Chapter 3

Triangulated persistence Fukaya categories

In this chapter we apply the theory developed in Chapter 2 to the case of Fukaya categories. The setup described before applies naturally to this context: under (significant) constraints, the derived Fukaya category naturally admits a TPC refinement, and this setting is ideal for approaching a variety of quantitative questions typical of symplectic topology.

We begin in Section 3.1 with the statements of the main symplectic applications in the memoir: Theorems 3.1 and 3.4 and Corollary 3.7. To prove these statements we first fix in Section 3.2 the basics of filtered A_∞-categories and associated TPCs, and we then discuss basic notions relative to filtered Floer theory. We describe how to proceed from Floer chain complexes to the Fukaya category. However, for technical reasons the construction leads only to a *weakly* filtered A_∞-category. In Section 3.3 we show that under certain restrictive conditions this construction can be adjusted to obtain a genuinely filtered A_∞-category. The main technical result of Chapter 3 appears in Theorem 3.12. The model for the Fukaya category that we construct in this case is based on clusters of punctured disks. While similar models have appeared before in the literature, we include enough details to justify the control of filtrations. In Section 3.4 we prove the statements from Section 3.1. In particular, we construct the metrics on the spaces of Lagrangians that were announced in the introduction of the memoir. The TPC formalism was inspired by earlier work on Lagrangian cobordism and it is useful to see how weighted triangles and operations with them appear geometrically in the cobordism setting. This is discussed in Section 3.5 together with some other geometric illustrations of some of the statements in Section 3.1.

3.1 Main symplectic topology applications

Let $(X, \omega = d\lambda)$ be a Liouville manifold, i.e., an exact symplectic manifold with a prescribed primitive λ of the symplectic structure ω and such that X is symplectically convex at infinity with respect to these structures. We will work here with pairs $L = (\overline{L}, h_L)$ consisting of a closed oriented exact Lagrangian submanifold $\overline{L} \subset X$ equipped with a function $h_L : \overline{L} \to \mathbb{R}$ that is a primitive of $\lambda|_{\overline{L}}$, i.e., $dh_L = \lambda|_{\overline{L}}$. We will refer to such a pair L as a marked Lagrangian submanifold and to $\overline{L}$ as its underlying Lagrangian.

Fix a collection of marked Lagrangians $\mathcal{X}$ in X. We assume that $\mathcal{X}$ is closed under shifts of the primitives, that is, if $L = (\overline{L}, h_L)$ is in $\mathcal{X}$, then for every $r \in \mathbb{R}$ and $k \in \mathbb{Z}$, the marked Lagrangian $\Sigma^r L := (\overline{L}, h_L + r)$ is also in $\mathcal{X}$. We will also assume that our marked Lagrangians are graded, in a sense recalled in Section 3.2.2.2. If we need to make the grading explicit, we write $L = (\bar{L}, h_L, \theta_L)$, and we assume the family $\mathcal{X}$ is also closed under translations $L[k] = (\bar{L}, h_L, \theta_L - k)$ of the grading.

Denote by $\bar{\mathcal{X}} = \{\bar{L} \mid L \in \mathcal{X}\}$ the collection of underlying Lagrangian submanifolds corresponding to the marked Lagrangians in $\mathcal{X}$. We will assume that the family $\bar{\mathcal{X}}$ is finite and that its elements are in general position, in the sense that any two distinct Lagrangians $L', L'' \in \bar{\mathcal{X}}$ intersect transversely, and for every three distinct Lagrangians $\bar{L}_0, \bar{L}_1, \bar{L}_2 \in \bar{\mathcal{X}}$ we have $\bar{L}_0 \cap \bar{L}_1 \cap \bar{L}_2 = \emptyset$.

As earlier in the memoir, algebraic considerations can be carried out over an arbitrary field $\mathbf{k}$. However, without additional assumptions on our Lagrangians, Floer theory works only over $\mathbf{k} = \mathbb{Z}_2$. We will therefore assume $\mathbf{k} = \mathbb{Z}_2$, but continue to denote the base field by $\mathbf{k}$, to indicate that under additional assumptions, our theory is expected to work over an arbitrary field $\mathbf{k}$. The marked Lagrangians in X are the objects of an A_∞-category, the Fukaya category $\mathcal{F}uk(X)$ of X, constructed as in Seidel's book [53]. The associated derived Fukaya category is denoted by $D\mathcal{F}uk(X)$. Its objects are the A_∞-modules over $\mathcal{F}uk(X)$ that belong to the triangulated completion of the Yoneda A_∞-modules, $\mathcal{Y}(L)$, where L is a marked Lagrangian. We denote by $\mathcal{F}uk(\mathcal{X})$ the A_∞-subcategory of $\mathcal{F}uk(X)$ with objects the Lagrangians in $\mathcal{X}$, and by $j_{\mathcal{X}} : \mathcal{F}uk(\mathcal{X}) \to \mathcal{F}uk(X)$ the inclusion. There are two Yoneda-type modules associated to the elements of $\mathcal{X}$: over the category $\mathcal{F}uk(X)$ and over the smaller category $\mathcal{F}uk(\mathcal{X})$. The two are related by applying the pullback $j^*_{\mathcal{X}}(-)$ and thus will be generally denoted by the same symbol.

We denote by $D\mathcal{F}uk(\mathcal{X})$ the associated derived category, consisting this time of modules over $\mathcal{F}uk(\mathcal{X})$ that belong to the triangulated completion of the Yoneda modules of the elements of $\mathcal{X}$. We emphasize that, with the terminology used in this memoir, a family $\mathcal{Z}$ of objects in a triangulated category $\mathcal{C}$ is a system of generators of $\mathcal{C}$ if the triangulated envelope of $\mathcal{Z}$ in $\mathcal{C}$ equals $\mathcal{C}$. In particular, the Yoneda modules of the elements of $\mathcal{X}$ form a system of generators of $D\mathcal{F}uk(\mathcal{X})$.

The following consequence of Theorem 3.12 is significant enough to be stated separately:

Theorem 3.1. *There exists a triangulated persistence category $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, independent up to TPC equivalence of the data used in its construction, such that:*

(i) *For any $L, L' \in \mathcal{X}$, there is a canonical isomorphism*

$$\hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty}(L, L') \cong \mathrm{HF}^0(L, L'),$$

where $\mathrm{HF}^0(-,-)$ is Floer homology in cohomological degree 0.

(ii) *$\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$ is triangulated equivalent to $D\mathcal{F}uk(\mathcal{X})$.*

(iii) *If the family $\mathcal{X}$ generates $D\mathcal{F}uk(X)$, then for each marked Lagrangian N that intersects transversely the family $\mathcal{X}$, the pullback $j^*_{\mathcal{X}}\mathcal{Y}(N)$ of $\mathcal{Y}(N)$*

– the Yoneda module of N over $\mathcal{F}uk(X)$ (defined with a convenient choice of perturbation data) – is quasi-isomorphic to an object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$.

We call $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ the *triangulated persistence Fukaya category* associated to $\mathcal{X}$.

Here we emphasize that the construction of $\mathcal{F}uk(\mathcal{X})$ always depends on some perturbation data $\mathcal{P}$; more precisely, $\mathcal{F}uk(\mathcal{X})$ is denoted by $\mathcal{F}uk(\mathcal{X};\mathcal{P})$. Theorem 3.12 guarantees that there always exists perturbation data $\mathcal{P}$ such that $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is a strict unital A_∞-category, together with filtered A_∞-functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}uk(\mathcal{X};\mathcal{P}_0) \to \mathcal{F}uk(\mathcal{X};\mathcal{P}_1)$ when changing the perturbation data from $\mathcal{P}_0$ to $\mathcal{P}_1$. An essential part of the proof of Theorem 3.1 (see Section 3.4.1.1) shows that the resulting triangulated persistence Fukaya category $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$, constructed from $\mathcal{F}uk(\mathcal{X};\mathcal{P})$, is in fact independent of the perturbation data up to a TPC equivalence. This justifies the notation $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ above without any reference to $\mathcal{P}$.

Point (ii) of Theorem 3.1 implies that if $\mathcal{X}$ generates $D\mathcal{F}uk(X)$, then $\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty$ is equivalent to $D\mathcal{F}uk(X)$. Point (iii) of Theorem 3.1 gives a bit more information and shows that one can use measurements in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ to study Lagrangians that do not necessarily belong to the finite family $\bar{\mathcal{X}}$. Nonetheless, it remains that the requirement that the family $\bar{\mathcal{X}}$ be finite is highly constraining. It is expected that this requirement can be dropped by using a more involved construction in place of the one used in the proof of Theorem 3.12.

Remark 3.2. As stated, Theorem 3.1 identifies $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ up to TPC equivalence (see Definition 2.25) but, while this equivalence is expected to be canonical, our methods do not quite give that. Still, the equivalences that appear here are not completely arbitrary. For example, their mapping on objects leaves the elements of $\mathcal{X}$ fixed. See Theorem 3.12 for more details.

The next result in this section will be formulated in terms of this TPC, $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, and will involve a notion of relative Gromov width that first appeared in [4] (see also [10]). Assume that L and L' are two Lagrangians, both possibly immersed. We define

$$\begin{aligned}\delta(L;L') = \sup\big\{\pi r^2 \mid \exists\, e : (B(r),\omega_0) \to (X,\omega) \text{ symplectic embedding,}\\ \text{such that } e^{-1}(L) = \mathbb{R}B(r),\ e(B(r)) \cap L' = \emptyset\big\},\end{aligned} \tag{3.1}$$

where $(B(r),\omega_0)$ is the standard closed ball of radius r in $(\mathbb{R}^{2n},\omega_0)$, and $\mathbb{R}B(r) = (\mathbb{R}^n \times \{0\}) \cap B(r)$ is its real part. A related measurement reflects the "quality" of the intersection points between L and L', relative to another subset. Assume that

L and L' intersect transversely and let $A \subset X$ be a subset. We define

$$\begin{aligned}\delta^{\cap}(L, L'; A) = \sup\big\{\pi r^2 \mid \forall x \in L \cap L',\ \exists\, e_x : (B(r), \omega_0) \to (X, \omega) \text{ symp. emb.}\\ \text{s.t. } e_x(0) = x,\ e_x^{-1}(L) = \mathbb{R}B(r),\ e_x^{-1}(L') = i\mathbb{R}B(r),\\ e_{x'}(B(r)) \cap e_{x''}(B(r)) = \emptyset \text{ whenever } x' \neq x'',\\ \text{and moreover } \forall x \in L \cap L',\ e_x(B(r)) \cap A = \emptyset\big\}.\end{aligned} \tag{3.2}$$

Here $i\mathbb{R}B(r) := (\{0\} \times \mathbb{R}^n) \cap B(r)$ is the "imaginary" part of the ball $B(r)$.

We will also need the spectral distance between two marked Lagrangians L and L'. We assume that L is Hamiltonian isotopic to L'. In this case, the Floer homology $\mathrm{HF}(L, L')$ is isomorphic to the singular homology $H_*(L; \mathbf{k})$ of L and there is a canonical class $o_{L,L'} \in \mathrm{HF}(L, L')$ corresponding to the fundamental class in $H_*(L; \mathbf{k})$. There is also a second class $pt_{L,L'} \in \mathrm{HF}(L, L')$ that corresponds to the point class in $H_*(L; \mathbf{k})$. Assume further that $L, L' \in \mathcal{X}$. In this case, given point (i) of Theorem 3.1, we have

$$\mathrm{HF}^0(L, L') = \big\{[f] \in \hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty}(L, L') \mid f \in \hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})}(L, L')\big\}.$$

Therefore these classes in $\mathrm{HF}(L, L')$ have spectral numbers $\sigma(-)$ as defined in (2.36). We define

$$\sigma(L, L') = \sigma(o_{L,L'}) - \sigma(pt_{L,L'}).$$

We extend the definition of σ to the case when L' is not Hamiltonian isotopic to L by setting $\sigma(L, L') = \infty$ in this case.

Remark 3.3. It is easily seen that this definition coincides with previous versions of spectral invariants introduced by Viterbo, Schwarz, and Oh, and later adjusted to the Lagrangian setting.

Pick a family $\mathcal{F} \subset \mathcal{X}$ that is invariant with respect to shift and translation. Fix an admissible perturbation data $\mathcal{P}$ and an associated triangulated persistence category $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})$. Consider the shift-invariant, persistence, fragmentation pseudometric $\widehat{d^{\mathcal{F}}}(-,-)$ associated to the persistence weight $\bar{w}$ on $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})_\infty$, as described in Section 2.4.3.1 and (2.41). Each such pseudometric is defined on the objects of $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, which contain the Yoneda modules of the Lagrangians in $\mathcal{X}$ but also additional A_∞-modules.

We now define a pseudometric on $\mathcal{X}$ by

$$D^{\mathcal{F}}(L, L') = \widehat{d^{\mathcal{F}}}(\mathcal{Y}(L), \mathcal{Y}(L')), \quad L, L' \in \mathcal{X}.$$

In the case $\mathcal{F} = \{0\}$, we write $D(-,-) = D^{\{0\}}(-,-)$. This is an upper bound for all the other fragmentation metrics $D^{\mathcal{F}}$. There is a slight abuse in notation here because the definition of $D^{\mathcal{F}}$ depends implicitly on the perturbation data $\mathcal{P}$, but this will be resolved in the next result.

Theorem 3.4. *Let $\mathcal{F} \subset \mathcal{X}$. In the setting above, the pseudometrics $D^{\mathcal{F}}$ are independent of the perturbation data $\mathcal{P}$ used for their definitions. Moreover:*

(i) *(Spectrality) Assume that the Lagrangians in $\mathcal{X}$ are graded; then for any $L, L' \in \mathcal{X}$ we have*
$$D(L, L') \leq 4\,\sigma(L, L').$$

(ii) *(Non-degeneracy) For all $L, L' \in \mathcal{X}$,*
$$\frac{\delta(L; L' \cup \bigcup_{F \in \mathcal{F}} F)}{8} \leq D^{\mathcal{F}}(L, L').$$

(iii) *(Persistence of intersections) Assume that $L, L', N \in \mathcal{X}$, $L' \notin \mathcal{F}$. If*
$$D^{\mathcal{F}}(L, L') < \frac{1}{16}\,\delta^{\cap}\big(N, L'; \bigcup_{F \in \mathcal{F}} F\big),$$
then
$$\#(L \cap N) \geq \#(L' \cap N).$$

(iv) *(Finiteness) If the family $\mathcal{F}$ generates $D\mathcal{F}uk(X)$, then the pseudometric $D^{\mathcal{F}}$ is finite.*

Compared to other metrics and measurements on spaces of Lagrangians, the key novelty here is that properties (i), (ii), (iii), and (iv) are valid for the same metric.

Remark 3.5. (a) Point (i) of Theorem 3.4 shows that all the fragmentation pseudometrics $D^{\mathcal{F}}$ are dominated by the spectral metric. In previous results involving metrics on spaces of Lagrangians, such as those based on the shadows of cobordisms in [10], the best one could do was to establish upper bounds on the metrics that are generally much harder to estimate, such as the Hofer distance. Further consequences of this point will be discussed in Section 3.4.1.

(b) Point (ii) of Theorem 3.4 can be read as a typical non-squeezing-type result: embeddings of large symplectic balls, as in the definition of δ, are obstructed by $D^{\mathcal{F}}$. Conversely, this point implies that if $D^{\mathcal{F}}(L, L') = 0$, then $L \subset L' \cup \bigcup_{F \in \mathcal{F}} F$. As a result, suppose that we fix a second family $\mathcal{F}' \subset \mathcal{X}$, obtained through a small Hamiltonian perturbation of the elements of $\mathcal{F}$. One can then consider $D^{\mathcal{F},\mathcal{F}'} = \max\{D^{\mathcal{F}}, D^{\mathcal{F}'}\}$ as in (2.42). This pseudometric is non-degenerate on $\overline{\mathcal{X}}$ in the sense that $D^{\mathcal{F},\mathcal{F}'}(L, L') = 0$ if and only if $\overline{L} = \overline{L}'$ (in other words, the two underlying Lagrangians involved coincide; obviously, the markings may differ). This type of argument first appeared in [10]. Various forms of the inequality in point (ii) appeared earlier in the literature, in particular in cases such as when $\mathcal{F} = \{0\}$, where the metric involved is the Hofer metric (see, e.g., [4]). However, it is useful to note that the pseudometrics $D^{\mathcal{F}}$ are, in general, smaller compared to the metrics in these earlier references. Note also that, even for $\mathcal{F} = \{0\}$, the inequalities obtained by combining (i) and (ii) appear to be new.

(c) We emphasize that in point (iii) of Theorem 3.4 the two Lagrangians L and L' are allowed to be very different. For instance, they can be in different smooth isotopy classes or even have different homeomorphism types, and still $D^{\mathcal{F}}(L, L')$ can be finite (this point is reinforced by the last part of the theorem). Therefore, this result shows a form of rigidity of Lagrangian intersections for perturbations that are small in this metric $D^{\mathcal{F}}$, but that can be very large (even infinite) in other metrics. The result extends earlier persistence-type statements in Morse and Floer theory (one of the earliest examples appearing in [23]), most of them expressed in terms of the Hofer distance, which is much larger than $D^{\mathcal{F}}$. Again, there is considerable interest in working with the algebraic metrics introduced here because for other metrics, such as the shadow metrics based on Lagrangian cobordism, the finiteness result in point (iv) is not known to hold.

(d) The fact that we have $j^*_{\mathcal{X}}\mathcal{Y}(N) \in \mathrm{Obj}(\mathcal{C}\mathcal{F}uk(\mathcal{X})_\infty)$ for each marked Lagrangian N that intersects transversely the elements of $\mathcal{X}$, as in Theorem 3.1, implies that we can define a pseudometric on the space of all marked Lagrangians in X by

$$\Delta^{\mathcal{F}}(N, N') = \limsup_{\epsilon \to 0} \widehat{\bar{d}^{\mathcal{F}}}\big(j^*_{\mathcal{X}}\mathcal{Y}(N_\epsilon), j^*_{\mathcal{X}}\mathcal{Y}(N'_\epsilon)\big),$$

where N_ϵ and N'_ϵ are ϵ-small (in the Hofer metric) Hamiltonian perturbations of N and N', respectively, both transverse to the elements of $\mathcal{X}$. This pseudometric is in general degenerate, as it does not "see" differences between N and N' away from the elements of $\mathcal{X}$.

(e) The constants providing the various bounds in Theorem 3.4 are very rough and can be improved in some cases, but we will not pursue this question here.

We will prove a consequence of Theorem 3.4, deduced by studying how the pseudometrics $D^{\mathcal{F}}$ change when the underlying set of marked Lagrangians $\mathcal{X}$ changes.

To state this consequence we need a global finiteness-type assumption on our Liouville manifold (X, ω). To formulate it, we denote by $\mathcal{L}ag(X)$ the set of exact, compact, graded, embedded Lagrangians in X, and by $\mathcal{L}ag(X)'$ the marked, exact Lagrangians in X (these are the elements of $\mathcal{L}ag(X)$ but with fixed primitives and grading choices). As before, the Fukaya category $D\mathcal{F}uk(X)$ is the derived category of the A_∞-category with objects the elements in $\mathcal{L}ag(X)'$. The category $D\mathcal{F}uk(X)$ is constructed as in [53]. In particular, the perturbation data depends only on the elements in $\mathcal{L}ag(X)$, and not on the choices of primitives and grading.

Definition 3.6. Let (X, ω) be a Liouville manifold. The *Fukaya rank* of (X, ω), $\mathrm{rank}\,\mathcal{F}uk(X, \omega)$, is the minimal cardinality of a family of Lagrangians $\bar{\mathcal{F}} \subset \mathcal{L}ag(X)$ such that the corresponding family of marked Lagrangians $\mathcal{F} \subset \mathcal{L}ag(X)'$, obtained from $\bar{\mathcal{F}}$ by adding all possible translates of the objects in $\mathcal{F}$ (in terms of grading), generates $D\mathcal{F}uk(X)$. (Note that the primitives have no effect here, since we are talking about generators in a non-filtered setting.)

We emphasize here that "generating" has the meaning of triangulated generation, as everywhere else in this memoir. The rank $\mathcal{D}$ can be defined similarly for any triangulated category $\mathcal{D}$. The terminology is justified by the fact that this quantity is an upper bound for the rank of the Grothendieck group $K(\mathcal{D})$.

Corollary 3.7. *Let (X, ω) be a Liouville manifold. Assume that* rank $\mathcal{F}uk(X, \omega) < \infty$. *Fix a family of generators $\mathcal{F} \subset \mathcal{L}ag(X)'$, invariant under shifts and translations, and such that the corresponding family $\bar{\mathcal{F}} \subset \mathcal{L}ag(X)$ obtained by forgetting the markings is finite and is in general position (in the sense defined at the beginning of Section 3.1). Then the set $\mathcal{L}ag(X)$ carries a finite pseudometric $\mathcal{D}^{\mathcal{F}}$ such that:*

(i) *(Spectrality) For any $L, L' \in \mathcal{L}ag(X)$, we have*

$$\mathcal{D}^{\mathcal{F}}(L, L') \leq 4\,\sigma(L, L').$$

(ii) *(Non-degeneracy) If $L, L' \in \mathcal{L}ag(X)$, then*

$$\frac{\delta(L; L' \cup \bigcup_{F \in \mathcal{F}} F)}{8} \leq \mathcal{D}^{\mathcal{F}}(L, L').$$

(iii) *(Persistence of intersections) Assume that $L, L', N \in \mathcal{L}ag(X)$ are in general position and $L' \notin \mathcal{F}$. If*

$$\mathcal{D}^{\mathcal{F}}(L, L') < \frac{1}{16}\,\delta^{\cap}\big(N, L'; \textstyle\bigcup_{F \in \mathcal{F}} F\big),$$

then

$$\#(L \cap N) \geq \#(L' \cap N).$$

In particular, if $\mathcal{F}'$ is another family obtained from $\mathcal{F}$ by generic Hamiltonian perturbations of the elements of $\bar{\mathcal{F}}$, then

$$\mathcal{D}^{\mathcal{F}, \mathcal{F}'} = \max\{\mathcal{D}^{\mathcal{F}}, \mathcal{D}^{\mathcal{F}'}\}$$

is a finite and non-degenerate metric on $\mathcal{L}ag(X)$ that satisfies the properties (i), (ii), and (iii) above.

Thus, under the hypothesis rank $\mathcal{F}uk(X, \omega) < \infty$, *the set $\mathcal{L}ag(M)$ has a metric space structure with respect to a metric satisfying the properties (i), (ii), and (iii).*

3.2 Filtrations in Floer homology and Fukaya categories

3.2.1 Filtered A_∞-categories and their associated TPCs

A filtered A_∞-category $\mathcal{A}$ is an A_∞-category over a given base field $\mathbf{k}$, such that the spaces of morphisms $\hom_{\mathcal{A}}(X, Y)$ between every two objects X, Y are filtered (with increasing filtrations) and *all* the composition maps μ_d, $d \geq 1$, respect the filtrations. We endow $\hom_{\mathcal{A}}(X, Y)$ with the differential μ_1 and view it as a filtered chain complex. We denote by $\hom^s_{\mathcal{A}}(X, Y)$, $s \in \mathbb{R}$, the level-s filtration subcomplex of $\hom_{\mathcal{A}}(X, Y)$. We refer the reader to [10] for more details on filtered A_∞-categories. Note, however, that in [10] the theory is developed for the more general case of *weakly* filtered A_∞-categories (the "genuinely" filtered case is obtained from the weakly filtered one by assuming the so-called "discrepancies" of $\mathcal{A}$, defined in [10], to vanish). Of course, the A_∞ considerations here are very similar to those for dg-categories in Section 2.5.1.

For simplicity, we make three further assumptions on our filtered A_∞-categories. The first is that $\mathcal{A}$ is strictly unital, with the units lying in persistence level 0. The second one is that for every two objects $X, Y \in \mathrm{Obj}(\mathcal{A})$, the space $\hom_{\mathcal{A}}(X, Y)$ is finite-dimensional over $\mathbf{k}$. The third assumption is that $\mathcal{A}$ is complete with respect to persistence shifts, in the sense that we have a shift "functor" consisting of a family of A_∞-functors $\Sigma = \{\Sigma^r : \mathcal{A} \to \mathcal{A}, r \in \mathbb{R}\}$ whose members satisfy the following conditions:

(1) Σ^r is strictly unital and the higher components $(\Sigma^r)_d$, $d \geq 2$, of Σ^r all vanish.

(2) $\Sigma^0 = \mathbb{1}$, $\Sigma^s \circ \Sigma^t = \Sigma^{s+t}$.

(3) We are given prescribed identifications $\hom^s_{\mathcal{A}}(\Sigma^r X, Y) \cong \hom^{s+r}_{\mathcal{A}}(X, Y)$ that are compatible with the inclusions $\hom^\alpha_{\mathcal{A}}(X, Y) \subset \hom^\beta_{\mathcal{A}}(X, Y)$ for $\alpha \leq \beta$. These identifications are considered as part of the structure of the shift functor Σ.

The assumption that $\mathcal{A}$ is complete with respect to shifts is merely a matter of convenience, in the sense that it is not essential to impose this condition in advance. Indeed, any filtered A_∞-category (satisfying all the above assumptions except the one on completeness with respect to shifts) can be completed with respect to shifts by adding suitable objects that will play the role of the shifted $\Sigma^r X$ objects and then defining the functors Σ^r accordingly. See again Section 2.5.1 for the similar case of filtered dg-categories.

We will generally use homological conventions in the context of A_∞-categories; however, for compatibility with the literature we will generally use cohomological grading. Whenever this is the case, we will denote the cohomological degrees by superscripts (e.g., H^0 will stand for the homology in cohomological degree 0, the units will be assumed to be in cohomological degree 0, and so on).

Given a filtered A_∞-category $\mathcal{A}$, one can form the category $\mathrm{Tw}(\mathcal{A})$ of twisted complexes over $\mathcal{A}$, which is itself a filtered A_∞-category (satisfying all the additional assumptions mentioned earlier). This can be done by following the construction in [53, Chapter I, Section (31)] and extending the filtrations from $\mathcal{A}$ to $\mathrm{Tw}(\mathcal{A})$ in the obvious way. The construction of the filtered $\mathrm{Tw}(\mathcal{A})$ in the case of dg-categories has been worked out in detail in Section 2.5.1, and the A_∞ case is very similar. There is a bit of abuse of notation in writing $\mathrm{Tw}(\mathcal{A})$, since the latter category carries additional structures (namely filtrations and a shift functor) beyond those of the unfiltered category of twisted complexes, which is denoted in the literature by the same notation $\mathrm{Tw}(\mathcal{A})$.

The filtered A_∞-category $\mathcal{A}$ embeds into $\mathrm{Tw}(\mathcal{A})$ in an obvious way, the embedding being a filtered A_∞-functor which is full and faithful (on the chain level). Moreover, $\mathrm{Tw}(\mathcal{A})$ is pre-triangulated in the filtered sense (which in particular means that it is closed under formation of filtered mapping cones). It follows that the homological category $H^0(\mathrm{Tw}(\mathcal{A}))$ is a TPC that contains the homological persistence category $H^0(\mathcal{A})$ of $\mathcal{A}$.

Another TPC associated with the A_∞-category $\mathcal{A}$ is provided by the category $\mathrm{Fmod}(\mathcal{A})$ of filtered A_∞-modules over $\mathcal{A}$. Weakly filtered modules are defined in [10] and the filtered definitions correspond to all discrepancies being 0. We will only consider strictly unital modules here. There is a natural shift functor on this category, $\Sigma : (\mathbb{R}, +) \to \mathrm{End}(\mathrm{Fmod}(\mathcal{A}))$. Given $r \in \mathbb{R}$ and a module $\mathcal{M} \in \mathrm{Fmod}(\mathcal{A})$, we define the filtered module $\Sigma^r \mathcal{M}$ by $(\Sigma^r \mathcal{M})^{\leq \alpha}(N) = \mathcal{M}^{\alpha - r}(N)$, endowed with the same μ_d-operations as $\mathcal{M}$.

Remark 3.8. In the cases of interest to us, namely the Fukaya category, this shift functor on $\mathrm{Fmod}(\mathcal{A})$ is compatible with a shift functor on $\mathcal{A}$.

The category $\mathrm{Fmod}(\mathcal{A})$ is in fact a filtered dg-category in the sense of Section 2.5.1 and it is pre-triangulated. Thus $H^0(\mathrm{Fmod}(\mathcal{A}))$ is a TPC, by Corollary 2.102. Of more interest to us is a subcategory of $\mathrm{Fmod}(\mathcal{A})$. First notice that, because $\mathcal{A}$ is filtered, the Yoneda functor $\mathcal{Y} : \mathcal{A} \to \mathrm{Fmod}(\mathcal{A})$ is filtered too. Moreover, our assumption of strict unitality of $\mathcal{A}$ implies that $\mathcal{Y}$ is homologically full and faithful. Furthermore, there exists a canonical map

$$\lambda : \mathcal{M}(X) \to \hom_{\mathrm{Fmod}}(\mathcal{Y}(X), \mathcal{M})$$

for all $X \in \mathrm{Obj}(\mathcal{A})$ and $\mathcal{M} \in \mathrm{Obj}(\mathrm{Fmod}(\mathcal{A}))$, as defined in [53, Chapter I, Section (11)]. Standard arguments show that λ is a *filtered quasi-isomorphism*, in the sense that it is filtered and induces an isomorphism between the persistence homologies of its domain and target filtered chain complexes.

We consider now the pre-triangulated closure $\mathcal{A}^\#$ of the Yoneda modules and their shifts: this is a full subcategory of $\mathrm{Fmod}(\mathcal{A})$ that has as objects all the iterated cones, over filtration-preserving morphisms, of shifts of Yoneda modules (thus of modules of the form $\Sigma^r \mathcal{Y}(X)$).

Finally, we denote by $\mathcal{A}^\nabla$ the smallest full subcategory of $\mathrm{Fmod}(\mathcal{A})$ that contains $\mathcal{A}^\#$ and all the modules (and all their shifts and translates) that are r-quasi-isomorphic to objects in $\mathcal{A}^\#$ for some $r \in [0, \infty)$. Here, a module $\mathcal{M}$ is called r-quasi-isomorphic to $\mathcal{N}$ if, in $H^0(\mathrm{Fmod}(\mathcal{A}))$, there is an r-isomorphism $\mathcal{M} \to \mathcal{N}$.

It is easy to see that $\mathcal{A}^{\nabla}$ remains pre-triangulated, carries the shift functor induced from $\mathrm{Fmod}(\mathcal{A})$, and thus $H^0(\mathcal{A}^{\nabla})$ is a TPC.

3.2.2 Persistence Floer homology

3.2.2.1 Filtered Floer complexes. Given a pair of marked Lagrangians L_0, L_1 as above and a choice of Floer data $\mathcal{D}_{L_0,L_1} = (H_{L_0,L_1}, J_{L_0,L_1})$, which consists of a (possibly time-dependent) Hamiltonian function and a choice of a compatible (time-dependent) almost complex structure, we can form the Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$. This is a $\mathbb{Z}_2$-graded chain complex (recall our Lagrangians are assumed to be oriented). It is generated by the Hamiltonian chords $x : [0,1] \to X$ of H_{L_0,L_1} with end points on the two Lagrangians, namely $x(0) \in L_0$, $x(1) \in L_1$. For simplicity, we work here with coefficients in $\mathbb{Z}_2$.

Moreover, $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is a filtered chain complex, where the filtration function is given by the action functional. More precisely, if $x \in \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is a generator (i.e., a Hamiltonian chord), its action is defined by

$$\mathcal{A}(x) := \int_0^1 H_{L_0,L_1}(t, x(t))\, dt - \int_0^1 \lambda(\dot{x}(t))\, dt + h_{L_1}(x(1)) - h_{L_0}(x(0)).$$

Remark 3.9. In case L_0 and L_1 intersect transversely and $H_{L_0,L_1} \equiv 0$, the Hamiltonian chords x that generate $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ are just the intersection points $L_0 \cap L_1$ and the action reduces to

$$\mathcal{A}(x) = h_{L_1}(x) - h_{L_0}(x) \quad \forall x \in L_0 \cap L_1. \tag{3.3}$$

Back to the general case, the homology of the filtered chain complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ gives rise to the persistence Floer homology $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$, which has the structure of a $\mathbb{Z}_2$-graded persistence module (over the field $\mathbb{Z}_2$). As a vector space, $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is independent of the auxiliary Floer data $\mathcal{D}_{L_0,L_1}$; however, as a persistence module, it does depend on that choice. More precisely, the persistence module structure of $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is independent of the choice of the almost complex structure J_{L_0,L_1} from $\mathcal{D}_{L_0,L_1}$; however, it depends strongly on the choice of the Hamiltonian H_{L_0,L_1}.

3.2.2.2 Grading. While $\mathbb{Z}_2$-grading is enough for our applications, one can obtain a $\mathbb{Z}$-graded theory if one makes additional assumptions on X and on the admissible class of Lagrangian submanifolds. The simplest such conditions are the following. Firstly, we assume that $2c_1(X) = 0$, where $c_1(X)$ stands for the first Chern class of the tangent bundle of X, viewed as a complex vector bundle by endowing X with any ω-compatible almost complex structure J. We now fix a nowhere vanishing quadratic complex n-form (where $n = \dim_{\mathbb{C}} X$), namely a nowhere vanishing section Θ of the bundle $\Omega^n(X, J)^{\otimes 2}$. The choice of Θ gives rise to a global phase map

$\det^2_\Theta : \mathcal{L}(T(X)) \to S^1$ defined on the Lagrangian Grassmannian bundle $\mathcal{L}(T(X))$ of X (see [52], [53, Chapter II, Section (11j)]). Given a Lagrangian $\overline{L} \subset X$, denote by $s_{\overline{L}} : \overline{L} \to \mathcal{L}(T(X))|_{\overline{L}}$ its Gauss map. A Lagrangian $\overline{L}$ is said to admit a grading if $\det^2_\Theta \circ s_{\overline{L}} : \overline{L} \to S^1$ can be lifted to a function $\theta_{\overline{L}} : \overline{L} \to \mathbb{R}$, and a choice of such a lift is called a grading on $\overline{L}$. In this case, by adding integral constants to $\theta_{\overline{L}}$ one obtains all possible gradings of $\overline{L}$.

Gradability of Lagrangians can be rephrased in cohomological terms. The map $\det^2_\Theta \circ s_{\overline{L}} : \overline{L} \to S^1$ gives rise to a cohomology class $\mu_{\overline{L}} \in H^1(\overline{L})$, which we call the Maslov class of $\overline{L}$. (There is a slight abuse of notation here since $\mu_{\overline{L}}$ actually depends on the homotopy class of Θ.) A Lagrangian $\overline{L}$ admits a grading if and only if $\mu_{\overline{L}} = 0$.

The relation between $\mu_{\overline{L}}$ and the more familiar Maslov index homomorphism $\mu_{X,\overline{L}} : H_2(X, \overline{L}) \to \mathbb{Z}$ is that $\mu_{X,\overline{L}}(A) = \langle \mu_{\overline{L}}, \partial_* A\rangle$ for every $A \in H_2(X, \overline{L})$, where $\partial_* : H_2(X, \overline{L}) \to H_1(\overline{L})$ is the connecting homomorphism. Note also that if the map $H_1(\overline{L}) \to H_1(X)$, induced by the inclusion $\overline{L} \subset X$, is trivial, then $\mu_{\overline{L}}$ is determined by $\mu_{X,\overline{L}}$ (hence in that case $\mu_{\overline{L}}$ is independent of the choice of Θ). This is because $\mu_{X,\overline{L}}(j(B)) = 2\langle c_1(X), B\rangle = 0$ for every $B \in H_2(B)$, where the map $j : H_2(X) \to H_2(X, \overline{L})$ is induced by the inclusion. Therefore $\mu_{X,\overline{L}}$ descends to $H_2(X, \overline{L})/j(H_2(X)) \cong H_1(\overline{L})$.

In the rest of the memoir, we optionally allow for a $\mathbb{Z}$-graded theory. Whenever this is desired, we will make the preceding assumptions on X, fix the auxiliary structure Θ, and consider only marked Lagrangians L that admit a grading (or, equivalently, $\mu_{\overline{L}} = 0$). Moreover, we extend the notion of a marked Lagrangian L to include also a choice of a grading denoted by θ_L, namely $L = (\overline{L}, h_L, \theta_L)$. However, below we will mostly suppress the choice θ_L from the notation, since it will not often be used explicitly.

Given a pair of marked Lagrangians L_0, L_1, their gradings induce an absolute $\mathbb{Z}$-grading on $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$, and therefore also on $\mathrm{HF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$. The effect of translating the grading functions on the Lagrangians is the following. Write $L[k] = (\overline{L}, h_L, \theta_L - k)$. Then, using cohomological and homological grading respectively, we have

$$\mathrm{CF}^i(L_0[k], L_1[l]; \mathcal{D}) \cong \mathrm{CF}^{i+k-l}(L_0, L_1; \mathcal{D}),$$
$$\mathrm{CF}_j(L_0[k], L_1[l]; \mathcal{D}) \cong \mathrm{CF}_{j+l-k}(L_0, L_1; \mathcal{D}).$$

3.2.3 Weakly filtered Fukaya categories

The above construction can be enhanced to an A_∞-category called the Fukaya category.

Fix a collection of marked Lagrangians $\mathcal{X}$ in X. We assume that $\mathcal{X}$ is closed under grading translations and shifts of the primitives, namely if $L = (\overline{L}, h_L, \theta_L)$ is in $\mathcal{X}$, then for every $r \in \mathbb{R}$ and $k \in \mathbb{Z}$, the marked Lagrangian $\Sigma^r L[k] := (\overline{L}, h_L + r,$

$\theta_L - k$) is also in $\mathcal{X}$. (Of course, in case $\mathcal{X}$ is not closed under shifts and translations we can easily fix this by adding to $\mathcal{X}$ all the shifts and translations of its objects.)

The Fukaya category $\mathcal{F}uk(\mathcal{X})$ associated to $\mathcal{X}$ is an A_∞-category whose objects are the elements of $\mathcal{X}$ and the complex of morphisms between a pair of objects from $\mathcal{X}$ is the Floer complex of that pair. In order to set up this A_∞-category one has to choose for every pair of objects (L_0, L_1) from $\mathcal{X}$ a regular Floer datum $\mathcal{D}_{L_0,L_1}$ and then extend this choice to a consistent choice of regular perturbation data $\mathcal{P}_{\mathcal{X}}$, which is defined for every tuple of Lagrangians $(L_0, \dots, L_d)$, $d \geq 1$, from the collection $\mathcal{X}$. (It is important that both the Floer data as well as the perturbation data associated to a tuple depend only on the underlying Lagrangians in that tuple, and not on the choice of primitives or gradings on the Lagrangians in the tuple.)

Once these choices are set, one defines

$$\hom_{\mathcal{F}uk(\mathcal{X})}(L_0, L_1) := \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1}),$$

endowed with the Floer differential μ_1. The higher-order operations μ_d for $d \geq 2$ are multilinear maps

$$\mu_d : \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1}) \otimes \cdots \otimes \mathrm{CF}(L_{d-1}, L_d; \mathcal{D}_{L_{d-1},L_d}) \to \mathrm{CF}(L_0, L_d; \mathcal{D}_{L_0,L_d})$$

of cohomological degree $2 - d$, which are defined for every tuple of Lagrangians $L_0, \dots, L_d$ from $\mathcal{X}$. They satisfy the A_∞-identities. The definition of μ_d goes by counting Floer $(d + 1)$-polygons in X with boundary conditions on the L_i's. These polygons satisfy a perturbed Cauchy–Riemann equation with perturbations prescribed by $\mathcal{P}_{\mathcal{X}}$. Note that the Fukaya category described above depends on the choice $\mathcal{P}_{\mathcal{X}}$ of the perturbation data, hence should in fact be denoted by $\mathcal{F}uk(\mathcal{X}; \mathcal{P}_{\mathcal{X}})$. However, it is well known that different choices of perturbation data lead to quasi-equivalent categories [53, Chapter II, Section 10]

Taking filtrations into account, as already mentioned in Section 3.2.2.1, the hom's of this category are filtered chain complexes. However, due to the perturbation data involved in defining the higher-order operations, the μ_d-operations for $d \geq 2$ do not preserve the action filtrations, but only do so up to an error (that depends on d). Consequently, the resulting A_∞-category is not filtered but only *weakly filtered*. Enhancing such a structure to a TPC, e.g., along the lines of the construction outlined in Section 3.2.1, seems like a non-trivial technical problem.

3.3 Genuinely filtered Fukaya categories

Here we will outline a construction that gives rise to a genuinely filtered Fukaya A_∞-category. This, however, will require very restrictive assumptions on the collection of objects $\mathcal{X}$, and some adjustments in the definition of the operations μ_d for certain tuples of Lagrangians.

Recall from the beginning of the section that we denote by $\overline{\mathcal{X}} = \{\overline{L} \mid L \in \mathcal{X}\}$ the collection of underlying Lagrangian submanifolds corresponding to the marked Lagrangians in $\mathcal{X}$. Recall also the assumptions on $\overline{\mathcal{X}}$: $\overline{\mathcal{X}}$ is finite; every two distinct Lagrangians $L', L'' \in \overline{\mathcal{X}}$ intersect transversely; for every three distinct Lagrangians $\overline{L}_0, \overline{L}_1, \overline{L}_2 \in \overline{\mathcal{X}}$ we have $\overline{L}_0 \cap \overline{L}_1 \cap \overline{L}_2 = \emptyset$. We also continue to assume, as before, that $\mathcal{X}$ is closed under shifts and translation in grading.

A more general approach, yielding genuinely filtered Fukaya categories, has been developed by Ambrosioni [1] after the first version of the current work had appeared. This approach does not impose any restrictions on the collection $\mathcal{X}$, besides the assumption that all Lagrangians in $\mathcal{X}$ are weakly exact (and possibly graded, if one wants a graded theory). In particular, no finiteness condition on $\overline{\mathcal{X}}$ is needed in that work, and no transversality assumption is made for distinct elements of this collection. On the other hand, the invariance properties of the filtered Fukaya categories from [1] are coarser than the ones provided by our approach (compare, e.g., Theorems 3.1 and 3.12 to [1, Theorem B]).

The construction outlined below is based on methods already well established in the literature, and we will therefore only provide a sketch of the construction, omitting quite a few technical details but emphasizing some points that are important for the control of filtration aspects.

3.3.1 Floer chain complexes redefined

We begin by redefining the Floer chain complexes in a way that will enable us to obtain a genuinely filtered Fukaya category.

Let $L_0, L_1 \in \mathcal{X}$. Assume first that $\overline{L}_0 \neq \overline{L}_1$ (hence they intersect transversely). In this case, we fix a Floer datum $\mathcal{D}_{L_0,L_1}$ of the type $(0, J_{L_0,L_1})$, i.e., its Hamiltonian term will be identically 0. (Once again, the choice of J_{L_0,L_1} is made such that it depends only on the underlying Lagrangians $\overline{L}_0$, $\overline{L}_1$.) We then define $CF(L_0, L_1; \mathcal{D}_{L_0,L_1})$ to be the standard Floer complex associated to the pair $(\overline{L}_0, \overline{L}_1)$ using the Floer data $\mathcal{D}_{L_0,L_1}$ chosen above. The grading is defined using the grading on the two marked Lagrangians L_0, L_1. The filtration on $CF(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is defined by using the action as a filtration function. Specifically, if $x \in \overline{L}_0 \cap \overline{L}_1$ is a generator of $CF(L_0, L_1; \mathcal{D}_{L_0,L_1})$, its action $\mathcal{A}(x)$ is defined by (3.3).

Assume now that $\overline{L}_0 = \overline{L}_1$ and denote by $\overline{L}$ this common Lagrangian. In this case, the Floer datum will be replaced by a choice of a Morse datum (which we continue to denote by $\mathcal{D}_{L_0,L_1}$), namely a pair $(f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ of a Morse function $f_{\overline{L}} : \overline{L} \to \mathbb{R}$ and a Riemannian metric $(\cdot,\cdot)_{\overline{L}}$ on the common underlying Lagrangian $\overline{L}$. We will further assume that all the Morse functions $f_{\overline{L}}$ have a unique (local) maximum (i.e., a unique critical point of index $n = \dim_{\mathbb{C}} X$). The purpose of this assumption is to ensure that the units in our Fukaya category will be strict (rather than only homology units).

The Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ is defined to be the Morse complex $\mathrm{CM}(\overline{L})$ of $\overline{L}$, associated to the Morse data $\mathcal{D}_{L_0,L_1} = (f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$. We filter this chain complex in the following way. We set the filtration level *for all* generators $x \in \mathrm{CF}(L_0, L_1; \mathcal{D}_{L_0,L_1})$ of this chain complex (which are critical points of $f_{\overline{L}}$) to be the constant $c \in \mathbb{R}$, where $c \equiv h_{L_1} - h_{L_0}$ is the difference of the primitive functions of the two markings of the Lagrangian $\overline{L}$. To keep notation uniform, we continue to denote the filtration level of x by $\mathcal{A}(x)$, in the same way as we have done for the action.

3.3.2 Clusters of punctured disks

To define the μ_d-operations we will use a hybrid model that combines Floer polygons with gradient Morse trees. The maps defining the μ_d-operations will be called clusters of Floer polygons. This approach is analogous to the cluster Floer homology theory initiated by Cornea–Lalonde [21], who also introduced the name "clusters" in this context. Further modifications and foundational work on the subject has been done in [15, 16]. The main difference between these works and what we will be doing below is the following. The cluster homology theory [21] deals with a single Lagrangian in the presence of pseudoholomorphic disks. The "clusters" in that work consist of Morse flow lines attached to pseudoholomorphic disks. In contrast, here we deal with Floer theory of many Lagrangians together (setting up a Fukaya category), but in the absence of pseudoholomorphic disks. Note that [1] too uses clusters of Floer polygons in much the same way as below, and contains a detailed account of the subject.

In order to describe clusters of Floer polygons, we first need to set up their domains, which we call clusters of punctured disks.

We begin with the notion of a *k-punctured disk*. By this we mean a Riemann surface S_k which is obtained from the closed 2-disk $D \subset \mathbb{C}$ by removing $k \geq 1$ distinct boundary points $z_1, \ldots, z_k \in \partial D$, ordered in the *clockwise* direction, together with the following additional data. The points z_i will be called punctures. Each puncture z_i is declared to be either an *entry* puncture or an *exit* puncture. We allow S_k to have at most one exit puncture. We will typically denote the entry punctures by a $+$ superscript (e.g., z^+) and the exit puncture by a $-$ superscript (e.g., z^-). See Figure 3.1. Note that the boundary ∂S_k consists of k arcs, which we typically denote by $C_1, \ldots, C_k$, where the arc C_j goes from z_j to z_{j+1} for $1 \leq j \leq k-1$, in the clockwise direction, and C_k goes from z_k to z_1.

For each punctured disk S_k we fix a choice of strip-like ends along each of its punctures, as in [53, Chapter II, Section 9]. These choices should be compatible with splitting and gluing, as will be described later on.

Next we need to consider certain types of trees, which we call admissible. An admissible tree T is a (connected) tree with a finite number of edges and with the

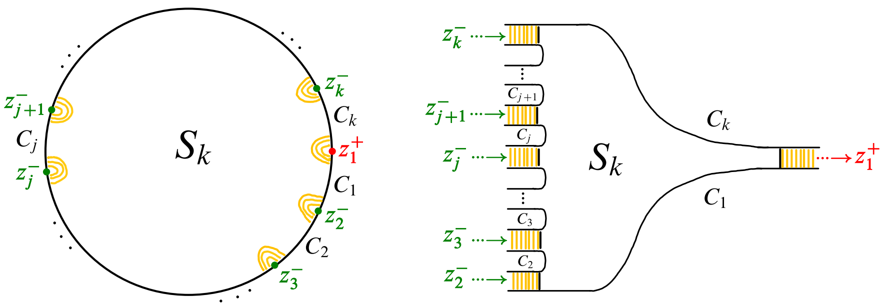

Figure 3.1. On the left: a k-punctured disk S_k with $k-1$ entries and one exit. The regions in yellow are the strip-like ends. On the right: a Riemann surface biholomorphic to a k-punctured disk, illustrating the strip-like ends modeled on $(-\infty, 0] \times [0, 1]$ for the entries and $[0, \infty) \times [0, 1]$ for the exit.

following properties and additional structures. In what follows, we will call all the end-vertices of T *leaves* (in particular, we will not distinguish between a possible root of the tree and the other end-vertices, and will just refer to all of them simply as leaves). We assume that all the leaves have valency 1 and all the other vertices of T (i.e., the internal ones) have valency 3. Moreover, the edges of T are oriented and these orientations satisfy the property that at every internal vertex (which by assumption has valency 3) there are precisely two incoming edges and one outgoing edge. The leaves of the tree T are divided into two types: E and A, where E stands for entry/exit leaves and A for attachment leaves. A leaf of type E will be called an *entry leaf* if the orientation on the edge connected to it goes from the leaf towards the rest of the tree. In the opposite case, i.e., when the orientation of that edge goes into the leaf, it will be called an *exit leaf*. The edges of T that are not connected to type-E leaves will be called *internal edges*. These consist of all edges that are not connected to any leaf, as well as those edges that are connected to leaves of type A. The other edges will be called external edges.

The edges of the trees are labeled by intervals in $\mathbb{R}$ as follows. The internal edges are labeled by intervals of the type $[0, R]$ (with possibly different values of $R > 0$ for different edges). The edges that connect a leaf of type E to an internal vertex are labeled either by $(-\infty, 0]$ or $[0, \infty)$, according to whether that leaf is an entry or an exit, respectively. If there is an edge connecting two leaves of type E (which happens if and only if the tree consists of exactly these two leaves and one edge connecting them), then this edge is labeled by the interval $(-\infty, \infty)$. Finally, we also fix an isotopy class of planar embeddings for the tree T. Note that, as a result, this fixes a cyclic clockwise order on the three edges connected to any given internal vertex. It also induces a cyclic clockwise order on the leaves of the entire tree. We illustrate a typical example of an admissible tree in Figure 3.2 below.

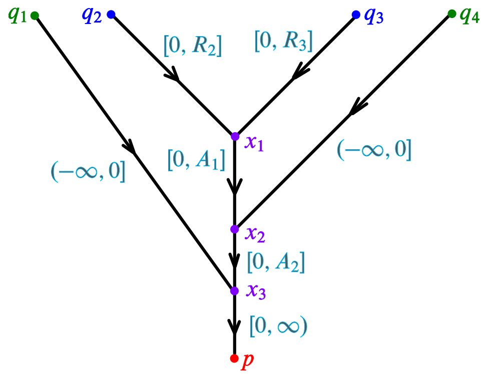

Figure 3.2. An admissible tree. The interval labelings are in dark gray. The leaves q_1, q_4 (in green) are type-E entry leaves, and p is a type-E exit leaf. The leaves q_2, q_3 (in blue) are of type A. The other vertices, x_1, x_2, x_3 (in purple), are interior vertices. The overall (cyclic) clockwise order on the leaves of this tree is q_1, q_2, q_3, q_4, p.

Having defined admissible trees, we now fix once and for all on each such tree T a collection of *orientation-preserving* identifications $\sigma_e^T : e \to I_e$ between each edge e and the interval I_e labeling it. Of course, in case the interval I_e is of the type $[0, \infty)$, $(-\infty, 0]$, or $(-\infty, \infty)$ (which happens when e is connected to vertices of type E) then the vertices corresponding to $\pm\infty$ are only asymptotically identified with $\pm\infty$. Note that every underlying tree has (infinitely) many different interval labelings for its edges (internal edges can be labeled by $[0, R]$ for different values of $R > 0$), leading to different admissible trees T. There is an obvious parametrization of these different interval labelings (basically by choosing the parameter R on each internal edge). We require the identifications σ^T to depend continuously on these parameters.

We are now in a position to introduce clusters of punctured disks. These are built from a collection $\mathcal{S} = \{S_{k_1}^{(1)}, \dots, S_{k_l}^{(l)}\}$, $l \geq 0$, of punctured disks and a collection of admissible trees $\mathcal{T} = \{T_1, \dots, T_r\}$, $r \geq 0$, which are attached to the punctured disks in $\mathcal{S}$ at their leaves of type A. The attachment of the trees is done as follows. Let $T \in \mathcal{T}$ and denote its leaves of type A by $a_1, \dots, a_{s_T} \in T$. For each $1 \leq i \leq s_T$, we identify the point $a_i \in T$ with a point lying on the boundary $\partial S_{k_j}^{(j)}$ of one of the punctured disks from $\mathcal{S}$. Here $j = j(T, i)$ depends on the tree T and the index i of the vertex a_i that is being attached. These attachments are subject to the following rule: each type-A leaf of a given tree $T \in \mathcal{T}$ is attached to one, and only one, punctured disk, and no two type-A leaves of the same tree $T \in \mathcal{T}$ are attached to the same punctured disk. There is no type-A leaf from the trees in $\mathcal{T}$ that is left unattached. We also require that, among the type-A leaves of all the trees in $\mathcal{T}$, there are no two leaves that are attached to the same point on the boundaries of the punctured disks.

We denote the space resulting from the above attachments by

$$\Sigma = \Big(\bigcup_{q=1}^{l} S^{(q)}_{k_q} \Big) \cup_A \Big(\bigcup_{p=1}^{r} T_p \Big), \tag{3.4}$$

where $\cup_A$ stands for the attachments described earlier. We will denote the part of Σ coming from the punctured disks (i.e., the leftward union in the right-hand side of (3.4)) by Σ_S, and the part coming from the trees (the rightward union in the right-hand side of (3.4)) by Σ_{tr}.

We now impose further restrictions on the previously described attachments. Consider the space obtained from Σ by collapsing each punctured disk $S^{(q)}_{k_q}$ from Σ_S to a (different) point:

$$\widetilde{\Sigma} := \Sigma / \big(S^{(q)}_{k_q} \sim \text{point}_q \ \forall q \big). \tag{3.5}$$

We require that the attachments of the tree described above be made in such a way that $\widetilde{\Sigma}$ is path-connected and, moreover, is a tree (hence, in particular, simply connected). We do not require this tree to be admissible.

Going back again to the space Σ, we note that it comes with a set of distinguished points: the punctures of the disks $S^{(q)}_{k_q}$ together with the leaves of the trees $T \in \mathcal{T}$ that are of type E. We call these points *external points* and denote by Σ_{ep} the set of all such points. The total number of external points of Σ will be called the *order* of Σ.

The external points Σ_{ep} are divided into two types: "entry points" and "exit points", regardless of a point being a puncture or a type-E leaf of some tree. We require that Σ_{ep} has precisely one exit point (which can be either an exit leaf or an exit puncture). We also require that Σ_{ep} has at least one entry point. Finally, in case there are only two such points (i.e., one entry and one exit), we require that either Σ is just a disk punctured at two points (with no trees attached) or that Σ has no punctured disks at all and it consists of just a tree with two vertices and one edge connecting them.

The last requirement on Σ is the following. Consider the tree $\widetilde{\Sigma}$ defined in (3.5). Note that, by construction, each edge of this tree is oriented (since the edges of all $T \in \mathcal{T}$ are oriented). Moreover, this tree has a distinguished vertex P_{exit}, namely the vertex that corresponds either to the punctured disk $S_k \subset \Sigma_S$ that contains the unique exit puncture, or to the unique exit leaf that belongs to one of the trees of Σ_{tr}. We require that the orientation on the tree $\widetilde{\Sigma}$ has the property that given any vertex $p \in \widetilde{\Sigma}$ there is a path from p to the distinguished vertex P_{exit} that is compatible with the orientation on $\widetilde{\Sigma}$. Figure 3.3 illustrates two examples of clusters of punctured disks (the labeling by Lagrangians $\overline{L}_i$ of the trees and the arcs in the ∂S_j's that appear in the picture should be ignored for the moment; these will be explained later on).

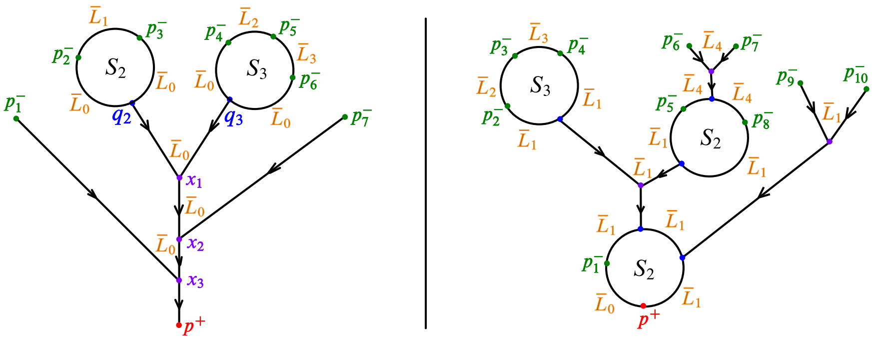

Figure 3.3. Two examples of clusters of punctured disks. The interval labeling of the edges of the trees is omitted here. The overall clockwise cyclic ordering of the external points on the left cluster is $(p_1^-, \dots, p_7^-, p^+)$ and on the right cluster $(p_1^-, \dots, p_{10}^-, p^+)$. The tuples describing the labeling by Lagrangians are $\bar{\mathcal{L}}_\Sigma = (\bar{L}_0, \bar{L}_0, \bar{L}_1, \bar{L}_0, \bar{L}_2, \bar{L}_3, \bar{L}_0, \bar{L}_0)$ for the left cluster and $\bar{\mathcal{L}}_\Sigma = (\bar{L}_0, \bar{L}_1, \bar{L}_2, \bar{L}_3, \bar{L}_1, \bar{L}_4, \bar{L}_4, \bar{L}_4, \bar{L}_1, \bar{L}_1, \bar{L}_1)$ for the cluster on the right.

Let Σ, Σ' be two spaces obtained as above from two pairs of collections $\mathcal{S}, \mathcal{T}$ and $\mathcal{S}', \mathcal{T}'$ of punctured disks and trees. We say that Σ and Σ' are equivalent if there is a homeomorphism $f : \Sigma \to \Sigma'$ with the following properties. The map f maps Σ_S biholomorphically to Σ'_S, and maps the trees Σ_{tr} to Σ'_{tr} by an isomorphism of trees (i.e., it maps vertices to vertices and edges to edges). Moreover, f intertwines all the other structures on $\Sigma_S, \Sigma_{\mathrm{tr}}$ with those on $\Sigma'_S, \Sigma'_{\mathrm{tr}}$. This means, in particular, that entry and exit punctures in Σ'_S correspond under f to the punctures of the same type in Σ_S; the same goes for the orientations on the edges of Σ_{tr}, Σ'_{tr}, the interval labeling, the identifications σ_e^T, and the classes of planar embeddings of the trees.

An equivalence class of spaces Σ as above (together with all the structures accompanying it) will be called a *cluster of punctured disks*. However, we will often use this name also for a specific representative Σ within a given equivalence class.

For a cluster of punctured disks, say represented by Σ, the orientation on the boundaries of the punctured disks in Σ_S and the classes of planar embeddings of the trees in Σ_{tr} induce a preferred clockwise cyclic order on the set of external points of Σ (recall that the external points consist of the entry and exit points, regardless of whether they are type-E leaves of the trees or punctures of the disks). Note that this ordering is preserved by the homeomorphisms defining the equivalence between different representatives Σ of the same class.

In what follows, it will be convenient to single out clusters of punctured disks of the following type. A cluster of punctured disks Σ is called *simple* if it consists of a single punctured disk without any trees attached.

We now turn to decorated clusters of punctured disks. Let Σ be a cluster of punctured disks. By a decoration of Σ by elements of $\bar{\mathcal{X}}$ we mean the following. We label each arc in the boundaries of Σ_S, as well as each edge in the trees of Σ_{tr}, by an element of $\overline{\mathcal{X}}$. The labeling is subject to the following restrictions. In each tree from Σ_{tr} all the edges are labeled by the same $\overline{L} \in \overline{\mathcal{X}}$ (alternatively, one can think of each tree $T \subset \Sigma_{\text{tr}}$ as being labeled by an element of $\overline{\mathcal{X}}$). The restriction on the labeling for the Σ_S-part of Σ is that in each punctured disk S_k from Σ_S there are no two consecutive arcs (i.e., two arcs with one puncture between them) that are labeled by the same element form $\overline{\mathcal{X}}$.

Once a cluster of punctured disks Σ is decorated by elements of $\overline{\mathcal{X}}$, we can form a tuple $\overline{\mathcal{L}}_\Sigma = (\overline{L}_0, \dots, \overline{L}_d)$ that encodes its decoration, where $d + 1 = |\Sigma_{\text{ep}}|$ is the order of Σ. The definition of $\overline{\mathcal{L}}_\Sigma$ goes as follows. Denote by $p^+, p_1^-, \dots, p_d^-$ the external points of Σ, ordered as explained earlier, where p^+ is the unique exit point and the p_j^-'s are all entry points. If p_j^- is a puncture of one of the disks $S_k \subset \Sigma_S$, we take $\overline{L}_j$ to be the Lagrangian that labels the arc on ∂S_k coming after the puncture (where "after" refers to the clockwise orientation on ∂S_k). If the entry p_j^- is a leaf of one of the trees $T \subset \Sigma_{\text{tr}}$, then we take $\overline{L}_j$ to be the Lagrangian that labels that tree. We define $\overline{L}_0$ in the same way, according to whether p^+ is a puncture or a leaf. Figure 3.3 shows two examples of decorated clusters of punctured disks.

We will now reverse in some sense the decoration construction. Namely, we fix a tuple $\overline{\mathcal{L}} = (\overline{L}_0, \dots, \overline{L}_d)$ of Lagrangians from $\overline{\mathcal{X}}$, and consider the space $\text{Clus}(\overline{\mathcal{L}})$ of all possible decorated clusters of punctured disks Σ with $\overline{\mathcal{L}}_\Sigma = \overline{\mathcal{L}}$. We call the elements of this space *$\overline{\mathcal{L}}$-decorated clusters of punctured disks*. As before, the elements of this space are equivalence classes of the spaces Σ, rather than the spaces Σ themselves. But it will often be convenient to work with an actual representative Σ of a given class.

Clearly every decorated cluster Σ belongs to a unique space $\text{Clus}(\overline{\mathcal{L}})$, since the tuple $\overline{\mathcal{L}}_\Sigma$ is uniquely defined by Σ.

Remark 3.10. Let $\overline{\mathcal{L}} = (\overline{L}_0, \dots, \overline{L}_d)$ be a tuple of Lagrangians from $\overline{\mathcal{X}}$.

(a) If $\overline{\mathcal{L}}$ has the property that $\overline{L}_i \neq \overline{L}_j$ for every $i \neq j$, then every cluster of punctured disks Σ that admit an $\overline{\mathcal{L}}$-decoration must be simple.

(b) The converse statement to point (a) above is obviously not true whenever $d \geq 3$. Namely, one can decorate a simple cluster of punctured disks Σ by a tuple $\overline{\mathcal{L}}$ whose entries do have repetitions. However, in such a case we must have $\overline{L}_i \neq \overline{L}_{i+1}$ for every $0 \leq i \leq d$ (where the indexing is to be understood cyclically $\text{mod}(d + 1)$).

(c) A tuple $\overline{\mathcal{L}}$ with $\overline{L}_i \neq \overline{L}_{i+1}$ for every $0 \leq i \leq d$ can decorate also non-simple clusters of punctured disks. However, if a non-simple Σ is decorated by such an $\overline{\mathcal{L}}$ then none of the trees in Σ_{tr} can have external leaves (which means that each leaf in Σ_{tr} is attached to some punctured disk in Σ_S).

3.3.3 Splitting and degeneration

Given a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ of Lagrangians from $\bar{\mathcal{X}}$, the space $\mathrm{Clus}(\bar{\mathcal{L}})$ of $\bar{\mathcal{L}}$-decorated clusters of punctured disks has the structure of a smooth manifold, analogous to the space of punctured disks from [53, Chapter II, Section 9]. This manifold admits a natural partial compactification into a manifold with corners. The top-dimensional strata of its boundary correspond to several types of degenerations of clusters of punctured disks, which we briefly describe below. Note that adding this boundary to $\mathrm{Clus}(\bar{\mathcal{L}})$ will still not make a full compactification of this space (hence the use of the words "partial compactification"); however, it will be enough for the purpose of establishing the A_∞-category identities. Below we will call those degenerations that lead to elements of this boundary *admissible degenerations* and their limiting objects *admissible degenerate clusters*.

Splitting within punctured disks. We begin by describing two variants of a degeneration that can occur to one punctured disk moving in a family. A family of punctured disks S_k^t (here k is the number of punctures and $t \in \mathbb{R}$ is parametrizing the family) can degenerate (or split) into two punctured disks $S'_{k'}$ and $S'_{k''}$, where $k' + k'' = k + 2$. The first punctured disk $S'_{k'}$ "inherits" $k' - 1$ of the entry punctures of S (placed in the same clockwise order as in S) and has one additional exit puncture z'_+. The other component, $S''_{k''}$, "inherits" all the other $k - (k' - 1) = k'' - 1$ punctures of S (again, in the same clockwise order) and has one additional entry puncture z''_-. Note that at the moment we do allow k' or k'' to take the values 1 or 2. This is in contrast to the more standard realizations of the Fukaya category, where each of the two disks in a splitting are required to have at least 3 punctures. However, later on, when viewing these disks as part of a degenerate cluster, more restrictions will be added in order to make such a degenerate configuration an admissible one.

Conversely, the two punctured disks $S'_{k'}$ and $S''_{k''}$ can be glued along the punctures z'_+, z''_- into a family of punctured disks S_k^t.

Depending on the context, in what follows we will sometimes view the preceding degeneration differently. Namely, regard the two punctures z'_+, z''_- as "removable" and view the degeneration of the family S_k^t as a splitting into two punctured disks $S'_{l'}$ and $S''_{l''}$ attached one to the other at a point (which is not a puncture) on their boundaries. Note that now we have $l' + l'' = k$, and similarly to the preceding case, we do allow l' or l'' to take the values 1 and 2.

Conversely, as before, the two punctured disks $S'_{l'}$ and $S''_{l''}$ can be glued into a family of punctured disks S_k^t.

Analytically, the two variants described above are the same; however, when taking decorations into account it is important to distinguish between them. More precisely, if the punctured disks S_k^t are decorated by the Lagrangians $(\bar{N}_0, \dots, \bar{N}_{k-1})$, the first variant of splitting corresponds to two punctured disks with decorations $(\bar{N}_r, \dots, \bar{N}_s)$

and $(\overline{N}_0, \ldots, \overline{N}_r, \overline{N}_s, \ldots, \overline{N}_{k-1})$, where $1 \le r < s \le k-1$, $\overline{N}_r \ne \overline{N}_s$, $s - r + 1 = k'$, and $k'' = k - (s - r) + 1$.

The second variant (i.e., where S^t_k degenerates into $S'_{l'}$ and $S''_{l''}$ attached at a common point, which is not a puncture, along their boundaries) corresponds to the case when the decoration $(\overline{N}_0, \ldots, \overline{N}_{k-1})$ has $\overline{N}_r = \overline{N}_s$ for some non-consecutive indices, $r < s - 1$, and the splitting yields the decorations $(\overline{N}_r, \ldots, \overline{N}_{s-1})$ and $(\overline{N}_0, \ldots, \overline{N}_r, \overline{N}_{s+1}, \ldots, \overline{N}_{k-1})$ on $S'_{l'}$ and $S''_{l''}$, respectively.

The two variants of splitting are illustrated in Figure 3.4.

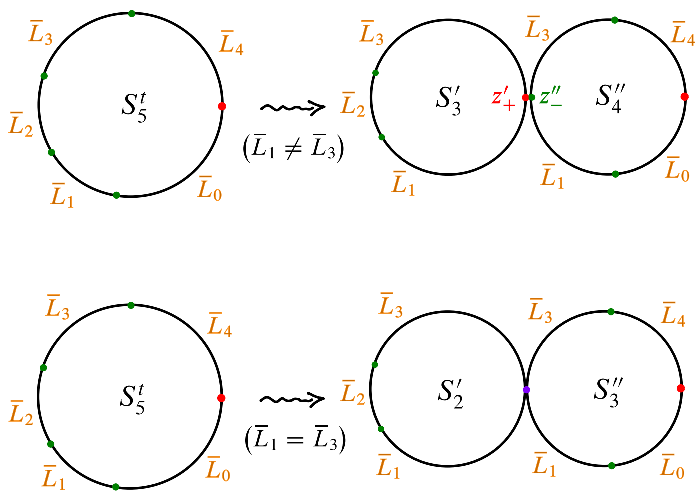

Figure 3.4. Two variants of splitting of decorated punctured disks. The points z'_+ (red) and z''_- (green) on the right-hand side of the upper figure are new punctures. The purple point on the right-hand side of the lower figure, where the two disks are attached, is *not* a puncture.

We now turn to clusters of punctured disks and describe their *admissible degenerations*. Fix a tuple $\overline{\mathcal{L}} = (\overline{L}_0, \ldots, \overline{L}_d)$ of Lagrangians. A family Σ^t of $\overline{\mathcal{L}}$-decorated clusters of punctured disks can converge to a decorated degenerate cluster Σ^∞ of punctured disks (which, strictly speaking, by our definitions, might not be a genuine cluster of punctured disks). The degeneration of Σ^t into Σ^∞ can be of several types. The first type is when one (or more) of the punctured disks in the clusters Σ^t degenerates in the way described earlier. Depending on the decoration $\overline{\mathcal{L}}$, one of the two variants mentioned above, or both, can occur. There is one slight exception to this rule. Namely, in both of the variants described above we view the degeneration as admissible only if each of the two punctured disks formed by the splitting contains at least three distinguished points. Here, by a distinguished point we mean either a puncture, a point attached to a tree, or (in case of the second variant) the point of attachment to the other punctured disk in the degenerate cluster.

Below we will describe another four types of admissible degenerations. Before we go into this, a quick remark about the decorations of the limit Σ^{∞} is in order.

Our conventions for decorations require the cluster to have an exit point (according to which we label the first entry in the decoration). However, the first variant of the degenerations described above yields two punctured disks S' and S'', where one of them has a (new) exit point and the other has a (new) entry point. The apparent problem is that one of these punctured disks might not have any exit point, hence there might be an ambiguity regarding the order in which we write its decoration. However, this ambiguity is fixed if we use the following conventions. The limit Σ^{∞} is divided into two components: the one that contains S' and the one that contains S''. The decorations are uniquely defined once we require that the exit point of the limit Σ^{∞} corresponds to the S'' part. A similar thing applies also to clusters in which one of the punctured disks degenerates according to the second variant described earlier.

We now proceed to describe four additional types of admissible degenerations.

Splitting within trees. Apart from degeneration of punctured disks in a cluster, there are several other types of degeneration that can occur within a family Σ^t of clusters of punctured disks. Part of these has to do with degeneration of the trees Σ^t_{tr} of Σ^t, and another part is related to how these trees are attached to Σ^t_S.

Shrinking of edges to a point. The first type of degeneration within trees is when an interior edge in one of the trees of Σ^t_{tr} shrinks to a point (this means that also its interval labeling and parametrization shrink to a point and a constant, respectively). The limit tree will now have one vertex less and will inherit from Σ^t_{tr} all the other structures (such as the labeling of the other edges, the class of planar embedding, etc.). See Figure 3.5. Note, however, that the limit tree will not be admissible (e.g., it might have vertices of valency 4, or a leaf of type A that becomes identified with an interior vertex).

Edge breaking. Another type of degeneration is when one of the interior edges e in a family of trees $T^t \subset \Sigma^t$ becomes of infinite length. We view the limit of the T^t's as a broken tree consisting of two trees T_1^{∞} and T_2^{∞}. These two trees are obtained from T^{∞} as follows. We delete the (interior of the) edge e from T^t and obtain two connected components: the part of T^t that appears "before" the edge e and the part that appears "after" that edge (here "before" and "after" refer to the orientation on T^t). Denote by p_1 the end-vertex of the first component (i.e., the entry vertex to the edge e) and by p_2 the new entry vertex of the second component (which corresponds to the exit vertex of e). We now take the first component and add to it a new edge e_1 emanating from p_1. The result is the tree T_1^{∞}. The other vertex q_1 of e_1 will now be a type-E leaf of T_1^{∞} and we regard it as an exit leaf. We label the edge e_1 by $[0, \infty)$, and the rest of the edges are labeled by the limiting labels of T^t as $t \to \infty$. The definition of T_2^{∞} is similar, only that now we add a new edge e_2 to the second component (i.e.,

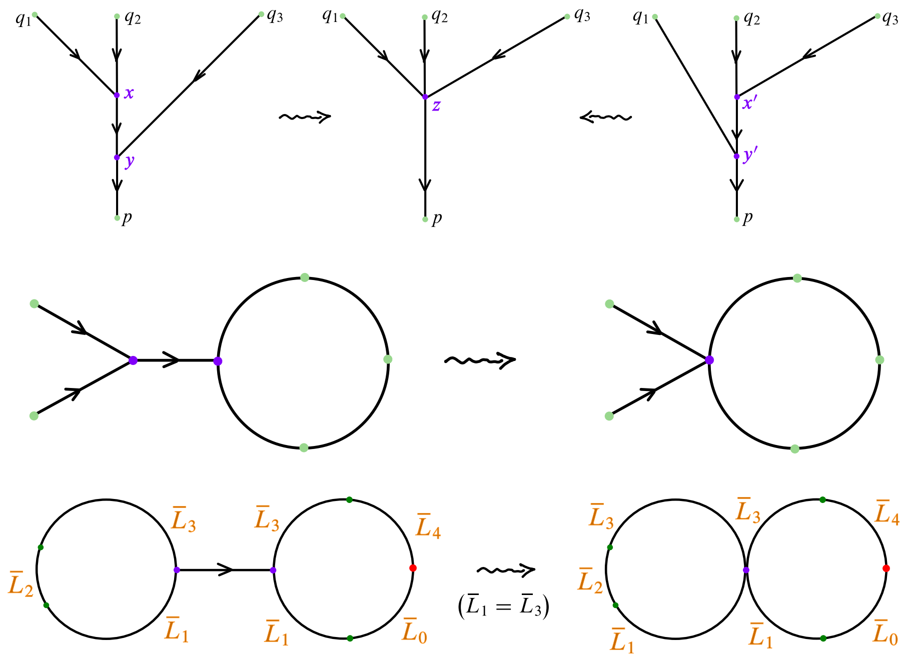

Figure 3.5. Examples of degenerations where an edge of a tree shrinks to a point, leading to a limit (inadmissible) tree, possibly being part of a degenerate cluster.

the one coming "after" the deleted edge e) attached at p_2. The resulting tree is T_2^∞. The new vertex (which is the entry to e_2) will be a type-E leaf of T_2^∞. The edge e_2 is labeled by $(-\infty, 0]$ and the rest of the edges are labeled by the limiting labels of T^t as $t \to \infty$. We refer to T^∞ as a "broken" tree with components T_1^∞ and T_2^∞. See Figure 3.6.

We add the following restriction on edge-breaking degenerations. A degeneration as described above is considered admissible only if neither T_1^∞ nor T_2^∞ is a tree with two vertices both of which are type-E leaves, connected by one edge. All other edge-breaking degenerations are considered admissible.

Collision of type-A leaves. The last type of admissible degeneration is when two type-A leaves (belonging to two different trees) that lie on the boundary of the same punctured disk $S_k \subset \Sigma_S$ collide. This means that two trees $T', T'' \subset \Sigma_{\mathrm{tr}}$ are grafted (or joined) at two of their type-A exit leaves. See Figure 3.7.

Remark 3.11. (a) The boundary of the compactification of the space $\mathrm{Clus}(\bar{\mathcal{L}})$ can be described by the types of degeneration described above. The points of the top-dimensional stratum of the boundary correspond to precisely one such degeneration. Of course, several instances of degeneration can occur simultaneously, but these

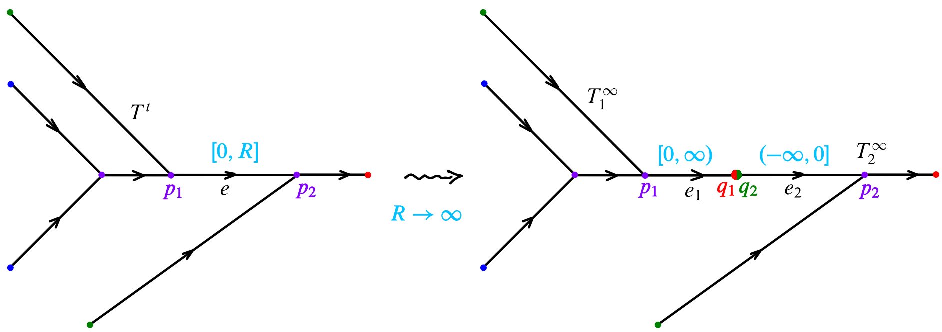

Figure 3.6. Breaking along an edge of a tree, leading to a broken tree with two components.

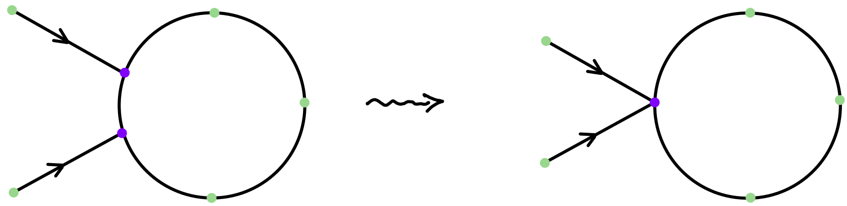

Figure 3.7. Collision of points on ∂S_k to which two different trees are attached.

instances correspond to the lower-dimensional strata of the boundary of $\mathrm{Clus}(\bar{\mathcal{L}})$. In particular, within generic 1-dimensional families of clusters of punctured disks, only one degeneration can occur at a given time.

(b) The converse to "splitting and degeneration" goes by the name gluing. Every degenerate configuration among the ones described above can be obtained as a limit of a family of clusters of punctured disks.

(c) Some of the limit configurations described above can occur as a result of *two different* degenerations. For example, collision of two type-A leaves of trees (along an arc of one of the punctured disks) leads to a configuration which is also the limit of another family of clusters in which one type-A leaf of a tree shrinks to a point. See Figure 3.7 versus the middle part of Figure 3.5.

Similarly, the second variant of splitting within a punctured disk (which leads to two punctured disks connected at a "non-puncture" point along their boundaries) also occurs as a limit of clusters in which two punctured disks are connected by a tree with one edge, and that edge shrinks to a point. See the lower part of Figure 3.4 versus the lower part of Figure 3.5.

The same thing happens with shrinking of interior edges of trees. Namely, each of the (inadmissible) trees that occur after the shrinking of interior edges appears as a limit of a different family of trees in which an edge shrank to a point. See the left- and right-hand sides of the upper part of Figure 3.5.

The fact that some limit configurations appear in pairs on the total boundary of the compactification of the Clus spaces is important for showing that the μ_d-operations in $\mathcal{F}uk(\overline{\mathcal{X}})$, as will be defined later, satisfy the A_∞-identities. Indeed, when one considers 1-parametric families of clusters of Floer polygons, with fixed entries and exit, some of the boundary points (that correspond to degeneration of the underlying clusters of punctured disks as described above) will appear in pairs and thus can be regarded as "interior" points inside extended families of clusters.

(d) Recall that the boundary of the partial compactification of $\mathrm{Clus}(\overline{\mathcal{L}})$ contains only admissible degenerate clusters. In particular, in the case of splitting of a punctured disk we require that the number of distinguished points on each component is at least 3. Of course, a splitting in which one of the disks has only two distinguished points can occur. The reason we do not add such a configuration to the boundary $\mathrm{Clus}(\overline{\mathcal{L}})$ is that disks with two marked points have a non-compact 1-dimensional group of automorphisms (isomorphic to $\mathbb{R}$). This is referred to in the literature as an unstable marked curve. The situation with the other inadmissible degenerate clusters, namely a broken tree with one component being a tree with one edge connecting between two type-E leaves, is similar. The latter component has an $\mathbb{R}$-action (acting by translation on the identifications between $(-\infty, \infty)$ and the edge of this tree).

In practice, not including these unstable configurations to the boundary of $\mathrm{Clus}(\overline{\mathcal{L}})$ will not cause any problems in showing that the μ_d-operations satisfy the A_∞-identities. The standard way to go about it in Floer and Morse theory is to compactify the space of clusters of *Floer polygons* in such a way that degenerations that correspond to the above unstable configurations are taken into account in the boundary of the latter spaces rather than in $\partial\,\mathrm{Clus}(\overline{\mathcal{L}})$. In terms of the A_∞-identities, these degenerations will contribute the terms in the identities that include μ_1's.

3.3.4 Perturbation data

We assume that Floer data has been chosen for every pair $L_0, L_1 \in \mathcal{X}$, as described at the beginning of Section 3.3.1 on page 109. The perturbation data for a decorated simple cluster (i.e., a cluster consisting of precisely one punctured disk and no trees) S_k is of the same type as in the standard theory, namely it consist of pairs (K, J), where K is a 1-form on S_k with values in the space of compactly supported Hamiltonian functions on the ambient manifold X. This 1-form is assumed to be compatible with the Floer data on each strip-like end of S_k, in the sense that on these ends we have $K \equiv 0$. (Note that we are dealing here with the case of one punctured disk without trees, which means that the decoration $\overline{\mathcal{L}}$ is such that every two consecutive Lagrangians in $\mathcal{L}$ have mutually transverse underlying Lagrangians. Recall also that for pairs of distinct underlying Lagrangians we have already set up the Floer data in advance to have 0 Hamiltonian terms.) Moreover, we require K to have compact support inside the interior of S_k. The second component of the Floer data is a family of ω-compatible almost complex structures $J = \{J_z\}$ that depends on $z \in S_k$ and coincides on each strip-like end with the almost complex structures chosen for the corresponding Floer data.

We now describe the perturbation data in the case of general decorated clusters of punctured disks Σ. The perturbation data in this case consists of two pieces of data. The first one is a choice of perturbation data (K, J) on each (decorated) punctured disk $S_k \subset \Sigma_S$. The second one is a choice of Morse data on each tree $T \subset \Sigma_{\text{tr}}$ of the cluster. Recall that each such tree T corresponds to an underlying Lagrangian $\overline{L}$ that appears in the decoration $\mathcal{L}$. Recall also that each edge e of T is parametrized by an interval $I_e \subset \mathbb{R}$ (where the intervals for the internal edges are closed of finite length and the ones corresponding to the edges that touch the type-E leaves are semi-infinite). The Morse data for T is a choice of a family $(f_\tau, (\cdot,\cdot)_\tau)$, $\tau \in I_e$, for every edge e of the tree, where for each τ, $f_\tau : \overline{L} \to \mathbb{R}$ is a smooth function and $(\cdot,\cdot)_\tau$ a Riemannian metric on $\overline{L}$. Here $\overline{L}$ is the underlying Lagrangian corresponding to the tree T. Moreover, we require that along the *ends* of the external edges e of T (i.e., the edges connected to the type-E leaves), the pair $(f_\tau, (\cdot,\cdot)_\tau)$ coincides with the Morse data $(f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ associated to $\overline{L}$ that has been fixed in advance. For example, if e is an edge connected to an entry leaf, then $I_e = (-\infty, 0]$, and the requirement is that $(f_\tau, (\cdot,\cdot)_\tau) = (f_{\overline{L}}, (\cdot,\cdot)_{\overline{L}})$ for $\tau \ll 0$. A similar choice of data is made also in case there is an exit leaf (which is the case when $\overline{L}_0 = \overline{L}_d$), only that now the edge e connected to the exit leaf is labeled by $I_e = [0, \infty)$.

There is only one slight exception to the above, namely when $\overline{\mathcal{L}} = (\overline{L}, \overline{L})$. In this case, the whole cluster consists of only one tree (and no punctured disks). This tree has two vertices and one edge e connecting them, which is modeled on the interval $I_e = (-\infty, \infty)$. The choice of Morse data here will be the same Morse data $(f_{\overline{L}}, (\cdot,\cdot))$ chosen in advance for $\overline{L}$, and it is required to be independent of the parameter $\tau \in I_e$.

For every tuple $\overline{\mathcal{L}}$ we make a continuous choice of perturbation data for all the clusters of punctured disks that are parametrized by $\text{Clus}(\overline{\mathcal{L}})$. We denote such a choice by $\mathcal{P}_{\overline{\mathcal{L}}}$ and denote by $\mathcal{P} = \{\mathcal{P}_{\overline{\mathcal{L}}}\}$ the collection of choices $\mathcal{P}_{\overline{\mathcal{L}}}$ made for all tuples $\overline{\mathcal{L}}$ of any length. We refer to $\mathcal{P}$ as a choice of perturbation data.

These choices of $\mathcal{P}$ are subject to being consistent with the splitting, degeneration, and gluing described in Section 3.3.3. This is crucial in order to establish the A_∞-identities among the μ_d-operations that will be described next.

3.3.4.1 The μ_d-operations. We now proceed to the A_∞-operations, taking also filtrations into account. The definition of μ_d, $d \geq 2$, is based on clusters of Floer polygons, which we now describe.

Let $\overline{\mathcal{L}} = (\overline{L}_0, \ldots, \overline{L}_d)$ be a tuple of Lagrangians in $\overline{\mathcal{X}}$. An $\overline{\mathcal{L}}$-decorated cluster of Floer polygons is a map $u : \Sigma \to X$ whose domain is an $\overline{\mathcal{L}}$-decorated cluster of punctured disks Σ. The restriction $u|_{S_k}$ of u to any of the punctured disks $S_k \subset \Sigma_S$ is a Floer polygon, exactly as in the standard theory of Fukaya categories [53, Chapter II, Section 9]. Namely, $u|_{S_k}$ satisfies the (generalized) Floer equation associated to the perturbation data that $\mathcal{P}_{\bar{\mathcal{L}}}$ assigns to S_k. The map $u|_{S_k}$ satisfies Lagrangian boundary conditions prescribed by the decoration. The punctures of S_k are sent by u to intersection points of pairs of Lagrangians, as prescribed by the decoration. In addition, we assume that the energy $E(u|_{S_k})$ of $u|_{S_k}$ is finite.

Next, the restriction $u|_T$ of u to any of the trees $T \subset \Sigma_{\mathrm{tr}}$ should satisfy the *negative* gradient equations corresponding to the Morse data specified along the edges of the trees. The interval labeling I_e on the edges e of T and the identifications $\sigma_e^T : e \to I_e$ are used in order to endow each interval with a "time parameter" for the negative gradient trajectories. Finally, the type-E leaves of each tree T from $\mathcal{T}$ are mapped by u to critical points of the functions $f_{\bar{L}}$, where $\bar{L}$ is the Lagrangian decorating the tree T.

Given the choices of Floer and perturbation data, the definition of the μ_d-operations, $d \geq 2$, is now done by counting decorated clusters of Floer polygons with specified boundary conditions and given entry and exit points. More specifically, let $\mathcal{L} = (L_0, \dots, L_d)$, $d \geq 1$, be a tuple of Lagrangians from $\mathcal{X}$, and denote by $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ the corresponding tuple of underlying Lagrangians. Define

$$\begin{aligned} &\mu_d : \mathrm{CF}(L_0, L_1) \otimes \cdots \otimes \mathrm{CF}(L_{d-1}, L_d) \to \mathrm{CF}(L_0, L_d), \\ &\mu_d(x_1, \dots, x_d) := \sum_y \#\mathcal{M}(x_1, \dots, x_d, y; \mathcal{P}) y. \end{aligned} \tag{3.6}$$

Here we have abbreviated $\mathrm{CF}(L', L'') := \mathrm{CF}(L', L''; \mathcal{D}_{L',L''})$ for any $L', L'' \in \mathcal{X}$. The sum in the second line of (3.6) goes as follows: y runs over all the generators of $\mathrm{CF}(L_0, L_d)$ of appropriate degree, and $\#\mathcal{M}(x_1, \dots, x_d, y; \mathcal{P})$ stands for the count (with values in $\mathbb{Z}_2$, or under additional assumptions in $\mathbf{k}$) of elements in the 0-dimensional component of the space $\mathcal{M}(x_1, \dots, x_d, y; \mathcal{P})$ of $\bar{\mathcal{L}}$-decorated clusters of Floer polygons with entry points $x_1, \dots, x_d$ and exit point y.

We denote by $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$ the collection of objects $\mathcal{X}$ together with the multilinear operations μ_d, $d \geq 1$, associated to the Floer and perturbation data $\mathcal{P}$. Notice that the perturbation data depends only on the geometric part of the marked Lagrangians, namely $\bar{\mathcal{X}}$.

Theorem 3.12. *For every finite collection of Lagrangians $\bar{\mathcal{X}}$ satisfying the conditions from the beginning of Section 3.2.1, there exists a (non-empty) space $\mathcal{B}(\bar{\mathcal{X}})$ of regular Floer and perturbation data, of the types described above, such that for every $\mathcal{P} \in \mathcal{B}(\bar{\mathcal{X}})$ the following holds:*

(i) *$\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, with the above μ_d-operations, $d \geq 1$, is a strictly unital A_∞-category. Moreover, with the filtrations defined in Section 3.3.1, this A_∞-category is genuinely filtered.*

(ii) *If one forgets the filtrations, then $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ as defined above is quasi-equivalent to the subcategory of the standard Fukaya category (e.g., as defined in* [53]*) whose collection of objects is $\mathcal{X}$. This quasi-equivalence can be assumed to be the identity map on the set of objects $\mathcal{X}$.*

Moreover, there exist filtered A_∞-functors

$$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}uk(\mathcal{X};\mathcal{P}_0) \to \mathcal{F}uk(\mathcal{X};\mathcal{P}_1),$$

defined for every $\mathcal{P}_0, \mathcal{P}_1 \in \mathcal{B}(\bar{\mathcal{X}})$, with the following properties:

(1) *$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is A_1-unital (see Section 3.3.6 for the definition).*

(2) *$\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is a filtered quasi-equivalence.*

(3) *The action of $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ on objects is the identity map. (Recall that all the categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ have the same set of objects $\mathcal{X}$.)*

(4) *For every $L', L'' \in \mathcal{X}$, the maps $(\mathcal{F}_1^{\mathcal{P}_1,\mathcal{P}_0})_* : \mathrm{HF}(L', L''; \mathcal{P}_0) \to \mathrm{HF}(L', L''; \mathcal{P}_1)$ induced by the first order components of $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ are the canonical continuation isomorphisms in Floer theory.*

(5) *$\mathcal{F}^{\mathcal{P},\mathcal{P}} = \mathbb{1}$.*

(6) *The composition $\mathcal{F}^{\mathcal{P}_2,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is isomorphic to $\mathcal{F}^{\mathcal{P}_2,\mathcal{P}_0}$ in the category $H^0(ffun(\mathcal{F}uk(\mathcal{X};\mathcal{P}_0), \mathcal{F}uk(\mathcal{X};\mathcal{P}_2)))_0$. Here, $ffun(\mathcal{F}uk(\mathcal{X};\mathcal{P}_0), \mathcal{F}uk(\mathcal{X};\mathcal{P}_2))$ is the category of filtered A_∞-functors from $\mathcal{F}uk(\mathcal{X};\mathcal{P}_0)$ to $\mathcal{F}uk(\mathcal{X};\mathcal{P}_2)$, $H^0(ffun(\cdots))$ is the persistence homological category of ffun in cohomological degree 0, and $H^0(ffun(\cdots))_0$ is its 0-level persistence subcategory. In other words, there exists an A_∞-natural transformation $T^{\mathcal{P}_2,\mathcal{P}_1,\mathcal{P}_0} : \mathcal{F}^{\mathcal{P}_2,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0} \to \mathcal{F}^{\mathcal{P}_2,\mathcal{P}_0}$ that preserves filtrations and is an isomorphism in the homological persistence category of filtered functors. Note that, in particular, this implies that $\mathcal{F}^{\mathcal{P}_0,\mathcal{P}_1} \circ \mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ is isomorphic to $\mathbb{1}_{\mathcal{F}uk(\mathcal{X};\mathcal{P}_0)}$ in the respective homological persistence category.*

Furthermore, the choice of the assignment $\bar{\mathcal{X}} \mapsto \mathcal{B}(\bar{\mathcal{X}})$ can be assumed to have the following property: if $\bar{\mathcal{X}}'$ is another finite collection of Lagrangians with $\bar{\mathcal{X}}' \supset \bar{\mathcal{X}}$ that (similarly to $\bar{\mathcal{X}}$) satisfies the conditions from the beginning of Section 3.2.1, then $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}} \subset \mathcal{B}(\bar{\mathcal{X}})$. Here, $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}}$ stands for the restriction of the perturbation data from $\mathcal{B}(\bar{\mathcal{X}}')$ to the spaces of clusters of punctured disks decorated by the elements of $\bar{\mathcal{X}}$.

Remark 3.13. The system of functors $\{\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}\}_{\mathcal{P}_1,\mathcal{P}_0 \in \mathcal{B}(\bar{\mathcal{X}})}$ and the natural transformations mentioned in Theorem 3.12 depend on a variety of choices and hence are not canonical in the strict sense of the word. The extent to which these structures are canonical will be briefly discussed later in Remark 3.24. We will refer to this system of functors as a *weakly coherent system* of comparison functors to emphasize that our construction does not produce canonical choices.

We proceed now to the proof of Theorem 3.12. The proof presented below is by no means complete and should be viewed as an outline only. We have left out quite a few technical details, especially concerning the analysis underlying the proof. However, these parts of the proof follow from rather standard and well-established ingredients in the analysis of Floer theory and Fukaya categories. As mentioned earlier, a more general approach to genuinely filtered Fukaya categories is worked out in a forthcoming paper by Ambrosioni [1], which will also contain a detailed proof of the construction.

3.3.5 Proof of Theorem 3.12, part 1

We will concentrate here on the second part of point (i) of the theorem (namely that $\mathcal{F}uk(\mathcal{X})$ can be made *genuinely filtered* for appropriate choices of Floer and perturbation data). The first part of point (i) and point (ii) are rather well known and have been addressed in the literature with various levels of rigor. The proofs of the statements concerning the functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$ will be outlined in Section 3.3.6 below.

Throughout the proof we will sometimes abbreviate $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ as $\mathcal{F}uk(\mathcal{X})$ in case $\mathcal{P}$ is clear from the context.

Fix a tuple of Lagrangians $\bar{\mathcal{L}} = (\bar{L}_0, \ldots, \bar{L}_d)$ and assume for simplicity that $\bar{L}_i \neq \bar{L}_j$ for every $i \neq j$. Fix a tuple of intersection points $x_1, \ldots, x_d, y$, where $x_i \in \bar{L}_{i-1} \cap \bar{L}_i$ and $y \in \bar{L}_0 \cap \bar{L}_d$. Let $u \in \mathcal{M}(x_1, \ldots, x_d, y; \mathcal{P})$ be an $\bar{\mathcal{L}}$-decorated cluster of Floer polygons. By our simplifying assumptions, the domain of u must be a simple cluster, namely just one punctured disk S_{d+1}. (See page 114 and Remark 3.10.) Thus $u : S_{d+1} \to X$ is a Floer polygon that sends the punctures of S_{d+1} to the points $x_1, \ldots, x_d, y$. For such a map u, denote by $A(u) = \int_{S_{d+1}} u^*\omega$ the symplectic area of u. We have

$$A(u) = \sum_{i=1}^{d} \mathcal{A}(x_i) - \mathcal{A}(y).$$

In order to prove that the μ_d-operation preserves filtrations, we need to show that for all u as above we have $A(u) \geq 0$. (In fact, we need to show the latter inequality holds for all clusters of Floer polygons u, not only the simple ones. However, as we will see below, the main difficulty is for punctured disks, and the generalization to more general clusters is straightforward.)

Before we go into the proof of the latter statement, let us explain the difficulties underlying it. Obviously, $A(u) \geq 0$ if we choose the perturbation data (K, J) with $K \equiv 0$, since then every Floer polygon will be J-holomorphic and hence of strictly positive area. However, for a variety of reasons it seems better to allow for non-trivial 1-forms K in the perturbation data, so we will not assume $K \equiv 0$. One of the reasons for allowing non-trivial perturbations is that it is easier to establish transversality for the spaces of Floer clusters with this extra parameter at hand. Another reason is that if one hopes to generalize the present approach to cases when not all pairs of Lagrangians in $\bar{\mathcal{X}}$ intersect transversely then Hamiltonian perturbations would definitely be needed. Other reasons have to do with compatibility of $\mathcal{F}uk(\mathcal{X})$ with other structures, such as the Floer (or symplectic) homology of the ambient manifold and maps relating $\mathcal{F}uk(\mathcal{X})$ to these invariants. These structures usually require Hamiltonian perturbations.

We may next try to show that $A(u) \geq 0$ once we take the 1-forms K in the perturbation data to be small enough. At first sight, this seems to work using a compactness argument as follows. Indeed, if this were not the case, then we would have a sequence of perturbation forms $K^{(l)}$ and a sequence $u_l : S^{(l)}_{d+1} \to X$ of corresponding Floer polygons whose punctures go to a fixed set of intersection points $x_1, \ldots, x_d, y$, such that $K^{(l)} \xrightarrow{C^1} 0$ but $A(u_l) \leq 0$ for every l. (We may assume that all the u_l's run between the *same* set of intersection points $x_1, \ldots, x_d, y$, because by our assumptions there is only a finite number of possible intersection points associated to the tuple $\bar{\mathcal{L}}$.) By compactness, passing to a subsequence of the u_l's we would then obtain a (possibly broken) limit polygon u which is non-constant and genuinely J-holomorphic, yet with $A(u) \leq 0$. A contradiction. Since, by assumption, the number of possible $d+1$ tuples $\bar{\mathcal{L}}$ is finite, it follows that, if we take the perturbation data small enough, then for all tuples $\bar{\mathcal{L}}$ of length $d+1$ with the properties from the beginning of the proof, and all $\bar{\mathcal{L}}$-decorated Floer polygons u, we have $A(u) > 0$. This easily extends also to decorated clusters of Floer polygons that are not necessarily simple, as well as to all decorations $\bar{\mathcal{L}}$ of *fixed* length $d+1$.

The problem with this argument is that, without further elaboration, it might create difficulties with obtaining a consistent choice of perturbation data $\mathcal{P}$. To explain this difficulty, let us rephrase the previous paragraph in more quantitative terms. From now on we will use the following more detailed notation. We denote the restriction of the perturbation data $\mathcal{P}$ to the space of all $(d+1)$-punctured disks by $\mathcal{P}_{d+1}$, and write $(K(\mathcal{P}_{d+1}), J(\mathcal{P}_{d+1}))$ for the two components of $\mathcal{P}_{d+1}$. Given an $\bar{\mathcal{L}}$-decorated punctured disk S_{d+1}, we denote by $(K(\mathcal{P}, S_{d+1}), J(\mathcal{P}, S_{d+1}))$ the restriction of $\mathcal{P}$ to S_{d+1}.

The previous argument shows that there exist numbers $\epsilon_{d+1} > 0$, $d \geq 2$, such that if the perturbation forms $K(\mathcal{P}_{d+1})$ satisfy $\|K(\mathcal{P}_{d+1})\| \leq \epsilon_{d+1}$ then for all tuples $\bar{\mathcal{L}}$ of length $d+1$ and all $\bar{\mathcal{L}}$-decorated Floer polygons u, we have $A(u) > 0$. Here $\|-\|$ is a suitable norm on the space of all perturbation 1-forms K_{d+1} (defined on the space of all possible $(d+1)$-punctured disks S_{d+1}). The value of $\|K_{d+1}\|$ involves the values of K_{d+1} and its first derivatives, both in the domain direction as well as in the direction of the manifold X (recall that the forms K_{d+1} have values in the space of compactly supported functions on X).

The problem that arises with the approach used so far has to do with the consistency of $\mathcal{P}$ with respect to gluing/splitting. A standard way to construct consistent perturbation data is to construct $K(\mathcal{P}_{d+1})$ (and the almost complex structures) by induction on d and to ensure that at each induction step the newly defined perturbation data is consistent with the data that have already been defined at earlier stages, with respect to gluing/splitting. Assume that $K(\mathcal{P}_{m+1})$ has already been defined for all $m \leq d_0$ in such a way that $\|K(\mathcal{P}_{m+1})\| \leq \epsilon_{m+1}$ for all $m \leq d_0$. Consider now $d = d_0 + 1$. Punctured disks of the type S_{d+1} can split into two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, with $d' + d'' = d + 1$ (and $d', d'' \geq 2$). See Figure 3.8.

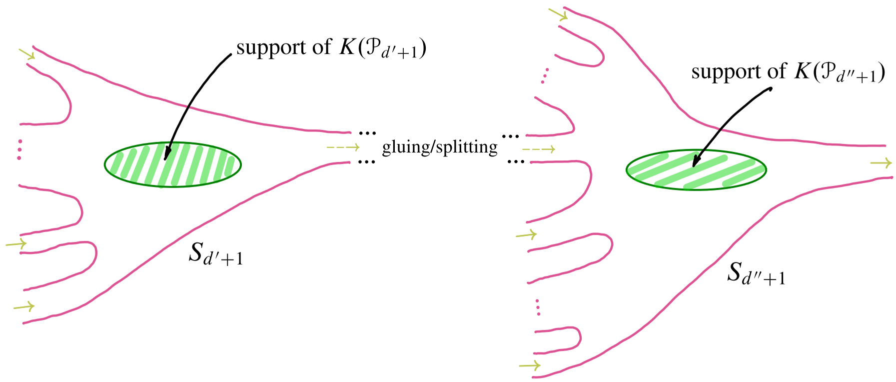

Figure 3.8. Gluing/splitting of punctured disks. Two punctured disks of the types $S_{d'+1}$ and $S_{d''+1}$, together with their perturbation 1-forms, are glued into a punctured disk of the type S_{d+1}, where $d + 1 = d' + d''$.

If it so happens that

$$\epsilon_{d+1} < \|K(\mathcal{P}_{d'+1})\| + \|K(\mathcal{P}_{d''+1})\|,$$

then the induction step will produce perturbation forms $K(\mathcal{P}_{d+1})$ that might not satisfy the condition $\|K(\mathcal{P}_{d+1})\| \leq \varepsilon_{d+1}$, which is required in order to have $A(u) > 0$ for all Floer polygons u. In other words, the construction described above does not go through at the induction step. Of course, the above argument shows that for every fixed d, we can choose consistent perturbation data such that all the Floer polygons involving no more than $d + 1$ intersection points have positive symplectic area. One could then easily modify this argument to handle more general clusters and consequently show that all the μ_k-operations in $\mathcal{F}uk(\mathcal{X})$ with $k \leq d$ will preserve filtrations. However, without any additional arguments, the above fails to prove that there is even one consistent perturbation data that will turn the μ_k-operations into filtration-preserving maps *for all* k. We will now refine the above arguments, showing how to achieve our goal by restricting further the type of perturbation data.

Recall that for all $\mathcal{L}$-decorated Floer polygons $u : S_{d+1} \to X$ we have the following energy-area identity:

$$E(u) = A(u) - \int_{S_{d+1}} R_{K(\mathcal{P}, S_{d+1})} \circ u, \tag{3.7}$$

where $E(u)$ is the energy of u and $R_{K(\mathcal{P},S_{d+1})}$ is the curvature associated to the 1-form $K(\mathcal{P}, S_{d+1})$. Recall that $R_{K(\mathcal{P},S_{d+1})}$ is a 2-form on S_{d+1} with values in the space of compactly supported functions on X. The expression $R_{K(\mathcal{P},S_{d+1})} \circ u$ is a real-valued 2-form on S_{d+1} obtained by composing the functions prescribed by the values of $R_{K(\mathcal{P},S_{d+1})}$ with the map u. We refer the reader to [53, Chapter II, Section (8g)] for more details on the definition of $R_{K(\mathcal{P},S_{d+1})}$ and the identity (3.7).

An important point about the curvature form $R_{K(\mathcal{P},S_{d+1})}$ is that it can be made arbitrarily small by choosing the perturbation form $K(\mathcal{P}, S_{d+1})$ to be small enough in an appropriate C^1-norm $\|-\|$. As alluded to above, $\|K(\mathcal{P}, S_{d+1})\|$ involves the sum of the L^1-norms of $K(\mathcal{P}, S_{d+1})$ and its derivatives, both in the direction of S_{d+1} and in the direction of the manifold X. (This norm can be viewed as a version of the Sobolev norm on $W^{1,1}$.) We therefore have

$$\left| \int_{S_{d+1}} R_{K(\mathcal{P},S_{d+1})} \circ v \right| \leq C_{d+1} \|K(\mathcal{P}, S_{d+1})\|$$

for all maps $v : S_{d+1} \to X$, where the constant C_{d+1} depends only on d (and not on the specific surface S_{d+1} or any other parameter involved in the integrand). Define

$$\begin{aligned} \|K(\mathcal{P}_{d+1})\| &:= \sup \|K(\mathcal{P}, S_{d+1})\|, \\ \nu(\mathcal{P}_{d+1}) &:= C_{d+1} \|K(\mathcal{P}_{d+1})\|, \end{aligned} \tag{3.8}$$

where the supremum in the first formula is taken over all $\bar{\mathcal{L}}$-decorated punctured disks S_{d+1}. Note that it is possible to choose the perturbation forms $K(\mathcal{P}, S_{d+1})$ such that they are all compactly supported inside the interior of each punctured disk S_{d+1}, and moreover we can control these supports so that the support of the entire family $K(\mathcal{P}_{d+1})$ is compact. (This poses no problems to having consistency with respect to gluing/splitting.) This implies that the supremum that appears in the definition of $\|K(\mathcal{P}_{d+1})\|$ in (3.8) can be assumed to be finite, and moreover can be made arbitrarily small by an appropriate choice of $K(\mathcal{P}_{d+1})$.

With the above notation we now have

$$A(u) \geq E(u) - \nu(\mathcal{P}_{d+1}) \tag{3.9}$$

for all $\bar{\mathcal{L}}$-decorated Floer polygons u.

We now add further restrictions to the perturbation data $\mathcal{P}$. Denote by $\bar{\mathcal{X}} = \{\bar{L} \mid L \in \mathcal{X}\}$ the set of all underlying Lagrangians from the collection $\mathcal{X}$. Denote by $I = \{p_1, \ldots, p_N\}$ the set of all intersection points between any two distinct Lagrangians $\bar{L}, \bar{L}' \in \bar{\mathcal{X}}$. By our assumption on $\bar{\mathcal{X}}$, I is a finite set. Moreover, every $p \in I$ corresponds to precisely one pair of distinct Lagrangians $\bar{L}, \bar{L}' \in \bar{\mathcal{X}}$ whose intersection contains p.

Denote by $B^{2n}(R) \subset \mathbb{R}^{2n}$ the closed $2n$-dimensional ball endowed with its standard symplectic structure, where $2n = \dim_{\mathbb{R}} X$. We claim that there exists a symplectic embedding of a disjoint union of N balls of some radius $R_{\overline{\mathcal{X}}} > 0$ into X,

$$\phi : \bigsqcup_{j=1}^{N} B^{2n}(R_{\overline{\mathcal{X}}}) \to X, \tag{3.10}$$

with the following properties:

- $\phi(0_j) = p_j$ for every j. Here and in what follows, we denote by B_j the jth ball in the disjoint union in (3.10), and by $0_j \in B_j$ the center of that ball.
- If $p_j \in \overline{L}' \cap \overline{L}''$ with $\overline{L}', \overline{L}'' \in \overline{\mathcal{X}}$ distinct, then

  $$\text{either } \phi(\mathbb{R}B_j) \subset \overline{L}',\ \phi(i\mathbb{R}B_j) \subset \overline{L}'',$$
  $$\text{or } \phi(\mathbb{R}B_j) \subset \overline{L}'',\ \phi(i\mathbb{R}B_j) \subset \overline{L}',$$

  where $\mathbb{R}B_j := B_j \cap (\mathbb{R}^n \times \{0\})$ is the "real" part of the ball B_j and $i\mathbb{R}B_j := B_j \cap (\{0\} \times \mathbb{R}^n)$ is its "imaginary" part. Moreover, $(\phi|_{B_j})^{-1}(\overline{L}' \cup \overline{L}'') \subset \mathbb{R}B_j \cup i\mathbb{R}B_j$, and $\phi(B_j)$ is disjoint from all the Lagrangians in $\overline{\mathcal{X}} \setminus \{\overline{L}', \overline{L}''\}$.

Clearly there exist an $R_{\overline{\mathcal{X}}} > 0$ and an embedding ϕ as above, and we fix both of them once and for all. Next we fix $0 < R < R_{\overline{\mathcal{X}}}$ and denote by $B_j(R) \subset B_j$ the smaller ball of radius R. Fix another parameter $0 < r \leq R$ and impose the following further restrictions on the perturbation data $\mathcal{P}_{d+1} = (K(\mathcal{P}_{d+1}), J(\mathcal{P}_{d+1}))$ for all d:

(1) $K|_{\phi(B_j(R))} \equiv 0$ for every j, i.e., the values of the 1-form K (which are compactly supported functions on X) vanish over the image of the restriction of ϕ to the smaller balls of radius R.

(2) $J_z|_{\phi(B_j(R))} = \phi_*(J_{\text{std}})$ for every j and every $z \in S_{d+1}$, $d \geq 2$. Here J_{std} is the standard complex structure on $B_j(R)$.

Note that the above additional restrictions on (K, J) do not contradict any of our previous assumptions on (K, J), and if we temporarily ignore the size of the perturbation forms K, then the above also do not pose any problems to the consistency of $\mathcal{P}$ with respect to gluing/splitting.

We now claim that there is a *consistent* choice of perturbation data $\mathcal{P}$ of the type described above such that for all Floer polygons u we have $A(u) > 0$.

To prove this, consider an $\overline{\mathcal{L}}$-decorated Floer polygon $u : S_{d+1} \to X$ associated with our perturbation data. Denote by $x_1, \dots, x_d, y$ the intersection points to which the punctures of S_{d+1} are mapped by u. By construction, u is genuinely holomorphic over image(ϕ) with respect to an almost complex structure that is diffeomorphic along that region to J_{std} via ϕ. Denote by $B_{j_1}, \dots, B_{j_d}, B_{j_{d+1}}$ the balls corresponding to the intersection points $x_1, \dots, x_d, y$ according to the construction made earlier. Applying a version of the monotonicity lemma (or, alternatively, a version of the

Lelong inequality) to u over $\phi(B_{j_k})$, $k = 1, \dots, d+1$, we obtain that there exists a constant C that does not depend on d (or on u or on any other parameter from the perturbation data) such that

$$E(u) \geq (d+1)CR^2. \tag{3.11}$$

Putting this together with (3.9) we obtain

$$A(u) \geq (d+1)CR^2 - \nu(\mathcal{P}_{d+1}). \tag{3.12}$$

Obviously, if we choose perturbation data $\mathcal{P}$ that are small enough such that

$$\nu(\mathcal{P}_{d+1}) < (d+1)Cr^2, \tag{3.13}$$

then $A(u) > 0$ for all Floer polygons u. (Recall that we fixed the parameter r such that $r \leq R$.) The main thing that needs to be verified now is that the condition (3.13) still enables a choice of perturbation data that are consistent with gluing/splitting. We address this point next.

We will choose the perturbation data $\mathcal{P}$ such that for all d we have $\nu(\mathcal{P}_{d+1}) < \alpha_{d+1}(d+1)Cr^2$ for some $0 < \alpha_{d+1} \leq 1$. Recall that punctured disks of the type S_{d+1} can split into two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, with $d' + d'' = d+1$ and $d', d'' \geq 2$. (Of course, splittings into more than two punctured disks are also possible; however, for the sake of obtaining a filtered A_∞-category, the top strata of the boundary of the space of disks matter, and these correspond to splitting into two punctured disks only.) And vice versa, by gluing two punctured disks of the type $S_{d'+1}$ and $S_{d''+1}$, we obtain punctured disks of the type $S_{d'+d''}$. Therefore, in order to make it possible to construct the perturbation data $\mathcal{P}_{d+1}$ by induction on d as indicated earlier, and to make them consistent with respect to gluing/splitting, we need to find a sequence of numbers α_{d+1}, $d \geq 2$, that satisfy the following set of inequalities:

$$\begin{aligned} &\alpha_{d'+1}(d'+1)Cr^2 + \alpha_{d''+1}(d''+1)Cr^2 \leq \alpha_{d'+d''}(d'+d'')Cr^2 \quad \forall\, d', d'' \geq 2, \\ &0 < \alpha_{d+1} \leq 1 \quad \forall\, d \geq 2. \end{aligned} \tag{3.14}$$

If such a sequence of numbers α_k does exist then we simply construct the perturbation data $\mathcal{P}_{d+1}$ by induction on d, where at each induction step we require that $\nu(\mathcal{P}_{d+1}) < \alpha_{d+1}(d+1)Cr^2$. The inequalities (3.14) will then ensure that the induction step goes through without posing problems to consistency with respect to gluing/splitting.

It remains to show that the inequalities in (3.14) admit solutions. Setting $\beta_k := k\alpha_k$ for every $k \geq 3$, the set of inequalities (3.14) can be simplified to

$$\begin{aligned} &\beta_{d'+1} + \beta_{d''+1} \leq \beta_{d'+d''} \quad \forall\, d', d'' \geq 2, \\ &0 < \beta_k \leq k \quad \forall\, k \geq 3. \end{aligned} \tag{3.15}$$

It is easy to see that this set of inequalities does have solutions. For example, for the sequence $\beta_k := (3k-6)B, k \geq 3$, where $B \leq \frac{1}{3}$, the first inequality in (3.15) becomes an equality and the second inequality is satisfied. One can also find sequences for which all inequalities in (3.15) become strict. This can be done as follows. Let β'_k be any sequence for which the first line in (3.15) is an equality and the second inequality holds (e.g., the preceding sequence $(3k-6)B$ with $B \leq \frac{1}{3}$). Let $\sigma : [3, \infty) \to \mathbb{R}$ be a strictly increasing function with $0 < \sigma(x) < 1$ for every x. Define now $\beta_k := \sigma(k)\beta'_k$ for all $k \geq 3$. A straightforward calculation shows that for this sequence all the inequalities in (3.15) become strict.

To conclude the proof regarding the preservation of filtrations of the μ_k-operations we need to address also the case of non-simple clusters of Floer polygons. The argument in this case is essentially the same, and below we will only outline it in the case of a cluster consisting of at most one Floer polygon with possibly several trees attached to its boundary.

Consider first a tuple $\bar{\mathcal{L}} = (\bar{L}_0, \dots, \bar{L}_d)$ in which not all the Lagrangians coincide. Fix a tuple $x_1, \dots, x_d, y$ of generators, with $x_i \in CF(L_{i-1}, L_i; \mathcal{D}_{L_{i-1}, L_i})$, $y \in CF(L_0, L_d; \mathcal{D}_{L_0, L_d})$. (Recall that some of these points are intersection points between the respective underlying Lagrangians, and some are critical points of the Morse functions prescribed by the Morse data.) Let Σ be an $\bar{\mathcal{L}}$-decorated cluster of punctured disks consisting of one punctured disk $S_{d'+1}$, where $d' \leq d$, and several trees. Let $u : \Sigma \to X$ be an $\bar{\mathcal{L}}$-decorated cluster of Floer polygons with entry points $x_1, \dots, x_d$ and exit point y. Denote by $u' := u|_{S_{d'+1}}$ the restriction of u to the underlying punctured disk $S_{d'+1}$ (which is a genuine Floer polygon with boundary conditions prescribed by a sub-tuple of $\bar{\mathcal{L}}$).

Recalling our filtration conventions for $CF(L', L'')$ in case $\bar{L}' = \bar{L}''$ (see at the end of Section 3.3.1), a simple calculation shows that

$$\sum_{i=1}^{d} \mathcal{A}(x_i) - \mathcal{A}(y) = A(u').$$

Now essentially the same argument as the one carried out earlier shows that the perturbation data can be chosen such that $A(u') \geq 0$.

Let us also consider the case when all the Lagrangians in $\bar{\mathcal{L}}$ coincide. In that case, a cluster of Floer polygons is just a collection of Morse trajectories modeled on a tree (without any actual polygons). A simple calculation shows that we have

$$\mathcal{A}(y) = \sum_{i=1}^{d} \mathcal{A}(x_i),$$

which implies that the filtration is preserved by the operations μ_d, $d \geq 2$, also in the case when the Lagrangians in $\bar{\mathcal{L}}$ all coincide.

As mentioned above, these arguments easily generalize to more complicated clusters, though the notation becomes more involved.

Transversality. We will briefly address now the topic of transversality. In order to show that $\mathcal{F}uk(\mathcal{X})$ is indeed an A_∞-category (i.e., that μ_d-operations satisfy the A_∞-identities, etc.), one needs to choose perturbation data $\mathcal{P}$ that satisfy various regularity properties. This would ensure that the spaces of clusters of Floer polygons involved in the definition of the μ_d-operations are smooth manifolds and have other desirable properties. Establishing the existence of regular perturbation data (or other auxiliary data) usually goes by the name "transversality", and is carried out via analytic techniques that have become standard in Floer theory. The typical result in this context is that the set of regular perturbation data is residual (in particular, dense) inside the space of all consistent perturbation data. However, in our case we work within a much more restricted space of perturbation data, as described above (e.g., specific choices of almost complex structures near the intersection points between pairs of distinct Lagrangians, perturbation forms that are compactly supported, etc.). Formally speaking, one would need to work out the transversality for our choices of perturbation data. While this does not formally follow from the general transversality theorems, it can still be achieved by rather standard arguments. For example, the fact that we restrict the almost complex structures to be constant on certain regions does not pose any problems (for achieving regularity), as long as the images of all the Floer polygons pass through regions of X in which we are allowed to perturb the almost complex structures without any restrictions. The same goes for the perturbation 1-forms K. In a similar vein, the fact that our perturbation forms K must be chosen to be small enough does not pose any transversality problem either. The only ingredient that requires slightly different transversality arguments is the part that uses the combination of Morse trees and Floer polygons. Transversality for spaces of Morse trees, as well as Morse trees mixed with holomorphic curves, has been worked out in various setups; see [15, 16, 32]. While we do not provide here details for these arguments, the case of clusters of Floer polygons follows from the source space description for clusters that appears in [15] and standard regularity arguments, as outlined above. Because all our Lagrangians are exact, the actual regularity arguments are much simpler compared to the ones developed in [15, 16].

The spaces $\mathcal{B}(\bar{\mathcal{X}})$. So far the proof above shows that there exist regular perturbation data $\mathcal{P}$ that turn $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ into a genuinely filtered, strictly unital, A_∞-category.

We now define the space $\mathcal{B}(\bar{\mathcal{X}})$. Recall that $R_{\bar{\mathcal{X}}}$ is the radius of the balls in the embedding ϕ. Given $r \leq R \leq R_{\bar{\mathcal{X}}}$, we denote by $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ the space of consistent, regular perturbation data $\mathcal{P}$, as defined earlier in the proof with the parameters r and R. The significance of the parameter R appears in conditions (1)–(2) on page 129 regarding the perturbation forms and almost complex structures on $\bigcup_j \phi(B_j(R))$. The parameter r plays a role in the size of the perturbations in condition (3.13) on page 130. To define $\mathcal{B}(\bar{\mathcal{X}})$ we specialize $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ to the case $r = R$, and take the union over all $R \leq R_{\bar{\mathcal{X}}}$. More precisely,

$$\mathcal{B}(\bar{\mathcal{X}}) := \bigcup_{0<R\leq R_{\bar{\mathcal{X}}}} \mathcal{B}(\bar{\mathcal{X}}; R, R).$$

This completes the outline of the proof of the second part of the first statement of Theorem 3.12. ■

Remark 3.14. (a) The reason for introducing the more general spaces $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ (instead of working with $r = R$ all the time) will become clear in Section 3.3.6 when we prove the results on the system of functors $\mathcal{F}^{\mathcal{P}_1,\mathcal{P}_0}$.

(b) If $R' \neq R''$ then neither $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ nor $\mathcal{B}(\bar{\mathcal{X}}; R'', R'')$ is a subspace of the other. In fact, if $R'' \leq R'$ then we have

$$\mathcal{B}(\bar{\mathcal{X}}; R', R') \cap \mathcal{B}(\bar{\mathcal{X}}; R'', R'') = \mathcal{B}(\bar{\mathcal{X}}; R', R'').$$

Moreover, if $R'' \leq R'$ and $r'' \geq r'$, then $\mathcal{B}(\bar{\mathcal{X}}; R', r') \subset \mathcal{B}(\bar{\mathcal{X}}; R'', r'')$.

(c) Obviously, the space $\mathcal{B}(\bar{\mathcal{X}})$ does not only depend on $\bar{\mathcal{X}}$ but also on the choice of the embedding ϕ, and there does not seem to be any preferred choice in this respect.

3.3.6 Proof of Theorem 3.12, part 2: Coherent systems

Here we will follow Seidel's approach to invariance of Fukaya categories based on coherent systems as described in [53, Chapter II, Section (10a)], but with several modifications needed to accommodate the filtered setting.

Before we go on, we briefly explain what is meant by a coherent system of *filtered* A_∞*-categories*. This is the filtered counterpart of the notion of coherent systems of A_∞-categories introduced in [53, Section (10a), pp. 133–135]. More specifically, let $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ be a family of filtered A_∞-categories. A coherent system consists of A_∞-functors $\mathcal{F}^{i_1,i_0} : \mathcal{A}^{i_0} \to \mathcal{A}^{i_1}$, defined for all $i_0, i_1 \in \mathcal{I}$, and with $\mathcal{F}^{i,i} = \mathbb{1}_{\mathcal{A}^i}$ for all i, as well as natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ for all $i_0, i_1, i_2 \in \mathcal{I}$. The functors $\mathcal{F}^{i,j}$ will be called comparison (or transition) functors. The functors $\mathcal{F}^{i,j}$ and natural transformations $T^{i,j,k}$ are required to satisfy a list of conditions explained in [53, Section (10a), p. 134]. In particular, the comparison functors $\mathcal{F}^{i_1,i_0}$ are all quasi-equivalences and the $T^{i,j,k}$ are quasi-isomorphisms. Turning to the filtered case, we require the following additional conditions to hold. All the functors $\mathcal{F}^{i_1,i_0}$ are required to be filtered (i.e., filtration preserving) and the T^{i_2,i_1,i_0} should be natural transformations of shift-0, i.e.,

$$T^{i_2,i_1,i_0} \in \hom^{\leq 0}_{\mathit{ffun}}(\mathcal{A}^{i_2}, \mathcal{A}^{i_0})(\mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0}, \mathcal{F}^{i_2,i_0}).$$

Here and in what follows, *ffun* is the (filtered) A_∞-category of filtered functors, and $\hom^{\leq 0}_{\mathit{ffun}}$ stands for the morphisms in that category that do not shift filtration, namely the natural transformations (between filtered functors) that preserve filtrations. Furthermore, all the cohomological identities from [53, Section (10a), p. 134] between these natural transformations should now hold in the 0-categories

$$(H^0(\mathit{ffun}(\mathcal{A}^{i'}, \mathcal{A}^{i''})))_0 \tag{3.16}$$

(i.e. persistence level 0) of the persistence homological categories $H^0(\mathit{ffun}(\mathcal{A}^{i'}, \mathcal{A}^{i''}))$. (In (3.16), the 0-superscript means cohomological degree 0 and the 0-subscript stands for the 0 persistence level subcategory.)

We will sometimes refer to $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ as a family of A_∞-categories over $\mathcal{I}$ and call $\mathcal{I}$ the base of the family. Similarly, in case we have a coherent system on $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$, we will call it a coherent system over $\mathcal{I}$.

One way to assemble a coherent system out of a family of A_∞-categories $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ is first to try to embed all of them into one total A_∞-category $\mathcal{A}^i \subset \mathcal{A}^{\text{tot}}$ by quasi-equivalences and then seek for suitable projection functors $\mathcal{A}^{\text{tot}} \to \mathcal{A}^j$ for all i. The functors $\mathcal{F}^{i_1,i_0} : \mathcal{A}^{i_0} \to \mathcal{A}^{i_1}$ participating in the coherent system will then be defined as the composition of the inclusions $\mathcal{A}^{i_0} \subset \mathcal{A}^{\text{tot}}$ with the projections $\mathcal{A}^{\text{tot}} \to \mathcal{A}^{i_1}$. We will soon adapt this scheme to the filtered framework. But first we need to introduce some relevant notions.

We begin with A_n-categories and functors. An A_n-category $\mathcal{A}$ is the same as an A_∞-category with the exception that now we have $\mu_k^{\mathcal{A}}$ only for $k = 1, \dots, n$. The $\mu_k^{\mathcal{A}}$'s are required to satisfy the subset of the A_∞-identities that involve only the $\mu_k^{\mathcal{A}}$'s with $1 \le k \le n$. In case the category $\mathcal{A}$ is clear from the context we will sometimes omit the superscript from $\mu_k^{\mathcal{A}}$.

Let $\mathcal{A}$ and $\mathcal{B}$ be two A_n-categories, and let $m \le n$. Similarly to A_∞-functors, we have A_m-functors $\mathcal{F} : \mathcal{A} \to \mathcal{B}$. They are defined in the same way as A_∞-functors, but now the higher-order components $\mathcal{F}_k$ of $\mathcal{F}$ are defined only for $1 \le k \le m$. The $\mathcal{F}_k$'s, $\mu_i^{\mathcal{A}}$'s, and $\mu_j^{\mathcal{B}}$'s are required to satisfy the same identities as for A_∞-functors, involving only indices with $1 \le k \le m$ and $i, j \le n$. A_n-functors can be composed in a similar way as A_∞-functors. Finally, the notions of pre-natural and natural transformations between A_∞-functors generalize to A_m-functors in a similar way.

Let $\mathcal{A}$ be an A_n-category, with $n \ge 3$. Let $X \in \text{Obj}(\mathcal{A})$ and $2 \le k \le n$. An element $e_X \in \hom_{\mathcal{A}}(X, X)$ is called a strict A_k-unit if it satisfies the following conditions: e_X is a cycle; for all $X' \in \text{Obj}(\mathcal{A})$, the maps

$$\mu_2(-, e_X) : \hom_{\mathcal{A}}(X', X) \to \hom_{\mathcal{A}}(X', X),$$
$$\mu_2(e_X, -) : \hom_{\mathcal{A}}(X, X') \to \hom_{\mathcal{A}}(X, X')$$

are the identity maps; and moreover $\mu_j(-, \dots, -, e_X, -, \dots, -) = 0$ for all $2 < j \le k$. An A_n-category is called strictly A_k-unital if we are given (as part of the structure) strict A_k-units e_X for every object $X \in \text{Obj}(\mathcal{A})$. If an A_n-category is strictly A_n-unital (i.e., $k = n$) we will simply say that it is strictly unital.

Let $\mathcal{A}$, $\mathcal{B}$ be two A_n-categories which are both strictly A_k-unital. An A_m-functor $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ is called strictly A_l-unital for $l \leq m$ if $\mathcal{F}_1(e_X) = e_{\mathcal{F}(X)}$ for all $X \in \mathrm{Obj}(\mathcal{A})$ and $\mathcal{F}_i(-, \ldots, -, e_X, -, \ldots, -) = 0$ for every $2 \leq i \leq l$.

Similarly to strict units, we also have the notion of strict isomorphisms. Let $\mathcal{A}$ be an A_n-category ($3 \leq n$) which is strictly A_l-unital ($l \leq n$), and denote its strict units by $e_Z \in \hom_{\mathcal{A}}(Z, Z)$, $Z \in \mathrm{Obj}(\mathcal{A})$. Let $X, Y \in \mathrm{Obj}(\mathcal{A})$ and let $u \in \hom_{\mathcal{A}}(X, Y)$ be a cycle. We say that u is a strict A_k-isomorphism (where $k \leq n$) if there exists a cycle $v \in \hom_{\mathcal{A}}(Y, X)$ such that

$$\mu_2(u, v) = e_X, \quad \mu_2(v, u) = e_Y, \quad \mu_j(-, \ldots, -, u_X, -, \ldots, -) = 0 \; \forall\, 2 < j \leq k.$$

If $k = n$, we will simply say that u is a strict isomorphism.

Remark 3.15. (a) In what follows, we will view A_∞-categories as a special case of A_n-categories by allowing $n = \infty$ (of course, one needs to slightly adjust the definition, since for $n = \infty$ the operations μ_k exist only for $1 \leq k < n$ and not for $1 \leq k \leq n$). The same remark applies also to A_n-functors, (pre-)natural transformations, and strict units, and we will view their A_∞-counterparts as a special case of the respective A_n-objects.

(b) A_n-categories are also $A_{n'}$-categories for all $1 \leq n' \leq n$; therefore, we will sometimes reduce n to the minimal value that is relevant in the context. A similar remark applies also to functors, (pre-)natural transformations, and strict units.

(c) Let $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ be an A_n-functor and $n' \leq n$. We denote by $\{\mathcal{F}\}_{n'}$ the $A_{n'}$-functor obtained from $\mathcal{F}$ by ignoring the terms of order greater than n' (i.e., the terms $\mathcal{F}_k$ for $n' < k \leq n$). We call $\{\mathcal{F}\}_{n'}$ the $A_{n'}$-reduction of $\mathcal{F}$.

In what follows, we will be mainly interested in A_2-functors $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ between two A_3-categories. Unwrapping the above definitions, in this case this means that $\mathcal{F}_1 : \hom_{\mathcal{A}}(X_0, X_1) \to \hom_{\mathcal{B}}(\mathcal{F} X_0, \mathcal{F} X_1)$ is a chain map and that $\mathcal{F}_1$ respects composition of morphisms up to a chain homotopy given by $\mathcal{F}_2$. In other words, $\mathcal{F}_2 : \hom_{\mathcal{A}}(X_0, X_1) \otimes \hom_{\mathcal{A}}(X_1, X_2) \to \hom_{\mathcal{A}}(X_0, X_2)$ defines a chain homotopy between $\mu_2 \circ (\mathcal{F}_1 \otimes \mathcal{F}_1)$ and $\mathcal{F}_1 \circ \mu_2$. Passing to homology, $\mathcal{F}$ induces a (non-unital) functor $H(\mathcal{F}) : H(\mathcal{A}) \to H(\mathcal{B})$ between the homological categories $H(\mathcal{A})$ and $H(\mathcal{B})$. Note that $H(\mathcal{A})$ and $H(\mathcal{B})$ are genuine categories, since $\mathcal{A}$ and $\mathcal{B}$ were assumed to be A_3-categories.

Assuming $\mathcal{A}$ and $\mathcal{B}$ to be A_2-unital, some of our A_2-functors $\mathcal{F}$ will be strictly A_1-unital, which means that $\mathcal{F}_1(e_X) = e_{\mathcal{F} X}$ for all objects X. In particular, this implies that $\mathcal{F}$ is homologically unital (i.e., the functor $H(\mathcal{F})$ is unital).

We now turn to the filtered setting. Filtered A_n-categories, A_m-functors and their pre-natural transformations are defined precisely in the same way as their filtered A_∞-counterparts. Strict A_k-units are required by definition to be in filtration level 0, and the same goes for strict isomorphisms. (Below we will not attach the adjective "filtered" to units/isomorphisms, implicitly assuming that whenever we are in the filtered setting these elements are in filtration level 0.)

If $\mathcal{A}$ is a filtered A_3-category which is strictly A_2-unital then $H(\mathcal{A})$ is a persistence category. If $\mathcal{F}$ is a filtered A_2-functor which is A_1-unital then $H(\mathcal{F})$ is a persistence functor.

Finally, we need the notion of a filtered quasi-equivalence. Let $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ be an A_2-functor between filtered A_3-categories. We call $\mathcal{F}$ a filtered quasi-equivalence if $\mathcal{F}$ is a filtered functor and its homological functor $H(\mathcal{F}) : H(\mathcal{A}) \to H(\mathcal{B})$ is an equivalence of persistence categories.

Next, we need to introduce persistence Hochschild cohomology. Let $\mathcal{A}$, $\mathcal{B}$ be A_3-categories which are strictly A_2-unital. Let $\mathcal{F}, \mathcal{G}$ be filtered A_2-functors which are A_1-unital. Given $X_0, \ldots, X_r \in \mathrm{Obj}(\mathcal{A})$, we will abbreviate $\mathcal{A}(X_0, \ldots, X_r) := \hom_{\mathcal{A}}(X_0, X_1) \otimes \cdots \otimes \hom_{\mathcal{A}}(X_{r-1}, X_r)$ and view it as a filtered chain complex in the standard way (and similarly for $\mathcal{B}$). The persistence Hochschild cochain complex associated to the above data is a cochain complex of persistence modules. The α-persistence level of this cochain complex in degree r is

$$\mathrm{PCC}^{r;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) := \prod_{X_0,\ldots,X_r} H^*\Big(\hom_{\mathbf{k}}^{\leq\alpha}\big(\mathcal{A}(X_0, \ldots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r)\big)\Big), \tag{3.17}$$

where the product runs over all r-tuples of objects in $\mathcal{A}$. Here we view

$$\hom_{\mathbf{k}}\big(\mathcal{A}(X_0, \ldots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r)\big)$$

as a filtered cochain complex, endowed with the standard differential, which will be denoted below by δ. The α-level filtration on this cochain complex, denoted here by $\hom_{\mathbf{k}}^{\leq\alpha}$, consists of those (graded) homomorphisms that shift filtration by at most α. Finally, H^* stands for persistence cohomology (in all degrees). Note that the spaces PCC are in fact bigraded since the $\hom_{\mathbf{k}}$ term is graded in itself. Thus a more detailed description of (3.17) would be to define $\mathrm{PCC}^{r,s}$ with the term H^* replaced by cohomology H^s in degree s. However, whenever it is not necessary we will ignore the s-degree.

The differential

$$\partial_{\mathrm{PCC}} : \mathrm{PCC}^{r;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) \to \mathrm{PCC}^{r+1;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G})$$

is defined as follows. Let $T \in \hom_{\mathbf{k}}^{\leq\alpha}(\mathcal{A}(X_0, \ldots, X_r), \mathcal{B}(\mathcal{F}X_0, \mathcal{G}X_r))$ be a δ-cocycle (i.e., a chain map). Then $\partial_{\mathrm{CC}}([T])$ is defined as the δ-cohomology class $[S]$ of the chain map

$$\begin{aligned} S(a_1,\dots,a_{r+1}) := {} & \epsilon' \mu_2(\mathcal{F}_1(a_1), T(a_2,\dots,a_{r+1})) \\ & + \epsilon'' \mu_2(T(a_1,\dots,a_r), \mathcal{G}_1(a_{r+1})) \\ & + \epsilon \sum_{k=0}^{r} T(a_1,\dots,\mu_2(a_{k+1},a_{k+2}),\dots,a_{r+1}), \end{aligned} \tag{3.18}$$

where $\epsilon', \epsilon'', \epsilon = \pm 1$ are signs that depend on the degrees of T, the a_i's, and r; see [53, Chapter I, Section (1f)]. Since we work over $\mathbf{k} = \mathbb{Z}_2$, these will play no role in our considerations anyway.

The cohomology of $(\mathrm{PCC}, \partial_{\mathrm{PCC}})$ is a graded persistence module, which we call the persistence Hochschild cohomology of $(\mathcal{A}, \mathcal{B})$ with respect to the functors $\mathcal{F}, \mathcal{G}$. We denote it by $\mathrm{PHH}^{*;\leq \alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G})$, $\alpha \in \mathbb{R}$.

Apart from PHH, there is a persistence variant of the classical Hochschild cohomology for the categories $H(\mathcal{A})$ and $H(\mathcal{B})$ with respect to the homological functors $[\mathcal{F}], [\mathcal{G}]$. This is defined in the same way as in [53, Chapter I, Section (1f)], only that we also take the persistence structure into account. We will use the following notation. We write

$$H^{\mathcal{A}}(X,Y) := H(\hom_{\mathcal{A}}(X,Y))$$

for the persistence homology (in all degrees) of $\hom_{\mathcal{A}}(X,Y)$ with respect to $\mu_1^{\mathcal{A}}$, and similarly for $\mathcal{B}$. Set also

$$H^{\mathcal{A}}(X_0,\dots,X_r) := H^{\mathcal{A}}(X_0,X_1) \otimes \cdots \otimes H^{\mathcal{A}}(X_{r-1},X_r),$$

viewed as a (graded) persistence module. Set $A := H(\mathcal{A})$ and $B := H(\mathcal{B})$, the homological categories of $\mathcal{A}$ and $\mathcal{B}$, and set $F := [\mathcal{F}]$ and $G := [\mathcal{G}]$ to be the homological functors corresponding to $\mathcal{F}$ and $\mathcal{G}$. The cochain complex $\mathrm{P}\overline{\mathrm{CC}}$ for the Hochschild cohomology of $H(\mathcal{A})$ and $H(\mathcal{B})$ has the following persistence module structure in degree r:

$$\mathrm{P}\overline{\mathrm{CC}}^{r;\leq\alpha}(A,B;F,G) := \prod_{X_0,\dots,X_r} \Big(\hom_{\mathrm{per}}^{\leq \alpha}\big(H^{\mathcal{A}}(X_0,\dots,X_r), H^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)\big)\Big), \tag{3.19}$$

where α denotes the persistence level and $\hom_{\mathrm{per}}^{\leq\alpha}$ stands for (graded) homomorphisms from the persistence module $H^{\mathcal{A}}(X_0,\dots,X_r)$ to the persistence module $H^{\mathcal{B}}(\mathcal{F}X_0,\mathcal{G}X_r)[\alpha]$. We are using the notation and conventions from Section 2.2.4 (see (2.15)). Namely, $H^{\mathcal{B}}(\mathcal{F}X_0,\mathcal{G}X_r)[\alpha]$ stands for $H^{\mathcal{B}}(\mathcal{F}X_0,\mathcal{G}X_r)$ shifted by α in terms of the persistence parameter. The structure map from level $\hom_{\mathrm{per}}^{\leq\alpha}$ to level $\hom_{\mathrm{per}}^{\leq\beta}$, whenever $\alpha \leq \beta$, is given by composing with the structure map indexed by α, β in the persistence module $H^{\mathcal{B}}(\mathcal{F}X_0,\mathcal{G}X_r)$. The differential $\partial_{\overline{\mathrm{CC}}}$ has a similar expression to (3.18), and an explicit formula can be found in [53, Chapter I, Section (1f), p. 13]. We denote the resulting persistence cohomology by $\mathrm{P}\overline{\mathrm{HH}}^{*;\leq\alpha}(A, B; F, G)$, $\alpha \in \mathbb{R}$.

Remark 3.16. (a) It is easy to see that the data required to define the second version, $\mathrm{P}\overline{\mathrm{HH}}$, of Hochschild cohomology (and in fact even $\mathrm{P}\overline{\mathrm{CC}}$) is entirely homological. In fact, this cohomology can be defined for every pair of persistence categories and a pair of persistence functors between them. In contrast, PCC and PHH seem to depend on some chain level information from $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{F}, \mathcal{G}$.

The two cochain complexes PCC and $\mathrm{P}\overline{\mathrm{CC}}$ are related to each other by a Künneth-type short exact sequence of persistence modules:

$$0 \longrightarrow E_{r-1}(H(\mathcal{A}), H(\mathcal{B})) \longrightarrow \mathrm{PCC}^r(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}) \longrightarrow \mathrm{P}\overline{\mathrm{CC}}^r(A, B; F, G) \longrightarrow 0, \tag{3.20}$$

where $E_{r-1}(-,-)$ is a derived functor of $\hom_{\mathrm{per}}$ (which is analogous to Ext^1) in the category of persistence modules. More specifically, when inserting the components of $X_0, \dots, X_r$ from (3.17) and (3.19) into (3.20), the term E_{r-1} involves only homologies of the type $H_*(\mathcal{A}(X_0, \dots, X_r))$ and $H_*^{\mathcal{B}}(\mathcal{F}X_0, \mathcal{G}X_r)$ of total degree $r-1$. The former term $H_*(\mathcal{A}(X_0, \dots, X_r))$ can also be related, via short exact sequences (a persistence analog of the universal coefficients theorem), to $H_*^{\mathcal{A}}(X_0, X_1) \otimes \cdots \otimes H_*^{\mathcal{A}}(X_{r-1}, X_r)$ and Tor_1-like derived functors associated to tensor products of persistence modules. We refer the reader to [13, Section 8] for the precise formulation of the Künneth and universal coefficients theorems for persistence homology (see also [49]).

(b) Let $3 \leq n \leq \infty$, $m \leq n$. Given two A_n-categories $\mathcal{A}$, $\mathcal{B}$ and two A_m-functors $\mathcal{F}, \mathcal{G} : \mathcal{A} \to \mathcal{B}$, the filtered natural transformations $\hom_{\mathit{ffun}(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ form a filtered chain complex (by filtered natural transformations we mean those that respect filtrations up to a bounded shift). The level-α filtration $\hom^{\leq\alpha}_{\mathit{ffun}(\mathcal{A},\mathcal{B})}$ consists of those natural transformations that shift filtration by at most α. Apart from this filtration, this space admits yet another filtration called the length filtration. This one is indexed by the natural numbers and is decreasing. Its p-level $F^p \hom^{\leq\alpha}_{\mathit{ffun}(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ is given by those natural transformations $T \in \hom^{\leq\alpha}_{\mathit{ffun}(\mathcal{A},\mathcal{B})}(\mathcal{F}, \mathcal{G})$ with $T_0 = \cdots = T_p = 0$. There is a spectral sequence of persistence modules associated to this filtration. A simple calculation shows that its first page, at persistence level α, is given by

$$E_1^{r,s;\leq\alpha} = \mathrm{PCC}^{r,s;\leq\alpha}(\mathcal{A}, \mathcal{B}; \mathcal{F}, \mathcal{G}).$$

Here we have used the second degree s on PCC as briefly explained earlier. See [53, Chapter I, Section (1f), p. 13] for more details in the non-filtered case. The filtration F^p is bounded in the case of A_m-functors with m finite. For A_∞-functors it is not bounded, and its associated spectral sequence might not converge; however, following [53] we will use this sequence as a tool for comparing the homologies of different $\hom_{\mathit{ffun}}$'s.

Let $\mathcal{A}, \mathcal{B}, \mathcal{C}$ be filtered A_n-categories ($3 \le n \le \infty$). Let $2 \le m \le n$, and denote by *nu-ffun*$(\mathcal{C}, \mathcal{A})$ and *nu-ffun*$(\mathcal{C}, \mathcal{B})$ the filtered A_m-categories of non-unital (or, better said, not necessarily unital) filtered A_m-functors $\mathcal{C} \to \mathcal{A}$ and $\mathcal{C} \to \mathcal{B}$, respectively. Let $\mathcal{G} : \mathcal{A} \to \mathcal{B}$ be a non-unital filtered A_m-functor and denote by

$$\mathcal{L}_{\mathcal{G}} : \textit{nu-ffun}(\mathcal{C}, \mathcal{A}) \to \textit{nu-ffun}(\mathcal{C}, \mathcal{B})$$

the functor induced by left composition with $\mathcal{G}$. Note that this is a (non-unital) filtered A_m-functor. An immediate consequence of Remark 3.16 is the following.

Lemma 3.17 (cf. [53, Lemma 1.7]). *If $\mathcal{G}$ is homologically full and faithful in the persistence sense, then so is $\mathcal{L}_{\mathcal{G}}$.*

The next lemma deals with invariance of persistence Hochschild cohomology under filtered quasi-equivalences. It is a persistence analog of a very special case of [53, Lemma 2.6], with several additional very strong assumptions made in order to accommodate the persistence case.

Lemma 3.18. *Let $\mathcal{A}$ be a filtered A_n-category ($3 \le n \le \infty$) with strict A_3-units, and let $\widetilde{\mathcal{A}} \subset \mathcal{A}$ be a full subcategory such that the inclusion $\mathcal{J} : \widetilde{\mathcal{A}} \to \mathcal{A}$ is a filtered quasi-equivalence. Suppose that $\mathcal{P} : \mathcal{A} \to \widetilde{\mathcal{A}}$ is a filtered A_2-functor which is A_1-unital, and assume that $\mathcal{P} \circ \mathcal{J} = \mathbb{1}_{\widetilde{\mathcal{A}}}$ (as A_2-functors). Assume further that for every $X \in \mathrm{Obj}(\mathcal{A})$ we have a strict isomorphism $u_X \in \hom_{\mathcal{A}}(X, \mathcal{P}(X))$. Then the map induced by the restriction*

$$\rho : \mathrm{PHH}(\mathcal{A}, \widetilde{\mathcal{A}}; \mathcal{P}, \mathcal{P}) \to \mathrm{PHH}(\widetilde{\mathcal{A}}, \widetilde{\mathcal{A}}; \mathbb{1}_{\widetilde{\mathcal{A}}}, \mathbb{1}_{\widetilde{\mathcal{A}}})$$

is a (bigraded) isomorphism of persistence modules.

We omit the proof since it is very similar to that of [53, Lemma 2.6]. The role of the strict isomorphisms u_X is to facilitate the definition, in the framework of PCC and PHH, of certain chain maps and chain homotopies that appear in the original proof. Note also that [53, Lemma 2.6] is stated for any two functors; here we only need it for the functor $\mathcal{P}$, which simplifies things further.

We now get to extending filtered A_n-functors from a subcategory to a larger one. The following lemma is a persistence analog, this time of a special case of [53, Lemma 1.10], and again with several additional assumptions.

Lemma 3.19. *Let $\mathcal{A}$, $\mathcal{B}$ be filtered A_n-categories ($3 \le n \le \infty$) and $\widetilde{\mathcal{A}} \subset \mathcal{A}$ a full subcategory. Let $\mathcal{P} : \mathcal{A} \to \mathcal{B}$ be a filtered A_2-functor and denote by $\widetilde{\mathcal{P}} := \mathcal{P}|_{\widetilde{\mathcal{A}}}$ its restriction to $\widetilde{\mathcal{A}}$. Assume that the map*

$$\rho : \mathrm{PHH}^r(\mathcal{A}, \mathcal{B}; \mathcal{P}, \mathcal{P}) \to \mathrm{PHH}^r(\widetilde{\mathcal{A}}, \mathcal{B}; \widetilde{\mathcal{P}}, \widetilde{\mathcal{P}})$$

induced by the restriction is an isomorphism of persistence modules for every r. Then every filtered A_n-functor $\widetilde{\mathcal{Q}} : \widetilde{\mathcal{A}} \to \mathcal{B}$ with $\{\widetilde{\mathcal{Q}}\}_2 = \widetilde{\mathcal{P}}$ can be extended to a filtered A_n-functor $\mathcal{Q} : \mathcal{A} \to \mathcal{B}$ with $\{\mathcal{Q}\}_2 = \mathcal{P}$. (See point (c) *of Remark 3.15.)*

We omit the proof again since it is very similar to that of [53, Lemma 1.10], with straightforward modifications needed for the persistence setting.

We are now ready to assemble a coherent system of filtered A_∞-categories out of a family of categories that are all included into one total category. Let $\{\mathcal{A}^i\}_{i\in\mathcal{I}}$ be a family of filtered, strictly unital A_∞-categories over $\mathcal{I}$. Suppose there is a filtered strictly unital A_∞-category $\mathcal{A}^{\mathrm{tot}}$ such that for every $i \in \mathcal{I}$, $\mathcal{A}^i$ is a full subcategory of $\mathcal{A}^{\mathrm{tot}}$. Denote by $\mathcal{J}^i : \mathcal{A}^i \to \mathcal{A}^{\mathrm{tot}}$ the inclusion functor and assume that $\mathcal{J}^i$ is a filtered quasi-equivalence. We will refer to a category $\mathcal{A}^{\mathrm{tot}}$ as above, together with the inclusion functors $\mathcal{J}^i$, as a (filtered, strictly unital) *total A_∞-category over $\mathcal{I}$*.

Assume that for every $i \in \mathcal{I}$ there is a filtered A_2-functor $\mathcal{P}r^i : \mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^i$ which is strictly A_1-unital and such that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{A}^i}$ as A_2-functors. Assume further that the following holds for every $i \in \mathcal{I}$: for every $X \in \mathrm{Obj}(\mathcal{A}^{\mathrm{tot}})$ there exists a strict isomorphism $u^i_X \in \hom_{\mathcal{A}^{\mathrm{tot}}}(X, \mathcal{P}r^i(X))$.

Proposition 3.20. *Under the above assumptions each of the A_2-functors $\mathcal{P}r^i$, $i \in \mathcal{I}$, can be extended to a filtered A_1-unital A_∞-functor $\mathcal{Q}^i : \mathcal{A}^{\mathrm{tot}} \to \mathcal{A}^i$ (i.e., $\{\mathcal{Q}^i\}_2 = \mathcal{P}r^i$). Moreover, the functors $\mathcal{F}^{j,i} := \mathcal{Q}^j \circ \mathcal{J}^i : \mathcal{A}^i \to \mathcal{A}^j$, $i, j \in \mathcal{I}$, are filtered, A_1-unital, and form a coherent system of filtered A_∞-categories over $\mathcal{I}$. The filtered natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ will be described in the proof.*

Proof. The restriction of the A_2-functors $\mathcal{P}r^i$ to the subcategory $\mathcal{A}^i \subset \mathcal{A}^{\mathrm{tot}}$ is the A_2-reduction of the identity A_∞-functor on $\mathcal{A}^i$. Using Lemmas 3.18 and 3.19 we can extend $\mathcal{P}r^i$ to the desired A_∞-functor $\mathcal{Q}^i$.

The construction of the natural transformations $T^{i_2,i_1,i_0} : \mathcal{F}^{i_2,i_1} \circ \mathcal{F}^{i_1,i_0} \to \mathcal{F}^{i_2,i_0}$ follows the same scheme as in [53, Chapter II, Section (10a), p. 134].

Consider the homological functor induced by left composition with $\mathcal{Q}^i$, viewed as a persistence functor:

$$H(\mathcal{L}_{\mathcal{Q}^i}) : H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^{\mathrm{tot}})) \to H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^i)).$$

Recall that $\mathcal{Q}^i$ is a filtered quasi-equivalence and, in particular, homologically full and faithful in the persistence sense. By Lemma 3.17, the action of the functor $H(\mathcal{L}_{\mathcal{Q}^i})$ on morphisms is an isomorphism of persistence modules:

$$H(\mathcal{L}_{\mathcal{Q}^i}) : \hom_{H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^{\mathrm{tot}}))}(\mathcal{J}^i \circ \mathcal{Q}^i, \mathbb{1}_{\mathcal{A}^{\mathrm{tot}}}) \to \hom_{H(\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^i))}(\mathcal{Q}^i, \mathcal{Q}^i).$$

Let $S^i \in \hom^{\leq 0}_{\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^{\mathrm{tot}})}(\mathcal{J}^i \circ \mathcal{Q}^i, \mathbb{1}_{\mathcal{A}^{\mathrm{tot}}})$ be a cycle whose homology class $[S^i]$ is sent by $H(\mathcal{L}_{\mathcal{Q}^i})$ to $[\mathbb{1}] \in H(\hom_{\mathit{ffun}(\mathcal{A}^{\mathrm{tot}},\mathcal{A}^i)}(\mathcal{Q}^i, \mathcal{Q}^i))$. Having defined S^i as above for all $i \in \mathcal{I}$, we define

$$T^{i_2,i_1,i_0} := \mathcal{L}_{\mathcal{Q}^{i_2}}(\mathcal{R}_{\mathcal{J}^{i_0}}(S^{i_1})),$$

where $\mathcal{R}_{\mathcal{J}^{i_0}}$ is the functor $\mathit{ffun}(\mathcal{A}^{\mathrm{tot}}, \mathcal{A}^{\mathrm{tot}}) \to \mathit{ffun}(\mathcal{A}^{i_0}, \mathcal{A}^{\mathrm{tot}})$ induced by right composition with $\mathcal{J}^{i_0}$. ■

3.3.7 Proof of Theorem 3.12, part 3: Coherent systems of Fukaya categories

In order to apply the algebraic statements from Section 3.3.6, in particular Proposition 3.20, to the case of Fukaya categories, we need two additional ingredients coming from geometry. First, we need to construct a filtered total Fukaya category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$ that contains all the Fukaya categories $\mathcal{F}uk(\mathcal{X}; i)$, constructed via various perturbation data i, as quasi-equivalent subcategories of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$. Second, we need to construct filtered A_2-functors $\mathcal{P}r^i : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}) \to \mathcal{F}uk(\mathcal{X}; i)$ that are left inverses of the inclusions $\mathcal{F}uk(\mathcal{X}; i) \subset \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X})$.

We begin with the first ingredient, namely the construction of $\mathcal{F}uk^{\mathrm{tot}}$. To this end, recall the spaces of perturbation data $\mathcal{B}(\bar{\mathcal{X}}; R, r)$ from page 132. We fix the symplectic embedding ϕ and the parameters $R_{\bar{\mathcal{X}}}$ and $r \leq R \leq R_{\bar{\mathcal{X}}}$. We will outline now a construction of a filtered strictly unital total A_∞-category over $\mathcal{B}(\bar{\mathcal{X}}; R, r)$. (See page 140 for the meaning of a total category over a base.) More specifically, we will construct a filtered strictly unital A_∞-category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ with the following property. For every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$, the filtered Fukaya category $\mathcal{F}uk(\mathcal{X}; i)$ is a full subcategory of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ and the inclusion $\mathcal{J}^i : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$ is a filtered quasi-equivalence.

The construction of $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r)$ follows steps similar to those of the construction introduced in [53, Chapter II, Section (10a), pp. 134–135], with some significant modifications necessary to accommodate the filtered setting. The objects of $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r)$ are pairs (L, i), where $L \in \mathcal{X}$ and $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$ is a choice of admissible perturbation data. The morphism space between (L_0, i_0) and (L_1, i_1) is defined to be the Floer complex $\mathrm{CF}(L_0, L_1; \mathcal{D}_{(L_0,i_0),(L_1,i_1)})$, where the Floer datum $\mathcal{D}$ is defined as in our earlier construction of filtered Fukaya categories, with the only restriction that if $i_0 = i_1$ then the Floer datum $\mathcal{D}_{(L_0,i_0),(L_1,i_1)}$ should agree with that of (L_0, L_1) in the category $\mathcal{F}uk(\mathcal{X}; i)$. Another important point is that we require the Floer data for pairs of the type $((L, i), (L, j))$ to continue to be of the same type as in our construction of filtered Fukaya categories. Namely, we take here a pair of a Morse function and a Riemannian metric, such that the Morse function has a unique local maximum (i.e., a unique critical point of index $n = \dim_{\mathbb{C}} X$).

The next step in the construction of the total category is to choose consistent perturbation data $\mathcal{P}^{\mathrm{tot}} = \mathcal{P}^{\mathrm{tot}}(\mathcal{X}; R, r)$ with the restriction analogous to the one imposed on the Floer data. Namely, whenever we have a cluster of punctured disks decorated entirely by pairs (L_k, i) with the same i, the value of $\mathcal{P}^{\mathrm{tot}}$ on such a cluster coincides with the one prescribed by the perturbation data i. Apart from that, there will be one important difference from the way we defined the perturbation data for each $\mathcal{F}uk(\mathcal{X}; i)$, which we now describe. Given a tuple $(i_0, \ldots, i_d)$ with $i_k \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$, we define the following two quantities:

$$\begin{aligned} R^{(i_0,\ldots,i_d)} &:= \sup\{\widetilde{R} \mid \widetilde{R} \le R_{\overline{\mathcal{X}}} \text{ and } i_k \in \mathcal{B}(\overline{\mathcal{X}}; \widetilde{R}, \widetilde{R}) \ \forall\, 0 \le k \le d\}, \\ r^{(i_0,\ldots,i_d)} &:= \inf\{\widetilde{r} \mid 0 \le \widetilde{r} \text{ and } i_k \in \mathcal{B}(\overline{\mathcal{X}}; \widetilde{r}, \widetilde{r}) \ \forall\, 0 \le k \le d\}. \end{aligned} \tag{3.21}$$

In other words, $R^{(i_0,\ldots,i_d)}$ measures the supremal radius of the sub-balls in the embedding ϕ on which the almost complex structures in all the perturbation data i_k are standard. The other quantity, $r^{(i_0,\ldots,i_d)}$, measures the infimal upper bound on the perturbation 1-forms in all the perturbation data i_k. Note that since $i_k \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$ for all k, we have

$$r^{(i_0,\ldots,i_d)} \le r \le R \le R^{(i_0,\ldots,i_d)}. \tag{3.22}$$

Turning to the definition of the perturbation data $\mathcal{P}^{\mathrm{tot}}$, we require it to satisfy the following conditions. Let S_{d+1} be a $(d+1)$-punctured disk, decorated by the tuple $((L_0, i_0), \ldots, (L_d, i_d))$, and denote by (K, J) the value of the perturbation data on S_{d+1}. Recall the symplectic embedding ϕ from (3.10) and the balls B_j (see page 129). We require that:

(1) $K|_{\phi(B_j(R^{(i_0,\ldots,i_d)}))} \equiv 0$ for every j. Here $B_j(R^{(i_0,\ldots,i_d)}) \subset B_j$ denotes the smaller ball of radius $R^{(i_0,\ldots,i_d)}$.

(2) $J_z|_{\phi(B_j(R^{(i_0,\ldots,i_d)}))} = \phi_*(J_{\mathrm{std}})$ for every j and every $z \in S_{d+1}$.

In addition to the above two conditions, we also require that

$$\nu(\mathcal{P}^{\mathrm{tot}}((L_0, i_0), \ldots, (L_d, i_d))) < C(d+1)(r^{(i_0,\ldots,i_d)})^2. \tag{3.23}$$

Here C is the constant from (3.11), and, similarly to (3.8),

$$\nu(\mathcal{P}^{\mathrm{tot}}((L_0, i_0), \ldots, (L_d, i_d))) := C_{d+1} \sup \|K(\mathcal{P}^{\mathrm{tot}}, S_{d+1})\|,$$

where now the supremum goes over all $(d+1)$-punctured disks S_{d+1} that are decorated by $((L_0, i_0), \ldots, (L_d, i_d))$. Finally, the above requirements extend to clusters of punctured disks in a similar way.

We claim that there exists a consistent choice of perturbation data $\mathcal{P}^{\mathrm{tot}}$ satisfying the above conditions. The proof is similar to that of the corresponding statement for the perturbation data $i \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$.

With a choice of perturbation data as above, one can define an A_∞-category in the same way as we defined our earlier Fukaya categories. We denote this category by $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; R, r)$, or sometimes by $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; \mathcal{P}^{\mathrm{tot}})$ when we want to emphasize the choice of the perturbation data $\mathcal{P}^{\mathrm{tot}}$ used to define it.

We claim that $\mathcal{F}uk^{\mathrm{tot}}(\overline{\mathcal{X}}; R, r)$ is a filtered and strictly unital A_∞-category. This follows by the same arguments we used for $\mathcal{F}uk(\mathcal{X}; i)$, $i \in \mathcal{B}(\overline{\mathcal{X}}; R, r)$, with minor modifications. The important points are that the analogs of inequalities (3.11), (3.12), and (3.13) will continue to hold with R and r replaced by $R^{(i_0,\ldots,i_d)}$ and $r^{(i_0,\ldots,i_d)}$

respectively, and $\nu(\mathcal{P}_{d+1})$ by $\nu(\mathcal{P}^{\mathrm{tot}}((L_0, i_0), \ldots, (L_d, i_d)))$. This completes the outline of the construction of the category $\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r)$.

For every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$ there is an obvious inclusion

$$\mathcal{J}^i : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r).$$

Clearly, this functor is filtered, and we claim that it is a filtered quasi-equivalence. To see the latter, first note that by construction $\mathcal{J}^i$ is full and faithful. Now, any object $(L, k) \in \mathrm{Obj}(\mathcal{F}uk^{\mathrm{tot}}(\bar{\mathcal{X}}; R, r))$ is isomorphic to (L, i) via an isomorphism that lies in $\hom^{\leq 0}$ of $\mathcal{F}uk^{\mathrm{tot}}$. This follows from Morse theory, since the hom between (L, k) and (L, i) is the Morse complex of L with respect to a Morse function with a unique local maximum. This shows that $\mathcal{J}^i$ is a filtered quasi-equivalence.

Next, for every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$, we construct a filtered A_2-functor

$$\mathcal{P}r^i : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r) \to \mathcal{F}uk(\mathcal{X}; i),$$

which is strictly A_1-unital and satisfies $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$.

Let $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$. The construction of $\mathcal{P}r^i$ goes as follows. Let (L, k) be an object of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, r)$, where $L \in \mathcal{X}$, $k \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$. We define $\mathcal{P}r^i((L, k)) = L$.

Next, we define the first-order part $\mathcal{P}r^i_1$ of $\mathcal{P}r^i$ on morphisms. This can be done by means of *Floer continuation maps*. Specifically, we need to define a filtration-preserving chain map

$$\mathcal{P}r^i_1 : \mathrm{CF}((L', k'), (L'', k''); \mathcal{D}_{(L',k'),(L'',k'')}) \to \mathrm{CF}(L', L''; \mathcal{D}_{L',L'',i}) \tag{3.24}$$

for every $L', L'' \in \mathcal{X}$, $k', k'' \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$, where $\mathcal{D}_{(L',k'),(L'',k'')}$ is the Floer data for the pair $((L', k'), (L'', k''))$ and $\mathcal{D}_{L',L'',i}$ is the one used in $\mathcal{F}uk(\mathcal{X}; i)$. Assume first that $\bar{L}' \neq \bar{L}''$ (which means that $\bar{L}' \pitchfork \bar{L}''$). By construction, the Hamiltonian terms in both Floer data $\mathcal{D}_{(L',k'),(L'',k'')}$ and $\mathcal{D}_{L',L'',i}$ are identically 0 (so the two Floer data may differ only in their almost complex structures). Denote by $J(\mathcal{P}^{\mathrm{tot}}; (L', k', L'', k''))$ the almost complex structure of $\mathcal{D}_{(L',k'),(L'',k'')}$ and by $J(i; L', L'')$ the one corresponding to $\mathcal{D}_{L',L'',i}$. Fix a generic homotopy J^{cont}_s, $s \in [0, 1]$, between these two almost complex structures (of the type admissible in $\mathcal{B}(\bar{\mathcal{X}}; R, r)$). We assume that J^{cont}_s coincides with $J(\mathcal{P}^{\mathrm{tot}}; (L', k', L'', k''))$ near $s = 0$ and with $J(i; L', L'')$ near $s = 1$. Extend this homotopy to $s \in \mathbb{R}$ by keeping it constant with respect to the s-parameter outside $[0, 1]$. Recall that, by construction, both Floer data $\mathcal{D}_{(L',k'),(L'',k'')}$ and $\mathcal{D}_{L',L'',i}$ have 0 Hamiltonian terms, hence $(0, J^{\mathrm{cont}}_s)$ defines a homotopy between the latter two Floer data. Standard Floer theory associates to this homotopy a filtration-preserving quasi-isomorphism, as claimed in (3.24), called *the Floer continuation map*. That $\mathcal{P}r^i_1$ preserves filtrations follows from standard arguments in Floer theory, using the assumption that the Hamiltonian terms in the preceding homotopy of Floer data are 0 for all times s.

For further use, we will add one more restriction on the definition of the continuation maps $\mathcal{P}r^i_1$. In case the two objects (L',k') and (L'',k'') correspond to the same perturbation data i, i.e., $k'=k''=i$, we will take the homotopy J^{cont}_s to be constant. As a result, the continuation map

$$\mathcal{P}r^i_1 : \mathrm{CF}((L',i),(L'',i);\mathcal{D}_{(L',i),(L'',i)}) \to \mathrm{CF}(L',L'';\mathcal{D}_{L',L'',i})$$

for such pairs will be the identity.

We now briefly address the case when $\overline{L}'=\overline{L}''$. Recall that in this case the Floer data on each of $\mathcal{D}_{(L',k'),(L'',k'')}$, $\mathcal{D}_{L',L'',i}$ consists of a Morse function and a Riemannian metric on $\overline{L}'$. In our model the corresponding CF's are just the Morse complexes on $\overline{L}'$ associated to these data. The map $\mathcal{P}r^i_1$ is now defined by means of standard Morse homology theory – it is just the continuation map between the two Morse complexes. The fact that $\mathcal{P}r^i_1$ is filtration preserving is automatic since, by definition, both Morse complexes $\mathrm{CF}((L',k'),(L'',k'');\mathcal{D}_{(L',k'),(L'',k'')})$ and $\mathrm{CF}(L',L'';\mathcal{D}_{L',L'',i})$ are concentrated at the same filtration level (which is a constant that depends on the difference between the primitives of the Liouville forms on L' and L'').

Similarly to the case $\overline{L}'\neq\overline{L}''$, here too we can arrange $\mathcal{P}r^i_1$ to be the identity map whenever $k'=k''=i$. This can be done by taking the homotopy between the two Morse data to be constant, and the resulting Morse continuation map will then be the identity.

Next we define the second-order part of $\mathcal{P}r^i$. This will be a map

$$\mathcal{P}r^i_2 : \mathrm{CF}((L_0,k_0),(L_1,k_1);\mathcal{D}_{\mathrm{tot}}) \otimes \mathrm{CF}((L_1,k_1),(L_2,k_2);\mathcal{D}_{\mathrm{tot}}) \to \mathrm{CF}(L_0,L_2;\mathcal{D}_i) \tag{3.25}$$

of cohomological degree -1. Here we have written $\mathcal{D}_{\mathrm{tot}}$ and $\mathcal{D}_i$ for the Floer data (for the corresponding pairs of Lagrangians) in the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$ and $\mathcal{F}uk(\mathcal{X};i)$, respectively.

Assume for simplicity that $\overline{L}_0$, $\overline{L}_1$, $\overline{L}_2$ are all distinct. To define the map (3.25) we will first need to introduce some new spaces of Floer-type polygons. Recall the homotopy J^{cont}_s from the definition of $\mathcal{P}r^i_1$ above. Below we will need a more precise notation, and from now on we will denote it by $J^{\mathrm{cont}}_{s,t}((L',k'),(L'',k''),(L',L'',i))$, where $s\in\mathbb{R}$ and $t\in[0,1]$. (Recall that all our almost complex structures are possibly time-dependent, and we denote here by t the time parameter.)

Denote by S_3 the 3-punctured disk. Recall that S_3 has two "entry" strip-like ends $St^-_{0,1}$, $St^-_{1,2}$ and one "exit" strip-like end $St^+_{0,2}$. We order them in the clockwise direction, $St^-_{0,1}$, $St^-_{1,2}$, $St^+_{0,2}$, according to the punctures they correspond to. We denote by (s,t) the coordinates on each of these strip-like ends. Thus $(s,t)\in(-\infty,0]\times[0,1]$ for $St^-_{0,1}$, $St^-_{1,2}$, and $(s,t)\in[0,\infty)\times[0,1]$ for $St^+_{0,2}$. We also fix a smooth positive decreasing function $A:(0,\delta)\to\mathbb{R}$ in a small neighborhood of 0 with $A(\tau)\to\infty$ as $\tau\to 0^+$.

Consider now a 1-parametric family $\mathcal{P}^{\tau}(\mathscr{Pr}) = (K^{\tau}(\mathscr{Pr}), J^{\tau}(\mathscr{Pr}))$, $\tau \in (0,1)$, of perturbation data on the 3-punctured disk S_3. We will write

$$\mathcal{P}^{\tau}_z(\mathscr{Pr}) = (K^{\tau}_z(\mathscr{Pr}), J^{\tau}_z(\mathscr{Pr}))$$

for the value of the perturbation data at the point $z \in S_3$. We will require the family $\mathcal{P}^{\tau}(\mathscr{Pr})$, $\tau \in (0,1)$, to satisfy the following conditions:

(1) When $0 < \tau < \delta$, for every point in $St^{-}_{0,1}$ with coordinates $(s,t) \in (-\infty, 0] \times [0,1]$ we have

$$J^{\tau}_{s,t}(\mathscr{Pr}) = J^{\mathrm{cont}}_{s+A(\tau)+1,t}((L_0,k_0),(L_1,k_1),(L_0,L_1,i)).$$

In other words, when $s \in [-A(\tau)-1, -A(\tau)]$, $J^{\tau}_{s,-}(\mathscr{Pr})$ coincides with the continuation homotopy J^{cont} after a suitable shift in the s-parameter. Note that $J^{\mathrm{cont}}_{s+A(\tau)+1,t}$ has been defined for all $s \in \mathbb{R}$.

(2) When $0 < \tau < \delta$, for every point in $St^{-}_{1,2}$ with coordinates $(s,t) \in (-\infty, 0] \times [0,1]$ we have

$$J^{\tau}_{s,t}(\mathscr{Pr}) = J^{\mathrm{cont}}_{s+A(\tau)+1,t}((L_1,k_1),(L_2,k_2),(L_1,L_2,i)).$$

(3) When $0 < \tau < \delta$, we require the perturbation data $\mathcal{P}^{\tau}(\mathscr{Pr})$ to coincide with the one assigned by the perturbation data i to the triple (L_0, L_1, L_2) along the complement $S_3 \setminus (St^{-}_{0,1} \cup St^{-}_{1,2})$. Note that this requirement is compatible with the previous two conditions.

(4) When $1-\delta < \tau < 1$, for every point in $St^{+}_{0,2}$ with coordinates $(s,t) \in [0,\infty) \times [0,1]$ we have

$$J^{\tau}_{s,t}(\mathscr{Pr}) = J^{\mathrm{cont}}_{s-A(1-\tau),t}((L_0,k_0),(L_2,k_2),(L_0,L_2,i)).$$

(5) When $1-\delta < \tau < 1$, we require the perturbation data $\mathcal{P}^{\tau}(\mathscr{Pr})$ to coincide with the one assigned by $\mathcal{P}^{\mathrm{tot}}$ to the triple $((L_0,k_0),(L_1,k_1),(L_2,k_2))$ along $S_3 \setminus St^{+}_{0,2}$. Again, this requirement is compatible with the previous one.

(6) For every $\tau \in [\delta, 1-\delta]$, the perturbation data $\mathcal{P}^{\tau}(\mathscr{Pr})$ coincides with the Floer datum of the pair $((L_0,k_0),(L_1,k_1))$ (as assigned by $\mathcal{P}^{\mathrm{tot}}$) along $St^{-}_{0,1}$ outside some compact subset. We require the analogous condition to hold also with respect to $((L_1,k_1),(L_2,k_2))$ along $St^{-}_{1,2}$.

(7) For every $\tau \in [\delta, 1-\delta]$, the perturbation data $\mathcal{P}^{\tau}(\mathscr{Pr})$ coincides with the Floer datum of the pair (L_0, L_2), as assigned by the perturbation data i, along $St^{+}_{0,2}$ outside some compact subset.

(8) For all $\tau \in (0,1)$, the almost complex structures $J^{\tau}(\mathscr{Pr})$ and perturbation forms $K^{\tau}(\mathscr{Pr})$ from $\mathcal{P}^{\tau}(\mathscr{Pr})$ are all of the types and sizes admissible in the construction of $\mathcal{P}^{\mathrm{tot}}$. In particular, they should satisfy the inequality (3.23)

(for $d+1=3$) and the two conditions on K and J that are listed before that inequality on page 142.

(9) In case the three pairs (L_0, k_0), (L_1, k_1), (L_2, k_2) all correspond to the perturbation data i, i.e., $k_0 = k_1 = k_2$, we take the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$ to be constant with respect to τ, and moreover to coincide with the perturbation data i. Note that this is compatible with the rest of the conditions above, since earlier we required each of the homotopies J_s^{cont} that appear in points (1), (2), and (4) above to be constant (with respect to s).

Fix three intersection points $x \in \overline{L}_0 \cap \overline{L}_1$, $y \in \overline{L}_1 \cap \overline{L}_2$, $w \in \overline{L}_0 \cap \overline{L}_2$. Denote by $\mathcal{M}^{\mathcal{P}r_2}(x, y, w; \{\mathcal{P}^\tau(\mathcal{P}r)\})$ the space of all pairs (η, u) where $\eta \in (0,1)$ and $u : S_3 \to X$ solves the generalized Floer equation with respect to the perturbation data $\mathcal{P}^\eta(\mathcal{P}r)$, with Lagrangian boundary conditions prescribed by $\overline{L}_0, \overline{L}_1, \overline{L}_2$ and with asymptotics at the ends being x, y, and w.

By choosing the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$ to be generic, $\mathcal{M}^{\mathcal{P}r_2}(x, y, w; \{\mathcal{P}^\tau(\mathcal{P}r)\})$ is a smooth manifold of dimension $d(x, y, w) = |w|' - |x|' - |y|' + 1$, where $|\cdot|'$ denotes cohomological degree. Moreover, standard arguments show that if $d(x, y, z) = 0$ then $\mathcal{M}^{\mathcal{P}r_2}(x, y, z; \{\mathcal{P}^\tau(\mathcal{P}r)\})$ is compact, hence consists of finitely many points. We then define

$$\mathcal{P}r_2^i(x, y) := \sum_z \#\mathcal{M}^{\mathcal{P}r_2}(x, y, z; \{\mathcal{P}^\tau(\mathcal{P}r)\})z,$$

where z runs over all points in $\overline{L}_0 \cap \overline{L}_2$ such that $|z|' = |x|' + |y|' - 1$ (hence $d(x, y, z) = 0$). As before, $\#\mathcal{M}^{\mathcal{P}r_2}(-)$ is the count (with values in $\mathbb{Z}_2$, or under additional assumptions in $\mathbf{k}$) of elements in the above space.

We claim that $\mathcal{P}r_1^i$ and $\mathcal{P}r_2^i$ form together an A_2-functor. This amounts to showing that $\mathcal{P}r_2^i$ satisfies the identity

$$\begin{aligned} &\mathcal{P}r_1^i(\mu_2^{\text{tot}}(x, y)) - \mu_2^{(i)}(\mathcal{P}r_1^i(x), \mathcal{P}r_1^i(y)) \\ &\quad = \epsilon_1 \mu_1^{(i)} \mathcal{P}r_2^i(x, y) + \epsilon_2 \mathcal{P}r_2^i(\mu_1^{\text{tot}}(x), y) + \epsilon_3 \mathcal{P}r_2^i(x, \mu_1^{\text{tot}}(y)), \end{aligned} \tag{3.26}$$

where $\epsilon_1, \epsilon_2, \epsilon_3 = \pm 1$ are signs that depend on the degrees of x, y, and z (but as we work with $\mathbf{k} = \mathbb{Z}_2$, the precise value of these signs is irrelevant). Here we have denoted by μ^{tot} the A_∞-operations in the category $\mathcal{F}uk^{\text{tot}}(\mathcal{X}; R, r)$ and by $\mu^{(i)}$ those from the category $\mathcal{F}uk(\mathcal{X}; i)$.

The proof of (3.26) is based on standard arguments in Floer theory. Fix w such that $d(x, y, w) = 1$. Then $\mathcal{M}^{\mathcal{P}r_2}(x, y, w; \{\mathcal{P}^\tau(\mathcal{P}r)\})$ is a 1-dimensional smooth manifold. Its compactification $\overline{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$ is a compact 1-dimensional smooth manifold with boundary. The boundary points of $\partial\overline{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$ consist of five types of broken trajectories, as depicted in Figure 3.9. The first two types correspond to $\tau \to 1^-$ and $\tau \to 0^+$, respectively, and the number of occurrences of each of them equals the coefficient of w in the first and second terms on the left-hand side of (3.26), respectively. The other three types of broken trajectories, occurring at instances of time $0 < \tau_0 < 1$, correspond to standard breaking along strip-like ends. The coefficient of w in each of the terms on the right-hand side of (3.26) equals the number of occurrences of broken trajectories of the corresponding type. The identity (3.26) now follows (with appropriate sings ϵ_i) since the signed number of boundary points in $\bar{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$ must be 0. This concludes the construction of the A_2-functor $\mathcal{P}r^i$ for every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, r)$.

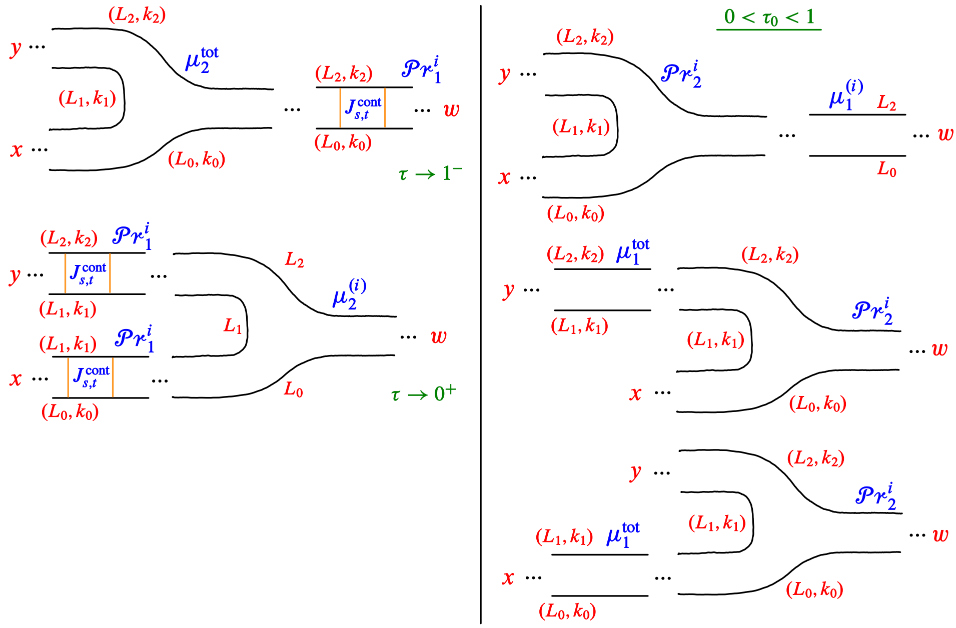

Figure 3.9. The five possible types of boundary points of $\partial\bar{\mathcal{M}}^{\mathcal{P}r_2}(x, y, w)$.

We now claim that $\mathcal{P}r^i$ preserves filtrations. We have already proved earlier that $\mathcal{P}r_1^i$ is filtration preserving, so it remains to deal with the second component $\mathcal{P}r_2^i$. The fact that $\mathcal{P}r_2^i$ is filtration preserving follows from condition (8) on page 145. Indeed, due to inequality (3.13), the proof from pages 127–131 extends with minor modifications to show that, for every $(\eta, u) \in \mathcal{M}^{\mathcal{P}r_2}(x, y, z; \{\mathcal{P}^\tau(\mathcal{P}r)\})$, the symplectic area of u satisfies $A(u) > 0$.

We will soon apply Proposition 3.20 with $\mathcal{A}^i = \mathcal{F}uk(\mathcal{X}; i)$, $i \in \mathcal{I} := \mathcal{B}(\bar{\mathcal{X}}; R, r)$, $\mathcal{A}^{\text{tot}} = \mathcal{F}uk^{\text{tot}}(\mathcal{X}; R, r)$. In order to do so, we still need to show several other properties of the A_2-functors $\mathcal{P}r^i$, as required by Proposition 3.20.

The first one is that $\mathcal{P}r^i$ is strictly A_1-unital. This follows immediately from Morse theory, since the continuation maps in Morse theory send the unique local maximum of the Morse function on (L,k) to the corresponding one on (L,i). Therefore $\mathcal{P}r^i_1$ sends strict units to strict units.

The next property is that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$ as A_2-functors. This will follow from the following two statements:

(1) For all pairs of the type $((L',i),(L'',i))$ we have $\mathcal{P}r^i_1 = \mathbb{1}$.

(2) For any three objects of the type (L_0,i), (L_1,i), (L_2,i) we have $\mathcal{P}r^i_2 = 0$.

The first statement has already been proved earlier. The second one follows from condition (9) in the definition of the family $\{\mathcal{P}^\tau(\mathcal{P}r)\}$, namely the requirement that this family is constant (in τ) and coincides with i for all τ. This implies that $\mathcal{M}^{\mathcal{P}r_2}(x,y,z)$ is the same space as the one defining the operation μ_2 (in the category $\mathcal{F}uk(\mathcal{X};i)$). A transversality/dimension argument now shows that whenever $|z|' = |x|' + |y|' - 1$ we have $\mathcal{M}^{\mathcal{P}r_2}(x,y,z) = \emptyset$. This proves the second statement and concludes the proof that $\mathcal{P}r^i \circ \mathcal{J}^i = \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$ as A_2-functors.

Finally, we claim that the A_2-functor $\mathcal{P}r^i$ has the property described just before the statement of Proposition 3.20. Namely, for every object (L,k) of $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$ there exists a strict isomorphism $u^i_{(L,k)} \in \hom^{\leq 0}_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)}((L,k),(L,i))$. Indeed, we can take $u^i_{(L,k)}$ to be the unique critical point of index $n = \dim_{\mathbb{C}} X$ for the Morse function in the Floer datum of the pair $((L,k),(L,i))$. Standard arguments in Morse theory then show that $u^i_{(L,k)}$ is a strict isomorphism.

We are now in a position to apply Proposition 3.20 with $\mathcal{A}^i = \mathcal{F}uk(\mathcal{X};i)$, $i \in \mathcal{I} := \mathcal{B}(\bar{\mathcal{X}};R,r)$, $\mathcal{A}^{\mathrm{tot}} = \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,r)$. By that proposition we obtain the structure of a coherent system of filtered A_∞-categories on the family $\{\mathcal{F}uk(\mathcal{X};i)\}_{i\in\mathcal{B}(\bar{\mathcal{X}};R,r)}$. Note that this holds for every $r \leq R$, hence in particular also for $r = R$. Note also that $\mathcal{B}(\bar{\mathcal{X}};R,r) \subset \mathcal{B}(\bar{\mathcal{X}};R,R)$, and we can arrange our choices (of $\mathcal{P}^{\mathrm{tot}}$) such that the coherent system over the larger base $\mathcal{B}(\bar{\mathcal{X}};R,R)$ restricts to the one over the smaller base $\mathcal{B}(\bar{\mathcal{X}};R,r)$.

The above construction gives us many coherent systems. Namely, one coherent system over each base $\mathcal{B}(\bar{\mathcal{X}};R,R)$ for all $0 < R \leq R_{\bar{\mathcal{X}}}$. We denote the comparison functors of the coherent system over $\mathcal{B}(\bar{\mathcal{X}};R,R)$ by $\mathcal{F}^{j,i}_R : \mathcal{F}uk(\mathcal{X};i) \to \mathcal{F}uk(\mathcal{X};j)$ for every $i,j \in \mathcal{B}(\bar{\mathcal{X}};R,R)$. Similarly, we denote by T^{i_2,i_1,i_0}_R the natural transformations of this system (relating $\mathcal{F}^{i_2,i_1}_R \circ \mathcal{F}^{i_1,i_0}_R$ to $\mathcal{F}^{i_2,i_0}_R$).

Remark 3.21. Ideally, we would have liked to construct one total category $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};\mathcal{B}(\bar{\mathcal{X}}))$ over the entire base $\mathcal{B}(\bar{\mathcal{X}})$. Unfortunately this is not so straightforward to achieve, at least not with our construction of the total categories. The difficulty has to do with establishing a set of perturbation data $\mathcal{P}^{\mathrm{tot}}$ (over the entire base $\mathcal{B}(\bar{\mathcal{X}})$, and in fact even over subspaces of it of the form $\mathcal{B}(\bar{\mathcal{X}};R_1,R_1) \cup \mathcal{B}(\bar{\mathcal{X}};R_2,R_2)$) that is both consistent (with respect to splitting/gluing) and at the same time yields filtration-preserving operations μ_d.

To understand better the difficulty, consider the case of the union $\mathcal{B}(\overline{\mathcal{X}}; R_1, R_1) \cup \mathcal{B}(\overline{\mathcal{X}}; R_2, R_2)$, where $R_2 < R_1 \leq R_{\overline{\mathcal{X}}}$. Let (i_0, i_1, i_2, i_3) be a tuple of perturbation data with

$$i_1, i_3 \in \mathcal{B}(\overline{\mathcal{X}}; R_1, R_1) \cap \mathcal{B}(\overline{\mathcal{X}}; R_2, R_2) = \mathcal{B}(\overline{\mathcal{X}}; R_1, R_2),$$
$$i_2 \in \mathcal{B}(\overline{\mathcal{X}}; R_1, R_1) \setminus \mathcal{B}(\overline{\mathcal{X}}; R_2, R_2), \quad i_0 \in \mathcal{B}(\overline{\mathcal{X}}; R_2, R_2) \setminus \mathcal{B}(\overline{\mathcal{X}}; R_1, R_1).$$

For some choices of i_0, i_1, i_2, i_3, we might have

$$R^{(i_0,i_1,i_2,i_3)} = R_2, \quad R^{(i_1,i_2,i_3)} = R_1, \quad R^{(i_0,i_1,i_3)} = R_2,$$
$$r^{(i_0,i_1,i_2,i_3)} = R_1, \quad r^{(i_1,i_2,i_3)} = R_1, \quad r^{(i_0,i_1,i_3)} = R_2,$$

where $R^{(\cdots)}$ and $r^{(\cdots)}$ are defined in (3.21). Now recall that in order to obtain a choice of perturbation data that preserves filtrations we used inequalities (3.22) and (3.23). Since $R_2 < R_1$, in our case we have $R^{(i_0,i_1,i_2,i_3)} < r^{(i_0,i_1,i_2,i_3)}$, so inequality (3.22) does not hold. There are also problems regarding the consistency of the perturbation data with respect to splitting and gluing. Indeed, a Floer polygon labeled by $((L_0, i_0), (L_1, i_1), (L_2, i_2), (L_3, i_3))$ can split into two polygons labeled by $((L_1, i_1), (L_2, i_2), (L_3, i_3))$ and by $((L_0, i_0), (L_1, i_1), (L_3, i_3))$. And vice versa, pairs of polygons of the latter type can be glued into polygons of the former type (assuming obvious matching assumptions). In our case we have $r^{(i_1,i_2,i_3)} > r^{(i_0,i_1,i_3)}$, so in order to achieve the consistency of the perturbation data while at the same time having filtration-preserving μ_d's, we would need to decrease the size of the perturbation forms of $\mathcal{P}^{\mathrm{tot}}((L_1, i_1), (L_2, i_2), (L_3, i_3))$ to at most $r^{(i_0,i_1,i_3)}$. In turn, this might not be compatible with the size of the perturbation forms of the data i_2 (which may be of size R_1). Similar problems arise also with the behavior of the measurement $R^{(i_0,\ldots,i_d)}$ with respect to splitting/gluing when the indices i_k are spread in spaces $\mathcal{B}(\overline{\mathcal{X}}; R, R)$ with different values of R.

Summing up, while our construction for total categories works over bases like $\mathcal{B}(\overline{\mathcal{X}}; R, R)$, it is not clear whether the construction can be extended over a base of the type $\mathcal{B}(\overline{\mathcal{X}}; R_1, R_1) \cup \mathcal{B}(\overline{\mathcal{X}}; R_2, R_2)$.

In view of the above, instead of constructing a total category over $\mathcal{B}(\overline{\mathcal{X}})$ we will simply try to extend the various coherent systems (coming from the total categories) over each $\mathcal{B}(\overline{\mathcal{X}}; R, R)$ to coherent systems over unions of such subspaces.

We proceed now with the extension of the system of comparison functors to spaces beyond the type $\mathcal{B}(\overline{\mathcal{X}}; R, R)$. Recall that if $R' \leq R \leq R_{\overline{\mathcal{X}}}$ then $\mathcal{B}(\overline{\mathcal{X}}; R, R) \cap \mathcal{B}(\overline{\mathcal{X}}; R', R') = \mathcal{B}(\overline{\mathcal{X}}; R, R')$. We claim that the construction of the total categories above and the coherent systems resulting from them can be carried out in such a way that the following holds for every $R' \leq R \leq R_{\overline{\mathcal{X}}}$: the two coherent systems, one over $\mathcal{B}(\overline{\mathcal{X}}; R, R)$ and the other over $\mathcal{B}(\overline{\mathcal{X}}; R', R')$, coincide over the overlap $\mathcal{B}(\overline{\mathcal{X}}; R, R')$. In other words, the functors $\mathcal{F}_R^{j,i}$ and $\mathcal{F}_{R'}^{j,i}$ coincide for every $i, j \in \mathcal{B}(\overline{\mathcal{X}}; R, R')$, and similarly for the natural transformations $T_R^{i_2,i_1,i_0}$, $T_{R'}^{i_2,i_1,i_0}$.

The proof is based on three steps. The first one is that we can choose the perturbation data for the total categories in such a way that the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ and $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R')$ coincide along the overlaps $\mathcal{B}(\bar{\mathcal{X}}; R, R')$ (and are equal along that overlap to $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$). The way we chose the perturbation data $\mathcal{P}^{\mathrm{tot}}$ makes this possible, since in case $i_0, \dots, i_d \in \mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R')$ the definitions of the parameters $R^{(i_0,\dots,i_d)}$ and $r^{(i_0,\dots,i_d)}$ are independent of whether we view the indices i_k as elements of $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ or of $\mathcal{B}(\bar{\mathcal{X}}; R', R')$.

The second step has to do with the A_2-functors $\mathcal{P}r^i$. To keep track of the domain of these functors, we will temporarily add a subscript R to their notation:

$$\mathcal{P}r^i_R : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i), \quad i \in \mathcal{B}(\bar{\mathcal{X}}; R, R);$$

we also have

$$\mathcal{P}r^i_{R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i), \quad i \in \mathcal{B}(\bar{\mathcal{X}}; R', R').$$

Recall that $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R') = \mathcal{B}(\bar{\mathcal{X}}; R, R')$ and that $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ is a subcategory of both $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ and $\mathcal{B}(\bar{\mathcal{X}}; R', R')$, which are the domains of $\mathcal{P}r^i_R$ and $\mathcal{P}r^i_{R'}$, respectively. We need to show that for every $i \in \mathcal{B}(\bar{\mathcal{X}}; R, R')$ the two A_2-functors $\mathcal{P}r^i_R$ and $\mathcal{P}r^i_{R'}$ coincide when restricted to the subcategory $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ of their respective domains. This can be achieved by arranging the 1-parametric families of the perturbation data $\{\mathcal{P}^\tau(\mathcal{P}r)\}_{\tau\in(0,1)}$ for $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ to coincide over their intersection.

The third and last step concerns the extensions of the preceding A_2-functors to A_∞-functors, as stated in Proposition 3.20. Here we need to show that the A_∞-extension $\mathcal{Q}^i_R$ of $\mathcal{P}r^i_R$ can be arranged to coincide with the A_∞-extension $\mathcal{Q}^i_{R'}$ of $\mathcal{P}r^i_{R'}$ over the common subcategory of their domains $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$ whenever i is in $\mathcal{B}(\bar{\mathcal{X}}; R, R')$.

This can be proved by means of the algebraic Lemmas 3.18 and 3.19 and Proposition 3.20. Note that for $i \in \mathcal{B}(\bar{\mathcal{X}}; R, R')$, the inclusion of $\mathcal{F}uk(\mathcal{X}; i)$ into any of the categories $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R')$, $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$, and $\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R')$ is a filtered quasi-equivalence.

We now claim that the four maps

$$\begin{aligned}
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R), \mathcal{F}uk(i); \mathcal{P}r^i_R, \mathcal{P}r^i_R\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R', R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R'}, \mathcal{P}r^i_{R'}\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R), \mathcal{F}uk(i); \mathcal{P}r^i_R, \mathcal{P}r^i_R\big)\\
&\qquad\qquad \to \mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big),\\
&\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R', R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R'}, \mathcal{P}r^i_{R'}\big)\\
&\qquad\qquad \to \mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big),
\end{aligned}$$

induced by the obvious restrictions, are all isomorphisms of persistence modules. Here we have omitted "$\mathcal{X}$" from all the $\mathcal{F}uk$-categories in an attempt to keep the formulas short. The A_2-functor $\mathcal{P}r^i_{R,R'} : \mathcal{F}uk^{\mathrm{tot}}(R, R') \to \mathcal{F}uk(i)$ appearing in the third and fourth maps above is just $\mathcal{P}r^i_{R,R'} := \mathcal{P}r^i_R|_{\mathcal{F}uk^{\mathrm{tot}}(R,R')} = \mathcal{P}r^i_{R'}|_{\mathcal{F}uk^{\mathrm{tot}}(R,R')}$.

Indeed, that the first two maps are isomorphisms follows from Lemma 3.18. For the third map, note that, by Lemma 3.18, for every $R' \leq R$ the map

$$\mathrm{PHH}\big(\mathcal{F}uk^{\mathrm{tot}}(R, R'), \mathcal{F}uk(i); \mathcal{P}r^i_{R,R'}, \mathcal{P}r^i_{R,R'}\big) \to \mathrm{PHH}(\mathcal{F}uk(i), \mathcal{F}uk(i); \mathbb{1}, \mathbb{1}) \tag{3.27}$$

(also induced by restriction) is an isomorphism of persistence modules. Now, the first of the above four maps factors as the composition of the third map with the map in (3.27). Since the latter map and the first map are both isomorphisms, it follows that the same holds for the third map. The proof that the fourth map is an isomorphism is similar.

We proceed now by extending $\mathcal{P}r^i_R|_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,R')} = \mathcal{P}r^i_{R'}|_{\mathcal{F}uk^{\mathrm{tot}}(\mathcal{X};R,R')}$ to an A_∞-functor

$$\mathcal{Q}^i_{R,R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R') \to \mathcal{F}uk(\mathcal{X}; i).$$

Next we extend $\mathcal{Q}^i_{R,R'}$ twice more: once to an A_∞-functor $\mathcal{Q}^i_R : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R) \to \mathcal{F}uk(\mathcal{X}; i)$ and another time to an A_∞-functor $\mathcal{Q}^i_{R'} : \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R', R') \to \mathcal{F}uk(\mathcal{X}; i)$. This concludes the third step.

Finally, it is possible to show that the systems of natural transformations T^{i_2,i_1,i_0}_R and $T^{i_2,i_1,i_0}_{R'}$ can be chosen to agree over the intersection $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R')$. This can be done by arguments similar to those above (based on the construction of the natural transformations from the proof of Proposition 3.20).

Recall that the comparison functors $\mathcal{F}^{j,i}_R$ in the coherent system over $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ are given by $\mathcal{F}^{j,i}_R = \mathcal{Q}^j_R \circ \mathcal{J}^i_R$, where $\mathcal{J}^i_R : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk^{\mathrm{tot}}(\mathcal{X}; R, R)$ is the inclusion functor. The comparison functors over $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ and $\mathcal{B}(\bar{\mathcal{X}}; R, R')$ have essentially the same expressions.

Before we proceed, let us summarize what we have proven so far.

Proposition 3.22. *The space of perturbation data $\mathcal{B}(\bar{\mathcal{X}})$ can be written as a union*

$$\mathcal{B}(\bar{\mathcal{X}}) = \bigcup_{R \leq R_{\bar{\mathcal{X}}}} \mathcal{B}(\bar{\mathcal{X}}; R, R)$$

of mutually overlapping subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$. The overlap between any two of these subspaces is given by $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R') = \mathcal{B}(\bar{\mathcal{X}}; R, R')$ for every $R' \leq R$.

The family of Fukaya categories $\{\mathcal{F}uk(\mathcal{X}; i)\}_{i \in \mathcal{B}(\bar{\mathcal{X}}; R, R)}$ over each subspace $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ can be endowed with a coherent system of comparison A_∞-functors that are filtered quasi-equivalences and are strictly A_1-unital. Moreover, for every

R, R', there exist coherent systems over $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and over $\mathcal{B}(\bar{\mathcal{X}}; R', R')$ that agree along the overlap $\mathcal{B}(\bar{\mathcal{X}}; R, R) \cap \mathcal{B}(\bar{\mathcal{X}}; R', R')$.

Finally, all the comparison functors in the various coherent systems above act as the identity maps on the objects of the respective categories. The first-order terms of the comparison functors induce the canonical continuation maps in Floer homology.

The structure of the subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and their overlaps is schematically depicted in Figure 3.10.

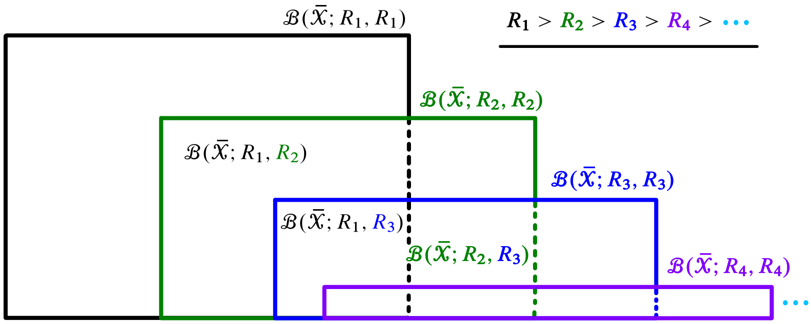

Figure 3.10. The subspaces $\mathcal{B}(\bar{\mathcal{X}}; R, R)$ and their overlaps.

Remark 3.23. The type of structure described in Proposition 3.22 can be defined abstractly for any collection of A_∞-categories (filtered or not) $\{\mathcal{A}^i\}_{i \in \mathcal{B}}$ over a base $\mathcal{B}$. An appropriate name for such a structure could be *a collection of coherent systems with overlaps*. We will not pursue this direction further in this memoir and will now proceed with our Fukaya categories.

We are ready now to describe the functors $\mathcal{F}^{\mathcal{P}_1, \mathcal{P}_0}$ claimed in Theorem 3.12. To simplify the notation, we continue to denote perturbation data by indices like i instead of $\mathcal{P}$, and the comparison functors will be denoted by $\mathcal{F}^{j,i}$ instead of $\mathcal{F}^{\mathcal{P}_1, \mathcal{P}_0}$. Note that the construction below is purely formal and can be applied to any family of A_∞-categories endowed with a collection of coherent systems with overlaps (as in Remark 3.23).

For every $R', R'' \leq R_{\bar{\mathcal{X}}}$ we choose a base point

$$l(R', R'') \in \mathcal{B}(\bar{\mathcal{X}}; \max\{R', R''\}, \min\{R', R''\}).$$

Given $i \in \mathcal{B}(\bar{\mathcal{X}})$, define $R^{(i)} := \sup\{\tilde{R} \mid \tilde{R} \leq R_{\bar{\mathcal{X}}} \text{ and } i \in \mathcal{B}(\bar{\mathcal{X}}; \tilde{R}, \tilde{R})\}$, which is a special case of the first parameter in (3.21).

Let $i, j \in \mathcal{B}(\bar{\mathcal{X}})$. If $i = j$, define $\mathcal{F}^{i,j} := \mathbb{1}_{\mathcal{F}uk(\mathcal{X};i)}$. If $i \neq j$, set $l_{i,j} := l(R^{(i)}, R^{(j)})$ and define

$$\mathcal{F}^{j,i} := \mathcal{F}^{j, l_{i,j}}_{R^{(j)}} \circ \mathcal{F}^{l_{i,j}, i}_{R^{(i)}} : \mathcal{F}uk(\mathcal{X}; i) \to \mathcal{F}uk(\mathcal{X}; j). \tag{3.28}$$

It is easy to see that the functors $\mathcal{F}^{j,i}$ satisfy all the properties claimed in Theorem 3.12.

Remark 3.24. The system of functors $\mathcal{F}^{j,i}$ defined by (3.28) is not canonical. The construction uses many different choices at different stages. However, one can show that the dependence on the choices of the base points $l_{R',R''}$ is somewhat controlled. If one replaces the base points $l_{R',R''}$ by a different set of choices $l'_{R',R''}$, then the resulting system of comparison functors $\mathcal{F}'^{j,i}$ will be naturally quasi-isomorphic to the system $\mathcal{F}^{j,i}$ by natural quasi-isomorphisms that preserve filtrations.

The applications in this memoir do not require the comparison functors or the natural transformations between them to be canonical. In fact, in what follows we just need to know that the different Fukaya categories $\mathcal{F}uk(\mathcal{X}; i)$ are filtered quasi-equivalent.

It remains to address the last statement of Theorem 3.12, concerning the collection of Lagrangians $\bar{\mathcal{X}}'$. This follows immediately from our construction. Indeed, let $\bar{\mathcal{X}}'$ be a collection of Lagrangians with $\bar{\mathcal{X}}' \supset \bar{\mathcal{X}}$, and assume that $\bar{\mathcal{X}}'$ satisfies all the conditions from the beginning of Section 3.2.1. We can use the same symplectic embedding ϕ from (3.10) but with a smaller radius $R_{\bar{\mathcal{X}}'} \leq R_{\bar{\mathcal{X}}}$. Clearly $\mathcal{B}(\bar{\mathcal{X}}')|_{\bar{\mathcal{X}}} \subset \mathcal{B}(\bar{\mathcal{X}})$.

This completes the proof of Theorem 3.12. ■

3.4 Proofs of the main symplectic applications

The first subsection is dedicated to the proofs of Theorems 3.1 and 3.4 and the second to the proof of Corollary 3.7.

3.4.1 Proofs of Theorems 3.1 and 3.4

We begin with Theorem 3.1, and then proceed to the proof of Theorem 3.4 in Section 3.4.1.2 on page 156.

3.4.1.1 Proof of Theorem 3.1. Recall from Theorem 3.12 that there are choices of Floer data and perturbation data $\mathcal{P}$ such that the resulting Fukaya category $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$ is filtered, strictly unital, and, without filtrations, quasi-equivalent to the subcategory $\mathcal{F}uk(\mathcal{X})$ of $\mathcal{F}uk(X)$ whose collection of objects is $\mathcal{X}$ (see Section 3.1 for the notation).

We will now apply the discussion in Section 3.2.1 to the filtered A_∞-category $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$. The first step is to discuss the shift functor. There is an obvious shift functor on the category $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$ which acts on objects by $\Sigma^r L = (\bar{L}, h_L + r, \theta_L)$ (see Section 3.1). This action on objects induces an A_∞-shift functor on $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, because the perturbation data $\mathcal{P}$ depends only on the geometric part of the marked Lagrangians. It is easy to see that this shift functor on $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is compatible with the shift functor defined on $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ as in Section 3.2.1: for a module $\mathcal{M} \in \mathrm{Fmod}(\mathcal{A})$, the filtered module $\Sigma^r \mathcal{M}$ is defined by $(\Sigma^r \mathcal{M})^{\leq \alpha}(N) = \mathcal{M}^{\alpha - r}(N)$, with the same μ_d-operations as $\mathcal{M}$. Indeed,

$$\Sigma^r(\mathcal{Y}(L)) = \mathcal{Y}(\Sigma^r L).$$

We now define

$$\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P}) = H^0(\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\nabla}).$$

Recall that $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\nabla}$ is constructed by first considering the category $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\#}$ consisting of the triangulated completion of the Yoneda modules of the elements in $\mathcal{X}$. Triangles are understood here to be of the form $\mathcal{M} \xrightarrow{f} \mathcal{N} \to \mathrm{Cone}(f)$, where the cone construction is in the sense of filtered A_∞-modules and the morphism f preserves filtration.

The category $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\nabla}$ is the full subcategory of $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ that contains $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\#}$, as well as all the modules in $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$, together with all their shifts and translates that are r-quasi-isomorphic to the objects in $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\#}$ for some $r \geq 0$. From the discussion above concerning the shift functor, it follows immediately that $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$ is indeed a TPC.

Remark 3.25. More explicitly, to obtain $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\nabla}$ from $\mathcal{F}uk(\mathcal{X};\mathcal{P})^{\#}$, we first add to the Yoneda modules, together with all possible shifts and translates, the r-acyclic modules for all $r \geq 0$; then we take the triangulated completion of all of them.

We now start the proof by discussing the independence of $\mathcal{C}\mathcal{F}uk(\mathcal{X};\mathcal{P})$ of the perturbation data $\mathcal{P}$, up to TPC equivalence. The argument is a direct consequence of the system of comparison functors for the categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$, as given in Theorem 3.12. Indeed, for two choices of admissible perturbation data $\mathcal{P}_1$ and $\mathcal{P}_2$, we have a filtered functor

$$\mathcal{F}^{1,2} : \mathcal{F}uk(\mathcal{X},\mathcal{P}_1) \to \mathcal{F}uk(\mathcal{X},\mathcal{P}_2).$$

The existence of the natural transformations that compare the compositions $\mathcal{F}^{1,2} \circ \mathcal{F}^{2,1}$, $\mathcal{F}^{2,1} \circ \mathcal{F}^{1,2}$ with the respective identities implies that the associated homological functor is full and faithful, in the sense that it induces an isomorphism of persistence modules $H(\mathcal{F}^{1,2}) : H^0(\hom^{\leq r}(X,Y)) \to H^0(\hom^{\leq r}(X,Y))$ for every $r \in \mathbb{R}$ and every two objects X, Y of $\mathcal{F}uk(\mathcal{X},\mathcal{P}_1)$. As in the unfiltered case, a consequence of the existence of these functors and the natural transformation relating them is that the pullback of filtered modules $[\mathcal{F}^{1,2}]^* : \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X},\mathcal{P}_2)) \to \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X},\mathcal{P}_1))$ sends each Yoneda module $\mathcal{Y}_{\mathcal{P}_2}(L)$ to a module 0-quasi-isomorphic to $\mathcal{Y}_{\mathcal{P}_1}(L)$ (in the sense that the two modules are related by a morphism that induces a 0-isomorphism in the homological category). The standard properties of the pullback of A_∞-modules imply that $[\mathcal{F}^{1,2}]^*$ respects the triangulated structure with respect to 0-weight triangles as well as shift functors. We deduce that this pullback also sends r-isomorphisms to r-isomorphisms, and thus sends $\mathcal{F}uk(\mathcal{X}, \mathcal{P}_2)^\nabla$ to $\mathcal{F}uk(\mathcal{X}, \mathcal{P}_1)^\nabla$. Finally, using the relevant natural transformations, we deduce that $H[\mathcal{F}^{1,2}]^*$ is an equivalence of TPCs.

We will denote the resulting TPC by $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ in place of $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, except where this may lead to confusion.

We now continue with points (i), (ii), and (iii) of Theorem 3.1. The first point is immediate because $\hom_{\mathcal{C}\mathcal{F}uk(\mathcal{X})}(L, L') \cong \mathrm{HF}(L, L')$ as persistence modules, with the persistence structure on $\mathrm{HF}(L, L')$ as described in Section 3.2.2.

The second point of the theorem claims that the ∞-level of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is equivalent to $D\mathcal{F}uk(\mathcal{X})$. The argument is the following. First, it is well known [53] that an alternative, equivalent model for $D\mathcal{F}uk(\mathcal{X})$ is given by the homological category of twisted complexes, $H^0(\mathrm{Tw}(\mathcal{F}uk(\mathcal{X})))$. Here $\mathcal{F}uk(\mathcal{X})$ is any A_∞-category that represents the Fukaya category with objects the Lagrangians in $\mathcal{X}$. Thus we can take in its place $\mathcal{F}uk_{uf}(\mathcal{X}; \mathcal{P})$, where the notation $(-)_{uf}$ means that we neglect the filtration. A variant of Lemma 2.96 applies also to our A_∞-setting, and it implies that each twisted complex in $\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X}; \mathcal{P}))$ can be viewed as a filtered twisted complex in $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))$ whose filtration is forgotten. Filtered twisted complexes are discussed in Section 3.2.1 (and in more detail in the dg case, in Section 2.5.1). By passing to homology, this means that $[H^0\,\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))]_\infty \cong H^0\,\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X}; \mathcal{P})) \cong D\mathcal{F}uk(\mathcal{X})$. In turn, $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))$ is easily seen to be equivalent to $\mathcal{F}uk(\mathcal{X}; \mathcal{P})^\#$ as TPC categories. Thus,

$$[\mathcal{F}uk(\mathcal{X}; \mathcal{P})^\nabla]_\infty \cong [\mathcal{F}uk(\mathcal{X}; \mathcal{P})^\#]_\infty \cong [H^0\,\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))]_\infty \cong D\mathcal{F}uk(\mathcal{X}).$$

We now turn to the third point of Theorem 3.1. The argument here makes again essential use of the systems of comparison functors provided by Theorem 3.12, which allows us to extend perturbation data from one set $\mathcal{X}$ to a larger one, $\mathcal{X}'$.

Let N be a marked Lagrangian that is in general position with respect to the family $\bar{\mathcal{X}}$. We add N, as well as all its shifts and translates, to the family $\mathcal{X}$, thus obtaining $\mathcal{X}'$. By Theorem 3.12 and the invariance of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ relative to perturbation data, we can now assume that the perturbation data $\mathcal{P}$ extends to a new one $\mathcal{P}'$ that defines the filtered A_∞-category $\mathcal{F}uk(\mathcal{X}', \mathcal{P}')$. Finally, we can extend $\mathcal{P}'$ to a perturbation data that defines an A_∞-category which is equivalent to $\mathcal{F}uk(X)$ if filtrations are forgotten. At this point, it is easier to pursue the argument using twisted complexes. As before, these will be of two types: filtered, those from $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))$, and unfiltered, those belonging to $\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X}'; \mathcal{P}'))$. The assumption that $\mathcal{X}$ generates $D\mathcal{F}uk(X)$ implies that there is an unfiltered quasi-isomorphism of twisted complexes $\phi : N \to C$, where C is a twisted module involving only elements of $\mathcal{X}$. Both N and C have structures of filtered twisted complexes, but the morphism ϕ, a priori, does not see the filtration.

However, using an argument similar to that in Lemma 2.96, we deduce that (after possibly shifting up N) we can view ϕ as a filtration-preserving morphism in $\hom_{\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}';\mathcal{P}'))}$. The mapping cone K of ϕ is acyclic as a twisted module (forgetting the filtration) because ϕ is a quasi-isomorphism. In other words, the identity 1_K of K is homologous to 0 in $\hom_{\mathrm{Tw}(\mathcal{F}uk_{uf}(\mathcal{X}';\mathcal{P}'))}(K, K)$. Using again a reasoning similar to that in Lemma 2.96, we deduce that 1_K vanishes in

$$H^0\big(\hom^{\leq r}_{\mathrm{Tw}(\mathcal{F}uk(\mathcal{X}';\mathcal{P}'))}(K, K)\big)$$

for some $r \geq 0$. This means that ϕ is an r-isomorphism. We can now reformulate the result in terms of modules, and we deduce that $j^*\mathcal{Y}(N)$ is r-isomorphic to a filtered module from $\mathcal{F}uk(\mathcal{X}, \mathcal{P})^{\#}$, where $j : \mathcal{F}uk(\mathcal{X}, \mathcal{P}) \to \mathcal{F}uk(\mathcal{X}', \mathcal{P}')$ is the inclusion. This concludes the proof. ∎

3.4.1.2 Proof of Theorem 3.4. Theorem 3.12 implies that for every two admissible perturbation data $\mathcal{P}_1$ and $\mathcal{P}_2$ there exists a functor

$$[\mathcal{F}^{1,2}]^* : \mathcal{F}uk(\mathcal{X}, \mathcal{P}_2)^{\nabla} \to \mathcal{F}uk(\mathcal{X}, \mathcal{P}_1)^{\nabla}$$

that induces an equivalence of TPCs in homology and is the identity map on the objects from $\mathcal{X}$. Given that $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P}_i) = H^0(\mathcal{F}uk(\mathcal{X}, \mathcal{P}_i)^{\nabla})$, this means that the pseudometric $D^{\mathcal{F}}$ defined on $\mathcal{X}$ using the TPC structure $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P}_2)$ is greater than or equal to the pseudometric defined using $\mathcal{C}\mathcal{F}uk(\mathcal{X}; \mathcal{P}_1)$. By using the functor $\mathcal{F}^{2,1}$ we deduce the opposite inequality, and we conclude that $D^{\mathcal{F}}$ is independent of the perturbation data used to define it.

We begin with the proof of point (i). The inequality relates the spectral distance $\sigma(L, L')$ to the simplest fragmentation metric, $D(-,-) = D^{\{0\}}(-,-)$. The proof is based on a simple consequence of the properties of the Yoneda embedding.

Let $L, L' \in \mathcal{X}$. Without loss of generality, we may assume that $\bar{L}$ and $\bar{L}'$ are Hamiltonian isotopic (otherwise $\sigma(L, L') = +\infty$). Recall also that all Lagrangians in $\mathcal{X}$ are assumed to be graded. Standard Floer theory shows that there is a canonical class $a = o_{L,L'} \in \mathrm{HF}(L, L')$ obtained as the image of the fundamental class $[L] \in H_n(L, \mathbf{k})$ through the PSS morphism $H_*(L, \mathbf{k}) \to \mathrm{HF}(L, L')$. There is a similar class $b = o_{L',L} \in \mathrm{HF}(L', L)$, and it is a simple consequence of the properties of the Yoneda embedding that these two classes have the property that $a * b = [L]$ and $b * a = [L']$. Here $*$ is the product induced in homology by the A_∞-composition μ_2, and $[L]$ and $[L']$ are the respective fundamental classes.

We start with an argument in which we neglect filtration issues. Suppose that there are two classes $a \in \mathrm{HF}(L, L')$ and $b \in \mathrm{HF}(L', L)$ such that $a * b = [e_{L,L}] = [L] \in \mathrm{HF}(L, L)$, where $e_{L,L}$ is the unit in $\mathrm{CF}(L, L)$. For coherence, we work here in homological notation (even if to keep track of signs it would be preferable to use cohomological notation as in [53]), so $e_{L,L}$ corresponds to the maximum of the Morse function $f_{\bar{L}}$ as in Section 3.3.1. This means that there are Floer cycles

$\alpha \in \mathrm{CF}(L, L')$, $\beta \in \mathrm{CF}(L', L)$ such that $\mu_2(\alpha, \beta)$ is homologous to $e_{L,L}$. However, recall that $\mathrm{CF}(L, L)$ is the Morse complex of the function $f_{\bar{L}}$. Thus, for degree reasons we obtain $\mu_2(\alpha, \beta) = e_{L,L}$. As explained in the construction of $\mathcal{F}uk(\mathcal{X}, \mathcal{P})$, $e_{L,L}$ is a strict unit.

For every $L_0, L_1 \in \mathcal{X}$ and any cycle $u \in \mathrm{CF}(L_0, L_1)$, we have a morphism of Yoneda modules $\phi^u : \mathcal{Y}(L_0) \to \mathcal{Y}(L_1)$ defined by

$$\phi^u_k(-, -, \ldots, x) = \mu_{k+1}(-, -, \ldots, x, u).$$

Returning to α and β above, we have that $\phi^{\mu_2(\alpha,\beta)} = \phi^{e_{L,L}} = \mathbb{1}_{\mathcal{Y}(L)}$. The composition $\phi^\beta \circ \phi^\alpha$ is not necessarily equal to the morphism $\phi^{\mu_2(\alpha,\beta)}$, but they are homologous as elements of $\hom_{\mathrm{mod}}(\mathcal{Y}(L), \mathcal{Y}(L))$. (See [53] for the explicit formulas for the composition of pre-module homomorphisms and for the differential μ_1^{mod} on $\hom_{\mathrm{mod}}(-, -)$.) We rewrite here this differential (neglecting signs) for a pre-module morphism $t : \mathcal{M} \to \mathcal{N}$, with components

$$t_k : \mathrm{CF}(L_1, L_2) \otimes \cdots \otimes \mathrm{CF}(L_{k-1}, L_k) \otimes \mathcal{M}(L_k) \to \mathcal{N}(L_1).$$

We have that

$$\begin{aligned}(\mu_1^{\mathrm{mod}} t)_m(-, \ldots, -, x) = &\sum_{r+s=m+1} \mu_r^{\mathcal{N}}(-, -, \ldots, t_s(-, -, \ldots, x)) \\ &+ \sum_{l+j=m+1} t_l(-, -, \ldots, \mu_j^{\mathcal{M}}(-, -, \ldots, x)) \\ &+ \sum_{k+g=m+1} t_k(-, \ldots, \mu_g(-, \ldots, -), -, \ldots, x),\end{aligned}$$

where μ_g is the operation in the A_∞-category, and $\mu_r^{\mathcal{M}}$ and $\mu_j^{\mathcal{N}}$ are the respective module operations. Consider the pre-module morphism $T : \mathcal{Y}(L) \to \mathcal{Y}(L)$ with the kth component given by

$$T_k(-, -, \ldots, -, x) = \mu_{k+2}(-, -, \ldots, -, x, \alpha, \beta).$$

Using the fact that $\mu_2(\alpha, \beta)$ is a strict unit and the A_∞-relations, it is easy to see that $\mu_1^{\mathrm{mod}}(T) = \mathbb{1}_{\mathcal{Y}(L)} - \phi^\beta \circ \phi^\alpha$.

In our case, the A_∞-category $\mathcal{F}uk(\mathcal{X}, \mathcal{P})$ is filtered and, if we assume that $\alpha \in \mathrm{CF}^{\leq r}(L, L')$ and $\beta \in \mathrm{CF}^{\leq s}(L, L')$, our discussion above shows that in the TPC $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ we have maps

$$\Sigma^{r+s}\mathcal{Y}(L) \xrightarrow{\Sigma^r \bar{\phi}^\alpha} \Sigma^s \mathcal{Y}(L') \xrightarrow{\bar{\phi}^\beta} \mathcal{Y}(L)$$

that compose to the "shift" map η_{r+s}. Here $\bar{\phi}^\beta : \Sigma^s \mathcal{Y}(L') \to \mathcal{Y}(L)$ is induced by ϕ^β and $\bar{\phi}^\alpha : \Sigma^r \mathcal{Y}(L) \to \mathcal{Y}(L')$ is induced by ϕ^α. The key ingredient here is that the homotopy T only shifts filtration by at most $r + s$.

Returning to our classes $a = [\alpha]$ and $b = [\beta]$, we now use $b * a = [e_{L',L'}]$ and obtain the existence of a similar diagram:

$$\Sigma^{r+s}\mathcal{Y}(L') \xrightarrow{\Sigma^s\bar{\phi}^\beta} \Sigma^r\mathcal{Y}(L) \xrightarrow{\bar{\phi}^\alpha} \mathcal{Y}(L').$$

We immediately deduce from this that $\mathcal{Y}(L)$ and $\mathcal{Y}(L')$ are $(r+s)$-interleaved in the sense of Definition 2.84 in the TPC $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ (one can do better, for instance by using $\max\{r,s\}$, but we will not try to improve the estimates). This means, by Corollary 2.86, that $D(L,L') \leq 4(r+s)$. By reviewing the definition of spectral invariants (2.36), this implies that $D(L,L') \leq 4(\sigma(o_{L,L'}) + \sigma(o_{L',L}))$. The inequality claimed at the first point of the theorem now follows because $\mathrm{HF}(L,L')$ and $\mathrm{HF}(L',L)$ are related by duality and $\sigma(o_{L',L}) = -\sigma(pt_{L,L'})$.

Remark 3.26. (a) The argument at the center of this proof has appeared before in the literature (in less generality, involving only μ_2 and μ_3 operations, instead of the full A_∞-structure) first in [41], as well as in [8]. All these arguments are based on properties of the Yoneda embedding from [53]. These were adjusted to the weakly-filtered setting in [10]. We can achieve a greater degree of generality here because we are able to work with genuinely filtered A_∞-categories which, moreover, are strictly unital.

(b) The fact that the modules $\mathcal{Y}(L)$ and $\mathcal{Y}(L')$ are interleaved by an order controlled by $\sigma(L,L')$ has numerous consequences, because it implies that persistence modules of the form $\mathrm{HF}(N,L)$ and $\mathrm{HF}(N,L')$ are also interleaved by the same order for all $N \in \mathcal{X}$. Viewing this through the prism of the bottleneck metric we easily deduce relations between various measurements of interest, such as boundary depth, spectral range, and so forth, associated to the complexes $\mathrm{CF}(N,L)$ and $\mathrm{CF}(N,L')$.

We now turn to the proof of point (ii) in Theorem 3.4. Assume that $D^{\mathcal{F}}(L,L') < r$. Then there exists a sequence of exact triangles in $[\mathcal{C}\mathcal{F}uk(\mathcal{X})]_\infty$ as below:

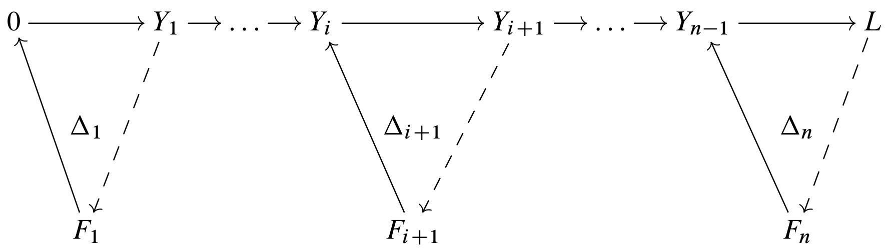

(3.29)

Here all the terms F_i are in $\mathcal{F}$, except for the jth term which is of the form $T^{-1}\Sigma^\alpha L'$ for some $\alpha \in \mathbb{R}$ (we write here N instead of $\mathcal{Y}(N)$ for the elements of $\mathcal{X}$); see (2.40). Moreover, by assumption we have $\sum_i \overline{w}(\Delta_i) \leq r$. We now appeal to the argument from the proof of Corollary 2.88 to associate to the sequence of exact triangles above another sequence of triangles:

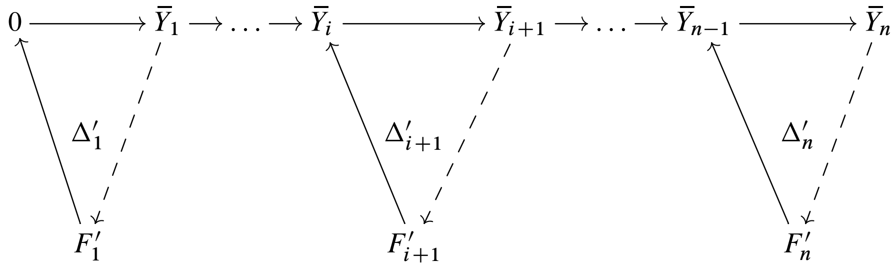

(3.30)

this time all exact in $[\mathcal{C}\mathcal{F}uk(\mathcal{X})]_0$, together with a $2r$-isomorphism $L \to \overline{Y}_n$. Here the F_i''s are shifts of the previous corresponding objects from (3.29). Given that $\mathcal{F}$ is closed under shifts, we have $F_i' \in \mathcal{F}$ except for the jth term, which is still a shift of $T^{-1}L'$. By possibly shifting $\overline{Y}_n$ down, we have a $4r$-isomorphism $\phi : \overline{Y}_n \to L$.

Each exact triangle Δ_i is the image of a cone-type triangle in the category of filtered modules:

$$F'_{i+1} \xrightarrow{\varphi_i} \overline{Y}_i \to \mathrm{Cone}(\varphi_i) = \overline{Y}_{i+1},$$

and there is a module map $\bar{\phi} : \overline{Y}_n \to L$ such that $K = \mathrm{Cone}(\bar{\phi})$ is $4r$-acyclic. In particular, $K(L)$ is a $4r$-acyclic filtered chain complex: the identity of this chain complex is homotopic to zero through a homotopy h of shift at most $4r$. The complex $K(L)$, viewed as a vector space, is the following sum:

$$\begin{aligned} K(L) &= \mathrm{CF}(L, L) \oplus \bigoplus_{i=1}^{n} \mathrm{CF}(L, F_i')[2] \\ &= \mathrm{CF}(L, L) \oplus \bigoplus_{i=1}^{j-1} \mathrm{CF}(L, F_i')[2] \oplus \mathrm{CF}(L, L')[1] \oplus \bigoplus_{i=j+1}^{n} \mathrm{CF}(L, F_i')[2]. \end{aligned} \tag{3.31}$$

Denote the maximum of the Morse function $f_{\overline{L}} : \overline{L} \to \mathbb{R}$ by m. We now make the following two assumptions:

(i) There exists an embedding of a standard symplectic ball $e' : B(s) \to X$ such that $e'(0) = m$, $(e')^{-1}(L) = \mathbb{R}B(s)$, and $e'(B(s)) \cap (L' \cup \bigcup_{F \in \mathcal{F}} F) = \emptyset$.

(ii) The almost complex structures on X that are part of the perturbation data $\mathcal{P}$ pull back to the standard almost complex structure on $B(s)$.

These assumptions are not restrictive because, as shown at the beginning of the proof, the pseudometric $D^{\mathcal{F}}$ is independent of the perturbation data used in its definition. Thus we can pick perturbation data $\mathcal{P}$ adapted to the embedding e', in the sense that the two points (i) and (ii) above are satisfied, and the sequences of exact triangles (3.29) and (3.30) still exist in this case, with the claimed properties.

Returning to the complex $K(L)$, denote by d_K its differential and consider the equation

$$d_K h(m) + h d_K(m) = m.$$

Now $\mathrm{CF}(L, L)$ is a subcomplex of $K(L)$, and thus $d_K(m) = 0$. Hence $d_K h(m) = m$, so m is a boundary in the complex $(K(L), d_K)$.

On the other hand, since K is an iterated cone of A_∞-Yoneda modules, the differential d_K has a particular shape with respect to the splitting from (3.31). This has been worked out in detail in [10, Section 2.6] (in particular, see Theorem 2.14 in that paper), and the relevant ingredients are as follows. The differential d_K can be described by a matrix $(a_{i,j})_{0 \leq i,j \leq n}$, where:

(1) $a_{i,j} = 0$ for $i > j$ (i.e., the matrix $(a_{i,j})$ is upper triangular);

(2) $a_{i,j} : \mathrm{CF}(L, F'_j) \to \mathrm{CF}(L, F'_i)$ for $1 \leq i \leq j \leq n$;

(3) $a_{0,j} : \mathrm{CF}(L, F'_j) \to \mathrm{CF}(L, L)$ for $j \geq 1$;

(4) $a_{0,0} : \mathrm{CF}(L, L) \to \mathrm{CF}(L, L)$ is the Floer differential on $\mathrm{CF}(L, L)$.

Here we have omitted reference to the grading on the F'_j's and L. Moreover, for $j \geq 1$, the maps $a_{0,j}$ can be written as follows. There exist Floer chains $c_{q,p} \in \mathrm{CF}(F'_q, F'_p)$ for every $q > p > 0$ and $c_{q,0} \in \mathrm{CF}(F'_q, L)$ for all $q \geq 1$, all at action levels ≤ 0, such that

$$a_{0,j}(-) = \sum_{2 \leq d, \underline{k}} \mu_d(-, c_{k_d,k_{d-1}}, \ldots, c_{k_2,0}), \tag{3.32}$$

where $\underline{k} = (k_2, \ldots, k_d)$ runs over all partitions $0 < k_2 < \cdots < k_{d-1} < k_d = j$ and μ_d is the dth-order operation in the Fukaya category $\mathcal{F}uk(\mathcal{X})$.

Since m is the maximum of the Morse function, and m is a boundary in the complex $(K(L), d_K)$, it follows from (3.32) that there is a J-holomorphic polygon with one edge on L and the others on the F'_i's and possibly L' (recall that $F'_j = T^{-1}L'$), that goes through m. (See [10, Section 5.1], and in particular pages 91–92 in that paper, for a detailed proof of a very similar statement.) The area of each such polygon is at least $\frac{1}{2}\pi s^2$. This means that the chain homotopy h increases filtration by at least $\frac{1}{2}\pi s^2$ and thus

$$4r \geq \frac{\pi s^2}{2},$$

which shows the claim.

We now proceed to the proof of point (iii) of Theorem 3.4. We will again make use of the fact that the pseudometric $D^{\mathcal{F}}$ is independent of the perturbation data $\mathcal{P}$.

Let $\delta = \delta^{\cap}(N, L'; \mathcal{F})$. For each intersection point $x \in N \cap L'$, fix standard ball embeddings $e_x : B(s) \to X$ with $e^{-1}(N) = \mathbb{R}B(s)$, $e^{-1}(L') = i\mathbb{R}B(s)$, $e(0) = x$, such that all these embeddings are disjoint from the family $\mathcal{F}$ and additionally $\pi s^2 = \delta - \epsilon$ for a small ϵ. We may assume that the almost complex structures that are part of the perturbation data $\mathcal{P}$ pull back to the standard almost complex structure through the embeddings e_x. This implies that if a Floer-type strip or polygon has an input or an output at a point of $N \cap L'$ and has boundary on N, L', and any other elements of the family $\mathcal{F}$, then its energy is at least $\delta' = \frac{\pi s^2}{4} - \epsilon' = \frac{\delta - \epsilon}{4} - \epsilon'$. Here the small ϵ' has to do with making the Hamiltonian (or 1-forms) part of the perturbation data small enough.

Assume now, as in the statement, that the condition $D^{\mathcal{F}}(L, L') < r < \frac{\delta}{16}$ holds. Then there exist a sequence of exact triangles in $\mathcal{CF}uk(\mathcal{X})_0$ as in (3.30) and a $2r$-isomorphism $\psi : L \to \bar{Y}_n$. By Lemma 2.85 this means that $d_{\text{int}}(L, Y_n) \leq 2r$. In particular, there exist maps $u : \Sigma^r \bar{Y}_n \to L$, $v : \Sigma^r L \to \bar{Y}_n$ such that $v \circ \Sigma^r u = \eta_{2r}$. Therefore, given that $N \in \mathcal{X}$, we obtain maps of filtered complexes

$$\bar{Y}_n(N) \xrightarrow{u'} \mathrm{CF}(N, L) \xrightarrow{v'} \bar{Y}_n(N),$$

each of shift at most $2r$, and whose composition is chain homotopic to the identity through a homotopy h that shifts filtration by at most $4r$.

Consider the differential of the complex $\bar{Y}_n(N)$. As a vector space, the complex $\bar{Y}_n(N)$ is a sum of the form $\mathrm{CF}(N, F_1') \oplus \cdots \oplus \mathrm{CF}(N, L') \oplus \mathrm{CF}(N, F_n')$, and the differential is represented by clustered polygons with boundaries on N, L', and the elements of the family $\mathcal{F}$.

Consider the composition $\Psi = p \circ v' \circ u' \circ i$, where $i : \mathrm{CF}(N, L') \to \bar{Y}_n(N)$ is the inclusion and $p : \bar{Y}_n(N) \to \mathrm{CF}(N, L')$ is the projection – in both cases as vector spaces. We claim that Ψ is injective. This would imply that $\dim_{\mathbf{k}} \mathrm{CF}(N, L') \leq \dim_{\mathbf{k}} \mathrm{CF}(N, L)$ and prove the statement in point (iii) of Theorem 3.4.

To show the injectivity of Ψ, we inspect the formula $dh + hd = \mathbb{1} - v' \circ u'$, and recall that the differential d, when restricted to $\mathrm{CF}(N, L')$, drops the filtration by at least $\delta' = \frac{\delta - \epsilon}{4} - \epsilon'$ while h raises filtration by at most $4r < \delta'$ (when ϵ and ϵ' are small enough). In other words, $p(hd(x))$ for each element $x \in \mathrm{CF}(N, L')$ is of strictly lower filtration than x. Similarly, $p(dh(x))$ is also of lower filtration than x. As a result, we deduce that Ψ can be written as $\mathbb{1}$ plus a map that strictly lowers the filtration level. It follows that Ψ is injective.

Finally, we discuss the last point in Theorem 3.4. We assume that $\mathcal{F}$ generates the usual (i.e., without persistence structure) derived Fukaya category $D\mathcal{F}uk(X)$ and we want to show that in this case the pseudometric $D^{\mathcal{F}}$ is finite. From Theorem 3.1 (iii) we deduce that $\mathcal{F}$ generates $\mathcal{CF}uk(\mathcal{X})_\infty$. This implies that any object A in $\mathcal{CF}uk(\mathcal{X})$ is r-isomorphic, for some r, to an object that can be written as an iterated cone with triangles of weight 0, of the form in (3.30). This means that $D^{\mathcal{F}}(A, 0) < \infty$ and implies the last claim in the statement of the theorem. ∎

3.4.2 Pseudometrics on $\mathcal{L}ag(X)$ and proof of Corollary 3.7

The purpose of this subsection is to use the results from Section 3.1 to construct the family of fragmentation metrics on the space $\mathcal{L}ag(X)$ of all closed exact, graded Lagrangians in (X, ω) and prove Corollary 3.7.

We assume the setting in Corollary 3.7, in particular that rank $\mathcal{F}uk(X, \omega) < \infty$. Thus $D\mathcal{F}uk(X)$ admits a finite set of triangular generators, in the sense that there is a family of triangular generators $\mathcal{F}$ such that the corresponding family $\bar{\mathcal{F}} \subset \mathcal{L}ag(X)$, obtained by forgetting the grading and the choices of primitives, is finite. We now fix such a family $\mathcal{F}$ of generators and assume that $\mathcal{F}$ is invariant under shifts and translations, and that the Lagrangians in $\bar{\mathcal{F}}$ are in general position (every two Lagrangians intersect transversely and there are no triple intersection points).

The proof of Corollary 3.7 is a consequence of Theorem 3.4 together with the invariance properties of the Fukaya TPC constructed earlier in this section, and is contained in the subsections below.

3.4.2.1 The Fukaya TPC revisited. Pick a shift- and translation-invariant family $\mathcal{X} \subset \mathcal{L}ag(X)'$ that contains $\mathcal{F}$. As before, we denote by $\bar{\mathcal{X}} \subset \mathcal{L}ag(X)$ the corresponding family of Lagrangians after forgetting the choices of primitives and grading. We assume that $\bar{\mathcal{X}}$ is finite and that its elements are in general position.

We recall the filtered Fukaya category $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$, where $\mathcal{P}$ is a choice of perturbation data, as in Theorem 3.12; see also Section 3.4.1.1. We already know that any two such categories, defined using two admissible perturbation data $\mathcal{P}$ and $\mathcal{P}'$, are filtered quasi-equivalent, in the sense that there are filtered A_∞-functors $\mathcal{F}uk(\mathcal{X}; \mathcal{P}) \to \mathcal{F}uk(\mathcal{X}; \mathcal{P}')$ that are the identity on objects and induce a (filtered) equivalence of the homological persistence categories.

As discussed in Section 3.2.1, there are two TPCs that one can associate to $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$. The first is $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ as in Theorem 3.1. This is obtained by considering the filtered modules $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))$ over $\mathcal{F}uk(\mathcal{X}; \mathcal{P})$. The category $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is the homological category

$$\mathcal{C}\mathcal{F}uk(\mathcal{X}) = H^0\big[\mathcal{F}uk(\mathcal{X}; \mathcal{P})^\nabla\big],$$

where $\mathcal{F}uk(\mathcal{X}; \mathcal{P})^\nabla$ is the smallest triangulated (with respect to weight-0 triangles) full subcategory of $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X}; \mathcal{P}))$ that contains the Yoneda modules $\mathcal{Y}(L)$ with $L \in \mathcal{X}$ and is closed under r-isomorphisms for all r, in the sense that if $j : M \to M'$ is an r-isomorphism of modules and $M \in \mathcal{F}uk(\mathcal{X}; \mathcal{P})^\nabla$, then $M' \in \mathcal{F}uk(\mathcal{X}; \mathcal{P})^\nabla$. Possibly more concretely, each object in this category is r-isomorphic, for some r, to a weight-0 iterated cone of Yoneda modules. Two such triangulated persistence categories, defined for different choices of perturbation data, are TPC equivalent (see also Remark 3.2), and thus we drop the reference to the perturbation data from the notation.

The second type of TPC will be denoted by $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ and is defined by

$$\mathcal{C}'\mathcal{F}uk(\mathcal{X}) = H^0[\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))], \tag{3.33}$$

where $\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$ is the category of filtered twisted complexes constructed from $\mathcal{F}uk(\mathcal{X};\mathcal{P})$; see Section 3.2.1. In this case too we can drop the reference to the choices of perturbation data, as any two such choices produce equivalent TPCs (again, these equivalences are not entirely canonical, but this will not have any impact on our further arguments).

There are filtered functors

$$\Theta : \mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P})) \to \mathrm{Fmod}[\mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))] \to \mathrm{Fmod}(\mathcal{F}uk(\mathcal{X};\mathcal{P})),$$

where the first arrow is the Yoneda embedding and the second is the pullback over the natural inclusion $\mathcal{F}uk(\mathcal{X};\mathcal{P}) \to \mathrm{Tw}(\mathcal{F}uk(\mathcal{X};\mathcal{P}))$. The composition Θ is a homologically full and faithful embedding, and it induces a full and faithful embedding of TPCs. The image of Θ lands inside $\mathcal{F}uk(\mathcal{X})^{\nabla}$ (actually inside the category denoted by $\mathcal{F}uk(\mathcal{X},\mathcal{P})^{\#}$ in Section 3.2.1), and thus we have an inclusion of TPCs:

$$\bar{\Theta} : \mathcal{C}'\mathcal{F}uk(\mathcal{X}) \hookrightarrow \mathcal{C}\mathcal{F}uk(\mathcal{X}).$$

By contrast with the unfiltered case, $\bar{\Theta}$ is not an equivalence of TPCs because, for it to be an equivalence of TPCs, each object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ needs to be 0-isomorphic to some object in $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ (see Definition 2.25) and, a priori, this might not happen (each object in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$ is r-isomorphic to some object in $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$, but possibly for $r > 0$).

3.4.2.2 Fragmentation metrics as in Theorem 3.4. We now focus on the pseudometric $D^{\mathcal{F}}$ as in Theorem 3.4. This is constructed by the general procedure described in the algebraic part of this memoir, by using the persistence triangular weight on $\mathcal{C}\mathcal{F}uk(\mathcal{X})_{\infty}$. Thus, this is a shift-invariant fragmentation pseudometric of the type $\widehat{\bar{d}^{\mathcal{F}}}(-,-)$, as described in Section 2.4.3.1, restricted to $\mathcal{X}$. To emphasize the relation of this pseudometric to the set $\mathcal{X}$ we will denote it by $D^{\mathcal{F}}_{\mathcal{X}}$. Because the pseudometric is shift invariant, it descends to $\bar{\mathcal{X}}$. As shown in Theorem 3.4, this pseudometric is independent of the choice of perturbation data $\mathcal{P}$.

Another way to construct a fragmentation-type pseudometric on $\bar{\mathcal{X}}$ is to apply the exact same construction to the triangulated persistence category $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ from (3.33). The resulting pseudometric, again defined on $\mathcal{X}$, will be denoted by $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$. It is again independent of the perturbation data used in its definition. If one is interested only in comparing objects in $\mathcal{X}$, this pseudometric is much more approachable from a computational point of view, because $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$ has fewer objects than $\mathcal{C}\mathcal{F}uk(\mathcal{X})$.

There is a simple relation between the two pseudometrics discussed above.

Lemma 3.27. *For any $L, L' \in \bar{\mathcal{X}}$, the pseudometrics $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$ and $D^{\mathcal{F}}_{\mathcal{X}}$ satisfy the inequalities*

$$\frac{1}{4}\,\bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq D^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq \bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L').$$

Proof. Indeed, the inequality on the right is obvious in view of the TPC inclusion $\bar{\Theta}$. The inequality on the left follows from the argument in Lemma 2.87. In fact, this argument shows that if we have a cone decomposition of $\mathcal{Y}(L)$ of total weight r in $\mathcal{C}\mathcal{F}uk(\mathcal{X})$, with linearization consisting of Yoneda modules of elements in $\mathcal{F}$ together with one term that is a shift/translate of $\mathcal{Y}(L')$, then there is another cone decomposition of $\mathcal{Y}(L)$ with a linearization of the same type such that all triangles are of weight 0 except the last triangle, which is of weight at most $4r$. But this means that all these triangles can be assumed to belong to $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$. This is clear for the weight-0 triangles because the image of Θ contains the Yoneda modules $\mathcal{Y}(L)$, $L \in \mathcal{X}$, and is closed with respect to taking cones over filtration-preserving maps. The last triangle corresponds to a strict exact triangle of weight $4r$ with the $4r$-isomorphism on the third term of the form $\phi : M_{n-1} \to \mathcal{Y}(L)$, where M_{n-1} is an object of $\mathcal{C}'\mathcal{F}uk(\mathcal{X})$. Because $\bar{\Theta}$ is full and faithful, we have that $\phi \in \hom_{\mathcal{C}'\mathcal{F}uk(\mathcal{X})}$. From this argument, it is easy to deduce the inequality on the left in the statement of the lemma, which concludes the proof. ■

3.4.2.3 Changing the set $\mathcal{X}$. Assume that $\mathcal{X}'$ is another family of marked Lagrangians which is invariant under shifts and translations and such that $\mathcal{X} \subset \mathcal{X}'$. Additionally, we assume that the family $\bar{\mathcal{X}}'$ is in general position.

Lemma 3.28. *Under the assumptions above, and for any two $L, L' \in \bar{\mathcal{X}}$, we have*

$$\bar{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') \geq \bar{D}^{\mathcal{F}}_{\mathcal{X}'}(L, L') \quad \text{and} \quad D^{\mathcal{F}}_{\mathcal{X}}(L, L') \leq D^{\mathcal{F}}_{\mathcal{X}'}(L, L').$$

Proof. We consider the category $\mathcal{F}uk(\mathcal{X}', \mathcal{P}')$ and we notice that the restriction of $\mathcal{P}'$ to $\mathcal{X}$ provides an allowable choice of perturbation data; see Theorem 3.12. Thus, there is an A_∞ filtered inclusion

$$\mathcal{F}uk(\mathcal{X}, \mathcal{P}') \to \mathcal{F}uk(\mathcal{X}', \mathcal{P}).$$

This inclusion induces a pullback TPC functor

$$qJ^* : \mathcal{C}\mathcal{F}uk(\mathcal{X}') \to \mathcal{C}\mathcal{F}uk(\mathcal{X}).$$

This is well defined because $\mathcal{F} \subset \mathcal{X} \subset \mathcal{X}'$ is a system of triangular generators for $D\mathcal{F}uk(X)$, which means that, in particular, any Yoneda module $\mathcal{Y}(L)$, $L \in \mathcal{X}'$, is r-isomorphic (in $\mathcal{C}\mathcal{F}uk(\mathcal{X}')$), for some $r \geq 0$, to a 0-weight iterated cone of Yoneda modules of elements from $\mathcal{F}$. This means that the pullback of $\mathcal{Y}(L)$ to $\mathrm{Fmod}(\mathcal{F}uk(\mathcal{X}, \mathcal{P}'))$ is an object of $\mathcal{C}\mathcal{F}uk(\mathcal{X})$.

The same A_∞ inclusion also induces a pushforward TPC functor

$$J_* : \mathcal{C}'\mathcal{F}uk(\mathcal{X}) \to \mathcal{C}'\mathcal{F}uk(\mathcal{X}'),$$

which is induced by the natural inclusion of twisted complexes.

As discussed before, our invariance statements imply that the pseudometrics $\bar{D}^{\mathcal{F}}_{\mathcal{X}}$ and $D^{\mathcal{F}}_{\mathcal{X}}$ do not depend on the choice of perturbation data. As a result, the fact that J_* is a TPC functor implies the first inequality in the lemma, and the fact that J^* is a TPC functor implies the second inequality. ∎

3.4.2.4 The pseudometric $\mathcal{D}^{\mathcal{F}}$ from Corollary 3.7. The construction of $\mathcal{D}^{\mathcal{F}}$ proceeds in two steps.

The first step is to consider again a family $\mathcal{X}$ as in the subsections above, together with two elements $L, L' \in \mathcal{L}ag(X)$. We define

$$\mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L') = \limsup_{\epsilon \to 0} D^{\mathcal{F}}_{\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}}(L_\epsilon, L'_\epsilon).$$

Here L_ϵ and L'_ϵ are Hamiltonian deformations of, respectively, L and L' through Hamiltonians of Hofer norm at most $\epsilon \geq 0$, such that the family $\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}$ is allowable for the definition of the Fukaya categories $\mathcal{F}uk(\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}; \mathcal{P})$. The lim sup is taken over all possible choices of such perturbations as ϵ goes to 0.

The second step is to put

$$\mathcal{D}^{\mathcal{F}}(L, L') = \sup_{\mathcal{X}} \mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L').$$

It is clear that $\mathcal{D}^{\mathcal{F}}$ is symmetric. We will see below, in Lemma 3.30, that $\mathcal{D}^{\mathcal{F}}$ is finite. However, before we get to that, we have the following result.

Lemma 3.29. *With the definition above, $\mathcal{D}^{\mathcal{F}}$ satisfies the triangle inequality.*

Proof. We fix three Lagrangians L, L', L'' in $\mathcal{L}ag(X)$. Fix also a family $\mathcal{X}$ as above. Consider L_n, L''_n in $\mathcal{L}ag(X)$ such that the family $\mathcal{X} \cup \bigcup_{n,m}\{L_n, L''_m\}$ is in general position in our usual sense: any couple of Lagrangians in the family intersect transversely and there are no triple intersection points (this choice is possible as there are only countable many transversality-type constraints). We also assume:

- $d_H(L, L_n) \leq \frac{1}{n}$ and $d_H(L'', L''_n) \leq \frac{1}{n}$, where $d_H(-,-)$ is the Hofer distance;
- $\lim_{n\to\infty} D^{\mathcal{F}}_{\mathcal{X} \cup \{L_n, L''_n\}}(L_n, L''_n) = \mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L'')$.

The lemma would follow if we prove that, for every $\delta > 0$,

$$\mathcal{D}^{\mathcal{F}}_{\mathcal{X}}(L, L'') \leq \mathcal{D}^{\mathcal{F}}(L, L') + \mathcal{D}^{\mathcal{F}}(L', L'') + 4\delta. \tag{3.34}$$

To show (3.34), we pick a sequence of Lagrangians L'_k such that the family $\mathcal{X} \cup \bigcup_{n,k,m}\{L_n, L'_k, L''_m\}$ is in general position and $d_H(L', L'_n) \leq \frac{1}{n}$. For any m, n, k, we have the inequalities

$$\begin{aligned} D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_m\}}(L_n, L''_m) &\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L_n, L''_m) \\ &\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L_n, L'_k) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_m\}}(L'_k, L''_m). \end{aligned}$$

The first inequality comes from Lemma 3.28 and the second is the triangle inequality for the fragmentation pseudometric $D^{\mathcal{F}}_{(-)}$. We will estimate separately the two terms on the right-hand side of this inequality.

Fix a natural number m_0. We can find $N_{m_0} \geq m_0$ such that for $n, k \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L_n, L'_k) \leq \mathcal{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L''_{m_0}\}}(L, L') + \delta \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta.$$

Thus, for $n, k \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L_n, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L'_k, L''_{m_0}),$$

and it remains to estimate the rightmost term. Using Lemma 3.28 again and the triangle inequality, we have

$$\begin{aligned} &D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0}\}}(L'_k, L''_{m_0}) \\ &\quad\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_{m_0}) \\ &\quad\leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_m) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L''_m, L''_{m_0}). \end{aligned}$$

All our fragmentation pseudometrics are bounded above by the Hofer norm and thus we have $D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L''_m, L''_{m_0}) \leq \frac{2}{m_0}$ as soon as $m \geq m_0$. We now consider $n \geq N_{m_0}$, and take k, m sufficiently large such that we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L'_k,L''_{m_0},L''_m\}}(L'_k, L''_m) \leq \mathcal{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L', L'') + \delta \leq \mathcal{D}^{\mathcal{F}}(L', L'') + \delta.$$

Putting things together, for our fixed (arbitrary) m_0 and any $n \geq N_{m_0}$ we have

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_n,L''_{m_0}\}}(L_n, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \delta + \mathcal{D}^{\mathcal{F}}(L', L'') + \delta + \frac{2}{m_0}. \tag{3.35}$$

An important remark is that inequality (3.35) applies to any fixed m_0 and any set $\mathcal{X}'$ that contains $\mathcal{X}$ (and is such that the family $\mathcal{X}' \cup \bigcup_{n,k,m}\{L_n, L'_k, L''_m\}$ is in general position), since the argument above applies verbatim to this situation. In this case the number N_{m_0} depends on m_0 but also, implicitly, on $\mathcal{X}'$. To make this dependence explicit we will denote it by $N_{m_0,\mathcal{X}'}$.

For any n, we have the triangle inequality:

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0}, L''_{m_0}) \\ \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n, L''_{m_0}) + D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0}, L_n).$$

Assuming $n \geq m_0$, the second term on the right-hand side is bounded above by $\frac{2}{m_0}$ and thus, using Lemma 3.28, we deduce

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L''_{m_0}\}}(L_{m_0}, L''_{m_0}) \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_{m_0}, L''_{m_0}) \\ \leq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n, L''_{m_0}) + \frac{2}{m_0}.$$

We now apply (3.35) to $\mathcal{X}' = \mathcal{X} \cup \{L_{m_0}\}$. We deduce that, for $n \geq N_{m_0,\mathcal{X}'}$,

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L_n,L''_{m_0}\}}(L_n, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \mathcal{D}^{\mathcal{F}}(L', L'') + 2\delta + \frac{2}{m_0}.$$

The last inequality implies

$$D^{\mathcal{F}}_{\mathcal{X}\cup\{L_{m_0},L''_{m_0}\}}(L_{m_0}, L''_{m_0}) \leq \mathcal{D}^{\mathcal{F}}(L, L') + \mathcal{D}^{\mathcal{F}}(L', L'') + 2\delta + \frac{4}{m_0}.$$

As this is true for an arbitrary choice of m_0, we deduce inequality (3.34). This concludes the proof. ∎

3.4.2.5 Properties of $\mathcal{D}^{\mathcal{F}}$. We know from Section 3.4.2.4 that $\mathcal{D}^{\mathcal{F}}$ is a pseudometric. In this subsection we will show that $\mathcal{D}^{\mathcal{F}}$ satisfies the other properties claimed in Corollary 3.7. The properties (i), (ii), and (iii), are in fact immediate consequences of the properties of the pseudometrics $D^{\mathcal{F}}_{\mathcal{X}}$ that appear in Theorem 3.4. Indeed, the estimates in this theorem do not depend on the set $\mathcal{X}$, and this easily implies the corresponding properties for $\mathcal{D}^{\mathcal{F}}$. A more delicate property is the following one.

Lemma 3.30. *The pseudometric $\mathcal{D}^{\mathcal{F}}$ is finite.*

Proof. This property follows by applying repeatedly Lemmas 3.27 and 3.28. Fix $L, L' \in \mathcal{L}ag(X)$, a family $\mathcal{X}$ as before, and perturbations L_ϵ, L'_ϵ such that $\mathcal{X} \cup \{L_\epsilon, L'_\epsilon\}$ is in general position and L_ϵ, L'_ϵ are ϵ-close to L and L', respectively, in the Hofer norm. We have

$$\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon, L'_\epsilon) \geq \bar{D}^{\mathcal{F}}_{\mathcal{X}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon, L'_\epsilon) \geq D^{\mathcal{F}}_{\mathcal{X}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon, L'_\epsilon).$$

Thus the argument reduces to showing that $\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L'_\epsilon\}}(L_\epsilon, L'_\epsilon)$ has a uniform upper bound when $\epsilon \to 0$. It is immediate to see that it is enough to find such a bound for

$$\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon, 0).$$

We fix some ϵ_0 and write, for $\epsilon \leq \epsilon_0$,

$$\begin{aligned}\frac{1}{4}\bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon,0) \leq D^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon\}}(L_\epsilon,0) &\leq D^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L_{\epsilon_0}\}}(L_\epsilon,0) \\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L_{\epsilon_0}\}}(L_\epsilon,0) \\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L_{\epsilon_0}\}}(L_{\epsilon_0},0) + \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L_{\epsilon_0}\}}(L_{\epsilon_0},L_\epsilon) \\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_\epsilon,L_{\epsilon_0}\}}(L_{\epsilon_0},0) + 2\epsilon_0 \\ &\leq \bar{D}^{\mathcal{F}}_{\mathcal{F}\cup\{L_{\epsilon_0}\}}(L_{\epsilon_0},0) + 2\epsilon_0.\end{aligned}$$

The first three inequalities come from Lemmas 3.27 and 3.28, and, the fourth from the triangle inequality. The next inequality is implied by the upper bound given by the Hofer norm. Finally, the last inequality comes from Lemma 3.28. ∎

The following result concludes the proof of Corollary 3.7.

Lemma 3.31. *If $\mathcal{F}'$ is a generic perturbation of $\mathcal{F}$ (in the sense that each element of $\mathcal{F}'$ is a small Hamiltonian perturbation of a corresponding element in $\mathcal{F}$ and the union $\mathcal{F} \cup \mathcal{F}'$ is in general position), then*

$$\mathcal{D}^{\mathcal{F},\mathcal{F}'} = \max\{\mathcal{D}^{\mathcal{F}}, \mathcal{D}^{\mathcal{F}'}\}$$

is non-degenerate.

Proof. The statement follows from point (ii) of Corollary 3.7, which has already been shown. From this point we deduce that $\mathcal{D}^{\mathcal{F},\mathcal{F}'}(L, L') = 0$ implies

$$\delta(L; L' \cup \textstyle\bigcup_{F\in\mathcal{F}} F) = 0 = \delta(L; L' \cup \bigcup_{F'\in\mathcal{F}'} F').$$

The definition of $\delta(-;-)$ is in (3.1). The first of the last two equalities means that there is no standard symplectic ball of any positive radius with its real part on L and that is disjoint from $L' \cup \bigcup_{F\in\mathcal{F}} F$. It follows that $L \subset L' \cup \bigcup_{F\in\mathcal{F}} F$. Of course the metric is symmetric, so we also have that $L' \subset L \cup \bigcup_{F\in\mathcal{F}} F$. Given that the same relations are valid for $\mathcal{F}'$, and that $\mathcal{F}'$ and $\mathcal{F}$ are in general position, we deduce that $L = L'$. ∎

3.5 The geometry behind TPCs

This section illustrates geometrically some of the TPC machinery. It contains two subsections. In Section 3.5.1 we explain how the theory of Lagrangian cobordism provides a concrete representation for the algebraic structures that are formalized in the language of TPCs. Some of the material presented in Section 3.5.1 is based on the theory developed in [7, 10], which is technically involved. Below, however, we have tried to avoid technicalities as much as possible in order to focus on the geometric ideas, at the expense of skipping some details. Most of these can be found in the above references.

In Section 3.5.2 we work out some examples where the estimates in Theorem 3.4 can be made concrete.

3.5.1 Lagrangian cobordism and weighted triangles

The theory of Lagrangian cobordism exhibits in a geometric way several key notions that are fundamental for the algebraic theory of TPCs. The purpose of this section is to provide some geometric interpretations of these notions – in particular, of weighted exact triangles – by using certain natural symplectic measurements associated to Lagrangian cobordisms. The geometric weights coming from geometry (such as those reflecting the shadows of cobordisms) are bigger than the algebraic weights discussed before in this memoir. The difficulty with using them in practice is that they depend on constructing specific cobordisms.

3.5.1.1 Background on Lagrangian cobordism. Let $(X, \omega = d\lambda)$ be a Liouville manifold, endowed with a given Liouville form λ. We endow $\mathbb{R}^2$ with the Liouville form $\lambda_{\mathbb{R}^2} = xdy$ and its associated standard symplectic structure $\omega_{\mathbb{R}^2} = d\lambda_{\mathbb{R}^2} = dx \wedge dy$. Let $\widetilde{X} := \mathbb{R}^2 \times X$, endowed with the Liouville form $\widetilde{\lambda} = \lambda_{\mathbb{R}^2} \oplus \lambda$ and the symplectic structure $\widetilde{\omega} = d(\widetilde{\lambda}) = \omega_{\mathbb{R}^2} \oplus \omega$. We denote by $\pi : \mathbb{R}^2 \times X \to X$ the projection.

Below we will assume known the notion of Lagrangian cobordism, as developed in [6, 7]. For simplicity, we will consider only *negative-ended* cobordisms $V \subset \mathbb{R}^2 \times X$, which means that V has only negative ends. Moreover, all the cobordisms considered below will be assumed to be exact with respect to the Liouville form $\widetilde{\lambda}$ and endowed with a given primitive $f_V : V \to \mathbb{R}$ of $\widetilde{\lambda}|_V$. Denote by $L_1, \ldots, L_k \subset X$ the Lagrangians corresponding to the ends of V and by $\ell_1, \ldots, \ell_k \subset \mathbb{R}^2$ the negative horizontal rays of V, so that V coincides at $-\infty$ with $(\ell_1 \times L_1) \coprod \cdots \coprod (\ell_k \times L_k)$. We remark that we adopt here the conventions from [7] regarding the ends of V, namely we always assume that the jth ray ℓ_j lies on the horizontal line $\{y = j\}$. Also, we allow some of the Lagrangians L_j to be void.

Note that $\lambda_{\mathbb{R}^2}|_{\ell_i} = 0$, hence $f_V|_{\ell_i \times L_i}$ is constant in the ℓ_i direction for all i. Therefore the Lagrangians $L_i \subset X$ are λ-exact and f_V induces well-defined primitives $f_{L_i} : L_i \to \mathbb{R}$ of $\lambda|_{L_i}$ for each i, namely $f_{L_i}(p) := f_V(z_0, p)$ for every $p \in L_i$, where z_0 is any point on ℓ_i.

3.5.1.2 Weakly filtered Fukaya categories and cobordism. As constructed in [8, 10], there is a weakly filtered Fukaya A_∞-category $W\mathcal{F}uk(X)$ of λ-exact Lagrangians whose objects are exact Lagrangians $L \subset X$ endowed with a primitive $f_L : L \to \mathbb{R}$

of $\lambda|_L$. The notation is ad hoc here to distinguish this category from the filtered versions constructed in Section 3.3.

Remark 3.32. The filtered A_∞ categories $\mathcal{F}uk(\mathcal{X};\mathcal{P})$ from Section 3.3 (see Theorem 3.12) are constructed under more restrictive assumptions than $W\mathcal{F}uk(X)$, since they are associated to only a finite number of geometric objects $\overline{\mathcal{X}}$. However, it is easy to see that, by choosing the perturbation data required to define the weakly filtered category $W\mathcal{F}uk(X)$ so that it extends the data $\mathcal{P}$, we have an embedding (of weakly filtered A_∞-categories)

$$\mathcal{F}uk(\mathcal{X};\mathcal{P}) \to W\mathcal{F}uk(X).$$

There is also a weakly filtered Fukaya A_∞-category of cobordisms $W\mathcal{F}uk(\mathbb{R}^2\times X)$ whose objects are negative-ended exact cobordisms $V \subset \mathbb{R}^2 \times X$ endowed with a primitive f_V of $\widetilde{\lambda}|_V$, as above.[1] We also have the dg-categories of weakly filtered A_∞-modules over each of the previous Fukaya categories, which we denote, respectively, by $\mathrm{mod}_{W\mathcal{F}uk(X)}$ and $\mathrm{mod}_{W\mathcal{F}uk(\mathbb{R}^2\times X)}$.

Below we will mostly concentrate on the chain complexes associated to various Lagrangians and modules, ignoring the higher-order A_∞-operations, and these are genuinely filtered. Thus, in this discussion the fact that the above categories are only weakly filtered rather than genuinely filtered will not a play an important role.

Let $\mathcal{Y} : W\mathcal{F}uk(X) \to \mathrm{mod}_{W\mathcal{F}uk(X)}$ be the Yoneda embedding (in the framework of weakly filtered A_∞-categories), and let $W\mathcal{F}uk(X)^\nabla \subset \mathrm{mod}_{W\mathcal{F}uk(X)}$ be the triangulated closure of the image of $\mathcal{Y}$. We denote by $\mathcal{C} = \mathcal{P}H(\mathcal{F}uk(X)^\nabla)$ the persistence homological category associated to $\mathcal{F}uk(X)^\nabla$. This is not a TPC due to the difference between “filtered” and “weakly filtered”; however, with this distinction kept in mind, its properties closely mimic those of a TPC. To understand the difference, while in a TPC the composition of two morphisms f of shift r and g of shift s is a morphism $f \circ g$ of shift $r + s$, in the weakly filtered case the composition $f \circ g$ has shift $r + s + \epsilon^{\mu_2}$, where the error term μ_2 is part of the structural data associated to the *weakly* filtered structure of $W\mathcal{F}uk(X)$.

By a slight abuse of notation we will denote the Yoneda module $\mathcal{Y}(L)$ of a Lagrangian $L \in \mathrm{Obj}(W\mathcal{F}uk(X))$ also by L.

There is also a Yoneda embedding $W\mathcal{F}uk(\mathbb{R}^2 \times X) \to \mathrm{mod}_{W\mathcal{F}uk(\mathbb{R}^2\times X)}$ and we will typically denote the Yoneda modules corresponding to cobordisms by calligraphic letters, e.g., the Yoneda module corresponding to $V \in \mathrm{Obj}(\mathcal{F}uk(\mathbb{R}^2 \times X))$ will be denoted by $\mathcal{V}$.

Under additional assumptions on X, on the Lagrangians taken as the objects of $W\mathcal{F}uk(X)$, and on the Lagrangian cobordisms of $\mathcal{F}uk(\mathbb{R}^2 \times X)$, one can set up

[1] For technical reasons one needs to enlarge the set of objects in $\mathcal{F}uk(\mathbb{R}^2 \times X)$ to contain also objects of the type $\gamma \times L$, where $\gamma \subset \mathbb{R}^2$ is a curve which, outside a compact set, is either vertical or coincides with horizontal ends with y-value being $l \pm \frac{1}{10}$, where $l \in \mathbb{Z}$.

a graded theory, endowing the morphisms in $\mathcal{F}uk(X)$ and $\mathcal{F}uk(\mathbb{R}^2 \times X)$ with a $\mathbb{Z}$-grading and the categories with grading-translation functors. See [53] for the case of $W\mathcal{F}uk(X)$ and [38] for grading in the framework of cobordisms. In what follows we will not work explicitly in a graded setting, but whenever possible we will indicate how grading fits in various constructions.

3.5.1.3 Iterated cones associated to cobordisms. Let $\gamma \subset \mathbb{R}^2$ be an oriented[2] plane curve that is the image of a proper embedding of $\mathbb{R}$ into $\mathbb{R}^2$. Viewing $\gamma \subset \mathbb{R}^2$ as an exact Lagrangian we fix a primitive f_γ of $\lambda_{\mathbb{R}^2}|_\gamma$. Given an exact Lagrangian $L \subset X$, consider the exact Lagrangian $\gamma \times L \subset \mathbb{R}^2 \times X$ and take $f_{\gamma,L} : \gamma \times L \to \mathbb{R}$, $f_{\gamma,L}(z, p) = f_\gamma(z) + f_L(p)$ for the primitive of $\widetilde{\lambda}|_{\gamma\times L}$. From now on we will make the following additional assumptions on γ. The ends of γ will be assumed to coincide with a pair of rays ℓ', ℓ'' each of which is allowed to be either horizontal or vertical. Moreover, in the case of a horizontal ray, we assume that its y-coordinate lies in $\mathbb{Z} \pm \frac{1}{10}$, and in the case of a vertical ray we assume that its x-coordinate is 0.

Below we will mainly work with the following two types of such curves. The first one is $\gamma^\uparrow = \{x = 0\} \subset \mathbb{R}^2$ (i.e., the x-axis with its standard orientation) and we take $f_{\gamma^\uparrow} \equiv 0$. Then for every exact Lagrangian $L \subset X$ we can identify $f_{\gamma^\uparrow,L}$ with f_L in the obvious way.

The second type is the curve $\gamma_{i,j}$, where $i \leq j$ are two integers, depicted in Figure 3.11 and oriented by going from the lower horizontal end to the upper horizontal end. Note that by taking $\gamma_{i,j}$ close enough to the dotted polygonal curve in Figure 3.11 we can assume that $\lambda_{\mathbb{R}^2}|_{\gamma_{i,j}}$ is very close to 0. We fix the primitive $f_{\gamma_{i,j}}$ to be the one that vanishes on the vertical part of $\gamma_{i,j}$.

Let $\gamma \subset \mathbb{R}^2$ and f_γ be as above. Following [7, Section 4.2] and [10, Section 3.6], there is a (weakly) filtered A_∞-functor $\mathcal{I}_\gamma : W\mathcal{F}uk(X) \to W\mathcal{F}uk(\mathbb{R}^2 \times X)$, called an inclusion functor, which sends the object $L \in \mathrm{Obj}(W\mathcal{F}uk(X))$ to the object $\gamma \times L \in \mathrm{Obj}(\mathcal{F}uk(\mathbb{R}^2 \times X))$. The first-order component of $\mathcal{I}_\gamma$ is a chain map

$$(\mathcal{I}_\gamma)_1 : \mathrm{CF}(N_0, N_1) \to \mathrm{CF}(\gamma \times N_0, \gamma \times N_1),$$

defined for all exact Lagrangians N_0, N_1, which induces an isomorphism in homology. Note that, since $\gamma \times N_0$ and $\gamma \times N_1$ do not intersect transversely (unless we have $N_0 \cap N_1 = \emptyset$), we need to use here Floer data with non-trivial Hamiltonians that also involve a component in the $\mathbb{R}^2$-direction. We skip these details here and refer the reader to [7, Section 4.2] and [10, pp. 68–69] for the precise details. The higher-order components $(\mathcal{I}_\gamma)_d$, $d \geq 2$, of $\mathcal{I}_\gamma$ are defined to be 0.

[2]The orientation on γ is necessary in order to set up a graded Floer theory and also to work with coefficients over rings of characteristic $\neq 2$. Here we work with $\mathbb{Z}_2$-coefficients; therefore, if one wants to ignore the grading then the orientation of γ becomes irrelevant.

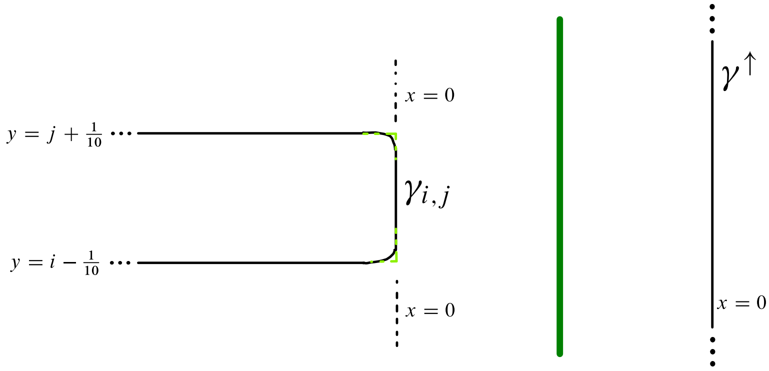

Figure 3.11. The curves $\gamma_{i,j}$ and $\gamma^{\uparrow}$.

Let $V \subset \mathbb{R}^2 \times X$ be a Lagrangian cobordism. Denote by $\mathcal{V}$ the (weakly) filtered Yoneda module of V and consider the pullback module $\mathcal{I}_\gamma^* \mathcal{V}$. Note that for every exact Lagrangian $N \subset X$ we have

$$\mathcal{I}_\gamma^* \mathcal{V}(N) = \mathrm{CF}(\gamma \times N, V)$$

as *filtered chain complexes*. (The filtrations are induced by f_V, $f_{\gamma,N}$, and by the Floer data in case it is not trivial.)

Assume that the ends of V are $L_1, \dots, L_k$ and, moreover, that V is cylindrical over $(-\infty, \delta] \times \mathbb{R}$ for some $\delta > 0$. (This can always be achieved by a suitable translation along the x-axis.) Fix $1 \le i \le k - 1$ and consider the curve $\gamma_{i,i+1}$ and the pullback (weakly) filtered module $\mathcal{I}_{\gamma_{i,i+1}}^* \mathcal{V}$. The cobordism V gives rise to a module homomorphism

$$\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$$

which preserves action filtrations and such that

$$\mathcal{I}_{\gamma_{i,i+1}}^* \mathcal{V} = T^{d_i} \mathrm{Cone}(L_{i+1} \xrightarrow{\Gamma_{V,\gamma_{i,i+1}}} L_i) \tag{3.36}$$

as (weakly) filtered modules. Here T stands for the grading-translation functor, and the amount $d_i \in \mathbb{Z}$ by which we translate depends only on i. We will be more precise about the values of d_i later on. The references for the construction of the map $\Gamma_{V,\gamma_{i,i+1}}$ are the following. For the unfiltered case, see [7, Section 4.4 and Proposition 4.4.1]. The map $\Gamma_{V,\gamma_{i,i+1}}$ is constructed on page 1805 of that paper, where it is denoted by ϕ_j; see also Proposition 4.4.3 there. The weakly filtered case is treated in [10]; see Proposition 3.5 and its proof, pp. 73–76. More relevant background material on inclusion functors and iterated cones can be found in Sections 3.6 and 3.7 of that paper. Note that here we are working in a strictly filtered setting (which is a special case of the weakly filtered case), and this simplifies many of the arguments from [10]. In addition to these references, we provide below in Section 3.5.1.4 an outline of the construction of $\Gamma_{V,\gamma_{i,i+1}}$, avoiding technicalities.

The map $\Gamma_{V,\gamma_{i,i+1}}$ is canonically defined by V and $\gamma_{i,i+1}$, up to a boundary in the chain complex $\hom^{\leq 0}_{\mathrm{mod}_{\mathcal{F}uk(M)}}(L_{i+1}, L_i)$. Therefore it gives rise to a well-defined morphism in the homological persistence category $\mathcal{C}_0 = H\big(\hom^{\leq 0}_{\mathrm{mod}_{\mathcal{F}uk(M)}}(L_{i+1}, L_i)\big)$, which by abuse of notation we still denote by $\Gamma_{V,\gamma_{i,i+1}} \in \hom_{\mathcal{C}_0}(L_{i+1}, L_i)$.

The above can be generalized to several consecutive ends in a row as follows. Fix $1 \leq i \leq j \leq k$. The pullback module $\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$ can be identified with an iterated cone of the type

$$\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V} = \mathrm{Cone}\big(L_j \longrightarrow \mathrm{Cone}(L_{j-1} \longrightarrow \ldots \longrightarrow \mathrm{Cone}(L_{i+1} \longrightarrow L_i)\ldots)\big), \quad (3.37)$$

where, similarly to the case $\Gamma_{V,\gamma_{i,i+1}}$, all the maps in the iterated cone are module homomorphisms that preserve filtrations. See Figure 3.12. The references given above for the construction of $\Gamma_{V,\gamma_{i,i+1}}$ are relevant also for the construction of (3.37).

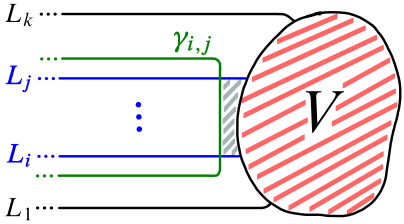

Figure 3.12. The module $\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$.

Note that there are some grading translations in (3.37) that we have ignored. We will be more precise about this point later on, when we consider iterated cones involving three objects.

Remark 3.33. If V has k ends, then for every $i \leq 1$ and $k \leq l$ we have $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V} = \mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$.

Finally, fix $1 \leq i \leq l < j \leq k$ and consider the two modules $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}$. There is a module homomorphism $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ which preserves filtrations. Note that we have

$$\mathcal{I}^*_{\gamma_{i,j}}\mathcal{V} = T^{d_{i,l,j}}\mathrm{Cone}\big(\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \xrightarrow{\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}} \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}\big)$$

for some $d_{i,l,j} \in \mathbb{Z}$. See Figure 3.13. While the construction of $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ does not explicitly appear in the references mentioned after (3.36), it can be easily deduced from the material in those papers. See also Section 3.5.1.4 below.

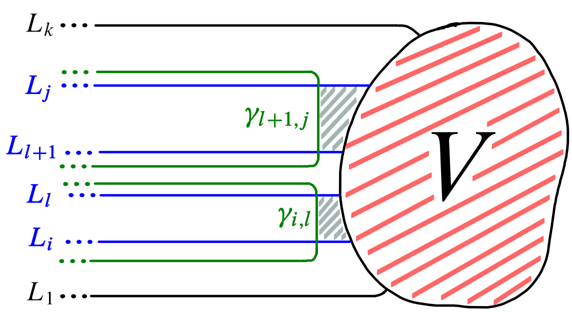

Figure 3.13. The modules $\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$.

3.5.1.4 Module maps induced by cobordisms. The purpose of this section is to outline the constructions of the module maps $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$ and $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}$ from Section 3.5.1.3. We will not give a fully rigorous account of the subject here, in an attempt to avoid technicalities as much as possible. Full details can be found in the references given after (3.36).

We begin with the map $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$. We will first explain how to construct the first-order component $(\Gamma_{V,\gamma_{i,i+1}})_1$ of this map.

Consider the pullback module $\mathcal{M} := \mathcal{I}^*_{\gamma_{i,i+1}}\mathcal{V}$ and Figure 3.14. For every exact Lagrangian N we have the following equalities of vector spaces:

$$\mathcal{M}(N) = \mathrm{CF}(\gamma_{i,i+1} \times N, V) = \mathrm{CF}(N, L_{i+1}) \oplus \mathrm{CF}(N, L_i). \tag{3.38}$$

In terms of Figure 3.14, the first summand corresponds to the intersection points $N \cap L_{i+1}$ lying above the point P, and the second summand to the intersection points $N \cap L_i$ lying above Q. For the sake of illustration, we have made here several simplifying assumptions (which cannot really be made in general). Namely, that N intersects both L_i and L_{i+1} transversely and that we can take the Floer and perturbation data for $W\mathcal{F}uk(\mathbb{R}^2 \times X)$ and $W\mathcal{F}uk(X)$ to have 0 Hamiltonian terms.

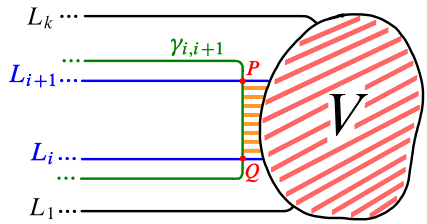

Figure 3.14. The module map $\Gamma_{V,\gamma_{i,i+1}} : L_{i+1} \to L_i$. The projection to $\mathbb{R}^2$ of the strips of type (PQ) is depicted as the orange horizontally-striped region.

Next we consider the differential $\mu_1^{\mathcal{M}}$ of this module. Again, for simplicity, assume that the almost complex structure J in the Floer data for $CF(\gamma_{i,i+1} \times N, V)$ is chosen so that the projection $\pi : \mathbb{R}^2 \times X \to \mathbb{R}^2$ is (J, i)-holomorphic, where i is the standard complex structure on $\mathbb{R}^2 \cong \mathbb{C}$. To describe $\mu_1^{\mathcal{M}}$ we need to consider Floer strips contributing to the differential of $CF(\gamma_{i,i+1} \times N, V)$. A priori, these are of four types:

(PP) Strips going from intersection points above P to points above P.

(QQ) Strips going from intersection points above Q to points above Q.

(PQ) Strips going from intersection points above P to points above Q.

(QP) Strips going from intersection points above Q to points above P.

Our assumptions on J and the Hamiltonians in the Floer data imply that all the PP and QQ strips have constant projection to $\mathbb{R}^2$, hence completely lie in $\{P\} \times X$ and $\{Q\} \times X$, respectively. Moreover, when viewed as strips in X, they are in one-to-one correspondence with the Floer strips that contribute to the differentials on $CF(N, L_{i+1})$ and $CF(N, L_i)$, respectively.

Standard arguments based on complex analysis in the plane (e.g., the open mapping theorem as used in [6, Section 4], [7, Sections 3-4]) imply that there are no Floer strips of type (QP). However, strips of type (PQ) may definitely exist and we write their contribution to the Floer complex as a linear map: $\phi_1 : CF(N, L_{i+1}) \to CF(N, L_i)$, which is based on counting strips emanating from an intersection point above P to an intersection point above Q.

Summing up, the differential $\mu_1^{\mathcal{M}}$ can be written as

$$\mu_1^{\mathcal{M}}(x^P, x^Q) = \big(\mu_1(x^P), \mu_1(x^Q) + \phi_1(x^P)\big).$$

Here, x^P, x^Q are intersection points corresponding to the first and second summands in (3.38), and μ_1 stands for the Floer differentials coming from $W\mathcal{F}uk(X)$. The fact that $\mu_1^{\mathcal{M}}$ is a differential implies that ϕ_1 is a chain map and, moreover, that $\mathcal{M}(N) = \mathrm{Cone}(L_{i+1} \xrightarrow{\phi_1} L_i)$ as chain complexes. We define the first-order component of $\Gamma_{V,\gamma_{i,i+1}}$ to be $(\Gamma_{V,\gamma_{i,i+1}})_1 := \phi_1$.

The construction of the higher-order components of $\Gamma_{V,\gamma_{i,i+1}}$ is similar, though technically more involved. To construct the maps $(\Gamma_{V,\gamma_{i,i+1}})_d$, $d \geq 2$, we need to analyze the μ_d-operations of the module $\mathcal{M}$. Fix d exact Lagrangians $N_0, \dots, N_{d-1}$ in X and $\underline{y} = (y_1, \dots, y_{d-1})$ with $y_k \in CF(N_{k-1}, N_k)$. Since $\mathcal{I}_{\gamma_{i,i+1}}$ has vanishing higher-order components, we have

$$\mu_d^{\mathcal{M}}(\underline{y}, (x^P, x^Q)) = \mu_d^{\mathbb{R}^2 \times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \dots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1}), (x^P, x^Q)\big), \tag{3.39}$$

where $\mu_d^{\mathbb{R}^2 \times X}$ is the dth-order A_∞-operation in $W\mathcal{F}uk(\mathbb{R}^2 \times X)$ and the points $x^P \in CF(N_0, L_{i+1})$, $x^Q \in CF(N_0, L_i)$ are as before.

The right-hand side of equation (3.39) counts Floer d-polygons with "edges" on the Lagrangians $\gamma_{i,i+1} \times N_0, \dots, \gamma_{i,i+1} \times N_{d-1}, V$. By arguments similar to those used for $d = 1$, one shows that there are no Floer polygons with entry points in $(\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \dots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1}), x_Q$ and exit points lying above P. Consequently, (3.39) has the shape

$$\mu_d^{\mathcal{M}}(\underline{y}, (x^P, x^Q)) = \Big(\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^P\big), \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^Q\big) + \phi_d\big(\underline{y}, x^P\big)\Big) \tag{3.40}$$

with respect to the splitting $\mathcal{M}(N_0) = \mathrm{CF}(N_0, L_{i+1}) \oplus \mathrm{CF}(N_0, L_i)$ used before. Here $(\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y})$ stands for $\big((\mathcal{I}_{\gamma_{i,i+1}})_1(y_1), \dots, (\mathcal{I}_{\gamma_{i,i+1}})_1(y_{d-1})\big)$. Thus $\phi_d(\underline{y}, x^P)$ counts Floer polygons in $\mathbb{R}^2 \times X$ with entry points $(\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y})$, x^P and an exit point over Q. The dth-order component $(\Gamma_{V,\gamma_{i,i+1}})_d$ of the desired map $\Gamma_{V,\gamma_{i,i+1}}$ is the multilinear map ϕ_d.

It remains to explain why $\mathcal{M} = \mathcal{I}^*_{\gamma_{i,i+1}}\mathcal{V}$ is the mapping cone (in the A_∞-sense) of the map $\Gamma_{V,\gamma_{i,i+1}}$. Consider the other two terms on the right-hand side of (3.40), namely

$$\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^P\big) \quad \text{and} \quad \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i+1}})_1(\underline{y}), x^Q\big).$$

These two terms can be identified with

$$\mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i+1,i+1}})_1(\underline{y}), x^P\big) \quad \text{and} \quad \mu_d^{\mathbb{R}^2\times X}\big((\mathcal{I}_{\gamma_{i,i}})_1(\underline{y}), x^Q\big),$$

respectively (note we are using now the curves $\gamma_{i+1,i+1}$ and $\gamma_{i,i}$, and not $\gamma_{i,i+1}$). In other words, the preceding two expressions can be identified with the μ_d-operations of the pullback modules $\mathcal{I}^*_{\gamma_{i+1,i+1}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,i}}\mathcal{V}$ applied to $(\underline{y}, x^P)$ and $(\underline{y}, x^Q)$, respectively.

Now, it follows from [7, Section 4.2] that the pullback modules $\mathcal{I}^*_{\gamma_{i+1,i+1}}\mathcal{V}$ and $\mathcal{I}^*_{\gamma_{i,i}}\mathcal{V}$ are quasi-isomorphic to the Yoneda modules of L_{i+1} and L_i, respectively. Therefore, up to grading translation, we obtain $\mathcal{M} = \mathrm{Cone}(L_{i+1} \xrightarrow{\phi} L_i)$, where $\phi = \{\phi_d\}_{d\geq 1}$, which proves (3.36). This concludes our rough outline of the construction of the map $\Gamma_{V,\gamma_{i,i+1}}$.

The definition of the maps $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is similar to the above and we will just go over the main points of the construction. Put $\mathcal{Q} = \mathcal{I}^*_{\gamma_{i,j}}\mathcal{V}$ and consider Figure 3.15. It is not hard to show that for every exact Lagrangian N we have

$$\mathcal{Q}(N) = \mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}(N) \oplus \mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}(N) \tag{3.41}$$

as vector spaces. Elements of the first summand can be written as $\underline{x}^P = (x^{P_j}, \dots, x^{P_{l+1}})$, with $x^{P_j} \in \mathrm{CF}(N, L_k)$ viewed as lying above the point P_k in Figure 3.15. Similarly, elements of the second summand of (3.41) can be written as $\underline{x}^Q = (x^{Q_i}, \dots, x^{Q_l})$. The differential $\mu_1^{\mathcal{Q}}$ of this module turns out to have the following shape:

$$\mu_1^{\mathcal{Q}}(\underline{x}^P, \underline{x}^Q) = \Big(\mu_1^{\mathcal{I}^*_{\gamma_{l+1,j}}\mathcal{V}}(\underline{x}^P), \mu_1^{\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}}(\underline{x}^Q) + \psi_1(\underline{x}^P)\Big),$$

where $\psi_1 : \mathcal{I}^*_{\gamma_{l+1,j}} \mathcal{V}(N) \to \mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}(N)$ is a linear map. The reason for this is similar to what has been explained earlier for the module $\mathcal{M} = \mathcal{I}^*_{\gamma_{i,i+1}} \mathcal{V}$. Namely, there cannot be any Floer strips connecting Q-type points with P-type points. The term $\psi_1(\underline{x}^P)$ counts Floer strips (in $\mathbb{R}^2 \times X$), with one boundary on $\gamma_{i,j} \times N$ and the other boundary on V, emanating from any entry x^{P_k} of $\underline{x}^P$ and going to some entry x^{Q_m} of $\underline{x}^Q$. In terms of Figure 3.15, the projections of, e.g., the strips that go from $x^{P_{l+1}}$ to x^{Q_l} are depicted in light blue (oblique stripes). The projections of the strips corresponding to $\mu_1^{\mathcal{I}_{\gamma_{l+1,j}}{}^*\mathcal{V}}$ are in orange (horizontal stripes), and those corresponding to $\mu_1^{\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}}$ $\mu_1^{\mathcal{I}^*_{\gamma_{i,l}}\mathcal{V}}$ are in purple (vertical stripes).

The first-order component of the desired map $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is defined to be the map ψ_1. The construction of the higher-order components of $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}}$ is analogous to the case of $\Gamma_{V,\gamma_{i,i+1}}$ discussed earlier.

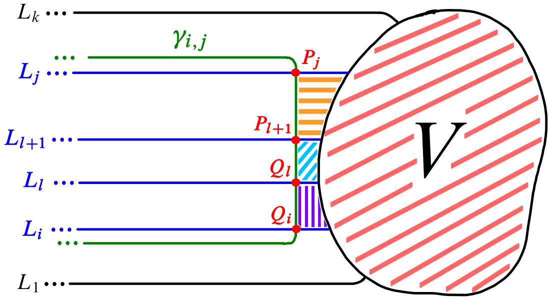

Figure 3.15. The module map $\Gamma_{V,\gamma_{i,l},\gamma_{l+1,j}} : \mathcal{I}^*_{\gamma_{l+1,j}} \mathcal{V} \to \mathcal{I}^*_{\gamma_{i,l}} \mathcal{V}$.

3.5.1.5 Shadow of cobordisms. Given a Lagrangian cobordism $V \subset \mathbb{R}^2 \times X$, we define its *outline* [24] as

$$\mathrm{out}(V) := \mathbb{R}^2 \setminus \mathcal{U},$$

where $\mathcal{U} \subset \mathbb{R}^2 \setminus \pi(V)$ is the union of all the *unbounded* connected components of $\mathbb{R}^2 \setminus \pi(V)$. An important measurement associated to V is its *shadow* [10, 24], $\mathcal{S}(V)$:

$$\mathcal{S}(V) := \mathrm{Area}(\mathrm{out}(V)).$$

Note that $\mathrm{out}(V) \subset \mathbb{R}^2$ is a measurable set, hence $\mathcal{S}(V)$ is well defined. The shadow plays a central role in defining cobordism-related metrics on spaces of Lagrangians [10].

For the purposes of this section, it will be easier to work with a slightly different variant of the shadow, which we call the *exterior shadow*. Fix a rectangle $Q \subset \mathbb{R}^2$ large enough that $\mathrm{out}(V) \setminus Q$ consists of only horizontal rays, and write

$$\mathrm{out}_Q(V) = Q \cap \mathrm{out}(V).$$

Define the *exterior shadow* of V to be

$$\mathcal{S}_e(V) = \inf\big\{A \mid \exists \text{ a smooth embedding } \varphi : B \to \mathbb{R}^2, \text{ with } \text{image}(\varphi) \supset \text{out}_Q(V) \text{ and } \text{Area}(\text{image}(\varphi)) \geq A\big\}.$$

Here $B \subset \mathbb{R}^2$ stands for the closed 2-dimensional unit disk. It is not hard to see that $\mathcal{S}_e(V)$ is independent of the choice of the rectangle Q.

Obviously, we have $\mathcal{S}_e(V) \geq \mathcal{S}(V)$ (because $\mathcal{S}(V) = \text{Area}(\text{out}_Q(V))$). But in fact, we actually have $\mathcal{S}_e(V) = \mathcal{S}(V)$. Since this statement has not been proved in full in the literature (though see [24, p. 33] for a related partial argument), we include in Section 3.5.1.11 below a sketch of a proof that $\mathcal{S}_e(V) = \mathcal{S}(V)$.

Previous papers on the subject used the shadow rather than the exterior shadow. However, for the rest of this section, whose purpose is mainly illustrative, we opt for the exterior shadow since it is more intuitive to work with.

3.5.1.6 *r*-acyclic objects. Let $V \subset \mathbb{R}^2 \times X$ be a cobordism with ends $L_1, \ldots, L_k$ such that V is cylindrical over $(-\infty, \delta] \times \mathbb{R}$ for some $\delta > 0$. Let $\mathcal{S}_e(V)$ be the exterior shadow of V and denote by σ the area of the region to the right of $\gamma^{\uparrow}$ enclosed between $\gamma^{\uparrow}$ and the projection to $\mathbb{R}^2$ of the non-cylindrical part of V. See Figure 3.16.

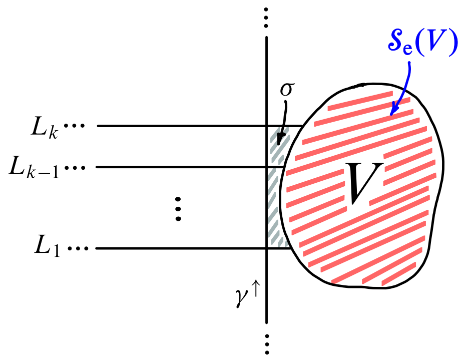

Figure 3.16. The projection of a cobordism to $\mathbb{R}^2$, its (exterior) shadow $\mathcal{S}_e(V)$ and the area σ.

The module $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V} = \mathcal{I}^*_{\gamma_{1,k}}\mathcal{V}$ (which also has the description (3.37) with $i = 1$, $j = k$) is r-acyclic, where $r := \mathcal{S}_e(V) + \sigma$. This can be easily seen from the fact that V can be deformed to a cobordism W which is disjoint from $\gamma^{\uparrow} \times X$, via a compactly supported Hamiltonian isotopy whose Hofer length is at most r. Standard Floer theory then implies that $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$ is acyclic of boundary depth at most r. In the terminology used in this memoir, this means that the object $\mathcal{I}^*_{\gamma^{\uparrow}}\mathcal{V}$ is r-acyclic.

Remark 3.34. The area summand σ that adds to the exterior shadow of V in the quantity r can be made arbitrarily small at the expense of applying appropriate shifts to each of the ends L_i of V. One way to do this is to replace the curve $\gamma^{\uparrow}$ by a curve $\gamma_V^{\uparrow}$ that coincides with $\gamma^{\uparrow}$ outside a compact subset and such that $\gamma_V^{\uparrow}$ approximates the shape of the projection of the non-cylindrical part of V in such a way that the area σ' enclosed between $\gamma_V^{\uparrow}$ and $\pi(V)$ is small. See Figure 3.17. One can apply a similar modification to the curves $\gamma_{i,j}$. Note that, in contrast to $f_{\gamma^{\uparrow}}$, the primitive $f_{\gamma_V^{\uparrow}}$ can no longer be assumed to be 0 (a similar remark applies to the primitives of the modifications of $\gamma_{i,j}$). As a result, the cone decompositions (3.37) associated to the pullback modules $\mathcal{I}^*_{\gamma_{i,j}}$ will have the same shape but each of the Lagrangians $L_i, \dots, L_j$ will gain a different shift in action. Note that this will also result in "tighter" weighted exact triangles than the ones we obtain below, in terms of the weights and the various shifts on the objects forming these triangles.

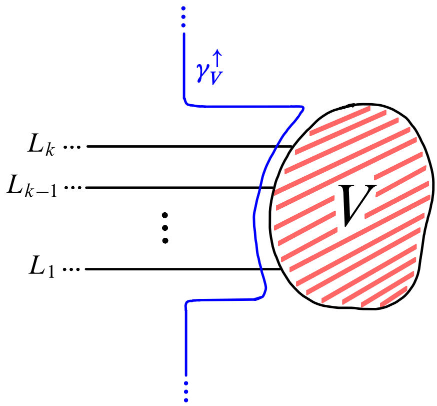

Figure 3.17. Replacing $\gamma^{\uparrow}$ by a curve $\gamma_V^{\uparrow}$ better approximating the shape of V.

To simplify the exposition, below we will not make these modifications and will stick to the curves $\gamma^{\uparrow}$ and $\gamma_{i,j}$ as defined above, at the expense of σ not being necessarily small and the weights of the triangles not being necessarily optimal.

3.5.1.7 *r*-isomorphisms. We begin by visualizing the canonical map $\eta_r^L : \Sigma^r L \to L$, where $L \subset X$ is an exact Lagrangian. Consider the curve $\gamma \subset \mathbb{R}^2$ depicted in Figure 3.18, and let r be the area enclosed between γ and $\gamma^{\uparrow}$. Let f_γ be the unique primitive of $\lambda_{\mathbb{R}^2}|_\gamma$ that vanishes along the lower end of γ. Note that $f_\gamma \equiv r$ along the upper end of γ. Let $V = L \times \gamma$ and set $f_V := f_{\gamma,L}$. Therefore, the primitives induced by V on its ends are as follows: the primitive on the lower end coincides with f_L, while the primitive on the upper end coincides with $f_L + r$. In other words, the cobordism V has ends L and $\Sigma^r L$. Moreover, the map $\Gamma_{V,\gamma_{1,2}} : \Sigma^r L \to L$ induced by V and $\gamma_{1,2}$ is precisely η_r^L.

Another source of geometric r-isomorphisms comes from Hamiltonian isotopies. Let ϕ_t^H, $t \in [0, 1]$, be a Hamiltonian isotopy and let $L \subset X$ be an exact Lagrangian.

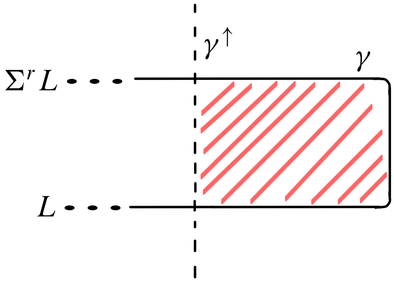

Figure 3.18. The cobordism inducing $\eta_r^L : \Sigma^r L \to L$.

The Lagrangian suspension construction gives rise to an exact Lagrangian cobordism between L and $\phi_1^H(L)$. After bending the ends of that cobordism so that they become negative, one obtains a Lagrangian cobordism whose negative ends are L and $\phi_1^H(L)$. See Figure 3.19. The primitive f_V on V is uniquely defined by the requirement that f_V coincides with f_L on the lower end of V. The exterior shadow $\mathcal{S}_e(V)$ of this cobordism equals the Hofer length of the isotopy $\{\phi_t^H\}_{t\in[0,1]}$, and we obtain an r-isomorphism $\phi_1^H(L) \to L$.

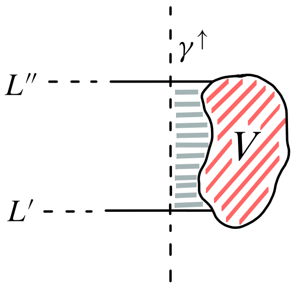

Figure 3.19. The Lagrangian suspension of a Hamiltonian isotopy, after bending the ends.

More generally, let V be a Lagrangian cobordism with ends $L_1, \dots, L_k$. Let $r = \mathcal{S}_e(V) + \sigma$ as above. Fix $1 \le l < k$. As explained above we have, up to an overall translation in grading,

$$\mathcal{I}^*_{\gamma^{\uparrow}} \mathcal{V} = \mathcal{I}^*_{\gamma_{1,k}} \mathcal{V} = \mathrm{Cone}\big(\mathcal{I}^*_{\gamma_{l+1,k}} \mathcal{V} \xrightarrow{\Gamma_{V,\gamma_{1,l},\gamma_{l+1,k}}} \mathcal{I}^*_{\gamma_{1,l}} \mathcal{V}\big),$$

and since $\mathcal{I}^*_{\gamma^{\uparrow}} \mathcal{V}$ is r-acyclic, the map $\Gamma_{V,\gamma_{1,l},\gamma_{l+1,k}} : \mathcal{I}^*_{\gamma_{l+1,k}} \mathcal{V} \to \mathcal{I}^*_{\gamma_{1,l}} \mathcal{V}$ is an r-isomorphism.

3.5.1.8 Weighted exact triangles. Let $V \subset \mathbb{R}^2 \times X$ be a Lagrangian cobordism with ends $L_1, \dots, L_k$. Let $V' \subset \mathbb{R}^2 \times X$ be the cobordism obtained from V by bending the upper end L_k clockwise around V so that it goes beyond the end L_1, as in Figure 3.20. To obtain a cobordism according to our conventions, we need to further shift V' upwards by one so that its lower end has y-coordinate 1 (instead of 0). Clearly, V' is also exact and $\mathcal{S}_e(V') = \mathcal{S}_e(V)$. We fix the primitive $f_{V'}$ for V' to be the unique one that coincides with f_V on the ends $L_1, \dots, L_{k-1}$. A simple calculation shows that $f_{V'}$ induces on the lowest end of V' the primitive $f_{L_k} - r$, where $r = \mathcal{S}_e(V) + \sigma + \epsilon$. (Here ϵ can be assumed to be arbitrarily small. It can be estimated from above by the area enclosed by the bent end corresponding to ℓ_k, the projection to $\mathbb{R}^2$ of the non-cylindrical part of V, ℓ_1, and $\gamma^{\uparrow}$. See Figure 3.20.)

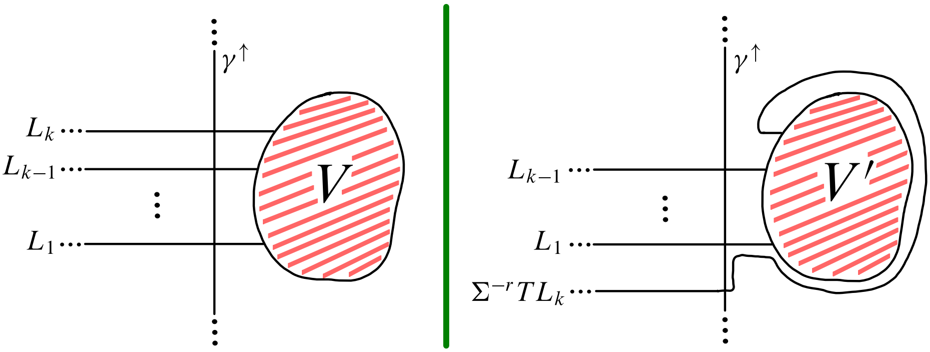

Figure 3.20. Bending the upper end L_k of a cobordism V.

Taking grading into account (in case V is graded in the sense of Floer theory), one can easily see that the grading on the lowest end of V' is translated by 1 in comparison to L_k. Summing up, the above procedure transforms a cobordism V with ends $L_1, \dots, L_k$ into a cobordism V' with ends $\Sigma^{-r}TL_k, L_1, \dots, L_{k-1}$.

Similarly, one can take V and bend its lowest end L_1 counterclockwise around V and obtain a new cobordism V'' with ends $L_2, \dots, L_k, \Sigma^r T^{-1} L_1$ and with $\mathcal{S}_e(V'') = \mathcal{S}_e(V)$.

We are now in a position to describe geometrically weighted exact triangles. Let V be a cobordism with three ends, which for compatibility with Definition 2.42 we denote by C, B, A (going from the lowest end upward). See Figure 3.21.

Let $r = \mathcal{S}_e(V) + \sigma + \epsilon$. Put

$$\bar{u} := \Gamma_{V,\gamma_{2,3}} : A \to B, \quad \bar{v} := \Gamma_{V,\gamma_{1,2}} : B \to C.$$

Consider also the counterclockwise rotation V'' of V whose ends are $B, A, \Sigma^r T^{-1} C$. Let

$$\bar{w} := \Sigma^{-r} T \Gamma_{V'',\gamma_{2,3}} = \Gamma_{V',\gamma_{1,2}} : C \to \Sigma^{-r} TA.$$

We claim that

$$A \xrightarrow{\bar{u}} B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r} TA \tag{3.42}$$

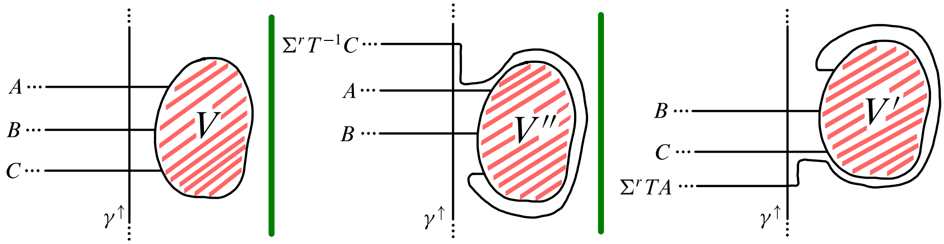

Figure 3.21. A cobordism leading to an exact triangle of weight r.

is a strict exact triangle of weight r. This triangle is based on the genuinely exact triangle from $\mathscr{C}_0$:

$$A \xrightarrow{u} B \xrightarrow{v} C' \xrightarrow{w} TA,$$

where $u = \bar{u}$, $C' = \mathcal{I}^*_{\gamma_{2,3}} V = \mathrm{Cone}(A \xrightarrow{u} B)$, $v : B \to C'$ is the standard inclusion, and $w : C' \to TA$ the standard projection. The r-isomorphism $\phi : C' \to C$ and its right r-inverse $\psi : \Sigma^r C \to C'$ are as follows:

$$\phi = \Gamma_{V,\gamma_{1,1},\gamma_{2,3}}, \quad \psi = T\Gamma_{V'',\gamma_{1,2},\gamma_{3,3}} : \Sigma^r C \to T\mathcal{I}^*_{\gamma_{1,2}} V'' = C'.$$

Note that $\mathcal{I}^*_{\gamma_{1,2}} V'' = T^{-1}\mathcal{I}^*_{\gamma_{2,3}} V = T^{-1}C' = T^{-1}\mathrm{Cone}(A \xrightarrow{u} B)$.

The fact that ψ is a right r-inverse of ϕ and that these maps fit into the diagram (2.20) follows from standard arguments in Floer theory. Note that these statements do not hold on the chain level, but only in $\mathscr{C}_0$.

3.5.1.9 Rotation of triangles. Let V be a cobordism with three ends C, B, A as in Section 3.5.1.8, and consider the exact triangle (3.42) of weight r. Let V' be the clockwise rotation of V, with ends $B, C, \Sigma^{-r}TA$. The exact triangle associated to V' is

$$B \xrightarrow{\bar{v}} C \xrightarrow{\bar{w}} \Sigma^{-r}TA \xrightarrow{\bar{u}'} \Sigma^{-r-\epsilon'}TB, \tag{3.43}$$

where ϵ' can be assumed to be arbitrarily small. It is not hard to see that in $\mathscr{C}_\infty$ (up to identifying objects with their shifts and ignoring signs in the maps) the exact triangle (3.43) is precisely the rotation of the exact triangle corresponding to (3.42) in $\mathscr{C}_\infty$.

The above shows that rotation of weighted exact triangles coming from cobordisms with three ends preserves weights (up to an arbitrarily small error). Interestingly, this is sharper than the case in a general TPC, described in Proposition 2.46, where the weight of a rotated triangle doubles. See also Remark 2.48.

3.5.1.10 Weighted octahedral property. The weighted octahedral formula from Proposition 2.49 admits too a geometric interpretation in the realm of cobordisms. We will not give the details of this construction here. Instead, we will briefly explain the cobordism counterpart of cone refinement and why it behaves additively with respect to weights, as described algebraically in Proposition 2.55. Note that weighted cone refinement is one of the main corollaries of the weighted octahedral property.

For simplicity, we focus here on the case described in Example 2.56 and ignore the grading translation T. Assume that we have two cobordisms: V with ends X, B, A, and U with ends A, F, E.

These cobordisms induce two weighted exact triangles of weights $r = \mathcal{S}_e(V) + \sigma_V + \epsilon$ and $s = \mathcal{S}_e(U) + \sigma_U + \epsilon$. By gluing the two cobordisms along the ends corresponding to A we obtain a new cobordism W with four ends X, B, F, E. See Figure 3.22. By the previous discussion, this exhibits X as an iterated cone with linearization (B, F, E), which corresponds precisely to the algebraic cone refinement of X with linearization (B, A) by A with linearization (F, E).

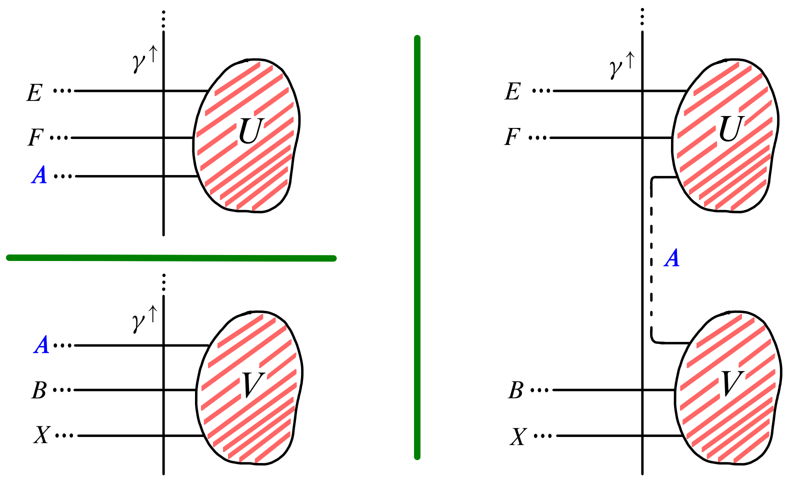

Figure 3.22. Cone refinement via gluing cobordisms along two ends.

Clearly the exterior shadow $\mathcal{S}_e(W)$ of W equals $\mathcal{S}_e(U) + \mathcal{S}_e(V)$, and the total weight of the cone decomposition of X associated to W is $r + s$. Note again that, by the procedure explained in Remark 3.34, one can make the errors e_V, e_U small at the expense of applying some shifts to the elements of the linearization.

Remark 3.35. It is well known that Lagrangian cobordism gives rise to a category with objects Lagrangian submanifolds and with morphisms certain Lagrangian cobordisms (see [7]). Combined with the discussion above, it is natural to wonder whether, by taking into account also the shadows of cobordisms, this category is naturally endowed with a TPC structure. The difficulty in achieving this is that one needs to have a triangulated structure that serves as the level 0 part of this expected TPC. Achieving geometrically such a triangulated structure is delicate, as it requires including immersed Lagrangians and cobordisms in the construction. To further produce a TPC structure, this construction needs to be combined with control of cobordism shadows, which makes the whole machinery even more complex. Only partial results in this direction are available at the moment, as in [9].

3.5.1.11 Exterior shadow equals shadow. Here we sketch a proof showing that $\mathcal{S}_{\mathrm{e}}(V) = \mathcal{S}(V)$. Throughout the proof below we will assume that V is connected and that the number of its ends is $l \geq 1$. (In any case, we do not consider cobordisms with no ends in this memoir.)

As at the beginning of 3.5.1.5, fix an open rectangle $Q \subset \mathbb{R}^2$ such that $\mathrm{out}(V) \setminus Q$ consists of only horizontal rays. Denote by $C_1, \ldots, C_l \subset \mathbb{R}^2$ the connected components of $\mathbb{R}^2 \setminus \mathrm{out}(V)$, ordered in counterclockwise order (e.g., with respect to a large enough circle surrounding Q). Set also $C_{Q,i} := C_i \cap Q$, $i = 1, \ldots, l$.

We first claim that each C_i is simply connected. Indeed, this follows easily from the fact that a connected open subset of $\mathbb{R}^2$ is simply connected if and only if every connected component of its complement is unbounded. It follows from the definition of $\mathrm{out}(V)$ that every connected component of $\mathbb{R}^2 \setminus C_i$ is unbounded, hence C_i is simply connected.

Since the C_i's are simply connected, so are the sets $C_{Q,i}$, $i = 1, \ldots, l$. It follows (e.g., by uniformization) that each of the open sets $C_{Q,i}$ is diffeomorphic to an open disk. Furthermore, by the Greene–Shiohama theorem [34] it follows that each $C_{Q,i}$ is in fact symplectomorphic to an open disk $\mathrm{Int}\, B^2(R_i)$ of some radius R_i, endowed with its standard symplectic structure. Fix such symplectomorphisms $\psi : B^2(R_i) \to C_{Q,i}$ for all i.

Assume for the moment that $l \geq 2$ (note that l is precisely the number of ends of V). Reduce the radii R_i slightly to $R_i' = R_i - \delta$ for small $\delta > 0$, and consider the corresponding domains $C'_{Q,i} = \psi(B^2(R_i'))$. Next, connect the boundary of $C'_{Q,i}$ to the boundary of $C'_{Q,i+1}$ by a small strip J_i that intersects $\mathrm{out}_Q(V)$ only along the areas where $\mathrm{out}_Q(V)$ consists solely of horizontal rays. If we smoothly (not symplectically) parametrize J_i as $[-\epsilon, \epsilon] \times [0, 1]$, we just embed J_i in $\mathbb{R}^2$ in such a way that:

(1) The area of J_i is very small and $J_i \subset Q$.

(2) $[-\epsilon, \epsilon] \times \{0\}$ is mapped to $\partial C'_{Q,i}$ near one of the horizontal rays, say E_i, of $\mathrm{out}_Q(V)$ that lies near $\partial C'_{Q,i}$.

(3) $[-\epsilon, \epsilon] \times \{1\}$ is mapped to $\partial C'_{Q_{i+1}}$ near the same horizontal ray E_i we have just used in (1) above.

(4) The rest of J_i intersects $\mathrm{out}(V)$ only along E_i.

We can think of the outcome of connecting $C'_{Q,i}$ to $C'_{Q,i+1}$ with J_i as the boundary connected sum of the closures of the domains $C'_{Q,i}$ and $C'_{Q,i+1}$.

We perform the above construction for all $1 \leq i \leq l-1$, and finally we connect $C'_{Q,l}$ back to $C'_{Q,1}$ in a similar manner, keeping the counterclockwise direction.

Denote by C'_Q the union of all the domains $C'_{Q,i}$ together with the connecting small strips J_i. The outcome $C'_Q \subset Q$ is a domain diffeomorphic to an annulus. Its inner boundary encircles a domain F which is diffeomorphic to the 2-dimensional disk B, and $\mathrm{out}(V) \setminus F$ consists of only horizontal rays. Moreover, by taking the parameter δ small enough we can assume that $\mathrm{Area}(Q \setminus F) = \mathrm{Area}(Q) - \mathrm{Area}(F)$ is arbitrarily close to $\mathrm{Area}(Q) - \mathrm{Area}(\mathrm{out}_Q(V)) = \mathrm{Area}(Q) - \mathcal{S}(V)$. It follows that $\mathrm{Area}(F)$ is arbitrarily close to $\mathcal{S}(V)$ and, at the same time,

$$\mathrm{Area}(F) \geq \mathcal{S}_{\mathrm{e}}(V) \geq \mathcal{S}(V).$$

This concludes the proof under the assumption that $l \geq 2$.

The case $l = 1$ is very similar, only that now we have just one domain $C'_{Q,1}$ and we form the annulus C'_Q by just gluing the small strip J to connect two portions of the boundary of the same domain $C'_{Q,1}$. ■

3.5.2 Some explicit estimates

We will illustrate here the statements in Theorem 3.4.

Our base manifold will be denoted here by W, and it is the plumbing of two copies of the disk cotangent bundle D^*S^1 of S^1, as in Figure 3.23. The family $\mathcal{F}$ has two elements, F_1 and F_2, as in this figure. They intersect at the single point P.

The primitives on both F_1 and F_2 are the functions identically equal to 0. The family $\mathcal{X}$ consists of F_1, F_2, and the Lagrangians Y, Z, X, and N from Figures 3.24 and 3.25.

The Lagrangian Y is constructed from the surgery $F_2 \# F_1$ at the point P (with a small handle), followed by a small Hamiltonian perturbation. It is easy to see that, for Y to be exact, the "small" gray triangle STP must have the same area as the "large" triangle with the same vertices (only the corners of the second triangle are grayed in the figure). We will denote the area of these triangles by A_Y.

Similarly, the Lagrangian X is constructed from the surgery $F_1 \# F_2$ at the point P (again with a small handle), followed by a small Hamiltonian perturbation. Again, the "small" gray triangle QRP has the same area as the "large" triangle with the same vertices, and we denote this area by A_X. The Lagrangian N is obtained from F_2 by a Hamiltonian perturbation that is large: its Hofer distance equals the area of the strip comprised between N and F_2 and the points x_1 and x_2 (there are two such strips, but they both have the same area). We will denote this area by A_N. Similarly, the Lagrangian Z is obtained from F_1 by a large Hamiltonian perturbation.

The first obvious remark is that X and Y are not smoothly isotopic, because homologically $[X] = [F_2] - [F_1]$ and $[Y] = [F_2] + [F_1]$, and these are not equal in $H_1(X, \mathbb{Z})$.

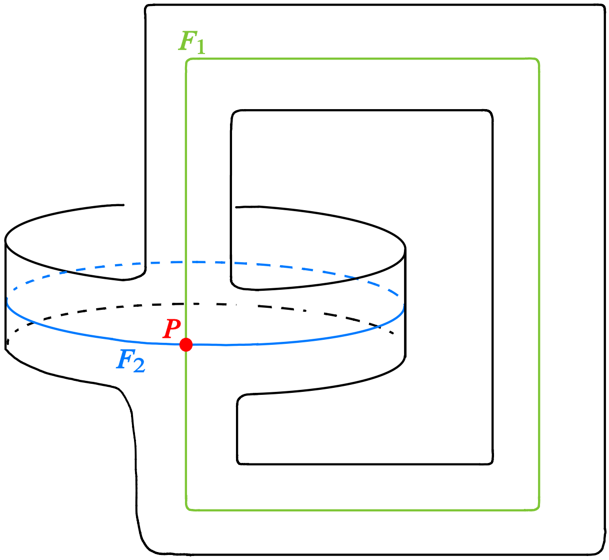

Figure 3.23. The manifold W and the Lagrangians F_1 and F_2.

We are interested in the distance $D^{\mathcal{F}}$.

Lemma 3.36. *We have the inequalities*

$$\frac{A_X}{4} \leq D^{\mathcal{F}}(X,0) \leq 2A_X \frac{A_Y}{4} \leq D^{\mathcal{F}}(Y,0) \leq 2A_Y, \tag{3.44}$$

$$\frac{\max\{A_Y, A_X\}}{4} \leq D^{\mathcal{F}}(X,Y) \leq 2A_X + 2A_Y. \tag{3.45}$$

Proof. We first show the upper bounds in (3.44). The cases of X and Y are entirely similar, and we focus on X. For this, we consider the cone of the map $F_1 \xrightarrow{P} F_2 \to K_1$ (constructed in terms of A_∞-Yoneda modules). We claim that the module K_1 can be mapped to the Yoneda module of X by a quasi-isomorphism. The simplest way to see this geometrically is the following: interpret the module K_1 as the Yoneda module of a marked immersed Lagrangian with one marked self-intersection point (marked in the order (F_1, F_2)). This type of Lagrangians is discussed in [9], for instance. The map we are looking for is of the form $\psi = \mu_2(-, R) : K_1 \to X$, with R the intersection point in Figure 3.25. Of course, once we "guess" this morphism, we can write it purely algebraically. It is easy to see that this is a quasi-isomorphism. For instance, applying it to the Lagrangian N in the picture, it sends x_i to y_i for $i = 1, 2$. Moreover, there is also a quasi-isomorphism $\phi = \mu_2(-, Q) : X \to K_1$, which is a quasi-inverse of the first (on N it is an actual inverse). We can fix the primitive on X that vanishes at the point Q, and thus the primitive on X has value A_X at R. In the terminology of

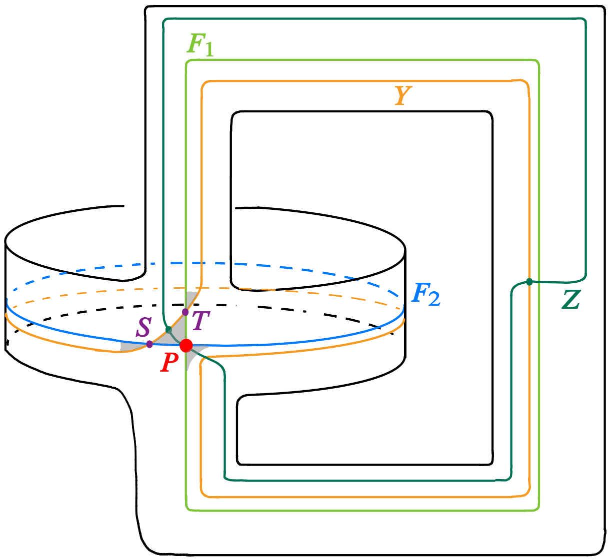

Figure 3.24. The Lagrangian Y obtained as a small perturbation of the surgery $F_2 \#_P F_1$.

this memoir, we have maps of filtered modules

$$\Sigma^{A_X} K_1 \to X \to K_1$$

whose composition agrees with the map η_{A_X} (in other words, the composition is the identity if the filtration is neglected, but once the filtration is taken into account, it shifts the filtration by A_X). We also have a similar identity in the opposite direction.

By applying the same argument as in the second part of Lemma 2.85, we deduce that ϕ and ψ are $2A_X$-isomorphisms, which implies an inequality for the half-distance $\bar{\delta}^{\mathcal{F}}(X, 0) \leq 2A_X$. The other inequality, for the second half-distance, is easy to obtain using the fact that the cone of $\phi : X \to K_1$ is $2A_X$-acyclic, and this implies our upper bound.

For the lower bound, notice that $\delta(X; F_1 \cup F_2) = 2A_X$, and thus the lower bound follows from Theorem 3.4 (ii) (here $\delta(-;-)$ is the relative Gromov width as in Section 3.1).

Clearly, in an entirely similar way we also have $D^{\mathcal{F}}(Y, 0) \leq 2A_Y$ and thus

$$D^{\mathcal{F}}(X, Y) \leq 2A_X + 2A_Y,$$

which is the upper bound in (3.45).

Remark 3.37. The first part of the argument is very similar to the one relating the spectral distance to the distance $D^{\mathcal{F}}$. Indeed, one can think about the two points R and Q as representing the point class and the fundamental class in $\mathrm{HF}(X, K_1)$, and then the first point of Theorem 3.4 implies $D^{\mathcal{F}}(X, K_1) \leq 4A_X$, which means that

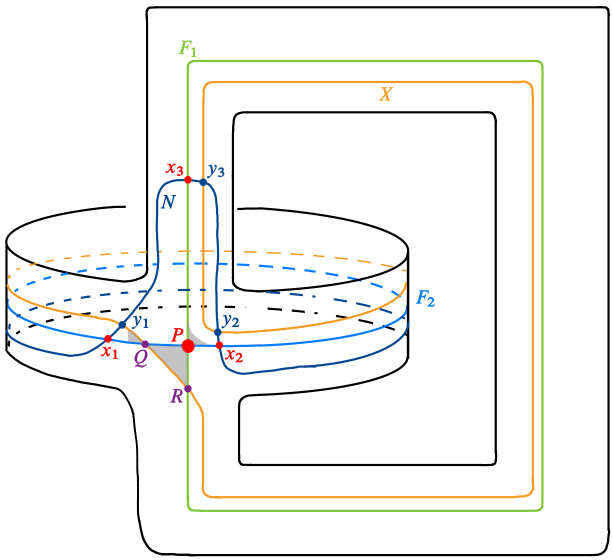

Figure 3.25. The Lagrangian X obtained as a small perturbation of the surgery $F_1 \#_P F_2$ and the Lagrangian N, which is a large Hamiltonian perturbation of F_2.

$D^{\mathcal{F}}(X, 0) \leq 4A_X$ because $D^{\mathcal{F}}(K_1, 0) = 0$. It is very likely that we actually have $D^{\mathcal{F}}(X, 0) = A_X$ and $D^{\mathcal{F}}(Y, 0) = A_Y$.

Finally, we discuss the lower bound in (3.45). For this purpose, we will use here point (ii) of Theorem 3.4. It can be easily shown that $\delta(X; Y \cup F_1 \cup F_2) \geq 2A_X$. Thus we get from point (ii) of Theorem 3.4 that $D^{\mathcal{F}}(X, Y) \geq \frac{A_X}{4}$, as claimed. By symmetry we also get $D^{\mathcal{F}}(X, Y) \geq \frac{A_Y}{4}$. ■

Remark 3.38. There is an alternative (and possibly more interesting) argument which however gives a slightly weaker inequality than the left-hand side of (3.45). Namely, it implies that

$$\frac{\max\{A_Y, A_X\}}{8} \leq D^{\mathcal{F}}(X, Y).$$

This argument is based on point (iii) of Theorem 3.4 and goes as follows. Consider the Lagrangian Z in Figure 3.24. It has three intersection points with Y and only one with X. By point (iii) of Theorem 3.4, we have $D^{\mathcal{F}}(X, Y) \geq \frac{1}{16}\delta^{\cap}(Z, Y; F_1 \cup F_2)$, where $\delta^{\cap}$ is the quantity defined in (3.2). So this time we need to estimate the number $\delta^{\cap}(Z, Y; F_1 \cup F_2)$. For this estimate it is useful to assume that Z cuts the triangle STP into two pieces of equal area. In this case we have that $\delta^{\cap}(Z, Y; F_1 \cup F_2) = 2A_Y$, and we deduce $D^{\mathcal{F}}(X, Y) \geq \frac{A_Y}{8}$. The inequality involving X follows in the same way, by choosing a deformation Z' of F_1 that this time intersects Y in a single point and X in three points.