An open string analogue of Viterbo functoriality

MOHAMMED ABOUZAID
PAUL SEIDEL

Liouville domains are a special type of symplectic manifolds with boundary (they have an everywhere defined Liouville flow, pointing outwards along the boundary). Symplectic cohomology for Liouville domains was introduced by Cieliebak–Floer–Hofer–Wysocki and Viterbo. The latter constructed a restriction (or transfer) map associated to an embedding of one Liouville domain into another.

In this preprint, we look at exact Lagrangian submanifolds with Legendrian boundary inside a Liouville domain. The analogue of symplectic cohomology for such submanifolds is called "wrapped Floer cohomology". We construct an A_∞–structure on the underlying wrapped Floer complex, and (under suitable assumptions) an A_∞–homomorphism realizing the restriction to a Liouville subdomain. The construction of the A_∞–structure relies on an implementation of homotopy direct limits, and involves some new moduli spaces which are solutions of generalized continuation map equations.

53D40

1 Introduction

Symplectic (co)homology is an invariant of symplectic manifolds with boundary, introduced by Cieliebak–Floer–Hofer–Wysocki [7; 3; 9; 4] and Viterbo [30; 29]. There are actually several flavours of the theory, which appear in different contexts; even when the domains of applicability overlap, their mutual relation is not always clear. In this paper, the symplectic manifolds we consider are *Liouville domains*, which are particularly simple to work with, since they combine two advantageous features (convex contact type boundary, and exactness of the symplectic form) and we use Viterbo's definition (except that we call cohomology what he calls homology). The notation will be $SH^*(M)$ for the symplectic cohomology of a Liouville domain M. Crucially, for every Liouville subdomain $M^{\mathrm{in}} \subset M$, Viterbo defined a restriction (in his terminology, transfer) homomorphism

$$SH^*(M) \longrightarrow SH^*(M^{\mathrm{in}}). \tag{1}$$

Published: 16 February 2010 DOI: 10.2140/gt.2010.14.627

The original motivation came from the Weinstein conjecture; this *Viterbo functoriality* property of symplectic cohomology also has many applications to Lagrangian embedding questions, especially for cotangent bundles. There are several surveys covering aspects of this material, for instance Oancea [24] and Seidel [26].

Symplectic cohomology has an open string analogue, where one considers a Lagrangian submanifold $L \subset M$ (whose boundary is Legendrian, and with suitable exactness conditions). The special case of cotangent fibres appears in Abbondandolo–Schwartz [1], and the general construction is known to specialists, even though it has not received wide attention. Following Fukaya–Seidel–Smith [14] we call the resulting invariant *wrapped Floer cohomology*, and denote it by $HW^*(L)$. In analogy with (1) one expects, under suitable assumptions, to have a map

$$HW^*(L) \longrightarrow HW^*(L^{\mathrm{in}}) \tag{2}$$

where $L^{\mathrm{in}} = L \cap M^{\mathrm{in}}$; indeed, a fairly straightforward translation of the argument in Viterbo [30] shows that this can be done. The aim of the present paper is refine that idea, by first constructing an A_∞–algebra structure on the chain complex underlying $HW^*(L)$, and then proving that (2) is induced by an A_∞–homomorphism. In fact, we also provide an extension to several Lagrangian submanifolds, where wrapped Floer cohomology turns into the *wrapped Fukaya A_∞–category* $\mathcal{W}(M)$. In that case, Viterbo functoriality takes on the form of an A_∞–functor defined on a suitable full subcategory of $\mathcal{W}(M)$, and mapping that to $\mathcal{W}(M^{\mathrm{in}})$.

We now enter a little more into the details. There are several equivalent versions of Viterbo's definition, but we have found one to be particularly effective. Choose a Hamiltonian function H whose vector field X equals the Reeb vector field along ∂M. Then set

$$SH^*(M) = \varinjlim_w HF^*(M; wH), \tag{3}$$

where the right hand side is ordinary (Hamiltonian) Floer cohomology, and the direct limit, over all $w = 1, 2, \dots$, is formed with respect to suitable continuation maps. In parallel, one can define $HW^*(L)$ to be the direct limit of Lagrangian Floer cohomologies $HF^*(L; wH)$. Recall that for each single w, the Floer cochain complex $CF^*(L; wH)$ is generated by chords

$$x\colon [0,1] \to M \text{ satisfying } dx/dt = wX \text{ and } x(0), x(1) \in L. \tag{4}$$

If one chooses H carefully, these generators will be in bijection with critical points of the Morse function $H|L$, together with Reeb chords of length $\leq w$ for the Legendrian submanifold $\partial L \subset \partial M$. Hence, wrapped Floer cohomology is a combination of ordinary Floer cohomology (which is isomorphic to the cohomology of L in the present context)

and linearized relative SFT (which counts Reeb chords, and sometimes goes under the name of linear Legendrian contact homology). Since direct limits are exact, one could in principle define the chain complex underlying $HW^*(L)$ to be the direct limit of the $CF^*(L; wH)$. However, a direct limit is in a sense a quotient construction, hence is unlikely to be strictly compatible with the A_∞–structure. Instead, we use a larger quasi-isomorphic chain complex $CW^*(L; H)$, which is the homotopy direct limit (the algebraic version of the telescope construction in homotopy theory). By definition, the differential on this complex includes both the ordinary Floer differentials (for all w), and continuation maps from w to $w+1$. When constructing the higher order components μ^d of the A_∞–structure, one must similarly include terms arising from solutions to continuation map equations on $(d+1)$–punctured discs. In order to write down such an equation, one needs a one-form γ on the disc, with suitable behaviour near the punctures, and which satisfies $d\gamma \leq 0$ everywhere. To organize the choices of one-forms, we introduce the notion of *popsicle*, which is a punctured disc equipped with additional geometric data, given by marked points which can move along special lines on the surface (intuitively, one imagines that $d\gamma$ is a smeared out version of the δ–distribution located at those points). Of course, there is a certain amount of necessary preliminary work, which concerns the moduli spaces of popsicles and their compactifications.

The construction of A_∞–homomorphisms underlying (2) has two main ingredients. The first idea, taken from Viterbo [30], is to compress the inner domain $M^{\mathrm{in}} \subset M$ by using the Liouville flow. This shrinks all symplectic areas by some factor ρ, and thereby also the values of the action functional on any chord (4) lying inside M^{in}. Viterbo exploits this by making a careful choice of ρ and of the Hamiltonian (both depending on w), which allows one to define the map $HF^*(M; wH) \to HF^*(M^{\mathrm{in}}; wH^{\mathrm{in}})$ by projecting to a quotient of the Floer cochain complex. This can be carried out in the Lagrangian case as well, but it is not really suitable for our purpose since we need to consider all w simultaneously. The second idea, used to overcome this difficulty, is adapted from work of Fukaya, Oh, Ohta and Ono [13, Section 19]. In order to prove the independence of Fukaya-type A_∞–structures from various choices, they considered parametrized moduli spaces arranged in trees, where the parameter is required to increase monotonically from the root to the leaves. We adapt their construction to the setting of popsicles, and call the resulting objects *cascades* (as far as the authors know, this is not related to the use of the same term in Frauenfelder [10]). The parameter is then precisely Viterbo's rescaling factor ρ, which we allow to become arbitrarily small, meaning that it varies freely in $(0, 1]$. We should point out that the resulting A_∞–homomorphism has (in general) infinitely many higher order terms, hence is not a projection in any sense.

The structure of the paper is as follows. Sections 2–5 contain the main thread of the argument. The first three of these are devoted to, respectively, the discussion of popsicles and other Riemann surface matters; the definition of the A_∞–algebra structure; and that of the A_∞–homomorphism. Section 5 contains some less central additional material, among it the extension of the previous constructions to wrapped Fukaya categories. The second part of the paper (Sections 6–9) is a sequence of appendices, covering details which were omitted from the main exposition. These concern the construction of popsicle moduli spaces; a priori estimates that enter into various compactness arguments; transversality issues; and coherent orientations (with the resulting signs).

As should be clear from this description, the present paper is entirely foundational. One could certainly think of applications to the chord conjecture, in line with what has been done for symplectic cohomology, and also to Lagrangian embeddings, in particular fleshing out the last part of Fukaya–Seidel–Smith [14]. However, for the authors of this paper, the main interest lies in potential applications to homological mirror symmetry, using Mikhalkin's generalized pair-of-pants decompositions (see Mikhalkin [23] and Abouzaid [2]). This relies crucially on having the entire A_∞–structure, rather than only the cohomology level theory $HW^*(L)$ (the same is true for [14] as well). Unfortunately, explaining these ideas is beyond our scope here.

Acknowledgments. We would like to thank Denis Auroux, Ivan Smith, and Katrin Wehrheim for valuable discussions. This research was conducted during the period the first author served as a Clay Research Fellow. The second author was partially supported by NSF grants DMS-0405516 and DMS-0652620.

2 Popsicles

We start this section by recalling one of the fundamental facts underlying the construction of Fukaya categories, namely the interpretation of Stasheff associahedra as moduli spaces of Riemann surfaces with boundary (see Fukaya [11]). Our main topic is an enriched version of this setup, where the Riemann surfaces carry holomorphic functions with suitable boundary conditions. We will explain the geometry of the resulting moduli spaces, relegating the more technical details to Section 6. As an important consequence, there is a natural class of one-forms on our Riemann surfaces, which will be used later to write down inhomogeneous $\overline{\partial}$–equations. As a final topic, we introduce yet another version of decorated Riemann surfaces, which come with additional real parameters (this time, the parameters have no geometric meaning, and adding them should be thought of as an abstract operadic-type construction).

2.1 Discs

Fix $d \geq 0$. A *disc with $d+1$ ends* is a Riemann surface S isomorphic to the closed unit disc D minus $d+1$ boundary points (called points at infinity), together with a distinguished choice of one of those points (the distinguished point is called negative, and the others positive). We usually denote the points at infinity by $\boldsymbol{\zeta} = \{\zeta^0, \dots, \zeta^d\}$. Here, the numbering convention is that ζ^0 is the negative point, and that the positive ones appear in the order determined by the orientation of ∂S. The obvious compactification of S is then written as $\overline{S} = S \cup \boldsymbol{\zeta} \cong D$.

Example 2.1 A simple example is the infinite strip $Z = \mathbb{R} \times [0,1] \subset \mathbb{C}$, with points at infinity $\zeta^0 = -\infty$, $\zeta^1 = +\infty$. We usually write the coordinates on Z as $z = s + it$. Of course, any other disc with two ends is isomorphic to Z.

The infinite strip is unstable, because it admits an infinite group of translational symmetries $(s,t) \mapsto (s+\sigma, t)$. Discs with $d+1 \geq 3$ ends are automatically stable, and in fact have no nontrivial automorphisms (preserving the distinguished point at infinity). In this stable range, it is easy to construct a universal family of discs, which we denote by $\mathcal{S}^{d+1} \to \mathcal{R}^{d+1}$. The base (the moduli space of stable discs) is a smooth manifold diffeomorphic to $\mathbb{R}^{d-2}$.

Consider a ribbon tree T with $d+1$ semi-infinite edges, together with a preferred choice of one such edge (the preferred one is called the root, and the others the leaves). A *broken disc* modelled on T is a disjoint union

$$S = \coprod_v S_v \tag{5}$$

indexed by vertices v of T, where each component S_v is a disc whose number of ends is given by the valency $|v|$ of the vertex. We assume from now on that stability holds, which means that $|v| \geq 3$ for all vertices of T. Let $\mathcal{R}^{d+1,T}$ be the moduli space of broken discs modelled on T. If the tree has a single vertex, this just agrees with $\mathcal{R}^{d+1}$, while in all other cases, it is a product of lower-dimensional moduli spaces $\mathcal{R}^{|v|}$. Take the disjoint union over all stable trees T with $d+1$ semi-infinite edges:

$$\overline{\mathcal{R}}^{d+1} = \coprod_T \mathcal{R}^{d+1,T}. \tag{6}$$

This can be equipped with a natural topology, and indeed the structure of a manifold with corners, which is such that the stratification (6) is the one into boundary and corner strata of various codimensions. In particular, the interior of $\overline{\mathcal{R}}^{d+1}$ is precisely $\mathcal{R}^{d+1}$. We call these compactifications the *moduli spaces of stable broken discs* (they are homeomorphic to the classical Stasheff associahedra).

At this point, it is convenient to introduce some more combinatorial terminology. Take a tree T as before. A flag (sometimes also called a half-edge) in T is a pair consisting of a vertex and an adjacent edge. In our case, there is a preferred way of numbering the flags adjacent to any given vertex by $\{0,\dots,|v|\}$. Namely, one starts with the unique flag (called negative) which points towards the root, and continues with the other ones (called positive) according to the ribbon structure. Hence, if S is a broken disc modelled on T, there is a natural correspondence between flags (adjacent to v) and points at infinity (belonging to the component S_v). Note also that the semi-infinite edges of T can be numbered in a preferred way by $\{0,\dots,d\}$, by starting with the root and proceeding in accordance with the cyclic ordering given by any plane embedding $T \hookrightarrow \mathbb{R}^2$ compatible with the ribbon structure. Since each semi-infinite edge gives rise to a unique flag, we have $d+1$ distinguished points at infinity in S.

With this in mind, one defines the compactification $\overline{S}$ of a broken disc by first taking the compactifications $\overline{S}_v$ of all components, and then identifying those pairs of points at infinity which correspond to flags belonging to the same interior edge of T. Equivalently, one takes the disjoint union of the discs S_v, and adds to that $d+1$ points at infinity (corresponding to the semi-infinite edges of T) as well as singular points (corresponding to the finite edges). We call the added points of both type *special points* of $\overline{S}$. There is a natural compactification $\overline{\mathcal{S}}^{d+1}$ of the universal family $\mathcal{S}^{d+1}$, coming with a map

$$\overline{\mathcal{S}}^{d+1} \longrightarrow \overline{\mathcal{R}}^{d+1}, \tag{7}$$

such that the fibre over any point of $\overline{\mathcal{R}}^{d+1}$ is precisely the compactification of the broken disc corresponding to that point. Everywhere except at the singular points of compactified broken discs, $\overline{\mathcal{S}}^{d+1}$ is a manifold with corners. At those singular points, we have a different local model, namely

$$\mathbb{R}^k \times [0,\infty)^l \times \{(\delta_+,\delta_-) \in \mathbb{C}^2 \,:\, \mathrm{im}(\delta_+) \geq 0, \mathrm{im}(\delta_-) \geq 0, \delta_+\delta_- \in [0,\infty)\} \tag{8}$$

where of course, $k+l = d-3$ for dimension reasons. Still, the transition functions between these local models are differentiable (more precisely, restrictions of differentiable maps, in the coordinates given above). Using that, one can make sense of the notion of smooth map from a manifold into $\overline{\mathcal{S}}^{d+1}$, and also in the inverse direction. For instance, the projection map in (7) is smooth; and the distinguished points at infinity give rise to a set of $d+1$ smooth, pairwise disjoint sections of that map. We will encounter the same situation again later on, see Section 2.3, and return to the details in Section 6.4.

Example 2.2 Figure 1 shows the three-dimensional space $\overline{\mathcal{S}}^4$. By definition, points of this space are pairs (S,z), where $z \in \overline{S}$ is arbitrary. The picture shows what such

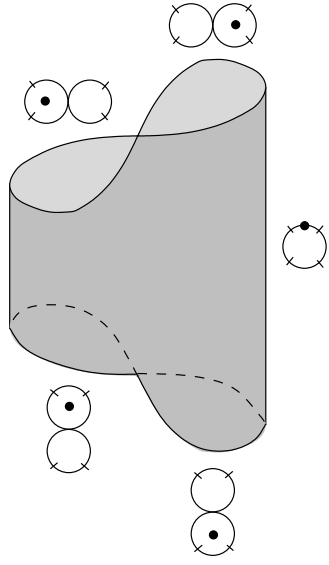

Figure 1

pairs look like on each of the five boundary faces of $\overline{\mathcal{S}}^4$. In the case of the "biggest" face, z can lie anywhere on $\partial\overline{S}$ (over the main stratum $\mathcal{R}^4 \cong \mathbb{R}$, where $\overline{S}$ is smooth, one gets a circle bundle; and as one gets to the limits, where the disc degenerates to a broken one, $\partial\overline{S}$ turns into a figure-eight).

2.2 Parametrizing the ends

Let $Z_- = \{s \le 0\}, Z_+ = \{s \ge 0\} \subset Z$ be the half-infinite strips. Take a disc S with $d+1$ ends. A *strip-like end* around a point at infinity ζ^k is a holomorphic embedding

$$\begin{cases} \epsilon\colon Z_- \longrightarrow S & \text{if } k = 0, \text{ or} \\ \epsilon\colon Z_+ \longrightarrow S & \text{if } k > 0, \end{cases} \tag{9}$$

satisfying $\epsilon^{-1}(\partial S) = Z_\pm \cap \partial Z$ and $\lim_{s\to\pm\infty} \epsilon(s,\cdot) = \zeta^k$. As usual, the first kind of end is called negative, and the other positive. A set of strip-like ends for S is a collection $\boldsymbol{\epsilon} = \{\epsilon^0, \dots, \epsilon^d\}$, where each ϵ^k is an end around ζ^k.

Strip-like ends are an auxiliary technical device, which enters into the standard gluing (or end connected sum) construction. To be specific, suppose that we have two discs $S_\pm$, a positive end $\epsilon_+\colon Z_+ \to S_+$, a negative end $\epsilon_-\colon Z_- \to S_-$, and a gluing parameter $\delta \in (-1, 0)$. The associated gluing length is set to be $\sigma = -\log(-\delta)$. One then defines

another disc

$$S = S_+ \# S_- \tag{10}$$

by first removing $\epsilon_+([\sigma, +\infty) \times [0,1])$ as well as $\epsilon_-((-\infty, -\sigma] \times [0,1])$ from $S_+ \sqcup S_-$, and then identifying the remaining finite pieces of the ends through the formula $\epsilon_+(s,t) \sim \epsilon_-(s-\sigma, t)$. The distinguished negative point at infinity in S is the one inherited from S_+.

This gluing process generalizes to families, and can be used to obtain collar neighbourhoods around the boundary strata of the compactified moduli space (6). Fix some tree T subject to the stability condition. Next, for each vertex v of that tree, choose a set of strip-like ends for the universal family $\mathcal{S}^{|v|} \to \mathcal{R}^{|v|}$ (this means a set of strip-like ends on each fibre of that family, varying smoothly over the base). Every finite edge e of the tree designates two vertices $v^\pm$ and the associated flags $(v^\pm, e)$, which in view of the considerations above, correspond to points at infinity in $\mathcal{S}^{|v_\pm|}$. One can then glue together the associated ends with some gluing parameter (say $\delta_e < 0$), or else leave them alone (which is the degenerate case $\delta_e = 0$). This gives rise to a gluing map on moduli spaces, of the form

$$\mathcal{R}^T \times \Big(\prod_e (-1, 0] \Big) \supset \mathcal{U}^T \longrightarrow \overline{\mathcal{R}}^{d+1}. \tag{11}$$

Here $\mathcal{U}^T$ is a subset where the gluing parameters are small, which means a neighbourhood of $\mathcal{R}^T \times \{0, \dots, 0\}$ (smallness of the parameter ensures that the various gluing processes don't interfere with each other). The main property of (11) is that it is smooth and, after shrinking $\mathcal{U}^T$ if necessary, a diffeomorphism onto a neighbourhood of $\mathcal{R}^T \subset \overline{\mathcal{R}}^{d+1}$. This is proved for instance in Seidel [27, Section 9].

A universal choice of strip-like ends (for discs) is the choice of a set $\boldsymbol{\epsilon}_S = \{\epsilon_S^k\}$ of ends for every stable disc S, which is compatible with isomorphisms of such discs, and smooth with respect to deformations. In practice, this means that one chooses ends for the universal family $\mathcal{S}^{d+1}$, varying smoothly over $\mathcal{R}^{d+1}$, and then obtains ends for any S by identifying it with a fibre of the universal family. It is useful to impose an additional *consistency* requirement, which makes such a choice compatible with gluing processes. Namely, assume that a universal choice has been made, and take the situation from (10), where both $S_\pm$ are assumed stable. As part of our global choice, these two surfaces come with sets of strip-like ends; and if the gluing parameter is small, S inherits a set of strip-like ends from that (basically, those ends which are not used up in the gluing process). The consistency condition says that for small values of the gluing parameter, these induced strip-like ends should be equivalent to those which are assigned to S as part of the universal choice. The existence of consistent

choices can be proved by an elementary inductive process, using coordinates (11) near the boundary strata (see [27, Section 9]; we will be using the same kind of construction repeatedly later on).

Addendum 2.3 The space of all possible strip-like ends around any point at infinity is contractible (in an appropriate topology) but infinite-dimensional, and this can be occasionally awkward to work with. To cut down on the amount of choice, one can restrict attention to *rational* strip-like ends $Z_\pm \to S$, which are the ones that extend to an isomorphism $\overline{Z} \cong \overline{S}$. The rational strip-like ends around a given point at infinity form a principal homogeneous space over the contractible group $Aut(D,1)$ of holomorphic automorphisms fixing a point. Since the gluing process preserves rationality, consistency still makes sense in this more restrictive context. The same argument as before can then be used to produce consistent choices of rational strip-like ends.

Addendum 2.4 The unstable case is problematic because strip-like ends cannot be made invariant under automorphisms of the Riemann surface. However, we can weaken the definition slightly, and then include at least the infinite strip. Call two strip-like ends *equivalent* if they differ by a translation of the domain $Z_\pm$. Equip Z itself with the tautological ends, which are the inclusions $Z_\pm \hookrightarrow Z$. From that, any $S \cong Z$ inherits a set of strip-like ends which are unique up to equivalence. Moreover, if we take a given consistent universal choice and enlarge it to $d = 1$ in this way, the compatibility with the gluing process is maintained (up to equivalence, of course).

2.3 The definition

Fix $d \geq 1$ and in addition, labels $p_f \in \{1, \dots, d\}$ indexed by some finite set F. We write $\mathbf{p} = \{p_f\}$ for a collection of such labels. A $\mathbf{p}$–*flavoured popsicle* is a disc S with $d+1$ ends, together with a collection $\boldsymbol{\phi} = \{\phi_f\}$ of holomorphic maps

$$\phi_f \colon S \longrightarrow Z, \tag{12}$$

each of which extends to an isomorphism $\overline{S} \cong \overline{Z}$, taking ζ^0 to $-\infty$ and ζ^{p_f} to $+\infty$. If we fix S, then the group $\mathbb{R}^F$ of translations acts simply transitively on the set of all possible popsicle structures $\boldsymbol{\phi}$. Additionally, one can permute the ϕ_f within each subset where p_f is constant, and that gives rise to an action of

$$Sym^{\mathbf{p}} = \{\text{permutations of } F \text{ stabilizing } \mathbf{p}\}. \tag{13}$$

Remark 2.5 One way to represent the additional data geometrically is as follows. Take some $k \in \{1, \dots, d\}$, and any holomorphic map $\phi \colon S \to Z$ which extends to

an isomorphism of compactifications, sending ζ^0 to $-\infty$ and ζ^k to $+\infty$. Consider the infinite line in S which is the preimage of $\{t = \frac{1}{2}\} \subset Z$. This line is obviously the same for all maps, but a specific choice of ϕ singles out a point on it, namely the preimage of $(0, \frac{1}{2}) \in Z$. In this way, a $\mathbf{p}$–flavoured popsicle structure can be encoded in a collection of points on S indexed by F (some of the points may coincide, and of course their position is not free, each being constrained to the infinite line determined by $k = p_f$). The group (13) acts by permuting those points.

Popsicles are stable, and in fact have no nontrivial automorphisms (compatible with the given structure), if and only if $d + |F| \geq 2$. In that range we have a universal family of popsicles, denoted by $\mathcal{S}^{d+1,\mathbf{p}} \to \mathcal{R}^{d+1,\mathbf{p}}$, whose base has dimension $d - 2 + |F|$. The structure of these moduli space is actually quite easy to understand. Assume first that $d \geq 2$. In that case, each fibre of the forgetful map

$$\mathcal{R}^{d+1,\mathbf{p}} \longrightarrow \mathcal{R}^{d+1} \tag{14}$$

can be identified with $\mathbb{R}^F$. This identification depends on a choice of origin in the fibre, but that can be chosen to vary smoothly with $\mathcal{R}^{d+1}$, giving rise to a non-canonical diffeomorphism $\mathcal{R}^{d+1,\mathbf{p}} \cong \mathcal{R}^{d+1} \times \mathbb{R}^F \cong \mathbb{R}^{d-2+F}$. The action of $\mathbb{R}^F$ is just translation on the fibres of (14), and (if one picks the origin to be invariant, which is always possible) $Sym^{\mathbf{p}}$ acts by permuting coordinates. In the remaining case $d = 1$, one can identify $\mathcal{R}^{d+1,\mathbf{p}}$ with the quotient $\mathbb{R}^F/\mathbb{R}$ (where $\mathbb{R}$ acts diagonally by translation on all factors). Note that in either case, any transposition in $Sym^{\mathbf{p}}$ acts orientation-reversingly.

Addendum 2.6 Suppose that we have a partition $\mathbf{P}$ of F into subsets $\{P_m\}$, such that the map $f \mapsto p_f$ is constant on each subset. The permutations of F which preserve $\mathbf{P}$ form a subgroup $Sym^{\mathbf{P}} \subset Sym^{\mathbf{p}}$. These are precisely the subgroups appearing as isotropy groups for the action of (13). The corresponding fixed point set in $\mathcal{R}^{d+1,\mathbf{p}}$ is the set of those $(S, \boldsymbol{\phi})$ such that for any two f belonging to the same P_m, the associated maps ϕ_f coincide.

Take a tree T of the same kind as before, and a decomposition $\mathbf{F} = \{F_v\}$ of F into subsets labeled by vertices of that tree. Moreover, the decomposition needs to be compatible with $\mathbf{p}$ in the following sense: for each $f \in F_v$, the vertex v must lie on the path from the root to the p_fth leaf. Note that $\mathbf{p}$ and the decomposition $\mathbf{F}$ determine numbers $\mathbf{p}_v = \{p_{v,f}\}$ for each vertex, where f runs through F_v and $p_{v,f} \in \{1, \dots, |v|\}$. This is because the path mentioned above leaves v in the direction of a specific flag, whose number is taken to be $p_{v,f}$. For instance, if T is the tree with just one vertex v, the compatibility condition implies that we necessarily have $\mathbf{p}_v = \mathbf{p}$. Then, a *broken popsicle* is a disjoint union (5), together with a $\mathbf{p}_v$–flavoured popsicle

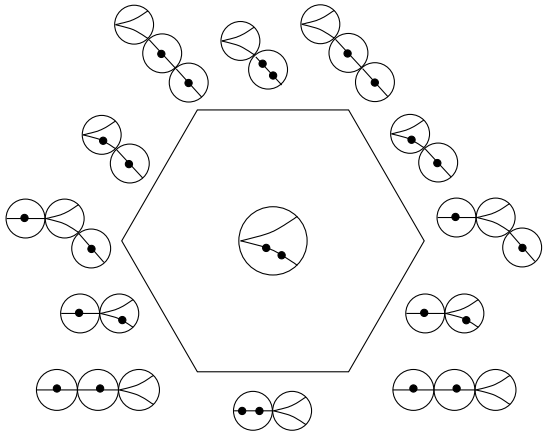

Figure 2

structure $\boldsymbol{\phi}_v$ on each S_v. This time, the stability condition is that two-valent vertices v are allowed as long as $F_v \neq \varnothing$. The moduli space of stable broken popsicles modelled on $(T, \mathbf{F})$ is a product

$$\mathcal{R}^{T,\mathbf{F}} \cong \prod_v \mathcal{R}^{|v|,\mathbf{p}_v}. \tag{15}$$

Take the disjoint union of those spaces,

$$\overline{\mathcal{R}}^{d+1,\mathbf{p}} = \coprod_{(T,\mathbf{F})} \mathcal{R}^{T,\mathbf{F}}. \tag{16}$$

We will see later (Corollary 6.5) that $\overline{\mathcal{R}}^{d+1,\mathbf{p}}$ carries a canonical structure of a smooth manifold with corners, such that (16) is the associated standard stratification, generalizing the corresponding statement for (6) (which would be the special case $F = \varnothing$). The forgetful map (14) (defined if $d \geq 2$), as well as the actions of $\mathbb{R}^F$ and $Sym^{\mathbf{p}}$, extend smoothly to the compactification. However, their geometry becomes a little more complicated: the compactified forgetful map is no longer a submersion, and the action of $\mathbb{R}^F$ on the compactification is not free for any value of d (both phenomena are due to the appearance of unstable components $S_v \cong Z$ in the boundary strata). Finally, we should point out that on the closure of a stratum (15), the action of $Sym^{\mathbf{p}}$ embeds into one of a potentially larger group, namely $\prod_v Sym^{\mathbf{p}_v}$.

Examples 2.7

(i) Take $d = 2$, $F = \{1, 2\}$, and $p_1 = p_2 = 1$. Then the compactified moduli space is a hexagon (Figure 2). The group of symmetries $Sym^{\mathbf{p}} = \mathbb{Z}/2$ acts by reflection along the vertical axis.

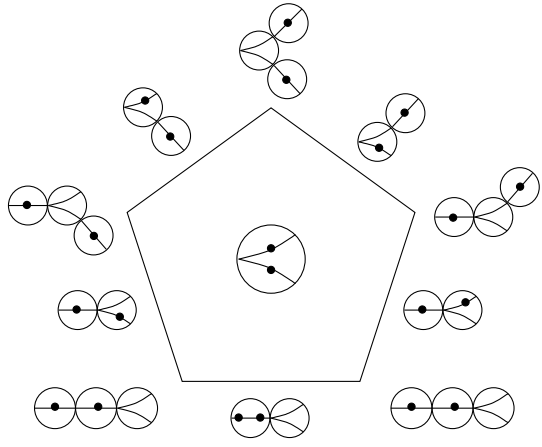

Figure 3

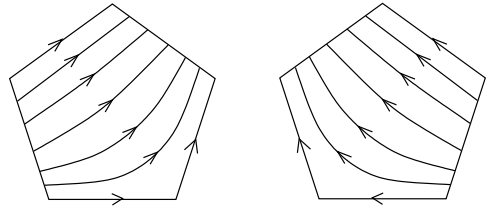

Figure 4

(ii) Take the same d and F, but now with $p_1 = 1$, $p_2 = 2$. Then the compactified moduli space is a pentagon (Figure 3). Note that while $Sym^{\mathbf{p}}$ is trivial in this case, one of the boundary edges (the bottom one in our picture) carries a nontrivial symmetry group $\prod_v Sym^{\mathbf{p}_v} = \mathbb{Z}/2$. In particular, the two corners adjacent to this edge are isomorphic moduli spaces. By this, we mean that the product expressions (15) for these two corner strata consist of the same factors, even though the strata themselves are distinguished by how the given F is partitioned into F_v.

It is also instructive to consider the $\mathbb{R}^2$–action on this moduli space. The orbits of the two vector fields generating the action are (schematically) drawn in Figure 4. Note that where orbits appear to be transverse to the boundary, their speed actually slows down to zero, and the limiting boundary edge consists of stationary points.

In exactly the same way as in (7), one can define compactifications of the universal families of popsicles,

$$\overline{\mathcal{S}}^{d+1,\mathbf{p}} \longrightarrow \overline{\mathcal{R}}^{d+1,\mathbf{p}}. \tag{17}$$

The total space has the same geometry as before, meaning that the local models are either the standard corners or (8). In addition to the $d+1$ canonical sections, the compactified family now comes with a collections of smooth maps to $\overline{Z}$. If we consider the fibre over a boundary point, which is the compactification of a broken popsicle, then each of these maps $\overline{\phi}_f\colon \overline{S} \to \overline{Z}$ is constant and equal to $\pm\infty$ on all components except one (the non-constant component $\overline{S}_v$ being precisely the one such that $f \in F_v$).

2.4 Popsicle sticks

The gluing process for popsicles is a little less straightforward than that for pointed discs. It seems difficult to work directly on the level of maps (12). We will therefore adapt the idea from Remark 2.5, replacing these maps by additional marked points lying on certain infinite lines. However, instead of taking those lines to be preimages of the line $\{t = 1/2\} \subset Z$, we will bend them slightly to achieve better compatibility with the strip-like ends, and hence with the gluing process.

Concretely, let S be a disc with $d+1$ ends, which comes equipped with a set $\boldsymbol{\epsilon}$ of strip-like ends. Fix some $k \in \{1, \dots, d\}$. A *popsicle stick* connecting ζ^0 to ζ^k is an infinite line $\mathbb{R} \cong Q \subset S$ satisfying

$$\begin{cases} \phi(Q) = \{t = \tau(s)\} \text{ for some } \tau\colon \mathbb{R} \longrightarrow (0,1), \\ Q = \epsilon^0(\{t = \frac{1}{2}\}) \text{ in a neighbourhood of } \zeta^0, \\ Q = \epsilon^k(\{t = \frac{1}{2}\}) \text{ in a neighbourhood of } \zeta^k. \end{cases} \tag{18}$$

Here, $\phi\colon S \to Z$ is any map as in Remark 2.5. To see that the three conditions in (18) are not mutually contradictory, one looks at the compositions $\rho_- = \phi \circ \epsilon^0\colon Z_- \to Z$ and $\rho_+ = \phi \circ \epsilon^k\colon Z_+ \to Z$. Either of these can be written as the sum of a translation and a remainder term which decays exponentially as $s \to \pm\infty$. Therefore, the image of $\{t = 1/2, \pm s \gg 0\}$ under $\rho_\pm$ is of the form $\{t = \tau_\pm(s)\}$. To build a popsicle stick, it is then sufficient to choose a function τ which at infinity agrees with the given $\tau_\pm$. The particular choice of ϕ is irrelevant for the definition of popsicle stick (and, incidentally, only the equivalence class of the strip-like ends matters). However, given Q, any choice of ϕ singles out a point $q \in Q$, namely the one which satisfies

$$\phi(q) \in \{0\} \times (0,1), \tag{19}$$

or equivalently in terms of (18), $\phi(q) = (0, \tau(0))$. This sets up a bijection between choices of ϕ and points q. We call those points *sprinkles*.

Define a set of popsicle sticks for S to be a collection $\mathbf{Q} = \{Q^1, \dots, Q^d\}$ where each Q^k connects ζ^0 to ζ^k. Clearly, given a set of popsicle sticks, there is a bijection between $\mathbf{p}$–flavoured popsicle structures and collections $\mathbf{q} = \{q_f\}$ of sprinkles, such that $q_f \in Q^{p_f}$ for all $f \in F$.

With this at hand, we return to the gluing construction. Take discs $S_\pm$, which come with sets of strip-like ends and popsicle sticks. For small values of the gluing parameter, the surface S defined as in (10) inherits the same additional structures. Moreover, if $S_\pm$ have $\mathbf{p}_\pm$–flavoured popsicle structures, then S inherits a $\mathbf{p}$–flavoured one, provided that $\mathbf{p}$ is related to $\mathbf{p}_\pm$ in a suitable way (which is straightforward to infer from our prescription for gluing popsicle sticks). As before, this process can be applied to families, and in particular to the boundary strata in (16). Pick a pair $(T, \mathbf{F})$ describing such a stratum. Assume that for each vertex v, we have chosen strip-like ends and popsicle sticks for the family $\mathcal{S}^{|v|,\mathbf{p}_v} \to \mathcal{R}^{|v|,\mathbf{p}_v}$. The resulting gluing map has the form

$$\mathcal{R}^{T,\mathbf{F}} \times \Big(\prod_e (-1, 0] \Big) \supset \mathcal{U}^{T,\mathbf{F}} \longrightarrow \overline{\mathcal{R}}^{d+1,p}. \tag{20}$$

As in the previous case (10), it will turn out that this is smooth, and a diffeomorphism onto a neighbourhood of $\mathcal{R}^{T,\mathbf{F}}$. Strictly speaking, we will prove this only under the additional assumption that the strip-like ends used are rational ones (Corollary 6.4). However, the general case is only slightly more involved, and interested readers (if any) should be easily able to extend the given argument.

Still proceeding in parallel with the previous discussion, we now introduce the relevant notions of universal choices. Suppose first that we have made a universal choice of strip-like ends for pointed discs, which should be consistent. Then, a universal choice of popsicle sticks is a choice of a set $\mathbf{Q}_S$ for every stable pointed disc S, varying smoothly over the moduli space. There is an appropriate notion of consistency, which as before means that the popsicle sticks are compatible with the gluing construction, for small values of the gluing parameter. Consistent choices can be constructed inductively. Moreover, there is an obvious way to include the case $d = 1$ as well: for $S \cong Z$, take the popsicle stick to be $\{t = \frac{1}{2}\}$. The final step is to pick, for every stable popsicle $(S, \boldsymbol{\phi})$ of some flavour $\mathbf{p}$, a set of strip-like ends $\boldsymbol{\epsilon}_{S,\boldsymbol{\phi}}$, which are equivalent to those previously chosen for the underlying pointed disc S if $d \geq 2$, and which in the remaining case $d = 1$ lie in the equivalence class from Addendum 2.4. As usual, the new ends should vary smoothly over the relevant moduli spaces. Additionally, they should be invariant under the action of (13). Given choices of strip-like ends and popsicle sticks, one then has a gluing construction for popsicle structures, and can

impose the relevant consistency condition on the $\epsilon_{S,\phi}$. As before, the structure of the maps (20) is crucial in proving that this condition can be met. We omit the details of this, since the argument is essentially the same as that for pointed discs.

Remark 2.8 The choice of $\mathbf{Q}_S$ itself only involves the underlying surface S. In other words, it is pulled back through the forgetful map (14). In principle, it would be possible to lift that restriction, making the popsicle sticks depend on ϕ as well, but there seems little point in doing so. The situation for strip-like ends is slightly different. While the equivalence class of $\epsilon_{S,\phi}$ depends only on S, the specific choice of end varies with the popsicle structure. This is necessary in order to achieve consistency, because there are stable popsicles (S,ϕ) with unstable underlying disc $S \cong Z$ (and moreover, even if one restricts to $d \geq 2$, such popsicles occur as components in the compactification of $\mathcal{R}^{d+1,\mathbf{p}}$).

We want to fix these auxiliary structures once and for all. To reiterate, this involves: first, a universal choice of strip-like ends for pointed discs (supposed to be consistent, and for technical reasons, also rational); second, a consistent universal choice of popsicle sticks; and third, a consistent universal choice of strip-like ends for popsicles. All these choices will be kept constant throughout the rest of the paper.

2.5 Closed one-forms

We now introduce a class of one-forms on discs, which can be seen as generalizations of $dt \in \Omega^1(Z)$. Namely, let S be a disc equipped with a set of strip-like ends. By a set of *basic closed one-forms* for S we mean a collection $\alpha = \{\alpha^1, \dots, \alpha^d\}$, where each $\alpha^k \in \Omega^1(S)$ is closed, with $\alpha^k|\partial S = 0$, and satisfies

$$(\epsilon^j)^*\alpha^k = \begin{cases} dt & \text{if } j = 0 \text{ or } k, \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

over the ends (more precisely, on a subset of $Z_\pm$ where $\pm s \gg 0$).

Basic closed one-forms behave well with respect to the gluing process (10). To spell this out, let $S_\pm$ be pointed discs with $d_\pm + 1$ ends, respectively, equipped with strip-like ends and basic closed one-forms. Suppose that we glue together the ends corresponding to the points at infinity ζ_-^0 and ζ_+^i, for some $i \in \{1, \dots, d_+\}$, using a small gluing parameter. The result is a disc S with $d + 1 = d_- + d_+$ ends. This will come with induced basic closed one-forms, which are constructed as follows:

$$\begin{cases} \text{in the case } k < i, \text{ take } \alpha_+^k \text{ and extend it trivially (by zero);} \\ \text{for } k \geq i + d_- \text{ do the same, but starting with } \alpha_+^{k-d_-+1}; \\ \text{for } i \leq k < i + d_-, \text{ take } \alpha_+^i \text{ and glue it together with } \alpha_-^{k-i+1}. \end{cases} \tag{22}$$

For a fixed S, it is clear that basic closed one-forms exist, and in fact form a contractible space. In the stable range $d+1 \geq 3$, we can choose such one-forms $\boldsymbol{\alpha}_S = \{\alpha_S^k\}$ for each pointed disc S, varying smoothly over the moduli space (and, of course, compatibly with our fixed universal choice of strip-like ends). Moreover, this can be done in such a way that the obvious consistency condition, for all gluing processes with small gluing parameter, is satisfied. Finally, returning to the original motivation, we extend this to $d = 1$ by setting $\alpha_Z^1 = dt$.

2.6 Sub-closed one-forms

Let $(S, \boldsymbol{\phi})$ be a $\mathbf{p}$–flavoured popsicle, equipped with strip-like ends. A set $\boldsymbol{\beta} = \{\beta_f\}$ of *basic sub-closed one-forms* consists of a one-form β_f for each $f \in F$, such that $\beta_f|\partial S = 0$, $d\beta_f = 0$ in a neighbourhood of ∂S, and $d\beta_f \leq 0$ everywhere (this is sub-closedness; the sign is with respect to the complex orientation of S). The analogue of (21) is the requirement that at infinity,

$$(\epsilon^j)^* \beta_f = \begin{cases} dt & \text{if } j = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{23}$$

This may seem puzzling at first sight, because f does not appear on the right hand side of (23); however, looking at the gluing process provides some clarification. Take two popsicles $S_\pm$ with flavours $\mathbf{p}_\pm$. These should be equipped with strip-like ends, popsicle sticks, and basic closed as well as sub-closed one-forms. Glue $S_\pm$ together as before, using the process from Section 2.4 to produce a popsicle structure on the resulting disc S. Recall that the flavour $\mathbf{p}$ of the new structure is indexed by the disjoint union $F = F_+ \sqcup F_-$. We then define sub-closed one-forms β_f on the surface obtained by gluing S_+ to S_- at the ith incoming marked point of S_+ as follows:

$$\begin{cases} \text{if } f \in F_+, \text{ take the given } \beta_{+,f} \text{ and extend it trivially;} \\ \text{if } f \in F_-, \text{ take } \beta_{-,f} \text{ and glue it together with } \alpha_+^i. \end{cases} \tag{24}$$

Take a universal choice of sub-closed one-forms $\{\boldsymbol{\beta}_{S,\boldsymbol{\phi}}\}$. We will always assume that this choice is invariant under the action of $Sym^{\mathbf{p}}$ (which is unproblematic, since sub-closed one-forms satisfying (23) form a convex subspace, where averages can be formed). Moreover, we can impose a consistency condition with respect to the gluing process introduced above (and our previous choice of strip-like ends as well as popsicle sticks); choices satisfying this condition can be constructed by applying the usual process, based on the recursive structure of the compactification (16).

Example 2.9 For ease of visualization, it's better to think of the two-forms $d\beta_f$, rather than the β_f themselves. For instance, consistency (compatibility with gluing)

has a straightforward meaning, because each $d\beta_f$ is compactly supported. Let's take a concrete look at the simplest situations. The trivial situation is when $d = 1$ and $|F| = 1$. The (unique) point in that moduli space is represented by a surface $S \cong Z$, and we can fix the isomorphism in such a way that the unique sprinkle gets mapped to $(0, 1/2)$. We then have a single two-form $d\beta \leq 0$ supported on some compact subset of $int(Z)$. In principle, this can be chosen arbitrarily, subject only to the condition that $\int_Z d\beta = -1$, but it is convenient to imagine its support lying in a small neighbourhood of $(0, 1/2)$. Now take $d = 1$ and $F = \{1, 2\}$. A $\mathbf{p}$–flavoured popsicle is a surface $S \cong Z$ equipped with two sprinkles, whose images in Z are $(s_1, 1/2)$ and $(s_2, 1/2)$. The modular parameter is $\sigma = s_1 - s_2$, giving an identification $\mathcal{R}^{2,\mathbf{p}} \cong \mathbb{R}$. On the subset where $|s_1 - s_2| \gg 0$, one starts with $d\beta_1$ and $d\beta_2$ both being copies of the previously introduced $d\beta$, translated away from each other exactly by the amount σ, so that each $d\beta_f$ is supported near the corresponding sprinkle. Over the rest of the moduli space, one interpolates between those given choices, while preserving the symmetry which exchanges both the sprinkles and the $d\beta_f$.

As before we end our discussion by fixing, once and for all, a consistent universal choice of closed basic forms $\boldsymbol{\alpha}_S$ and of sub-closed basic forms $\boldsymbol{\beta}_{S,\boldsymbol{\phi}}$.

2.7 Weighted popsicles

Take some $d \geq 1$ and $\mathbf{p} = \{p_f\}$, $f \in F$. A set of *weights* is a collection of positive integers $\mathbf{w} = \{w^0, \dots, w^d\}$ satisfying

$$w^0 = w^1 + \cdots + w^d + |F|. \tag{25}$$

The primary use of this concept is as follows. Suppose that $(S, \boldsymbol{\phi})$ is a $\mathbf{p}$–flavoured popsicle, which also comes with strip-like ends $\boldsymbol{\epsilon}$, basic closed one-forms $\boldsymbol{\alpha}$, and sub-closed one-forms $\boldsymbol{\beta}$. Given weights, one then defines the *total sub-closed one-form* $\gamma \in \Omega^1(S)$ to be

$$\gamma = \sum_{k=1}^{d} w^k \alpha^k + \sum_{f \in F} \beta_f. \tag{26}$$

By construction, see (21), (23) and (25), this has the following properties:

(27) $\gamma|\partial S = 0$. Moreover, $d\gamma$ is non-positive everywhere, and vanishes in a neighbourhood of ∂S. Finally, on the strip-like ends we have $(\epsilon^k)^*\gamma = w^k dt$, at all points (s, t) with $\pm s \gg 0$.

Recall that in Section 2.6 we made a universal choice of basic closed and sub-closed one-forms. This determines, for each stable popsicle with weights, a total sub-closed

one-form (26), which we denote by $\gamma_{S,\boldsymbol{\phi},\mathbf{w}}$. Note that by construction, our choice of total sub-closed one-form is compatible with the action of (13) on the moduli space. Concretely, this means that if we have two weighted popsicles which differ only by a reordering of the maps ϕ_f, then they carry the same total sub-closed one-form. Finally, in the unstable case where $S \cong Z$ and $F = \varnothing$, hence $\mathbf{w} = \{w^0 = w^1 = w\}$, we set

$$\gamma_{Z,\mathbf{w}} = w\, dt. \tag{28}$$

Because of the previously imposed consistency conditions on the one-forms α and β, their sums γ are compatible with gluing processes (assuming that any two ends which are glued together carry the same weight; and that the gluing parameters remain small).

We write $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$ for the moduli space of stable popsicles with weight $\mathbf{w}$. Of course, this is just a copy of $\mathcal{R}^{d+1,\mathbf{p}}$, but the separate notation is still useful. For instance, if we take the compactification $\overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}$ defined as in (16), then the boundary strata are canonically products of lower-dimensional moduli spaces of weighted popsicles. To see this explicitly, fix a set of weights $\mathbf{w}$, and consider a pair $(T,\mathbf{F})$ which labels one of the boundary strata. Recall from Section 2.1 that pairs (v,k), where v is a vertex and k is a number in $\{0,\dots,|v|-1\}$, correspond canonically to flags of T. We can then define induced weights $\mathbf{w}_v = \{w_v^k\}$ for every vertex v, by the following conditions: first, if (v,k) corresponds to the jth semi-infinite edge of T, then $w_v^k = w^j$; secondly, two flags belonging to the same finite edge should carry the same weight; and finally, the analogue of (25) should be satisfied at each vertex, meaning that $w_v^0 = w_v^1 + \cdots + w_v^{|v|-1} + |F_v|$. In these terms, the boundary stratum of $\overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}$ indexed by $(T,\mathbf{F})$ is isomorphic to the product of the spaces $\overline{\mathcal{R}}^{|v|,\mathbf{p}_v,\mathbf{w}_v}$.

2.8 Cascades

What follows now is a bad case of mixed metaphor. Fix $(d,\mathbf{p})$. Take a pair $(T,\mathbf{F})$ as in the definition of broken popsicle, and consider the induced flavours $\{\mathbf{p}_v\}$. A *cascade* modelled on $(T,\mathbf{F})$ is a collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v)\}$, where each $(S_v,\boldsymbol{\phi}_v)$ is a $\mathbf{p}_v$–flavoured popsicle, equipped with an additional parameter value $\rho_v \in (0,1]$. These parameters are subject to the following *causal ordering* condition:

(29) If a finite edge of T has endpoints $v_\pm$, with v_+ being closer to the root than v_-, then $\rho_{v_+} < \rho_{v_-}$.

We can impose a stability condition, which is the same as for broken popsicles, and take the disjoint union of the resulting moduli spaces for all stable choices $(T,\mathbf{F})$. We call the resulting space *moduli space of stable cascades* and denote it by $\mathcal{Q}^{d+1,\mathbf{p}}$. Note that even though the notation is analogous to (16), the structure of the moduli space

is entirely different, since each $(T, \mathbf{F})$ defines a connected component of the same dimension $d-1+|F|$.

There is also a natural partial compactification, which is a manifold with corners denoted by $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$. The strata of this space are indexed by combinatorial data

$$(T, \mathbf{F}, \bar{E}, E) \tag{30}$$

where $(T, \mathbf{F})$ is as before, and $\bar{E} \subset E$ are subsets of the set of finite edges of T. A point of such a stratum is again given by a collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v)\}$, with the difference that we now require the parameter values to satisfy (29) only if the edge connecting $v_\pm$ is not in E, whereas for the other edges the equality $\rho_{v_+} = \rho_{v_-}$ should hold. The codimension of the stratum is then precisely $|E|$. To understand the topology of the partial compactification, note that one can associate to each datum (30) a collapsed tree $\overline{T}$, obtained by contracting the edges in $\bar{E}$ to points. For any vertex $\overline{v}$ of $\overline{T}$, let $\bar{F}_{\overline{v}}$ be the union of F_v over all vertices v mapped to $\overline{v}$ under the collapsing map. Denote the resulting decomposition of F by $\overline{\mathbf{F}}$. Reversing the direction of the argument, take a fixed $(\overline{T}, \overline{\mathbf{F}})$, with its associated collection of labels $\{\overline{\mathbf{p}}_{\overline{v}}\}$, and consider all choices of (30) which yield the given pair after collapsing. The union of all the associated strata in $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$ is a copy of

$$\Delta^{\overline{T}} \times \prod_{\overline{v}} \overline{\mathcal{R}}^{|\overline{v}|, \overline{\mathbf{p}}_{\overline{v}}}, \tag{31}$$

where $\Delta^{\overline{T}}$ is the set of all parameter values $\{\overline{\rho}_{\overline{v}}\}$ indexed by vertices, such that in the situation analogous to (29), the weak inequality $\overline{\rho}_{\overline{v}_+} \leq \overline{\rho}_{\overline{v}_-}$ holds. Then (31), with its natural topology, is a connected component of $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$. As a direct consequence, the map

$$\begin{aligned} &\overline{\mathcal{Q}}^{d+1,\mathbf{p}} \longrightarrow (0,1], \\ &\{(\rho_v, S_v, \boldsymbol{\phi}_v)\} \longmapsto \rho_{v_0}, \end{aligned} \tag{32}$$

where v_0 is the vertex closest to the root, is proper. This means that the non-compactness of $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$ only reflects that of the parameter interval $(0,1]$.

It is important to note that a given collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v)\}$ can appear more than once in $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$. The reason is that, while E is determined by the relation between parameters ρ_v, the choice of $\bar{E} \subset E$ is free. In particular, codimension one faces defined by the equality of exactly two parameters occur in isomorphic pairs (the two faces in such a pair belong to different connected components of $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$).

Example 2.10 In the simplest example, T is a tree with two vertices $v_\pm$, and we have corresponding triples $\{(\rho_\pm, S_\pm, \boldsymbol{\phi}_\pm)\}$ with $\rho_+ = \rho_- = \rho$, so that E consists

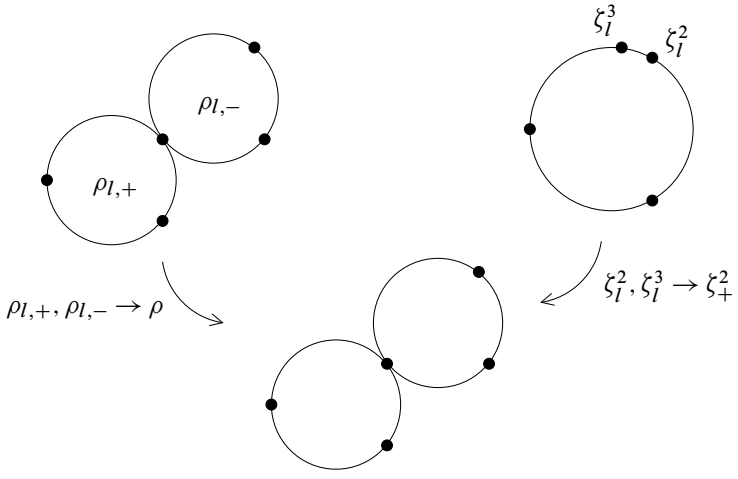

Figure 5

of the unique finite edge of the tree. If one then sets $\overline{E} = \varnothing$, the resulting point in $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$ is the limit of a sequence $\{(\rho_{l,\pm}, S_{l,\pm}, \boldsymbol{\phi}_{l,\pm})\}$ with the same combinatorial data $(\overline{T}, \overline{\mathbf{F}}) = (T, \mathbf{F})$, and where $\rho_{l,\pm}$ become equal in the limit $l \to \infty$. The other possibility is $\overline{E} = E$. In that case $\overline{T}$ has a single vertex, so the resulting point in $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$ is the limit of a sequence $(\rho_l, S_l, \boldsymbol{\phi}_l)$ in which $\rho_l \to \rho$, and the popsicle $(S_l, \boldsymbol{\phi}_l)$ degenerates into the broken popsicle given by $\{(S_\pm, \boldsymbol{\phi}_\pm)\}$ (see Figure 5 for a schematic picture).

Finally, $\mathcal{Q}^{d+1,\mathbf{p}}$ carries an action of $Sym^{\mathbf{p}}$. If we fix a connected component, and restrict to the subgroup of permutations which preserve each subset F_v, then this is just the product of the previously defined actions on the factors $\mathcal{R}^{|v|,\mathbf{p}_v}$, plus trivial actions on $(0,1]$. Outside that subgroup, the action exchanges different connected components. The same applies to the partial compactification.

3 Wrapped Floer cohomology

In this section, we construct the wrapped Floer cochain complex and its A_∞–structure. As mentioned before, this should be thought of as a kind of direct limit, although the definition itself is not explicitly formulated in such terms. We emphasize again that, even though the cochain level construction itself is new, the underlying cohomology level theory is well-known to specialists, as one will see by looking at Abbondandolo–Schwarz [1]. Our focus will be on the geometry of the relevant moduli spaces of

(deformed) pseudo-holomorphic maps. The most interesting aspect is that the contribution of most of these spaces is zero, because of cancellations induced by orientation-reversing symmetries; in fact, this cancellation phenomenon is crucial in order to obtain the A_∞–associativity relations. Analytic issues, which do not go beyond the standard technological level, are relegated to Sections 7–9. To conclude, we should mention that, in principle, there is an alternative and more direct definition of wrapped Floer cohomology using a single Hamiltonian function with unboundedly increasing slope instead of a sequence of functions with constant slope. However, the alternative approach has its own technical quirks and is not as suitable for discussing Viterbo functoriality.

3.1 Basic geometry

Let M be a Liouville domain. By definition, M is a compact $2n$–dimensional manifold with boundary, together with a one-form θ such that $\omega = d\theta$ is symplectic, and the dual Liouville vector field Z, uniquely determined by the requirement that $i_Z\omega = \theta$, points strictly outwards along ∂M. This implies that $\theta|\partial M$ is a contact one-form. By flowing inward from the boundary along Z, one obtains a collar embedding $(0,1]\times\partial M \hookrightarrow M$ modelled on the small half of the symplectization of ∂M. One can then complete M by attaching the big half:

$$\widehat{M} = M \cup_{\partial M} ([1,\infty)\times\partial M). \tag{33}$$

The piece $[1,\infty)\times\partial M$ of (33) is called the *infinite cone*. The completion carries a natural one-form $\widehat{\theta}$ such that $d\widehat{\theta} = \widehat{\omega}$ is symplectic, and an associated Liouville vector field $\widehat{Z}$. On $M \subset \widehat{M}$ these restrict to the previously given data, while on the overlapping piece $(0,\infty)\times\partial M \subset \widehat{M}$ we have $\widehat{\theta} = r(\theta|\partial M)$ and $\widehat{Z} = r\partial_r$; here r is the variable in $(0,\infty)$. We will use a particular class of Hamiltonian functions $H \in C^\infty(M,\mathbb{R})$, namely those for which:

(34) $H > 0$ everywhere. Moreover, H should admit a smooth extension $\widehat{H}$ to the completion $\widehat{M}$ such that $\widehat{H}(r,y) = r$ on the infinite cone.

Let $\widehat{X}$ be the Hamiltonian vector field of $\widehat{H}$. On the infinite cone we have $\widehat{X} = (0, R)$, where R is the Reeb vector field associated to $\theta|\partial M$. By restricting, one sees that the Hamiltonian vector field X of H satisfies $X|\partial M = R$.

Let $L \subset M$ be a Lagrangian submanifold intersecting ∂M transversally, and which has the following property:

(35) $\theta|L \in \Omega^1(L)$ is exact, $\theta|L = df$. Moreover, $\theta|L \in \Omega^1(L)$ vanishes to infinite order along the boundary $\partial L = L \cap \partial M$.

The second part implies that ∂L is a Legendrian submanifold of ∂M, and also that, by attaching an infinite cone to that boundary, one can extend L to a smooth noncompact Lagrangian submanifold

$$\widehat{L} = L \cup_{\partial L} ([1,\infty) \times \partial L) \subset \widehat{M}. \tag{36}$$

Recall that a Reeb chord of length $w > 0$ is a trajectory $x\colon [0,1] \to \partial M$ of wR, such that $x(0), x(1) \in \partial L$. An integer Reeb chord is one which has integer length. We will make the following technical assumption:

(37) There are no integer Reeb chords.

In parallel with the previous terminology, an *integer X–chord* is a trajectory $x\colon [0,1] \to M$ of wX, for some positive integer w, such that $x(0), x(1) \in L$. In these terms, (37) says that all such trajectories must lie in the interior of M. Additionally, we assume:

(38) All integer X–chords are nondegenerate.

(39) No point of L is at once a starting point of an integer X–chord, and an endpoint of one (either the same chord or a different one).

Note that as a consequence of either of these two assumptions, no critical point of H may lie on L. Besides that, we will also impose certain topological conditions which are traditional in Floer theory:

(40) The relative Chern class $2c_1(M, L) \in H^2(M, L; \mathbb{Z})$ and the second Stiefel–Whitney class $w_2(L) \in H^2(L; \mathbb{Z}/2)$ both vanish.

It seems appropriate to briefly discuss the significance of these conditions. There are at least two alternative versions of (35). One is stronger:

(41) $\theta|L$ is exact and vanishes on a neighbourhood of ∂L.

The second part implies that the Liouville flow is parallel to L near its boundary, making it locally into a cone over ∂L. The other version is weaker:

(42) $\theta|L$ is exact and the restriction $\theta|\partial L \in \Omega^1(\partial L)$ is zero.

The second part of this is equivalent to saying that ∂L is Legendrian. While (42) may be the most natural assumption, it is not directly suitable for our construction. On the other hand, simply by allowing isotopies which are constant on the boundary, the difference between any of these versions disappears.

Lemma 3.1 *Let L_0 be a Lagrangian submanifold satisfying (42). Then there is a Hamiltonian isotopy rel ∂M, which deforms L_0 into another submanifold L_1 satisfying (41).*

Proof We start with a construction which is local near ∂L_0. By assumption, $\theta|L_0$ vanishes along ∂L_0, hence is exact nearby. This means that we can find a function $h \in C^\infty(M, \mathbb{R})$ vanishing along ∂M, such that $(\theta + dh)|L_0 \in \Omega^1(L_0)$ is zero in a neighbourhood of ∂L_0. The Hamiltonian vector field Y of h is parallel to ∂M. Hence, the Liouville vector field of $\theta_t = \theta + (1-t)dh$, which is $Z_t = Z + (1-t)Y$, still points outwards along ∂M. Take ∂L_0 and push it inwards using $-Z_t$, for times in some interval $[0, \infty)$. The orbits of each such flow yield a local Lagrangian submanifold L_t^{loc} with boundary $\partial L_t^{\mathrm{loc}} = \partial L_0$. More precisely, we consider L_t^{loc} as a germ defined near the boundary, hence reserve the right to shrink it whenever necessary. For $t = 0$ this germ agrees with that of L_0, since Z_0 is parallel to that submanifold; on the other hand, for $t = 1$ we have $\theta|L_1^{\mathrm{loc}} = 0$ by construction.

This local Lagrangian isotopy is described by a family of closed one-forms $\beta_t \in \Omega^1(L_t^{\mathrm{loc}})$, each of which is zero along the boundary. Take functions $g_t \in C^\infty(M, \mathbb{R})$ vanishing along the boundary, such that $dg_t|L_t^{\mathrm{loc}} = \beta_t$. By construction, the Hamiltonian isotopy generated by the family g_t is constant along the boundary, and moves L_0 to a Lagrangian submanifold L_1 which extends the previously given germ L_1^{loc}. □

Next, among the more technical conditions, (37) can be satisfied by a small rescaling of ω and θ, simply because the periods of all Reeb chords form a measure zero subset of $\mathbb{R}$ (by a suitable application of Sard's theorem). Suppose that this has been done, and additionally that $2n = \dim M \geq 4$. Then, a generic choice of H ensures that (38) and (39) hold (Lemmas 8.1 and 8.2, respectively). Unfortunately, in the two-dimensional case the second of these conditions is not generic, and a workaround is needed (Section 5.2). Finally, one can drop (40) altogether, at the cost of the usual loss of Floer-theoretic structure. This would mean working with coefficients in $\mathbb{Z}/2$ and obtaining only an ungraded A_∞–structure, since degrees and signs can no longer be assigned consistently.

3.2 Almost complex structures

We begin by recalling a classical definition from Hofer [17]. Suppose that we are given an almost complex structure on the contact hyperplane field $\xi = \ker(\theta|\partial M) \subset T(\partial M)$, compatible with its symplectic structure. From that, one constructs an almost complex structure on the symplectization $(0, \infty) \times \partial M$ as follows: write

$$T((0, \infty) \times \partial M) = \mathbb{R}^2 \oplus \xi, \tag{43}$$

where the first factor is spanned by $r\partial_r$ and $(0, R)$. Take the direct sum of the standard complex structure on the first summand, meaning that $r\partial_r \mapsto (0, R)$, and of the given one on the second summand. The almost complex structures on the symplectization

constructed in this way are called *of contact type*. Any such structure is invariant under the flow of $r\partial_r$.

Take a $\widehat{\omega}$–compatible almost complex structure $\widehat{I}$ on $\widehat{M}$, whose restriction to the infinite cone is of contact type. The restriction $I = \widehat{I}|M$ is called an almost complex structure *of contact type at the boundary*. Note that I determines $\widehat{I}$ uniquely. Denote the space of all such structures by $\mathcal{J}(M)$. This is an infinite-dimensional manifold in a loose sense. Its tangent space $T\mathcal{J}(M)$ at any point I consists of endomorphisms K of TM which anti-commute with I, which have the property that $\omega(\cdot, K\cdot)$ is symmetric, and which moreover are restrictions of suitable endomorphisms of $T\widehat{M}$ (we omit the details of the last condition, which are straightforward but lengthy). We call these endomorphisms *infinitesimal deformations* of the almost complex structure. One can get an actual deformation by exponentiating K, which means setting

$$J = I\,\exp(-IK). \tag{44}$$

Let S be a disc equipped with a set of strip-like ends $\boldsymbol{\epsilon}$. Let $\mathbf{I} = \{I_z\}$ be a family of almost complex structures in $\mathcal{J}(M)$ parametrized by points $z \in S$. We say that this family is *strictly compatible with the strip-like ends* if for each $k \in \{0, \dots, d\}$ there is a family of such structures $\mathbf{I}^k = \{I^k_t\}$ parametrized by $t \in [0,1]$, such that

$$I_{\epsilon^k(s,t)} = I^k_t \tag{45}$$

when $\pm s \gg 0$. In other words, if we go sufficiently far along any end, $\mathbf{I}$ becomes independent of s, hence compatible with translations in the s direction.

A universal choice of almost complex structures consists of the following data. For each stable weighted popsicle $(S, \boldsymbol{\phi}, \mathbf{w})$, choose a family $\mathbf{I}_{S,\boldsymbol{\phi},\mathbf{w}}$ parametrized by S, and strictly compatible with (the previous universal choice of) strip-like ends. We require that these structures should vary smoothly over the moduli spaces $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$, and be invariant with respect to the action of $Sym^{\mathbf{p}}$. Additionally, for each w choose a family $\mathbf{I}_w$ parametrized by $[0,1]$. This can actually be viewed as an extension of the previous choice to the unstable case, invariant under automorphisms:

$$I_{Z,\mathbf{w},s+it} = I_{w,t}, \tag{46}$$

where $\mathbf{w} = \{w^0 = w^1 = w\}$. A universal choice is called consistent if the following two conditions are satisfied. First of all, over the kth strip-like end the behaviour of $\mathbf{I}_{S,\boldsymbol{\phi},\mathbf{w}}$, in the sense of (45), is actually governed by $\mathbf{I}_{w^k}$. Secondly, the choice is compatible with gluing together weighted popsicles, for small values of the gluing parameter. Here, the gluing process is defined by our previous universal choices of strip-like ends and popsicle sticks, as in Section 2.4 (and carrying over the weights in the obvious way).

Remark 3.2 The consistency requirements ensure the existence of well-behaved compactifications for the moduli spaces of holomorphic maps to be defined later on. However, they are not convenient from the point of view of transversality results, because the behaviour of the almost complex structure on parts of the Riemann surface is restricted (in a non-local way, even). To see an example of this, suppose that $(S, \boldsymbol{\phi}, \mathbf{w})$ is obtained from $(S_+, \boldsymbol{\phi}_+, \mathbf{w}_+)$ and two copies of $(S_-, \boldsymbol{\phi}_-, \mathbf{w}_-)$ (glued into different positive ends of S_+, which necessarily should have the same weight). Then, consistency dictates that the family of almost complex structures associated to S should be equal on the two pieces coming from S_-. To avoid these problems, we will introduce perturbations of the almost complex structures which allow the conditions to be relaxed, while retaining enough asymptotic control for the purposes of compactification.

Take a disc with a set of strip-like ends, and suppose that we have chosen $\mathbf{I}$ on it satisfying (45). An infinitesimal deformation of $\mathbf{I}$ is a family $\mathbf{K} = \{K_z\}$, where each K_z lies in the tangent space to $T\mathcal{J}(M)$ at I_z, and with the property that as $z = \epsilon^k(s,t)$ goes to infinity along a strip-like end, K_z and all its derivatives go to zero faster than any exponential function $\exp(-C|s|)$. Equivalently, let $\overline{\mathbf{K}}$ be the trivial extension of $\mathbf{K}$ to $\overline{S}$; the condition is that this extension is smooth, and vanishes to infinite order at the points at infinity. Then, the perturbed family $\mathbf{J} = \{J_z\}$ obtained by applying (44) at every point of S agrees asymptotically with the original $\mathbf{I}$, with good control over the rate of convergence.

We now consider the corresponding notion for families. Assume we have already made a universal choice of almost complex structures, assumed to be consistent. A universal infinitesimal deformation consists of an infinitesimal deformation $\mathbf{K}_{S,\boldsymbol{\phi},\mathbf{w}}$ of each $\mathbf{I}_{S,\boldsymbol{\phi},\mathbf{w}}$, which varies smoothly with respect to the moduli, and is compatible with the action of $Sym^{\mathbf{p}}$. In slightly different terms, the universal infinitesimal deformation is given by a $Sym^{\mathbf{p}}$–invariant section of the pullback bundle

$$End(TM) \longrightarrow \mathcal{S}^{d+1,\mathbf{p},\mathbf{w}} \times M \tag{47}$$

for each $(d, \mathbf{p}, \mathbf{w})$. We say that the universal infinitesimal deformation is *asymptotically consistent* if the sections extend smoothly to the compactifications $\overline{\mathcal{S}}^{d+1,\mathbf{p},\mathbf{w}} \times M$, and the following two additional properties hold.

For the first condition, fix a pair $(T, \mathbf{F})$ which represents a boundary stratum $\mathcal{R}^{T,\mathbf{F},\mathbf{w}} \subset \overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}$. For every vertex v of T, with its associated $\mathbf{p}_v$ and weights $\mathbf{w}_v$, we take the projection $\pi_v\colon \mathcal{R}^{T,\mathbf{F},\mathbf{w}} \longrightarrow \mathcal{R}^{|v|,\mathbf{p}_v,\mathbf{w}_v}$, use that to pull back the universal family, then take the disjoint union over all v. By definition of the compactification, there is a

natural embedding

$$\coprod_v \pi_v^* \mathcal{S}^{|v|,\mathbf{p}_v,\mathbf{w}_v} \longrightarrow \overline{\mathcal{S}}^{d+1,\mathbf{p},\mathbf{w}}. \tag{48}$$

Given a universal infinitesimal deformation, there are two ways of making a section of $End(TM) \to \coprod_v \pi_v^* \mathcal{S}^{|v|,\mathbf{p}_v,\mathbf{w}_v} \times M$. We could take the sections attached to the smaller infinitesimal families $\mathcal{S}^{|v|,\mathbf{p}_v,\mathbf{w}_v}$ and combine them; on the other hand, we could take the section attached to $\mathcal{S}^{d+1,\mathbf{p},\mathbf{w}}$, extend it to the compactification, and then pull it back via (48). The condition then says that these two constructions yield the same result, for all $(T,\mathbf{F})$.

For the second condition, take a point of the compactified moduli space $\overline{\mathcal{S}}^{d+1,\mathbf{p},\mathbf{w}}$ which is represented by a pair $(\overline{S}, z)$, where z is either a point at infinity of S or (in the case of a broken popsicle) a singular point. Then, our section should vanish when restricted to that point (times M), and so should all its derivatives.

From now on, we will work with a fixed consistent universal choice of almost complex structures, as well as a fixed asymptotically consistent infinitesimal deformation. By applying (44) at each point of S, one can then form perturbed families $\mathbf{J}_{S,\boldsymbol{\phi},\mathbf{w}}$. By definition, this kind of perturbation does not affect the limits of our almost complex structures over the ends of S, so we set

$$\mathbf{J}_w = \mathbf{I}_w, \ \ \mathbf{J}_{Z,\mathbf{w}} = \mathbf{I}_{Z,\mathbf{w}} \text{ for } \mathbf{w} = (w^0 = w^1 = w). \tag{49}$$

Remark 3.3 To understand the concrete implications of asymptotic consistency, one has to look at the geometry of the compactified universal family, and more specifically at the behaviour of functions with similar vanishing conditions at special points. This is done in Section 6.4, and the outcome of that discussion can be summarized as follows. As one approaches a point in the boundary of $\overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}$, the fibre of the universal family acquires "necks" which are finite strips of increasing length. On each of these finite strips, as well as on the strip-like ends, the infinitesimal deformation $\mathbf{K}_{S,\boldsymbol{\phi},\mathbf{w}}$ decays faster than exponentially, in the sense of (141), (142). Hence, $\mathbf{J}_{S,\boldsymbol{\phi},\mathbf{w}}$ converges everywhere to the structure which is inherited from the components of the limiting broken popsicle, and moreover, the rate of convergence is precisely controlled on the finite strips and strip-like ends.

3.3 Pseudo-holomorphic maps

For each w, let $\mathcal{X}_w$ be the set of integer X–chords of length w. From (38) we know that each set $\mathcal{X}_w$ is finite.

Take a weighted popsicle $(S, \boldsymbol{\phi}, \mathbf{w})$, equipped with strip-like ends, as well as sets of basic closed and sub-closed one-forms. Let γ be the induced total sub-closed one-form (26). In addition, we assume that we have a family $\mathbf{I}$ which is compatible with the strip-like ends, as well as an infinitesimal deformation $\mathbf{K}$, which together define the perturbed family $\mathbf{J}$. Pick $\mathbf{x} = \{x^0, \dots, x^d\}$, where each x^k is an element of $\mathcal{X}_{w^k}$. Consider the inhomogeneous $\overline{\partial}$–equation

$$\begin{cases} u\colon S \longrightarrow M, \qquad u(\partial S) \subset L, \\ \lim_{s\to\pm\infty} u(\epsilon^k(s,\cdot)) = x^k(\cdot), \\ (du_z - X_{u(z)} \otimes \gamma_z) \circ j + J_{z,u(z)} \circ (du_z - X_{u(z)} \otimes \gamma_z) = 0. \end{cases} \tag{50}$$

Here j is the complex structure on S; $X \otimes \gamma$ is the section of $\mathrm{Hom}(TS, u^*TM)$ obtained by composing $\gamma \in C^\infty(TS^*)$ with $u^*X \in C^\infty(u^*TM)$; and the overall condition says that when we subtract that inhomogeneous term from du, the outcome is of type $(1,0)$ with respect to j and J_z. We will sometimes abbreviate this equation by

$$(du - X \otimes \gamma)^{0,1} = 0. \tag{51}$$

Fix $d \geq 1$ and $\mathbf{p} = \{p_f\}_{f\in F}$, assuming first that these satisfy the stability condition $d + |F| \geq 2$. Choose weights $\mathbf{w}$ satisfying (25), and limits $\mathbf{x}$ as before. The *space of stable popsicle maps*, denoted by

$$\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}) \quad (\text{or } \mathcal{R}^{d+1,\mathbf{w}}(\mathbf{x}) \text{ if } F = \varnothing), \tag{52}$$

is the set of triples $(S, \boldsymbol{\phi}, u)$ of the following kind. $(S, \boldsymbol{\phi})$ is a $\mathbf{p}$–flavoured popsicle, representing a point of $\mathcal{R}^{d+1,\mathbf{p}} = \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$. As such, it carries additional data specified by our previous choices, in particular: ends $\boldsymbol{\epsilon}_{S,\boldsymbol{\phi}}$, a one-form $\gamma_{S,\boldsymbol{\phi},\mathbf{w}}$, and a family of almost complex structures $\mathbf{J} = \mathbf{J}_{S,\boldsymbol{\phi},\mathbf{w}}$. Using all that, one can write down (50) on S, and u is then a solution of that equation, with limits given by $\mathbf{x}$. Note that (52) carries an induced action of $Sym^{\mathbf{p}}$, which consists of permuting the sprinkles ϕ_f while keeping everything else the same.

It remains to extend the definition of the moduli spaces to the case where the underlying popsicle is unstable, namely $d = 1$ and $F = \varnothing$. Equip $S = Z$ with (28) as well as the family $\mathbf{J} = \mathbf{J}_{Z,\mathbf{w}}$ of almost complex structures taken from (49), (46). Then (50) reduces to the classical equation for Floer trajectories:

$$\begin{cases} u\colon Z \longrightarrow M, \qquad u(\partial Z) \subset L, \\ \lim_{s\to-\infty} u(s,\cdot) = x^0(\cdot), \qquad \lim_{s\to+\infty} u(s,\cdot) = x^1(\cdot), \\ \partial_s u + J_{t,u(t)}(\partial_t u - wX) = 0. \end{cases} \tag{53}$$

As usual, one defines $\mathcal{R}^{2,\mathbf{w}}(\mathbf{x})$ to be the space of non-stationary solutions of that equation (solutions for which $\partial_s u$ does not vanish identically), divided by the translational $\mathbb{R}$–action.

Convention 3.4 From now on, we will lump both cases together. This means that, even though we will mostly use the notation $(S, \boldsymbol{\phi}, u)$ which is appropriate for stable popsicles, it is tacitly understood that the discussion includes the case of Floer moduli spaces, with the necessary minor modifications.

3.4 Smoothness

Let $(S, \boldsymbol{\phi}, u)$ be a point in (52). The virtual dimension of the moduli space at that point is

$$\operatorname{vdim} \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}) = \dim \mathcal{R}^{d+1,\mathbf{p}} + \deg(x^0) - \deg(x^1) - \cdots - \deg(x^d). \tag{54}$$

Here, $\deg(x)$ is the Maslov index of $x \in \mathcal{X}_w$ (in general, this depends on a choice of grading). In the case of Floer's equation, the dimension is set to "$\dim \mathcal{R}^2 = -1$", recovering the standard formula.

Theorem 3.5 *Suppose that the almost complex structures $\mathbf{I}_w$ have been chosen generically. Then, for a generic choice of infinitesimal deformations $\mathbf{K}_{S,\boldsymbol{\phi},\mathbf{w}}$, subject to asymptotic compatibility, all the moduli spaces $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ are regular, hence smooth manifolds of the expected dimension (54).*

The action of $Sym^{\mathbf{p}}$ on $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ has the same structure as the underlying action on the moduli space of popsicles. By this, we mean that the isotropy groups which appear are the subgroups $Sym^{\mathbf{P}}$ from Addendum 2.6. The virtual codimension of the corresponding fixed point set is

$$\operatorname{vcodim} \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{Sym^{\mathbf{P}}} = \sum_m (|P_m| - 1). \tag{55}$$

Theorem 3.6 *Generically, in the same sense as in Theorem 3.5, the subspaces of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ fixed by any subgroup $Sym^{\mathbf{P}}$ are themselves regular, hence smooth submanifolds of codimension (55).*

The proofs of Theorems 3.5 and 3.6 are pretty much standard transversality theory (see Section 8.3 for details). From now on, we will assume throughout our discussion that all the almost complex structures involved have been chosen generically, so that the conclusions of these theorems hold.

3.5 Compactness

The spaces (52) admit compactifications whose structure is modelled on (16). Suppose that we have a pair $(T, \mathbf{F})$ as in the definition of broken popsicle. Define flavours $\mathbf{p}_v$ as in Section 2.3, and weights $\mathbf{w}_v$ as in Section 2.7. Moreover, for each v we want to have a collection $\mathbf{x}_v = \{x_v^k\}$, where each $x_v^k \in \mathcal{X}_{w_v^k}$, with the following properties:

(56) If the pair (v,k) correspond to the j th semi-infinite edge of T, then $x_v^k = x^j$ is one of the given X–chords. Otherwise, given a finite edge of T, and the two associated flags (v_+,k) and $(v_-,0)$, we require that $x_{v_+}^k = x_{v_-}^0$.

Finally, for every v we want to have an element

$$(S_v, \boldsymbol{\phi}_v, u_v) \in \mathcal{R}^{|v|,\mathbf{p}_v,\mathbf{w}_v}(\mathbf{x}_v). \tag{57}$$

Then, the collection $\{(S_v, \boldsymbol{\phi}_v, u_v)\}$, which we call a broken popsicle map, defines a point in the compactified moduli space. As the terminology indicates, this space is compact in an appropriately defined Gromov topology. A basic role in the compactness argument is played by an priori estimate for the energy $||du - X \otimes \gamma||_{L^2}$, derived from an inequality between that norm and the action functional; see Section 7.2, or more specifically (149) and (151). Other important players are the various consistency conditions; these ensure that, when the domain degenerates into a broken popsicle, the limiting map satisfies (50) on each component of that popsicle, with the correct almost complex structure and inhomogeneous term. The proof of compactness starts with these observations, and then proceeds exactly as in the case of ordinary Fukaya categories. We will not comment on it further.

The simplest application is when the virtual dimension is zero. Then the boundary of the compactification is empty, hence $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ itself is a finite set. Next, suppose that the virtual dimension is one. Then $\overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ is a one-dimensional compact manifold with boundary. Boundary points correspond to broken pseudo-holomorphic maps, where T is a tree with two vertices, and each component of the broken map is itself part of a zero-dimensional moduli space. Part of this statement requires a suitable gluing theorem, which we again take for granted, given its similarity with established results for Fukaya categories. Another point, which is more specific to the geometric situation here, is the following one: in principle, one could think that there might be other boundary points of one-dimensional moduli spaces, related to the fact that our pseudo-holomorphic maps can hit the boundary of M. It turns out that these points are actually not boundary points, essentially due the convexity of ∂M (see Section 7.4 for a proof; alternatively one could use the maximum principle, as in Oancea [24] or Khovanov–Seidel [20]).

3.6 Algebraic relations

At this point, we fix a coefficient field $\mathbb{K}$. We also choose an orientation of each popsicle moduli space $\mathcal{R}^{d+1,\mathbf{p}}$. Finally, we make a choice of grading and *Pin* structure for our Lagrangian submanifold. The first choice fixes the Maslov indices $\deg(x) \in \mathbb{Z}$ of all $x \in \mathcal{X}_w$, and the second one allows one to associate to any x a one-dimensional $\mathbb{K}$–vector space, the so-called $\mathbb{K}$–*normalized orientation space* $|o_x|_{\mathbb{K}}$ (see Section 9.1 for a brief recapitulation). Moreover, for every point $(S, \boldsymbol{\phi}, u)$ in a zero-dimensional moduli space $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$, we get a preferred isomorphism

$$|o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}|_{\mathbb{K}} \colon |o_{x^d}|_{\mathbb{K}} \otimes \cdots \otimes |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}}. \tag{58}$$

This is, in a slightly more abstract formulation than usual, the sign with which $(S, \boldsymbol{\phi}, u)$ contributes to the algebraic count of points in the moduli space. Denote by $m^{d,\mathbf{p},\mathbf{w}}(\mathbf{x})$ the sum of the homomorphisms (58)

$$m^{d,\mathbf{p},\mathbf{w}}(\mathbf{x}) = \sum_{(S,\boldsymbol{\phi},u)} |o_{x^d}|_{\mathbb{K}}.$$

When $F = \varnothing$ we may omit $\mathbf{p}$ and just write $m^{d,\mathbf{w}}(\mathbf{x})$, in parallel with the notation for the spaces (52) themselves.

Lemma 3.7 *If $Sym^{\mathbf{p}}$ is nontrivial, $m^{d,\mathbf{p},\mathbf{w}}(\mathbf{x})$ vanishes.*

To see why this is the case, recall that Theorem 3.6 implies that all strata of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ consisting of points with nontrivial isotropy group have positive codimension. Hence, in the zero-dimensional case, $Sym^{\mathbf{p}}$ necessarily acts freely on the moduli space. Now, any two points related by a transposition in that group actually contribute with opposite sign, hence cancel out (for details, see Lemma 9.1).

The vanishing argument outlined above applies exactly when the map $\mathbf{p}\colon F \to \{1, \dots, d\}$ fails to be injective (or to put it more intuitively, if there is a popsicle stick carrying more than one sprinkle). We now turn to the remaining injective case. In that situation, it is more natural think of F as a subset of $\{1, \dots, d\}$, with $\mathbf{p}$ being the inclusion, and we correspondingly replace $\mathbf{p}$ with F everywhere in the notation. Fix $d+1$, F, $\mathbf{w}$ and $\mathbf{x}$ such that the moduli space $\mathcal{R}^{d+1,F,\mathbf{w}}(\mathbf{x})$ is one-dimensional.

Definition 3.8 An *admissible cut* of F consists of $d_+, d_- \geq 1$ such that $d_- + d_+ = d+1$, a number $i \in \{1, \dots, d_+\}$, and subsets $F_\pm \subset \{1, \dots, d_\pm\}$ satisfying $|F_-| + |F_+| = |F|$,

with the following property:

(59) F contains all $k \in F_+$ satisfying $k < i$, the numbers $k + d_- - 1$ for all $k \in F_+$ with $k > i$, and the numbers $k + i - 1$ for $k \in F_-$. If $i \notin F_+$, this completely describes F. Otherwise, F has one element other than the ones we have already given, which lies in the range $\{i, \dots, i + d_- - 1\}$.

Note that these subsets automatically come with canonical injective maps $\iota_\pm\colon F_\pm \to F$, which together cover the whole of F. One of them is simply $\iota_-(k) = k + i - 1$. The other one satisfies $\iota_+(k) = k$ for $k < i$, $\iota_+(k) = k + d_- - 1$ for $k > i$, and maps $i \in F_+$ (if that is the case) to the unique element in F specified above.

Given such a cut, set $w^{new} = w^i + \cdots + w^{i+d_- -1} + |F_-|$. Choose any $x^{new} \in \mathcal{X}_{w^{new}}$ such that $\deg(x^{new}) = \deg(x^i) + \cdots + \deg(x^{i+d_- -1}) + 2 - d_- - |F_-|$. Then, define weights $\mathbf{w}_\pm = \{w_\pm^0, \dots, w_\pm^{d_\pm}\}$ and collections of chords $\mathbf{x}_\pm = \{x_\pm^0, \dots, x_\pm^{d_\pm}\}$ as follows:

$$(w_+^k, x_+^k) = \begin{cases} (w^k, x^k) & k < i, \\ (w^{new}, x^{new}) & k = i, , \\ (w^{k+d_- -1}, x^{k+d_- -1}) & k > i \end{cases} \tag{60}$$

$$(w_-^k, x_-^k) = \begin{cases} (w^{new}, x^{new}) & k = 0, \\ (w^{k+i-1}, x^{k+i-1}) & k > 0. \end{cases}$$

In this terminology, $\mathcal{R}^{d_+ +1, F_+, \mathbf{w}_+}(\mathbf{x}_+) \times \mathcal{R}^{d_- +1, F_-, \mathbf{w}_-}(\mathbf{x}_-)$ appears as one of the boundary strata in the compactification of $\mathcal{R}^{d, F, \mathbf{w}}(\mathbf{x})$. In fact, this construction describes precisely those strata which have a trivial group of symmetries. The contributions from the other boundary strata cancel, for a reason parallel to that in Lemma 3.7 (see Section 9.3 for the precise argument). We therefore get a relation

$$\sum (-1)^{\S}\, m^{d_+, F_+, \mathbf{w}_+}(\mathbf{x}_+) \circ (\mathrm{id}^{\otimes d_+ - i} \otimes m^{d_-, F_-, \mathbf{w}_-}(\mathbf{x}_-) \otimes \mathrm{id}^{\otimes i-1}) = 0, \tag{61}$$

where the sum is over all admissible cuts $(d_\pm, i, F_\pm)$ and all possible x^{new}. The sign depends on the choice of orientation of the underlying spaces $\mathcal{R}^{d+1, F}$ (see Section 9.4 for a detailed discussion). For a specific such choice, it turns out to be

$$\begin{aligned} \S = {} & d_- i + i + 1 + d_+ |F_-| \\ & + (d_- + |F_-|)(\deg(x^{i+d_-}) + \cdots + \deg(x^d)) \\ & + |\{(k_+, k_-) \in F_+ \times F_- \,:\, \iota_+(k_+) < \iota_-(k_-)\}|. \end{aligned} \tag{62}$$

Examples 3.9

(i) Take $d = 1$ and $F = \{1\}$. In that case, the admissible cuts necessarily have $d_\pm = 1$. One of the two sets $F_\pm$ is $\{1\}$ while the other is empty. Omitting

weights for the sake of brevity, we find that (61) specializes to

$$\begin{aligned}(63)\qquad &\textstyle\sum m^{1,F}(x^0, x^{new}) \circ m^1(x^{new}, x^1) \\ &\textstyle -\sum m^1(x^0, x^{new}) \circ m^{1,F}(x^{new}, x^1) = 0,\end{aligned}$$

where the sum is over all x^{new} with the appropriate Maslov index. Geometrically, this kind of degeneration corresponds to an ordinary Floer trajectory bubbling off at the ends $\pm\infty \in \overline{Z}$.

(ii) Take $d = 2$ and $F = \{1, 2\}$, which is the situation from Examples 2.7(ii). Out of the five boundary edges in $\overline{\mathcal{R}}^{3,F}$, four give nontrivial contributions to (61), while the remaining one cancels for symmetry reasons. Again omitting weights, these contributions are

$$\begin{aligned}(64)\qquad &-\textstyle\sum m^{1,\{1\}}(x^0, x^{new}) \circ m^{2,\{1\}}(x^{new}, x^1, x^2) \\ &+\textstyle\sum m^{1,\{1\}}(x^0, x^{new}) \circ m^{2,\{2\}}(x^{new}, x^1, x^2) \\ &+\textstyle\sum m^{2,\{1\}}(x^0, x^1, x^{new}) \circ (m^{1,\{1\}}(x^{new}, x^2) \otimes \mathrm{id}) \\ &-\textstyle\sum m^{2,\{2\}}(x^0, x^{new}, x^2) \circ (\mathrm{id} \otimes m^{1,\{1\}}(x^{new}, x^1)).\end{aligned}$$

There are three more summands coming from bubbling off of Floer trajectories as in (i):

$$\begin{aligned}(65)\qquad &\textstyle\sum m^1(x^0, x^{new}) \circ m^{2,F}(x^{new}, x^1, x^2) \\ &\textstyle -\sum m^{2,F}(x^0, x^1, x^{new}) \circ (m^1(x^{new}, x^2) \otimes \mathrm{id}) \\ &\textstyle +(-1)^{\deg(x^2)+1} \sum m^{2,F}(x^0, x^{new}, x^2) \circ (\mathrm{id} \otimes m^1(x^{new}, x^1)).\end{aligned}$$

Remark 3.10 The $m^{d,F,\mathbf{w}}$ and their relations are all the analytic input underlying our geometric construction. In particular, moduli spaces with nontrivial symmetry are actually relevant only insofar as they appear in the compactification of other spaces. The crucial fact is that whenever $\overline{\mathcal{R}}^{d+1,F,\mathbf{w}}(\mathbf{x})$ one-dimensional, those of its boundary strata which do not correspond to cuts admit orientation-reversing free involutions, hence do not appear in (61). With that in mind, the various consistency conditions imposed on the additional geometric data living on $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$ could be somewhat relaxed when $Sym^{\mathbf{p}}$ is nontrivial (but nothing substantial is gained in doing that).

3.7 Wrapped Floer cohomology

For any w, one defines the Floer cochain group (actually a finite-dimensional graded $\mathbb{K}$–vector space) of L with respect to wH as a direct sum of $\mathbb{K}$–normalized orientation

spaces, placed in the degree indicated by the Maslov index:

$$CF^*(L; wH) = \bigoplus_{x \in \mathcal{X}_w} |o_x|_{\mathbb{K}}[-\deg(x)]. \tag{66}$$

This group carries the usual differential δ defined by counting Floer gradient trajectories. In our terminology, the $\mathbf{x} = \{x^0, x^1\}$ matrix coefficient of δ is

$$\delta_{(x^0,x^1)} = m^{1,\{w,w\}}(x^0, x^1)\colon |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}}.$$

There is one other nontrivial linear operation which appears naturally in our framework. Namely, take $d = 1$ with $F = \{1\}$, so $\mathbf{w} = \{w^0 = w + 1, w^1 = w\}$. Counting points in these moduli spaces yields maps $\kappa\colon CF^*(L; wH) \to CF^*(L; (w+1)H)$, whose matrix coefficients are

$$\kappa_{(x^0,x^1)} = m^{1,\{1\},\{w+1,w\}}(x^0, x^1)\colon |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}}. \tag{67}$$

Examples 3.9(i) translates into the statement that κ is a chain map (in more standard terminology, this is the continuation map associated to the homotopy from wH to $(w+1)H$; this is a monotone homotopy since $H \geq 0$ by assumption).

Let q be a formal variable of degree -1, with $q^2 = 0$. Take the space $CF^*(L; wH)[q] = CF^*(L; wH) \oplus q\, CF^*(L; wH)$, whose elements are formal sums $c = a + bq$. We consider this as a module over the one-dimensional exterior algebra $\mathbb{C}[\partial_q]$, where ∂_q has degree $+1$ and acts by formally differentiating in q–direction: $\partial_q(a + qb) = b$. The wrapped Floer cochain space is the infinite direct sum

$$CW^*(L; H) = \bigoplus_{w=1}^{\infty} CF^*(L; wH)[q] \tag{68}$$

We equip it with a differential which commutes with ∂_q:

$$\mu^1(a + qb) = (-1)^{\deg(a)}\delta(a) + (-1)^{\deg(b)}(q\delta(b) + \kappa(b) - b). \tag{69}$$

Visually (and with signs omitted), the resulting chain complex looks like this:

(70)

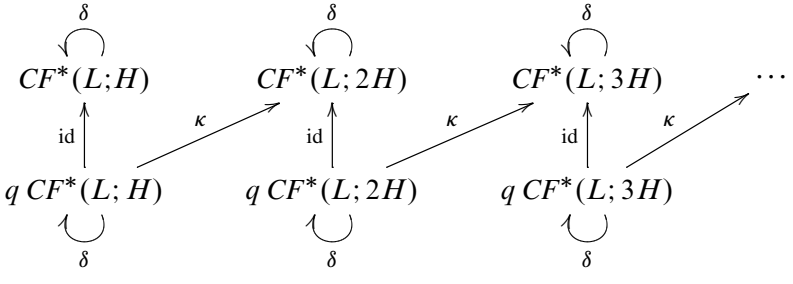

The cohomology of (69) is called *wrapped Floer cohomology* and denoted by $HW^*(L)$ (we omit the function H from the notation because the cohomology groups are independent of it, even though we won't give a proof of that fact here).

Lemma 3.11 *For every ν, the subcomplex C^ν consisting of those summands in (68) with $w \geq \nu$ is quasi-isomorphic to $CW^*(L;H)$.*

Proof The C^ν form a decreasing filtration of $CW^*(L;H)$. Each quotient $C^\nu/C^{\nu+1}$ is the mapping cone of the identity map on $CF^*(L;\nu H)$, hence acyclic. This implies that the inclusions $C^{\nu+1} \to C^\nu$ are quasi-isomorphisms. □

Lemma 3.12 *There is a canonical isomorphism*

$$HW^*(L) \cong \varinjlim_w HF^*(L;wH), \tag{71}$$

where the maps in the direct system are induced by κ.

Proof Consider the exhausting increasing filtration of $CW^*(L;H)$ by subcomplexes $C_w = CF^*(L;H)[q] \oplus \cdots \oplus CF^*(L;(w-1)H)[q] \oplus CF^*(L;wH)$, where the last summand has no q component. Obviously, this means that

$$HW^*(L;H) \cong \varinjlim_w H^*(C_w). \tag{72}$$

In turn, each C_w itself carries a finite decreasing filtration, $C_w^\nu = CF^*(L;\nu H)[q] \oplus \cdots \oplus CF^*(L;(w-1)H)[q] \oplus CF^*(L;wH)$. For $\nu \neq w$, the quotients $C_w^\nu/C_w^{\nu+1}$ are acyclic, for the same reason as in Lemma 3.11. Hence, the inclusion $C_w^w \to C_w$ is a quasi-isomorphism. Moreover, the diagram

$$\begin{array}{ccc} C_w^w & \xrightarrow{\kappa} & C_{w+1}^{w+1} \\ \downarrow & & \downarrow \\ C_w & \hookrightarrow & C_{w+1} \end{array} \tag{73}$$

commutes up to a chain homotopy, which is $a \mapsto (-1)^{\deg(a)} qa$. Therefore, if in (72) we replace the cohomology of C_w by that of its quasi-isomorphic subcomplex C_w^w, the connecting maps in the resulting direct system are the ones induced by κ. Now, C_w^w agrees with $CF^*(L;wH)$ up to a sign change in the differential, which does not affect cohomology. □

Remark 3.13 $CW^*(L;H)$ should be regarded as the homotopy direct limit of the $CF^*(L;wH)$. Indeed, this is just the chain level version of a well-known construction

in classical homotopy theory (see Hatcher [16, page 457], for instance). Note also that, while wrapped Floer cohomology carries an induced $\mathbb{C}[\partial_q]$–module structure, that structure is actually trivial. This is an easy consequence of the proof above: every class in wrapped Floer cohomology can be represented by a cocycle in C^w_w, which is obviously killed by ∂_q.

3.8 The A_∞–structure

We will now equip $CW^*(L;H)$ with the structure of an A_∞–algebra, generalizing the differential μ^1 introduced above. For each $(d, F, \mathbf{w})$ we will in fact define a multilinear map

$$CF^*(L; w^d H)[q] \otimes \cdots \otimes CF^*(L; w^1 H)[q] \xrightarrow{\mu^{d,F,\mathbf{w}}} CF^*(L; w^0 H)[q] \tag{74}$$

of degree $2-d$, which encodes all the $m^{d,F,\mathbf{w}}(\mathbf{x})$. We will first define the first summand of (74), namely the one which takes values in $CF^*(L; w^0 H)$. That summand vanishes unless each entry lies in $q^{i^k} CF^*(L; w^k H)$, where by definition $i^k = 1$ if $k \in F$, and $i^k = 0$ otherwise. The remaining coefficients are

$$\begin{aligned} &(-1)^* m^{d,F,\mathbf{w}}\colon q^{i^d} |o_{x^d}|_{\mathbb{K}} \otimes \cdots \otimes q^{i^1} |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}}, \\ &* = \textstyle\sum_j j \deg(x_j) + \sum_{\substack{j \in F \\ k>j}} (\deg(x_k) - 1). \end{aligned} \tag{75}$$

Having done that, there is a unique way to define the second summand, which is the one taking values in $q\, CF^*(L; w^0 H)$, such that the outcome commutes with the action of ∂_q when considering reduced degrees $\overline{\deg}(c) = \deg(c) - 1$. By this we mean that

$$\begin{aligned} &\partial_q \mu^{d,F,\mathbf{w}}(c^d, \dots, c^1) = \\ &\qquad \textstyle\sum_{k=1}^d (-1)^{\overline{\deg}(c^{k+1}) + \cdots + \overline{\deg}(c^d)} \mu^{d,F,\mathbf{w}}(c^d, \dots, \partial_q c^k, \dots, c^1). \end{aligned} \tag{76}$$

Finally, we extend $\mu^{d,F,\mathbf{w}}$ trivially (by zero) to a multilinear map defined on the whole of $CW^*(L;H)$, and let μ^d be the sum of those maps over all F and $\mathbf{w}$. For $d = 1$, this reproduces (69) except for the last term $qb \mapsto (-1)^{\deg(b)+1} b$, which one needs to add by hand.

Example 3.14 Suppose that we have fixed isomorphisms $|o_x|_{\mathbb{K}} \cong \mathbb{K}$ for each x. Then $CF^*(L; wH)$ becomes a graded vector space with distinguished basis indexed by $x \in \mathcal{X}_w$, and each $m^{d,F,\mathbf{w}}(\mathbf{x})$ is just an element of $\mathbb{K}$. In those terms (and omitting

weights for brevity, as in Examples 3.9) we have

(77)
$$\begin{aligned}
\mu^2(x^2, x^1) &= (-1)^{\deg(x^1)} \textstyle\sum_{x^0} m^2(x^0, x^1, x^2)\, x^0, \\
\mu^2(qx^2, x^1) &= (-1)^{\deg(x^1)} \textstyle\sum_{x^0} \big(m^2(x^0, x^1, x^2)\, qx^0 + m^{2,\{2\}}(x^0, x^1, x^2)\, x^0\big), \\
\mu^2(x^2, qx^1) &= (-1)^{\deg(x^1)+\deg(x^2)-1} \textstyle\sum_{x^0} \big(m^2(x^0, x^1, x^2)\, qx^0 \\
&\qquad + m^{2,\{1\}}(x^0, x^1, x^2)\, x^0\big), \\
\mu^2(qx^2, qx^1) &= (-1)^{\deg(x^1)+\deg(x^2)-1} \textstyle\sum_{x^0} \big(-m^{2,\{2\}}(x^0, x^1, x^2)\, qx^0 \\
&\qquad + m^{2,\{1\}}(x^0, x^1, x^2)\, qx^0 + m^{2,\{1,2\}}(x^0, x^1, x^2)\, x^0\big).
\end{aligned}$$

Proposition 3.15 *The maps $\mu^1, \mu^2, \cdots$ satisfy the A_∞–associativity equations. This means that for every $d \geq 1$ and $c^1, \dots, c^d \in CW^*(L; H)$, we have*

(78)
$$\begin{aligned}
\textstyle\sum (-1)^{\overline{\deg}(c^1)+\cdots+\overline{\deg}(c^{i-1})} \mu^{d_+}(c^d, \dots, c^{i+d_-}, \\
\mu^{d_-}(c^{i+d_--1}, \dots, c^i), c^{i-1}, \dots, c^1) = 0,
\end{aligned}$$

where the sum is over all $d_+ + d_- = d + 1$ and all i. Hence, $CW^*(L; H)$ equipped with these maps is an A_∞–algebra.

This follows from (61) and the definition of μ^d by a direct (if somewhat tedious, because of the signs) computation. The only terms in (78) which do not have direct counterparts in (61) come from last term in μ^1, but those cancel out because of (76).

4 Viterbo functoriality

This section sets up the A_∞–homomorphisms $CW^*(L; H) \longrightarrow CW^*(L^{\text{in}}; H^{\text{in}})$ associated to an embedding $(M^{\text{in}}, L^{\text{in}}) \subset (M, L)$. In the corresponding construction for symplectic cohomology [30], Viterbo's strategy was to shrink the interior domain M^{in} by a conformal symplectic factor $\rho \ll 1$. This makes all energies inside that domain small compared to those outside, and when carried out in a carefully designed way, allows one to identify the Floer complex $CF^*(L^{\text{in}}; wH^{\text{in}})$ with a quotient complex of $CF^*(L; wH)$. One takes the map on cohomology induced by this projection, which under the direct limit $w \to \infty$ yields the desired homomorphism between symplectic cohomology groups. In our case, the chain complexes for different values of w are tied together by the A_∞–structure. It seems that there is no single value of ρ which achieves the desired properties for all moduli spaces simultaneously, and as a consequence, we can't define $\mathcal{F}$ as a projection. Instead, we let ρ vary in $(0, 1]$, and look at how the

moduli spaces change depending on that parameter, which then yields the correction terms that need to be added to the naive projection map. The relevant formalism of parametrized moduli spaces, and its relation to the A_∞–homomorphism equation, is an adaptation of that in Fukaya–Oh–Ohta–Ono [13, Section 19] (see also Seidel [27, Section 10]). In the interest of simplicity and brevity, we give a full description of $\mathcal{F}^1$, but deal in a slightly more abbreviated way with the higher order components $\mathcal{F}^d$, $d \geq 2$, in particular omitting sign issues.

4.1 Geometric setup

Let M be a Liouville domain, and M^{in} a Liouville subdomain. This means that $M^{\mathrm{in}} \subset M \setminus \partial M$ is a compact submanifold with boundary (of equal dimension), such that Z points outwards along ∂M^{in}. In particular, $\theta^{\mathrm{in}} = \theta|M^{\mathrm{in}}$ turns M^{in} itself into a Liouville domain. We call ∂M^{in} the *dividing hypersurface*, and write $M^{\mathrm{out}} = (M \setminus M^{\mathrm{in}}) \cup \partial M^{\mathrm{in}}$ for the part outside that hypersurface. By integrating Z starting from the boundary, respectively from the dividing hypersurface, one gets collars, which we denote by

$$\begin{aligned} \kappa &\colon (0,1] \times \partial M \longrightarrow M, \\ \kappa^{\mathrm{in}} &\colon (0, 1+\epsilon) \times \partial M^{\mathrm{in}} \longrightarrow M \end{aligned} \tag{79}$$

(for some $\epsilon > 0$). In the case of κ^{in}, points (r, y) with $r \leq 1$ get mapped to M^{in}, and those with $r \geq 1$ to M^{out}. Let $L \subset M$ be a Lagrangian submanifold, intersecting ∂M as well as ∂M^{in} transversally, and write $L^{\mathrm{in}} = L \cap M^{\mathrm{in}}$, $L^{\mathrm{out}} = L \cap M^{\mathrm{out}}$. We make the following assumption:

(80) $\theta|L^{\mathrm{out}}$ vanishes on a neighbourhood of the boundary $\partial L^{\mathrm{out}} = \partial L \cup \partial L^{\mathrm{in}}$. Moreover, one can write $\theta|L = dh$, where $h|\partial L \cup \partial L^{\mathrm{in}}$ is zero.

This requirement merits some discussion. The first part of (80) says that Z is parallel to L^{out} near its boundaries. To put it more explicitly in terms of (79), there is some small $\lambda > 0$ such that

$$\begin{aligned} \kappa^{-1}(L) \cap ([1-\lambda, 1] \times \partial M) &= [1-\lambda, 1] \times \partial L, \\ (\kappa^{\mathrm{in}})^{-1}(L) \cap ([1, 1+\lambda] \times \partial M^{\mathrm{in}}) &= [1, 1+\lambda] \times \partial L^{\mathrm{in}}. \end{aligned} \tag{81}$$

This is similar to (41), hence marginally stronger than its counterpart in (35), but that is just for technical convenience. A more interesting aspect is the condition on h, which amounts to a strengthened form of exactness. In more algebro-topological terms, the statement is that the relative class $[\theta|L] \in H^1(L, \partial L^{\mathrm{in}} \cup \partial L; \mathbb{R})$ is zero.

Lemma 4.1 *Let $L_0 \subset M$ be a Lagrangian submanifold satisfying the following weaker property:*

(82) $\theta|L_0 = dh_0|L_0$, *where $h_0 \in C^\infty(M, \mathbb{R})$ is a function which is locally constant on $\partial M \cup \partial M^{\mathrm{in}}$.*

Then there is a Hamiltonian isotopy (stationary on $\partial M \cup \partial M^{\mathrm{in}}$) which deforms L_0 into a submanifold L_1 such that $[\theta|L_1] \in H^1(L_1, \partial L_1^{\mathrm{in}} \cup \partial L_1; \mathbb{R})$ is zero.

Proof We may assume that $\theta|L_0$ vanishes in a neighbourhood of $\partial L_0^{\mathrm{in}} \cup \partial L_0$ (since that can be achieved by the same argument as in Lemma 3.1). In that case, we may further assume that h_0 is locally constant on a neighbourhood of $\partial M^{\mathrm{in}} \cup \partial M \subset M$. Let Y be the Hamiltonian vector field of h_0. Then $L_Y\theta$ is a closed one-form vanishing near $\partial M^{\mathrm{in}} \cup \partial M$. The relative cohomology class of this form is

$$[L_Y\theta] = [i_Y\omega] = [-dh_0] \in H^1(M, \partial M \cup \partial M^{\mathrm{in}}; \mathbb{R}). \tag{83}$$

Integrating this out, we find that if ξ^t is the flow of Y, then $[(\xi^1)^*\theta] = [\theta - dh_0]$ becomes zero when restricted to $H^1(L_0, \partial L_0 \cup \partial L_0^{\mathrm{in}}; \mathbb{R})$, which of course means that $[\theta]$ itself has the same property with respect to $L_1 = \xi^1(L_0)$. This provides the desired isotopy. □

Note that unlike Lemma 3.1, the isotopy constructed here is not necessarily local near $\partial M \cup \partial M^{\mathrm{in}}$. On the other hand, the following example shows that exactness of L, which would be vanishing of $[\theta|L] \in H^1(L; \mathbb{R})$, is not sufficient to construct meaningful Viterbo functoriality maps.

Example 4.2 Take $M = D^*S^1$ to be the annulus, which is the disc cotangent bundle of S^1 with fibres of some radius r. Let L be the union of two different cotangent fibres $L_i = D^*_{q_i}S^1$ $(i = 1, 2)$. Take the zero-section and perturb it locally near q_1 in a Hamiltonian way, such that it intersects L_1 transversally in three points. A neighbourhood of this perturbed zero-section looks like a smaller cotangent disc bundle $M^{\mathrm{in}} = D^*S^1$ with $r^{\mathrm{in}} \ll r$. One can arrange that the intersections $L_i^{\mathrm{in}} = L_i \cap M^{\mathrm{in}}$ are unions of cotangent fibres (three and one, respectively). Moreover, the canonical one-forms on our two cotangent disc bundles will satisfy $\theta|M^{\mathrm{in}} = \theta^{\mathrm{in}} - dh$ for some function h. Take h and (shrinking r^{in} first, if necessary) extend it to a function on the whole of M, which vanishes near ∂M. Then, if we replace the original θ by $\theta + dh$, we will actually have $\theta^{\mathrm{in}} = \theta|M^{\mathrm{in}}$ as in the general framework set out above. The Lagrangian L satisfies (81), since $\theta|L = dh|L$ and $\theta^{\mathrm{in}}|L^{\mathrm{in}} = 0$. On the other hand, let Δ is the shaded region in Figure 6, and write $\partial_l\Delta$ and $\partial_r\Delta$ for its left and right boundary arcs, and x_u, x_d for its upper and lower corners. Then

$$0 < \textstyle\int_\Delta \omega = \int_{\partial_r\Delta} \theta + \int_{\partial_l\Delta} \theta = \int_{\partial_r\Delta} \theta + h(x_d) - h(x_u), \tag{84}$$

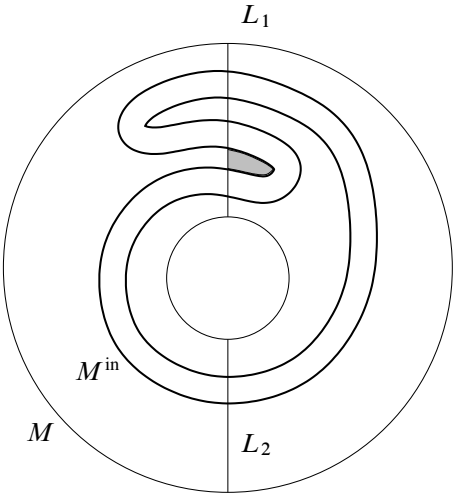

Figure 6

and by construction $\theta|\partial M^{\text{in}} > 0$, which in view of the orientations occurring means that $\int_{\partial_r \Delta} \theta < 0$. Hence, the values of h at the two corners are different, so that (80) as well as the weaker version (82) are violated.

In this situation, the wrapped Floer cohomologies of L and L^{in} are well-defined. However, as we will now show, no Viterbo-type restriction homomorphism (with reasonable properties) can exist. Write $R = \mathbb{K}[t, t^{-1}]$, which is the wrapped Floer cohomology of a single cotangent fibre in D^*S^1. The wrapped Floer cohomology of L is a matrix algebra $R^{2\times 2}$. More geometrically, we write this as

$$HW^*(L) = \begin{pmatrix} R & R \\ R & R \end{pmatrix}, \tag{85}$$

where the (j,i)th piece in the splitting is represented by X–chords going from L_i to L_j. Similarly, the wrapped Floer cohomology of L^{in} is a matrix algebra $R^{4\times 4}$ (actually, it is now nontrivially graded, but we will not discuss that since it's irrelevant for our purpose). In terms of the pieces L_i^{in} this has a block matrix decomposition

$$HW^*(L^{\text{in}}) = \begin{pmatrix} R^{3\times 3} & R^3 \\ R^3 & R \end{pmatrix}. \tag{86}$$

Suppose that we had a restriction map $HW^*(L) \to HW^*(L^{\text{in}})$ with the following properties: it is a map of rings preserving the unit elements of the ring structure and is compatible with the splittings (85), (86). In particular, multiplication with the images of $x_{12} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ and $x_{21} = \left(\begin{smallmatrix} 0 & 0 \\ 1 & 0 \end{smallmatrix}\right)$ on the left and right would yield a ring isomorphism

between two subalgeras of $HW^*(L^{\mathrm{in}})$ which are isomorphic to $R^{3\times 3}$ and R, which is of course impossible (this example finds a more natural interpretation in the framework of wrapped Fukaya categories, where L is the direct sum of two isomorphic objects, while L^{in} is the sum of two non-isomorphic ones).

4.2 Hamiltonian functions

We will consider functions $H \in C^\infty(M, \mathbb{R})$ of the following type:

(87) Both H and its restriction $H^{\mathrm{in}} = H|M^{\mathrm{in}}$ satisfy (34); equivalently, H is everywhere positive, and both $H(\kappa(r, y)) - r$ and $H(\kappa^{\mathrm{in}}(r, y)) - r$ vanish to infinite order along $r = 1$. In addition to that, we assume that the restriction of $H - dH(Z)$ to $M^{\mathrm{out}} \setminus \partial M^{\mathrm{out}}$ is positive.

Example 4.3 Suppose that we make the following ansatz:

$$\text{(88)} \quad \begin{cases} H = c^{\mathrm{in}} \text{ on } M^{\mathrm{in}} \setminus \kappa^{\mathrm{in}}([1-\delta, 1] \times \partial M^{\mathrm{in}}), \\ H(\kappa^{\mathrm{in}}(r, y)) = \chi^{\mathrm{in}}(r) \text{ for } (r, y) \in [1-\delta, 1+\delta] \times \partial M^{\mathrm{in}}, \\ H = c^{\mathrm{out}} \text{ on } M^{\mathrm{out}} \setminus (\kappa^{\mathrm{in}}([1, 1+\delta] \times \partial M^{\mathrm{in}}) \cup \kappa([1-\delta, 1] \times \partial M)), \\ H(\kappa(r, y)) = \chi^{\mathrm{out}}(r) \text{ for } (r, y) \in [1-\delta, 1] \times \partial M. \end{cases}$$

Here $\delta > 0$ is small; $c^{\mathrm{in}}, c^{\mathrm{out}} > 0$ are constants; $\chi^{\mathrm{in}} \in C^\infty(\mathbb{R}, (0, \infty))$ is a function satisfying $\chi^{\mathrm{in}}(r) = c^{\mathrm{in}}$ for $r \le 1-\delta$, $\chi^{\mathrm{in}}(r) = c^{\mathrm{out}}$ for $r \ge 1+\delta$, and such that $\chi^{\mathrm{in}}(r) - r$ vanishes to infinite order at $r = 1$; similarly, χ^{out} should satisfy $\chi^{\mathrm{out}}(r) = c^{\mathrm{out}}$ for $r \le 1-\delta$, and $\chi^{\mathrm{out}}(r)$ agrees with r to infinite order when $r = 1$ (see Figure 7 for a schematic picture). Then

$$\text{(89)} \quad H - dH(Z) = \begin{cases} \chi^{\mathrm{in}} - r \cdot d\chi^{\mathrm{in}}/dr \text{ at points } \kappa^{\mathrm{in}}(r, y),\ r \in (1, 1+\delta], \\ c^{\mathrm{out}} \text{ on } M^{\mathrm{out}} \setminus (\kappa^{\mathrm{in}}([1, 1+\delta] \times \partial M^{\mathrm{in}}) \cup \kappa([1-\delta, 1] \times \partial M)), \\ \chi^{\mathrm{out}} - r \cdot d\chi^{\mathrm{out}}/dr \text{ at points } \kappa(r, y),\ r \in [1-\delta, 1). \end{cases}$$

It is easy to choose our functions and constants so that this is > 0 (note the elementary calculus fact that $\chi - r \cdot d\chi/dr$ is where the tangent line to the graph of χ at r hits the y–axis).

From this point onwards, we will assume that (37) holds for both $\partial L \subset \partial M$ and $\partial L^{\mathrm{in}} \subset \partial M^{\mathrm{in}}$. Since by assumption, the Hamiltonian vector field X of H restricts to the Reeb fields $R \in C^\infty(T(\partial M))$ and $R^{\mathrm{in}} \in C^\infty(T(\partial M^{\mathrm{in}}))$, it follows that these two hypersurfaces contain no integer X–chords. We can then can split the set $\mathcal{X}_w$ of X–chords of length w into two parts

$$\text{(90)} \quad \mathcal{X}_w = \mathcal{X}_w^{\mathrm{in}} \cup \mathcal{X}_w^{\mathrm{out}}$$

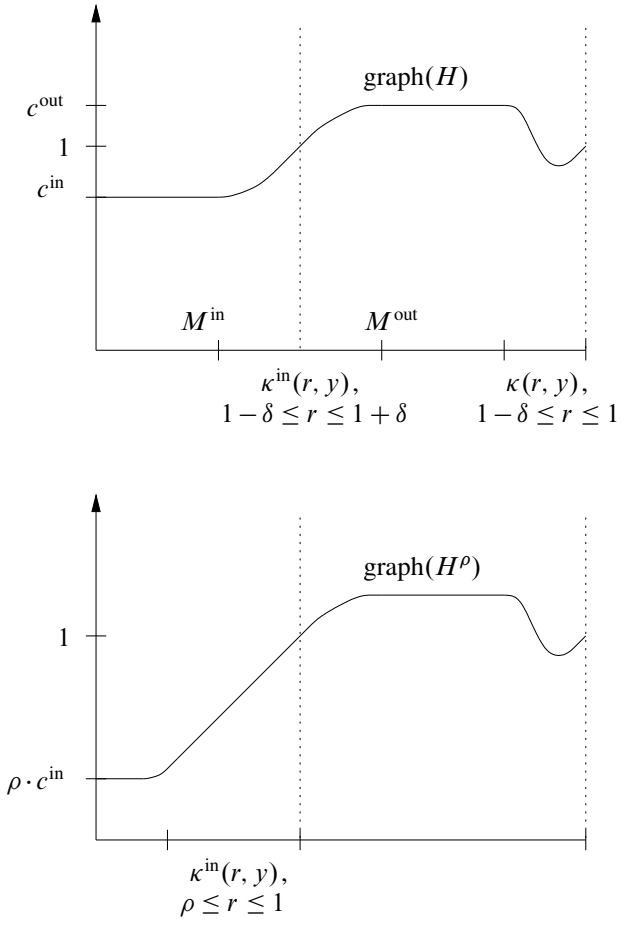

Figure 7

containing the integer X–chords lying in $M^{\text{in}} \setminus \partial M^{\text{in}}$ and $M^{\text{out}} \setminus \partial M^{\text{out}}$, respectively.

4.3 Rescaling the inner part

Somewhat unconventionally, we will write ψ^ρ for the flow of Z at time $\log(\rho)$. By definition, this shrinks all the symplectic data by a factor of ρ. Our strategy will be to use this flow, with $\rho \le 1$, to shrink the inner Liouville domain to

$$M^{\text{in},\rho} = \psi^\rho(M^{\text{in}}). \tag{91}$$

By definition, $\psi^\rho(\kappa^{\text{in}}(r, y)) = \kappa^{\text{in}}(\rho r, y)$. In particular, the boundary of (91) is $\partial M^{\text{in},\rho} = \kappa^{\text{in}}(\{\rho\} \times \partial M^{\text{in}})$. Similarly, we define Lagrangian submanifolds $L^{\text{in},\rho} =$

$\psi^\rho(L^{\mathrm{in}})$ and complete them to submanifolds $L^\rho \subset M$ as follows:

$$\begin{cases} L^\rho \cap M^{\mathrm{in},\rho} = L^{\mathrm{in},\rho}, \\ L^\rho \cap \kappa^{\mathrm{in}}((\rho,1)\times \partial M^{\mathrm{in}}) = (\rho,1)\times \partial L^{\mathrm{in}}, \\ L^\rho \cap M^{\mathrm{out}} = L \cap M^{\mathrm{out}}; \end{cases} \tag{92}$$

There is a natural choice of function h^ρ with $dh^\rho = \theta|L^\rho$, namely

$$h^\rho = \begin{cases} \rho\cdot(h\circ\psi^{1/\rho}) \text{ on } M^{\mathrm{in},\rho}, \\ 0 \text{ at points } \kappa^{\mathrm{in}}(r,y) \text{ with } r\in(\rho,1), \\ h \text{ on } M^{\mathrm{out}}, \end{cases} \tag{93}$$

which makes sense because of (80). Finally, given any Hamiltonian function H as in (87) and any $\rho\in(0,1]$, we define the *rescaled function* H^ρ as follows:

$$H^\rho = \begin{cases} \rho\cdot(H\circ\psi^{1/\rho}) \text{ on } M^{\mathrm{in},\rho}, \\ r \text{ at points } \kappa^{\mathrm{in}}(r,y) \text{ with } r\in(\rho,1), \\ H \text{ on } M^{\mathrm{out}}. \end{cases} \tag{94}$$

This is smooth since, along $\partial M^{\mathrm{in},\rho}$, the function $\rho\cdot(H\circ\psi^{1/\rho})$ agrees to infinite order with $\rho\cdot(r\circ\psi^{1/\rho}) = \rho\cdot\rho^{-1}\cdot r = r$ (Figure 7 shows what happens if we apply this rescaling process to one of the functions from Example 4.3). Correspondingly, the associated Hamiltonian vector field is

$$X^\rho = \begin{cases} \psi^\rho_*(X) \text{ on } M^{\mathrm{in},\rho}, \\ \kappa^{\mathrm{in}}_*(0,R^{\mathrm{in}}) \text{ at points } \kappa^{\mathrm{in}}(r,y) \text{ with } r\in(\rho,1), \\ X \text{ on } M^{\mathrm{out}}. \end{cases} \tag{95}$$

There is a natural bijection between X–chords with boundary in L and X^ρ–chords with boundary in L^ρ, which we denote by

$$\begin{aligned} \mathcal{X}_w &\xrightarrow{\cong} \mathcal{X}^\rho_w, \\ x &\longmapsto x^\rho. \end{aligned} \tag{96}$$

For chords lying in M^{out}, the map is a simple equality $x^\rho = x$, while for the ones in M^{in} it is $x^\rho = \psi^\rho(x)$. As usual, we end our discussion by fixing once and for all a Hamiltonian H as in (87), and which in addition should satisfy (38) and (39). These conditions are still generic for $\dim(M)\geq 4$, even within the present smaller class of functions (one can see that by inspecting the argument from Section 8.1).

We also need to relate our choices of almost complex structures on M and M^{in}, but here the connection will be made in a somewhat looser sense. Suppose we are given $I \in \mathcal{J}(M)$ and $I^{\mathrm{in}} \in \mathcal{J}(M^{\mathrm{in}})$. Then, an *interpolating family* of almost complex structures is a family $I^\rho \in \mathcal{J}(M)$ parametrized by $\rho \in (0,1]$, such that $I^1 = I$ and:

(97) There is a $\delta > 0$ such that for $\rho \leq \delta$, the following properties hold: $I^\rho | M^{\mathrm{in},\rho} = \psi_*^\rho I^{\mathrm{in}}$, and $(\kappa^{\mathrm{in}})^*(I^\rho)$ is of contact type on some neighbourhood of $\{\rho\} \times \partial M^{\mathrm{in}} \subset [\rho, 1] \times \partial M^{\mathrm{in}}$.

Note that the first part already implies that I^ρ is of contact type along the boundary of $M^{\mathrm{in},\rho}$; the second part then extends that control slightly to a piece of the cone on the outside of that boundary. The same idea applies to infinitesimal deformations. Namely, assume that we have I, I^{in}, I^ρ as well as infinitesimal deformations $K \in T\mathcal{J}(M)_I$, $K^{\mathrm{in}} \in T\mathcal{J}(M^{\mathrm{in}})_{I^{\mathrm{in}}}$. Then, an interpolating family consists of infinitesimal deformations K^ρ of I^ρ, satisfying the analogue of (97). Obviously, this leads to the same relationship between the perturbed almost complex structures which we saw in Equation (44).

Remark 4.4 Usually, we will be dealing with families of almost complex structures parametrized by points on a Riemann surface, and possibly additional moduli. In that case, we always assume that the constant δ and the neighbourhood of $\{\rho\} \times \partial M^{\mathrm{in}}$ in (97) should be the same throughout the family.

4.4 Parametrized moduli spaces

Consider the moduli spaces (52) for the given $L \subset M$. In addition to various data which concerns only the domains (strip-like ends, popsicle sticks, one-forms), the construction of those spaces involves choices of almost complex structures on M. Namely, for each w we have a family $\mathbf{J}_w$ parametrized by $[0,1]$, and for every stable weighted popsicle we have another family $\mathbf{J}_{S,\boldsymbol{\phi},\mathbf{w}}$ parametrized by points of S. Choose analogous families $\mathbf{J}^{\mathrm{in}}_w$ and $\mathbf{J}^{\mathrm{in}}_{S,\boldsymbol{\phi},\mathbf{w}}$ on M^{in}, and introduce the corresponding moduli spaces, which will be denoted by

$$\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\mathrm{in}} \quad \text{for } \mathbf{x} = (x^0, \dots, x^d),\ x^k \in \mathcal{X}^{\mathrm{in}}_{w^k}. \tag{98}$$

To relate the two constructions, we will then need interpolating families $\mathbf{J}^\rho_w$ and $\mathbf{J}^\rho_{S,\boldsymbol{\phi},\mathbf{w}}$, in the sense defined above. To ensure that the interpolating families are asymptotically consistent in a suitable sense, one needs to copy in more detail the construction from Section 3.2, by which we mean the following. Recall that $\mathbf{J}_{S,\boldsymbol{\phi},\mathbf{w}}$ was obtained by taking an initial choice $\mathbf{I}_{S,\boldsymbol{\phi},\mathbf{w}}$, which satisfied consistency in a strict sense, and then perturbing it by exponentiating an infinitesimal deformation $\mathbf{K}_{S,\boldsymbol{\phi},\mathbf{w}}$. The same applies to the corresponding data on M^{in}. Similarly, one should first choose interpolating

families $\mathbf{I}^{\rho}_{S,\boldsymbol{\phi},\mathbf{w}}$ for the initial choices of data, satisfying strict consistency, and then perturb them using a suitable $\mathbf{K}^{\rho}_{S,\boldsymbol{\phi},\mathbf{w}}$ which is asymptotically consistent. We omit the details, which are straightforward.

For any $\rho \in (0,1]$, define $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho}$ to be the moduli space of popsicle maps analogous to (52), but where L is replaced by L^{ρ}, X by X^{ρ}, the limits of the map u are the $x^{k,\rho}$, and the almost complex structures belong to the interpolating families we have just constructed. For $\rho = 1$ this just reproduces the original moduli space $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$. On the other hand, if ρ is sufficiently small and the x^k all lie in $\mathcal{X}^{\mathrm{in}}_{w^k}$, it contains a copy of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\mathrm{in}}$, obtained by rescaling the map, $u = \psi^{\rho} \circ u^{\mathrm{in}}$. In the converse direction we have the following statement, proved in Lemma 7.4:

(99) There is a positive integer ν, such that for any fixed $(d, \mathbf{p}, \mathbf{w})$ with $w^0, \dots, w^d \geq \nu$, the following holds. There is a δ such that all points in $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho}$ with $x^0 \in \mathcal{X}^{\mathrm{in}}_{w^0}$ and $\rho \leq \delta$ are of the form $(S, \boldsymbol{\phi}, u = \psi^{\rho} \circ u^{\mathrm{in}})$ for some $(S, \boldsymbol{\phi}, u^{\mathrm{in}}) \in \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\mathrm{in}}$. In particular, this means that $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho} = \varnothing$ unless all the other x^k, $k > 0$, lie in $\mathcal{X}^{\mathrm{in}}_{w^k}$ as well.

Convention 3.4 applies here, meaning that while the statement is formulated in terms which make sense for stable popsicles $(S, \boldsymbol{\phi})$, the analogous result holds for $d + |F| = 1$ as well. We will use this property heavily from now on, hence

(100) we require throughout the subsequent discussion that all weights should be $\geq \nu$.

The first step is to suitably package our moduli spaces for varying values of ρ. The most straightforward way is to introduce *parametrized moduli spaces*

$$\mathcal{P}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}) = \big\{(\rho, S, \boldsymbol{\phi}, u) \,:\, \rho \in (0,1],\ (S, \boldsymbol{\phi}, u) \in \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho}\big\}. \tag{101}$$

Adding the parameter ρ modifies some analytic aspects of the theory, but only in a relatively straightforward way. For instance, if one considers the standard compactification of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho}$ for any ρ, the disjoint union of these compactifications provides a partial compactification of the parametrized moduli space; more precisely, this is a properification with respect to the projection map

$$(\rho, S, \boldsymbol{\phi}, u) \longmapsto \rho. \tag{102}$$

Next, suppose that the original almost complex structures on M, as well as their counterparts on M^{in}, have been chosen generically, so that both Theorems 3.5 and 3.6 apply. In that case, the parametrized moduli spaces are regular near their boundary points, which by definition are the points $(\rho, S, \boldsymbol{\phi}, u)$ with $\rho = 1$. Moreover, at each such point, the derivative of the projection (102) is onto, so $\rho = 1$ is a regular value of that map.

Theorem 4.5 *For a generic choice of interpolating families, satisfying the required asymptotic consistency conditions, the following properties will hold. First, all moduli spaces $\mathcal{P}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ are regular, hence smooth manifolds with boundary of the expected dimension. Secondly, for every fixed value of ρ, there is at most one point $(\rho, S, \boldsymbol{\phi}, u)$ which belongs to a zero-dimensional parametrized moduli space. Finally, if ρ is such that there is one such point, then all moduli spaces $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})^{\rho}$ for various choices of $(\mathbf{p}, \mathbf{w}, \mathbf{x})$ are regular everywhere except at that point.*

This is the counterpart of Theorem 3.5 for parametrized moduli spaces, and is proved by essentially the same methods (see Section 8.4). There is also a stratified version, which means an analogue of Theorem 3.6 applying to the isotropy strata of the $Sym^{\mathbf{p}}$–action on the parametrized moduli spaces. We leave its formulation to the reader. From now on, assume that the interpolating families have been chosen in such a way that the conclusions of Theorem 4.5, and of its stratified version, hold. Points of the zero-dimensional parametrized moduli spaces will be called *rigid*; they represent behaviour which is non-generic for a fixed value of ρ, but which is unavoidable in the parametrized context.

4.5 Cascade maps (naive version)

Fix $(d, \mathbf{p})$ as well as $(\mathbf{w}, \mathbf{x})$, where it is assumed throughout that

$$x^0 \in \mathcal{X}^{\mathrm{in}}_{w^0}. \tag{103}$$

Let $(T, \mathbf{F})$, $\{\mathbf{w}_v\}$, $\{\mathbf{p}_v\}$ and $\{\mathbf{x}_v\}$ be as in the definition of the Gromov compactification in Section 3.5. A stable *cascade map* modelled on $(T, \mathbf{F})$ is a collection of points

$$(\rho_v, S_v, \boldsymbol{\phi}_v, u_v) \in \mathcal{P}^{|v|,\mathbf{p}_v,\mathbf{w}_v}(\mathbf{x}_v) \tag{104}$$

satisfying the causal ordering condition (29). We denote the moduli space of cascade maps by $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$. Each component of this space, corresponding to a particular choice of $(T, \mathbf{F})$ and $\mathbf{x}_v$, is an open subset of a product of parametrized moduli spaces $\mathcal{P}^{|v|,\mathbf{p}_v,\mathbf{w}_v}(\mathbf{x}_v)$, hence regular by Theorem 4.5. There is also a natural partial compactification $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$, whose structure is analogous to that for moduli spaces of cascades considered in Section 2.8. Namely, in addition to $(T, \mathbf{F})$, assume that we have subsets $\overline{E} \subset E$ of the set of internal edges of T, as in (30). Points of the resulting stratum in the partial compactification are collections (104), with the causal ordering condition replaced by equality $\rho_{v_+} = \rho_{v_-}$ for those edges contained in E.

Theorem 4.6 *The map $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}) \to (0, 1]$ which associates the value ρ_{v_1} to any cascade $\{(\rho_v, S_v, \boldsymbol{\phi}_v, u_v)\}$, where v_1 is the vertex closest to the root, is proper.*

This is the counterpart of (32). The strategy of proof is as follows. Fixing $\mathbf{x}$ gives an a priori bound on the energy available for all components of the cascade. The next step is to bound the complexity of $(T, \mathbf{F})$. Because the number of semi-infinite edges of T is $d+1$, and $|F| = \sum_v |F_v|$, this comes down to bounding the length of chains of Floer trajectories. However, a standard monotonicity argument, based on the fact that integer X^ρ–chords never merge as ρ varies, shows that each such trajectory consumes at least some fixed amount ϵ of energy, as long as the parameter values stay in some interval $[\delta, 1]$. This achieves the desired bound, and parametrized Gromov compactness theory takes care of bubbling processes.

4.6 The zero-dimensional case

In the case where the moduli space of cascade maps is zero-dimensional, its partial compactification is regular, which means that all strata of positive codimension are actually empty. To see that, consider a point lying in such a stratum. Since the codimension is $|E|$, there is at least one edge belonging to that subset. Take the two vertices $v_\pm$ connected by that edge, with the usual convention according to which v_+ is supposed to lie closer to the root. For simplicity, denote the associated components by $(\rho_\pm = \rho, S_\pm, \boldsymbol{\phi}_\pm, u_\pm) \in \mathcal{P}^{d_\pm+1, \mathbf{p}_\pm, \mathbf{w}_\pm}(\mathbf{x}_\pm)$. For index reasons, each component must be a rigid point in its parametrized moduli space. On the other hand, Theorem 4.5 tells us that at most one rigid point exists for any fixed value of the parameter ρ. Hence, the two components must necessarily agree. Suppose first that these components are both Floer trajectories. We then know that their limits satisfy $x_-^0 = x_+^1$ (this is part of our general condition on the $\mathbf{x}_v$) as well as $x_-^1 = x_+^1$ (because the two components agree). But this is impossible; a Floer trajectory has non-negative energy and hence cannot have the same limits at both ends. In the remaining case where $d_\pm + |F_\pm| \geq 2$, the argument is parallel but based on weights instead of the limits. Suppose that v_- is obtained by leaving v_+ in direction of its kth flag. Then $w_-^0 = w_+^k$ (by definition of the $\mathbf{w}_v$), but also $w_-^k = w_+^k$ (because the two components are the same). But by (25) we necessarily have $w_-^0 > w_-^k$ for any $k > 0$, which is again contradictory.

Corollary 4.7 *The zero-dimensional moduli spaces $\mathcal{Q}^{d+1, \mathbf{p}, \mathbf{w}}(\mathbf{x})$ are finite sets.*

To see why that is true, take a point of a zero-dimensional moduli space, and look at the component associated to the vertex $v = v_1$ nearest to the root. By (103), $x_v^0 = x^0 \in \mathcal{X}^{\mathrm{in}}_{w^0}$. If ρ_v was sufficiently small, (99) would ensure that this component comes from a suitable $\mathcal{R}^{\mathrm{in}}$ moduli space, hence is regular in the unparametrized sense, which is in contradiction to its rigidity. To be a little more precise, one has to consider all the possibilities for $(|v|, \mathbf{p}_v, \mathbf{w}_v, \mathbf{x}_v)$, but since there are only finitely many, the resulting

lower bound for ρ_v is uniform. In view of the previous discussion, $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ agrees with its partial compactification, so applying Theorem 4.6 yields the desired result.

Addendum 4.8 At this point we tweak the existing definitions in a small, but useful, way. Consider a zero-dimensional space $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ where $d = 1$, $F = \varnothing$, $\mathbf{w} = \{w^0 = w, w^1 = w\}$ and $x^0 = x^1 = x \in \mathcal{X}^{\mathrm{in}}_w$. We formally add a point, denoted by $\varnothing_x$, to this space. One can consider that as the degenerate case of an empty cascade (whose underlying tree would be an infinite line, with no vertices). Of course, doing that does not affect Corollary 4.7.

4.7 The one-dimensional case

Now consider the partial compactifications of one-dimensional moduli spaces of cascade maps. Points in such a space have the property that one component belongs to a one-dimensional parametrized moduli space (we will refer to it as the *moveable component*), while the remaining ones are all rigid. As before, one can show that the vertices belonging to an edge in E cannot both carry rigid components. In particular, if $|E| = 1$ then equality of parameter values can only happen between the moveable component and a single adjacent rigid component. In that case, another look at Theorem 4.5 shows that the moveable component is a regular point in its unparametrized moduli space, and that implies that the codimension one strata of $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ are regular. We now turn to the remaining case $|E| \geq 2$, where the associated strata should be empty. A slight generalization of the previous argument shows that no two rigid components can be connected by a sequence of edges in E in which each step goes further away from the root. The only remaining possibility is the following one: E is a set of edges starting at a fixed vertex, which is the one where the moveable component is located, and going in direction of the leaves towards other vertices; and moreover, all of those other vertices have the same (rigid) component associated to them. In this situation, which is unfortunately not ruled out by Theorem 4.5, regularity fails. We will have to return to this issue later on.

Corollary 4.9 *Suppose that $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ is one-dimensional. Then its partial compactification $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ contains finitely many additional points.*

This is essentially the same argument as in Corollary 4.7. The only case requiring a little extra thought is when the vertex closest to the root is the one carrying the moveable component. But then, the rigid components tied to it by edges in E must have the same parameter value, and if that value is small, (99) applies, leading to the same contradiction as before.

With that in mind, we can describe the ends of the partially compactified one-dimensional moduli spaces. Take a partition of d into $d_{-,1}+\cdots+d_{-,l}$, and set $d_+ = l$. This should come with a decomposition of F into subsets F_+ and $F_{-,j}$ $(j = 1,\dots,l)$, and associated flavours $\mathbf{p}_+$, $\mathbf{p}_{-,j}$, such that the following holds:

(105) For each $f \in F$ such that $d_{-,1}+\cdots+d_{-,j-1} < p_f \le d_{-,1}+\cdots+d_{-,j}$, either $f \in F_+$ and $p_{+,f} = j$, or $f \in F_{-,j}$ and $p_{-,j,f} = p_f - d_{-,1} - \cdots - d_{-,j-1}$.

Choose weights $\mathbf{w}_+$, $\mathbf{w}_{-,j}$ and collections of chords $\mathbf{x}_+$, $\mathbf{x}_-$ as follows:

$$\text{(106)}\qquad \begin{cases} x^k_{-,j} = x^{d_{-,1}+\cdots+d_{-,j-1}+k},\ w^k_{-,j} = w^{d_{-,1}+\cdots+d_{-,j-1}+k} & \text{for } k>0;\\ x^0_{-,j} = x^j_+,\ w^0_{-,j} = w^j_+ = w^1_{-,j}+\cdots+w^{d_{-,j}}_{-,j} + |F_{-,j}|; \\ x^0_+ = x^0,\ w^0_+ = w^0. \end{cases}$$

Corollary 4.10 *Suppose that $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ is one-dimensional. Then, its partial compactification has finitely many ends, and these are modelled on the disjoint union of products*

$$\text{(107)}\qquad (0,\delta] \times \mathcal{R}^{d_+ +1,\mathbf{p}_+,\mathbf{w}_+}(\mathbf{x}_+)^{\mathrm{in}} \times \prod_j \mathcal{Q}^{d_{-,j}+1,\mathbf{p}_{-,j},\mathbf{w}_{-,j}}(\mathbf{x}_{-,j}).$$

The combinatorial data is as described in (105), (106), with the degrees of the x^+_j chosen so that all the moduli space factors in (107) are zero-dimensional.

As usual, the situation is clearer in graphical terms (see Figure 8). Each end (107) consists of cascades where the vertex nearest to the root carries the moveable component. For that moveable component, the parameter can become arbitrarily small, and then (99) applies, identifying it with a point in $\mathcal{R}^{\mathrm{in}}$. Note that it is crucial to include the additional points $\varnothing_x$ in the $\mathcal{Q}$ factors in (107), since they cover the cases where some of the positive ends of the moveable components do not have additional components attached to them (in the simplest case there is only one component, which must be moveable).

4.8 The linear restriction homomorphism

Temporarily, let's concentrate on the moduli spaces of cascade maps with $d = 1$. This gets rid of the transversality issue described in the discussion preceding Corollary 4.9 (since that issue assumed the existence of a vertex with valency $|v|-1 \ge |E| \ge 2$). Of course, the situation also simplifies in other respects, which makes it possible to shed some notational ballast. The maps $\mathbf{p}$ are constantly equal to 1, so we replace them in the notation everywhere by F. The trees T reduce to chains of two-valent vertices,

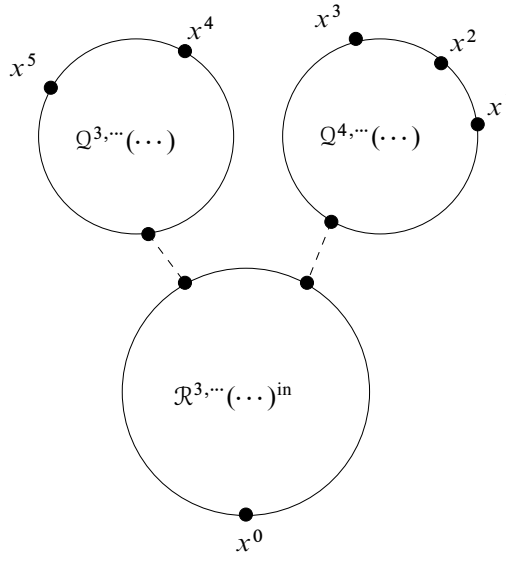

Figure 8

which we number as $\{v_1, \dots, v_l\}$ starting from the one closest to the root. We also write

$$\begin{aligned} &F_j = F_{v_j} \text{ for the subsets forming the decomposition } \mathbf{F}, \\ &\{w^0_{v_j} = w_{j-1}, w^1_{v_j} = w_j\} \text{ where } w_j = w^1 + |F_{j+1}| + \cdots + |F_l|, \\ &\{x^0_{v_j} = x_{j-1}, x^1_{v_j} = x_j\}, \text{ where } x_j \in \mathcal{X}_{w_j},\ x_0 = x^0, x_l = x^1. \end{aligned} \tag{108}$$

Finally, we identify each surface S_{v_j} with Z. With that in mind, the components of a cascade map in our moduli space are of the form

$$(\rho_j, Z, \boldsymbol{\phi}_j, u_j) \in \mathcal{P}^{2,F_j,\{w_{j-1},w_j\}}(x_{j-1}, x_j). \tag{109}$$

More precisely, the parametrized moduli space is the space of $(\rho_j, Z, \boldsymbol{\phi}_j, u_j)$ divided by the diagonal action of $\mathbb{R}$ on the last two factors, reflecting the fact that the identification $S_{v_j} \cong Z$ is not unique.

Let's begin translating the geometry of the moduli spaces into algebra. For this purpose, assume that L satisfies (40), and choose a grading as well as a *Pin* structure. These structures carry over to L^{in} (by restriction), and to L^ρ (extending them over the obvious isotopy). To any point $\{(\rho_j, Z, \boldsymbol{\phi}_j, u_j)\}$ in a zero-dimensional moduli space $\mathcal{Q}^{2,F,\mathbf{w}}(\mathbf{x})$ one can associate a map

$$|o^{\text{red}}_{\{(\rho_j,Z,\boldsymbol{\phi}_j,u_j)\}}|_{\mathbb{K}} \colon |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}} \tag{110}$$

analogous to (58). The orientation issues underlying this definition are discussed in Section 9.6. In the case of the formal point $\varnothing_x$ of the moduli space, we set (110) to be the identity on $|o_x|_{\mathbb{K}}$. Let $q^{1,F,\mathbf{w}}(\mathbf{x})$ be the sum of (110) over all points in the zero-dimensional moduli space. A symmetry argument parallel to Lemma 3.7 (see again Section 9.6 for more details) shows that

Lemma 4.11 *If $|F| > 1$ then $q^{1,F,\mathbf{w}}(\mathbf{x})$ vanishes.*

Of the two remaining cases, consider first that where $F = \varnothing$. In that situation, we omit F from the notation for moduli spaces and the resulting maps between orientation spaces, following the model of (52). Consider a partially compactified one-dimensional space $\overline{\mathcal{Q}}^{2,\mathbf{w}}(\mathbf{x})$, where $\mathbf{w} = \{w^0 = w, w^1 = w\}$ and $\mathbf{x} = \{x^0, x^1\}$. The ends of such spaces were described in Corollary 4.10. In this specific situation, they are modelled on

$$(0,\delta] \times \mathcal{R}^{2,\{w,w\}}(x^0, x^{new})^{\mathrm{in}} \times \mathcal{Q}^{2,\{w,w\}}(x^{new}, x^1), \tag{111}$$

where $x^{new} \in \mathcal{X}_w^{\mathrm{in}}$ satisfies $\deg(x^{new}) = \deg(x^0) - 1 = \deg(x^1)$. On the other hand, we have three kinds of boundary points. The first kind are those already contained in $\mathcal{Q}^{2,\mathbf{w}}(\mathbf{x})$ itself, namely those where the last component is moveable, and reaches time $\rho_l = 1$. The set of such boundary points is the union of the spaces

$$\mathcal{Q}^{2,\{w,w\}}(x^0, x^{new}) \times \mathcal{R}^{2,\{w,w\}}(x^{new}, x^1) \tag{112}$$

over all $x^{new} \in \mathcal{X}_w$ with $\deg(x^{new}) = \deg(x^0) = \deg(x^1) + 1$. As in the case of (107) and (111), this works because the left hand factor may contain a point $\varnothing_x$, which represents the case where the boundary point is a chain of total length $l = 1$. The remaining kinds of boundary points lie in the partial compactification, where the set E in (30) consists of a single edge, connecting v_j to v_{j+1} for some j, and where $\bar{E}$ is either $\varnothing$ or E. It turns out that for each collection $\{(\rho_j, Z, \boldsymbol{\phi}_j, u_j)\}$ satisfying $\rho_j = \rho_{j+1}$, the contributions coming from the two choices of $\bar{E}$ cancel each other, so that the total contribution is zero. In other words, one could glue together these boundary points in pairs, in a way compatible with the orientations of the one-dimensional moduli spaces induced from isomorphisms $o_{x^0} \cong \mathbb{R}$ and $o_{x^1} \cong \mathbb{R}$. Hence, the remaining relation is the one which counts points of the form (111) and (112). Taking into account signs, which will be elaborated on in Section 9.6, one finds that

$$\begin{aligned} &\textstyle\sum m^{1,\{w,w\},in}(x^0, x^{new}) \circ q^{1,\{w,w\}}(x^{new}, x^1) \\ &\textstyle- \sum q^{1,\{w,w\}}(x^0, x^{new}) \circ m^{1,\{w,w\}}(x^{new}, x^1) = 0, \end{aligned} \tag{113}$$

where m and m^{in} are the maps defined in Section 3.6 on M and M^{in}, respectively, and the sum is over all x^{new} with the appropriate Maslov index. The entire discussion

carries over to the remaining case $F = \{1\}$, $\mathbf{w} = \{w^0 = w+1, w^1 = w\}$ with minor modifications. In considering the analogue of (111), there are two possibilities, depending on whether the component carrying the popsicle is the first one or not. This leads to two types of ends,

$$\begin{aligned}&(0,\delta] \times \mathcal{R}^{2,\{w+1,w+1\}}(x^0, x^{new})^{\mathrm{in}} \times \mathcal{Q}^{2,\{1\},\{w+1,w\}}(x^{new}, x^1),\\ &(0,\delta] \times \mathcal{R}^{2,\{1\},\{w+1,w\}}(x^0, x^{new})^{\mathrm{in}} \times \mathcal{Q}^{2,\{w,w\}}(x^{new}, x^1).\end{aligned} \tag{114}$$

Similarly, there are two possibilities for the boundary points corresponding to (112),

$$\begin{aligned}&\mathcal{Q}^{2,\{w+1,w+1\}}(x^0, x^{new}) \times \mathcal{R}^{2,\{1\},\{w,w+1\}}(x^{new}, x^1),\\ &\mathcal{Q}^{2,\{1\},\{w+1,w\}}(x^0, x^{new}) \times \mathcal{R}^{2,\{w,w\}}(x^{new}, x^1).\end{aligned} \tag{115}$$

The upshot are the algebraic relations

$$\begin{aligned}&-\textstyle\sum m^{1,\{w+1,w+1\},in}(x^0, x^{new}) \circ q^{1,\{1\},\{w+1,w\}}(x^{new}, x^1)\\ &+\textstyle\sum m^{1,\{1\},\{w+1,w\},in}(x^0, x^{new}) \circ q^{1,\{w,w\}}(x^{new}, x^1)\\ &-\textstyle\sum q^{1,\{w+1,w+1\}}(x^0, x^{new}) \circ m^{1,\{1\},\{w+1,w\}}(x^{new}, x^1)\\ &-\textstyle\sum q^{1,\{1\},\{w+1,w\}}(x^0, x^{new}) \circ m^{1,\{w,w\}}(x^{new}, x^1) = 0.\end{aligned} \tag{116}$$

For each $w \geq \nu$, let $\gamma\colon CF^*(L; wH) \to CF^*(L^{\mathrm{in}}; wH^{\mathrm{in}})$ and $\lambda\colon CF^*(L; wH) \to CF^*(L^{\mathrm{in}}; (w+1)H^{\mathrm{in}})[-1]$ be the maps whose matrix coefficients are $q^{1,\{w,w\}}(x^0, x^1)$ and $q^{1,\{1\},\{w+1,w\}}(x^0, x^1)$, respectively. We want to combine them into a single ∂_q–linear map between suitable wrapped cochain groups. More precisely, if C^ν is the subcomplex from Lemma 3.11, we define

$$\begin{aligned}&\mathcal{F}^1\colon C^\nu \longrightarrow CW^*(L^{\mathrm{in}}; H^{\mathrm{in}}),\\ &\mathcal{F}^1(a + qb) = \gamma(a) + q\gamma(b) + \lambda(b).\end{aligned} \tag{117}$$

From (113), (116) and the definition of the wrapped differential (69), it follows that $\mathcal{F}^1$ is a chain homomorphism.

Remark 4.12 The decomposition (90) induces a projection from $CW^*(L; H) \to CW^*(L^{\mathrm{in}}; H^{\mathrm{in}})$, whose restriction to C^ν is precisely the summand in $\mathcal{F}^1$ coming from the formal points $\varnothing_x$ in our moduli spaces. One can therefore think of $\mathcal{F}^1$ as the projection map plus "instanton corrections" coming from rigid points in the parametrized moduli spaces $\mathcal{P}^{2,F,\mathbf{w}}(\mathbf{x})$ with $|F| \leq 1$. Of course, the projection by itself would generally not be a chain map.

4.9 Almost complex structures on cascades

In the original definition, moduli spaces of cascade maps were (open subsets of) products of parametrized moduli spaces. There is no intrinsic need for maintaining

this product structure, as long as the necessary relations between codimension one boundary faces are preserved. This allows us to enlarge the space of perturbations, which solves the previously encountered transversality issues.

Fix $d, \mathbf{p}, \mathbf{w}$, and let $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}$ be the partially compactified moduli space of weighted cascades. As in the case of popsicles, this is just a copy of the space $\overline{\mathcal{Q}}^{d+1,\mathbf{p}}$ defined in Section 2.8, but where we wish to emphasize that on each $(T, \mathbf{F}, E, \bar{E})$ stratum, the vertices of the tree come with weights $\mathbf{w}_v$ as well as flavours $\mathbf{p}_v$. Associated to any point in this stratum, represented by a collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v)\}$, we want to have families of almost complex structures of the following kind:

(118) For each edge e of T, a family $\mathbf{J}_e$ of almost complex structures. If the edge is finite and connects two vertices $v_\pm$, this family is parametrized by $(\rho, t) \in [\rho_{v_+}, \rho_{v_-}] \times [0,1]$. If the edge is semi-infinite and connects v to one of the leaves, $\mathbf{J}_e$ is parametrized by $[\rho_v, 1] \times [0,1]$; similarly, if it connects the root to v, $\mathbf{J}_e$ is parametrized by $(0, \rho_v] \times [0,1]$.

(119) For each vertex v of the tree, we want to have a family $\mathbf{J}_v$ of almost complex structures parametrized by S_v, which is asymptotically compatible with the strip-like ends (in the usual sense: as $s \to \pm\infty$, $J_{v,\epsilon^k(s,t)}$ converges to a family depending only on t, and this convergence is faster than exponential in any C^r norm).

It is crucial that $\mathbf{J}_e$ and $\mathbf{J}_v$ may depend on the whole collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v)\}$, and not just on the part of the data associated to that particular vertex or edge. This should really be reflected in the notation, but we avoid doing that in order to keep things reasonably brief. It is implicit in our definition that the dependency on the points in the moduli space of cascades is smooth. In fact, as one goes to the boundary strata in connected components (31) where some popsicles $(S_v, \boldsymbol{\phi}_v)$ degenerate into broken popsicles, one wants the same convergence behaviour as in Remark 3.3. Moreover, we require all our choices to be invariant under the action of $Sym^{\mathbf{p}}$.

The additional requirements on (118), (119) are of four kinds. The *large parameter value* restriction says the following. Let e be an edge, and suppose that the weight on the flags associated to it is w. Then, the restriction of $\mathbf{J}_e$ to $\rho = 1$, if it is nonempty, agrees with the family $\mathbf{J}_w$ previously chosen as part of the construction of wrapped Floer cohomology of L. Similarly, suppose that $\rho_v = 1$ for some vertex v. Then $\mathbf{J}_v = \mathbf{J}_{S_v,\boldsymbol{\phi}_v,w_v}$ should agree with the almost complex structure on that particular stable popsicle. Finally, note that the vertices where $\rho_v = 1$ necessarily form a finite union of subtrees inside T, each of which reaches up into the leaves. If we remove those subtrees, the result is another tree T', which comes with its own data $(\mathbf{p}', \mathbf{F}', E', \bar{E}')$. We then require that our families of almost complex structures, when restricted to the

edges and vertices of T', should depend only on that part of the cascade, and should in fact agree with the ones the cascade carries as part of our universal choice.

Secondly, we have a *small parameter value* requirement, whose structure is roughly similar to the previous one. Namely, there should be a $\delta > 0$ (depending on $d, \mathbf{p}, \mathbf{w}$) such that the restriction of $\mathbf{J}_e$ to any $\rho \leq \delta$, if nonempty, has properties analogous to those for an interpolating family. Namely, on $M^{\mathrm{in},\rho}$ it should agree with $\psi_*^\rho \mathbf{J}^{\mathrm{in}}_w$, as previously chosen when defining the wrapped Floer cohomology of L^{in}. Moreover, in some neighbourhood of $\partial M^{\mathrm{in},\rho} \subset \kappa^{\mathrm{in}}([\rho, 1] \times \partial M^{\mathrm{in}})$ it should be of contact type. Similarly, if $\rho_v \leq \delta$ for some vertex v, then $\mathbf{J}_v | M^{\mathrm{in},\rho}$ is the pushforward of the family $\mathbf{J}^{\mathrm{in}}_{S_v, \boldsymbol{\phi}_v, w_v}$ associated to that stable popsicle for M^{in}, and we have the same local condition near $\partial M^{\mathrm{in},\rho}$ as before. Slightly more interestingly, the vertices where $\rho_v \leq \delta$ form a subtree of T, reaching down to the root. Removing T yields a finite disjoint union of similar connected trees T'_k, and we require that on the part of our cascade corresponding to each of these, the almost complex structures should be those given by considering only the subtree.

Third, there is an *unstable consistency* condition. Pick a vertex v, a number $k \in \{0, \dots, |v| - 1\}$, and let e be the edge belonging to the kth flag adjacent to v. Then the limit of $\mathbf{J}_v$ over the kth end must agree with the family $\mathbf{J}_e$ specialized to $\rho = \rho_v$. Finally, we have *stable consistency*, which says that all the almost complex structures depend only on E, and are independent of $\bar{E}$. In particular, if we consider pairs of isomorphic codimension one boundary faces defined by $|E| = 1$ and $\bar{E} = \varnothing$ or E, then the almost complex structures on both copies of such a face must coincide.

One can build families of almost complex structures satisfying these conditions by gradually increasing the complexity, say by induction on $d + |F|$. For any given $(T, \mathbf{F})$, one first considers the large and small parameter value requirements. In those situations, the almost complex structures are completely determined by those for smaller trees, which had been chosen in a previous step. Note that even though both situations can apply simultaneously, the resulting combined condition is non-contradictory. Having done that, one then extends the choices over the entire moduli space, while fulfilling the relatively weak consistency conditions over the boundary strata, and finally averages to obtain $Sym^{\mathbf{p}}$–invariance.

4.10 Cascade maps (corrected versions)

Suppose now that (118), (119) have been chosen, satisfying all the requirements stated above. Fix $(d, \mathbf{p}, \mathbf{w})$. Consider pairs $(T, \mathbf{F})$ as before, but where now two-valent vertices with empty subsets $F_v = \varnothing$ are allowed, violating the previous stability

condition. A *cascade map* modelled on $(T, \mathbf{F})$ is a collection $\{(\rho_v, S_v, \boldsymbol{\phi}_v, u_v)\}$, where the maps u_v have to satisfy certain $\overline{\partial}$–equations, which we will now specify.

Suppose that we erase each unstable vertex of T, directly sewing together the two edges which meet there, and correspondingly forget all components $(\rho_v, S_v, \boldsymbol{\phi}_v)$ associated to such vertices. The outcome is a stabilized tree, and a stable cascade modelled on it, which defines a point of $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}$. Take the almost complex structures (118), (119) associated to that cascade. Consider first a stable vertex v of T, which survives to a vertex of the stabilized tree. In that case, we require that u_v should be a solution of (50) for the almost complex structure $\mathbf{J}_v$ belonging to that vertex. On the other hand, an unstable vertex v gets mapped to an edge of the stabilized tree, and in that case u_v should be a Floer trajectory for the almost complex structure $\mathbf{J}_e$ associated to that edge. As usual, two such trajectories which differ by a translation are considered to be the same.

From now on, $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ denotes the moduli space of cascade maps in the modified sense we have just explained. Essentially by construction, the definition and structure of the partial compactification $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ developed in Section 4.5 carries over, and so does the transversality theory in so far as it yielded positive results. The breakdown in transversality for one-dimensional partially compactified moduli spaces, discussed in Section 4.7, is resolved for the following simple reason. Take two vertices $v_{1,-}$ and $v_{2,-}$ of T, both of which are connected to the same vertex v_+ by edges pointing towards the root. Suppose first that the $v_{k,-}$ are stable (v_+ is automatically stable, since it is at least trivalent by assumption). By the previous discussion, we only need to consider the situation when $\rho_{v_+} = \rho_{v_-,1} = \rho_{v_-,2} = \rho$, and the weighted popsicles $(S_{v_-,k}, \boldsymbol{\phi}_{v_-,k}, \mathbf{w}_{v_-,k})$ are isomorphic. In the framework of (119), the families of almost complex structures on $S_{v_-,k}$ can usually be chosen independently of each other, which means that generically, rigid solutions of the $\overline{\partial}$–equation will not appear for the same parameter value, thus removing the obstruction to transversality. There are two exceptions to the statement about independence. If $\rho = 1$, the almost complex structures on both $v_{-,k}$ components are necessarily the same. But in that case, transversality for the original construction of wrapped Floer cohomology of L ensures that the components $\{(\rho_v, S_v, \boldsymbol{\phi}_v, u_v)\}$ for $v = v_+, v_{k,-}$ must all be moveable, which can never happen in a one-dimensional moduli space. The other, parallel but slightly more complicated, exception occurs when $\rho \leq \delta$. Supposing that δ has been chosen sufficiently small, (99) will then apply all components of our popsicle map lying between $v_{-,k}$ and the root, which implies that these components must be moveable, leading to the same dimensional contradiction as before (this is not a circular argument: the upper bounds on parameters ρ provided by Lemma 7.4 depend only on action considerations, hence are independent of the choice of almost complex structures).

Finally, one also needs to consider the situation when both $\rho_{v_-,k}$ are unstable, but the argument there is essentially the same.

Remark 4.13 To summarize, we have addressed the transversality issue by enlarging the class of allowed perturbations, while preserving enough consistency to retain the basic relations between the boundary strata of one-dimensional moduli spaces. There are of course other possibilities. For instance, Fukaya, Oh, Ohta and Ono [13] use Kuranishi structures induced by virtual perturbations for the same purpose (the version of Kuranishi space theory described by Joyce [18] seems particularly relevant here, since we need to deal with moduli spaces for all d at the same time).

4.11 The restriction homomorphism

Assume now that generic choices of (118) and (119) have been made, subject to all the conditions above, making the spaces $\mathcal{Q}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ and their partial compactifications regular. As usual we consider only the case when $Sym^{\mathbf{p}}$ is trivial, and replace $\mathbf{p}$ by $F \subset \{1,\dots,d\}$ in the notation. Write

$$q^{d+1,F,\mathbf{w}}(\mathbf{x})\colon |o_{x^d}|_{\mathbb{K}} \otimes \cdots \otimes |o_{x^1}|_{\mathbb{K}} \longrightarrow |o_{x^0}|_{\mathbb{K}} \tag{120}$$

for the map obtained by counting the contributions from a zero-dimensional moduli space $\mathcal{Q}^{d+1,F,\mathbf{w}}(\mathbf{x})$ (as announced at the beginning of the section, we won't properly discuss the signs that enter into this definition). For $d = 1$ this reproduces the definition from Section 4.8, in a slightly generalized version since the choices of almost complex structures considered there form a strict subset of the ones allowed here. We assemble the (120) for all $\mathbf{x}$ into a multilinear map

$$CF^*(L; w^d H)[q] \otimes \cdots \otimes CF^*(L; w^1 H)[q] \xrightarrow{\mathcal{F}^{d,F,\mathbf{w}}} CF^*(L^{\mathrm{in}}; w^0 H^{\mathrm{in}})[q] \tag{121}$$

of degree $1-d$. More precisely, the first summand of (121), which takes values in $CF^*(L^{\mathrm{in}}; w^0 H)$, has coefficients (120) with suitable signs, and we define the second summand to make $\mathcal{F}^{d,F,\mathbf{w}}$ a ∂_q–linear map in the same sense as in (76). In the case where $d = 1$ and $F = \varnothing$, we also add a trivial term which is projection $CF^*(L; wH)[q] \to CF^*(L^{\mathrm{in}}; wH^{\mathrm{in}})[q]$, compare Remark 4.12.

Recall from (100) that all weights that we are considering are larger than a fixed weight $\nu > 0$. In particular, we extend each (121) by zero to a map defined on the whole of $(C^\nu)^{\otimes d} \to CW^*(L; H)^{\otimes d}$, and taking values in $CW^*(L^{\mathrm{in}}; H^{\mathrm{in}})$. Define $\mathcal{F}^d$ to be the sum of all these maps for fixed d but varying $(F, \mathbf{w})$.

Theorem 4.14 *The maps $\mathcal{F}^d$ satisfy the A_∞–homomorphism equations*

$$\begin{aligned}(122)\qquad &\sum_{d_1+\cdots+d_k=d} \mu^{\mathrm{in},k}(\mathcal{F}^{d_k}(c_d,\dots,c_{d-d_k+1}),\dots,\mathcal{F}^{d_1}(c_{d_1},\dots,c_1)) \\ &= \sum_k (-1)^{\overline{\deg}(c^1)+\cdots+\overline{\deg}(c^{i-1})} \mathcal{F}^{d+}(c^d,\dots,c^{i+d-}, \\ &\qquad\qquad \mu^{d-}(c^{i+d-1},\dots,c^i),c^{i-1},\dots,c^1).\end{aligned}$$

Moreover, they can be extended to an A_∞–homomorphism $CW^*(L)\to CW^*(L^{\mathrm{in}})$, in a way which is unique up to homotopy.

Proof The underlying geometry is very simple: the left hand side of the equation collects contributions from the ends of one-dimensional spaces $\overline{\mathcal{Q}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$, where the component for the vertex next the root is the moveable one, and (99) applies to it; the right hand side takes into account those boundary points where, for one of the vertices nearest to the leaves, the parameter value becomes equal to 1. All other boundary strata cancel in pairs, just as in the corresponding argument in Section 4.8.

We know from Lemma 3.11 that the inclusion $C^\nu\to CW^*(L;H)$ is a quasi-isomorphism of A_∞–algebras. It is a general algebraic fact that there is always a quasi-isomorphism in inverse direction, $CW^*(L;H)\to C^\nu$, whose restriction to C^ν is the identity, and this is unique up to homotopy (see for instance Keller [19, Section 3.7] and the references there). By composing the given $\mathcal{F}$ with such an inverse, one can get an A_∞–homomorphism defined on the whole of $CW^*(L;H)$. □

Remark 4.15 The reason why $\mathcal{F}$ is initially only defined on C^ν comes down to Lemma 7.3. There is a simpler alternative, in the spirit of Lemmas 3.1 and 4.1. Namely, one extends h to a function on M, which vanishes near $\partial M\cup\partial M^{\mathrm{in}}$, and then makes a replacement $\theta\mapsto\theta-dh$. One effect of this is that the h terms in (150) become absorbed into the $\int\theta$, which removes the previous obstruction. Of course, it simultaneously changes Z, hence also our family of rescaled Lagrangian submanifolds L^ρ.

5 Complements

Our first, and rather straightforward, task in this section is to extend the previous construction to wrapped Fukaya categories. After that, we return to the two-dimensional case, which could not be treated previously due to the failure of (39) to be generic, and outline a simple workaround, which is to increase dimensions by taking the product with a standard factor (other solutions are possible; see Remark 5.1).

5.1 Wrapped Fukaya categories

Let $\mathbf{L} = \{L_i\}_{i\in I}$ be a collection of Lagrangian submanifolds of M, which is at most countable. Suppose that each of them satisfies (35) as well as (40), and choose gradings as well as *Pin* structures. An integer Reeb chord from ∂L_{i_0} to ∂L_{i_1} is a map $x\colon [0,1] \to \partial M$ satisfying $dx/dt = wR$ for some positive integer w, such that $x(0) \in \partial L_{i_0}$, $x(1) \in \partial L_{i_1}$. Integer X–chords going from L_{i_0} to L_{i_1} are defined in the same way. Generalizing (37), we assume that there are no integer Reeb chords between the boundaries of any two Lagrangian submanifolds in our collection. Having that, we choose some Hamiltonian H as in (34), such that the appropriate generalizations of (38) and (39) hold. As before, these are generic assumptions (true after a generic rescaling of θ, and small isotopies of the L_i) provided that $\dim(M) \geq 4$. Let $CF^*(L_{i_0}, L_{i_1}; wH)$ be the Floer cochain complex generated by $|o_x|_{\mathbb{K}}$, where x runs over all X–chords of length w going from L_{i_0} to L_{i_1}. Generalizing (68), we define the wrapped complex by

$$CW^*(L_{i_0}, L_{i_1}; H) = \bigoplus_{w=1}^{\infty} CF^*(L_{i_0}, L_{i_1}; wH)[q] \tag{123}$$

The wrapped Fukaya category $\mathcal{W}(\mathbf{L})$ has the L_i as objects, (123) as morphism spaces, and an A_∞–structure defined by extending the construction from Section 3, in an essentially straightforward way to multiple Lagrangians.

Now suppose that we have a Weinstein subdomain $M^{\mathrm{in}} \subset M$, and that each L_i satisfies (80) with a uniform lower bound on the size of the neighbourhood of the boundary of L_i where the restriction of θ vanishes. In this situation, one can define an A_∞–functor $\mathcal{W}(\mathbf{L}) \to \mathcal{W}(\mathbf{L}^{\mathrm{in}})$, where $\mathbf{L}^{\mathrm{in}} = \{L_i^{\mathrm{in}} = L_i \cap M^{\mathrm{in}}\}$ by following the same procedure as before. To be slightly more precise, for every pair of elements in I there will be a constant ν_{i_0,i_1} such that the analogue of Lemma 7.3 holds with $w \geq \nu_{i_0,i_1}$. By inspecting the proof of Lemma 7.3, we find that we may choose ν_{i_0,i_1} so that for any collection $(L_{i_0}, \cdots, L_{i_d})$, we have

$$\nu_{i_0,i_d} \leq \sum_k \nu_{i_k,i_{k+1}}. \tag{124}$$

In particular, the collection of morphism spaces which are sums (123) restricted to $w \geq \nu_{i_0,i_1}$ forms a subcategory on which we can define a restriction functor. However, this subcategory is quasi-isomorphic to the whole thing by the same argument as before, and one can use that to formally extend the functor.

5.2 Stabilization

Consider $D^*S^1 = [-1,1] \times S^1$. We use standard coordinates (p,q), so that the tautological one-form is $p\,dq$. Let M be a Liouville domain, and $k\colon M \times S^1 \to \mathbb{R}$ a function which vanishes in a neighbourhood of $\partial M \times S^1$. One can then consider the product $M \times D^*S^1$, with the one-form

$$\widetilde{\theta} = \theta + C\ p\,dq - dk \tag{125}$$

for some constant C (pullbacks by projection to various factors are implicit in this definition). The associated Liouville field is $\widetilde{Z} = Z + C\ p\partial_p + Y_q + (\partial k/\partial q)\partial_p$, where Y_q is the Hamiltonian vector field of $k(\cdot,q)$ on M. This will point strictly outwards along both boundary faces, provided that $C \gg 0$ is big enough. Hence, after rounding off the corners suitably, we get a Liouville domain $\widetilde{M} \subset M \times D^*S^1$.

Suppose that $L \subset M$ is a Lagrangian submanifold satisfying (80). We take the product of L with some cotangent fibre $D^*_a S^1 = [-1,1] \times \{a\}$. Assume that the function k appearing above has been chosen in such a way that

$$dk(\cdot,a)|L = \theta|L. \tag{126}$$

In that case, the Lagrangian submanifold $\widetilde{L} = (L \times D^*_a S^1) \cap \widetilde{M}$ satisfies $\widetilde{\theta}|\widetilde{L} = 0$, hence our previous construction of the wrapped A_∞–algebra, with a suitably chosen Hamiltonian function $\widetilde{H}$, applies to it (gradings and *Pin* structures for L induce the same kind of structures on $\widetilde{L}$). Recall that generators of $CF^*(\widetilde{L}; w\widetilde{H})$ correspond naturally to $\widetilde{X}$–chords of length w, with boundary on $\widetilde{L}$. The image of any such chord $\widetilde{x}$ under projection $\widetilde{M} \to S^1$ is a closed loop, whose winding number we denote by $\alpha(\widetilde{x})$. Obviously, the moduli spaces $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\widetilde{\boldsymbol{x}})$ can be nonempty only if $\alpha(\widetilde{x}^0) = \alpha(\widetilde{x}^1) + \cdots + \alpha(\widetilde{x}^d)$. Hence, the subcomplex of $CW^*(\widetilde{L}; \widetilde{H})$ consisting only of those generators with zero winding number is in fact an A_∞–subalgebra. For $2n = dim(M) > 2$, one can prove that it is quasi-isomorphic to the previously defined $CW^*(L; H)$. We will not prove this here, but with this motivation in mind, one can use the subalgebra as a replacement for the wrapped complex in the two-dimensional case. It is easy to see that the Viterbo restriction maps also preserve these subalgebras.

Slightly more generally, given a finite collection $\mathbf{L} = \{L_i\}_{i \in I}$ satisfying (80), one can take $\widetilde{L}_i = (L_i \times D^*_{a_i} S^1) \cap \widetilde{M}$ for pairwise distinct a_i, and choose k in such a way that the analogue of (126) is satisfied for all i. In addition, fix preimages of all the a_i under the projection $\mathbb{R} \to \mathbb{R}/\mathbb{Z} = S^1$. For any chord $\widetilde{x}$ which connects $\widetilde{L}_{i_0}$ to $\widetilde{L}_{i_1}$, one then has an integer $\alpha(\widetilde{x})$ measuring the difference in winding numbers between the projection of $\widetilde{x}$ and that of any path in $\mathbb{R}$ connecting the preimages of a_{i_0}, a_{i_1}. The subcomplexes of the $CW^*(\widetilde{L}_{i_0}, \widetilde{L}_{i_1}; \widetilde{H})$ generated by chords with $\alpha(\widetilde{x}) = 0$ form

an A_∞–subcategory of $\mathcal{W}(\widetilde{\mathbf{L}})$. As before, this can be used as a replacement for $\mathcal{W}(\mathbf{L})$ in the two-dimensional case.

Remark 5.1 The stabilization construction is actually a version of the approach outlined in Remark 4.15, and correspondingly comes with some limitations. Most substantially, it only works for finitely many Lagrangian submanifolds at once. In principle, one can overcome that particular restriction by using the universal cover $D^*\mathbb{R} = [-1,1] \times \mathbb{R}$ as a stabilizing factor, but that no longer strictly falls into the previously considered framework, hence requires some additional work (in particular, in the choice of Hamiltonian functions, to avoid problems arising from the noncompactness of the $\mathbb{R}$ factor). Among the other possible approaches which avoid stabilization entirely, the most obvious one, since we are dealing with a transversality problem, is to use Kuranishi structures. Alternatively, along more classical lines, one could try to replace the multiples wH by a more general sequence of Hamiltonian functions, or make the Hamiltonians time-dependent (which has the disadvantage of making action computations less straightforward; we have not considered this in detail).

6 Popsicles and stable maps

This section revisits the compactified moduli spaces $\overline{\mathcal{R}}^{d+1,\mathbf{p}}$, filling in some gaps left in the original discussion (Section 2). Instead of starting with the combinatorial description of these spaces as union of strata, as in (16), we construct them from a complex geometry perspective, based on the existing theory of stable maps (this follows the strategy used by Fukaya and Oh [12] for Stasheff polyhedra).

6.1 A complex analogue

Fix $d \geq 1$ and $\mathbf{p} = \{p_f\}$, $f \in F$. Consider the projective line $\mathbb{C}P^1 = \mathbb{C} \cup \{\infty\}$, and let $Y = (\mathbb{C}P^1)^F$ be the product of such lines indexed by F, with coordinates y_f. For $k \in \{0, \dots, d\}$, define subsets $H^k \subset Y$ by

$$H^k = \begin{cases} \text{the point where all } y_f = 0 & \text{if } k = 0, \\ \text{the subset where } y_f = \infty \text{ for all } f \text{ with } p_f = k & \text{if } k \in \{1, \dots, d\}. \end{cases} \tag{127}$$

A $\mathbf{p}$–flavoured *lollipop* is a closed genus zero Riemann surface C, equipped with a collection of $d+1$ distinct ordered marked points $\boldsymbol{\zeta} = \{\zeta^0, \dots, \zeta^d\}$, together with a holomorphic map $\boldsymbol{\psi}\colon C \to Y$ of degree $(1, \dots, 1)$ such that

$$\boldsymbol{\psi}(\zeta^k) \in H^k \quad \text{for all } k. \tag{128}$$

Equivalently, the components $\psi_f\colon C \to \mathbb{C}P^1$ are isomorphisms satisfying $\psi_f(\zeta^0) = 0$, $\psi_f(\zeta^{p_f}) = \infty$. More generally, define a *broken lollipop* to be a genus zero surface C with nodes, equipped with a collection $\boldsymbol{\zeta}$ of smooth points, and with a map satisfying the same incidence condition (128) as before. The stability condition says that on each irreducible component of C which carries less than three special points (nodes or marked points), $\boldsymbol{\psi}$ may not be constant; in our case, stability and the nature of (128) rule out having just one special point on a component, so there are always at least two. For $d + |F| \geq 2$, we then have the moduli space of lollipops $\mathcal{M}^{d+1,\mathbf{p}}$ and its compactification $\overline{\mathcal{M}}^{d+1,\mathbf{p}}$, the space of stable broken lollipops, considered as a subspace of the space of stable maps with target Y. The compactification has a decomposition into strata $\mathcal{M}^{T,\mathbf{F}}$ indexed by data similar to that in (16): the difference is that instead of a ribbon structure, T only comes with a labeling of its semi-infinite edges by $\{0,\dots,d\}$. The pair $(T,\mathbf{F})$ fixes the topological type of the broken lollipop (the structure of the nodal curve, including the choice of irreducible components on which the marked points lie; and the degrees of the map $\boldsymbol{\psi}$ on each irreducible component). There is a natural action of $(\mathbb{C}^*)^F$ on Y by rescaling coordinates, and therefore an induced action on our moduli spaces. Similarly, we have an action of $Sym^{\mathbf{p}}$, which comes from permuting the factors in the target space Y. In both cases, the important point is that each H^k remains invariant.

Lemma 6.1 *$\overline{\mathcal{M}}^{d+1,\mathbf{p}}$ is a smooth compact complex manifold.*

Proof The moduli space of all genus zero stable maps with target Y and degree $(1,\dots,1)$ is known to be smooth and compact: compactness is a general property, and smoothness follows from the convexity of Y (see Kontsevich and Manin [22, Definition 2.4.2]). By definition, $\overline{\mathcal{M}}^{d+1,\mathbf{p}}$ is a closed subspace, hence itself compact. Smoothness does not follow directly, but the proof proceeds along well-established lines. First of all, the infinitesimal deformation space $\mathcal{T}$ of any stable broken lollipop sits in a long exact sequence

$$(129)\quad 0 \longrightarrow \mathrm{Ext}^0(\Omega^1_C(\zeta^0+\cdots+\zeta^d),\mathcal{O}_C) \longrightarrow H^0(C,\mathcal{E}) \longrightarrow \mathcal{T} \to \\ \longrightarrow \mathrm{Ext}^1(\Omega^1_C(\zeta^0+\cdots+\zeta^d),\mathcal{O}_C) \longrightarrow H^1(C,\mathcal{E}),$$

where $\mathcal{E} \subset \boldsymbol{\psi}^* TY$ is the subsheaf of vector fields which at each marked point ζ^k are tangent to H^k. More specifically, $\mathcal{E}$ is a direct sum of line bundles $\mathcal{E}_f \subset \psi_f^* T\mathbb{C}P^1$. Each $\psi_f^* T\mathbb{C}P^1$ itself is of degree 2 on exactly one irreducible component of C, and trivial on all others. Moreover, $\mathcal{E}_f \subset \psi_f^* T\mathbb{C}P^1$ is defined by imposing vanishing conditions at the points ζ^0, ζ^{p_f}. From this and the fact that our curve has genus zero,

it is easy to see that $H^0(C, \mathcal{E}_f) \cong \mathbb{C}$, $H^1(C, \mathcal{E}_f) = 0$. Hence the obstruction term, which is the last one in (129), vanishes, ensuring smoothness. □

This computation has another useful consequence. Fix a stable broken lollipop, modelled on a tree T. Nodes of C then correspond canonically to interior edges e of T. Choose, for each node, local holomorphic coordinates $\{z_{e,\pm}\}$ near its preimages in the normalization. Then there is a deformation

$$\mathcal{C} \longrightarrow \mathcal{U} \tag{130}$$

of the underlying surface with marked points $(C, \{\zeta^0, \dots, \zeta^d\})$, parametrized by some neighbourhood of the origin $\mathcal{U} \subset \prod_e \mathbb{C}$. The local model for the deformation is the standard smoothing

$$z_{e,-} z_{e,+} = \delta_e, \tag{131}$$

applied to each node. Possibly after shrinking $\mathcal{U}$, the map $\boldsymbol{\psi}$ extends to $\Psi : \mathcal{C} \to Y$, in such a way that (128) is satisfied on each fibre. This is a general fact about deformations of lollipops, which follows directly from the Kodaira–Spencer deformation theory of holomorphic maps: in the notation from Lemma 6.1, the obstruction to such an extension lies in $H^1(C, \mathcal{E})$, which we know to be zero. Of course, the extension is not unique, but the ambiguity is easy to understand: the group of maps $\mathcal{U} \to (\mathbb{C}^*)^F$ which equal $(1, \dots, 1)$ at the origin acts simply transitively on the set of all possible choices. Once we have chosen Ψ, the resulting family is classified by a holomorphic map $\mathcal{U} \to \bar{\mathcal{M}}^{d+1,\mathbf{p}}$ (this uses the universality property of the moduli space, which in turn is based on the fact that stable broken lollipops have no nontrivial automorphisms).

Instead of starting with a single stable broken lollipop, one can also take a family, in particular that corresponding to one of the strata $\mathcal{M}^{T,\mathbf{F}}$. Suppose that for each surface in the family, we have chosen local coordinates near the nodes, as in the previous discussion. Given that, carry out the smoothing process, and find suitable maps Ψ. As before, this yields a classifying map, which now takes a form analogous to (11):

$$\mathcal{M}^{T,\mathbf{F}} \times \prod_e \mathbb{C} \supset \mathcal{U}^{T,\mathbf{F}} \longrightarrow \bar{\mathcal{M}}^{d+1,\mathbf{p}}. \tag{132}$$

Lemma 6.2 *Possibly after shrinking $\mathcal{U}^{T,\mathbf{F}}$, (132) is a diffeomorphism onto a neighbourhood of $\mathcal{M}^{T,\mathbf{F}} \subset \bar{\mathcal{M}}^{d+1,\mathbf{p}}$.*

Proof Take a stable broken lollipop $(C, \boldsymbol{\psi})$ representing a point of $\mathcal{M}^{T,\mathbf{F}}$. Let $\widetilde{C} \to C$ be the normalization, and $\{\nu_{e,\pm}\}$ the preimages of the nodes in $\widetilde{C}$. Concerning the

tangent space of our stratum at this point, we have an exact sequence

$$0 \longrightarrow \mathrm{Ext}^0(\Omega^1_C(\zeta^0 + \cdots + \zeta^d), \mathcal{O}_C) \longrightarrow H^0(C, \mathcal{E}) \longrightarrow T\mathcal{M}^{T,\mathbf{F}} \to \\ \longrightarrow \mathrm{Ext}^1(\Omega^1_{\widetilde{C}}(\textstyle\sum_k \zeta^k + \sum_e \nu_{e,+} + \sum_e \nu_{e,-}), \mathcal{O}_{\widetilde{C}}) \longrightarrow 0. \tag{133}$$

By comparing this with (129), one sees that $T\mathcal{M}^{d+1,\mathbf{p}}/T\mathcal{M}^{T,\mathbf{F}}$ can be identified with the quotient $\mathrm{Ext}^1(\Omega^1_C(\sum_k \zeta^k), \mathcal{O}_C)/\mathrm{Ext}^1(\Omega^1_{\widetilde{C}}(\sum_k \zeta^k + \sum_e \nu_{e,+} + \sum_e \nu_{e,-}), \mathcal{O}_{\widetilde{C}})$. On the other hand, there is a well-known short exact sequence (see Harris and Morrison [15, page 100]) which identifies that quotient with

$$\bigoplus_e T\widetilde{C}_{\nu_{e,+}} \otimes T\widetilde{C}_{\nu_{e,-}}. \tag{134}$$

Geometrically, projection from $\mathrm{Ext}^1(\Omega^1_C(\sum_k \zeta^k), \mathcal{O}_C)$ to (134) measures the amount of smoothing that is applied to the nodes (the reduction to first order of the parameter δ_e from (131), essentially). This means that the differential of (132) induces an isomorphism from the normal bundle of $\mathcal{M}^{T,\mathbf{F}}$ to (134). □

Addendum 6.3 In the construction of (132), we have implicitly assumed that the local coordinates used in the gluing process are not only holomorphic on each single surface, but also depend holomorphically on the moduli. In our intended applications, the dependance is in fact only differentiable, which makes it difficult to apply the universality property directly. One way to get around this problem is to restrict attention to *rational* local coordinates, which are ones that extend to isomorphisms from $\mathbb{C}P^1$ to the relevant component of the broken lollipop. There is an enhanced moduli space $\widehat{\mathcal{M}}^{T,\mathbf{F}}$ of broken lollipops equipped with such coordinates for all the nodes, and this comes with a canonical holomorphic gluing map

$$\widehat{\mathcal{M}}^{T,\mathbf{F}} \times \prod_e \mathbb{C} \supset \widehat{\mathcal{U}}^{T,\mathbf{F}} \longrightarrow \overline{\mathcal{M}}^{d+1,\mathbf{p}}. \tag{135}$$

On the other hand, the forgetful map $\widehat{\mathcal{M}}^{T,\mathbf{F}} \to \mathcal{M}^{T,\mathbf{F}}$ is a holomorphic principal bundle. A choice of rational local coordinates, depending differentiably on moduli, is simply a smooth section of that bundle. One then defines the generalization of (132) by composing this section and (135). This is obviously differentiable, and it is easy to see that Lemma 6.2 continues to hold.

6.2 Real structures

The complex conjugate of a lollipop is the complex conjugate $\overline{C}$ of the underlying curve, equipped with the map $\overline{\boldsymbol{\psi}}$. A real lollipop is one which is invariant under this

operation. This means that it has an anti-holomorphic involution $\iota\colon C \to C$ preserving each ζ^k, and such that $\boldsymbol{\psi} \circ \iota = \bar{\boldsymbol{\psi}}$. The same applies to broken lollipops. Note that in the stable case, a real structure (if it exists) is necessarily unique. Complex conjugation induces an anti-holomorphic involution of the compactified moduli space, whose fixed point set $\bar{\mathcal{M}}^{d+1,\mathbf{p}}_{\mathbb{R}}$ parametrizes real broken lollipops.

We now revisit (130) in the presence of real structures. Start with $(C, \boldsymbol{\psi})$ which is real. For each node of C, choose local holomorphic coordinates around its preimages in the normalization, and assume that these coordinates are compatible with the anti-holomorphic involution. In that case, the family $\mathcal{C}$ inherits a real structure as well, and the map Ψ on it can be chosen to be compatible with that structure. This is still not unique, but we can make a more specific choice based on additional geometric data, as follows. Suppose that $(C, \boldsymbol{\psi})$ is modelled on $(T, \mathbf{F})$. This means in particular that the irreducible components C_v are indexed by vertices of T, and that for each $f \in F_v$, the restriction $\psi_f|C_v\colon C_v \to \mathbb{C}P^1 = \mathbb{C} \cup \{\infty\}$ is an isomorphism. Choose, for each f, an additional marked point q_f lying in $\psi_f^{-1}(S^1) \subset C_v$, which should not be a node. Since the deformation $\mathcal{C}$ is defined by a local process near the nodes, these points will survive into the nearby fibres, yielding sections $\kappa_f\colon \mathcal{U} \to \mathcal{C}$ of (130). For any choice of Ψ compatible with the real structure, we get an associated collection of real analytic functions

$$|\Psi_f \circ \kappa_f|\colon \mathcal{U}_{\mathbb{R}} \longrightarrow \mathbb{R}^*, \tag{136}$$

each of which sends the origin to 1. If we change Ψ_f by a a real analytic map $\mathcal{U} \to \mathbb{C}^*$, the effect is to multiply (136) with the restriction of the same map to $\mathcal{U}_{\mathbb{R}}$. Hence, possibly after shrinking $\mathcal{U}_{\mathbb{R}}$, there is a unique choice of Ψ such that all the (136) become constant equal to 1. We call the resulting $(\mathcal{C}, \Psi)$ a *geometric deformation* of the real broken lollipop $(C, \boldsymbol{\psi})$. Such deformations depend on the choice of local holomorphic coordinates near the nodes, and on the points q_f.

Take a stratum $\mathcal{M}^{T,\mathbf{F}}_{\mathbb{R}}$ in the real part of the lollipop moduli space. Suppose that along this stratum, we choose rational local coordinates around the nodes, as well as additional marked points as described above. For the same reason as in Addendum 6.3, it is sufficient to have these coordinates and points vary differentiably over the moduli space. The geometric deformation construction then yields a gluing map, whose real part is a diffeomorphism onto a neighbourhood of our stratum:

$$\mathcal{M}^{T,\mathbf{F}}_{\mathbb{R}} \times \prod_e \mathbb{R} \supset \mathcal{U}^{T,\mathbf{F}}_{\mathbb{R}} \longrightarrow \bar{\mathcal{M}}^{d+1,\mathbf{p}}_{\mathbb{R}}. \tag{137}$$

6.3 Popsicles vs lollipops

Take, as before, a tree T with $d+1$ semi-infinite edges, labeled by $\{0,\dots,d\}$. We say that T is of ribbon type if it can be embedded into the plane in such a way that the given numbering of the semi-infinite edges is compatible with the counterclockwise cyclic ordering. If such an embedding exists, it is essentially unique, so we do indeed get a distinguished ribbon structure in the ordinary sense. Suppose from now on that T is of this type; and let $(C,\boldsymbol{\psi})$ be a stable broken lollipop modelled on $(T,\mathbf{F})$, which carries a real structure. From the ribbon structure of T, we get a canonical ordering of the flags adjacent to any given vertex v (compare Section 2.3), hence of the special points (nodes and marked points) on each irreducible component C_v. We say that the real structure is admissible if there is some orientation of the real parts $C_{v,\mathbb{R}} \cong S^1$ which is compatible with the ordering of the special points, and such that for each $f \in F_v$, the diffeomorphism $\psi_{f,\mathbb{R}}\colon C_{v,\mathbb{R}} \to \mathbb{R}\cup\{\infty\} = Y_{f,\mathbb{R}}$ becomes orientation-preserving. Again, there can be at most one such orientation. Supposing that one exists, one can construct a popsicle modelled on $(T,\mathbf{F})$ as follows. For each v, take $\bar{S}_v$ to be the disc in C_v whose boundary is $C_{v,\mathbb{R}}$ with the positive orientation, and then remove all special points to get a pointed disc S_v. The popsicle maps are defined by

$$\phi_f = \tfrac{1}{\pi}\log(\psi_f) \mid S_v. \tag{138}$$

Conversely, the broken popsicle with components S_v and maps ϕ_f, $f \in F_v$, determines the original (C,ψ). Now suppose that we are given popsicle sticks on each S_v, which allow one to convert the information contained in the popsicle maps ϕ_f into sprinkles $q_f \in S_v \subset C_v$. Because of (138), these points necessarily satisfy $\psi_f(q_f) \in S^1$. Hence, the geometric deformation introduced above actually reproduces the gluing process from Section 2.4. More precisely, geometric deformation with non-positive gluing parameter corresponds to gluing of popsicles, whereas a positive gluing parameter results in a non-admissible real structure.

To conclude, $\bar{\mathcal{R}}^{d+1,\mathbf{p}}$ can be identified with the subset of $\bar{\mathcal{M}}_{\mathbb{R}}^{d+1,\mathbf{p}}$ consisting of those broken lollipops which have admissible real structures. Consider some boundary stratum $\mathcal{R}^{T,\mathbf{F}}$, and suppose that for each family $\mathcal{S}^{|v|,\mathbf{p}_v} \to \mathcal{R}^{|v|,\mathbf{p}_v}$, we have made choices of rational strip-like ends as well as popsicle sticks. The associated gluing map (20) is a restriction of (137), which means that

Corollary 6.4 *Possibly after making $\mathcal{U}^{T,\mathbf{F}}$ smaller, (20) is a diffeomorphism onto a neighbourhood of $\mathcal{R}^{T,\mathbf{F}}$ inside $\bar{\mathcal{R}}^{d+1,\mathbf{p}} \subset \bar{\mathcal{M}}_{\mathbb{R}}^{d+1,\mathbf{p}}$.*

Corollary 6.5 *$\bar{\mathcal{R}}^{d+1,\mathbf{p}} \subset \bar{\mathcal{M}}_{\mathbb{R}}^{d+1,\mathbf{p}}$ is a smooth submanifold with corners.*

6.4 Compactifying the universal family

We now discuss, in slightly less detail, the construction of the compactified total space $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$. Consider stable maps $\boldsymbol{\psi}\colon C \to Y$ of the same genus and degree as before, but where C now carries $d+2$ marked non-special points $\{\zeta^0,\dots,\zeta^d,z\}$. The first $d+1$ of these should satisfy the usual incidence condition (128), and the last one can be mapped freely to Y. On one hand, the computation from Lemma 6.1 shows that such maps form a smooth moduli space $\mathcal{C}^{d+1,\mathbf{p}}$. On the other hand, it is part of the general theory of stable maps (the prototypical case of stable curves was treated by Knudsen [21]) that there is a natural holomorphic forgetful map

$$\overline{\mathcal{C}}^{d+1,\mathbf{p}} \longrightarrow \overline{\mathcal{M}}^{d+1,\mathbf{p}}, \tag{139}$$

and that this can in fact be identified with the universal family of stable curves over the moduli space.

Consider the subset $\overline{\mathcal{S}}^{d+1,\mathbf{p}} \subset \overline{\mathcal{C}}^{d+1,\mathbf{p}}$ consisting of those points $(C,\boldsymbol{\zeta}\cup\{z\},\boldsymbol{\psi})$ which satisfy the following two conditions. First, the image under (139) should lie in $\overline{\mathcal{R}}^{d+1,\mathbf{p}}$. Secondly, if it is the case that $(C,\boldsymbol{\zeta},\boldsymbol{\psi})$ is stable (note that we are forgetting z here), we require additionally that z should lie in the subset of C given by the associated compactified broken popsicle. One checks easily that (as a set) this can be identified with the compactification of the universal family (17). The case where forgetting z makes the map unstable deserves special mention. This occurs whenever $\boldsymbol{\psi}$ is constant on an irreducible component of C which contains z as well as

(140)
- (i) exactly one of the points ζ^k, and exactly one nodal point; or
- (ii) none of the ζ^k, and exactly two nodal points.

In terms of (17), this corresponds to having the compactification $\overline{S}$ of an (ordinary or broken) popsicle, together with a point $z\in\overline{S}$ which either (i) agrees with one of the points at infinity, or (ii) is a singular point. The correspondence is shown schematically in Figure 9.

It is now easy to construct local models for $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$. Near any point where $(C,\boldsymbol{\zeta},\psi)$ is stable, this is a manifold with corners, the codimension of the corner being given as usual by the number of double points. In the remaining case (i), if we consider the node lying on the same component as z, the relevant gluing parameter does not have to be real, but (with suitable conventions) it should lie in the upper half-plane; otherwise, the outcome of the gluing process does not lie in the subspace $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$. Hence, the local model is a corner whose codimension is the number of double points of the image of our point under (139). Finally, in case (ii) one has two gluing parameters $\delta_\pm$ associated

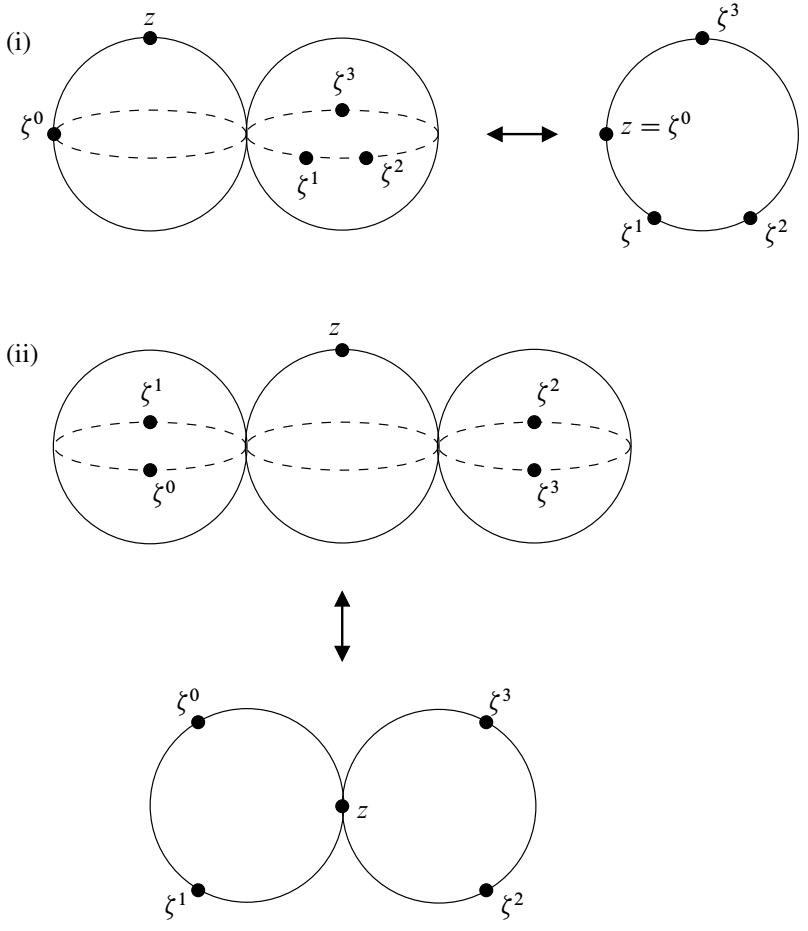

Figure 9

to the nodes which lie on the same component as z. Both of them must lie in the upper half-plane, and moreover, their product needs to be real and positive, so we get a singularity of type (8); this is particularly easy to see if one thinks of $\arg(\delta_\pm)$ as an angle of rotation in Figure 9.

As we have briefly mentioned before, even though $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$ is singular, one can define the notion of smooth function on that space, either by being smooth with respect to the local coordinates mentioned above, or (equivalently) as the restriction of a smooth function on $\overline{\mathcal{C}}^{d+1,\mathbf{p}}$. For our purposes, only a very special class of such functions will be relevant. Namely, suppose that $h\colon \mathcal{S}^{d+1,\mathbf{p}} \to \mathbb{R}$ is a smooth function with the following property: it admits a smooth extension $\overline{h}$ to the compactification, which vanishes to arbitrary order at special points (points at infinity and double points). From

the more concrete viewpoint of gluing parameters, this property has the following implications. First, consider a single popsicle $(S, \boldsymbol{\phi})$, identified with some fibre of the universal family $\mathcal{S}^{d+1,\mathbf{p}}$, and let h_S be the restriction of h to that fibre. Take a strip-like end, written for simplicity as an inclusion $Z_\pm \hookrightarrow S$. Then, for every $C > 0$ there is a constant $D > 0$ such that on the strip-like end,

$$|h(s,t)| < De^{\pm Cs}. \tag{141}$$

In words, h decays faster than exponentially near the points at infinity. To get (141) from the given vanishing condition, one simply notices that $\exp(-\pi(\pm s \pm it))$ is part of the local coordinate system on $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$ centered at a point at infinity; therefore, h grows smaller than any of its powers.

Next, suppose that we have a one-parameter family of popsicles $(S_l, \boldsymbol{\phi}_l)$, obtained by gluing together two fixed popsicles $(S_\pm, \boldsymbol{\phi}_\pm)$ with gluing length $l \gg 0$. By definition, S_l contains a finite strip $[0,l] \times [0,1]$. On that strip, the functions $h_{S_l} = h|S_l$ must have the following behaviour. For any $C > 0$ there is a constant $D > 0$ (independent of l) such that

$$|h_{S_l}(s,t)| \le \min\{De^{-Cs}, De^{C(s-l)}\}. \tag{142}$$

This means that h decays faster than exponentially (uniformly in l), whether one counts from the left end or from the right end of the strip. The reason is essentially the same as in (141): in terms of (8), the point (s,t) on the neck of the surface S_l has local coordinates $\delta_- = -\exp(-\pi(s+it))$, $\delta_+ = -\exp(-\pi(l-s-it))$. We know from the assumption that locally near $\delta_- = \delta_+ = 0$, $|h|$ is less than any $|\delta_-|^C$, $|\delta_+|^C$. This yields (142) on a piece $[L, l-L] \times [0,1]$, where the constant L is independent of l. On the other hand, if one restricts $h_{S_l}(s,t)$ to $[0,L] \times [0,1]$, it converges uniformly as $l \to \infty$, and the same is true for $h_{S_l}(s-l,t)$ on $[-L,0] \times [0,1]$. Using that fact, one easily extends the given bound to the whole of $[0,l] \times [0,1]$.

Bounds similar to (141), (142) hold for the derivatives of $h_{S_l}(s,t)$. There are also analogues where one considers families of popsicles, as well as more complicated gluing processes.

Remark 6.6 We have used $\overline{\mathcal{S}}^{d+1,\mathbf{p}}$ because it is the most straightforward way of compactifying the universal family. There is at least one other possibility, which would be to consider triples $(C, \boldsymbol{\zeta} \cup \{z, \overline{z}\}, \phi)$, admitting a real involution which exchanges the two additional marked points. The compactification obtained in this way is somewhat larger, but has the advantage of being a genuine manifold with corners. If we restrict to a single popsicle $(S, \boldsymbol{\phi})$ representing a fibre of $\mathcal{S}^{d+1,\mathbf{p}}$, what this

alternative compactification does is to add an *interval at infinity* to each end, rather than a single point.

7 A priori estimates

This section collects various elementary analytic arguments. Action and energy considerations encode what, in more standard terms, is known as the technique of monotone homotopies, going back at least to Floer and Hofer [7]. We also need an elementary SFT-type convexity result, which constrains pseudo-holomorphic curves to certain parts of the target manifold; this replaces the maximum principle and Monotonicity Lemma arguments used for the same purpose by Viterbo [30].

7.1 Preliminaries

Let's start by summarizing the geometric data that enters into the formulation of our inhomogeneous $\overline{\partial}$–equations (the following considerations will be largely independent of how that data is chosen in the specific context of wrapped Fukaya categories). On the target space side, let M be a Weinstein domain, $L \subset M$ a Lagrangian submanifold satisfying (35) as well as (37), and H a function as in (34), (38). On the source side, we have a pointed disc S, which comes with strip-like ends $\boldsymbol{\epsilon}$ as well as weights $\mathbf{w}$. Besides that, S carries a one-form γ satisfying (27). By applying Stokes one sees that

$$-\textstyle\int_S d\gamma = w^0 - w^1 - \cdots - w^d. \tag{143}$$

In particular, $d\gamma = 0$ if and only if w^0 is the sum of the other weights. Finally, we have a family $\mathbf{J} = \{J_z\}$ of almost complex structures parametrized by points of S, each of which is of contact type at the boundary. Over the strip-like ends $\epsilon^k(s,t)$, this family should converge to one which depends only on k and t, and the convergence is assumed to be faster than exponential in any C^r topology. Given that, we consider solutions of (50), which (to save the reader the trouble of looking it up) means maps

$$\begin{cases} u\colon S \longrightarrow M, \\ u(\partial S) \subset L, \\ \lim_{s\to\pm\infty} u(\epsilon^k(s,\cdot)) = x^k \in \mathcal{X}_{w^k}, \\ (du - X \otimes \gamma)^{0,1} = 0. \end{cases} \tag{144}$$

If we choose local holomorphic coordinates $z = s + it$ around some point of S, and write accordingly $\gamma = p(s,t)ds + q(s,t)dt$, the $\overline{\partial}$–equation in (144) becomes

$$\partial_t u - q(s,t)X = J_{s,t}(\partial_s u - p(s,t)X). \tag{145}$$

This class of equations is compatible with certain changes of variables, given by the action of the flow ϕ of X. More precisely, if we make the ansatz

$$u(s,t) = \phi^{l(s,t)}(\widetilde{u}(s,t)), \tag{146}$$

then $\widetilde{u}$ satisfies the same equation with modified data

$$\begin{aligned} \widetilde{J}_{s,t} &= \big(\phi^{l(s,t)}\big)^* J_{s,t}, \\ \widetilde{p}(s,t) &= p(s,t) - \partial_s l(s,t), \\ \widetilde{q}(s,t) &= q(s,t) - \partial_t l(s,t). \end{aligned} \tag{147}$$

In a less coordinate-bound way, what this means is that one has replaced γ with $\widetilde{\gamma} = \gamma - dl$. Suppose that the point we are considering is an interior point of S, and that $d\gamma = 0$ in a neighbourhood of that point. If we then take l such that $dl = \gamma$, the associated transformation turns (145) into a standard pseudo-holomorphic map equation (still with almost complex structure which depends on z). Similarly, near boundary points (where $d\gamma$ always vanishes, by assumption) one has $\gamma = dl$, where l is constant along the boundary, and that transforms our equation into a standard pseudo-holomorphic map equation with Lagrangian boundary conditions.

In the general case where $d\gamma \neq 0$, one can still get rid of the inhomogeneous term by applying the traditional Gromov trick. Namely, consider $S \times M$ with the symplectic form $\omega - d(H\gamma) + (\textit{a large positive two-form on } S)$. There is a unique compatible almost complex structure such that the equation for pseudo-holomorphic sections $S \to S \times M$ becomes equivalent to $(du - X \otimes \gamma)^{0,1} = 0$. One consequence of this reduction, and of standard pseudo-holomorphic curve theory, is the following:

Lemma 7.1 *Two different solutions of (144) can only agree on a discrete set of points.*

7.2 Energy and action

The geometric and topological energies of a solution of (144) are defined by

$$\begin{aligned} E^{\mathrm{geom}}(u) &= \int_S \tfrac{1}{2}\|du - X \otimes \gamma\|^2 = \int_S u^*\omega - u^* dH \wedge \gamma, \\ E^{\mathrm{top}}(u) &= \int_S u^*\omega - d(u^* H \cdot \gamma) = E^{\mathrm{geom}}(u) - \int_S u^* H \cdot d\gamma. \end{aligned} \tag{148}$$

Since $H > 0$ and $d\gamma \leq 0$, it follows that

$$0 \leq E^{\mathrm{geom}}(u) \leq E^{\mathrm{top}}(u). \tag{149}$$

Moreover, the first inequality is an equality if and only if $du = X \otimes \gamma$ (see Section 8.2 for further discussion of these special solutions), while the second one is an equality

if and only if $d\gamma = 0$. On the other hand, the topological energy can be computed in terms of action differences by a direct application of Stokes's theorem. Namely, define the action of $x \in \mathcal{X}_w$ to be

$$A_{wH}(x) = \int_0^1 -x^*\theta + wH(x(t))\,dt + h(x(1)) - h(x(0)), \tag{150}$$

where h is as in (35). Then, for any solution of (144) with limits x^k,

$$E^{\mathrm{top}}(u) = A_{w^0 H}(x^0) - \sum_{k=1}^{d} A_{w^k H}(x^k). \tag{151}$$

7.3 A convexity argument

Suppose that we have a manifold with boundary M, together with a one-form θ such that $d\theta = \omega$ is symplectic, but where now the Liouville field Z points strictly inwards along ∂M. This is sometimes called a *concave* contact type boundary, and the natural completion of such a manifold is to attach the small part of the symplectization $(0, 1] \times \partial M$ to the boundary. Let $L \subset M$ be a Lagrangian submanifold satisfying (35), where we assume additionally that

$$h|\partial L = 0. \tag{152}$$

Suppose also that we have a function H satisfying the analogue of (34). Finally, there is a natural notion of almost complex structure which is of contact type on ∂M, analogous to the one in Section 3.2.

Take a connected compact Riemann surface S with corners. The corners divide ∂S into pieces which are closed intervals or circles. In our case, each piece should carry an n or l label, with the property that two pieces ending at the same corner must be labeled differently. We denote by $\partial S = \partial_n S \cup \partial_l S$ the resulting decomposition, and assume that $\partial_n S \neq \varnothing$. Suppose also that our surface comes with some $\gamma \in \Omega^1(S)$, satisfying $\gamma|\partial_l S = 0$ and $d\gamma \leq 0$. Additionally, we want to have a family $\mathbf{J}$ of almost complex structures, parametrized by points of S, which are of contact type on ∂M. Given this, consider the equation

$$\begin{cases} u\colon S \longrightarrow M, \\ u(\partial_l S) \subset L, \\ u(\partial_n S) \subset \partial M, \\ (du - X \otimes \gamma)^{0,1} = 0. \end{cases} \tag{153}$$

Lemma 7.2 *There are no solutions of (153) unless $d\gamma = 0$. Even then, all solutions must satisfy $du = X \otimes \gamma$ and $u(S) \subset \partial M$.*

Proof We apply Stokes to the integral defining the topological energy (148):

$$E^{\mathrm{top}}(u) = \int_{\partial S} u^*\theta - u^* H \cdot \gamma. \tag{154}$$

By assumption, we have $H = \theta(X) = 1$ along ∂M, and moreover $\theta|L = dh$, where h vanishes along ∂L. Using this, one can rewrite (154) as

$$\begin{aligned} E^{\mathrm{top}}(u) &= \int_{\partial_n S} \theta \circ (du - X \otimes \gamma) \\ &= \int_{\partial_n S} (\theta \circ J_z) \circ (du - X \otimes \gamma) \circ (-j). \end{aligned} \tag{155}$$

We know that along ∂M, $\theta \circ J_z$ is the radial one-form (written as dr in terms of the coordinates on the cone). In particular, since $X|\partial M$ is the Reeb flow, $(\theta \circ J_z)(X) = 0$. If ξ is a vector tangent to $\partial_n S$, and positive with respect to the boundary orientation, then $j\xi$ points inwards, hence $du(j\xi)$ does not point outwards along ∂M, whence $(\theta \circ J_z \circ du)(j\xi) \geq 0$. Integrating that, we find that $E^{\mathrm{top}}(u) \leq 0$. This contradicts (149) and the nonnegativity of the geometric energy, unless $E^{\mathrm{geom}} = 0$ and $E^{\mathrm{top}} = E^{\mathrm{geom}}$, which means that $du = X \otimes \gamma$ and $d\gamma = 0$. Since du is always a multiple of X, it follows that the image of u lies inside a single orbit of that vector field, which in turn must be contained in ∂M since $\partial_n S \neq \varnothing$. □

Note that Lemma 7.2 still holds if we replace H by cH for some constant $c > 0$, since that can be compensated by rescaling γ. This will be occasionally useful.

7.4 Solutions escaping to infinity

We now return to the original setup of equation (144), and modify it by completing all the target space structures. This means that we consider

$$\begin{cases} u\colon S \longrightarrow \widehat{M}, \\ u(\partial S) \subset \widehat{L}, \\ \lim_{s \to \pm\infty} u(\epsilon^k(s, \cdot)) = x^k \in \mathcal{X}_{w^k}, \\ (du - \widehat{X} \otimes \gamma)^{0,1} = 0. \end{cases} \tag{156}$$

Here, $\widehat{X}$ is the Hamiltonian vector field of the natural extension $\widehat{H}$ from (34). Similarly, we use the extended almost complex structures $\widehat{J}_z$ which are of contact type on the infinite cone. By assumption (37), all the possible limits x^k still lie in $M \setminus \partial M \subset \widehat{M}$.

The infinite cone $[1,\infty)\times\partial M$ is a manifold with a concave contact type boundary. Suppose first that u intersects $\{1\}\times\partial M$ transversally. Then, one can apply Lemma 7.2 to the part of u that gets mapped to the infinite cone. Note that the additional condition (152) holds automatically here, since $\widehat{\theta}$ vanishes along $[1,\infty)\times\partial L$. The outcome is that u does not actually penetrate into the interior of the cone, hence stays in $M\subset\widehat{M}$ after all. To remove the technical transversality assumption, take some $\epsilon\geq 0$ such that u intersects $\{1+\epsilon\}\times\partial M$ transversally. Lemma 7.2 (or rather, the slight generalization mentioned at the end of Section 7.3) applies to $[1+\epsilon,\infty)\times\partial M$, showing that the image of u remains inside $M\cup([1,1+\epsilon]\times\partial M)$. Since ϵ can be chosen arbitrarily small, we then come to the same conclusion as before.

To see the specific implications of this, suppose that we have a one-dimensional regular moduli space $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ and a point in that space, represented by a triple $(S_0,\boldsymbol{\phi}_0,u_0)$, such that $u_0^{-1}(\partial M)\neq\varnothing$. Because of regularity and the dimension assumption, this point sits in a family $(S_r,\boldsymbol{\phi}_r,u_r)$, where a priori the u_r could leave M and penetrate into $(1,\infty)\times\partial M\subset\widehat{M}$. However, the argument above shows that this does not happen (the family of surfaces $u_r(S_r)$ becomes tangent to ∂M at $r=0$, but does not actually cross the boundary). Hence, $(S_0,\boldsymbol{\phi}_0,u_0)$ is not a boundary point of the original moduli space.

7.5 Solutions escaping from a subdomain

On the target space side, let's suppose that we have M together with a Liouville subdomain M^{in} and a Lagrangian submanifold L, as in Section 4.1, and assuming that (37) holds for both $L\subset M$ and $L^{\mathrm{in}}\subset M^{\mathrm{in}}$. Take a Hamiltonian H of the form (87). In particular, for each w we have the splitting of the set of integer X–chords of length w into ones lying in the interior and exterior of the dividing hypersurface, as in (90).

Lemma 7.3 *There is a constant ν such that for all $w\geq\nu$ and all $x\in\mathcal{X}_w^{\mathrm{out}}$, the action $A_{wH}(x)$ is positive.*

Proof By assumption (80), we can extend $h|L^{\mathrm{out}}$ to a function defined on the whole of M^{out}, which vanishes in a neighbourhood of $\partial M^{\mathrm{out}}=\partial M^{\mathrm{in}}\cup\partial M$. Note that on both hypersurfaces we have $H=1$, with nontrivial derivative in transverse direction. Hence, we can choose our neighbourhood to be bounded by level sets of H on either side, which means that it will be invariant under the flow of X.

Take $x\in\mathcal{X}_w^{\mathrm{out}}$, and write the action as

$$A_{wH}(x)=\int_0^1 w(H(x(t))-dH(Z)_{x(t)})\,dt+h(x(1))-h(x(0)). \tag{157}$$

If x lies in either of the two neighbourhoods defined above, the h term in (157) vanishes, and the integral is over a positive function by (87), which yields the desired result for any w. On the other hand, outside those neighbourhoods $H - dH(Z)$ is bounded below by some positive constant, and of course h is itself bounded, which means that (157) will necessarily be positive if w is sufficiently large. □

Consider the rescaled family H^ρ as in (94), whose Hamiltonian vector field X^ρ satisfies (95). The crucial point about this rescaling process is that it affects the action values on the inside and outside differently. Namely, for x^ρ as in (96),

$$A_{wH^\rho}(x^\rho) = \begin{cases} A_{wH}(x) & \text{if } x \in \mathcal{X}^{\mathrm{out}}_w, \\ \rho \cdot A_{wH}(x) & \text{if } x \in \mathcal{X}^{\mathrm{in}}_w. \end{cases} \tag{158}$$

Take $(S, \boldsymbol{\epsilon}, \mathbf{w}, \gamma)$ as in Section 7.1. For the almost complex structures, we want to have the following: a family $\mathbf{J} = (J_z)$ in $\mathcal{J}(M)$ parametrized by S, which asymptotically over the ends depends only on t; a similar family $\mathbf{J}^{\mathrm{in}} = (J^{\mathrm{in}}_z)$ in $\mathcal{J}(M^{\mathrm{in}})$ of the same kind; and a third family interpolating between the two, in the sense of Section 4.3, again with the same kind of asymptotic behaviour. For any $\rho \in (0, 1]$, we then consider solutions to the parametrized equation

$$\begin{cases} u\colon S \longrightarrow M, \\ u(\partial S) \subset L^\rho, \\ \lim_{s \to \pm\infty} u(\epsilon^k(s, \cdot)) = x^{k,\rho}, \\ (du_z - X^\rho_{u(z)} \otimes \gamma_z) \circ j + J^\rho_{z,u(z)} \circ (du_z - X^\rho_{u(z)} \otimes \gamma_z) = 0. \end{cases} \tag{159}$$

Clearly, for $\rho = 1$ this reduces to (144). Moreover, if u^{in} is a solution of the corresponding equation on M^{in} (using $H^{\mathrm{in}} = H|M^{\mathrm{in}}$, and the almost complex structures $\mathbf{J}^{\mathrm{in}}$), then for each sufficiently small ρ, the rescaled map $\psi^\rho \circ u^{\mathrm{in}}$ is a solution of (159). The last observation has a partial converse:

Lemma 7.4 *Suppose that all weights are $\geq \nu$, ρ is sufficiently small, and $x^0 \in \mathcal{X}^{\mathrm{in}}_{w^0}$. Then every solution of (159) is of the form $\psi^\rho \circ u^{\mathrm{in}}$. In particular, its other asymptotics are $x^k \in \mathcal{X}^{\mathrm{in}}_{w^k}$, and the image $u(S)$ is entirely contained inside $M^{\mathrm{in},\rho}$.*

Proof Combining (149) with (151) and (158), one finds that

$$0 \leq E^{\mathrm{top}}(u) = \rho \cdot A_{wH}(x^0) - \rho \sum_{\substack{k>0 \\ x^k \in \mathcal{X}^{\mathrm{in}}_{w^k}}} A_{wH}(x^k) - \sum_{\substack{k>0 \\ x^k \in \mathcal{X}^{\mathrm{out}}_{w^k}}} A_{Hw}(x^k). \tag{160}$$

In the last term on the right hand side, all the actions are positive by Lemma 7.3, and in fact bounded below by a positive constant which is independent of ρ. The other terms are bounded by ρ times a constant, which depends only on the weights (the maximum of $|A_{wH}(x)|$ among the finite sets $\mathcal{X}_{w^0}, \dots, \mathcal{X}_{w^d}$). This leads to a contradiction for $\rho \ll 1$ unless we assume that the last sum in (160) is empty, which means that $x^k \in \mathcal{X}^{\mathrm{in}}_{w^k}$ for all k.

Having excluded the possibility that some input lies in M^{out}, we now prove that all solutions to (159) lie in $M^{\mathrm{in},\rho}$. Consider $M^{\mathrm{out}} \cup \kappa^{\mathrm{in}}([\rho, 1) \times \partial M^{\mathrm{in}})$. This has $\partial M^{\mathrm{in},\rho}$ as a concave contact type boundary (of course, there is another boundary component, namely ∂M, but that won't be relevant). Moreover, for all $\rho \ll 1$, the almost complex structures J^ρ_z are of contact type in a neighbourhood of $\partial M^{\mathrm{in},\rho} \subset \kappa^{\mathrm{in}}([\rho, 1) \times \partial M^{\mathrm{in}})$. From the previous argument, we know that all limits $x^{k,\rho}$ lie in the interior of $M^{\mathrm{in},\rho}$. Suppose first that u intersects $\partial M^{\mathrm{in},\rho}$ transversally, so that the preimage $u^{-1}(M^{\mathrm{out}} \cup \kappa^{\mathrm{in}}([\rho, 1) \times \partial M^{\mathrm{in}}))$ is a compact surface with corners. Applying Lemma 7.2 then shows that $u(S) \subset M^{\mathrm{in},\rho}$. To remove the transversality assumption, one argues as before by moving the hypersurface slightly, then taking the limit. □

8 Transversality

We now revisit in more detail the transversality claims made throughout the main part of the paper. The proofs are basically routine, in spite of the presence of a finite symmetry group (see Remark 8.10 for an informal explanation of why this group does not cause any problems).

8.1 Nondegeneracy of chords

Let M be a compact symplectic manifold possibly with boundary, $L \subset M$ a Lagrangian submanifold, and $H \in C^\infty(M, \mathbb{R})$ a function, with its associated Hamiltonian vector fields X and flow ϕ. We will be interested in the generic behaviour of integer X–chords. Here, generic is understood in the sense of Baire, with respect to perturbations of H. Actually, we will sometimes find it useful to perturb L as well, in a Hamiltonian (exact Lagrangian) way. However, as far as X–chords are concerned, passing from L to $\psi(L)$ is equivalent to going from H to $H \circ \psi^{-1}$ while keeping L fixed. Hence, all the results ultimately translate into statements about perturbing H alone, for arbitrary fixed L.

As a first step, after a perturbation of H, we may assume that: first, at any stationary point x, the derivative $D\phi^1_x$ has no eigenvalues which are roots of unity; secondly, if x is a periodic point whose minimal period is some integer $w > 0$, then $D\phi^w_x$ has X

as the only eigenvector for the eigenvalue 1, and no other eigenvalues which are roots of unity. As a consequence of these properties, if x belongs to any (possibly stationary of multiply covered) periodic orbit of integer period w, the image of

$$\mathrm{id} - D\phi_x^w \colon TM_x \longrightarrow TM_x, \tag{161}$$

is precisely the symplectic orthogonal complement of X_x. In addition, we can assume that the stationary points are disjoint from L, and that the non-stationary integer periodic orbits intersect L transversally. Fix an H with these properties. We want to know if, by a small generic Hamiltonian isotopy of L, one can make the intersections $\phi^w(L) \cap L$ transverse. By a standard Sard-theory argument, this is possible if for every $x \in \phi^w(L) \cap L$, the map

$$\begin{aligned} Def(L) &\longrightarrow TM_x/(TL_x + T(\phi^w(L))_x), \\ \xi &\longmapsto \xi_x - D\phi^w(\xi_{\phi^{-w}(x)}) \end{aligned} \tag{162}$$

is surjective. Here, $Def(L)$ is the space of infinitesimal deformations of L, seen as normal vector fields. If $x \neq \phi^{-w}(x)$, the map (162) subtracts values of ξ at two different points, so surjectivity is automatic. In the remaining case, we have $X_x \notin TL_x$ by assumption, which implies that TL_x is not contained in the symplectic orthogonal complement of X_x. From this and the previous remark concerning (161), it follows that (162) is still onto. Hence:

Lemma 8.1 *For generic H, all integer X–chords are nondegenerate.* □

Next, consider triple intersections $L \cap \phi^w(L) \cap \phi^v(L)$, for integers $0 < w < v$. In that case, the map corresponding to (162) is

$$\begin{aligned} Def(L) \oplus TM_x &\longrightarrow TM_x/TL_x \oplus TM_x/T(\phi^w(L))_x \oplus TM_x/T(\phi^v(L))_x, \\ (\xi, \delta) &\longmapsto (\xi_x + \delta, D\phi^w(\xi_{\phi^{-w}(x)}) + \delta, D\phi^v(\xi_{\phi^{-v}(x)}) + \delta). \end{aligned} \tag{163}$$

This is obviously surjective if the points $(x, \phi^w(x), \phi^v(x))$ are all different. If exactly two of them coincide, it is surjective by the previous argument. In the remaining situation, where all three points coincide, surjectivity fails. This happens whenever L intersects an integer periodic orbit of X, which doesn't happen in dimensions 4 and higher due to our previous transverse intersection assumption. The upshot is:

Lemma 8.2 *Assume that $\dim M \geq 4$. For generic H, starting points of integer X–chords are never endpoints of integral X–chords.*

8.2 Trivial solutions

Consider an equation (144), in the setup described there. A *trivial solution* of that equation is a map u such that

$$du = X \otimes \gamma. \tag{164}$$

These are precisely the solutions whose geometric energy is zero, or more intuitively, those maps which satisfy (144) for all possible J at the same time. In local coordinates as in (145), this means that

$$\partial_s u - p(s,t)X = 0, \quad \partial_t u - q(s,t)X = 0. \tag{165}$$

The simplest examples are constant maps $u(z) = x$, where $X_x = 0$.

Lemma 8.3 *Suppose that $d\gamma$ is not identically zero. The only trivial solutions are the constant ones taking values at stationary points of X.*

Proof This follows from the integrability condition for (165). Namely, taking ∇ to be any torsion-free connection on TM, we have

$$\begin{aligned} 0 &= \nabla_{\partial_s u}\partial_t u - \nabla_{\partial_t u}\partial_s u \\ &= q \cdot \nabla_{\partial_s u} X + \partial_s q \cdot X - p \cdot \nabla_{\partial_t u} X - \partial_t p \cdot X \\ &= q \cdot \nabla_{pX} X - p \cdot \nabla_{qX} X + (\partial_s q - \partial_t p) \cdot X = d\gamma(\partial_s, \partial_t) \cdot X. \end{aligned} \tag{166}$$

Since $d\gamma$ is not identically zero by assumption, there is some $z \in S$ such that $u(z)$ is a stationary point of X. But then, (164) and elementary ODE theory show that u must be constant. □

Lemma 8.4 *Suppose that $d\gamma = 0$, and that H satisfies (39). The only trivial solutions are the ones with domain $S \cong Z$.*

Proof Since S is simply-connected we can write $\gamma = dl$, and then integrate (164), which yields

$$u(z) = \phi^{l(z)}(x) \tag{167}$$

for some x which is independent of z. Equivalently, one could view this as the result of applying a transformation (146) globally, and getting a $\widetilde{u}$ which is a holomorphic curve with energy zero. The function l is locally constant on ∂S, and over the strip-like ends it satisfies

$$l(\epsilon^k(s,t)) = \begin{cases} w^0 t & k = 0, \\ w^1 + \cdots + w^{k-1} + w^k t & k > 0. \end{cases} \tag{168}$$

For the limits $\mathbf{x} = \{x^k\}$ this means that

$$\begin{aligned} x^1(0) &= x^0(0), \\ x^2(0) &= x^1(1) = \phi^{w_1}(x^1(0)), \\ &\vdots \\ x^d(0) &= x^{d-1}(1) = \phi^{w_{d-1}}(x^{d-1}(0)), \\ x^1(1) &= x^d(1). \end{aligned} \tag{169}$$

In particular, if $d > 1$ then $x^2(0)$ would be both an endpoint and a starting point of an integral X–chord, which contradicts (39). □

Lemma 8.5 *Suppose that $d\gamma = 0$, and let u be a solution of (144) which is not a trivial solution. Then the set of points where $du = X \otimes \gamma$ holds is a discrete subset of S.*

Proof Locally near each point of S we can apply a transformation (146), which takes (144) into an ordinary pseudo-holomorphic curve equation, and transforms trivial solution into constant maps. Using standard pseudo-holomorphic curve techniques (the Carleman similarity principle; see Floer, Hofer and Salamon [8, Corollary 2.3] for interior points, and Oh [25, Corollary 2.2] for boundary points), one sees that the subset of U consisting of points at which $du - X \otimes \gamma$ vanishes to infinite order is both open and closed. The rest is straightforward. □

Lemma 8.6 *Suppose that $d\gamma \neq 0$, and that H satisfies (39). Let $U \subset S$ be a connected open subset, containing ∂S as well as neighbourhoods of all the points at infinity, and such that $d\gamma|U = 0$ (such a subset always exists, because of the assumptions on γ). Then, for any solution u of (144), the set of points in U where $du = X \otimes \gamma$ holds is a discrete subset of U.*

Proof Suppose that the contrary is true. Applying the same argument as in Lemma 8.5, one then sees that $du = X \otimes \gamma$ holds on all of U. As in Lemma 8.4, it follows that the limits $\mathbf{x}$ of our map satisfy $x^1(0) = x^0(0)$. In view of (143), the assumption $d\gamma \neq 0$ implies that $w^0 > w^1$. Therefore $x^1(1) = \phi^{w^1}(x^1(0)) = \phi^{w^1 - w^0}(x^0(1))$ is both starting point and endpoint of an integral X–chord, contradicting the other assumption. □

8.3 Linearization

We now consider the linearization of (144) with respect to infinitesimal changes of $\mathbf{J}$ as well as as u. This is given by an operator

$$\begin{aligned} &\mathcal{Y} \oplus W^{1,r}(u^*TM) \longrightarrow L^r(\mathrm{Hom}^{0,1}(TS, u^*TM)), \\ &\qquad (\mathbf{Y}, X) \longmapsto \tfrac{1}{2} Y_{z,u(z)} \circ (du_z - X_{u(z)} \otimes \gamma_z) \circ j_z + D_u X. \end{aligned} \tag{170}$$

Here, $\mathcal{Y}$ is the space of infinitesimal deformations of our given family of almost complex structure, which have superexponential decay over the ends of S. The second component is the standard linearized $\overline{\partial}$–operator D_u, taking the Sobolev completion for some $r > 2$.

Lemma 8.7 *Suppose that u satisfies (144), and is not a trivial solution of that equation. Then the map (170) is surjective.*

Proof This is a standard argument. Since D_u itself is Fredholm, we only need to show that the image of (170) has zero orthogonal complement. Suppose that $T \neq 0$ lies in the kernel of the adjoint D_u^*. Then it is automatically smooth, and vanishes only on a discrete subset of points in S. If $d\gamma = 0$, one can use Lemma 8.5 to find a point $z \in S$ where both T and $du - X \otimes \gamma$ are nonzero, and a $\mathbf{Y}$ such that

$$\langle Y_{z,u(z)} \circ (du_z - X_{u(z)} \otimes \gamma_z) \circ j_z, T_z \rangle \neq 0. \tag{171}$$

Strictly speaking, our almost complex structures are constrained to be of contact type along the boundary. Hence, to ensure that (171) can be achieved, we need $u(z)$ to lie in the interior of M. But since we are assuming that (37) holds, all limits x lie in the interior, hence so does $u(z)$ outside a compact subset of S. Having found (171), one multiplies $\mathbf{Y}$ with a function supported very close to z, and gets a contradiction to the assumption that T lies in the orthogonal complement to the image of (170). In the case where $d\gamma \neq 0$, the same argument works except that one has to take z in a subset $U \subset S$ of the kind considered in Lemma 8.6. □

Now consider two solutions u, v of (144), with possibly different limits at infinity. We can then consider the linearization of the equation at both solutions, with respect to a common variation of the family of almost complex structures. In parallel with (170), this is given by an operator

$$\begin{aligned} &\mathcal{Y} \oplus W^{1,r}(u^*TM) \oplus W^{1,r}(v^*TM) \longrightarrow \\ &\qquad \to L^r(\mathrm{Hom}^{0,1}(TS, u^*TM)) \oplus L^r(\mathrm{Hom}^{0,1}(TS, v^*TM)). \end{aligned} \tag{172}$$

One can now argue as before, varying the almost complex structure near a point $z \in S$ such that $(du - X \otimes \gamma)_z \neq 0$, $(dv - X \otimes \gamma)_z \neq 0$, and $u(z) \neq v(z)$. With the obvious exceptions, the existence of sufficiently many points with that property is ensured by Lemma 7.1. The outcome is:

Lemma 8.8 *Suppose that neither u nor v are trivial solutions, and that $u \neq v$. Then (172) is surjective.*

Addendum 8.9 Inspection of the proofs above shows that one can restrict to the subspace of those $\mathbf{Y}$ which vanish outside some open subset of M, as long as that subset contains all integer X–chords.

8.4 Application

We will now explain how these arguments apply to the moduli spaces $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$. It is a classical fact that a generic choice of $\mathbf{I}_w$ ensures that the moduli spaces of solutions to Floer's equation (53) are regular. With that done, let's turn to the case of stable weighted popsicles $(S, \boldsymbol{\phi}, \mathbf{w})$. Recall that in (50) we use families of almost complex structures obtained by taking a fixed universal choice $\mathbf{I} = \mathbf{I}_{S,\boldsymbol{\phi},\mathbf{w}}$ and then changing it in the way indicated by an infinitesimal deformation $\mathbf{K} = \mathbf{K}_{S,\boldsymbol{\phi},\mathbf{w}}$. The latter is restricted by $Sym^{\mathbf{p}}$–invariance as well as the asymptotic consistency requirement. Take a dense Banach subspace $\mathcal{K}$ of the space of all possible choices of infinitesimal deformations, and consider the universal moduli space, which is the space of solutions of (50) with $\mathbf{K}$ considered as an additional variable. The linearization of this equation is a map

$$\mathcal{K} \times T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}} \times W^{1,r}(u^*TM) \longrightarrow L^r(\mathrm{Hom}^{0,1}(TS, u^*TM)), \tag{173}$$

whose first and third components are as in (170), and where the tangent space to $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$ is taken at the point corresponding to $(S, \boldsymbol{\phi})$. The implications of asymptotic consistency were spelled out in Remark 3.3. Inspection of that discussion shows that for any given S the choice of $\mathbf{K}$ is free, subject only to the conditions of superexponential decay over the ends. It therefore follows from Lemma 8.7 (with $\mathbf{Y}$ restricted to a dense subspace of $\mathcal{Y}$, which does not affect the result) that (173) is onto. By a standard Sard–Smale argument, a generic choice of $\mathbf{K} \in \mathcal{K}$ will ensure that for every $(S, \boldsymbol{\phi}, u) \in \mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ the restriction of (173) to fixed almost complex structure, which is an operator

$$T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}} \times W^{1,r}(u^*TM) \longrightarrow L^r(\mathrm{Hom}^{0,1}(TS, u^*TM)), \tag{174}$$

is again onto. This is precisely the statement of Theorem 3.5, and the kernel of (174) is the tangent space of our moduli space. The same argument, applied to the stratification of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}$ by isotropy groups of the $Sym^{\mathbf{p}}$–action, proves Theorem 3.6.

Remark 8.10 In view of the generally problematic nature of equivariant transversality, it behoves us to explain why in this specific instance, the action of $Sym^{\mathbf{p}}$ does not pose any problems. As a toy model, suppose we have a manifold B and a vector bundle E over it, both of which come with compatible actions of a finite group Γ. Crucially, suppose that *for every point $b \in B$, the isotropy group Γ^b acts trivially on E_b*. In that case, for every given $\eta \in E_b$ one can find an equivariant section s such that $s(b) = \eta$, simply by averaging. In other words, the evaluation map on the space of equivariant sections,

$$B \times C^\infty(B, E)^\Gamma \longrightarrow E \tag{175}$$

is a submersion. Therefore, a generic $s \in C^\infty(B, E)^\Gamma$ is transverse to the zero-section. Returning to our real-world case, the assumption on isotropy groups corresponds to the fact that if $(S, \boldsymbol{\phi})$ lies in the subset of $\mathcal{R}^{d+1,\mathbf{p}}$ fixed by some subgroup of $Sym^{\mathbf{p}}$, then that subgroup acts trivially on the range of (174) for any map $u\colon S \to M$ satisfying (144).

To conclude, we consider parametrized moduli spaces as defined in (101). As before, we can allow variations of the associated interpolating family of almost complex structures (taking suitable technical precautions, so that the variations form a Banach space), leading to an infinite-dimensional universal moduli space, which will be a smooth Banach manifold with boundary. Moreover, projection to the parameter is a submersion from the universal moduli space to $(0, 1]$. This follows from Lemma 8.7 and Addendum 8.9, where the latter makes up for the restrictions imposed on the behaviour of interpolating families for small ρ. As before, this implies regularity of the actual parameter moduli space for generic choices of the interpolating family, which is the first part of Theorem 4.5.

Now consider the product of two parametrized moduli spaces, and its universal version (where the almost complex structures on both factors are varied in the same way). Lemma 8.8 shows that the universal space is smooth, and that the projection map to $(0, 1]^2$ is a submersion away from the diagonal. This implies the second part of Theorem 4.5, and the final part is similar. Of course, these arguments strictly speaking only apply in the range $d + |F| \geq 2$ where the underlying popsicles are stable, but the corresponding results in ordinary Floer theory are well-known.

9 Signs

The issues involved in constructing orientations of moduli spaces in Floer type theories, and the role played by *Pin* structures in the open string case, are well-understood (see

Floer–Hofer [6], Fukaya–Oh–Ohta–Ono [13] and de Silva [28]). Therefore, we only give an outline the general theory, and then concentrate on two phenomena where contributions coming from different points in the moduli space cancel. The first of these comes from the action of $Sym^{\mathbf{p}}$ on popsicle moduli spaces. The second one occurs in the context of cascades, where the same boundary component appears several times, see Example 2.10. We will also discuss the origin of the signs appearing in various of our formulae. This mainly comes down to clarifying the conventions, and the choices of orientations for $\mathcal{R}^{d+1,\mathbf{p}}$.

9.1 Degrees and orientation spaces

Let $L \subset M$ be a Lagrangian submanifold satisfying (40). One can then associate to any $x \in \mathcal{X}_w$ a Maslov index $\deg(x) \in \mathbb{Z}$. This is easiest to explain if one assumes that $[x] \in \pi_1(M, L)$ is trivial, which is when x can be filled in by a half-disc. More precisely, we want to think of this half-disc as a map $w\colon H \to M$, where $H = \mathbb{R} \times [0, \infty) \subset \mathbb{C}$ is the upper half-plane, satisfying $w(\partial H) \subset L$ and $\lim_{s\to-\infty} w(se^{-i\pi t}) = x(t)$. Such a w comes with a linear $\overline{\partial}$–operator D_w, and one defines

$$\deg(x) = \operatorname{index} D_w. \tag{176}$$

Vanishing of $2c_1(M, L)$ ensures that this integer is independent of the choice of w. In fact, one can allow higher genus surfaces as well and thereby extend the treatment to the case where only the homology class $[x] \in H_1(M, L)$ is trivial. However, the answer in full generality is not as canonical and requires a slightly different approach. Fix a compatible almost complex structure on M, and let

$$\mathcal{K}^{-2} = (\Lambda^{\mathrm{top}}_{\mathbb{C}} TM)^{\otimes 2} \tag{177}$$

be the bundle with first Chern class $2c_1(M)$. Along L this bundle has a preferred trivialization, given by taking an orthonormal basis $\{\xi_k\}$ of TL at a point, and considering the nonzero element $(\xi_1 \wedge \cdots \wedge \xi_n)^{\otimes 2}$ in the fibre of $\mathcal{K}^{-2}$ at that point, which is independent of the choice of basis. The vanishing of $2c_1(M, L)$ allows us to equip L with a *grading*, which is an extension of the given trivialization to the whole of M. The choice of grading in turn singles out, for each $x \in \mathcal{X}_w$, a class of $\overline{\partial}$–operators on H which are abstract generalizations of the D_w introduced above. Those operators, called *orientation operators*, are then used to extend (176) (one reference is Seidel [27, Section 11], but the idea goes back to Floer [5]). Changing the homotopy class of the grading by some $c \in H^1(M, L; \mathbb{Z})$ will affect the degree of each x by $\langle c, [x] \rangle$, which explains why for trivial $[x]$ one can get away without choosing a grading.

Next, we want to associate to each $x \in \mathcal{X}_w$ a one-dimensional real vector space o_x, the so-called *orientation space*. If one considers only those x that can be filled in by a

half-disc w, and assumes additionally that $H^1(M;\mathbb{Z}/2) \to H^1(L;\mathbb{Z}/2)$ is onto, one can simply define

$$o_x = \det D_w \tag{178}$$

to be the determinant line of the operator D_w. The additional assumption we just made ensures that if we have another choice of half-disc $\widetilde{w}$, the loop formed by the boundaries of w and $\widetilde{w}$ is contractible in L. A choice of contraction of that loop yields an isomorphism $\det D_w \cong \det D_{\widetilde{w}}$; and since $w_2(L) = 0$, that isomorphism is independent of the specific contraction, up to multiplication with a positive constant. The general procedure is slightly different: one chooses a *Pin* structure on L, and then replaces D_w by abstract orientation operators which are compatible with that structure in a suitable way. Twisting the *Pin* structure by a real line bundle λ has the effect of changing each o_x to $o_x \otimes \lambda_{x(0)} \otimes \lambda_{x(1)}$.

We assume from now that a grading and *Pin* structure have been chosen, so that $\deg(x)$ and o_x are unambiguously defined for all x. Additionally, we need the following notation. Let ξ be a one-dimensional real vector space, and denote by $\sigma^{\pm}$ its two possible orientations. For a fixed coefficient field $\mathbb{K}$, the $\mathbb{K}$*–normalization* of ξ is the one-dimensional $\mathbb{K}$–vector space

$$|\xi|_{\mathbb{K}} = \mathbb{K}\sigma^+ \oplus \mathbb{K}\sigma^-/\mathbb{K}(\sigma^+ + \sigma^-). \tag{179}$$

By construction, $|\xi|_{\mathbb{K}}$ is identified with $\mathbb{K}$ in a way which is canonical up to sign. A choice of identification is precisely the same as an orientation of ξ. Similarly, any isomorphism between one-dimensional real vector spaces induces an isomorphism of their $\mathbb{K}$–normalizations, which is ± 1 if one chooses identifications of those spaces with $\mathbb{K}$.

9.2 Zero-dimensional moduli spaces

Consider a point $(S, \boldsymbol{\phi}, u)$ in a moduli space $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$. If we take the linearized operator D_u associated to u, and glue it at the positive ends to orientation operators for $x^1, \dots, x^d$, the outcome is an operator on H which is homotopic to an orientation operator for x^0. The gluing relation for determinant lines of $\overline{\partial}$–operators therefore yields an isomorphism

$$o_{x^o} \cong \det D_u \otimes o_{x^d} \otimes \cdots \otimes o_{x^1}. \tag{180}$$

To be more precise, one should say that by choosing the homotopy of operators to be compatible with the given *Pin* structures, one gets an isomorphism of one-dimensional vector spaces which is canonical up to multiplication by a positive constant (the same

interpretation should be applied to similar statements made later on). Assume from now on that all moduli spaces are regular, as in Theorem 3.5, and also that $d+|F|\geq 2$, which means that the underlying popsicle $(S,\boldsymbol{\phi})$ is stable. Then, because the tangent space to the moduli space is the kernel of an operator of the form (174), we have

$$\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}))_{S,\boldsymbol{\phi},u}\cong\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p}})_{S,\boldsymbol{\phi}}\otimes\det D_u, \tag{181}$$

Plugging those two relations into each other, we get an isomorphism

$$o_{S,\boldsymbol{\phi},u}\colon \Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}))_{S,\boldsymbol{\phi},u}\otimes o_{x^d}\otimes\cdots\otimes o_{x^1}\cong\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p}})_{S,\boldsymbol{\phi}}\otimes o_{x^0}. \tag{182}$$

At this point, let's choose an orientation of each moduli space $\mathcal{R}^{d+1,\mathbf{p}}$ (which is possible since they are contractible). This allows us to trivialize the $T\mathcal{R}$ term in (182). Assume additionally that the moduli space of popsicle maps is zero-dimensional, so that $\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x}))$ is canonically trivial. We then get an isomorphism

$$o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}\colon o_{x^d}\otimes\cdots\otimes o_{x^1}\cong o_{x^0}. \tag{183}$$

Since this is unique up to multiplication by a positive constant, the induced map on $\mathbb{K}$–normalizations, denoted by $|o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}|_{\mathbb{K}}$, is canonically well-defined. By definition, this is the contribution (58) of $(S,\boldsymbol{\phi},u)$ to the algebraic count of points in our moduli space.

Lemma 9.1 *Take $\sigma\in Sym^{\mathbf{p}}\subset Sym^{|F|}$, acting on the zero-dimensional moduli space $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ by permuting popsicle maps, $(S,\boldsymbol{\phi},u)\mapsto(S,\sigma(\boldsymbol{\phi}),u)$. Then, up to a positive constant,*

$$o^{\mathrm{red}}_{S,\sigma(\boldsymbol{\phi}),u}=\mathrm{sign}(\sigma)\,o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}. \tag{184}$$

This is clear from the construction above. Since $\mathcal{R}^{d+1,\mathbf{p}}$ projects to $\mathcal{R}^{d+1}$ with fibers $\mathbb{R}^F$, and $Sym^{\mathbf{p}}\subset Sym^F$ acts fibrewise by the standard representation of permutations on $|F|$ letters, σ preserves orientations if and only if σ is an even permutation. Hence, if we compare $(S,\sigma(\boldsymbol{\phi}),u)$ with (S,σ,u), the associated maps (180) agree, but the trivializations $\det D_u\cong\mathbb{R}$ obtained from (181) differ by $\mathrm{sign}(\sigma)$, and so do the maps (183).

9.3 One-dimensional moduli spaces

Consider a point in the boundary of a one-dimensional moduli space $\overline{\mathcal{R}}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$, represented by a broken map with two components $(S_\pm,\boldsymbol{\phi}_\pm,u_\pm)$. As before, we assume stability of the underlying popsicles, which means that $(S_\pm,\boldsymbol{\phi}_\pm,u_\pm)\in\mathcal{R}^{d_\pm+1,\mathbf{p}_\pm,\mathbf{w}_\pm}(\mathbf{x}_\pm)$ with $d_\pm+|F_\pm|\geq 2$. Gluing them together, one gets a nearby point $(S,\boldsymbol{\phi},u)$ in

the interior $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$. The gluing parameter yields a preferred local orientation for the moduli space (pointing towards the boundary), hence an isomorphism $\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})) \cong \mathbb{R}$ at $(S,\boldsymbol{\phi},u)$. With that in mind, and an additional choice an orientation of $\mathcal{R}^{d+1,\mathbf{p}}$, one can turn $o_{S,\boldsymbol{\phi},u}$ into a map $o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}$ of the same form as in (183). The elementary fact that boundary points come in pairs implies that the sum of the normalizations $|o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}|_{\mathbb{K}}$ over all $(S_\pm,\boldsymbol{\phi}_\pm,u_\pm)$ vanishes. Of course, this is not strictly speaking true, since we are ignoring the case where a Floer trajectory bubbles off, but we will make up for that omission later (Section 9.5).

On the other hand, we shall presently see that $o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}$ can be decomposed as follows:

(185)
$$\begin{aligned}
&o_{x^d} \otimes \cdots \otimes o_{x^1} \\
&\cong o_{x^d} \otimes \cdots \otimes o_{x^{i+d_-}} \otimes \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_-+1,\mathbf{p}_-})_{S_-,\boldsymbol{\phi}_-} \otimes o_{x^{new}} \otimes o_{x^{i-1}} \otimes \cdots \otimes o_{x^1} \\
&\cong \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_-+1,\mathbf{p}_-})_{S_-,\boldsymbol{\phi}_-} \otimes o_{x^d} \otimes \cdots \otimes o_{x^{new}} \otimes \cdots \otimes o_{x^1} \\
&\cong \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_++1,\mathbf{p}_+})_{S_+,\boldsymbol{\phi}_+} \otimes \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_-+1,\mathbf{p}_-})_{S_-,\boldsymbol{\phi}_-} \otimes o_{x^0} \\
&\cong \Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p}})_{S,\boldsymbol{\phi}} \otimes o_{x^0} \\
&\cong o_{x^0}.
\end{aligned}$$

The first step consists of applying $o_{S_-,\boldsymbol{\phi}_-,u_-}$, keeping in mind that $\mathcal{R}^{d_-+1,\mathbf{p}_-,\mathbf{w}_-}(\mathbf{x}_-)$ is necessarily of dimension zero; here, $x^{new} = x^0_- = x^i_+$ is the X–chord at the ends where our two components are glued together. The second step is just a permutation of tensor factors (with suitable Koszul signs given by dimensions and degrees; this is natural, since each factor encodes orientations of a vector space). Thirdly, we apply $o_{S_+,\boldsymbol{\phi}_+,u_+}$, and exchange the first two factors of the resulting expression. Next, recall that $(S,\boldsymbol{\phi})$ is a point of $\mathcal{R}^{d+1,\mathbf{p}}$ close to the boundary stratum to which the broken popsicle $(S_\pm,\boldsymbol{\phi}_\pm)$ belongs. Using any collar neighbourhood of that stratum, we get an isomorphism

$$\Lambda^{\mathrm{top}}(T\mathcal{R}^{d+1,\mathbf{p}})_{S,\boldsymbol{\phi}} \cong \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_++1,\mathbf{p}_+})_{S_+,\boldsymbol{\phi}_+} \otimes \Lambda^{\mathrm{top}}(T\mathcal{R}^{d_-+1,\mathbf{p}_-})_{S_-,\boldsymbol{\phi}_-}, \tag{186}$$

which is used in the fourth step in (185). The fifth and last one is just to put in the chosen orientation of $\mathcal{R}^{d+1,\mathbf{p}}$. After unwinding the definition of (182), the fact that the sequence of operations in Equation (185) agrees with the original definition of $o^{\mathrm{red}}_{S,\boldsymbol{\phi},u}$ reduces to applying the gluing isomorphism $\det D_u \cong \det D_{u_+} \otimes \det D_{u_-}$ and using its associativity properties.

From now on, we always assume that $Sym^{\mathbf{p}}$ is trivial, and accordingly think of $F \subset \{1,\dots,d\}$ as in Section 3.6. The same property then holds for $\mathbf{p}_-$, but not necessarily for $\mathbf{p}_+$, compare Examples 2.7(ii). Given $\sigma \in Sym^{\mathbf{p}_+}$, let $(\widetilde{S},\widetilde{\boldsymbol{\phi}},\widetilde{u})$ be the element

of $\mathcal{R}^{d+1,\mathbf{p},\mathbf{w}}(\mathbf{x})$ obtained by gluing together $(S_+, \sigma(\boldsymbol{\phi}_+), u_+)$ and $(S_-, \boldsymbol{\phi}_-, u_-)$. We can see the effect of σ on (185) by looking at the diagram of isomorphisms

(187)

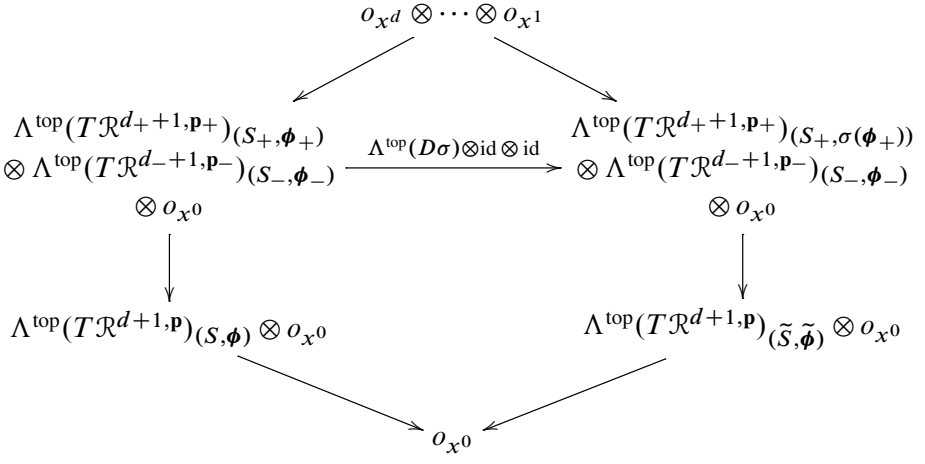

Going down the left or right sides yields the contributions of $(S, \boldsymbol{\phi}, u)$ and $(\widetilde{S}, \widetilde{\boldsymbol{\phi}}, \widetilde{u})$ to the total count of boundary points, respectively. The upper triangle of the diagram commutes. The lower pentagon depends only on the orientations of the various popsicle moduli spaces, and can be described as follows. On the left side, we compare some orientation of the boundary stratum (186) with the chosen orientation of the interior $\mathcal{R}^{d+1,\mathbf{p}}$. On the right hand side we do the same, but where the orientation of the boundary stratum is changed by the action of σ, so the difference is $\mathrm{sign}(\sigma)$. In particular, whenever $\mathrm{sign}(\sigma) = -1$, the contributions of $(S, \boldsymbol{\phi}, u)$ and $(\widetilde{S}, \widetilde{\boldsymbol{\phi}}, \widetilde{u})$ cancel out.

With that out of the way, we turn to the case where $Sym^{\mathbf{p}\pm}$ are both trivial, whose combinatorics is captured by the notion of admissible cut (Definition 3.8). In that case, the outcome of our discussion can be summarized as follows. Given a choice of orientation of all the popsicle moduli spaces, let $(-1)^{\triangle}$ be the sign measuring the difference between the product orientation of (186) and that of $\mathcal{R}^{d+1,\mathbf{p}}$. Then what (185) says is that

$$
\begin{aligned}
& o^{\mathrm{red}}_{S,\boldsymbol{\phi},u} = (-1)^{\S}\, o^{\mathrm{red}}_{S_+,\boldsymbol{\phi}_+,u_+} \circ (\mathrm{id}^{\otimes d_+ - i} \otimes o^{\mathrm{red}}_{S_-,\boldsymbol{\phi}_-,u_-} \otimes \mathrm{id}^{\otimes i-1}), \\
& \S = \triangle + (\dim \mathcal{R}^{d_- +1,F_-})(\dim \mathcal{R}^{d_+ +1,F_+}) \\
& \qquad + (\dim \mathcal{R}^{d_- +1,F_-})(\deg x^{i+d_-} + \cdots + \deg x^{d}).
\end{aligned} \tag{188}
$$

9.4 Explicit orientations

So far, the choice of orientations of the popsicle moduli spaces has been arbitrary. For those cases where $Sym^{\mathbf{p}}$ is nontrivial, there is no need to restrict that freedom, since their contributions cancel out anyway. In contrast, for the cases with trivial symmetry, we do want to make an explicit choice, so as to get a concrete expression for $\triangle$ in (188).

The standard way of choosing an orientation of $\mathcal{R}^{d+1}$ is as follows. For every pointed disc S there is a unique isomorphism $\overline{S} \cong D$ which takes $\zeta^0, \zeta^1, \zeta^2$ to $-1, -i, +1 \in \partial S$, and hence all other points at infinity ζ^k to points in the upper semicircle. This yields an identification

$$\mathcal{R}^{d+1} \cong \{0 < t_1 < \cdots < t_{d-2} < \pi\} \subset \mathbb{R}^{d-2}, \tag{189}$$

and one pulls back the standard orientation from the right hand side. Next, consider the moduli spaces $\mathcal{R}^{d+1,F}$ for $d \geq 2$ and $F \subset \{1, \dots, d\}$. These come with diffeomorphisms to $\mathcal{R}^{d+1} \times \mathbb{R}^F$, obtained by splitting the forgetful map (14). The splitting is not canonical, but the induced orientation is well-defined (to make things precise, we should say that the identification of the fibre of the forgetful map with $\mathbb{R}^F$ is such that increasing the coordinates has the effect of moving the sprinkles away from $\zeta^0 = -\infty$; moreover, the orientation of $\mathbb{R}^F$ is the one obtained by ordering the elements of F in the obvious way). We extend this to the case where $d = 1$, $F = \{1\}$ by taking the *opposite* of the tautological orientation of $\mathcal{R}^{1,\{1\}} = \text{point}$. An elementary computation shows that with these conventions, the orientation difference associated to a given cut is

$$\begin{aligned} \triangle &= d_- d_+ + d_- i + i + 1 \\ &\quad + |F_+| d_- + |\{(k_+, k_-) \in F_+ \times F_- \,:\, \iota_+(k_+) > \iota_-(k_-)\}|. \end{aligned} \tag{190}$$

After substituting this, as well as the obvious formulae for the dimensions of the moduli spaces, into (188), one gets signs as originally stated in (62). In principle, any other choice of orientation is equally viable; the only difference is that the sign in (75) would have to be adjusted accordingly, so as to recover the A_∞–associativity relations in the desired form (78).

9.5 The unstable case

We will now describe a slightly different version of the formalism above, for $d = 1$, which is more easily extended to $F = \varnothing$. Take $(S, \boldsymbol{\phi}, u) \in \mathcal{R}^{2,\mathbf{p},\mathbf{w}}(\mathbf{x})$. As in Section 4.8 we identify $S = Z$, so that $(\boldsymbol{\phi}, u)$ are maps defined on Z, determined up to translation.

In that case, the tangent space of the moduli space can be identified with a kernel of an operator

$$\mathbb{R}^F \times W^{1,r}(u^*TM) \longrightarrow \mathbb{R} \times L^r(u^*TM). \tag{191}$$

Here, increasing any of the entries in $\boldsymbol{\sigma} = \{\sigma_f\} \in \mathbb{R}^F$ has the effect of moving the corresponding sprinkle away from $-\infty$, while leaving the map u constant. The second component of (191) is the linearization of (50), with the standard trivialization of TZ taken into account. The first component is the L^2 pairing with the infinitesimal generator of the translational $\mathbb{R}$–action, concretely

$$\begin{aligned} \mathbb{R}^F \times W^{1,r}(u^*TM) &\longrightarrow \mathbb{R}, \\ (\{\sigma_f\}, X) &\longmapsto -\textstyle\sum_f \sigma_f + \langle \partial_s u, X \rangle. \end{aligned} \tag{192}$$

From (191) we get an isomorphism

$$\Lambda^{\mathrm{top}}(T\mathcal{R}^{2,\mathbf{p},\mathbf{w}}(\mathbf{x})) \otimes o_{x^1} \cong \mathbb{R}^\vee \otimes \Lambda^{\mathrm{top}}(\mathbb{R}^F) \otimes o_{x^0}. \tag{193}$$

If we choose an ordering of F, hence an orientation of $\mathbb{R}^F$, the right hand side of (193) can be identified with o_{x^0}. Using that, one then associates to any point in a zero-dimensional moduli space a map $o^{\mathrm{red}}_{(Z,\boldsymbol{\phi},u)} : o_{x^1} \to o_{x^0}$. In the special case where $|F| = 1$, (191) takes on the form

$$\begin{pmatrix} -\mathrm{id}_{\mathbb{R}} & \langle \partial_s u, \cdot \rangle \\ D_u(\partial_s u) & D_u \end{pmatrix}, \tag{194}$$

and one can deform it through invertible operators to $-\mathrm{id}_{\mathbb{R}} \oplus D_u$. In this way, one sees that the sign contribution here agrees with the convention in Section 9.4.

Now consider a point on the boundary of a one-dimensional space $\mathcal{R}^{2,\mathbf{p},\mathbf{w}}(\mathbf{x})$. In parallel with Section 9.3 we denote the components of that point by $(Z, \boldsymbol{\phi}_\pm, u_\pm)$, and a nearby point in the interior by $(Z, \boldsymbol{\phi}, u)$. The preferred orientation at the latter point, going towards the boundary, determines a map $o^{\mathrm{red}}_{(Z,\boldsymbol{\phi},u)} \colon o_{x^1} \to o_{x^0}$. A gluing argument as in (185) shows that

$$\begin{aligned} o^{\mathrm{red}}_{(Z,\boldsymbol{\phi},u)} &= (-1)^{\#} o^{\mathrm{red}}_{(Z,\boldsymbol{\phi}_+,u_+)} \circ o^{\mathrm{red}}_{(Z,\boldsymbol{\phi}_-,u_-)}, \\ \# &= |F_+| + 1 + \{(f_+, f_-) \in F_+ \times F_- \,:\, f_+ > f_-\}. \end{aligned} \tag{195}$$

In this formula, the first summand comes from swapping the $\Lambda^{\mathrm{top}}(\mathbb{R}^{F_+})$ factor with an adjacent $\mathbb{R}^\vee$, the second one arises because the standard generator of $\mathbb{R}^\vee$ moves the popsicle to the right, whereas in the gluing process, going to the boundary moves popsicles in opposite ways, and the last one comes from our chosen ordering of F and the induced orderings of the subsets $F_\pm$. Of course, in practice only the case $|F| = 1$ is important, in which case the sign is just $(-1)^{|F_-|}$. This explains why in the two

terms in Examples 3.9(i) appear with opposite signs, thereby justifying the relevant case of (62). A similar argument explains the other instances of (61) where one of the two operations involved is $m^{1,\varnothing,\mathbf{w}}$.

9.6 Parametrized moduli spaces

We now consider the parametrized moduli spaces and moduli spaces of cascades maps, always assuming that $d = 1$. The notation will generally follow that from Section 4.8. In the parametrized case, variations of the parameter ρ have to be included into the operator whose kernel describes the tangent space to the moduli space. Hence, the analogue of (193) for a point $(\rho, Z, \boldsymbol{\phi}, u) \in \mathcal{P}^{2,F,\mathbf{w}}(\mathbf{x})$ is a canonical isomorphism

$$\Lambda^{\mathrm{top}}(T\mathcal{P}^{2,F,\mathbf{w}}(\mathbf{x})) \otimes o_{x^1} \cong \mathbb{R} \otimes \mathbb{R}^\vee \otimes \Lambda^{\mathrm{top}}(\mathbb{R}^F) \otimes o_{x^0}, \tag{196}$$

where the $\mathbb{R}$ factor is the tangent space to the parameter space $(0, 1]$.

Now suppose that we have chosen an ordering of F, as well as identifications $o_{x^1} \cong \mathbb{R}$, $o_{x^0} \cong \mathbb{R}$. After trivializing the $\mathbb{R}$ and $\mathbb{R}^\vee$ factors in (196) in the standard way, we then simply get an orientation of the parametrized moduli space. Consider a codimension one boundary stratum in the partial compactification $\overline{\mathcal{P}}^{2,F,\mathbf{w}}(\mathbf{x})$, which is a fibre product

$$\mathcal{P}^{2,F_+,\mathbf{w}_+}(x^0, x^{new}) \times_{(0,1]} \mathcal{P}^{2,F_-,\mathbf{w}_-}(x^{new}, x^1). \tag{197}$$

After taking the orderings of $F_\pm$ induced from the one on F, and choosing some identification $o_{x^{new}} \cong \mathbb{R}$, one gets orientations of both factors in (197), hence an orientation of the fibre product. Note that since $o_{x^{new}}$ appears in either factor, its trivialization does not really affect the fibre product orientation.

Remark 9.2 To explain the convention here, consider Euclidean spaces $\mathbb{R}^{n_\pm}$ with coordinates $t_\pm$. Take projections to the first coordinate, $t_{\pm,1} \colon \mathbb{R}^{n_\pm} \to \mathbb{R}$. Then, the diffeomorphism

$$\begin{aligned} \mathbb{R}^{n_+} \times_{\mathbb{R}} \mathbb{R}^{n_-} &\longrightarrow \mathbb{R}^{n_+ + n_- - 1}, \\ (t_+, t_-) &\longmapsto (t_{+,2}, \dots, t_{+,n_+}, t_-) \end{aligned} \tag{198}$$

takes the fibre product orientation to the standard orientation of $\mathbb{R}^{n_+ + n_- - 1}$. In particular, if we consider the fibre product as the boundary of the open set $\{(t_+, t_-) : t_{+,1} < t_{-,1}\} \subset \mathbb{R}^{n_+} \times \mathbb{R}^{n_-}$, the induced boundary orientation agrees with the fibre product orientation.

Again arguing in the style of (185), one finds the following:

Lemma 9.3 *The boundary orientation of the stratum (197), which is induced from that of $\mathcal{P}^{2,F,\mathbf{w}}(\mathbf{x})$, differs from the fibre product orientation by $(-1)^{\triangle}$, where*

$$\triangle = |\{(f_+, f_-) \in F_+ \times F_- \subset F \times F : f_+ > f_-\}| + |F_-|(\deg(x^{new}) - \deg(x^0)). \tag{199}$$

Now we turn to the moduli spaces of cascade maps $\mathcal{Q}^{2,F,\mathbf{w}}(\mathbf{x})$ from Section 4.8. Recall that a point in such a space consists of a finite chain of $l \geq 1$ components, each of which is of the form (109). As before, we choose an ordering of F, equip the subsets F_j with the induced orderings, and consider the associated isomorphisms

$$\Lambda^{\mathrm{top}}(T\mathcal{P}^{2,F_j,(w_{j-1},w_j)}(x_{j-1}, x_j)) \otimes o_{x_j} \longrightarrow o_{x_{j-1}}. \tag{200}$$

Now $\mathcal{Q}^{2,F,\mathbf{w}}(\mathbf{x})$ is locally equal to the product of the $\mathcal{P}^{2,F_j,(w_{j-1},w_j)}(x_{j-1}, x_j)$. Hence, by composing all the (200), one gets an isomorphism $\Lambda^{\mathrm{top}}(T\mathcal{Q}^{2,F,\mathbf{w}}(\mathbf{x})) \otimes o_{x^1} \to o_{x^0}$. We twist this isomorphism by $(-1)^{\star}$, where

$$\begin{aligned} \star = {} & |\{(f_+, f_-) \in F_j \times F_k \subset F \times F \,:\, j > k,\ f_+ > f_-\}| \\ & + \sum_j |F_j|\,(\deg(x_{j-1}) - \deg(x_0)) \ + \ l, \end{aligned} \tag{201}$$

and denote the outcome by $o^{\mathrm{red}}_{\{(\rho_j, Z, \boldsymbol{\phi}_j, u_j)\}}$. Suppose that we change the ordering of F by swapping two elements. If both of these lie in the same subset F_j, the associated isomorphism (200) changes sign by definition, while in the other case, each component still gives rise to the same map, but $\star$ changes. In both cases, the result is that the sign of $o^{\mathrm{red}}_{\{(\rho_j, Z, \boldsymbol{\phi}_j, u_j)\}}$ gets flipped, which explains the cancellations which result in Lemma 4.11.

We now concentrate on the case where $|F| \leq 1$, when the choice of ordering of F, hence also the first line in (201), become trivial. For simplicity, assume as before that we have trivialized o_{x^0} and o_{x^1}, so that the moduli spaces of cascade maps connecting these two chords are canonically oriented. Suppose that $\mathcal{Q}^{2,F,\mathbf{w}}(\mathbf{x})$ is one-dimensional, and consider a boundary point in its partial compactification, where two components $\rho_{j-1} = \rho_j$ occur for the same parameter value. As in the model case of Example 2.10, such points can arise in two different ways: either $\bar{E} = \varnothing$, which means that equality of parameter values is achieved simply as a degenerate limiting case of the usual inequality $\rho_{j-1} < \rho_j$; or else $\bar{E} = E$, which means that our boundary point is the limit of a sequence in which one component is divided into two parts by (the standard Floer-theoretic) bubbling. In the first case, the orientation of the boundary point is just

the fibre product orientation of $\mathcal{Q}^{2,F_+,\mathbf{w}_+}(\mathbf{x}_+) \times_{(0,1]} \times \mathcal{Q}^{2,F_-,\mathbf{w}_-}(\mathbf{x}_-)$, where

$$\begin{aligned} F_+ &= F_1 \cup \cdots \cup \cdots F_{j-1}, & F_- &= F_j \cup \cdots \cup F_l, \\ \mathbf{w}_+ &= (w^0, w^{new}), & \mathbf{w}_- &= (w^{new}, w^1), \\ \mathbf{x}_+ &= (x^0, x^{new}), & \mathbf{x}_- &= (x^{new}, x^1) \\ \text{with } w^{new} &= w^0 - |F_+| = w^1 + |F_-|, & x^{new} &= x_j. \end{aligned} \tag{202}$$

Let's temporarily forget the signs (201) and use the untwisted orientations of the moduli spaces of cascade maps. Then, Lemma 9.3 would tell us that the boundary orientation in the first and second cases differ from each other by $|F_-|(\deg(x^{new} - \deg(x^0))$. Turning the twist by (201) back on cancels out that difference, up to an overall -1 which comes from the difference in l (obvious, since bubbling raises the number of components by one). Hence, we have shown that with the conventions used in defining $o^{\mathrm{red}}_{\{(\rho_j, Z, \boldsymbol{\phi}_j, u_j)\}}$, the contributions from the two boundary points to the standard counting argument will vanish, as claimed in Section 4.8.

The remaining boundary points of $\overline{\mathcal{Q}}^{2,F,\mathbf{w}}(\mathbf{x})$ are those where $\rho_l = 1$. In we take $j = l$ and use notation as in (202), the boundary orientation differs from that of the product $\mathcal{Q}^{2,F_+,\mathbf{w}_+}(\mathbf{x}_+) \times \mathcal{R}^{2,F_-,\mathbf{w}_-}(\mathbf{x}_-)$ by $(-1)^\times$, where the sign comes from (201) and is

$$\times = |F_-|(\deg(x^{new}) - \deg(x^0)) + 1. \tag{203}$$

The other relevant geometric phenomenon are the ends of the partial compactification, which are modelled on $(0, \delta] \times \mathcal{R}^{2,F_+,\mathbf{w}_+}(\mathbf{x})^{\mathrm{in}} \times \mathcal{Q}^{2,F_-,\mathbf{w}_-}(\mathbf{x}_-)$ with data as in (202) for $j = 1$. In this case, pointing towards infinity along such an end corresponds to decreasing the parameter. Hence, if we cut off the end, the boundary orientation differs from the product orientation by $(-1)^{\times+1}$. Specializing to $F = \varnothing$ and $F = \{1\}$, and taking into account the fact that the component of each boundary point belong to zero-dimensional moduli spaces, then yields (113) and (116), respectively.

References

[1] **A Abbondandolo**, **M Schwarz**, *On the Floer homology of cotangent bundles*, Comm. Pure Appl. Math. 59 (2006) 254–316 MR2190223

[2] **M Abouzaid**, *Homogeneous coordinate rings and mirror symmetry for toric varieties*, Geom. Topol. 10 (2006) 1097–1157 MR2240909

[3] **K Cieliebak**, **A Floer**, **H Hofer**, *Symplectic homology II: A general construction*, Math. Z. 218 (1995) 103–122 MR1312580

[4] **K Cieliebak**, **A Floer**, **H Hofer**, **K Wysocki**, *Applications of symplectic homology II: Stability of the action spectrum*, Math. Z. 223 (1996) 27–45 MR1408861

[5] **A Floer**, *Symplectic fixed points and holomorphic spheres*, Comm. Math. Phys. 120 (1989) 575–611 MR987770

[6] **A Floer**, **H Hofer**, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z. 212 (1993) 13–38 MR1200162

[7] **A Floer**, **H Hofer**, *Symplectic homology I: Open sets in $\mathbb{C}^n$*, Math. Z. 215 (1994) 37–88 MR1254813

[8] **A Floer**, **H Hofer**, **D Salamon**, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. 80 (1995) 251–292 MR1360618

[9] **A Floer**, **H Hofer**, **K Wysocki**, *Applications of symplectic homology I*, Math. Z. 217 (1994) 577–606 MR1306027

[10] **U Frauenfelder**, *The Arnold–Givental conjecture and moment Floer homology*, Int. Math. Res. Not. (2004) 2179–2269 MR2076142

[11] **K Fukaya**, *Morse homotopy, A^∞–category, and Floer homologies*, from: "Proceedings of GARC Workshop on Geometry and Topology '93 (Seoul, 1993)", Lecture Notes Ser. 18, Seoul Nat. Univ., Seoul (1993) 1–102 MR1270931

[12] **K Fukaya**, **Y-G Oh**, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, Asian J. Math. 1 (1997) 96–180 MR1480992

[13] **K Fukaya**, **Y-G Oh**, **H Ohta**, **K Ono**, *Lagrangian intersection Floer theory: anomaly and obstruction*, AMS/IP Studies in Advanced Mathematics 46, American Mathematical Society, Providence, RI (2009) MR2553465

[14] **K Fukaya**, **P Seidel**, **I Smith**, *Exact Lagrangian submanifolds in simply-connected cotangent bundles*, Invent. Math. 172 (2008) 1–27 MR2385665

[15] **J Harris**, **I Morrison**, *Moduli of curves*, Graduate Texts in Mathematics 187, Springer, New York (1998) MR1631825

[16] **A Hatcher**, *Algebraic topology*, Cambridge University Press, Cambridge (2002) MR1867354

[17] **H Hofer**, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, Invent. Math. 114 (1993) 515–563 MR1244912

[18] **D Joyce**, *Kuranishi bordism and Kuranishi homology* arXiv:0707.3572

[19] **B Keller**, *Introduction to A–infinity algebras and modules*, Homology Homotopy Appl. 3 (2001) 1–35 MR1854636

[20] **M Khovanov**, **P Seidel**, *Quivers, Floer cohomology, and braid group actions*, J. Amer. Math. Soc. 15 (2002) 203–271 MR1862802

[21] **F F Knudsen**, *The projectivity of the moduli space of stable curves II: The stacks $M_{g,n}$*, Math. Scand. 52 (1983) 161–199 MR702953

[22] **M Kontsevich**, **Y Manin**, *Gromov–Witten classes, quantum cohomology, and enumerative geometry*, Comm. Math. Phys. 164 (1994) 525–562 MR1291244

[23] **G Mikhalkin**, *Decomposition into pairs-of-pants for complex algebraic hypersurfaces*, Topology 43 (2004) 1035–1065 MR2079993

[24] **A Oancea**, *A survey of Floer homology for manifolds with contact type boundary or symplectic homology*, from: "Symplectic geometry and Floer homology", Ensaios Mat. 7, Soc. Brasil. Mat., Rio de Janeiro (2004) 51–91 MR2100955

[25] **Y-G Oh**, *On the structure of pseudo-holomorphic discs with totally real boundary conditions*, J. Geom. Anal. 7 (1997) 305–327 MR1646780

[26] **P Seidel**, *A biased view of symplectic cohomology*, from: "Current developments in mathematics, 2006", Int. Press, Somerville, MA (2008) 211–253 MR2459307

[27] **P Seidel**, *Fukaya categories and Picard–Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich (2008) MR2441780

[28] **V de Silva**, *The index of a family of pseudo-holomorphic disks*, preprint (2000)

[29] **C Viterbo**, *Functors and computations in Floer homology with applications II*, preprint (1996)

[30] **C Viterbo**, *Functors and computations in Floer homology with applications I*, Geom. Funct. Anal. 9 (1999) 985–1033 MR1726235

Department of Mathematics, Massachusetts Institute of Technology
77 Massachusetts Ave, Cambridge MA 02139, USA

abouzaid@math.mit.edu, seidel@math.mit.edu

Proposed: Yasha Eliashberg Received: 8 January 2008
Seconded: Simon Donaldson, Leonid Polterovich Revised: 15 August 2009