A Morse-theoretic approach to family Floer homology

Keeley Hoek

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
in Mathematics

Harvard University
Cambridge, Massachusetts

April 2025

Typeset in Palatino by $\TeX$ and $\LaTeX\,2_\varepsilon$.

Dissertation Advisor: Denis Auroux *Author: Keeley Hoek*

A Morse-theoretic approach to family Floer homology

Abstract

This dissertation introduces a new model of the family Floer approach to Kontsevich's homological mirror symmetry conjecture constructed via Morse theoretic technology. Homological mirror symmetry (HMS) asserts a derived equivalence between the Fukaya category of a symplectic manifold X and the category of coherent sheaves on its mirror $X^\vee$. On the other hand, the family Floer program gives a modern reinterpretation of the construction of a Strominger–Yau–Zaslow (SYZ) mirror, and this mirror space typically comes equipped with a functor from the Fukaya category of X into coherent sheaves on $X^\vee$ which can be used to prove HMS as asserted.

In order to give an analogous presentation of this story, we define the Morse–Fukaya algebra $\mathcal{A}$ associated to a suitable class of SYZ fibrations $\pi : X \to B$; this is a curved A_∞-algebra determined by a Morse function on the total space X, taking coefficients in analytic functions on its rigid analytic mirror space. For an appropriate choice of Morse function, $\mathcal{A}$ can be understood as a (suitably deformed) algebra of Čech cochains valued in polyvector fields on $X^\vee$. We then construct an A_∞-functor from (a suitable subcategory of) the Fukaya category of X into the category mod-$\mathcal{A}$ of modules over $\mathcal{A}$ implementing the expected correspondence. Along the way we record comparison maps which together witness invariance of our constructions under a change of auxiliary technical choices.

Contents

Acknowledgments

I cannot overstate how much of a privilege it has been to be Denis' student, and I also cannot imagine having had a more generous PhD advisor. Denis has ceaselessly aided me through treacherous mathematical waters whenever called upon, his door always open, and I always found him there armed with an uncanny ability to make time for me where there seemed plainly to be none. Over these last five years I was almost certain to be visited by a colossal torrent of new and exciting mathematics, but I was lucky indeed to also be gifted that shining example of academic mentorship, which I will no doubt carry with me for the rest of my life. Thank you.

I am grateful to Mohammed Abouzaid and Chris Woodward for a number of useful comments and suggestions which shaped the final version of this dissertation. Thank you to Ben Gammage and Peter Kronheimer for kindly agreeing to serve on my dissertation committee.

During my time here the department has shown me only kindness, which I have felt from its members since my first day in Cambridge (we don't talk about 2020). I would like to thank you all for your unreasonably good-natured, inviting, and gracious demeanors. I have made many friends.

It is quite a dangerous gambit to attempt to list those who have found themselves owed gratitude accrued over such a long time, for fear of immortalizing some grave oversight. But let me indulge in that imprudence anyway, and exalt with particular distinction for one reason or another the following (possibly honorary) graduate students. Whether given for thoughtful conversations or a polite favor, you all know what you've done: Niven Achenjang, Amanda Burcroff, Jaewon Chang, Clair Dai, Alex Epelde, Elliot Glazer, Alexandra Hoey, Maxim Jeffs, Taeuk Nam, Carly Proctor, Wyatt Reeves, Rafael Saavedra, Natalie Stewart, and Jane Tan. Chief among them is my office mate Grant Barkley, whom I have had the honor of working alongside on the tiniest blackboard in the whole department for four years now. Best wishes to you all, and I hope to see you again sooner rather than later.

I must also thank those others both near and far who at one time or another have offered me their great counsel, or good humor. I beg your forgiveness, for the list which bears your names is too long to include here.

Finally I thank my family, whose support for me is so unwavering and complete that I will never be able to repay them. I love you all so much.

* * *

This work was partially supported by NSF grant DMS-2202984 via Denis Auroux.

Chapter 0

Overview

The Fukaya category of a symplectic manifold is an algebraic package—an example of an A_∞-category—which encapsulates its Lagrangian Floer theory. The purpose of this dissertation is to introduce a new model of the family Floer approach to Kontsevich's homological mirror symmetry conjecture, constructed via Morse theoretic technology, for the purpose of proving Theorem A.

Theorem A.[1] *There is a functor of curved A_∞-categories*

$$\mathcal{C} : \mathcal{F}_{\mathrm{sec}} \to \text{mod-}\mathcal{A}$$

from the Fukaya category of Lagrangian sections of a (suitable) SYZ fibration $\pi : X \to Q$ into the category of modules over the Morse–Fukaya algebra $\mathcal{A}$ of the fibration π.

Below we briefly summarize the approach to mirror symmetry taken in this dissertation, recalling some of the context in which it resides along the way.

0.1 Synopsis

The Strominger–Yau–Zaslow (SYZ) conjecture [SYZ96], and homological mirror symmetry (HMS) as originally set out by Kontsevich [Kon95b], are conjectures spawned from the mirror symmetry phenomenon observed by string theorists. Roughly speak-

[1] This is Theorem 4.27 below.

ing and as originally understood, SYZ mirror symmetry begins with a fibration of some kind of Kähler manifold by Lagrangian tori, and builds from this information a dual torus fibration via a geometric recipe. The difficulty is that the original fibration is allowed to have singular fibers (and generally will), and so the constructed mirror must be deformed accordingly.

On the other hand, homological mirror symmetry asserts a derived equivalence between the Fukaya category of a symplectic manifold and the category of coherent sheaves on its mirror [KS01]. Though in general it must be decided—as HMS is extended much beyond its initial incarnation comparing Calabi–Yau manifolds with their honest Calabi–Yau mirrors—what precisely is meant by Fukaya category, and whether to replace the derived category of coherent sheaves with for instance a non-commutative analogue.

The family Floer program [Fuk02; Abo14; Abo17; Tu14; Yua20] gives a modern reinterpretation of the construction of the SYZ mirror of $\pi : X \to Q$ as a moduli space of objects of the Fukaya category of X supported on the fibers of the fibration. The resulting object, technically a rigid analytic mirror $X^\vee$ of X, comes equipped with a functor from the Fukaya category of X into coherent sheaves on $X^\vee$ which can then, as an application, be used to prove HMS as asserted.

The Morse–Fukaya algebra $\mathcal{A}$ of the fibration $\pi : X \to Q$ is an A_∞-algebra determined by a Morse function on the total space X, taking coefficients in analytic functions on its mirror. For an appropriate choice of Morse function, $\mathcal{A}$ can be understood as a (suitably deformed) algebra of Čech cochains valued in polyvector fields on the SYZ mirror of X. The functor we construct in Theorem A then gives an analogous presentation of this story; here the category mod-$\mathcal{A}$ plays the role of a category of coherent sheaves as we explain below.

0.1.1 The Morse–Fukaya algebra of an SYZ fibration

Throughout fix a Kähler manifold (X, ω, J) with symplectic form ω and almost complex structure J. Let $\pi : X - D \to B$ be a fibration of X by Lagrangian tori, where $D \subset X$ is a complex hypersurface representing the anticanonical class of X. The standard way to proceed is to first construct a mirror to the complement $X^0 = X - D$.

As an act of technical expediency let us fix any compact, simply connected subset $B^0 \subset B$ which is disjoint from the critical values of π, and let $X^{00} = \pi^{-1}(B^0) \subset X^0$ be the corresponding restriction of the total space. For now also assume that $f : X^{00} \to \mathbb{R}$ is any Morse–Smale function with respect to a choice of metric on X^{00}.

The basic objects we consider are *(pseudoholomorphic) treed disks*; geometrically, these are continuous maps $u : \Delta \to X^{00}$ of decorated domains Δ built from the complex unit disk via an inductive gluing procedure, and which satisfy relations determined by their decorations (as developed by Charest–Woodward in [CW22] and originally Cornea–Lalonde [CL06]). Namely, an additional copy C of the unit disk may be glued into Δ by attaching one endpoint of a new line segment to the boundary ∂C, and the other endpoint to the boundary of a disk already in Δ. We call the image of each such disk C in Δ a *disk component*. It is also desirable to permit the attachment of semi-infinite line segments (rays) to disk boundaries, and to always remember the orientation of line segments we attach (whether finite or semi-infinite). Figure 1 depicts a schematic diagram of a treed disk domain built from two complex unit disks and four line segments (three of the segments having open ends).

Write $S \subset \Delta$ for the interior of the disks in Δ (the *surface part*) and $T \subset \Delta$ for the interior of the attached line segments (the *tree part*); then u restricts to maps $u_S(x)$ and $u_T(t)$ defined on S and T respectively. Let j be the complex structure on S induced by the standard complex structure on the unit disk.

Definition 0.2. A continuous map $u : \Delta \to X^{00}$ with treed disk domain Δ is *pseudo-holomorphic* if we have both

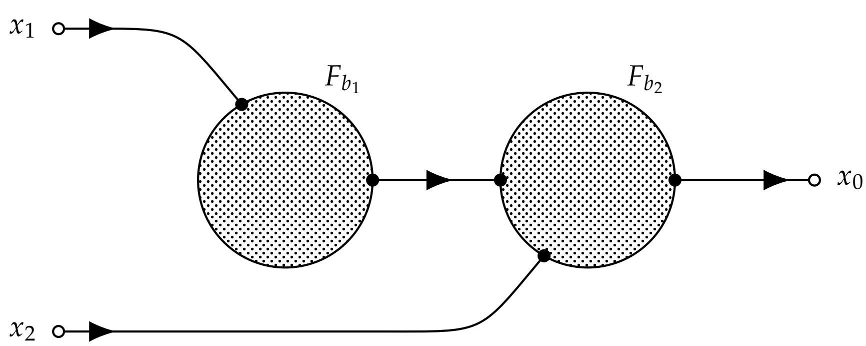

Figure 1: A schematic diagram of a pseudoholomorphic treed disk.

(0.1.I) u is *pseudoholomorphic* on the surface part: $J \circ \mathrm{D}\, u_S = \mathrm{D}\, u_S \circ j$, and

(0.1.II) u is a *Morse gradient flow line* on the tree part: $\dfrac{\mathrm{d}u_T}{\mathrm{d}t} = \nabla f \circ u_T$.

In other words, u must consist of a family of pseudoholomorphic disks attached along their boundaries, according to the edges of a tree, via Morse gradient flow lines.

Note that we have already suppressed several technical details; for example, in practice we perturb the pseudoholomorphic curve equation (0.1.I) due to transversality issues which arise while setting up the theory. In general, we allow J to be a domain-dependent almost complex structure determined by a background system of perturbation data, and similarly for (0.1.II). This data is chosen and managed consistently via an extension of the scheme of Charest–Woodward [CW22] (using stabilizing divisors) to the family setting. Relatedly, it is often convenient to equip points of treed disk domains with certain combinatorial bookkeeping labels, but we suppress these here as well.

Pseudoholomorphic treed disks give rise to algebraic operations via fixing a family of domains, prescribing boundary conditions, and then taking signed counts of their zero dimensional (compact) moduli spaces. These algebraic operations act on a complex with Morse theoretic generators, and this complex takes coefficients in a sheaf

$\mathcal{O}_{\mathrm{an}}$ obtained from the analytic functions on the *uncorrected mirror* of X^0. Let us now describe each of these details.

First, it is not difficult for us to arrange that the Morse function f on X^{00} lifts a Morse function on B^0 (so that their critical points coincide under π), and that f restricted to each fiber is perfect (i.e. gives a minimal Morse model for the n-torus). Let P be a cellular decomposition of B^0 with the property that each k-cell $\sigma \in P^{[k]}$ contains in its interior a unique index k critical point b_σ, and that σ is itself the closure of the stable manifold of b_σ.

Writing $F_b = \pi^{-1}(b)$ for the fibers, the uncorrected mirror then has underlying set of points simply the disjoint union [KS01; Abo14]

$$X^{\vee 0} = \bigsqcup_{b \in B^0} \mathrm{H}^1(F_b, U_\Lambda)$$

with $U_\Lambda = \mathrm{val}^{-1}(0)$ the unitary subgroup of the Novikov field Λ we work over.[2] Identify the groups $\pi_2(X, F_b)$ via isotoping fibers; then for each $\beta \in \pi_2(X, F_b)$ we naturally obtain a function

$$z^\beta = T^{\omega(\beta)}\,\mathrm{hol}(\partial\beta)$$

on $X^{\vee 0}$. In the definition of this *Floer-theoretic weight*, $\omega(\beta)$ is the symplectic area of β and $\mathrm{hol}(\partial\beta)$ denotes[3] evaluation of points of $X^{\vee 0}$ on the class $\partial\beta$. The set $X^{\vee 0}$ is naturally endowed with the structure of a rigid analytic space (having $(\Lambda^*)^n$ as a local model) for which the functions z^β are analytic. A chart is furnished by $(z^{\beta_i})_{1 \le i \le n}$, with the classes β_i chosen so that the $\partial\beta_i$ form a basis of $\mathrm{H}_1(F_b;\mathbb{Z})$.

[2] The Novikov field with coefficients in the field $\Bbbk$ (of characteristic zero, which we fix throughout) consists of series in the formal variable T of the form

$$\Lambda = \left\{ \textstyle\sum_i c_i T^{x_i} : c_i \in \Bbbk, x_i \in \mathbb{R}, x_i \to \infty \right\},$$

and comes equipped with a valuation map $\mathrm{val} : \sum_i c_i T^{x_i} \mapsto \min\{x_i : c_i \neq 0\}$.

[3] This notation is due to the fact that $X^{\vee 0}$ is realized as a moduli space of fibers of π equipped with a unitary rank 1 local system—since points of which are determined by their holonomy map, they equivalently belong to $\mathrm{H}^1(F_b, U_\Lambda)$ for some $b \in B$.

There is a natural projection $\pi^\vee : X^{\vee 0} \to B^0$ and, after suitably refining P by perturbing f, the collection of analytic functions on $(\pi^\vee)^{-1}(\mathrm{star}(\sigma))$ for each $\sigma \in P$ assemble into a sheaf $\mathcal{O}_{\mathrm{an}} = \pi^\vee_*(\mathcal{O}_{X^{\vee 0}})$ of universal weights.

Definition 0.3. The *Morse–Fukaya algebra* $\mathcal{A} = \mathrm{CM}^\bullet(\pi, f; \mathcal{O}_{\mathrm{an}})$ is the module freely generated by the critical points of f, with coefficients taken in $\mathcal{O}_{\mathrm{an}}$, equipped with algebraic operations μ^d for $d \geq 0$ we outline below.

Fix $d+1$ points $x_0, x_1, \ldots, x_d \in \mathrm{crit}\, f$. We say that a pseudoholomorphic treed disk $u : \Delta \to X^{00}$ has d *inputs and* 1 *output* if, respecting orientations, in constructing Δ we attached exactly d copies of the ray $(-\infty, 0]$ and 1 copy of the ray $[0, \infty)$. Call $x \in X^{00}$ the input (resp. output) of the ray $R = (-\infty, 0] \subset \Delta$ (resp. $R = [0, \infty) \subset \Delta$) whenever $\lim_{|t| \to \infty} u|_R(t) = x$. Thus Figure 1 depicts a pseudoholomorphic treed disk with 2 inputs x_1 and x_2 and 1 output x_0. Note that a treed disk domain Δ with 1 output has a canonical ordering on its inputs induced by the orientation of the disk components of Δ.

For each $\beta \in \pi_2(X, F_{b_0})$ we may form the moduli space $\overline{\mathcal{M}}_{d+1}(x_0, \ldots, x_d; \beta)$ from all (suitably perturbed) pseudoholomorphic treed disks $u : \Delta \to X^{00}$ with:

- Disk boundaries lying on fibers—each disk component $C \subset \Delta$ satisfies $u(\partial C) \subset F_b$ for some $b \in B$ (all possibly different).
- Representing class β—each disk component $C \subset \Delta$ gives rise to a class $[u|_C] \in \pi_2(X, F_b)$ hence in $\pi_2(X, F_{b_0})$, and we demand that all such classes sum to β.
- Correct I/O—we require that u has d inputs $x_1, \ldots, x_d$ and 1 output x_0.
- Stable components—the map u obeys a family of straightforward technical conditions[4] which ensure we obtain a compact Hausdorff moduli space with the correct dimension.

[4] For instance, we require that each disk component on which u is constant must meet at least 3 line segments.

We must of course also take care to develop a consistent scheme to orient these moduli spaces, though we do not elaborate further here on these technical details [WW15; Fuk+10b].

Now given $\mathbf{x} = (x_1, \dots, x_d) \in \operatorname{crit} f$ we set

$$\mu^d(\mathbf{x}) := \sum_{\beta, x_0} \# \overline{\mathcal{M}}_{d+1}(x_0, x_1, \dots, x_d; \beta) \cdot z^\beta x_0, \tag{0.1.III}$$

where # is the signed count of oriented points, and the sum is taken over classes β and critical points x_0 for which the expected dimension of $\overline{\mathcal{M}}_{d+1}(x_0, x_1, \dots, x_d; \beta)$ is zero. In accordance with the $\mathbb{Z}_2$-grading induced by Morse index mod 2, upon declaring that each μ^d is $\mathcal{O}_{\mathrm{an}}$-linear we obtain a family of graded multiplication maps $\mu^d : \mathcal{A}^{\otimes d} \to \mathcal{A}[2-d]$.

Theorem 0.4. *The operations μ^d endow the Morse–Fukaya algebra $\mathcal{A}$ with the structure of a curved A_∞-algebra [Aur23]. In other words, for each $d > 0$ and homogeneous elements $a_1, \dots, a_d \in \mathcal{A}$ of respective degrees $|a_i|$ we have the identity [Sei08a]*

$$0 = \sum_{m+n \leq d} (-1)^{\heartsuit} \mu^{d-n+1}(a_1, \dots, a_m, \mu^n(a_{m+1}, \dots, a_{m+n}), a_{m+n+1}, \dots, a_d), \tag{0.1.IV}$$

with $\heartsuit = m + \sum_{i=1}^m |a_i|$.

One obtains a proof of Theorem 0.4 by a careful analysis of the boundary components of the higher dimensional strata of the moduli spaces $\overline{\mathcal{M}}_{d+1}(x_0, \dots, x_d; \beta)$ we have just introduced; ultimately, the signed count of points on the boundary of a 1-dimensional oriented compact moduli space is zero. For example, when a Morse gradient flow line in a treed disk "breaks" (on the boundary of a moduli space) through an intermediate critical point, the treed disk naturally decomposes as the composition of two less complex treed disks, one stacked upon the other. All such possible decompositions appear as terms in (0.1.IV).

0.1.2 An HMS comparison functor

We associate to the Lagrangian fibration $\pi : X^{00} \to B^0$ a full subcategory $\mathcal{F}_{\text{sec}}$ of the Fukaya category of X, whose objects are Lagrangian sections L of π over B^0. Each such section is naturally equipped with a Morse function f_L via restriction of the global Morse function on X^{00}. For simplicity, fix a finite collection $\{L_i\} \subset \mathcal{F}_{\text{sec}}$ intersecting pairwise transversely. If $L_i \neq L_j$ we let $\text{Hom}(L_i, L_j)$ be freely generated by the points of $L_i \cap L_j$ with coefficients in Λ. If instead $L_i = L_j$ we substitute the Fukaya–Morse algebra of the ordinary Lagrangian $L = L_i$ as defined by Charest–Woodward [CW22] (i.e. $\text{Hom}(L, L)$ is generated by critical points of f_L with coefficients in Λ—the algebra operations are as above, except that we now require all Morse flow lines and disk component boundaries to lie wholly in L). The composition of, for example, morphisms $p \in \text{Hom}(L_1, L_2)$ and $q \in \text{Hom}(L_2, L_3)$ between distinct Lagrangian sections is the familiar multiplication in the Fukaya category; we count pseudoholomorphic strips with boundary on $L_1 \cup L_2 \cup L_3$ meeting p, q, and all possible third points of $\text{Hom}(L_1, L_3)$, in the usual way.

We are now in a position to see how the comparison functor $\mathcal{C} : \mathcal{F}_{\text{sec}} \to \text{mod-}\mathcal{A}$ of Theorem *A* is defined. First, on objects we set

$$L \in \mathcal{F}_{\text{sec}} \mapsto \mathcal{C}(L) := \text{CM}^\bullet(\pi, f_L; \mathcal{O}_{\text{an}}),$$

this module being generated by the points of crit f_L with coefficients in $\mathcal{O}_{\text{an}}$. The object $\mathcal{C}(L)$ carries a family of A_∞-module action maps $\triangleleft^d : \mathcal{C}(L) \otimes \mathcal{A}^{\otimes d} \to \mathcal{C}(L)$ which, according to the natural analogue of (0.1.III), now count moduli spaces of treed disks of the kind for example schematically depicted in Figure 2a. The key modification is that we now allow the boundary of disk components to lie on the union of a particular fiber and some number of Lagrangian sections; Morse gradient flow lines are in turn suitably constrained to either a fiber or particular sections.

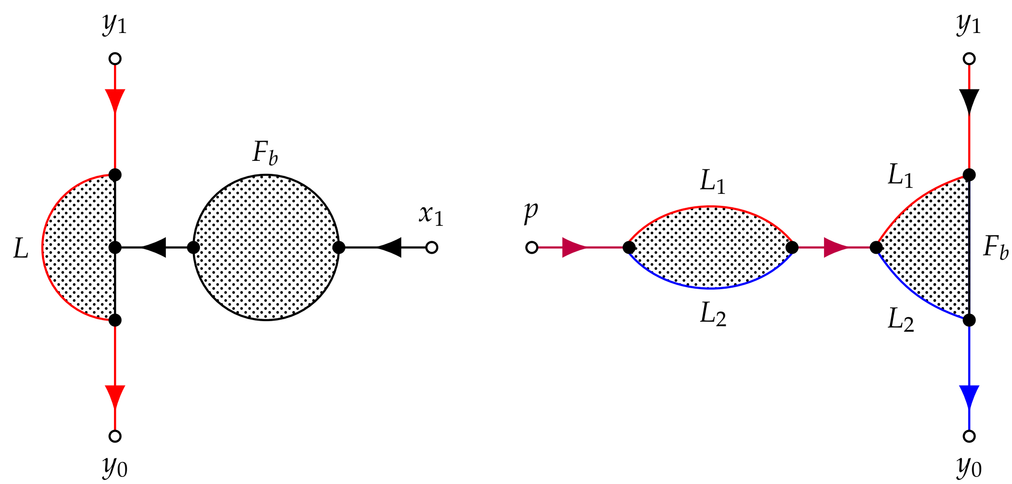

(a) Coefficient of y_0 in $y_1 \triangleleft^1 x_1$ (b) Coefficient of y_0 in $C^1(p)(y_1)$

Figure 2: A pair of schematic diagrams of treed disks captioned with the coefficient to which their counts contribute in the associated module action or morphism. Boundaries and line segments constrained to a Lagrangian section are shown in color.

Similarly, a morphism $p \in L_1 \cap L_2 = \mathrm{Hom}(L_1, L_2)$ gives rise to a morphism $C(L_1) \to C(L_2)$ of A_∞-modules by counting treed disks such as those schematically of the type depicted in Figure 2b or, for the higher order terms of the A_∞-module homomorphism, analogous configurations with additional inputs. Note that in this particular example both of the horizontal line segments in Figure 2b are designated Morse flow lines wholly contained in the transverse intersection $L_1 \cap L_2$—hence their image in any pseudoholomorphic treed disk must be constant. The requisite A_∞-relations for both the module actions and module morphisms hold by essentially the same analysis as in the previous section; we again carefully consider the several ways treed disks such as those schematically depicted in Figure 2 can break.

0.2 Layout

The technical foundations necessary to construct the algebra $\mathcal{A}$ are considerable. In Chapter 1 we recall fundamental facts from the theory of pseudoholomorphic curves and Morse theory, principally as a vehicle to introduce our basic terminology. In Chapter 2 we treat the moduli spaces of the fundamental objects we study—so-called *pseudoholomorphic treed disks*, and prove the necessarily generic regularity results of general type required for Floer theory. In Chapter 3 we actually construct the Morse–Fukaya algebra $\mathcal{A}$, in particular proving compactness of the relevant moduli spaces and therefore well-definedness of the multiplication law, before discussing invariance of the construction under all choices. In Chapter 4 we define the operations specifying the family Floer functor $\mathcal{C}$, in the process making the necessary enhancements of the constructions of Chapters 2 and 3, finally proving that this data assembles into an actual A_∞-functor.

Chapter 1

Preliminaries

Throughout fix a Kähler manifold $(\overline{X}, \omega, J)$ with symplectic form ω, almost complex structure J, and metric g. Let $\pi : \overline{X} - D \to \overline{Q}$ be a fibration of $\overline{X}$ by Lagrangian tori, where $D \subset \overline{X}$ is a complex hypersurface representing the anticanonical class of $\overline{X}$. The standard way to proceed is to construct a mirror to the complement $X^0 := \overline{X} - D$ of this divisor [Aur23].

In order to avoid at the outset issues of convergence near singular fibers $F_q \subset X^0$ of the fibration, fix a compact subset $Q \subset \overline{Q}$ disjoint from the critical values of π, and let $X := \pi^{-1}(Q) \subset X^0$ be the corresponding restriction of the total space. The perturbations and auxiliary structures which we construct in the sequel will be defined directly on X.

1.1 Morse theory

We begin by recalling basic facts from Morse theory, introducing our terminology and notation as we go. Let $f : X \to \mathbb{R}$ be a smooth function.

Definition 1.1. Denote by $\operatorname{crit} f$ the collection of *critical points* of f. A critical point $x \in \operatorname{crit} f$ is *non-degenerate* if the Hessian $\operatorname{Hess}(f)$ is invertible at x. The *index* $I(x)$ of a non-degenerate critical point is its Hessian's number of negative eigenvalues.

The gradient $\nabla^g f$ defines a *gradient flow* on X. For any $[a, b] \subset [-\infty, \infty]$, a continuous map $u(t) : [a, b] \to X$ which is smooth on (a, b) and obeys $u'(t) = \nabla^g_{u(t)} f$ is a *Morse* *gradient flow trajectory from $u(a)$ to $u(b)$*. For each $x \in \operatorname{crit} f$ we have an *ascending set* $W^{\uparrow}(x)$ and *descending set* $W^{\downarrow}(x)$ defined by

$$W^{\uparrow}(x) := \left\{ I(b) \mid u : [-\infty, \quad b] \to X \textit{ is a gradient flow line with } u(-\infty) = x \right\} \text{ and}$$
$$W^{\downarrow}(x) := \left\{ I(a) \mid u : [\quad a, +\infty] \to X \textit{ is a gradient flow line with } u(+\infty) = x \right\}.$$

Theorem 1.2 (Stable manifold theorem)**.** *If f has no degenerate critical points then $W^{\uparrow}(x)$ and $W^{\downarrow}(x)$ are each smooth manifolds and we have*

$$\dim W^{\uparrow}(x) = \dim X - I(x) \quad \textit{and} \quad \dim W^{\downarrow}(x) = I(x).$$

Definition 1.3. The function $f : X \to \mathbb{R}$ is *Morse–Smale* (or simply *Morse*) if all critical points of f are non-degenerate and for all $x_0, x_1 \in \operatorname{crit} f$ the manifolds $W^{\downarrow}(x_0)$ and $W^{\uparrow}(x_1)$ meet transversely.[1]

Henceforth assume that f is Morse.

Definition 1.4. Each space $W^{\downarrow}(x_0)$ and $W^{\uparrow}(x_1)$ carries an action of $\mathbb{R}$ by the gradient flow, or equivalently for Morse gradient flow trajectories with domain $[-\infty, \infty]$, by domain translation. This action is compatible with passing to the intersection $W^{\downarrow}(x_0) \cap W^{\uparrow}(x_1)$. We define the *moduli space of Morse gradient flow trajectories from x_0 to x_1*, denoted by $\mathcal{M}(x_0, x_1)$, as the quotient of $W^{\downarrow}(x_0) \cap W^{\uparrow}(x_1)$ by this $\mathbb{R}$-action.

Theorem 1.5. *Each moduli space $\mathcal{M}(x_0, x_1)$*

- *is a smooth manifold of dimension $I(x_0) - I(x_1) - 1$,*
- *has a compactification $\overline{\mathcal{M}}(x_0, x_1)$ with boundary strata consisting of the broken Morse gradient flow trajectories*[2] *from x_1 to x_0 (see Figure 1.1), and*

[1] Thus each intersection $W^{\downarrow}(x_0) \cap W^{\uparrow}(x_1)$ is again a smooth manifold of dimension $I(x_0) - I(x_1)$.

[2] Precise models for these compactified moduli spaces appear for example in [Fuk93]. The general formalism of our Chapter 2 subsumes the concept as a special case of Definition 2.16.

- *there exist coherent choices of orientations of the descending and ascending manifolds* $W^{\downarrow}(x_0)$ *and* $W^{\uparrow}(x_1)$ *which induce orientations on the moduli spaces* $\overline{\mathcal{M}}(x_0, x_1)$ *compatible with restriction to boundary strata.*

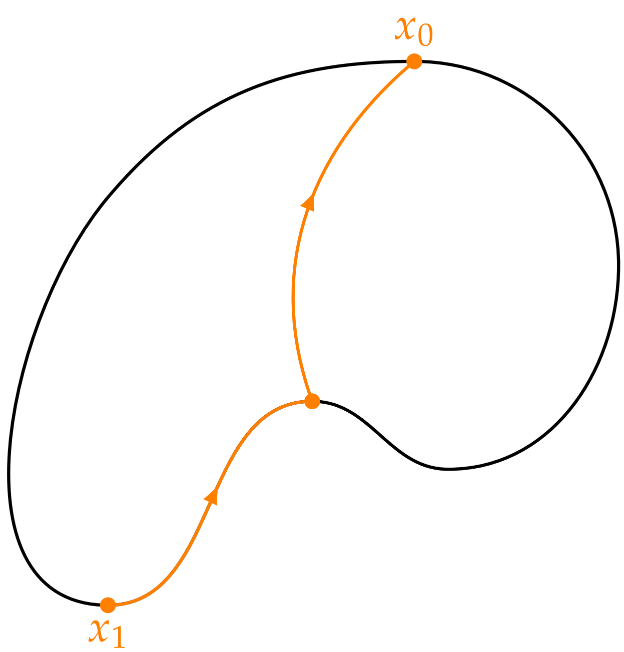

Figure 1.1: A breaking of a Morse gradient flow trajectory on the sphere at the boundary of the moduli space $\mathcal{M}(x_0, x_1) \subset \overline{\mathcal{M}}(x_0, x_1)$ into a broken Morse trajectory through the point x_2.

As in [Fuk93], there is also the generalized notion of a *Morse gradient flow tree* beginning at multiple *input* points $x_1, x_2, \ldots, x_n$ simultaneously and again flowing to a single *output* x_0 (see Figure 1.2). Arranging via some strategy that transversality concerns are assuaged—such as in Theorem 2.45 below—we obtain the following theorem.

Theorem 1.6 ([Fuk93], [Maz22, Theorem I.9])**.** *The Morse complex* $CM^{\bullet}(f) := \mathbb{R}\langle \mathrm{crit}\, f \rangle$*, freely generated by the critical points of* f *and graded by index* I*, is equipped with* n*-ary* *multiplications*

$$\mu^n : CM^\bullet(f)^{\otimes n} \to CM^\bullet(f)[2-n]$$

endowing $CM^\bullet(f)$ *with the structure of a flat* A_∞*-algebra. Each multiplication* μ^n *is respectively defined by counting Morse gradient flow trees with* n *inputs.*

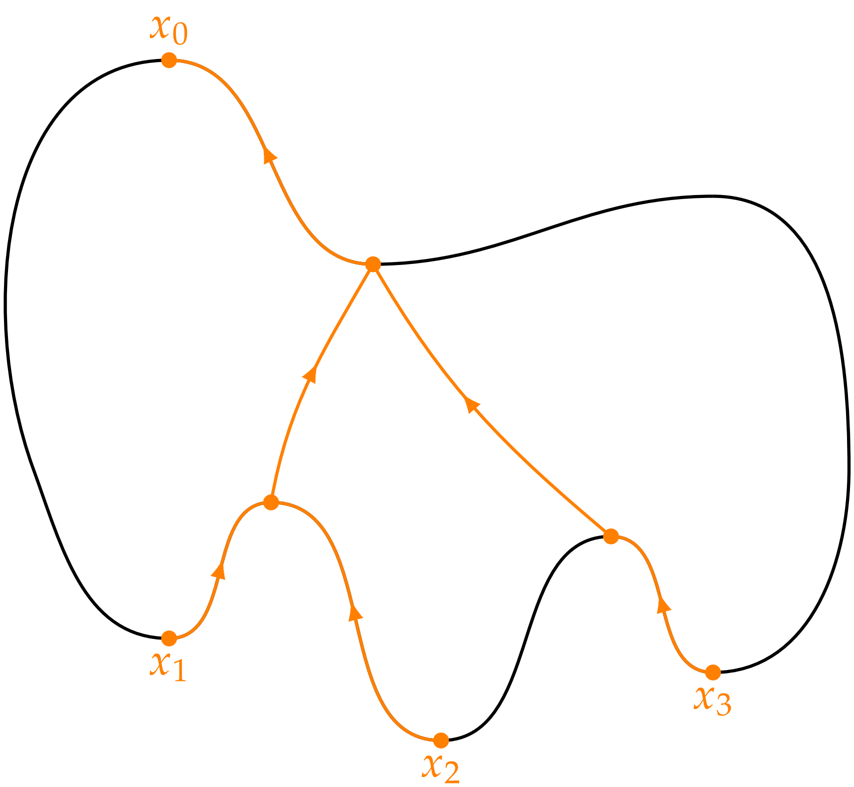

Figure 1.2: A Morse gradient flow tree with 3 inputs x_1, x_2, and x_3, and output x_0.

1.2 Pseudoholomorphic disks and curves

There is a related theory of compactified moduli spaces of pseudoholomorphic curves, essentially arising as the infinite-dimensional analogue of the Morse case. In order to make sense of the statement of the main theorem, we introduce the key tool upon which essentially all of our constructions fundamentally rely for well-definedness.

Definition 1.7. An bounded linear map $\Phi : V \to W$ between Banach spaces is *Fredholm* if $\dim\ker\Phi < \infty$ and $\dim\operatorname{coker}\Phi < \infty$. The *index* of Φ is

$$\operatorname{ind}\Phi := \dim\ker\Phi - \dim\operatorname{coker}\Phi.$$

A C^q-map $f : \mathcal{B} \to \mathcal{E}$ between connected Banach manifolds is *C^q-Fredholm* if $Df : T\mathcal{B} \to \mathcal{E}$ is pointwise Fredholm. The index of Df is constant on $\mathcal{B}$, and gives *index* $\operatorname{ind} f$ *of* f.

Theorem 1.8 (Sard–Smale [Sma65])**.** *Let $f : \mathcal{B} \to \mathcal{E}$ be a C^q-Fredholm map between separable Banach manifolds with $q > \max(\operatorname{ind} f, 0)$. The set of regular values of f is comeager in $\mathcal{E}$.*

Denote by j the canonical complex structure on both the complex unit disk $\mathbb{D}$ and sphere $\mathbb{C}\mathrm{P}^1$.

Definition 1.9. A smooth map $u : \Sigma \to \overline{X}$ is respectively a *J-holomorphic disk* (if $\Sigma = \mathbb{D}$) or *curve* (if $\Sigma = \mathbb{C}\mathrm{P}^1$) when $\mathrm{D}_z u \circ j = J \circ \mathrm{D}_z u$ for all $z \in \Sigma$. Equivalently, $u : C \to \overline{X}$ is J-holomorphic when u is a zero of the operator

$$\overline{\partial}_J u := \frac{1}{2}(\mathrm{D}u + J \circ \mathrm{D}u \circ j). \tag{1.2.I}$$

A J-holomorphic disk or curve u is *simple* if it is not multiply covered. Denote by $\mathcal{M}^*(\Sigma, \beta)$ the *moduli space of J-holomorphic maps with domain Σ and representing the homology class*[3] β.

The following transversality theorem is the analogue of the first part of Theorem 1.5.

[3]When $\Sigma = \mathbb{D}$ we require that u takes $\partial\mathbb{D}$ into a fixed Lagrangian $L \subset X$.

Theorem 1.10 ([MS12, Theorem 3.1.6, Theorem A.3.3]). *The operator $\overline{\partial}_J$ of (1.2.I) may be recast, parameterized over the space of all ω-tame almost complex structures of class C^l on X, as a section of a particular separable Banach vector bundle.*

Interpreted this way, the derivative D_u of this section at an almost complex structure J and J-holomorphic map $u : \Sigma \to \overline{X}$ is a Cauchy–Riemann operator in the sense of [MS12, Appendix C], hence a Fredholm operator of index[4]

$$I(\beta) := \operatorname{ind} D_u = \begin{cases} \dim L \ + 2c_1(u^*\mathrm{T}X, u|^*_{\partial\mathbb{D}}\mathrm{T}L) & \Sigma = \mathbb{D} \\ \dim X + 2c_1(u^*\mathrm{T}X) & \Sigma = \mathbb{C}\mathrm{P}^1 \end{cases}. \tag{1.2.II}$$

In particular, for a comeager set of J the moduli space $\mathcal{M}^*(\Sigma, \beta)$ is a smooth manifold of dimension

$$\dim \mathcal{M}^*(\Sigma, \beta) = I(\beta).$$

The remainder of the analogy of Theorem 1.5 is accounted for by the following theorem of Gromov [Gro85], subsequently recast here in more modern terms due to Kontsevich [Kon95a]. The key idea is that, by suitably enhancing the notion of pseudoholomorphic disks and curves by permitting nodal unions of the same, and recording distinguished marked points in their domains (see Figure 1.3), there is a well-defined notion of *stable pseudoholomorphic disk and spheres*.[5]

Theorem 1.11 (Gromov compactness [MS12, Theorem 5.3.1]). *Let (J_k) be a sequence of ω-tame almost complex structures C^∞-converging to J, for which each element of the sequence $(u_k : \Sigma \to \overline{X})$ of maps of disks or spheres are each respectively J_k-holomorphic. If the symplectic area of the sequence (u_k) is uniformly bounded, then there is a subsequence which converges in a suitable sense to a stable pseudoholomorphic disk or sphere u.*

---

[4] Here $2c_1(u^*\mathrm{T}X, u|^*_{\partial\mathbb{D}}\mathrm{T}L)$ denotes twice the relative Chern class of the bundle pair, i.e. the *Maslov class* of $\beta = [u]$.

[5] We give a precise model in Chapter 2 of particular objects—namely *pseudoholomorphic treed disks*—which subsume this concept.

Moreover, there is a moduli space $\overline{\mathcal{M}}(\Sigma, \beta)$ of all stable maps representing a fixed class β, equipped with a natural topology known as the Gromov topology compatible with this notion of convergence. For each $E > 0$, every subspace of stable maps with total symplectic area at most E is compact.

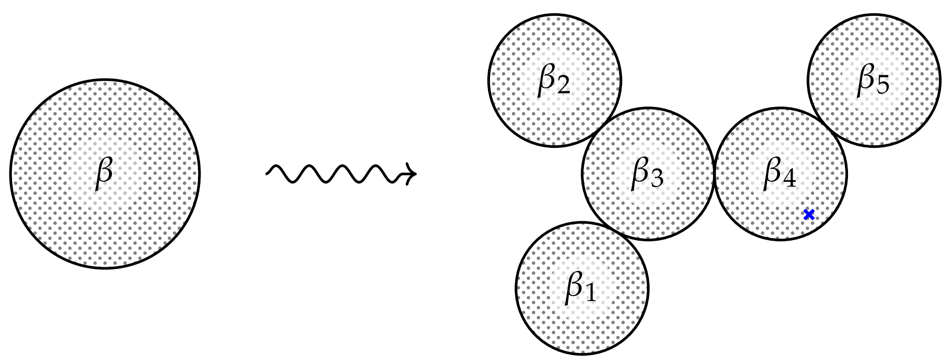

Figure 1.3: An example stable limiting configuration of a sequence of pseudoholomorphic disks of class $\beta \in H_2(\overline{X}, L)$ degenerating a nodal union of disks of respective classes $\beta_i \in H_2(\overline{X}, L)$ at the boundary of the compactified moduli space. Here $\beta = \sum_{i=1}^{5} \beta_i$) and the distinguished marked point stabilizing a constant component (with $\beta_4 = 0$) is indicated by a blue cross.

Chapter 2

Pseudoholomorphic treed disks

In this chapter we lay the necessary technical foundations, defining the basic objects we study—so-called "pseudoholomorphic treed disks"—building directly on the ideas of Charest–Woodward [CW22] and Venugopalan–Woodward–Xu [VWX20]. In various forms, these constructions (incorporating Fukaya's *Morse gradient flow trees* [Fuk93] and pseudoholomorphic disks) have appeared as the *clusters* of Cornea–Lalonde [CL06] and e.g. in Seidel's [Sei08c] subsequent variation. Our description culminates in a proof of the fundamental regularity theorem for the moduli spaces we consider.

2.1 Treed disks and decorations

Pseudoholomorphic treed disks $u : \underline{\Delta} \to \overline{X}$ are continuous maps from decorated domains $\underline{\Delta}$ which satisfy relations determined by the decorations. We first define the combinatorial data specifying these domains.

Definition 2.1. A *(marked) treed disk* Δ consists of:

- A finite collection $\mathrm{Vert}(\Delta)$ of *vertices*, each $v \in \mathrm{Vert}(\Delta)$ being a point, or a *component* homeomorphic to either the complex unit disk $\mathbb{D}$ or sphere $\mathbb{CP}^1$; hence admitting a partition

  $$\mathrm{Vert}(\Delta) = \mathrm{Vert}^+(\Delta) \sqcup \mathrm{Vert}^\bullet(\Delta) \sqcup \mathrm{Vert}^\circ(\Delta)$$

  into point, disk, and sphere components respectively.

- A distinguished *root* vertex $v_0 \in \mathrm{Vert}^+(\Delta)$.

- A unique directed *edge* $e \in \mathrm{Edge}(\Delta) := \mathrm{Vert}(\Delta) - \{v_0\}$ emanating from a point on each non-root vertex as specified by *head* and *tail* maps $h, t : \mathrm{Edge}(\Delta) \to \bigsqcup \mathrm{Vert}(\Delta)$, each with a respective *length* given by $l : \mathrm{Edge}(\Delta) \to [0, \infty]$. In particular, each head $h(e)$ or tail $t(e)$ is a point inside another vertex $v \in \mathrm{Vert}(\Delta)$.

- A finite family $\mathcal{K}$ of possible *flavors*, and for each $\kappa \in \mathcal{K}$ a finite collection $\mathrm{Mark}_\kappa(\Delta) = \{d_{\kappa,i}\}$ of *marked points*, each $d_{\kappa,i}$ lying in the interior of a sphere or disk component.

In order that the edges of Δ give acceptable rules gluing the vertices of Δ into essentially a geometrically realized rooted tree, we require that:

(2.1.a) A point z of a component $v \in \mathrm{Vert}(\Delta)$ which is the either the head or tail of an edge, or is a marked point, is called a *joint*. All joints whatever their kind must be distinct, with the exception that marked points of different flavors may coincide.

(2.1.b) Each point $v^+ \in \mathrm{Vert}^+(\Delta)$ is the head or tail of at most one edge in each case. A point v^+ with valence two is called a *breaking*. Otherwise v^+ has valence one and is either the tail of an edge, in which case it is called an *input*, or else is the head of an edge, in which case $v^+ = v_0$ (the root) is the (unique) *output*. We denote the number of inputs by $n(\Delta)$.

(2.1.c) There is a unique *root edge* e_0 with $h(e_0) = v_0$, so that in particular there are at least two vertices.

(2.1.d) Each edge $e \in \mathrm{Edge}(\Delta)$ is *infinite* if $h(e)$ or $t(e)$ lies on a point, or otherwise is *combinatorially finite*; this gives a decomposition

$$\mathrm{Edge}(\Delta) = \mathrm{Edge}_\to(\Delta) \sqcup \mathrm{Edge}_-(\Delta)$$

into infinite and combinatorially finite edges respectively.

(2.1.e) The unique point of any $v^+ \in \mathrm{Vert}^+(\Delta)$ and the boundary points of any disk component $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$ are said to be of *boundary type*.

For each edge e we have that $t(e)$ is of boundary type if and only if $h(e)$ is of boundary type. Moreover if $t(e) \in v$ for some $v \in \mathrm{Vert}^\bullet(\Delta)$, then $t(e) \in \partial v$. This gives a further refinement

$$\mathrm{Edge}_{\_}(\Delta) = \mathrm{Edge}^\circ_{\_}(\Delta) \sqcup \mathrm{Edge}^\partial_{\_}(\Delta)$$

into the collections of *interior* and *boundary* combinatorially finite edges; namely consisting of those edges e for which $t(e)$ hence $h(e)$ lies in the interior of some component, or otherwise, respectively.

(2.1.f) The length of each edge $e \in \mathrm{Edge}(\Delta)$ must obey the requirement that

$$l(e) \in \begin{cases} \{\infty\} & e \in \mathrm{Edge}_{\to}(\Delta) \\ \{0\} & e \in \mathrm{Edge}^\circ_{\_}(\Delta) \\ [0,\infty) & e \in \mathrm{Edge}^\partial_{\_}(\Delta) \end{cases} .$$

Remark 2.2. Of course, the head and tail maps respectively descend to a pair of maps $\underline{h}, \underline{t} : \mathrm{Edge}(\Delta) \to \mathrm{Vert}(\Delta)$ which equip $\mathrm{Vert}(\Delta)$ with the structure of a directed tree with root v_0. The boundary axiom (2.1.e) prevents the existence of sphere components when there are no disk components.

In the sequel, it will be convenient to denote by $\mathrm{Joint}(\Delta)$ the entire set of joints, and to write $\mathrm{Joint}^\wedge(\Delta) \subset \mathrm{Joint}(\Delta)$ for the subset consisting of all breakings. For notational convenience we adopt the convention that a point vertex $v^+ \in \mathrm{Vert}^+(\Delta)$ is identified with the unique point it contains.

The data of a treed disk Δ gives rise in straightforward fashion to, and essentially decorates, an honest topological space $\underline{\Delta}$. To build $\underline{\Delta}$, we associate the line segment

$L_e = [0, l(e)]$ to each combinatorially finite edge $e \in \mathrm{Edge}_{-}(\Delta)$. Similarly, to each infinite edge $e \in \mathrm{Edge}_{\rightarrow}(\Delta)$ we associate

$$L_e = \begin{cases} [-\infty,\ 0] & \underline{t}(e) \in \mathrm{Vert}^{+}(\Delta) \text{ and } \underline{h}(e) \notin \mathrm{Vert}^{+}(\Delta) \\ [\ \ \ 0, \infty] & \underline{t}(e) \notin \mathrm{Vert}^{+}(\Delta) \text{ and } \underline{h}(e) \in \mathrm{Vert}^{+}(\Delta) \\ [-\infty, \infty] & \text{otherwise} \end{cases},$$

and write w_e^- and w_e^+ for the respective left and right endpoints of L_e. Of course, all of these spaces are diffeomorphic, but we understand them as recording a distinguished parameterization by an interval subset of $\mathbb{R}$.

Now form $\underline{\Delta}$ as the quotient by $\sim$ of the disjoint union of

$$\bigsqcup_{v \in \mathrm{Vert}(\Delta)} v \quad \text{and} \quad \bigsqcup_{e \in \mathrm{Edge}(\Delta)} L_e,$$

where for all $e \in \mathrm{Edge}(\Delta)$ we declare $t(e) \sim w_e^-$ and $h(e) \sim w_e^+$.

Writing S_Δ for the image of the components $v \in \mathrm{Vert}^{\bullet}(\Delta) \cup \mathrm{Vert}^{\circ}(\Delta)$ in $\underline{\Delta}$ (the *surface part*), and T_Δ for the image of the line segments L_e (the *tree part*), by construction $\underline{\Delta}$ is the union $S_\Delta \cup T_\Delta$. Figure 2.1 schematically depicts a representative pair of treed disk configurations. Note that inclusion of the union of all components into $\underline{\Delta}$ need not be injective since a zero-length edge will identify points of distinct components, creating a *node*, as depicted in Figure 2.1a (for a disk and sphere) and Figure 2.1b (for a pair of disks).

Definition 2.3. Let Δ, Δ' be treed disks. A homeomorphism $\psi : \underline{\Delta} \to \underline{\Delta}'$ is an *isomorphism of treed disks* if

- the map $\phi|_{\mathrm{int}(S_\Delta)}$ is a biholomorphism $\mathrm{int}(S_\Delta) \cong \mathrm{int}(S_{\Delta'})$ on the interior of the surface part,
- the map $\phi|_{\mathrm{int}(T_\Delta)}$ is an isometry $\mathrm{int}(T_\Delta) \cong \mathrm{int}(T_{\Delta'})$ on the interior of the tree part,

(a)

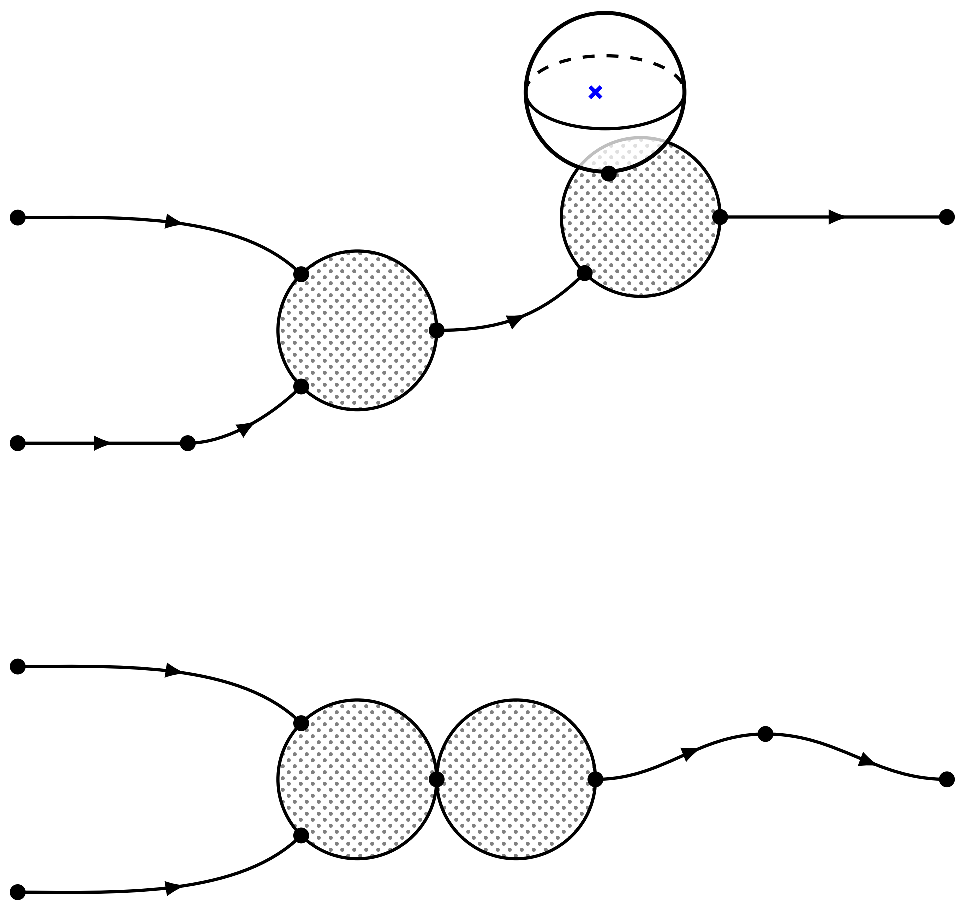

(b)

Figure 2.1: Schematic diagrams of a pair of treed disks.

and

- the map $\psi|_{\mathrm{Mark}_\kappa(\Delta)}$ is a bijection onto $\mathrm{Mark}_\kappa(\Delta')$ for each flavor $\kappa \in \mathcal{K}$.

Of course, an isomorphism of treed disks is equivalently the data of an appropriate identification of each individual $v \in \mathrm{Vert}(\Delta)$, L_e for $e \in \mathrm{Edge}(\Delta)$, and joint subject to compatibility with the head and tail maps, and flavors.

Definition 2.4. An *ordered* treed disk is a treed disk Δ equipped with an orientation of each component $v \in \mathrm{Vert}(\Delta)$, hence—given the existence of the distinguished root vertex—an induced ordering of the (boundary) joints lying on each ∂v, along with a choice of ordering of each flavor κ's corresponding set $\mathrm{Mark}_\kappa(\Delta)$ of interior joints.

An isomorphism of ordered treed disks is an isomorphism of underlying treed disks which respects orientation and all orderings.

From now on, we will assume that all of our treed disks are ordered.

Definition 2.5. The *combinatorial type* $\Gamma = \Gamma(\Delta)$ of a treed disk Δ is an enhancement of its "underlying tree" as in Remark 2.2; this is the directed tree with vertex set $\mathrm{Vert}(\Gamma) := \mathrm{Vert}(\Delta)$ and edge set $\mathrm{Edge}(\Gamma) := \mathrm{Edge}(\Delta)$, where we interpret $e \in \mathrm{Edge}(\Gamma)$ as an edge from the vertex containing $t(e)$ to the vertex containing $h(e)$. We continue to decorate the graph Γ with orderings, both of the flavors of marked points and of the edges incident to each vertex. We also remember the subsets $\mathrm{Edge}_-^0(\Delta)$ and $\mathrm{Edge}_-^{(0,\infty)}(\Delta)$ of $\mathrm{Edge}_-(\Delta)$ consisting of those edges with respective zero and positive finite lengths.

Remark 2.6. As Charest–Woodward [CW22, Section 4.2] explain, the type $\Gamma(\Delta)$ of a treed disk is essentially the data of a so-called metric ribbon tree, except that we have not ordered the entire collection of edges incident on each vertex. Note that in their approach interior markings are recorded as the attachment points of additional "fake" infinite edges (so-called *interior leaves*)—we deviate from this convention here.

Though it is not our particular focus, for consistency with the setup of [CW22] we now introduce weights attached to some of the edges in our treed disks. In the sequel they will be used to produce an explicit strict unit for the A_∞-operations we construct, via ensuring compatibility of our perturbations with forgetting inputs on-the-nose (an adaption of the methods of [Gan12, Section 4]).

Definition 2.7. A *weighted* treed disk is a treed disk Δ equipped with a function $\rho : \mathrm{Edge}_{\rightarrow}(\Delta) \to [0, \infty]$ which assigns edge *weights*. There is then a subset $\mathrm{Edge}^{{\color{gray}\blacktriangledown}}(\Delta) := \rho^{-1}((0, \infty)) \subset \mathrm{Edge}_{\rightarrow}(\Delta)$ of *weighted* edges. Unweighted edges $e \in \mathrm{Edge}_{\rightarrow}(\Delta)$ must either have $\rho(e) = 0$ in which case we call them *unforgettable*, else $\rho(e) = \infty$ and they are *forgettable* edges. In particular there is a decomposition

$$\mathrm{Edge}_{\rightarrow}(\Delta) = \mathrm{Edge}^{\blacktriangledown}(\Delta) \sqcup \mathrm{Edge}^{{\color{gray}\blacktriangledown}}(\Delta) \sqcup \mathrm{Edge}^{\triangledown}(\Delta)$$

into unforgettable, weighted, and forgettable edges, respectively.[1] Finally, we require that if $h(e) = t(e')$ is the unique point of a breaking then $\rho(e) = \rho(e')$.

We correspondingly enhance the combinatorial type $\Gamma = \Gamma(\Delta)$ of a weighted treed disk Δ in order to record the "weight class" of each of its edges; we remember each of the collections $\mathrm{Edge}^{\blacktriangledown}(\Delta)$, $\mathrm{Edge}^{{\color{gray}\blacktriangledown}}(\Delta)$, and $\mathrm{Edge}^{\triangledown}(\Delta)$ and associate them to Γ.

Definition 2.8. When the root edge e_0 is unweighted, an *isomorphism of weighted treed disks* $\psi : \underline{\Delta} \to \underline{\Delta}'$ is simply an isomorphism of treed disks which preserves edge weights and the collection of weighted edges. If instead e_0 is weighted, then rather than demanding that ψ preserves weights on-the-nose, we instead require that there there exists $c \in (0, \infty)$ such that $\rho(\psi(e)) = c \cdot \rho(e)$ for each $e \in \mathrm{Edge}(\Delta)$.

Definition 2.9. A weighted treed disk Δ is *stable* if both the domain constraints:

[1] To ease interoperability between texts, we adopt the general convention of [CW22; VWX20] that data is "fully active" when its corresponding weight function reports 0, and that data "fades out" as its weight increases until becoming fully disabled at ∞.

- Each sphere component $v \in \mathrm{Vert}^\circ(\Delta)$ contains at least 3 joints.
- Each disk component $v \in \mathrm{Vert}^\bullet(\Delta)$ contains at least 3 boundary joints, or at least 1 boundary joint and 1 interior joint.
- No point $v^+ \in \mathrm{Vert}^+(\Delta)$ is adjacent in the tree to any breaking.

and weight constraints:

- If the root edge e_0 is weighted then there exists at least one weighted input.
- If there are no disk components, hence no sphere components (c.f. Remark 2.2), then $\mathrm{Edge}(\Delta) = \{e_0, e_1\}$ with e_0 unweighted and e_1 weighted.

are satisfied.

2.2 Moduli spaces of treed disks

One of the great advantages of our present setting is that the universal treed disks for the moduli spaces we will consider are themselves honest treed disks (in contrast to other cases where it would be necessary to deal with orbifold singularities) [VWX20, Section 2.2].

Definition 2.10. The *universal treed disk of type* Γ is the space $\mathcal{U}_\Gamma$ consisting of isomorphism classes of pairs (Δ, x) where Δ is a stable weighted treed disk with combinatorial type Γ and $x \in \underline{\Delta}$ is a distinguished point. We write $[\Delta, x]$ for such an isomorphism class.

Forgetting the distinguished point and retaining only the isomorphism class of Δ defines a map $[\Delta, z] \mapsto [\Delta]$ from $\mathcal{U}_\Gamma$ onto $\mathcal{M}_\Gamma$, the *moduli space of stable weighted treed disks of type* Γ.

By construction the moduli space $\mathcal{M}_\Gamma$ of isomorphism classes of stable weighted treed disks of combinatorial type Γ has dimension

$$\begin{aligned}\dim \mathcal{M}_\Gamma = n(\Gamma) - 2 + 2\sum_{\kappa\in\mathcal{K}} \#\,\mathrm{Mark}_\kappa(\Gamma([u])) - \#\,\mathrm{Edge}^\circ_-(\Gamma) - \#\,\mathrm{Joint}^\wedge(\Gamma) \\ \underbrace{+\,\#\,\mathrm{Edge}^\triangledown(\Gamma)}_{\substack{\text{free weighting}\\ \text{parameter}}}\ \underbrace{-\,\#\,\mathrm{Edge}^0_-(\Gamma)}_{\substack{\text{no length}\\ \text{parameter}}} \qquad (2.2.\mathrm{I})\end{aligned}$$

whenever $e_0 \in \mathrm{Edge}(\Gamma)$ is unweighted. If $e_0 \in \mathrm{Edge}(\Gamma)$ is weighted, then the dimension is two less.

Remark 2.11. These (unweighted) spaces have a natural cell structure closely related to Stasheff associhedra [CW22], for which orientations may be readily constructed by hand.

The moduli spaces $\mathcal{M}_\Gamma$ are related by maps $\Gamma' \to \Gamma$ induced by degenerations of the combinatorial type Γ in various senses. In each case, "un-doing" a degeneration of a weighted treed disk Δ' of type Γ' produces a treed disk Δ of type Γ, and a corresponding map $\Gamma' \to \Gamma$ of the edges and vertices of the underlying combinatorial types in an obvious way (this map is usually surjective). We enumerate these below, and each is correspondingly depicted in Figure 2.2 or Figure 2.3.

(2.2.a) *An edge breaking*: Let $e \in \mathrm{Edge}(\Gamma)$ be an edge with $l(e) \in (0, \infty)$. The type Γ' is obtained from Γ by deleting the edge e, creating a new (breaking) vertex $v^+ \in \mathrm{Vert}^+(\Gamma')$, and creating two new edges $e_{\mathrm{in}}, e_{\mathrm{out}} \in \mathrm{Edge}(\Gamma')$ defined by $\underline{t}(e_{\mathrm{in}}) = \underline{t}(e)$, $\underline{h}(e_{\mathrm{in}}) = v^+ = \underline{t}(e_{\mathrm{out}})$, and $\underline{h}(e_{\mathrm{out}}) = \underline{h}(e)$, necessarily with lengths $l(e_{\mathrm{in}}) = l(e_{\mathrm{out}}) = \infty$.

(2.2.b) *A node spawning*: Let $e \in \mathrm{Edge}(\Gamma)$ be an edge with $l(e) \in (0, \infty)$. The type Γ' is obtained from Γ by declaring $l(e) := 0$ (and thereby creating a node).

(2.2.c) *A component bubbling off*: Let $v \in \mathrm{Vert}^\bullet(\Gamma) \cup \mathrm{Vert}^\circ(\Gamma)$ be a disk or sphere component. The type Γ' is obtained from Γ by creating a new sphere component, or possibly a disk component if v is a disk component, $v' \in \mathrm{Vert}^\bullet(\Gamma') \cup \mathrm{Vert}^\circ(\Gamma')$ and new edge $e \in \mathrm{Edge}(\Gamma')$ with $\underline{t}(e) = v'$, $\underline{h}(e) = v$, and $l(e) = 0$.

(2.2.d) *An edge becomes forgettable or unforgettable*: Let $e \in \mathrm{Edge}(\Gamma)$ be an edge with $\rho(e) \in (0, \infty)$. The type Γ' is obtained from Γ by declaring either $\rho(e) := \infty$ (forgettable) or $\rho(e) := 0$ (unforgettable).

(2.2.e) *Forgetting an input*: Let $v^+ \in \mathrm{Vert}^+(\Gamma)$ be an input, let $e \in \mathrm{Edge}(\Gamma)$ be the unique edge with $\underline{t}(e) = v^+$, and suppose that $\rho(v^+) = \infty$. The type Γ' is obtained from Γ by deleting v^+ and e.

After performing the operation (2.2.e) on a type Γ it may be the case that the resulting type Γ' is no longer stable. In this case we will replace Γ' with its *stabilization* $(\Gamma')^{\mathrm{stab}}$, formed by recursively collapsing unstable sphere or disk components and breakings (in the sense of Definition 2.9) into their parent wherever this is possible; whenever edges are collapsed together, their lengths add. The resulting type is again stable, and there is a map of vertex and edge sets $\Gamma' \to (\Gamma')^{\mathrm{stab}}$. In summary, the result of forgetting an input of a type Γ and then stabilizing is in general related to the original type Γ by a cospan $(\Gamma')^{\mathrm{stab}} \leftarrow \Gamma' \to \Gamma$ of maps of vertex and edge sets.

Together the operations (2.2.a)–(2.2.e), possibly after stabilizing, generate a partial order on stable combinatorial types by complexity; we declare $(\Gamma')^{\mathrm{stab}} > \Gamma$ whenever there is an operation with corresponding morphism $\Gamma' \to \Gamma$, and take the transitive closure.

Remark 2.12. At this point it is also convenient to note a closely related operation, which has a more algebraic rather than topological flavor; an *edge cut*.

Let Γ be a combinatorial type with $v^+ \in \mathrm{Vert}^+(\Gamma)$ a breaking. Partition $\mathrm{Vert}(\Gamma) - \{v^+\}$ into the subcollections V_{v^+} and V_{-v^+} of vertices which according to the orientation of

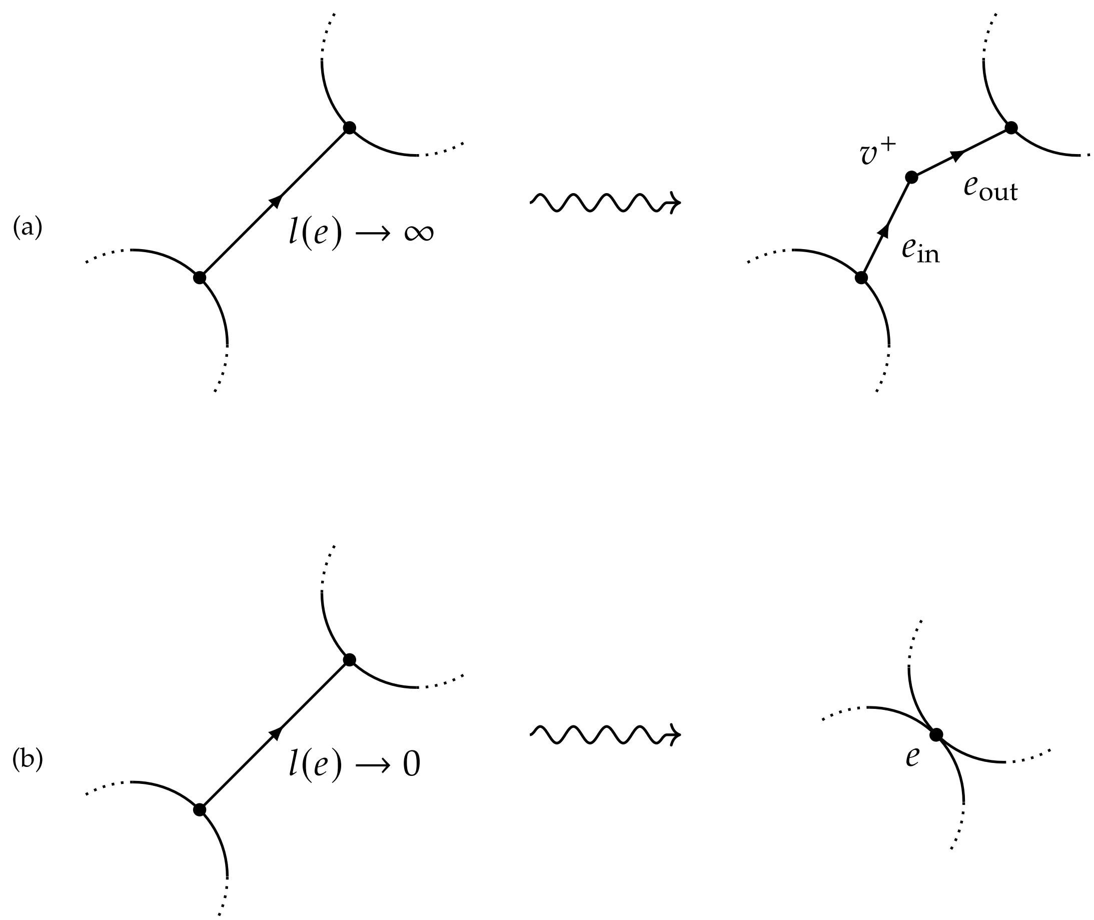

Figure 2.2: Schematic diagrams of the enumerated degenerations (2.2.a)-(2.2.b) of combinatorial types.

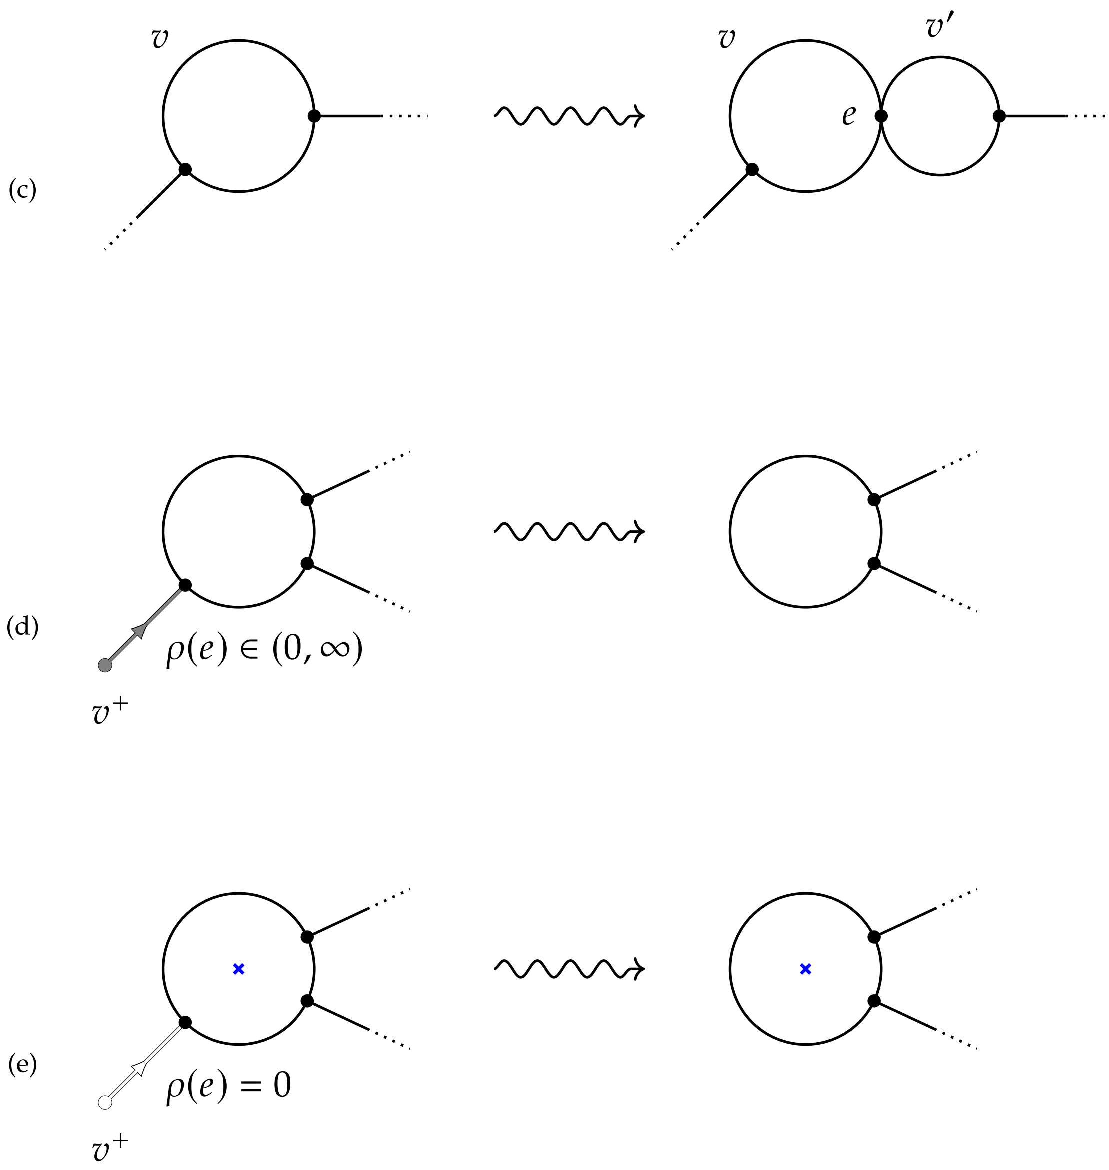

Figure 2.3: Schematic diagrams of the enumerated degenerations (2.2.c)-(2.2.e) of combinatorial types. A distinguished marked point is labeled by a cross.

Γ are, and are not, children of v^+ respectively. Define types Γ_{v^+} and Γ_{-v^+} by declaring $\mathrm{Vert}(\Gamma_v) := \{v^+\} \cup V_{v^+}$ and $\mathrm{Vert}(\Gamma_{-v}) := \{v^+\} \cup V_{-v^+}$, and restricting the head, tail, flavor, length, and weight maps of the original type Γ in each case. We obtain a surjective morphism $\Gamma \to \Gamma_v \sqcup \Gamma_{-v}$ which will be fundamental when subsequently considering the Morse–Fukaya algebra's composition operations.

The compactification $\overline{\mathcal{M}}_\Gamma$ assembles as

$$\overline{\mathcal{M}}_\Gamma = \bigcup_{\substack{\Gamma \leq \Gamma' \\ n(\Gamma) = n(\Gamma')}} \mathcal{M}_{\Gamma'}, \tag{2.2.II}$$

i.e. as the union of all strata corresponding to types Γ' of complexity greater than or equal to that of Γ, and with the same number of inputs. In other words, the union is taken allowing Γ' to vary over all types obtained from Γ by applying the operations (2.2.a)–(2.2.d) and not (2.2.e). This is a finite union; all operations (2.2.a)–(2.2.e) decrease dimension as computed by (2.2.I). By the natural extension of Theorem 1.11 to this setting, there is a Gromov topology on the compactification $\overline{\mathcal{M}}_\Gamma$ for which each inclusion $\mathcal{M}_{\Gamma'} \hookrightarrow \overline{\mathcal{M}}_\Gamma$ of each $\mathcal{M}_{\Gamma'}$ in (2.2.II) is an embedding. Similarly, there is in turn a universal treed disk $\overline{\mathcal{U}}_\Gamma \to \overline{\mathcal{M}}_\Gamma$ which is compatibility topologized.

In order to define a perturbation datum it will be useful to distinguish the surface- and tree-parts of $\overline{\mathcal{U}}_\Gamma$; the (not necessarily disjoint) decomposition $\underline{\Delta} = S_\Delta \cup T_\Delta$ gives rise to subspaces

$$\overline{\mathcal{S}}_\Gamma = \bigcup_{\substack{[\Delta] \in \overline{\mathcal{M}}_\Gamma \\ z \in S_\Delta}} [\Delta, z] \quad \text{and} \quad \overline{\mathcal{T}}_\Gamma = \bigcup_{\substack{[\Delta] \in \overline{\mathcal{M}}_\Gamma \\ z \in T_\Delta}} [\Delta, z]$$

which of course have union $\overline{\mathcal{U}}_\Gamma$. Once and for all we fix a compact subset $\mathcal{S}_\Gamma^{\circledast} \subset \overline{\mathcal{S}}_\Gamma$ on which to perturb the almost complex structure, which has interior meeting (i.e. having nonempty intersection with) each disk and sphere component, and is disjoint from $\mathcal{T}_\Gamma$ and $\partial\mathcal{S}_\Gamma$. Likewise in order to implement perturbations of the global Morse function choose a compact $\mathcal{T}_\Gamma^\bullet \subset \mathcal{T}_\Gamma$ which has interior meeting the interior of each edge and is disjoint from a neighborhood of all vertices (points and disk/sphere components).

Remark 2.13. Each degeneration (2.2.a)–(2.2.d) of a combinatorial type Γ into a type $\Gamma' > \Gamma$ induces a corresponding restriction map of compact moduli spaces: edges breaking, components bubbling off, nodes or infinite edges spawning, and edges becoming forgettable/unforgettable each correspond to an inclusion of a boundary stratum $\overline{\mathcal{M}}_{\Gamma'} \hookrightarrow \overline{\mathcal{M}}_\Gamma$ of codimension 1 or 2. The edge cutting operation $\Gamma \to \Gamma_{v^+} \sqcup \Gamma_{-v^+}$ of Remark 2.12 similarly gives rise to an isomorphism of moduli spaces $\overline{\mathcal{M}}_\Gamma \cong \overline{\mathcal{M}}_{\Gamma_{v^+}} \times \overline{\mathcal{M}}_{\Gamma_{-v^+}}$. The universal treed disk corresponding decomposes as the union $p_+^*\overline{\mathcal{U}}_{\Gamma_{v^+}} \cup p_-^*\overline{\mathcal{U}}_{\Gamma_{-v^+}}$ of pullbacks along the respective projections $p_\pm$ of $\overline{\mathcal{M}}_\Gamma$ onto $\overline{\mathcal{M}}_{\Gamma_{v^+}}$ and $\overline{\mathcal{M}}_{\Gamma_{-v^+}}$.

If instead Γ' is obtained from Γ by forgetting an input $e \in \mathrm{Edge}(\Gamma)$ as in (2.2.e), there is a projection $p : \overline{\mathcal{M}}_\Gamma \to \overline{\mathcal{M}}_{\Gamma'}$ fibered over the closed interval parameterizing all possible positions of $h(e)$ consistent with the order of the input edges of Γ. Correspondingly, there is an inclusion of the pullback $p^*\overline{\mathcal{U}}_{\Gamma'}$ into $\overline{\mathcal{U}}_\Gamma$.

2.3 Perturbed pseudoholomorphic maps

In this section fix a Kähler manifold $(\overline{X}, \omega, J)$ equipped with a (possibly singular) Lagrangian torus fibration $\pi : \overline{X} \to \overline{Q}$. Let $Q \subset \overline{Q}$ be a compact subset disjoint from the critical values of π. We write $F_q := \pi^{-1}(q)$ for each fiber. Given an underlying treed disk Δ, a pseudoholomorphic treed disk is a particular kind of map $u : \underline{\Delta} \to \overline{X}$. It will not be much use for us to consider ordinary, unperturbed, pseudoholomorphic treed disks because of transversality issues which obstruct a well-defined theory of their counts. Instead, we build into their definition a global scheme for making suitable domain-dependent perturbations of the pseudoholomorphic curve equation and Morse gradient flow equation. We begin by specifying the perturbation datum to be associated to each fixed type.

Definition 2.14. Let Γ be a combinatorial type of weighted treed disks. A *perturbation datum P_Γ for Γ modeled on (J, f)* is a choice of domain-dependent perturbations:

(2.14.a) Of the almost complex structure J, given by a smooth map

$$J_\Gamma : \overline{\mathcal{S}}_\Gamma \to \mathcal{J}_\tau(\overline{X}, \omega),$$

which restricts to J on $\overline{\mathcal{S}}_\Gamma - \mathcal{S}_\Gamma^{\circledast}$.

(2.14.b) Of the Morse function f, given by a smooth map

$$f_\Gamma : \overline{\mathcal{T}}_\Gamma \to C^\infty(X \to \mathbb{R}),$$

which restricts to f on $\overline{\mathcal{T}}_\Gamma - \mathcal{T}_{\Gamma,\circ}^{\circledast}$.

(2.14.c) Which are each local, satisfying a technical gluing condition which we defer to Section 2.4 (c.f. Definition 2.42), in which we prove that satisfactory choices of perturbation data actually exist.

Definition 2.15. Let $\mathbf{\Gamma}$ be a collection of combinatorial types such that $\Gamma \in \mathbf{\Gamma}$ and $\Gamma' > \Gamma$ implies $\Gamma' \in \mathbf{\Gamma}$, and whenever $\Gamma \in \mathbf{\Gamma}$ and $v^+ \in \mathrm{Vert}^+(\Gamma)$ is a breaking then $P_{\Gamma_{v^+}}, P_{\Gamma_{-v^+}} \in \mathbf{\Gamma}$.

A family $\mathbf{P} = \{P_\Gamma\}_{\Gamma \in \mathbf{\Gamma}}$ of perturbation data chosen for each combinatorial type $\Gamma \in \mathbf{\Gamma}$ is a *perturbation system* if (recalling the induced maps of moduli spaces of Remark 2.13):

(2.15.a) The family respects degenerations—if $\Gamma' > \Gamma$ and $\iota : \overline{\mathcal{M}}_{\Gamma'} \hookrightarrow \overline{\mathcal{M}}_\Gamma$ is the induced inclusion of a boundary stratum, the restriction $i^* P_\Gamma$ of P_Γ along ι is equal to $P_{\Gamma'}$.

(2.15.b) The family respects cuts—if $v_+ \in \mathrm{Vert}^+(\Gamma)$ is a breaking, the induced edge cut isomorphism $\overline{\mathcal{M}}_\Gamma \cong \overline{\mathcal{M}}_{\Gamma_{v^+}} \times \overline{\mathcal{M}}_{\Gamma_{-v^+}}$ realizes P_Γ as the product $P_{\Gamma_{v^+}} \times P_{\Gamma_{-v^+}}$.

(2.15.c) The family contains all lower-dimensional types—if Γ' is any type with[2] $n(\Gamma') < n(\Gamma)$, or if $n(\Gamma') = n(\Gamma)$ and $\#\,\mathrm{Mark}(\Gamma') < \#\,\mathrm{Mark}(\Gamma)$, then $\Gamma' \in \mathbf{\Gamma}$.

Definition 2.16. Let $P_\Gamma = (J_\Gamma, f_\Gamma)$ be a perturbation datum for a stable type Γ. A *P_Γ-perturbed pseudoholomorphic* treed disk is a weighted treed disk Δ along with a continuous map $u : \underline{\Delta} \to \overline{X}$ which:

(2.16.a) Is pseudoholomorphic on the surface part, in that u restricts to a map $u : S_\Delta \to \overline{X}$ such that

$$\mathrm{D}_z u \circ j_z = J_\Gamma([\Delta, z])_{u(z)} \circ \mathrm{D}_z u \quad \text{for all } z \in \mathrm{int}(S_\Delta)$$

with j the canonical complex structure on S_Δ.

(2.16.b) Is a Morse gradient flow trajectory on the boundary tree part, in that for each $e \in \mathrm{Edge}(\Delta)$ with $t(e)$ (hence $h(e)$, c.f. the boundary axiom (2.1.e)) lying on a point or the boundary of a disk component and recalling the corresponding corresponding $L_e \subset [-\infty, \infty]$, the map u restricts along the inclusion $\iota : L_e \hookrightarrow T_\Delta$ such that

$$(u \circ \iota)'(t) = \nabla^g_{u(\iota(t))} f_\Gamma([\Delta, \iota(t)]) \quad \text{for all } t \in \mathrm{int}(L_e).$$

(2.16.c) Has each disk boundary constrained to a Lagrangian fiber, in that there is a function $b^u : \pi_0(\partial S_\Delta - \mathrm{Joint}(\Delta)) \to \{F_q : q \in Q\}$ such that

$$u(z) \in b^u([z]) \quad \text{for all } z \in \partial S_\Delta - \mathrm{Joint}(\Delta).$$

Remark 2.17. The connected components of $\partial S_\Delta - \mathrm{Joint}(\Delta)$ are exactly the boundary arcs of each disk component of S_Δ as punctured by the respective head and tail points of the incident edges. Of course, by continuity of $u : \underline{\Delta} \to \overline{X}$ the map b^u of (2.16.c) lifts to a locally constant function defined on all of ∂S_Δ, i.e. so that the entirety of each

[2]Recall that $n(\Gamma)$ denotes the number of inputs of the type Γ.

disk boundary is sent by u into the same fiber of π. However, the present somewhat-redundant formulation will prove convenient in the sequel, when it will be necessary to extend the collection $\{F_q : q \in Q\}$ to contain other Lagrangian submanifolds of X (c.f. Definition 4.4).

Remark 2.18. Let $u : \underline{\Delta} \to \overline{X}$ be a pseudoholomorphic treed disk. As a consequence of the axioms of a treed disk, if $v^\circ \in \mathrm{Vert}^\circ(\Delta)$ is a sphere component, upon iteratively replacing v° with its parent $\underline{(}h)(v^\circ)$ we must eventually obtain a disk component $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$. By Remark 2.17, under u the component $v^\bullet$ bounded by a unique fiber F_q; despite not having any boundary itself, we say that v° is *bounded* by the same fiber F_q.

Definition 2.19. A (P_Γ-perturbed) pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ is *stable* when

- if u is constant on a sphere component $v \in \mathrm{Vert}^\circ(\Delta)$ then v contains at least 3 joints,
- if u is constant on a disk component $v \in \mathrm{Vert}^\bullet(\Delta)$ then v contains at least 1 boundary joint and 1 interior joint, and
- if $e \in \mathrm{Edge}_\rightarrow(\Delta)$ is an infinite edge then u is not constant on L_e.

In order to overcome the multiple-cover problem and stabilize our moduli spaces of stable pseudoholomorphic treed disks we adapt the scheme of Cieliebak–Mohnke [CM07] to our setting. The key ingredient is the following theorem, which is our analogue of [VWX20, Lemma 2.10]—though in the present setting where X is Kähler the result essentially goes back to [Gue99].

Theorem 2.20. *Suppose that the class $[\omega]$ is integral. For each finite collection $\mathcal{L} = \{L_i \subset X\}$ of rational Lagrangian submanifolds*[3] *of X and all $k \gg 1$ there exists a codimension 2 complex*

[3]A Lagrangian submanifold $L \subset X$ is *rational* if the class $[\omega]$ is rational on $\pi_2(X, L)$, i.e. $[\omega] \in H^2(X, L; \mathbb{Q})$.

hypersurface $D \subset X$ disjoint from each L_i and of degree k such that

- *every nonconstant holomorphic disk in X with boundary contained in $\bigcup \mathcal{L}$ intersects D at least once,*
- *every nonconstant holomorphic sphere intersects D at least three times, and*
- *the collection $\mathcal{L}$ is jointly exact on the complement $X - D$.*

Proof. The proof is an application of Donaldson [Don96] and Auroux–Gayet–Mohsen's [AGM01] theory of stabilizing divisors. □

Definition 2.21. A hypersurface D satisfying the conclusions of Theorem 2.20 for the family $\mathcal{L} = \{L_i\}$ is called a *stabilizing divisor for $\mathcal{L}$*.

Corollary 2.22. *Let D be a stabilizing divisor for the family $\mathcal{L}$ furnished by Theorem 2.20. If $F_q \in \mathcal{L}$ for some $q \in Q$ then there exists an open neighborhood $U_q \subset Q$ of q such that D is also a stabilizing divisor for F_p for all $p \in U_q$.*

Proof. By compactness of X we may find a neighborhood $U_q \subset Q$ small enough that, for all $p \in U_q$, the divisor D is disjoint from all F_p and the fibers F_p and F_q are isotopic by some isotopy ψ_p avoiding D. Now, given $u : \mathbb{D} \to \overline{X}$ any holomorphic disk in X with boundary on F_p for some $p \in U_q$, the boundary $u(\partial\mathbb{D}) \subset F_p$ is isotopic via ψ_p to a subset of F_q. This yields a (no longer holomorphic) map $u' : \mathbb{D} \to \overline{X}$ bounded by F_q; suitably further restricting U_q we may ensure that all such maps u' obtained from this process have positive symplectic area. Therefore the intersection number of u' (hence u) with D is positive, i.e. u meets D in at least one point. □

Remark 2.23. In the sequel we will find that it is undesirable that, in the setup of Corollary 2.22, there may exist nonconstant pseudoholomorphic spheres u which meet the divisor D in infinitely many places; namely, u may be contained in D. In order to prohibit this, we make the following compromise: suppose that finitely many families $\mathcal{L}_\kappa$ and corresponding respective stabilizing divisors D_κ have been chosen. Then we may deform the complex structure on X into an almost complex structure J_0 for which the each D_κ is still stabilizing, is now a J_0-holomorphic hypersurface, and is regular for the moduli space of simple J_0-spheres. If the hypersurface D_κ satisfy an a priori minimum bound on their degree, the expected dimension of the moduli space of such J_0-spheres in each D_κ is negative, and therefore each such space is empty.

Definition 2.24. A *system of stabilizing divisors for* $\pi : X \to Q$ is a family of divisors $\mathbf{D} = \{D_\kappa \subset X\}$, equipped with *activity* functions $\alpha_\kappa : Q \to [0, \infty]$ such that:

(2.24.a) Each divisor D_κ is a stabilizing divisor for all Lagrangian fibers F_q with $q \in \alpha_\kappa^{-1}([0, \infty))$.

(2.24.b) For each $q \in Q$ there is at least one D_κ such that $\alpha_\kappa(q) < 1$ and there are only finitely many divisors $D_{\kappa'}$ such that $\alpha_{\kappa'}(q) < \infty$.

Theorem 2.25. *There exists a system of stabilizing divisors for* $\pi : X \to Q$.

Proof. By Corollary 2.22, for each fiber F_q of $\pi : X \to Q$ there exists an open neighborhood $U_q \subset Q$ of q and divisor $D_q \subset X$ such that D_q is stabilizing for each $F_{q'}$ with $q' \in U_q$. For each $q \in Q$ choose a smooth function $\alpha_q : Q \to [0, \infty]$ such that $\alpha_q(q) = 0$ and $\alpha_q|_{Q-U_q} = \infty$ identically. There is a finite subset $S \subset Q$ such that $\{\alpha_q^{-1}([0, \infty)) : q \in S\}$ covers Q.

Let $\eta_R(t) : \mathbb{R} \times [0, \infty] \to [0, \infty]$ be a smooth family of monotone reparameterizations of $[0, \infty]$ such that $\eta_R(\infty) = \infty$ and $\eta_R|_{[0,R]} = 0$ identically. For $R > 0$ sufficiently large it must be the case that the set of divisors $\{D_q : q \in S\}$ equipped with the respective activity functions $\alpha_q \circ \eta_R$ have the desired properties. □

Henceforth, once we have fixed a system of stabilizing divisors $\mathbf{D}$ and an almost complex structure J_0 as in Remark 2.23, we will set $\mathcal{K} := \{\kappa : D_\kappa \in \mathbf{D}\}$ as the universe of all possible flavors. In addition, it will be desirable to equip marked points of our pseudoholomorphic treed disks with *multiplicity* labels, so that $d_{\kappa,i} \in \mathrm{Mark}_\kappa(\Delta)$ meets D_κ with multiplicity $m_{d_{\kappa,i}} \in \mathbb{N}_{\geq 1}$.

We are about to specify the classes of pseudoholomorphic treed disks which we actually intend to count. It is necessary to impose constraints to both stabilize the treed disks and to enforce correct boundary conditions. We handle each of these in turn; first, it will be convenient to restrict our considerations to a class of Morse functions on X which are particularly compatible with the fibration $\pi : X \to Q$.

Definition 2.26. A Morse function $f : X \to \mathbb{R}$ is a *perfect lift* of a Morse function $\breve{f} : Q \to \mathbb{R}$ if

- we have $\pi(\mathrm{crit}\, f) = \mathrm{crit}\, \breve{f}$, and
- for all $q \in \mathrm{crit}\, \breve{f}$ we have that the gradient flow of f is tangent to all points of F_q and f restricted to F_q is perfect[4].

A perfect lift f of any Morse function $\breve{f} : Q \to \mathbb{R}$ exists by lifting $\breve{f}$ to X on-the-nose and then locally making small perturbations.

Each Morse function $\breve{f}$ determines a cellular decomposition of Q, for which each $q \in \mathrm{crit}_k\, \breve{f}$ corresponds to the unique $(\dim Q - k)$-cell equal to the ascending manifold $W^{\uparrow}(q)$.

Definition 2.27. A polyhedral cover Θ of Q is *confining* for a perfect lift f of a Morse function $\breve{f} : Q \to \mathbb{R}$ if each $q \in \mathrm{crit}_k\, \breve{f}$ corresponds to a unique $\theta_q \in \Theta$ containing the star of the cell $W^{\uparrow}(q)$ in the cellular decomposition induced by $\breve{f}$.

Figure 2.4 depicts Morse flow trajectories in Q superimposed onto a confining polyhedral cover of Q.

Definition 2.28. Let $\mathbf{D}$ be a system of divisors and let f be a perfect lift of a Morse function $\breve{f} : Q \to \mathbb{R}$. For each $\theta \in \Theta$ let $\mathbf{D}|_\theta \subset \mathbf{D}$ be the subset consisting of all divisors

[4]A Morse function is *perfect* if it induces a minimal Morse model—in this case, of each n-torus fiber. Equivalently, the Morse differential ∂ is zero on the entire Morse complex.

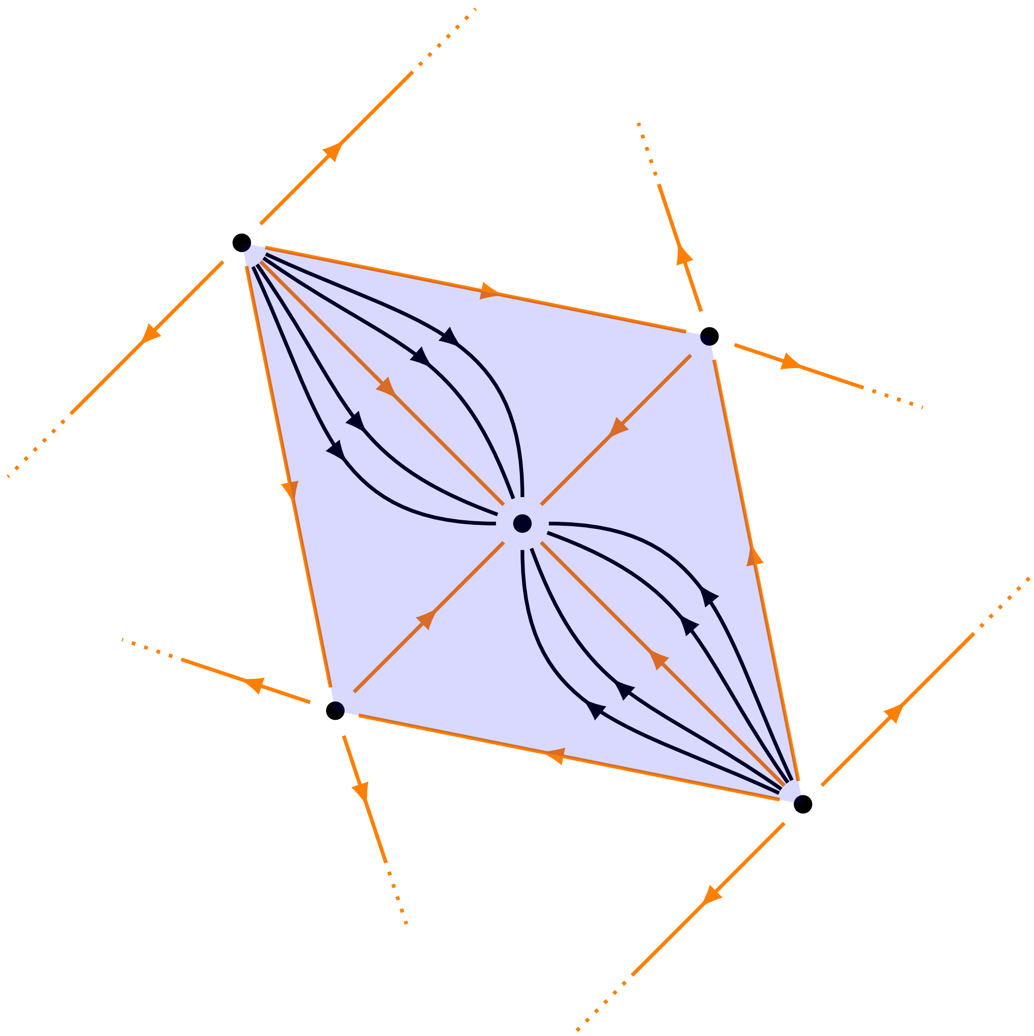

Figure 2.4: A representative diagram depicting the Morse flow trajectories for $\check{f}$ in Q, along with an indicative cell of a confining polyhedral cover.

D_κ for which there exists some $q \in \theta$ such that $\alpha_\kappa(q) < \infty$. We call a system of divisors $\mathbf{D}$ and perfect lift f *compatible* if for all $\theta \in \Theta$ and $D_\kappa \in \mathbf{D}|_\theta$ the divisor D_κ is stabilizing for each $q \in \theta$.

The collections $\mathbf{D}|_\theta \subset \mathbf{D}$ record the divisors which are active somewhere in θ. Compatibility (as in Definition 2.28) in particular ensures for us that it is safe to track all points at which a pseudoholomorphic treed disk component in θ meets each $D_\kappa \in \mathbf{D}|_\theta$, in the sense that (ignoring constant components) the number of intersection points with D_κ is necessarily finite.

Henceforth we fix a particular perfect lift f of a Morse function $\breve{f}$ on Q, with $\breve{f}$ chosen to be outward transverse to the boundary of Q, equipped with a corresponding confining polyhedral cover Θ of Q. We will assume that the systems of divisors $\mathbf{D}$ we consider are always compatible with f.

Definition 2.29. A (P_Γ-perturbed) pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ is *adapted* to a system of divisors $\mathbf{D} = \{D_\kappa\}$ compatible with f if, whenever $v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)$ is a component bounded (in the sense of Remark 2.18) by a fiber lying over some $\theta \in \Theta$, we have that $\mathrm{Mark}_\kappa(\Delta) \cap v$ intersects each connected component of $u|_v^{-1}(D_\kappa)$ for all $D_\kappa \in \mathbf{D}|_\theta$.

Remark 2.30. In Definition 2.29 we, roughly speaking, record all intersection points of a disk or sphere component bounded by a fiber F_q with all divisors active anywhere in any cell $\theta \in \Theta$ containing q. This is a convenient bookkeeping device; however, in order to achieve stability of our treed disks and coherent perturbations it is actually only necessary to record intersections with divisors D_κ having $\alpha_\kappa(q) < \infty$.

In this equivalent formulation, we would declare two pseudoholomorphic treed disks $u, u' : \underline{\Delta} \to \overline{X}$ adapted to $\mathbf{D}$ equivalent whenever they become isomorphic upon deleting all marked points $d \in \mathrm{Mark}_{(\kappa,m)}(\Delta)$ with $\alpha_\kappa(\pi(d)) = \infty$ (i.e. wherever the underlying divisor is disabled).

Modulo a technical modification the valid possible input or output points of pseudoholomorphic treed disks—which will be promoted to generators of the Morse–Fukaya algebra subsequently—are the critical points of the global Morse function f. More precisely, write $x^{\blacktriangledown}_{\mathsf{m},i} \in X$ for the ith minimum of f under a fixed choice of ordering. For each such critical point let $x^{\color{gray}\blacktriangledown}_{\mathsf{m},i}$ and $x^{\triangledown}_{\mathsf{m},i}$ be a pair of additional formal generators, and correspondingly extend the definition of index by declaring $I(x^{\color{gray}\blacktriangledown}_{\mathsf{m},i}) = -1$ and $I(x^{\triangledown}_{\mathsf{m},i}) = 0$. Now define the set of generators

$$\operatorname{gen} f := \operatorname{crit} f \sqcup \bigcup_i \{x^{\color{gray}\blacktriangledown}_{\mathsf{m},i}, x^{\triangledown}_{\mathsf{m},i}\}. \tag{2.3.I}$$

These formal generators will be used to implement a strict unit for the Morse–Fukaya algebra, with each formal generator $x^{\color{gray}\blacktriangledown}_{\mathsf{m},i}$ of index -1 witnessing the homological equality of $x^{\blacktriangledown}_{\mathsf{m},i}$ and $x^{\triangledown}_{\mathsf{m},i}$ (c.f. Theorem 3.12).

Definition 2.31. Let Γ be a combinatorial type with 1 output and n inputs, with marked points labeled with flavors determined by a system of divisors $\mathbf{D} = (D_\kappa)_{\kappa\in\mathcal{K}}$. Recall that by construction the inputs $v_i \in \operatorname{Vert}^+(\Gamma)$ are canonically ordered, that there is a the root edge e_0 with $h(e_0) = v_0$, and that for each $i > 0$ there is a unique $e_i \in \operatorname{Edge}(\Gamma)$ such that $t(e_i) = v_i$.

A pseudoholomorphic treed disk *spec(ification)* for the type Γ is a choice of tuple $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ consisting of:

- Multiplicity labels $\mathbf{m} = (m_{\kappa,d} \in \mathbb{N}_{\geq 1})_{\kappa\in\mathcal{K}, d\in\operatorname{Mark}_\kappa(\Delta)}$.
- Input/output labels $\mathbf{x} = (x_i \in \operatorname{gen} f)_{0\leq i\leq n}$ satisfying

$$x_i \in \begin{cases} \operatorname{crit} f & \rho(e_i) = 0 \\ \{x^{\color{gray}\blacktriangledown}_{\mathsf{m}}\} & \rho(e_i) \in (0,\infty) \\ \{x^{\triangledown}_{\mathsf{m}}\} & \rho(e_i) = \infty \end{cases}.$$

- Boundary classes $\boldsymbol{\beta} = (\beta_v)_{v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)}$ satisfying[5]

$$\beta_v \in \begin{cases} H_2(X, \bigcup b^u(v)) & v \in \mathrm{Vert}^\bullet(\Delta) \\ H_2(X) & v \in \mathrm{Vert}^\circ(\Delta) \end{cases}.$$

A pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ with $\Gamma = \Gamma(\Delta)$ and adapted to $\mathbf{D}$ *obeys the spec* $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ if:

- For each $\kappa \in \mathcal{K}$ and $d \in \mathrm{Mark}_\kappa(\Delta)$, if $u(d)$ does not lie on a constant component then $u(d)$ meets D_κ with order of tangency $m_{\kappa,d}$.

- For the unique output $v_0 \in \underline{\Delta}$ and inputs $v_i \in \underline{\Delta}$ we have $u(v_i) = x_i$ for all $0 \le i \le n$.

- For each component $v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)$ we have that $u|_v$ represents the homology class β_v.

Definition 2.32. Given a combinatorial type Γ labeled by a background system of divisors $\mathbf{D}$, fix a perturbation datum P_Γ. For each spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ we may form the *moduli space* $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *of stable* P_Γ*-perturbed pseudoholomorphic treed disks with domain combinatorial type* Γ *of spec* $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$.

The expected dimension of $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is computed in (2.4.II) via a Fredholm index formula. Substituting the (essentially combinatorial) domain dimension formula (2.2.I), and recalling the Morse and Maslov indices and first Chern numbers of

[5] Recall the locally constant function b^u of (2.16.c). As above, this formulation is currently somewhat redundant in that the union $\bigcup b^u(v)$ must be equal to the unique fiber F_q such that $u(\partial v) = \{q\}$. When the function b^u is permitted to take values in other Lagrangian submanifolds of X in Chapter 4 this will no longer be the case.

Chapter 1, we arrive at the expected dimension

$$\begin{aligned}\dim \mathcal{M}_\Gamma(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) = I(x_0) - \sum_{i=1}^{n} I(x_i) + \sum_{v \in \mathrm{Vert}^\bullet(\Gamma)} I(\beta_v) + 2 \sum_{v \in \mathrm{Vert}^\circ(\Gamma)} c_1(\beta_v) \\ + n(\Gamma) - 2 - 2 \sum_{d \in \mathrm{Mark}_\kappa(\Gamma)} (m_{\kappa,d} - 1) \\ - \#\,\mathrm{Edge}_-^\circ(\Gamma) - \#\,\mathrm{Joint}^\wedge(\Gamma) - \#\,\mathrm{Edge}_-^0(\Gamma)\end{aligned} \tag{2.3.II}$$

whenever the root edge $e_0 \in \Gamma$ is unweighted. When e_0 is weighted, the dimension $\dim \mathcal{M}_\Gamma(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is two less.

2.4 Regularization and uncrowding

In this section we define the specific class of perturbation systems which enjoy all of the necessary technical properties required to successfully construct the Morse–Fukaya algebra, and then prove that they exist and are plentiful.

Before handling the requisite functional analysis, we must first overcome a general technical problem which arises in our approach (adapting the methods of [CW22] and as originally formulated in [CM07]) when a pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ meets a particular stabilizing divisor D_κ at a component $v^\circ \in \mathrm{Vert}^\bullet(\Delta)$ on which u is constant; according to Definition 2.19 the map u may be stable with an arbitrarily large number of marked points $d_{\kappa,i} \in \mathrm{Mark}_\kappa(\Delta)$ labeling v° or adjacent constant components. Supposing that it were possible to stabilize such configurations the expected dimension formula (2.4.II) we will shortly obtain could not hold, since if Δ' is a (stable) treed disk obtained from Δ by replacing a marked point with a constant sphere component stabilized by the creation of two additional marked points, then $\underline{\Delta} = \underline{\Delta}'$ and the same underlying map u defines a pseudoholomorphic treed disk $\underline{\Delta}' \to \overline{X}$ belonging to a moduli space with a strictly lower expected dimension.

Thus, we must delineate the combinatorial types which Cieliebak–Mohnke perturbation theory may treat directly.

Definition 2.33. A *maximal ghost sphere tree* in a pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ is a maximal connected subtree of $\Gamma(\Delta)$ with vertices $G = \{v_i^\circ\} \subset \mathrm{Vert}^\circ(\Delta)$ for which each restriction $u|_{v_i^\circ}$ is constant. Observe that if u has spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ then $u|_{v_i^\circ}$ is constant if and only if $\beta_{v_i^\circ} = 0$.

A stable pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is *crowded* if, for any active flavor κ, we have that

(2.33.a) the vertices of any maximal ghost sphere tree $G \subset \mathrm{Vert}^\circ(\Delta)$ together contain in total more than a single marked point $d_{\kappa,i} \in \mathrm{Mark}_\kappa(\Delta)$, or

(2.33.b) we have $m_{d_{\kappa,i}} > 2$ for any $d_{\kappa,i} \in \mathrm{Mark}_\kappa(\Delta)$.

Similarly, we will say that a spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for the type Γ is *crowded* if pseudoholomorphic treed disks of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ are crowded. Observe that since the set of active divisors is determined by the critical points $\mathbf{x}$, crowdedness of $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is determined completely by combinatorial information and labels.

Remark 2.34. The requirement (2.33.b) prohibits entirely legitimate configurations which are not degenerate on a constant component in the sense of their combinatorial decorations. Nonetheless, in practice we will only need to consider moduli spaces of expected dimension 0 or 1, and their exclusion simplifies our subsequent arguments.

In order to define the class of perturbations which enjoy all properties necessary to construct a suitable theory of Floer operations, we introduce one final compatibility condition between the chosen divisors and candidate perturbations of the underlying almost complex structure. Essentially, the requirement is that the energy continues to sufficiently-well control pseudoholomorphic spheres for the prescribed perturbed almost complex structures.

Definition 2.35. Let Γ be a combinatorial type. Consider the relation on the vertices of Γ generated defined by declaring $\underline{t}(e) \sim \underline{h}(e)$ whenever an edge $e \in \mathrm{Edge}(\Gamma)$

- is an edge of length 0 (necessarily between components of Γ), or
- is an edge of length ∞ with $\underline{h}(e)$ a component (necessarily the vertex $\underline{t}(e)$ is a point).

This relation partitions Γ into finitely many subtrees Γ_i called the *maximal nodal trees* of Γ.

Each component $v \in \mathrm{Vert}^\bullet(\Gamma) \cup \mathrm{Vert}^\circ(\Gamma)$ is contained in a unique maximal nodal tree $\Gamma(v) = \Gamma_i$. If Δ is a treed disk then similarly there is a canonical decomposition $\underline{\Delta} = \bigcup_i \underline{\Delta_i}$ (with $\Gamma(\Delta_i) = \Gamma_i$ for each i), and we denote by $\Delta(v)$ the treed disk in this union with combinatorial type $\Gamma(v)$.

If $u : \underline{\Delta} \to \overline{X}$ is a pseudoholomorphic treed disk then, since under u the maximal nodal trees Δ_i of Δ are connected by Morse gradient flow trajectories, the image of each restriction $u|_{\partial S_{\Delta_i}}$ lies wholly in a polyhedron $\theta_{u,\Delta_i} \in \Theta$. Then by definition each divisor $D_\kappa \in \mathbf{D}|_{\theta_{u,\Delta_i}}$ is stabilizing for the fiber bounding $u|_{\partial S_{\Delta_i}}$.

Definition 2.36. The *energy* $\omega(u)$ of a P_Γ-perturbed pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ is the sum of the energies of its components, i.e.

$$\omega(u) := \sum_{v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)} \int_v u^* \omega.$$

Lemma 2.37. *Let $u : \underline{\Delta} \to \overline{X}$ be a P_Γ-perturbed pseudoholomorphic treed disk, with $\mathbf{D} = \{D_\kappa\}_{\kappa \in \mathcal{K}}$ a background system of stabilizing divisors. Decompose Δ into maximal nodal trees Δ_i, and write $l_\kappa \in \mathbb{N}$ for the degree of D_κ. We define*

$$E(u) := \sum_i \min_{D_\kappa \in \mathbf{D}|_{\sigma_{u,\Delta_i}}} \frac{\#\,\mathrm{Mark}_\kappa(\Gamma(\Delta_i)) + 1}{l_\kappa}. \tag{2.4.I}$$

Each P_Γ-perturbed pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ obeys

$$\omega(u) \leq E(u).$$

Proof. This lemma is the family version of [CW22, Proposition 4.19]; the claim follows in the case of a single divisor since each D_κ provided by Theorem 2.20 is exact away from $\theta_{u,\Delta_i} \in \Theta$. □

Definition 2.38. Let $E > 0$. An almost complex structure J' is E-stabilizing if, whenever $u : \mathbb{C}\mathrm{P}^1 \to \overline{X}$ is a nonconstant J'-pseudoholomorphic sphere and $E(u) \leq E$ then u intersects each active divisor D_κ at least three and at most finitely many points.

A perturbation datum P_Γ is *compatible* with the system of stabilizing divisors $\mathbf{D}$ if for all $[\Delta, z] \in \overline{\mathcal{U}}_\Gamma$, whenever $z \in \underline{\Delta}(v)$ we have that $J_\Gamma([\Delta, z])$ is $E(\Gamma(v))$-stabilizing.

The following lemma ensures that the space of almost complex structures yielding compatible perturbation data is plentiful.

Lemma 2.39 ([VWX20, Lemma 2.17])**.** *For each $E > 0$ there exists an open neighborhood of the ambient almost complex structure J_0 (of Remark 2.23) with respect to the C^∞-topology consisting of almost complex structures J' which are E-stabilizing.*

Remark 2.40. In fact, there is a comeager set of almost complex structures which possess the property of Lemma 2.39 for all $E > 0$ simultaneously. The purpose of Lemma 2.39 is to guarantee for us a good space of compatible perturbation data (c.f. Definition 2.38) nearby J_0 in which we may perturb (this space ultimately then being realized as a finite intersection of open neighborhoods of the point J_0 in the Banach manifold of possible choices).

Finally, before stating the characterization of consistent systems of perturbation data that we require, we introduce the operation which enables us to handle crowded types.

Definition 2.41. Let Γ be a (possibly crowded) stable combinatorial type, and let $G \subset \mathrm{Vert}^\circ(\Gamma)$ be a subset of sphere components which together are the vertices of a connected subtree of Γ. The *uncrowding* Γ_G *of* Γ *at* G is the (possibly still crowded) stable combinatorial type obtained by deleting all marked points of $\bigcup G$ except for the unique greatest marked point of each flavor κ and then stabilizing.

Uncrowding at G induces a map of universal curves $\mathcal{U}_\Gamma \to \mathcal{U}_{\Gamma_G}$, and if $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is a spec of type Γ then there is a corresponding spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})_G$ of type Γ_G.

Definition 2.42. A perturbation datum P_Γ is *local* if the domain-dependent choices (J_Γ, f_Γ) of perturbed almost complex structure and Morse function both factor through all maps $\mathcal{U}_\Gamma \to \mathcal{U}_{\Gamma_G}$ induced by uncrowding. Local perturbation data of type Γ and of type Γ_G are therefore in canonical correspondence.

Recall that by condition (2.14.c) all of our perturbation data is local.

Definition 2.43. A perturbation system $\mathbf{P} = \{P_\Gamma\}_{\Gamma \in \mathbf{\Gamma}}$ *works* if:

- For each spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ of type $\Gamma \in \mathbf{\Gamma}$ we have that the moduli space $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ equipped with its natural topology is a smooth manifold.
- For each uncrowding Γ_G of $\Gamma \in \mathbf{\Gamma}$ with $G \subset \mathrm{Edge}^\circ(\Delta)$ contained within some maximal ghost sphere tree in u we have $\Gamma_G \in \mathbf{\Gamma}$ and the perturbation datum P_{Γ_G} is induced by P_Γ (in the sense of Definition 2.42).
- Each perturbation datum P_Γ is compatible with $\mathbf{D}$.[6]

Ultimately, we will establish the existence of working perturbation systems by appeal to the Sard–Smale theorem for (separable, infinite-dimensional) Banach manifolds (c.f. Theorem 1.8)—a variation of an argument tracing its way back to [Flo88]. We will exhibit a functional-analytic framework in which to encode the pseudoholomorphic/Morse gradient flow constraints, and the natural way to achieve this is to equip a

[6]This last condition yields compactness of the moduli spaces of expected dimension zero, see Theorem 3.8 below.

suitable Banach vector bundle with a section which measures their failure. However, we will see from the proof of the main result Theorem 2.45 that the space of possible solutions is plentiful and no sharp choice of functional-analytic parameters (e.g. determining the Sobolev class to which each belongs or other auxiliary constraints) is necessary. The picture will become only slightly more complicated in Chapter 4 (c.f. Section 82 4.2), in which we will be required to introduce some more exotic function spaces.

First, we produce a space of perturbation systems which are suitably small perturbations of the background data: given a background tame almost complex structure $J_0 \in \mathcal{J}_\tau(X, \omega)$ and Morse function $f \in C^\infty(X \to \mathbb{R})$ by Lemma 5.1 of [Flo88] there exist norms on neighborhoods of J_0 and f respectively, so that in the space of all compatible perturbation data P_Γ modeled on (J_0, f) the subset

$$\mathcal{P}_\Gamma = \{(J_\Gamma, f_\Gamma) : \|J_\Gamma - J_0\| + \|f_\Gamma - f\| < \infty\}$$

is a separable Banach manifold [VWX20].

Definition 2.44. Let Γ be a combinatorial type and suppose that $\mathbf{P} = \{P_{\Gamma'}\}_{\Gamma' \in \mathbf{\Gamma'}}$ is a perturbation system. Denote by $\mathcal{P}_\Gamma(\mathbf{P})$ the subset of $\mathcal{P}_\Gamma$ containing all P_Γ such that $\{P_\Gamma\} \cup \mathbf{P}$ is again a perturbation system. In other words, $\mathcal{P}_\Gamma(\mathbf{P})$ is the space of extensions (inside $\mathcal{P}_\Gamma$) of $\mathbf{P}$ over the combinatorial type Γ; viewing elements of $\mathcal{P}_\Gamma(\mathbf{P})$ this way, there is also a subset $\mathcal{P}_\Gamma^{\checkmark}(\mathbf{P}) \subset \mathcal{P}_\Gamma(\mathbf{P})$ of working perturbation systems.

The central object we consider is the dependent product of the space of perturbation data of fixed type with the moduli space of perturbed pseudoholomorphic treed disks respectively determined by each datum; this is the universal moduli space of possible choices

$$\mathcal{M}_\Gamma(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) := \Big\{(P_\Gamma, [u]) : P_\Gamma \in \mathcal{P}_\Gamma, [u] \in \mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})\Big\}.$$

Locally, each point $[\Delta, z] \in \mathcal{U}_\Gamma$ of the universal treed disk is contained in a trivial neighborhood $\mathcal{U}_{\Gamma,\Delta} \cong \mathcal{M}_{\Gamma,\Delta} \times \underline{\Delta}$ for some $\mathcal{M}_{\Gamma,\Delta} \subset \mathcal{M}_\Gamma$. The universal moduli space restricts along this local piece yielding a subspace $\mathcal{M}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ which consists of all $(P_\Gamma, [u : \underline{\Delta}' \to \overline{X}])$ for which $[\Delta'] \in \mathcal{M}_{\Gamma,\Delta}$. It will be convenient for us to work on these local pieces.

Thus consider the space $\mathrm{Map}_\Delta(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ of all continuous maps $u : \underline{\Delta} \to \overline{X}$ such that u has spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ in the sense of Definition 2.31 and obeys (2.16.c) (but nonetheless need not obey the pseudoholomorphic or Morse gradient flow equations (2.16.a) and (2.16.b)). Denote by $\mathrm{Map}^{k,p}_\Delta(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ the subspace consisting of those maps of Sobolev class $W^{k,p}$. This latter space is a Banach manifold, with local charts furnished (as in [CW22]) by the geodesic exponential map for a metric on X for which the fibers of $\pi : X \to Q$ are totally geodesic. We will also need the weak space $\mathcal{P}^l_\Gamma$ of perturbations of type Γ for which each $P_\Gamma \in \mathcal{P}^l_\Gamma$ is exactly as in Definition 2.14 with the exception that smoothness of data is relaxed to the requirement that it is of class C^l. Similarly, we have a weak moduli space of choices

$$\mathcal{M}^{l,k,p}_\Gamma(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) = \left\{(P_\Gamma, [u]) : P_\Gamma \in \mathcal{P}^l_\Gamma, [u] \in \mathcal{M}^{k,p}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})\right\},$$

for which trivialization of $\mathcal{U}_\Gamma$ over $\mathcal{M}_{\Gamma,\Delta}$ yields that a local piece decomposes as a subspace of a product

$$\mathcal{M}^{l,k,p}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) \subset \mathcal{B}^{l,k-1,p}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) := \mathcal{P}^l_\Gamma \times \mathcal{M}_{\Gamma,\Delta} \times \mathrm{Map}^{k,p}_\Delta(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}).$$

Implicitly restricting to the interior of the surface part S_Δ, as in Theorem 1.10 the operator $d_S = D - J_\Gamma \circ D \circ j$ measures the failure of $u : \underline{\Delta} \to \overline{X}$ to be P_Γ-perturbed pseudoholomorphic, and lifts to an operator on $\mathcal{B}^{l,k-1,p}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ naturally taking values in $u|^*_{\mathrm{int}(S_\Delta)}\mathrm{T}X$-valued $(0,1)$-forms on $\mathrm{int}(S_\Delta)$ of class $W^{k-1,p}$. Similarly, on the interior of the tree part T_Δ the operator $d_t = \frac{\mathrm{d}}{\mathrm{d}t} - \nabla f$ measures failure of the Morse gradient flow equation, and is valued in $u|^*_{\mathrm{int}(T_\Delta)}\mathrm{T}X$-valued 1-forms on $\mathrm{int}(T_\Delta)$ again of class $W^{k-1,p}$.

The direct sum $d_S \oplus d_T$ is then a section of the C^q-Banach vector bundle (whenever $q < l - k$)

$$\mathcal{E}^{l,k,p}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) \to \mathcal{B}^{l,k-1,p}_{\Gamma,\Delta}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}),$$

namely that subbundle consisting of all pairs of forms $\theta_S \oplus \theta_T$ with θ_S having a multiplicity $m_{\kappa,d} - 1$ zero at each $d \in \mathrm{Mark}_\kappa(\Delta) \hookrightarrow \underline{\Delta}$.

In the setting of uncrowded pseudoholomorphic treed disks $u : \underline{\Delta} \to \overline{X}$ of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$, the linearization $\widetilde{D}_u$ of $d_S \oplus d_T$ is Fredholm, of index

$$\begin{aligned} \operatorname{ind} \widetilde{D}_u = \dim \mathcal{M}_\Gamma - I(x_0) + \sum_{i=1}^{n} I(x_i) + \sum_{v \in \mathrm{Vert}^\bullet(\Delta)} I(\beta_v) + 2 \sum_{v \in \mathrm{Vert}^\circ(\Delta)} c_1(\beta_v) \\ - 2 \sum_{\substack{\kappa \in \mathcal{K} \\ d \in \mathrm{Mark}_\kappa(\Delta)}} m_{\kappa,d} - \#\,\mathrm{Edge}^\triangledown(\Gamma). \end{aligned} \tag{2.4.II}$$

The expected dimension formula (2.3.II) follows.

The following theorem of standard type encapsulates the necessary functional analysis required to obtain transversality.

Theorem 2.45 (Generic transversality)**.** *For each combinatorial type* Γ *and working perturbation system* $\mathbf{P}$ *the subset* $\mathcal{P}^{\checkmark}_{\Gamma}(\mathbf{P}) \subset \mathcal{P}_{\Gamma}(\mathbf{P})$ *is comeager.*

Before proving Theorem 2.45, let us establish its key immediate consequence.

Corollary 2.46. *There exists a complete working perturbation system* $\mathbf{P}$*.*

Proof. We construct $\mathbf{P}$ inductively. The only difficulty is that while all boundary compatibility constraints between choices of P_Γ and $P_{\Gamma'}$ are recorded by the order relation $\Gamma' > \Gamma$ on combinatorial types, this order is far from total. Thus repeatedly making naïve choices from the collections $\mathcal{P}^{\checkmark}_{\Gamma}(\mathbf{P})$ may eventually result in a contradiction (i.e. $\mathcal{P}_\Gamma(\mathbf{P})$ may become empty at some stage).

The straightforward trick is to observe that all types Γ have a unique minimal $\Gamma^{\mathrm{simp}} \leq \Gamma$ such that $\mathrm{Vert}^\bullet(\Gamma^{\mathrm{simp}}) = \mathrm{Vert}^\bullet(\Gamma)$—therefore identical collections of marked points on all disk vertices—and that a choice of $P_{\Gamma^{\mathrm{simp}}}$ determines a compatible choice of P_Γ. Assuming that perturbation data $P_{\Gamma^{\mathrm{simp}}}$ have been compatibly chosen for all $(\Gamma')^{\mathrm{simp}} < \Gamma^{\mathrm{simp}}$ the space $\mathcal{P}_{\Gamma^{\mathrm{simp}}}(\mathbf{P})$ is never empty and we may choose a working perturbation system $P_{\Gamma^{\mathrm{simp}}} \in \mathcal{P}^{\checkmark}_{\Gamma^{\mathrm{simp}}}(\mathbf{P})$ by Theorem 2.45. Repeating this process for all minimal combinatorial types Γ^{simp} according to the order relation $<$ yields the desired complete system $\mathbf{P}$. □

Proof of Theorem 2.45. First observe that if Γ is a crowded type then, under the correspondence of local perturbation data for uncrowdings (c.f. Definition 2.42), the set $\mathcal{P}^{\checkmark}_\Gamma(\mathbf{P})$ is the intersection of $\mathcal{P}^{\checkmark}_{\Gamma_G}(\mathbf{P})$ over all (finitely many) possible uncrowdings Γ_G of G. Therefore we may assume that the type Γ is uncrowded.

In this case, by the Sard–Smale Theorem 1.8, it in turn suffices to show that the linearization of the section $d_S \oplus d_T$ is surjective everywhere on its zero locus. Thus let $\eta_S \oplus \eta_T$ be orthogonal to the kernel of the linearization of $d_S \oplus d_T$ at

$$((J_\Gamma, f_\Gamma), \psi : \Delta' \cong \Delta, u : \underline{\Delta} \to \overline{X}) \in \mathcal{P}^l_\Gamma \times \mathcal{M}_{\Gamma,\Delta} \times \mathrm{Map}^{k,p}_\Delta(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}).$$

This condition on $d_S \oplus d_T$ is extremely strong.

First, for each $e \in \mathrm{Edge}(\Delta)$ either $l(e) = 0$, or the support $\mathcal{T}^{\circledast}_\Delta$ of the Morse function perturbation contains an open subset of the interior of $L_e \subset T_\Delta$; in either case–since u may not be constant on the tree part—we deduce that $\eta_T = 0$ on L_e.

Second, if u is nonconstant on any component $v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)$, then the support $\mathcal{S}^{\circledast}_\Delta$ of the almost complex structure contains an open subset of the interior of v, and pseudoholomorphicity (c.f. [MS12, Remark 3.2.3]) implies that $\eta_S = 0$ on the entirety of v.

Finally, let $G \subset \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)$ be the vertices of a maximal nodal tree on which u is constant. If G consists a single vertex without marked points then it is well-known that d_S is surjective there (e.g. [MS12, Chapter 3]), so it remains to check that local solutions on single components $v \in G$ assemble into a global solution. However, this follows immediately from the facts that: first, the components of v are glued according to edges determined by a tree, and second, uncrowdedness of Γ implies that $\bigcup G$ contains at most one marked point of each flavor. $\square$

Chapter 3

The Morse–Fukaya algebra

In this chapter we construct the A_∞-algebra structure maps for the Morse–Fukaya algebra $\mathscr{A}$ by taking corrected counts of various zero-dimensional moduli spaces of pseudoholomorphic treed disks. We prove that these operations satisfy the corresponding (curved) A_∞-relations via exhibiting corresponding 1-dimensional moduli spaces and analyzing the boundaries of their respective compactifications. Finally, we address the invariance of our constructions under changes of all choices.

3.1 Universal coefficients

The Morse–Fukaya algebra is generated by the critical points of a suitable Morse function on the total space of $\pi : X \to Q$. Its structure maps are defined by making a (signed) count of the cardinalities of moduli spaces of pseudoholomorphic treed disks of the kind studied in Chapter 2 and applying appropriate *Floer-theoretic weights* $z^{\boldsymbol{\beta}}$ determined by the respective homology classes $\boldsymbol{\beta}$ the treed disks represent.

Consider the moduli space $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ of Definition 2.32 for a choice of complete working perturbation system $\mathbf{P}$, fixing $\mathbf{x}$ (i.e. with $\mathbf{m}$, $\boldsymbol{\beta}$, and the type Γ allowed to vary). In general there may be infinitely many vectors $\boldsymbol{\beta}$ so that the moduli space of P_Γ-perturbed pseudoholomorphic treed disks with spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ is nonempty, regardless of potential additional stipulations prohibiting degeneracy of the type Γ. Correspondingly, our formulas for the multiplication laws of the Morse–Fukaya algebra will contain infinite sums. Fixing a base field $\Bbbk$ of characteristic zero over which to work, in order to avoid all questions of analytic convergence (in the classical sense) it is convenient to introduce the *Novikov field*

$$\Lambda = \left\{ \sum_{i=1}^{\infty} c_i T^{x_i} : c_i \in \Bbbk, x_i \in \mathbb{R}, \lim_{i \to \infty} x_i = \infty \right\}$$

of formal series in the variable T with real exponent converging to ∞. The ring Λ is a nonarchimedean field, and is equipped with a valuation map

$$\mathrm{val} : \Lambda^* \to \mathbb{R} \quad \text{defined by} \quad \sum_{i=1}^{\infty} c_i T^{x_i} \mapsto \min\{x_i : c_i \neq 0\}.$$

We write $U_\Lambda := \mathrm{val}^{-1}(0) \subset \Lambda^*$ for its unitary subgroup.

First proposed by [KS01] and employed in [Abo14; Abo17] to define a family Floer functor, the *uncorrected rigid analytic mirror* $X^\vee$ of $\pi : X \to Q$ has underlying set the disjoint union

$$X^\vee = \bigsqcup_{q \in Q} H^1(F_q, U_\Lambda) \tag{3.1.I}$$

taken over the fibers F_q of π, and comes equipped with an obvious projection map $\pi^\vee : X^\vee \to Q$ which forgets all but the fiber label. Under mirror symmetry, this construction has the straightforward interpretation that each $H^1(F_q, U_\Lambda)$ is the torus in $X^\vee$ dual to F_q. The modern perspective is that $X^\vee$ should be realized as the moduli space of Lagrangian fibers of π equipped with a unitary rank 1 local system [Fuk05]—i.e. objects of a suitable Fukaya category of X supported on a particular torus fiber. Since points of this moduli space are determined by their holonomy map $\mathrm{hol} : H_1(F_q) \to U_\Lambda$, they equivalently belong to the disjoint union (3.1.I).

Fix $q \in Q$ and a simply connected domain $P \subset Q$ containing q. Then there is a canonical homotopy class of paths $[\psi_p(t)]$ from q to any $p \in P$. For any loop $\gamma \in H_1(F_q)$, parallel transport along the path $\psi_p(t)$ takes the class γ to a class $\psi_{p*}\gamma \in H_1(X, F_p)$.

While doing so γ sweeps out a cylinder, and hence yields a class $\beta = \psi_{q\to p*}\gamma \in H_2(X, F_q \cup F_p)$. Since the fibers of π are Lagrangian, allowing $p \in P$ to vary we obtain a well-defined function

$$z^\beta = T^{\omega(\beta)}\,\mathrm{hol}(\partial\beta) : X^\vee|_P \to \Lambda^*$$

where $\omega(\beta) = \int_\beta \omega$ is the symplectic area of $\beta = \psi_{q\to p*}\gamma$ and $\mathrm{hol}(\partial\beta) = \mathrm{hol}(\psi_{p*}\gamma)$ denotes application of the evaluation map $H^1(F_p, U_\Lambda) \times H_1(F_p) \to U_\Lambda$ to $\psi_{p*}\gamma$.

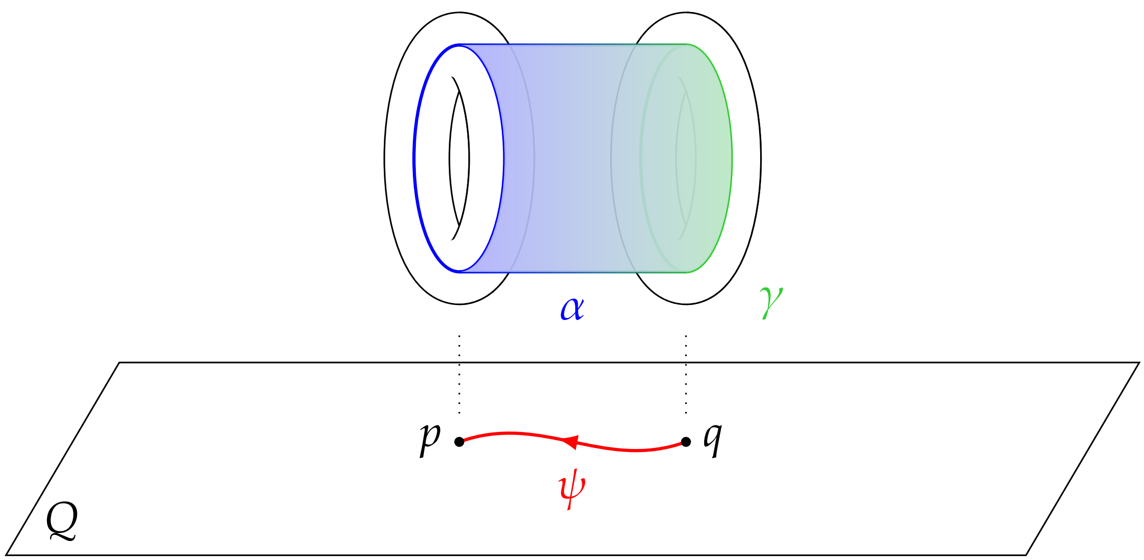

Figure 3.1: A schematic diagram depicting the parallel transport of a class $\gamma \in H_1(F_q)$ by $\psi_p : [0, 1] \to Q$ yielding a class $\beta = \psi_{q\to p*}\gamma$.

Let $\gamma_1, \ldots, \gamma_n$ be loops in F_q which together represent a basis of $H_1(F_q)$, and denote by $\beta_1, \ldots, \beta_n$ the corresponding classes (depending on $p \in P$) in $H_2(X, F_q \cup F_p)$. Then the map $(z^{\beta_1}, \ldots, z^{\beta_n}) : X^\vee|_P \to (\Lambda^*)^n$ fits into a commutative square

$$\begin{tikzcd} X^\vee|_P \arrow[r] \arrow[d, "\pi^\vee"'] & (\Lambda^*)^n \arrow[d, "\mathrm{val}"] \\ P \arrow[r] & \mathbb{R}^n \end{tikzcd} \tag{3.1.II}$$

and defines an analytic chart endowing $X^\vee|_P$ with the structure of a rigid analytic space. The analytic functions $\mathcal{O}_{X^\vee|_P}$ on this local piece are precisely the Laurent series in the variables $z_1, \ldots, z_n \in \Lambda^*$ furnished by the chart (3.1.II) which T-adically converge over the entirety of $X^\vee|_P$; this means that each $g \in \mathcal{O}_{X^\vee|_P}$ is of the form

$$g = \sum_{(i_1,\ldots,i_n)\in\mathbb{Z}^n} g_{i_1,\ldots,i_n} z_1^{i_1} \cdots z_n^{i_n} \quad \text{for } g_{i_1,\ldots,i_n} \in \Lambda$$

such that

$$\lim_{\substack{k\to\infty \\ |i_1|+\ldots+|i_n|=k}} \left(\mathrm{val}(g_{i_1,\ldots,i_n}) + \sum_{j=1}^{n} i_j \omega(\psi_{q\to p*}\gamma_j) \right) = \infty \quad \text{for all } p \in P.$$

In other words, $\mathcal{O}_{X^\vee|_P}$ is a particular completion of the ring of Laurent polynomials $\Lambda[H_1(F_q)]$.

In effect, for each $p \in P$ we have produced an element of $H^1(F_p;\mathbb{R}) \cong H^1(F_q;\mathbb{R})$ defined by $\gamma_i \mapsto \mathrm{val}\, z^{\beta_i} = \omega(\psi_{q\to p*}\gamma_i)$. Allowing p to vary yields a local identification $P \hookrightarrow H^1(F_q;\mathbb{R})$.

Definition 3.1. A subset $P \subset Q$ is a *special affine subset*[1] if there exist finitely many $[\gamma_i] \in H_1(F_q)$ and constants $\lambda_i \in \mathbb{R}$ such that, under the identification of elements of P and $H^1(F_q;\mathbb{R})$ just described, we have

$$P = \{p \in H^1(F_q;\mathbb{R}) : p([\gamma_i]) \le \lambda_i \text{ for all } i\}.$$

[1] This is Definition 7.1 of [Tat71].

As in [Abo14], if $P \subset Q$ is a special affine subset, then P and $X^{\vee}|_P$ are then each both *affinoid domains* in the sense of rigid analytic geometry. In order to assemble a sheaf of universal coefficients for all of X, we will appeal to Tate's acylicity theorem:

Theorem 3.2 (Gerritzen–Grauert–Tate [GG15] [BGR84, 7.3.5 Theorem 1])**.** *Every finite covering of an affinoid space by affinoid domains may be refined to a finite covering by rational domains. Finite rational coverings, hence finite affinoid coverings, are acylic for Čech cohomology.*

Proposition 3.3. *There exists a perfect lift f of a Morse function $\check{f} : Q \to \mathbb{R}$, a compatible system of divisors $\mathbf{D}$, and a confining polyhedral cover Θ of Q by special affine subsets.*

Proof. Lemma 5.1 of [Abo17] implies that, by choosing $\check{f}$ such that the induced cellular decomposition is fine enough, we are guaranteed a cover Θ of Q which is confining for any perfect lift f of $\check{f}$. Choosing $\check{f}$ so that the induced cellular decomposition is also fine enough to support a compatible system of divisors provided by Theorem 2.25 proves the claim. ◻

Definition 3.4. Allowing $P \subset X$ to vary over the affinoid cover we have just described, by Theorem 3.2 the sheaves $\mathcal{O}_{X^{\vee}|_P}$ glue along each inclusion $P_i \hookrightarrow P_j$ yielding a sheaf $\mathcal{O}_{\mathrm{an}} := \mathcal{O}_{\mathrm{an}}(f, \Theta)$ of *universal coefficients* for family Floer theory on $\pi : X \to Q$.

Definition 3.5. The *Morse–Fukaya algebra* is the $\mathbb{Z}_2$-graded $\mathcal{O}_{\mathrm{an}}$-module

$$\mathcal{A} := \mathrm{CM}^{\bullet}(f; \mathcal{O}_{\mathrm{an}}) = \bigoplus_{x \in \operatorname{gen} f} \mathcal{O}_{\mathrm{an}}|_{\theta_{\pi(x)}}[I(x)]$$

i.e. freely generated by the finitely many generators $\operatorname{gen} f$ of (2.3.I) and graded by index I (modulo 2), where it is understood in the definition that $\pi(x^{\triangledown}_{\mathsf{m},i}) = \pi(x^{\blacktriangledown}_{\mathsf{m},i}) = \pi(x^{\blacktriangle}_{\mathsf{m},i})$.

It remains to prescribe well-defined structure maps for $\mathcal{A}$.

3.2 Compactness and boundary strata

In this section we verify that the moduli spaces $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ of low expected dimension are compact in the appropriate sense, and in doing so characterize their boundary strata. Figure 3.2 depicts a 1-parameter family in a 1-dimensional moduli space with boundary points we will consider subsequently.

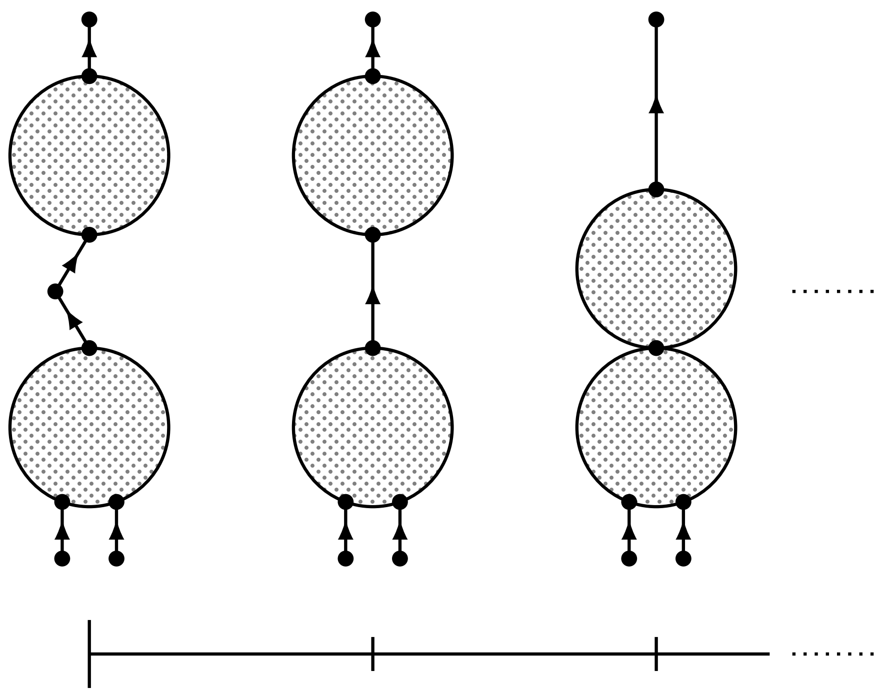

Figure 3.2: A representative diagram depicting a 1-parameter family meeting the boundary of the moduli space in which it resides.

Definition 3.6. A combinatorial type Γ with $n(\Gamma)$ inputs is *minimally degenerate*[2] if whenever $\Gamma' < \Gamma$ then $n(\Gamma') < n(\Gamma)$. For each $n > 0$ define

$$\mathcal{M}_{\mathbf{P},n}(\mathbf{x}) := \bigcup_{\substack{n(\Gamma)=n \text{ and } \Gamma \text{ is stable} \\ \text{and minimally degenerate}}} \ \bigcup_{\substack{(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \text{ is a spec} \\ \text{for the type } \Gamma}} \mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}). \tag{3.2.I}$$

[2]Note that this is a coarser notion than that of Γ^{simp} introduced in the proof of Corollary 2.46. Of course, minimally degenerate types cannot have sphere components and therefore all specs of their type are uncrowded.

Each $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})$ decomposes as the union of strata $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})_d$ of expected dimension d (according to (2.3.II)). Replacing each stratum of (3.2.I) with its respective compactification, we obtain natural compactifications[3] $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})$ and $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_d$ of the respective moduli spaces $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})$ and $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})_d$.

Lemma 3.7. *Each inclusion of a stratum $\mathcal{M}_{P_\Gamma}(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \subset \mathcal{M}_{\mathbf{P},n}(\mathbf{x})_d$ extends to an embedding of a tubular neighborhood.*

Proof. This follows from standard constructions, such as in [Sch16], gluing nodal pseudoholomorphic disks and broken Morse gradient flow trajectories. □

It is immediate from Lemma 3.7 that all of the spaces $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})_d$, $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_d$, $\mathcal{M}_{\mathbf{P},n}(\mathbf{x})$, and $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})$ are smooth manifolds. Recalling the energy bound $E(\Gamma)$ of (2.4.I), for each $E > 0$ we also have submanifolds $\mathcal{M}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_d \subset \mathcal{M}_{\mathbf{P},n}(\mathbf{x})_d$ of all strata enumerated in (3.2.I) over types Γ with $E(\Gamma) \leq E$. Each in turn again has a natural compactification $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_d$. In this language, our main compactness theorem takes the following form.

Theorem 3.8. *Let* $\mathbf{P}$ *be a complete working perturbation system. For all* $\mathbf{x} \in (\operatorname{gen} f)^{n+1}$ *we have that:*

(3.8.a) for each $E > 0$ the spaces $\mathcal{M}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_0$ and $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1$ are compact, and

(3.8.b) the boundary of $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1$ is the (disjoint) union of all

$$\mathcal{M}_{P_{\Gamma'}}(\mathbf{m}',\mathbf{x}',\boldsymbol{\beta}') \quad \text{taken over all strata} \quad \mathcal{M}_{P_\Gamma}(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \subset \overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1,$$

where $(\mathbf{m}',\mathbf{x}',\boldsymbol{\beta}')$ is a spec for the type Γ' obtained from the spec $(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ for Γ by a single application of the operation (2.2.a) or (2.2.d).

Theorem 3.8 is proved by studying each $\mathcal{M}_{P_\Gamma}(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \subset \overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1$ individually. Certain operations (namely (2.2.a) and (2.2.d)) survive the gluing of these strata to give

[3] For $d = 1$ we will see that these spaces are in general infinite unions of compact 1-dimensional components.

boundary points of all of $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1$. Others (namely (2.2.b) and (2.2.c)) give common points of boundary strata which glue together.

Theorem 3.9. *Let* **P** *be a complete working perturbation system, and suppose that* $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *is a spec for a stable and minimally degenerate type* Γ*. Denote by* d *the expected dimension of* $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *according to* (2.3.II)*. We have that*

(3.9.a) if $d = 0$ *then the space* $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *is compact, and*

(3.9.b) if $d = 1$ *then the boundary of the space* $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *is the union of the spaces* $\mathcal{M}_{P_{\Gamma'}}(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ *taken over all specs* $(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ *of type* Γ' *obtained from* $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ *by a single application of one of the operations* (2.2.a)–(2.2.d)*. Moreover, in the operation* (2.2.c)*, the creation of sphere bubbles is prohibited.*

Proof. Let $(u_i : \Delta_i \to \overline{X})$ be a sequence of pseudoholomorphic treed disks representing a sequence $([u_i])$ in $\mathcal{M}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for $0 \leq d \leq 1$. For each of the finitely many component vertices $v \in \mathrm{Vert}^\bullet(\Gamma) \cup \mathrm{Vert}^\circ(\Gamma)$ we may form the restriction $(u_i|_v)$, which yields in each case a sequence of honest pseudoholomorphic disks or spheres; by successively passing to subsequences, Gromov compactness (c.f. Theorem 1.11) implies that $(u_i|_v)$ converges to a pseudoholomorphic treed disk $u_v : \underline{\Delta_v} \to \overline{X}$ for each fixed v. Similarly, compactness of the space of Morse gradient flow trajectories (c.f. Theorem 1.5) yields, again by passing to subsequences, that the restriction $u_i|_{L_e}$ to each edge $e \in \mathrm{Edge}(\Delta)$ converges to some $u_e : \underline{\Delta_e} \to X$ with Δ_e having 1 input and containing no non-constant components (a broken Morse gradient flow trajectory).

Gluing these limits u_v and u_e at their inputs and outputs, we may therefore assume that (u_i) converges to a pseudoholomorphic treed disk $u : \Delta \to \overline{X}$ of some spec $(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ for the type Γ'. In order show that Γ' is again stable, let $v' \in \mathrm{Vert}^\bullet(\Gamma') \cup \mathrm{Vert}^\circ(\Gamma')$ be any disk or sphere component; we must check that v' contains enough joints (in the sense of Definition 2.9). If $u|_{v'}$ is constant then there is nothing to check, because the map u is itself stable (c.f. Definition 2.19).

Otherwise, assume $u|_{v'}$ is nonconstant. Since u was assembled as a gluing, the vertex v' also belongs to $\mathrm{Vert}(\Gamma(\Delta_v))$ for some $v \in \mathrm{Vert}(\Gamma)$. By conservation of symplectic energy in the limit we have $\omega(u|_{v'}) \leq E(u|_{\underline{\Delta_v}}) \leq E(\Gamma(v), \mathbf{x})$. As the perturbation datum P'_Γ is assumed compatible (c.f. Definition 2.38) with $\mathbf{D}$, the perturbed almost complex structure for which $u|_{v'}$ is pseudoholomorphic is $E(\Gamma(v), \mathbf{x})$-stabilizing. It follows immediately that v' contains enough joints in this case as well.

We must also verify the spec $(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ is again uncrowded; let Γ'_G be the uncrowding at a subset $G \subset \mathrm{Vert}^\circ(\Gamma')$. By restriction, G determines a spec $(\mathbf{m}'_G, \mathbf{x}', \beta'_G)$ of type Γ'_G for which the map $u : \underline{\Delta} \to \overline{X}$ gives rise to an element $[u_G]$ of $\mathcal{M}_G := \mathcal{M}_{P_{\Gamma'_G}}(\mathbf{m}'_G, \mathbf{x}', \beta'_G)$. Choosing G so that $(\mathbf{m}'_G, \mathbf{x}', \boldsymbol{\beta}'_G)$ is uncrowded we have that the expected dimension formula (2.3.II) holds for $\mathcal{M}_G$, and since $d < 2$, it predicts that $\mathcal{M}_G$ is empty (i.e. of negative dimension) unless $\Gamma'_G = \Gamma'$. Therefore this is the only possibility, and the spec $(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ was already uncrowded.

This all shows that the expected dimension formula (2.3.II) holds for the stratum $\mathcal{M}_{P_\Gamma}(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ of $\overline{\mathcal{M}}_{P_\Gamma}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ in which u lies. If $d = 0$ this implies that $\Gamma' = \Gamma$ and verifies compactness. If $d = 1$ this implies admissibility of the operations claimed in (3.9.b), and prohibits the creation of any sphere component since once again after any single application of this operation the expected dimension of $\mathcal{M}_{P_\Gamma}(\mathbf{m}', \mathbf{x}', \boldsymbol{\beta}')$ would be reduced to $-1 = d - 2 < 0$. □

With this characterization in hand, Theorem 3.8 now follows easily.

Proof of Theorem 3.8. By Theorem 3.9 each space $\mathcal{M}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_0$ and $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\mathbf{x})_1$ is an honest compact topological space, for the reason that there are finitely many vectors $\boldsymbol{\beta}$ with $E(\boldsymbol{\beta}) \leq E$ for which pseudoholomorphic treed disks u of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ may exist. This follows by an elementary argument relying on the fact that $H_2(X, F_q)$ is finitely generated and that the total homology class is preserved in any limit (or alternatively, as a particular consequence of Gromov compactness Theorem 1.11).

The claimed characterization of boundary strata of $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_1$ is a result of the pairing-up of 1-dimensional strata along 0-dimensional common boundaries respectively obtained from an application of the operation (2.2.b) or (2.2.c)—this is depicted in Figure 3.3. In one direction, a pair of disk components meeting at a common node may be deformed into one another (thereby viewing the initial nodal configuration as a result of disk bubbling à la (2.2.c)). In the other direction, the nodal configuration may be resolved by making the length of the zero-length edge encoding the nodal joint positive (à la (2.2.b)).

For completeness[4], we note that the boundary-normal direction of a boundary stratum arising from the operation (2.2.a) points in the direction which corresponds to the resolution of a broken edge into an edge of finite length; this is a semi-infinite and not infinite family of deformations because a broken Morse gradient flow line cannot stably "break more" at a breaking. Finally, operation (2.2.d) yields points of a 1-dimensional stratum of $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_1$ which manifestly lie on the boundary, the entire stratum to which they each respectively belong in each case being itself parameterized by the weight $\rho(e) \in [0, \infty]$ of a particular edge. □

3.3 Multiplication maps and grading

An arbitrary spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for a combinatorial type Γ in particular records classes β_v, each belonging to homology groups $H_2(X, F_q)$ depending on the fiber F_q which bounds each respective component v. However, parallel transport over the simply connected base Q allows all of these classes to be transported to a class $\beta'_v \in H_2(X, F_{\pi(x_0)})$ bounded by the output fiber. Thus we may interpret $z^{\boldsymbol{\beta}} x_0 := z^{\sum_v \beta'_v} x_0$ as a well-defined element of $\mathcal{A}|_{\theta_{x_0}}$.

[4]Technically, in order to verify the A_∞-relations we subsequently enumerate, we need only know that honest boundary points of $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_1$ do not arise from operations (2.2.b) and (2.2.c); the other possibilities are themselves accounted for by the algebraic relations themselves.

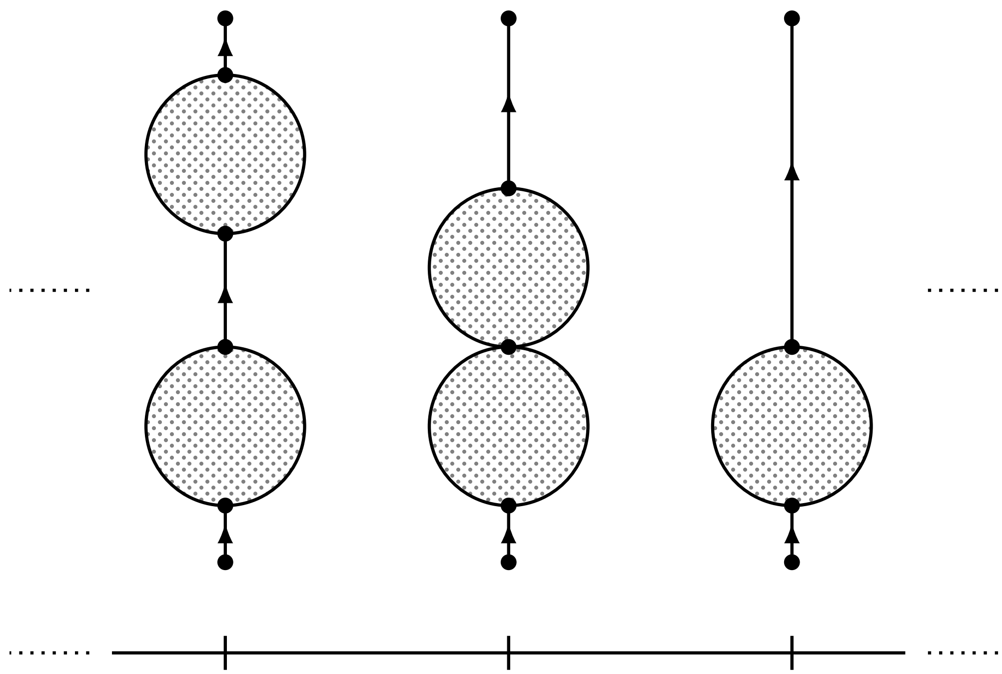

Figure 3.3: A schematic diagram depicting the nodal domain configuration lying at the common boundary point where two 1-dimensional strata of $\overline{\mathcal{M}}_{\mathbf{P},n}(\mathbf{x})_1$ are glued, this point obtained from domain configurations arising from a single application of the operation (2.2.b) or (2.2.c) respectively.

We are now finally in a position to define the A_∞-multiplication maps $\mu^n : \mathcal{A}^{\otimes n} \to \mathcal{A}[2-n]$ for the Morse–Fukaya algebra $\mathcal{A}$. Given generators $\mathbf{x} = (x_1, \ldots, x_n) \in (\operatorname{gen} f)^n$ and $x_0 \in \operatorname{gen} f$, write $(x_0, \mathbf{x})$ for the concatenation $(x_0, x_1, \ldots, x_n)$. We declare

$$\mu^n(\mathbf{x}) := (-1)^{\heartsuit} \sum_{\substack{x_0 \in \operatorname{gen} f \\ [u] \in \mathcal{M}_{\mathbf{P},n}(x_0,\mathbf{x})_0}} \frac{o([u])}{\prod_{\kappa \in \mathcal{K}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} \cdot z^{\beta([u])} x_0 \quad \text{for all } n \geq 0, \tag{3.3.I}$$

with $\heartsuit = \sum_{k=1}^n kI(x_k)$ and extend $\mathcal{O}_{\mathrm{an}}$-linearly (with the understanding that $o([u]) = \pm 1$ denotes the orientation of $[u]$ and if u is of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ then $\boldsymbol{\beta}([u]) = \boldsymbol{\beta}$).

We prove well-definedness of this sum via an elementary technique known as *Fukaya's trick* (observed in [Fuk10], see also [Abo14; Abo17]).

Lemma 3.10. *If the polyhedral cover Θ is sufficiently fine, then the formula (3.3.I) is well-defined for every* $\mathbf{x}$.

Proof. We must arrange a polyhedral cover Θ such that for each fixed $x_0 \in \operatorname{gen} f$ and minimally degenerate stable type Γ the infinite sum

$$\sum_{\substack{(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \text{ is a spec} \\ \text{for the type } \Gamma}} \frac{\#\mathcal{M}_{P_\Gamma}(\mathbf{m},\mathbf{x},\boldsymbol{\beta})}{\prod_{\kappa\in\mathcal{K}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} \cdot z^{\boldsymbol{\beta}} \tag{3.3.II}$$

is a well-defined element of $\mathcal{O}_{\mathrm{an}}|_{\theta_{x_0}}$, i.e. converges with respect the valuation val at all points of $(\pi^\vee)^{-1}(\theta_{x_0})$—note that, by construction, each $\mathcal{O}_{\mathrm{an}}|_{\theta_{x_i}}$ is naturally a subring of $\mathcal{O}_{\mathrm{an}}|_{\theta_{x_0}}$ whenever there exists a broken Morse gradient flow trajectory from x_i to x_0.

Recall that (X, ω, J) is a Kähler manifold. Now, the symplectic area of a J-holomorphic sphere or disk $u : C \to \overline{X}$ is computed as

$$\omega(u) = \int_C u^*\omega = \frac{1}{2}\int_C \|\mathrm{d}u\|_J^2 \,\mathrm{dvol}$$

via the metric, and hence for a perturbed almost complex structure J' (such as J_0 of Remark 2.23) chosen C^1-close to J we have (for the induced deformed metric $\|\cdot\|_{J'}$)

$$\frac{1}{2}\omega(u) < \frac{1}{2}\int_C \|\mathrm{d}u\|_{J'}^2 \,\mathrm{dvol} < 2\omega(u). \tag{3.3.III}$$

So long as each cell $\theta \in C$ is sufficiently small, for each $p, q \in C$ we can find a diffeomorphism $\Phi : Q \to Q$ isotopic to the identity for which $\Phi(q) = p$, such that the further deformed complex structure $J_q := \Phi^* J_0$ is ω-tame, and with Φ chosen C^1-close enough to the identity such that (3.3.III) continues to hold for J_q.

But now (3.3.III) implies that that the valuation of each term of (3.3.II) obeys $\frac{1}{2}\omega(\boldsymbol{\beta}) \le \operatorname{val} z^{\boldsymbol{\beta}} \le 2\omega(\boldsymbol{\beta})$. Convergence of the sum (3.3.III) follows from the consequence of ordinary Gromov compactness Theorem 1.11 that there are only finitely many isomorphic classes of treed disks with boundary on a single Lagrangian with energy below any finite bound. □

Theorem 3.11. *The multiplication law (3.3.I) endows the Morse–Fukaya algebra $\mathcal{A}$ with the structure of an A_∞-algebra. That is, for each $n \geq 0$ and homogeneous elements $a_i \in \mathcal{A}$ of respective degrees $|a_i|$ we have the A_∞-identity [Sei08a]*

$$0 = \sum_{j+k \leq n} (-1)^{j+\sum_{i=1}^{j}|a_i|} \mu^{n-k+1}(a_1, \ldots, a_j, \mu^k(a_{j+1}, \ldots, a_{j+k}), a_{j+k+1}, \ldots, a_n). \tag{3.3.IV}$$

Proof. Recall the subspaces $\mathcal{M}^{\leq E}_{\mathbf{P},n}(x_0, \mathbf{x})_1 \subset \mathcal{M}_{\mathbf{P},n}(x_0, \mathbf{x})_1$ of all strata enumerated in (3.2.I) over types Γ with $E(\Gamma) \leq E$. By Theorem 3.8 each such space is an honest compact oriented 1-manifold with boundary. Therefore the signed count of boundary points of $\mathcal{M}^{\leq E}_{\mathbf{P},n}(x_0, \mathbf{x})_1$ is zero, and hence allowing $E > 0$ to vary we obtain the identity

$$0 = \sum_{[u] \in \partial \overline{\mathcal{M}}_{\mathbf{P},n}(x_0, \mathbf{x})_1} \frac{o([u])}{\prod_{\kappa \in \mathcal{K}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} z^{\beta([u])} x_0 \tag{3.3.V}$$

for all energies simultaneously. Observe that by construction we have multiplicativity of the weights z^β, in the sense that if $[u]$ is obtained by gluing $[u_2]$ to $[u_1]$ at an input we have, after parallel transport to the final output x_0 of $[u]$, the identity $z^{\beta([u])} x_0 = z^{\beta([u_1])} z^{\beta([u_2])} x_0$. In the absence of weighted edges, the boundary characterization of Theorem 3.8 now implies the claim, up to sign; a detailed analysis of compatibility of orientation signs for constructions of this kind appears in [WW15; MWW18].

Finally, in order to handle weighted edges, consider a connected component of $\overline{\mathcal{M}}_{\mathbf{P},n}(x_0, \mathbf{x})_1$ with boundary components arising from the operation (2.2.d) (an edge weight becoming zero or infinite). By the orientation convention described above, together these boundary components contribute a difference $x^{\triangledown}_{\mathsf{m},i} - x^{\blacktriangledown}_{\mathsf{m},i}$ to (3.3.V), precisely as desired. □

Let $e := \sum_i x^{\triangledown}_{\mathsf{m},i}$ be the sum of all of the finitely many formal generators corresponding to minima $x^{\blacktriangledown}_{\mathsf{m},i}$ of f (c.f. (2.3.I)).

Theorem 3.12. *The element $e \in \mathcal{A}$ is a strict unit, in that for all $n > 2$ we have $\mu^n(\dots, e, \dots) = 0$ identically and*

$$\mu^2(e, a) = (-1)^{|a|}\mu^2(a, e) = a \quad \textit{for all homogeneous } a \in \mathcal{A}.$$

Proof. First let $n > 2$, consider an arbitrary product $\mu^n(\mathbf{x})$ determined by the vector $\mathbf{x} = (x_1, \dots, x_{j-1}, x^{\triangledown}_{\mathsf{m},i}, x_{j+1} \dots, x_n)$ for some j, and denote by $\widehat{\mathbf{x}}$ the vector obtained from $\mathbf{x}$ by deleting the jth component. By construction, since the perturbation datum P_Γ respects operation (2.2.e) (forgetting a forgettable edge), each space $\mathcal{M}_{P_\Gamma}(\mathbf{m}, (x_0, \mathbf{x}), \boldsymbol{\beta}) \subset \mathcal{M}_{P_\Gamma}(x_0, \mathbf{x})_0$ is in canonical bijection with $\mathcal{M}_{P_\Gamma}(\mathbf{m}, (x_0, \widehat{\mathbf{x}}), \boldsymbol{\beta})$. The dimension of this latter moduli space is at least 1 less than the expected dimension of a generic stratum of $\mathcal{M}_{P_\Gamma}(x_0, \mathbf{x})_0$, which is zero. It follows that both such spaces must actually be empty, from which we immediately conclude $\mu^n(\mathbf{x}) = 0$.

Now suppose that $n = 2$. In this case each $[u] \in \mathcal{M}_{P_\Gamma}(\mathbf{m}, (x_0, \mathbf{x}), \boldsymbol{\beta})$ is, modulo sphere components, an honest broken Morse gradient flow tree necessarily with inputs $x^{\triangledown}_{\mathsf{m},i}$ and x (in some order) and output $x_0 = x$. But generic points of X lie in the ascending flow manifold of a unique minimum, and by construction the perturbation datum specifying the perturbation of the Morse function on edges labeled by $x^{\triangledown}_{\mathsf{m},i}$ preserves the property that (generically) every point lies in a unique flowline from a minimum. This completes the proof. □

3.4 Refinement and invariance

The Morse–Fukaya algebra $\mathcal{A}$ we have constructed depends on compatible choices of Morse function f on X, background system of stabilizing divisors $\mathbf{D}$, and confining (integral affine) polyhedral cover Θ. Changes of these choices induce comparison maps $\mathcal{A}_0 \to \mathcal{A}_1$ between the algebras that we would respectively construct in each case. We explain how to define these comparison maps below; they will permit us to establish a suitable notion of invariance of choices for $\mathcal{A}$, in addition to providing a convenient ingredient in one formulation of the family Floer functor we will define in Chapter 4.

The natural (and geometric) way to proceed is to exhibit a suitable Morse function, stabilizing divisors, and integral affine cover for $[0,1] \times X$, so that we may directly apply our existing technology to obtain an algebra $\mathcal{A}_*$ interpolating between $\mathcal{A}_0$ and $\mathcal{A}_1$. However, this approach suffers from some technical issues (for example, due to enlarging $\mathcal{O}_{\mathrm{an}}$ with the addition of another variable) which make implementing the strategy in practice unnecessarily cumbersome.

Instead we will prefer to deal with what is essentially a formal version of this construction, combining ideas of [Maz22] and [Sei08b]. Avoiding basically all of the technical issues which would otherwise arise, we pay the simple price that we cannot literally re-use the moduli spaces and associated results produced above. Below we explain the necessary small enhancement of both our constructions and of the corresponding arguments required to complete this strategy.

First, in order to introduce an auxiliary time-tracking parameter key to the construction, denote by $\mathbb{I} = \bigsqcup_{n \in \mathbb{N}} E_n/\sim$ the gluing of disjoint copies $E_n = [-\infty, \infty]$ of the extended real line at their endpoints when laid sequentially end-to-end; writing $\pm\infty_n \in E_n$ to distinguish each such pair of endpoints, we declare $\infty_n = -\infty_{n+1}$ for all $n \in \mathbb{N}$. We think of $\mathbb{I}$ as *time*, divided into an infinite sequence of *epochs* E_n. Of course, $\mathbb{I}$ is topologically again a closed interval, but we view this decomposition as providing a distinguished parametrization of (and indeed metric on) each epoch in $\mathbb{I}$.

Definition 3.13. A *timed* P_Γ-perturbed pseudoholomorphic treed disk (u, τ) is a pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ equipped with a continuous function $\tau : \underline{\Delta} \to \mathbb{I}$ such that

- the background data J and f, stabilizing divisor activities α_κ, and perturbations P_Γ all depend in addition on the *current time* τ,
- for all $v \in \mathrm{Vert}^\bullet(\Delta) \cup \mathrm{Vert}^\circ(\Delta)$ the restriction $\tau|_v$ is constant (i.e. disk and sphere components are constant in time), and
- for all $e \in \mathrm{Edge}^\partial(\Delta)$ and $t \in \mathrm{int}(L_e)$ we have that $(\tau|_{L_e})'(t) = 1$ (i.e. time increases along Morse flow trajectories at uniform speed). .

In particular, given Morse functions f_0 and f_1 on X, we will choose a smooth interpolating family $f_t : \mathbb{I} \times X \to \mathbb{R}$ (not necessarily through Morse functions) for which we have $f_t = f_0$ for all $t \in E_n$ with $n < 0$, and conversely $f_t = f_1$ for all $t \in E_n$ with $n > 0$.

Definition 3.14. A family of maps $h^n : \mathcal{A}^{\otimes n} \to \mathcal{A}'[1-n]$ between A_∞-algebras $\mathcal{A}$ and $\mathcal{A}'$ assemble into an *A_∞-morphism* $h : \mathcal{B} \to C$ precisely when for each $n \geq 0$ and homogeneous elements $a_i \in \mathcal{A}$ of respective degrees $|a_i|$ we have the identity

$$\sum_{j+k\leq n} (-1)^{\heartsuit} h^{n-k+1}(a_1, \ldots, a_j, \mu^k_{\mathcal{A}}(a_{j+1}, \ldots, a_{j+k}), a_{j+k+1}, \ldots, a_n)$$
$$= \sum_{i_1+\ldots+i_d=n} \mu^d_{\mathcal{A}'}(h^{i_1}(a_1, \ldots, a_{i_1}), \ldots, h^{i_d}(a_{i_1+\ldots+i_{d-1}+1}, \ldots, a_n)) \tag{3.4.I}$$

with $\heartsuit = j + \sum_{i=1}^{j} |a_i|$.

Theorem 3.15. *By considering moduli spaces of timed P_Γ-perturbed pseudoholomorphic treed disks, whenever the cover Θ' used to define an instance $\mathcal{A}'$ of the Morse–Fukaya algebra is a refinement of another cover Θ used to define an instance $\mathcal{A}$, there exists a* refinement morphism *$r : \mathcal{A} \to \mathcal{A}'$ of A_∞-algebras.*

Proof sketch. We avoid explicitly enumerating each of our previous constructions in order to make the straightforward small modification needed to incorporate time-dependence in each case, and instead simply explain the new consequences. Existence of applicable working systems of perturbation data follows from precisely the same argument as in the proof of Theorem 2.45. Moduli spaces of timed P_Γ-perturbed treed disks have expected dimension one more than their ordinary counterparts due to presence of the time parameter, and therefore the corresponding counts according to (3.3.I) define an operation of degree one less (in accordance with (3.4.I)). In particular, so long as the cover Θ' is a refinement of Θ and f_t is suitably chosen the formula (3.3.I) continues to be convergent in this setting. The classification theorem corresponding to Theorem 3.8 asserts once again that the only honest boundary components of the 1-dimensional moduli spaces we consider arise from either a Morse gradient flow lines breaking, or forgetting an input. The presence of the time function τ causes the existence of a breaking (necessarily lying on a Morse flow line of infinite length) to imply, after cutting the treed disk at the breaking, that one of the two resulting treed disks has time τ strictly contained in E_n for either all $n < 0$ or all $n > 0$. But, since $f_t = f_0$ or $f_t = f_1$ in either region, the A_∞-relation (3.4.I) exactly enumerates all broken timed treed disks of this form; the left hand side counting the former case, and the right hand side counting the latter. The claim then reduces to a verification of coherence of orientation signs, which one checks directly. $\square$

It is left to the reader to establish a suitable coherence result for the refinement morphisms of Theorem 3.15 (applicable to the curved, filtered, gapped A_∞-algebraic setting, c.f. [Fuk+10a]) which adequately captures the notion of homotopy equivalence that they witness.

Chapter 4

The family Floer functor

In this chapter we construct a family Floer A_∞-functor

$$C : \mathcal{F}_{\mathrm{sec}} \to \text{mod-}\mathcal{A},$$

defined on the *Fukaya category of Lagrangian sections* $\mathcal{F}_{\mathrm{sec}}$ of π and valued in the category mod-$\mathcal{A}$ of *A_∞-modules for $\mathcal{A}$*. That is, $\mathcal{F}_{\mathrm{sec}}$ is the full A_∞-subcategory of the Fukaya category of $\overline{X}$ with objects

$$\mathrm{ob}\,\mathcal{F}_{\mathrm{sec}} = \{L \subset \overline{X} \mid L \cap X \text{ is a section of } \pi : X \to Q\},$$

and morphisms and composition maps determined by the Floer theory of these sections with one another. We begin by explaining how to upgrade the technology of Chapter 2 to produce a model for this category, and for which we may subsequently produce from each $L \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ a Λ-module carrying a natural action of $\mathcal{A}$.

4.1 Bookkeeping upgrades

First, in order to model treed configurations of disks and spheres bounded by a collection of Lagrangian sections of π (possibly also in addition to fibers F_q) it is necessary to enhance the class of perturbed pseudoholomorphic treed disk-like objects we consider. No change to the domains $\underline{\Delta}$ or associated combinatorial types $\Gamma = \Gamma(\Delta)$

they specify is necessary on-the-nose, but it will greatly simplify our subsequent constructions if from the outset we annotate each of the combinatorial types Γ we consider with a particular vector $\boldsymbol{\lambda}$ of Lagrangian section boundary labels associated to the output and each input of Γ.

In order to motivate and then specify these labels precisely, for the remainder of this section fix a background finite collection $\mathbf{L} \subset \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ of Lagrangian sections of π. Define $\mathcal{B}(\mathbf{L}) := \{F_q : q \in Q\} \cup \mathbf{L}$ as a convenient shorthand. The starting point is that we will shortly relax the condition (2.16.c) in the definition of a perturbed pseudoholomorphic treed disk u so that b^u is a function

$$b^u : \pi_0(\partial S_\Delta - \mathrm{Joint}(\Delta)) \longrightarrow \mathcal{B}(\mathbf{L}). \tag{4.1.I}$$

In other words, boundary components of the surface part $S_\Delta - \mathrm{Joint}(\Delta)$ of a treed disk Δ will be allowed to lie on Lagrangian sections $L \in \mathbf{L}$.

Remark 4.1. We note a consequence of this, and thereby extend the definition of b^u to boundary edges: let $e \in \mathrm{Edge}^{\partial}(\Delta)$ be an edge of boundary type, i.e. possibly meeting only disk component boundaries and point vertices. Suppose that the edge e meets $\partial v^\bullet$ at $t(e)$ for some disk component $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$, so that e is adjacent on either side to two (possibly identical) boundary components $c^-, c^+ \subset \partial v^\bullet - \mathrm{Joint}(\Delta)$. This is depicted in Figure 4.1a. These components are canonically ordered by the orientation on $v^\bullet$ and together determine an ordered pair of *edge boundary labels*

$$b^u_t(e) := (b^u(c^-), b^u(c^+)) \in \mathcal{B}(\mathbf{L})^2.$$

Similarly, if the edge e meets components c^- and c^+ of $\partial v^\bullet$ at $h(e)$ for some disk component $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$, then there is again a canonically determined collection of boundary labels $b^u_h(e) := (b^u(c^-), b^u(c^+))$.

When the image of an edge $e \in \mathrm{Edge}^{\partial}(\Delta)$ under u meets a disk boundary component mapping to a Lagrangian section (as detected by b^u_t and b^u_h), we would like to constrain e to be a Morse flow trajectory lying wholly in that section (with respect to the restriction of the global Morse function on X). Otherwise, the image of e is free to be a gradient flow trajectory inside the entirety of X. Also, we would like $b^u_t(e)$ and $b^u_h(e)$ to agree whenever either of their components lies on a Lagrangian section.

As a convenient artifice to enforce both of these constraints, define $\widetilde{\mathcal{B}}(\mathbf{L}) := \{X\} \cup \mathbf{L}$ and let $\rho : \mathcal{B}(\mathbf{L}) \to \widetilde{\mathcal{B}}(\mathbf{L})$ be the natural projection defined on fibers F_q and Lagrangian sections L of π respectively by

$$F_q \mapsto X \quad \text{and} \quad L \mapsto L.$$

The functions b^u_t and b^u_h above descend to maps $\widetilde{b}^u_t$ and $\widetilde{b}^u_h$ valued in $\widetilde{\mathcal{B}}(\mathbf{L})$.

In order to keep consistent track of the boundary labels defined above, we introduce an auxiliary function

$$\widetilde{b} = (\widetilde{b}^-, \widetilde{b}^+) : \mathrm{Edge}^\partial(\Delta) \to \widetilde{\mathcal{B}}(\mathbf{L})^2$$

defined on all boundary edges whatever. Consistency of Lagrangian section boundary labels across edges of $u : \underline{\Delta} \to \overline{X}$ (respecting the orientation of disk components) now amounts to the requirement that $\widetilde{b}^u_t(e) = \widetilde{b}(e)$ and $\widetilde{b}^u_h(e) = \widetilde{b}(e)$ whenever either of these equations make sense (see Figure 4.1b). To ensure consistency along infinite (broken) edges, whenever there is $v^+ \in \mathrm{Vert}^+(\Delta)$ with $h(e_{\mathrm{in}}) = v^+ = t(e_{\mathrm{out}})$ we require that $\widetilde{b}^u(e_{\mathrm{in}}) = \widetilde{b}^u(e_{\mathrm{out}})$ as in Figure 4.1c.

Observe that once the value of b^u has been prescribed on each component of $\partial S_\Delta - \mathrm{Joint}(\Delta)$, the value of the boundary label $\widetilde{b}^u(e)$ of all edges $e \in \mathrm{Edge}^\partial(\Delta)$ is completely determined.

Importantly, the converse of the final conclusion of Remark 4.1 also holds.

Definition 4.2. For $1 \leq i \leq n$ let $v_i \in \mathrm{Vert}^+(\Delta)$ be the ith input of Δ and write e_i for the corresponding unique edge such that $t(e_i) = v_i$. Similarly, write e_0 for the root edge.

The consistency conditions imposed on $\widetilde{b}$ above across all boundary edges whatever imply that the value of $\widetilde{b}(e)$ for $e \in \mathrm{Edge}^{\partial}(\Delta)$ is completely determined by knowledge of $(\lambda_i^-, \lambda_i^+) := \widetilde{b}(e_i)$ for all $0 \leq i \leq n$.

The labels $\lambda_i^{\pm}$ assemble into a vector $\boldsymbol{\lambda} \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+1}$ specifying *(I/O) boundary labels* for the type Γ. After prescribing a particular choice of boundary labels $\boldsymbol{\lambda}$, write $\widetilde{b}^{\Gamma,\boldsymbol{\lambda}} : \pi_0(\partial S_\Delta - \mathrm{Joint}(\Delta)) \to \widetilde{\mathcal{B}}(\mathbf{L})$ for the induced labels on all of the boundary components of $\partial S_\Delta - \mathrm{Joint}(\Delta)$. It will often be convenient to refer to the pair $(\Gamma, \boldsymbol{\lambda})$ as a *labeled (combinatorial) type*.

In order that our general perturbation scheme may support maps from treed disks with the weakened boundary conditions we wish to consider, we introduce an additional parameter to our perturbation data in order to establish transversality for sections. The use of Hamiltonian perturbations for this purpose is very standard; [Sei08b] is a main reference. Thus for all treed disks Δ and boundary label $\boldsymbol{\lambda}$ vector for the type $\Gamma(\Delta)$—in particular inducing a boundary label on each component of $\partial S_\Delta - \mathrm{Joint}(\Delta)$ as in Definition 4.2—choose once and for all a compact subset $\mathcal{S}^{\circledast,H}_{\Delta,\boldsymbol{\lambda}} \subset \overline{\mathcal{S}}_\Delta$ of the surface part of the universal treed disk on which our Hamiltonian perturbations will be supported, such that $\mathcal{S}^{\circledast,H}_{\Delta,\boldsymbol{\lambda}}$

- meets a neighborhood of each boundary joint with adjacent boundary components both labeled by Lagrangian sections $(L, L') \in \widetilde{\mathcal{D}}(\mathbf{L})^2$,
- is disjoint from each sphere component, from a neighborhood of all boundary components labeled by $X \in \widetilde{\mathcal{D}}(\mathbf{X})$, and from the the entirety of each disk component for which all boundary labels are $X \in \widetilde{\mathcal{D}}(\mathbf{L})$.

Recalling the compact subset $\mathcal{S}^{\circledast}_{\Delta} \subset \overline{\mathcal{S}}_\Delta$ of Section 2.2 on which perturbations of the almost complex structure are supported, write $\mathcal{S}^{\circledast,J}_{\Delta} := \mathcal{S}^{\circledast}_{\Delta}$ to avoid confusion.

Let $(\Gamma, \boldsymbol{\lambda})$ be a labeled type and suppose that the treed disk Δ has $\Gamma(\Delta) = \Gamma$. As in [Sei08b] we also choose once and for all a *strip-like end* for a neighborhood of each

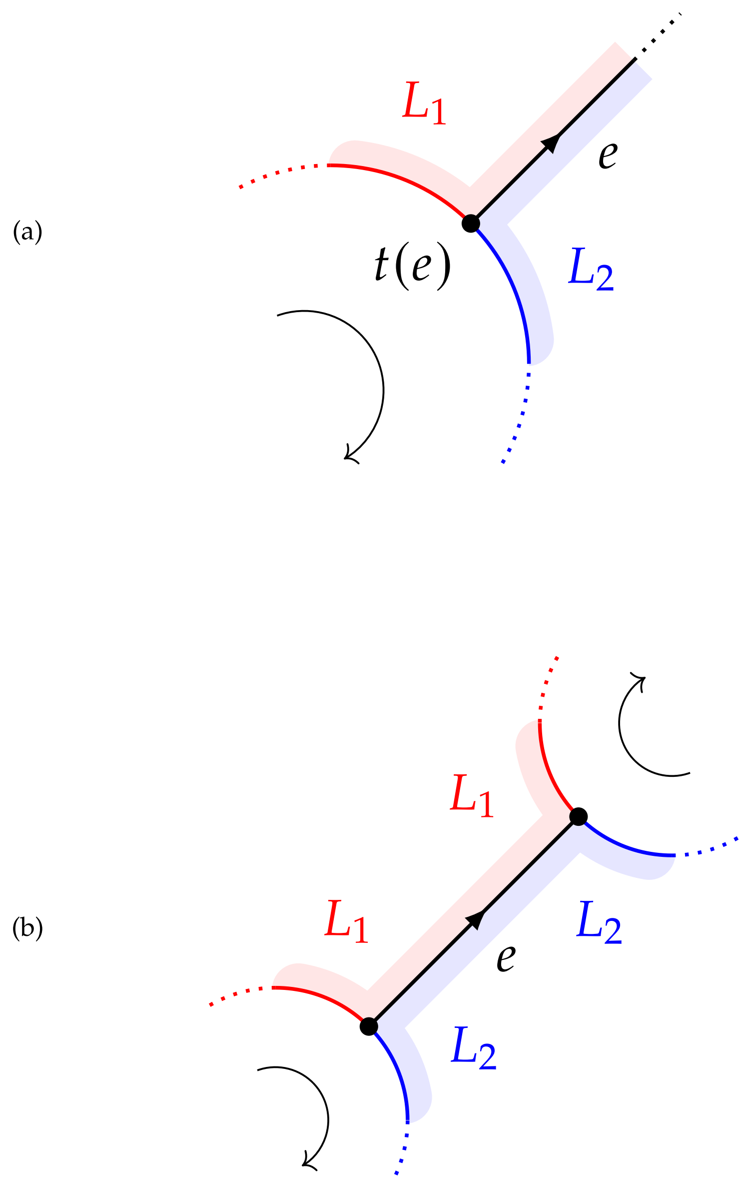

Figure 4.1: Schematic diagrams of the consistency conditions determining the boundary labels $b^u(e)$ of boundary edges $e \in \mathrm{Edge}^\partial(\Delta)$.

(c)

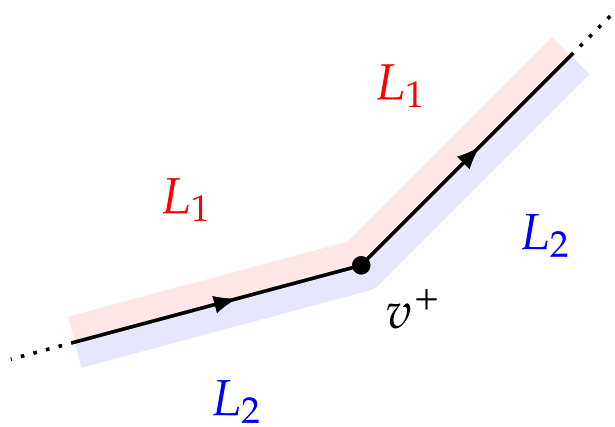

Figure 4.1 (cont.): Schematic diagrams of the consistency conditions determining the boundary labels $b^u(e)$ of boundary edges $e \in \mathrm{Edge}^\partial(\Delta)$.

boundary joint in Δ adjacent to a pair of Lagrangian labels, i.e. a biholomorphism $\psi(s+ti)$ from $[0,\infty]\times[0,1]$ to a neighborhood of each boundary joint $z \in \mathrm{Joint}^\partial(\Delta)$ belonging to some $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$ with adjacent to components $c^\pm \subset \partial v^\bullet$ such that $\widetilde{b^u}(c^\pm) = L^\pm \in \mathbf{L}$. We also choose a model Hamiltonian $H_{L^\pm}$ defined on this neighborhood, depending only on the (ordered) boundary labels $L^\pm$, with $H_{L^\pm}$ a function of t only with respect to the coordinates $s + ti$ induced by ψ.

Definition 4.3. Let $\mathbf{L} \subset \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ be a subset. An $\mathbf{L}$-*perturbation datum* $P_{\Gamma,\lambda}$ *for the labeled type* (Γ,λ) is a perturbation datum for Γ (c.f. Definition 2.14) along with, in addition, a choice of *Hamiltonian perturbations* (writing $\Omega^1_{\mathrm{vert}}(\overline{\mathcal{S}}_\Gamma)$ for the vertical 1-forms[1] on $\overline{\mathcal{S}}_\Gamma$)

$$H_{\Gamma,\lambda} \otimes \alpha_{\Gamma,\lambda} : \overline{\mathcal{S}}_\Gamma \to C^\infty(Q \to \mathbb{R}) \otimes \Omega^1_{\mathrm{vert}}(\overline{\mathcal{S}}_\Gamma)$$

[1] These may be identified with 1-forms on S_Δ for some Δ with $\Gamma(\Delta) = \Gamma$.

which restrict to zero on $\overline{\mathcal{S}}_\Gamma - \mathcal{S}^{\circledast,H}_{\Gamma,\lambda}$. We require that in each strip-like end $H_{\Gamma,\lambda}$ converges to the prechosen Hamiltonian $H_{L^\pm}$ associated to the ordered pair (L^-, L^+), and similarly that $\alpha_{\Gamma,\lambda} = \mathrm{d}t$ in each strip-like end and $\alpha_{\Gamma,\lambda}$ vanishes at the disk boundary. We correspondingly also extend the notion of locality of perturbation data (as in Definition 2.42) to require that each $H_{\Gamma,\lambda} \otimes \alpha_{\Gamma,\lambda}$ similarly factors through the uncrowding construction in the natural way.

An $\mathbf{L}$-perturbation system $\mathbf{P} = \{P_{\Gamma,\lambda}\}_{(\Gamma,\lambda)\in\mathbf{\Gamma}}$ is a collection of $\mathbf{L}$-perturbation data obeying the axioms of Definition 2.15; of course, the degeneration and cut operations naturally extend to operations on labeled types $(\Gamma, \boldsymbol{\lambda})$.

Generalizing the the terminology of [VWX20], we now define the extended kind of treed maps we consider.

Definition 4.4. A *$P_{\Gamma,\lambda}$-perturbed pseudoholomorphic treed (*$\mathbf{L}$*-)polygon* is a pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ (c.f. Definition 2.16), with the exceptions that:

- The boundary label function is extended as in (4.1.I) and (2.16.c) holds with this modification.

- On the tree part u is consistent with the boundary labels in that

$$u(t) \in \widetilde{b}^u_-(e) \cap \widetilde{b}^u_+(e) \quad \text{for all } e \in \mathrm{Edge}^\partial(\Delta) \text{ and } t \in L_e.$$

  This implies, for example, that if in the image of u the head of an edge e meets the boundary of a disk component with adjacent boundaries labeled by sections $L, L' \in \mathbf{L}$ which happen to intersect transversely, then $u|_{L_e}$ is a Morse gradient flow trajectory contained in $L \cap L'$—and is therefore constant.

- The pseudoholomorphic curve equation (2.16.a) no longer holds on-the-nose, and instead is perturbed by the chosen Hamiltonian perturbation $H_{\Gamma,\lambda}$ so that if $\overline{\partial}$ is the operator $\frac{1}{2}(D - J \circ D \circ j)$ and $X_{H_{\Gamma,\lambda}}$ is the corresponding Hamiltonian vector field we have

$$\left(\overline{\partial} u - X_{H_{\Gamma,\lambda}} \otimes \alpha\right)^{0,1}_{J,j} = 0.$$

- The Morse gradient flow equation (2.16.b) holds on each edge L_e for $e \in \mathrm{Edge}^{\partial}(\Delta)$ of the tree part, with respect to a Morse function f_e depending on the boundary label $\widetilde{b}^u(e) \in \widetilde{\mathcal{B}}(\mathbf{L})^2$:
  1. If $\widetilde{b}^u(e) = (X, X)$ then $f_e = f$ is Morse function on all of X.
  2. If $\widetilde{b}^u(e) \in \{(L, X), (X, L), (L, L)\}$ for $L \in \mathbf{L}$ then $f_e = f|_L$ is a Morse function on $L \cap X$.
  3. Otherwise $\widetilde{b}^u(e) = (L, L')$ for $L, L' \in \mathbf{L}$ with $L \neq L'$ and f_e is a Morse function on the set of Floer trajectories of the Hamiltonian perturbation $X_{H_{\Gamma,\lambda}}\, \mathrm{d}t$ from L to L'. Since L and L' are perturbed to meet transversely, this is a Morse function on a discrete set and therefore is a vacuous constraint.

The notion of stability of $P_{\Gamma,\lambda}$-perturbed pseudoholomorphic treed polygons is unchanged (c.f. Definition 2.19), but in order that the moduli spaces we will construct have the expected boundary strata, it will be necessary to treat the stabilizing divisors themselves with some care.

Definition 4.5. A *complete family of systems of stabilizing divisors for* $\pi : X \to Q$ is a choice of, for each finite collection $\mathbf{L} \subset \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$, a system of stabilizing divisors $\mathbf{D}_{\mathbf{L}} = \{D_{\mathbf{L},\kappa}\}_{\kappa \in \mathcal{K}_{\mathbf{L}}}$ for π with respective activity functions $\alpha_{\mathbf{L},\kappa}$ (c.f. Definition 2.24), such that in addition whenever $q \in \alpha^{-1}_{\mathbf{L},\kappa}([0, \infty))$ the divisor $D_{\mathbf{L},\kappa}$ is stabilizing for the entire family $\{F_q\} \cup \mathbf{L}$.

We have the following analogue of Theorem 2.25.

Theorem 4.6. *There exists a complete family of systems of stabilizing divisors for* $\pi : X \to Q$.

Proof. The argument is same as in the proof of Theorem 2.25, recalling that Corollary 2.22 holds for all finite families, and nothing that since our Hamiltonian perturbations are supported away from all marked points, pseudoholomorphicity in a neighborhood of each guarantees positivity of intersection with the applicable chosen divisors. □

Thus fix a background complete family of systems of stabilizing divisors $(\mathbf{D_L})_{\mathbf{L}\subset \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}}$. Despite the choice of an entire family of stabilizing divisors, we will only need to work with P_Γ-perturbed pseudoholomorphic treed polygons $u : \underline{\Delta} \to \overline{X}$ adapted to one system $\mathbf{D_L}$ at a time.

Definition 4.7. A pseudoholomorphic treed $\mathbf{L}$-polygon $u : \underline{\Delta} \to \overline{X}$ is *adapted* to $\mathbf{D_L}$ if u satisfies Definition 2.29 (i.e. is adapted as a treed disk), with the exception that disks $v^\bullet \in \mathrm{Vert}^\bullet(\Delta)$ which have at least 3 boundary marked points and have boundary components labeled (according to $\widetilde{b}^u$) by both L and X need not meet a divisor in $\mathbf{D_L}$.

Perturbed pseudoholomorphic treed polygons are sufficiently general to model all of the moduli spaces we will subsequently need to consider; first, we produce our model of the Fukaya category $\mathcal{F}_{\mathrm{sec}}$. Thus fix $L^-, L^+ \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$. We denote by gen_{L^-,L^+} the time 1 chords of the flow induced by the Hamiltonian $H_{L^\pm}$ chosen at the time we selected our strip-like ends.

Definition 4.8. The *morphism spaces of* $\mathcal{F}_{\mathrm{sec}}$ are defined for all $L, L' \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ by

$$\mathcal{F}_{\mathrm{sec}}(L, L') := \Lambda\langle \mathrm{gen}_{L,L'}\rangle,$$

i.e. the Λ-module freely generated by the (perturbed) intersection points of L and L'.

Definition 4.9. Let $(\Gamma, \boldsymbol{\lambda})$ be a labeled combinatorial type of treed disks with n inputs. A pseudoholomorphic treed $\mathbf{L}$-polygon *(spec)ification* for the labeled type $(\Gamma, \boldsymbol{\lambda})$ is a a spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ of pseudoholomorphic treed disks for which the generators $\mathbf{x}$ and classes $\boldsymbol{\beta}$ are compatible with the boundary labels of Definition 4.2. In other words, we respectively have:

- Write $\mathbf{x} = (x_0, x_1, \ldots, x_n)$. For each $1 \leq i \leq n$ denote the ith input vertex by $v_i \in \mathrm{Vert}^+(\Delta)$ and write e_i for the unique edge such that $t(e_i) = v_i$. Similarly, let e_0 denote the root edge. The induced edge label function $\widetilde{b} : \mathrm{Edge}^{\partial}(\Delta) \to \widetilde{\mathcal{B}}(\mathbf{L})$ determined by $\boldsymbol{\lambda}$ satisfies $\widetilde{b}(e_i) = \lambda_i^{\pm}$ for all $0 \leq i \leq n$. We require that $x_i \in \mathrm{gen}_{\lambda_i^-, \lambda_i^+}$ for each $0 \leq i \leq n$.

- Recall the induced boundary label function $\widetilde{b} : \pi_0(S_\Delta - \mathrm{Joint}(\Delta)) \to \widetilde{\mathcal{B}}(\mathbf{L})$. For all $v \in \mathrm{Vert}^{\bullet}(\Gamma)$ we require that the class β_v belongs to $H_2(X, B)$ for B_v the union of finitely many fibers of π with $\widetilde{b}(c)$ for all $c \in \pi_0(v - \mathrm{Joint}(\Delta))$.

A $(P_{\Gamma,\boldsymbol{\lambda}})$-perturbed pseudoholomorphic treed polygon $u : \underline{\Delta} \to \overline{X}$ adapted to $\mathbf{D_L}$ *obeys the spec* $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ if:

- the map u obeys $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ as in Definition 2.31, and

- the function $\widetilde{b} : \mathrm{Edge}^{\partial}(\Delta) \to (\widetilde{\mathcal{B}}(\mathbf{L}))^2$ is compatible with the actual boundary component labels $b^u : \pi_0(\partial S_\Delta - \mathrm{Joint}(\Delta)) \to \mathcal{B}(\mathbf{L})$ as in Remark 4.1.

The *moduli space of* $(P_{\Gamma,\boldsymbol{\lambda}})$*-perturbed pseudoholomorphic treed polygons* is denoted $\mathcal{M}_{P_{\Gamma,\boldsymbol{\lambda}}}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$.

Because the input and output labels x_i determined by a choice of $\mathbf{x}$ play different algebraic roles depending on whether they label a generator associated to one or more Lagrangian sections, or to the fibration π itself, we will generally distinguish the former with labels $y_0, y_1, \ldots$ from the latter with labels $x_0, x_1, \ldots$ (consistent with the notation of Chapters 2 and 3).

All of the boundary labels $\boldsymbol{\lambda}$ we will consider will be determined systematically based on schematic diagrams corresponding to the particular Floer-theoretic operation in question. Since it is inconvenient to then continuously specify such data componentwise, for notational expediency we make the following definition.

Definition 4.10. For each tuple $(L_0, \ldots, L_n) \subset \mathbf{L}^{n+1}$ the corresponding $\mathcal{F}_{\mathrm{sec}}$-*shaped boundary labels* $\boldsymbol{\lambda}^{(L_i)_{0\le i\le n}} = (\lambda_i^{\pm}) \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+1}$ are defined by specifying that $\lambda_0^- = L_n$, $\lambda_0^+ = L_0$, and for all $1 \le i \le n$ we have $\lambda_i^- = L_{i-1}$ and $\lambda_i^+ = L_i$.

Figure 4.2 schematically depicts a pseudoholomorphic treed polygon obeying a spec $(\mathbf{m}, \mathbf{y}, \boldsymbol{\beta})$ with $\mathbf{y} = (y_0, y_1, y_2)$ for a labeled type with $\mathcal{F}_{\mathrm{sec}}$-shaped boundary labels $\boldsymbol{\lambda}^{(L_0,L_1,L_2)}$. It is precisely configurations of this kind which we will count in order to define the composition maps in the Fukaya category (c.f. Definition 4.15).

Indeed, in order to systematically count configurations obeying specs with $\mathcal{F}_{\mathrm{sec}}$-shaped boundary labels, or the other boundary shapes which we will consider later, we now turn our attention to establishing the appropriate generalizations of transversality (c.f. Corollary 2.46) and compactness/the classification of boundary strata (c.f. Theorem 3.8) for pseudoholomorphic treed polygons.

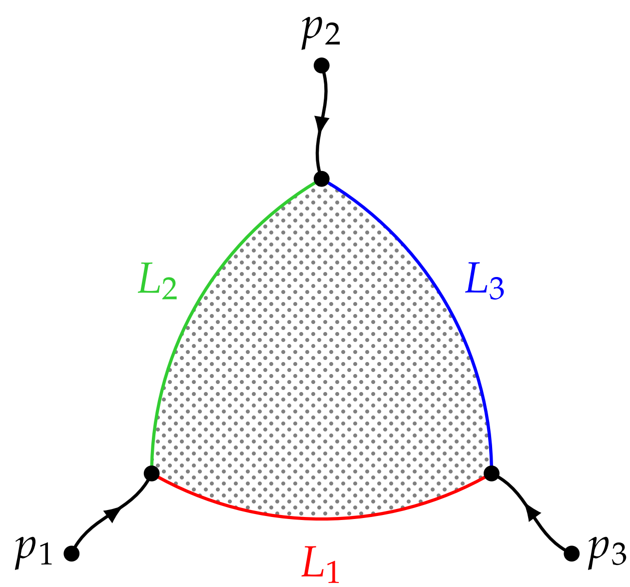

Figure 4.2: A schematic diagram of a pseudoholomorphic treed polygon obeying a spec $(\mathbf{m}, \mathbf{y}, \boldsymbol{\beta})$ with $\mathbf{y} = (y_0, y_1, y_2)$ for a labeled type with $\mathcal{F}_{\mathrm{sec}}$-shaped boundary labels $\boldsymbol{\lambda}^{(L_0,L_1,L_2)}$, belonging to the moduli space $\mathcal{M}_{\mathbf{P},2}(\boldsymbol{\lambda}^{(L_0,L_1,L_2)}, \mathbf{y})_0$ and therefore contributing to the composition law in the Fukaya category (c.f. Definition 4.15).

4.2 Transversality and compactness

Essentially the same arguments as in Section 2.4 produce regular moduli spaces of (suitably constrained) pseudoholomorphic treed polygons—we will just describe the necessary modifications. First, we must slightly alter the space $\mathrm{Map}^{k,p}_{\Delta,\boldsymbol{\lambda}}(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ of class $W^{k,p}$ maps obeying the conditions of Definition 4.4 (with the exception of the perturbed pseudoholomorphic and Morse gradient flow equations)—the issue is that the boundary label $b^u : \pi_0(S_\Delta - \mathrm{Joint}(\Delta)) \to \mathcal{B}(\mathbf{L})$ of some $u \in \mathrm{Map}^{k,p}_{\Delta,\boldsymbol{\lambda}}(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ may suddenly change as we pass around ∂S_Δ over a boundary joint. The remedy, as in [VWX20], is to choose a decay constant $\delta > 0$ sufficiently small and restrict to the subclass of maps which are, near a sufficiently small neighborhood of any joint over which the boundary label changes, given by a constant plus the exponential of a vector field of class $W^{k,p,\delta}$ (i.e. are of class $W^{k,p,\delta}$ for a cylindrical-type metric on the domain).

Definition 4.11. A perturbation system $\mathbf{P} = \{P_{\Gamma,\boldsymbol{\lambda}}\}_{(\Gamma,\boldsymbol{\lambda})\in\boldsymbol{\Gamma}}$ for perturbed treed $\mathbf{L}$-polygons *works* if:

- For each spec $(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ of labeled type $(\Gamma,\boldsymbol{\lambda}) \in \Gamma$ we have that the moduli space $\mathcal{M}_{P_{\Gamma,\boldsymbol{\lambda}}}(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ equipped with its natural topology is a smooth manifold.
- For each $(\Gamma,\boldsymbol{\lambda}) \in \boldsymbol{\Gamma}$, the uncrowding Γ_G of Γ with $G \subset \mathrm{Edge}^\circ(\Delta)$ contained within some maximal ghost sphere tree in u satisfies $(\Gamma_G,\boldsymbol{\lambda}) \in \boldsymbol{\Gamma}$ and the perturbation datum $P_{\Gamma_G,\boldsymbol{\lambda}}$ is induced by $P_{\Gamma,\boldsymbol{\lambda}}$ (in the sense of Definition 2.42).
- Each perturbation datum $P_{\Gamma,\boldsymbol{\lambda}}$ is compatible with $\mathbf{D}_{\mathbf{L}}$.

Defining the subcollection $\mathcal{P}^{\checkmark}_{\Gamma}(\mathbf{P}) \subset \mathcal{P}_{\Gamma}(\mathbf{P})$ of working $\mathbf{L}$-perturbation data as before (as a subset of an honest Banach manifold), we have the following analogue of Theorem 2.45 and its immediate corollary.

Theorem 4.12. *For each combinatorial type Γ and working $\mathbf{L}$-perturbation system $\mathbf{P}$ the subset $\mathcal{P}^{\checkmark}_{\Gamma}(\mathbf{P}) \subset \mathcal{P}_{\Gamma}(\mathbf{P})$ is comeager.*

Proof. We just explain the necessary modification of the proof of Theorem 2.45; the situation is that we again must verify that an arbitrary element $\eta_S \oplus \eta_T$ of the orthogonal complement of the image of the linearization of the section $d_S \oplus d_T$ (now with the addition of a Hamiltonian perturbation term) at some

$$((J_\Gamma, f_\Gamma, H_{\Gamma,\lambda}), \psi : \Delta' \cong \Delta, u : \underline{\Delta} \to \overline{\mathrm{X}}) \in \mathcal{P}^l_\Gamma \times \mathcal{M}_{\Gamma,\Delta} \times \mathrm{Map}^{k,p,\delta}_{\Delta,\lambda}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$$

in the zero locus of $d_S \oplus d_T$ is identically zero.

Recall that, by definition the Hamiltonian perturbation $H_{\Gamma,\lambda}$ furnished by any **L**-perturbation datum P_Γ is zero on all sphere components. Therefore precisely the same argument as in the treed disk case shows that η_S and η_T are zero respectively on $L_e \subset \underline{\Delta}$ for all edges $e \in \mathrm{Edge}(\Delta)$ and on all sphere components $v \in \mathrm{Vert}^\circ(\Delta)$ independent of whether u is constant there.

If u is nonconstant on a disk component $v \in \mathrm{Vert}^\bullet(\Delta)$ then since $\mathcal{S}^{\circledcirc,J}_\Delta$ meets v in an open set disjoint from the support $\mathcal{S}^{\circledcirc,H}_\Delta$ of the Hamiltonian perturbation we again conclude $\eta_S|_v = 0$. Our previous methods fail when u is constant on a disk component v with Lagrangian section boundary, but in precisely this situation v must meet $\mathcal{S}^{\circledcirc,H}_\Delta$ in an open set (on which the Hamiltonian perturbation may be varied), and pseudoholomorphicity implies $\eta_S|_v = 0$ identically on the entirety of v once again. □

Corollary 4.13. *There exists a complete working* **L***-perturbation system* **P***.*

Proof. The proof is identical to the proof of Corollary 2.46. □

As the direct analogy of (3.2.I) we define

$$\mathcal{M}_{\mathbf{P},n}(\lambda, \mathbf{x}) := \bigcup_{\substack{n(\Gamma) = n \text{ and } \Gamma \text{ is stable} \\ \text{and minimally degenerate}}} \; \bigcup_{\substack{(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}) \text{ is a spec} \\ \text{for } (\Gamma, \lambda)}} \mathcal{M}_{P_{\Gamma,\lambda}}(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta}).$$

Let $\mathcal{M}_{\mathbf{P},n}(\lambda, \mathbf{x})_d \subset \mathcal{M}_{\mathbf{P},n}(\lambda, \mathbf{x})$ be the subspace of expected dimension d, and let

$\overline{\mathcal{M}}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})_d$ be its natural compactification. Let $\mathcal{M}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})$ and $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})$ be the corresponding subspaces with energy at most E.

All of the operations on specs $(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ induced by the domain operations (2.2.a)–(2.2.d) preserve the boundary labels $\boldsymbol{\lambda}$. Therefore Theorem 3.9 applies (substituting only the correct notion of spec for pseudoholomorphic treed $\mathbf{L}$-polygons), and we immediately obtain the following generalization of Theorem 3.8.

Theorem 4.14. *Let* $\mathbf{P}$ *be a complete working* $\mathbf{L}$*-perturbation system. For all* $\mathbf{x} \in (\mathrm{gen}_{\mathbf{L}} f)^{n+1}$ *and compatible boundary labels* $\boldsymbol{\lambda} \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+1}$ *we have that:*

(4.14.a) for each $E > 0$ *the spaces* $\mathcal{M}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})_0$ *and* $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})_1$ *are compact, and*

(4.14.b) the boundary of $\overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})_1$ *is the (disjoint) union of all*

$$\mathcal{M}_{P_{\Gamma',\lambda'}}(\mathbf{m}',\mathbf{x}',\boldsymbol{\beta}') \quad \textit{taken over all strata} \quad \mathcal{M}_{P_{\Gamma,\lambda}}(\mathbf{m},\mathbf{x},\boldsymbol{\beta}) \subset \overline{\mathcal{M}}^{\leq E}_{\mathbf{P},n}(\boldsymbol{\lambda},\mathbf{x})_1,$$

where $(\mathbf{m}',\mathbf{x}',\boldsymbol{\beta}')$ *is a spec for the type* (Γ',λ') *obtained from the spec* $(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ *for* (Γ,λ) *by a single application of the operation* (2.2.a) *or* (2.2.d)*.*

As a first application of these constructions, we define the familiar composition law in the Fukaya category.

Definition 4.15. For $L_0,\dots,L_n \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ define a map

$$\mu^n : \mathcal{F}_{\mathrm{sec}}(L_0,L_1)\otimes\cdots\otimes\mathcal{F}_{\mathrm{sec}}(L_{n-1},L_n) \to \mathcal{F}_{\mathrm{sec}}(L_0,L_n)[2-n]$$

on each $\mathbf{y} = (y_1,\dots,y_n)$ by the assignment

$$\mu^n(\mathbf{y}) := (-1)^{\heartsuit} \sum_{\substack{y_0\in \mathrm{gen}_{L_0,L_n} f \\ [u]\in\mathcal{M}_{\mathbf{P},n}(\lambda^{(L_0,\dots,L_n)},(y_0,\mathbf{y}))_0}} \frac{o([u])}{\prod_{\kappa\in\mathcal{K}_{\mathbf{L}}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} \cdot T^{\omega(u)} y_0 \tag{4.2.I}$$

with $\heartsuit = n + \sum_{k=1}^{n} kI(y_k)$.

Theorem 4.16. *Formula* (4.2.I) *endows $\mathcal{F}_{sec}$ with the structure of an A_∞-category* [Sei08a]*, in that for each $n \geq 0$ and homogeneous elements $a_i \in \mathcal{F}_{\mathrm{sec}}(L_{i-1}, L_i)$ of respective degrees $|a_i|$ we have the A_∞-identity*

$$0 = \sum_{j+k \leq n} (-1)^{\heartsuit} \mu^{n-k+1}(a_1, \ldots, a_j, \mu^k(a_{j+1}, \ldots, a_{j+k}), a_{j+k+1}, \ldots, a_n), \tag{4.2.II}$$

with $\heartsuit = (-1)^{j+\sum_{i=1}^{j}|a_i|}$.

Proof. As in the proof of Theorem 3.11, the claim follows from combining Theorem 4.14 with the fact that all operations on **L**-polygon specs induced by (2.2.a)–(2.2.d) preserve boundary labels. □

Remark 4.17. Via Morse-theoretic constructions for a single Lagrangian L Charest–Woodward [CW22] produce a strictly unital model for the algebra $\mathcal{F}_{\mathrm{sec}}(L, L)$ which is essentially a special case of our machinery and is thereby trivially compatible. With minimal effort one may substitute this model $\mathcal{F}_{\mathrm{sec}}(L, L)$ above and below, and produce a strictly unital A_∞-category $\mathcal{F}_{\mathrm{sec}}$ together with a strictly unital family Floer functor.

4.3 Fukaya mirror modules

In order to define the module $C(L)$ associated to each $L \in \mathrm{ob}\, \mathcal{F}_{\mathrm{sec}}$ by the family Floer functor we introduce the notion of an "anchor path". For this purpose, choose a distinguished section $L_* \in \mathrm{ob}\, \mathcal{F}_{\mathrm{sec}}$ selecting a basepoint of each fiber of $\pi : X \to Q$.

Definition 4.18. An *anchor* $[\gamma]$ *on* $L \in \mathrm{ob}\, \mathcal{F}_{sec}$ is a homotopy class of paths represented by $\gamma : [0, 1] \to F_q$, contained in a single fiber F_q and constrained so that

- $\gamma(0) \in \mathrm{crit}\, f|_L$ is a critical point of f restricted to L and
- $\gamma(1) \in L_*$ is the basepoint of F_q.

Denote the collection of all anchors on L by $\mathrm{gen}_{L\to L_*} f$, and when the Lagrangian L to which $x \in \mathrm{crit}\, f|_L$ belongs is understood, write $\mathrm{gen}_{x\to L_*} f \subset \mathrm{gen}_{L\to L_*} f$ for the subset of anchors $[\gamma]$ on L at x (i.e. for which $\gamma(0) = x$). Figure 4.3 depicts a schematic example.

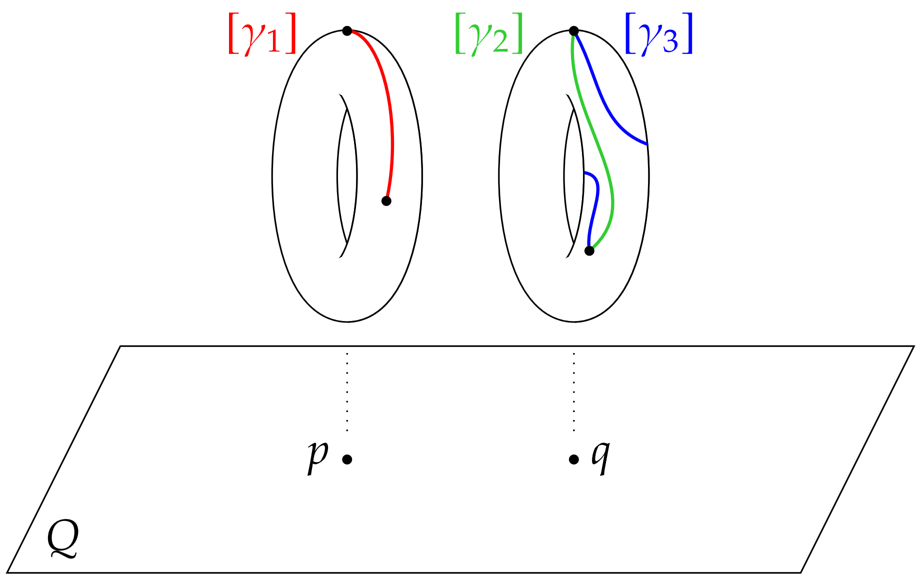

Figure 4.3: A schematic diagram depicting three distinct anchors on the same Lagrangian section $L \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$.

As a building block toward the Fukaya mirror module $C(L)$ to L, observe that for any such $x \in \mathrm{crit}\, f|_L$ we may form the free Λ-module $M_{L,x} := .\Lambda\langle \mathrm{gen}_{x\to L_*} f\rangle$. Now, $M_{L,x}$ carries a natural action of $H_1(F_{\pi(x)})$ arising from the fact that each formal difference $[\gamma] - [\gamma']$ of elements of $\mathrm{gen}_{x\to L_*} f$ determines an element of $H_1(F_{\pi(x)})$ by concatenating γ with the reverse of γ'. This action extends to a Λ-linear action on $M_{L,x}$, and as in Section 3.1 completes to a free rank 1 $\mathcal{O}_{\mathrm{an}}|_{\theta_{\pi(x)}}$-module $\widehat{M}_{L,x}$.

We now define the family Floer functor on $L \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ by letting $C(L)$ be

$$C(L) := \bigoplus_{x \in \mathrm{gen}\, f|_L} \widehat{M}_{L,x}.$$

In order to furnish $C(L)$ with the structure of an A_∞-module, we must provide module action maps

$$\triangleleft^n : C(L) \otimes \mathcal{A}^{\otimes n} \to C(L),$$

and for this purpose we define a new class of boundary shapes.

Definition 4.19. For each tuple $(L_0, \ldots, L_n) \subset \mathbf{L}^{n+1}$ the corresponding $\mathcal{F}_{\mathrm{sec}}$*-shaped boundary labels* $\boldsymbol{\lambda}^{(L_i)_{0\le i\le n}} = (\lambda_i^{\pm}) \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+1}$ are defined by specifying that $\lambda_0^- = L_n$, $\lambda_0^+ = L_0$, and for all $1 \le i \le n$ we have $\lambda_i^- = L_{i-1}$ and $\lambda_i^+ = L_i$.

Definition 4.20. For each $L \in \mathbf{L}$ and $n \in \mathbb{N}$ the corresponding *module-shaped boundary labels* $\boldsymbol{\lambda}^{L|n} = (\lambda_i^{\pm}) \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+2}$ are defined by $(\lambda_0^-, \lambda_0^+) = (\lambda_1^-, \lambda_1^+) = (L, X)$ and $(\lambda_i^-, \lambda_i^+) = (X, X)$ for all $2 \le i \le n+2$. A schematic configuration of this type is depicted in Figure 4.4.

Roughly speaking, configurations of the type depicted in Figure 4.4 contribute to the coefficient of y_0 in $y_1 \triangleleft^2 (x_1, x_2)$. However, y_0 and y_1 as specified are not anchors; in general a pseudoholomorphic treed disk $u : \underline{\Delta} \to \overline{X}$ of spec $(\mathbf{m}, (y_0, y_1, \mathbf{x}), \boldsymbol{\beta})$ for the labeled type $(\Gamma, \boldsymbol{\lambda})$ induces a canonical map from anchors $[\gamma] \in \mathrm{gen}_{L\to L_*} f$ with $\gamma(0) = y_1$ to anchors $[\gamma']$ with $\gamma(0) = y_0$, as we now describe.

Remark 4.21. Let $u : \underline{\Delta} \to \overline{X}$ be a pseudoholomorphic treed polygon of spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for $(\Gamma, \boldsymbol{\lambda})$. The image of the tree part $T_\Delta \subset \underline{\Delta}$ under $\pi \circ u$ is a tree consisting of the union[2] of Morse gradient flow trajectories in Q.

Suppose that an input vertex $v_i^+ \in \mathrm{Vert}^+(\Delta)$, numbered with respect to the canonical ordering, is such that $u(v_i^+)$ and $u(v_0)$ lie in the same connected component of $u(T_\Delta)$. In this case there is a (unique) homotopy class of paths from $\pi(x_i)$ to $\pi(x_0)$ contained in $(\pi \circ u)(T_\Delta) \subset Q$, which we call the *base path associated to the input* v_i^+. Denote a choice of representing path by $\psi_i^u : [0,1] \to Q$.

[2]This is precisely a disjoint union of broken Morse gradient flow trees in Q in the sense of Theorem 1.5.

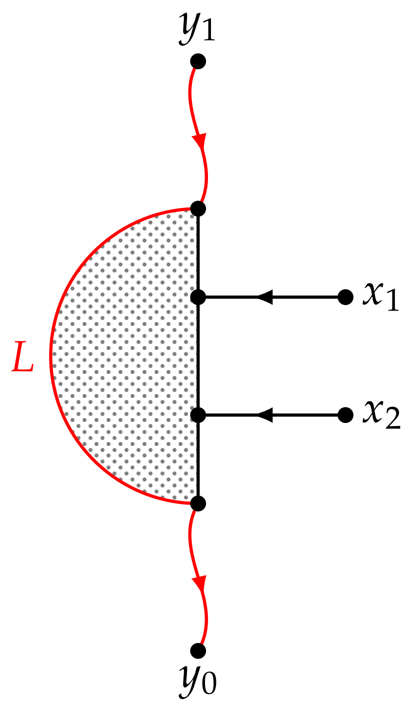

Figure 4.4: A schematic diagram of a pseudoholomorphic treed polygon obeying a spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for a labeled type with module-shaped boundary labels $\boldsymbol{\lambda}^{L|2}$, belonging to the moduli space $\mathcal{M}_{\mathbf{P},3}(\boldsymbol{\lambda}^{L|2}, \mathbf{x})_0$ and therefore contributing to the to the module action map for $C(L)$ (c.f. Definition 4.22).

Each of the shapes of boundary labels we will consider have the following property: whenever we have that precisely one of $\lambda_i^- = X$ or $\lambda_i^+ = X$ holds, then $u(v_i^+)$ and $u(v_0)$ lie in the same connected component of $u(T_\Delta)$. Moreover, in this case the path ψ_i^u admits a unique lift, up to homotopy in $u(\partial S_\Delta \cup T_\Delta)$, to a path $\widehat{\psi}_i^u : [0,1] \to X$ that traverses only components $c \in \pi_0(\partial S_\Delta - \mathrm{Joint}(\Delta))$ labeled by a fiber $b^u(c) = F_q$ of $\pi : X \to Q$. A schematic diagram is depicted in Figure 4.5.

Thus given a pseudoholomorphic treed polygon $u : \underline{\Delta} \to \overline{X}$ obeying a spec $(\mathbf{m}, (y_0, y_1, \mathbf{x}), \boldsymbol{\beta})$ with boundary labels of module type and an anchor $[\gamma] \in \mathrm{gen}_{L \to L_*} f$ with $\gamma(0) = y_1$ we canonically obtain an anchor $[\gamma']$ by parallel transport of γ along the path $\widehat{\psi}_1^u$ in Q. This parallel transport sweeps out a rectangle α_1 bounded by L_* and

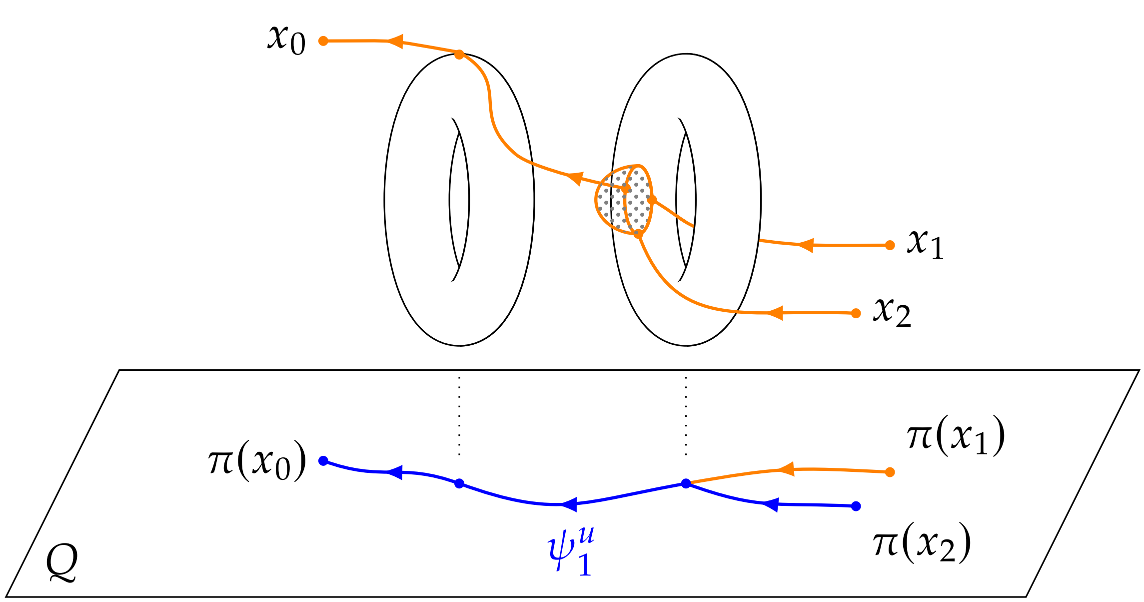

Figure 4.5: A schematic diagram of a base path associated to a pseudoholomorphic treed polygon with multiple inputs.

the image of $\widehat{\psi}_1^u$ of symplectic area $\omega(\alpha_1)$ (as depicted in Figure 4.6). As a convenient shorthand, we define

$$[\gamma] \triangleleft \widehat{\psi}_1^u := T^{\omega(\alpha)}[\gamma'].$$

Definition 4.22. For $L \in \operatorname{ob} \mathcal{F}_{\mathrm{sec}}$ and $n \geq 1$ define a map

$$\triangleleft^{n-1} : C(L) \otimes \mathcal{A}^{\otimes n-1} \to C(L)[2-n]$$

on each $[\gamma] \in \operatorname{gen}_{L \to L_*} f$ and $\mathbf{x} = (x_1, \ldots, x_{n-1}) \in (\operatorname{gen} f)^{n-1}$ by the assignment

$$[\gamma] \triangleleft^{n-1} \mathbf{x} := (-1)^{\heartsuit} \sum_{\substack{y_0 \in \operatorname{gen} f|_L \\ [u] \in \mathcal{M}_{\mathbf{P},n}(\lambda^{L|n}, (y_0, \gamma(0), \mathbf{x}))_0}} \frac{o([u])}{\prod_{\kappa \in \mathcal{K}_{\mathbf{L}}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} \cdot z^{\beta([u])}([\gamma] \triangleleft \widehat{\psi}_1^u) \quad (4.3.\mathrm{I})$$

with $\heartsuit = I(\gamma(0)) + \sum_{k=1}^{n-1} kI(x_i)$.

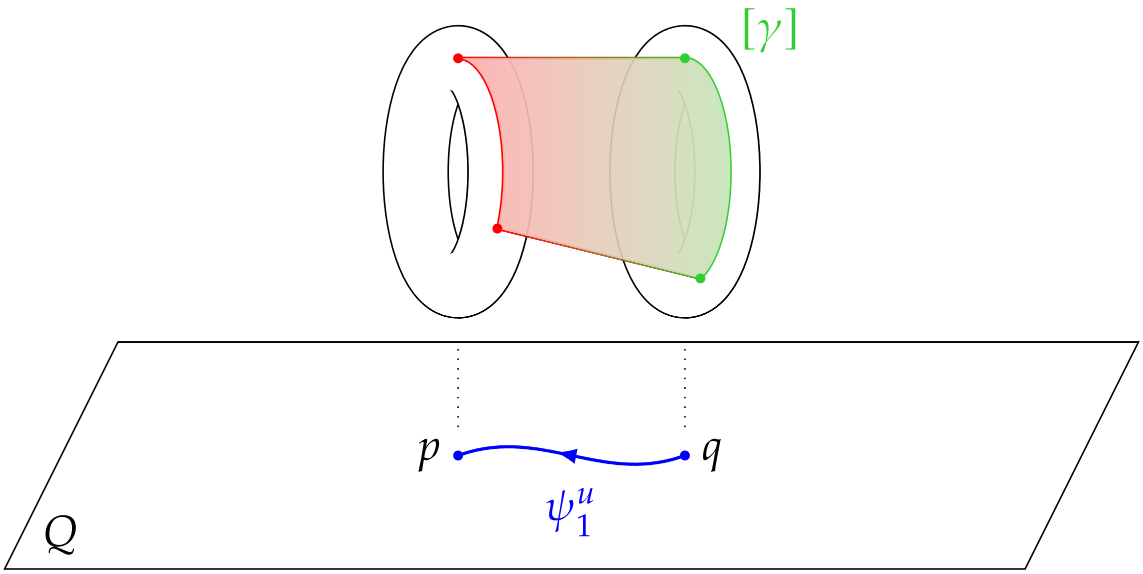

Figure 4.6: A schematic diagram of the symplectic area in X swept out by the parallel transport of the anchor $[\gamma]$ along $\widehat{\psi}_1^u$.

Recall that this operation essentially counts treed polygon configurations such as those schematically depicted in Figure 4.4.

Theorem 4.23. *Suppose that for some ω-compatible J' the section L does not bound nonconstant J'-holomorphic disks. Formula (4.3.I) endows $C(L)$ with the structure of an A_∞-module for $\mathcal{A}$, in that for each $n \geq 1$ and homogeneous elements $m \in C(L)$ and $a_i \in \mathcal{A}$ of respective degrees $|a_i|$ we have the A_∞-identity*

$$\sum_{j+k\leq n-1} (-1)^{\heartsuit} m \triangleleft^{n-k} (a_1, \ldots, a_j, \mu^k(a_{j+1}, \ldots, a_{j+k}), a_{j+k+1}, \ldots, a_{n-1})$$
$$= \sum_{j\leq n-1} \left(m \triangleleft^j (a_1, \ldots, a_j)\right) \triangleleft^{n-j-1} (a_{j+1}, \ldots, a_{n-1}) \tag{4.3.II}$$

with $\heartsuit = (-1)^{1+j+\sum_{i=1}^{j}|a_i|}$.

Proof. Once again, we appeal to Theorem 4.14; we consider all possible results of applying the operation (2.2.a) to a spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for a labeled type $(\Gamma, \lambda^{L|n})$ contributing to $\mathcal{M}_{\mathbf{P},n}(\boldsymbol{\lambda}^{L|n}, \mathbf{x})_1$. Terms on the left hand side of (4.3.II) correspond precisely those obtained by an application of the breaking operation (2.2.a) to an edge $e \in \mathrm{Edge}^{\partial}(\Gamma)$ with label $\widetilde{b}(e) = (X, X)$, while terms on the right hand side arise whenever $\widetilde{b}(e) = (L, X)$. All such breakings are enumerated above, and by the consistency relation which $\widetilde{b}$ must obey along all boundary edges, given the boundary labels $\boldsymbol{\lambda}^{L|n}$ of module-shape the only other possible edge label is $\widetilde{b}(e) = (L, L)$. By the hypothesis that L does not bound disks of positive symplectic area the claim now follows up to a careful check of orientations signs as in [WW15; MWW18]. □

Remark 4.24. In general, without the hypothesis that $L \in \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ does not bound disks of positive symplectic area, the action map (4.3.I) equips $C(L)$ with the structure of a curved $\mathcal{A}$-module. Indeed, there is a natural $(\mathcal{F}_{\mathrm{sec}}(L, L), \mathcal{A})$-bimodule structure on $C(L)$ (independent of any hypothesis). From this perspective one sees that curvature of $\mathcal{F}_{\mathrm{sec}}(L, L)$ prevents us from restricting the bimodule structure yielding an honest right $\mathcal{A}$-module structure, since the corresponding A_∞-module relation no longer holds on-the-nose as a special case of the A_∞-bimodule relation.

4.4 Functoriality and invariance

Having defined the family Floer functor on objects, we next specify its action on composable tuples of morphisms. This necessitates the definition of one final boundary shape.

Definition 4.25. For each $L_0, \ldots, L_n \in \mathbf{L}$ and $m \in \mathbb{N}$ the corresponding *morphism-shaped boundary labels* $\boldsymbol{\lambda}^{(L_0,\ldots,L_n)|m} = (\lambda_i^{\pm}) \in \left(\widetilde{\mathcal{B}}(\mathbf{L})^2\right)^{n+m+2}$ are defined by $(\lambda_0^-, \lambda_0^+) = (L_n, X)$, $(\lambda_i^-, \lambda_i^+) = (L_{n-i}, L_{n-i+1})$ for all $1 \leq i \leq n$, $(\lambda_{n+1}^-, \lambda_{n+1}^+) = (L_0, X)$, and $(\lambda_{n+i}^-, \lambda_{n+1+i}^+) = (X, X)$ for all $1 \leq i \leq m$. A schematic configuration of this type is depicted in Figure 4.7.

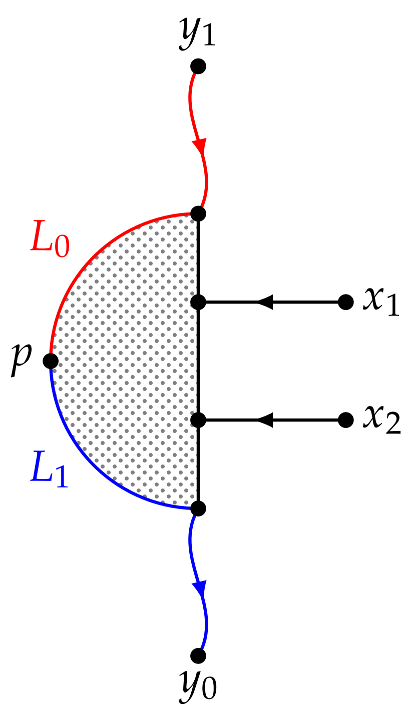

Figure 4.7: A schematic diagram of a pseudoholomorphic treed polygon obeying a spec $(\mathbf{m}, \mathbf{x}, \boldsymbol{\beta})$ for a labeled type with morphism-shaped boundary labels $\boldsymbol{\lambda}^{(L_0,L_1)|2}$, belonging to the moduli space $\mathcal{M}_{\mathbf{P},4}(\boldsymbol{\lambda}^{(L_0,L_1)|2}, \mathbf{x})_0$ and therefore contributing to the morphism maps of the family Floer functor (c.f. Definition 4.26).

Definition 4.26. For $L_0, \ldots, L_n \in \mathbf{L}$, $p_i \in \mathcal{F}_{\mathrm{sec}}(L_{i-1}, L_i)$ for all $1 \leq i \leq n$, and $n' \geq 1$ define a map

$$C^n(p_1, \ldots, p_n)^{n'-1} : C(L_0) \otimes \mathcal{A}^{\otimes n'-1} \to C(L')[1-n-n']$$

on each $[\gamma] \in \mathrm{gen}_{L_0 \to L_*} f$ and $\mathbf{x} = (x_1, \ldots, x_{n'-1}) \in (\mathrm{gen}\, f)^{n-1}$ by the assignment

$$C^n(p_1, \ldots, p_n)^{n'-1}([\gamma], \mathbf{x}) := (-1)^{\heartsuit} \sum_{\substack{y_0 \in \mathrm{gen}\, f|_{L_n} \\ [u] \in \mathcal{M}_{\mathbf{P},n}(\lambda, (y_0, p_1, \ldots, p_n, \gamma(0), \mathbf{x}))_0}} \frac{o([u])}{\prod_{\kappa \in \mathcal{K}_{\mathbf{L}}} \#\mathrm{Mark}_\kappa(\Gamma([u]))!} \cdot z^{\beta([u])}([\gamma] \triangleleft \widehat{\psi}^u_1)$$

with $\heartsuit = I(\gamma(0)) + \sum_{k=1}^{n-1} kI(x_i)$ and $\boldsymbol{\lambda} = \boldsymbol{\lambda}^{(L_0,\ldots,L_n)|n'-1}$.

Theorem 4.27. *Formula (4.26) defines components C^n which assemble into an A_∞-functor $C : \mathcal{F}_{\mathrm{sec}} \to \mathrm{mod}\text{-}\mathcal{A}$, in that for each $n \geq 0$ and homogeneous elements $b_i \in \mathcal{F}_{\mathrm{sec}}(L_{i-1}, L_i)$ for all $1 \leq i \leq n$ of respective degrees $|b_i|$ we have the identity*

$$\sum_{j+k\leq n} (-1)^{\heartsuit} C^{n-k+1}(b_1, \ldots, b_j, \mu^k(b_{j+1}, \ldots, b_{j+k}), b_{j+k+1}, \ldots, b_n)$$
$$= \sum_{i_1+\ldots+i_d=n} \circ^d(C^{i_1}(b_1, \ldots, b_{i_1}), \ldots, C^{i_d}(b_{i_1+\ldots+i_{d-1}+1}, \ldots, b_n)) \tag{4.4.I}$$

with $\heartsuit = (-1)^{j+\sum_{i=1}^{j}|b_i|}$ and $\circ^d$ denoting the d-ary composition of morphisms in $\mathrm{mod}\text{-}\mathcal{A}$.

Proof. The equation (4.4.I) is an equality of pre-$\mathcal{A}$-module morphisms; thus fix homogeneous elements $m \in C(L_0)$ and $a_i \in \mathcal{A}$ for $1 \leq i \leq n'$ of respective degrees $|m|$ and $|a_i|$. Unrolling the definition of the composition in mod-$\mathcal{A}$, since $\circ^k = 0$ identically for all $k > 2$ we have the equivalent identity [Sei08b, Subsection 1j]

$$\sum_{j+k\leq n} (-1)^{\heartsuit} C^{n-k+1}(b_1, \ldots, b_j, \mu^k(b_{j+1}, \ldots, b_{j+k}), b_{j+k+1}, \ldots, b_n)$$
$$= \circ^1(C^n(b_1, \ldots, b_n)) + \sum_{j\leq n} \circ^2(C^j(b_1, \ldots, b_j), C^{n-j}(b_{j+1}, \ldots, b_n)),$$

or explicitly (acting on the tuple $(m, a_1, \ldots, a_{n'})$)

$$\sum_{j+k\leq n} (-1)^{\heartsuit} C^{n-k+1}(b_1, \ldots, b_j, \mu^k(b_{j+1}, \ldots, b_{j+k}), b_{j+k+1}, \ldots, b_n)^{n'}(m, a_1, \ldots, a_{n'})$$
$$= \circ^1(C^n(b_1, \ldots, b_n))^{n'}(m, a_1, \ldots, a_{n'})$$
$$+ \sum_{j\leq n} \circ^2(C^j(b_1, \ldots, b_j), C^{n-j}(b_{j+1}, \ldots, b_n))^{n'}(m, a_1, \ldots, a_{n'}).$$

Now, by definition we have

$$\begin{aligned}
&\circ^1\,(C^n(b_1,\dots,b_n))^{n'}(m,a_1,\dots,a_{n'})\\
&\quad=\sum_{j+k\le n'} C^n(b_1,\dots,b_n)^{n'-k+1}(m,a_1,\dots,a_j,\mu^k(a_{j+1},\dots,a_{j+k}),a_{j+k+1},\dots,a_{n'})\\
&\qquad\qquad+\sum_{j\le n'}(C^n(b_1,\dots,b_n)^j(m,a_1,\dots,a_j))\triangleleft^{n'-j}(a_{j+1},\dots,a_{n'})\\
&\qquad\qquad+\sum_{j\le n'} C^n(b_1,\dots,b_n)^{n'-j}(m\triangleleft^j(a_1,\dots,a_j),a_{j+1},\dots,a_{n'})
\end{aligned}$$

and

$$\begin{aligned}
&\sum_{j\le n}\circ^2(C^j(b_1,\dots,b_j),C^{n-j}(b_{j+1},\dots,b_n))^{n'}(m,a_1,\dots,a_{n'})\\
&\quad=\sum_{j\le n}\sum_{k\le n'} C^{n-j}(b_{j+1},\dots,b_n)^{n'-k}(C^j(b_1,\dots,b_j)^k(m,a_1,\dots,a_k),a_{k+1},\dots,a_{n'}).
\end{aligned}$$

In summary, we must verify that all of the terms of the five sums

$$\begin{aligned}
&\sum_{j+k\le n}(-1)^{\heartsuit}C^{n-k+1}(b_1,\dots,b_j,\mu^k(b_{j+1},\dots,b_{j+k}),b_{j+k+1},\dots,b_n)^{n'}(m,a_1,\dots,a_{n'})\\
&\quad=\sum_{j\le n}\sum_{k\le n'} C^{n-j}(b_{j+1},\dots,b_n)^{n'-k}(C^j(b_1,\dots,b_j)^k(m,a_1,\dots,a_k),a_{k+1},\dots,a_{n'})\\
&\qquad\qquad+\sum_{j+k\le n'} C^n(b_1,\dots,b_n)^{n'-k+1}(m,a_1,\dots,a_j,\mu^k(a_{j+1},\dots,a_{j+k}),a_{j+k+1},\dots,a_{n'})\\
&\qquad\qquad+\sum_{j\le n'}(C^n(b_1,\dots,b_n)^j(m,a_1,\dots,a_j))\triangleleft^{n'-j}(a_{j+1},\dots,a_{n'})\\
&\qquad\qquad+\sum_{j\le n'} C^n(b_1,\dots,b_n)^{n'-j}(m\triangleleft^j(a_1,\dots,a_j),a_{j+1},\dots,a_{n'}) \qquad (4.4.\mathrm{II})
\end{aligned}$$

are accounted for.

As in the proof of Theorem 4.23, each sum in (4.4.II) corresponds to a family of possible breakings (i.e. the operation (2.2.a)) on $(\mathbf{m},\mathbf{x},\boldsymbol{\beta})$ for a labeled type $(\Gamma,\boldsymbol{\lambda}^{(L_0,\dots,L_n)|n'})$.

In each case an edge $e \in \mathrm{Edge}^{\partial}(\Gamma)$ which breaks at a boundary stratum may have possible boundary labels $\widetilde{b}(e) = (X, X)$, $\widetilde{b}(e) = (L_i, X)$, and $\widetilde{b}(e) = (L_i, L_j)$.

The second sum in the right hand side of (4.4.II) exactly account for all possible ways edges labeled with $\widetilde{b}(e) = (X, X)$ can break, thus introducing an intermediary multiplication in $\mathcal{A}$. All possible ways edges labeled with $\widetilde{b}(e) = (L_i, X)$ can break are accounted for by the first, third, and fourth terms of the right hand side of (4.4.II); the module operations $\triangleleft$ arise precisely when $i = 0$ (for the fourth sum) and when $i = n$ (the third sum). Finally, the sum on the left hand side of (4.4.II) similarly accounts for all possible breaking of edges with label $\widetilde{b}(e) = (L_i, L_j)$, thus introducing an intermediary multiplication in $\mathcal{F}_{\mathrm{sec}}$. This completes the proof. □

Remark 4.28. Up until this point we have dealt exclusively with a fixed finite collection $\mathbf{L} \subset \mathrm{ob}\,\mathcal{F}_{\mathrm{sec}}$ of possible Lagrangian section boundary labels, and hence have constructed A_∞-operations of various kinds within the corresponding universe of such. All of our constructions admit a natural extension to arbitrary countable families of sections $\mathbf{L}$. We explain two possible approaches which encode the same idea via an algebraic or a geometric strategy.

From the algebraic perspective, on one hand we have the Morse–Fukaya algebra $\mathcal{A}$ as defined in Chapter 3 (independent of any choice of finite family $\mathbf{L}$). On the other hand, after fixing such a family $\mathbf{L}$ in Chapter 4 we essentially constructed an algebra $\mathcal{A}_{\mathbf{L}}$ with defining choices conveniently compatible with $\mathbf{L}$ and subsequently produced $\mathcal{A}_{\mathbf{L}}$-modules. Allowing the finite family $\mathbf{L}$ to vary, for each finite $\mathbf{L}'$ containing $\mathbf{L}$ we may with little effort inductively construct $\mathcal{A}_{\mathbf{L}'}$ making sure that our defining choices yield that there is a refinement morphism $r_{\mathbf{L},\mathbf{L}'} : \mathcal{A}_{\mathbf{L}} \to \mathcal{A}_{\mathbf{L}'}$ as in Section 3.4. Each such map induces a restriction functor $r^*_{\mathbf{L},\mathbf{L}'} : \text{mod-}\mathcal{A}_{\mathbf{L}'} \to \text{mod-}\mathcal{A}_{\mathbf{L}}$. Functor C on any countable family may therefore be computed as a colimit over the poset of its countable subsets.

The second possible approach geometrically implements the changing of divisors in the moduli spaces we consider all at once, at the price of an equivalence relation identifying pseudoholomorphic treed polygons up to the presence of marked points by inactive divisors. By measuring the total distance (according to the lengths of Morse trajectories) of each disk component from the closest copy of each Lagrangian section boundary label L (in every other component in the tree), the stabilizing divisor activity weights may be modulated so that components which are infinitely far away do not consider one another's divisors. Essentially, the refinement morphisms $r_{\mathbf{L},\mathbf{L}'}$ have been integrated into the moduli spaces themselves, and we swap out stabilizing data as we drift further along Morse trajectories.

Bibliography

[Abo14] Mohammed Abouzaid. "Family Floer cohomology and mirror symmetry". *Proceedings of the International Congress of Mathematicians*. Vol. II. Seoul, Korea, Aug. 2014, pp. 813–836.

[Abo17] Mohammed Abouzaid. "The family Floer functor is faithful". *Journal of the European Mathematical Society* 19.7 (2017), pp. 2139–2217.

[AGM01] Denis Auroux, Damien Gayet, and Jean-Paul Mohsen. "Symplectic hypersurfaces in the complement of an isotropic submanifold". *Mathematische Annalen* 321 (2001), pp. 739–754.

[Aur23] Denis Auroux. "Holomorphic discs of negative Maslov index and extended deformations in mirror symmetry" (2023). arXiv: 2309.13010.

[BGR84] Siegfried Bosch, Ulrich Güntzer, and Reinhold Remmert. *Non-archimedean analysis*. Vol. 261. Springer Berlin, 1984.

[CL06] Octav Cornea and François Lalonde. "Cluster homology: an overview of the construction and results". *Electronic Research Announcements of the American Mathematical Society* 12.1 (2006), pp. 1–12.

[CM07] Kai Cieliebak and Klaus Mohnke. "Symplectic hypersurfaces and transversality in Gromov-Witten theory". *Journal of Symplectic Geometry* 5.3 (2007). ISSN: 15275256.

[CW22] François Charest and Chris Woodward. *Floer cohomology and flips*. Vol. 279. 1372. American Mathematical Society, 2022.

[Don96] S. K. Donaldson. "Symplectic submanifolds and almost-complex geometry". *Journal of Differential Geometry* 44.4 (1996), pp. 666–705. DOI: 10.4310/jdg/1214459407.

[Flo88] Andreas Floer. "The unregularized gradient flow of the symplectic action". *Communications on Pure and Applied Mathematics* 41.6 (1988), pp. 775–813.

[Fuk+10a] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian Intersection Floer Theory: Anomaly and Obstruction, Part I*. Vol. 1. American Mathematical Soc., 2010.

[Fuk+10b] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian Intersection Floer Theory: Anomaly and Obstruction, Part II*. Vol. 2. American Mathematical Soc., 2010. Chap. 8.

[Fuk02] Kenji Fukaya. "Floer homology for families—a progress report". *Integrable Systems, Topology, and Physics*. Contemporary Mathematics. Conference held at the University of Tokyo, Japan, July 17–21, 2000. Providence, RI: American Mathematical Society, 2002.

[Fuk05] Kenji Fukaya. "Multivalued Morse theory, asymptotic analysis and mirror symmetry". *Graphs and patterns in mathematics and theoretical physics, Proceedings of the Conference dedicated to Dennis Sullivan on his 60th birthday, Stony Brook Univ., 2001*. Amer. Math. Soc. 2005, pp. 205–278.

[Fuk10] Kenji Fukaya. "Cyclic symmetry and adic convergence in Lagrangian Floer theory". *Kyoto Journal of Mathematics* 50.3 (Jan. 2010), pp. 521–590. DOI: 10.1215/0023608X-2010-004.

[Fuk93] Kenji Fukaya. "Morse homotopy, A_∞-category, and Floer homologies". *Proc. of the GARC Workshop on Geometry and Topology'93, Seoul, 1993*. Seoul Nat. Univ. 1993, pp. 1–102.

[Gan12] Sheel Ganatra. "Symplectic cohomology and duality for the wrapped Fukaya category". PhD thesis. Massachusetts Institute of Technology, 2012.

[GG15] Lothar Gerritzen and Hans Grauert. "Die Azyklizitat der affinoiden tlberdeckungen". *Global Analysis: Papers in Honor of K. Kodaira (PMS-29)* (2015), pp. 159–184.

[Gro85] Mikhael Gromov. "Pseudo holomorphic curves in symplectic manifolds". *Inventiones mathematicae* 82.2 (1985), pp. 307–347.

[Gue99] Vincent Guedj. "Approximation of currents on complex manifolds". *Mathematische Annalen* 313 (1999), pp. 437–474.

[Kon95a] Maxim Kontsevich. "Enumeration of rational curves via torus actions". *The moduli space of curves*. Springer, 1995, pp. 335–368.

[Kon95b] Maxim Kontsevich. "Homological algebra of mirror symmetry". *Proceedings of the International Congress of Mathematicians: August 3–11, 1994, Zürich, Switzerland*. 1995, pp. 120–139.

[KS01] Maxim Kontsevich and Yan Soibelman. "Homological mirror symmetry and torus fibrations". *Symplectic geometry and mirror symmetry (Seoul, 2000)*. World Sci. Publ., River Edge, NJ, 2001, pp. 203–263.

[Maz22] Thibaut Mazuir. "Morse theory and higher algebra of A-infinity algebras". PhD thesis. Sorbonne Université, 2022.

[MS12] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and symplectic topology*. 2nd ed. Vol. 52. American Mathematical Society, 2012.

[MWW18] Sikimeti Ma'u, Katrin Wehrheim, and Chris Woodward. "A_∞ functors for Lagrangian correspondences". *Selecta Mathematica* 24.3 (2018), pp. 1913–2002.

[Sch16] Felix Schmäschke. "Floer homology of Lagrangians in clean intersection" (2016). arXiv: 1606.05327.

[Sei08a] Paul Seidel. "A_∞-subalgebras and natural transformations". *Homology, Homotopy and Applications* 10.2 (2008), pp. 83–114.

[Sei08b] Paul Seidel. *Fukaya Categories and Picard–Lefschetz Theory*. Zurich Lectures in Advanced Mathematics. Zürich: European Mathematical Society, 2008. ISBN: 978-3-03719-063-0.

[Sei08c] Paul Seidel. "Homological mirror symmetry for the genus two curve" (2008). arXiv: 0812.1171.

[Sma65] S Smale. "An Infinite Dimensional Version of Sard's Theorem". *American Journal of Mathematics* 87.4 (1965), pp. 861–866.

[SYZ96] Andrew Strominger, Shing-Tung Yau, and Eric Zaslow. "Mirror symmetry is T-duality". *Nuclear Physics B* 479.1-2 (1996), pp. 243–259.

[Tat71] John Tate. "Rigid analytic spaces". *Inventiones Mathematicae* 12.4 (1971), pp. 257–289.

[Tu14] Junwu Tu. "On the reconstruction problem in mirror symmetry". *Advances in Mathematics* 256 (2014), pp. 449–478.

[VWX20] Sushmita Venugopalan, Chris T Woodward, and Guangbo Xu. "Fukaya categories of blowups" (2020). arXiv: 2006.12264.

[WW15] Katrin Wehrheim and Chris Woodward. "Orientations for pseudoholomorphic quilts" (2015). arXiv: 1503.07803.

[Yua20] Hang Yuan. "Family Floer program and non-archimedean SYZ mirror construction" (2020). arXiv: 2003.06106.