

HOMOLOGICAL MIRROR SYMMETRY WITHOUT CORRECTION

MOHAMMED ABOUZAID

ABSTRACT. Let X be a closed symplectic manifold equipped a Lagrangian torus fibration over a base Q . A construction first considered by Kontsevich and Soibelman produces from this data a rigid analytic space Y , which can be considered as a variant of the T -dual introduced by Strominger, Yau, and Zaslow. We prove that the Fukaya category of tautologically unobstructed graded Lagrangians in X embeds fully faithfully in the derived category of coherent sheaves on Y , under the technical assumption that $\pi_2(Q)$ vanishes (all known examples satisfy this assumption). The main new tool is the construction and computation of Floer cohomology groups of Lagrangian fibres equipped with topological infinite rank local systems that correspond, under mirror symmetry, to the affinoid rings introduced by Tate, equipped with their natural topologies as Banach algebras.

CONTENTS

1. Introduction	1
2. Cohomological constructions	10
3. Higher moduli spaces	58
4. Chain level constructions	80
Appendix A. Reverse isoperimetric inequalities	89
Appendix B. Tate's acyclicity theorem	96
Appendix C. Computations of cohomology groups with twisted coefficients	99
Appendix D. Invariance of the Family Floer functor	107
References	109

1. INTRODUCTION

1.1. Statement of the main result. Let X be a closed symplectic manifold. One of the key tools used to understand the symplectic topology of X is its Fukaya category, placing the computation of this category as a central problem in the subject. Such computations would ideally rely on geometric features of X . For example, the fact that cotangent bundles are fibered in Lagrangian planes ultimately accounts for the computation of their (derived) Fukaya categories as categories of modules over the chains on the based loop space of the base [2].

Date: October 4, 2022.

The author was supported by NSF grants DMS-1308179, DMS-1609148, and DMS-1564172, the Simons Foundation through its “Homological Mirror Symmetry” Collaboration grant, the Erik Ellentuck Fellowship, and the IAS Fund of Math.

The case of a symplectic manifold equipped with a Lagrangian torus fibration has been the focus of much interest because of its relevance to mirror symmetry via the Strominger-Yau-Zaslow conjecture [28]. Kontsevich and Soibelman [23] were the first to propose a mathematically precise conjecture in this context: under the assumption that the torus fibration admits a Lagrangian section, they showed that one can associate a rigid analytic space Y (in the sense of Tate [30]) to every such symplectic manifold, and conjectured that the Fukaya category is equivalent to the category of (rigid analytic) coherent sheaves on Y . They also took the first step in this direction by assigning a mirror line bundle on Y to each Lagrangian section of X , and comparing the multiplication on the Floer cohomology groups of Lagrangian sections with the composition of (derived) maps between the mirror line bundles.

The next step was taken by Fukaya in [14], showing in complete generality that the Floer cohomology of Lagrangian submanifolds varies *analytically* with respect to the natural local coordinates on the space of Lagrangians modulo Hamiltonian isotopy coming from the flux homomorphism. Fukaya's work was phrased in terms of the self-Floer cohomology of Lagrangians, but a minor adaptation shows that one can assign to an object L of the Fukaya category of X a *sheaf \mathcal{L}_L of coherent analytic complexes* on the mirror space [5] (the notation \mathcal{L} refers to the fact that this complex will arise for us as a *left module* over a category associated to a cover of the mirror).

To avoid technical difficulties, we shall from now on restrict attention to Lagrangian torus fibrations whose bases have vanishing second homotopy group, and consider only the subcategory of the Fukaya category consisting of *tautologically unobstructed Lagrangians*, i.e. those for which there is a choice of almost complex structure so that all the holomorphic discs which they bound are constant. Since the homotopy groups of the base of a fibration are isomorphic to the relative homotopy groups of the inclusion of each fibre in the total space, the first condition implies that $\pi_2(X, X_q)$ vanishes for each fibre X_q , hence that any holomorphic disc bounded by such fibres is constant. Since the main constructions of this paper rely on studying the Floer theory of such torus fibres, this condition implies that we may construct Floer cochain complexes without having to choose bounding cochains as in the work of Fukaya, Oh, Ohta, and Ono [17].

In this context, the family Floer functor was shown to be faithful in [6]. More precisely, associated to a Lagrangian torus fibration is a gerbe β on Y , which is classified by an element of $H^2(Y, \mathcal{O}^*)$. An A_∞ functor from the Fukaya category of X to the derived category of β -twisted coherent sheaves on Y was constructed, and the corresponding family Floer map

$$(1.1.1) \quad HF^*(L, L') \rightarrow H^*(Y; \text{Hom}_{\mathcal{O}_Y}(\mathcal{L}_{L'}, \mathcal{L}_L))$$

was proved to be injective for every pair of Lagrangians. Here, and in contrast to the rest of the paper, we use the notation $H^*(Y; \cdot)$ to indicate the cohomology of Y with coefficients in a sheaf, noting that the morphisms between β -twisted sheaves form an ordinary sheaf on Y . In order to state the main result of this paper, we observe that the structure of a Lagrangian torus fibration equips the total space X with a canonical grading in the sense of Seidel [25], which allows us to restrict our attention to graded Lagrangians, i.e. those which define objects of the \mathbb{Z} -graded Fukaya category:

Theorem 1.1. *If X is a closed symplectic manifold equipped with a Lagrangian torus fibration $X \rightarrow Q$ whose base has vanishing second homotopy group, then the family Floer map in Equation (1.1.1) is surjective for graded Lagrangians.*

Using the main result of [6], which asserts the injectivity of the Family Floer map, we conclude:

Corollary 1.2. *There is a fully faithful embedding of the Fukaya category of tautologically unobstructed and graded Lagrangians in X in the β -twisted derived category of coherent sheaves on Y .* \square

Remark 1.3. The restriction to graded Lagrangians is required for technical reasons, in particular in Section 4.3. We expect that it can be eliminated at the cost of more carefully using the natural filtration of operations in Floer theory by energy, and of the structure sheaf of Y by the valuation of functions.

The main deficiencies of this result are its restrictive assumptions that (1) all Lagrangians are tautologically unobstructed (and that $\pi_2(Q) = 0$), and (2) that the ambient symplectic manifold is equipped with a *non-singular* fibration. Removing the first assumption would require the use of a package of virtual fundamental chains; the one developed by Fukaya-Oh-Ohta-Ono [17] would be sufficient for the task at hand. The decision not to use it to prove a theorem for general Lagrangians amounts to the desire not to add another layer of complexity to the paper. On the other hand, admitting singular Lagrangians as fibres, e.g. immersed Lagrangians, will require some new insights about Floer theory in families, though the first steps have been taken by Fukaya in his announced results about mirror symmetry for $K3$ surfaces.

Remark 1.4. The construction of the family Floer functor, which is implemented in the Appendix of [6], goes beyond the construction of the chain level analogue of Equation (1.1.1), and requires a comparison between composition in the two categories. Since the Fukaya category is an A_∞ category, there is a tower of maps, generalising the chain level analogue of Equation (1.1.1), which are required to specify the functor, but these are not discussed in this paper (except in Appendix D) since the assertion that a functor is fully faithful is a statement only about Equation (1.1.1) being an isomorphism.

Remark 1.5. A similar result holds if we twist the Fukaya category of X with a background or bulk class arising from Q : letting Λ_+ denote the elements of the Novikov field with strictly positive valuation (see Equation (1.2.2)), each class $\alpha \in H^2(X; \mathbb{Z}_2) \oplus H^2(X; \Lambda_+)$ determines a deformation of the Fukaya category of X . If this class is obtained by pullback from Q , it lifts to a class in $H^2(Y; \mathcal{O}^*)$, which we also denote α . The α -twisted Fukaya category of X can be shown by our methods to embed fully faithfully in the $\alpha + \beta$ twist of the derived category of coherent sheaves on Y .

Our methods are currently insufficient to address the case of more general twists: these give rise to locally non-commutative deformations of the rings of functions on Y , and we do not currently know how to formulate some of the required algebraic results in this context (in particular, Tate's acyclicity theorem as discussed in Appendix B).

1.2. Affinoid rings and local systems. In order to introduce the new ideas in this paper, we begin by explaining its background in more detail: for specificity, we fix a field \mathbf{k} , and denote by Λ the corresponding Novikov field consisting of series

$$(1.2.1) \quad \sum_{\lambda \in \mathbb{R}} c_\lambda T^\lambda, \quad c_\lambda \in \mathbf{k}$$

with the property that the set of exponents λ for which the coefficients c_λ do not vanish is discrete and bounded below. The Fukaya category that we consider will be defined over Λ , in the sense that morphism are Λ -vector spaces and all operations are linear over Λ . The

Novikov ring admits a valuation

$$(1.2.2) \quad \text{val}\left(\sum_{\lambda \in \mathbb{R}} c_\lambda T^\lambda\right) = \lambda_0$$

where λ_0 is the smallest exponent whose coefficient does not vanish (we set $\text{val} 0 = +\infty$).

The points of the space Y are in bijective correspondence with (isomorphism classes) of simple objects of the Fukaya category of X , supported on the fibres of the projection to Q . As discussed above, our assumption that $\pi_2(Q)$ vanishes implies that the fibres do not bound any non-constant holomorphic disc, and hence that such isomorphism classes naturally correspond to isomorphism classes of unitary rank-1 local systems on the fibres.

In order to describe Y as an analytic space, one must describe a topology and a sheaf of functions. The simplest description is to observe that the base of a Lagrangian torus fibration is equipped with an integral affine structure; in particular Q is obtained by gluing integral affine polytopes (i.e. bounded subsets of \mathbb{R}^n defined by affine equations whose linear terms are integral), along integral affine transformations of \mathbb{R}^n (i.e. compositions of translations and elements of $GL(n, \mathbb{Z})$). In much of this paper, we shall express such a decomposition by considering a partially ordered set Σ whose elements index a polyhedral cover $\{P_\sigma\}_{\sigma \in \Sigma}$ of Q . Under the valuation map from $(\Lambda \setminus \{0\})^n$ to \mathbb{R}^n , the inverse image of each such polytope $P_\sigma \subset \mathbb{R}^n$ is an affinoid domain Y_σ in the sense of Tate [30]. Since the valuation map intertwines the operations

$$(1.2.3) \quad (z_1, \dots, z_n) \mapsto (T^{\lambda_1} \prod z_j^{a_{1j}}, \dots, T^{\lambda_n} \prod z_j^{a_{nj}})$$

$$(1.2.4) \quad (u_1, \dots, u_n) \mapsto (\lambda_1 + \sum a_{1j} u_j, \dots, \lambda_n + \sum a_{nj} u_j)$$

on $(\Lambda \setminus \{0\})^n$ and \mathbb{R}^n , which are associated to each collection $(\lambda_1, \dots, \lambda_n)$ of real numbers and integral matrix (a_{ij}) , the gluing of polytopes which yields Q lifts to a gluing of the affinoid domains $\{Y_\sigma\}_{\sigma \in \Sigma}$ which yields a rigid analytic space which we denote Y , and this space is independent of the choice of cover. By construction, the space Y is equipped with a projection map to Q , and the fibres are easily seen, via the above description, to correspond to unitary local systems on the fibres of the Lagrangian torus fibration over Q .

The starting point of the results of [6] is that affinoid rings are related to local systems on torus fibres as follows: picking a basepoint $q_\sigma \in P_\sigma$, we can identify the ring Γ^{P_σ} of rigid analytic functions on Y_σ as a completion of the group ring of the fundamental group of X_{q_σ} , which is isomorphic to $H_1(X_{q_\sigma}; \mathbb{Z})$ because X_{q_σ} is a torus. The basic idea is that P_σ can be identified with a polytope in $H^1(X_{q_\sigma}; \mathbb{R})$, so it makes sense to complete the group ring of $H_1(X_{q_\sigma}; \mathbb{Z})$ by allowing infinite series with the property that norm of the pairing of the terms in the series with all elements of P_σ converges to 0 (i.e. the valuation goes to $+\infty$). The relationship with the familiar idea of studying Floer theory with coefficients in rank-1 local systems is that elements of $H_1(X_{q_\sigma}; \mathbb{Z})$ define regular functions on the space of rank-1 local systems (over the Novikov ring). The set of rank-1 local systems is parametrised by the first cohomology with coefficients in $\Lambda \setminus \{0\}$, and the valuation defines a natural map

$$(1.2.5) \quad H^1(X_{q_\sigma}; \Lambda \setminus \{0\}) \rightarrow H^1(X_{q_\sigma}; \mathbb{R}).$$

The completion associated to the polytope P_σ considered as a subset of $H^1(X_{q_\sigma}; \mathbb{R})$ thus corresponds to the ring of analytic functions on the inverse image in $H^1(X_{q_\sigma}; \Lambda \setminus \{0\})$. We thus conclude that the spaces Y_σ can be thought of as spaces of rank-1 local systems with constraints on their valuation.

1.3. From Lagrangians to coherent sheaves. In order to prove Homological mirror symmetry, we must specify a model for the derived category of Y . As in [6] we shall use the affinoid covering of Y to express this category in terms of a Čech model. Indeed, a classical result of Kiehl [22] implies that the derived category of coherent sheaves on each affinoid domain is equivalent to the derived category of finitely presented modules over the corresponding ring of functions which we denoted Γ^{P_σ} , so that derived morphisms between sheaves can be computed as Čech cohomology groups, as is done for example in [9]. In particular the Čech cohomology groups are independent of the choice of affinoid cover.

We shall use the differential graded model of the derived category of coherent sheaves on Y given by a subcategory of the category of A_∞ -modules over a category \mathcal{F} with objects elements of the cover, with morphisms from τ to σ vanishing unless τ precedes σ in the partial order, in which case they are given by the affinoid ring Γ^{P_τ} : the subcategory of modules over \mathcal{F} that corresponds to coherent sheaves consists of those A_∞ modules whose associated cohomological modules assign to each element σ of the cover a finitely generated graded module over Γ^{P_σ} , and for which the restriction map associated to an arrow $\tau \rightarrow \sigma$ in \mathcal{F} induces an isomorphism on cohomology after tensoring over Γ^{P_τ} with Γ^{P_σ} . In the most basic case of a space covered by two affinoid domains, this encodes the idea that a (coherent) sheaf on the ambient space can be thought of as a pair of modules over the rings of functions on the two elements of the cover, a module over the ring of functions on the intersection, together with an isomorphism between the module associated to the intersection and the restrictions of the modules associated to the two elements of the cover. This is equivalent, though less economical, than the data of the two modules together with an isomorphism between their restrictions, but it proves to be much more convenient when formulating the corresponding notion at the level of cochain complexes.

Returning to the point of view that Γ^{P_σ} is a completion of the group ring of the fundamental group of X_{q_σ} , the usual correspondence between modules over the group ring and local systems associates to σ a local system U^σ on the Lagrangian X_{q_σ} , which yields an object of an enlargement of the Fukaya category in which infinite rank local systems are allowed. We can then associate to each (tautologically unobstructed) Lagrangian L , and each element of the cover a cochain complex

$$(1.3.1) \quad \mathcal{L}_L(\sigma) := CF^*(L, (X_{q_\sigma}, U^\sigma))$$

via Floer theory. This complex admits a natural action (on the right) by the ring Γ^{P_σ} , and the collection of all such cochain complexes underlies the left module \mathcal{L}_L over \mathcal{F} , which is the image of L under the Family Floer functor. In particular, given an arrow $\sigma \rightarrow \tau$, we have a restriction map

$$(1.3.2) \quad \mathcal{L}_L(\tau) \rightarrow \mathcal{L}_L(\sigma)$$

induced by the map of local systems $U^\sigma \rightarrow U^\tau$.

Having introduced the functor at the level of objects, we now consider morphisms: given a pair (L, L') of Lagrangians, the pair of pants product defines a map

$$(1.3.3) \quad CF^*(L', (X_{q_\sigma}, U^\sigma)) \otimes_\Lambda CF^*(L, L') \rightarrow CF^*(L, (X_{q_\sigma}, U^\sigma)).$$

Identifying the Floer cochains with coefficients in U^σ with the values of the modules \mathcal{L}_L and $\mathcal{L}_{L'}$ at σ , this defines a Λ -linear map

$$(1.3.4) \quad CF^*(L, L') \rightarrow \text{Hom}_\Lambda(\mathcal{L}_L(\sigma), \mathcal{L}_{L'}(\sigma)).$$

The usual arguments for associativity of multiplication in Floer theory imply that this map is naturally compatible with the above restriction map (and the module action by the affinoid

rings associated to our cover), yielding a map

$$(1.3.5) \quad CF^*(L, L') \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_L, \mathcal{L}_{L'}),$$

which is the cochain level model that we use for Equation (1.1.1). As discussed in Remark 1.4, one can go further and construct an A_∞ functor from the Fukaya category to the category of modules over \mathcal{F} , of which this is the first order map, but this extension will only be enter our discussion in Appendix D.

1.4. A summary of the proof. The proof of Theorem 1.1 is based on the following idea: in addition to the left module \mathcal{L}_L considered above, Floer theory associates to each Lagrangian L a right module \mathcal{R}_L over \mathcal{F} , given by

$$(1.4.1) \quad \mathcal{R}_L(\sigma) := CF^*((X_{q\sigma}, U^\sigma), L).$$

There is a natural map, which was essentially already used in [6], from the derived tensor product of these modules to the Lagrangian Floer cochains

$$(1.4.2) \quad \mathcal{R}_{L'} \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow CF^*(L, L').$$

The surjectivity of Equation (1.1.1) can thus be deduced from the sujectivity of the composition

$$(1.4.3) \quad \mathcal{R}_{L'} \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow CF^*(L, L') \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \mathcal{L}_L)$$

at the level of cohomology; in fact, we shall show that this composition is a quasi-isomorphism.

The holomorphic discs giving rise to this composition are illustrated on the right of Figure 1, where the label by polytopes indicates that we are considering fibres equipped with the corresponding local systems, the arrows associated to inputs point towards the interior of the disc, and the outward pointing arrow corresponds to the output. Each disc is labelled by the corresponding map or diagram. One should interpret Figure 1 as illustrating a cobordism

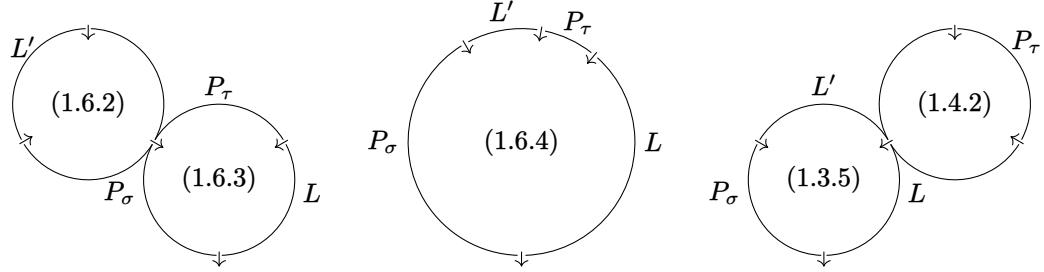


FIGURE 1. An informal representation of the moduli spaces giving rise to two compositions in Diagram (1.6.4), together with the moduli space giving rise to the homotopy between them. The arrows indicate inputs (pointing towards the disc), or outputs.

between the moduli spaces of discs used to define the two maps we are considering and a moduli space of discs which we shall use to prove that the composition is an isomorphism. The essential problem is to give a meaning, both in geometry and algebra, to the picture shown on the left. At the level of geometry, the idea used in [6] of considering families of fibres over the two polytopes runs into a transversality problem, because two fibres intersect if and only if they are equal, and it seems difficult to arrange for such a parametrised

problem (in which the intersections of Lagrangians change), to give moduli spaces of the correct dimension.

As is standard in Floer theory, we resolve this issue by introducing Hamiltonian perturbations, and associate to the node in Figure 1 with label (P_τ, P_σ) the set of intersections between a fibre over P_τ , and the image of a fibre over P_σ under a Hamiltonian isotopy. This leads to a new problem, which is responsible for the first real innovation of this paper: we need the Floer cohomology groups

$$(1.4.4) \quad HF^*((X_{q_\tau}, U^\tau), (X_{q_\sigma}, U^\sigma)),$$

associated to a perturbed pair to have appropriate meaning under mirror symmetry. At the least, we need to know that the perturbed Floer cohomology of the pair $P_\sigma = P_\tau$ is isomorphic to the affinoid ring associated to the corresponding subset of Y . This will require extending Floer theory to a setting in which Lagrangians are equipped with topologised local systems.

1.5. Topological local systems in Floer theory. The reason that the standard definition of Floer cohomology with local coefficients, as discussed e.g. in [4] in the infinite rank case, does not give an adequate definition of the group in Equation (1.4.4) is that this homology group is far larger than the desired one. This is analogous to the statement that, if we consider a power series ring $\mathbf{k}[[x]]$ as a module over the polynomial ring $\mathbf{k}[x]$ in the natural way, then the map

$$(1.5.1) \quad \mathbf{k}[[x]] \rightarrow \text{Hom}_{\mathbf{k}[x]}(\mathbf{k}[[x]], \mathbf{k}[[x]])$$

is far from being an isomorphism. A solution is offered by the fact that a power series ring acquires a natural topology from its description as an inverse limit, with respect to which it is complete. If we consider instead *continuous* morphisms, then the above map becomes an isomorphism.

Remark 1.6. While the foundations of rigid analytic geometry, following Tate [30], relies upon equipping rings of affinoid functions with a topology, subsequent developments in the subject largely avoid using topological methods after the first initial steps. The approach we shall therefore take goes against the standard framework in rigid analytic geometry, though it ties with the recent introduction of completions with respect to the action filtration in Floer theory [27, 19, 32, 31].

We are thus led to consider the Floer theory of Lagrangians equipped with *topological* local systems, taking the topologies into account when defining Floer cohomology groups; the reader may find Section 2.2, where we consider the simplified setting of Morse theory, to be a good place to first read about how one precisely takes the topology into account. The idea of incorporating topological vector spaces in the study of Floer theory probably goes back to Fukaya [13], who intended to use it to study Lagrangian foliations. We do not develop the general theory, limiting ourselves to those properties required for the proof of the main theorem. Among the indications that this approach gives the right answer is that we succeed in proving that the self-Floer cohomology of a Lagrangian with such a local system can be computed using either a Morse-theoretic model or appropriate, though not all, Hamiltonian perturbations (see Section 2.7.2).

The fact that the Floer theory that we consider is not invariant under all Hamiltonian isotopies will be essential to the success of our approach to proving Theorem 1.1, since it will allow us to produce an isomorphism in Equation (1.4.4) even when the fibres are distinct, as long as they are sufficiently close, and the polytopes are sufficiently small. To understand

the underlying ideas, recall that a basic consequence of Gromov's compactness theorem is that Floer homology for Lagrangians equipped with unitary local systems is invariant under Hamiltonian isotopies, but that Lagrangians equipped with non-unitary local systems do not have well-defined Floer theory because the sums required to define the differential need not converge. Fukaya observed in [14] that, if the local systems are sufficiently close to being unitary, then in fact the sums do converge.

Unfortunately, Fukaya's indirect approach, which uses the openness of the space of tame almost complex structures, does not seem to be adapted to the problem we face. The solution we adopt uses instead the reverse isoperimetric inequality of Gromov and Solomon [20] (with a simplified proof by DuVal [11]). In its simplest form, it states:

Theorem 1.7. *If $L \subset M$ is an embedded Lagrangian, and J is a tame almost complex structure on M , then there exists a constant C such that, for each holomorphic map $u: D^2 \rightarrow M$ with boundary mapping to L , the length $\ell(\partial u)$ of the boundary curve and the energy $E(u)$ of the map satisfy the inequality*

$$(1.5.2) \quad \ell(\partial u) \leq CE(u).$$

□

To understand the relationship with Floer theory, we let $[\partial u] \in H_1(L; \mathbb{Z})$ denote the homology class of the boundary curve, and note that, if we consider a norm on $H_1(L; \mathbb{Z})$ so that the norm of a homology class is smaller than the minimal length of a representative, then the above reverse isoperimetric inequality implies

$$(1.5.3) \quad |[\partial u]| \leq CE(u).$$

The contribution of any such curve to a Floer theoretic operation with respect to a rank-1 local system classified by $\alpha \in H^1(L, \Lambda \setminus \{0\})$ has valuation given by

$$(1.5.4) \quad E(u) + \langle \text{val } \alpha, [\partial u] \rangle.$$

Assuming that the valuation of α is bounded by $1/2C$, we thus find that the valuation of this contribution is bounded by $E(u)/2$. Gromov compactness then implies that the sum of the contributions of all holomorphic curves converges (in the adic topology), and hence that the Floer homology of L equipped with such a local system is well defined, for the given choice of almost complex structure. Passing from points to affinoid domains, the same argument implies that Floer theory with coefficients in a local system associated to a polytope $P \subset H^1(L; \mathbb{R})$ is well-defined whenever such a polytope lies within the ball of radius $1/2C$ about the origin.

In this paper, we shall be concerned not just with the Floer theory of a single Lagrangian L , but also with the Floer theory of the pair (L, X_q) for a point $q \in Q$, and more generally with operations involving L together with multiple fibres. For any such operation, there is a variant of the reverse isoperimetric inequality with multiple Lagrangian boundary conditions which specifies a maximal diameter of polytopes which give rise to convergent Floer operation. Since there are infinitely many fibres, and infinitely many operations, there may a priori be no constant which is sufficiently large so that all the desired operations converge. In Section 3, we ensure that this problem does not arise by constructing a compact parameter space for all the choices used in the paper, which allows us to establish a reverse isoperimetric inequality in this context using the results of Appendix A.

1.6. The geometric diagonal of rigid analytic spaces. Having introduced the relevant Floer groups, we would like to interpret the picture on the left of Figure 1 as a factorisation

$$(1.6.1) \quad \mathcal{R}_{L'} \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta}) \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \mathcal{L}_L),$$

where $\overline{\Delta}$ is a bimodule over \mathcal{F} , which is equipped with a pair of module maps

$$(1.6.2) \quad \mathcal{R}_{L'} \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta})$$

$$(1.6.3) \quad \overline{\Delta} \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow \mathcal{L}_L.$$

In particular, the morphism space $\text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta})$ in the middle uses the left module structure on $\overline{\Delta}$, and is equipped with a residual right action arising from the right action on $\overline{\Delta}$.

These various maps are related to each other by the following result, which is proved in Section 4.2:

Proposition 1.8. *Given pairs $(L, L') \in \mathcal{A}$, there is a homotopy commutative diagram*

$$(1.6.4) \quad \begin{array}{ccc} \mathcal{R}_{L'} \otimes_{\mathcal{F}} \mathcal{L}_L & \xrightarrow{(1.4.2)} & CF^*(L, L') \\ \downarrow (1.6.2) & & \downarrow (1.3.5) \\ \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta}_{\mathcal{F}}) \otimes_{\mathcal{F}} \mathcal{L}_L & & \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \mathcal{L}_L) \\ & \searrow (1.6.3) & \swarrow (1.6.3) \\ & \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta}_{\mathcal{F}} \otimes_{\mathcal{F}} \mathcal{L}_L) & \end{array}$$

Having arranged for Figure 1 to correspond to a cobordism of moduli spaces of curves, and for the left hand side to have an algebraic interpretation, we are left with the problem of showing that this alternate factorisation of the composition is a quasi-isomorphism. Showing that the middle arrow in this composition is by far the easiest: it is a purely algebraic consequence of Lemma 4.1, which asserts that \mathcal{L}_L is a perfect module over \mathcal{F} .

We encounter the final difficulty when trying to prove that the other two arrows in this composition are isomorphisms: the intuition is that the bimodule $\overline{\Delta}$ corresponds to the geometric diagonal of Y , but the model of sheaves on Y provided by (a subcategory of) modules over \mathcal{F} is not adequate to define $\overline{\Delta}$ as the diagonal bimodule of \mathcal{F} . The key problem is that we have set morphisms in \mathcal{F} from τ to σ to vanish unless τ precedes σ in the ordering of the set indexing the cover, but the corresponding Floer cohomology groups only vanish if the polytopes are disjoint. On the algebraic side, the intuition is that Floer theory recovers a model for an enlargement of the category of coherent sheaves on Y , which admits as objects pushforwards of the structure sheaves of affinoid sub-domains. In algebraic geometry, the analogous problem of enlarging the category of coherent sheaves to admit pushforwards of the structure sheaves of open subsets (in the Zariski topology) is usually handled by considering the category of quasi-coherent sheaves, but the analogue in the analytic setting is poorly understood. In some sense, the paper's main contribution is thus the insight that, despite the fact that such the rigid analytic geometry literature lacks a satisfactory notion of quasi-coherence in general, Floer theory provides a model for morphisms between pushforwards of structure sheaves in the context that we are studying, and it is not difficult to show that this model gives the expected answer for inclusions, as discussed in Appendix C.

While it would be possible to proceed along this route, and construct a category Po in which morphisms are not artificially set to vanish for non-inclusions, we choose instead a shortcut: we define $\overline{\Delta}$ as the bimodule over \mathcal{F} which would correspond to the pullback of the

diagonal of the hypothetical category Po under a natural embedding $\mathcal{F} \rightarrow \text{Po}$. This approach has the advantage of allowing us to study only the local analogue of the category Po , which is the subject of Sections 2.3 and 4.4.1. To understand the basic ideas behind the proof that these local categories allow us to show that the desired maps are quasi-isomorphism, we refer the reader to Section 4.1. The discussion in Section 2.1 below, which is limited to the cohomological aspects of this construction, may also be useful in understanding the main ideas.

1.7. Outline of the paper. In Section 4.1, we reduce the proof of Theorem 1.1 to algebraic properties of the modules \mathcal{L}_L and $\mathcal{R}_{L'}$ and of the bimodule $\bar{\Delta}_{\mathcal{F}}$. The proof that these properties hold accounts for the remainder of Section 4. These algebraic arguments in turn rely on the studying families of pseudo-holomorphic curve equations, constructed compatibly over various abstract moduli spaces, as explained in the preceding Section 3.

Because of the inherent combinatorial complexity of such constructions, we have chosen to start the paper with Section 2, which constructs all the operations at the cohomological level, and in fact proves a weak version of Theorem 1.1 for Lagrangian sections. We have organised this section to proceed from the simplest construction to the most complicated one, highlighting at each step the additional constraints on the choices we make that are required in order to carry out the next step.

Given the content of Section 2 and the three Appendices on reverse isoperimetric inequalities, Tate's acyclicity results, and the computation of Floer cohomology groups with coefficients in completions of the homology of the group ring of the torus, filling in the rest of the paper is just a matter of lifting cohomological constructions to the chain level.

Acknowledgments. The writing of this paper took a long and meandering route, including drafts that explored different strategies for the Floer-theoretic and algebraic parts. Discussions with Brian Conrad, Pierre Deligne, Kenji Fukaya, David Hansen, Johan de Jong, Thomas Kragh, Amnon Neeman, Paul Seidel and Katrin Wehrheim, were crucial at various points to abandon potential strategies as hopeless or too technically difficult, or to confirm that Floer-theoretic ideas could have mirror analogues.

I would like to thank the anonymous referees for extensive and helpful comments; in particular for pointing out that it is inadequate in the statement of Theorem 1.1 to impose the weaker condition that the total space X have vanishing second homotopy group, as was the case in the first version of this paper.

The author was supported by the Simons Foundation through its “Homological Mirror Symmetry” Collaboration grant SIMONS 385571, and by NSF grants DMS-1308179, DMS-1609148, and DMS-1564172. He was also partially supported by the Erik Ellentuck Fellowship and the IAS Fund of Math during the “Homological Mirror Symmetry” program at the Institute for Advanced Study.

2. COHOMOLOGICAL CONSTRUCTIONS

2.1. Statement of the cohomological results. Let $X \rightarrow Q$ be a Lagrangian n -torus fibration and consider a (finite) collection \mathcal{A} of Lagrangians in X , which are graded in the sense of [25] (see Section 2.4.1 below). In order to bypass difficulties involving the foundations of holomorphic curve theory, we shall assume that each $L \in \mathcal{A}$ is tautologically unobstructed in the sense that we are given

- (2.1.1) a tame almost complex structure J_L on X , with respect to which all holomorphic discs with boundary on L are constant.

As discussed in the introduction, we also assume that $\pi_2(Q) = 0$ which implies that, for any compatible almost complex structure, each fibre X_q of this fibration over a point $q \in Q$ can only bound constant holomorphic discs. Moreover, we shall assume for simplicity that all Lagrangians in \mathcal{A} are mutually transverse.

Since Q is the base of a Lagrangian torus fibration, the Arnol'd-Liouville theorem implies the existence of a lattice $T_q^{*,\mathbb{Z}}Q \subset T_q^*Q$ locally spanned by closed 1-forms. We define an *integral affine polytope* $P \subset Q$ to be the embedded image of a polytope in \mathbb{R}^n , defined by inequalities $\langle _, \alpha_i \rangle \geq \lambda_i$ with $\alpha_i \in \mathbb{Z}^n$ and $\lambda_i \in \mathbb{R}$, under a map which pulls back $T_q^{*,\mathbb{Z}}Q$ to the standard lattice \mathbb{Z}^n in \mathbb{R}^n . Note that our conditions allow for $\dim P$ to be strictly less than $\dim Q$. As recalled in Section 2.2.1 below, we may associate to each such polytope an *affinoid ring* Γ^P in the sense of Tate [30]; such a ring has appeared in other contexts in symplectic topology, most notably in Fukaya, Oh, Ohta, and Ono's work [18] on toric mirror symmetry, where it is denoted $\Lambda\langle y, y^{-1} \rangle^P$.

We consider a finite set A indexing a cover of Q by integral affine polytopes, and the associated cover $\{P_\sigma\}_{\sigma \in \Sigma}$ by the partially ordered set Σ whose elements are the subsets of A labelling non-empty intersections, and which we order by reverse inclusion, so that $\tau \leq \sigma$ implies that $P_\sigma \subseteq P_\tau$. We require that

(2.1.2) the maximal length of a totally ordered subset of Σ is $n+1$, and all polytopes P_σ have non-empty interior.

In order to achieve this condition, we appeal to Lemma 2.3 of [6]: one may construct a triangulation of Q into integral affine simplices. Let A denote the set of simplices. By induction on the dimension of simplices, we pick a cover indexed by A , consisting of integral affine polytopes which only intersect for nested simplices, and so that all intersections have non-empty interior (the base case is to choose sufficiently small non-intersecting polygonal neighbourhoods of all vertices, and the next step is to remove these neighbourhood from all edges, and take non-intersecting polygonal neighbourhoods of the resulting intervals). In this way, totally ordered subsets of Σ correspond to sequences of nested simplices in a triangulation, whose length is bounded as above because Q is a manifold of dimension n . We shall later require that the cover $\{P_\sigma\}_{\sigma \in \Sigma}$ be sufficiently fine with respect to increasingly strong conditions (see Condition (2.3.1) where we impose the first constraint), and that we can pick basepoints q_σ in P_σ so that the fibres X_{q_σ} are transverse to our chosen collection \mathcal{A} of Lagrangians (see Section 2.4.1).

For convenience, we introduce the notation Σ_q for the partially ordered subset of Σ corresponding to elements of the cover which contain a given point $q \in Q$. This set has a maximal element, which corresponds to the intersection of all the elements of our cover which contain q .

Remark 2.1. We are working with a restricted class of covers to facilitate the appeal to the results of Appendix B concerning Tate's acyclicity theorem. A more natural setting to consider is that of covers indexed by finite categories, such that the nerve of the subcategory of polytopes containing any point is contractible. The bound by $n+1$ in Condition (2.1.2) can be replaced by any choice of finite constant, as long as it is possible to construct arbitrarily fine covers while maintaining the same bound. The assumption that the polytopes P_σ have non-empty interior is made for convenience, and makes it easier to achieve transversality e.g. in Section 2.7.1.

In this section, which can be read as an extended introduction, we construct the following structures at the cohomological level which we shall later lift to the A_∞ level:

- A category \mathbf{HF} whose objects are elements $\sigma \in \Sigma$, with $\mathbf{HF}(\tau, \sigma)$ vanishing unless $\tau \leq \sigma$, and otherwise given by a Floer cohomology group $HF^*(P_\tau, P_\sigma)$ which is isomorphic to the affinoid ring Γ^{P_σ} .
- For each Lagrangian $L \in \mathcal{A}$ and $\sigma \in \Sigma$, Floer cohomology groups $HF^*(P_\sigma, L)$ and $HF^*(L, P_\sigma)$ which give rise to right and left modules $H\mathcal{R}_L$ and $H\mathcal{L}_L$ over \mathbf{HF} (see Section 2.4).
- For each pair of Lagrangians (L, L') in \mathcal{A} , a map

$$(2.1.3) \quad HF^*(L, L') \rightarrow \text{Hom}_\Lambda(HF^*(L', P), HF^*(L, P))$$

inducing a map into the space of left-module maps:

$$(2.1.4) \quad HF^*(L, L') \rightarrow \text{Hom}_{\mathbf{HF}}(H\mathcal{L}_{L'}, H\mathcal{L}_L).$$

We define as well a map

$$(2.1.5) \quad HF^*(P_\sigma, L') \otimes_\Lambda HF^*(L, P_\sigma) \rightarrow HF^*(L, L')$$

which is compatible with the action of morphisms in \mathbf{HF} , in the sense that we have an induced map

$$(2.1.6) \quad H\mathcal{R}_{L'} \otimes_{\mathbf{HF}} H\mathcal{L}_L \rightarrow HF^*(L, L').$$

At the cohomological level, the structures listed above are the same as those considered in [6]. Our cochain-level construction however will be different as we shall need to consider the following additional structures:

- A *geometric diagonal* bimodule $H\overline{\Delta}_{\mathbf{HF}}$ over \mathbf{HF} , given for a pair (τ, σ) by a Floer cohomology group $HF^*(P_\tau, P_\sigma)$. The key point here is that these groups may be non-vanishing even if the condition $\tau \leq \sigma$ does not hold.
- For each $L \in \mathcal{A}$, and for each pair (τ, σ) of polytopes, a map

$$(2.1.7) \quad HF^*(L, P_\sigma) \otimes_\Lambda HF^*(P_\tau, L) \rightarrow HF^*(P_\tau, P_\sigma),$$

which descends to a map of bimodules:

$$(2.1.8) \quad H\mathcal{L}_L \otimes_\Lambda H\mathcal{R}_L \rightarrow H\overline{\Delta}_{\mathbf{HF}}.$$

- Finally, we construct a map

$$(2.1.9) \quad HF^*(P_\tau, P_\sigma) \otimes_\Lambda HF^*(L, P_\tau) \rightarrow HF^*(L, P_\sigma),$$

which gives rise to a map of left modules

$$(2.1.10) \quad \overline{\Delta}_{\mathbf{HF}} \otimes_{\mathbf{HF}} H\mathcal{L}_L \rightarrow H\mathcal{L}_L.$$

For the statement of the main result of this section, it will be convenient to interpret Equation (2.1.8) as a map of right modules:

$$(2.1.11) \quad H\mathcal{R}_L \rightarrow \text{Hom}_{\mathbf{HF}}(H\mathcal{L}_L, \overline{\Delta}_{\mathbf{HF}}).$$

Proposition 2.2. *Given Lagrangians $(L, L') \in \mathcal{A}$, the following diagram commutes:*

$$(2.1.12) \quad \begin{array}{ccc} H\mathcal{R}_{L'} \otimes_{\mathbf{HF}} H\mathcal{L}_L & \xrightarrow{(2.1.6)} & HF^*(L, L') \\ \downarrow (2.1.11) & & \downarrow (2.1.4) \\ \text{Hom}_{\mathbf{HF}}(H\mathcal{L}_{L'}, H\overline{\Delta}_{\mathbf{HF}}) \otimes_{\mathbf{HF}} H\mathcal{L}_L & & \text{Hom}_{\mathbf{HF}}(H\mathcal{L}_{L'}, H\mathcal{L}_L) \\ & \searrow & \nearrow (2.1.10) \\ & \text{Hom}_{\mathbf{HF}}(H\mathcal{L}_{L'}, H\overline{\Delta}_{\mathbf{HF}} \otimes_{\mathbf{HF}} H\mathcal{L}_L). & \end{array}$$

This result is proved in Section 2.8.2, and is illustrated in Figure 1. One key point is that the statement of the above proposition is at the level of cohomological categories. This means in particular that the tensor product in the top left of Diagram (2.1.12) is computed as a quotient

$$\bigoplus_{\sigma, \tau} HF^*(P_\sigma, L') \otimes_{\Lambda} HF(\tau, \sigma) \otimes_{\Lambda} HF^*(L, P_\tau) \rightarrow \bigoplus_{\tau} HF^*(P_\tau, L') \otimes_{\Lambda} HF^*(L, P_\tau) \rightarrow H\mathcal{R}_{L'} \otimes_{\mathcal{H}\mathcal{F}} H\mathcal{L}_L,$$

while the morphism space in the middle right is computed as a kernel

$$\text{Hom}_{\mathcal{H}\mathcal{F}}(H\mathcal{L}_{L'}, H\mathcal{L}_L) \hookrightarrow \prod_{\sigma} \text{Hom}_{\Lambda}(HF^*(L', P_\sigma), HF^*(L, P_\sigma)) \rightarrow \prod_{\sigma, \tau} \text{Hom}_{\Lambda}(HF^*(L', P_\sigma), HF^*(L, P_\sigma))$$

With the above in mind, once all structure maps are constructed, the commutativity of Diagram (2.1.12) follows from proving that the diagram

$$\begin{array}{ccc} HF^*(P_\tau, L') \otimes_{\Lambda} HF^*(L, P_\tau) & \xrightarrow{\hspace{10cm}} & HF^*(L, L') \\ \downarrow & & \downarrow \\ \text{Hom}_{\Lambda}(HF^*(L', P_\sigma), HF^*(P_\tau, P_\sigma)) \otimes_{\Lambda} HF^*(L, P_\tau) & \xrightarrow{\hspace{10cm}} & \text{Hom}_{\Lambda}(HF^*(L', P_\sigma), HF^*(L, P_\sigma)) \\ & \searrow & \swarrow \\ & \text{Hom}_{\Lambda}(HF^*(L', P_\sigma), HF^*(P_\tau, P_\sigma) \otimes_{\Lambda} HF^*(L, P_\tau)) & \end{array} \tag{2.1.13}$$

commutes for each pair (τ, σ) in Σ . When we revisit this result in Section 4, such an approach will not be sufficient, as the cochain-level version of the above diagram will only commute up to homotopy, and the analogue of Diagram (2.1.12) will involve a tower of (homotopy) commutative diagrams.

We now state the main assumption under which we can prove that Floer homology is isomorphic to morphisms of modules at the level of cohomological categories; even though this is in some sense a special case of Theorem 1.1, we shall discuss the proofs, because it serves as a model for the A_∞ analogue which holds in general:

Proposition 2.3. *If L and L' are Lagrangian sections, Equation (2.1.4) is surjective.*

Given Proposition 2.2, it is clear that the above result would follow from the statement that each of the three arrows which compose counterclockwise in Diagram (2.1.12) is an isomorphism. Starting with the middle arrow, we have the following result, which is proved in Section 2.8 below:

Lemma 2.4. *If L and L' are Lagrangian sections, the map*

$$(2.1.14) \quad \text{Hom}_{\mathcal{H}\mathcal{F}}(H\mathcal{L}_{L'}, H\bar{\Delta}_{\mathcal{H}\mathcal{F}}) \otimes_{\mathcal{H}\mathcal{F}} H\mathcal{L}_L \rightarrow \text{Hom}_{\mathcal{H}\mathcal{F}}(H\mathcal{L}_{L'}, H\bar{\Delta}_{\mathcal{H}\mathcal{F}} \otimes_{\mathcal{H}\mathcal{F}} H\mathcal{L}_L),$$

is an isomorphism.

Remark 2.5. The fact that Equation (2.1.14) is an isomorphism amounts to the statement that the module $H\mathcal{L}_L$ behaves with respect to the functors $\text{Hom}_{\mathcal{H}\mathcal{F}}$ and $\otimes_{\mathcal{H}\mathcal{F}}$ in much the same way as a projective module over a ring. It is therefore not surprising that this property holds whenever L is a Lagrangian section because a section gives rise, under mirror symmetry, to a line bundle.

Regarding the remaining arrows in the counterclockwise composition arising in Diagram (2.1.12), we use the fact that the Floer cohomology groups $HF^*(P_\tau, P_\sigma)$ vanish whenever

the inputs are disjoint. We can formalise the corresponding property of the bimodule $\overline{\Delta}_{H\mathcal{F}}$ as follows:

Definition 2.6. *A bimodule \mathcal{P} over $H\mathcal{F}$ is local if $\mathcal{P}(\tau, \sigma)$ vanishes whenever P_τ and P_σ are disjoint.*

For computations involving local bimodules, it is convenient to consider a local category associated to each element of Σ . For concreteness, we introduce some terminology:

Definition 2.7. *The star of any element $\sigma \in \Sigma$ is the partially ordered set $\Sigma_\sigma \subseteq \Sigma$ consisting of elements $\tau \in \Sigma$ whose associated polytope intersects P_σ . We denote the full subcategory of $H\mathcal{F}$, whose set of objects is Σ_σ , by $H\mathcal{F}_\sigma$.*

We are interested in computing morphisms and tensor products locally, i.e. using the restriction to $H\mathcal{F}_\sigma$. More specifically, given any left module \mathcal{L} and bimodule \mathcal{P} , we have natural comparison maps

$$(2.1.15) \quad \text{Hom}_{H\mathcal{F}}(\mathcal{L}, \mathcal{P}) \rightarrow \text{Hom}_{H\mathcal{F}_\sigma}(\mathcal{L}, \mathcal{P})$$

$$(2.1.16) \quad \mathcal{P} \otimes_{H\mathcal{F}_\sigma} \mathcal{L} \rightarrow \mathcal{P} \otimes_{H\mathcal{F}} \mathcal{L}$$

or right and left modules over $H\mathcal{F}_\sigma$, where we abuse notation by using the same symbol for a module on \mathcal{F} and its restriction to \mathcal{F}_σ . The values of these modules on general objects of $H\mathcal{F}_\sigma$ are difficult to control, but the following result asserts that the value on the object corresponding to σ is the same:

Lemma 2.8. *For each left module \mathcal{L} over $H\mathcal{F}$, and each local bimodule \mathcal{P} , the natural maps*

$$(2.1.17) \quad \text{Hom}_{H\mathcal{F}}(\mathcal{L}, \mathcal{P})(\sigma) \rightarrow \text{Hom}_{H\mathcal{F}_\sigma}(\mathcal{L}, \mathcal{P})(\sigma)$$

$$(2.1.18) \quad (\mathcal{P}_{H\mathcal{F}} \otimes_{H\mathcal{F}_\sigma} \mathcal{L})(\sigma) \rightarrow (\mathcal{P} \otimes_{H\mathcal{F}} \mathcal{L})(\sigma)$$

are isomorphisms.

Proof. The two arguments are entirely analogous; we explain the second. We can describe the value at σ of the tensor product over $H\mathcal{F}$ as the cokernel of the map

$$(2.1.19) \quad \bigoplus_{\tau_{-1}, \tau_{-2}} \mathcal{P}(\tau_{-1}, \sigma) \otimes_{\Lambda} H\mathcal{F}(\tau_{-2}, \tau_{-1}) \otimes_{\Lambda} \mathcal{L}(\tau_{-2}) \rightarrow \bigoplus_{\tau} \mathcal{P}(\tau, \sigma) \otimes_{\Lambda} \mathcal{L}(\tau),$$

with the direct sums being taken over objects and pairs of objects in $H\mathcal{F}$, and the arrow being given by the difference between the two possible compositions. The key fact is that, whenever $H\mathcal{F}(\tau_{-2}, \tau_{-1}) \neq 0$, the polytope $P_{\tau_{-1}}$ is contained in $P_{\tau_{-2}}$, so that τ_{-1} being an object of $H\mathcal{F}_\sigma$ implies the same for τ_{-2} . In particular, both direct sums are in fact taken over objects of $H\mathcal{F}_\sigma$, which implies the desired isomorphism. \square

Because of the above result, whose A_∞ generalisation holds with the same proof, verifying that Equations (2.1.11) and (2.1.10) are isomorphisms is a *local* computation, in the sense that the corresponding computations involve the each of the categories $H\mathcal{F}_\sigma$ separately.

To perform these local computations, it is useful to introduce a category $H\text{Po}_\sigma$, whose objects are the polytopes of Q which are contained in a polytopal neighbourhood of P_σ that itself contains all polytopes P_τ belonging to the star of σ . The morphisms in $H\text{Po}_\sigma$ are again given by Floer cohomology groups, and we have a faithful embedding

$$(2.1.20) \quad j: H\mathcal{F}_\sigma \rightarrow H\text{Po}_\sigma,$$

such that for each pair of elements τ and ρ of Σ_σ for which $\rho \leq \tau$, the induced map

$$(2.1.21) \quad H\mathcal{F}_\sigma(\rho, \tau) \rightarrow H\text{Po}_\sigma(P_\rho, P_\tau),$$

is an isomorphism. We prove in Section 2.7 that the pullback of the diagonal bimodule of HPo_σ is naturally isomorphic to the restriction of $\bar{\Delta}_{\text{HF}}$ to HF_σ (see Lemma 2.51). Furthermore we show in Section 2.8 that, for each $L \in \mathcal{A}$ there are left and right modules $H\mathcal{L}_{L,\sigma}$ and $H\mathcal{R}_{L,\sigma}$ over HPo_σ whose pullbacks to HF_σ are naturally isomorphic to the restrictions of $H\mathcal{L}_L$ and $H\mathcal{R}_L$ (see Proposition 2.55).

This allows us to reduce local computations to the category HPo_σ in certain special cases, and to the corresponding A_∞ category in general. For example, the proof that Equation (2.1.10) is an isomorphism reduces to the statement that the map of the tensor product of pullbacks of left and right modules

$$(2.1.22) \quad j^* \Delta_{\text{HPo}_\sigma}(\sigma, -) \otimes_{\text{HF}_\sigma} j^* (H\mathcal{L}_L) \rightarrow j^* H\mathcal{L}_L(\sigma)$$

is an isomorphism. We shall prove that the above follows from Tate acyclicity assuming that L meets X_{q_σ} at a single point (see Section 2.4.4).

Remark 2.9. The sign conventions on the cohomological constructions implemented in this section are those corresponding to A_∞ algebras with trivial differential and higher products. One can translate to the usual *dg* conventions as discussed in [26, Section I.1.1a].

We complete the introduction to this section by reducing its main conclusion to the results proved in the remainder of the section:

Proof of Proposition 2.3. As noted earlier, Proposition 2.2 is proved in Section 2.8.2, and reduces the desired result to proving that the unlabeled arrow in Diagram (2.1.12), as well as Equations (2.1.10) and (2.1.11) are isomorphisms. The first isomorphism is established by Lemma 2.4 above, and the remaining two isomorphisms are reduced to their local versions by Lemma 2.8 above, and Proposition 2.55 from Section 2.8 below. The corresponding local statements are proved in Proposition 2.36 and Lemma 2.41. \square

2.2. Morse theory of local systems associated to polytopes.

2.2.1. Affinoid rings associated to polytopes. The first step to construct the ring associated to an integral affine polytope is to observe that each point $p \in \mathbb{R}^n$ determines a non-archimedean valuation on the ring of Laurent polynomials with coefficients in the Novikov field Λ , which we think of as the group ring of \mathbb{Z}^n , given by

$$(2.2.1) \quad \Lambda[\mathbb{Z}^n] \rightarrow \mathbb{R} \cup \{+\infty\}$$

$$(2.2.2) \quad \text{val}_p \left(\sum c_\beta z^\beta \right) := \min_\beta (\text{val}(c_\beta) + \langle \beta, p - q \rangle).$$

Here, $\text{val}(c_\beta)$ denotes the valuation of this element of Λ as defined in Equation (1.2.2). Taking the minimum over all elements of an integral affine polytope P in \mathbb{R}^n , we obtain the valuation

$$(2.2.3) \quad \text{val}_P \left(\sum c_\beta z^\beta \right) := \min_{p \in P} \text{val}_p \left(\sum c_\beta z^\beta \right),$$

and define the ring Γ^P to be the corresponding completion of the Laurent polynomial ring (i.e. consider series such that the number of terms with valuation bounded above by any fixed constant is finite). Elements of this ring can be written uniquely as series

$$(2.2.4) \quad \sum_{\beta \in \mathbb{Z}^n} c_\beta z^\beta$$

satisfying the condition

$$(2.2.5) \quad \lim_{|\beta| \rightarrow +\infty} \text{val}_P(c_\beta z^\beta) = +\infty.$$

If P is an integral affine polytope in Q , we can relate this affinoid ring to the fundamental group ring of a fibre X_q as follows: the integral affine structure on Q defines an *affine exponential map*

$$(2.2.6) \quad H^1(X_q, \mathbb{R}) \cong T_q Q \dashrightarrow Q$$

given on a neighbourhood of the origin, which induces an isomorphism of lattices at each point where it is defined. A path from q to P determines a lift of P to $H^1(X_q, \mathbb{R})$. Each point $p \in P$ thus determines a valuation on the group ring $\Lambda[H_1(X_q; \mathbb{Z})]$ of the first homology of the fibre X_q , and we obtain a valuation associated to the polytope P by taking the minimum as in Equation (2.2.3). Given that the integral affine structure on Q and the embedding of P in \mathbb{R}^n give an isomorphism $\mathbb{Z}^n \cong H_1(X_q; \mathbb{Z})$, we conclude:

Lemma 2.10. *A choice of homotopy class of paths from q to P determines an isomorphism between Γ^P and the completion of $\Lambda[H_1(X_q; \mathbb{Z})]$ at the valuation associated to the image of P in $H^1(X_q, \mathbb{R}) \cong T_q Q$.* \square

2.2.2. Loops, paths, and local systems. To proceed further, it is convenient to eliminate the choice of homotopy class by fixing a contractible subset $\mathcal{N} \subseteq Q$, which will contain both the basepoint q and the polytope P . We equip \mathcal{N} with a Lagrangian section

$$(2.2.7) \quad \iota: \mathcal{N} \rightarrow X_{\mathcal{N}}$$

of the restriction of the projection $X \rightarrow Q$ to the inverse image $X_{\mathcal{N}}$ of \mathcal{N} . Given this data, we shall associate to each integral affine polytope $P \subseteq \mathcal{N}$, and each point $q \in \mathcal{N}$, a local system U_ι^P on the Lagrangian X_q with fibre isomorphic to Γ^P .

To start, let U_ι denote the local system on X_q whose value at a point x is the free Λ -module

$$(2.2.8) \quad U_{\iota,x} := \Lambda[\pi_0(\Omega_{\iota(q),x} X_q)],$$

freely generated by the components of the space $\Omega_{\iota(q),x} X_q$ of paths in X_q from $\iota(q)$ to x . This space is naturally a module over the loop space of X_q based at $\iota_{\mathcal{N}}(q)$, so that $U_{\iota,x}$ is a module over the group ring of the fundamental group of X_q . Since X_q is a torus, the fundamental group is isomorphic to $H_1(X_q; \mathbb{Z})$, so we conclude that $U_{\iota,x}$ is a module over the group ring of $H_1(X_q; \mathbb{Z})$.

On the other hand, Lemma 2.10 implies that the affinoid ring Γ^P is naturally a module over the group ring of $H_1(X_q; \mathbb{Z})$. The tensor product of U_ι over $\Lambda[H_1(X_q; \mathbb{Z})]$ with Γ^P defines a local system U_ι^P , with fibre

$$(2.2.9) \quad U_{\iota,x}^P := \Lambda[\pi_0(\Omega_{\iota(q),x} X_q)] \otimes_{\Lambda[H_1(X_q; \mathbb{Z})]} \Gamma^P.$$

Note that $U_{\iota,x}$ is a free rank-1 module over $\Lambda[H_1(X_q; \mathbb{Z})]$, with a generator corresponding to a choice of homotopy class of paths from $\iota(q)$ to x . Choosing such a path, we obtain a norm on $U_{\iota,x}$, and $U_{\iota,x}^P$ is the completion with respect to this norm.

Next, given a polytope in Q , we would like to compare the associated local systems on different Lagrangian fibres. For applications, we shall more precisely need to compare the values at points which are defined as the image of a Lagrangian section; since the Lagrangian sections we encounter are seldom globally defined, we introduce a contractible subset B of \mathcal{N} , which will be the domain of definition of a Lagrangian section

$$(2.2.10) \quad x: B \rightarrow X_B.$$

Given $q \in B$, we write x^q for this intersection of this section with the fibre X_q .

Lemma 2.11. *If q and q' lie in B , there is an identification*

$$(2.2.11) \quad U_{\iota, x^q}^P \cong U_{\iota, x^{q'}}^P$$

for all points $q, q' \in B$, with the property that, for a triple q, q' and q'' , the following diagram commutes:

$$(2.2.12) \quad \begin{array}{ccc} U_{\iota, x^q}^P & \longrightarrow & U_{\iota, x^{q''}}^P \\ \downarrow & \nearrow & \\ U_{\iota, x^{q'}}^P & & \end{array}$$

Proof. Pick a class γ in $\pi_0(\Omega_{\iota(q), x} X_q) \cong \pi_0(\Omega_{\iota(q'), x} X_{q'})$. Since B is contractible, there is a unique homotopy class of maps

$$(2.2.13) \quad u: [0, 1]^2 \rightarrow X_B$$

such that the restrictions to $\{0\} \times [0, 1]$ and $\{1\} \times [0, 1]$ map to X_q and $X_{q'}$ and represent the class γ , and the restrictions to $[0, 1] \times \{0\}$ and $[0, 1] \times \{1\}$ map to the sections ι and x . We define

$$(2.2.14) \quad F_\gamma(q, q') = \int_{[0, 1]^2} u^* \omega,$$

and note that this integral is independent of the choice of representative of the given homotopy class because all boundary conditions are Lagrangian. We then define the map from U_{ι, x^q}^P to $U_{\iota, x^{q'}}^P$ by

$$(2.2.15) \quad z^\gamma \mapsto T^{F_\gamma(q, q')} z^\gamma.$$

□

As the proof shows, the identification in Lemma 2.11 is compatible with parallel transport maps in the following sense: we can associate to a homotopy class γ of paths with endpoints (x_0, x_1) on X_q , and to $q' \in B$ a real number $F_\gamma(q, q')$ which is the *flux* of this path. Under the identification of Equation (2.2.11), the parallel transport maps associated to this homotopy class of paths in X_q and the corresponding homotopy class in $X_{q'}$ differ by multiplication by $T^{F_\gamma(q, q')}$. In the setting of the above Lemma, we shall therefore omit the superscript from x^q unless it is required for clarity of exposition.

A key aspect of the constructions of this paper will be the need to verify the T -adic convergence of operators constructed using holomorphic curves. Consider points $x, y \in X_q$, a path γ with endpoints x and y , and the induced parallel transport map

$$(2.2.16) \quad z^{[\gamma]}: U_{\iota, x}^P \rightarrow U_{\iota, y}^P$$

for a polytope $P \subseteq Q$ which is in the image of the affine exponential map based at q . Fixing paths from x and y to the intersection of X_q with the section associated to \mathcal{N} , we obtain valuations on the fibres of these local systems. These data also assign a homology class $[\gamma] \in H_1(X, \mathbb{Z})$ to γ , and the valuation of $z^{[\gamma]}$ is then bounded by the product of the norm of $[\gamma]$ with the distance from the origin to the inverse image of P in $T_q Q$, where

$$(2.2.17) \quad |[\gamma]| := \min_{[\gamma']=[\gamma]} \ell(\gamma'),$$

and ℓ is the length of the loop γ' , with respect to the flat metric on X_q induced by a choice of Riemannian metric on Q .

Lemma 2.12. *If the lift of P to $T_q Q$ is contained in the ball of radius ϵ , the valuation of the parallel transport map $z^{[\gamma]}$ on U_ι^P is bounded below by $-\epsilon \cdot |[\gamma]|$, up to adding a constant which is independent of γ . \square*

It is convenient to replace the condition about the lift of P under the exponential map by a condition in Q . To this end, we shall assume from now on that:

(2.2.18) the convexity radius of the chosen metric on Q is greater than 2. Moreover, for all $q \in Q$, the affine exponential map restricts to an embedding on the ball of radius 2 in $T_q Q$ with distortion bounded by 2.

We recall that the convexity radius of a Riemannian metric is the smallest constant r for which geodesic balls of radius smaller than r , centered at any point in Q , have the property that any two points are connected by a unique minimal geodesic which is contained in the given ball. This constant is non-zero for any closed Riemannian manifold, as discussed e.g. in [8, Section 6.5.3]. The key property we shall use is that the intersection of any balls of radius smaller than the convexity radius is again convex, and hence contractible.

The following result will play a key role in the proof of convergence of various Floer theoretic constructions:

Corollary 2.13. *Let $0 < \delta$ be a positive real number, and let ϵ be a constant which is smaller than the minimum of $1/2$ and $\delta/2$. Assume that $\{\gamma_i\}_{i=0}^\infty$ is a sequence of paths with endpoints x and y in X_q and $\lambda_i \in \mathbb{R}$ is a sequence going to $+\infty$ such that*

$$(2.2.19) \quad \delta \cdot |[\gamma_i]| \leq \lambda_i + \text{a constant independent of } i.$$

Whenever P is contained in the ball of radius ϵ about q in Q , the map

$$(2.2.20) \quad \sum_{i=1}^{\infty} T^{\lambda_i} z^{[\gamma_i]}: U_{\iota,x}^P \rightarrow U_{\iota,y}^P$$

converges in the topology induced by P .

Proof. The valuation of $T^{\lambda_i} z^{[\gamma_i]}$ is given by

$$(2.2.21) \quad \lambda_i + \text{val}_P z^{[\gamma_i]} \geq \lambda_i - 2\epsilon \cdot |\gamma_i| \geq (1 - 2\epsilon/\delta)\lambda_i + \text{a constant independent of } i.$$

The assumptions that $2\epsilon < \delta$ and $\lambda_i \rightarrow \infty$ imply that this valuation also goes to $+\infty$. \square

Remark 2.14. The reader should have in mind that the constant δ which appears in the statement is obtained from applying the reverse isoperimetric inequality to a collection of holomorphic curves with boundaries γ_i , and energies λ_i .

2.2.3. Morse homology groups associated to polytopes. Given a pair $(\mathcal{N}_0, \mathcal{N}_1)$ of contractible subsets of Q , with contractible intersection, and equipped with Lagrangian sections ι_0 and ι_1 , consider a point $q \in \mathcal{N}_0 \cap \mathcal{N}_1$ together with a Morse function

$$(2.2.22) \quad f_{0,1}: X_q \rightarrow \mathbb{R},$$

whose critical locus we denote $\text{Crit}(0, 1)$.

In order to later be able to study Floer theory over a field of characteristic different from 2, we pick a Pin structure on the restriction of TQ to \mathcal{N}_i for each $\sigma \in \Sigma$; as \mathcal{N}_i is contractible, such a choice is unique up to isomorphism. Via the canonical identification, for $q \in Q$, of the tangent space of X_q with the trivial vector bundle with fibre $T_q^* Q$, we obtain a

Pin structure on X_q , varying continuously over $q \in \mathcal{N}_i$. We then denote by $\lambda_{0,1}$ the rank-1 free abelian group of isomorphism classes of families of Pin structures on $T_q Q$, parametrised by $[0, 1]$ and twisted by the orientation line of $T_q Q$ as in [26, Equation (11.32)], which agree with the chosen structures associated to \mathcal{N}_0 and \mathcal{N}_1 at the endpoints. Given a critical point $x \in \text{Crit}(0, 1)$, define

$$(2.2.23) \quad \delta_x := \lambda_{0,1} \otimes \det_x,$$

where \det_x is the orientation line of the stable manifold of x . Let $\deg(x)$ denote the degree of this graded line.

The choices ι_i of Lagrangian sections over \mathcal{N}_i yield local systems U_i on X_q . Given polytopes $P_i \subset T_q^* Q$, we obtain local systems $U_i^{P_i}$ which are constructed by completion, and define

$$(2.2.24) \quad CM^*(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda) := \bigoplus_{x \in \text{Crit}(0, 1)} \text{Hom}_\Lambda^c(U_{0,x}^{P_0}, U_{1,x}^{P_1}) \otimes \delta_x$$

where $\text{Hom}_\Lambda^c(U_{0,x}^{P_0}, U_{1,x}^{P_1})$ is the space of continuous homomorphisms, i.e. those with finite valuation

$$(2.2.25) \quad \text{val } \phi = \inf_{f \in U_{0,x}^{P_0} \setminus \{0\}} \text{val}(\phi(f)) - \text{val}(f).$$

We now recall the construction of the differential: let \mathcal{I} denote the interval $(-\infty, +\infty)$ for which we use the parameter t . We define

$$(2.2.26) \quad \mathcal{T}(0, 1) = \left\{ \gamma: \mathcal{I} \rightarrow X_q \mid \frac{d\gamma}{dt} = \nabla f_{0,1} \right\} / \mathbb{R}$$

where the gradient is taken with respect to a Riemannian metric on X_q , and the \mathbb{R} action is by translation. We have a standard compactification $\mathcal{T}(0, 1) \subset \overline{\mathcal{T}}(0, 1)$ consisting of broken gradient trajectories, and an evaluation map

$$(2.2.27) \quad \overline{\mathcal{T}}(0, 1) \rightarrow \text{Crit}(0, 1) \times \text{Crit}(0, 1),$$

which on $\mathcal{T}(0, 1)$ is given by taking the limits at $-\infty$ and $+\infty$. We write $\overline{\mathcal{T}}(x_0, x_1)$ for the fibre over critical points x_0 and x_1 .

For a generic choice of metric, the moduli space $\overline{\mathcal{T}}(x_0, x_1)$ for $x_0 \neq x_1$ is a compact manifold with boundary of dimension

$$(2.2.28) \quad \dim \overline{\mathcal{T}}(x_0, x_1) = \deg(x_0) - \deg(x_1) - 1,$$

whose direct sum with the tangent space of the interval is oriented relative the tensor product $\det_{x_0}^\vee \otimes \det_{x_1}$ in the sense that we have a canonical isomorphism

$$(2.2.29) \quad \text{o}_{\overline{\mathcal{T}}(x_0, x_1)} \otimes \text{o}_{\mathcal{I}} \otimes \det_{x_0}^\vee \otimes \det_{x_1} \cong \mathbb{Z},$$

where we use o_M to denote the orientation line of a manifold M (i.e. the top exterior power of its tangent space). Fixing the standard trivialisation of the tangent space of the interval, we obtain, whenever

$$(2.2.30) \quad \deg(x_0) = \deg(x_1) + 1$$

a natural isomorphism of degree 1

$$(2.2.31) \quad \det_\gamma: \det_{x_1} \rightarrow \det_{x_0}.$$

associated to each element of $\overline{\mathcal{T}}(x_0, x_1)$. Parallel transport also induces an isomorphism of topological vector spaces

$$(2.2.32) \quad \Pi_\gamma: \text{Hom}_\Lambda^c(U_{0,x_1}^{P_0}, U_{1,x_1}^{P_1}) \rightarrow \text{Hom}_\Lambda^c(U_{0,x_0}^{P_0}, U_{1,x_0}^{P_1})$$

$$(2.2.33) \quad \psi \mapsto z^{[\gamma]} \circ \psi \circ z^{-[\gamma]}.$$

Taking the tensor product of these two maps with the identity on $\lambda_{0,1}$, we define

$$(2.2.34) \quad \mu^1: CM^*\left(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda\right) \rightarrow CM^*\left(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda\right)$$

$$(2.2.35) \quad \mu^1 := \sum_{[\gamma] \in \overline{\mathcal{T}}^0(0,1)} (-1)^{\deg(x_1)} \Pi_\gamma \otimes \text{id}_\lambda \otimes \det_\gamma,$$

where $\overline{\mathcal{T}}^0(0,1)$ is the space of rigid gradient flow lines, i.e. those of virtual dimension 0, which in this case corresponds to the union of components $\overline{\mathcal{T}}(x_0, x_1)$ for critical points x_0 and x_1 satisfying Equation (2.2.30). Since the sum is necessarily finite, this map is continuous, and defines a differential by the usual analysis of the boundary of the moduli space of flow lines of dimension 1, and the fact that the parallel transport maps are invariant under homotopy, and compatible with concatenation of paths.

Choosing a basepoint $q_{0,1} \in \mathcal{N}_0 \cap \mathcal{N}_1$, we define

$$(2.2.36) \quad CF^*((\mathcal{N}_0, P_0), (\mathcal{N}_1, P_1)) := CM^*\left(X_{q_{0,1}}, \text{Hom}_\Lambda^c(U^{P_0}, U^{P_1}) \otimes \lambda\right).$$

We now prove that, if we choose a different basepoint, we can arrange Morse data so that the resulting complex is isomorphic. Recall that the choice of section ι_0 induces an identification of symplectic manifolds

$$(2.2.37) \quad X_{\mathcal{N}_0} \cong T^* \mathcal{N}_0 / T^{*,\mathbb{Z}} \mathcal{N}_0$$

over the base \mathcal{N}_0 , hence diffeomorphisms $X_q \cong X_{q'}$ for all pairs $(q, q') \in \mathcal{N}_0$, which are compatible for triples. We shall leave these diffeomorphisms implicit in our notation. With this in mind, the function $f_{0,1}$ appearing in the right hand side of Equation (2.2.24) can therefore be thought of as a Morse function on *any fibre* over $\mathcal{N}_0 \cap \mathcal{N}_1$, and the metric on X_q induces a metric on $X_{q'}$. Since the base is contractible, Lemma 2.11 provides an isomorphism of local systems, so we obtain an isomorphism of cochain complexes

$$(2.2.38) \quad CM^*(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda) \rightarrow CM^*(X_{q'}, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda)$$

which is compatible with composition for triples (q, q', q'') . This establishes that this Morse complex is indeed independent of the choice of basepoint in \mathcal{N}_0 .

Remark 2.15. Note that we broke symmetry and chose ι_0 to identify fibres over $\mathcal{N}_0 \cap \mathcal{N}_1$. We could have chosen ι_1 , or in fact any other trivialisation. In cochain-level constructions, we shall need to verify compatibility between various trivialisations, by ensuring that they arise from a contractible set of choices.

2.2.4. Morse theoretic product. Consider a triple $\{\mathcal{N}_i\}_{i=0}^2$ of contractible subsets of Q , equipped with sections $\{\iota_i\}_{i=0}^2$ of the Lagrangian torus fibration. We assume that the triple intersection, which we denote \mathcal{N} , is contractible. We use the above discussion to identify the functions $f_{i,j}$ as Morse functions on a single fibre over a point $q \in \mathcal{N}$. Morse theory thus induces a product on the Floer cohomology groups for pairs, whose construction, while standard, we now recall in order to set up notation for future use.

Consider the semi-infinite intervals

$$(2.2.39) \quad \mathcal{I}_+ := [0, \infty) \text{ and } \mathcal{I}_- := (-\infty, 0].$$

We define the space of Morse data

$$(2.2.40) \quad \mathcal{V}_\pm(i, j) \subset C^\infty(\mathcal{I}_\pm, C^\infty(X_q, TX_q))$$

to consist of families of vector fields on X_q , parametrised by \mathcal{I}_\pm , which agree with $\nabla f_{i,j}$ outside a compact set. The gradient flow is taken with respect to the metric chosen in the construction of the Floer complex for the pair (i, j) .

Given $\xi_{ij}^\pm \in \mathcal{V}_\pm(i, j)$, we then define

$$(2.2.41) \quad \mathcal{T}_\pm(i, j) \subset C^\infty(\mathcal{I}_\pm, X_q)$$

to be the set of perturbed gradient flow lines, i.e. maps γ from \mathcal{I}_\pm to X_q satisfying

$$(2.2.42) \quad \frac{d\gamma}{dt} = \xi_{ij}^\pm.$$

The limit of γ at $\pm\infty$ yields a natural evaluation map

$$(2.2.43) \quad \mathcal{T}_\pm(i, j) \rightarrow \text{Crit}(i, j).$$

We define the spaces of *broken semi-infinite perturbed gradient flow lines*, to be

$$(2.2.44) \quad \overline{\mathcal{T}}_\pm(i, j) := \mathcal{T}_\pm(i, j) \cup \mathcal{T}_\pm(i, j) \times_{\text{Crit}(i, j)} \overline{\mathcal{T}}(i, j).$$

where we use the evaluation map at $\mp\infty$ for $\overline{\mathcal{T}}(i, j)$. An element of this fibre product can be thought of as a broken flow line together with a semi-infinite flow line, with matching asymptotic limits.

There is a natural evaluation map

$$(2.2.45) \quad \overline{\mathcal{T}}_\pm(i, j) \rightarrow X_q \times \text{Crit}(i, j),$$

given by evaluation at 0 and of the limit of γ at $\pm\infty$. We denote the coequaliser of the three evaluation maps from the product to X_q by

$$(2.2.46) \quad \overline{\mathcal{T}}(0, 1, 2) := \text{Coeq}(\overline{\mathcal{T}}_-(0, 2) \times \overline{\mathcal{T}}_+(1, 2) \times \overline{\mathcal{T}}_+(0, 1) \rightarrow X_q),$$

This is the space of flow lines which map 0 to the same point.

Given critical points $x_i \in \text{Crit}(i, i+1)$ (with index counted modulo 3), we define

$$(2.2.47) \quad \overline{\mathcal{T}}(x_0, x_2, x_1)$$

to be the inverse image of (x_1, x_2, x_0) under the evaluation map

$$(2.2.48) \quad \overline{\mathcal{T}}(0, 1, 2) \rightarrow \prod_{i=0}^2 \text{Crit}_{i, i+1}.$$

For a generic choice of triples of Morse perturbations $(\xi_{01}^+, \xi_{12}^+, \xi_{02}^-)$, the space $\overline{\mathcal{T}}(x_0, x_2, x_1)$ is a compact topological manifold with boundary, naturally oriented relative

$$(2.2.49) \quad \det_{x_0}^\vee \otimes \det_{x_2} \otimes \det_{x_1}.$$

Dualising, we obtain, for each rigid element $\gamma \in \overline{\mathcal{T}}(x_0, x_2, x_1)$, a natural map

$$(2.2.50) \quad \det_\gamma: \det_{x_2} \otimes \det_{x_1} \rightarrow \det_{x_0}.$$

Let us now assume that we are given polytopes $P_i \subseteq \mathcal{N}_i$. Parallel transport (and composition) also induces a map

$$(2.2.51) \quad \Pi_\gamma: \text{Hom}_\Lambda^c(U_{1, x_2}^{P_1}, U_{2, x_2}^{P_2}) \otimes_\Lambda \text{Hom}_\Lambda^c(U_{0, x_1}^{P_0}, U_{1, x_1}^{P_1}) \rightarrow \text{Hom}_\Lambda^c(U_{0, x_0}^{P_0}, U_{2, x_0}^{P_2})$$

which can be defined as follows: let γ_{ij} denote the path from x_i to x_j determined by γ . We have

$$(2.2.52) \quad \Pi_\gamma (\psi_{12} \otimes \psi_{01}) = z^{[\gamma_{20}]} \circ \psi_{12} \circ z^{[\gamma_{12}]} \circ \psi_{01} \circ z^{[\gamma_{01}]}.$$

We also have a natural map

$$(2.2.53) \quad \lambda_{0,1,2} : \lambda_{1,2} \otimes \lambda_{0,1} \rightarrow \lambda_{0,2}$$

induced by concatenating paths.

Taking the sum, over all triples of critical points, of the tensor products of Equations (2.2.50), (2.2.51), and (2.2.53), we obtain a map

$$(2.2.54) \quad \begin{aligned} CM^* \left(X_q, \text{Hom}_\Lambda^c(U_1^{P_1}, U_2^{P_2}) \otimes \lambda \right) \otimes CM^* \left(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_1^{P_1}) \otimes \lambda \right) \\ \rightarrow CM^* \left(X_q, \text{Hom}_\Lambda^c(U_0^{P_0}, U_2^{P_2}) \otimes \lambda \right) \end{aligned}$$

which is given by the formula

$$(2.2.55) \quad \mu^2 := \bigoplus_{x_0^\vee \in \text{Crit}(0,2)} \sum_{\substack{x_1 \in \text{Crit}(0,1) \\ x_2 \in \text{Crit}(1,2)}} (-1)^{\deg(x_1)} \Pi_\gamma \otimes \lambda_{0,1,2} \otimes \delta_\gamma.$$

Since the sum is finite, this is necessarily a continuous map. Composing the left and right hand sides with the isomorphisms of Equation (2.2.38), we obtain the product

$$(2.2.56) \quad \begin{aligned} \mu^2 : CF^*((\mathcal{N}_1, P_1), (\mathcal{N}_2, P_2)) \otimes_\Lambda CF^*((\mathcal{N}_0, P_0), (\mathcal{N}_1, P_1)) \\ \rightarrow CF^*((\mathcal{N}_0, P_0), (\mathcal{N}_2, P_2)). \end{aligned}$$

The fact that this is a cochain map is again a combination of the standard Morse-theoretic description of the boundary of the moduli space of gradient trees, with the invariance of parallel transport maps under homotopies, and their compatibility with concatenation.

2.3. The cohomological category of polytopes. We now return to the setting discussed at the beginning of Section 2.1, and consider a cover of Q by integral affine polytopes $\{P_\sigma\}_{\sigma \in \Sigma}$, satisfying Condition (2.1.2). In addition, we assume that the cover is sufficiently small so that we may choose a contractible neighbourhood \mathcal{N}_σ of P_σ for each element of this cover with the property that:

$$(2.3.1) \quad \begin{aligned} & \text{all non-empty intersections among the sets } \{\mathcal{N}_\sigma\}_{\sigma \in \Sigma} \text{ are contractible, and} \\ & \text{such that } \mathcal{N}_\sigma \text{ contain } P_\tau \text{ whenever } \tau \in \Sigma_\sigma, \text{ i.e. whenever } P_\sigma \cap P_\tau \text{ is non-} \\ & \text{empty.} \end{aligned}$$

To see that such a cover exists, start by picking an arbitrary cover $\{\mathcal{N}_i\}_{i=1}^N$ by the interiors of integral affine polytopes so that all non-empty intersections are contractible. Then pick a simplicial triangulation Σ which is finer than this cover in the sense that the open star of any simplex lies in some \mathcal{N}_i . Then define \mathcal{N}_σ , for $\sigma \in \Sigma$, to be some choice of \mathcal{N}_i which contains the open star of σ .

2.3.1. Definition of the cohomological categories. Let us fix, for each $\sigma \in \Sigma$, a local Lagrangian section

$$(2.3.2) \quad \iota_\sigma : \mathcal{N}_\sigma \rightarrow X$$

of the projection $X \rightarrow Q$. We pick a Pin structure on \mathcal{N}_σ , as in the previous section, as well as a Morse function $f_{\tau, \sigma}$ on the fibre over a basepoint $q_{\tau, \sigma} \in \mathcal{N}_\sigma \cap \mathcal{N}_\tau$. Whenever $\tau = \sigma$, we

assume for consistency with the previous discussion that this basepoint agrees with $q_\sigma \in P_\sigma$. For each pair (P_0, P_1) of polytopes in $\mathcal{N}_\sigma \cap \mathcal{N}_\tau$, we then define

$$(2.3.3) \quad CF^*((\tau, P_0), (\sigma, P_1)) := CF^*((\mathcal{N}_\tau, P_0), (\mathcal{N}_\sigma, P_1)).$$

The most important case of the above construction is for the pair (P_τ, P_σ) , which allows us to define a category with morphisms

$$(2.3.4) \quad \mathcal{F}(\tau, \sigma) := \begin{cases} CF^*((\tau, P_\tau), (\sigma, P_\sigma)) & \tau \leq \sigma \\ 0 & \text{otherwise.} \end{cases}$$

Letting $\text{HF}(\tau, \sigma)$ denote the corresponding cohomology group, we obtain a category HF with compositions given as in Section 2.2.4. We omit the verification that the associativity conditions hold, as this will follow from the construction of an A_∞ category in Section 4. We denote by HF_σ the full subcategory with objects given by $\tau \in \Sigma_\sigma$.

Remark 2.16. In [6], the product was twisted by an explicit term obtained from a Čech cocycle representing the obstruction to the existence of a Lagrangian section of $X \rightarrow Q$ which is equipped with a Pin structure. To see that the construction of this paper is equivalent, note that the obstruction to a consistent trivialisation of the local systems $\lambda_{0,1}$ is exactly $w_2(Q) \in H^2(Q, \mathbb{Z}_2)$, while resolving the ambiguity in the construction of the local systems U^P over all basepoints corresponds to the choice of a global Lagrangian section of $X \rightarrow Q$.

In order to formulate and prove the twisted version of Theorem 1.1 discussed in Remark 1.5, one would choose a Čech cochain representative of the deformation class, which would give a deformed category HF^α by the methods of [6].

For each $\sigma \in \Sigma$ and pair of polytopes $P_0, P_1 \subseteq \mathcal{N}_\sigma$, we also define

$$(2.3.5) \quad \text{Po}_\sigma(P_0, P_1) := CF^*((\sigma, P_0), (\sigma, P_1)).$$

Letting $\text{HPO}_\sigma(P_0, P_1)$ denote the corresponding Floer cohomology group, we obtain a category HPO_σ , which will be the local (cohomological) category of polytopes associated to an element of Σ .

Remark 2.17. In Section 4.4, we shall find it useful to redefine Po_σ to add the choice of a basepoint $q_i \in P_i$ to each object. The additional choice gives a category with many more objects, but it will be clear from the computations of this section that objects corresponding to different choices of basepoints on the same polytope are quasi-isomorphic.

The fact that Floer cohomology groups are independent of the all auxiliary choices yields a faithful embedding $\text{HF}_\sigma \rightarrow \text{HPO}_\sigma$. One way to make this embedding explicit is as follows: Fix a homotopy of sections between the restrictions of ι_τ and ι_σ to $\mathcal{N}_\sigma \cap \mathcal{N}_\tau$. This induces an isomorphism of local systems:

$$(2.3.6) \quad U_\tau^{P_\tau} \rightarrow U_\sigma^{P_\tau}.$$

Taking the sum of these isomorphisms over all maxima of the Morse function f_{σ_0, σ_1} , we obtain a *continuation* element

$$(2.3.7) \quad \kappa \in CF^0((\tau, P_\tau), (\sigma, P_\tau))$$

which is closed and whose cohomology class is canonical. Given a pair (τ_0, τ_1) of objects of HF_σ , the left and right products with the corresponding continuation elements induce a map

$$(2.3.8) \quad \mathcal{F}(\tau_0, \tau_1) \cong CF^0((\tau_0, P_{\tau_0}), (\tau_1, P_{\tau_1})) \rightarrow CF^0((\sigma, P_{\tau_0}), (\sigma, P_{\tau_1})) \cong \text{Po}_\sigma(P_{\tau_0}, P_{\tau_1}).$$

Passing to cohomology, we obtain the functor

$$(2.3.9) \quad \mathrm{HF}_\sigma \rightarrow \mathrm{HPO}_\sigma.$$

Remark 2.18. The standard way of constructing a map of Morse complexes for different choices of Morse functions is to consider a 1-parameter family of vector fields interpolating between the two gradients, and counting solutions of the corresponding flow lines. Keeping in mind that a generic point in a manifold lies on a unique negative gradient flow line starting at a maximum, one sees that the count of pairs of gradient trees defining the product is the same count that defines a composition of continuation maps

$$(2.3.10) \quad \mathrm{CF}^0((\tau_0, P_{\tau_0}), (\tau_1, P_{\tau_1})) \rightarrow \mathrm{CF}^0((\sigma, P_{\tau_0}), (\sigma, P_{\tau_1})).$$

To show that this is a fully faithful embedding one may construct a continuation element in $\mathrm{CF}^0((\sigma, P_\tau), (\tau, P_\tau))$ as in Equation (2.3.7), and show that the product

$$(2.3.11) \quad \mathrm{CF}^0((\tau, P_\tau), (\sigma, P_\tau)) \otimes_\Lambda \mathrm{CF}^0((\sigma, P_\tau), (\tau, P_\tau)) \rightarrow \mathrm{CF}^0((\sigma, P_\tau), (\sigma, P_\tau))$$

maps the tensor products of the two continuation elements to the multiplicative unit.

2.3.2. Computation of morphisms in the category. To understand the categories HF and HPO_σ , we summarise some computations established in the Appendices. The first result is a computation for the Floer cohomology groups associated to inclusions: given a point q in \mathcal{N}_σ , and a critical point x of $f_{\sigma,\sigma}$, we have a natural map

$$(2.3.12) \quad \Lambda[H_1(X_q; \mathbb{Z})] \rightarrow \mathrm{Hom}_\Lambda^c(U_{\sigma,x}, U_{\sigma,x}).$$

If P_0 and P_1 are polytopes lying in \mathcal{N}_σ , so that $P_1 \subseteq P_0$, the above map has a unique continuous extension to a map

$$(2.3.13) \quad \Gamma^{P_1} \rightarrow \mathrm{Hom}_\Lambda^c(U_{\sigma,x}^{P_0}, U_{\sigma,x}^{P_1}).$$

The following result immediately follows from the first half Proposition C.1 proved in Appendix C.3:

Proposition 2.19. *If $P_1 \subseteq P_0$, the sum of the homomorphisms (2.3.13) over all minima of $f_{\sigma,\sigma}$ defines a natural quasi-isomorphism*

$$(2.3.14) \quad \Gamma^{P_1} \rightarrow \mathrm{CF}^*((\sigma, P_0), (\sigma, P_1)).$$

□

Combining this with continuation maps, we have:

Corollary 2.20. *If $P_1 \subseteq P_0 \subseteq \mathcal{N}_\sigma \cap \mathcal{N}_\tau$, there is an isomorphism*

$$(2.3.15) \quad \Gamma^{P_1} \rightarrow \mathrm{HF}^*((\sigma, P_0), (\tau, P_1)).$$

In particular, the right hand side vanishes except in degree 0. □

Remark 2.21. We remind the reader that the above isomorphism is not natural; it depends on the choice of path between the sections associated to σ and τ . In particular, it may not be compatible with products, and this is ultimately the reason for the appearance of the gerbe β , twisting the derived category of coherent sheaves, in our statement of mirror symmetry.

The above result implies that the morphism spaces in $H\mathcal{F}$ are given by

$$(2.3.16) \quad H\mathcal{F}(\tau, \sigma) = \begin{cases} \Gamma^{P_\tau} & \tau \leq \sigma \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, Equation (2.3.4) stipulates that the only morphisms in $H\mathcal{F}$ arise as Floer cohomology groups for pairs of nested polytopes, which agree with affinoid rings by the above result. In particular, $H\mathcal{F}$ is isomorphic to the category denoted \mathcal{F} in [6], allowing us to tie the constructions of the two papers.

Next, we consider a polytope $P \subseteq \mathcal{N}_\sigma$ for $\sigma \in \Sigma$, and a cover $\{P_\alpha\}_{\alpha \in A}$ of P , indexed by a finite ordered set. The natural (restriction) map $\Gamma^{P_\alpha} \rightarrow \Gamma^{P_\alpha \cap P_\beta}$ gives rise to a map of local systems $U_\sigma^{P_\alpha} \rightarrow U_\sigma^{P_\alpha \cap P_\beta}$, allowing us to form the Čech complex

$$(2.3.17) \quad \check{U}(P; A) := \bigoplus_{\alpha_0 \in A} U_\sigma^{P_{\alpha_0}} \rightarrow \bigoplus_{\alpha_0 < \alpha_1 \in A} U_\sigma^{P_{\alpha_0} \cap P_{\alpha_1}} \rightarrow \bigoplus_{\alpha_0 < \alpha_1 < \alpha_2 \in A} U_\sigma^{P_{\alpha_0} \cap P_{\alpha_1} \cap P_{\alpha_2}} \rightarrow \dots$$

as a complex of (topological) local systems over X_{q_σ} . Note that this is a finite direct sum of topological local systems, and thus there is no ambiguity in the construction of the topology on $\check{U}(P; A)$. Moreover, we have a canonical map of local systems

$$(2.3.18) \quad U_\sigma^P \rightarrow \check{U}(P; A)$$

given by the restriction to $U_\sigma^{P_\alpha}$ for all $\alpha \in A$.

We now consider an enlargement of the category of polytopes, including objects $\check{T}(P; A)$ which correspond to the complex of local systems $\check{U}(P; A)$. Explicitly, for each $P' \subseteq \mathcal{N}_\sigma$, we can define cochain groups

$$(2.3.19) \quad \text{Po}_\sigma(\check{T}(P; A), P') := CM^*(X_q, \text{Hom}_\Lambda^c(\check{U}(P; A), U_\sigma^{P'})) \otimes \lambda$$

$$(2.3.20) \quad \text{Po}_\sigma(P', \check{T}(P; A)) := CM^*(X_q, \text{Hom}_\Lambda^c(U_\sigma^{P'}, \check{U}(P; A))) \otimes \lambda$$

equipped with the sum of the Morse differential and the internal differential of $\check{U}(P; A)$.

Because $\check{U}(P; A)$ is built in finitely many steps, the Floer complex $\text{Po}_\sigma(P', \check{T}(P; A))$ is isomorphic to the complex

$$(2.3.21) \quad \bigoplus_{\alpha_0 \in A} \text{Po}_\sigma(P', P_{\alpha_0}) \rightarrow \bigoplus_{\alpha_0 < \alpha_1 \in A} \text{Po}_\sigma(P', P_{\alpha_0}) \rightarrow \bigoplus_{\alpha_0 < \alpha_1 < \alpha_2 \in A} \text{Po}_\sigma(P', P_{\alpha_0}) \rightarrow \dots,$$

and similarly for $\text{Po}_\sigma(\check{T}(P; A), P')$. This implies that these complexes are isomorphic to the ones discussed in Appendix B, where we describe $\check{T}(P; A)$ as a twisted complex in the category Po_σ . We can now state an immediate consequence of Proposition B.1, which is the version of Tate acyclicity which we shall use for computations:

Lemma 2.22. *The map from U_σ^P to $\check{T}(P; A)$ is a quasi-isomorphism of complexes of local systems, and induces quasi-isomorphisms*

$$(2.3.22) \quad \text{Po}_\sigma(\check{T}(P; A), P') \rightarrow \text{Po}_\sigma(P, P')$$

$$(2.3.23) \quad \text{Po}_\sigma(P', P) \rightarrow \text{Po}_\sigma(P', \check{T}(P; A))$$

for all $P' \subseteq \mathcal{N}_\sigma$. □

The above result allows us to reduce global computations to local computations. To fully make use of locality, we need the fact, proved in Appendix C.3, that the Floer complex $CF^*((\sigma_0, P_0), (\sigma_1, P_1))$ is acyclic whenever P_0 and P_1 are disjoint.

Corollary 2.23. *Let (P, P', P'') be polytopes contained in \mathcal{N}_σ . If the intersections of P' and P'' with an open neighbourhood of P agree, there are natural isomorphisms*

$$(2.3.24) \quad \text{HPo}_\sigma^*(P, P') \cong \text{HPo}_\sigma^*(P, P'')$$

$$(2.3.25) \quad \text{HPo}_\sigma^*(P', P) \cong \text{HPo}_\sigma^*(P'', P).$$

Proof. By taking the intersection of P' and P'' , it suffices to prove the result under the assumption that $P'' \subseteq P'$. In this case, we can extend P'' to a cover of P' with the property that all other elements of the cover intersect P trivially, hence the Floer cohomology of all other elements of the cover with P vanish. From Equation (2.3.21), we obtain an isomorphism

$$(2.3.26) \quad \text{HPo}_\sigma(\check{T}(P; A), P') \cong \text{HPo}_\sigma(P, P'').$$

The result then follows from Lemma 2.22. \square

By applying the above result twice, we conclude:

Corollary 2.24. *The morphisms between objects in HPo_σ are local in the sense that they only depend on a neighbourhood of the intersection of the corresponding polytopes.* \square

2.4. Local cohomological modules. In this section, we shall assign, to each Lagrangian L in \mathcal{A} , left and right modules over the categories HPo_σ , and use Tate acyclicity to compute these modules whenever L is Hamiltonian isotopic to a Lagrangian meeting X_{q_σ} at a point.

We begin by imposing a condition which can be achieved by a small Hamiltonian perturbation of L :

$$(2.4.1) \quad \text{the intersection of } L \text{ with every fibre } X_q \text{ is contained in the interior of a disjoint union of closed balls.}$$

Remark 2.25. In fact, after a generic perturbation, the intersection of L with any fibre is finite: to prove this, observe that the problem is equivalent to proving that the fibres of the projection to the base are generically finite. The study of projection of Lagrangian submanifolds to the base of Lagrangian fibration is the subject of Lagrangian singularities [7]. We shall only need a very small part of the theory: by a generic Hamiltonian perturbation, one may achieve transversality for the projection map $L \rightarrow Q$ with respect to any jet condition. This implies that, at any point and in any direction, there is some sufficiently high derivatives of the projection map which does not vanish, and hence that the intersection with all fibres are isolated.

For the remainder of the paper, we fix the disjoint union of balls appearing in Equation (2.4.1), assuming further that the distance between these balls is bounded above uniformly over $q \in Q$. For each element L of \mathcal{A} , we fix a compact codimension 0 submanifold $\nu_X L \subset X$ of L so that the intersection $\nu_X L \cap X_q$ is a finite union of path connected components which lie in this union of balls (the fact that this is a finite union is a consequence of the much stronger fact that the fibres of a generic smooth map are triangulable, which is the main result of [33]). We denote by \mathcal{J}_L the space of tame almost complex structures which agree with J_L away from a fixed compact subset of the interior of $\nu_X L$, where J_L is the almost complex structure with respect to which we have assumed that L does not bound any non-constant holomorphic disc.

We now impose constraints on the diameter of \mathcal{N}_σ , which will be essential in ensuring that, for each $\sigma \in \Sigma$, we can define a Floer cohomology group for pairs $L \in \mathcal{A}$ and $P \in \mathcal{N}_\sigma$. Consider the quotient X_{q_σ} / \sim by the equivalence relation which collapses each component of $X_{q_\sigma} \cap \nu_X L$ to a point. Since $\nu_X L$ is compact, the quotient space is compact and Hausdorff.

We shall presently see that a Riemannian metric on X_{q_σ} induces a length metric on X_{q_σ}/\sim . We recall that the length of a path with target any metric space may be defined as the supremum of the distance between successive points in a partition of the domain interval (see e.g. [10, Definition I.1.18]). A *length metric* is one with the property that the distance between two points is given by the infimum of the lengths of paths between them [10, Definition I.3.1]. Since the spaces we shall consider are compact, the Hopf-Rinow theorem holds and the infimum is achieved [10, Proposition I.3.7].

Lemma 2.26. *Let M be a compact Riemannian manifold, K a compact subset all of whose components are path components, and \sim the equivalence relation given by collapsing each component of K to a point. There is a length metric on X/\sim , determined by the property that, for each path $\gamma: I \rightarrow X$, the length $\ell(\gamma/\sim)$ of the image in X/\sim is given by the sum of the lengths of the subpaths which lie away from K .*

Proof. Writing χ_K for the characteristic function of K (which vanishes on K and is unity away from it), we can write the formula for the length as

$$(2.4.2) \quad \ell(\gamma/\sim) = \int_0^1 \left| \frac{d\gamma}{dt} \right| \cdot (1 - \chi_K(\gamma(t))) dt.$$

The main claim is therefore that, for points x and y in X/\sim the formula

$$(2.4.3) \quad d(x, y) := \inf_{\gamma: [0, 1] \rightarrow M} \ell(\gamma/\sim)$$

defines a metric, where the infimum is taken over all paths whose endpoints project to x and y . Symmetry is obvious, and the fact that d separates points follows from the assumption that K is closed, hence the Hausdorff distance between any two components is strictly positive. The triangle inequality follows from the assumption that each component is path connected, since this allows us to connect any two paths with endpoints in a component K_y via a path lying entirely in K_y , and this procedure does not change the length of the projection as we have defined it. \square

We now return to our specific setting, in which $M = X_{q_\sigma}$ and K is the intersection with $\nu_X L$. Using the fact that the quotient of a manifold by a ball does not change the homeomorphism type and the inclusion of the components of $\nu_X L \cap X_{q_\sigma}$ in a disjoint union of balls, we obtain a retraction

$$(2.4.4) \quad X_{q_\sigma} \rightarrow X_{q_\sigma}/\sim \rightarrow X_{q_\sigma}$$

so the map on first homology groups

$$(2.4.5) \quad H_1(X_{q_\sigma}; \mathbb{Z}) \rightarrow H_1(X_{q_\sigma}/\sim; \mathbb{Z})$$

is injective. The metric on X_{q_σ}/\sim equips the right hand side with a norm given by the minimal length of a representative. The first homology of X_{q_σ} inherits a norm satisfying the following property: for any loop γ in X_{q_σ} the length $\ell(\gamma/\sim)$ of the projection is greater than or equal to the norm $||[\gamma]||$ of the associated homology class.

We can therefore equip X_{q_σ}/\sim with the structure of a metric space so that the norm $||[\gamma]||$ of the homology class of any loop in X_{q_σ} is bounded above by the length $\ell(\gamma/\sim)$ of the projection.

For each J_L -holomorphic curve u from a strip to X , with boundary conditions on a fixed compact subset of $\nu_X L$ along $t = 0$, and X_{q_σ} at $t = 1$ (the coordinates on the strip $\mathbb{R} \times [0, 1]$ are (s, t)), we consider the energy $E(u) = \int u^* \omega$ and the length $\ell(\partial u/\sim)$ of the projection to X_{q_σ}/\sim of the boundary component of the strip labelled X_{q_σ} . According to Corollary A.2,

we may choose a constant C independent of u such that this length is bounded by $CE(u)$. We require that

(2.4.6) the diameter of \mathcal{N}_σ is bounded by $1/4C$.

By construction, we note that the constant C depends on the collection \mathcal{A} of chosen Lagrangians and the complex structure J_L (as well at the fibration $X \rightarrow Q$).

Remark 2.27. This is the first of many places where we impose a condition on the diameters of the covers $\{\mathcal{N}_\sigma\}$ and $\{P_\sigma\}$, and we will repeatedly be imposing conditions of this nature. The only difficulty with this idea is that the reverse isoperimetric inequality depends on the choice of Lagrangian boundary conditions, and that imposing further conditions on the cover entails changing which fibres are associated to elements of the cover (since we require $q_\sigma \in \mathcal{N}_\sigma$). As discussed in the introduction, the solution implemented in Section 3 is to establish a reverse isoperimetric inequality for compact families of choices, and then pick a cover which is sufficiently fine with respect to this uniform constant. In the case of the condition in Equation (2.4.6), we note that the constant C can be bounded from the distance between the balls which cover $X_q \cap \nu_X L$, minimised over all $q \in Q$.

The reason that the inductive procedure can ultimately be implemented is that the ordering of choices proceeds by increasing complexity: the algebraic constructions we perform at any given stage depend on reverse isoperimetric constants for some moduli space of holomorphic curves and all moduli spaces which appear in its boundary; these constants dictate how fine the cover has to be for the given algebraic construction to converge. Later constructions consider different moduli spaces of holomorphic curves which may include previous ones in their boundary, and simply impose additional constraints on the cover.

In order to construct a module over a field of characteristic different from 2, we assume that each $L \in \mathcal{A}$ is equipped with a Pin structure. In order for the categories we construct to be \mathbb{Z} graded, we use the fact that the (canonical up to homotopy) trivialisation of the square of the top exterior power of TQ induces a trivialisation of the square of the top (complex) exterior power of TX with respect to any compatible almost complex structure, i.e. a complex quadratic volume form. We then assume that the Lagrangians are graded with respect to the chosen quadratic volume form on X (in the sense of [25]). Note that all fibres X_q are canonically graded in this sense.

Remark 2.28. In the twisted setting of Remark 1.5, one would equip L with the data of a Pin structure relative the pull-back of a vector bundle on the 3-skeleton of Q , which represents the given class in $H^2(Q; \mathbb{Z}_2)$. We would also add to our assumptions the vanishing of the pullback of the chosen class in $H^2(Q; \Lambda_+)$ to L .

2.4.1. Floer complexes between polytopes and Lagrangians. In order to construct the modules associated to elements of \mathcal{A} , we begin by requiring that

(2.4.7) for each $\sigma \in \Sigma$ and $L \in \mathcal{A}$, the Lagrangian pair (X_{q_σ}, L) is transverse.

We note that this condition can be achieved by an arbitrarily small Hamiltonian perturbation, in particular by a perturbation taking place in $\nu_X L$, and that the condition that L be tautologically unobstructed is preserved by such a perturbation, since we may replace J_L by its pushforward under the chosen Hamiltonian isotopy. Let $\text{Crit}(\sigma, L)$ and $\text{Crit}(L, \sigma)$ denote the intersection of L with X_{q_σ} . The chosen Pin structures on X_{q_σ} and L give rise to an assignment δ_x of a \mathbb{Z}_2 -graded free abelian group of rank 1 associated to each element x of $\text{Crit}(\sigma, L)$ or $\text{Crit}(L, \sigma)$. The canonical grading of the fibre with respect to the standard

quadratic complex volume form on X , together with a choice of grading on L determine a \mathbb{Z} -grading on δ_x .

With this in mind, we define, for each $P \subseteq \mathcal{N}_\sigma$ the Floer complex

$$(2.4.8) \quad CF^*(L, (\sigma, P)) := \bigoplus_{x \in \text{Crit}(\sigma, L)} U_{\sigma, x}^P \otimes \delta_x.$$

A choice of paths connecting the endpoints of orbits to a basepoint on X_{q_σ} induces a (complete) norm on these complexes, and the corresponding topology is independent of choice.

In order to define the differential, pick a family of almost complex structures

$$(2.4.9) \quad J(L, \sigma) : [0, 1] \rightarrow \mathcal{J}_L$$

which restricts at 0 to J_L . We obtain a moduli space $\overline{\mathcal{R}}(L, \sigma)$ of finite energy $J(L, \sigma)$ stable holomorphic strips with boundary conditions L along $t = 0$, and X_{q_σ} at $t = 1$ (the coordinates on the strip $\mathbb{R} \times [0, 1]$ are (s, t)).

There is a natural evaluation

$$(2.4.10) \quad \overline{\mathcal{R}}(L, \sigma) \rightarrow \text{Crit}(\sigma, L) \times \text{Crit}(\sigma, L)$$

given by the asymptotic conditions at $\pm\infty$. We denote the fibre over (x_0, x_1) by $\overline{\mathcal{R}}(x_0, x_1)$. Choosing $J(L, \sigma)$ generically ensures that this is a topological manifold of dimension

$$(2.4.11) \quad \dim \overline{\mathcal{R}}(x_0, x_1) = \deg(x_0) - \deg(x_1) - 1,$$

whose boundary is covered by codimension-1 strata corresponding to breaking of strips:

$$(2.4.12) \quad \partial \overline{\mathcal{R}}(x_0, x_1) = \bigcup_{x \in \text{Crit}(\sigma, L)} \overline{\mathcal{R}}(x_0, x) \times \overline{\mathcal{R}}(x, x_1).$$

The output of Floer theory is that, whenever $\overline{\mathcal{R}}(x_0, x_1)$ has dimension 0, each element induces a map:

$$(2.4.13) \quad \delta_u : \delta_{x_1} \rightarrow \delta_{x_0}.$$

In order to define the differential in Equation (2.4.8), we recall from Section 2.2.2 that a path from x_0 to x_1 induces a parallel transport map from U_{σ, x_1}^P to U_{σ, x_0}^P . The boundary of an element $u \in \overline{\mathcal{R}}(x_0, x_1)$ gives rise to such a path which we denote ∂u , so that we have a parallel transport map $z^{[\partial u]}$ as in Equation (2.2.16). For the statement of the next result, we recall that the energy $E(u)$ of a holomorphic curve is its area.

Lemma 2.29. *Assuming that the cover satisfies Condition 2.4.6, there is a constant A , independent of u , such that whenever $P \subseteq \mathcal{N}_\sigma$, we have*

$$(2.4.14) \quad E(u) + \text{val}_P z^{[\partial u]} \geq E(u)/2 + A.$$

Proof. It suffices to bound $\text{val}_q z^{[\partial u]}$ for any $q \in \mathcal{N}_\sigma$. The condition that the distortion is bounded by 2 implies that the image of P in $T_q X$ is contained in the ball of radius $2 \text{diam } \mathcal{N}_\sigma$. Thus

$$(2.4.15) \quad \text{val}_q z^{[\partial u]} \geq -2 \text{diam } \mathcal{N}_\sigma |[\partial u]| \geq -\frac{|[\partial u]|}{2C},$$

where the second inequality follows from Equation (2.4.6). At the cost of introducing an additive constant, the choice of metric fixed in the discussion preceding (2.4.6) allows us to replace $|[\partial u]|$ by the length $\ell([\partial u]/\sim)$. We can then apply the reverse isoperimetric inequality: the key point is that, according to Corollary A.2 the reverse isoperimetric inequality for J_L -holomorphic with boundary conditions on $\nu_X L$ and X_{q_σ} applies (with the same constant) to

$J(L, \sigma)$ holomorphic curves, because the two almost complex structures agree by assumption away from a fixed compact subset of the interior of $\nu_X L$. The result thus follows. \square

The bound in Equation (2.4.14), which can of course be simplified, should be read as follows: the right hand side is the valuation of the operation $T^{E(u)} z^{[\partial u]}$, and the left hand side goes to $+\infty$ as the energy of u goes to infinity. From Gromov compactness, and as in Corollary 2.13, we therefore conclude:

Corollary 2.30. *For each pair (x_0, x_1) of intersection points, the expression*

$$(2.4.16) \quad \sum_{u \in \bar{\mathcal{R}}_q(x_0, x_1)} T^{E(u)} z^{[\partial u]}$$

converges with respect to the topology defined by P , hence induces a map from U_{σ, x_1}^P to U_{σ, x_0}^P . \square

We conclude that, for each pair (x_0, x_1) of intersection points, the expression

$$(2.4.17) \quad \partial_{x_0, x_1} := \sum_{u \in \bar{\mathcal{R}}_q(x_0, x_1)} T^{E(u)} z^{[\partial u]} \otimes \delta_u$$

gives a well-defined map from $U_{\sigma, x_1}^P \otimes \delta_{x_1}$ to $U_{\sigma, x_0}^P \otimes \delta_{x_0}$.

Definition 2.31. *The differential on $CF^*(L, (\sigma, P))$ is given by*

$$(2.4.18) \quad \bigoplus_{x_0} \sum_{x_1} (-1)^{\deg(x_1)+1} \partial_{x_0, x_1}.$$

By the previous discussion, this differential is a continuous operator with respect to the natural topology on Floer complexes (i.e. bounded with respect to the norm induced by a choice of homotopy classes of paths to the basepoint).

Remark 2.32. Note that the sign in Equation (2.4.18) differs by one from the sign in Equation (2.2.24). The reason for this choice is that the sign conventions for modules are more intuitive if they are based on unreduced gradings.

Reversing the rôles of the Lagrangian and the polytope, we construct a complex

$$(2.4.19) \quad CF^*((\sigma, P), L) := \bigoplus_{x \in \text{Crit}(\sigma, L)} \text{Hom}_{\Lambda}^c(U_{\sigma, x}^P, \Lambda) \otimes \delta_{x^q},$$

using moduli spaces of holomorphic strips $\bar{\mathcal{R}}(\sigma, L)$ with boundary X_{q_σ} along $t = 0$ and L along $t = 1$. To simplify the discussion later, we define $\bar{\mathcal{R}}(\sigma, L)$ using the t -dependent family of almost complex structure

$$(2.4.20) \quad J(\sigma, L)(t) = J(L, \sigma)(1 - t).$$

Following the above procedure, and using parallel transport along the boundary arc labelled X_{q_σ} , we obtain the differential on $CF^*((\sigma, P), L)$. Given the choices we have made, this complex is naturally isomorphic to the Λ -linear dual of $CF^*(L, (\sigma, P))$.

2.4.2. *Modules over the local categories.* For each $\sigma \in \Sigma$, $L \in \mathcal{A}$, and $P \subseteq \mathcal{N}_\sigma$, we define

$$(2.4.21) \quad \mathcal{L}_{L,\sigma}^*(P) := CF^*(L, (\sigma, P))$$

$$(2.4.22) \quad \mathcal{R}_{L,\sigma}^*(P) := CF^*((\sigma, P), L).$$

The differentials $\mu_{\mathcal{L}_{L,\sigma}}^{1|0}$ and $\mu_{\mathcal{R}_{L,\sigma}}^{0|1}$ are given by Equation (2.4.18) and its analogue when the boundary conditions are reversed.

We now construct the module structure on $H\mathcal{L}_{L,\sigma}$, i.e. the left action of morphism spaces in HPo_σ . The construction of the right module action is entirely similar, as we shall explain at the end.

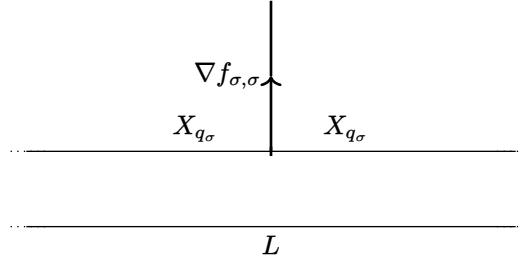


FIGURE 2. An element of the moduli space $\overline{\mathcal{TR}}(L, \sigma, \sigma)$ defining the left module over HPo_σ .

Define $\overline{\mathcal{R}}(L, \sigma, \sigma)$ to be the moduli space of strips in $\overline{\mathcal{R}}(L, \sigma)$, with an additional marked point along the segment mapping to X_{q_σ} . We have a natural evaluation map

$$(2.4.23) \quad \overline{\mathcal{R}}(L, \sigma, \sigma) \rightarrow \text{Crit}(L, \sigma) \times X_{q_\sigma} \times \text{Crit}(L, \sigma).$$

Let $\overline{\mathcal{TR}}(L, \sigma, \sigma)$ be the fibre product over X_{q_σ} of $\overline{\mathcal{R}}(L, \sigma, \sigma)$ with the space $\overline{\mathcal{T}}_+(\sigma, \sigma)$ of (perturbed) positive half-gradient flow lines of the function $f_{\sigma, \sigma}$. We shall call such moduli spaces *mixed moduli spaces*, as they consist of gradient flow lines and pseudo-holomorphic discs with matching evaluation maps to Lagrangians in X . Their use is standard in constructions combining Morse and Floer theory, going all the way back to [16].

We have a natural evaluation map

$$(2.4.24) \quad \overline{\mathcal{TR}}(L, \sigma, \sigma) \rightarrow \text{Crit}(L, \sigma) \times \text{Crit}(\sigma, \sigma) \times \text{Crit}(L, \sigma),$$

with the ordering given counterclockwise around the boundary starting at the outgoing end. Denote by E the set of ends and marked points, which we decompose into $E^{\text{inp}} = \{e_1, e_2\}$ and $E^{\text{out}} = \{e_0\}$, with E^{inp} consisting of the incoming (positive) end and the marked point, and E^{out} consisting of the singleton output. We denote the fibre at a triple $x = \{x_e\}_{e \in E}$ in the left hand side of Equation (2.4.24) by $\overline{\mathcal{TR}}(x)$. By parallel transport along the gradient flow line and the part of the boundary mapping to X_{q_σ} , we obtain a map

$$(2.4.25) \quad \text{Hom}_\Lambda^c(U_{\sigma, x_{e_2}}, U_{\sigma, x_{e_2}}) \otimes_\Lambda U_{\sigma, x_{e_1}} \rightarrow U_{\sigma, x_{e_0}}.$$

For generic choices of perturbations, the fibre product defining $\overline{\mathcal{TR}}(x)$ is transverse, so that it is a topological manifold with boundary of dimension

$$(2.4.26) \quad \deg(x_{e_0}) - \sum_{e \in E^{\text{inp}}} \deg(x_e),$$

and which is naturally oriented relative

$$(2.4.27) \quad \delta_{x_{e_0}}^\vee \otimes \bigotimes_{e \in E^{\text{inp}}} \delta_{x_e}.$$

A rigid element $u \in \overline{\mathcal{TR}}(\Upsilon)$ thus induces a map

$$(2.4.28) \quad \bigotimes_{e \in E^{\text{inp}}} \delta_{x_e} \rightarrow \delta_{x_{e_0}}.$$

Given a triple $\Upsilon = (L, P_1, P_2)$, with P_1 and P_2 contained in \mathcal{N}_σ , we combine Equations (2.4.25) and (2.4.28) to obtain a map

$$(2.4.29) \quad \text{Po}_\sigma(P_1, P_2) \otimes_\Lambda \mathcal{L}_{L, \sigma}(P_1) \rightarrow \mathcal{L}_{L, \sigma}(P_2)$$

which, upon twisting by $(-1)^{\deg(x_1)+1}$, gives rise to the structure map $\mu_{\mathcal{L}_{L, \sigma}}^{1|1}$. Passing to cohomology, we obtain

$$(2.4.30) \quad \text{HPo}_\sigma(P_1, P_2) \otimes_\Lambda H\mathcal{L}_{L, \sigma}(P_1) \rightarrow H\mathcal{L}_{L, \sigma}(P_2).$$

The construction of the right module map proceeds as follows: we construct a moduli space $\overline{\mathcal{R}}(\sigma, \sigma, L)$ by considering strips in $\overline{\mathcal{R}}(\sigma, L)$ with an additional marked point on the boundary with label X_{q_σ} , then define $\overline{\mathcal{TR}}(\sigma, \sigma, L)$ to be the fibre product over X_{q_σ} with the moduli space of semi-infinite gradient flow lines of $f_{\sigma, \sigma}$ (see Figure 3).

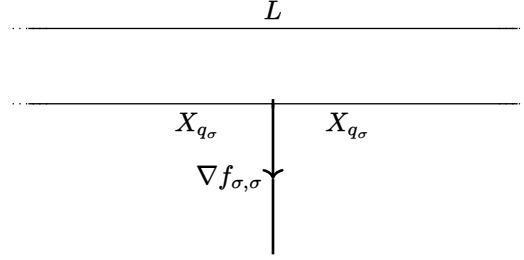


FIGURE 3. An element of the moduli space $\overline{\mathcal{TR}}(\sigma, \sigma, L)$ defining the right module over HPo_σ .

The count of rigid elements of these moduli spaces defines a map $\mu_{\mathcal{R}_{L, \sigma}}^{1|1}$ on Floer cochains

$$(2.4.31) \quad \mathcal{R}_{L, \sigma}(P_{-1}) \otimes_\Lambda \text{Po}_\sigma(P_{-2}, P_{-1}) \rightarrow \mathcal{R}_{L, \sigma}(P_{-2})$$

for each pair of polytopes (P_{-1}, P_{-2}) in \mathcal{N}_σ . At the level of cohomology, we obtain the desired map:

$$(2.4.32) \quad H\mathcal{R}_{L, \sigma}(P_{-1}) \otimes_\Lambda \text{HPo}_\sigma(P_{-2}, P_{-1}) \rightarrow H\mathcal{R}_{L, \sigma}(P_{-2}).$$

2.4.3. Computation of the module structure over HF_σ . Equation (2.3.16) gives a particularly simple description of the category HF_σ , with morphisms given by affinoid algebras. Identifying Γ^P with the completion of the homology of the based loop space of X_{q_σ} , the construction of the Floer complexes yields natural maps

$$(2.4.33) \quad \Gamma^{P'} \otimes_\Lambda \mathcal{L}_L(P) \rightarrow \mathcal{L}_L(P')$$

$$(2.4.34) \quad \mathcal{R}_L(P') \otimes_\Lambda \Gamma^{P'} \rightarrow \mathcal{R}_L(P),$$

whenever $P' \subseteq P$, arising from the map

$$(2.4.35) \quad U_{\sigma,x}^P \otimes_{\Lambda} \Gamma^{P'} \rightarrow U_{\sigma,x}^P \otimes_{\Gamma^P} \Gamma^{P'} \cong U_{\sigma,x}^{P'}.$$

Passing to cohomology, we conclude that the groups $H\mathcal{L}_{L,\sigma}(P_\tau)$ and $H\mathcal{R}_{L,\sigma}(P_\tau)$ give rise to modules over $H\mathcal{F}_\sigma$. This construction, which does not use any gradient trees or holomorphic discs in the module structure maps, was used in [6]. In this section, we show:

Lemma 2.33. *The pullbacks of $H\mathcal{L}_{L,\sigma}$ and $H\mathcal{R}_{L,\sigma}$ under the inclusion of $H\mathcal{F}_\sigma$ in $H\mathcal{P}_{\sigma}$ are naturally isomorphic to the modules constructed from Equation (2.4.35).*

The key point is the following result:

Lemma 2.34. *All contributions to the maps*

$$(2.4.36) \quad \text{Po}^0(P_1, P_2) \otimes_{\Lambda} \mathcal{L}_L^k(P_1) \rightarrow \mathcal{L}_L^k(P_2)$$

$$(2.4.37) \quad \mathcal{R}_L^k(P_{-1}) \otimes_{\Lambda} \text{Po}^0(P_{-2}, P_{-1}) \rightarrow \mathcal{R}_L^k(P_{-2})$$

are given by configurations whose holomorphic component is a constant strip.

Proof. By construction, the holomorphic component of such an element is a strip with endpoints x and y of equal Maslov index, hence has Fredholm index 0. However, the fact that the almost complex structure on the strip is translation invariant and our assumption that the moduli space is regular implies that the minimal Fredholm index of a non-constant disc is 1. \square

The identification between the two ways of constructing left and right modules can be derived as follows:

Proof of Lemma 2.33. By Lemma 2.34, if $\max x$ denotes the unique maximum of $f_{\sigma,\sigma}$ which is the endpoint of a (perturbed) gradient flow line starting at x , the left and right module actions are given by the sums of the maps

$$(2.4.38) \quad \text{Po}^0(P_1, P_2) \otimes_{\Lambda} \mathcal{L}_L^k(P_1) \rightarrow \bigoplus_{x \in \text{Crit}(L, P)} \text{Hom}_{\Lambda}^c(U_{\sigma, \max x}^{P_1}, U_{\sigma, \max x}^{P_2}) \otimes_{\Lambda} U_{\sigma,x}^{P_1} \rightarrow U_{\sigma,x}^{P_2}$$

$$(2.4.39) \quad \mathcal{R}_L^k(P_{-1}) \otimes_{\Lambda} \text{Po}^0(P_{-2}, P_{-1}) \rightarrow \bigoplus_{x \in \text{Crit}(L, P)} \text{Hom}_{\Lambda}^c(U_{\sigma,x}^{P_{-1}}, \Lambda) \otimes_{\Lambda} \text{Hom}(U_{\sigma, \max x}^{P_{-2}}, U_{\sigma, \max x}^{P_{-1}}) \rightarrow \text{Hom}_{\Lambda}^c(U_{\sigma,x}^{P_{-2}}, \Lambda)$$

where the second map (in both cases) is induced by parallel transport along the flow line from x to $\max x$. Using the inclusion $\Gamma^{P_2} \subseteq \text{Po}^0(P_1, P_2)$, we obtain the desired result. \square

We end this section by noting that the above discussion, together with the result stated in Proposition 2.19, implies:

Corollary 2.35. *If $P_2 \subseteq P_1$, the pullback of $\mathcal{L}_L^k(P_1)$ to $\Gamma^{P_2} \subseteq \text{Po}^0(P_1, P_2)$ is a free module of rank equal to the number of elements of $\text{Crit}(L, \sigma)$ of degree k .* \square

2.4.4. Computation for sections. In this section, we prove that Equation (2.1.22) is an isomorphism for sections, i.e.:

Proposition 2.36. *If L meets the fibre over q_σ transversely at a single point, the map*

$$(2.4.40) \quad j^* \Delta_{H\mathcal{P}_{\sigma}}(\sigma, -) \otimes_{H\mathcal{F}_\sigma} j^*(H\mathcal{L}_{L,\sigma}) \rightarrow j^* H\mathcal{L}_{L,\sigma}(\sigma)$$

is an isomorphism.

The tensor product on the left of Equation (2.1.22) is the cokernel of the map

$$(2.4.41) \quad \bigoplus_{\rho_0 \leq \rho_1 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \otimes_\Lambda \mathrm{HF}(P_{\rho_0}, P_{\rho_1}) \otimes_\Lambda H\mathcal{L}_\sigma(P_{\rho_0}) \rightarrow \bigoplus_{\rho_0 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_0}, P_\sigma) \otimes_\Lambda H\mathcal{L}_\sigma(P_{\rho_0}).$$

If we restrict to $\rho_0 = \rho_1$ in the above complex, the cokernel computes the tensor product over $\Gamma^{P_{\rho_0}} = \mathrm{HF}(P_{\rho_0}, P_{\rho_0})$ of the modules $\mathrm{HPo}_\sigma(P_{\rho_0}, P_\sigma)$ and $H\mathcal{L}_{L,\sigma}(P_{\rho_0})$. Thus the cokernel is the same as that of the map

$$(2.4.42) \quad \bigoplus_{\rho_0 < \rho_1 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \otimes_\Lambda \mathrm{HF}(P_{\rho_0}, P_{\rho_1}) \otimes_\Lambda H\mathcal{L}_\sigma(P_{\rho_0}) \rightarrow \bigoplus_{\rho_0 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_0}, P_\sigma) \otimes_{\Gamma^{P_{\rho_0}}} H\mathcal{L}_{L,\sigma}(P_{\rho_0}).$$

The above map factors through the direct sum of the surjections

$$(2.4.43) \quad \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \otimes_\Lambda \mathrm{HF}(P_{\rho_0}, P_{\rho_1}) \otimes_\Lambda H\mathcal{L}_\sigma(P_{\rho_0}) \rightarrow \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \otimes_{\Gamma^{P_{\rho_1}}} \mathrm{HF}(P_{\rho_0}, P_{\rho_1}) \otimes_{\Gamma^{P_{\rho_0}}} H\mathcal{L}_\sigma(P_{\rho_0})$$

Since $\mathrm{HF}(P_{\rho_0}, P_{\rho_1})$ is a free rank-1 module over $\Gamma^{P_{\rho_1}}$, the flatness of the map of rings $\Gamma^{P_{\rho_0}} \rightarrow \Gamma^{P_{\rho_1}}$ (see e.g. [30, Lemma 8.6]) yields an isomorphism

$$(2.4.44) \quad \mathrm{HF}(P_{\rho_0}, P_{\rho_1}) \otimes_{\Gamma^{P_{\rho_0}}} H\mathcal{L}_\sigma(P_{\rho_0}) \rightarrow H\mathcal{L}_\sigma(P_{\rho_1}).$$

It thus suffices to compute the cokernel of the map

$$(2.4.45) \quad \bigoplus_{\rho_0 < \rho_1 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \otimes_{\Gamma^{P_{\rho_1}}} H\mathcal{L}_\sigma(P_{\rho_1}) \rightarrow \bigoplus_{\rho_0 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_0}, P_\sigma) \otimes_{\Gamma^{P_{\rho_0}}} H\mathcal{L}_{L,\sigma}(P_{\rho_0}).$$

So far, the discussion has been completely general. We now use the assumption that L meets X_{q_σ} at a point: let x denote the intersection of L with X_{q_σ} . For each $P \subseteq \mathcal{N}_\sigma$, we have a canonical isomorphism

$$(2.4.46) \quad \mathcal{L}_{L,\sigma}(P) = U_{\sigma,x}^P,$$

with trivial differential. Using Corollary 2.35, we conclude that each module $H\mathcal{L}_{L,\sigma}(P)$ is free of rank 1 over Γ^{P_ρ} , and we are thus reduced to proving:

Lemma 2.37. *The following sequence*

$$(2.4.47) \quad \bigoplus_{\rho_0 < \rho_1 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_1}, P_\sigma) \rightarrow \bigoplus_{\rho_0 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_0}, P_\sigma) \rightarrow \mathrm{HPo}_\sigma(P_\sigma, P_\sigma) \rightarrow 0$$

is right exact, i.e. the middle arrow is an isomorphism from the cokernel of the first map to $\mathrm{HPo}_\sigma(P_\sigma, P_\sigma)$.

Proof. Let P be a polytope containing P_σ in its interior, and contained in the union of polytopes P_τ with $\tau \in \Sigma_\sigma$. By Corollary 2.23, the restriction

$$(2.4.48) \quad \mathrm{HPo}_\sigma(P_\tau \cap P, P_\sigma) \rightarrow \mathrm{HPo}_\sigma(P_\tau, P_\sigma),$$

is an isomorphism, so that the Lemma follows from the right exactness of the complex

$$(2.4.49) \quad \bigoplus_{\rho_0 < \rho_1 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_1} \cap P, P_\sigma) \rightarrow \bigoplus_{\rho_0 \in \Sigma_\sigma} \mathrm{HPo}_\sigma(P_{\rho_0} \cap P, P_\sigma) \rightarrow \mathrm{HPo}_\sigma(P_\sigma, P_\sigma) \rightarrow 0.$$

This is a consequence of Tate acyclicity: Lemma 2.22 implies that the map from the Čech complex

$$(2.4.50) \quad \left(\bigoplus_{0 \leq k} \bigoplus_{\rho_0 < \dots < \rho_k} \mathrm{HPo}_\sigma(P_{\rho_1} \cap P, P_\sigma)[-k], \check{\delta} \right)$$

to $\text{HPo}_\sigma(P_\sigma, P_\sigma)$ is a quasi-isomorphism. Passing to cohomology, the quotient complex corresponding to $1 \leq k$ yields a map of cohomology groups which fits in a commutative diagram

(2.4.51)

$$\begin{array}{ccc} \bigoplus_{\rho_0 < \rho_1 \in \Sigma_\sigma} \text{HPo}_\sigma(P_{\rho_1} \cap P, P_\sigma) & \longrightarrow & H^* \left(\bigoplus_{1 \leq k} \bigoplus_{\rho_0 < \dots < \rho_k} \text{HPo}_\sigma(P_{\rho_1} \cap P, P_\sigma)[-k], \check{\delta} \right) \\ & \searrow & \swarrow \\ & \bigoplus_{\rho_0 \in \Sigma_\sigma} \text{HPo}_\sigma(P_{\rho_0}, P_\sigma). & \end{array}$$

Moreover, the images of these two maps agree. Applying the long exact sequence on cohomology and using the fact that $\text{HPo}_\sigma(P_\sigma, P_\sigma)$ vanishes except in degree 0 implies the desired result. \square

Proposition 2.36 is an immediate consequence of the above result.

2.5. Local module duality. In this section, we fix $\sigma \in \Sigma$, and construct a map of bimodules

(2.5.1)

$$H\mathcal{L}_{L,\sigma} \otimes_{\Lambda} H\mathcal{R}_{L,\sigma} \rightarrow \Delta_{\text{HPo}_\sigma},$$

with target the diagonal bimodule of HPo_σ . We then prove that, if L is a Lagrangian section, this map defines an isomorphism between the values at P_σ of $H\mathcal{R}_{L,\sigma}$ and the bimodule dual of $H\mathcal{L}_{L,\sigma}$, i.e. the right module defined as the space of left module maps from $H\mathcal{L}_{L,\sigma}$ to the diagonal bimodule.

Remark 2.38. The results of this section will be used again when considering general Lagrangians in Section 4. The key idea is that such a Lagrangian gives rise to a filtered module whose associated graded module is isomorphic to a (finite) direct sum of modules corresponding to local Lagrangian sections. The argument given here can then be used for each of the summands of this associated graded module, and the statement that the left and right A_∞ modules associated to a Lagrangian are bimodule duals then follows by homological algebra (see Section 4.5).

2.5.1. A map to the local diagonal bimodule. Consider the triple $\Upsilon = (\sigma, L, \sigma)$. We define $\overline{\mathcal{R}}(\Upsilon)$ to consist of elements of $\overline{\mathcal{R}}(\sigma, L)$ equipped with a marked point along the boundary component labelled σ , and define $\overline{\mathcal{T}}\overline{\mathcal{R}}(\Upsilon)$ to be the fibre product of $\overline{\mathcal{R}}(\Upsilon)$ with the space $\overline{\mathcal{T}}_-(\sigma, \sigma)$ of (perturbed) negative half-gradient flow lines. We have a natural evaluation map

$$(2.5.2) \quad \overline{\mathcal{T}}\overline{\mathcal{R}}(\Upsilon) \rightarrow \text{Crit}(\sigma, \sigma) \times \text{Crit}(\sigma, L) \times \text{Crit}(L, \sigma).$$

For each pair (P, P_-) of polytopes contained in \mathcal{N}_σ , the count of rigid elements of this moduli space thus defines a map

$$(2.5.3) \quad \mathcal{L}_{L,\sigma}(P) \otimes_{\Lambda} \mathcal{R}_{L,\sigma}(P_-) \rightarrow \text{Po}_\sigma(P_-, P)$$

which, on co-homology, induces a map

$$(2.5.4) \quad H\mathcal{L}_{L,\sigma}(P) \otimes_{\Lambda} H\mathcal{R}_{L,\sigma}(P_-) \rightarrow \text{HPo}_\sigma(P_-, P).$$

We omit, as usual, the straightforward verification that this arises from a map of bimodules. Dualising, we obtain

$$(2.5.5) \quad H\mathcal{R}_{L,\sigma} \rightarrow \text{Hom}_{\text{HPo}_\sigma}(H\mathcal{L}_{L,\sigma}, \Delta_{\text{HPo}_\sigma}).$$

We have the following variant of Lemma 2.34:

Lemma 2.39. *All contributions to the maps*

$$(2.5.6) \quad \mathcal{L}_L^{n-k}(P) \otimes_{\Lambda} \mathcal{R}_L^k(P_-) \rightarrow \text{Po}_{\sigma}^n(P_-, P)$$

are given by configurations whose holomorphic component is a constant strip. \square

Letting $\min x$ denote the unique minimum of $f_{\sigma, \sigma}$ which is the endpoint of a (perturbed) gradient flow line starting at x , we conclude

Corollary 2.40. *In degree n , the map in Equation 2.5.4 is given by the sum of the maps*

$$(2.5.7)$$

$$\mathcal{L}_L^{n-k}(P) \otimes_{\Lambda} \mathcal{R}_L^k(P_-) \rightarrow \bigoplus_{x \in \text{Crit}(L, \sigma)} U_{\sigma, x}^P \otimes_{\Lambda} \text{Hom}_{\Lambda}^c(U_{\sigma, x}^{P_-}, \Lambda) \rightarrow \text{Hom}(U_{\sigma, \min x}^{P_-}, U_{\sigma, \min x}^P)$$

associated to each critical point x , where the second map is induced by parallel transport along the flow line from x to $\min x$. \square

2.5.2. *Computing the right module for sections.* In this section, we prove:

Lemma 2.41. *The restriction of Equation (2.5.5) to P_{σ}*

$$(2.5.8) \quad H\mathcal{R}_{L, \sigma}(P_{\sigma}) \rightarrow \text{Hom}_{\text{HPo}_{\sigma}}(H\mathcal{L}_{L, \sigma}, \Delta_{\text{HPo}_{\sigma}}(P_{\sigma}, -))$$

is an isomorphism whenever L meets $X_{q_{\sigma}}$ transversely at one point.

Before giving the proof, we need the following result, which should be thought of as the dual of Proposition 2.19, and is a restatement of the second half of Proposition C.1:

Proposition 2.42. *If P_0 and P_1 are polytopes so that P_0 is contained in the interior of P_1 , a choice of orientation of P_0 determines a trace*

$$(2.5.9) \quad \text{tr}: \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \rightarrow \Lambda.$$

Composing the trace with the module action, we obtain a map

$$(2.5.10) \quad CF^n((\sigma, P_0), (\sigma, P_1)) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Lambda)$$

$$(2.5.11) \quad \theta \mapsto \text{tr} \circ \theta \circ m$$

(where n is the dimension of Q), which induces an isomorphism

$$(2.5.12) \quad HF^*((\sigma, P_0), (\sigma, P_1)) \cong \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Lambda).$$

\square

We now prove the Lemma stated above:

Proof of Lemma 2.41. The argument is formally dual to the one given in Section 2.4.4. We start by noting that the right hand side is the kernel of the map

$$(2.5.13)$$

$$\prod_{\rho_0 \in \Sigma_{\sigma}} \text{Hom}_{\Lambda}(H\mathcal{L}_{L, \sigma}(P_{\rho_0}), \text{HPo}_{\sigma}(P_{\sigma}, P_{\rho_0})) \rightarrow \prod_{\rho_0 \leq \rho_1 \in \Sigma_{\sigma}} \text{Hom}_{\Lambda}(H\mathcal{F}(P_{\rho_0}, P_{\rho_1}) \otimes_{\Lambda} H\mathcal{L}_{L, \sigma}(P_{\rho_0}), \text{HPo}_{\sigma}(P_{\sigma}, P_{\rho_1})).$$

We use the isomorphism $\Gamma^{P_{\rho_0}} = H\mathcal{F}(P_{\rho_0}, P_{\rho_0})$, the fact that $H\mathcal{F}(P_{\rho_0}, P_{\rho_1})$ is a free rank-1 module over $\Gamma^{P_{\rho_1}}$, and the flatness of the map $\Gamma^{P_{\rho_0}} \rightarrow \Gamma^{P_{\rho_1}}$ to rewrite this as the kernel of the map

$$(2.5.14)$$

$$\prod_{\rho_0 \in \Sigma_{\sigma}} \text{Hom}_{\Gamma^{P_{\rho_0}}}(H\mathcal{L}_{L, \sigma}(P_{\rho_0}), \text{HPo}_{\sigma}(P_{\sigma}, P_{\rho_0})) \rightarrow \prod_{\rho_0 < \rho_1 \in \Sigma_{\sigma}} \text{Hom}_{\Gamma^{P_{\rho_1}}}(H\mathcal{L}_{L, \sigma}(P_{\rho_1}), \text{HPo}_{\sigma}(P_{\sigma}, P_{\rho_1})).$$

Assuming that L meets X_{q_σ} at one point implies that $H\mathcal{L}_{L,\sigma}(P_\rho)$ is free of rank one over Γ^{P_ρ} , hence simplifying the above to:

$$(2.5.15) \quad \prod_{\rho_0 \in \Sigma_\sigma} \text{HPo}_\sigma(P_\sigma, P_{\rho_0}) \rightarrow \prod_{\rho_0 < \rho_1 \in \Sigma_\sigma} \text{HPo}_\sigma(P_\sigma, P_{\rho_1}).$$

We now introduce a polytope P , containing P_σ in its interior, and covered by its intersection with P_τ . By Corollary 2.23, the restriction

$$(2.5.16) \quad \text{HPo}_\sigma(P_\sigma, P_\tau) \rightarrow \text{HPo}_\sigma(P_\sigma, P \cap P_\tau)$$

is an isomorphism. The desired computation then follows from the fact that the complex

$$(2.5.17) \quad 0 \rightarrow \text{HPo}_\sigma(P_\sigma, P) \rightarrow \prod_{\rho_0 \in \Sigma_\sigma} \text{HPo}_\sigma(P_\sigma, P \cap P_{\rho_0}) \rightarrow \prod_{\rho_0 < \rho_1 \in \Sigma_\sigma} \text{HPo}_\sigma(P_\sigma, P \cap P_{\rho_1})$$

is left exact, which is a consequence of Tate acyclicity.

Since P_σ is contained in the interior of P , the computations of Appendix C, summarised above in Proposition 2.42, together with the definition of $\mathcal{R}_{L,\sigma}(P_\sigma)$ yield an isomorphism

$$(2.5.18) \quad \text{HPo}_\sigma(P_\sigma, P) \cong \text{Hom}_\Lambda^c(U_{\sigma,x}^{P_\sigma}, \Lambda) \cong H\mathcal{R}_{L,\sigma}(P_\sigma).$$

As in Section 2.4.4, this isomorphism is compatible with module action, implying Lemma 2.41. \square

2.6. Global cohomological modules. Given $L \in \mathcal{A}$ and $\sigma \in \Sigma$, we define

$$(2.6.1) \quad \mathcal{L}_L^*(\sigma) := CF^*(L, (\sigma, P_\sigma))$$

$$(2.6.2) \quad \mathcal{R}_L^*(\sigma) := CF^*((\sigma, P_\sigma), L),$$

with differentials $\mu_{\mathcal{L}_L}^{1|0}$ and $\mu_{\mathcal{R}_L}^{0|1}$ obtained from Equation (2.4.18) and its dual. In this section, we show that the corresponding cohomology groups give rise to left and right modules $H\mathcal{L}_L$ and $H\mathcal{R}_L$ over $H\mathcal{F}$. In Section 2.6.5 below, we show that the restrictions of these modules to $H\mathcal{F}_\sigma$ are isomorphic to the pullbacks of the corresponding local modules, while in Section 2.6.6, we construct the maps relating the Floer cohomology of Lagrangians in \mathcal{A} with the corresponding left and right modules. All these constructions are minor variants of those appearing in [6].

2.6.1. Continuation maps with non-Hamiltonian boundary conditions. Consider a triple

$$(2.6.3) \quad \Upsilon = (\sigma_{-2}, \sigma_{-1}, L) \text{ or } \Upsilon = (L, \sigma_0, \sigma_1),$$

such that $P_{\sigma_{-2}} \cap P_{\sigma_{-1}} \neq \emptyset$, or $P_{\sigma_0} \cap P_{\sigma_1} \neq \emptyset$. The continuation maps between modules associated to L and the two elements of the cover will be defined by a count of pseudo-holomorphic strips, so we begin by labelling the positive end e_{inp} of the strip $S = \mathbb{R} \times [0, 1]$ with the pair $\Upsilon_{e_{\text{inp}}}$ of labels (σ_0, L) or (L, σ_{-1}) , and the negative end e_{out} with the pair $\Upsilon_{e_{\text{out}}}$ given by (σ_1, L) or (L, σ_{-2}) . We thus obtain assignments $(\sigma_e, q_{\Upsilon_e})$ of elements of Σ and Q for each end e . Fix disjoint neighbourhoods ν_{Se} of these ends, and write $\partial_\Sigma S$ for the boundary component of S which is labelled by elements of Σ .

Consider the straight path q_Υ from $q_{\sigma_{e_{\text{inp}}}}$ to $q_{\sigma_{e_{\text{out}}}}$ determined by the affine structure; we parametrise this path by $\partial_\Sigma S$, with the condition that it agree with q_{σ_e} in ν_{Se} . The quotient of $\partial_\Sigma S$ by the intersection with ν_{Se} is a closed interval.

If we choose a family

$$(2.6.4) \quad J(\Upsilon): S \rightarrow \mathcal{J}$$

of almost complex structures which agree with $J(\Upsilon_e)$ along the end e , and whose restriction to the boundary $t = 1$ (the coordinates are (s, t) on the strip) agrees with J_L we obtain a moduli space $\overline{\mathcal{R}}(\Upsilon)$ of stable $J(\Upsilon)$ -holomorphic strips with boundary conditions given by the path X_{qr} along the boundary $t = 0$ (ordered compatibly with the counter-clockwise orientation), and by L along $t = 1$. By evaluation along the boundary $t = 0$, we associate to each element u of $\overline{\mathcal{R}}(\Upsilon)$ a path

$$(2.6.5) \quad \partial_Q u: \mathbb{R} \rightarrow X_S,$$

where the right hand side denotes the fibre bundle over a closed interval, obtained from the Lagrangian boundary conditions by collapsing the inverse image of the neighbourhoods ν_S to the fibres $X_{q_{\sigma_e}}$.

Consider as in Section 2.4 the quotient X_S/\sim by the equivalence relation which collapses the components of the intersection of $\nu_X L$ with the boundaries $X_{q_{\sigma_{e_{\text{inp}}}}}$ and $X_{q_{\sigma_{e_{\text{out}}}}}$ to points. According to Lemma A.3, there is a constant C such that, up to an additive constant, the length of $\partial_Q u/\sim$ is bounded by $CE^{geo}(u)$, where the *geometric energy* is given by

$$(2.6.6) \quad E^{geo}(u) = \int |du|^2 = \int u^* \omega$$

with the norm taken with respect to the metric induced by the almost complex structure. Using the fact that we have an isomorphism $H_1(X_q, \mathbb{Z}) \cong H_1(X_{qr}, \mathbb{Z})$ for any q in the path qr , we arrange as before for the metric on X_{qr}/\sim to have the property that for any $q \in qr$ the norm of the homology class in $H_1(X_q, \mathbb{Z})$ associated to a loop in X_{qr} is bounded by the length of the image of this loop in X_{qr}/\sim . We then require that

$$(2.6.7) \quad \text{the diameter of } \mathcal{N}_\sigma \text{ is bounded by } 1/8C.$$

Remark 2.43. We are slightly abusing notation by using C for the constant C appearing above, instead of using a symbol that distinguishes it from the constant in Condition (2.4.6). One justification for the abuse of notation is that Condition (2.4.6) holds for a given constant C whenever Condition (2.6.7) holds for the same constant: the former arises from a reverse isoperimetric inequality for holomorphic strips, and the latter for continuation maps, and the moduli spaces of strips arise as boundary strata in the moduli spaces of continuation maps.

2.6.2. Energy of continuation maps. The moduli space $\overline{\mathcal{R}}(\Upsilon)$ is slightly unusual because the moving Lagrangian boundary conditions along $t = 0$ do not form a Hamiltonian family. The standard description of the Gromov-Floer bordification (in terms of breaking of strips at the ends, bubbling of discs at the boundary and of spheres at the interior) applies with unmodified proof, as does the treatment of regularity.

The main difference with the standard case of Hamiltonian families is that Gromov compactness fails in general. We now explain that it remains valid in our setting: the section $\iota_{\sigma_{e_{\text{inp}}}}$ determines an identification of $X_{\mathcal{N}_{\sigma_{e_{\text{inp}}}}}$ with a neighbourhood of the 0-section in the cotangent bundle of the fibre over the basepoint q_{σ_e} , in such a way that the section corresponds to a cotangent fibre; let $\lambda_{\sigma_{e_{\text{inp}}}}$ denote the (locally defined) Liouville 1-form obtained from this identification. Since the boundary conditions are not Lagrangian, the expression for geometric energy in Equation 2.6.6 is not invariant under homotopies of the map u . This leads us to define the (topological) energy of an element $u \in \overline{\mathcal{R}}(\Upsilon)$ as

$$(2.6.8) \quad E(u) := \int u^* \omega - \int (\partial_Q u)^* \lambda_{\sigma_{e_{\text{inp}}}}.$$

Lemma 2.44. *The (topological) energy of u depends only on its relative homotopy class; i.e. if $\{u_r\}_{r \in [0,1]}$ is a 1-parameter family of strips with boundary conditions given by the path X_{q_Υ} along the boundary $t = 0$, and by L along $t = 1$, with constant asymptotic conditions at the ends, then $E(u_0) = E(u_1)$.*

Proof. This is a straightforward application of Stokes's theorem: consider the associated map

$$(2.6.9) \quad [0, 1] \times \mathbb{R} \times [0, 1] \rightarrow X.$$

Since ω is closed, and we have assumed that the asymptotic conditions are constant, the integral of ω over the boundary vanishes. This sum decomposes in two 4 terms; the integrals over the faces corresponding to the boundary of the first factor give rise to the terms involving ω in $E(u_0)$ and $E(u_1)$. For the integral over the face corresponding to $t = 0$, we note that the image of this face is contained in $X_{\mathcal{N}_{\sigma_{e_{\text{inp}}}}}$, and that the choice of primitive $\lambda_{\sigma_{e_{\text{inp}}}}$ allows us to apply Stokes's theorem to yield the two corresponding terms in $E(u_i)$, using the fact that contributions from the end $s = \pm\infty$ vanish because the asymptotic conditions are constant. Finally, the restriction of the map to the face corresponding to $t = 1$ factors through L , so the Lagrangian condition implies that the corresponding term vanishes. Having accounted for all the terms in $E(u_0) - E(u_1)$, the result follows. \square

We now prove that Gromov compactness holds, under Condition (2.6.7):

Lemma 2.45. *For $u \in \overline{\mathcal{R}}(\Upsilon)$, we have*

$$(2.6.10) \quad E(u) \geq \frac{3}{4} E^{\text{geo}}(u) + \text{a constant independent of } u.$$

In particular, for any positive real number E , the subset of $\overline{\mathcal{R}}(\Upsilon)$ consisting of curves of topological energy bounded by E is compact.

Proof. Gromov's argument [21] shows that the geometric energy $\int u^* \omega$ defines a proper map $\overline{\mathcal{R}}(\Upsilon) \rightarrow [0, \infty)$, so that the first statement implies the second, and it suffices to bound the difference between the two energies.

Identifying the path q_Υ with a path in $H^1(X_{q_{\sigma_{e_{\text{inp}}}}}; \mathbb{R})$, and lifting $\partial_Q u$ to a path $\widetilde{\partial_Q u}$ in $H_1(X_{q_{\sigma_{e_{\text{inp}}}}}; \mathbb{R})$, we can express the integral of the 1-form $\lambda_{\sigma_{e_{\text{inp}}}}$ as

$$(2.6.11) \quad \int (\partial_Q u)^* \lambda_{\sigma_{e_{\text{inp}}}} = \int_{-\infty}^{\infty} \langle q_\Upsilon(s), \partial_s \widetilde{\partial_Q u} \rangle ds$$

$$(2.6.12) \quad = \langle q_\Upsilon(-\infty), \widetilde{\partial_Q u}(-\infty) \rangle - \int_{-\infty}^{\infty} \langle \partial_s q_\Upsilon(s), \widetilde{\partial_Q u} \rangle ds$$

$$(2.6.13) \quad \leq \frac{1}{8C} \cdot \ell(\partial_Q u) + \ell(\partial_Q u) \int_{-\infty}^{\infty} |\partial_s q_\Upsilon(s)| ds$$

$$(2.6.14) \quad \leq \frac{\ell(\partial_Q u)}{4C}.$$

Above, we have used Equation (2.6.7) to bound the length of q_Υ . The reverse isoperimetric inequality thus implies that

$$(2.6.15) \quad \left| \int (\partial_Q u)^* \lambda_{\sigma_{e_{\text{inp}}}} \right| \leq E^{\text{geo}}(u)/4.$$

We conclude the desired result. \square

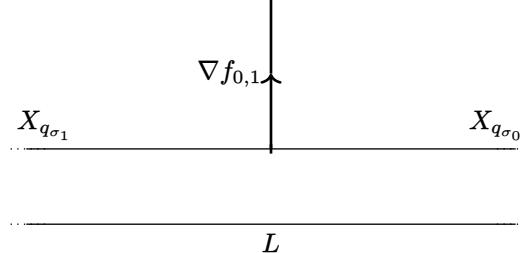


FIGURE 4. An element of the moduli space $\overline{\mathcal{TR}}(L, \sigma_0, \sigma_1)$ defining the left module over \mathcal{HF} .

We now restrict to the union of components corresponding to a pair of intersection points (x_-, x_+) . Choosing a path, for each end e , from these points to the basepoint in $X_{q_{\sigma_e}}$ obtained by intersecting with the section associated to $\sigma_{e_{\text{inp}}}$, we obtain a homology class $[\partial_Q u]$ in $H_1(X_{q_{\sigma_e}}; \mathbb{Z})$ associated to $u \in \overline{\mathcal{R}}(x_-, x_+)$. Using Condition (2.6.7), we find that

$$(2.6.16) \quad |[\partial_Q u]| \leq CE^{geo}(u) \leq 3C/4E(u) + \text{a constant independent of } u.$$

2.6.3. Mixed moduli spaces. We equip the strip with the marked point $e = (0, 0)$ in the case $\Upsilon = (\sigma_{-2}, \sigma_{-1}, L)$, and $e = (0, 1)$ in the case $\Upsilon = (L, \sigma_0, \sigma_1)$. We assign to this marked point the label $\Upsilon_e = (\sigma_{-2}, \sigma_{-1})$ in the first case and $\Upsilon_e = (\sigma_0, \sigma_1)$ in the second.

This choice of marked point equips the moduli space $\overline{\mathcal{R}}(\Upsilon)$ from the previous section with a natural evaluation map

$$(2.6.17) \quad \overline{\mathcal{R}}(\Upsilon) \rightarrow X_{q_{\Upsilon_e}}$$

where we use the trivialisation associated to the first element of Υ_e to identify the two fibres. The manifold $X_{q_{\Upsilon_e}}$ carries a Morse function f_{Υ_e} from Section 2.2.3, and a choice of family of vector fields parametrised by $[0, \infty)$, which agree with the gradient flow outside a compact set, yields a moduli space $\overline{\mathcal{T}}_+(\Upsilon_e)$, which also admits an evaluation map to $X_{q_{\Upsilon_e}}$. We define the mixed moduli space as a fibre product (see Figure 4)

$$(2.6.18) \quad \overline{\mathcal{TR}}(\Upsilon) := \overline{\mathcal{R}}(\Upsilon) \times_{X_{q_{\Upsilon_e}}} \overline{\mathcal{T}}_+(\Upsilon_e).$$

As before, this space admits an evaluation map to the product of the intersection points associated to the triples $\{e, e_{\text{out}}, e_{\text{inp}}\}$, and we denote the fibre over a triple x of generators by $\overline{\mathcal{TR}}(x)$

2.6.4. The global left and right multiplications. If $\Upsilon = (L, \sigma_0, \sigma_1)$, an element of $\overline{\mathcal{TR}}(x)$ defines a map

$$(2.6.19) \quad \text{Hom}(U_{\sigma_0, x_e}^{P_{\sigma_0}}, U_{\sigma_1, x_e}^{P_{\sigma_1}}) \otimes_{\Lambda} U_{\sigma_0, x_{e_{\text{inp}}}}^{P_{\sigma_0}} \rightarrow U_{\sigma_1, x_{e_{\text{out}}}}^{P_{\sigma_1}},$$

where we have used the canonical identifications from Lemma 2.11 to omit the superscript from the intersection points of Lagrangians. The map is defined as follows: given an element u in this moduli space represented by a pair (v, γ) , the trivialisation $\iota_{\sigma_{e_{\text{inp}}}}$, allows us to identify x_e and $x_{e_{\text{out}}}$ with points in $X_{q_{e_{\text{inp}}}}$. Applying this trivialisation to the path along the boundary of the elements of v from the positive end to $v(e)$ and the gradient flow line γ , we obtain a parallel transport map

$$(2.6.20) \quad U_{\sigma_0, x_{e_{\text{inp}}}}^{P_{\sigma_0}} \rightarrow U_{\sigma_0, x_e}^{P_{\sigma_0}}.$$

Given an element of the left hand side of Equation (2.6.19), the composition of this isomorphism with the given linear map in $\text{Hom}_\Lambda^c(U_{\sigma_0, x_e}^{P_{\sigma_0}}, U_{\sigma_1, x_e}^{P_{\sigma_1}})$ yields an element of $U_{\sigma_1, x_e}^{P_{\sigma_1}}$ which we transport, along the image in $X_{q_{e_{\text{inp}}}}$ of the gradient flow line and the boundary of the holomorphic strip to the fibre of $U_{\sigma_1}^{P_{\sigma_1}}$ at $x_{e_{\text{out}}}$, which is the right hand side of Equation (2.6.19).

Remark 2.46. As explained in the discussion following Lemma 2.11, one must specify the fibre along which one performs parallel transport because the identifications of the local systems over different fibres are only compatible with parallel transport maps up to multiplication by the exponential of the flux.

Multiplying the tensor product of the isomorphism in Equations (2.4.28) and that in Equation (2.6.19) by $(-1)^{\deg(x_1)+1}$, and restricting to continuous morphisms from $U_{\sigma_0, x_e}^{P_{\sigma_0}}$ to $U_{\sigma_1, x_e}^{P_{\sigma_1}}$ we obtain, for u a rigid element in $\overline{\mathcal{TR}}(\Upsilon)$, a map

$$(2.6.21) \quad \mu_u: \mathcal{F}(\sigma_0, \sigma_1) \otimes_\Lambda \mathcal{L}_L(\sigma_0) \rightarrow \mathcal{L}_L(\sigma_1).$$

Lemma 2.47. *Under assumption (2.6.7), the sum*

$$(2.6.22) \quad \sum_{u \in \overline{\mathcal{TR}}(\sigma_1, \sigma_0, L)} T^{E(u)} \mu_u$$

is convergent.

Proof. The key point is to bound the valuation of $T^{E(u)} \mu_u$. The map μ_u is a composition of two parallel transport maps with an evaluation map. The evaluation map has trivial valuation and the valuation of the two parallel transport maps is bounded by the sum of the norms of the two corresponding loops obtained by concatenating with fixed paths to a basepoint. Up to an additive constant independent of u , the sum of norms of these loops is bounded by the length of the projection of $\partial_Q u$ to $X_{q_{\Upsilon}}/\sim$. Applying the reverse isoperimetric inequality, we conclude that the valuation of the parallel transport maps is bounded by $CE^{geo}(u)$. By Condition (2.6.7), the polytopes P_{σ_i} are contained in the ball of radius $1/4C$ in $T_{q_{e_{\text{inp}}}} Q$, so that applying Lemma 2.12 yields

$$(2.6.23) \quad \text{val } \mu_u \geq -E^{geo}(u)/4 + \text{a constant independent of } u.$$

From Lemma 2.45, we conclude that the valuation of $T^{E(u)} \mu_u$ is bounded above by $E(u)/2$, up to a constant term independent of u . Gromov compactness (for the geometric energy) thus implies that the sum is convergent. \square

Remark 2.48. Note that the definitions of $E(u)$ and μ_u are asymmetric, since they rely on distinguishing the fibre $X_{q_{\text{inp}}}$ corresponding to the incoming end. However, the product $T^{E(u)} \mu_u$ is independent of this choice.

We denote the sum in Lemma 2.47 by

$$(2.6.24) \quad \mu_{\mathcal{L}_L}^{1|1}: \mathcal{F}^*(\sigma_0, \sigma_1) \otimes_\Lambda \mathcal{L}_L^*(\sigma_0) \rightarrow \mathcal{L}_L^*(\sigma_1).$$

Applying Lemma 2.44 (and using the fact that the topological energy is additive for broken curves), the standard description of the boundary of $\overline{\mathcal{TR}}(\Upsilon)$ implies that $\mu_{\mathcal{L}_L}^{1|1}$ is a cochain map, hence that it induces a map on cohomology

$$(2.6.25) \quad H\mathcal{F}^*(\sigma_0, \sigma_1) \otimes_\Lambda H\mathcal{L}_L^*(\sigma_0) \rightarrow H\mathcal{L}_L^*(\sigma_1),$$

which makes the groups $H\mathcal{L}_L$ into a left module over the category $H\mathcal{F}$. The last statement requires the construction of a homotopy corresponding, in the usual language of A_∞ modules to the operation $\mu_{\mathcal{L}_L}^{2|1}$, and hence will be subsumed by later constructions.

In the case $\Upsilon = (\sigma_{-2}, \sigma_{-1}, L)$, elements of the moduli space $\overline{\mathcal{TR}}(x_{e_{\text{inp}}}; x_{e_{\text{out}}}, x_e)$ give rise to a map

$$(2.6.26) \quad \text{Hom}_\Lambda^c(U_{\sigma_{-1}, x_{e_{\text{inp}}}}^{P_{\sigma_{-1}}}, \Upsilon) \otimes_\Lambda \text{Hom}_\Lambda^c(U_{\sigma_{-2}, x_e}^{P_{\sigma_{-2}}}, U_{\sigma_{-1}, x_e}^{P_{\sigma_{-1}}}) \rightarrow \text{Hom}_\Lambda^c(U_{\sigma_{-2}, x_{\text{out}}}^{P_{\sigma_{-2}}}, \Upsilon).$$

The sum of maps associated to this moduli space defines the cochain-level product

$$(2.6.27) \quad \mu_{\mathcal{R}_L}^{1|1} : \mathcal{R}_L^*(\sigma_{-1}) \otimes_\Lambda \mathcal{F}^*(\sigma_{-2}, \sigma_{-1}) \rightarrow \mathcal{R}_L^*(\sigma_{-2}).$$

which induces the structure map of the right module $H\mathcal{R}_L$ at the cohomological level.

2.6.5. Comparison of local and global modules. Recall that we denote by $H\mathcal{F}_\sigma$ the subcategory of $H\mathcal{F}$ with objects the elements $\tau \in \Sigma$ such that $P_\tau \cap P_\sigma \neq \emptyset$. We have a faithful embedding $H\mathcal{F}_\sigma \rightarrow H\mathcal{P}_\sigma$. In this section, we prove that the pullback of $H\mathcal{L}_{L,\sigma}$ under this functor is naturally isomorphic to the restriction of $H\mathcal{L}_L$.

We could define such a map using previously constructed moduli spaces, but it is convenient for later purposes to have a variant, in which we replace the boundary marked point by the data of a map to the strip: we denote by $\overline{\mathcal{R}}_{2,1}$ the moduli space of degree 1 stable maps to the disc D^2 whose domain is a pre-stable disc with two boundary marked points, one of which is marked as incoming and required to map to $+1$, and the other as outgoing, which is required to map to -1 .

It is straightforward to see that $\overline{\mathcal{R}}_{2,1}$ is a single point: the assumption that the map is stable forces the domain to consist of a single component, which is identified with the target D^2 by the given map. We find it convenient to label such a domain with an extra interior marked point, corresponding to the inverse image of the origin under this map, since the position of this marked point determines the map.

Remark 2.49. We could describe the construction of this section in terms of maps whose domain is D^2 , specifying that the data we use is not invariant under reparametrisation. The convenience of the additional formalism (which we called *popsicles* in [1]) lies in its ability to efficiently handle additional boundary marked points as discussed in Section 3.1.

Consider a basepoint $\sigma \in \Sigma$, and a pair $\Upsilon = (L, \tau)$, with $\tau \in \Sigma_\sigma$. Pick a parametrisation of the path $q_{\sigma, \tau}$ by the boundary component $\mathbb{R} \times \{1\}$ of the strip, which agrees with q_σ near the positive end and with q_τ near the other end. We equip the strip with boundary conditions given by the corresponding path of fibres, and by the Lagrangian L along the boundary $\mathbb{R} \times \{0\}$.

Choosing a family of almost complex structures as before, we obtain a moduli space $\overline{\mathcal{R}}(\Upsilon)$. Choosing a homotopy between the sections ι_σ and ι_τ , the count of rigid elements of this space defines a map

$$(2.6.28) \quad \mathcal{L}_{L,\sigma}(P_\tau) := CF^*(L, (\sigma, P_\tau)) \rightarrow CF^*(L, (\tau, P_\tau)) := \mathcal{L}_L(\tau),$$

which is a quasi-isomorphism. Swapping the roles of the boundary conditions at 0 and 1, we obtain

$$(2.6.29) \quad \mathcal{R}_{L,\sigma}(P_\tau) := CF^*((\sigma, P_\tau), L) \rightarrow CF^*((\tau, P_\tau), L) := \mathcal{R}_L(\tau).$$

These constructions give rise to maps of cohomological modules by the use of mixed moduli spaces as in previous sections.

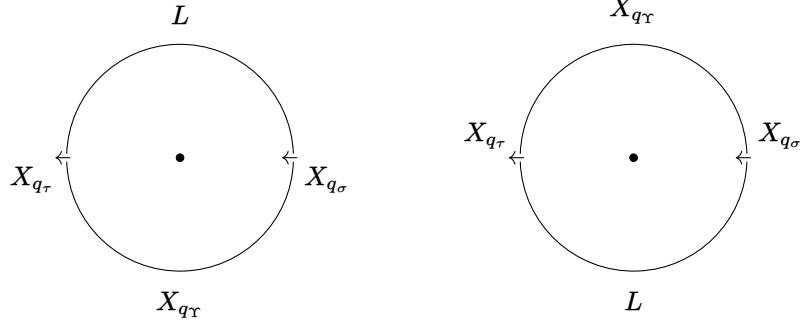


FIGURE 5. The boundary conditions of elements of $\overline{\mathcal{R}}(\Upsilon)$ for $\Upsilon = (\sigma, \tau, L)$ and $\Upsilon = (L, \sigma, \tau)$. The interior marked point corresponds to the origin, and indicates the fact that the data we choose on the strip is not invariant under translation.

2.6.6. *Floer cohomology of Lagrangians, and global modules.* We now proceed to construct the top horizontal and right vertical maps in Diagram (2.1.12); this is a minor variant of the maps considered in [6]. Given $L \neq L' \in \mathcal{A}$, and $\sigma \in \Sigma$, we consider the triple $\Upsilon = (L, L', \sigma)$ or (L, σ, L') . Let S be a disc with 3 punctures. The triple Υ induces a unique labelling of the boundary components of the complement of the punctures which respects the ordering counterclockwise around the boundary; we write Υ_e for the pair associated to each end e . We pick a strip-like end on each puncture, which is positive if the corresponding edge is incoming, and negative otherwise.

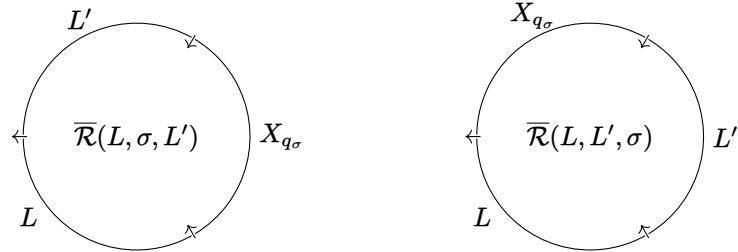


FIGURE 6. The boundary conditions for elements of the moduli spaces $\overline{\mathcal{R}}(L, L', \sigma)$ and $\overline{\mathcal{R}}(L, \sigma, L')$.

We define

$$(2.6.30) \quad \mathcal{J}(\Upsilon) \subset C^\infty(S, \mathcal{J})$$

to be the space of almost complex structures parametrised by S which respectively agree with J_L and $J_{L'}$ along the boundary components labelled by these Lagrangians, and whose restriction to each end e agrees, under a choice of strip-like ends, with the families $J(\Upsilon_e)$ assigned earlier to pairs of labels containing an element of Σ , and given by a fixed regular choice of path from J_L to $J_{L'}$ for $\Upsilon_e = (L, L')$.

If we impose constant Lagrangian conditions X_{q_σ} along the remaining boundary component, the choice of an element $J(\Upsilon)$ of $\mathcal{J}(\Upsilon)$ determines a moduli space $\overline{\mathcal{R}}(\Upsilon)$ of holomorphic

triangles equipped with an evaluation map

$$(2.6.31) \quad \overline{\mathcal{R}}(\Upsilon) \rightarrow \prod_e \text{Crit}(\Upsilon_e),$$

where the product is taken over all the ends. Assuming that the family of almost complex structure is generic, and after imposing a further constraint on the diameter of \mathcal{N}_σ in terms of the corresponding reverse isoperimetric inequality, the count of rigid elements of $\overline{\mathcal{R}}(\Upsilon)$ (together with parallel transport) induces maps

$$(2.6.32) \quad \mathcal{L}_{L'}(\sigma) \otimes_{\Lambda} CF^*(L, L') \rightarrow \mathcal{L}_L(\sigma)$$

$$(2.6.33) \quad \mathcal{R}_{L'}(\sigma) \otimes_{\Lambda} \mathcal{L}_L(\sigma) \rightarrow CF^*(L, L'),$$

depending on whether $\Upsilon = (L, L', \sigma)$ or (L, σ, L') .

Dualising Equation (2.6.32) gives rise to a map

$$(2.6.34) \quad CF^*(L, L') \rightarrow \text{Hom}_{\Lambda}^c(\mathcal{L}_{L'}^*(\sigma), \mathcal{L}_L^*(\sigma)).$$

Using a mixed moduli space as in Section 2.6.3, and a family of moving Lagrangian boundary conditions, we find that this map commutes with the action of morphisms groups in $H\mathcal{F}$, and obtain the map from Equation (2.1.4).

We also obtain a commutative diagram

$$(2.6.35) \quad \begin{array}{ccc} H\mathcal{R}_{L'}^*(\sigma) \otimes_{\Lambda} H\mathcal{F}^*(\sigma_-, \sigma) \otimes_{\Lambda} H\mathcal{L}_L^*(\sigma_-) & \longrightarrow & H\mathcal{R}_{L'}^*(\sigma) \otimes_{\Lambda} H\mathcal{L}_L^*(\sigma) \\ \downarrow & & \downarrow \\ H\mathcal{R}_{L'}^*(\sigma_-) \otimes_{\Lambda} H\mathcal{L}_L^*(\sigma_-) & \longrightarrow & HF^*(L, L'), \end{array}$$

which proves that we have defined a map from the tensor product of left and right modules over $H\mathcal{F}$ to the Floer cohomology of the pair (L, L') .

2.7. The perturbed diagonal. The constructions of Section 2.6 generalise those of Section 2.4 to the global category. In this section, we extend the results of Section 2.5: the key point is to construct a bimodule over $H\mathcal{F}$ whose restriction to $H\mathcal{F}_\sigma$, for each $\sigma \in \Sigma$, agrees with the pullback of the diagonal bimodule of $H\text{Po}_\sigma$.

The construction will be implemented via a perturbed Lagrangian Floer cohomology group of fibres. To this end, we fix a Hamiltonian diffeomorphism $\phi: X \rightarrow X$ which is generic in the sense that

$$(2.7.1) \quad \text{for all pairs } (q_-, q) \in Q^2, X_{q_-} \cap \phi X_q \text{ is discrete,}$$

which can be achieved by jet transversality. We also fix a neighbourhood $\nu(X \cap \phi X)$ in $Q^2 \times X$ of the set

$$(2.7.2) \quad \coprod_{(q_-, q) \in Q^2} X_{q_-} \cap \phi X_q \subset Q^2 \times X.$$

We denote by $\nu_X(X_{q_-} \cap \phi X_q)$ the inverse image over $(q_-, q) \in Q^2$, and require that this subset be sufficiently small so that

$$(2.7.3) \quad \text{for all pairs } (q_-, q), \text{ the intersections of } \nu_X(X_{q_-} \cap \phi X_q) \text{ with } X_{q_-} \text{ and } \phi X_q \text{ are respectively contained in disjoint unions of closed balls in } X_{q_-} \text{ and } \phi X_q.$$

Given a path $J(q_-, q)$ of almost complex structures on X which are constant away from a fixed compact subset of the interior of $\nu_X(X_{q_-} \cap \phi X_q)$, Lemma A.1 provides a reverse isoperimetric constant C for $J(q_-, q)$ holomorphic strips with boundary conditions X_{q_-} and ϕX_q , with the length measured in the quotient of these Lagrangians by the equivalence

relation which identifies two points in the same component of the intersection with $\nu_X(X_{q_-} \cap \phi X_q)$, and which is uniform in the choice of points in Q^2 , because we have assumed that Q is compact. As before, the reverse isoperimetric inequality is independent of the restriction of $J(q_-, q)$ to the chosen compact subset of the interior of $\nu_X(X_{q_-} \cap \phi X_q)$, and of the intersection of the Lagrangians with this region. This will allow us to pick perturbations in order to achieve transversality.

With the isoperimetric constant from the above paragraph, we impose an additional constraint on the cover by requiring that

$$(2.7.4) \quad \text{diam } \mathcal{N}_\sigma < 1/4C,$$

where we note that the constant C now depends additionally on the Hamiltonian diffeomorphism ϕ , since it incorporates a reverse isoperimetric constraint for the pair of Lagrangians X_{q_-} and ϕX_q .

Remark 2.50. We note that we do not need to impose any assumption on ϕ other than genericity, since the existence of a uniform reverse isoperimetric constant is a very general fact about families of Lagrangians. Indeed, the key geometric quantity controlling the reverse isoperimetric constant is the distance between the two Lagrangians. This implies that any family of pairs of Lagrangians parametrised by a compact manifold admits a C^2 -small perturbation for which there is a uniform reverse isoperimetric constant. In particular, we need not require that ϕ be a small perturbation of the identity.

2.7.1. Perturbed Floer cohomology for pairs of polytopes. Let Σ^ϕ denote the set $\Sigma \times \{\phi\}$; in other words, we think of Σ^ϕ as another copy of Σ , whose elements we shall denote σ^ϕ as a shorthand for (σ, ϕ) . This set parametrises the same cover of Q as Σ , but we associate to it the Lagrangian fibration obtained by composing the projection $X \rightarrow Q$ with ϕ . Given a pair $\Upsilon = (\sigma_-, \sigma^\phi) \in \Sigma \times \Sigma^\phi$, we introduce the notation

$$(2.7.5) \quad \nu_X \Upsilon := \nu_X(X_{q_{\sigma_-}} \cap \phi X_{q_\sigma}).$$

We shall impose the following additional condition on the basepoints q_σ :

$$(2.7.6) \quad \text{for each pair } (\sigma_-, \sigma) \text{ of elements of } \Sigma^2, \text{ the Lagrangians } X_{q_\sigma} \text{ and } \phi X_{q_\sigma} \text{ are transverse.}$$

Having assumed that each polytope P_σ has non-empty interior, Sard's theorem implies that this condition is achieved for generic choices of basepoints. We write

$$(2.7.7) \quad \text{Crit}(\Upsilon) := X_{q_{\sigma_-}} \cap \phi X_{q_\sigma}$$

for the set of intersection points between these Lagrangians.

By assumption, this is a discrete space, and, given the choice of Pin structures from Section 2.2.3, we obtain a 1-dimensional free abelian group δ_x associated to each element x of $\text{Crit}(\Upsilon)$. Given a choice $J(\Upsilon)$ of an almost complex structure on X , which agrees with J away from $\nu_X \Upsilon$, we obtain a moduli space

$$(2.7.8) \quad \overline{\mathcal{R}}(\Upsilon) \rightarrow \text{Crit}(\Upsilon)^2$$

of stable finite energy $J(\Upsilon)$ -holomorphic strips with boundary condition $X_{q_{\sigma_-}}$ along $t = 0$ and ϕX_{q_σ} along $t = 1$, equipped with its natural evaluation map at the ends to intersection points of the boundary Lagrangians. Choosing the almost complex structure generically, the fibre $\overline{\mathcal{R}}(x_0, x_1)$ over a pair (x_0, x_1) has the expected dimension, and whenever the moduli space is rigid, we obtain a map δ_u of orientation lines for each $u \in \overline{\mathcal{R}}(x_0, x_1)$.

Given a curve $u \in \overline{\mathcal{R}}(x_0, x_1)$, let $\partial_\Sigma u$ and $\partial_\Sigma^\phi u$ denote the boundary components labelled by σ_- and σ^ϕ . Parallel transport along these boundaries defines maps

$$(2.7.9) \quad z^{[\partial_\Sigma^\phi u]}: U_{\sigma, x_1}^{P_\sigma} \rightarrow U_{\sigma, x_0}^{P_\sigma}$$

$$(2.7.10) \quad z^{[\partial_\Sigma u]}: U_{\sigma_-, x_0}^{P_{\sigma_-}} \rightarrow U_{\sigma_-, x_1}^{P_{\sigma_-}}.$$

We define the Floer complex

$$(2.7.11) \quad CF^*((\sigma_-, P_{\sigma_-}), (\sigma^\phi, P_\sigma)) := \bigoplus_{x \in \text{Crit}(\Upsilon)} \text{Hom}_\Lambda^c(U_{\sigma_-, x}^{P_{\sigma_-}}, U_{\sigma, x}^{P_\sigma}) \otimes \delta_x,$$

with differential $\mu_\phi^{0|1|0}$ given by

$$(2.7.12) \quad \mu_\phi^{0|1|0} |\psi \otimes \delta_{x_0, x_1}| := \sum_{u \in \mathcal{R}(x_0, x_1)} (-1)^{\deg(x_1)+1} T^{E(u)} z^{[\partial_\Sigma^\phi u]} \cdot \psi \cdot z^{[\partial_\Sigma u]} \otimes \delta_u.$$

Condition (2.7.4) implies that P_σ and P_{σ_-} are contained in the ball of radius $1/4C$ about q_σ and q_{σ_-} . As in Corollary 2.13, we conclude that this differential is well-defined and continuous with respect to the topology on these Floer complexes.

To define a bimodule over HF , we set

$$(2.7.13) \quad \overline{\Delta}(\sigma_-, \sigma) := CF^*((\sigma_-, P_{\sigma_-}), (\sigma^\phi, P_\sigma))$$

and denote the cohomology by $H\overline{\Delta}(\sigma_-, \sigma)$.

As it is completely analogous to the discussion from Section 2.6.4, we omit for the moment the details of the construction of the bimodule structure maps (they will appear in greater generality when we discuss the A_∞ refinement).

2.7.2. Map from the diagonal bimodule. The main result of this section is the existence of a map from the diagonal bimodule on HF to the perturbed diagonal. It will require a further constraint on the cover, to be specified in Condition (2.7.19) below:

Lemma 2.51. *If Σ labels a sufficiently fine cover, then for each triple (σ, τ, τ_-) in Σ such that τ and τ_- both lie in Σ_σ , there is a natural quasi-isomorphism*

$$(2.7.14) \quad \text{Po}_\sigma(P_{\tau_-}, P_\tau) \rightarrow \overline{\Delta}(\tau_-, \tau).$$

In particular, the restriction of $H\overline{\Delta}$ to HF_σ is isomorphic to the pullback of the diagonal bimodule of HPo_σ .

Fix a basepoint $\sigma \in \Sigma$, and consider elements τ and τ_- of Σ_σ . Let $\Upsilon = \{\tau_-, \tau^\phi\}$, and denote by $\overline{\mathcal{R}}_\Upsilon$ a copy of $\overline{\mathcal{R}}_{2,1}$, which we think of as parametrising a disc S with a puncture (we denote the corresponding end by e), and a boundary marked point, together with a stabilising map to the strip D^2 mapping the puncture to -1 and the marked point to $+1$. The basic idea is that the puncture corresponds to the perturbed bimodule, the marked point at 1 to the diagonal of HPo_σ , and the interior marked point to the fact that we shall interpolate between the data defining these two bimodules. We denote the intervals from 1 to -1 along the upper and lower part of the boundary by $\partial_\Sigma^\phi S$ and $\partial_\Sigma S$.

We begin by strengthening the conditions imposed at the beginning of Section 2.7 by imposing the relevant isoperimetric constraint on an additional moduli space. We pick

maps from ∂_Σ^ϕ and $\partial_\Sigma S$ to the interval $[0, 1]$, which respectively map neighbourhoods of the points 1 and -1 to the endpoints 0 and 1 of the interval. We obtain paths $q_{\tau_-, \sigma}$ and $q_{\tau, \sigma}$

$$(2.7.15) \quad q_{\underline{\Upsilon}}: \partial_\Sigma S \rightarrow Q$$

$$(2.7.16) \quad q_{\underline{\Upsilon}}^\phi: \partial_\Sigma^\phi S \rightarrow Q$$

which agree with q_σ near the marked point 1, and have value (q_{τ_-}, q_τ) near the end. We also obtain a map

$$(2.7.17) \quad \phi_{\underline{\Upsilon}}: \partial_\Sigma^\phi S \rightarrow \text{Ham}(X),$$

which is the identity near 1, and agrees with ϕ near the end.

Given this data, we form Lagrangian boundary conditions given by the path $X_{q_{\underline{\Upsilon}}}$ on $\partial_\Sigma S$ and $\phi_{\underline{\Upsilon}} X_{q_{\underline{\Upsilon}}^\phi}$ on $\partial_\Sigma^\phi S$ (see Figure 7); by construction, these boundary conditions agree with X_{q_σ} at the marked point. Let X_S to be the fibre bundle of Lagrangian boundary conditions over the quotient of the boundary of S by the two components of its intersection with a fixed neighbourhood $\nu_{S\ell}$ of the negative end over which these paths are locally constant. Define X_S/\sim to be the quotient by the equivalence relation that collapses the intersections of $\nu_X \Upsilon_e$ with the fibres $X_{q_{\tau_-}}$ and ϕX_{q_τ} over the endpoints. We can associate to each map u from S to X with such boundary conditions a map

$$(2.7.18) \quad \partial u: \partial S \rightarrow X_S/\sim.$$

We equip this quotient with a metric so that, for a loop in X_S , the norm of the homology class in $H_1(X_{q_\sigma}, \mathbb{Z})$ is bounded by the length in the quotient.

Let $J(\underline{\Upsilon})$ be a family of almost complex structures on the strip which agree with J outside of $\nu_X \Upsilon_e$ and agree with $J(\Upsilon_e)$ near the end e .

Lemma 2.52. *There is a constant C , depending on the collection \mathcal{A} of Lagrangians and the choices of almost complex structures, but independent of the choice of cover as long as it is sufficiently fine, so that for each pseudoholomorphic map u with the above boundary conditions, the length of the path $\partial u/\sim$ is bounded by the product of the energy with C .*

Proof. This is a direct consequence of Lemma A.3, for an appropriate parametrised problem that restricts to the above moduli space for each sufficiently fine cover indexed by Σ , and for each choice of basepoints q_σ , q_τ , and q_{τ_-} associated to an element σ of Σ and a pair of elements τ and τ_- of Σ_σ .

We can use the same strategy described before imposing the constraint in Equation (2.7.4): once a Hamiltonian isotopy Φ from the identity to a diffeomorphism ϕ satisfying Condition (2.7.1) is fixed, we may consider the family of boundary conditions on S parametrised by a triple $\vec{q} = (q, q', q'_-) \in Q^3$ such that the distance from q to q' and q'_- is less than or equal to 1: the boundary conditions are given by the unique short affine linear path from q to q' along one boundary segment, and by applying Φ to the path from q to q'_- along the other. As above, we obtain an equivalence relation σ on the boundary conditions $X_{S, \vec{q}}$, by collapsing the intersection of $\nu_X(X_{q'_-} \cap \phi X_{q'})$ with the boundary conditions over the endpoints.

According to Lemma A.3, there is a constant C independent of u , such that, whenever $u: S \rightarrow X$ is a J -holomorphic curve with any of the above boundary conditions, the length of the path $\partial u/\sim$ is bounded by the product of the energy with C , up to an additive constant which is also independent of u . Moreover, the same reverse isoperimetric constant holds after perturbing the almost complex structure, along the end, as long as the perturbation takes place in $\nu_X(X_{q'_-} \cap \phi X_{q'})$. \square

Given the above result, we may require our choice of cover to satisfy the conditions that

$$(2.7.19) \quad \text{diam } \mathcal{N}_\sigma < 1/8C$$

for the constant C in Lemma 2.52, in addition to all previous imposed constraints. This means in particular that C depends on the collection \mathcal{A} of Lagrangian, on the fibration $X \rightarrow Q$, on the Hamiltonian diffeomorphism ϕ , as well as on the path Φ (in addition to the auxiliary choices of almost complex structures).

Let $\overline{\mathcal{R}}(\underline{\Upsilon})$ denote the moduli space of finite energy stable $J(\underline{\Upsilon})$ holomorphic strips, with moving Lagrangian boundary conditions given by the paths $X_{q_{\underline{\Upsilon}}}$ and $\Phi_{\underline{\Upsilon}} X_{q_{\underline{\Upsilon}}}^\phi$. We have a natural evaluation map at ± 1 :

$$(2.7.20) \quad \overline{\mathcal{R}}(\underline{\Upsilon}) \rightarrow \text{Crit}(\underline{\Upsilon}) \times X_{q_\sigma}.$$

By construction, the reverse isoperimetric inequality from Equation (2.7.19) applies to this moduli space.

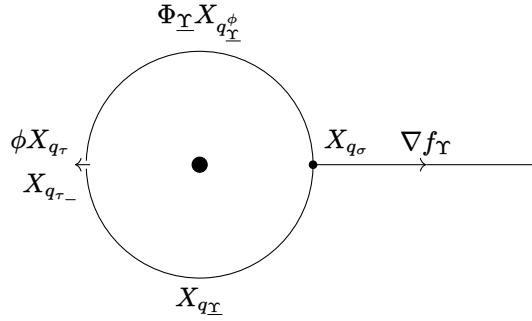


FIGURE 7. A representation of an element of the moduli space $\overline{\mathcal{TR}}(\underline{\Upsilon})$, for $\underline{\Upsilon} = \{\tau_-, \tau^\phi\}$, and basepoint σ .

We now consider the fibre product

$$(2.7.21) \quad \overline{\mathcal{TR}}(\underline{\Upsilon}) := \overline{\mathcal{R}}(\underline{\Upsilon}) \times_{X_{q_\sigma}} \overline{\mathcal{T}}_+(\sigma, \sigma)$$

where $\overline{\mathcal{T}}_+(\sigma, \sigma)$ is the moduli space of perturbed half-gradient flow lines introduced in Section 2.2.4. We have a natural evaluation map

$$(2.7.22) \quad \overline{\mathcal{TR}}(\underline{\Upsilon}) \rightarrow \text{Crit}(\tau_-, \tau^\phi) \times \text{Crit}(\sigma, \sigma),$$

whose fibre at a pair (x_0, x_1) we denote

$$(2.7.23) \quad \overline{\mathcal{TR}}(x_0, x_1).$$

Choosing paths connecting the intersections of X_{q_σ} with the sections associated to τ , τ_- , and σ yields isomorphisms of local systems $U_{\tau_-}^{P_{\tau_-}} \cong U_\sigma^{P_{\tau_-}}$ and $U_\tau^{P_\tau} \cong U_\sigma^{P_\tau}$. Using this together with parallel transport along the boundary of an element u of this moduli space yields a map

$$(2.7.24) \quad \text{Hom}_\Lambda^c(U_{\sigma, x_1}^{P_{\tau_-}}, U_{\sigma, x_1}^{P_\tau}) \rightarrow \text{Hom}_\Lambda^c(U_{\tau_-, x_0}^{P_{\tau_-}}, U_{\tau, x_0}^{P_\tau})$$

$$(2.7.25) \quad \phi \mapsto z^{[\partial_\Sigma u]} \cdot \psi \cdot z^{[\partial_\Sigma^\phi u]},$$

where $\partial_\Sigma u$ and $\partial_\Sigma^\phi u$ are the restrictions of u to the boundary components $\partial_\Sigma S$ and $\partial_\Sigma^\phi S$.

For generic choices of the almost complex structure and the Morse perturbation, $\overline{\mathcal{TR}}(x_0, x_1)$ is a manifold with boundary of dimension $\deg(x_0) - \deg(x_1)$, and rigid elements of this moduli space induce an isomorphism δ_u of orientation lines. Tensoring these with the parallel transport maps, we define

$$(2.7.26) \quad \kappa: \text{Po}_\sigma(P_{\tau_-}, P_\tau) \rightarrow \overline{\Delta}(\tau_-, \tau)$$

to be given by the sum

$$(2.7.27) \quad \kappa|\phi \otimes \delta_{x_1} := - \sum_{u \in \mathcal{TR}_q^1(x_0, x_1)} T^{E(u)} z^{[\partial u]} \cdot \psi \cdot z^{[\partial_- u]} \otimes \delta_u.$$

The reverse isoperimetric inequality for moduli problems with moving Lagrangian boundary conditions, together with Condition (2.7.19), implies that this map is well-defined and continuous.

Remark 2.53. The minus sign above accounts for the fact that we use reduced gradings to define the differential on $\text{Po}_\sigma(P_{\tau_-}, P_\tau)$, but unreduced gradings for $\overline{\Delta}(\tau_-, \tau)$.

To prove that this map is a cochain equivalence, we construct a two-sided homotopy inverse: Given a sequence $\Upsilon = (\tau_-, \tau^\phi)$, we denote by $\overline{\mathcal{R}}_{\Upsilon}$ a copy of $\overline{\mathcal{R}}_{2,1}$ which we now think of as equipped with a marked point at -1 and a puncture at 1 . By reflecting the data chosen above, we obtain moving Lagrangian boundary conditions which agree with X_{q_σ} near -1 and with the pair $(X_{q_{\tau_-}}, \phi X_{q_\tau})$, near the puncture. We obtain a moduli space $\overline{\mathcal{R}}(\Upsilon)$ with an evaluation map to $X_{q_\sigma} \times \text{Crit}(\Upsilon_e)$. Taking the fibre product over X_{q_σ} with the moduli space $\mathcal{T}_-(\sigma, \sigma)$ of perturbed flow lines, we obtain the moduli space $\overline{\mathcal{TR}}(\Upsilon)$, equipped with an evaluation map to the product of $\text{Crit}(\sigma, \sigma)$ at the negative end, and $\text{Crit}(\tau_-, \tau^\phi)$ at the positive end. Counts of rigid elements of this moduli space define a map

$$(2.7.28) \quad \overline{\Delta}(\tau_-, \tau) \rightarrow \text{Po}_\sigma(P_{\tau_-}, P_\tau).$$

We omit the proof that this is a left and right homotopy inverse to Equation (2.7.26), noting only that this requires a further reverse isoperimetric constraint on the cover. Passing to cohomology proves Lemma 2.51.

2.7.3. Computing the bimodule: distant polytopes. In this section, we compute the group $H\overline{\Delta}(\tau_-, \tau)$ associated to the Hamiltonian diffeomorphism ϕ chosen at the beginning of Section 2.7. We fix as well the Hamiltonian isotopy Φ from the identity to ϕ chosen in the previous section:

Lemma 2.54. *If the cover Σ is sufficiently fine, and $\text{diam } P_\sigma \ll \text{diam } \mathcal{N}_\sigma$ for all elements $\sigma \in \Sigma$, then the group $H\overline{\Delta}(\tau_-, \tau)$ vanishes whenever P_τ and P_{τ_-} are not contained in a common chart \mathcal{N}_σ .*

Sketch of proof: The difference in scale between P_σ and \mathcal{N}_σ will be set by a reverse isoperimetric condition. Consider the space of discs with punctures at ± 1 , and two interior marked points, which lie on the real axis and are equidistant from the origin. This moduli space has two boundaries, corresponding to the two marked points lying at the origin, and to the breaking of the domain into two components. We equip the curves over the moduli space with Lagrangian boundary conditions as follows (see Figure 8): for each of these curves, we equip the boundary segment from -1 to $+1$, going along the lower semi-circle, with the constant boundary condition $X_{q_{\tau_-}}$. Along the other end, we take a family of paths from ϕX_{q_τ} to itself, which is the constant path when the two marked points agree, and with the concatenation of the path ΦX_{q_τ} and its inverse when the disc breaks into two components.

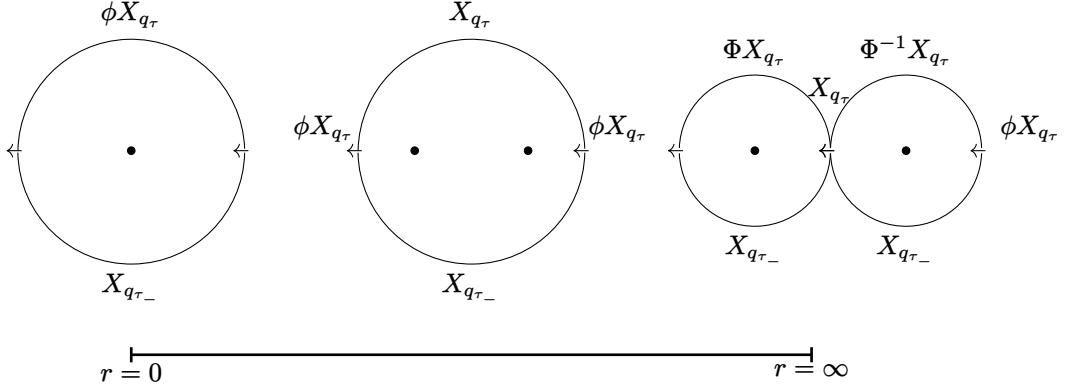


FIGURE 8. The boundary conditions, near from the ends, of the moduli space which induces a null-homotopy of Floer complexes with $\Upsilon = (\tau_-, \tau)$. The notation Φ^{-1} denotes the path Φ composed with the orientation reversing involution $t \rightarrow 1 - t$ of the parametrising interval.

We now pick almost complex structures for this family of Lagrangian boundary conditions: for the constant boundary conditions, we use the translation-invariant almost complex structure $J(\tau_-, \tau^\phi)$ chosen in Section 2.7.1, while near the broken curve, we assume that the almost complex structure is obtained by gluing, and agrees with $J(\tau_-, \tau^\phi)$ at the ends. Applying Lemma A.4, we obtain a reverse isoperimetric constant for holomorphic curves in this family.

Using the fact that X_{q_τ} and $X_{q_{\tau_-}}$ are disjoint, we find that the composition of maps associated to the broken curve vanishes, since it corresponds to a map factoring through a Floer complex that is the 0 group by definition. Choosing the diameter of the elements of the cover $\{P_\sigma\}_{\sigma \in \Sigma}$ to be much smaller than that of the cover $\{\mathcal{N}_\sigma\}_{\sigma \in \Sigma}$, we obtain a null homotopy for the identity map on $\overline{\Delta}(\tau_-, \tau)$, implying Lemma 2.54. \square

2.8. Global bimodule maps. Given a Lagrangian $L \in \mathcal{A}$ we have defined left and right modules over $H\mathcal{F}$. In Section 2.8.1, we construct a map of bimodules

$$(2.8.1) \quad H\mathcal{L}_L \otimes_{\Lambda} H\mathcal{R}_L \rightarrow \overline{\Delta}$$

by using a moduli space of holomorphic triangles, with one moving Lagrangian boundary condition. We also construct a map of left modules

$$(2.8.2) \quad \overline{\Delta} \otimes_{H\mathcal{F}} H\mathcal{L}_L \rightarrow H\mathcal{L}_L.$$

These are the missing ingredients in the statement of Proposition 2.2, which we proceed to prove in Section 2.8.2.

In order to compare these constructions with the local ones, we recall that Lemma 2.8 shows that the computation of tensor products with $\overline{\Delta}$ and morphisms of left module to $\overline{\Delta}$ are local. The main results of this section are summarised by the following, for which we recall that we denote by j the embedding of the category $H\mathcal{F}_\sigma$ in HPo_σ , as in Equation (2.1.20):

Proposition 2.55. *For each $\sigma \in \Sigma$, there are commutative diagrams of bimodules over $H\mathcal{F}_\sigma$*

$$(2.8.3) \quad \begin{array}{ccc} j^* H\mathcal{L}_{L,\sigma} \otimes_{\Lambda} j^* H\mathcal{R}_{L,\sigma} & \longrightarrow & j^* \Delta_{H\mathcal{P}_\sigma} \\ \downarrow & & \downarrow \\ H\mathcal{L}_L \otimes_{\Lambda} H\mathcal{R}_L & \longrightarrow & H\overline{\Delta} \end{array}$$

and of left $H\mathcal{F}_\sigma$ -modules:

$$(2.8.4) \quad \begin{array}{ccc} j^* H\Delta_{H\mathcal{P}_\sigma} \otimes_{H\mathcal{F}_\sigma} j^* H\mathcal{L}_{L,\sigma} & \longrightarrow & j^* H\mathcal{L}_{L,\sigma} \\ \downarrow & & \downarrow \\ H\overline{\Delta} \otimes_{H\mathcal{F}_\sigma} H\mathcal{L}_L & \longrightarrow & H\mathcal{L}_L. \end{array}$$

Note that the first commutative diagram in the above Proposition is equivalent to

$$(2.8.5) \quad \begin{array}{ccc} j^* H\mathcal{R}_{L,\sigma} & \longrightarrow & \text{Hom}_{H\mathcal{F}_\sigma}(j^* H\mathcal{L}_{L,\sigma}, j^* \Delta_{H\mathcal{P}_\sigma}) \\ \downarrow & & \downarrow \\ H\mathcal{R}_L & \longrightarrow & \text{Hom}_{H\mathcal{F}_\sigma}(j^* H\mathcal{L}_{L,\sigma}, H\overline{\Delta}) \\ & & \uparrow \\ & & \text{Hom}_{H\mathcal{F}_\sigma}(H\mathcal{L}_L, H\overline{\Delta}). \end{array}$$

Our previous results establish that, whenever L meets X_{q_σ} at one point, all the vertical arrows in Diagrams (2.8.4) and (2.8.5) are isomorphisms if we restrict to the object $\sigma \in H\mathcal{F}_\sigma$. Together with Lemma 2.8, we conclude

Lemma 2.56. *If L meets X_{q_σ} at one point, there are natural isomorphisms*

$$(2.8.6) \quad \overline{\Delta}(\sigma, -) \otimes_{H\mathcal{F}} H\mathcal{L}_L \rightarrow H\mathcal{L}_L(\sigma)$$

$$(2.8.7) \quad H\mathcal{R}_L(\sigma) \rightarrow \text{Hom}_{H\mathcal{F}}(H\mathcal{L}_L, \overline{\Delta}(-, \sigma)).$$

□

At this stage, we also note that we can establish the following result using the techniques developed in Section 2.4.4:

Proof of Lemma 2.4. Choosing a generator for $H\mathcal{F}(P_\tau, P_\rho)$ and $H\mathcal{L}_L(\rho)$ as rank-1 modules over Γ^{P_ρ} , we can write morphisms from $H\mathcal{L}_{L'}$ to any module M as the kernel of the map

$$(2.8.8) \quad \bigoplus_{\sigma_0 \in \Sigma} M(\sigma_0) \rightarrow \bigoplus_{\sigma_0 < \sigma_1 \in \Sigma} M(\sigma_1).$$

This implies that $\text{Hom}_{H\mathcal{F}}(H\mathcal{L}_{L'}, -)$ is right exact. Identifying $\overline{\Delta}_{H\mathcal{F}} \otimes_{H\mathcal{F}} H\mathcal{L}_L$, as a left module, with the cokernel of

$$(2.8.9) \quad \bigoplus_{\rho_0 < \rho_1 \in \Sigma} \overline{\Delta}_{H\mathcal{F}}(\rho_0, -) \otimes_{\Gamma^{\rho_0}} \Gamma^{\rho_1} \rightarrow \bigoplus_{\rho_0 \in \Sigma} \overline{\Delta}_{H\mathcal{F}}(\rho_0, -)$$

we obtain the desired result.

□

Remark 2.57. Combining the results of Section 4.3 with the discussion of Corollary 4.2, gives a less computational proof of the A_∞ refinement of this statement. The reader is invited to adapt that proof to the cohomological level, or alternatively adapt this computational proof to the A_∞ refinement.

2.8.1. Construction of the maps. Let S be a thrice-punctured disc, equipped with a fixed neighbourhood $\nu_S e$ of each end e . Let Υ be a set of labels for the boundary components which is either (i) $(\sigma_-, L, \sigma^\phi)$, or (ii) $(L, \sigma_-, \sigma^\phi)$, with $L \in \mathcal{A}$, and $\sigma, \sigma_- \in \Sigma$. Pick a map

$$(2.8.10) \quad \Phi_\Upsilon: \partial S \rightarrow \text{Ham}(X)$$

which agrees with the identity on the components labelled σ_- and L , and with a parametrisation of the path Φ along on the component labelled by σ^ϕ (specifically, we require that it interpolate between the identity along the end corresponding to the label L and the map ϕ on the end labelled by σ_-). Applying these Hamiltonians to the Lagrangians L , X_{q_σ} , and $X_{q_{\sigma_-}}$ yields moving Lagrangian boundary conditions on S (see Figure 9).

In order to straightforwardly appeal to the reverse isoperimetric inequality, we require that the restriction of the path Φ_Υ to each end $\nu_S e$ agree away from a fixed compact subset of the interior of $\nu_S \Upsilon_e$ with (i) ϕ along the boundary component labelled by σ^ϕ and (ii) the identity on every other component. This restriction is exactly the same as the one made in the choice of perturbations used to define the Floer cohomology groups of fibres with L , or in Section 2.7 for pairs for fibres. We choose a family of almost complex structures on X parametrised by S , whose restriction to $\nu_S e$ is obtained from $J(\Upsilon_e)$ by a choice of strip-like ends. We denote by $\overline{\mathcal{R}}(\Upsilon)$ the corresponding moduli space of stable holomorphic discs.

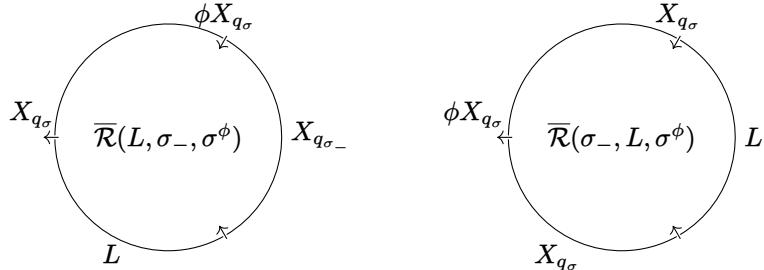


FIGURE 9. The boundary conditions for elements of the moduli spaces $\overline{\mathcal{R}}(\sigma_-, \sigma^\phi, L)$ and $\overline{\mathcal{R}}(\sigma_-, L, \sigma^\phi)$.

We write X_S for the total space of the boundary conditions labelled σ and σ^ϕ , modulo the relation which identifies each component of the inverse image of $\nu_S e$ to the corresponding fibre; this is a fibre bundle over a pair of closed intervals. We consider the equivalence relation \sim , which is the identity in the interior, and which is given, at a boundary Lagrangian labelled by an end e , by collapsing the components of the intersection with $\nu_X \Upsilon_e$ to a point. We equip this quotient with a metric such that norm of the homology class of a loop is bounded by the length of the projection.

At this stage, we need to impose a constraint on the diameter of the polytopes labelled by the elements of Σ . As before, the key point is that the moduli space $\overline{\mathcal{R}}(\Upsilon)$ can be described as the fibre over the pair $(q_{\sigma_-}, q_{\sigma_-})$ of a moduli space parametrised by a neighbourhood of the diagonal in $Q \times Q$, consisting of pairs of points whose distance is smaller than the convexity radius. Lemma A.3 provides a reverse isoperimetric constant C for elements of

this parametrised moduli space, which allows us to impose the condition that the our chosen cover of Q satisfy the constraint that

$$(2.8.11) \quad \text{diam } \mathcal{N}_\sigma < \frac{1}{8C}$$

for each $\sigma \in \Sigma$.

This condition ensures the convergence of the count of rigid elements of the moduli space $\overline{\mathcal{R}}(\Upsilon)$, yielding maps

$$(2.8.12) \quad \overline{\Delta}(\sigma_-, \sigma) \otimes_\Lambda \mathcal{L}_L(\sigma_-) \rightarrow \mathcal{L}_L(\sigma)$$

$$(2.8.13) \quad \mathcal{L}_L(\sigma) \otimes_\Lambda \mathcal{R}_L(\sigma_-) \rightarrow \overline{\Delta}(\sigma_-, \sigma).$$

Passing to cohomology, we obtain the maps

$$(2.8.14) \quad H\overline{\Delta}(\sigma_-, \sigma) \otimes_\Lambda H\mathcal{L}_L(\sigma_-) \rightarrow H\mathcal{L}_L(\sigma)$$

$$(2.8.15) \quad H\mathcal{L}_L(\sigma) \otimes_\Lambda H\mathcal{R}_L(\sigma_-) \rightarrow H\overline{\Delta}(\sigma_-, \sigma).$$

2.8.2. Proof of Proposition 2.2. Consider the moduli space $\overline{\mathcal{R}}_4$ of holomorphic discs with 4 boundary punctures, one of which is distinguished as outgoing. Given a quadruple $\Upsilon = (L, \sigma_-, L', \sigma^\phi)$, with $\sigma, \sigma_- \in Q$ and $L, L' \in \mathcal{A}$, we denote by $\overline{\mathcal{R}}_\Upsilon$ a copy of $\overline{\mathcal{R}}_4$ with the corresponding boundary labels, and by

$$(2.8.16) \quad \overline{\mathcal{S}}_\Upsilon \rightarrow \overline{\mathcal{R}}_\Upsilon$$

the universal curve over this moduli space. Recall that $\overline{\mathcal{R}}_\Upsilon$ is homeomorphic to a closed interval, with boundary given by the two possible configurations of stable discs consisting of two components each of which is disc with 3 punctures. These configurations are distinguished by the fact that the node is labelled by the pair (L, L') in one case and by (σ_-, σ^ϕ) in the second.

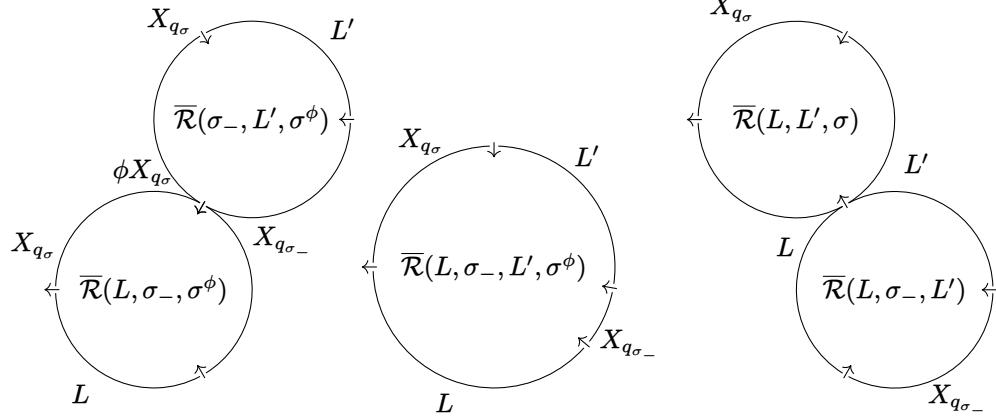


FIGURE 10. The boundary conditions for different elements of $\overline{\mathcal{R}}(L, \sigma_-, L', \sigma^\phi)$.

We choose families of strip-like ends for all surfaces over $\overline{\mathcal{R}}_\Upsilon$, which we assume are compatible with those made in Sections 2.6.6, 2.7.1, and 2.8.1 over the boundary strata. The choices made in these sections also determine moving Lagrangian boundary conditions on the fibres over the boundary. We extend these choices to arbitrary surfaces representing elements of $\overline{\mathcal{R}}_\Upsilon$, compatibly with gluing near the boundary of the moduli space, in such a

way that the moving Lagrangian boundary conditions are (i) constant (with value $X_{q_{\sigma_-}}$, L and L') along the boundary segments with label σ_- , L and L' , and (ii) given by a family of paths with endpoints on $X_{q_{\sigma}}$, interpolating between the constant path and the concatenation of the path $\Phi X_{q_{\sigma}}$ from $X_{q_{\sigma}}$ to $\phi(X_{q_{\sigma}})$ and its inverse along the remaining boundary. Our previous choices also determine families of almost complex structures on X parametrised by the fibres over the endpoints of the moduli space; we extend them to a family J_S over each curve S in $\overline{\mathcal{R}}_{\Upsilon}$ in such a way that the almost complex structure along the boundary components labelled L and L' respectively agree with J_L and $J_{L'}$, and the restrictions to the ends agree with the choices made in the construction of the Floer cohomology groups.

Each such surface is equipped with a decomposition which we call the thick-thin decomposition; for a surface far from the boundary, the thin parts are neighbourhoods of the ends where the boundary conditions are locally constant and the almost complex structure is obtained by pull-back from an interval by a choice of strip-like end. For a surface near the boundary, there is an additional component coming from the gluing region. Each component Θ of the thin part has a corresponding subset $\nu_X \Theta$ of X equipped with a compact subset of the interior away from which the restrictions to Θ of the boundary conditions and the almost complex structures are constant.

For each surface $S \in \overline{\mathcal{S}}_{\Upsilon}$, let X_S denote the fibre bundle over the pair of closed intervals obtained from the subset of ∂S labelled by σ^{ϕ} and σ_- by collapsing the components of the intersection with the thin part. We denote by X_S/\sim the quotient of X_S by the equivalence relation which, in each fibre corresponding to a component Θ of the thin part, identifies the components of the intersection with $\nu_X \Theta$ to points. As before, we pick a metric on this quotient so that the lengths of loops in the quotient provide a bound for the norm of the homology class in a fibre.

By the results of Appendix A, we can find a constant C which uniformly controls the length of the projection to X_S/\sim of the boundary of J_S -holomorphic curves in X with the above boundary conditions. We then require that

$$(2.8.17) \quad \text{diam } \nu_X \sigma < \frac{1}{8C}.$$

By construction, we have an evaluation map

$$(2.8.18) \quad \overline{\mathcal{R}}(\Upsilon) \rightarrow \text{Crit}(L, \sigma_-) \times \text{Crit}(\sigma_-, L') \times \text{Crit}(L', \sigma) \times \text{Crit}(L, \sigma).$$

Choosing the data generically, the count of rigid elements of these moduli spaces thus induces a map

$$(2.8.19) \quad \mathcal{R}_{L'}(\sigma_-) \otimes_{\Lambda} \mathcal{L}_L(\sigma_-) \rightarrow \text{Hom}_{\Lambda}(\mathcal{L}_{L'}(\sigma), \mathcal{L}_L(\sigma))$$

which gives a homotopy for the following diagram

$$(2.8.20) \quad \begin{array}{ccc} \mathcal{R}_{L'}(\sigma_-) \otimes_{\Lambda} \mathcal{L}_L(\sigma_-) & \longrightarrow & CF^*(L, L') \\ \downarrow & & \downarrow \\ \text{Hom}_{\Lambda}(\mathcal{L}_{L'}(\sigma), \overline{\Delta}(\sigma_-, \sigma)) \otimes_{\Lambda} \mathcal{L}_L(\sigma_-) & \longrightarrow & \text{Hom}_{\Lambda}(\mathcal{L}_{L'}(\sigma), \mathcal{L}_L(\sigma)), \end{array}$$

because the boundary strata of $\overline{\mathcal{R}}(\Upsilon)$ over the boundary of $\overline{\mathcal{R}}_{\Upsilon}$ correspond to the two compositions. Passing to cohomology, we obtain the commutativity of Diagram (2.1.13).

2.8.3. Compatibility of the maps of left modules. We now prove the commutativity of Diagram (2.8.4). Consider pre-stable discs equipped with two boundary punctures, one of which is distinguished as outgoing and the other as incoming, and with one boundary marked point.

Let $\overline{\mathcal{R}}_{3,2}$ denote the moduli space of stable maps from such discs to $(D^2)^2$ such that the projection to each factor has degree 1, the outgoing end maps to -1 in both factors, the incoming end maps to $+1$ in the first factor, and the boundary marked point maps to $+1$ in the second factor. This moduli space is a 2 dimensional manifold with corners which is homeomorphic as a topological manifold with boundary to a closed disc.

We can label each element of this moduli space by a disc equipped with a pair of interior marked points (in addition to the given boundary punctures and marked points), corresponding to the inverse image of the origin under the two maps: there is a unique point on the boundary of this moduli space where these marked points agree, lying on the stratum where both maps to the strip are supported on a component of the domain which is collapsed to the outgoing end by the forgetful map to $\overline{\mathcal{R}}_3$ (see the left side of Figure 11). This corresponds to an embedding

$$(2.8.21) \quad \overline{\mathcal{R}}_{2,1} \times \overline{\mathcal{R}}_3 \subset \overline{\mathcal{R}}_{3,2},$$

and gives rise to the composition of the top and right arrows in Diagram (2.8.4).

We have as well an embedding

$$(2.8.22) \quad \overline{\mathcal{R}}_3 \times (\overline{\mathcal{R}}_{2,1})^2 \subset \overline{\mathcal{R}}_{3,2},$$

shown on the right of Figure 11, and corresponding to the corner stratum of the right hand side given by the configuration where the components carrying the degree 1 maps project (under the forgetful map to $\overline{\mathcal{R}}_3$) to the boundary marked point and to the incoming puncture. For the appropriate choices of Lagrangian boundary conditions and families of almost complex structures, this embedding corresponds to the composition of the bottom and left arrows in Diagram (2.8.4).

To construct a homotopy in Diagram (2.8.4), we fix a path

$$(2.8.23) \quad \overline{\mathcal{R}}_{3,2} \subset \overline{\mathcal{R}}_{3,2}$$

interpolating between these two strata.

Let σ be a basepoint in Σ , L a Lagrangian in \mathcal{A} , and τ, τ_- elements of Σ_σ . Let Υ denote the sequence (L, τ_-, τ^ϕ) , and denote by $\overline{\mathcal{R}}_\Upsilon$ a copy of the space $\overline{\mathcal{R}}_{3,2}$, with corresponding boundary labels on the complement of the boundary marked point. Let $\overline{\mathcal{S}}_\Upsilon$ denote the universal curve over $\overline{\mathcal{R}}_\Upsilon$.

Over the boundary boundary stratum of $\overline{\mathcal{R}}_\Upsilon$ labelled by a configuration with two disc components, the boundary labels are given by (L, τ^ϕ) on the disc carrying an interior marked point, and (L, τ_-, τ^ϕ) on the other. The first case was already considered in Section 2.6.5, where we chose a path of Lagrangians between the pairs (L_σ, X_{q_σ}) and (L_τ, X_{q_τ}) . In the second case, we use the triple of constant Lagrangians $(L_\sigma, X_{q_\sigma}, X_{q_\sigma})$.

Over the other boundary stratum of $\overline{\mathcal{R}}_\Upsilon$ the choice of moving Lagrangian boundary conditions is immediately obtained from the choices made in Sections 2.6.5 for the disc with labels (L, τ_-^ϕ) , in Section 2.7.2 for the disc with interior marked point with labels (τ_-, τ^ϕ) , and in Section 2.8.1 for the disc with labels (L, τ_-, τ^ϕ) .

We now pick families of moving Lagrangian boundary conditions, which are Hamiltonian along the boundary labelled L , and almost complex structures on the universal curve over $\overline{\mathcal{R}}_\Upsilon$, which near the endpoints of this interval, are obtained by gluing, up to a perturbation term supported in the thick part away from the boundaries in which we apply the reverse isoperimetric inequality. In this way, we achieve transversality using standard methods for the corresponding moduli space $\overline{\mathcal{R}}(\Upsilon)$, while still being able to appeal to Lemma A.4 to ensure the existence of a uniform reverse isoperimetric constant for elements of this moduli

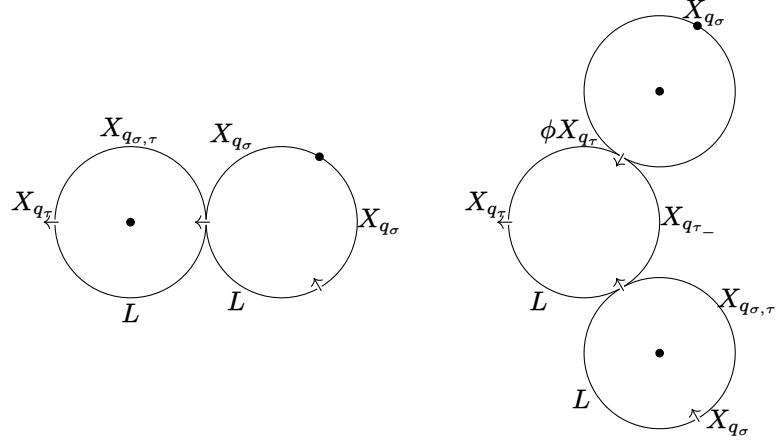


FIGURE 11. The boundary of the moduli space giving rise to the homotopy between the two compositions in Diagram (2.8.4). On the right, the component carrying the output corresponds to the left side of Figure 9 while the other two components correspond to the right side of Figure 5 and to Figure 7. On the left, the component carrying the output corresponds to the right side of Figure 7, while the other component corresponds to the right side of Figure 5. The two interior marked points on the right have collided on the left.

space. To avoid bubbling problems, we assume that the family of almost complex structures at each point z along the boundary condition labelled L agrees with J_L . We also assume that the boundary condition at the marked point with label (τ_-, τ^ϕ) agrees with X_{q_σ} , so that we have an evaluation map

$$(2.8.24) \quad \overline{\mathcal{R}}(\underline{\Upsilon}) \rightarrow X_{q_\sigma}.$$

Assuming that the diameter of $\nu_Q \sigma$ is sufficiently small relative to the reverse isoperimetric constant for this moduli space, the count of elements of the fibre product

$$(2.8.25) \quad \overline{\mathcal{T}}\mathcal{R}(\underline{\Upsilon}) := \overline{\mathcal{R}}(\underline{\Upsilon}) \times_{X_{q_\sigma}} \overline{\mathcal{T}}_+(\sigma, \sigma)$$

space defines a homotopy for the diagram

$$(2.8.26) \quad \begin{array}{ccc} \text{Po}_\sigma(P_{\tau_-}, P_\tau) \otimes_\Lambda \mathcal{L}_L(P_{\tau_-}) & \longrightarrow & \mathcal{L}_L(P_\tau) \\ \downarrow & & \downarrow \\ \overline{\Delta}(\tau_-, \tau) \otimes_\Lambda \mathcal{L}_L(\tau) & \longrightarrow & \mathcal{L}_L(\tau). \end{array}$$

Passing to cohomology yields the commutativity of Diagram (2.8.4).

2.8.4. Comparison with the map to the perturbed diagonal. In this section, we prove the commutativity of Diagram (2.8.5); most of the arguments are essentially the same as in the previous section, except that we shall encounter a moduli space with a boundary facet containing a node for which the boundary label is given by a pair of Lagrangians which agree. This requires us to add to the moduli space an additional component consisting of pairs of discs connected by a gradient flow line of varying length. This type of additional

cobordism appears throughout the literature when combining Morse-theoretic and Floer-theoretic moduli spaces.

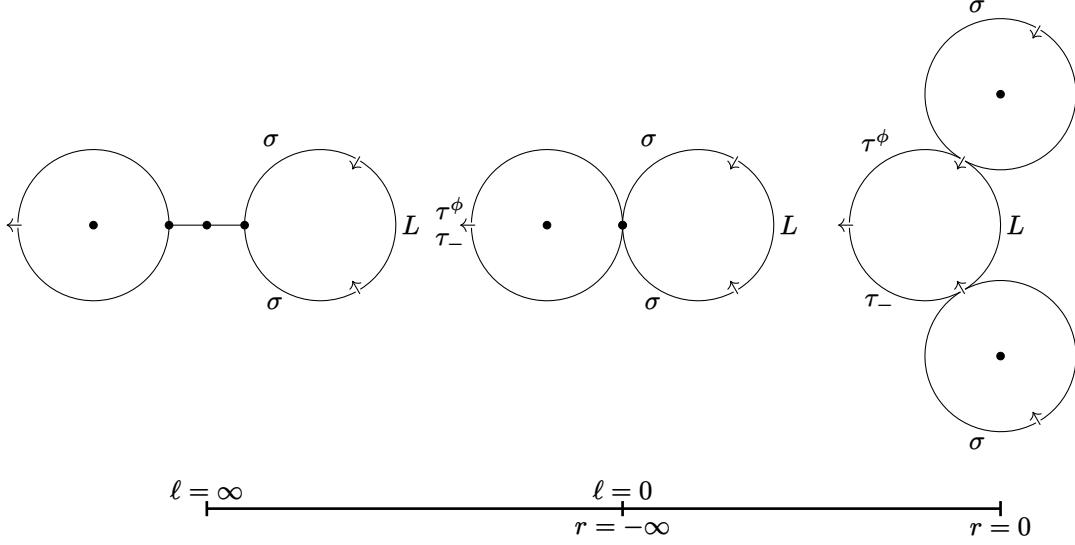


FIGURE 12. The boundary of the moduli space giving rise to the homotopy between the two compositions in Diagram (2.8.5).

Given a triple $\Upsilon = (\tau_-, L, \tau)$ with $L \in \mathcal{A}$ and $\tau, \tau_- \subseteq \mathcal{N}_\sigma$ as before, we shall construct a moduli space interpolating between the two moduli spaces defining the compositions being compared. To this end, we let $\bar{\mathcal{R}}_\Upsilon$ denote a copy of the space $\bar{\mathcal{R}}_{3,2}$, which we now think of as parametrising discs with 3 boundary punctures and 2 interior marked points whose position is constrained, as illustrated in the middle and right of Figure 12.

Over the boundary stratum of $\bar{\mathcal{R}}_\Upsilon$ labelled by a configuration consisting of three discs, the moving Lagrangian boundary conditions are given by the choices made in Section 2.6.5 for the two discs carrying marked points, and by those in Section 2.8.1 for the thrice-punctured disc.

Over the other boundary stratum of $\bar{\mathcal{R}}_\Upsilon$, the choice of basepoint $\sigma \in \Sigma$ yields moving Lagrangian conditions on the boundary of the disc carrying the interior marked point, which were fixed in Section 2.7.2. We forget the labels on the second component, and consider the constant boundary conditions $(X_{q_\sigma}, L, X_{q_\sigma})$. Note that these two conditions are compatible because, near the marked point, the boundary conditions considered in Section 2.7.2 are constant with value X_{q_σ} ; this is why we label the middle of Figure 12 with a node connecting the two components (instead of a pair of ends with matching labels).

As in the previous section we interpolate between these two boundary conditions, and the corresponding choices of almost complex structures, to obtain a moduli space $\bar{\mathcal{R}}(\Upsilon)$, equipped with an evaluation map:

$$(2.8.27) \quad \bar{\mathcal{R}}(\Upsilon) \rightarrow \text{Crit}(L, \sigma) \times \text{Crit}(\sigma, L) \times \text{Crit}(\tau_-, \tau^{\phi}).$$

There is a reverse-isoperimetric constant for these moduli space, with respect to which we impose the condition that the cover be sufficiently small.

The stratum of the boundary of this moduli space corresponding to discs with 3 components evidently gives rise to the composition

$$(2.8.28) \quad j^* H\mathcal{R}_{L,\sigma} \rightarrow H\mathcal{R}_L \rightarrow \text{Hom}_{H\mathcal{F}_\sigma}(j^* H\mathcal{L}_{L,\sigma}, \overline{\Delta}),$$

around the left and bottom of Diagram (2.8.5). The other boundary does not, however, correspond to the other composition, because it should involve a Floer cohomology group of local systems over X_{q_σ} , which has been defined using Morse theory.

We therefore introduce a space $\overline{\mathcal{T}}_{[0,\infty]}(\sigma, \sigma)$ of (perturbed) gradient flow lines of $f_{\sigma,\sigma}$ with arbitrary length, which is equipped with a natural evaluation map

$$(2.8.29) \quad \overline{\mathcal{T}}_{[0,\infty]}(\sigma, \sigma) \rightarrow X_{q_\sigma}^2$$

corresponding to the two ends of the flow line. The boundary is covered by a copy of X_{q_σ} on which the evaluation map is the diagonal (corresponding to the gradient flow line of length 0), and the closure of the codimension 1 stratum

$$(2.8.30) \quad \mathcal{T}_+(\sigma, \sigma) \times_{\text{Crit}(\sigma, \sigma)} \mathcal{T}_-(\sigma, \sigma)$$

which corresponds to gradient lines of infinite length. The compactification is obtained by allowing broken flow lines with multiple components, i.e. by adding the strata:

$$(2.8.31) \quad \mathcal{T}_+(\sigma, \sigma) \times_{\text{Crit}(\sigma, \sigma)} \overline{\mathcal{T}}(\sigma, \sigma) \times_{\text{Crit}(\sigma, \sigma)} \mathcal{T}_-(\sigma, \sigma).$$

We now define $\overline{\mathcal{T}\mathcal{R}}(\underline{\Upsilon})$ to be the union of $\overline{\mathcal{R}}(\underline{\Upsilon})$ with the fibre product

$$(2.8.32) \quad \overline{\mathcal{R}}(\underline{\Upsilon}_{e_{\text{out}}}) \times_{X_{q_\sigma}} \overline{\mathcal{T}}_{[0,\infty]}(\sigma, \sigma) \times_{X_{q_\sigma}} \overline{\mathcal{R}}(\sigma, L, \sigma),$$

where $\underline{\Upsilon}_{e_{\text{out}}} = (\tau_-, \tau^\phi)$. By construction, we have an evaluation map

$$(2.8.33) \quad \overline{\mathcal{T}\mathcal{R}}(\underline{\Upsilon}) \rightarrow \text{Crit}(L, \sigma) \times \text{Crit}(\sigma, L) \times \text{Crit}(\tau_-, \tau^\phi).$$

For generic choices of Floer and Morse data, this moduli space gives rise to a homotopy for the diagram

$$(2.8.34) \quad \begin{array}{ccc} \mathcal{L}_L(P_\tau) \otimes_\Lambda \mathcal{R}_L(\tau_-) & \longrightarrow & \text{Po}_\sigma(\tau_-, \tau) \\ \downarrow & & \downarrow \\ \mathcal{L}_L(\tau) \otimes_\Lambda \mathcal{R}(\tau_-) & \longrightarrow & \overline{\Delta}(\tau_-, \tau). \end{array}$$

Passing to cohomology, and using adjunction, we obtain the commutativity of Diagram (2.8.5).

3. HIGHER MODULI SPACES

In this section, we consider the families of holomorphic curves which will be used in the A_∞ refinement of the constructions of Section 2. We separate considerations of convergence and transversality by proceeding in two steps: (i) we first write families of equations with moving Lagrangian boundary conditions for which we state reverse isoperimetric inequalities that are robust under a certain class of perturbations, then (ii) we perturb these equations to achieve transversality among the Lagrangian boundary conditions and regularity for the moduli spaces of holomorphic curves, within the allowable class.

3.1. Abstract moduli spaces of discs. A *stable tree* is a finite tree T with edges $E(T)$ and vertices $V(T)$ such that the valency of every vertex is larger than or equal to 3; we write $F(T)$ for the set of flags (v, e) (i.e. a pair consisting of an edge and an adjacent vertex). We allow edges which are adjacent to a single vertex, and which are called *external* and form a set denoted $E^{\text{ext}}(T)$. Given a vertex v of T , we write T_v for the tree consisting of those edges adjacent to v ; this is a tree with a single vertex, and with edge set that we denote $E_v(T)$.

A stable *ribbon* tree is a tree as above, equipped with a cyclic ordering of the edges $E_v(T)$ for each vertex. We shall more precisely consider *rooted ribbon trees*, i.e. ribbon trees equipped with a distinguished external edge (the root) which we call outgoing, forming the singleton $E^{\text{out}}(T) \subset E(T)$. The remaining external edges are called incoming and are denoted $E^{\text{inp}}(T)$. The path from any vertex v to the root determines a unique outgoing edge e_v^{out} of T_v ; we call the remaining edges adjacent to v incoming, and we denote the sets of incoming and outgoing edges at a vertex v by $E_v^{\text{inp}}(T)$ and $E_v^{\text{out}}(T)$. The cyclic ordering and the choice of outgoing edge give rise to a unique ordering of $E_v(T)$ compatible with the cyclic structure with the property that e_v^{out} is the last element.

The cyclic ordering also determines an isotopy class of proper embeddings of the corresponding topological tree in the plane, and the choice of outgoing edge determines an ordering of the components of $\mathbb{R}^2 \setminus T$ given counterclockwise starting with the component which is to the left of the outgoing edge when directed outwards. Given a set A , each finite sequence Υ of elements of A consisting of $E^{\text{ext}}(T)$ elements induces a map

$$(3.1.1) \quad \pi_0(\mathbb{R}^2 \setminus T) \rightarrow A.$$

We assign to each edge e of T the ordered pair Υ_e consisting of the labels of the two regions adjacent to e , with the convention that, if e is oriented towards the outgoing edge, the label which is to the left appears first. We denote by Υ_v the induced label on the tree T_v associated to each vertex v .

Let us fix a Riemannian metric on Q satisfying Condition (2.2.18). In this section, we shall consider as labels (finite) sequences Υ of elements $Q \amalg Q^\phi \amalg \mathcal{A}$ (where $Q^\phi = Q \times \{\phi\}$) such that

- (3.1.2) (i) there are at most $\dim Q + 1$ distinct elements of Q (respectively Q^ϕ), and at most 2 elements of \mathcal{A} appearing in Υ (ii) the elements of Q (respectively Q^ϕ) are consecutive in Υ and are contained in a ball of radius 1, and (iii) any elements of Q^ϕ appear last.

In other words, the sequence is of the form

$$(3.1.3) \quad (L_1, \dots, L_j, q_{-r}, \dots, q_{-1}, L'_1, \dots, L'_k, q_1^\phi, \dots, q_\ell^\phi)$$

where $0 \leq j+k \leq 2$ and $0 \leq \ell, r \leq \dim Q + 1$.

Given such a label Υ , let $\bar{\mathcal{R}}_\Upsilon$ denote the moduli space of stable holomorphic discs with (i) a boundary marked point for each cyclically successive pair of elements of Q or Q^ϕ in Υ and (ii) a boundary puncture for each other cyclically successive pair of elements of Υ . We require that the cyclic ordering induced by the ordering of Υ correspond to the counterclockwise ordering around the boundary of the disc. There are thus $|\Upsilon|$ punctures and marked points in total, and the boundary components of the complement of the marked points are labelled by the elements of Υ . In addition, there is a distinguished puncture corresponding to the first and last element of Υ , which we call outgoing. For each tree T labelled by Υ , we then

define

$$(3.1.4) \quad \overline{\mathcal{R}}_{\Upsilon}^T := \prod_{v \in V(T)} \overline{\mathcal{R}}_{\Upsilon_v}.$$

We distinguish the subset

$$(3.1.5) \quad V^{\deg}(T) \subseteq V(T)$$

consisting of vertices with *degenerate labels*, i.e. such that the label is contained in Q or Q^ϕ .

Each map $T \rightarrow T'$ which collapses internal edges induces an inclusion

$$(3.1.6) \quad \overline{\mathcal{R}}_{\Upsilon}^T \rightarrow \overline{\mathcal{R}}_{\Upsilon'}^{T'}$$

and the tree with a unique vertex corresponds to the moduli space $\overline{\mathcal{R}}_{\Upsilon}$. Moreover, if $\Upsilon \rightarrow \Upsilon'$ is a map of ordered sets which collapses successive elements that are equal, we obtain a (forgetful) map of moduli spaces

$$(3.1.7) \quad \overline{\mathcal{R}}_{\Upsilon} \rightarrow \overline{\mathcal{R}}_{\Upsilon'}.$$

Let $\overline{\mathcal{S}}_{\Upsilon}$ denote the universal curve over $\overline{\mathcal{R}}_{\Upsilon}$. For $L \in \mathcal{A}$, $q \in Q$, and $q^\phi \in Q^\phi$, we denote by $\partial_L \overline{\mathcal{S}}_{\Upsilon}$, $\partial_q \overline{\mathcal{S}}_{\Upsilon}$, and $\partial_q^\phi \overline{\mathcal{S}}_{\Upsilon}$ the corresponding boundary segment of the complement of the marked points. We write $\partial_Q \overline{\mathcal{S}}_{\Upsilon}$ and $\partial_Q^\phi \overline{\mathcal{S}}_{\Upsilon}$ for the union of boundary components labelled by elements of Q or Q^ϕ .

For each tree T , we write $\overline{\mathcal{S}}_{\Upsilon}^T$ for the fibre over the stratum $\overline{\mathcal{R}}_{\Upsilon}^T$. This space decomposes as a disjoint union of components labelled by the vertices of T :

$$(3.1.8) \quad \overline{\mathcal{S}}_{\Upsilon}^T := \coprod_{v \in V(T)} \overline{\mathcal{S}}_v$$

$$(3.1.9) \quad \overline{\mathcal{S}}_v := \overline{\mathcal{S}}_{\Upsilon_v} \times_{\overline{\mathcal{R}}_{\Upsilon_v}} \overline{\mathcal{R}}_{\Upsilon}^T.$$

We shall need to consider certain moduli spaces of discs with conformal constraints, which depend on a choice of basepoint q_* on Q . Consider ordered subsets Υ of $Q \amalg Q^\phi \amalg \mathcal{A}$ of the form

$$(3.1.10) \quad (L, q_1, \dots, q_\ell)$$

$$(3.1.11) \quad (q_{-r}, \dots, q_{-1}, L)$$

$$(3.1.12) \quad (q_{-r}, \dots, q_{-1}, q_1^\phi, \dots, q_\ell^\phi)$$

$$(3.1.13) \quad (L, q_{-r}, \dots, q_{-1}, q_1^\phi, \dots, q_\ell^\phi)$$

$$(3.1.14) \quad (q_{-r}, \dots, q_{-1}, L, q_1^\phi, \dots, q_\ell^\phi),$$

where all elements of Q and Q^ϕ are within distance 1 of q_* .

In the first three cases above, we define $\overline{\mathcal{R}}_{\Upsilon}$ to be the moduli space of stable degree 1 maps to D^2 whose domain is a pre-stable disc with boundary marked points labelled by successive pairs of elements of the sequence Υ both of which lie in Q or Q^ϕ , and a puncture for the remaining two cyclically successive elements of Υ (again, we require that the counterclockwise ordering on the boundary of the disc corresponds to the order of the elements of Υ). We only consider maps so that the outgoing end maps to -1 , and the incoming puncture labelled by (L, q_1) , (q_{-1}, L) , or (q_{-1}, q_1^ϕ) maps $+1$.

In the last two cases, we define $\overline{\mathcal{R}}_{\Upsilon}$ to be the inverse image of the 1-dimensional submanifold $\overline{\mathcal{R}}_{3,2} \subset \overline{\mathcal{R}}_{3,2}$ fixed in Section 2.8.3, in the moduli space of stable maps from a pre-stable

disc with boundary marked points or punctures given as before by cyclically successive elements of Υ to $(D^2)^2$, so that the image of the outgoing end under both factors is -1 , and the ends labelled by (L, q_{-r}) and (q_{-1}, q_1^ϕ) in Equation (3.1.13) (respectively by (q_{-1}, L) and (q_1^ϕ, q_2^ϕ) in Equation (3.1.14)) map to $+1$ under the two factors. The map to $\bar{\mathcal{R}}_{3,2}$ is thus obtained by forgetting all marked points labelled by pairs consisting of two elements of Q or two elements of Q^ϕ .

Remark 3.1. There is a minor difference between the moduli spaces $\bar{\mathcal{R}}_{\underline{\Upsilon}}$ for Υ given by Equations (3.1.13) and (3.1.14). In the first case, we consider $\bar{\mathcal{R}}_{3,2}$ as parametrising discs with one boundary marked point, and two punctures, while in the second, we have three boundary punctures. The difference will be apparent when we discuss the boundary conditions that will be imposed on these moduli spaces.

The boundary strata of $\bar{\mathcal{R}}_{\underline{\Upsilon}}$ are also labelled by trees. To describe them, we begin by associating to each stratum of $\bar{\mathcal{R}}_{2,1}$ or $\bar{\mathcal{R}}_{3,2}$ a rooted ribbon tree T equipped with a distinguished subset of vertices denoted $M(T)$ (*marked vertices*) consisting of those vertices which correspond to disc components carrying a non-trivial map to a disc. For $\bar{\mathcal{R}}_{2,1}$, we thus obtain a tree with a unique bivalent node, and no other vertices. In the same way, each boundary stratum of $\bar{\mathcal{R}}_{\underline{\Upsilon}}$ corresponds to a ribbon tree with a distinguished set of vertices $M(T) \subset V(T)$ which are allowed to be bivalent, and such that the tree obtained by forgetting all incoming edges with labels pairs of elements in Q or Q^ϕ corresponds to a tree labelling a boundary stratum of $\bar{\mathcal{R}}_{2,1}$ or $\bar{\mathcal{R}}_{3,2}$.

We distinguish the set of vertices $V^*(T) \subset V(T)$ with the property that they separate an incoming Floer edge from all nodes. Writing Υ_v^* for the sequence obtained from Υ_v by replacing all points in Q or Q^ϕ by q_* , we have natural product decomposition

$$(3.1.15) \quad \bar{\mathcal{R}}_{\underline{\Upsilon}}^T := \prod_{v \in M(T)} \bar{\mathcal{R}}_{\Upsilon_v} \times \prod_{v \in V^*(T)} \bar{\mathcal{R}}_{\Upsilon_v^*} \times \prod_{v \notin M(T) \cup V^*(T)} \bar{\mathcal{R}}_{\Upsilon_v}.$$

Remark 3.2. The distinction between $\bar{\mathcal{R}}_{\Upsilon_v^*}$ and $\bar{\mathcal{R}}_{\Upsilon_v}$ is entirely formal, as the two spaces are naturally homeomorphic. As indicated by the labelling of Figure 13, we shall use these moduli spaces to extend the construction of Section 2.7.2, which should indicate to the reader why we introduce this notation.

As in Equation (3.1.5), we also define the set $V^{\deg}(T) \subset V(T) \setminus M(T)$ to consist of those vertices with label Υ_v contained in Q or Q^ϕ .

By construction, elements of $\bar{\mathcal{R}}_{\underline{\Upsilon}}$ correspond to curves with at most two incoming ends. We set the labels on these incoming ends to be

$$(3.1.16) \quad \Upsilon_e = (L, q_*) \text{ or } \Upsilon_e = (q_*, L),$$

while for all other ends we use the ribbon embedding from Equation (3.1.1) to set the label.

Let $\bar{\mathcal{S}}_{\underline{\Upsilon}}$ denote the universal curves over $\bar{\mathcal{R}}_{\underline{\Upsilon}}$, obtained by restriction of the universal curve over $\bar{\mathcal{R}}_{\Upsilon}$. As before, given a tree T , the fibre $\bar{\mathcal{S}}_{\underline{\Upsilon}}^T$ over $\bar{\mathcal{R}}_{\underline{\Upsilon}}^T$ can be written as a union of components labelled by vertices of T , and which we denote $\bar{\mathcal{S}}_{\underline{\Upsilon}}^v$. We have

$$(3.1.17) \quad \bar{\mathcal{S}}_{\underline{\Upsilon}}^v := \begin{cases} \bar{\mathcal{S}}_{\Upsilon_v} \times_{\bar{\mathcal{R}}_{\Upsilon_v}} \bar{\mathcal{R}}_{\underline{\Upsilon}}^T & v \in M(T) \\ \bar{\mathcal{S}}_{\Upsilon_v^*} \times_{\bar{\mathcal{R}}_{\Upsilon_v^*}} \bar{\mathcal{R}}_{\underline{\Upsilon}}^T & v \in V^*(T) \\ \bar{\mathcal{S}}_{\Upsilon_v} \times_{\bar{\mathcal{R}}_{\Upsilon_v}} \bar{\mathcal{R}}_{\underline{\Upsilon}}^T & \text{otherwise.} \end{cases}$$

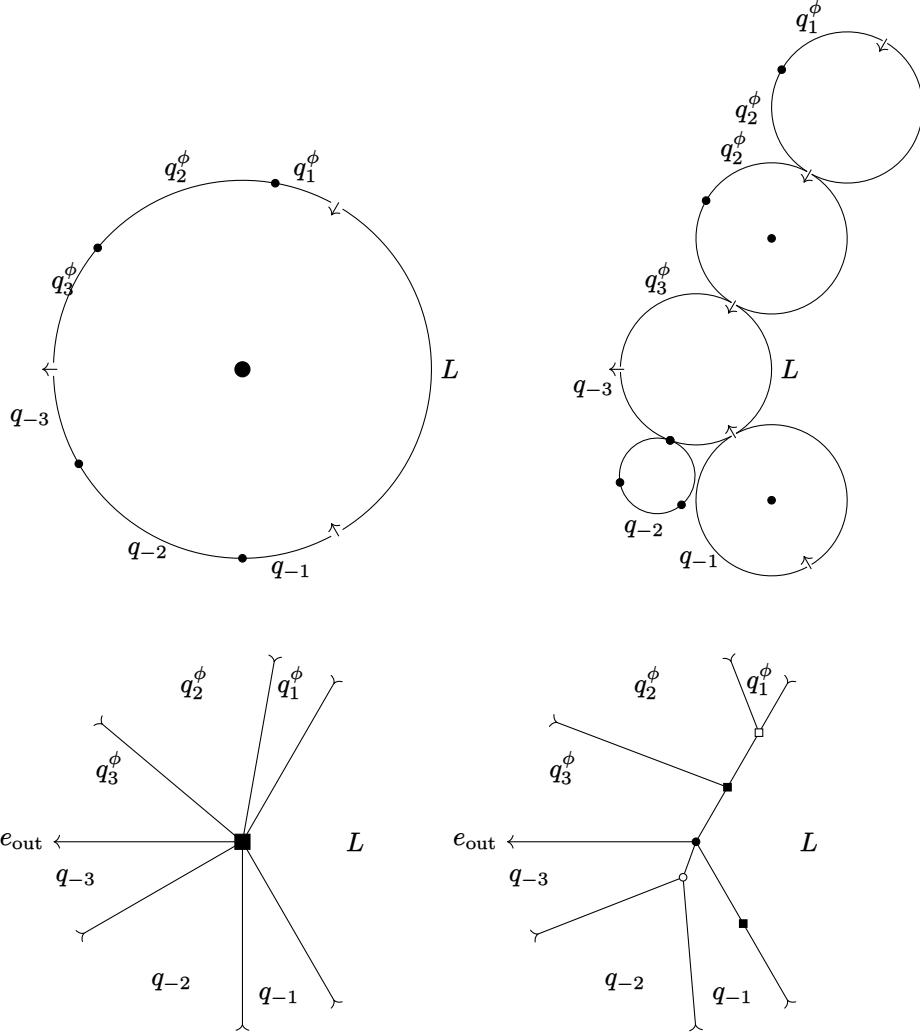


FIGURE 13. A representation of two elements of the same moduli space $\bar{\mathcal{R}}(\Upsilon)$, and the trees labeling the strata they lie on. The solid square vertices are in $M(T)$, the open square vertex in $V^*(T)$, and the open circle in $V^{\deg}(T)$.

3.1.1. *Strip-like ends, gluing charts, and thin-thick decompositions.* Recall that a positive (respectively negative) strip-like end on a Riemann surface S is a holomorphic map

$$(3.1.18) \quad [0, \infty) \times [0, 1] \rightarrow S \text{ or } (-\infty, 0] \times [0, 1] \rightarrow S$$

which is a biholomorphism in a neighbourhood of a puncture. We shall equip the outgoing puncture of a curve in $\bar{\mathcal{R}}_\Upsilon$ with a negative strip-like end, and all other punctures with positive strip-like ends. As in [26, Section (9f)], we make this choice consistently for all curves in $\bar{\mathcal{R}}_\Upsilon$ and for all labels Υ . The key consistency condition is that, near the boundary strata of $\bar{\mathcal{R}}_\Upsilon$, the strip-like ends are compatible with the gluing maps which are constructed as follows:

for each tree T labelling a stratum of $\overline{\mathcal{R}}_\Upsilon$, the choices of ends induces an open embedding (the gluing map)

$$(3.1.19) \quad \overline{\mathcal{R}}_\Upsilon^T \times (0, \infty]^{E(T)} \rightarrow \overline{\mathcal{R}}_\Upsilon.$$

This open embedding lifts to the level of universal curves, so that the choice of strip-like ends on curves parametrised by $\overline{\mathcal{R}}_\Upsilon^T$ induces strip-like ends on curves in the image of the gluing map, and we require that the choice of ends be induced by the gluing map whenever the gluing parameters are sufficiently large.

Fixing the identification of $(0, \infty]$ with $(-1, 0]$ which takes R to $-e^{-R}$, we obtain charts

$$(3.1.20) \quad \overline{\mathcal{R}}_\Upsilon^T \times (-1, 0]^{E(T)} \rightarrow \overline{\mathcal{R}}_\Upsilon$$

which equip the moduli space with the structure of a smooth manifold with corners (the easiest way to prove that this is a smooth structure is to realise these moduli spaces are submanifolds of Deligne-Mumford spaces of closed genus 0 Riemann surfaces with marked points).

The gluing map lifts at the level of universal curves to a map

$$(3.1.21) \quad \overline{\mathcal{S}}_\Upsilon^T \times (-1, 0]^{E(T)} \rightarrow \overline{\mathcal{S}}_\Upsilon,$$

which is surjective over the image of the gluing map at the level of moduli spaces. The same discussion applies to the moduli spaces $\overline{\mathcal{R}}_\Upsilon$: choices of strip-like ends induce gluing charts near the boundary strata, which lift to the universal curve.

The gluing charts equip each curve S in $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_\Upsilon$ with a thick-thin decomposition: the thin part will be the union of the image under the gluing maps of (i) the strip-like ends, (ii) the curves corresponding to degenerate vertices, and (iii) curves which become unstable upon omitting repeated elements of Q or Q^ϕ . The thick part is the complement. We associate to each component Θ of the thin part the subset $\nu_X \Theta$ of X given by (i) the empty set if all labels appearing on the intersection of the boundary with Θ are contained in Q or Q^ϕ , (ii) the set $\nu_X \Upsilon_e$ if Θ contains the image of an end e under gluing. Here, we set $\nu_X \Upsilon$ to be (i) X if Υ is degenerate and (ii) the set $\nu_X L$ if Υ consists of an element of $Q \cup Q^\phi$ and $L \in \mathcal{A}$, and (iii) the set $\nu_X (X_{q_-} \cap \phi X_q)$ if $\Upsilon = (q_-, q^\phi)$.

By construction, any two ends with images in Θ have the same label, so this definition is consistent (it is important here to distinguish ends from marked points).

Finally, for each S in $\overline{\mathcal{S}}_\Upsilon$, we fix neighbourhoods $\nu \partial_Q S$ and $\nu \partial_Q^\phi S$ in S of the corresponding boundary components of S , which are compatible with gluing, only intersect each other in the thin part, and whose union forms an open set in the universal curve. If S contains a component with degenerate labels, we assume that such a component is contained in $\nu \partial_Q S$ or $\nu \partial_Q^\phi S$ as long as there is a part of the boundary with the corresponding label. We write $\nu \partial_Q \overline{\mathcal{S}}_\Upsilon$ and $\nu \partial_Q^\phi \overline{\mathcal{S}}_\Upsilon$ for the union of such neighbourhoods over all points $S \in \overline{\mathcal{R}}_\Upsilon$.

We also choose such neighbourhoods of the boundary components of $S \in \overline{\mathcal{S}}_\Upsilon$; the only difference is that the neighbourhood $\nu \partial_Q S$ and $\nu \partial_Q^\phi S$ should intersect in the union of the thin part with a neighbourhood of any marked point with label in $Q \amalg Q^\phi$.

3.2. Unperturbed families of equations. We continue to work with a fixed metric on the base Q satisfying Condition (2.2.18). The key property which shall use is that, for such a metric, the intersection of any two geodesic balls of radius less than 2 is either empty or contractible since it is geodesically convex.

3.2.1. *Paths of fibres.* For each non-degenerate sequence Υ of the form given in Equations (3.1.10)–(3.1.14), we choose maps

$$(3.2.1) \quad q_\Upsilon: \partial_Q \bar{\mathcal{S}}_\Upsilon \rightarrow Q$$

$$(3.2.2) \quad q_\Upsilon^\phi: \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon \rightarrow Q,$$

whose value along each segment labelled by $q \in Q$ lies in the ball of radius 1 about this point, and varying smoothly in the choice of labels. We require these maps to satisfy the following additional properties:

- (1) (Values along the ends) The restrictions of q_Υ and q_Υ^ϕ to the intersection of each end with $\partial_q \bar{\mathcal{S}}_\Upsilon$ and $\partial_q^\phi \bar{\mathcal{S}}_\Upsilon$ are constant with value q .
- (2) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_\Upsilon^T$, the restrictions of q_Υ and q_Υ^ϕ to a neighbourhood of $\bar{\mathcal{R}}_\Upsilon^T$ are obtained by gluing in the sense that the following diagram commutes in a neighbourhood of the origin:

$$(3.2.3) \quad \begin{array}{ccc} \left(\partial_Q \bar{\mathcal{S}}_\Upsilon^T \amalg \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon^T \right) \times (-1, 0]^{E(T)} & \longrightarrow & \partial_Q \bar{\mathcal{S}}_\Upsilon \amalg \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon \\ & \searrow q_\Upsilon \amalg q_\Upsilon^\phi & \downarrow q_\Upsilon \amalg q_\Upsilon^\phi \\ & & Q. \end{array}$$

- (3) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_\Upsilon$ and vertex $v \in V^{\deg}(T)$, the restriction of q_Υ and q_Υ^ϕ to $\bar{\mathcal{S}}_\Upsilon^v$ is constant (note that such components are disjoint from the ends, so this condition is independent of Condition 1 above). If v is not a degenerate vertex, these data are compatible with $(q_{\Upsilon_v}, q_{\Upsilon_v}^\phi)$ in the sense that we have a commutative diagram

$$(3.2.4) \quad \begin{array}{ccc} \partial_Q \bar{\mathcal{S}}_\Upsilon^v \amalg \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon^v & \longrightarrow & \partial_Q \bar{\mathcal{S}}_{\Upsilon_v} \amalg \partial_Q^\phi \bar{\mathcal{S}}_{\Upsilon_v} \\ & \searrow q_\Upsilon \amalg q_\Upsilon^\phi & \downarrow q_{\Upsilon_v} \amalg q_{\Upsilon_v}^\phi \\ & & Q. \end{array}$$

- (4) (Forgetful maps) If Υ is obtained from a sequence Υ' by repeating elements of Q or Q^ϕ , the choices of paths are compatible with the induced forgetful maps in the sense that the following diagrams commute:

$$(3.2.5) \quad \begin{array}{ccc} \partial_Q \bar{\mathcal{S}}_\Upsilon & \longrightarrow & \partial_Q \bar{\mathcal{S}}_{\Upsilon'} \\ & \searrow & \downarrow \\ & & Q. \end{array} \quad \begin{array}{ccc} \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon & \longrightarrow & \partial_Q^\phi \bar{\mathcal{S}}_{\Upsilon'} \\ & \searrow & \downarrow \\ & & Q. \end{array}$$

- (5) (Forgetting ϕ) If Υ is an ordered non-degenerate subset of $Q^\phi \amalg \mathcal{A}$, and π_Υ the corresponding ordered subset of $Q \amalg \mathcal{A}$, we have a commutative diagram

$$(3.2.6) \quad \begin{array}{ccc} \partial_Q^\phi \bar{\mathcal{S}}_\Upsilon & \longrightarrow & \partial_Q \bar{\mathcal{S}}_{\pi_\Upsilon} \\ & \downarrow & \swarrow \\ & & Q. \end{array}$$

The construction of such paths proceeds by induction on the number of elements of Υ , and the bound on the lengths of the paths ensures the contractibility of the corresponding space

of choices, ensuring that there is no obstruction to extending the choices from the boundary strata to the interior of the moduli space.

Next, we consider the moduli spaces associated to a choice of basepoint $q_* \in Q$. For each sequence Υ of the form given in Equations (3.1.10)–(3.1.14), we choose maps

$$(3.2.7) \quad q_{\Upsilon}: \partial_Q \bar{\mathcal{S}}_{\Upsilon} \rightarrow Q$$

$$(3.2.8) \quad q_{\Upsilon}^{\phi}: \partial_Q^{\phi} \bar{\mathcal{S}}_{\Upsilon} \rightarrow Q$$

varying smoothly in the choice of labels (and in the choice of basepoint q_* , which we leave implicit in the notation), and whose restriction to each boundary segment labelled by $q \in Q$ lies in the ball of radius 1 about this point. The key condition is that:

- (1) The maps have value q_* near any marked point or node with label given by one element in Q and one in Q^{ϕ} .

We require these maps to satisfy the following additional properties, which are entirely analogous to those listed above:

- (2) (Values along the ends) The restrictions of q_{Υ} and q_{Υ}^{ϕ} to the intersection of each end e with $\partial_q \bar{\mathcal{S}}_{\Upsilon}$ and $\partial_q^{\phi} \bar{\mathcal{S}}_{\Upsilon}$ are constant with value q if e is an outgoing end, and q_* if e an incoming end.
- (3) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_{\Upsilon}$, the restrictions of q_{Υ} and q_{Υ}^{ϕ} to a neighbourhood of $\bar{\mathcal{S}}_{\Upsilon}^T$ are obtained by gluing.
- (4) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_{\Upsilon}$ and vertex $v \in V(T)$, the restrictions of the maps q_{Υ} and q_{Υ}^{ϕ} to $\bar{\mathcal{S}}_{\Upsilon}^v$ agree with (i) constant functions if v is degenerate, (ii) the maps q_{Υ_v} and $q_{\Upsilon_v}^{\phi}$ if v is a node, (iii) the constant function q_* if $v \in V^*(T)$, and (iv) the functions q_{Υ_v} and $q_{\Upsilon_v}^{\phi}$ in the remaining cases.
- (5) (Forgetful maps) If $\Upsilon \rightarrow \Upsilon'$ is a map of labels, then q_{Υ} and $q_{\Upsilon'}^{\phi}$ are obtained from $q_{\Upsilon'}$ and $q_{\Upsilon'}^{\phi}$ by composition with the forgetful map.

Note that Condition (2) above is compatible with Equation (3.1.16), which sets the labels for incoming ends to consist of an element of \mathcal{A} and the basepoint q_* .

3.2.2. Deformation of the diagonal. Recall that we chose a Hamiltonian diffeomorphism ϕ in Section 2.7. We now pick a map

$$(3.2.9) \quad \Phi_{\Upsilon}: \partial_Q^{\phi} \bar{\mathcal{S}}_{\Upsilon} \rightarrow \text{Ham}(X)$$

such that the following properties hold:

- (1) (Values along the ends) the restriction to an end e agrees with ϕ if $\Upsilon_e = (\sigma_-, \sigma^{\phi})$ and with the identity otherwise.
- (2) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_{\Upsilon}$ the restriction of Φ_{Υ} to a neighbourhood of $\bar{\mathcal{R}}_{\Upsilon}^T$ is obtained by gluing in the sense that the following diagram commutes near the origin:

$$(3.2.10) \quad \begin{array}{ccc} \partial_Q^{\phi} \bar{\mathcal{S}}_{\Upsilon}^T \times (-1, 0]^{E(T)} & \xrightarrow{\quad} & \partial_Q^{\phi} \bar{\mathcal{S}}_{\Upsilon} \\ & \searrow \Phi_{\Upsilon} & \downarrow \Phi_{\Upsilon} \\ & & \text{Ham}(X). \end{array}$$

(3) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\overline{\mathcal{R}}_T$ and vertex $v \in V(T)$ the map Φ_T is (i) constant if v is degenerate and (ii) otherwise agrees with Φ_{T_v} in the sense that we have a commutative diagram

$$(3.2.11) \quad \begin{array}{ccc} \partial_Q^\phi \overline{\mathcal{S}}_T^v & \longrightarrow & \partial_Q^\phi \overline{\mathcal{S}}_{T_v} \\ \Phi_T \searrow & & \downarrow \Phi_{T_v} \\ & & \text{Ham}(X). \end{array}$$

(4) (Forgetful maps) For each forgetful map $\Upsilon \rightarrow \Upsilon'$, the following diagrams commute:

$$(3.2.12) \quad \begin{array}{ccc} \partial_Q^\phi \overline{\mathcal{S}}_\Upsilon & \longrightarrow & \partial_Q^\phi \overline{\mathcal{S}}_{\Upsilon'} \\ \searrow & & \downarrow \\ & & \text{Ham}(X). \end{array}$$

(5) (Forgetting ϕ) If Υ is an ordered subset of $Q^\phi \amalg \mathcal{A}$, Φ_Υ is constant with value the identity map.

Again, the construction proceeds by induction on the number of elements of Υ ; while the space of Hamiltonian diffeomorphisms may not be contractible, it is easy to lift all choices to the space of paths based at the identity in $\text{Ham}(X)$. In this way, all obstructions to extending choices from boundary strata to the interior vanish.

Similarly we consider a family

$$(3.2.13) \quad \Phi_{\underline{\Upsilon}}: \partial_Q^\phi \overline{\mathcal{S}}_{\underline{\Upsilon}} \rightarrow \text{Ham}(X),$$

such that the following properties hold:

- (1) The restriction to a neighbourhood of a node or marked point with label in $Q \times Q^\phi$ is constant with value the identity.
- (2) (Values along the ends) the restriction to each end e agrees with ϕ if Υ_e contains an element of Q , and with the identity otherwise.
- (3) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\overline{\mathcal{R}}_T$ the restriction of Φ_T to a neighbourhood of $\overline{\mathcal{S}}_T^T$ is obtained by gluing.
- (4) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\overline{\mathcal{R}}_T$ and vertex $v \in V(T)$, the restriction of Φ_T to $\overline{\mathcal{S}}_T^v$ agrees with (i) a constant function if v is degenerate, (ii) the map Φ_{T_v} if v is a node, (iii) the identity if $v \in V^*(T)$, and (iv) the map Φ_{T_v} otherwise.
- (5) (Forgetful maps) Whenever Υ is obtained from Υ' by repeating elements of Q or Q^ϕ , $\Phi_{\underline{\Upsilon}}$ is obtained from $\Phi_{\underline{\Upsilon}'}$ by composition with the forgetful map.

3.2.3. Choices of almost complex structures. Let \mathcal{J} denote the space of ω -tame almost complex structures on X . For simplicity, fix an almost complex structure J_ϕ on X which we will use to define the Floer theory of the perturbed diagonal (i.e. of fibres with their image under ϕ), and recall that we have chosen an almost complex structure J_L associated to each $L \in \mathcal{A}$, and subsets $\nu_X \Upsilon$ of X for each pair Υ of labels. For each sequence Υ as in the preceding sections, we consider maps

$$(3.2.14) \quad J_\Upsilon^\nu: \nu \partial_Q \overline{\mathcal{S}}_\Upsilon \cup \nu \partial_Q^\phi \overline{\mathcal{S}}_\Upsilon \rightarrow \mathcal{J}$$

$$(3.2.15) \quad J_{\underline{\Upsilon}}^\nu: \nu \partial_Q \overline{\mathcal{S}}_{\underline{\Upsilon}} \cup \nu \partial_Q^\phi \overline{\mathcal{S}}_{\underline{\Upsilon}} \rightarrow \mathcal{J}$$

satisfying the following conditions:

- (1) (Values along the ends) Away from a fixed compact subset of the interior of $\nu_X \Upsilon_e$, the restriction to each end e is constant, with value J_L if $L \in \Upsilon_e$, and J_ϕ otherwise.
- (2) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$ the restrictions to a neighbourhood of $\overline{\mathcal{R}}_T$ or $\overline{\mathcal{R}}_{\underline{T}}$ are obtained by gluing.
- (3) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$ and non-degenerate vertex $v \in V(T)$, the restriction of J_Υ^v or $J_{\underline{\Upsilon}}^v$ to $\overline{\mathcal{S}}_\Upsilon^v$ or $\overline{\mathcal{S}}_{\underline{\Upsilon}}^v$ agrees with the almost complex structure associated to Υ_v .
- (4) (Forgetful maps) The choice of almost complex structure is compatible with forgetful maps associated to repeating elements of Q^ϕ or Q .
- (5) (Forgetting ϕ) The choice of almost complex structure is compatible with the projection $Q^\phi \rightarrow Q$ if Υ contains no element of Q .

3.2.4. *Reverse isoperimetric constant.* For each curve S in $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$, the paths q_Υ and q_Υ^ϕ or $q_{\underline{\Upsilon}}$ and $q_{\underline{\Upsilon}}^\phi$ determine a closed manifold with boundary X_S , homeomorphic to two copies of the product of a fibre with $[0, 1]$ in the first case, and one such copy in the second case, and obtained from the boundary conditions over $\partial_Q S$ and $\partial_Q^\phi S$ by identifying the fibres over each component $\Theta \subset S$ of the thin part. We define the equivalence relation \sim on X_S to collapse the intersection of the fibre over Θ with each component of $\nu_X \Theta$ to a point. The assumption that this intersection is contained in a disjoint union of balls allows us (as in Lemma 2.26) to fix a metric on the quotient so that, for each loop in X_S , the length of the projection to X_S/\sim bounds the norm of the corresponding homology class in $H_1(X_q, \mathbb{Z})$ for any q on the path q_Υ or q_Υ^ϕ .

Given a curve S in $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$, let $u: S \rightarrow X$ be a map such that $u(z)$ lies (i) in $\nu_X L$ if $z \in \partial_L S$, (ii) in $X_{q_\Upsilon(z)}$ or $X_{q_{\underline{\Upsilon}}(z)}$ if $z \in \partial_Q S$, and (iii) in $\Phi_\Upsilon(z)X_{q_\Upsilon^\phi(z)}$ or $\Phi_{\underline{\Upsilon}}(z)X_{q_{\underline{\Upsilon}}^\phi(z)}$ if $z \in \partial_Q^\phi S$. Each such curve has an evaluation map

$$(3.2.16) \quad \partial u / \sim: \partial_Q S \amalg \partial_Q^\phi S \rightarrow X_S / \sim.$$

Let $\ell(\partial u / \sim)$ denote the length of this curve.

Lemma 3.3. *There is a constant C (depending on the collection of Lagrangians and the choices of almost complex structures) such that, for each choice of labels Υ satisfying Condition (3.1.2), we have*

$$(3.2.17) \quad \ell(\partial u / \sim) \leq CE^{geo}(u) + \text{a constant independent of } u$$

for each curve u with boundary conditions as above, whose restriction to $\nu \partial_Q S$ and $\nu \partial_Q^\phi S$ are holomorphic with respect to the complex structures fixed in Section 3.2.3. Moreover, for each components Θ of the thin part of S , this constant is independent of the restriction of the complex structure and the Lagrangian boundary conditions to $\Theta \times \nu_X \Theta$.

Proof. This is a consequence of Proposition A.5, which concerns reverse isoperimetric constants for families. Indeed, the existence of such a constant for each curve S follows from Lemma A.3. The assumption that the Floer data are compatible with gluing and boundary strata implies that we can use Lemma A.4 to conclude that the constant can be chosen uniformly for S in $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$ for any finite collection of labels. The fact that Q is compact implies that the constant may be chosen uniformly if the number of labels on the boundary is bounded. The assumption that the data are compatible with forgetful maps, and the

constraint that Υ contains at most $n + 1$ distinct elements of Q or Q^ϕ implies that the constant may be chosen uniformly. \square

3.3. Perturbed equations. We now return to the context of Section 2, and consider a finite cover Q by integral affine polygons with the property that the intersection of any collection consisting of more than $n + 1$ element is empty. Letting Σ denote the set of collections of elements of the cover with non-empty intersection, which is partially ordered by inclusion, we obtain a cover $\{P_\sigma\}_{\sigma \in \Sigma}$ of Q which has dimension n (i.e. the maximal length of any totally ordered subset of Σ is $n + 1$). We assume that this cover is sufficiently fine that we can pick a contractible set \mathcal{N}_σ for each element of Σ , which includes P_τ whenever P_σ intersects P_τ , and such that

$$(3.3.1) \quad \text{diam } \mathcal{N}_\sigma \leq \min(1, \frac{1}{8C}),$$

where the constant C is the one from Lemma 3.3. For each $\sigma \in \Sigma$, we choose a basepoint $q_\sigma \in P_\sigma$, which we think of as a map

$$(3.3.2) \quad \Sigma \rightarrow Q.$$

As before, we write Σ^ϕ for $\Sigma \times \{\phi\}$, whose elements are equipped with the same choices of basepoints. We assume that these basepoints are chosen so that the Lagrangian fibre X_{q_σ} is transverse to each element L of \mathcal{A} , and that the pair $(X_{q_\sigma}, \phi X_{q_\tau})$ is transverse for each pair $(\sigma, \tau) \in \Sigma$.

Let Υ denote an ordered subset of $\Sigma \amalg \Sigma^\phi \amalg \mathcal{A}$ such that

- (i) all elements of Σ (respectively Σ^ϕ) are consecutive and are increasing with respect to the partial order on Σ , (ii) all elements of Σ^ϕ appear last, and (iii) all elements of \mathcal{A} are distinct.

In other words, the sequence is of the form

$$(3.3.4) \quad (L_1, \dots, L_j, \sigma_{-r}, \dots, \sigma_{-1}, L'_1, \dots, L'_k, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$$

where $\sigma_i \leq \sigma_{i+1}$. In this paper, we shall only consider the cases $j + k \leq 2$.

We have a map of labelling sets

$$(3.3.5) \quad \Sigma \amalg \Sigma^\phi \amalg \mathcal{A} \rightarrow Q \amalg Q^\phi \amalg \mathcal{A},$$

so that any curve $S \in \overline{\mathcal{R}}_\Upsilon$ is equipped with moving Lagrangians boundary conditions along the boundary components $\partial_\Sigma S$ and $\partial_\Sigma^\phi S$ by the construction of Section 3.2. By abuse of notation, we write

$$(3.3.6) \quad q_\Upsilon: \partial_\Sigma \overline{\mathcal{S}}_\Upsilon \rightarrow Q$$

$$(3.3.7) \quad q_\Upsilon^\phi: \partial_\Sigma^\phi \overline{\mathcal{S}}_\Upsilon \rightarrow Q$$

for these paths.

Next, we fix an element $\sigma \in \Sigma$, and consider an ordered subset Υ of $\Sigma_\sigma \amalg \Sigma_\sigma^\phi \amalg \mathcal{A}$, such that

- the image $\pi\Upsilon$ under the map $\Sigma \rightarrow Q$ is one of the sequences in Equations (3.1.10)–(3.1.14), and such that the corresponding ordered subsets of Σ and Σ^ϕ respect the partial ordering.

We write $\overline{\mathcal{R}}_\Upsilon$ for $\overline{\mathcal{R}}_{\pi\Upsilon}$, and shall set $q_* = q_\sigma$ in all future constructions. We have the same product decomposition as in Equation (3.1.15), replacing Υ_v^* by Υ_v^σ , i.e. the sequence obtained by replacing all elements of Σ and Σ^ϕ by σ . We also impose the analogue of

Equation (3.1.16), and set the labels of the incoming ends to be either $\Upsilon_e = (L, \sigma)$ or $\Upsilon_e = (\sigma, L)$.

3.3.1. Families of almost complex structures. For each non-degenerate pair Υ in $\Sigma \amalg \Sigma^\phi \amalg \mathcal{A}$, we fix a map

$$(3.3.9) \quad J(\Upsilon): [0, 1] \rightarrow \mathcal{J}$$

which is constant outside of $\nu_X \Upsilon$ with value J_Υ , and agrees with J_L on an endpoint corresponding to $L \in \mathcal{A}$. Recall that we set $\nu_X \Upsilon = X$ if Υ is a subset of \mathcal{A} , so that there is no obstruction to choosing such a family for a pair (L, L') .

We pick families of almost complex structures

$$(3.3.10) \quad J_\Upsilon: \bar{\mathcal{S}}_\Upsilon \rightarrow \mathcal{J}$$

$$(3.3.11) \quad J_{\underline{\Upsilon}}: \bar{\mathcal{S}}_{\underline{\Upsilon}} \rightarrow \mathcal{J}$$

subject to the following constraints:

- (1) (Restriction to the boundary) In the thick part, the restriction to the neighbourhoods of the boundary components labelled by elements of $\Sigma \amalg \Sigma^\phi$ agree with those fixed in Equations (3.2.14)-(3.2.15).
- (2) (Restriction to the thin parts) Over each component Θ of the thin part, $J(\Upsilon)$ or $J(\underline{\Upsilon})$ are constant away from $\nu_X \Theta$.
- (3) (Restriction to the ends) Sufficiently far along each end e , $J(\Upsilon)$ and $J(\underline{\Upsilon})$ agree with J_{Υ_e} under a choice of strip-like end (it is important here to keep Equation (3.1.16) into account).
- (4) (Compatibility with gluing) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_\Upsilon$ or $\bar{\mathcal{R}}_{\underline{\Upsilon}}$, the restriction of J_Υ to a neighbourhood of the corresponding boundary stratum asymptotically agrees with the map obtained by gluing. In the case of $\bar{\mathcal{R}}_\Upsilon$, this means that the diagram

$$(3.3.12) \quad \begin{array}{ccc} \partial \bar{\mathcal{S}}_\Upsilon^T \times (-1, 0]^{E(T)} & \longrightarrow & \partial \bar{\mathcal{S}}_\Upsilon \\ & \searrow J_\Upsilon & \downarrow J_\Upsilon \\ & & \mathcal{J} \end{array}$$

commutes up to a perturbation which is supported in a compact subset, and which vanishes to infinite order at the origin.

- (5) (Compatibility with boundary) For each tree T labelling a boundary stratum of $\bar{\mathcal{R}}_\Upsilon$ or $\bar{\mathcal{R}}_{\underline{\Upsilon}}$ and vertex $v \in V(T)$ such that Υ_v is non-degenerate, the restriction of $J(\Upsilon)$ or $J(\underline{\Upsilon})$ to $\bar{\mathcal{S}}_v^\Upsilon$ or $\bar{\mathcal{S}}_v^{\underline{\Upsilon}}$ agree with the pullbacks of $J(\Upsilon_v)$, $J(\Upsilon_v^\sigma)$, or $J(\underline{\Upsilon}_v)$.
- (6) (Forgetful maps) $J(\Upsilon)$ and $J(\underline{\Upsilon})$ are compatible with the forgetful maps associated to repeating elements of Σ or Σ^ϕ .
- (7) (Forgetting ϕ) For each sequence Υ in $\Sigma^\phi \amalg \mathcal{A}$, $J(\Upsilon)$ agrees with $J(\pi \Upsilon)$.

3.3.2. Moduli spaces of pseudoholomorphic discs. Given a choice of Hamiltonian paths H_Υ and almost complex structure J_Υ , we define the moduli space $\bar{\mathcal{R}}(\Upsilon)$ as the space of isomorphism classes consisting of an element $S \in \bar{\mathcal{R}}_\Upsilon$ and a finite energy J_Υ -holomorphic maps u from a pre-stable curve whose stabilisation is S to X , with (moving) Lagrangian boundary conditions given by (i) L along $\partial_L S$, (ii) $\Phi_\Upsilon(X_{q_\Upsilon})$ along $\partial_\Sigma S$, and (iii) $\Phi_\Upsilon(X_{q_\Upsilon^\phi})$ along $\partial_\Sigma^\phi S$.

Returning to the context of Section 3.2.4, we now consider the (topological) energy $E(u)$ defined as in Section 2.6.2. The following result follows immediately from Lemma 3.3 and the proof of Lemma 2.45:

Corollary 3.4. *Every element u in $\overline{\mathcal{R}}(\Upsilon)$ or $\overline{\mathcal{R}}(\underline{\Upsilon})$ satisfies the reverse isoperimetric inequality*

$$(3.3.13) \quad \ell(\partial u / \sim) \leq \frac{3C}{4} E(u) + \text{a constant independent of } u.$$

□

The finite energy condition implies that we have a natural evaluation map

$$(3.3.14) \quad \overline{\mathcal{R}}(\Upsilon) \rightarrow \prod \text{Crit}(\Upsilon_e)$$

where the product is taken over all ends of elements of $\overline{\mathcal{R}}_\Upsilon$, i.e. over all external edges e of the underlying tree T whose label Υ_e is non-degenerate (i.e. not contained in Σ or Σ^ϕ). We have a similar map for $\overline{\mathcal{R}}(\underline{\Upsilon})$, and denote the fibre over a collection $x_e \in \text{Crit}(\Upsilon_e)$ of intersection points by $\overline{\mathcal{R}}(\{x_e\})$ in either case. The following result is then a standard consequence of regularity theory for holomorphic curves; for its statement, we fix the constant codim which vanishes for the moduli spaces $\overline{\mathcal{R}}(\Upsilon)$, and equals the codimension of $\overline{\mathcal{R}}_\Upsilon \subseteq \overline{\mathcal{R}}_\Upsilon$ for the moduli spaces $\overline{\mathcal{R}}(\underline{\Upsilon})$.

Lemma 3.5. *For generic choices of Floer data $(H_\Upsilon, J(\Upsilon))$ and $(H_{\underline{\Upsilon}}, J(\underline{\Upsilon}))$ the moduli space $\overline{\mathcal{R}}(\Upsilon)$ and $\overline{\mathcal{R}}(\underline{\Upsilon})$ are regular. In particular, $\overline{\mathcal{R}}(\{x_e\})$ is a topological manifold with boundary of dimension*

$$(3.3.15) \quad |\Upsilon| - 3 - \text{codim} + \deg(x_{e^{\text{out}}}) - \sum_{e \neq e^{\text{out}}} \deg(x_e),$$

which is stratified by the topological type of the underlying curve, and the interior of each stratum is a smooth manifold.

This space is naturally oriented relative the tensor product of the orientation lines of $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$ with $\delta_{x_{e^{\text{out}}}} \otimes \bigotimes_{e \in E^{\text{inp}}(T_v)} \delta_{x_e}^\vee$. The boundary is covered by the inverse image of the boundary strata of $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$, together with the union, over all ends e of elements of $\overline{\mathcal{R}}_\Upsilon$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$, of the images of the fibre products

$$(3.3.16) \quad \overline{\mathcal{R}}(\Upsilon) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{R}}(\Upsilon_e) \text{ or } \overline{\mathcal{R}}(\underline{\Upsilon}) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{R}}(\Upsilon_e).$$

□

3.4. Parametrised Morse moduli spaces. In Section 2, we constructed various moduli spaces as fibre products of moduli spaces of discs and gradient flow lines, along common evaluation maps to Lagrangian fibres (see [12] and [15]). This construction is not sufficiently flexible in general, as iterated fibre products may not be transverse. The standard solution is to take perturbations of the Floer equations and the gradient flow equations which depend on the abstract configuration being considered. For our applications, it suffices to perturb the gradient flow lines, so we adopt this simplified context.

3.4.1. Morse and Floer edges. Let T be a rooted ribbon tree with a label Υ by elements of $\Sigma \amalg \Sigma^\phi \amalg \mathcal{A}$ as in the previous section. We define a decomposition

$$(3.4.1) \quad E(T) = E^{\text{Fl}}(T) \amalg E^{\text{Mo}}(T)$$

into *Floer* and *Morse* edges as follows:

$$(3.4.2) \quad \begin{aligned} \text{any edge labelled by a pair in } \Sigma \text{ or in } \Sigma^\phi \text{ lies in } E^{\text{Mo}}(T). \text{ All other edges lie in} \\ E^{\text{Fl}}(T). \end{aligned}$$

We have a corresponding decomposition of the set of flags of T into Floer and Morse flags:

$$(3.4.3) \quad F(T) = F^{\text{Fl}}(T) \amalg F^{\text{Mo}}(T).$$

To edges in $E^{\text{Mo}}(T)$, we shall associate (perturbed) Morse flow lines in a Lagrangian fibre of $X \rightarrow Q$. To this end, we set the space of allowable lengths for each edge e of T to be

$$(3.4.4) \quad \overline{\mathcal{T}}_e = \begin{cases} [0, \infty] & e \in E^{\text{int}}(T) \\ \{\infty\} & e \in E^{\text{ext}}(T). \end{cases}$$

For each vertex v , recall that Υ_v denotes the ordered subset of Υ obtained from the labels of the components adjacent to v , starting again with the component to the left of the outgoing edge e_v^{out} . Note that we have a canonical identification between the punctures of an element of $\overline{\mathcal{R}}_{\Upsilon_v}$ and the Floer edges of T_v , and between the boundary marked points and the Morse edges. We define

$$(3.4.5) \quad \overline{\mathcal{TR}}_{T, \Upsilon} := \prod_{v \in V(T)} \overline{\mathcal{R}}_{\Upsilon_v} \times \prod_{e \in E^{\text{Mo}}(T)} \overline{\mathcal{T}}_e$$

Similarly, given a fixed element $\sigma \in \Sigma$, and consider an ordered subset Υ of $\Sigma_\sigma \amalg \Sigma_\sigma^\phi \amalg \mathcal{A}$ satisfying Condition (3.3.8). Referring back to Section 3.1, and setting $V^\sigma(T)$ to be $V^*(T)$ in order to account for the fact that we have chosen σ as a basepoint, we similarly define

$$(3.4.6) \quad \overline{\mathcal{TR}}_{T, \underline{\Upsilon}} := \prod_{v \in M(T)} \overline{\mathcal{R}}_{\Upsilon_v} \times \prod_{v \in V^\sigma(T)} \overline{\mathcal{R}}_{\Upsilon_v^\sigma} \times \prod_{v \notin V^\sigma(T) \cup M(T)} \overline{\mathcal{R}}_{\Upsilon_v} \times \prod_{e \in E^{\text{Mo}}(T)} \overline{\mathcal{T}}_e.$$

Identifying $\overline{\mathcal{T}}_e$ with the interval $[-1, 0]$ via the map $R \mapsto -e^{-R}$, we obtain the structure of a manifold with corners on this space. The corner strata are products of the corner strata of $\overline{\mathcal{R}}_{\Upsilon_v}$, $\overline{\mathcal{R}}_{\Upsilon_v^\sigma}$, and $\overline{\mathcal{R}}_{\Upsilon_v}$ according to the topological type of the corresponding stable disc, and the stratification of $\overline{\mathcal{T}}_e$ according to whether the length is 0, non-zero and finite, or infinite. To have a better description of the boundary strata of $\overline{\mathcal{TR}}_{T, \underline{\Upsilon}}$, we write

$$(3.4.7) \quad \overline{\mathcal{TR}}_{T, \Upsilon}^\sigma := \prod_{v \in V(T)} \overline{\mathcal{R}}_{\Upsilon_v^\sigma} \times \prod_{e \in E^{\text{Mo}}(T)} \overline{\mathcal{T}}_e,$$

for an isomorphic copy of $\overline{\mathcal{TR}}_{T, \Upsilon}$ whenever Υ is a sequence of the form

$$(3.4.8) \quad (L, \sigma_0, \dots, \sigma_\ell)$$

$$(3.4.9) \quad (\sigma_{-\ell}, \dots, \sigma_{-1}, L)$$

$$(3.4.10) \quad (\sigma_{-\ell}, \dots, \sigma_{-1}, \sigma_0^\phi, \dots, \sigma_\ell^\phi),$$

and all elements of Σ above lie in Σ_σ .

Remark 3.6. We stress the difference between $\overline{\mathcal{TR}}_{T, \Upsilon}^\sigma$ and $\overline{\mathcal{TR}}_{T, \Upsilon^\sigma}$. In the first case, the labels for Morse edges are given by pairs of elements of Υ , and hence can include elements of Σ which differ from the basepoint σ . In the second case, all Morse edges have label (σ, σ) .

For each map $T' \rightarrow T$ which collapses internal edges, and v a vertex of T , let T'_v denote the subtree of T' whose vertices are those mapping to v , and whose edges all are edges adjacent to such vertices. We have a natural identification of strata

$$(3.4.11) \quad \overline{\mathcal{TR}}_{T', \Upsilon} \supset \overline{\mathcal{TR}}_{T', \Upsilon}^T \subset \overline{\mathcal{TR}}_{T, \Upsilon}$$

given on one side by the locus where the lengths of all collapsed edges vanish and on the other by the inclusion $\overline{\mathcal{R}}_{\Upsilon_v}^{T_v} \subset \overline{\mathcal{R}}_{\Upsilon_v}$. We obtain the space

$$(3.4.12) \quad \overline{\mathcal{TR}}_{\Upsilon} := \bigcup_{|E^{\text{ext}}(T)|=|\Upsilon|} \overline{\mathcal{TR}}_{T,\Upsilon}/\sim$$

by gluing the spaces $\overline{\mathcal{TR}}_{T,\Upsilon}$ along their common strata. Restricting to trees labelling the boundary strata of $\overline{\mathcal{R}}_{\Upsilon}$, we similarly obtain:

$$(3.4.13) \quad \overline{\mathcal{TR}}_{\underline{\Upsilon}} := \bigcup_{|E^{\text{ext}}(T)|=|\Upsilon|} \overline{\mathcal{TR}}_{T,\underline{\Upsilon}}/\sim.$$

We can stratify the boundary of $\overline{\mathcal{TR}}_{\Upsilon}$ in such a way that the strata are indexed by trees T labelled by Υ :

$$(3.4.14) \quad \overline{\mathcal{TR}}_{\Upsilon}^T := \prod_{v \in V(T)} \overline{\mathcal{TR}}_{T_v, \Upsilon_v}.$$

Such a stratum corresponds to collections where each Floer edge of T can be identified with the node of a broken disc, and where the Morse edges of T have infinite length. In particular, the codimension 1 boundary strata are given by a choice of tree T with $|\Upsilon|$ external edges, together with a distinguished internal edge $e \in E$ which is either (i) the unique internal Floer edge or (ii) the unique internal Morse edge whose length is specified to equal ∞ . Such a stratum is naturally homeomorphic to the product

$$(3.4.15) \quad \overline{\mathcal{TR}}_{T_e^-, \Upsilon_e^-} \times \overline{\mathcal{TR}}_{T_e^+, \Upsilon_e^+},$$

where T_e^- and T_e^+ are the trees whose vertices are those vertices of T whose path to the outgoing edge respectively passes through e or avoids it; the labels Υ_e^\pm are inherited from Υ , and we note that there are distinguished external edges

$$(3.4.16) \quad e_\pm \in E^{\text{ext}}(T_e^\pm)$$

which is the outgoing edge of T_e^- , and the incoming edge of T_e^+ corresponding to e .

We have a similar description for the moduli spaces $\overline{\mathcal{TR}}_{T,\underline{\Upsilon}}$ with constraints, whose boundary strata are given by

$$(3.4.17) \quad \overline{\mathcal{TR}}_{\underline{\Upsilon}}^T := \prod_{v \in M(T)} \overline{\mathcal{TR}}_{T_v, \Upsilon_v} \times \prod_{v \in V^\sigma(T)} \overline{\mathcal{TR}}_{T_v, \Upsilon_v}^\sigma \times \prod_{v \notin V^\sigma(T) \cup M(T)} \overline{\mathcal{TR}}_{T_v, \Upsilon_v},$$

where T is required to label a boundary stratum of $\overline{\mathcal{R}}_{\underline{\Upsilon}}$. With the exception of strata corresponding to the boundary of $\overline{\mathcal{R}}_{3,2}$ in the cases corresponding to Equations (3.1.13) and (3.1.14), the boundary strata again have either a unique Floer edge, or a unique Morse edge of infinite length.

Remark 3.7. Since we have chosen to give a uniform description, our construction in the case Υ is a subset of Q is slightly more complicated than strictly necessary for the intended application of construction an A_∞ category with objects given by a sufficiently fine covering of Q : the moduli space $\overline{\mathcal{R}}_{\Upsilon}$ can be thought of as the union of the moduli space of discs with points marked by successive elements of Υ with collars of all its boundary strata. By collapsing the moduli spaces of discs to a point, one obtains a map from $\overline{\mathcal{TR}}_{\Upsilon}$ to the corresponding Stasheff associahedron labelling metric ribbon trees (the metric corresponds to the collar coordinate). Having assumed that the Lagrangian fibres do not bound any

holomorphic discs, the construction of the A_∞ category can proceed entirely by Morse-theoretic methods, as in [15, 3]. As a consequence of Remark 3.12 below, our construction ultimately yields an isomorphic category (i.e. with A_∞ operations that are equal).

We associate to each Morse edge of T a universal interval over $\overline{\mathcal{TR}}_{T,\Upsilon}$: first, we define $\overline{\mathcal{I}}_r$ for $r \in [0, \infty)$ to be the interval $[0, r]$. We extend this to $r = \infty$, by setting

$$(3.4.18) \quad \overline{\mathcal{I}}_\infty := \mathcal{I}_+ \amalg \mathcal{I}_-.$$

We equip the space

$$(3.4.19) \quad \overline{\mathcal{I}}_{[0, \infty]} := \coprod_{r \in [0, \infty]} \overline{\mathcal{I}}_r$$

with the natural topology away from $r = \infty$, extended to the latter by the requirement that, for each positive real number R , the sets

$$(3.4.20) \quad \coprod_{R < r \leq \infty} [0, R) \text{ and } \coprod_{R < r \leq \infty} (-R, 0]$$

be open, where on the right side we identify the interval $(-R, 0]$ with $(r - R, r]$ over $r \in (R, \infty]$. We then define

$$(3.4.21) \quad \overline{\mathcal{I}}_e := \begin{cases} \overline{\mathcal{I}}_{[0, \infty]} & e \in E^{\text{int}}(T) \\ [0, \infty) & e \in E^{\text{inp}}(T) \\ (-\infty, 0] & e \in E^{\text{out}}(T) \end{cases}$$

and note that we have a natural projection map

$$(3.4.22) \quad \overline{\mathcal{I}}_e \rightarrow \overline{\mathcal{I}}$$

for all Morse edges of T .

Pulling back the universal interval over $\overline{\mathcal{I}}_e$ yields the universal interval associated to the edge e :

$$(3.4.23) \quad \overline{\mathcal{I}}_{T,\Upsilon}^e \rightarrow \overline{\mathcal{TR}}_{T,\Upsilon}.$$

and the union over all Morse edges is the *universal interval over $\overline{\mathcal{TR}}_{T,\Upsilon}$*

$$(3.4.24) \quad \overline{\mathcal{I}}_{T,\Upsilon} \rightarrow \overline{\mathcal{TR}}_{T,\Upsilon},$$

which is again equipped with a natural smooth structure. The two universal intervals over the codimension 1 strata of $\overline{\mathcal{TR}}_{T,\Upsilon}$ are naturally isomorphic. We have an entirely analogous construction of a universal interval

$$(3.4.25) \quad \overline{\mathcal{I}}_{T,\underline{\Upsilon}} \rightarrow \overline{\mathcal{TR}}_{T,\underline{\Upsilon}}.$$

3.4.2. Morse data and moduli spaces for pairs. Given a pair of elements $\Upsilon = (\Upsilon_0, \Upsilon_1)$ of Σ or Σ_ϕ , let $\nu_Q \Upsilon$ denote the intersection $\nu_Q \Upsilon_0 \cap \nu_Q \Upsilon_1$. Recall that $X_{\nu_Q \Upsilon}$ denotes the restriction of the torus bundle $X \rightarrow Q$ to $\nu_Q \Upsilon$; we pick a distinguished fibre X_{q_Υ} which we equip with a metric, and with a Morse–Smale function

$$(3.4.26) \quad f_\Upsilon: X_{q_\Upsilon} \rightarrow \mathbb{R}.$$

Pick in addition a Lagrangian section over $\nu_Q \Upsilon$, as in Remark 2.15, which induces a trivialisation of $X_{\nu_Q \Upsilon}$ by parallel transport. We write

$$(3.4.27) \quad \psi_\Upsilon^{q,q'}: X_q \rightarrow X_{q'}$$

for the induced map of fibres over points $q, q' \in \nu_q \Upsilon$. For simplicity of notation, we may sometimes omit the subscript from the above maps.

We extend the definition of Morse data given in Equation (2.2.40) as follows: for a non-negative real number r , we define

$$(3.4.28) \quad \mathcal{V}_r(\Upsilon) = C^\infty(\bar{\mathcal{I}}_r, C^\infty(X_{q_r}, TX_{q_r})).$$

For the case $r = \infty$, we define

$$(3.4.29) \quad \mathcal{V}_\infty(\Upsilon) = \mathcal{V}_+(\Upsilon) \times \mathcal{V}_-(\Upsilon),$$

and recall that we require that elements of $\mathcal{V}_\pm(\Upsilon)$ correspond to a family of vector fields parametrised by $\bar{\mathcal{I}}_\pm$, which agree with the gradient vector field outside a compact set. We then define

$$(3.4.30) \quad \mathcal{V}_{[0, \infty]}(\Upsilon) := \coprod_{r \in [0, \infty]} \mathcal{V}_r(\Upsilon).$$

Exactly as in Section 2.2.4, we define $\mathcal{T}_\bullet(\Upsilon, \mathcal{V})$, for $\bullet \in \{\pm\} \amalg [0, \infty)$, to be the set of pairs (γ, ξ) , with γ a path in X_{q_r} with domain $\bar{\mathcal{I}}_\bullet$ and $\xi \in \mathcal{V}_\bullet(\Upsilon)$ satisfying the perturbed gradient flow equation $\frac{d\gamma}{dt} = \xi$. We also define

$$(3.4.31) \quad \mathcal{T}_\infty(\Upsilon, \mathcal{V}) = \mathcal{T}_-(\Upsilon, \mathcal{V}) \times_{\text{Crit}_\Upsilon} \mathcal{T}_+(\Upsilon, \mathcal{V}),$$

and obtain the union

$$(3.4.32) \quad \mathcal{T}_{[0, \infty]}(\Upsilon, \mathcal{V}) = \coprod_{r \in [0, \infty]} \mathcal{T}_r(\Upsilon).$$

By adding to $\mathcal{T}_\bullet(\Upsilon, \mathcal{V})$ the fibre products at the ends with the space $\bar{\mathcal{T}}(\Upsilon)$ of gradient flow lines, we obtain a fibrewise compactification of $\mathcal{T}_\bullet(\Upsilon, \mathcal{V})$ over \mathcal{V}_\bullet , which we denote $\bar{\mathcal{T}}_\bullet(\Upsilon, \mathcal{V})$. In the case $\bullet = \infty$, the boundary stratum is given by

$$(3.4.33) \quad \mathcal{T}_-(\Upsilon, \mathcal{V}) \times_{\text{Crit}_\Upsilon} \bar{\mathcal{T}}(\Upsilon) \times_{\text{Crit}_\Upsilon} \mathcal{T}_+(\Upsilon, \mathcal{V}).$$

Taking the union over $r \in [0, \infty]$, we obtain the space

$$(3.4.34) \quad \mathcal{T}_{[0, \infty]}(\Upsilon, \mathcal{V}) \subset \bar{\mathcal{T}}_{[0, \infty]}(\Upsilon, \mathcal{V})$$

which maps to $\mathcal{V}_{[0, \infty]}(\Upsilon)$ with compact fibres.

3.4.3. Morse moduli spaces for labelled trees. Let T be a tree labelled as in Section 3.1. For each edge in $E^{\text{Mo}}(T)$, we define

$$(3.4.35) \quad \mathcal{V}_e(\Upsilon) := \begin{cases} \mathcal{V}_\pm(\Upsilon_e) \times \bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon} & \text{if } e \text{ is external} \\ \mathcal{V}_{[0, \infty]}(\Upsilon_e) \times_{\bar{\mathcal{T}}_e} \bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon} & \text{otherwise.} \end{cases}$$

In particular, for any edge e , there is thus a natural projection map

$$(3.4.36) \quad \mathcal{V}_e(\Upsilon) \rightarrow \bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon}.$$

The fibre product of these spaces over $\bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon}$ for all $e \in E^{\text{Mo}}(T)$ defines the space of Morse data

$$(3.4.37) \quad \mathcal{V}_T(\Upsilon) \rightarrow \bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon}.$$

We can define spaces $\mathcal{V}_T^\sigma(\Upsilon) \rightarrow \bar{\mathcal{T}}\mathcal{R}_{T, \Upsilon}^\sigma$ and $\mathcal{V}_T(\underline{\Upsilon}) \rightarrow \bar{\mathcal{T}}\mathcal{R}_{T, \underline{\Upsilon}}$ in exactly the same way.

A section of $\mathcal{V}_T(\Upsilon)$ thus consists of a choice of perturbed gradient flow equation for each edge $e \in E^{\text{Mo}}(T)$, metrised according to the image of this point in $\overline{\mathcal{T}}_T$, and hence corresponds to a map

$$(3.4.38) \quad \overline{\mathcal{I}}_e \rightarrow C^\infty(X_{q_{\Upsilon_e}}, TX_{q_{\Upsilon_e}}).$$

for each Morse edge e . We shall assume that such a map is smooth, and moreover agrees to infinite order, along the boundary of the moduli space, with the map obtained in gluing.

We shall moreover pick the data consistently in the following sense: identifying the elements of $E(T) \setminus \{e\}$ and $E(T/e)$, we assume that the following diagram commutes for each Morse edge e' of T :

$$(3.4.39) \quad \begin{array}{ccc} \overline{\mathcal{I}}_{T,\Upsilon}^{e'} | \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^{T/e} & \longrightarrow & \overline{\mathcal{I}}_{T/e,\Upsilon}^{e'} | \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^{T/e} \\ & \searrow & \downarrow \\ & & C^\infty(X_{q_{\Upsilon_{e'}}}, TX_{q_{\Upsilon_{e'}}}). \end{array}$$

In addition, for each Morse edge e of T , and for each edge $e' \neq e_\pm$ of T_e^\pm (with corresponding edges $e' \in E(T)$), we require the commutativity of the diagram:

$$(3.4.40) \quad \begin{array}{ccc} \overline{\mathcal{I}}_{T,\Upsilon}^{e'} | \overline{\mathcal{T}}\mathcal{R}_{T_e^-, \Upsilon_e^-} \times \overline{\mathcal{T}}\mathcal{R}_{T_e^+, \Upsilon_e^+} & \longrightarrow & \overline{\mathcal{I}}_{T_e^\pm, \Upsilon_e^\pm}^{e'} | \overline{\mathcal{T}}\mathcal{R}_{T_e^\pm, \Upsilon_e^\pm} \\ & \searrow & \downarrow \\ & & C^\infty(X_{q_{\Upsilon_{e'}}}, TX_{q_{\Upsilon_{e'}}}). \end{array}$$

We impose the analogous condition that the restrictions of the Morse data associated to e agree, upon restriction to this boundary stratum, with the pullbacks of the data associated to e_\pm on the respective components.

Such sections determine moduli spaces

$$(3.4.41) \quad \overline{\mathcal{T}}_e(\Upsilon) \rightarrow \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon},$$

for each Morse edge e , which compactify the parametrised moduli space of (perturbed) gradient solutions corresponding to e . The same construction yields moduli spaces

$$(3.4.42) \quad \overline{\mathcal{T}}_e^\sigma(\Upsilon) \rightarrow \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^\sigma$$

$$(3.4.43) \quad \overline{\mathcal{T}}_e(\underline{\Upsilon}) \rightarrow \overline{\mathcal{T}}\mathcal{R}_{T,\underline{\Upsilon}},$$

whenever e is a Morse edge of a tree labelling a boundary stratum of $\overline{\mathcal{R}}_\Upsilon^\sigma$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$.

Remark 3.8. Note that, while we have an identification $\overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^\sigma \cong \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon^\sigma}$, the spaces $\overline{\mathcal{T}}_e^\sigma(\Upsilon)$ and $\overline{\mathcal{T}}_e(\Upsilon^\sigma)$ are spaces of perturbed gradient flow lines for Morse functions which are a priori different, over fibres which themselves may be different.

For each flag f containing e , we have an evaluation map

$$(3.4.44) \quad \overline{\mathcal{T}}_e(\Upsilon) \rightarrow X_{q_{\Upsilon_e}}^f,$$

while for each infinite end of e , we have an evaluation map

$$(3.4.45) \quad \overline{\mathcal{T}}_e(\Upsilon) \rightarrow \text{Crit}_{\Upsilon_e},$$

so we obtain a decomposition of $\overline{\mathcal{T}}_e(\Upsilon)$ into components $\overline{\mathcal{T}}_e(x; \Upsilon)$ labelled by such critical points.

Lemma 3.9 (c.f. Theorem 1.4 of [15]). *For generic choices of admissible Morse data, the moduli spaces $\overline{\mathcal{T}}_e(\Upsilon)$, $\overline{\mathcal{T}}_e^\sigma(\Upsilon)$, and $\overline{\mathcal{T}}_e(\underline{\Upsilon})$ associated to an internal Morse edge e are manifolds with boundary of dimension equal to*

$$(3.4.46) \quad n + |\Upsilon| - 3 - \text{codim}$$

where codim is as in Lemma 3.5. This manifold is naturally oriented relative the tensor product of the orientation line

$$(3.4.47) \quad |X_{q_{\Upsilon_e}}| := \Lambda^n TX_{q_{\Upsilon_e}}$$

of $X_{q_{\Upsilon_e}}$ with the underlying abstract moduli spaces, and has boundary given by the inverse image of their boundaries under the evaluation map.

If e is a external edge, then $\overline{\mathcal{T}}_e(x; \Upsilon)$, $\overline{\mathcal{T}}_e^\sigma(x; \Upsilon)$, and $\overline{\mathcal{T}}_e(x; \underline{\Upsilon})$ are manifolds with boundary of dimension

$$(3.4.48) \quad \begin{cases} n - \deg(x) + |\Upsilon| - 3 - \text{codim} & e \in E^{\text{inp}}(T) \\ \deg(x) + |\Upsilon| - 3 - \text{codim} & e \in E^{\text{out}}(T) \end{cases}$$

which are naturally oriented relative the tensor product of the orientation line of the underlying moduli space with $|X_{q_{\Upsilon_e}}| \otimes \det_x^\vee$ in the first case, and \det_x in the second case (see Equation (2.2.23)). The inverse image of each stratum of the underlying abstract moduli space is again naturally a submanifold, of the same codimension. There is an additional boundary stratum of codimension 1, given for $\overline{\mathcal{T}}(\Upsilon)$ by

$$(3.4.49) \quad \overline{\mathcal{T}}(\Upsilon_e) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{T}}_e(\Upsilon)$$

and similarly for $\overline{\mathcal{T}}_e^\sigma(\underline{\Upsilon})$ and $\overline{\mathcal{T}}_e(\underline{\Upsilon})$.

Sketch of proof. Whenever e is an internal edge, the interior of $\overline{\mathcal{T}}_e(\Upsilon)$, $\overline{\mathcal{T}}_e^\sigma(\Upsilon)$, and $\overline{\mathcal{T}}_e(\underline{\Upsilon})$ is by construction the product of $X_{q_{\Upsilon_e}}$ with the (interior of) moduli space $\overline{\mathcal{TR}}_{T, \Upsilon}$, and hence has dimension given by Equation (3.4.46): the identification goes via picking a section of the universal interval over the interior of $\overline{\mathcal{TR}}_{T, \Upsilon}$, which allows us to describe the moduli space as the space of solutions to a family of ODE's (prescribed by a point in $\overline{\mathcal{TR}}_{T, \Upsilon}$) with initial conditions given by a point in $X_{q_{\Upsilon_e}}$. Whenever the chosen Morse function is Morse-Smale, transversality of ascending and descending manifolds implies that the inverse images of each codimension k boundary stratum of $\overline{\mathcal{T}}_e(\Upsilon)$, $\overline{\mathcal{T}}_e^\sigma(\Upsilon)$, and $\overline{\mathcal{T}}_e(\underline{\Upsilon})$ is a manifold of dimension $n + |\Upsilon| - 3 - (\text{codim} + k)$. Gluing of (half)-gradient flow lines establishes that the union of the interior with these boundary strata is a topological manifold with boundary (in fact, it can be equipped with the structure of a smooth manifold with corners, but we shall not use this).

The argument for external edges is entirely analogous, keeping in mind that \det_x is the orientation line of the stable manifold of the critical point x , whose dimension is given by $\deg(x)$, and that the direct sums of the tangent spaces of the stable and unstable manifolds of x is naturally isomorphic to $TX_{q_{\Upsilon_e}}$, so that the unstable manifold has orientation line isomorphic to $|X_{q_{\Upsilon_e}}| \otimes \det_x^\vee$. Equation (3.4.49) accounts for the boundary stratum corresponding to breaking of gradient trajectories along either end. \square

Consistency of the choices of data implies that the two moduli spaces over each codimension 1 boundary stratum of the boundary are naturally identified: in particular if e is an internal edge, then the inverse image in $\overline{\mathcal{T}}_e(\Upsilon)$ of the stratum $\ell(e) = \infty$, which is

homeomorphic to the product $\overline{\mathcal{TR}}_{T_e^-, \Upsilon_e^-} \times \overline{\mathcal{TR}}_{T_e^+, \Upsilon_e^+}$, is naturally identified with the fibre product

$$(3.4.50) \quad \overline{\mathcal{T}}_{e^-}(\Upsilon_e^-) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{T}}_{e^+}(\Upsilon_e^+),$$

along the evaluation maps associated to the edges e_\pm (see Equation (3.4.16)). We have the same type of decomposition for boundary strata associated to edges of $\overline{\mathcal{T}}^\sigma(\underline{\Upsilon})$ and $\overline{\mathcal{T}}_e(\underline{\Upsilon})$.

3.5. Mixed moduli spaces. Let T be a rooted ribbon tree labelled by Υ . Given a choice of Floer data as in Section 3.3 and Morse data as in Section 3.4, we obtain moduli spaces of holomorphic discs for all vertices and of gradient flow lines for all Morse edges. We shall define mixed moduli spaces as fibre products with respect to evaluation maps to the fibres. In order to specify these evaluation maps, we need to make some additional choices:

For any Morse flag f of a tree T labelling a boundary stratum of $\overline{\mathcal{R}}_\Upsilon$, $\overline{\mathcal{R}}_\Upsilon^\sigma$ or $\overline{\mathcal{R}}_{\underline{\Upsilon}}$, we pick maps

$$(3.5.1) \quad q_\Upsilon^f: \overline{\mathcal{TR}}_{T, \Upsilon} \rightarrow \nu_Q \Upsilon_f$$

$$(3.5.2) \quad q_\Upsilon^{f, \sigma}: \overline{\mathcal{TR}}_{T, \Upsilon}^\sigma \rightarrow \nu_Q \Upsilon_f$$

$$(3.5.3) \quad q_{\underline{\Upsilon}}^f: \overline{\mathcal{TR}}_{T, \underline{\Upsilon}} \rightarrow \nu_Q \Upsilon_f$$

subject to the following constraints:

- (1) (Value at marked points) For each vertex v such that Υ_v is degenerate, the values of q_Υ^f , $q_\Upsilon^{f, \sigma}$, or $q_{\underline{\Upsilon}}^f$ for flags containing v agree. If Υ_v is non-degenerate, $q_\Upsilon^{f, \sigma}$ is constant with value q_σ , while q_Υ^f (or $q_{\underline{\Upsilon}}^f$) agrees with the value of the path q_{Υ_v} or $q_{\Upsilon_v}^\phi$ (resp. $q_{\underline{\Upsilon}_v}$ or $q_{\underline{\Upsilon}_v}^\phi$) at the corresponding marked point of the surface in $\overline{\mathcal{S}}_{\Upsilon_v}$ (resp. $\overline{\mathcal{S}}_{\underline{\Upsilon}_v}$).
- (2) (Compatibility with boundary) The two values of the maps q_Υ^f and $q_{\underline{\Upsilon}}^f$ over each boundary stratum of $\overline{\mathcal{TR}}_{T, \Upsilon}$, $\overline{\mathcal{TR}}_{T, \Upsilon}^\sigma$, or $\overline{\mathcal{TR}}_{T, \underline{\Upsilon}}$ agree.

To clarify the last condition, recall that we have given a description of each boundary stratum as either a product of lower dimensional moduli spaces, corresponding to the breaking of a disc giving rise to a Floer edge or of a Morse edge having infinite length, or as a boundary stratum common to two different moduli spaces, corresponding to a Morse edge having length 0, or the breaking of a disc giving rise to a Morse edge.

Let X_Υ^f denote the pullback of X to a bundle over $\overline{\mathcal{TR}}_\Upsilon$ under the map q_Υ^f . For each Morse flag $f = (v, e)$, we have a natural evaluation map

$$(3.5.4) \quad \overline{\mathcal{T}}_e(\Upsilon) \rightarrow X_\Upsilon^f$$

obtained by applying ψ_Υ in order to identify the fibres of X_Υ^f with $X_{q_{\Upsilon_e}}$. We also have a map from $\overline{\mathcal{R}}(\Upsilon_v) \times_{\overline{\mathcal{R}}_{\Upsilon_v}} \overline{\mathcal{TR}}_{T, \Upsilon}$ to the same space, obtained by evaluation at the marked point associated to e . Let X_Υ^T define the fibre product over $\overline{\mathcal{TR}}_\Upsilon$ of the spaces X_Υ^f for all Morse flags over $\overline{\mathcal{TR}}_\Upsilon$. We then define the *mixed moduli space* as the fibre product

$$(3.5.5) \quad \begin{array}{ccc} \overline{\mathcal{TR}}_T(\Upsilon) & \longrightarrow & \prod_{e \in E^{\text{Mo}}(T)} \overline{\mathcal{T}}_e(\Upsilon) \times \prod_{e \in E^{\text{Fl}}(T)} \text{Crit}(\Upsilon_e) \\ \downarrow & & \downarrow \\ \prod_{v \in V(T)} \overline{\mathcal{R}}(\Upsilon_v) \times_{\overline{\mathcal{R}}_{\Upsilon_v}} \overline{\mathcal{TR}}_{T, \Upsilon} & \longrightarrow & X_\Upsilon^T \times \prod_{f \in F^{\text{Fl}}(T)} \text{Crit}(\Upsilon_f). \end{array}$$

Unwinding the definition of the fibre product, we find that an element of $\overline{\mathcal{TR}}_T(\Upsilon)$ consists of (i) a point in $\overline{\mathcal{TR}}_{T,\Upsilon}$, (ii) a gradient flow line in $\overline{\mathcal{T}}_e$ over this point in $\overline{\mathcal{TR}}_{T,\Upsilon}$ for each Morse edge e in T , (iii) an intersection point of the corresponding Lagrangians at each Floer edge e of T , and (iv) a holomorphic map in $\overline{\mathcal{R}}(\Upsilon_v)$ over the projection to $\overline{\mathcal{R}}_{\Upsilon_v}$. At each Morse flag $v \in e$, we require that the evaluations of the gradient flow line in $\overline{\mathcal{T}}_e(\Upsilon)$ (at the end of e) and of the curve in $\overline{\mathcal{R}}(\Upsilon_v)$ (at the marked point corresponding to e) agree as points in X_Υ^f , while at each Floer flag $v \in e$ we require the asymptotic conditions for the elements of $\overline{\mathcal{R}}(\Upsilon_v)$ to be given by the intersection point of Lagrangians chosen for e .

Taking the same fibre product construction yield moduli spaces $\overline{\mathcal{TR}}_T^\sigma(\Upsilon) \rightarrow \overline{\mathcal{TR}}_{T,\Upsilon}^\sigma$, as well as moduli spaces with constraints

$$(3.5.6) \quad \overline{\mathcal{TR}}_T(\Upsilon) \rightarrow \overline{\mathcal{TR}}_{T,\Upsilon}.$$

Taking the evaluation map at all external edges yields the map

$$(3.5.7) \quad \overline{\mathcal{TR}}_T(\Upsilon) \rightarrow \prod_{e \in E^{\text{ext}}(T)} \text{Crit}(\Upsilon_e),$$

whose fibre at a collection $x = \{x_e\}$ of critical and intersection points we denote $\overline{\mathcal{TR}}_T(x)$, and similarly for the other two moduli spaces.

Lemma 3.10. *For generic Floer and Morse data, the moduli spaces $\overline{\mathcal{TR}}_T(x)$, $\overline{\mathcal{TR}}_T^\sigma(\Upsilon)$, and $\overline{\mathcal{TR}}_T(\Upsilon)$ are manifolds with boundary of dimension*

$$(3.5.8) \quad |\Upsilon| - 3 - \text{codim} + \deg(x_{e^{\text{out}}}) - \sum_{e \in E^{\text{inp}}(T)} \deg(x_e),$$

where codim is as in Lemma 3.5. This space is naturally oriented relative the tensor product of the tangent space of the underlying moduli space with

$$(3.5.9) \quad \delta_{x_{e^{\text{out}}}} \otimes \bigotimes_{e \in E^{\text{inp}}(T_v)} \delta_{x_e}^\vee.$$

The codimension 1 boundary strata of $\overline{\mathcal{TR}}_T(\Upsilon)$, $\overline{\mathcal{TR}}_T^\sigma(\Upsilon)$, and $\overline{\mathcal{TR}}_T(\Upsilon)$ are given by the inverse images of the boundary strata of the corresponding abstract moduli spaces $\overline{\mathcal{TR}}_{T,\Upsilon}$, $\overline{\mathcal{TR}}_{T,\Upsilon}^\sigma$, and $\overline{\mathcal{TR}}_{T,\Upsilon}$, together with (i) boundary strata associated to each external edge $e \in E^{\text{ext}}(T)$

$$(3.5.10) \quad \begin{cases} \overline{\mathcal{TR}}_T(\Upsilon) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{T}}(\Upsilon_e) & e \in E^{\text{Mo}}(T) \cap E^{\text{ext}}(T) \\ \overline{\mathcal{TR}}_T(\Upsilon) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{R}}(\Upsilon_e) & e \in E^{\text{Fl}}(T) \cap E^{\text{ext}}(T) \end{cases}$$

and (ii) boundary strata associated to each internal flag (e, v) of Floer type

$$(3.5.11) \quad \overline{\mathcal{TR}}_{T_e^-}(\Upsilon_e^-) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{R}}(\Upsilon_e) \times_{\text{Crit}(\Upsilon_e)} \overline{\mathcal{TR}}_{T_e^+}(\Upsilon_e^+).$$

We analogous descriptions for the boundary strata of $\overline{\mathcal{TR}}_T^\sigma(\Upsilon)$ and $\overline{\mathcal{TR}}_T(\Upsilon)$.

Sketch of proof. We have already established that each of the moduli spaces in the bottom right and top left corner of Diagram (3.5.5) are generically topological manifolds with boundary, equipped with smooth structure in the interior of each stratum; to prove the first statement, it suffices to show that one may in addition assume that their fibre product is transverse as stratified-smooth manifolds, i.e. that the restriction to the interior of each stratum is transverse.

The key point is that, in each fibre product over X_Υ^f , one of the factors is a Morse gradient flow line, and that the inhomogeneous data for flow lines is allowed to vary with

respect to the parameter in the abstract moduli space of stable discs and metric trees. The description of the boundary strata for generic choices is then a consequence of the description of the boundary strata of each constituent moduli space: the moduli space $\overline{\mathcal{R}}(\Upsilon_v)$ has boundary strata associated to breaking of pseudo-holomorphic strips at the ends, and the moduli space $\overline{\mathcal{T}}_e(\Upsilon)$ has boundary strata associated to breaking of gradient trajectories. In the case of external edges, the corresponding stratum of the moduli space of maps incorporates a matching condition at one of the ends of the strip or the gradient flow line, corresponding to the fibre product over Crit_{Υ_e} . In the case of internal Morse edges, the breaking of gradient trajectories is incorporated into the length parameter of the corresponding flow line becoming infinite, and so does not give rise to a new boundary stratum. Finally, in the case of an internal Floer edge, we consider the decomposition of the tree T into the trees T_e^- and T_e^+ as in Equation (3.4.15), and obtain a boundary stratum consisting of elements of the moduli spaces $\overline{\mathcal{T}}\mathcal{R}_{T_e^-}(\Upsilon_e^-)$ and $\overline{\mathcal{T}}\mathcal{R}_{T_e^+}(\Upsilon_e^+)$ together with a pseudo-holomorphic strip which matches the asymptotic conditions of these moduli spaces; there are two such boundary strata, corresponding to whether we consider the pseudo-holomorphic strip as having bubbled from the moduli space of discs associated to either of the two vertices to which it is adjacent. \square

Since the data are chosen consistently for different trees T , there are codimension 1 strata which appear in pairs: given a Morse edge e , we have a map

$$(3.5.12) \quad \overline{\mathcal{T}}\mathcal{R}_T(\Upsilon) \mid \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^{T/e} \rightarrow \overline{\mathcal{T}}\mathcal{R}_{T/e}(\Upsilon) \mid \overline{\mathcal{T}}\mathcal{R}_{T,\Upsilon}^{T/e}$$

corresponding on the left to locus where this Morse edge has length 0, and on the right to the locus where it appears from a breaking of holomorphic discs.

Taking the union along the identification induced by Equation (3.5.12) over all T with $d+1$ external edges and no interior edges of Floer type, we obtain a moduli space $\overline{\mathcal{T}}\mathcal{R}(\Upsilon)$ which admits an evaluation map to $\text{Crit}(\Upsilon_e)$ for each external edge as before. The same construction yields moduli spaces $\overline{\mathcal{T}}\mathcal{R}^\sigma(\Upsilon)$ and $\overline{\mathcal{T}}\mathcal{R}(\Upsilon)$. We denote by $\overline{\mathcal{T}}\mathcal{R}(x_0, x_d, \dots, x_1)$ the fibre over points in $\prod_{e \in E^{\text{ext}}} \text{Crit}(\Upsilon_e)$, with x_0 corresponding to the output. For consistency of notation, whenever Υ is a pair of elements of our labelling set, and (x, y) is a pair in $\text{Crit} \Upsilon$, we write

$$(3.5.13) \quad \overline{\mathcal{T}}\mathcal{R}(x, y) := \begin{cases} \overline{\mathcal{T}}(x, y) & \text{if } x \text{ and } y \text{ are critical points of a Morse function} \\ \overline{\mathcal{R}}(x, y) & \text{if } x \text{ and } y \text{ are intersection points of Lagrangians.} \end{cases}$$

Lemma 3.11. *If the asymptotic conditions (x_0, x_d, \dots, x_1) satisfy*

$$(3.5.14) \quad d - 2 - \text{codim} + \deg(x_0) - \sum_{i=1}^d \deg(x_i) = 1,$$

then $\overline{\mathcal{T}}\mathcal{R}(x_0, x_d, \dots, x_1)$ is a 1-dimensional manifold with boundary. The boundary decomposes into strata labelled by the codimension 1 boundary strata of $\overline{\mathcal{T}}\mathcal{R}_\Upsilon$, $\overline{\mathcal{T}}\mathcal{R}_\Upsilon^\sigma$, or $\overline{\mathcal{T}}\mathcal{R}_\Upsilon$, together with the boundary strata

$$(3.5.15) \quad \overline{\mathcal{T}}\mathcal{R}(x_0, y_0) \times \overline{\mathcal{T}}\mathcal{R}(y_0, x_d, \dots, x_1)$$

$$(3.5.16) \quad \overline{\mathcal{T}}\mathcal{R}(x_0, x_d, \dots, y_i, \dots, x_1) \times \overline{\mathcal{T}}\mathcal{R}(y_i, x_i).$$

\square

Remark 3.12. As a special case of the above construction, we consider a subset Υ of Σ . The moduli spaces above will be used to define the A_∞ structure on a Fukaya category with

objects indexed by the elements of Σ . Note that all vertices of a tree labelling a stratum of this moduli space lie in $V^{\deg}(T)$ (see Equation (3.1.5)), which implies, according to the conditions imposed in the construction, that all pseudo-holomorphic discs with boundary on a fibre that appear in the construction of $\overline{\mathcal{TR}}(\Upsilon)$ in this situation are ordinary holomorphic discs with trivial inhomogeneous term, and constant boundary conditions. Since we have assumed that all such discs are constant, we conclude that rigid elements of the moduli space $\overline{\mathcal{TR}}(\Upsilon)$ can only contain holomorphic discs with three marked points, as any moduli space with more than three marked points passing through any point would have higher dimension. Since the moduli spaces of constant discs with three marked points form a copy of the fibre, they can be forgotten from the description of the rigid moduli spaces $\overline{\mathcal{TR}}(\Upsilon)$, which thus agree with the corresponding moduli spaces of rigid gradient trees.

4. CHAIN LEVEL CONSTRUCTIONS

The cover Σ of Q chosen in Section 3.3 will be fixed for the remainder of this paper. Our first task is to refine the global structures from Section 2 to the A_∞ level. We then proceed to the local case, and establish the equivalence between local and global invariants. One minor difference with the cohomological construction is that the local category will have objects labelled by additional choices of basepoints. This leads to a quasi-equivalent category with many more objects, which is easier to make into the target of our functor. The classification of objects up to quasi-isomorphism will be the same as in Section 2.

4.1. Statement of the cochain-level results. We shall begin, in Section 4.2, by proving Proposition 1.8: this means that we first construct the category \mathcal{F} with objects given by elements of Σ , and morphisms given by Equation (2.3.4), from which we see that cohomological category of \mathcal{F} is naturally identified with the category $H\mathcal{F}$ introduced in Section 2. As discussed in Remark 3.12, the structure constants of this category can be expressed entirely in terms of Morse theory, because we have assumed that the fibres bound no non-constant holomorphic discs. Next, associated to an element L of \mathcal{A} , the left and right A_∞ modules \mathcal{L}_L and \mathcal{R}_L over \mathcal{F} have underlying cochain complexes that were introduced in Section 2.4.2, and we shall construct the map from Equation (1.3.5), which we are trying to prove induces an isomorphism on cohomology. As described in the introduction, our strategy for proving that Equation (1.3.5) is an isomorphism proceeds by showing that its composition with Equation (1.4.2) factors as in Diagram (1.6.4). The cochain complexes, morphisms, and the homotopy in this diagram, are also constructed in Section 4.2, before the proof of Proposition 1.8 is given.

It is clear from the statement of Proposition 1.8 that Theorem 1.1 would follow by showing that each of the three arrows which can be composed counterclockwise in Diagram (1.6.4) are isomorphisms. It is easiest to see this for the middle arrow:

Lemma 4.1. *For each $L \in \mathcal{A}$, the left module \mathcal{L}_L admits a filtration whose associated graded modules are isomorphic to left Yoneda modules over \mathcal{F} .*

The above result is proved in Section 4.3. Using the tensor product of A_∞ modules (see, e.g. [24, Section 2] for a discussion that's compatible with our sign conventions), we conclude:

Corollary 4.2. *The natural map*

$$(4.1.1) \quad \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta}) \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \overline{\Delta} \otimes_{\mathcal{F}} \mathcal{L}_L)$$

is a quasi-isomorphism.

Proof. This is a special case of the following general situation: we have a pair \mathcal{L}_0 and \mathcal{L}_1 of left modules over an A_∞ category \mathcal{C} and a bimodule \mathcal{P} , and consider the map

$$(4.1.2) \quad \text{Hom}_{\mathcal{C}}(\mathcal{L}_0, \mathcal{P}) \otimes_{\mathcal{C}} \mathcal{L}_1 \rightarrow \text{Hom}_{\mathcal{C}}(\mathcal{L}_{L_0}, \mathcal{P} \otimes_{\mathcal{C}} \mathcal{L}_1).$$

The Yoneda Lemma implies that the above map is an isomorphism whenever \mathcal{L}_1 is a Yoneda module. A filtration argument then implies that the above map is an isomorphism whenever \mathcal{L}_1 admits a filtration with associated graded modules which are quasi-isomorphic to Yoneda modules, which proves the desired result. \square

As in Section 2.1, the remaining two arrows are shown to be isomorphisms by reducing the global computation to a local one: given an element $\sigma \in \Sigma$, let \mathcal{F}_σ denote the full subcategory of \mathcal{F} with objects $\tau \in \Sigma_\sigma$. We introduce as in Section 2.1 the notion of a *local* A_∞ bimodule, which is an A_∞ bimodule \mathcal{P} over \mathcal{F} with the property that

$$(4.1.3) \quad \mathcal{P}(\sigma, \tau) \text{ is acyclic whenever } P_\tau \cap P_\sigma = \emptyset.$$

By the computations of Sections 2.7.2 and 2.7.3, the left and right modules $\bar{\Delta}(\sigma, -)$ and $\bar{\Delta}(-, \sigma)$ satisfy this assumption. We have the following variant of Lemma 2.8:

Lemma 4.3. *For each left module \mathcal{L} over \mathcal{F} , and each local bimodule \mathcal{P} , the natural maps*

$$(4.1.4) \quad \text{Hom}_{\mathcal{F}}(\mathcal{L}, \mathcal{P})(\sigma) \rightarrow \text{Hom}_{\mathcal{F}_\sigma}(\mathcal{L}, \mathcal{P})(\sigma)$$

$$(4.1.5) \quad \mathcal{P} \otimes_{\mathcal{F}_\sigma} \mathcal{L}(\sigma) \rightarrow \mathcal{P} \otimes_{\mathcal{F}} \mathcal{L}(\sigma)$$

are quasi-isomorphisms.

Proof. We consider the map in Equation (4.1.5): the length filtration of the bar complex, with associated graded complex given by the direct sum of complexes

$$(4.1.6) \quad \mathcal{P}(\sigma_d, \sigma) \otimes_{\Lambda} \mathcal{F}(\sigma_{d-1}, \sigma_d) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_0, \sigma_1) \otimes_{\Lambda} \mathcal{L}(\sigma_0).$$

By assumption, $\mathcal{P}(\sigma_d, \sigma)$ is acyclic unless $\sigma_d \in \Sigma_\sigma$. On the other hand, if $\tau \in \Sigma_\sigma$, and there is a non-trivial morphism $\rho \rightarrow \tau$ in \mathcal{F} , then $\rho \in \Sigma_\sigma$. Thus the only way for the above complex to not be acyclic is if all other elements of the sequence $(\sigma_d, \dots, \sigma_0)$ to lie in Σ_σ , so that the inclusion of bar complexes induces an isomorphism on the cohomology of associated graded groups. The argument for $\text{Hom}_{\mathcal{F}}(\mathcal{L}, \mathcal{P})$ is entirely analogous. \square

Having reduced the main result to a local computation, we now proceed, as suggested by the results of Section 2, by comparing the global constructions to local ones. To this end, we introduce a category Po_σ whose objects are pairs (q, P) consisting of a polytope P contained in the neighbourhood $\nu_Q \sigma$ of P_σ fixed in Section 3, and a point $q \in P$, and whose morphisms are constructed using (perturbed) gradient flow lines, of Morse functions that depend on the choice of pair (q, P) . We construct this category in such a way that we have a strict A_∞ embedding

$$(4.1.7) \quad j: \mathcal{F}_\sigma \rightarrow \text{Po}_\sigma,$$

which at the level of objects assigns to P_τ the pair (q_τ, P_τ) , and at the level of morphisms and compositions is achieved by copying the choice of Morse functions (and perturbations) used in the construction of the non-zero morphisms of \mathcal{F}_σ when constructing the corresponding morphisms and compositions arising in the image of j .

We then construct modules $\mathcal{L}_{L,\sigma}$ and $\mathcal{R}_{L,\sigma}$, as well as a map of right modules

$$(4.1.8) \quad \mathcal{R}_{L,\sigma} \rightarrow \text{Hom}_{\mathcal{F}_\sigma}(\mathcal{L}_{L,\sigma}, \Delta_{\text{Po}_\sigma}).$$

The comparison between local and global constructions is summarised by the following result, from Section 4.4 below:

Proposition 4.4. *There are quasi-isomorphisms of modules over \mathcal{F}_σ*

$$(4.1.9) \quad j^* \Delta_{\text{Po}_\sigma} \rightarrow \overline{\Delta}$$

$$(4.1.10) \quad j^* \mathcal{L}_{L,\sigma} \rightarrow \mathcal{L}_L$$

$$(4.1.11) \quad j^* \mathcal{R}_{L,\sigma} \rightarrow \mathcal{R}_L$$

which fit in homotopy commutative diagrams of right \mathcal{F}_σ modules

$$(4.1.12) \quad \begin{array}{ccccc} j^* \mathcal{R}_{L,\sigma} & \longrightarrow & \text{Hom}_{\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}) & & \\ \downarrow & & \searrow & & \\ \mathcal{R}_L & \longrightarrow & \text{Hom}_{\mathcal{F}_\sigma}(\mathcal{L}_L, \overline{\Delta}) & \longrightarrow & \text{Hom}_{\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, \overline{\Delta}) \end{array}$$

and of left \mathcal{F}_σ modules

$$(4.1.13) \quad \begin{array}{ccc} j^* \Delta_{\text{Po}_\sigma} \otimes_{\mathcal{F}_\sigma} j^* \mathcal{L}_{L,\sigma} & \longrightarrow & j^* \mathcal{L}_{L,\sigma} \\ \downarrow & & \downarrow \\ \overline{\Delta} \otimes_{\mathcal{F}_\sigma} \mathcal{L}_L & \longrightarrow & \mathcal{L}_L. \end{array}$$

This reduces the computations of the global modules over the local category, to computations of the local modules. In Section 4.5, we prove the A_∞ analogues of Proposition 2.36 and Lemma 2.41:

Lemma 4.5. *The natural maps*

$$(4.1.14) \quad j^* \mathcal{R}_{L,\sigma}(\sigma) \rightarrow \text{Hom}_{\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}(\sigma, -))$$

$$(4.1.15) \quad j^* \Delta_{\text{Po}_\sigma}(-, \sigma) \otimes_{\mathcal{F}_\sigma} j^* \mathcal{L}_{L,\sigma} \rightarrow j^* \mathcal{L}_{L,\sigma}(\sigma)$$

are quasi-isomorphisms.

This is the final ingredient which we need in order to prove the main result of this paper:

Proof of Theorem 1.1. The results of [6] imply that the functor is faithful, so that it suffices to show that Equation (1.3.5) is surjective on cohomology. Proposition 1.8 and Corollary 4.2 reduce this to showing that the bi-module $\overline{\Delta}$ behaves like the diagonal bimodule when tensored with \mathcal{L}_L , and that \mathcal{R}_L is the space of left module maps from \mathcal{L}_L to $\overline{\Delta}$. Both statements are reduced by Lemma 4.3 to each local category \mathcal{F}_σ , and reduced further by Proposition 4.4 to the corresponding statement for the local modules $\mathcal{L}_{L,\sigma}$ and $\mathcal{R}_{L,\sigma}$, and the pullback of the diagonal of Po_σ . Lemma 4.5 thus completes the argument. \square

4.2. Global constructions. We start by constructing the category \mathcal{F} : given a sequence of objects $\Upsilon = \{\sigma_i\}_{i=0}^d$, the A_∞ operation

$$(4.2.1) \quad \mu^d: \mathcal{F}(\sigma_{d-1}, \sigma_d) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \rightarrow \mathcal{F}(\sigma_0, \sigma_d)$$

is defined by counting rigid elements of $\overline{\mathcal{TR}}(\Upsilon)$ as follows: given a collection $x = (x_0; x_1, \dots, x_d) \in \text{Crit}(\sigma_0, \dots, \sigma_d)$, we note that every element of $\overline{\mathcal{TR}}(x)$ induces a map

$$(4.2.2) \quad \text{Hom}_\Lambda^c(U_{\sigma_{d-1}, x_d}^{P_{\sigma_{d-1}}}, U_{\sigma_d, x_d}^{P_{\sigma_d}}) \otimes_\Lambda \cdots \otimes_\Lambda \text{Hom}_\Lambda^c(U_{\sigma_0, x_1}^{P_{\sigma_0}}, U_{\sigma_1, x_1}^{P_{\sigma_1}}) \rightarrow \text{Hom}_\Lambda^c(U_{\sigma_0, x_0}^{P_{\sigma_0}}, U_{\sigma_d, x_0}^{P_{\sigma_d}})$$

obtained by parallel transport along all gradient flow line components (the boundaries of all discs are constant in this case). Tensoring with the map on orientation lines induced by Equation (3.5.9), we obtain the map

$$(4.2.3) \quad \mu_u: \mathcal{F}(\sigma_{d-1}, \sigma_d) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \rightarrow \mathcal{F}(\sigma_0, \sigma_d).$$

We define the higher products as the finite sum

$$(4.2.4) \quad \mu^k := \sum_{u \in \overline{\mathcal{TR}}(x)} (-1)^{\mathfrak{X}} \mu_u,$$

where $\mathfrak{X} = 2 - d + \sum \deg x_i$, as in [26]. The proof that these operations satisfy the A_∞ relation goes back to Fukaya and Oh's work [16, 15] (for a discussion incorporating our choices of signs and perturbations, see [3]). We recall that the sign conventions in [26] are associated to assigning to each generator its *reduced* grading.

Given a sequence $\Upsilon = (L, \sigma_1, \dots, \sigma_\ell)$, with $L \in \mathcal{A}$ and $\sigma_i \in \Sigma$, the module structure on \mathcal{L}_L is defined by counting elements of $\overline{\mathcal{TR}}(\Upsilon)$: for each $x = (x_0; x_1, \dots, x_\ell) \in \text{Crit}(\Upsilon)$ and $u \in \overline{\mathcal{TR}}(x)$, we obtain a map

$$(4.2.5) \quad \text{Hom}_\Lambda^c(U_{\sigma_{\ell-1}, x_\ell}^{P_{\sigma_{\ell-1}}}, U_{\sigma_\ell, x_\ell}^{P_{\sigma_\ell}}) \otimes_\Lambda \cdots \otimes_\Lambda \text{Hom}_\Lambda^c(U_{\sigma_1, x_2}^{P_{\sigma_1}}, U_{\sigma_2, x_2}^{P_{\sigma_2}}) \otimes_\Lambda U_{\sigma_1, x_1}^{P_{\sigma_1}} \rightarrow U_{\sigma_\ell, x_0}^{P_{\sigma_\ell}}$$

by parallel transport along the boundary of the disc components as well as along all constituent gradient flow lines of u . Together with the map on orientation lines given by Equation (3.5.9), a rigid element thus induces a map

$$(4.2.6) \quad \mu_u: \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_1, \sigma_2) \otimes_\Lambda \mathcal{L}_L(\sigma_1) \rightarrow \mathcal{L}_L(\sigma_\ell).$$

We define the higher left module structure maps as the sum

$$(4.2.7) \quad \mu_{\mathcal{L}_L}^{\ell-1|1} := \sum_{\substack{y \in \overline{\mathcal{TR}}(x) \\ x \in \text{Crit}(\Upsilon)}} (-1)^{\mathfrak{X}+1} T^{E(u)} \mu_u$$

where $E(u)$ is the sum of the (topological) energies of the underlying strip, and the sign $\mathfrak{X}+1$ accounts for the fact that x_0 and x_1 are assigned their usual gradings, and all other generators their reduced grading.

Lemma 4.6. *The sum in Equation (4.2.7) converges.*

Proof. Fix a path connecting a basepoint on $X_{q_{\sigma_1}}$ with the intersection point of this Lagrangian with each section ι_{σ_i} . This induces a norm on each Floer complex appearing in Equation (4.2.6), and we shall prove that there are only finitely many elements u of the moduli space such that valuation of μ_u is uniformly bounded by any given constant. As μ_u is a composition of maps associated to gradient flow lines and to a holomorphic strip, and the valuation of the maps associated to gradient flow lines is bounded, the valuation of μ_u is thus bounded, up to a constant independent of u , by the product of the diameter of elements of the cover $\{P_\sigma\}_{\sigma \in \Sigma}$ with the length of the boundary of the strip. The latter is bounded by the reverse isoperimetric inequality as in Lemma 3.3, and the result thus follows from Gromov compactness. \square

For subsequent use, we note the following result, which relies essentially on the assumption that the Lagrangians that we consider are graded: Proposition 2.19 implies that the endomorphism A_∞ algebra of each object of \mathcal{F} has cohomology supported in degree 0, and moreover than we have a natural quasi-isomorphism

$$(4.2.8) \quad H\mathcal{F}(\sigma, \sigma) \rightarrow \mathcal{F}(\sigma, \sigma).$$

This implies that, for each A_∞ module over \mathcal{F} , the filtration by degree of the restriction to σ is a filtration by submodules with respect to the action of $H\mathcal{F}(\sigma, \sigma)$.

Lemma 4.7. *For each $\sigma \in \Sigma$, and each Lagrangian L , the associated graded $H\mathcal{F}(\sigma, \sigma)$ -module of $\mathcal{L}_L(\sigma)$ splits as a direct sum of free modules, indexed by the intersection points of L with X_{q_σ} .*

Proof. Since we are considering a single element $\sigma \in \Sigma$, the only holomorphic discs that appear in the computation of the module structure of $\mathcal{L}_L(\sigma)$ over $\mathcal{F}(\sigma, \sigma)$ are constant discs; this implies that the action of $H\mathcal{F}(\sigma, \sigma)$ respects the direct sum decomposition, and the proof that the module is free then follows from Corollary 2.35. \square

We now briefly outline how the remaining operations are constructed: if we consider instead a sequence $\Upsilon = (L, L', \sigma_1, \dots, \sigma_\ell)$, with $L, L' \in \mathcal{A}$ and $\sigma_i \in \Sigma$, then the count of elements of $\overline{\mathcal{TR}}(\Upsilon)$ induces a map

$$(4.2.9) \quad CF^*(L, L') \rightarrow \text{Hom}(\mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_1, \sigma_2) \otimes_\Lambda \mathcal{L}_{L'}(\sigma_1), \mathcal{L}_L(\sigma_\ell)).$$

Regarding signs, we use the convention that the generators of $CF^*(L, L')$ and $\mathcal{L}_L(\sigma_\ell)$ are equipped with unreduced gradings, and all other are equipped with reduced grading.

Fixing L and L' , and considering an arbitrary sequence of elements of Σ , we obtain the map

$$(4.2.10) \quad CF^*(L, L') \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \mathcal{L}_L).$$

Letting $\Upsilon = (\sigma_{-r}, \dots, \sigma_{-2}, L)$, the count of rigid elements of $\overline{\mathcal{TR}}(\Upsilon)$ yields the structure map

$$(4.2.11) \quad \mu_{\mathcal{R}_L}^{1|r-1} : \mathcal{R}_L(\sigma_{-2}) \otimes_\Lambda \mathcal{F}(\sigma_{-3}, \sigma_{-2}) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \rightarrow \mathcal{R}_L(\sigma_{-r})$$

for the right module \mathcal{R}_L . Our sign conventions are again dictated by the requirement that generators of $\mathcal{R}_L(\sigma)$ are equipped with the unreduced grading; this corresponds to the fact that the moduli space $\overline{\mathcal{TR}}(x_{-1}; x_{-r+1}, \dots, x_{-2})$ contributes with sign

$$(4.2.12) \quad \boxplus + 1 + \sum_{i=-r+1}^{-1} \deg(x_i)$$

in the definition of $\mu_{\mathcal{R}_L}^{1|r}$. Considering the case $\Upsilon = (L, \sigma_1, \dots, \sigma_r, L')$ yields the map

$$(4.2.13) \quad \mathcal{R}_{L'}(\sigma_r) \otimes_\Lambda \mathcal{F}(\sigma_{r-1}, \sigma_r) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_1, \sigma_2) \otimes_\Lambda \mathcal{L}_L(\sigma_1) \rightarrow CF^*(L', L).$$

The structure map $\mu_{\overline{\Delta}}^{\ell|1|r}$ of the bimodule $\overline{\Delta}$

$$(4.2.14) \quad \begin{aligned} \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \otimes_\Lambda \overline{\Delta}(\sigma_{-1}, \sigma_0) \otimes_\Lambda \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \\ \rightarrow \overline{\Delta}(\sigma_{-r}, \sigma_\ell). \end{aligned}$$

is obtained by the count of rigid elements of $\overline{\mathcal{TR}}(\Upsilon)$, with $\Upsilon = (\sigma_{-r}, \dots, \sigma_{-2}, \sigma_{-1}, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$. If we consider instead a sequence $\Upsilon = (\sigma_{-r}, \dots, \sigma_{-2}, \sigma_{-1}, L, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$, we obtain a map

$$(4.2.15) \quad \begin{aligned} \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \otimes_\Lambda \mathcal{L}_L(\sigma_0) \otimes_\Lambda \mathcal{R}_L(\sigma_{-1}) \otimes_\Lambda \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \\ \rightarrow \overline{\Delta}(\sigma_{-r}, \sigma_\ell). \end{aligned}$$

The maps associated to all such sequences for fixed L yield the map of bimodules:

$$(4.2.16) \quad \mathcal{L}_L \otimes_\Lambda \mathcal{R}_L \rightarrow \overline{\Delta}.$$

Finally, the sequence $\Upsilon = (L, \sigma_{-s}, \dots, \sigma_{-2}, \sigma_{-1}, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$ yields the map

$$(4.2.17) \quad \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_0, \sigma_1) \otimes_{\Lambda} \overline{\Delta}(\sigma_{-1}, \sigma_0) \otimes_{\Lambda} \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \otimes_{\Lambda} \mathcal{L}_L(\sigma_{-r}) \rightarrow \mathcal{L}_L(\sigma_\ell).$$

Fixing L , we obtain the map of left modules:

$$(4.2.18) \quad \mathcal{L}_L \otimes_{\mathcal{F}} \overline{\Delta} \rightarrow \mathcal{L}_L.$$

We now prove the commutativity of the main diagram:

Proof of Proposition 1.8. Given a sequence $\Upsilon = (L, \sigma_{-r}, \dots, \sigma_{-2}, \sigma_{-1}, L', \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$, the count of rigid elements of $\overline{\mathcal{TR}}(\Upsilon)$ defines a map

$$(4.2.19) \quad \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_0, \sigma_1) \otimes_{\Lambda} \mathcal{L}_{L'}(\sigma_0) \otimes_{\Lambda} \mathcal{R}_{L'}(\sigma_{-1}) \otimes_{\Lambda} \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \otimes_{\Lambda} \mathcal{L}_L(\sigma_{-r}) \rightarrow \mathcal{L}_L(\sigma_\ell),$$

which we rewrite by adjunction as a map

$$(4.2.20) \quad \begin{aligned} \mathcal{R}_{L'}(\sigma_{-1}) \otimes_{\Lambda} \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \otimes_{\Lambda} \mathcal{L}_L(\sigma_{-r}) \\ \rightarrow \text{Hom}(\mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_{\Lambda} \cdots \otimes_{\Lambda} \mathcal{F}(\sigma_0, \sigma_1) \otimes_{\Lambda} \mathcal{L}_{L'}(\sigma_0), \mathcal{L}_L(\sigma_\ell)). \end{aligned}$$

Fixing L and L' , and letting the sequences of elements in Σ vary, we obtain a map

$$(4.2.21) \quad \mathcal{R}_{L'} \otimes_{\mathcal{F}} \mathcal{L}_L \rightarrow \text{Hom}_{\mathcal{F}}(\mathcal{L}_{L'}, \mathcal{L}_L)$$

which is a homotopy between the two compositions in Diagram 1.8. \square

4.3. Perfectness of left modules. The left module \mathcal{Y}_σ associated to each element σ of Σ is given by

$$(4.3.1) \quad \mathcal{Y}_\sigma(\tau) := \mathcal{F}(\sigma, \tau).$$

Because Σ is a partially ordered set, and the morphisms in \mathcal{F} from σ to τ vanish unless σ precedes τ with respect to the partial order, we can associate to each of its objects a different left module \mathcal{M}_σ given by

$$(4.3.2) \quad \mathcal{M}_\sigma(\tau) := \begin{cases} \mathcal{F}(\sigma, \sigma) & \tau = \sigma \\ 0 & \text{otherwise,} \end{cases}$$

where the module structure of $\mathcal{F}(\sigma, \sigma)$ is the obvious one. The next result asserts that these modules can be built from Yoneda modules. For the proof, we introduce the notion of a module being *supported at* $\sigma \in \Sigma$ if its cohomology vanishes at every other object of \mathcal{F} . By definition, the module \mathcal{M}_σ is supported at σ .

Lemma 4.8. *There is an iterated extension of Yoneda modules which is quasi-isomorphic to \mathcal{M}_σ .*

Proof. The proof is by induction on the number of elements larger than σ in the partial ordering of Σ . By definition, the maximal elements correspond to objects σ of \mathcal{F} such that $\mathcal{F}(\sigma, \tau)$ vanishes whenever $\tau \neq \sigma$. This establishes the Lemma in the base case of maximal elements. By induction, we therefore assume that the Lemma holds for all τ larger than a fixed element σ .

The fact that Σ is partially ordered provides a filtration of every left module with subquotient direct sums of modules supported on elements of Σ . Applying this to the Yoneda module \mathcal{Y}_σ , the quasi isomorphism between $\mathcal{F}(\sigma, \tau)$ and $\mathcal{F}(\tau, \tau)$ implies that \mathcal{Y}_σ is filtered by modules quasi-isomorphic to \mathcal{M}_τ for $\sigma \leq \tau$. Applying the induction hypothesis, we conclude that the Yoneda module \mathcal{Y}_σ admits a filtration so that one subquotient is quasi-isomorphic

to \mathcal{M}_σ , and all others are iterated extensions of Yoneda modules. This implies that \mathcal{M}_σ is itself obtained as an iterated extension of Yoneda modules, proving the result. \square

From the above, we conclude:

Proof of Lemma 4.1. Consider the filtration of \mathcal{L}_L associated to the partial order on Σ . The subquotient associated to each object $\sigma \in \Sigma$ is the complex $\mathcal{L}_L(\sigma)$ as a module over $\mathcal{F}(\sigma, \sigma)$. By Lemma 4.7, $\mathcal{L}_L(\sigma)$ admits a filtration as a module over Γ^{P_σ} , with associated graded quotient that is free. Since the inclusion map $\Gamma^{P_\sigma} \rightarrow \mathcal{F}(\sigma, \sigma)$, is a quasi-isomorphism, the conclusion is immediate. \square

4.4. Local constructions. In this section, we revisit the local constructions from Section 2, and lift them to the A_∞ level. We start in Section 4.4.1 by constructing the category Po_σ associated to an element of Σ , which admits \mathcal{F}_σ as an embedded subcategory (recall that this is the full subcategory of \mathcal{F} with objects $\tau \in \Sigma_\sigma$). We then briefly indicate how the constructions of Section 4.2 can be adapted to produce left and right modules over Po_σ associated to each Lagrangian $L \in \mathcal{A}$, and the local duality map from Equation (4.1.8). We then use the moduli spaces constructed in Section 3.5 to prove Proposition 1.8.

4.4.1. The local category of (based) polytopes. We shall give a relatively convoluted definition of the category Po_σ , which ensures that there is a strict A_∞ embedding $\mathcal{F}_\sigma \subset \text{Po}_\sigma$. The basic idea is that we can think of the structure maps of Po_σ as counting either (perturbed) gradient flow segments on X_{q_σ} , glued together along choices of diffeomorphisms with specified isotopies to the identity, or as collections of (perturbed) gradient flow segments on fibres X_q over different points in \mathcal{N}_σ , glued together along identifications of these fibres. The diffeomorphisms $\psi_\sigma^{q, q'}$ associated to the chosen section of X over \mathcal{N}_σ will be essential in comparing these two points of view. In order to achieve an embedding of \mathcal{F}_σ , we also choose a homotopy between the Lagrangian sections associated to each $\tau \in \Sigma_\sigma$, and that associated to σ itself.

The objects of the category Po_σ are pairs (q, P) , where $q \in P \subseteq \mathcal{N}_\sigma$. The choice of basepoint q is of no real consequence; the objects corresponding to all such choices will be quasi-isomorphic, and the choice is only included to give us enough flexibility to produce the desired strict A_∞ embedding.

For each pair $\Upsilon = ((q_0, P_0), (q_1, P_1))$, we choose a fibre q_Υ in \mathcal{N}_σ , a metric g_Υ on X_{q_Υ} , and a Morse–Smale function f_Υ on X_{q_Υ} . We assume that these choices are subject to the following constraint:

(4.4.1) If $(q_0, q_1) = (q_{\sigma_0}, q_{\sigma_1})$, the data $(q_\Upsilon, g_\Upsilon, f_\Upsilon)$ agree with the data used to define the Floer complex $CF^*(\sigma_0, P_0), (\sigma_1, P_1)$ in Equation (2.2.36).

To be more precise, since the cover Σ that we are considering arises from the construction of Section 3, the data that we consider in this situation agree with the choice of Morse data specified in Section 3.4.2.

Remark 4.9. We implicitly assume that, whenever $\tau \neq \rho$, the basepoints q_ρ and q_τ are distinct. Of course, this condition can be easily achieved by generic choices, but more importantly, we can always enlarge the category Po_σ so that we can associate distinct objects of Po_σ to distinct elements Σ_σ , as is required in order for Condition (4.4.1) to hold.

The morphisms are then given by the Morse complexes

$$(4.4.2) \quad \text{Po}_\sigma(\Upsilon) := CM^*(X_{q_\Upsilon}, \text{Hom}_\Lambda^c(U_\sigma^{P_0}, U_\sigma^{P_1}) \otimes \lambda).$$

To define the compositions, we pick, for each sequence Υ of objects of Po_σ , metric tree T labelled by Υ , and discs with marked points $S_v \in \bar{\mathcal{R}}_{\Upsilon_v}$ for each vertex, the following data: (i) perturbed gradient equations on $X_{q_e} := X_{q_{\Upsilon_e}}$ for each edge of T , (ii) a point $q_v \in \mathcal{N}_\sigma$ for each vertex of T , and (iii) a diffeomorphism

$$(4.4.3) \quad \psi^f: X_{q_e} \rightarrow X_{q_v}$$

for each flag $f = (v \in e)$, with a specified homotopy to $\psi_\sigma^{q_e, q_v}$. We again require the perturbations for sequences of objects of Po_σ which are the images of totally ordered subset of Σ_σ to agree with the choices made in Section 3.4.

Remark 4.10. The choice of discs associated to vertices is completely unnecessary, and only appears here to make the construction exactly the same as that for the construction of the A_∞ structure on \mathcal{F} .

Because the path space is contractible, there is no obstruction to making such choices in families, compatibly with degenerations of discs and breaking of edges. In addition, to ensure compatibility with the construction of the category \mathcal{F}_σ , we require that:

(4.4.4) If $\Upsilon = ((q_{\sigma_0}, P_0), \dots, (q_{\sigma_n}, P_n))$, with $\sigma_i \in \Sigma_\sigma$ such that $\sigma_0 \leq \dots \leq \sigma_n$, the choice for any tree labelled by Υ agrees with that in Sections 3.4.3 and 3.5.

The specified homotopy for sequences coming from objects of Σ arises from the isotopy, fixed above, between the Lagrangian sections associated to elements of Σ_σ , and the section for σ itself.

It is now entirely straightforward to see that the counts of such rigid configurations equip Po_σ with the structure of an A_∞ category. We obtain an embedding of \mathcal{F}_σ as a subcategory of Po_σ , corresponding to the objects (q_σ, P_σ) , by using for each pair $\Upsilon = (\tau_0, \tau_1)$ of objects of \mathcal{F}_σ , the isomorphism

$$(4.4.5) \quad CM^* \left(X_{q_\tau}, \text{Hom}_\Lambda^c(U_{\tau_0}^{P_{\tau_0}}, U_{\tau_1}^{P_{\tau_1}}) \otimes \lambda \right) \cong CM^* \left(X_{q_\tau}, \text{Hom}_\Lambda^c(U_\sigma^{P_{\tau_0}}, U_\sigma^{P_{\tau_1}}) \otimes \lambda \right)$$

induced by the choice of isotopy of sections associated to σ and τ_i .

4.4.2. Modules and structure maps. Given $L \in \mathcal{A}$, we have already fixed a Hamiltonian isotopic Lagrangian L_σ which is transverse to X_{q_σ} . Given an element (q, P) of Po_σ , we define

$$(4.4.6) \quad \mathcal{L}_L(q, P) := CF^*(L, (\sigma, P))$$

$$(4.4.7) \quad \mathcal{R}_L(q, P) := CF^*((\sigma, P), L)$$

as in Section 2.4.2. In other words, all the groups are defined using holomorphic strips with Lagrangian boundary conditions (L, X_{q_σ}) .

Let Υ be a sequence of the form

$$(4.4.8) \quad (L, (q_0, P_0), \dots, (q_\ell, P_\ell))$$

$$(4.4.9) \quad ((q_{-r}, P_{-r}), \dots, (q_{-1}, P_{-1}), L)$$

$$(4.4.10) \quad (q_{-r}, P_{-r}), \dots, (q_{-1}, P_{-1}), L, (q_0, P_0), \dots, (q_\ell, P_\ell)).$$

As in Section 3.3, let Υ^σ denote the sequence obtained from Υ by replacing all objects of Po_σ by the element σ of Σ . We have a corresponding moduli space $\bar{\mathcal{R}}(\Upsilon^\sigma)$ of holomorphic discs boundary conditions L and X_{q_σ} , with two ends (adjacent to the segment labelled L), equipped with a collection of boundary marked points.

As in Section 3.4.1, we consider a moduli space $\bar{\mathcal{T}}\bar{\mathcal{R}}_\Upsilon^\sigma$ of discs and metric trees, which has a top-dimensional stratum of the form $\bar{\mathcal{R}}_{\Upsilon^\sigma}^T \times \prod_{e \in E^{\text{Mo}}(T)} \bar{\mathcal{T}}_e$ for each tree T with label Υ . We

then proceed to inductively choose Morse data on all edges of the tree T , compatibly with the choice of Morse data on the trees which define the A_∞ structure. As in Equation (3.5.5), we obtain a mixed moduli space $\overline{\mathcal{TR}}_T^\sigma(\Upsilon)$ as a fibre product of moduli spaces $\overline{\mathcal{R}}(\Upsilon^\sigma)$ and the moduli spaces of (perturbed) gradient trajectories over the common evaluation maps to Lagrangian torus fibres. For sequences Υ of the form

$$(4.4.11) \quad (L, (q_{\sigma_0}, P_{\sigma_0}), \dots, (q_{\sigma_\ell}, P_{\sigma_\ell}))$$

$$(4.4.12) \quad ((q_{\sigma_{-r}}, P_{\sigma_{-r}}), \dots, (q_{\sigma_{-1}}, P_{\sigma_{-1}}), L)$$

$$(4.4.13) \quad (q_{\sigma_{-r}}, P_{\sigma_{-r}}), \dots, (q_{\sigma_{-1}}, P_{\sigma_{-1}}), L, (q_{\sigma_0}, P_{\sigma_0}), \dots, (q_{\sigma_\ell}, P_{\sigma_\ell}),$$

with $\sigma_0 \leq \dots \leq \sigma_\ell$, and $\sigma_{-r} \leq \dots \leq \sigma_{-1}$ ordered subsets of Σ , we assume that the choices agree with those made in Section 3.

The count of rigid elements of the moduli space associated to Equation (4.4.8) defines a map

$$(4.4.14)$$

$$\text{Po}_\sigma((q_{\ell-1}, P_{\ell-1}), (q_\ell, P_\ell)) \otimes_\Lambda \dots \otimes_\Lambda \text{Po}_\sigma((q_0, P_0), (q_1, P_1)) \otimes_\Lambda \mathcal{L}_{L,\sigma}(q_0, P_0) \rightarrow \mathcal{L}_{L,\sigma}(q_\ell, P_\ell),$$

which makes $\mathcal{L}_{L,\sigma}$ into a left module over Po_σ . We similarly obtain the structure of a right module on $\mathcal{R}_{L,\sigma}$ by considering sequences as in Equation (4.4.9).

We note that these modules allow us to formulate a generalisation of Lemma 4.7: namely, consider the cohomological category $H\mathcal{F}_\sigma$, with objects given by elements $\tau \in \Sigma_\sigma$. The morphism spaces in this category were computed in Equation (2.3.16) to be given by affinoid rings, and in particular to be supported in degree 0. Since the morphism spaces in \mathcal{F} are supported in non-negative degree, we have a natural quasi-isomorphism $H\mathcal{F}_\sigma \rightarrow \mathcal{F}_\sigma$, so that we can pull back $\mathcal{L}_{L,\sigma}$ to a module over $H\mathcal{F}_\sigma$. With this in mind, the next result follows from Corollary 2.35:

Lemma 4.11. *The associated graded $H\mathcal{F}_\sigma$ -module of the filtration of $\mathcal{L}_{L,\sigma}$ by degree is quasi-isomorphic to the direct sum, over all intersection points of L with X_{q_σ} , of shifts of the module $\tau \mapsto \Gamma^{P_\tau}$.* \square

Finally, the sequence in Equation (4.4.10) gives rise to a map

$$(4.4.15)$$

$$\begin{aligned} & \text{Po}_\sigma((q_{\ell-1}, P_{\ell-1}), (q_\ell, P_\ell)) \otimes_\Lambda \dots \otimes_\Lambda \text{Po}_\sigma((q_0, P_0), (q_1, P_1)) \otimes_\Lambda \mathcal{L}_{L,\sigma}(q_0, P_0) \otimes_\Lambda \mathcal{R}_{L,\sigma}(q_{-1}, P_{-1}) \\ & \otimes_\Lambda \text{Po}_\sigma((q_{-2}, P_{-2}), (q_{-1}, P_{-1})) \otimes_\Lambda \dots \otimes_\Lambda \text{Po}_\sigma((q_{-r}, P_{-r}), (q_{-r+1}, P_{-r+1})) \rightarrow \text{Po}_\sigma((q_{-r}, P_{-r}), (q_\ell, P_\ell)). \end{aligned}$$

Fixing L , and letting r and ℓ vary, we obtain the map of bimodules

$$(4.4.16) \quad \mathcal{L}_{L,\sigma} \otimes_\Lambda \mathcal{R}_{L,\sigma} \rightarrow \Delta_{\text{Po}_\sigma}.$$

4.4.3. The global-to-local comparison. The purpose of this section is to complete the reduction of the global computations to local ones. The argument is an entirely straightforward use of the moduli spaces $\overline{\mathcal{TR}}(\Upsilon)$ from Section 3.5.

Proof of Proposition 4.4. Fix an element $\sigma \in \Sigma$. We first begin by considering sequences Υ of the form $(L, \sigma_0, \dots, \sigma_\ell)$ or $(\sigma_{-r}, \dots, \sigma_{-1}, L)$, where $\sigma_i \in \Sigma_\sigma$. The counts of rigid elements of the corresponding moduli spaces define maps

$$(4.4.17) \quad \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \dots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \otimes_\Lambda \mathcal{L}_L(q_{\sigma_0}, P_{\sigma_0}) \rightarrow \mathcal{L}_L(\sigma_\ell)$$

$$(4.4.18) \quad \mathcal{R}_L(q_{\sigma_{-1}}, P_{\sigma_{-1}}) \otimes_\Lambda \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_\Lambda \dots \otimes_\Lambda \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \rightarrow \mathcal{R}_L(\sigma_{-r})$$

which are the structures maps of left and right module homomorphisms $j^* \mathcal{L}_L \rightarrow \mathcal{L}_L$ and $j^* \mathcal{R}_L \rightarrow \mathcal{R}_L$. Here, we use the identification

$$(4.4.19) \quad \mathcal{F}(\sigma_i, \sigma_j) \cong \text{Po}_\sigma((q_{\sigma_i}, P_{\sigma_i}), (q_{\sigma_j}, P_{\sigma_j}))$$

giving rise to the A_∞ embedding noted at the end of Section 4.4.1.

Considering instead a sequence $\Upsilon = (\sigma_{-r}, \dots, \sigma_{-2}, \sigma_{-1}, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$, we obtain a map

$$(4.4.20) \quad \begin{aligned} \mathcal{F}(\sigma_{\ell-1}, \sigma_\ell) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_0, \sigma_1) \otimes_\Lambda \text{Po}_\sigma((q_{\sigma_{-1}}, P_{\sigma_{-1}}), (q_{\sigma_0}, P_{\sigma_0})) \\ \otimes_\Lambda \mathcal{F}(\sigma_{-2}, \sigma_{-1}) \otimes_\Lambda \cdots \otimes_\Lambda \mathcal{F}(\sigma_{-r}, \sigma_{-r+1}) \rightarrow \overline{\Delta}(\sigma_{-r}, \sigma_\ell), \end{aligned}$$

which defines the map of bimodules $j^* \Delta_{\text{Po}_\sigma} \rightarrow \overline{\Delta}$.

The homotopies in Diagrams (4.1.12) and (4.1.13) follow in exactly the same way by counting rigid elements of $\overline{\mathcal{TR}}(\Upsilon)$ for sequences:

$$(4.4.21) \quad (L, \sigma_{-r}, \dots, \sigma_{-1}, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi)$$

$$(4.4.22) \quad (\sigma_{-r}, \dots, \sigma_{-1}, L, \sigma_0^\phi, \sigma_1^\phi, \dots, \sigma_\ell^\phi).$$

□

4.5. Local computations. We conclude the main part of the paper by performing to the necessary local computations:

Proof of Lemma 4.5. The two arguments are entirely analogous; we explain the proof of the duality isomorphism in Equation (4.1.14). First, using the quasi-isomorphism $H\mathcal{F}_\sigma \rightarrow \mathcal{F}_\sigma$, we reduce the problem so showing that the map

$$(4.5.1) \quad j^* \mathcal{R}_{L,\sigma}(\sigma) \rightarrow \text{Hom}_{H\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}(\sigma, -))$$

is a quasi-isomorphism. By Lemma 4.11, the decomposition of the intersection points of L with X_{q_σ} by degree provides a filtration of $j^* \mathcal{L}_{L,\sigma}$ as a left module over $H\mathcal{F}_\sigma$, hence of $\text{Hom}_{H\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}(\sigma, -))$ as a cochain complex. It also provides a filtration of $j^* \mathcal{R}_{L,\sigma}(\sigma)$ as a cochain complex. It suffices to prove that the map at the level of each associated graded group is a quasi-isomorphism. Since each associated graded summand of $j^* \mathcal{L}_{L,\sigma}$ is the module which assigns to τ the affinoid ring Γ^{P_τ} , the complex $\text{Hom}_{H\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}(\sigma, -))$ is thus quasi-isomorphic to the Čech complex

$$(4.5.2) \quad \left(\bigoplus_{\tau_0 < \dots < \tau_k} \text{Hom}_\Lambda^c(\Gamma^{P_\sigma}, \Gamma^{\tau_k})[-k], \check{\delta} \right).$$

As in Section 2.5.2, we may replace P_{τ_k} in the above expression by $P_{\tau_k} \cap P$ for a polytope P containing P_σ in its interior, and covered by the elements of Σ_σ . By Tate acyclicity, we conclude that each associated graded group of $\text{Hom}_{H\mathcal{F}_\sigma}(j^* \mathcal{L}_{L,\sigma}, j^* \Delta_{\text{Po}_\sigma}(\sigma, -))$ is quasi-isomorphic to the complex $CF^*(P, P^\sigma)$ whose cohomology is quasi-isomorphic to the linear dual of Γ^{P_σ} by Proposition 2.42. □

APPENDIX A. REVERSE ISOPERIMETRIC INEQUALITIES

The purpose of this appendix is to extend the reverse isoperimetric inequalities established by Groman-Solomon [20] and Duval [11] for holomorphic curves with Lagrangian boundary conditions, to the case of moving Lagrangian boundary conditions. Since our main result requires a rather long list of technical hypotheses, we postpone its statement in the form of Proposition A.5 to the end, and develop the basic ideas one step at a time.

A.1. The basic inequality. Let X be a closed symplectic manifold, S a Riemann surface obtained from a closed Riemann surface with boundary by removing finitely many boundary punctures. Let $K \subset X$ be a closed codimension 0 submanifold with boundary, and assume that we have a labelling

$$(A.1.1) \quad \Upsilon_S: \pi_0(\partial S) \rightarrow \{K\} \cup \mathcal{L}$$

of ∂S by $\{K\} \cup \mathcal{L}$, where \mathcal{L} is a collection of Lagrangians in X . Given $i \in \pi_0(\partial S)$, we write $\partial_i S$ for the corresponding component of the boundary. We write $\partial_K S$ for the union of components labelled by K , and L_i for the Lagrangian labelled by a component i . We assume that all intersections among Lagrangians appearing in \mathcal{L} are contained in the interior of K , and choose another closed subset K' containing this intersection, and contained in the interior of K :

$$(A.1.2) \quad K' \Subset K \subset X.$$

In addition, we fix a Riemannian metric g_X with respect to which we shall compute all norms.

Given an almost complex structure J , consider a family

$$(A.1.3) \quad J_S: S \rightarrow \mathcal{J}$$

of almost complex structures which agree with J away from K' . There is an associated moduli space

$$(A.1.4) \quad \mathcal{M}(\Upsilon_S, J_S)$$

of J_S -holomorphic curves with boundary conditions given by Υ_S in the sense that a point $z \in \partial_K S$ maps to K' , while a point $z \in \partial_i S$ on a component labelled by a Lagrangian maps to L_i .

Recall that the geometric energy $E^{geo}(u)$ of any element u of the moduli space is the area

$$(A.1.5) \quad E^{geo}(u) = \int \|du\|^2.$$

Moreover, given a component i of ∂S , we define

$$(A.1.6) \quad \ell_K(\partial_i u)$$

to be the length, with respect to g_X , of the collection of curves obtained removing $u^{-1}K$ from $\partial_i S$.

Lemma A.1. *For each choice of labels Υ_S , there exists a constant C , depending on J but not on the family J_S , such that for each element $u \in \mathcal{M}(\Upsilon_S, J_S)$ and each component $i \in \pi_0(\partial S)$ which is labelled by a Lagrangian, we have*

$$(A.1.7) \quad \ell_K(\partial_i u) \leq C E^{geo}(u).$$

This constant is independent of the restriction of L_i to K .

Proof. This is a minor modification of the results of Groman-Solomon [20, Theorem 1.1] and Duval [11]. We briefly indicate how to adapt Duval's proof:

We choose a tubular neighbourhood $\nu_X L_i$ of L_i , which is equipped with a non-negative weakly plurisubharmonic function ρ_i (with respect to J) that vanishes on the intersection with $K' \cup L_i$, and does not vanish away from the intersection with $K \cup L_i$. Moreover, away from K , we require that ρ_i be (strictly) plurisubharmonic with weakly plurisubharmonic square root. Near $\partial K' \cap L_i$, such a function can be locally modelled after the function

$$(A.1.8) \quad \chi(x_1) (|y_1|^2 + \cdots + |y_n|^2)$$

for the Lagrangian $[0, \infty) \times \mathbb{R}^{n-1} \subset \mathbb{C}^n$, where χ is a smooth function vanishing for $x_1 \leq 0$. For simplicity, we assume that g_X everywhere dominates the (semi)-metric induced by ρ and J .

Let $\ell_{\rho_i}(\partial u)$ denote the length of ∂u with respect to the (semi)-metric induced by $dd^c \rho_i$, and E_{ρ_i} the integral of $dd^c \rho_i$ over the part of u with image in $\nu_X L_i$. There is a constant C_0 such that

$$(A.1.9) \quad \ell_{\rho_i}(\partial u) \leq C_0 E_{\rho_i}(u).$$

The basic idea is that the function $\frac{E_{\rho_i}(r, u)}{r^2}$ which measures the area of u in the domain $\rho_i \leq r$ is monotonic, and by Fubini's theorem, has limit bounded above by a constant multiple $\ell_{\rho_i}(\partial u)$ (see [11] for details); the fact that ρ_i vanishes to order greater than 1 along K' implies that this part of boundary does not contribute to $\ell_{\rho_i}(\partial u)$.

Away from K , the metric induced by $dd^c \rho_i$ is uniformly comparable to g_X , so we may assume the existence of a constant C_1 such that

$$(A.1.10) \quad \ell_K \partial_i u \leq C_1 \ell_{\rho_i}(\partial u)$$

On the other hand, the assumption that g_X dominates $dd^c \rho_i$ implies that

$$(A.1.11) \quad E_{\rho_i}(u) \leq E^{geo}(u).$$

The result follows from combining these inequalities. \square

For applications, it will be convenient to state the estimate in a different way: define L_i/\sim to be the quotient of L_i by the relation which identifies points in the same component of $K \cap L_i$. Given a length metric on L_i/\sim (see the discussion preceding Lemma 2.26), we have:

Corollary A.2. *There exists a constant C such that, for each element $u \in \mathcal{M}(\Upsilon_S, J_S)$, we have*

$$(A.1.12) \quad \ell(\partial_i u / \sim) \leq C E^{geo}(u).$$

\square

A.2. Moving Lagrangian boundary conditions. It will be necessary to have a generalisation for moving boundary conditions, and families of almost complex structures. Let S be a Riemann surface equipped with a collection of subsets $\Theta \subset S$, which we call the components of the thin part, that include neighbourhoods of all punctures. By abuse of notation, we will refer to this as a thick-thin decomposition, even though the thin part that we consider is much larger than the intrinsic one coming from hyperbolic geometry.

Let Υ_S be moving boundary conditions on S , i.e. a smooth assignment of a subset of X to each point on ∂S , which we assume is locally constant on the thin-part. We shall only consider the situation in which the restriction to each component is either a moving family of Lagrangians, or is a constant family given by a compact subset $K \subset X$ as in the previous section. We assume that any component labelled by K is contained in a component of the thin part.

For each component $i \in \pi_0(\partial S)$ labelled by a moving family of Lagrangians, we denote the graph by $\tilde{L}_i \subset \partial_i S \times X$. By construction, this map is independent of the first factor in the thin part, so that we obtain a Lagrangian $L_{i, \Theta}$ in X for each component Θ of the thin part.

Assume that we are given, for each such component, an almost complex structure J_Θ on X as well as a pair of codimension 0 manifolds with boundary

$$(A.2.1) \quad \nu'_X \Theta \subseteq \nu_X \Theta \subseteq X$$

which contain the intersection of any pair of Lagrangians labelled by boundary components of S which meet Θ . In addition, we assume that $K \subset \nu_X \Theta$ if Θ intersects a component labelled by the subset K . We pick a family J_S of almost complex structures on X parametrised by S , such that

$$(A.2.2) \quad \begin{aligned} &\text{the restriction of } J_S \text{ to the product of a component } \Theta \text{ of the thin part with} \\ &X \setminus \nu'_\Theta X \text{ agrees with } J_\Theta. \end{aligned}$$

Given this data, we shall consider the moduli space

$$(A.2.3) \quad \mathcal{M}(\Upsilon_S, J_S)$$

of J_S holomorphic curves with boundary conditions Υ_S : it is convenient to introduce the almost complex structure \tilde{J}_S on $S \times X$ induced by J_S , and describe this moduli space as the space of \tilde{J}_S holomorphic sections with boundary conditions given by \tilde{L}_i over $\partial_i S$.

Our goal is to prove a reverse isoperimetric inequality for these moduli spaces. The basic idea is that Lemma A.1 provides an estimate for the moduli spaces holomorphic strips corresponding to each end; we shall extend this estimate to S , at the cost of a possibly weaker proportionality constant, as well as the addition of a constant term. The main point is to provide a proof that extends to families of (broken) holomorphic curves.

We shall compare the geometric energy of each curve u

$$(A.2.4) \quad E^{geo}(u) = \int |du|^2 = \int u^* \omega$$

to the length of the boundary, which we formulate as follows: we denote by \tilde{L}_i/\sim the quotient of \tilde{L}_i by the equivalence relation which

$$(A.2.5) \quad \begin{aligned} &\text{(i) collapses each component of the inverse image of } \Theta \text{ in } \tilde{L}_i \text{ to a single fibre,} \\ &\text{and (ii) identifies points in the same component of the intersection of } \tilde{L}_i \\ &\text{with } \nu_X \Theta. \end{aligned}$$

In particular, \tilde{L}_i/\sim maps to the quotient of $\partial_i S$ by the components of the intersection with the thin part, agrees with \tilde{L}_i away from the inverse image of these components, and agrees with the quotient considered in the previous section over each component of the thin part.

Lemma A.3. *There exists a constant C such that, for each $u \in \mathcal{M}(\Upsilon_S, J_S)$ and for each component i of ∂S , we have*

$$(A.2.6) \quad \ell(\partial_i u / \sim) \leq C E^{geo}(u) + \text{a constant independent of } u.$$

The constant depends only on the restriction of the Lagrangian boundary conditions and the almost complex structures to a neighbourhood of $\partial_i S$, but is independent of their restriction to $\Theta \times \nu'_X \Theta$ for each component Θ of the thin part.

Proof. The graph \tilde{L}_i is a totally real submanifold with respect to the almost complex structure \tilde{J}_S which at every point in $z \in S$ is the product of the complex structure on S with the value of the family J_S at x . This almost complex structure is compatible with the symplectic form

$$(A.2.7) \quad \omega_X + \omega_S,$$

where ω_S is a symplectic form on S of total area 1. If we define $E(\tilde{u})$ to be the area with respect to such a symplectic form, we have

$$(A.2.8) \quad E(\tilde{u}) = E^{geo}(u) + 1,$$

which will be one origin for the constant term in the statement of the Lemma. The remainder of the proof proceeds in essentially the same way as that of Lemma A.1:

For each component Θ of the thin part which is adjacent to i , equipped with the family of almost complex structures J_Θ , fix the neighbourhood $\nu_X L_{i,\Theta}$ and the function $\rho_{i,\Theta}: \nu_X L_{i,\Theta} \rightarrow \mathbb{R}$ considered in the proof of Lemma A.1.

By projection to the second factor, we obtain a function on $\Theta \times X$ which we denote $\rho_{i,\Theta}$. We consider a neighbourhood $\nu_{S \times X} \tilde{L}_i$ of \tilde{L}_i which contains the product $\Theta \times L_{i,\Theta}$ for all components Θ of the thin part which meet \tilde{L}_i , and is equipped with a weakly plurisubharmonic function

$$(A.2.9) \quad \rho_i: \nu_{S \times X} \tilde{L}_i \rightarrow [0, \infty)$$

such that the following properties hold

- (1) Over $\Theta \subset S$, ρ_i agrees with $\rho_{i,\Theta}$.
- (2) Away from $\Theta \times \nu_X \Theta$, the function ρ_i is strictly plurisubharmonic, and only vanishes on \tilde{L}_i .
- (3) The square root of ρ_i is everywhere weakly plurisubharmonic.

The existence of such a function follows from the usual patching argument for plurisubharmonic functions as in [11]. For simplicity, we also assume that there is a finite diameter metric on S whose product with g_X everywhere dominates the metric induced by ρ_i .

Given a section \tilde{u} corresponding to an element of $\mathcal{M}(\Upsilon_S, J_S)$, let $\ell_{\rho_i}(\partial_i \tilde{u})$ denote the length of $\partial_i \tilde{u}$ with respect to the (semi)-metric induced by $dd^c \rho_i$, and E_{ρ_i} the integral of $dd^c \rho_i$ over the part of \tilde{u} with image in $\nu_{S \times X} \tilde{L}_i$. As before, there is a constant C_0 such that

$$(A.2.10) \quad \ell_{\rho_i}(\partial_i \tilde{u}) \leq C_0 E_{\rho_i}(u).$$

Away from a small neighbourhood of the inverse image of Θ , the metric induced by $dd^c \rho_i$ is uniformly comparable to the product of g_X with a metric on S . Because ρ_i is non-degenerate in the fibre direction over Θ , we therefore obtain a constant C_1 such that

$$(A.2.11) \quad \ell(\partial_i u / \sim) \leq C_1 \ell_{\rho_i}(\partial_i \tilde{u}) + \text{a constant independent of } u,$$

where the constant term accounts for the fact that the derivatives of ρ_i in the base direction vanish identically on Θ_i .

On the other hand, the assumption that g_X and the metric on S dominate $dd^c \rho_i$ implies that

$$(A.2.12) \quad E_{\rho_i}(u) \leq E(\tilde{u}).$$

The result follows from combining this inequality with Equations (A.2.8), (A.2.10), and (A.2.11). \square

A.3. Compatibility with gluing. We shall require a uniform estimate for families of Riemann surfaces. Let us therefore consider a rooted ribbon tree T , and a collection of Riemann surfaces with punctures S_v for each vertex v of T , with an identification of the ends of S_v with the edges adjacent to v . We write $E^{\text{int}}(T)$ for the edges of T which are adjacent to two vertices. Given gluing parameters

$$(A.3.1) \quad R: E^{\text{int}}(T) \rightarrow [0, \infty]$$

and a choice of disjoint strip-like ends for each end of S_v , we obtain a Riemann surface $S_{T,R}$ by gluing, which is the image of a surjective map with domain the disjoint union of the surfaces S_v . Assuming that each S_v is equipped with a thick-thin decomposition, we obtain a thick-thin decomposition of $S_{T,R}$. We allow the possibility that the decomposition of a component S_v be degenerate, in the sense that the thin part consists of the entire surface.

The ribbon structure determines an embedding $T \subset \mathbb{R}^2$ up to isotopy; assume that we have moving boundary conditions Υ_{S_v} , which are locally constant on all components of the thin part, and are consistent for adjacent vertices, in the sense that the subsets of X assigned to the two ends corresponding to each interior edge of T agree. We write $\Upsilon_{S_{T,R}}$ for the moving boundary conditions on $S_{T,R}$ obtained by gluing. As in the previous section, we pick, for each component Θ of the thin part of a curve S_v , closed subsets $\nu'_X \Theta \subseteq \nu_X \Theta$ of X , which contains the intersections of all labels, and an almost complex structure J_Θ on X ; we assume that these choices are the same for the thin parts meeting the two punctures corresponding to a given edge in T .

Finally, we pick families of almost complex structures J_{S_v} which agree with J_Θ on $\Theta \times X \setminus \nu'_X \Theta$, and are also compatible across the edges of T . Let $J_{S_{T,R}}$ be a family of almost complex structure on $S_{T,R}$ which

(A.3.2) agrees with the family obtained by gluing (i) in the thick part near each boundary stratum, and (ii) in the thin part away from $\Theta \times \nu'_X \Theta$ for each component Θ of the thin part.

Given this data, we shall consider the moduli space

$$(A.3.3) \quad \mathcal{M}(\Upsilon_{S_{T,R}}, J_{S_{T,R}})$$

which is given as in the previous section if all gluing parameters are finite, and is given for infinite gluing parameters by the fibre products, along the evaluation maps for the edges, of the moduli spaces $\mathcal{M}(\Upsilon_{S_v}, H_{S_v}, J_{S_v})$ corresponding to the vertices.

For each $i \in \pi_0(\mathbb{R}^2 \setminus T)$ with Lagrangian label, and finite gluing parameter R , consider the quotient $\tilde{L}_{i,R}/\sim$ of the moving Lagrangian boundary condition over the component of the boundary corresponding to i , by the relation considered in the previous section: collapse each component of the inverse image of the thin part, then identify points in the same component of $\nu'_X \Theta$ in the fibres over a component Θ of the thin part. By construction, the spaces we obtain for different choices of gluing parameters are naturally homeomorphic; we write \tilde{L}_i/\sim for any of these spaces, and note that we have an evaluation map

$$(A.3.4) \quad \partial_i u / \sim : \partial_i S_{T,R} \rightarrow \tilde{L}_i / \sim$$

for any choice of gluing parameter.

Lemma A.4. *There exists a constant C such that, for each $R \in [0, \infty]^{E^{\text{int}}(T)}$ and $u \in \mathcal{M}(\Upsilon_{S_{T,R}}, J_{S_{T,R}})$, we have*

$$(A.3.5) \quad \ell(\partial_i u / \sim) \leq CE^{\text{geo}}(u) + \text{a constant independent of } u \text{ and } R.$$

This constant depends on the restriction of $J_{S_{T,R}}$ to a neighbourhood of the corresponding boundary component, and is independent of the restriction of $(\tilde{L}_i, J_{S_{T,R}})$ to the product of each component Θ of the thin part with $\nu'_X \Theta$.

Proof. It suffices to check that the constants in the proof of Lemma A.3 can be chosen in terms of the corresponding constants for the moduli spaces $\mathcal{M}(\Upsilon_{S_v}, J_{S_v})$ and independently of R . For Equation (A.2.8), this follows from the fact that the choice of symplectic form on $X \times S_V$ induces a symplectic form on $X \times S_{T,R}$. For Equation (A.2.10), we note that

a choice of function ρ_{i,S_v} for all v induces, by gluing, a function $\rho_{S_{T,R}}$, for which we have the same estimate. The same argument applies to Equation (A.2.11), which completes the argument. \square

A.4. Reverse isoperimetric inequality for families. We now assemble the previous discussion into a statement for families of Riemann surfaces: consider a toplogical space P parametrising a family of pre-stable Riemann surfaces with boundary marked points. This means that we are given a space $\bar{S}_P \rightarrow P$, whose fibre \bar{S}_p for $p \in P$ is a compact pre-stable Riemann surface with boundary \bar{S}_p , equipped with finitely many distinct marked points on $\partial\bar{S}_p$. We assume that the conformal structure on \bar{S}_p as a Riemann surface with marked points varies continuously with respect to p . Note that P inherits a stratification by the number of components of \bar{S}_p , and we shall impose the technical condition that each such stratum $\partial^n P$ admits a neighbourhood deformation retract

$$(A.4.1) \quad \pi_n: \nu\partial^n P \rightarrow \partial^n P$$

on which the family of Riemann surfaces is obtained by gluing, in the sense that we are given a map

$$(A.4.2) \quad \nu\partial^n P \rightarrow [0, \infty]^n$$

and we can choose strip-like ends on either side of each node of \bar{S}_p for $p \in \partial^n P$, so that \bar{S}_q for $q \in \nu\partial^n P$ agrees with the result of gluing the components of $\bar{S}_{\pi_n(q)}$ with gluing parameters given by Equation (A.4.2).

We denote by $S_P \rightarrow P$ the space obtained by removing all nodes and all marked points, and fix a collection of subsets $\Theta \subset S_P$ such that each puncture of a fibre S_p lies in one of these subsets, and we require in addition that, whenever p lies in $\nu\partial^n P$, each component of the gluing region is contained in some subset Θ .

Given the above setup, we consider a moving family Υ_{S_P} of boundary conditions on the family S_P , i.e. a map from ∂S_P to the subsets of a closed symplectic manifold X . We assume that we are given a fixed a closed codimension 0 subset $K \subset X$, such that the restriction of Υ_{S_P} to each component of ∂S_P is either constant with value K , or a moving family of Lagrangians. In the second case, we assume that the family of Lagrangians is obtained by gluing near $\partial^n P$.

For each component Θ of the thin part, we consider a continuously varying family of almost complex structures J_{Θ_P} on X , parametrised by P , as well as open subsets

$$(A.4.3) \quad \nu'_X \Theta \subseteq \nu_X \Theta \subseteq X \times P,$$

which, over a point $p \in P$, contain the intersection of the Lagrangian labels of the components of ∂S_p meeting Θ . We finally pick a family J_{S_P} of almost complex structures on X , parametrised by S_P , such that Condition (A.2.2) holds for the restriction to S_p for each $p \in P$.

We now define the moduli space

$$(A.4.4) \quad \mathcal{M}(\Upsilon_{S_P}, J_{S_P}) \equiv \coprod_{p \in P} \mathcal{M}(\Upsilon_{S_p}, J_{S_p}),$$

to be the disjoint union over all elements of p of the moduli spaces of J_{S_p} holomorphic maps with moving Lagrangian boundary conditions considered in Equation (A.2.3). This space can be equipped with a natural topology as a parametrised moduli space, but we shall not require it in our discussion.

We have an assignment of a geometric energy $E^{geo}(u)$ to each element of this moduli space (see Equation (A.2.4)), as well as a reduced length $\ell(\partial_i u / \sim)$ of each boundary component,

obtained by collapsing the part of the boundary lying in the thin part as in Equation (A.2.5). The compatibility of our proof of the reverse isoperimetric inequality in Lemma A.3 with gluing (as discussed in Lemma A.4) implies:

Proposition A.5. *If P is compact, there exists a constant C such that, for each $u \in \mathcal{M}(\Upsilon_{S_P}, J_{S_P})$ and for each component i of ∂S_P , we have*

$$(A.4.5) \quad \ell(\partial_i u / \sim) \leq CE^{geo}(u) + \text{a constant independent of } u.$$

The constant depends only on the restriction of the Lagrangian boundary conditions and the almost complex structures to a neighbourhood of $\partial_i S_P$, but is independent of their restriction to $\Theta \times \nu'_X \Theta$ for each component Θ of the thin part. \square

APPENDIX B. TATE'S ACYCLICITY THEOREM

Given an affinoid covering of an affinoid domain, Tate showed in [30] that the augmented Čech complex of the rings of functions of the domains of the cover is acyclic, and more generally for the Čech complex with coefficients in a complex of coherent sheaves. Tate's argument starts by constructing a null-homotopy for Laurent coverings (see Section B.3 below), then proceeds to use standard tools of homological algebra to conclude the general case.

B.1. Statement of the main result. In this section, we imitate the strategy of Tate's proof in order to be able to use Čech methods to compute morphism spaces in the analytic Fukaya category. Since the construction is completely local, we consider a torus \mathbb{T}^n equipped with a basepoint and Morse function. Let Po denote the A_∞ -category with objects integral affine polytopes $P \subset H_1(\mathbb{T}^n, \mathbb{R})$, morphisms for a pair (P_0, P_1) given by the complex

$$(B.1.1) \quad CF^*(P_0, P_1) := CM^*(X_q, \text{Hom}_\Lambda^c(U^{P_0}, U^{P_1}) \otimes \lambda)$$

and A_∞ operations as in Section 4.4.1. We do not require that the objects of P have non-empty interior, which implies that the objects of Po are closed under intersections.

Let A be a finite partially ordered set indexing a cover $\{P_a\}_{a \in A}$ of a polytope P , with the property that the polytopes P_a and P_b are disjoint whenever a and b are not comparable (note that the construction outlined after Equation (2.1.2) satisfies this property). Given a totally ordered subset $\tau \subseteq A$, we denote by P_τ the intersection of the polytopes P_a for $a \in \tau$. If τ is a subset of σ , the inclusion $P_\sigma \subseteq P_\tau$ gives rise to a canonical element

$$(B.1.2) \quad \delta_\tau^\sigma \in CF^0(P_\tau, P_\sigma)$$

whose construction is recalled in Section B.2 below.

We obtain a twisted complex

$$(B.1.3) \quad \check{T}(\{P_\alpha\}) := \left(\bigoplus_{\emptyset \neq \sigma \subset A} P_\sigma[-|\sigma|], \delta \right),$$

where δ is the Čech differential. Explicitly, if σ equals (a_1, \dots, a_m) as an ordered set, then the restriction of the differential to P_σ is given by

$$(B.1.4) \quad \sum_{a \in A \setminus \sigma} (-1)^i \delta_\sigma^{\sigma \cup \{a\}}.$$

Proposition B.1. *The natural map*

$$(B.1.5) \quad P \rightarrow \check{T}(\{P_\alpha\})$$

induced by $\bigoplus_{a \in A} \delta_\emptyset^a$ is a quasi-isomorphism in the category of twisted complexes over P_0 .

The proof of this proposition is given in Section B.4 below.

B.2. Restriction maps on Floer cochains. We begin by defining the map in Equation (B.1.2) more precisely: if $P_1 \subseteq P_0$, the restriction map $\Gamma^{P_0} \rightarrow \Gamma^{P_1}$ induces a map of local systems

$$(B.2.1) \quad U^{P_0} \rightarrow U^{P_1},$$

which is a continuous inclusion (we point out again that, unlike in Section 2.2.2, we do not decorate local systems by elements of an underlying cover of Q). These maps are natural in sense that given a triple $P_2 \subseteq P_1 \subseteq P_0$ the map $U^{P_0} \rightarrow U^{P_2}$ is given by composition.

Taking the sum of these elements over all maxima of the Morse function on \mathbb{T}^n , we obtain

$$(B.2.2) \quad \delta_{P_0}^{P_1} \in CF^0(P_0, P_1).$$

To tie back to Equation (B.1.2), we define

$$(B.2.3) \quad \delta_\tau^\sigma := \delta_{P_\tau}^{P_\sigma}.$$

B.3. Tate's null-homotopy. Following Tate, we begin by considering an integral affine function u on $H_1(\mathbb{T}^n, \mathbb{R})$, which we assume is primitive in the sense that the class of $du \in H^1(\mathbb{T}^n, \mathbb{Z})$ is not divisible.

If P is an integral affine polytope, let P_+ and P_- denote the subsets of P where u is non-negative, respectively non-positive, and let P_0 denote their intersection. In the language of rigid geometry, this corresponds to a *Laurent cover* with two terms (Tate calls these *two term special affine coverings*, see [30, Lemma 8.3]).

Consider the two-term Čech complex $\check{C}^*(\Gamma^P; \{+, -\})$

$$(B.3.1) \quad \Gamma^{P_+} \oplus \Gamma^{P_-} \xrightarrow{\hat{d}} \Gamma^{P_0}.$$

This complex is natural in the sense that every inclusion $P \subseteq P'$ induces a natural map of complexes

$$(B.3.2) \quad \check{C}^*(\Gamma^{P'}; \{+, -\}) \rightarrow \check{C}^*(\Gamma^P; \{+, -\}).$$

Lemma B.2. *A choice of complementary subspace to the line spanned by du in $H^1(\mathbb{T}^n, \mathbb{Z})$ determines a continuous null-homotopy for the augmented Čech complex*

$$(B.3.3) \quad \Gamma_P \rightarrow \check{C}^*(\Gamma^P; \{+, -\})$$

which is natural in P .

Proof. Applying a change of coordinates, we may assume that u is a coordinate function on $H_1(\mathbb{T}^n, \mathbb{R})$, corresponding to a monomial z in the group ring Γ . We formally write every element of the ring of functions on P_0 , P_+ , P_- , and P as

$$(B.3.4) \quad F(z, w) = \sum_{i=-\infty}^{+\infty} z^i f_i(w)$$

where $w = (w_2, \dots, w_n)$ are the other coordinates, and define

$$(B.3.5) \quad F_+(z, w) := \sum_{i=1}^{+\infty} z^i f_i(w)$$

$$(B.3.6) \quad F_-(z, w) := \sum_{i=-\infty}^0 z^i f_i(w).$$

Evidently, $F_- + F_+ = F$. On Γ^{P_0} , the null homotopy is provided by

$$(B.3.7) \quad \Gamma^{P_+} \oplus \Gamma^{P_-} \leftarrow \Gamma^{P_0} : \hat{h}$$

$$(B.3.8) \quad (F_+, -F_-) \leftrightarrow F$$

The valuation of this map is non-negative because the minimal valuation of F_\pm on P_\pm is achieved on P_\pm , and $\hat{d} \circ \hat{h}|_{\Gamma^{P_0}}$ is the identity, where \hat{d} is the differential on the augmented Čech complex. On $\Gamma^{P_+} \oplus \Gamma^{P_-}$, the null homotopy is

$$(B.3.9) \quad \Gamma^P \leftarrow \Gamma^{P_+} \oplus \Gamma^{P_-} : \hat{h}$$

$$(B.3.10) \quad F_- + G_+ \leftrightarrow (F, G),$$

which again has non-negative valuation. The reader may easily compute that $\hat{h} \circ \hat{d}|_{\Gamma^P}$ is the identity. On $\Gamma^{P_+} \oplus \Gamma^{P_-}$, we have

$$(B.3.11) \quad \begin{array}{ccccc} & & (F, G) & & \\ & \swarrow & & \searrow & \\ F_- + G_+ & & & & F - G \\ & \searrow & & \swarrow & \\ & (F_- + G_+ + (F - G)_+, F_- + G_+ - (F - G)_-) & & & \end{array}$$

from which the equation for a null-homotopy follows.

Naturality with respect to restriction maps is automatic from the fact that we did not appeal to any property of P in constructing \hat{h} , and the only choice we made is the choice of complementary subspace that determines the decomposition in Equation (B.3.4) (this decomposition clearly does not depend on the actual coordinates (w_2, \dots, w_n) , but only on subspace of $H^1(\mathbb{T}^n, \mathbb{Z})$ spanned by their derivatives). \square

Given a topological vector space V , we consider the complex

$$(B.3.12) \quad \text{Hom}_\Lambda^c(\check{C}^*(\Gamma^P; \{+, -\}), V) := \text{Hom}_\Lambda^c(\Gamma^{P_0}, V) \rightarrow \text{Hom}_\Lambda^c(\Gamma^{P_+}, V) \oplus \text{Hom}_\Lambda^c(\Gamma^{P_-}, V),$$

where Hom_Λ^c is the space of continuous maps. Composing with \tilde{h} , we obtain a null-homotopy for the augmented complex

$$(B.3.13) \quad \text{Hom}_\Lambda^c(\check{C}^*(\Gamma^P; \{+, -\}), V) \rightarrow \text{Hom}_\Lambda^c(\Gamma^P, V)$$

which is natural in P and V .

B.4. Acyclicity of the augmented complex. We next consider a general Laurent cover: let $\{u_m\}_{m \in M}$ be a collection of integral affine functions on $H_1(\mathbb{T}^n, \mathbb{R})$ indexed by a finite set M . Given a polytope P , we associate to each element of $M \times \{+, -\}$ the polytope $P_{m, \pm}$ given by the subset where u_m is non-negative (or non-positive). We say that the elements $P_{m, \pm}$ are the *Laurent cover* associated to $M \times \{+, -\}$, so that we obtain a twisted complex $\check{T}(P, M \times \{+, -\})$ on Po given by Equation (B.1.3).

Let (P, P') be a pair of polytopes. By definition, the space of morphisms in the category of twisted complexes on Po from $\check{T}(P, M \times \{+, -\})$ to P' is given by

$$(B.4.1) \quad \left(\bigoplus_{\sigma \subset M \times \{+, -\}} CF^*(P_\sigma, P')[-n], \delta \right)$$

with the differential induced by restriction.

Lemma B.3. *For each pair of polytopes (P, P') , the natural map*

$$(B.4.2) \quad \left(\bigoplus_{\sigma \subset M \times \{+, -\}} CF^*(P_\sigma, P')[-n], \delta \right) \rightarrow CF^*(P, P')$$

is a quasi-isomorphism.

Proof. Filtering by the degree of critical points of the Morse function, it suffices to prove that the augmented complex

$$(B.4.3) \quad \bigoplus_{\sigma \subset M \times \{+, -\}} \text{Hom}_\Lambda^c(\Gamma^{P_\sigma}, \Gamma^{P'}) \rightarrow \text{Hom}_\Lambda^c(\Gamma^P, \Gamma^{P'})$$

is a quasi-isomorphism. As in [30, Lemma 8.4], induction on the number of elements of M reduces this to the case of a singleton, which follows immediately from Lemma B.2. \square

Corollary B.4. *If $M \times \{+, -\}$ indexes a Laurent cover of a polytope P , there is a natural quasi-isomorphism*

$$(B.4.4) \quad P \rightarrow \check{T}(P, M \times \{+, -\}).$$

\square

We now prove the main result of this section:

Proof of Proposition B.1. We essentially follow the method introduced by Tate in [30, Section 8]: every cover Σ admits a refinement by a Laurent cover $M \times \{+, -\}$ obtained by considering the functions defining all boundary facets of polytopes appearing in the cover. The naturality of the construction of the complexes \check{T} implies that we can write the map from P to the Čech twisted complex associated to $M \times \{+, -\}$ as a composition:

$$(B.4.5) \quad P \rightarrow \check{T}(P, \Sigma) \rightarrow \check{T}(P, M \times \{+, -\}),$$

where the first arrow is the map which we would like to show is a quasi-isomorphism. By Corollary B.4, it suffices to show that the second map is a quasi-isomorphism. Filtering by the number of elements of a subset $\sigma \in \Sigma$, this follows by applying Corollary B.4 to

$$(B.4.6) \quad P_\sigma \rightarrow \check{T}(P_\sigma, M \times \{+, -\}).$$

\square

APPENDIX C. COMPUTATIONS OF COHOMOLOGY GROUPS WITH TWISTED COEFFICIENTS

The main goal of this section is to prove Propositions 2.19 and 2.42. As in Appendix B, we consider the local situation: given a Morse function on the torus, we define a Floer group $CF^*(P_0, P_1)$ for each pair of polytopes in $H_1(\mathbb{T}^n; \mathbb{R})$.

Proposition C.1. *If $P_1 \subseteq P_0$, the Floer group $HF^*(P_0, P_1)$ is supported in degree 0, and we have a natural quasi-isomorphism*

$$(C.0.1) \quad \Gamma^{P_1} \rightarrow CF^*(P_0, P_1)$$

given by the inclusion of Γ^{P_1} in $CF^0(P_0, P_1)$. If P_0 is contained in the interior of P_1 , the Floer group $HF^(P_0, P_1)$ is supported in degree n , and a choice of orientation of P_0 determines a quasi-isomorphism*

$$(C.0.2) \quad CF^*(P_0, P_1) \cong \text{Hom}_\Lambda^c(\Gamma^{P_0}, \Lambda)$$

given by the trace map on $CF^n(P_0, P_1)$.

Along the way, we shall prove that morphisms between disjoint polytopes vanish.

Remark C.2. A version of the results of this section hold for the completions of the homology of the based loops space of any topological space having the homotopy type of a finite CW complex, where the polytopes are integral affine subsets of first cohomology. The proof would take us too far afield, so we give a computational and explicit proof in the case of tori.

Throughout this section, we shall write $\text{Hom}(A, B)$ and $A \otimes B$ for the group of homomorphisms and for the tensor product of two abelian groups A and B . If A and B are equipped with (semi)-norms, we write $\text{Hom}^c(A, B)$ for the group of bounded homomorphisms, and $A \hat{\otimes} B$ for the completed tensor product (i.e. for the completion of $A \otimes B$ with respect to the induced semi-norm). When R is a (commutative) ring, and A and B are modules over R , we write $\text{Hom}_R(A, B)$ and $A \otimes_R B$ for the R -homomorphisms, and the tensor product of modules. Finally, assuming that R is a normed ring, and A and B are R -modules equipped with (semi)-norms, we write $\text{Hom}_R^c(A, B)$ for the group of bounded R -module homomorphisms, and $A \hat{\otimes}_R B$ for the completion of $A \otimes_R B$.

C.1. The 1-dimensional case. Consider the circle equipped with the standard Morse function f with a unique minimum and maximum, and let U denote the local system corresponding to the regular representation of the fundamental group. Identifying the space of paths from the minimum to the maximum with the space of based loops at the maximum via the choice of a path connecting these two points, and the homology of the latter with the Laurent polynomial ring $\Gamma = \mathbb{Z}[z, z^{-1}]$, the differential in the Morse complex

$$(C.1.1) \quad CM^*(f; \text{Hom}(U, U))$$

with coefficients in $\text{Hom}(U, U)$ can be expressed as the map

$$(C.1.2) \quad \text{Hom}(\Gamma, \Gamma) \rightarrow \text{Hom}(\Gamma, \Gamma)$$

$$(C.1.3) \quad \phi \mapsto \phi - z \cdot \phi \cdot z^{-1},$$

and the kernel of this differential is naturally isomorphic to

$$(C.1.4) \quad \Gamma \cong \text{Hom}_\Gamma(\Gamma, \Gamma) \subset \text{Hom}(\Gamma, \Gamma).$$

Consider the map

$$(C.1.5) \quad \text{Hom}(\Gamma, \Gamma) \leftarrow \text{Hom}(\Gamma, \Gamma) : h,$$

which is recursively defined by the formula

$$(C.1.6) \quad \psi(z^i) + z(h\psi)(z^{i-1}) = (h\psi)(z^i)$$

which we normalise by setting $(h\psi)(1) = 0$. Note in particular that the above formula implies that $(h\psi)(z) = \psi(z)$, and $(h\psi)(z^{-1}) = -z^{-1}\psi(1)$.

Lemma C.3. *The map h defines a homotopy from the identity of $CM^*(f; \text{Hom}(U, U))$ to the composition*

$$(C.1.7) \quad CM^*(f; \text{Hom}(U, U)) \rightarrow \Gamma \rightarrow CM^*(f; \text{Hom}(U, U)),$$

where the first map is the projection

$$(C.1.8) \quad \text{Hom}(\Gamma, \Gamma) \rightarrow \text{Hom}_\Gamma(\Gamma, \Gamma) := \Gamma$$

$$(C.1.9) \quad \phi \mapsto \phi(1)$$

from $CM^0(f; \text{Hom}(U, U))$ to Γ , and the second map is the inclusion in Equation (C.1.4). \square

For later purposes, it is convenient to derive this complex, and the corresponding null-homotopy, from a version of the Koszul complex: namely, consider

$$(C.1.10) \quad \Gamma \otimes \Gamma \rightarrow \Gamma \otimes \Gamma$$

with differential

$$(C.1.11) \quad d(f \otimes g) = f \otimes g - z^{-1} \cdot f \otimes g \cdot z.$$

The complex $CM^*(f; \text{Hom}(U, U))$ is naturally isomorphic to the complex of Γ -module maps from Equation (C.1.10) to Γ , and the homotopy h arises from a null homotopy of the extended complex

$$(C.1.12) \quad \Gamma \otimes \Gamma \rightarrow \Gamma \otimes \Gamma \rightarrow \Gamma,$$

where the last map sends $f \otimes g$ to the product $f \times g$.

Remark C.4. Note that the choice of Koszul complex depends on a choice of decomposition of Laurent polynomials into positive and negative powers. In particular, the differential d_- associated to swapping the rôles of z and z^{-1} is related to the above differential by the equation

$$(C.1.13) \quad d_- = -z \cdot d \cdot z^{-1}.$$

The minus sign above accounts for the need to choose an orientation when we define the trace map on the dual Floer cohomology group.

C.2. The standard Morse complex on the torus. We consider the torus $\mathbb{T}^n := S^1 \times \dots \times S^1$, and denote by Γ the group ring of $H_1(\mathbb{T}^n; \mathbb{Z})$. The product decomposition, and an orientation of each factor, induce an isomorphism

$$(C.2.1) \quad \Gamma \cong \mathbb{Z}[z_1^\pm, \dots, z_n^\pm].$$

It is convenient to switch back and forth between this notation, and the notation wherein we write elements of Γ as z^α for $\alpha \in \mathbb{Z}^n$.

Consider the standard Morse function, i.e. the sum of the standard Morse functions on each factor, having a unique minimum and maximum, and pick on each factor a path from the minimum to the maximum. Consider the universal local system whose fibre at a point is the space of paths to a basepoint, which we choose to be the maximum. Our choice of paths identifies the Morse complex with coefficients in the endomorphisms of this local system with

$$(C.2.2) \quad \text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z}),$$

with a differential given by

$$(C.2.3) \quad \partial(\phi \otimes \alpha) = \sum_i \partial_i(\phi \otimes \alpha) := \sum_i (\phi - z_j \cdot \phi \cdot z_j^{-1}) \otimes b_j \wedge \alpha,$$

where $\{b_j\}_{j=1}^n$ is the standard basis of $H^1(\mathbb{T}^n; \mathbb{Z})$.

There is a natural subcomplex of Equation (C.2.2) given by the inclusion

$$(C.2.4) \quad \Gamma \rightarrow \text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z})$$

whose image lies is $\text{Hom}_\Gamma(\Gamma, \Gamma) \otimes H^0(\mathbb{T}^n; \mathbb{Z})$. We shall construct an explicit retraction from the right to the left hand side.

To this end, we define a map

$$(C.2.5) \quad h_j: \text{Hom}(\Gamma, \Gamma) \rightarrow \text{Hom}(\Gamma, \Gamma)$$

$$(C.2.6) \quad h_j \psi(z^\alpha) = \begin{cases} 0 & \text{if } \alpha_j = 0 \\ \psi(z^\alpha) + z_j \cdot \psi(z^{\alpha-e_j}) & \text{otherwise} \end{cases}$$

where $\alpha \in \mathbb{Z}^n$, and e_j is the j^{th} basis element. Note that this is a recursive definition of $h_j \psi$, and that the explicit formula is

$$(C.2.7) \quad h_j \psi(z^\alpha) = \begin{cases} \sum_{i=0}^{\alpha_j-1} z^{ie_j} \psi(z^{\alpha-ie_j}) & 0 \leq \alpha_j \\ -\sum_{i=\alpha_j}^{-1} z^{ie_j} \psi(z^{\alpha-ie_j}) & \alpha_j < 0. \end{cases}$$

We then define

$$(C.2.8) \quad h: \text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z}) \rightarrow \text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z})$$

$$(C.2.9) \quad h = \sum_{j=1}^n h_j \otimes \iota_j,$$

where ι_j is the slant product

$$(C.2.10) \quad H^*(\mathbb{T}^n; \mathbb{Z}) \rightarrow H^{*-1}(\mathbb{T}^n; \mathbb{Z})$$

with the basis element of $e_j \in H_1(\mathbb{T}^n; \mathbb{Z})$.

Lemma C.5. *The map h is a homotopy between the identity on $\text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z})$ and the composition of the inclusion $\Gamma \subset \text{Hom}(\Gamma, \Gamma) \otimes H^0(\mathbb{T}^n; \mathbb{Z})$ with the projection*

$$(C.2.11) \quad \text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z}) \mapsto \Gamma,$$

which vanishes on the graded components of strictly positive degree, and is given on the degree 0 component by the map

$$(C.2.12) \quad \phi \otimes 1 \mapsto \phi(1).$$

Proof. The complex $\text{Hom}(\Gamma, \Gamma) \otimes H^*(\mathbb{T}^n; \mathbb{Z})$ is naturally isomorphic to the complex of Γ -homomorphisms from the n -fold tensor product of Equation (C.1.10) to Γ , and the homotopy to the projection is induced from the corresponding homotopy in Equation (C.1.6). \square

C.3. Construction of the homotopy for inclusions: I. Our goal is to extract from Lemma C.5 a bounded homotopy for the completions. To this end, we now use Γ to denote the ring of Laurent polynomials over the Novikov field Λ .

Let P_0 and P_1 be integral affine polytopes in $H^1(\mathbb{T}^n; \mathbb{R})$. The standard Morse complex computing morphisms between the corresponding local systems is given by

$$(C.3.1) \quad CF^*(P_0, P_1) \cong \text{Hom}_\Lambda^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^*(\mathbb{T}^n; \mathbb{Z}),$$

with differential given by Equation (C.2.3). Moreover, we have the inclusion of $\text{Hom}_\Gamma^c(\Gamma^{P_0}, \Gamma^{P_1})$ in degree 0.

Lemma C.6. *If $P_1 \subseteq P_0$, the homotopy h is continuous, hence induces a retraction*

$$(C.3.2) \quad \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^*(\mathbb{T}^n; \mathbb{Z}) \rightarrow \text{Hom}_{\Gamma}^c(\Gamma^{P_0}, \Gamma^{P_1}) = \Gamma^{P_1}$$

Proof. It suffices to prove that each map h_j is continuous, in which case it suffices to bound

$$(C.3.3) \quad \text{val}(h_j \psi) - \text{val } \psi$$

for any (continuous) map ψ from Γ^{P_0} to Γ^{P_1} . A straightforward computation reduces this to proving that

$$(C.3.4) \quad \text{val}_{P_0} z_j \leq \text{val}_{P_1} z_j$$

which follows immediately from the inclusion $P_1 \subseteq P_0$. \square

C.4. Construction of the homotopy for disjoint sets. Let us now consider the case where P_0 and P_1 are disjoint integral affine polytopes. By a change of coordinates, we may assume that the first lies in the region where the first coordinate is strictly positive, and the second in the region where it is strictly negative. We then define

$$(C.4.1) \quad h: \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^*(\mathbb{T}^n) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^*(\mathbb{T}^n)$$

$$(C.4.2) \quad \psi \otimes v \mapsto \sum_{i=1}^{\infty} z_1^i \cdot \psi \cdot z_1^{-i} \otimes \iota_1 v.$$

We note that the infinite series on the right hand side is convergent because

$$(C.4.3) \quad \text{val}(z_1^i \cdot \psi \cdot z_1^{i-1}) = \text{val}(\psi) + i(\text{val}_{P_0}(z) + \text{val}_{P_1}(z^{-1}))$$

and the assumptions on P_0 and P_1 respectively imply that

$$(C.4.4) \quad 0 < \text{val}_{P_0} z \text{ and } 0 < \text{val}_{P_1} z^{-1}.$$

Lemma C.7. *The map h defines a null-homotopy of $\text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^*(\mathbb{T}^n)$.*

Proof. Let us write $H^*(\mathbb{T}^n)$ as the direct sum $H^0(S^1) \otimes H^*(\mathbb{T}^{n-1}) \oplus H^1(S^1) \otimes H^*(\mathbb{T}^{n-1})$. The compositions

$$h \circ \partial_1: \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^0(S^1) \otimes H^*(\mathbb{T}^{n-1}) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^0(S^1) \otimes H^*(\mathbb{T}^{n-1})$$

$$\partial_1 \circ h: \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^1(S^1) \otimes H^*(\mathbb{T}^{n-1}) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_1}) \otimes H^1(S^1) \otimes H^*(\mathbb{T}^{n-1}),$$

both agree with the identity by an explicit computation. On the other hand, h commutes with each differential ∂_i for $i \neq 1$. The result follows. \square

Corollary C.8. *If $P_0 \cap P_1 = \emptyset$, the cohomology group $HF^*(P_0, P_1)$ vanishes.* \square

C.5. Computation of morphisms for inclusions: II. Consider the 1-dimensional case, with bounded closed intervals $P_1 \Subset P_0$ (i.e. so that P_1 is included in the interior of P_0). Recall that

$$(C.5.1) \quad \text{val}_{P_0} z < \text{val}_{P_1} z \text{ and } \text{val}_{P_0} z^{-1} < \text{val}_{P_1} z^{-1},$$

thus, there is a constant c such that

$$(C.5.2) \quad \text{val}_{P_1} z^i - \text{val}_{P_0} z^i > c|i|.$$

For the next statement, recall that $\hat{\otimes}$ denote the completed tensor product.

Lemma C.9. *The natural inclusions*

$$(C.5.3) \quad \text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_0}) \leftarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_0}) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_1})$$

factor through the inclusions

$$(C.5.4) \quad \Gamma^{P_i} \hat{\otimes} \text{Hom}_{\Lambda}^c(\Gamma^{P_i}, \Lambda) \subset \text{Hom}_{\Lambda}^c(\Gamma^{P_i}, \Gamma^{P_i}).$$

Proof. We can formally write any element $\phi \in \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_0})$ as

$$(C.5.5) \quad \sum_{i,j} \phi_i(z^j) z^i \otimes \rho_j,$$

where $\phi_i(z^j)$ is the coefficient of z^i in $\phi(z^j)$, and ρ_j is the homomorphism which assigns 1 to z^j and 0 to every other monomial basis element of Γ . The assumption that ϕ is bounded implies that there exists a constant K such that

$$(C.5.6) \quad \min_{i,j} (\text{val } \phi_i(z^j) + \text{val}_{P_0} z^i - \text{val}_{P_1} z^j) \geq K.$$

The result now follows from Equation (C.5.2), since replacing $\text{val}_{P_0} z^i$ by $\text{val}_{P_1} z^i$ allows us to add $c|i|$ to the right hand side, so that, when considered as an element of $\text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_1})$, there are only finitely many terms with valuation bounded above by any given number. Replacing $\text{val}_{P_1} z^j$ by $\text{val}_{P_0} z^j$ implies the same for $\text{Hom}_{\Lambda}^c(\Gamma^{P_0}, \Gamma^{P_0})$. \square

We conclude that there is a natural trace

$$(C.5.7) \quad \text{tr}: \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_0}) \rightarrow \Lambda,$$

given by the composition of the inclusions into $\Gamma^{P_i} \hat{\otimes} \text{Hom}_{\Lambda}^c(\Gamma^{P_i}, \Lambda)$ with the evaluation map

$$(C.5.8) \quad \Gamma^{P_i} \hat{\otimes} \text{Hom}_{\Lambda}^c(\Gamma^{P_i}, \Lambda) \rightarrow \Lambda$$

$$(C.5.9) \quad f \otimes \rho \mapsto \rho(f).$$

Explicitly, the trace can be written as

$$(C.5.10) \quad \psi \mapsto \sum_{i=-\infty}^{+\infty} \psi(z^i)_i,$$

where the subscript records the coefficient of z^i . This in particular shows that the traces defined via endomorphisms of Γ^{P_0} and Γ^{P_1} agree.

We now consider the map

$$(C.5.11) \quad \epsilon: \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_0}) \rightarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Lambda)$$

$$(C.5.12) \quad \epsilon\psi(f) := \text{tr}(\psi \circ f).$$

It is easy to see that ϵ composes trivially with the differential on $CF^*(P_1, P_0)$ hence defines a cochain map to $\text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Lambda)$. Consider the map

$$(C.5.13) \quad \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma^{P_0}) \leftarrow \text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Lambda): \delta$$

$$(C.5.14) \quad \rho(f) = \delta(\rho)(f).$$

Lemma C.10. *The composition $\epsilon \circ \delta$ is the identity on $\text{Hom}_{\Lambda}^c(\Gamma^{P_1}, \Gamma)$.*

Proof. Consider the map $\delta(\rho) \cdot z^i$. We compute that

$$(C.5.15) \quad \delta(\rho)(z^i \cdot z^j) = \rho(z^{i+j}).$$

Thus the trace of this map is given by setting $j = 0$, and is equal to $\rho(z^i)$. \square

We now define a map \hbar which will serve as a homotopy between the identity and the composition $\delta \circ \epsilon$, and which is determined by the expression

$$(C.5.16) \quad \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Gamma^{P_0}) \leftarrow \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Gamma^{P_0})$$

$$(C.5.17) \quad \psi + \hbar(z^{-1} \circ \psi \circ z) = \hbar(\psi) \leftarrow \psi,$$

which we normalise by requiring that $\hbar(\psi)$ vanish if the image of ψ is contained in the image of δ . We obtain an explicit expression for this map as follows: formally write

$$(C.5.18) \quad \psi = \sum_{i \in \mathbb{Z}} \psi_i$$

with ψ_i having image in the line spanned by z^i . We have

$$(C.5.19) \quad \hbar\psi_i = \begin{cases} \sum_{j=0}^{i-1} z^{-j} \circ \psi_i \circ z^j & 1 \leq i \\ 0 & i = 0 \\ \sum_{j=i}^{-1} z^{-j} \circ \psi_i \circ z^j & i \leq -1, \end{cases}$$

and we (formally) define

$$(C.5.20) \quad \hbar\psi = \sum_{i \in \mathbb{Z}} \hbar\psi_i$$

$$(C.5.21) \quad = \sum_{j \leq -1} z^{-j} \circ \psi_{\leq j} \circ z^j + \sum_{0 \leq j} z^{-j} \circ \psi_{j+1 \leq} \circ z^j,$$

where $\psi_{\leq j}$ is the sum of all components ψ_i with $i \leq j$, and $\psi_{j+1 \leq}$ is the sum of all components with $j+1 \leq i$.

Lemma C.11. *The expression in Equation (C.5.21) is convergent, and the map \hbar is continuous.*

Proof. We compute that, for j strictly negative, the valuation of $z^{-j} \cdot \psi_{\leq j} \cdot z^j$ is given by the infimum over all Laurent polynomials f of

$$(C.5.22) \quad \text{val}_{P_0} z^{-j} \cdot \psi_{\leq j} \cdot z^j f - \text{val}_{P_1} f = j \text{val}_{P_0} z^{-1} + \text{val}_{P_0} \psi_{\leq j} \cdot z^j f - \text{val}_{P_1} z^j f$$

$$(C.5.23) \quad + \text{val}_{P_1} z^j f - \text{val}_{P_1} f$$

$$(C.5.24) \quad \geq j \text{val}_{P_0} z^{-1} + \text{val} \psi_{\leq j} + \text{val}_{P_1} z^j$$

$$(C.5.25) \quad \geq -j (\text{val}_{P_1} z^{-1} - \text{val}_{P_0} z^{-1}) + \text{val} \psi.$$

The desired bound in this case thus follows from Equation (C.5.2). The case of positive monomials is similar, and the result immediately follows. \square

Lemma C.12. *The map \hbar defines a homotopy between the identity on $CF^*(P_1, P_0)$, and the projection $\delta \circ \epsilon$.*

Proof. The fact that $\hbar \circ d$ is the identity follows trivially from Equation (C.5.17). The equality $d \circ \hbar = id - \delta \circ \epsilon$ follows from a computation using the formal decomposition used above. In particular, for ϕ with image in the line spanned by z^i with i positive, we have

$$(C.5.26) \quad d\hbar\phi = d(\phi + z^{-1}\phi z + \cdots + z^{-i+1}\phi z^{i-1})$$

$$(C.5.27) \quad = \phi - z^{-1}\phi z + z^{-1}\phi z - z^{-2}\phi z^2 + \cdots - z^{-i}\phi z^i$$

$$(C.5.28) \quad = \phi - z^{-i}\phi z^i.$$

The result thus follows from the equality $\delta \circ \epsilon(\phi) = z^{-i}\phi z^i$. \square

We now consider the higher dimensional situation: let $P_1 \Subset P_0$ be nested integral affine polytopes in \mathbb{R}^n , with the property that the interior of P_0 contains P_1 . As before, we have a map

$$(C.5.29) \quad \epsilon: \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Gamma^{P_0}) \rightarrow \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Lambda)$$

$$(C.5.30) \quad \epsilon\psi(f) = \text{tr}(\psi \circ f).$$

with one-sided inverse given by the inclusion

$$(C.5.31) \quad \delta: \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Lambda) \rightarrow \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Gamma^{P_0})$$

whose image consists of maps factoring through $\Lambda \cdot 1 \subseteq \Gamma^{P_0}$.

Consider the maps \hbar_j given by

$$(C.5.32) \quad \text{Hom}(\Gamma^{P_1}, \Gamma^{P_0}) \leftarrow \text{Hom}(\Gamma^{P_1}, \Gamma^{P_0})$$

$$(C.5.33) \quad \psi + \hbar_j(z_j^{-1} \circ \psi \circ z_j) = \hbar_j(\psi) \leftarrow \psi,$$

which we normalise by requiring that $\hbar_j(\psi)$ vanish if the image of ψ is contained in the space spanned by monomials with trivial power of z_j (i.e. convergent Laurent series in the variables z_i for $i \neq j$). As in Section C.3, we set

$$(C.5.34) \quad \hbar = \sum_{j=1}^n \hbar_j \otimes \iota_j.$$

This map provides a homotopy between the identity on $CF^*(P_1, P_0)$ and the projection to $\text{Hom}_\Lambda^c(\Gamma^{P_1}, \Lambda)$ given by the composition

$$(C.5.35) \quad CF^n(P_1, P_0) \xrightarrow{\epsilon} \text{Hom}_\Lambda^c(\Gamma^{P_1}, \Lambda) \xrightarrow{\delta} CF^n(P_1, P_0),$$

on the degree n graded component of the Floer complex, which we extend by 0 to the other graded components.

Proof of Proposition C.1. For Morse functions on \mathbb{T}^n which are products of the standard Morse function on S^1 , the first part of Proposition C.1 follows from Lemma C.6, and the second from the homotopy \hbar constructed above. To conclude the result for general Morse functions, we use the standard argument that continuation maps in Morse theory induce homotopy equivalences of Morse complexes; these are automatically continuous, because each continuation map and each homotopy of continuation maps is a sum of finitely many terms. \square

Remark C.13. The constructions of this section are formally dual to those of Section C.3, in the sense that the formulae we use can be derived from those of that section by dualising with respect to the pairing

$$(C.5.36) \quad \Gamma \otimes_\Lambda \Gamma \rightarrow \Lambda$$

$$(C.5.37) \quad f \otimes g \mapsto \text{Res} \left(fg \frac{dz}{z} \right)$$

where the symbol $\text{Res}(hdz)$ assigns 1 to the monomial $h = z^{-1}$, and 0 to every non-trivial monomial. One can thus link the discussion of this section with the theory of residues via Tate's approach [29].

APPENDIX D. INVARIANCE OF THE FAMILY FLOER FUNCTOR

In this Appendix, we briefly outline an argument for the independence of the family Floer functor on the choices we have made (after placing further constraints on them). For simplicity, and as in [6, Appendix A] where the A_∞ functor is constructed, we consider the functor on a single tautologically unobstructed immersed Lagrangian L equipped with a compatible almost complex structure J_L so that L bounds no holomorphic discs and no holomorphic 1-gons; the case of a finite collection of mutually transverse tautologically unobstructed Lagrangians is only more complicated for notational reasons.

Using the methods of this paper, the essential choice in the construction of the functor is that of a *system of paths of fibres* as in Equation (3.2.1): for each holomorphic disc Σ with two punctures and a choice of $\ell \leq \dim Q + 1$ boundary marked points on one of the two components of the boundary, together with a labelling of the segments between the marked points by a collection (q_1, \dots, q_ℓ) of points in Q which lie in a ball of radius 1, we choose a smooth path q_T in Q , parametrised by the part of $\partial\Sigma$ carrying the marked points, and subject to the conditions listed in Section 3.2.4. We write \mathcal{S} for such a system of paths, to which we can assign according to Lemma 3.3, a reverse isoperimetric constant C_S for the families of moduli spaces of strips with one boundary condition along L , and the other along the paths of Lagrangian fibres prescribed by \mathcal{S} , parametrised by the choices of points (q_1, \dots, q_ℓ) .

We obtain from these moduli spaces a module \mathcal{L}_L over the category \mathcal{F}_A from Section 4.2, whenever A is a cover of Q whose associated cover Σ_A by the partially ordered set of non-empty intersections satisfies Condition (2.1.2), and which is sufficiently fine with respect to the constant C_S so that Condition (3.3.1) holds. Moreover, we have a map

$$(D.0.1) \quad CF^*(L, L) \rightarrow \text{Hom}_{\mathcal{F}_A}(\mathcal{L}_L, \mathcal{L}_L).$$

Following the methods of [6, Appendix A], we can extend this map to an A_∞ -homomorphism from the category \mathcal{F}_L with one object (with endomorphism $CF^*(L, L)$) to the category of modules on \mathcal{F}_A . The key point is that this does not require any further constraint on the cover because the higher terms of the functor are defined using moduli spaces of strips with marked points along both boundary segments; we use the system of paths \mathcal{S} along one, and boundary condition L along the other. The reverse isoperimetric constant which we used for the construction of the linear part of the functor thus applies to the higher parts as well.

Proposition D.1. *If \mathcal{S}_1 and \mathcal{S}_2 are systems of paths, there is a commutative diagram*

$$(D.0.2) \quad \begin{array}{ccccc} & & \mathcal{F}_L & & \\ & \swarrow & \downarrow & \searrow & \\ \text{mod}_{\mathcal{F}_{A_1}} & \longleftarrow & \text{mod}_{\mathcal{F}_{A_1 \amalg A_2}} & \longrightarrow & \text{mod}_{\mathcal{F}_{A_2}}, \end{array}$$

whenever A_1 and A_2 are sufficiently fine covers, in which the arrows from \mathcal{F}_L to the categories of modules over \mathcal{F}_{A_1} and \mathcal{F}_{A_2} are respectively defined using \mathcal{S}_1 and \mathcal{S}_2 .

Sketch of proof: Consider a system of paths \mathcal{S}_{12} , defined for sequences

$$(D.0.3) \quad (L, q_0^1, \dots, q_\ell^1, q_0^{12}, \dots, q_k^{12})$$

$$(D.0.4) \quad (L, q_0^2, \dots, q_\ell^2, q_0^{12}, \dots, q_k^{12})$$

with all points q_i^1 and q_j^{12} (or q_i^2 and q_j^{12}) lying within a ball of radius 1, and where ℓ is bounded by $\dim Q$ and k by $2\dim Q$. We require that these satisfy again the conditions listed in Section 3.2.4, and that they extend the choices of paths \mathcal{S}_1 and \mathcal{S}_2 in the sense

that they agree with them whenever there are no points labelled q_j^{12} . According to Lemma 3.3, there is a reverse isoperimetric constant $C_{S_{12}}$ for the moduli spaces of strips with one boundary condition along L , and the other along the paths prescribed by S_{12} .

We now consider covers A_1 and A_2 with the property that any sequence of nested non-empty intersections of elements of these sets has length smaller than or equal to $\dim Q + 1$. This implies that any sequence of nested non-empty intersections of elements of $A_1 \cup A_2$ has length smaller than or equal to $2 \dim Q + 2$. These non-empty intersections are the objects of the category $\mathcal{F}_{A_1 \sqcup A_2}$, which is equipped with natural embeddings of \mathcal{F}_{A_1} and \mathcal{F}_{A_2} .

Using the system of paths S_{12} , we may construct a functor

$$(D.0.5) \quad \mathcal{F}_L \rightarrow \mathcal{F}_{A_1 \sqcup A_2}.$$

The key point is to use the points labelled q_i^1 and q_i^2 as basepoints of the objects of $\mathcal{F}_{A_1 \sqcup A_2}$ which respectively come from \mathcal{F}_{A_1} and \mathcal{F}_{A_2} , and q_j^{12} for the remaining objects. This choice ensures that Diagram (D.0.2) commutes. The main results of this paper and of [6] imply that the functors we have constructed from \mathcal{F}_L to these categories of modules are all fully faithful embeddings (upon imposing a further constraint on the size of the covers). \square

In order to conclude from the above result that the family Floer functor is independent of choices, we need as well a comparison between the categories of modules over $\mathcal{F}_{A_1 \sqcup A_2}$ and \mathcal{F}_{A_i} . While these categories are not equivalent, Lemma 4.1 and the fact that continuation maps induce equivalences of family Floer homology groups (c.f. [5, Corollary 3.13]) together imply that the image of this functor lies in our model for the derived category of (twisted) coherent sheaves (c.f. [6, Definition 2.7]): this subcategory consists of those modules \mathcal{M} for which (i) the value on any object σ is a perfect module over $\mathcal{F}(\sigma, \sigma)$, and (ii) the induced map

$$(D.0.6) \quad \mathcal{M}(\sigma) \otimes_{\mathcal{F}(\sigma, \sigma)} \mathcal{F}(\sigma, \tau) \rightarrow \mathcal{M}(\tau)$$

is an equivalence.

Lemma D.2. *The functor $\text{mod}_{\mathcal{F}_{A_1 \sqcup A_2}} \rightarrow \text{mod}_{\mathcal{F}_{A_i}}$ restricts to a fully faithful embedding on the subcategory corresponding to complexes of (twisted) coherent sheaves.*

Sketch of proof: Let $\mathcal{F}_{A_1}^\perp$ denote the subcategory of $\mathcal{F}_{A_1 \sqcup A_2}$ consisting of objects which do not lie in \mathcal{F}_{A_1} . By construction, there are no morphisms from objects of $\mathcal{F}_{A_1}^\perp$ to those of \mathcal{F}_{A_1} (i.e. this is a semi-orthogonal decomposition). The morphism space between modules \mathcal{M} and \mathcal{N} over $\mathcal{F}_{A_1 \sqcup A_2}$ is thus naturally isomorphic to the cone of the map

$$(D.0.7) \quad \text{Hom}_{\mathcal{F}_{A_1}}(\mathcal{M}, \mathcal{N}) \oplus \text{Hom}_{\mathcal{F}_{A_1}^\perp}(\mathcal{M}, \mathcal{N}) \rightarrow \text{Hom}_{\mathcal{F}_{A_1}^\perp}(\mathcal{F}_{A_1}^\perp \Delta_{\mathcal{F}_{A_1}} \otimes_{\mathcal{F}_{A_1}} \mathcal{M}, \mathcal{N}),$$

where $\mathcal{F}_{A_1}^\perp \Delta_{\mathcal{F}_{A_1}}$ is the bimodule associated to the semi-orthogonal decomposition of $\mathcal{F}_{A_1 \sqcup A_2}$, and we omit any notation for the restriction of the modules to \mathcal{F}_{A_1} and $\mathcal{F}_{A_1}^\perp$. This isomorphism identifies the kernel of the restriction map (on cohomology) with the cohomology of the cone of

$$(D.0.8) \quad \text{Hom}_{\mathcal{F}_{A_1}^\perp}(\mathcal{M}, \mathcal{N}) \rightarrow \text{Hom}_{\mathcal{F}_{A_1}^\perp}(\mathcal{F}_{A_1}^\perp \Delta_{\mathcal{F}_{A_1}} \otimes_{\mathcal{F}_{A_1}} \mathcal{M}, \mathcal{N}).$$

The finiteness assumptions on \mathcal{M} and \mathcal{N} , imply that this complex is quasi-isomorphic to

$$(D.0.9) \quad \text{Hom}_{\mathcal{F}_{A_1}^\perp}(\text{Cone}(\mathcal{F}_{A_1}^\perp \Delta_{\mathcal{F}_{A_1}} \otimes_{\mathcal{F}_{A_1}} \mathcal{M} \rightarrow \mathcal{M}), \mathcal{N}).$$

Equation (D.0.6) implies that the cone of the map $\mathcal{F}_{A_1}^\perp \Delta_{\mathcal{F}_{A_1}} \otimes_{\mathcal{F}_{A_1}} \mathcal{M} \rightarrow \mathcal{M}$ vanishes on every object of $\mathcal{F}_{A_1}^\perp$ except those which are objects of \mathcal{F}_{A_1} . We shall use this to show that the complex in Equation (D.0.9) is acyclic.

The filtration of the category $\mathcal{F}_{A_1}^\perp$ by the number of elements A_2 appearing in the nested intersection associated to each object reduces the result to the following situation: P is an integral affine polytope (i.e. a polytope associated to an object of \mathcal{F}_{A_2}), A_P is a cover of P by integral affine polytopes (i.e. the cover obtained from intersecting the elements of A_1 with P), and \mathcal{C} is a module over the category obtained from \mathcal{F}_{A_P} by adding one object σ_P corresponding to P (with endomorphism algebra Γ^P), so that \mathcal{C} vanishes on every object other than σ_P , and whose value on σ_P is a perfect complex. Given another module \mathcal{N} whose value on σ_P is perfect, and such that Equation (D.0.6) holds, our goal is then to show that the space of morphisms from \mathcal{C} to \mathcal{N} is acyclic. Given the vanishing assumption on \mathcal{C} and Equation (D.0.6), this complex is quasi-isomorphic to the cone of the map

$$(D.0.10) \quad \text{Hom}_{\Gamma^P}(\mathcal{C}(\sigma_P), \mathcal{N}(\sigma_P)) \rightarrow \text{Hom}_{\mathcal{F}_{A_P}}(\mathcal{F}_{A_P} \Delta_{\mathcal{F}(\sigma_P, \sigma_P)} \otimes_{\Gamma^P} \mathcal{C}(\sigma_P), \mathcal{F}_{A_P} \Delta_{\Gamma^P} \otimes_{\Gamma^P} \mathcal{N}(\sigma_P)).$$

The result now follows from Tate acyclicity using the perfectness of $\mathcal{C}(\sigma_P)$ and $\mathcal{N}(\sigma_P)$. \square

REFERENCES

- [1] M. Abouzaid and P. Seidel. An open string analogue of Viterbo functoriality. *Geom. Topol.*, 14:627–718, 2010.
- [2] Mohammed Abouzaid. A cotangent fibre generates the Fukaya category. *Adv. Math.*, 228(2):894–939, 2011.
- [3] Mohammed Abouzaid. A topological model for the Fukaya categories of plumbings. *J. Differential Geom.*, 87(1):1–80, 2011.
- [4] Mohammed Abouzaid. Nearby Lagrangians with vanishing maslov class are homotopy equivalent. *Inventiones mathematicae*, 189:251–313, 2012.
- [5] Mohammed Abouzaid. Family Floer cohomology and mirror symmetry. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. II*, pages 813–836. Kyung Moon Sa, Seoul, 2014.
- [6] Mohammed Abouzaid. The family Floer functor is faithful. *J. Eur. Math. Soc. (JEMS)*, 19(7):2139–2217, 2017.
- [7] V. I. Arnol'd. *Singularities of caustics and wave fronts*, volume 62 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1990.
- [8] Marcel Berger. *A panoramic view of Riemannian geometry*. Springer-Verlag, Berlin, 2003.
- [9] S. Bosch, U. Güntzer, and R. Remmert. *Non-Archimedean analysis*, volume 261 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1984. A systematic approach to rigid analytic geometry.
- [10] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [11] Julien Duval. On a result by Y. Groman and J. P. Solomon. *Math. Ann.*, 364(3-4):1361–1363, 2016.
- [12] K. Fukaya. Morse homotopy, A_∞ -categories, and Floer homologies. In H. J. Kim, editor, *Proceedings of GARC workshop on Geometry and Topology*. Seoul National Univ., 1993.
- [13] K. Fukaya. Floer homology of Lagrangian foliation and non-commutative mirror symmetry. Preprint, 1998.
- [14] K. Fukaya. Cyclic symmetry and adic convergence in Lagrangian Floer theory. *Kyoto J. Math.*, 50:521–590, 2010.
- [15] K. Fukaya and Y.-G. Oh. Zero-loop open strings in the cotangent bundle and Morse homotopy. *Asian J. Math.*, 1:196–180, 1998.
- [16] Kenji Fukaya. Morse homotopy and its quantization. In *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 409–440. Amer. Math. Soc., Providence, RI, 1997.
- [17] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.
- [18] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Lagrangian Floer theory and mirror symmetry on compact toric manifolds. *Astérisque*, (376):vi+340, 2016.

- [19] Yoel Groman. Floer theory and reduced cohomology on open manifolds, 2015, arXiv:1510.04265.
- [20] Yoel Groman and Jake P. Solomon. A reverse isoperimetric inequality for J -holomorphic curves. *Geom. Funct. Anal.*, 24(5):1448–1515, 2014.
- [21] M. Gromov. Pseudoholomorphic curves in symplectic manifolds. *Invent. Math.*, 82(2):307–347, 1985.
- [22] Reinhardt Kiehl. Theorem A und Theorem B in der nichtarchimedischen Funktionentheorie. *Invent. Math.*, 2:256–273, 1967.
- [23] M. Kontsevich and Y. Soibelman. Homological mirror symmetry and torus fibrations. In *Symplectic geometry and mirror symmetry*, pages 203–263. World Scientific, 2001.
- [24] P. Seidel. A_∞ -subalgebras and natural transformations. *Homology, Homotopy Appl.*, 10:83–114, 2008.
- [25] Paul Seidel. Graded Lagrangian submanifolds. *Bull. Soc. Math. France*, 128(1):103–149, 2000.
- [26] Paul Seidel. *Fukaya categories and Picard-Lefschetz theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008.
- [27] Paul Seidel. Some speculations on pairs-of-pants decompositions and Fukaya categories. In *Surveys in differential geometry. Vol. XVII*, volume 17 of *Surv. Differ. Geom.*, pages 411–425. Int. Press, Boston, MA, 2012.
- [28] Andrew Strominger, Shing-Tung Yau, and Eric Zaslow. Mirror symmetry is T -duality. *Nuclear Phys. B*, 479(1-2):243–259, 1996.
- [29] John Tate. Residues of differentials on curves. *Ann. Sci. École Norm. Sup. (4)*, 1:149–159, 1968.
- [30] John Tate. Rigid analytic spaces. *Invent. Math.*, 12:257–289, 1971.
- [31] Umut Varolgunes. Mayer-vietoris property for relative symplectic cohomology, 2018, arXiv:1806.00684.
- [32] Sara Venkatesh. Rabinowitz Floer homology and mirror symmetry. *J. Topol.*, 11(1):144–179, 2018.
- [33] Andrei Verona. *Stratified mappings—structure and triangulability*, volume 1102 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1984.