an isomorphism, and the map

$$\pi_0 \text{Symp}(Z, \Omega_{irr}^\delta) \rightarrow \text{Auteq}(\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per})$$

lands in the subgroup Auteq_{HH} by Lemma 8.3. Combining this with Corollary 8.4, we therefore obtain a map

$$\mathbb{F}_N \rightarrow \pi_0 \text{Symp}(Z, \Omega_{irr}^\delta) \rightarrow \text{Auteq}_{HH}(\mathcal{F}(Z, \Omega_{irr}^\delta; -q^\varepsilon)^{per}) \rightarrow I_2 \cong \mathbb{F}_\infty$$

which one can compose with the quotient map $q : \mathbb{F}_\infty \rightarrow \mathbb{F}_N$ which kills all but the finite set of chosen generators (twists in all other separating simple closed curves). Chasing through the stages, the composite map is the natural inclusion of a finite rank free subgroup of the infinitely generated free group, or the identity after composition with q . It follows that $\pi_0 \text{Symp}(Z, \Omega_{irr}^\delta)$ surjects onto a free group of rank $N = N(\delta)$ which tends to infinity as $\delta \rightarrow 0$. This completes the proof of Theorem 1.3.

REFERENCES

1. Mohammed Abouzaid, *A geometric criterion for generating the Fukaya category*, Publ. Math. Inst. Hautes Études Sci. (2010), no. 112, 191–240.
2. ———, *Family Floer cohomology and mirror symmetry*, Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. II, Kyung Moon Sa, Seoul, 2014, pp. 813–836.
3. ———, *The family Floer functor is faithful*, J. Eur. Math. Soc. (JEMS) **19** (2017), no. 7, 2139–2217.
4. Mohammed Abouzaid, Denis Auroux, Alexander I. Efimov, Ludmil Katzarkov, and Dmitri Orlov, *Homological mirror symmetry for punctured spheres*, J. Amer. Math. Soc. **26** (2013), no. 4, 1051–1083.
5. Mohammed Abouzaid, Denis Auroux, and Ludmil Katzarkov, *Lagrangian fibrations on blowups of toric varieties and mirror symmetry for hypersurfaces*, Publ. Math. Inst. Hautes Études Sci. **123** (2016), 199–282.
6. Mohammed Abouzaid and Paul Seidel, *An open string analogue of Viterbo functoriality*, Geom. Topol. **14** (2010), no. 2, 627–718.
7. Mohammed Abouzaid and Ivan Smith, *Exact Lagrangians in plumbings*, Geom. Funct. Anal. **22** (2012), no. 4, 785–831.
8. Miguel Abreu and Dusa McDuff, *Topology of symplectomorphism groups of rational ruled surfaces*, J. Amer. Math. Soc. **13** (2000), no. 4, 971–1009.
9. Yves André, *Différentielles non commutatives et théorie de Galois différentielle ou aux différences*, Ann. Sci. École Norm. Sup. (4) **34** (2001), no. 5, 685–739.
10. D. Auroux, *Asymptotically holomorphic families of symplectic submanifolds*, Geom. Funct. Anal. **7** (1997), no. 6, 971–995.
11. Lev A. Borisov and R. Paul Horja, *On the K-theory of smooth toric DM stacks*, Snowbird lectures on string geometry, Contemp. Math., vol. 401, Amer. Math. Soc., Providence, RI, 2006, pp. 21–42.
12. Michel Brion, *Some structure theorems for algebraic groups*, Algebraic groups: structure and actions, Proc. Sympos. Pure Math., vol. 94, Amer. Math. Soc., Providence, RI, 2017, pp. 53–126.
13. Igor Burban and Yuriy Drozd, *Coherent sheaves on rational curves with simple double points and transversal intersections*, Duke Math. J. **121** (2004), no. 2, 189–229.
14. Hsien-Chih Chang, Jeff Erickson, David Letscher, Arnaud de Mesmay, Saul Schleimer, Eric Sedgwick, Dylan Thurston, and Stephan Tillmann, *Tightening curves on surfaces via local moves*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2018, pp. 121–135.
15. Sheel Ganatra, *Automatically generating Fukaya categories and computing quantum cohomology*, Preprint, available at arXiv:1605.07702.
16. ———, *Symplectic Cohomology and Duality for the Wrapped Fukaya Category*, Preprint, available at arXiv:1304.7312.
17. Sheel Ganatra, John Pardon, and Vivek Shende, *Sectorial descent for wrapped Fukaya categories*, Preprint, available at arXiv:1809.03427.
18. M. Gromov, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347.
19. F. Haiden, L. Katzarkov, and M. Kontsevich, *Flat surfaces and stability structures*, Publ. Math. Inst. Hautes Études Sci. **126** (2017), 247–318.
20. J. Hanselman, J. Rasmussen, and L. Watson, *Bordered Floer homology for manifolds with torus boundary via immersed curves*, Preprint, available at arXiv:1604.03466.
21. Joel Hass and Peter Scott, *Intersections of curves on surfaces*, Israel J. Math. **51** (1985), no. 1-2, 90–120.
22. ———, *Shortening curves on surfaces*, Topology **33** (1994), no. 1, 25–43.
23. Yusuf Baris Kartal, *Dynamical invariants of mapping torus categories*, Preprint, available at arXiv:1809.04046.
24. Heather Lee, *Homological mirror symmetry for open Riemann surfaces from pair-of-pants decompositions*, Preprint, available at arXiv:1608.04473.
25. L. A. Nazarova and A. V. Roĭter, *A certain problem of I. M. Gel'fand*, Funkcional. Anal. i Priložen. **7** (1973), no. 4, 54–69.
26. Alexander F. Ritter and Ivan Smith, *The monotone wrapped Fukaya category and the open-closed string map*, Selecta Math. (N.S.) **23** (2017), no. 1, 533–642.

27. Paul Schmutz Schaller, *Mapping class groups of hyperbolic surfaces and automorphism groups of graphs*, *Compositio Math.* **122** (2000), no. 3, 243–260.
28. Paul Seidel, *Fukaya categories and Picard–Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society, Zürich, 2008.
29. ———, *Homological mirror symmetry for the genus two curve*, *J. Algebraic Geom.* **20** (2011), no. 4, 727–769.
30. ———, *Lagrangian homology spheres in (A_m) Milnor fibres via \mathbb{C}^* -equivariant A_∞ -modules*, *Geom. Topol.* **16** (2012), no. 4, 2343–2389.
31. ———, *Abstract analogues of flux as symplectic invariants*, *Mém. Soc. Math. Fr. (N.S.)* (2014), no. 137, 135.
32. ———, *Lectures on categorical dynamics and symplectic topology.*, Unpublished manuscript, available at <http://math.mit.edu/~seidel/>, 2014.
33. ———, *Homological mirror symmetry for the quartic surface*, *Mem. Amer. Math. Soc.* **236** (2015), no. 1116, vi+129.
34. ———, *Picard–Lefschetz theory and dilating \mathbb{C}^* -actions*, *J. Topol.* **8** (2015), no. 4, 1167–1201, with an Appendix available at arXiv:1403.7571v2.
35. Nick Sheridan, *Homological mirror symmetry for Calabi–Yau hypersurfaces in projective space*, *Invent. Math.* **199** (2015), no. 1, 1–186.
36. Dmytro Shklyarov, *Hirzebruch–Riemann–Roch-type formula for DG algebras*, *Proc. Lond. Math. Soc. (3)* **106** (2013), no. 1, 1–32.
37. I. Smith, *Floer cohomology and pencils of quadrics*, *Invent. Math.* **189** (2012), no. 1, 149–250.
38. Ernst Steinitz, *Polyeder und Raumeinteilungen*, 1916, Encyclopädie der mathematischen Wissenschaften, Band 3 (Geometries) (IIIAB12), pp. 1–139.
39. Ernst Steinitz and Hans Rademacher, *Vorlesungen über die Theorie der Polyeder unter Einschluss der Elemente der Topologie*, Springer-Verlag, Berlin-New York, 1976, Reprint der 1934 Auflage, Grundlehren der Mathematischen Wissenschaften, No. 41.
40. Charles Weibel, *The negative K-theory of normal surfaces*, *Duke Math. J.* **108** (2001), no. 1, 1–35.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, CAMBRIDGE MA 02138, USA.

Email address: auroux@math.harvard.edu

IVAN SMITH, CENTRE FOR MATHEMATICAL SCIENCES, UNIVERSITY OF CAMBRIDGE, CB3 0WB, ENGLAND.

Email address: is200@cam.ac.uk