A Lagrangian pictionary

Paul Biran and Octav Cornea

Abstract We describe a dictionary of geometry $\longleftrightarrow$ algebra in Lagrangian topology. As a by-product, we obtain a tautological (in a particular sense which we explain) proof of a folklore conjecture (sometimes attributed to Kontsevich) claiming that the objects and structure of the derived Fukaya category can be represented through immersed Lagrangians. Our construction is based on certain Lagrangian cobordism categories endowed with a structure called *surgery models*.

Contents

1. Introduction 399
2. Cobordism categories with surgery models 406
3. Unobstructed classes of Lagrangians and the category $\mathsf{C}ob^*(M)$ 430
4. From geometry to algebra and back 461
5. Some further remarks 484
References 490

1. Introduction

Our purpose here is to describe a dictionary of

$$\textit{geometry} \longleftrightarrow \textit{algebra}$$

in Lagrangian topology. As a by-product, we obtain a tautological (in a sense we explain below) proof of a folklore conjecture (sometimes attributed to Kontsevich) claiming that the objects and structure of the derived Fukaya category can be represented through immersed Lagrangians. We also shed some new light on the rigidity/flexibility dichotomy that is at the heart of symplectic topology.

We focus on a symplectic manifold $(M^{2n}, \omega = d\lambda)$ with M being a Liouville manifold complete at infinity, or a compact Liouville domain, and we consider classes $\mathcal{L}ag^*(M)$ of (possibly immersed) Lagrangian submanifolds $L \hookrightarrow M$ carrying also a variety of additional structures. Throughout this article, the Lagrangians are exact and the choices of primitives are part of the structure.

Kyoto Journal of Mathematics, Vol. 61, No. 2 (2021), 399–493
DOI 10.1215/21562261-2021-0007X,
Received April 8, 2020. Revised October 4, 2020. Accepted October 26, 2020.
First published online April 16, 2021.
2020 Mathematics Subject Classification: Primary 53D12; Secondary 53D37, 57R90.

1.1. Outline of the dictionary

1.1.1. Geometry

The geometric side of our dictionary, which does not appeal at any point to J-holomorphic curves, consists of a class $\mathcal{L}ag^*(M)$ of Lagrangians, possibly immersed and marked (and endowed with additional structures). For an immersed Lagrangian $j_L : L \to M$, in generic position, a marking consists of a choice of a set of self-intersection points $(P_-, P_+) \in L \times L$ with $j_L(P_-) = j_L(P_+)$. The Lagrangians in $\mathcal{L}ag^*(M)$ are organized as a cobordism category $\mathsf{C}ob^*(M)$, whose morphisms are represented by Lagrangian cobordisms belonging to a class $\mathcal{L}ag^*(\mathbb{C} \times M)$. These are Lagrangians in $\mathbb{C} \times M$ whose projection onto $\mathbb{C}$ is as in Figure 3. To such a category we associate a number of geometric operations. We then list a set of axioms governing them. If the category $\mathsf{C}ob^*(M)$ obeys these axioms, then we say that it has *surgery models*. The reason for this terminology is that the cobordisms associated to the classical operation of Lagrangian surgery (see, e.g., [34], [43], [9]) are in some sense the simplest morphisms in our category (see also Figure 9). There is a natural equivalence relation $\sim$, called *cabling equivalence*, on the morphisms of $\mathsf{C}ob^*(M)$. The fact that the category has surgery models means that each morphism is equivalent to one coming from surgery (together with some other properties that we skip now). To understand this condition through an analogy, consider the path category of a Riemannian manifold, with objects the points of the manifold and with morphisms the Moore paths. In this case, the equivalence relation is homotopy and the place of surgery morphisms is taken by geodesics. If $\mathsf{C}ob^*(M)$ has surgery models, then the quotient category

$$\widehat{\mathsf{C}}ob^*(M) = \mathsf{C}ob^*(M)/\sim$$

is triangulated. The triangulated structure appears naturally in this context because a surgery cobordism, resulting from the surgery operation, has three ends.

Cobordisms are endowed with a natural measurement provided by their shadow—roughly the area of the connected "filling" of their projection onto $\mathbb{C}$ (see [23]). If $\mathsf{C}ob^*(M)$ has surgery models, then the triangulated structure of $\widehat{\mathsf{C}}ob^*(M)$ together with the weight provided by the shadow of cobordisms leads to the definition of a class of pseudometrics on $\mathcal{L}ag^*(M)$ called (shadow) *fragmentation pseudometrics* and denoted by $d^{\mathcal{F}}$ (see [12]). Here $\mathcal{F}$ is a family of objects in $\mathsf{C}ob^*(M)$ (it can be imagined as a family of generators with respect to the triangulated structure). Intuitively, the pseudodistance $d^{\mathcal{F}}(L, L')$ between two objects L and L' infimizes the weight of iterated cone-decompositions expressing L in terms of L' and of objects from the family $\mathcal{F}$. The category $\mathsf{C}ob^*(M)$ is called (strongly) *rigid* (or *nondegenerate*) in case, under some natural constraints on $\mathcal{F}$ and a second such family $\mathcal{F}'$ (that can be viewed as a generic perturbation of $\mathcal{F}$), the pseudometrics $d^{\mathcal{F},\mathcal{F}'} = d^{\mathcal{F}} + d^{\mathcal{F}'}$ are nondegenerate in the sense that they separate any two objects of $\mathsf{C}ob^*(M)$ that have different geometric images.

1.1.2. Algebra

The algebraic side of the dictionary is mainly given by the derived Fukaya category $D\mathcal{F}uk^*(M)$ associated to the *embedded* Lagrangians $\mathcal{L}ag_e^*(M) \subset \mathcal{L}ag^*(M)$. The definition of this triangulated category [48] is an algebraically sophisticated application of J-holomorphic curve techniques. The construction is only possible for Lagrangian submanifolds that are *unobstructed* (see [28]). Starting with the definition of Floer theory in [25], a variety of constraints, often of topological type (e.g., exactness, monotonicity, and so forth), have been identified that are sufficient to ensure unobstructedness. Further unobstructedness conditions are now understood for Lagrangian immersions following the work of Akaho and Joyce in [1] and [2].

1.2. Main result

In this framework, the central point of our dictionary is the following.

THEOREM A

There are certain classes $\mathcal{L}ag^*(M)$, $\mathcal{L}ag^*(\mathbb{C} \times M)$ of exact, marked, unobstructed, immersed *Lagrangian submanifolds and, respectively, cobordisms such that the category $\mathsf{C}ob^*(M)$ has surgery models and is rigid. The classes $\mathcal{L}ag^*(M)$, $\mathcal{L}ag^*(\mathbb{C} \times M)$ can be assumed to contain all the exact embedded Lagrangian submanifolds and exact embedded cobordisms. Moreover, the quotient category $\widehat{\mathsf{C}ob}^*(M)$ is triangulated and its subcategory generated by the embedded Lagrangians is triangulated, isomorphic to $D\mathcal{F}uk^*(M)$.*

A more precise version of this statement is Theorem 4.1.1 in Section 4.1. Classes $\mathcal{L}ag^*(L)$, $\mathcal{L}ag^*(\mathbb{C} \times M)$ as above can be found, in particular, for all Liouville domains. One way to look at this result is that if the class $\mathcal{L}ag^*(M)$ is small enough, then the metric structure is nondegenerate, but defining a triangulated structure on $\widehat{\mathsf{C}ob}^*(M)$ requires more objects than those in $\mathcal{L}ag^*(M)$. On the other hand, if the class $\mathcal{L}ag^*(M)$ is too big, then the triangulated structure is well defined, but the expected metric is degenerate. However, the unobstructed class in the statement satisfies both properties—it is flexible enough so that triangulation is well defined and, at the same time, it exhibits enough rigidity so that the metric is geometrically nondegenerate. Additionally, in this case the subcategory generated by embedded objects agrees with the derived Fukaya category and, thus, most algebraic structures typical in Lagrangian Floer theory can be reduced to understanding how unobstructedness behaves with respect to the geometric operations appearing in our axioms.

Among the immediate consequences of this result, we see that the natural cobordism group generated by the embedded Lagrangians, modulo the relations given by the (possibly immersed) cobordisms in $\mathcal{L}ag^*(\mathbb{C} \times M)$, is isomorphic to the Grothendieck group $K_0(D\mathcal{F}uk^*(M))$. In a different direction, because the Lagrangians in $\mathcal{L}ag^*(M)$ are unobstructed, there is a Donaldson category

$\mathcal{D}on(\mathcal{L}ag^*(M))$ associated to this family and this category is isomorphic to $\widehat{\mathsf{C}}ob^*(M)$.

The proof of Theorem A relates geometric constructions in $\widehat{\mathsf{C}}ob^*(M)$ to corresponding operations in $D\mathcal{F}uk^*(M)$. All objects and morphisms in $D\mathcal{F}uk^*(M)$ are represented geometrically and all exact triangles in $D\mathcal{F}uk^*(M)$ are represented by surgery. There are also relations between the pseudometrics $d^{\mathcal{F},\mathcal{F}'}$ and corresponding pseudometrics defined on the algebraic side by making use of action-filtered A_∞ machinery.

This language also allows one to formulate what today seems a very strong conjecture (or, maybe, just an intriguing question) that is a formal version in this context of Yasha Eliashberg's claim that "there is no rigidity beyond J-holomorphic curves."

CONJECTURE/QUESTION 1

If $\mathsf{C}ob^*(M)$ has surgery models and is rigid, then it consists of unobstructed Lagrangians; in particular, the morphisms in $\widehat{\mathsf{C}}ob^*(M)$ are Floer homology groups.

1.3. Some details

We discuss briefly here some of the main points in our approach.

1.3.1. Triangulation

As mentioned before, there is a simple reason why triangulated categories are the correct framework for organizing Lagrangian submanifolds. The simplest operation with such submanifolds—that is, surgery—produces a Lagrangian cobordism with three ends, two being the Lagrangians that are inserted in the construction and the third being the result of the surgery. By "shuffling" appropriately the three ends (see Figure 5) one immediately extracts a triple of maps that can be expected to give rise, in the appropriate setting, to a distinguished triangle. This simple remark implies that, to aim toward building a triangulated category, the class of Lagrangians considered has to be closed under the particular type of surgery used. For two Lagrangians L and L' that intersect transversely, surgery is possible along any subset $c \subset L \cap L'$ of their intersection points. However, it is clear that without some constraints on these sets c the resulting category will not be rigid. So the question becomes how to constrain the choices of these c's allowed in the surgeries so that the output is rigid. The notion of unobstructedness comes to the forefront at this point: if L and L' are unobstructed, then the Floer chain complex $CF(L, L')$ is defined and one can chose c to be a cycle in this complex. One expects in this case that the output of the surgery is unobstructed and that the surgery cobordism itself is unobstructed. This process can then be iterated. Further along the line, if all these expectations are met, comparison with the relevant Fukaya category becomes credible as well as the rigidity of the resulting category.

1.3.2. Immersed Lagrangians

There are a number of difficulties with this approach. The first is that surgery in only some of the intersection points of L and L' produces an immersed Lagrangian, even if both L and L' are embedded. As a result, one is forced to work with unobstructed immersed Lagrangians. Unobstructedness for immersed objects depends on choices of data and, therefore, this leads to the complication that the objects in the resulting category cannot be simply immersed Lagrangian submanifolds L but rather pairs $(L, \mathcal{D}_L)$ consisting of such L together with these data choices $\mathcal{D}_L$. This turns out to be just a formal complication but, more significantly, to prove unobstructedness of the result of the surgery and of the surgery cobordism, one needs to compare moduli spaces of J-holomorphic curves before and after surgery in a way that is technically very delicate.

1.3.3. Marked Lagrangians

In this paper, we pursue a different approach that bypasses this comparison. We further enrich our objects (thus the immersed Lagrangians L, assumed having only transverse double points) with one additional structure, a marking $\mathbf{c}$, which consists of a set of double points of L. Therefore, now our objects are triples $(L, \mathbf{c}, \mathcal{D}_L)$. The advantage of markings is the following. Given two Lagrangians L and L' (possibly immersed and endowed with markings) and a set of intersection points $c \subset L \cap L'$, we can define a new immersed, marked Lagrangian by the union $L \cup L'$ endowed with a marking consisting of c together with the markings of L and L'. Geometrically, this corresponds to a surgery at the points in c but using 0-*size handles* in the process. Once markings are introduced, all the usual definitions of Floer complexes and the higher μ_k operations have to be adjusted. While usual operations for immersed Lagrangians count J-holomorphic curves that are not allowed to jump branches at self-intersection points, in the marked version such jumps are allowed but only if they occur at points that are part of the marking. Experts will probably recognize that these markings are, in fact, just particular examples of bounding chains as in [2] (see Section 5.2.2 for more details). Keeping track of the markings and the data $\mathcal{D}_L$ through all the relevant operations requires roughly the same type of choices and the same effort as needed to build coherent, regular perturbations in the construction of the Fukaya category. Indeed, the process produces geometric representatives of the modules in $\mathcal{DF}uk^*(M)$ that are unions of embedded Lagrangian submanifolds $L_1 \cup \cdots \cup L_k$ (together with some marking given by relevant intersection points) and the data associated to such a union is similar to the choice of data associated to the family $L_1, \ldots, L_k$ in the construction of the Fukaya category $\mathcal{F}uk^*(M)$. This is why we call *tautological* our solution to the conjecture mentioned at the beginning of the Introduction. At the same time, the use of markings has also some conceptual, intrinsic advantages (see also Section 5.2.1).

1.3.4. Immersed cobordisms

There are two other additional difficulties that need to be addressed. The first is that cobordisms with immersed ends do not have isolated double points, at best they have clean self-intersections that contain some half-lines of double points. As a result, we need to consider some appropriate deformations that put them in generic position. This leads to additional complications with respect to composing the morphisms in our category. In particular, because the morphisms are represented by cobordisms together with the relevant deformations, the objects themselves have to be equipped with one additional piece of data that comes down to a germ of this deformation along the ends.

1.3.5. Cabling

The second difficulty has to do with the equivalence relation $\sim$ (cabling equivalence) that is defined on the morphisms of $\mathsf{C}ob^*(M)$. This relation is defined through an operation called *cabling* that is applied to two cobordisms $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$ and $V' : L \rightsquigarrow (L'_1, \ldots, L'_s, L')$ (see Figure 3) and produces a new Lagrangian $\mathcal{C}$ in $\mathbb{C} \times M$ pictured in Figure 12. Both $V, V' \in \mathcal{L}ag^*(\mathbb{C} \times M)$ are viewed in our category as morphisms in $\mathrm{Mor}(L, L')$ and they are cabling-equivalent if their cabling also belongs to $\mathcal{L}ag^*(\mathbb{C} \times M)$. It is easy to see that $\mathcal{C}$ is generally immersed (even if V and V' are both embedded) and that there are, in general, holomorphic curves $u : D^2 \backslash \{1\} \to \mathbb{C} \times M$ with boundary on $\mathcal{C}$ and with the boundary puncture that is sent to one of the self-intersection points of $\mathcal{C}$. Such curves are called *teardrops*. In the unobstructed setting, the counts of these teardrops, at each self-intersection point, need to vanish. The difficulty is to ensure that the relevant moduli spaces are regular so that these counts are possible. We deal with this issue by using appropriate interior marked points (having to do with the planar configurations associated to cabling) and standard regularity arguments.

1.3.6. A functor from geometry to algebra

The proof of Theorem A is based on constructing a functor

$$\Theta : \mathsf{C}ob^*(M) \to H\big(\mathrm{mod}\big(\mathcal{F}uk^*(M)\big)\big),$$

where $\mathrm{mod}(\mathcal{F}uk^*(M))$ is the dg-category of A_∞-modules over the (embedded) Fukaya category $\mathcal{F}uk^*(M)$ (which is assumed to be well defined now). We then show that, by passing to the triangulated quotient $\widehat{\mathsf{C}ob}^*(M)$, there is an induced triangulated functor $\widehat{\Theta}$ with the properties listed in Theorem 4.1.1. The construction of Θ follows the approach in [9] and [10], but there are additional difficulties as outlined earlier in this subsection. For the reader's convenience, we provide here an overview of the construction of Θ (without mention of any of the technical details).

We start with Figure 1 below. It contains a cobordism V with three ends, one positive, L, and two negative, L' and L''. Such a cobordism is viewed as

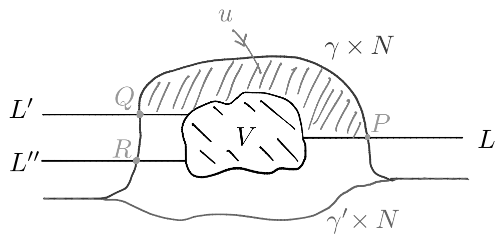

Figure 1. (Color online) A cobordism V with three ends and planar curves γ and γ'.

a morphism $V : L \rightsquigarrow L'$. We also consider a curve γ as in Figure 1. The functor Θ has a simple definition on the objects of $\mathsf{C}ob^*(M)$: it associates to each such object L the Yoneda module $\mathcal{Y}(L)$. To define Θ on morphisms, we associate to a cobordism $V : L \rightsquigarrow L'$ a morphism, denoted $\phi_V : \mathcal{Y}(L) \to \mathcal{Y}(L')$, and we put $\Theta(V) = [\phi_V] \in H(\hom_{\mathrm{mod}(\mathcal{F}uk^*(M))}(\mathcal{Y}(L), \mathcal{Y}(L')))$. The morphism ϕ_V is easy to describe when applied to another object $N \in \mathcal{O}b(\mathsf{C}ob^*(M))$. It is defined by counting elements in a 0-dimensional moduli space $\mathcal{M}_{\gamma \times N, V}(x, y)$ of Floer strips $u : [0,1] \times \mathbb{R} \to \mathbb{C} \times M$ with boundary conditions $u(\{0\} \times \mathbb{R}) \subset \gamma \times N$, $u(\{1\} \times \mathbb{R}) \subset V$ and asymptotic conditions $\lim_{t \to -\infty} u([0,1] \times \{t\}) = \{P\} \times \{x\}$, $\lim_{t \to +\infty} u([0,1] \times \{t\}) = \{Q\} \times \{y\}$, where $P, Q \in \mathbb{C}$ are the respective points that appear in Figure 1, and $x \in N \cap L$, $y \in N \cap L'$ are generators of the Floer complexes $CF(N, L)$ and $CF(N, L')$, respectively. We assume here that N intersects transversely both L and L'. The projection of such a strip u covers the region in green in Figure 1.

By choosing convenient, almost-complex structures, one sees, as in [9], that the formula $\phi_V^N(x) = \sum \#\mathcal{M}_{\gamma \times N, V}(x, y) y$ defines, after extension by linearity, a chain map

$$\phi_V^N : CF(N, L) \to CF(N, L').$$

It is easy to see in this context why the different morphisms associated to a cobordism V as before, with only three ends, fit in an exact triangle. There are two additional cobordisms obtained from V by planar rotation, as in Figure 5, and there are corresponding maps $\phi_{RV}^N : CF(N, L') \to CF(N, L'')$, $\phi_{R^{-1}V}^N : CF(N, L'') \to CF(N, L)$. The relations among these maps are then seen as a consequence of the structure of the differential of the "global" Floer complex $CF(\gamma \times N, V)$ which is identified—as a vector space—with the sum $CF(N, L) \oplus CF(N, L') \oplus CF(N, L'')$ together with the fact that $CF(\gamma \times N, V)$ is acyclic. This acyclicity follows from the invariance of Floer homology with respect to (horizontal) Hamiltonian deformation which implies $HF(N \times \gamma, V) \cong HF(N \times \gamma', V) = 0$, where γ' is as in Figure 1.

Of course, one also needs to show that the construction of ϕ_V^N is sufficiently functorial in N such as to extend to an A_∞-module morphism and similarly for the triangulation properties of the construction. One can prove this in a direct way, by checking that the construction is compatible with the module multiplication and proceeding to deduce the other properties. More conceptually, one can consider a Fukaya category of cobordisms $\mathcal{F}uk^*(\mathbb{C}\times M)$ and the associated A_∞ Yoneda module $\mathcal{Y}(V)$ over this category. There is an inclusion functor $i_\eta : \mathcal{F}uk^*(M) \to \mathcal{F}uk^*(\mathbb{C}\times M)$ that is defined by $L \to \eta \times L$, for any planar curve η, flat at infinity. One then considers the pullback module $i_\gamma^*(\mathcal{Y}(V))$ with γ as in Figure 1 and shows that it has the structure of an iterated cone of A_∞-modules

$$\mathrm{Cone}\big(\mathcal{Y}(L) \to \big(\mathrm{Cone}\big(\mathcal{Y}(L') \to \mathcal{Y}(L'')\big)\big)\big).$$

The differential in this cone contains a module morphism $\mathcal{Y}(L) \to \mathcal{Y}(L')$ and this is then identified (in homology) with ϕ_V. As in the simplified setting with a fixed N, the triangulation properties of the construction follow from the description of $i_\gamma^*(\mathcal{Y}(V))$ and the fact that $i_\gamma^*(\mathcal{Y}(V))$ and $i_{\gamma'}^*(\mathcal{Y}(V)) = 0$ are quasi-isomorphic modules.

1.4. Organization of the article

In the second section, we review the geometric cobordism category and list the axioms mentioned above. We also prove that, when the axioms are satisfied, the resulting quotient category is triangulated. We also introduce the shadow fragmentation pseudometrics. All the constructions in this part are "soft." In the third section, we discuss unobstructedness for both Lagrangians and cobordisms and we also define the classes $\mathcal{L}ag^*(M)$ and $\mathcal{L}ag^*(\mathbb{C}\times M)$ that appear in the statement of Theorem A—the constructions here are "hard." In the fourth section, we pursue the correspondence *geometry* $\leftrightarrow$ *algebra* and prove Theorem A. In the last section, we discuss some further technical points and open questions.

2. Cobordism categories with surgery models

2.1. Objects, morphisms, and basic operations

The definition of Lagrangian cobordism that we use, as in [9], is a variant of a notion first introduced by Arnold [5], [6].

We fix on $\mathbb{R}^2$ the symplectic form $\omega_0 = dx \wedge dy$, $(x,y) \in \mathbb{R}^2$ and endow $\mathbb{R}^2 \times M$ with the symplectic structure $\omega_0 \oplus \omega$. We denote by $\pi : \mathbb{R}^2 \times M \to \mathbb{R}^2$ the projection. For a subset $V \subset \mathbb{R}^2 \times M$ and $S \subset \mathbb{R}^2$, we let $V|_S = V \cap \pi^{-1}(S)$. In what follows, we identify $\mathbb{R}^2 \cong \mathbb{C}$ by $(x,y) \longleftrightarrow x+iy$.

DEFINITION 2.1.1

Let $(L_i)_{1\le i\le k_-}$ and $(L'_j)_{1\le j\le k_+}$ be two tuples of closed Lagrangian submanifolds of M. We say that these two tuples are Lagrangian cobordant, $(L_i) \simeq (L'_j)$, if there exist a smooth compact cobordism $(V; \coprod_i L_i, \coprod_j L'_j)$ and a Lagrangian embedding $V \subset ([0,1]\times\mathbb{R}) \times M$ such that for some $\epsilon > 0$ we have

$$\begin{aligned} V|_{[0,\epsilon)\times\mathbb{R}} &= \coprod_i \big([0,\epsilon)\times\{i\}\big)\times L_i, \\ V|_{(1-\epsilon,1]\times\mathbb{R}} &= \coprod_j \big((1-\epsilon,1]\times\{j\}\big)\times L'_j. \end{aligned} \tag{1}$$

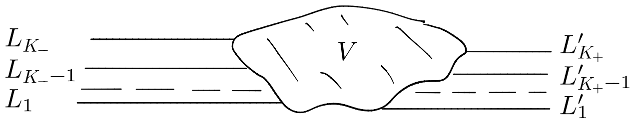

Figure 2. (Color online) A cobordism $V : (L'_j) \rightsquigarrow (L_i)$ projected on $\mathbb{R}^2$.

The manifold V is called a Lagrangian cobordism from the Lagrangian tuple (L'_j) to the tuple (L_i). We denote such a cobordism by $V : (L'_j) \rightsquigarrow (L_i)$. See Figure 2.

We will allow below for both the cobordism V and its ends L_i, L'_j to be immersed manifolds. The connected components of the ends L_i, L'_j will be assumed to have only isolated, transverse double points. The connected components of the cobordism V will be assumed to have double points that are of two types, either isolated transverse double points or 1-dimensional clean intersections (possibly diffeomorphic to semiclosed intervals). Such intervals of double points are associated to the double points of the ends.

It is convenient to extend the ends of such a cobordism horizontally in both directions and thus to view V as a (noncompact) submanifold of $\mathbb{C} \times M$. A cobordism V is called *simple* if it has a single positive end and a single negative end. Such a cobordism $V : L \rightsquigarrow L'$ can be rotated by using a 180° rotation in $\mathbb{C}$ and this provides a new simple cobordism denoted by $\bar{V} : L' \rightsquigarrow L$. The two most common examples of cobordisms are Lagrangian surgery (discussed in more detail later) and Lagrangian suspension. This is a simple cobordism associated to a Hamiltonian isotopy ϕ, $V^{\phi} : L \rightsquigarrow \phi(L)$.

The category $\mathsf{C}ob^*(M)$ has appeared before in [10] written as $S\mathcal{C}ob$ that we shorten here for ease of notation (that paper also included a slightly different category denoted $\mathcal{C}ob^*(M)$ that will not play a role here). There are a number of variants of this definition so we give a few details here.

The objects of this category are closed Lagrangian submanifolds in M, possibly immersed, in a fixed class $\mathcal{L}ag^*(M)$. We allow $\emptyset \in \mathcal{L}ag^*(M)$. We also fix a class of cobordisms $\mathcal{L}ag^*(\mathbb{C} \times M)$ with the property that the ends of the cobordisms in this class belong to $\mathcal{L}ag^*(M)$. To fix ideas, the classes in $\mathcal{L}ag^*(M)$, $\mathcal{L}ag^*(\mathbb{C} \times M)$ consist of immersed Lagrangian submanifolds endowed with certain decorations (such as, possibly, an orientation, a spin structure, or possibly other such choices) and subject to certain constraints (such as exactness, monotonicity, and so forth). An important type of less well-known decoration (for nonembedded Lagrangians) is a marking, as given in Definition 2.1.5.

The class $\mathcal{L}ag^*(\mathbb{C} \times M)$ is closed with respect to Hamiltonian isotopy horizontal at infinity and it contains all cobordisms of the form $\gamma \times L$, where $L \in \mathcal{L}ag^*(M)$ and γ is a curve in $\mathbb{C}$ horizontal (and of nonnegative integer heights) at $\pm\infty$. This also implies that Lagrangian suspension cobordisms for Lagrangians in $\mathcal{L}ag^*(M)$ are in $\mathcal{L}ag^*(\mathbb{C} \times M)$ (the reason is that they can be obtained by Hamiltonian isotopies of the trivial cobordism). We will also assume that both classes $\mathcal{L}ag^*(M)$ and $\mathcal{L}ag^*(\mathbb{C} \times M)$ are closed under (disjoint) finite unions.

REMARK 2.1.2

It is useful to be more explicit at this point concerning the properties of the class of Lagrangian immersions $\mathcal{L}ag^*(M)$. We assume that all the immersions $j_L : L \to M$ in $\mathcal{L}ag^*(M)$ are such that the restriction of j_L to any connected component of L only has isolated, transverse double points and that L has a finite number of connected components. Moreover, as mentioned above, the class $\mathcal{L}ag^*(M)$ (as well as $\mathcal{L}ag^*(\mathbb{C} \times M)$) is required to be closed under disjoint union. In other words, given two immersions $j_{L_1} : L_1 \hookrightarrow M$, $j_{L_2} : L_2 \hookrightarrow M$ in the class $\mathcal{L}ag^*(M)$, the obvious immersion defined on the disjoint union $j_{L_1} \sqcup j_{L_2} : L_1 \sqcup L_2 \hookrightarrow M$ is also in the class $\mathcal{L}ag^*(M)$. This operation gives rise to Lagrangians with nonisolated double points (e.g., by "repetition" of the same immersion $j_L : L \to M$) but, clearly, the restriction of the resulting immersion to the connected components of the domain only has isolated, transverse double points if j_{L_i}, $i = 1, 2$ have this property.

A similar definition applies to the class $\mathcal{L}ag^*(\mathbb{C} \times M)$ that consists of immersed cobordisms such that each connected component of such an immersion has only isolated, transverse self-intersection points or, possibly, clean intersections of dimension 1.

The morphisms in $\mathsf{C}ob^*(M)$ between two objects L and L' are the cobordisms $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$ (with m arbitrary), $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$ modulo Hamiltonian isotopy horizontal at infinity (see again [10]), as in Figure 3 (and possibly modulo some additional identifications, having to do with the decorations included in $*$). When viewing a cobordism V as before as a morphism $V : L \to L'$, we will refer to the ends $L_1, \ldots, L_m$ as the *secondary ends* of the morphism V.

The composition of two morphisms $V : L \to L'$, $V' : L' \to L''$ is given by the composition of the cobordism V with V' by gluing the upper "leg" of V to the input of V', as in Figure 4. We assume that the composition of the cobordisms in the class $\mathcal{L}ag^*(\mathbb{C} \times M)$ along *any* possible "leg"—not only on the top leg of V—preserves the class $\mathcal{L}ag^*(\mathbb{C} \times M)$.

As mentioned before, the connected immersed objects in $\mathcal{L}ag^*(M)$ are supposed to only have (isolated) self-transverse double points and the connected immersed objects in $\mathcal{L}ag^*(\mathbb{C} \times M)$ have clean intersections with the double points along manifolds (possibly with boundary) of dimension at most 1.

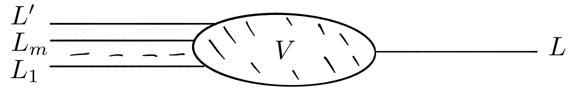

Figure 3. A morphism $V : L \to L'$.

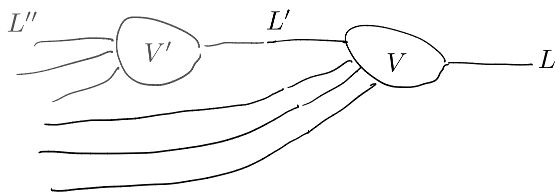

Figure 4. (Color online) A cobordism $V : (L'_j) \rightsquigarrow (L_i)$ projected on $\mathbb{R}^2$. The composition $V' \circ V$ of $V : L \to L'$ and $V' : L' \to L''$.

REMARK 2.1.3

It seems that under additional assumptions on the ambient manifold M and a further restriction of the class of Lagrangians, the category $\mathsf{C}ob^*(M)$ can be endowed with a grading and a shift functor. More specifically, a typical assumption on M would be that it is Calabi–Yau and the Lagrangian immersions L making the objects of $\mathsf{C}ob^*(M)$ are endowed with a grading in the sense of [47], that is, a lift of the Gauss map $L \longrightarrow S^1$, to an $\mathbb{R}$-valued function called a grading. Similarly, the cobordisms that serve as morphisms in this category are assumed to be endowed with a grading too. Due to the cylindrical behavior at infinity, a graded cobordism induces a grading on its ends. The shift functor acts, as usual, by adding an appropriate integral constant to the grading function (and similarly for the grading on the cobordisms which are the morphisms in this category). We refer the reader to [47] for more details on graded Lagrangian submanifolds. See also [3] and [2] for grading in the framework of immersed Lagrangians, and [32] for the context of Lagrangian cobordism.

Endowing $\mathsf{C}ob^*(M)$ with a grading really becomes useful when one considers the triangulated structure on the associated category $\widehat{\mathsf{C}}ob^*(M)$ introduced in Section 2.3, and the related functors to Fukaya categories as in Theorem 4.1.1. To illustrate this idea, consider for example the shuffling operation depicted in Figure 5, taking grading into account. Denote by T the previously mentioned shift functor. Shuffling the ends of a graded cobordism V as in the upper right side of Figure 5 results in a new graded cobordism $T^{-1}L' \rightsquigarrow (T^{-1}L, L'')$. Shifting the grading on this cobordism results in a graded cobordism $L' \rightsquigarrow (L, TL'')$, representing the morphism Ru. Similarly, the lower part of the same figure depicts a cobordism $TL'' \rightsquigarrow (L', TL)$, which after a negative shift becomes a cobordism $L'' \rightsquigarrow (T^{-1}L', L)$, representing $R^{-1}u$. Putting the three morphisms $R^{-1}u$, u, and Ru together, we obtain a distinguished triangle

$$L'' \longrightarrow L \longrightarrow L' \longrightarrow TL'',$$

as expected (from the grading viewpoint) in triangulated categories.

Of course, there are many more details to be worked out in order to establish a graded theory for the remainder of the present article, and we have not verified them in full detail. Therefore, throughout the following, we will always work in an ungraded setting.

We next discuss a few basic notions related to this framework.

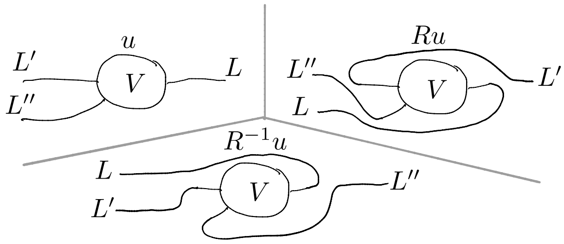

Figure 5. The three morphisms $(R^{-1}u, u, Ru)$ in a distinguished triangle.

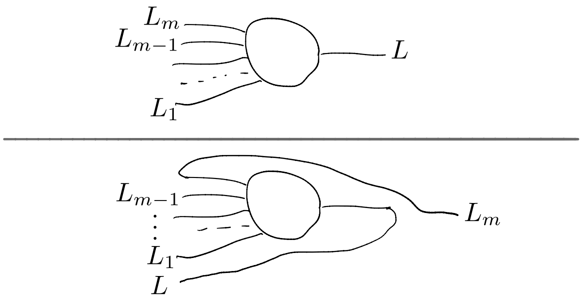

Figure 6. The morphism v on top and Rv at the bottom.

2.1.1. Distinguished triangles

Cobordisms $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$ with three ends $V : L \rightsquigarrow (L'', L')$ play an important role here. Let $u : L \to L'$ be the morphism represented by V. There are two other morphisms $Ru : L' \to L''$ and $R^{-1}u : L'' \to L$ that are obtained by "shuffling" the ends of V as in Figure 5.

The triple $(u, Ru, R^{-1}u)$ will be called a *distinguished triangle* in $\mathsf{C}ob^*(M)$, and we refer to the shuffling R as a *rotation* (operator).

REMARK 2.1.4

For a cobordism $V : L \to (L'', L')$ that represents a morphism $u : L \to L'$ as above, we have $RR^{-1}u = u$, $R^3u = u$ (recall that two cobordisms induce the same morphism if they are horizontally Hamiltonian isotopic). In particular, an alternative writing of an exact triangle $(u, Ru, R^{-1}u)$ as above is (u, Ru, R^2u).

For a general cobordism $V : L \to (L_1, \ldots, L_m)$ inducing a morphism $v : L \to L_m$, we let $Rv : L_m \to L_{m-1}$ be the morphism induced by the cobordism obtained from V by shuffling the L, L_m, and L_{m-1} ends as in Figure 6 (after shuffling L becomes the bottom negative end; L_m is the unique positive end; L_{m-1} is the top negative end). We have obvious relations similar to those in Remark 2.1.4. In this case too, for instance, $R^{m+1}v = v$. More generally, given a cobordism V as above there are unique morphisms (defined up to horizontal Hamiltonian isotopy)

$L_{j+1} \to L_j$ and $L_1 \to L$ that are associated to V and defined as an appropriate iterated rotation of v. To simplify terminology, we will refer in the following to such a morphism $L_{j+1} \to L_j$ or $L_1 \to L$ associated to V as the *rotation of v with domain* L_{j+1} (resp., L_1). For instance, in Figure 5 at the right we have the rotation of v with domain L' and, at the bottom, the rotation with domain L''.

We assume that the class $\mathcal{L}ag^*(\mathbb{C} \times M)$ is closed with respect to these rotations.

2.1.2. Surgery and marked Lagrangians

Lagrangian surgery, as introduced by Lalonde–Sikorav [34] for surfaces and Polterovich [43] in full generality, is one of the basic operations with Lagrangian submanifolds. It was remarked in [9] that the trace of a Lagrangian surgery is itself a Lagrangian cobordism. We will use the conventions in [9] and will apply the construction in a situation a bit more general than [9] to two—possibly immersed—Lagrangians L_1 and L_2. We assume that L_i, $i = 1, 2$ has only self-transverse double points and that L_1 and L_2 intersect transversely at smooth points. The surgery will be performed on a subset c of the intersection points of L_1 and L_2 and the outcome will be denoted by $L_1 \#_{c,\epsilon} L_2$ and is again a (generally) immersed Lagrangian with self-transverse double points. The ϵ that appears in the notation is a parameter that is used to keep track of the size of the handle used in the surgery (it equals the area of the region encompassed by the handle in $\mathbb{C}$). We consider here only surgeries using the same model handle for all double points in c.

The projection on $\mathbb{C}$ of the trace of a surgery is as in Figure 7. The model for the surgery in the neighborhoods of the points where the surgery handles are attached is as described in [9]. The projection of this handle attachment is (up to a small smoothing) as the left drawing in Figure 8 below. On regions away from the surgered intersection points the local model is as in the drawing at the right in Figure 8. To glue the two local models, a further perturbation is required in the direction of the horizontal and vertical axes which leads to a projection as in Figure 7.

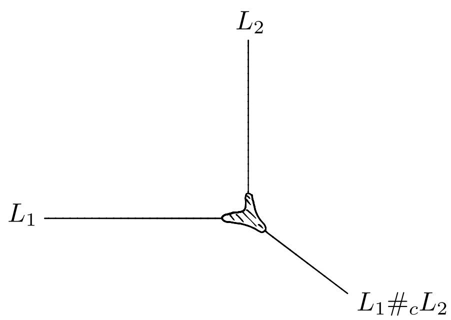

Figure 7. The projection of the trace of a surgery onto $\mathbb{C}$.

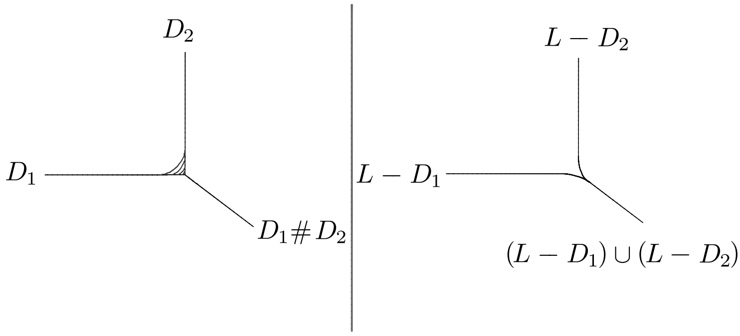

Figure 8. (Color online) The local models for the trace of the surgery around points where the surgery takes place, at the left, and where nothing happens, at the right. The sets $D_1 \subset L_1$ and $D_2 \subset L_2$ are small disks around the point of surgery.

In particular, it is clear that even if L_1, L_2 are embedded but c does not consist of all intersection points $L_1 \cap L_2$, then the trace of the surgery has double points along an interval of type $[0, \infty)$.

2.1.3. 0-*Size surgeries*

Given that we are working with immersed Lagrangians, it is natural to consider also "formal" 0-size surgeries $L_1 \#_c L_2 := L_1 \#_{c,0} L_2$ that correspond to $\epsilon = 0$. In this case $L_1 \#_c L_2$ is, by definition, the immersed Lagrangian $L_1 \cup L_2$ together with the set c viewed as an additional decoration. This seemingly trivial operation will play an important role later in the article and we formalize the structure of $L_1 \#_c L_2$ in the next definition.

For an immersed Lagrangian $j_L : L \to M$ with only transverse, isolated double points, we denote by $I_L \subset L \times L$ the set of double points of L:

$$I_L = \big\{(x, y) \in L \times L : j_L(x) = j_L(y)\big\}.$$

DEFINITION 2.1.5

A marked Lagrangian immersion (L, c) in M consists of a Lagrangian immersion $j_L : L \to M$ with isolated, transverse double points and a choice of a subset $c \subset I_L$ called a *marking* of L.

REMARK 2.1.6

We will always assume that only one of the pairs (x, y) or (y, x) belongs to a subset c as above.

The 0-size surgery $L_1 \#_c L_2$ is a marked Lagrangian in this sense, the relevant immersion is $L_1 \sqcup L_2 \hookrightarrow M$, and the marking c consists of the family of intersection points $c \subset L_1 \cap L_2$ with each point $P \in c$ lifted back to L_1 and L_2 as a couple $(P_-, P_+) \in L_1 \times L_2 \subset (L_1 \sqcup L_2) \times (L_1 \sqcup L_2)$. Once the class of Lagrangians is extended as before to marked ones, we may also consider surgery as an operation with marked Lagrangians. The construction discussed above extends trivially to this case. For instance, given two marked Lagrangians (L_1, c_1), (L_2, c_2), and $c \subset (L_1, L_2)$ a subset of intersection points of L_1 and L_2, the 0-size surgery $(L_1, c_1)\#_c(L_2, c_2)$ is the union $L_1 \cup L_2$ marked by $c_1 \cup c \cup c_2$. We also allow for ϵ-surgeries $(L_1, c_1)\#_{c,\epsilon}(L_2, c_2)$ with $\epsilon > 0$ in which case the new marking is simply $c_1 \cup c_2$. A nonmarked Lagrangian is a special case of a marked Lagrangian whose marking is the void set.

Recall from Remark 2.1.2 that the elements of $\mathcal{L}ag^*(M)$ are unions of immersions with isolated, transverse double points. A marking of such an element $L \in \mathcal{L}ag^*(M)$ consists of a set $c = \{(P_-, P_+)\} \subset L \times L$ such that if $L_\pm \subset L$ is, respectively, the component of L containing, respectively, $P_\pm$, then L_- intersects L_+ transversely at the points corresponding to $P_\pm$ (or self-transversely in case $L_- = L_+$). In other words, the class $\mathcal{L}ag^*(M)$ consists of (possibly) marked Lagrangians in the sense of Definition 2.1.5 and their disjoint unions.

To define the surgery operation for two elements $(L_1, c_1), (L_2, c_2) \in \mathcal{L}ag^*(M)$, we impose the condition that the set $c \neq \emptyset$ (that is used to define ϵ-size surgery, $\epsilon \geq 0$) consist of a finite set of double points $(P_-, P_+) \in (L_1, L_2)$ such that the connected component of L_1 containing P_- intersects transversely (and away from the double points of L_i, $i = 1, 2$) the connected component of L_2 containing P_+. Notice that with this convention the disjoint union $j_{L_1} \sqcup j_{L_2}$ can be viewed as marked 0-size surgery with $c = \emptyset$ and it does not require any transversality among the components of L_1 and L_2.

The trace of a 0-size surgery projects globally as the drawing on the right in Figure 8. This trace is obviously an immersed Lagrangian such that each connected component has clean intersections having 1-dimensional double point manifolds corresponding to the intersection points $L_1 \cap L_2$. In the following, we will also need to use marked cobordisms. Due to the presence of these 1-dimensional double point sets, the definition of this notion is a bit more complicated, and we postpone it to Section 3.2. The trace of the 0-size surgery $L_1\#_cL_2$ is marked in this sense in a particularly simple way (see Section 4.3.1). Both $L_1\#_cL_2$ as well as the trace of this surgery can obviously be viewed as limits of $L_1\#_{c,\epsilon}L_2$ and, respectively, of the corresponding surgery trace, when $\epsilon \to 0$.

REMARK 2.1.7

(a) To include the 0-size surgery and its trace inside a cobordism category $\mathsf{C}ob^*(M)$ a bit more elaboration is needed. In essence, in this case, the class $\mathcal{L}ag^*(M)$ needs to contain marked Lagrangians (as always, subject to further constraints included in the condition $*$), similarly the elements of $\mathcal{L}ag^*(\mathbb{C} \times M)$ are marked cobordisms. A marked cobordism (V, v) between two tuples of marked Lagrangians $\{(L_i, c_i)\}_i$ and $\{(L'_j, c'_j)\}_j$ is a cobordism as in Definition 2.1.1 with the additional property that the restriction of the marking v to the ends L_i, L'_j coincides, respectively, with the markings c_i, c'_j. The composition of marked cobordisms V, V' is again a marked cobordism in the obvious way, by concatenating the markings to obtain a marking on the gluing of V and V'. Full details appear in Section 3.2. Obviously, the class of usual "nonmarked" Lagrangians and cobordisms is included in this setup by simply associating to them the empty marking. Whenever discussing/using 0-size surgeries below we will always assume that the classes of Lagrangians involved are as described above.

(b) While markings behave in a way similar to other structures that are suitable to incorporation in cobordism categories such as choices of orientation or spin structure, there is a significant distinction: immersed marked Lagrangians are *more general* objects compared to nonmarked ones. It is useful to imagine them as the result of a two-stage extension process *embedded* $\to$ *immersed* $\to$ *marked.*

The trace of a surgery as in Figure 7 above will be viewed as a morphism in a cobordism category by twisting the ends as in Figure 9 below. This cobordism (and the associated morphism) will be denoted by $S_{L_2,L_1;c,\epsilon}$. We will refer to $S_{L_2,L_1;c,\epsilon}$ as the *surgery morphism* from L_2 to L_1 relative to c and to a handle of area ϵ. The same conventions and notation apply to 0-size surgeries, under the assumptions in Remark 2.1.7. The morphism associated to a 0-size surgery associated to c is denoted by $S_{L_2,L_1;c}$. In the figures below we will not distinguish the case $\epsilon = 0$ and $\epsilon > 0$.

Notice that the set of all 0-size surgery morphisms from L_2 to L_1 as before are in bijection with the choices of the subset c of intersection points of the two Lagrangians. In turn, we can view c as an element in the $\mathbb{Z}/2$-vector space $\mathbb{Z}/2\langle L_1 \cap L_2\rangle$. In short, we will identify the set of all 0-size surgery morphisms from L_2 to L_1, $\bar{S}_{L_2,L_1}$, with the vector space

$$\bar{S}_{L_2,L_1} \equiv \mathbb{Z}/2\langle L_1 \cap L_2\rangle. \tag{2}$$

This identification endows $\bar{S}_{L_2,L_1}$ with a linear structure. Assuming fixed choices of handle models, small ϵ, and fixed choices of Darboux charts around the intersection points of L_1 and L_2, the same applies to the set of ϵ-size surgery morphisms $S_{L_2,L_1;c,\epsilon}$.

In some of the constructions below we will consider Lagrangians in $\mathbb{C} \times M$ whose projection onto $\mathbb{C}$ will contain a picture as in Figure 10 below. In other words, the projection will contain two segments γ_1 and γ_2 that intersect transversely at one point $P \in \mathbb{C}$. Their preimage in $\mathbb{C} \times M$ coincides with $\gamma_1 \times L_1$ (resp., $\gamma_2 \times L_2$), with L_1, L_2 two Lagrangian submanifolds in M that intersect

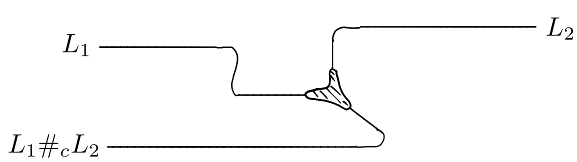

Figure 9. The morphism $S_{L_2,L_1;c,\epsilon}$ associated to the trace of the surgery of L_1 and L_2 along $c \subset L_1 \cap L_2$.

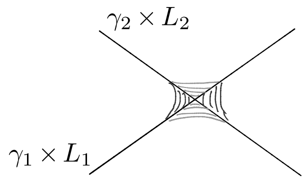

Figure 10. (Color online) The crossing of $\gamma_1 \times L_1$ and $\gamma_2 \times L_2$ and a positive surgery—in blue—and a negative surgery—in red—at the crossing point.

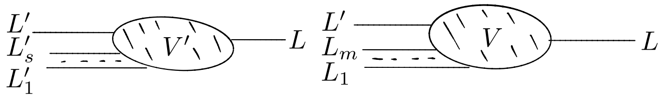

Figure 11. Two cobordisms $V : L \rightsquigarrow (L_1, \dots, L_m, L')$, $V' : L \rightsquigarrow (L'_1, \dots, L'_s, L')$.

transversely. For any point $x \in L_1 \cap L_2$, we can operate two types of surgeries in $\mathbb{C} \times M$ between $\gamma_1 \times L_1$ and $\gamma_2 \times L_2$ at the point $P \times x$. A *positive* surgery whose handle has a projection like the region in blue in the figure or a *negative* surgery whose projection is as the region in red. More generally, we will talk about a $+$ surgery at P (resp., a $-$ surgery at P) to refer to simultaneous positive (resp., negative) surgeries along a subset $c \subset L_1 \cap L_2$. Such a surgery will be denoted by $\#^+_{c,\epsilon}$ (resp., $\#^-_{c,\epsilon}$) when the size of the handles is ϵ. Again, we allow for $\epsilon = 0$ and the notation in that case is $\#^{\pm}_c$. In terms of the relevant marking, the difference between a positive 0-size surgery and a negative one at some point $P \in (\gamma_2 \times L_2) \cap (\gamma_1 \times L_1)$ is simply the order of the two points $P_- \in \gamma_1 \times L_1$, $P_+ \in \gamma_2 \times L_2$: the marking is (P_-, P_+) for a positive surgery and (P_+, P_-) for a negative one. For positive surgeries, we will generally omit the superscript $^+$.

2.1.4. Cabling

We now consider a simple construction, based on Lagrangian surgery, that will play an important role in our framework.

Consider two cobordisms $V : L \rightsquigarrow (L_1, \dots, L_m, L')$ and $V' : L \rightsquigarrow (L'_1, \dots, L'_s, L')$ as in Figure 11, and assume that L and L' intersect transversely.

A *cabling* $\mathcal{C}(V, V'; c, \epsilon)$ of V and V' relative to $c \in L \cap L'$ (and relative to a handle of size ϵ) is a cobordism $\mathcal{C}(V, V'; c, \epsilon) : L_1 \to (L_2, \dots, L_m, L'_1, \dots, L'_s)$ whose projection is as in Figure 12. More explicitly, the positive end of V is bent, above V, toward the negative direction and is glued to the positive end of V'. Similarly, the top negative end of V' is bent above V' toward the positive direction and is glued to the top negative end of V. The extensions of these two ends intersect above a point P in the plane in a configuration similar to that in Figure 10. At this point P are performed negative surgeries along the subset $c \subset L \cap L'$ (more precisely, $c \subset L \times L' \subset (V \sqcup V') \times (V \sqcup V')$). Finally, the L_1 end is bent below V in the positive direction. Again, we allow for $\epsilon = 0$ which leads to a marked cobordism with the marking associated to negative surgeries of size 0 along c.

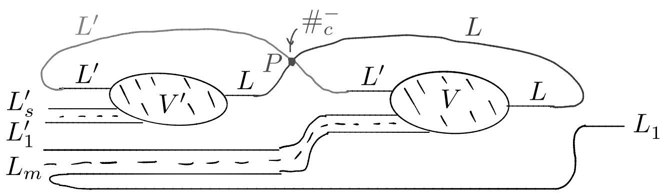

Figure 12. (Color online) The cabling $\mathcal{C}(V, V'; c)$ of V and V' relative to c.

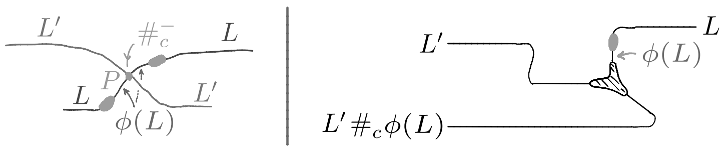

Figure 13. (Color online) The perturbed cabling around P using an isotopy ϕ such that $\phi(L)$ intersects L' transversely. Similarly, using a perturbation to define the surgery morphism $S_{L,L';c}$.

The notation in this case is $\mathcal{C}(V, V'; c)$. In case L and L' are themselves marked Lagrangians, then the marking of the cabling $\mathcal{C}(V, V'; c)$ at the point P includes the markings of L and of L'.

REMARK 2.1.8

(a) As mentioned before, the surgery operation (and thus also cabling) depends not only on the choice of the subset c of intersection points but also on the choice and size of surgery handles, on the choice of Darboux charts around the point of surgery and, in the case of cabling, on the choice of the bounded regions inside the bent curves, on one side and the other of P, in Figure 12.

(b) Surgery as well as cabling have been defined before only if the relevant Lagrangians are intersecting transversely. However, if this is not the case, then one can use a small Hamiltonian isotopy to achieve transversality. For instance, in the cabling case, assume that ϕ is a small Hamiltonian isotopy such that $\phi(L)$ intersects L' transversely. In this case the region around P in Figure 12 is replaced with the configuration on the left in Figure 13. The regions in gray are projections of small Lagrangian suspension cobordisms, the one to the left of P associated to ϕ and the one to the right of P associated to ϕ^{-1}. In this case $c \subset \phi(L) \cap L'$. We still use the notation $\mathcal{C}(V, V'; c)$ to denote a cabling in this case even if it obviously also depends on the choice of ϕ. A similar perturbation can be used for the surgery morphisms, as represented in the drawing on the right in Figure 13. Obviously, in both these perturbed cases $c \subset \phi(L) \cap L'$.

The cabling construction leads to the following definition that is central for us.

DEFINITION 2.1.9

Let V, V' be two cobordisms in $\mathcal{L}ag^*(\mathbb{C} \times M)$. We say that V and V' are related by cabling and write $V \sim V'$ if there exists a cabling $\mathcal{C}(V, V'; c, \epsilon) \in \mathcal{L}ag^*(\mathbb{C} \times M)$.

The meaning of this definition when L and L' are not intersecting transversely is that there should exist a perturbation ϕ as above and a c associated to it such that the resulting cabling $\mathcal{C}$ belongs to the class $\mathcal{L}ag^*(\mathbb{C} \times M)$. In case we need to indicate the precise cabling in question, we will write

$$\mathcal{C}(V, V'; c, \epsilon) : V \sim V'.$$

We allow for ϵ to be 0, and adopt the notation

$$\mathcal{C}(V, V'; c) : V \sim V'$$

in this case. As discussed in Remark 2.1.7, this requires that all the classes of Lagrangians and cobordisms under consideration be marked. In case the context allows for it, we sometimes refer to a cabling simply by $\mathcal{C}(V, V') : V \sim V'$ with the understanding that this represents a cabling of V and V' associated to an $\epsilon \geq 0$ and a c as above.

2.2. Axioms and surgery models for $\mathsf{C}ob^*(M)$

Recall the cobordism category $\mathsf{C}ob^*(M)$ whose morphisms are Lagrangian cobordisms modulo horizontal isotopy. We will say that the cobordism category $\mathsf{C}ob^*(M)$ *has surgery models* if it satisfies the following five axioms.

AXIOM 1 (Homotopy)

The cabling relation descends to $\mathrm{Mor}_{\mathsf{C}ob^*(M)}$ *and it gives an equivalence relation* $\sim$ *that is preserved by the composition of cobordisms along* any *end, by the union of cobordisms as well as by the insertion of the void end among the negative ends of a cobordism.*

The first part of this axiom means that two cobordisms that are Hamiltonian isotopic by an isotopy horizontal at infinity are cabling equivalent and that the resulting relation on $\mathrm{Mor}_{\mathsf{C}ob^*(M)}$ is an equivalence. We intend to make more explicit the further properties of this relation as implied by this axiom. To simplify notation, from now on we will no longer distinguish between a morphism $U \in \mathrm{Mor}_{\mathsf{C}ob^*(M)}(L, L')$ and the underlying cobordism $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$.

The relation $\sim$ is preserved by the composition of cobordisms in the following sense. First, if $V \sim V'$ and U, W are morphisms in $\mathrm{Mor}_{\mathsf{C}ob^*(M)}$ such that the relevant compositions are defined, then $U \circ V \sim U \circ V'$ and $V \circ W \sim V' \circ W$. Moreover, let $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$, and consider another cobordism $K : L_i \rightsquigarrow (L'_1, \ldots, L'_s)$ such that the composition $V \circ_i K : L \to (L_1, \ldots, L_{i-1}, L'_1, \ldots, L'_s, L_{i+1}, \ldots, L_m, L')$ is defined by gluing V to K along the (secondary) L_i end of V, then $V \circ_i K$ viewed as a morphism $L \to L'$ has the property that $V \sim V \circ_i K$. The union statement means that if $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$

can be written as a union of cobordisms $V' : L \rightsquigarrow (L_{j_1}, \ldots, L')$ and a second cobordism $V'' : \emptyset \rightsquigarrow (L_{s_1}, \ldots, L_{s_r})$, where the ends L_{j_i}, L_{s_k} put together recover all the ends $L_1, \ldots, L_m$, then $V \sim V'$. Finally, the last property claims that two cobordisms $V : L \rightsquigarrow (L_1, \ldots, L_k, L')$ and $V' : L \rightsquigarrow (L_1, \ldots, L_i, \emptyset, L_{i+1}, \ldots, L_k, L')$ that differ only by the insertion of a void end (and the corresponding rearrangement of the other ends), are equivalent.

AXIOM 2 (Cabling reduction)
Suppose that V and V' are two cobordisms such that the cabling $\mathcal{C}(V, V'; c, \epsilon)$ is defined and belongs to $\mathcal{L}ag^*(\mathbb{C} \times M)$. Then $V, V' \in \mathcal{L}ag^*(\mathbb{C} \times M)$.

By inspecting Figure 12, this means that, assuming $\mathcal{C}(V, V'; c, \epsilon)$ is in $\mathcal{L}ag^*(\mathbb{C} \times M)$, by erasing in that picture the point P as well as the blue and green arcs, we are left with a configuration that remains in $\mathcal{L}ag^*(\mathbb{C} \times M)$.

For any $L \in \mathcal{L}ag^*(M)$, we denote by $\mathrm{id}_L : L \to L$ the cobordism given by $\mathbb{R} \times L \subset \mathbb{C} \times M$.

AXIOM 3 (Existence of inverse)
If $V \in \mathrm{Mor}_{\mathcal{C}ob^*(M)}(L, L')$ is a simple cobordism, then

$$\bar{V} \circ V \sim \mathrm{id}_L.$$

Recall that $\bar{V}$ is the cobordism obtained by rotating V using a 180° rotation in the plane.

AXIOM 4 (Surgeries)
For any $V \in \mathrm{Mor}_{\mathcal{C}ob^*(M)}(L, L')$, there exists a surgery morphism $S_{L,L';c,\epsilon} \in \mathcal{L}ag^*(\mathbb{C} \times M)$ such that $V \sim S_{L,L';c,\epsilon}$. Moreover, the sum of surgeries preserves $\mathcal{L}ag^*(\mathbb{C} \times M)$ and is compatible with $\sim$.

As in Remark 2.1.8(b), small Hamiltonian perturbations may be required for the surgery morphisms used here. In this case, the meaning of the axiom is that, for any Hamiltonian perturbation that puts L and L' in generic position, there exist both c and an associated surgery morphism that is $\sim$ equivalent to V. This type of perturbation is already necessary to associate to id_L an equivalent surgery morphism. The second part of this axiom claims that the set $\bar{S}^*_{L,L'} = \bar{S}_{L,L'} \cap \mathcal{L}ag^*(\mathbb{C} \times M)$ of those surgeries that are in the class $\mathcal{L}ag^*(\mathbb{C} \times M)$ is a vector subspace of $\bar{S}_{L,L'}$ (this depends on choices of perturbations if L and L' are not in general position). In particular, as $0 = 2b$ for any $b \in \bar{S}_{L,L'}$ this means that the 0-surgery (i.e., the "absence" of surgery) is in $\bar{S}^*_{L,L'}$. The compatibility with $\sim$ means that the set of equivalence classes $\bar{S}^*_{L,L'}/\sim$ is a $\mathbb{Z}/2$-vector space (and even if L and L' do not intersect transversely, this implies that the relevant vector space is independent of perturbations up to isomorphism). In particular, this allows one to define an operation $(V, V') \to V + V'$ well defined up to the equivalence $\sim$.

REMARK 2.2.1

In the main example discussed in a later section, ϵ in Axiom 4 can be taken to be 0. However, we do not impose this condition in the axiom because this requirement would automatically imply that any category with surgery models necessarily consists of marked Lagrangians.

The next axiom is considerably more difficult to unwrap.

AXIOM 5 (Naturality)

Fix $V : L \to L'$, $V' : K \to L'$, $V'' : L \to K'$, $S : K \to L$, $S' : L' \to K'$ *in* $\mathrm{Mor}_{\mathcal{C}ob^*(M)}$ *such that* S *and* S' *have exactly three ends. If* $\mathcal{C}(V \circ S, V') : V \circ S \sim V'$ *belongs to* $\mathcal{L}ag^*(\mathbb{C} \times M)$, *then*

$$\mathcal{C}(V \circ S, V') \circ RS \sim RV' \circ V. \tag{3}$$

Similarly, if $\mathcal{C}(V'', S' \circ V) : V'' \sim S' \circ V$ *belongs to* $\mathcal{L}ag^*(\mathbb{C} \times M)$, *then*

$$R^{-1}S' \circ \mathcal{C}(V'', S' \circ V) \sim V \circ R^{-1}V''. \tag{4}$$

We geometrically explain this axiom in Figures 14, 15, 16, and 17.

To better understand the meaning of this axiom, see Remark 2.3.3(b).

In the axioms above we allow 0-size surgeries and cabling as long as the respective classes $\mathcal{L}ag^*(M)$ and $\mathcal{L}ag^*(\mathbb{C} \times M)$ are enlarged to include marked Lagrangians and cobordisms (see also Remark 2.1.7).

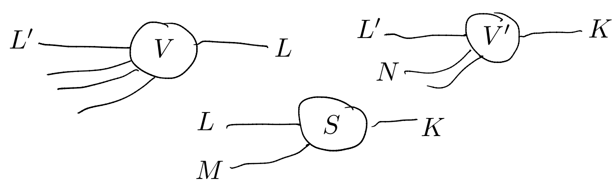

Figure 14. The three cobordisms $V : L \to L'$, $V' : K \to L'$, and $S : K \to L$.

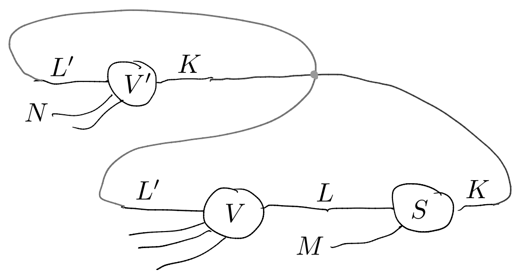

Figure 15. (Color online) The cabling $\mathcal{C}(V \circ S, V')$.

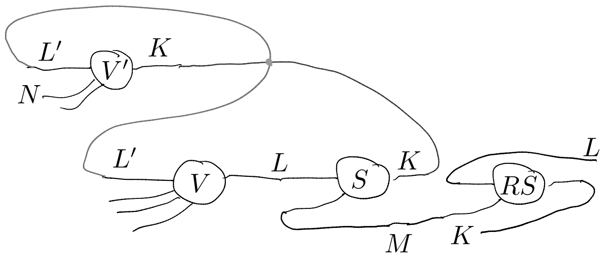

Figure 16. (Color online) The composition $C = \mathcal{C}(V \circ S, V') \circ RS$ viewed as a morphism $C : L \to N$.

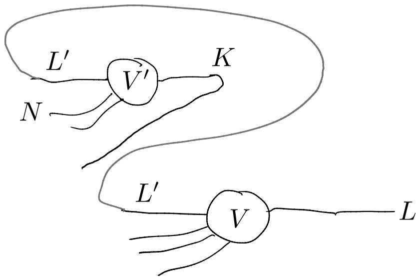

Figure 17. (Color online) The cobordism $C' = RV' \circ V$ can be seen as obtained from C by erasing from C the blue arc as well as S and RS. Axiom 5 claims that C is equivalent to C'.

2.3. The quotient category $\widehat{\mathsf{C}ob}^*(M)$ is triangulated

We assume that $\mathsf{C}ob^*(M)$ has surgery models so that, in particular, by Axiom 1, the cabling relation $\sim$ is an equivalence. We define $\widehat{\mathsf{C}ob}^*(M)$ to be the quotient category of $\mathsf{C}ob^*(M)$ by this equivalence relation. As usual, the objects of $\widehat{\mathsf{C}ob}^*(M)$ are the same as the objects of $\mathsf{C}ob^*(M)$ and the morphisms are the classes of equivalence of the morphisms in $\mathsf{C}ob^*(M)$.

Here is the main property of the structures described above.

THEOREM 2.3.1

If the cobordism category $\mathsf{C}ob^*(M)$ has surgery models, then the quotient category $\widehat{\mathsf{C}ob}^*(M)$ is triangulated with the exact triangles represented by surgery distinguished triangles as in Section 2.1.1.

Proof

The proof is elementary. To start, we remark that in the somewhat naive setting here the suspension functor T is taken to be the identity.

(a) $\widehat{\mathsf{C}ob}^*(M)$ *is additive.* We start by noting that, in view of Axiom 4, the set of morphisms between L and L' in $\widehat{\mathsf{C}ob}^*(M)$ has the structure of a $\mathbb{Z}/2$-vector space. The sum of objects is given by the union $L \oplus L' = L \sqcup L'$, endowed with the obvious immersion $L \sqcup L' \hookrightarrow M$ that restricts to the immersions of L and L' (see also Remark 2.1.2).

(b) *Exact triangles.* We define an exact triangle in $\widehat{\mathsf{C}}ob^*(M)$ as a triangle isomorphic to a triangle represented by a distinguished triangle—as defined in Section 2.1.1—in $\mathsf{C}ob^*(M)$. Clearly, as any morphism in $\mathsf{C}ob^*(M)$ is equivalent to a surgery morphism which is itself part of an obvious distinguished triangle, we deduce that any morphism in $\widehat{\mathsf{C}}ob^*(M)$ can be completed to an exact triangle.

REMARK 2.3.2

To further understand the interplay of the different axioms, we shall now describe another way of completing a morphism $V : L \to L'$ associated to a cobordism $V : L \rightsquigarrow (L_1, \ldots, L_m, L')$ to an exact triangle. We will assume that $m = 3$ and discuss the construction by following Figures 18, 19, and 20. We start with the cobordism V in Figure 18. By Axiom 4, the rotation R^3V with domain L_2 is equivalent to a surgery morphism $S_{L_2,L_1;c}$. We use this surgery morphism to construct the cabling of R^3V and $S_{L_2,L_1;c}$ and this produces (after another rotation) the cobordism V' in Figure 19. Here and below, we omit $\epsilon \geq 0$ from the notation.

We now reapply the same argument, but this time for the bottom two negative ends of V', obtaining the cobordism V'' in Figure 20 that represents a morphism $V'' : L \to L'$ and has three ends and thus fits directly into an exact triangle.

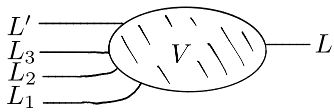

Figure 18. The initial cobordism V.

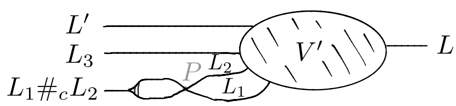

Figure 19. (Color online) The cobordism V' resulting from a cabling of a rotation of V with a surgery of the ends L_2 and L_1.

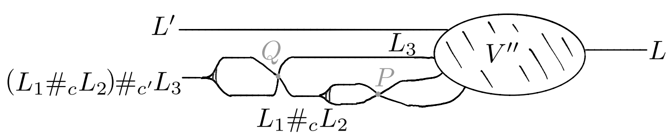

Figure 20. (Color online) The final cobordism V'' obtained by iterating the operation before one more step.

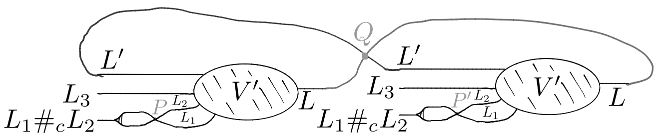

Figure 21. (Color online) The cabling of V' and V'.

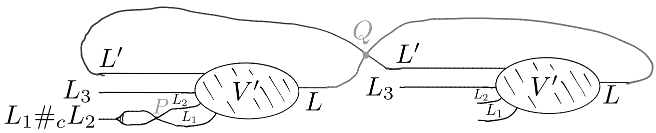

Figure 22. (Color online) The cabling of V' and V.

To end this discussion, we want to remark that the morphism $V'' : L \to L'$ is equivalent to V. Given that the relation $\sim$ is an equivalence, this follows if we can show that $V \sim V'$ where V' is also viewed as a morphism $V' : L \to L'$. To show this, we proceed as follows. By symmetry, $V' \sim V'$ so that a cabling as in Figure 21 belongs to $\mathcal{L}ag^*(\mathbb{C} \times M)$. In this cabling we denote by P' the crossing point corresponding to one of the copies of V' and by P the other such point.

We now use Axiom 2 to "erase" the point P' and remove the part of the picture at the left of P'—as in Figure 22—and to deduce that the remaining portion, which is precisely the cabling of V' and V, belongs to $\mathcal{L}ag^*(\mathbb{C} \times M)$ and thus $V \sim V'$.

For further reference, notice that the type of braiding of two consecutive ends as in Figures 19 and 20 produces a Lagrangian representing the cone of the morphism relating the two ends. For instance, the cone of the morphism R^3V is (isomorphic) to the Lagrangian $L_1 \#_c L_2$ in Figure 19. Using also Axiom 3, it is easy to see that, up to isomorphism, the cone depends only on the equivalence class of a morphism.

(c) *Naturality of exact triangles.* We now consider the following diagram in $\mathsf{C}ob^*(M)$.

$$
\begin{tikzcd}
L \arrow[r, "S"] \arrow[d, "V"'] & M \arrow[r, "T"] \arrow[d, "V'"'] & K \arrow[r, "U"] \arrow[d, dotted, "V''"'] & L \arrow[d, "V"'] \\
L' \arrow[r, "S'"] & M' \arrow[r, "T'"] & K' \arrow[r, "U'"] & L'
\end{tikzcd}
\tag{5}
$$

The morphisms S, V, V', and S' are given and $V' \circ S \sim S' \circ V$. Each of these morphisms can be completed to exact triangles corresponding to the horizontal solid lines. Explicitly, this means that the cobordisms S, V, V', and S' can each be replaced by cobordisms with exactly three ends that represent the same respective

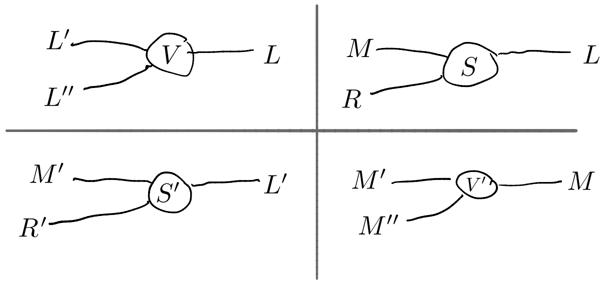

Figure 23. The morphisms in the left square of (5).

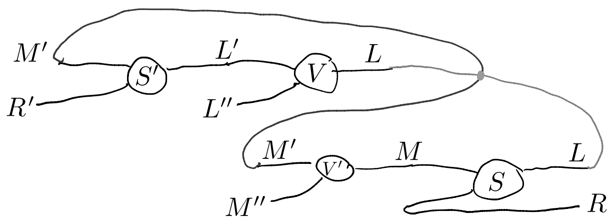

Figure 24. (Color online) The cabling of $S' \circ V$ and $V' \circ S$ viewed as a morphism $V'' : K \to K'$.

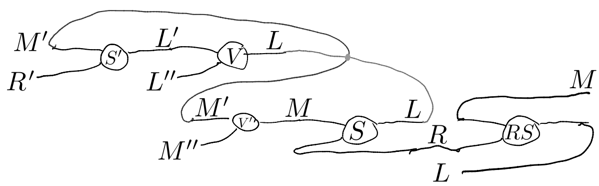

Figure 25. (Color online) The composition of $V'' : K \to K'$ and $T = RS : M \to K$.

morphisms in $\widehat{\mathsf{C}ob}^*(M)$. Notice that there are many ways to do this replacement (we have indicated two such ways, equivalence with a surgery morphism and braiding, above). Assuming from now on that S, V, V', and S' are each given by cobordisms with three ends, we have $T = RS$, $T' = RS'$ and $U = R^2 S$, $U' = R^2 S'$.

We want to show that there exists a morphism $V'' : K \to K'$ such that $V'' \circ T \sim T' \circ V'$ and $V \circ U \sim U' \circ V''$. We will again use a series of figures to pursue this construction. We start with Figure 23, where we fix the four morphisms V, S, V', S' as above.

The next step is in Figure 24, which describes the cabling of $S' \circ V$ and $V' \circ S$—this exists in $\mathcal{L}ag^*(\mathbb{C} \times M)$ because $V' \circ S \sim S' \circ V$—and views it as a morphism $V'' : K \to K'$.

For the next step, we will make use of Axiom 5 to show the commutativity $V'' \circ T \sim T' \circ V'$. We start in Figure 25 by representing the composition $V'' \circ T$.

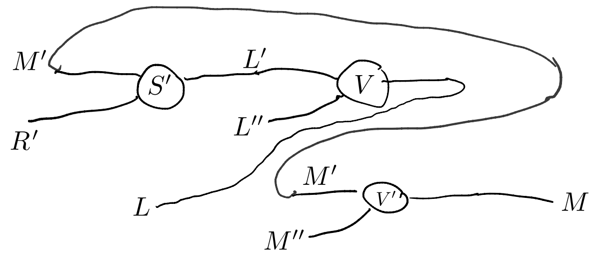

Figure 26. (Color online) Here is the cobordism V''' obtained from $V'' \circ T$ after erasing the green arc as well as S and RS (and bending the L end in the positive direction).

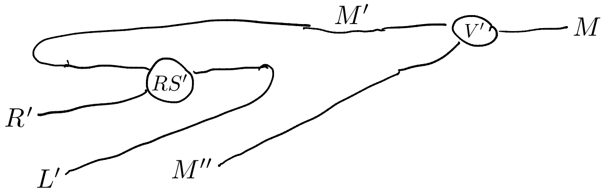

Figure 27. The morphism $T' \circ V'$.

We now apply Axiom 5 to "erase" the green arc and the cobordisms S and RS in the picture thus obtaining a new morphism $V''' \sim V'' \circ T$ as in Figure 26.

The composition of $T' \circ V'$ is represented in Figure 27. We notice that V''' is equal to the composition of $T' \circ V'$ with a rotation of V along a secondary leg (with end L'). In view of Axiom 1, this means that V''' is equivalent to the cobordism $T' \circ V'$, which concludes the argument. The argument required to show that $V \circ U \sim U' \circ V''$ is perfectly similar, appealing to the second identity in Axiom 5.

REMARK 2.3.3

(a) It follows from the naturality of exact triangles that if $V : L \rightsquigarrow (K, L')$, then $RV \circ V \sim 0$. This is in fact an instance of a more general fact, valid in any triangulated category—the composition of any two consecutive morphisms in an exact triangle always vanishes.

(b) Some special choices in diagram (5), as in diagram (6) below, indicate that *naturality* is an adequate name for Axiom 5.

(6)
$$\begin{tikzcd} K \arrow[r, "S"] \arrow[d, "\mathrm{id}"'] & L \arrow[r, "RS"] \arrow[d, "V"'] & M \arrow[d, dotted, "V''"'] \\ K \arrow[r, "V'"] & L' \arrow[r, "RV'"] & N \end{tikzcd}$$

In this setting, notice that the definition of V'' given above is precisely the cabling $\mathcal{C}(V \circ S, V')$ and thus the first identity in Axiom 5 simply claims that the square on the right in diagram (6) commutes in $\widehat{\mathsf{C}}ob^*(M)$. The second identity has a similar interpretation. We consider the next diagram (we assume that all cobordisms here have only three ends) which is again obtained from (5) by adjusting the notation so that it fits with the statement of Axiom 5.

$$\begin{tikzcd} L \arrow[r, "V''"] \arrow[d, "V"'] & K' \arrow[r, "RV''"] \arrow[d, "\mathrm{id}"'] & M \arrow[r, "R^{-1}V''"] \arrow[d, dotted, "\mathcal{C}"'] & L \arrow[d, "V"'] \\ L' \arrow[r, "S'"] & K' \arrow[r, "RS'"] & N \arrow[r, "R^{-1}S'"] & L' \end{tikzcd} \tag{7}$$

The map $\mathcal{C}$ is the cabling coming from $V'' \sim S' \circ V$. With this notation, the second statement from Axiom 5 claims the commutativity of the square on the right in (7).

(d) *The octahedral axiom.* We consider the following diagram:

$$\begin{tikzcd} L \arrow[r, "S"] \arrow[d, "\mathrm{id}"'] & M \arrow[r, "T"] \arrow[d, "V'"'] & K \arrow[r, "U"] \arrow[d, "V''"'] & L \arrow[r, "S"] \arrow[d, "\mathrm{id}"] & M \\ L \arrow[r, "S'"] & M' \arrow[r, "T'"] \arrow[d, "W'"'] & K' \arrow[r, "U'"] \arrow[d, dotted, "W''"'] & L & \\ & M'' \arrow[r, "\mathrm{id}"] \arrow[d, "R^2V'"'] & M'' \arrow[d, dotted, "R^2V''"'] & & \\ & M \arrow[r, "T"] & K & & \end{tikzcd} \tag{8}$$

Notice that, compared to (5), the place of the morphism V is taken by $\mathrm{id} : L \to L$. The assumption is that the two horizontal rows are exact triangles and that the left vertical column is also an exact triangle. Moreover, as before, $W' = RV'$, $T = RS$, $T' = RS'$ (and we assume that S, S', V' have only three ends). We first need to show the existence of an exact triangle like the column on the right and that it makes the bottom square on the right commutative. We will take V'' to be defined as in point (c) above. This implies the commutativity of V'' with the connectants $U = R^2S$ and $U' = R^2S'$. Compared to Figure 24, the morphism V'' is simpler in our case because instead of V we can insert the identity cobordism of L. We thus get the morphism in Figure 28 below.

This morphism has three ends and thus directly gives rise to an exact triangle. We put $W'' = RV'' : K' \to M''$. We now check that $W'' \circ T' \sim W'$.

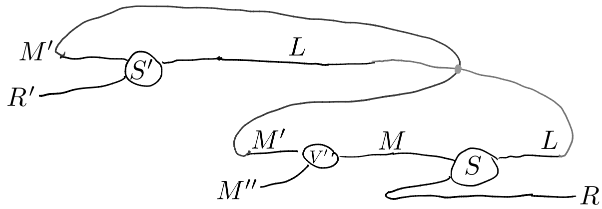

Figure 28. (Color online) The morphism $V'' : K \to K'$ when $V = \mathrm{id}|_L$.

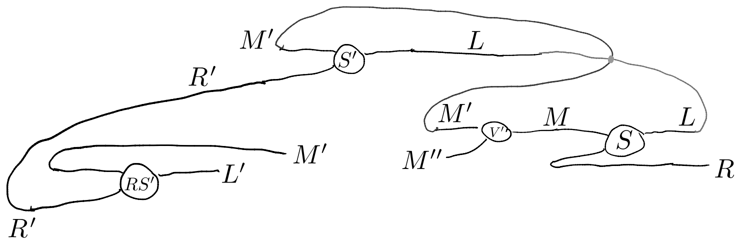

Figure 29. (Color online) The composition $W'' \circ RS'$ with $W'' = RV'' : K' \to M''$ and $T' = RS'$.

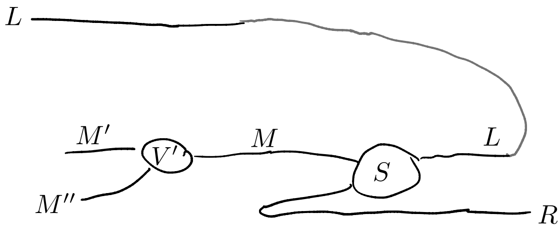

Figure 30. (Color online) Erasing the blue arc as well as RS' and S' from Figure 29 leads, by Axiom 5, to the equivalent morphism above which, in turn, is equivalent to $RV' = W'$.

This is done again in Figures 29 and 30: the first represents the composition $RV'' \circ RS'$ (recall that $T' = RS'$) and the second reflects again an application of Axiom 5 to conclude that $W' \sim W'' \circ RS'$.

To conclude establishing the octahedral axiom, we need to also remark that $R^2V' \circ W'' \sim S \circ U'$ or, in other words, $R^2V' \circ RV'' \sim S \circ R^2S'$. This follows from a similar argument, by making use of the fact that the cobordism V' has three ends and applying the second identity in Axiom 5. $\square$

REMARK 2.3.4

Fix a cobordism $V : L \rightsquigarrow (L_1, \ldots, L_m)$ in $\mathcal{L}ag^*(\mathbb{C} \times M)$. By successively braiding the ends $L_1, \ldots, L_m$, as in Remark 2.3.2 (but including in the braiding all the negative ends), we replace V by a new cobordism $V' : L \rightsquigarrow L'$, where $L' = L_m \#_{c_m} (L_{m-1} \# \ldots (L_2 \#_{c_1} L_1) \ldots))$. In view of Axiom 3, V' is an isomorphism and L' is given in $\widehat{\mathsf{C}}ob^*(M)$ as an iterated cone $L' = \mathcal{C}one(L_m \to \mathcal{C}one(L_{m-1} \to \cdots \mathcal{C}one(L_2 \to L_1)\ldots))$.

2.4. Shadows and fragmentation metrics

Given a Lagrangian cobordism $V : L \rightsquigarrow (L_1, \ldots, L_m)$, we define its shadow (see [23]), $\mathcal{S}(V)$, by

$$\mathcal{S}(V) = \mathrm{Area}(\mathbb{R}^2 \setminus \mathcal{U}), \tag{9}$$

where $\mathcal{U} \subset \mathbb{R}^2 \setminus \pi(V)$ is the union of all the *unbounded* connected components of $\mathbb{R}^2 \setminus \pi(V)$. Here $\pi : \mathbb{R}^2 \times M \longrightarrow \mathbb{R}^2$ is the projection.

The shadow of a surgery morphism is bounded by the "size" of the handle used in the surgery. Therefore, assuming that $\mathsf{C}ob^*(M)$ has surgery models, a natural additional requirement to impose is to ask that $\mathsf{C}ob^*(M)$ have *small* surgery models—in the sense that each morphism is equivalent to a surgery cobordism of arbitrarily small shadow. Similarly, in the definition of cabling equivalence, Definition 2.1.9, the requirement for $V \sim V'$ becomes that for each ϵ sufficiently small there exists a cabling $\mathcal{C}(V, V'; c, \epsilon) \in \mathcal{L}ag^*(M)$. In the marked context, this includes the case $\epsilon = 0$ and, in fact, $V \sim V'$ in this setting if and only if there exists a cabling $\mathcal{C}(V, V'; c) \in \mathcal{L}ag^*(M)$.

We will say that the category $\mathsf{C}ob^*(M)$ is *rigid with surgery models* if it has small surgery models in the sense above and if additionally it satisfies the next axiom.

AXIOM 6 (Rigidity)

For any two Lagrangians L and L', there exists a constant $\delta(L, L') \geq 0$ that vanishes only if $j_L(L) = j_{L'}(L')$ such that for any simple cobordism $V : L \rightsquigarrow L'$, $V \in \mathrm{Mor}_{\mathsf{C}ob^*(M)}$ we have

$$\mathcal{S}(V) \geq \delta(L, L'). \tag{10}$$

REMARK 2.4.1

An important point concerning this axiom is that, while the objects L of $\mathsf{C}ob^*(M)$ are actually Lagrangian immersions $j_L : L \to M$ together with a variety of additional structures (such as markings, choices of almost-complex structures, perturbations, and so forth), the constant $\delta(L, L')$ vanishes if and only if the immersions corresponding to L and to L' have the same image in the ambient manifold M.

We now discuss the meaning of this axiom. We start by noting that the notion of shadow easily leads to a family of pseudometrics, called *fragmentation shadow pseudometrics*, as introduced in [12]. They are defined as follows. Fix a family of objects $\mathcal{F}$ in $\mathcal{L}ag^*(M)$ and, for two objects $L, L' \in \mathcal{L}ag^*(M)$, put

$$\begin{aligned} d^{\mathcal{F}}(L, L') = \inf\{&\mathcal{S}(V) : V : L \rightsquigarrow (F_1, \ldots, F_i, L', F_{i+1}, \ldots, F_m), \\ &F_i \in \mathcal{F}, V \in \mathcal{L}ag^*(\mathbb{C} \times M)\}. \end{aligned} \tag{11}$$

This is clearly symmetric and satisfies the triangle inequality; thus it defines a (possibly infinite) pseudometric.

If $\mathcal{F}'$ is another family of objects in $\mathcal{L}ag^*(M)$, then we can consider the average of the two pseudometrics which is again a pseudometric:

$$d^{\mathcal{F},\mathcal{F}'}(L,L') = \frac{d^{\mathcal{F}}(L,L') + d^{\mathcal{F}'}(L,L')}{2}.$$

There is an obvious order among fragmentation pseudometrics induced by the order on pairs of families $\mathcal{F}$, $\mathcal{F}'$ given by inclusion. In particular, $d^{\emptyset,\emptyset} \geq d^{\mathcal{F},\mathcal{F}'}$.

There is a specific sense in which we are interested in the nondegeneracy of such pseudometrics: we will say that a pseudometric is *geometrically nondegenerate* if $d^{\mathcal{F},\mathcal{F}'}(L,L') = 0$ if and only if $j_L(L) = j_{L'}(L')$. A priori, none of the pseudometrics above are nondegenerate, even in this geometric sense. It is natural to define the rigidity of a cobordism category $\mathsf{C}ob^*(M)$ by imposing this requirement, as below. We will only be concerned with some special pairs $\mathcal{F}$, $\mathcal{F}'$ (see Corollary 4.4.2).

DEFINITION 2.4.2

A cobordism category $\mathsf{C}ob^*(M)$ is called *strongly rigid* if, for any two families $\mathcal{F}$ and $\mathcal{F}'$ such that the intersection

$$\overline{\bigcup_{F\in\mathcal{F}} F} \cap \overline{\bigcup_{F'\in\mathcal{F}'} F'}$$

is totally discrete, the pseudometric $d^{\mathcal{F},\mathcal{F}'}$ is nondegenerate on the quotient set obtained from $\mathcal{O}b(\mathsf{C}ob^*(M))$ by identifying any two objects with the same image in M.

This definition of strong rigidity does not require $\mathsf{C}ob^*(M)$ to have surgery models, it is a purely geometric constraint. Similarly, the condition in Axiom 6 is also purely geometric and can be formulated independently of the existence of surgery models. With our definitions, the inequality (10) is equivalent to

$$d^{\emptyset,\emptyset}(L,L') \geq \delta(L;L'), \tag{12}$$

and thus (10) is equivalent to the fact that $d^{\emptyset,\emptyset}$ is geometrically nondegenerate. In view of displacement energy – width inequalities, relations such as (10) are naturally to be expected in nonflexible settings.

We will see that unobstructed classes of Lagrangians, in the sense to be made precise later, have small surgery models and are strongly rigid. In fact, as will be noticed in Corollary 4.4.2, in this case the geometric nondegeneracy of $d^{\emptyset,\emptyset}$ implies the geometric nondegeneracy of all the other pseudometrics from Definition 2.4.2 (for this result, our objects need to be immersed Lagrangians).

REMARK 2.4.3

(a) To further unwrap the meaning of Axiom 6, notice that, because Lagrangian suspensions of objects in $\mathcal{L}ag^*(M)$ are in $\mathcal{L}ag^*(\mathbb{C}\times M)$, the pseudometric $d^{\emptyset,\emptyset}$ is bounded from above by the Hofer distance on Lagrangian submanifolds. Thus, Axiom 6 implies the nondegeneracy of this distance too. In a different direction, the (pseudo)metrics $d^{\mathcal{F},\mathcal{F}'}$ are finite much more often compared to $d^{\emptyset,\emptyset}$, and thus much larger classes of Lagrangians are endowed with a meaningful geometric structure.

(b) In case the category $\mathsf{C}ob^*(M)$ has surgery models (and thus $\widehat{\mathsf{C}}ob^*(M)$ is triangulated), there is a more general context that fits the construction of this shadow pseudometrics that we briefly recall from [12]. Let $\mathcal{X}$ be a triangulated category, and recall that there is a category denoted by $T^S\mathcal{X}$ that was introduced in [9] and [10]. This category is monoidal and its objects are finite ordered families $(K_1,\dots,K_r)$ with $K_i\in\mathcal{O}b(\mathcal{X})$ with the operation given by concatenation. In essence, the morphisms in $T^S\mathcal{X}$ parameterize all the cone-decompositions of the objects in $\mathcal{X}$. Composition in $T^S\mathcal{X}$ comes down to refinement of cone-decompositions. Assume given a weight $w:\mathrm{Mor}_{T^S\mathcal{X}}\to[0,\infty]$ with the property

$$w(\bar{\phi}\circ\bar{\psi})\le w(\bar{\phi})+w(\bar{\psi}),\qquad w(\mathrm{id}_X)=0,\quad \forall X. \tag{13}$$

Fix also a family $\mathcal{F}\subset\mathcal{X}$. We can then define a measurement on the objects of $\mathcal{X}$:

$$s^{\mathcal{F}}(K',K)=\inf\bigl\{w(\bar{\phi})\bigm|\bar{\phi}:K'\to(F_1,\dots,K,\dots,F_r),F_i\in\mathcal{F},\forall i\bigr\}. \tag{14}$$

This satisfies the triangle inequality but is generally nonsymmetric. However, as noted in [12], $s^{\mathcal{F}}$ can be symmetrized and it leads to what is called in [12] a *fragmentation pseudometric*. Our remark here is that if $\mathcal{X}=\widehat{\mathsf{C}}ob^*(M)$, then the category $T^S\mathcal{X}$ can be identified with (a quotient of) the general cobordism category $\mathcal{C}ob^*(M)$ as introduced in [9]. This category has as objects families of Lagrangians $(L_1,\dots,L_k)$ and as morphisms families of cobordisms $V:L\rightsquigarrow(L_1,\dots,L_k)$. The difference compared to the category $\mathsf{C}ob^*(M)$ is that such a V is viewed in $\mathcal{C}ob^*(M)$ as a morphism from the family formed by a single element L to the family $(L_1,\dots,L_k)$, and in $\mathsf{C}ob^*(M)$ the same V is viewed as a morphism from L to L_k. The relation between $T^S\widehat{\mathsf{C}}ob^*(M)$ and $\mathcal{C}ob^*(M)$ is not at all surprising because, as noted in Remark 2.3.4, each cobordism $V:L\rightsquigarrow(L_1,\dots,L_k)$ induces an iterated cone-decomposition in $\widehat{\mathsf{C}}ob^*(M)$. Finally, the fragmentation pseudometric $s^{\mathcal{F}}$, defined on the objects of $\widehat{\mathsf{C}}ob^*(M)$ by using as weight the shadow of cobordisms, coincides with $d^{\mathcal{F}}$.

2.5. Cobordism groups

Given a cobordism category $\mathsf{C}ob^*(M)$, there is a natural associated cobordism group $\Omega^*(M)$ given as the free abelian group generated by the Lagrangians in $\mathcal{L}ag^*(M)$ modulo the subgroup of relations generated by the expressions $L_1+\dots+L_m=0$ whenever there exists a cobordism $V:\emptyset\rightsquigarrow(L_1,\dots,L_m)$ with $V\in\mathcal{L}ag^*(\mathbb{C}\times M)$.

In our setting, where orientations are neglected, this group is a $\mathbb{Z}_2$-vector space.

COROLLARY 2.5.1

If the category $\mathsf{C}ob^*(M)$ *has surgery models, then there is an isomorphism*

$$\Omega^*(M) \cong K_0\big(\widehat{\mathsf{C}}ob^*(M)\big).$$

Proof

Recall that $K_0\mathcal{C}$ is the Grothendieck group of the triangulated category $\mathcal{C}$. It is generated by the objects of $\mathcal{C}$ subject to the relations $B = A + C$ for each exact triangle $A \to B \to C$. It immediately follows from Remark 2.3.4 that the relations defining the cobordism group $\Omega^*(M)$ belong to the relations subgroup giving K_0. Conversely, as all exact triangles in $\widehat{\mathsf{C}}ob^*(M)$ are represented by surgeries it follows that all relations defining K_0 are also cobordism relations. Thus, we have the isomorphism claimed. □

3. Unobstructed classes of Lagrangians and the category $\mathsf{C}ob^*(M)$

The purpose of this section is to discuss in detail the classes of Lagrangians $\mathcal{L}ag^*(M)$ and cobordisms $\mathcal{L}ag^*(\mathbb{C} \times M)$ that are used to define the category $\mathsf{C}ob^*(M)$. This category will be shown later in the article to have rigid surgery models, leading to a proof of Theorem A.

To fix ideas, the Lagrangians L belonging to the class $\mathcal{L}ag^*(M)$ are exact, generic immersions $i_L : L \to M$, endowed with some additional structures. In short, these additional decorations consist of a primitive, a marking c (in the sense of Definition 2.1.5—thus, a choice of double points of i_L), and some perturbation data associated to L, $\mathcal{D}_L$, providing good control on the pseudoholomorphic (marked) curves with boundary on L and such that, with respect to these data, L is unobstructed. The cobordisms $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$ satisfy similar properties: they are exact, marked, immersions $i_V : V \to \mathbb{C} \times M$, and they carry perturbation data of a similar nature as $\mathcal{D}_L$ in a way that extends the data of the ends. An additional complication is that the cobordisms come with an additional decoration consisting of a choice of perturbation required to transform such cobordisms with immersed ends—therefore having nonisolated double points along the ends—into generic immersions.

This section is structured as follows. In Section 3.1, we first discuss the types of curves with boundary that appear in our constructions as well as the notion of unobstructed, marked Lagrangians in M. In Section 3.2, we discuss similar notions for cobordisms before specializing the discussion to the exact case in Section 3.3. Finally, in Section 3.4, after some other adjustments, we define the category $\mathsf{C}ob^*(M)$.

3.1. J-Holomorphic curves and unobstructed Lagrangians

Immersed Lagrangians have been considered from the point of view of Floer theory starting with the work of Akaho [1] and Akaho–Joyce [2] as well as other authors such as Alston–Bao [3], Fukaya [27], and Palmer–Woodward [40] to name a few. A variety of unobstructedness-type conditions appear in all these works. We discuss here only the aspects that are relevant for our approach.

While in the embedded case unobstructedness can be deduced from certain topological constraints (such as exactness or monotonicity) and thus it is independent of generic choices of almost-complex structures and other such data, for immersed marked Lagrangians this condition is much more delicate as it requires certain counts of J-holomorphic curves to vanish (mod2). For these counts to be well defined, the relevant moduli spaces need to be regular and, further, the counts themselves depend on the choices of data. This makes the definition of an unobstructed, immersed, marked Lagrangian considerably more complicated as one can see in Definition 3.1.4 below. For immersed marked cobordisms, the relevant definitions are even more complex (as seen in Definition 3.2.4) because one needs to deal with the additional problem that double points of cobordisms with immersed ends are not isolated.

3.1.1. J-Holomorphic curves with boundary along marked Lagrangians

We consider here marked, immersed Lagrangians (L, c) with $j_L : L \to M$ an immersion with only transverse double points, as in Definition 2.1.5. Recall that the set $c \subset L \times L$ is a collection of ordered double points of j_L; in particular, for each $(P_-, P_+) \in c$ we have $j_L(P_-) = j_L(P_+)$.

Consider an n-tuple of such Lagrangians (L_i, c_i), $1 \le i \le n$, and assume for the moment that, for all $i \ne j$, the intersections of L_i and L_j are transverse and distinct from all self-intersection points of the L_i's. We denote by $\mathbf{c}$ the family $\mathbf{c} = \{c_i\}_i$.

A J-holomorphic polygon with boundary on the (L_i, c_i)'s is a map $u : D^2 \to M$ with $u(\partial D^2) \subset \bigcup_i j_{L_i}(L_i)$ such that (see Figure 31):

(i) The curve u is continuous and smooth on $D^2 \setminus \bigcup_i \{a_1^i, \ldots a_{s_i}^i\}$ and satisfies $\bar{\partial}_J u = 0$ inside D^2. Here $a_k^i \in S^1$, $1 \le k \le s_i$, and this family $\mathbf{a} = \{a_k^i\}$ is ordered as

$$(a_1^1, a_2^1, \ldots, a_{s_1}^1, a_1^2, a_2^2, \ldots, a_{s_2}^2, \ldots, a_{s_n}^n).$$

With this order, the points in $\mathbf{a}$ are placed around the circle S^1 in clockwise order. We denote by C_k^i the (closed) arc of S^1 that starts at a_k^i and stops at the next point in $\mathbf{a}$ (in cyclic order).

(ii) For each i, k, the restriction $u|_{C_k^i}$ has a continuous lift $\hat{u}_k^i : C_k^i \to L_i$.

(iii) The curve u has asymptotic corners (in the usual sense of [48]) at each of the a_k^i's such that each of the a_k^i's is mapped to a self-intersection point of L_i for $1 < k \le s_i$, and to an intersection point of $L_i \cap L_{i+1}$ for a_1^{i+1}, $i \ge 1$ and, for a_1^1, to an intersection point in $L_n \cap L_1$. If the boundary conditions consist of a single

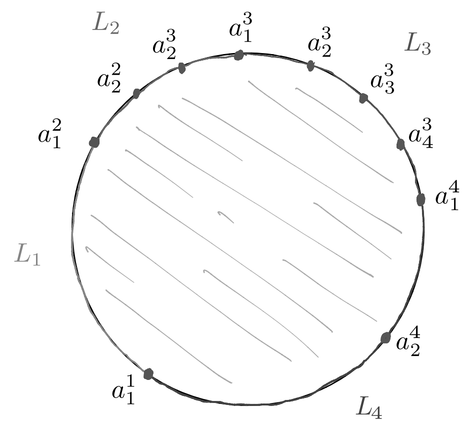

Figure 31. (Color online) An example of a curve u with boundary conditions along L_1, L_2, L_3, L_4.

Lagrangian L_1, then we assume that a_1^1 is also mapped to a self-intersection point of L_1.

With the usual conventions for the orientation of $D^2 \subset \mathbb{C}$, the punctures a_k^i correspond to "entry points" in the disk. To be more precise, viewing a puncture a_k^i as an entry or an exit point is equivalent to making a choice of a class of strip-like end coordinates around the puncture: 0-side of the strip before the 1-side in clockwise order corresponds to an entry and the other way around for an exit. The convention is that each time we use strip-like coordinates around a puncture, they are assumed to be of the corresponding type as soon as we fix the type of puncture to be either an entrance or an exit. To relate these configurations to the operations typical in Floer theory, it is often useful to consider all punctures as entries except possibly for one for which the strip-like coordinates are reversed so that it becomes an exit point. We will fix the convention that the exit point, if it exists, is associated to a_1^1.

We denote the moduli spaces of curves as above by $\mathcal{M}_{J;L_1,\dots,L_n}(x_1,\dots,x_m; y)$, where each x_i is either one of the self-intersection points of an L_j viewed as a couple in $((x_i)_-,(x_i)_+) \in L_j \times L_j$, or an intersection point of $L_j \cap L_{j+1}$ that can also be viewed as a pair of points $((x_i)_-,(x_i)_+) \in (L_j, L_{j+1})$, in such a way that the punctures a_k^i are sent in order to the x_j's, starting with the exit a_1^1 that is sent to y. Moreover, we fix conventions such that the path $\hat{u}_k^i(C_k^i)$, followed clockwise around the circle, starts at a point $(x_j)_+$ (or at y_- in the case of C_1^1) and ends at $(x_{j+1})_-$ (or at y_+ for $C_{s_n}^n$).

It is also useful to consider the case when there is no exit point; thus, even a_1^1 is an entry. The corresponding moduli space is denoted by $\mathcal{M}_{J;L_1,\dots,L_n}(x_1,\dots,x_m; \emptyset)$. In the case of a single boundary condition L_1, we also allow for the case of no punctures ($\mathbf{a} = \emptyset$) and notice that, in this case, u is a J-holomorphic disk with boundary on L_1 with the moduli space denoted $\mathcal{M}_{J,L_1}$. We omit some of the subscripts, such as the boundary conditions $L_1,\dots,L_n$, if they are clear from the context.

Moduli spaces of this type appear often in Floer-type machinery, for instance, in [28] and [48] and, in the case of immersed Lagrangians, in [3], [4] as well as [2].

A special class of moduli spaces as above plays a particular role for us. We will say that u is a **c**-marked J-holomorphic curve if it is a curve as above with the following additional constraint:

(iv) If $n \geq 2$, for each $i \in \{1, \ldots, n\}$ and $1 \leq k < s_i$ we have $(\hat{u}^i_k(a^i_{k+1}), \hat{u}^i_{k+1}(a^i_{k+1})) \in c_i$. In the case $n = 1$ and if u has no exit, then we also assume that $(\hat{u}^1_{s_1}(a^1_1), \hat{u}^1_1(a^1_1)) \in c_1$.

In other words, a **c**-marked curve u is a J-holomorphic polygon with asymptotic corners at the punctures given by the points a^i_k, such that it switches branches along the Lagrangian L_i *only* at points belonging to c_i and it switches from L_i to L_{i+1} (or, respectively, from L_n to L_1) at the puncture point a^{i+1}_1 (resp., a^1_1). Notice that if there is a single boundary condition L_1, and a^1_1 is an exit, then the switch at the point a^1_1 is not required to belong to c_1. On the other hand, if there is a single boundary condition L_1, and all punctures are entries, then we require for a^1_1 to also correspond to a point in c_1. We denote the moduli spaces of **c**-marked polygons by $\mathcal{M}_{J,\mathbf{c};L_1,\ldots,L_n}(x_1, \ldots, x_m; y)$. With respect to this notation, the condition relevant to **c** is, in summary, that all the switching of branches at entry points corresponds to x_k's that belong to some c_i.

In the literature, most Floer-type algebraic constructions associated to immersed Lagrangians are defined through counts of curves that correspond, in our language, to the case $\mathbf{c} = \emptyset$. In other words, these are curves that do not switch branches along the immersed Lagrangians. In our case, the invariants that we discuss in a later section here are associated to counts of **c**-marked curves, thus *switching of branches along L_i is allowed as long as it takes place at self-intersection points belonging to c_i.*

A few special cases are worth mentioning: Floer strips, in the usual sense, correspond to the case when there are just two boundary conditions L_1, L_2 and only two punctures a^1_1—the exit—and a^2_1—the entry (i.e., $s_1 = 1$, $s_2 = 1$); curves with boundary conditions again along L_1 and L_2 but with s_1, s_2 not necessarily equal to 1, but still with an exit at a^1_1, will be referred to as **c**-marked Floer strips; another important special case appears when there is a single boundary condition and one puncture a^1_1 which is viewed as an exit—these curves are called *teardrops*. Such a curve with boundary condition along L_2 appears in Figure 32. A more general case of a **c**-marked curve is also of interest to us: in case there is a single boundary condition L_1 and $s_1 > 1$, with a^1_1 an exit and a^i_k entries for all $k > 1$ and, as discussed above, all entries are associated to elements in c_1. These type of curves will be referred to as **c**-*marked teardrops.* (Of course, there are also *nonmarked* teardrops. These are (holomorphically) parameterized by a disk with only one boundary puncture. Note that, according to our formalism, a nonmarked teardrop is not regarded as a **c**-marked one, unless $\mathbf{c} = \emptyset$.) Finally, one last case: again there is just one boundary condition L_1, but in this case all

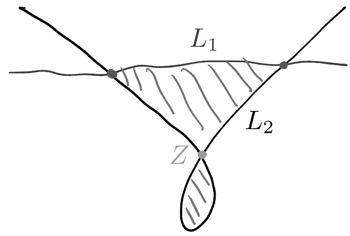

Figure 32. (Color online) A teardrop at Z (along L_2) and a triangle—a curve with two boundary conditions L_1, L_2 and $s_1 = 1$, $s_2 = 2$.

a_k^1's are viewed as entries. If all these entries (including a_1^1) are included in c_1, then we call the curve a $\mathbf{c}$-marked J-disk with boundary on L_1.

REMARK 3.1.1

(a) With our conventions some geometric curves can appear in more than a single category. For instance, a teardrop with an exit (P_-, P_+) such that (P_+, P_-) belongs to the marking c_1 can also be viewed as a $\mathbf{c}$-marked disk. However, once the nature of the punctures is fixed there is no such ambiguity.

(b) Recall from Remark 2.1.2 that the objects in $\mathcal{L}ag^*(M)$ are unions of immersions with isolated, transverse double points. The definition of the moduli spaces above extends naturally to this situation because of condition (ii) at the beginning of this section, and because the markings are self-intersection points of immersions with isolated, transverse double points.

For a curve u as above, we denote by $|u|$ the number of punctures of u, $|u| = s_1 + s_2 + \cdots + s_n$.

While in the discussion above we have only considered J-holomorphic curves, in practice we also need to include in our consideration perturbed such curves. We will discuss the system of perturbations in greater detail below but, to fix ideas, we mention that for all moduli spaces of curves u with $|u| \geq 2$, the choice of these perturbations follows Seidel's scheme in his construction of the Fukaya category in [48] (we assume here familiarity with his approach).

Fix one marked, immersed Lagrangian (L, c). Put $\mathbf{c} = \{c\}$. (We will sometimes use this convention when there is only one Lagrangian under consideration.)

DEFINITION 3.1.2

A regular, coherent system of perturbation data $\mathcal{D}_L$ for (L, c) consists of the following:

(i) A time-dependent almost-complex structure $J = J_t$, $t \in [0, 1]$, with $J_0 = J_1$, such that:

(a) all moduli spaces of J_0-holomorphic $\mathbf{c}$-*marked* disks with boundary on L are void,

(b) the moduli spaces of nonmarked J_0-holomorphic teardrops are all void.

(ii) A system of perturbations parameterized by the universal associahedron, as described in [48], such that all the resulting moduli spaces of polygons u with boundary on L, with $|u| \geq 2$ and with an exit are regular and satisfy the following additional properties:

(a) On the boundary of each polygon u (with $|u| \geq 1$) the almost complex structure coincides with J_0 and the Hamiltonian perturbation is trivial.

(b) For each polygon u with $|u| = 2$, with one input and one exit, the respective perturbation has a trivial Hamiltonian term and its almost complex structure part coincides with J_t in strip-like coordinates $(s,t) \in \mathbb{R} \times [0,1]$.

(c) For each polygon u with $|u| \geq 3$, there are strip-like end coordinates around the punctures and in these coordinates the perturbation has vanishing Hamiltonian term and its almost-complex part coincides with J_t.

(d) The choices of perturbations are coherent with respect to gluing and splitting for all curves u with $|u| \geq 2$.

We refer to the almost-complex structure J_0 as the *base* almost complex structure of the coherent system of perturbations $\mathcal{D}_L$. In some of the considerations we will discuss here, we will also need to use curves with some interior marked points (see also Remark 3.1.3(b) below), but for the moment we will limit our discussion to the types of curves introduced above.

REMARK 3.1.3

(a) Condition (i)(a) in the definition above requires not only that there be no J_0-holomorphic disks with boundary on L, but also that there be no $\mathbf{c}$-marked polygons with boundary on L and without an exit.

(b) It is easy to formally relax condition (i)(b) to only require that moduli spaces of (nonmarked), possibly perturbed, teardrops be regular and most further definitions and constructions work under this weaker assumption. However, this regularity is hard to achieve in practice in full generality. One of the reasons is that teardrops carry a single puncture and thus their domains are unstable. As a result, Seidel's recursive scheme in choosing perturbations for polygons with more and more corners does not automatically work. Alternative ways to address this in some settings appear in [3]; another possibility is to implement the Kuranishi method as in [2]. Yet another different approach, initiated by Lazzarini [35] for disks and pursued for polygons by Perrier [42], is to attempt to use more special but still generic classes of autonomous almost-complex structures in the whole construction.

Throughout the present article, we will mostly assume as in Definition 3.1.2(i)(b) that *all moduli spaces of (nonmarked) teardrops are void.* Later on, when dealing with the moduli spaces of curves with boundaries along cobordisms, there are special situations when teardrops can be stabilized by using an interior marked point and in that case Seidel's method for picking coherent perturbations continues to apply.

As we will see later on, there are situations in the exact setting in which this limited approach is still sufficient for our purposes.

(c) Notice that we require regularity of moduli spaces of all polygons with boundary on L and not only of the **c**-marked polygons. The reason is that, even if the various μ_k operations which we will define below use only **c**-marked polygons, the proofs of relations of type $\mu \circ \mu = 0$ also make use of non-**c**-marked polygons as these appear through an application of Gromov compactness. However, the regularity of all polygons can be relaxed in Definition 3.1.2(ii) to the regularity of all the moduli spaces of **c**-marked polygons as well as that of the moduli spaces appearing in compactifications of these.

3.1.2. Moduli spaces of marked polygons

Given systems of coherent perturbations $\mathcal{D}_L$ for each (L,c) in some class $Lag^*(M)$, these can be extended in the sense of [48] to a coherent system of perturbations for curves with boundary conditions along all the $(L,c,\mathcal{D}_L)$'s. In full generality, this requires picking for each pair (L,c), (L',c') Floer data $H_{L,L'}$, $J_{L,L'}$ where, in general, $H_{L,L'}$ is not vanishing so that the construction can handle Lagrangians with nontransverse intersections. Moreover, $J_{L,L'}$ is nonautonomous and picked in such a way as to agree with $(J_0)_L$ and, respectively, with $(J_0)_{L'}$ for t close to 0 (resp., close to 1).

The resulting moduli spaces involving these systems of coherent perturbations will be denoted by $\mathcal{M}_{\mathcal{D};L_1,\ldots,L_n}(x_1,\ldots,x_m;y)$ for boundary conditions along $L_1,\ldots,L_n$. Notice that, in the nontransverse intersection case the punctures of a curve u in these moduli spaces correspond to two types of ends: self intersection points $x_i \in I_{L_i}$ and time-1 Hamiltonian chords associated to the Hamiltonians $H_{L_i,L_{i+1}}$. We denote these chords by $\mathcal{P}(L_i,L_{i+1})$. To simplify the discussion, we focus below on the transverse case with the understanding that the adjustments of the relevant notions to the nontransverse situation are immediate. Similarly, in the marked case, we have the moduli spaces $\mathcal{M}_{\mathcal{D},\mathbf{c};L_1,\ldots,L_n}(x_1,\ldots,x_m;y)$.

It is useful to regroup various components of these moduli spaces in the following way. We fix intersection points (or, more generally, Hamiltonian chords) $h_i \in \mathcal{P}(L_i,L_{i+1})$ and $y \in \mathcal{P}(L_1,L_n)$. We denote by $\overline{\mathcal{M}}_{\mathcal{D};L_1,\ldots,L_n}(h_1,\ldots,h_{n-1};y)$ the union of the curves $u \in \mathcal{M}_{\mathcal{D};L_1,\ldots,L_n}(x_1,\ldots,x_m;y)$, where, in order, $x_1,\ldots,x_{s_1}$ are self-intersection points of L_1, $x_{s_1+1} = h_1$, the next s_2-points x_i are self-intersection points of L_2, the next point equals h_2 and so forth around the circle (see again Figure 31). In essence, the spaces $\overline{\mathcal{M}}$ group together curves with fixed "corners" at the chords $h_i \in \mathcal{P}(L_i,L_{i+1})$, but with variable numbers of additional corners at self-intersection points along each one of the L_i's. We will also use the notation $\overline{\mathcal{M}}_{\mathcal{D},\mathbf{c};L_1,\ldots,L_n}(h_1,\ldots,h_{n-1};y)$ for the corresponding **c**-marked moduli spaces which are defined, as usual, by requiring that all self-intersection corners belong to $c_i \subset I_{L_i}$. For instance, with this notation, the moduli space of all **c**-marked teardrops at a self-intersection point $y \in I_{L_1}$ is written as $\overline{\mathcal{M}}_{\mathcal{D}_{L_1},c_1;L_1}(\emptyset;y)$.

3.1.3. Unobstructed marked Lagrangians in M

The purpose of this subsection is to define the notion of an unobstructed marked Lagrangian that appears in the proof of Theorem 4.1.1.

DEFINITION 3.1.4

A triple $(L, c, \mathcal{D}_L)$ with (L, c) a marked Lagrangian and $\mathcal{D}_L$ a coherent system of regular perturbations as in Definition 3.1.2 is called *unobstructed* if, for each self-intersection point $y \in I_L$, the 0-dimensional part of the moduli space of c-marked teardrops $\overline{\mathcal{M}}_{\mathcal{D}_L,c;L}(\emptyset; y)$ is finite and the mod2 number of its elements vanishes.

REMARK 3.1.5

In the terminology of [28] and [2], the choice of a marking c as in the definition above is a particular case of a *bounding chain*.

Over $\mathbb{Z}/2$, the definitions of Floer and Fukaya category type algebraic structures generally require that the relevant moduli spaces satisfy certain compactness conditions ensuring that the relevant 0-dimensional spaces are compact and that the 1-dimensional moduli spaces have appropriate compactifications. This follows from Gromov compactness as long as a priori energy bounds are available. We integrate this type of condition in our definition of an unobstructed class of Lagrangians.

DEFINITION 3.1.6

A class $Lag^*(M)$ of marked, immersed Lagrangians $(L, c, \mathcal{D}_L)$, each of them unobstructed in the sense of Definition 3.1.4, is *unobstructed with respect to a system of coherent perturbations $\mathcal{D}$ that extends the $\mathcal{D}_L$'s* if:

(i) $\mathcal{D}$ is regular,

(ii) for any finite family (L_i, c_i), $1 \leq i \leq m$, there exists a constant $E_{L_1,\ldots,L_n;\mathcal{D}}$ such that we have $E(u) \leq E_{L_1,\ldots,L_n;\mathcal{D}}$ for any $h_i \in \mathcal{P}(L_i, L_{i+1})$, $y \in \mathcal{P}(L_1, L_n)$, and $u \in \overline{\mathcal{M}}_{\mathcal{D},\mathbf{c};L_1,\ldots,L_n}(h_1, \ldots, h_{n-1}; y)$.

Here $E(u)$ is the energy of the curve u, $E(u) = \frac{1}{2}\int_S \|du - Y\|^2$ (where S is the punctured disk, domain of u, and Y is the 1-form with values in Hamiltonian vector fields used in the perturbation term; see [48, Section (8f)]). Regularity of $\mathcal{D}$ means that all the moduli spaces of $\mathbf{c}$-marked (and nonmarked) polygons with boundaries along families picked from $Lag^*(M)$ are regular. We refer to condition (ii) as the *energy bounds* condition relative to $\mathcal{D}$.

REMARK 3.1.7

(a) Through an application of Gromov compactness, the energy bounds condition implies that the number of corners at self-intersection points of all

$\mathbf{c}$-marked polygons with boundaries along the family $(L_1, \ldots, L_n)$ is uniformly bounded.

(b) Standard methods show that, starting with a family $\{(L, c, \mathcal{D}_L)\}$ of unobstructed marked Lagrangians, it is possible to extend the perturbations $\mathcal{D}_L$ to a coherent, regular system of perturbations $\mathcal{D}$, as required in Definition 3.1.6(i). However, in the absence of some a priori method to bound the energy of the curves in the resulting moduli spaces—such as estimates involving primitives of the Lagrangians involved, monotonicity arguments, and so on—defining invariants over $\mathbb{Z}/2$ is not possible even for this regular $\mathcal{D}$. Moreover, the energy bounds condition itself is not stable with respect to small perturbations and constraints such as exactness or monotonicity are lost through surgery. As a consequence, we made here the unusual choice to integrate the energy bounds condition in the definition.

3.2. Unobstructed marked cobordisms

The next step is to explain the sense in which a class of marked, immersed, cobordisms $Lag^*(\mathbb{C} \times M)$ is unobstructed. In essence, the condition is the same as in Definition 3.1.4, but there are additional subtleties in this case that have to do with the fact that a cobordism with immersed ends has double points that are not isolated.

We start by making more precise the type of immersions included in the class $Lag^*(\mathbb{C} \times M)$. These are unions of Lagrangian cobordism immersions $j_V : V \to \mathbb{C} \times M$ such that the immersion j_V has singular points of (potentially) two types: the first type consists of isolated, transverse double points; the second type are self-intersections modeled on clean intersections along a submanifold Σ_V of dimension 1 in $\mathbb{C} \times M$ (possibly with boundary, and not necessarily compact or connected) such that the projection

$$\pi : \mathbb{C} \times M \to \mathbb{C}$$

restricts to an embedding on each connected component of Σ_V. Moreover, each noncompact component of Σ_V corresponds to an end of the cobordism. More precisely, if a cobordism V has an end L (say, positive) that is immersed with a double point (P_-, P_+), then, by definition, a component of Σ_V includes an infinite semiaxis of the form $[a, +\infty) \times \{k\} \times \{P\}$, where $P = j_L(P_-) = j_L(P_+)$. We assume that all noncompact components of Σ_V correspond to a double point from the fiber in this sense. We continue to denote by I_V the double point set of such an immersion:

$$I_V = \big\{(P_-, P_+) \subset V \times V : j_V(P_-) = j_V(P_+)\big\}.$$

We denote the i-dimensional subspace of I_V by I_V^i, $i \in \{0, 1\}$. In particular, Σ_V is the image of I_V^1 through j_V.

As in the case of the class $Lag^*(M)$, it is sufficient to focus here on the cobordisms V as above as the various arguments extend trivially to finite disjoint unions of such objects.

3.2.1. A class of perturbations for cobordisms with singularities along 1*-dimensional clean intersections*

Dealing analytically with clean self-intersections along intervals with boundary is probably possible but it requires some new ingredients (due to the boundary points) that we prefer to avoid (for Floer theory for Lagrangians with clean intersections, see [44], [46]). For this purpose, we will describe below a class of perturbations that transforms such a cobordism V into a Lagrangian V_h immersed in $\mathbb{C} \times M$ which is no longer a cobordism but has isolated, generic double points and whose behavior at ∞ presents "bottlenecks" in the sense of [10]. These perturbations will also be needed to even define the notion of a marked cobordism.

Fix a cobordism V immersed in $\mathbb{C} \times M$, as discussed above, with clean intersections along a 1-dimensional manifold Σ_V. All perturbations to be considered here can be described as follows. Consider $U \subset T^*V$ a small neighborhood of the 0-section. We can take U small enough so that the immersion $i_V : V \to \mathbb{C} \times M$ extends to an immersion (still denoted by i_V) of U which is symplectic and such that i_V is an embedding on each $U_x = U \cap T^*_x V$. We then consider a small Hamiltonian perturbation (that will be made more precise below) of the 0-section, $V_h \subset U$. The perturbed immersion that we are looking for is given by the restriction of i_V to V_h.

We now describe more precisely the perturbation V_h of the 0-section. It will be written as the time-1 image of V through a Hamiltonian isotopy induced by a Hamiltonian h. To make this h explicit, notice that the 1-dimensional manifold Σ_V has some compact components as well as some noncompact ones. We focus now on one of the noncompact components, Σ'. This corresponds to one of the ends L of the cobordism and, more precisely, to a double point $(P_-, P_+) \in I_L$ of the immersion j_L. We consider two small disks D_- and D_+ in L around, respectively, the points P_- and P_+. To fix ideas, we imagine these two disks to be of unitary radius. Inside these disks we consider smaller disks D'_- and D'_+ of half-radius as well as closed shells C_-, C_+ that are, respectively, the closure of the complements of the $D'_\pm$'s inside the disks $D_\pm$.

By assumption, outside of a compact set, Σ' is of the form $[a, +\infty) \times \{k\} \times \{P\}$ and we have the inclusions $[a, +\infty) \times \{k\} \times (D_- \sqcup D_+) \subset [a, +\infty) \times \{k\} \times L \subset V \subset U$. Outside of $[a, +\infty) \times \{k\} \times (D_- \cup D_+)$, we take the perturbation h to be zero and thus V_h is equal to V there. We now describe the construction on the region $[a, +\infty) \times \{k\} \times (D'_- \cup D'_+)$. We consider a Hamiltonian $h_\pm$ defined in a small neighborhood of the 0-section of T^*V given as a composition $h_\pm = \hat{h}_\pm \circ \mathrm{Re}(\pi \circ j_V)$, where $\hat{h}_\pm : [a, \infty) \to \mathbb{R}$ has a nondegenerate critical point at $a+2$, is linear for $t \geq a+4$ (increasing for h_+ and decreasing for h_-), and vanishes for $t \in [a, a+1]$ (see Figure 33).

The perturbation h is required to agree with $h_\pm$ on $[a, +\infty) \times \{k\} \times (D'_- \cup D'_+)$. We now describe h on $[a, +\infty) \times \{k\} \times (C_- \cup C_+)$. We write each $C_\pm$ as a cylinder $S_\pm \times [0, 1]$ with the 0 end corresponding to the inner boundary of $C_\pm$ and the 1 end corresponding to the outer boundary. Denote by s the variable associated to the height in these cylinders. We define $h : [a, +\infty) \times \{k\} \times C_\pm \to \mathbb{R}$

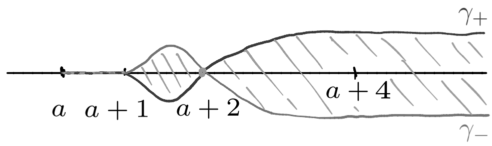

Figure 33. (Color online) The projection onto $\mathbb{C}$ of the L-end of V_h, after replacing the semiaxis of double points $\Sigma' = [a, \infty) \times \{k\} \times \{P\}$ with a compact clean intersection along $[a, a+1]$ and a double point over $\{a+2\} \times \{k\}$.

to be an interpolation $h^s_\pm$, $s \in [0,1]$ between $h_\pm$ and 0, equal to $h_\pm$ for s close to 0 and equal to 0 for s close to 1. More explicitly, $h^s_\pm = \hat{h}^s_\pm \circ \mathrm{Re}(\pi \circ j_V)$ with $\hat{h}^s_\pm = (1-s)\hat{h}_\pm$. The image of V through the time-1 Hamiltonian diffeomorphism associated to h, V_h, has a projection onto $\mathbb{C}$ (along the end discussed here) as in Figure 33. Assuming that $h_\pm$ is sufficiently small, we see that, inside the set of double points of $j_V|_{V_h}$, the set Σ' has been replaced by a union of a single, generic, double point that lies over $\{a+2\} \times \{k\}$ (and, in the fiber, it is close to $P \in M$) and a clean intersection, compact component along $\Sigma'_h = [a, a+1] \times \{k\} \times \{P\}$. We will say that the double point over $\{a+2\} \times \{k\}$ is a bottleneck at $a+2$. This construction can be repeated for all noncompact ends. There is an obvious adjustment for the negative ends where the bottlenecks will be over $-a-2$. Notice that in this case too, the crossing pattern of the two curves γ_- and γ_+ over the bottleneck is assumed to be the same as in Figure 33. Once this construction is completed, we are left with double points that are either generic and isolated or along clean intersections along compact 1-manifolds. The construction of V_h ends by a generic perturbation supported in a neighborhood of the preimage through j_V of these compact components such that these components are also replaced by a union of isolated double points.

REMARK 3.2.1

For each double point P of an end L of the cobordism V, the perturbation above requires a choice of a lift of P to I_L. We will refer to a perturbation h as before to be *positive* relative to a point $(P_-, P_+) \in I_L$ if h has the profile in Figure 33, in other words, if V_h contains $(\gamma_+ \times P_+) \cup (\gamma_- \times P_-)$. It is negative relative to (P_-, P_+) in the opposite case. We denote all the intersection points where h is positive by $(I_L)^+_h$.

3.2.2. Unobstructed cobordisms

The definition of unobstructed marked cobordism below is quite complicated. This being said, we require a minimal list of conditions to attain the following aims: first, such a cobordism V should allow the definition of Floer homology $HF(V', V)$ for $V' = \gamma \times L$ with $\gamma \subset \mathbb{C}$ an appropriate planar curve and $L \in Lag^*_e(M)$ (recall that the subscript e stands for the subclass of embedded Lagrangians); second, the data associated to V should restrict appropriately to the ends; finally, cabling in the sense of Section 2 should be possible within this class. In practice, these conditions will become much simpler in the exact setting that we will mainly focus on, as indicated in Section 3.3.2.

Assume that $\{(L'_j, c'_j)\}$ and $\{(L_i, c_i)\}$ are, respectively, marked Lagrangians in M.

DEFINITION 3.2.2

A *marked cobordism* (V, h, c) between these two families of marked Lagrangians is an immersed Lagrangian cobordism $V : (L'_j) \rightsquigarrow (L_i)$ as in Section 3.2.1, with double points either transverse and isolated or clean along 1-dimensional intersections, together with a perturbation V_h as above and a choice of $c \in I_{V_h}$ such that

(i) the restriction of c over the positive bottlenecks coincides with c'_j for the end L'_j; similarly, the restriction of c over the negative bottlenecks coincides with c_i for the end L_i;

(ii) the perturbation h is positive relative to each point of c that corresponds to an end, as in point (i) above.

Defining unobstructedness for cobordisms requires consideration of some additional moduli spaces of (marked) J-holomorphic curves. The reason is that cabling produces teardrops naturally as one can see from Figure 12. In turn, this leads to difficulties with regularity, as mentioned before in Remark 3.1.3(b).

Consider a marked cobordism (V, h, c) as above. To define the new moduli spaces, we first fix some points in the plane $P_1, \ldots, P_k$ such that $(\{P_i\} \times M) \cap V_h = \emptyset$. The new moduli spaces consist of marked, perturbed J-holomorphic polygons u (just as in the definitions in Section 3.1.1) with boundary along V_h but with additional distinct (moving) interior marked points $b_1, b_2, \ldots, b_r \in \mathrm{int}(D^2)$ with the property that $u(b_i) \in \{P_j\} \times M$ for some j. Around each point P_j, we fix closed disks D_j, D'_j such that $D_j \subset D'_j \subset \mathbb{C}\backslash\pi(V_h)$ and $P_j \in \mathrm{int}(D_j) \subset D_j \subset \mathrm{int}(D'_j)$. We will assume that each bounded, connected component of $\mathbb{C}\backslash\pi(V_h)$ contains *at most* one point P_j. Moreover, we assume that, for each P_j, there is *at most one* of the marked points b_j that is sent to $P_j \times M$. See Figure 34.

REMARK 3.2.3

The technical assumption that each interior marked point b_i is sent to a different $P_j \times M$ is very restrictive. Given that (M, ω) is assumed to be exact and that

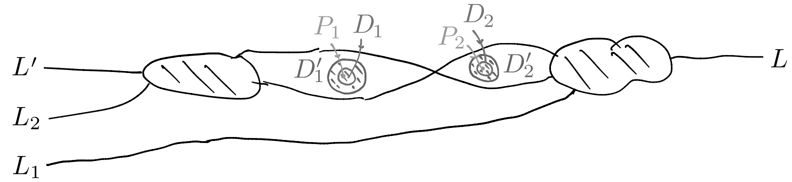

Figure 34. (Color online) The projections of a perturbed cobordism V_h onto $\mathbb{C}$ together with the points $P_1, P_2 \in \mathbb{C}$ used to stabilize teardrops.

the complex hypersurfaces $\{P_j\} \times M$ are mutually disjoint, the compactifications of our moduli spaces do not contain configurations where the points b_i collide (i.e., bubbling off of spheres, and even of "ghost" spheres, resulting from the collision of the b_i's does not occur). As will be discussed briefly below, there are ways to deal with such collisions and thus this technical condition could be relaxed. However, these methods would considerably increase the technical complexity of the present paper, and they generally also require working over $\mathbb{Z}$ or $\mathbb{Q}$ and not $\mathbb{Z}/2$. Thus, our approach is based on the idea that the actual teardrops (without boundary $\mathbf{c}$-marking) in our setting only appear in relation to the cabling construction and fit in this setting.

We will consider moduli spaces of curves u with marked points along the boundary and with interior marked points b_i that satisfy a Cauchy–Riemann-type equation such that the allowable perturbations appearing in the (perturbed) Cauchy–Riemann-type equation associated to these u contain a perturbation supported in small neighborhoods of the points b_i, in the domain of u, and defined in terms of almost-complex structures and Hamiltonian perturbations supported inside $D_i \times M$. Over the region $D'_i \backslash D_i$, the Hamiltonian perturbation term is trivial and the almost-complex structure is of the form $i \times J_0$ (outside of the region $D'_i \times M$, we allow for the possibility of other perturbation terms; in short, the support of our perturbations is in $(\mathbb{C} \backslash D'_i) \cup D_i$). The resulting moduli spaces will be denoted by $\mathcal{M}_{(\mathcal{D},P_1,\ldots,P_r);\mathbf{c};V}(x_1,\ldots,x_m;y)$. Notice that the choices available for a marked point b_i as before add two parameters to the relevant moduli space and the condition $u(b_i) \in \{P_i\} \times M$ subtracts two parameters.

Our choices mean that, for any point P_i and any curve u in one of these moduli spaces, there is a well-defined degree of the map $\pi \circ u$ at any of the points P_i and, because the composition $\pi \circ u$ is holomorphic over $D'_i \backslash D_i$, this degree is positive or null. Moreover, the degree is null only if the curve u does not go through P_i.

We will say that a curve u is of *positive degree* if one of these degrees is nonvanishing and is *plane-simple* if each of these degrees is 0 or 1. All the curves that we will need to use here will be plane-simple.

We denote by $\|u\|$ the total valence of u—the number given as $|u|+$ twice the number of internal marked points of u (recall that $|u|$ is the number of boundary punctures of u). By convention, we will assume that an element $u \in \mathcal{M}_{(\mathcal{D},P_1,\ldots,P_s);\mathbf{c};V}(x_1,\ldots,x_m;y)$ has a positive degree at each of the P_i's that appear as subscripts. As mentioned before, in our examples, the degree of the curves u at each point P_j is at most 1 and there is precisely a single marked point b_j associated to each P_j for which this degree is nonvanishing.

We will view below the choice of points $P_1,\ldots,P_k$ and the relevant perturbations to be included in the perturbation data $\mathcal{D}_{V_h}$. We call the points P_j the *pivots* associated to V and we will denote the associated hypersurface $\bigcup_j P_j \times M$ by H_V. In terms of terminology, we will refer to $\mathbf{c}$-marked polygons, teardrops, and so on without generally distinguishing between the corresponding objects with interior marked points and those without them: thus (nonmarked) teardrops are curves with $|u|=1$ and $\|u\|\geq 1$. We emphasize that there is a slight abuse of language at this point as *nonmarked* teardrops might be curves with a single boundary marking which is an output (and thus $|u|=1$) but that have some interior markings (sent to hypersurfaces associated to pivots P_i) and thus $\|u\|\geq 3$. In case more precision is needed, we will talk about such curves as non-**c**-marked teardrops.

In this setting, we define unobstructed cobordisms as follows.

DEFINITION 3.2.4

A cobordism (V,h,c), as in Definition 3.2.2, is *unobstructed* if:

(i) there is a system of coherent perturbation data $\mathcal{D}_{V_h}$ containing a base almost-complex structure, denoted $\bar{J}_0$, defined on $\mathbb{C}\times M$ and such that:

(a) the moduli spaces of $\bar{J}_0$-holomorphic **c**-marked disks are void and the moduli spaces of nonmarked $\bar{J}_0$-holomorphic teardrops are also void except for those moduli spaces of teardrops that are of positive degree and plane-simple,

(b) the regularity and coherence conditions in Definition 3.1.2 apply but only to those moduli spaces of curves with $\|u\|\geq 2$ consisting of **c**-marked polygons and the moduli spaces appearing in their compactification (as in Remark 3.1.3(c)),

(c) the 0-dimensional moduli space of teardrops with $\|u\|\geq 3$ associated to $(V_h,c,\mathcal{D}_{V_h})$ at any point $y\in I(V_h)$ is finite and the number of its elements $\bmod 2$ vanishes (as in Definition 3.1.4);

(ii) the data $\mathcal{D}_{V_h}$ extends the corresponding data for the ends in the following sense:

(a) there are coherent systems of perturbations $\mathcal{D}_{L_i}$ and $\mathcal{D}_{L'_j}$ for all the ends of V such that the base almost-complex structure J_0 (as in Definition 3.1.2(i)) is the same for all L_i, L'_j;

(b) the triples $(L_i,c_i,\mathcal{D}_{L_i})$ and $(L'_j,c'_j,\mathcal{D}_{L'_j})$ are unobstructed;

(c) the base almost-complex structure $\bar{J}_0=\bar{J}_{V_h}$ has the property that $\bar{J}_0=i\times J_0$ over a region of the form $(-\infty,-a-\frac{3}{2}]\times\mathbb{R}\cup[a+\frac{3}{2},\infty)\times\mathbb{R}$;

(d) the coherent system of perturbations $\mathcal{D}_{V_h}$ has the property that it restricts to the $\mathcal{D}_{L_i}$ (resp., $\mathcal{D}_{L'_j}$) over all negative (resp., positive) bottlenecks;

(e) the perturbation h and the data $\mathcal{D}_{L_i}$, $\forall i$, have the property that the moduli spaces of polygons in M with boundary on L_i and with entries all belonging to $(I_{L_i})^+_h$ are nonvoid only if the exit of the polygons also belongs to $(I_{L_i})^+_h$. This is similarly the case for $\mathcal{D}_{L'_j}$ and L'_j, $\forall j$.

We will denote an unobstructed cobordism in this sense by $(V,h,c,\mathcal{D}_{V_h})$.

REMARK 3.2.5

(a) Assumptions (i) and (ii) in Definition 3.2.2 and assumptions (ii)(a)–(d) in Definition 3.2.4 are basically what is expected from compatibility with the ends. We have explained after Definition 3.2.2, at the beginning of this section, the reason for allowing the existence of teardrops as well as the use of interior marked points in achieving regularity.

The remaining parts of Definition 3.2.4 that warrant explanation are assumptions (i)(b) and (ii)(e). These two requirements are a reflection of the following remark. Consider a moduli space $\mathcal{M}$ of polygons in M with boundary on an end L_i, and assume that it is regular in M. The same moduli space can be viewed as a moduli space $\widehat{\mathcal{M}}$ of polygons in $\mathbb{C} \times M$, with boundary along V_h, but included in the fiber over a bottleneck. In general, $\widehat{\mathcal{M}}$ is regular only if h is positive at all the corners of the polygons. There might exist nonvoid moduli spaces of polygons in M with corners that are not all positive for h. When these moduli spaces are viewed as moduli spaces in $\mathbb{C} \times M$, they are no longer regular and further perturbations to remedy this situation interfere with condition (ii)(c). We deal with this difficulty by using condition (ii)(e) to ensure that moduli spaces of **c**-marked polygons, included in the fiber over a bottleneck, remain regular in $\mathbb{C} \times M$ and that in the compactification of such moduli spaces only appear moduli spaces with corners where h is positive, and thus their regularity is also achieved. This is sufficient for our purposes as all algebraic structures of interest to us are defined in terms of **c**-marked polygons. It is not completely trivial to see that, under the positivity assumption mentioned above, the fiber moduli spaces remain regular in $\mathbb{C} \times M$. In [10] this point is discussed explicitly, but the setting there is simpler as the branches meeting at intersection points belong to different connected components (each embedded) and the choices of perturbations are simpler. Our case, however, looks similar to [10, Section 4.3.1], at least locally, in that near the double point that lies over the bottleneck at $\{a+2\} \times \{k\}$ the perturbed immersion V_h looks like $(\gamma_+ \times D_+) \cup (\gamma_- \times D_-)$, which is compatible with the splitting of the almost-complex structure in (ii)(c).

While we will not give a rigorous argument to prove this regularity statement, here is the intuition underlying it. We start with the unperturbed (also called *Morse–Bott*) analytic setup (in a neighborhood of the bottleneck). The "clean" self-intersection takes place along a segment passing through the bottleneck. The analysis in this context already appears in the literature, for instance, in [3] and [46]. Starting with generic Morse–Bott data, the perturbation h can be viewed as a "Morsification." Standard arguments in this case imply (similarly, to usual Morse–Bott theory) that the moduli spaces in the fiber are regular as soon as their virtual indexes in M and in $\mathbb{C} \times M$ agree. This happens when the inputs of the respective polygons are positive for h.

Obviously, given fixed ends for a cobordism together with associated coherent systems of data, showing that a perturbation h as before exists is nontrivial. However, this turns out to be rather easy to do under the exactness assumptions that we will use in our actual constructions.

(b) Inserting the data h in the definition of a marked cobordism has the disadvantage that one needs to be precise on how gluing of cobordisms behaves with respect to these perturbations. In particular, if (V, h) and (V', h') need to be glued along a common end L, then the perturbations h and h' have to be positive exactly for the same double points of L. On the other hand, there is no obvious way to avoid including this h in the definition because associating the marking c directly to connected components of I_V is not sufficient to canonically induce a marking to a deformation V_h and, at the same time, for Floer homology type considerations the actual object in use is precisely the marked Lagrangian V_h.

(c) In our arguments, whenever non-$\mathbf{c}$-marked teardrops contribute to the definition of algebraic structures they will carry interior punctures b_j and the relevant moduli spaces will be regular (condition (i)(b)). At the same time, if such teardrops exist, then there are also unperturbed teardrops, with the same degree relative to each pivot P_i, for which regularity is not necessarily true (they can be seen to exist by making perturbations $\to 0$). This explains the formulation in (i)(a): these nonmarked $\bar{J}_0$ (unperturbed) teardrops are allowed to exist even if in all arguments only perturbed curves u with $\|u\| \geq 3$ will be used.

Unobstructed, marked, Lagrangian cobordisms in the sense above can be assembled in a class $Lag^*(\mathbb{C} \times M)$ as in Definition 3.1.6. We will denote by $\bar{\mathcal{D}}$ the perturbation data for this family. More details on how to pick this perturbation data will be given further below, in Section 3.3.3.

3.3. Exact marked Lagrangians and cobordisms

In this subsection, we specialize the discussion in Sections 3.1 and 3.2 to the case of exact Lagrangians and cobordisms. Our purpose is to prepare the ground to set up the category $\mathsf{C}ob^*(M)$ that appears in Theorem A.

We fix the context for the various considerations to follow. We assume $(M, \omega = d\lambda)$ to be either a compact Liouville domain or a Liouville manifold complete at infinity. Henceforth the Liouville form λ will be fixed. We will consider Lagrangian immersions $j_L : L \to M$ that are exact and *endowed with a primitive* $f_L : L \longrightarrow \mathbb{R}$ of $j_L^*\lambda$, that is, $j_L^*\lambda = df_L$. We will also make two additional assumptions that are generic:

(i) j_L only has double points that are isolated, self-transverse intersections;
(ii) if $(P_-, P_+) \in I_L$, then $f_L(P_-) \neq f_L(P_+)$.

For such an exact Lagrangian, we denote by $I_L^{<0}$ the set of all *action-negative* double points $(P_-, P_+) \in I_L$. By definition, these are the double points (P_-, P_+) such that $f_L(P_+) < f_L(P_-)$. Similarly, we denote by $I_L^{>0}$ the complement of $I_L^{<0}$ in I_L. We will only consider marked Lagrangians of type (L, c), where L is exact with a fixed primitive as above and

(iii) $c \subset I_L^{<0}$.

Notice that, in general, the Lagrangian L is not assumed to be connected. Condition (iii) will play an important role in a number of places in our construction, in particular, in relation to unobstructedness and when composing cobordisms together with their perturbations. Condition (ii) is of use, among other places, to properly define the composition of cobordisms. Notice that if j_L is an embedding, then $c = \emptyset$ and L trivially satisfies the unobstructedness condition in Definition 3.1.4. Moreover, the class of all such embedded Lagrangians is easily seen to satisfy Definition 3.1.6(ii).

Finally, we consider an unobstructed class $Lag^*(M)$ consisting of finite unions of *exact*, marked Lagrangians $(L, c, \mathcal{D}_L)$ (as above) in the sense of Definition 3.1.6; in particular, they satisfy Definition 3.1.2(i)(b) (i.e., all moduli spaces of non-marked teardrops are void). In practice, at the center of our construction of such a class $Lag^*(M)$ with surgery models are Lagrangians of the following type: a sequence of embeddings, $e_1 : L_1 \to M, \ \ldots, e_k : L_k \to M$ that are in generic position and a marking (constructed iteratively) for the disjoint union $L_1 \cup L_2 \cup \cdots \cup L_k$.

We denote by $\mathcal{D}$ the perturbation data for this class and by $\mathcal{D}_L$ the relevant perturbation data for each individual Lagrangian L, as in Definition 3.1.6. For embedded Lagrangians L, the data $\mathcal{D}_L$ reduces to fixing one almost-complex structure J on M. We will assume that

(iv) the collection $Lag^*(M)$ contains all exact embedded Lagrangians in M; more precisely, for a certain almost-complex structure J—called the *ground a.c.s.*—on M, all exact embedded Lagrangians L (with all choices of primitives) and with $\mathcal{D}_L = J$ belong to $Lag^*(M)$,

$$(L, \emptyset, J) \in Lag^*(M). \tag{15}$$

REMARK 3.3.1

The same geometric immersion $j_L : L \to M$ may appear numerous times inside the class $Lag^*(M)$: first, it can appear with more than a single primitive f_L; with various markings c; even if both the primitive and the marking are fixed, there are possibly more choices of data $\mathcal{D}_L$ leading to different triples $(L, c, \mathcal{D}_L)$. At the same time, the class $Lag^*(M)$ does not necessarily contain all the unobstructed triples as before, but it does contain all embedded Lagrangians in the sense described above.

3.3.1. Fukaya category of unobstructed marked Lagrangians

In view of Definition 3.1.6, we apply Seidel's method to construct an A_∞-category associated with objects in the class $Lag^*(M)$ defined as above with perturbation data $\mathcal{D}$. We use homological notation here but, in matters of substance, there is a single difference in our setting compared to the construction in [48]. In our case, as we deal with *marked* Lagrangians (L, c), the operations μ_k, $k \geq 1$,

$$\begin{aligned} \mu_k : CF\big((L_1, c_1), (L_2, c_2)\big) \otimes \cdots \otimes CF\big((L_k, c_k), (L_{k+1}, c_{k+1})\big) \\ \longrightarrow CF\big((L_1, c_1), (L_{k+1}, c_{k+1})\big) \end{aligned}$$

are defined by counting elements u in 0-dimensional moduli spaces of $\mathbf{c}$-*marked* polygons

$$u \in \overline{\mathcal{M}}_{\mathcal{D};L_1,\dots,L_{k+1}}(h_1,\dots,h_k;y)$$

as defined in Section 3.1.2. The conditions in the unobstructedness Definitions 3.1.4 and 3.1.6 imply that the sums involved in the definitions of the μ_k's are well defined and finite and that the usual A_∞-relations are true in this setting. At the same time, for the reader's convenience, we recall that even in this exact setting the absence of teardrops depends on a choice of *base* almost-complex structure J_0.

REMARK 3.3.2

A slightly delicate point is to show that this category is (homologically) unital. One way to do this is to use a Morse–Bott-type description of the complex $CF(L,L)$, where $j_L : L \to M$ is an unobstructed, marked immersion in our class (with L connected), through a model that consists of a complex having as generators the critical points of a Morse function f on L together with two copies of each of the self-intersection points of L. The differential has a part that is given by the usual Morse trajectories of f, another part that consists of marked J-holomorphic strips joining self-intersection points of L, as well as an additional part that consists of (marked) teardrops starting (or arriving) at a self-intersection point joined by a Morse flow line to a critical point of f (this model is in the spirit of [22] and appears in the nonmarked case in [4]). Of course, one also needs to construct the operations μ_k in this setting. We omit the details here but mention that, with this model, it is easily seen that the unit is represented by the class of the maximum, as in the embedded case.

This leads to an A_∞-category $\mathcal{F}uk^*_i(M;\mathcal{D})$, where i indicates that the objects considered here are in general immersed, marked Lagrangians. This category contains the subcategory $\mathcal{F}uk^*(M;\mathcal{D})$ whose objects are the embedded Lagrangians $\mathcal{L}ag^*_e(M) \subset Lag^*(M)$. There are also associated derived versions, $D\mathcal{F}uk^*_i(M;\mathcal{D})$ and $D\mathcal{F}uk^*(M;\mathcal{D})$. The Yoneda functor

$$\mathcal{Y} : \mathcal{F}uk^*_i(M;\mathcal{D}) \to \operatorname{mod}\big(\mathcal{F}uk^*_i(M;\mathcal{D})\big)$$

continues to be well defined and the Yoneda lemma from [48] still applies. (Here it is important to recall that $\mathcal{F}uk^*_i(M;\mathcal{D})$ is homologically unital.) In particular, there is a quasi-isomorphism:

$$\operatorname{Mor}_{\mathrm{mod}}\big(\mathcal{Y}(L),\mathcal{Y}(L')\big) \simeq \operatorname{Mor}_{\mathcal{F}uk}(L,L'). \tag{16}$$

3.3.2. The class $Lag^*(\mathbb{C}\times M)$

In this subsection, we more precisely delineate the type of cobordisms that will appear in the class $Lag^*(\mathbb{C}\times M)$. This class can be viewed as a close approximation to the actual class $\mathcal{L}ag^*(\mathbb{C}\times M)$ from Theorem A. Our conventions are that $(\mathbb{C},\omega_0)$ is exact with primitive $\lambda_0 = xdy$, $(x,y)\in\mathbb{R}\oplus\mathbb{R}\cong\mathbb{C}$ and, similarly, the primitive of $\omega_0\oplus\omega$ is $\lambda_0\oplus\lambda\in\Omega^1(\mathbb{C}\times M)$.

Consider a cobordism

$$V : (L_i) \rightsquigarrow (L'_j).$$

We fix some additional notation. Given some data defined for the Lagrangian V_h, obtained by perturbing V as described in Section 3.2.1, we denote by $(-)|_{L_i}$ the restriction of this data to the end L_i in the fiber (of $\pi : \mathbb{C} \times M \to \mathbb{C}$) that lies over the bottleneck corresponding to L_i, and similarly for the negative ends L'_j.

We consider an unobstructed family $Lag^*(\mathbb{C} \times M)$ of cobordisms $(V, c, h, \mathcal{D}_{V_h})$, $V : (L_i) \rightsquigarrow (L'_j)$ as in Section 3.2.2, subject to the following additional conditions:

(i) We have $L_i, L'_j \in Lag^*(M)$, $\forall i, j$.

(ii) V is exact and the primitive f_V on V_h has the property that $f_V|_{L_i} = f_{L_i}$, $f_V|_{L'_j} = f_{L'_j}$, $\forall i, j$.

(iii) The double points of V_h have the property that $f_V(P_-) \neq f_V(P_+)$ whenever $(P_-, P_+) \in I_V$ and, moreover, $c \subset I^{<0}_{V_h}$.

(iv) The perturbation h is positive exactly for those points $(P_-, P_+) \in I_{V_h}|_{L_i, L'_j}$ such that $(P_-, P_+) \in I^{<0}_{V_h}$.

REMARK 3.3.3

(a) Notice that if V is exact, then its deformation V_h is also exact. In case V is embedded, the perturbation h is taken to be 0 and all restrictions over bottlenecks simply correspond to restrictions to the fiber over the corresponding point (where V is horizontal).

(b) Conditions (i), (ii), and (iii) are the immediate analogues of the conditions already imposed on $Lag^*(M)$. Condition (iv) ties the perturbation h to the "action negativity" of the self-intersection points of the ends L_i, L_j of V in a natural way. The impact of this condition is further discussed below.

(c) In the setting above, we can revisit the regularity constraint in Definition 3.2.4(i) and condition (ii)(e) in the same definition. The key point is that, at least when the Hamiltonian terms in the perturbation data $\bar{\mathcal{D}}$ are sufficiently small (which we will implicitly assume from now on), a nonconstant polygon $u \in \mathcal{M}_{\mathcal{D}_{V_h}; V_h}(x_1, \ldots, x_m; y)$ has energy

$$0 < E(u) \approx \sum_{i=1}^{m} \big(f_V\big((x_i)_+\big) - f_V\big((x_i)_-\big)\big) - \big(f_V(y_+) - f_V(y_-)\big),$$

where $x_i \in I_{V_h}$ is written as $x_i = ((x_i)_-, (x_i)_+) \in V_h \times V_h$. (The sign $\approx$ is due to the curvature terms that we do not make explicit here. The important point is that by taking the perturbations to be small enough we can make the error in the approximation $\approx$ arbitrarily small. Note also that the error in $\approx$ above depends on m, and in general cannot be bounded uniformly in m; however, this will not interfere with our purposes since we will consider only a finite number of polygons u at each time.) Thus, if $x_i \in I^{<0}_{V_h}$ (which, we recall, means that $f_V((x_i)_-) > f_V((x_i)_+)$, then we also have $y \in I^{<0}_{V_h}$. In particular, this means that Definition 3.2.4(ii)(e) is automatically satisfied in our context because condition

(iv) above requires that the perturbation h be positive precisely at those self-intersection points (along the ends) that are action-negative. Moreover, recall that we assume (at (iii) above) that the marking $c \subset I_{V_h}$ is also action-negative in the sense that $c \subset I^{<0}_{V_h}$. This implies that a c-marked polygon with boundary on V_h has an exit $y \in I^{<0}_{V_h}$ and thus there are no c-marked J-disks in this setting, thus the disk part in Definition 3.2.4(i)(a) is automatically satisfied. Further, it is easy to see that all moduli spaces appearing in the compactification of moduli spaces of c-marked polygons with boundary on a *single* V_h, as above, only have action-negative inputs and an action-negative exit (this follows from the energy estimate above which implies, as mentioned above, that if all the inputs are action-negative, then the output also has to be action-negative; we are not defining here the operations of an A_∞-category but only focusing on some configurations associated to V_h). In short, the regularity constraint in Definition 3.2.4 can be ensured by requiring that all moduli spaces $\mathcal{M}_{\mathcal{D}_{V_h};V_h}(x_1, \ldots, x_m; y)$ with $x_i, y \in I^{<0}_{V_h}, \forall i$ be regular (see also Remark 3.2.5).

In view of the above remark, the regularity requirements become simpler in our exact setting and we adjust them slightly here to a form that is easier to use in practice. We first notice that the discussion in Remark 3.3.3(c) above also applies in an obvious way to the Lagrangians $L \in \mathcal{L}ag^*(M)$ and to their systems of coherent perturbations $\mathcal{D}_L$. Thus, in our context, regularity of moduli spaces of polygons with boundary on $L \subset M$ (resp., $V \subset \mathbb{C} \times M$) is only necessary for those moduli spaces such that all inputs are action-negative (see also Remark 3.1.3(c)), as long as the Hamiltonian perturbation terms in the respective systems of perturbations are sufficiently small. We will work from now on with this understanding of regularity. We formulate it more precisely below.

Given all the various definitions and conditions contained in the previous subsections, it is useful to list here all the types of assumptions that operate in our setting. First, we recall that we work with Lagrangians $(L, c, \mathcal{D}_L)$ and cobordisms $(V, c, h, \mathcal{D}_{V_h})$ that are exact, marked, endowed with fixed primitives such that each self-intersection point of L and V_h are action-positive or action-negative (but nonzero) and also with fixed coherent perturbation data $\mathcal{D}_L$ and, respectively, $\mathcal{D}_{V_h}$. The perturbations $\mathcal{D}_{V_h}$ associated to cobordisms also include choices of points in the plane and associated domain-dependent perturbations that are used to deal with teardrops of positive degree (see Section 3.2.2).

In view of the discussion above, the classes $\mathcal{L}ag^*(M)$, $\mathcal{L}ag^*(\mathbb{C} \times M)$ satisfy the following.

ASSUMPTION (I) (Regularity)

(i) The Hamiltonian perturbation terms for $\mathcal{D}_L$, $\mathcal{D}_{V_h}$ are sufficiently small so that moduli spaces with action-negative inputs and an action-positive output (or no output) are void.

(ii) The condition in Definition 3.1.2(ii) applies to only those moduli spaces with action-negative inputs.

(iii) The condition in Definition 3.2.4(i)(b) applies to all moduli spaces with action-negative inputs.

There are six other assumptions (each detailed in the various points in Definitions 3.1.2, 3.1.4, 3.1.6, 3.2.2, 3.2.4):

(II) (Negativity of markings) All markings for Lagrangians and cobordisms are action-negative.

(III) (Restriction to ends) Cobordism data restricts to end data over the bottlenecks. In particular, the ends of $V \in \mathit{Lag}^*(\mathbb{C} \times M)$ are in $\mathit{Lag}^*(M)$.

(IV) (Positivity of perturbations) For cobordisms (V, h, c), the perturbation h is positive precisely at the action-negative self-intersection points of the ends.

(V) (Unobstructedness) The number of $\mathbf{c}$-marked teardrops through any self-intersection point for Lagrangians as well as cobordisms is $0 \bmod 2$. The moduli spaces of (nonmarked) teardrops are void for the elements of $\mathit{Lag}^*(M)$. The moduli spaces of (nonmarked) teardrops for the elements of $\mathit{Lag}^*(\mathbb{C} \times M)$ are void, except possibly for those that are of positive degree and plane-simple.

(VI) (Regularity of data for families) The data $\mathcal{D}$ for $\mathit{Lag}^*(M)$ (resp., $\bar{\mathcal{D}}$ for $\mathit{Lag}^*(\mathbb{C} \times M)$) extends the data $\mathcal{D}_L$ (resp., $\mathcal{D}_{V_h}$), and is regular.

(VII) (Nontriviality) All embedded exact Lagrangians are contained in $\mathit{Lag}^*(M)$ and all embedded exact cobordisms are contained in $\mathit{Lag}^*(\mathbb{C} \times M)$.

In our context, as mentioned above, Definition 3.2.4(ii)(e) is redundant and the energy bounds condition in Definition 3.1.6(ii) is automatically satisfied. Moreover, because all the markings are action-negative, there are no $\mathbf{c}$-marked J-holomorphic disks. The choice of perturbation data $\bar{\mathcal{D}}$ will be made explicit below. The part of Assumption VII that applies to cobordisms means that we assume that the class $\mathit{Lag}^*(\mathbb{C} \times M)$ contains all exact *embedded* Lagrangian cobordisms V as quadruples $(V, \emptyset, 0, i \times J)$, where J is the ground almost-complex structure from (15). An analogue of Remark 3.3.1 remains valid in the case of immersed cobordisms $j_V : V \to \mathbb{C} \times M$: such a cobordism can appear multiple times in the class $\mathit{Lag}^*(\mathbb{C} \times M)$ due to different primitives f_V, different perturbations h, different markings c, and different perturbation data $\mathcal{D}_{V_h}$.

In reference to Assumption V above, if some moduli space of nonmarked teardrops associated to an element $V \in \mathit{Lag}^*(\mathbb{C} \times M)$ is nonvoid, then we say that *V carries nonmarked teardrops*. If all these moduli spaces are void, then we say that V is *without nonmarked teardrops*. The purpose of the distinction is that the elements of $\mathit{Lag}^*(\mathbb{C} \times M)$ without nonmarked teardrops form a Fukaya category, while the other unobstructed cobordisms can be used in Floer-type arguments but with more care due to the presence of the pivots in the plane (see also Section 3.2.2). We denote the subset of $\mathit{Lag}^*(\mathbb{C} \times M)$ consisting of cobordisms without nonmarked teardrops by $\mathit{Lag}^*_0(\mathbb{C} \times M)$ (of course, this class depends on the choice of almost-complex structure J).

3.3.3. The Fukaya category of cobordisms

We denote the class of embedded cobordisms by $Lag_e^*(\mathbb{C}\times M)$. Embedded cobordisms are the objects of a Fukaya category $\mathcal{F}uk^*(\mathbb{C}\times M)$ as constructed in [10]. We extend this construction here, namely, we define a Fukaya category of immersed, marked, unobstructed cobordisms $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ that contains $\mathcal{F}uk^*(\mathbb{C}\times M)$ as a full subcategory. Along the way, we make more explicit the choice of the perturbation data $\bar{\mathcal{D}}$. The objects of $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ are the elements of $Lag_0^*(\mathbb{C}\times M)$, in other words, those unobstructed cobordisms without nonmarked teardrops. To fix ideas, we emphasize that, because the objects of $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ do not carry any non-**c**-marked teardrops, there are no pivots P_j or interior marked points b_i for the moduli spaces considered in this construction. Because of this the construction of the Fukaya category $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ is very similar to the construction in the embedded case. However, the cobordisms obtained from cabling and other such constructions are not objects in this category.

The construction is an almost identical replica of the construction of the embedded cobordism category $\mathcal{F}uk^*(\mathbb{C}\times M)$ in [10] with a few modifications that we indicate now.

(a) Profile functions and cobordism perturbations

Recall from Section 3.2.1 that the immersed marked cobordisms (V,c,h) have bottlenecks at points of real coordinates $a+2$ for the positive ends and $-a-2$ for the negative ends (here $a\in(0,\infty)$ is some large constant). We will make the additional assumption that all the perturbations h associated to marked immersed cobordisms (see Figure 33) have the property that the projection of the perturbed Lagrangians along one end is included inside the strip of width 2ϵ (for a fixed small ϵ) along the respective horizontal axis and, more precisely, inside the region drawn in Figure 35 below.

We assume that the angle α in Figure 35 is equal to $\frac{\pi}{8}$. We will assume that all embedded cobordisms in $\mathcal{L}ag_e^*(\mathbb{C}\times M)$ are cylindrical outside the region $[-a-1,a+1]\times\mathbb{R}$. The construction of $\mathcal{F}uk^*(\mathbb{C}\times M)$ makes use of a profile function that we will denote here by $\hbar$ (this appears in [10, pp. 1760–1761] with

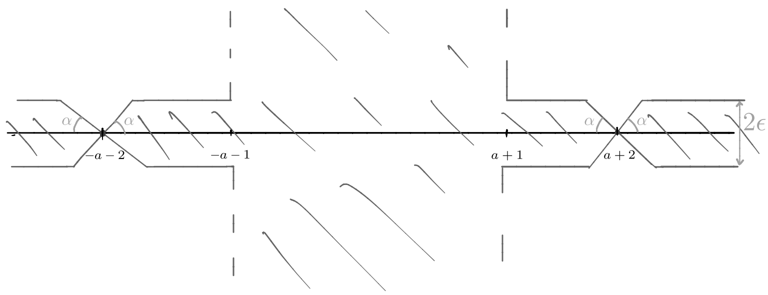

Figure 35. (Color online) The region containing the projection of all perturbations h associated to marked, immersed cobordisms (V,c,h).

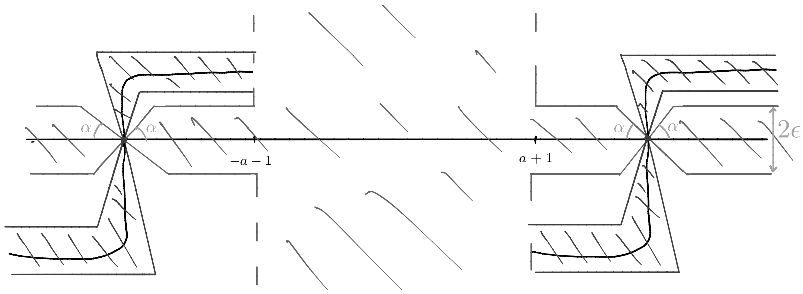

Figure 36. (Color online) The red region containing the transformation through $(\phi_1^{\hbar})^{-1}$ of the blue region.

the notation h; we use $\hbar$ here to distinguish this profile function from the perturbations associated to immersed marked cobordisms). We recall that, in essence, the role of the profile function $\hbar : \mathbb{R}^2 \to \mathbb{R}$ is to disjoin the ends of cobordisms at ∞ by using the associated Hamiltonian flow $\phi_t^{\hbar}$. This profile function also has some bottlenecks, as they appear in [10, Figures 8 and 9]. The real coordinate of the bottlenecks is fixed in [10] at $\frac{5}{2}$ for the positive ones and $-\frac{3}{2}$ for the negative (of course, it is later shown that this is just a matter of convention). Adjusting that profile function to our case, our first assumption here is that the positive bottlenecks of $\hbar$ have real coordinate $a+2$ and that the negative one has real coordinate $-a-2$. Moreover, we will assume that the (inverse) Hamiltonian isotopy $(\phi_1^{\hbar})^{-1}$ associated to $\hbar$ transforms the region in Figure 35 as in Figure 36 below.

(b) Choices of almost-complex structures

Once the profile functions are chosen as above, the construction proceeds as in Section 3 of [10] with an additional requirement having to do with the choice of almost-complex structures (see points ii and iii* in [10, pp. 1763–1764]) associated to a family of marked, unobstructed cobordisms $\{(V_i, c_i, h_i, \mathcal{D}_{V_{h_i}})\}_{1\leq i\leq k+1}$ (we denote $V_{h_i} = (V_i)_{h_i}$). The requirement is that the choice of (domain-dependent) almost-complex structure $\mathbf{J} = \{J_z\}_z$ on $\mathbb{C}\times M$, where $z \in S_r$ with r varying in the relevant family of pointed disks, have the property that it agrees with the $\bar{J}_{V_i}$—this is the base almost structure of $\mathcal{D}_{V_{h_i}}$ (see Definition 3.2.4(ii)(c))—along each boundary of S_r that lies on the cobordism V_{h_i}. This condition is added to point iii* in [10, p. 1764]. All the relevant perturbation data is collected in the set $\bar{\mathcal{D}}$.

(c) Definition of the μ_k operations

With the data fixed as above, the operations μ_k are defined as in Section 3.3.1 (see also Section 3.1.2 for the definition of the relevant moduli spaces). Namely, they are given through counts of $\mathbf{c}$-marked polygons:

$$u \in \overline{\mathcal{M}}_{\bar{\mathcal{D}};V_{h_1},\ldots,V_{h_{k+1}}}(h_1,\ldots,h_k;y).$$

The moduli space $\overline{\mathcal{M}}_{\bar{\mathcal{D}};V_{h_1},\ldots,V_{h_{k+1}}}(h_1,\ldots,h_k;y)$ contains in this case all the **c**-marked polygons u (see also Section 3.2.2). Recall from Definition 3.2.4 that the almost-complex structures $\bar{J}_{V_h}$ are split outside a finite width vertical band in $\mathbb{C}$ that does not contain the bottlenecks. As a consequence, all the arguments related to naturality and compactness of moduli spaces in [10, Section 3.3] remain valid. It follows that the operations μ_k are well defined and that they satisfy the A_∞-relations.

From the discussion above, we conclude that the A_∞-category $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ is now well defined. Moreover, this category contains as a full subcategory $\mathcal{F}uk^*(\mathbb{C}\times M)$. In particular, to each unobstructed, marked cobordism without non-**c**-marked teardrops $(V,c,h,\mathcal{D}_{V_h})$ we may associate an A_∞ Yoneda module $\mathcal{Y}(V)$ over $\mathcal{F}uk_i^*(\mathbb{C}\times M;\bar{\mathcal{D}})$ and, by restriction, also an A_∞-module over $\mathcal{F}uk^*(\mathbb{C}\times M)$.

*3.3.4. Algebraic structures associated to cobordisms that carry non-**c**-marked teardrops*

Given a cobordism $V\in Lag^*(\mathbb{C}\times M)$ that carries non-**c**-marked teardrops, it is still possible to define certain Floer-type invariants associated to it. As previously mentioned, these cobordisms appear in the cabling construction and, thus, we need to take them into account too, even if we will not include them as objects in a Fukaya category.

To see how to associate such structures to V, consider a curve γ in the plane $\mathbb{C}$ such that γ only intersects $\pi(V)$ transversely along two fixed successive horizontal ends of $\pi(V)$ and is away from the pivots P_j associated to V. Let $U_\gamma\subset\mathbb{C}$ be a small neighborhood of γ that intersects $\pi(V)$ only at the same two ends (along small intervals) and all pivots $P_j\notin U_\gamma$. Consider the full subcategory $\mathcal{F}uk_{i,V}^*(\mathbb{C}\times M)$ of $\mathcal{F}uk_i^*(\mathbb{C}\times M)$ that has as objects those elements W of $\mathcal{F}uk_i^*(\mathbb{C}\times M)$ such that $\pi(W)\subset U_\gamma$.

With appropriate choices of perturbations (which we fix from now on), the cobordism $V=(V,h,c,\mathcal{D}_{V_h})$ induces a module, denoted $\mathcal{Y}(V)$ over $\mathcal{F}uk_{i,V}^*(\mathbb{C}\times M)$. In the definition of the different operations we only count curves that are plane-simple (i.e., such that the degree of these curves at each pivot $P_j\in\mathbb{C}$ is at most 1) and, moreover, such that for each point P_j where the degree is 1, there is a marked point b_j in the domain of the curve. The choice of relevant regular perturbations is possible by Seidel's scheme, by induction on $\|u\|$.

Compared to the construction of $\mathcal{F}uk_i^*(\mathbb{C}\times M)$, there is one additional delicate point to note: whenever two plane-simple curves can be glued, the outcome needs again to be plane-simple. Indeed, in case this condition—that we will refer to as *plane coherence*—is not satisfied, we would need to allow counts of curves where more than a single interior marked point is sent to the same $P_j\times M$ and this creates problems with our counting over $\mathbb{Z}/2$. We want to avoid this additional complication here and this is the purpose of the constraint on the objects of $\mathcal{F}uk_{i,V}^*(\mathbb{C}\times M)$ as it ensures that this type of plane-simple coherence is satisfied. Indeed, in this case, we see that the breaking for any 1-dimensional moduli space of the type that appears in checking the A_∞ identities for $\mathcal{Y}(V)$ has necessarily a component formed by curves that project inside U_γ and thus are of degree 0 with respect to any of the pivots P_j.

More generally, the same construction works for any other full subcategory of $\mathcal{F}uk_i^*(\mathbb{C}\times M)$, as long as each object continues to be disjoint from H_V and if this plane coherence condition continues to be satisfied. (Recall from page 442 that $H_V=\bigcup_j P_j\times M$, where the points P_j are the pivots associated to V.) The resulting module will again be denoted by $\mathcal{Y}(V)$.

REMARK 3.3.4

The prototype for the category $\mathcal{F}uk_{i,V}^*(\mathbb{C}\times M)$ is the image of an inclusion functor $i_\gamma:\mathcal{F}uk^*(M)\to\mathcal{F}uk_i^*(\mathbb{C}\times M)$ that associates to an object $L\subset M$ the object $\gamma\times L\subset\mathbb{C}\times M$ (see Section 4.2.1).

3.4. The category $\mathsf{C}ob^*(M)$ revisited

We now consider unobstructed classes of Lagrangians $Lag^*(M)$ and cobordisms $Lag^*(\mathbb{C}\times M)$ subject to Assumptions I–VII in Section 3.3.2. To define $\mathsf{C}ob^*(M)$, additional modifications of these classes are needed. These modifications have to do with the composition of morphisms in $\mathsf{C}ob^*(M)$ as well as with the identity morphisms and, as will be seen later on, with the definition of cabling.

Concerning composition, recall that for two immersed cobordisms $V:L\rightsquigarrow(L_1,\dots,L_k,L')$ and $V':L'\rightsquigarrow(L_1',\dots,L_s',L'')$, the composition $V'\circ V$ is defined as in Figure 4. However, the actual objects in $Lag^*(\mathbb{C}\times M)$ are immersed cobordisms together with perturbations (as well as additional decorations). This means that, to properly define the composition in this setting, we need to actually describe a specific perturbation of $V'\circ V$ (together with its other decorations) and we also need to show that the resulting object is unobstructed.

A similar issue has to do with the identity morphism. Assuming that L is immersed, we can consider the cobordism $V=\mathbb{R}\times L$ that should play the role of the identity. However, in our formalism we need to indicate a specific perturbation of this V.

3.4.1. An additional decoration for objects

Given an immersed marked Lagrangian

$$(L,c,\mathcal{D}_L)\in Lag^*(M),$$

the additional decoration we require is an ϵ-*deformation germ* for L. To make this notion precise, assume that L is the kth positive end of a marked, immersed cobordism $(V,c,h,\mathcal{D}_{V_h})$. The restriction to the band $[a+2-\epsilon,a+2+\epsilon]\times\mathbb{R}$ of the part of V_h corresponding to the end L can be viewed as a specific perturbation h_L of the (trivial) cobordism $[-\epsilon,\epsilon]\times\{1\}\times L$ with a bottleneck at 0, as defined in Section 3.2.1, with the profile of the picture in Figure 33 (around the point $a+2$), but translated so that the bottleneck is moved at the point $\{a+2\}\times\{k\}$. The same notion also applies for negative ends and, because of our constraints involving the positivity of perturbations (Assumption IV), the same perturbation h_L can appear as a restriction to a negative end (in the band $[-a-2-\epsilon,-a-2+\epsilon]\times\mathbb{R}$).

From now on, we will view the objects in $Lag^*(M)$ as quadruples $(L,c,h_L,\mathcal{D}_L)$, where h_L is a perturbation of $[-\epsilon,\epsilon]\times\{1\}\times L$ with a bottleneck at 0, in the sense above. Two such objects $(L,c,\mathcal{D}_L,h_L)$, $(L,c,\mathcal{D}_L,h'_L)$ will be identified if for a sufficiently small ϵ the perturbations h_L and h'_L agree on $[-\epsilon,\epsilon]$. We call h_L a *perturbation germ* associated to L. If L is embedded, then the perturbation germ h_L is null.

Given this notion, Assumption III in Section 3.3.2 is strengthened to require that if an immersed marked cobordism $(V,c,h,\mathcal{D}_{V_h})\in Lag^*(\mathbb{C}\times M)$ has an end $(L,c,h_L,\mathcal{D}_L)$, then h_L is the restriction of h.

REMARK 3.4.1

The same triple $(L,c,\mathcal{D}_L)$ can appear multiple times in the class $Lag^*(M)$ with different perturbation germs h_L. However, as we will see below, we will restrict the class of allowable perturbation germs to only sufficiently small perturbations (in a sense to be explained), and with this restriction any two perturbation germs will become equivalent from our perspective.

3.4.2. Definition of composition of immersed marked cobordisms

The composition of two cobordisms $V,V'\in Lag^*(\mathbb{C}\times M)$ along an end $L'\in Lag^*(M)$ (with the decorations listed in Section 3.4.1) follows the scheme in Figure 4, but by taking into account the perturbations h (resp., h') associated to V (resp., V') as in Figure 37.

In other words, if the two objects to be composed are $(V,c,h,\mathcal{D}_{V_h})$, $(V',c',h',\mathcal{D}_{V'_{h'}})$, then we define the resulting object $(V'\circ V,c'',h'',\mathcal{D}'')$ by letting h'' be the perturbation of the cobordism $V'\circ V$ (as defined in Figure 4) obtained by splicing together h' and h in a neighborhood of type $[b-\epsilon,b+\epsilon]\times[k-\delta,k+\delta]\times L'$ (with convenient choices of $b\in\mathbb{R}$, $k\in\mathbb{N}$, $\delta>0$). We can also splice together the perturbation data $\mathcal{D}_{V_h}$ and $\mathcal{D}_{V'_{h'}}$ and extend the result over $\mathbb{C}$ to get the perturbation data $\mathcal{D}''$. In particular, we assume that the almost complex structure in a band $[b-\epsilon,b+\epsilon]\times\mathbb{R}$ is a product, and

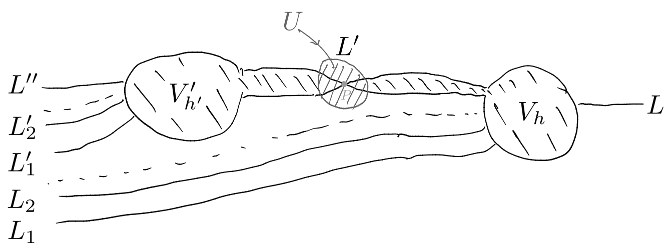

Figure 37. (Color online) The composition of $V,V'\in Lag^*(\mathbb{C}\times M)$. This composition restricts to the perturbation $h_{L'}$ over the set U. The bottleneck is the point P (the perturbations of V and V' along the ends different from L' are omitted).

we also assume that the same is true in a region under the "bulk" of $V'_{h'}$ that contains the extended ends of V_h ($L_1, L_2, \ldots$ in the picture). More details on the definition of $\mathcal{D}''$ are given in Section 3.4.3 below. Finally, the marking c'' is simply the union of c and c', where the marking of L'—in the fiber over the bottleneck P—is common to both c and c'.

3.4.3. *Unobstructedness of the composition*

To discuss this point, we start by taking a more detailed look at the following aspect related to the composition of two cobordisms. This concerns the fact that when operating the splicing described above it is convenient to extend the end L' as well as the secondary ends $L_1, L_2, \ldots$ in Figure 37. This type of operation needs to be implemented with some care as the specific aspect of the bottleneck needs to be preserved, even if the bottleneck itself is translated to the left or right. The argument is simple and makes use of a horizontal Hamiltonian isotopy as in Figure 38.

Such a Hamiltonian isotopy also acts on the relevant additional data, c, h, $\mathcal{D}_{V_h}$ and it preserves all the properties of the initial cobordism V_h.

With this understanding of the extension to the left of the ends $L_1, L_2, \ldots$ in Figure 37, the argument to show that the composition $V' \circ V$ has the required properties reduces to a couple of observations. First, due to a simple open mapping argument, no J-holomorphic polygons can pass from one side of P to the other (the two sides are to the right and left of the vertical axis through P): this is clear (as in [10]) if the polygon does not switch branches at P but, even if it does, such a curve would still be forced to pass through an unbounded quadrant at the point P, which is not possible. As a result of this remark, the relevant moduli spaces are regular.

It remains to check that the number of **c**-marked teardrops through any self-intersection point is 0 mod 2. The key remark here is that due to our assumption on the action-negativity of the marking **c**, a **c**-marked teardrop can only have an exit puncture that is action-negative. At the same time, in view of the positivity of the perturbation $h_{L'}$ at the action-negative points (see Assumption IV) it follows that all the teardrops with an exit point belonging to the fiber over the bottleneck P are completely included in this fiber. Thus they are teardrops with

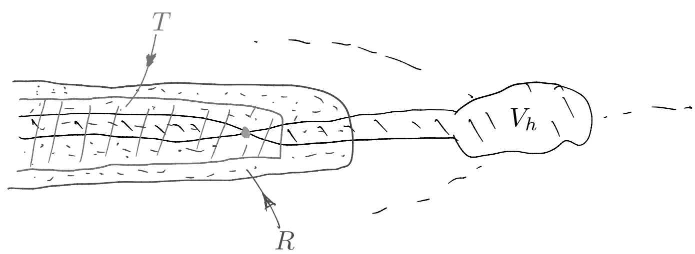

Figure 38. (Color online) The horizontal Hamiltonian isotopy that acts as a translation (to the left) inside the region T and is supported inside R.

boundary along L' and, as L' is unobstructed, their count mod2 vanishes. All the other teardrops are only associated either with V_h or with $V'_{h'}$ and their counts also vanish because V, V' are unobstructed.

3.4.4. Small perturbation germs and identity morphisms

To define the category $\mathsf{C}ob^*(M)$, some further modifications of the classes $Lag^*(M)$ and $Lag^*(\mathbb{C}\times M)$ are needed. In essence, they come down to a notion of "sufficiently small" perturbation germs h_L. Such a notion will be used to define identity morphisms in our category. Additionally, it will serve to view any two objects $(L, c, h_L, \mathcal{D}_L)$ and $(L, c, h'_L, \mathcal{D}_L)$ as equivalent as soon as h_L and h'_L are sufficiently small.

First, notice that the perturbation germs—as defined through the perturbations in Section 3.2.1—can be compared to the 0-perturbation in the C^k-norm ($k = 2$ is sufficient for us). Given two perturbation germs h_L and h'_L for the same marked, exact, Lagrangian $(L, c, \mathcal{D}_L)$, we may consider the immersed Lagrangian V_h drawn in Figure 39 below which is a perturbation of $\gamma \times L$ with $\gamma = \mathbb{R}$ such that h restricts to h_L close to the positive bottleneck and to h'_L close to the negative bottleneck. The definition of the actual perturbations is as in Section 3.2.1, Figure 33, but one uses instead of γ_+ there the upper blue curve exiting P in the positive direction, on top in Figure 39, and instead of γ_- the blue curve exiting P in the positive direction, at the bottom. In particular, there is no perturbation of $\gamma \times L$ away from some disks around the self-intersection points of L and, along $\mathbb{R}\times(D_- \cup D_+)$, where D_-, D_+ are, respectively, small disks around the components of a self-intersection point $(P_-, P_+) \in I_L^{<0} \subset L \times L$, V_h equals $\gamma_- \times D_- \cup \gamma_+ \times D_+$.

There is an *inverted* bottleneck at 0 (this is inverted in the sense that, this time, the points in $\{0\} \times I_L^{<0}$ are points where the perturbation h is negative; we will refer to the usual type of bottleneck as being *regular*). The marked points of V_h are over the positive and negative bottlenecks and they coincide there with the markings on L; the self-intersection points over $\{0\} \in \mathbb{C}$ are not marked. We consider an additional curve γ' as in the picture and a similarly defined

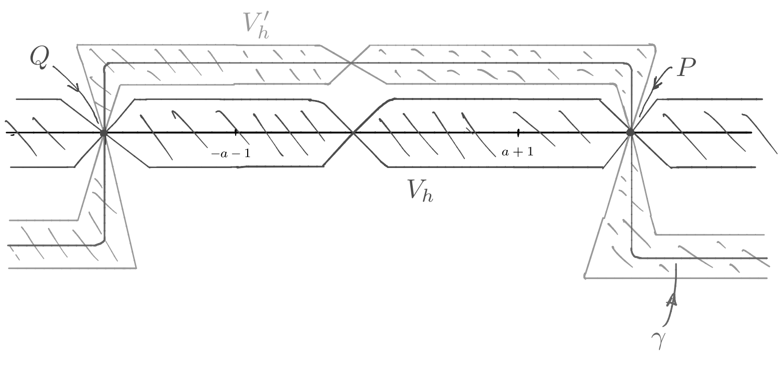

Figure 39. (Color online) Two small perturbations V_h and V'_h of $\mathbb{R} \times L$ and $\gamma' \times L$, respectively.

perturbation V'_h of $\gamma' \times L$ so that again the germs over the positive/negative bottlenecks are h_L and h'_L, respectively. The cobordism V'_h is obtained through a horizontal Hamiltonian isotopy (that is independent of L, h_L, h'_L) from the cobordism $V_h - i$ (this is simply V_h translated by $-i$). The bottlenecks of V'_h and V_h are positioned as in Figure 39.

We fix the perturbation data $\mathcal{D}_L$ on L (as well as the marking c). We consider the perturbation data $\mathcal{D}_{V_h}$ on V_h satisfying Assumptions I–VII and that is a small perturbation of the perturbation $\mathcal{D}_L$. "Smallness" here is understood in the sense that if h_L, h'_L tend to 0, then V_h tends to $\mathbb{R} \times L$ and this perturbation data tends to the product data $i \times D_L$. Moreover, the data $\mathcal{D}_{V_h}$ is defined in such a way that, on the boundary of polygons, the perturbation vanishes (i.e., the almost-complex structure is the product $i \times J_0$, and the Hamiltonian term is 0—as in Definition 3.1.2(ii)(a)). We define the data $\mathcal{D}_{V'_h}$ by appropriately transporting the data of V_h through a Hamiltonian isotopy induced from an isotopy in $\mathbb{C}$.

In view of the fact that there are no teardrops with boundary on L, we will see below that if V_h is sufficiently close to the product $\mathbb{R} \times L$ and the data $\mathcal{D}_{V_h}$ is sufficiently close to $i \times \mathcal{D}_L$, then V_h is unobstructed (and similarly for V'_h).

LEMMA 3.4.2

If V_h is sufficiently close to $\mathbb{R} \times L$ (and thus h_L is small too) and if $\mathcal{D}_{V_h}$ is sufficiently close to the product $i \times \mathcal{D}_L$, then V_h is unobstructed.

Proof

Let X denote the inverted bottleneck (at 0). Given that L is unobstructed, and in view of the shape of the curves γ_+, γ_-, only (marked) teardrops with boundary on V_h can have inputs in points either all in the fiber $\{Q\} \times M$ or all in $\{P\} \times M$ and the output has to be in a point in $\{X\} \times M$. A teardrop cannot pass from one side to the other of X (due to the fact that the perturbation of the data chosen here vanishes on the boundary of the perturbed J-holomorphic polygons). The key remark is that for V_h sufficiently close to $\mathbb{R} \times L$ and $\mathcal{D}_{V_h}$ sufficiently close to the product $i \times \mathcal{D}_L$, the only marked teardrops in 0-dimensional moduli spaces are of a very special type: they only have one input and one output ($|u| = 2$), with the input some point of the form $\{Q\} \times \{x\}$ or $\{P\} \times \{x\}$ and the output $\{X\} \times \{x\}$ with $x \in c$. If this were not the case, then there would exist marked teardrops u_n (with the same asymptotic conditions) for V_h approaching $\mathbb{R} \times L$ and data approaching $i \times \mathcal{D}_L$, and by taking the limit we would get a curve u whose projection v onto M would also be a marked teardrop with boundary on L. These curves are regular by the assumption on $\mathcal{D}_L$. But it is easy to see, due to the fact that X is an inverted bottleneck, that the dimension of the moduli space containing v drops by at least 1 compared to the dimension of the moduli spaces containing the curves u_n. Thus, such curves v do not exist, except if they are constant, in which case there is a single output and a single input and they coincide, which implies the claim.

We now fix V_h and $\mathcal{D}_{V_h}$ so that the only teardrops in 0-dimensional moduli spaces are as explained above. We claim that for each point $x \in c$ both possibilities occur, in the sense that both types of marked teardrops appear—the ones starting over P and the ones starting over Q—and in each case their number $(\bmod 2)$ is 1. This follows by a cobordism argument, based on the fact that for the product structure $i \times \mathcal{D}_L$ these two types of teardrops exist (they project to the point x on M), are regular, and there is precisely one starting over P and one starting over Q.

The conclusion is that (mod2) the number of teardrops at each such point $\{X\} \times \{x\}$ vanishes and this implies that V_h is unobstructed. $\square$

REMARK 3.4.3

This argument does not directly work over coefficients different from $\mathbb{Z}/2$. However, we expect an appropriate weight system associated to markings to allow one to deduce a similar result over more general coefficient rings.

We now select Floer data for the couple (V'_h, V_h), which is also close to $i \times \mathcal{D}_L$ in the sense above.

We consider the resulting chain map $\phi : CF(L, L) \to CF(L, L)$ given by counting elements in 0-dimensional moduli spaces of (marked) Floer strips with boundary on V'_h and V_h that go from Hamiltonian chords lying over the point P to chords in the fiber over Q (the definition of this map is standard in embedded cobordism type arguments as in [9], and the fact that it is a chain map is a simple application of the unobstructedness conditions). Keeping the curves γ and γ' fixed as well as the Lagrangian L, another application of a Gromov compactness argument shows that there exist V_h, V'_h possibly even closer to products, as well as data also possibly closer to $i \times \mathcal{D}_L$ such that the map ϕ is the identity.

We will call a cobordism V_h endowed with the data $\mathcal{D}_{V_h}$ as before a *sufficiently small* deformation of the cobordism $\gamma \times L$ if there exists V'_h as above (thus, in particular, the associated map ϕ is the identity). Similarly, the respective perturbation germs h_L (that appear at the ends of V_h) will be called *sufficiently small* perturbation germs.

In summary, the deformation V_h of the trivial cobordism $\gamma \times L$ is sufficiently small if it is unobstructed and it admits itself a deformation V'_h, positioned as in Figure 39, also unobstructed and such that the morphism ϕ relating the Floer homologies over P and Q is the identity.

3.4.5. Final definition of the objects and morphisms in $\mathsf{C}ob^*(M)$

We first adjust the classes $Lag^*(M)$ and $Lag^*(\mathbb{C} \times M)$ (that satisfy the Assumptions I–VII from Section 3.3.2) in accordance with the discussion in Section 3.4.4. Thus, we will only consider the unobstructed, marked Lagrangians $(L, c, \mathcal{D}_L, h_L) \in Lag^*(M)$ with h_L a sufficiently small perturbation germ, in the sense of Section 3.4.4. An additional important constraint is that we will assume that the base almost-complex structure included in the data $\mathcal{D}_L$ is the *same* J_0 for all Lagrangians L. This constraint is particularly important to be able to define cabling properly, in this context.

We denote the resulting class by $\mathcal{L}ag^*(M)$. In particular, all Lagrangians in $\mathcal{L}ag^*(M)$ are unobstructed with respect to the same base almost-complex structure J_0 on M, that we fix from now on (see Definition 3.1.2). In our setting, this means that there are no J_0-holomorphic (non-**c**-marked) teardrops with boundary on any $L \in \mathcal{L}ag^*(M)$ and that, for the perturbation data $\mathcal{D}_L$, the number of c-marked teardrops at any self-intersection point of L is 0.

The class of cobordisms $\mathcal{L}ag^*(\mathbb{C} \times M)$ will refer from now on to those cobordisms

$$(V, c, h, \mathcal{D}_{V_h}) \in Lag^*(\mathbb{C} \times M)$$

with ends among the elements of $\mathcal{L}ag^*(M)$. To recall notation, the immersions associated to these Lagrangians are denoted by j_L, j_V and the primitives are f_L, f_V.

Finally, we define the cobordism category $\mathsf{C}ob^*(M)$ with objects the elements in $\mathcal{L}ag^*(M)$. The morphisms are the elements in $\mathcal{L}ag^*(\mathbb{C} \times M)$ modulo an equivalence relation that identifies cobordisms up to:

(i) horizontal Hamiltonian isotopy (of course, these isotopies act not only on V, c, and h but, in an obvious way, on $\mathcal{D}_{V_h}$ as well),

(ii) any two small enough deformations V_L, V'_L, of the cobordism $\mathbb{R} \times L$, $\forall L$, are equivalent and, for any other cobordism V, the compositions $V \circ V_L$ and $V'_L \circ V$ (if defined) are both equivalent to V.

Property (ii) allows for the definition of an identity morphism for each object L of our category: this is the equivalence class of any sufficiently small deformation of the identity cobordism $\mathbb{R} \times L$.

We complete the set of elements of $\mathcal{L}ag^*(\mathbb{C} \times M)$ with respect to composition, as defined in Section 3.4.2. We denote by $\mathcal{L}ag_0^*(\mathbb{C} \times M)$ the class of objects in $\mathcal{L}ag^*(\mathbb{C} \times M)$ that do not carry teardrops (see also Section 3.3.3).

REMARK 3.4.4

In the definition of the composition it is possible to make choices such that the resulting composition operation defined on the morphisms of $\mathsf{C}ob^*(M)$ is associative, but we will not further detail this point here.

To simplify notation, we will continue to denote by $\mathcal{F}uk_i^*(M)$ and, respectively, $\mathcal{F}uk_i^*(\mathbb{C} \times M)$, the Fukaya categories associated to the classes $\mathcal{L}ag^*(M)$ and $\mathcal{L}ag_0^*(\mathbb{C} \times M)$ even if, properly speaking, these are full subcategories of the Fukaya categories introduced in Sections 3.3.1 and 3.3.3.

4. From geometry to algebra and back

4.1. The main result

With the preparation in Section 3 and, in particular, with the category $\mathsf{C}ob^*(M)$ defined in Section 3.4 and with the machinery of surgery models introduced in Section 2.2, the purpose of this section is to prove the main result, Theorem A. The constructions in this proof can be viewed as establishing a number of "lines" in the dictionary of *geometry* $\leftrightarrow$ *algebra* mentioned at the beginning of the article. The passage from geometry to algebra is, in essence, a direct extension of the results concerning Lagrangian cobordisms from [9] and [10].

We reformulate Theorem A in a more precise fashion as follows.

THEOREM 4.1.1

The category $\mathsf{C}ob^*(M)$, as defined in Section 3.4, has the following properties:

(i) *$\mathsf{C}ob^*(M)$ has rigid surgery models, where Axiom 4 on page 418 is understood here with 0-size surgeries (i.e., $\epsilon = 0$).*

(ii) *There exists a triangulated functor*

$$\widehat{\Theta} : \widehat{\mathsf{C}ob}^*(M) \to H\big(\mathrm{mod}\big(\mathcal{F}uk^*(M)\big)\big)$$

that restricts to a triangulated equivalence $\widehat{\Theta}_e : \widehat{\mathsf{C}ob}^*_e(M) \to D\mathcal{F}uk^*(M)$, where $\widehat{\mathsf{C}ob}^*_e(M)$ is the triangulated subcategory of $\widehat{\mathsf{C}ob}^*(M)$ generated by $\mathcal{L}ag^*_e(M)$.

(iii) *The category $\widehat{\mathsf{C}ob}^*(M)$ is equivalent, as an additive category, to the Donaldson category associated to the Lagrangians in $\mathcal{L}ag^*(M)$.*

(iv) *For any family $\mathcal{F} \subset \mathcal{L}ag^*_e(M)$, the functor $\widehat{\Theta}_e$ is noncontracting with respect to the shadow pseudometric $d^{\mathcal{F}}$ on $\widehat{\mathsf{C}ob}^*_e(M)$ and the pseudometric $s^{\mathcal{F}}_a$ on $D\mathcal{F}uk^*(M)$.*

REMARK 4.1.2

(a) Allowing Lagrangians to be immersed is necessary, in general, for surgery models to exist. However, our proof also makes essential use of the fact that we allow for *marked* immersed Lagrangians. It is possible that this extension to marked Lagrangians is not necessary for the construction of categories with surgery models, but, as we will explain later in the paper, restricting to non-marked Lagrangians adds significant complications. It is an interesting open question at this time whether these can be overcome.

(b) As mentioned before in Remark 3.1.5, the specific type of marked immersed Lagrangians that appear in our construction is a variant of the notion of an immersed Lagrangian endowed with a bounding chain in the sense of [28] and [2].

(c) We recall here that the triangulated structure on $\widehat{\mathsf{C}ob}^*(M)$ follows, through Theorem 2.3.1, from the fact that $\mathsf{C}ob^*(M)$ has surgery models.

The proof of the theorem is contained in the next subsections. Here is an outline of the argument. We first move from geometry to algebra. There are a couple of main steps in this direction. First, in Section 4.2.1 we show the existence of a functor

$$\Theta : \mathsf{C}ob^*(M) \to H\big(\mathrm{mod}\big(\mathcal{F}uk^*(M)\big)\big). \tag{17}$$

We then revisit (in our setting) the cabling construction from Section 2.1.4 and show in Section 4.2.2 that the cabling equivalence of two cobordisms V, V' in $\mathcal{L}ag^*(\mathbb{C} \times M)$ comes down to the equality $\Theta(V) = \Theta(V')$. We then proceed in the reverse direction, from algebra to geometry. We start by associating, in Section 4.3.1, to each module morphism ϕ an appropriate 0-surgery cobordism V such that $\Theta(V) \simeq \phi$. Then, in Section 4.3.2, we observe that, as a consequence of the properties of the functor Θ, $\mathsf{C}ob^*(M)$ has surgery models, and that Θ induces $\widehat{\Theta}$ with the properties in the statement. The various parts of the proof are put together in Section 4.4 where we also address the rigidity part of the statement.

4.2. From geometry to algebra

4.2.1. The functor Θ

The definition and properties of this functor are similar to those of the functor Θ constructed for embedded Lagrangians and cobordisms in [10]. In this subsection, we first define Θ and then proceed to describe some of its main properties. Namely, we give a homological expression for Θ in Lemmas 4.2.1 and 4.2.3. We then show in Lemma 4.2.5 that Θ behaves functorially. We then provide in Lemma 4.2.7 an alternative description of Θ involving Hamiltonian isotopies, at least for simple cobordisms, and deduce the behavior of Θ with respect to inverting simple cobordisms in Corollary 4.2.9. Finally, we discuss the impact of changes of perturbation data for a given geometric Lagrangian L in Lemma 4.2.11.

Recall from Section 3.3.3 that there is a category $\mathcal{F}uk_i^*(\mathbb{C} \times M)$ (we neglect the data $\bar{\mathcal{D}}$ in this writing) with objects the cobordisms in $\mathcal{L}ag_0^*(\mathbb{C} \times M)$.

The first step is to consider the so-called *inclusion functor* associated to a curve $\gamma : \mathbb{R} \to \mathbb{C}$ that is horizontal at infinity (we view such a curve as a cobordism between two points, embedded in $\mathbb{C}$ itself):

$$i_\gamma : \mathcal{F}uk^*(M) \to \mathcal{F}uk^*(\mathbb{C} \times M) \to \mathcal{F}uk_i^*(\mathbb{C} \times M).$$

This is defined as a composition, with the second map being the embedding of the category of embedded cobordisms in the category of immersed, marked cobordisms:

$$\mathcal{F}uk^*(\mathbb{C} \times M) \hookrightarrow \mathcal{F}uk_i^*(\mathbb{C} \times M).$$

The first map $\mathcal{F}uk^*(M) \to \mathcal{F}uk^*(\mathbb{C} \times M)$ is the inclusion functor constructed in [10] (and denoted there also by i_γ). To fix ideas, we recall that, on objects, $i_\gamma(L) = \gamma \times L$.

Thus, to a cobordism $V \in \mathcal{O}b(\mathcal{F}uk_i^*(\mathbb{C} \times M))$ we can associate a module over $\mathcal{F}uk^*(M)$ defined as the pullback $i_\gamma^*(\mathcal{Y}(V))$, where $\mathcal{Y}(V)$ is the Yoneda module of V. This module is invariant, up to quasi-isomorphism, with respect to the usual horizontal isotopies of the curve γ (horizontal isotopies of γ are assumed here to leave the curve γ horizontal outside of the region $[-a-1, a+1] \times \mathbb{R}$; see Figure 35).

The second step is to consider a particular curve γ that intersects $\pi(V_h)$ (recall that V_h is the deformation of the cobordism V) only over the bottlenecks (in particular, it is away from the points P_i in case V carries non-**c**-marked teardrops), at points of real coordinates $-a-2$ for the negative ends and $a+2$ for the positive ones, and give a description of the module $i_\gamma^*(\mathcal{Y}(V))$ as an iterated cone of the Yoneda modules of the ends. Assuming that the curve γ intersects the bottlenecks as in Figure 36 (i.e., the curve γ intersects the blue horizontal region only at the bottlenecks themselves), the arguments from [10] adjust to this case without difficulty. For instance, consider $V : L \rightsquigarrow (L_1, \ldots, L_k, L')$ and γ as in Figure 40 below, such that γ starts below the positive end then crosses it to pass over the nonhorizontal part of V and descends by crossing the negative ends $L', L_k, \ldots, L_{s+1}$, becoming again horizontal between the ends L_s and L_{s-1}. In this case, the iterated cone $i_\gamma^*(\mathcal{Y}_e(V))$ is quasi-isomorphic to

$$\begin{aligned}\text{Cone}\big(\mathcal{Y}_e(L) \to &\text{Cone}(\mathcal{Y}_e(L') \to \\ &\text{Cone}(\mathcal{Y}_e(L_k) \to \cdots \to \text{Cone}(\mathcal{Y}_e(L_{s+1}) \to \mathcal{Y}_e(L_s))\ldots))\big).\end{aligned} \tag{18}$$

The module morphisms corresponding to the arrows in (18) above are associated in a geometric fashion to the cobordism V. Here we have denoted by $\mathcal{Y}_e(N)$ the pullback to $\mathcal{F}uk^*(M)$ (the category of embedded objects) of the Yoneda module $\mathcal{Y}(N)$ which is defined over the category of immersed Lagrangians $\mathcal{F}uk_i^*(M)$.

Of course, the J-holomorphic curves being counted in the different structures considered here are marked. The use of positive action markings and the fact that the perturbation h is itself positive at these intersection points of the ends implies that precisely the correct Yoneda module is associated to the intersection of γ with each end of V_h.

To conclude, the construction of the inclusion functors i_γ and the description of the Yoneda modules $i_\gamma^*(\mathcal{Y}(V))$ carry over to the more general setting here without significant modifications when restricting to modules over the category of embedded objects. The definition of Θ is a particular application of this construction, as follows. On objects Θ associates to a Lagrangian $L \in \mathcal{L}ag^*(M)$ the Yoneda module $\mathcal{Y}_e(L)$. To define Θ on morphisms, consider first a cobordism $V \in \mathcal{L}ag_0^*(\mathbb{C} \times M)$ (thus V does not carry (nonmarked) teardrops) and a curve

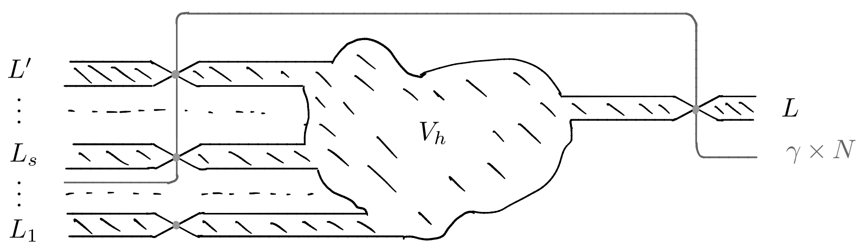

Figure 40. (Color online) Position of the curve γ relative to the (deformed) cobordism V_h.

γ similar to the one in Figure 40 such that γ only crosses the ends L and L'. In this case,

$$i_\gamma^*\big(\mathcal{Y}(V)\big) \simeq \mathrm{Cone}\big(\mathcal{Y}_e(L) \xrightarrow{\phi_V} \mathcal{Y}_e(L')\big) \tag{19}$$

and we put $\Theta(V) = [\phi_V] \in \mathrm{Mor}(H(\mathrm{mod}(\mathcal{F}uk^*(M))))$.

For the particular curve γ used to define Θ, the discussion above also applies (with obvious small modifications) to V's that carry nonmarked teardrops because γ avoids the pivots $P_i \in \mathbb{C}$ associated to V and plane-simple coherence (see Section 3.3.4) is easy to check in this case. The resulting module is invariant under horizontal homotopies of the curve γ that avoid the points P_i and continue to intersect V only along its two ends L and L'.

To show that Θ so defined is indeed a functor, we need to check that $\Theta(V)$ does not depend on the choice of curve γ (with the fixed behavior at the ends); that it is left invariant through horizontal Hamiltonian isotopies of V; that it associates the identity to each sufficiently small deformation of the identity (recall these notions from Section 3.4.4); and that it behaves functorially with respect to composition. The last two points are key. They follow from an alternative description of $\Theta(V)$ that will play an important role for us later on, and thus we will first give this description.

LEMMA 4.2.1

In homology, the morphism ϕ_V admits the following description. Let W be a sufficiently small deformation of $\gamma \times L$, where $\gamma \subset \mathbb{C}$ is as in Figure 41 below. There is an associated morphism $\phi_V^W : CF(L,L) \to CF(L,L')$ such that if we put $a_V^W := \phi_V^W([L])$, where $[L] \in HF(L,L)$ is the unit, then we have the identity

$$[\phi_V] = a_V^W. \tag{20}$$

Here we have used the identification $H_*(\mathrm{hom}_{\mathrm{mod}}(\mathcal{Y}_e(L), \mathcal{Y}_e(L'))) \cong HF(L,L')$. In other words, the morphism ϕ_V coincides with $\mu_2(-, a_V^W)$ in homology.

The morphism ϕ_V in (20) is defined through equation (19) and the identity (20) is written, as will be explained below, in the homological category of modules over the immersed Fukaya category $\mathcal{F}uk_i^*(M)$. Therefore, more precisely, equation (20) means $[\phi_V] = i_e^*[\mu_2(-, a_V^W)]$ where $i_e : \mathcal{F}uk^*(M) \to \mathcal{F}uk_i^*(M)$ is the obvious embedding.

Proof

Recall the existence of the immersed Fukaya categories $\mathcal{F}uk_i^*(M)$, $\mathcal{F}uk_i^*(\mathbb{C} \times M)$. We start by noting that there is an extension of the morphism ϕ_V as a module morphism $\bar{\phi}_V : \mathcal{Y}(L) \to \mathcal{Y}(L')$ over the category $\mathcal{F}uk_i^*(M)$. To define this more general $\bar{\phi}_V$, we consider for each Lagrangian $N \in \mathcal{L}ag^*(M)$ a certain sufficiently small deformation, W_N, of $\gamma \times N$, where γ is the fixed curve in Figure 41 (W is one such example for $N = L$). Once these choices are fixed, the same method as in [10] shows that this data leads to the extension of the morphism ϕ_V to

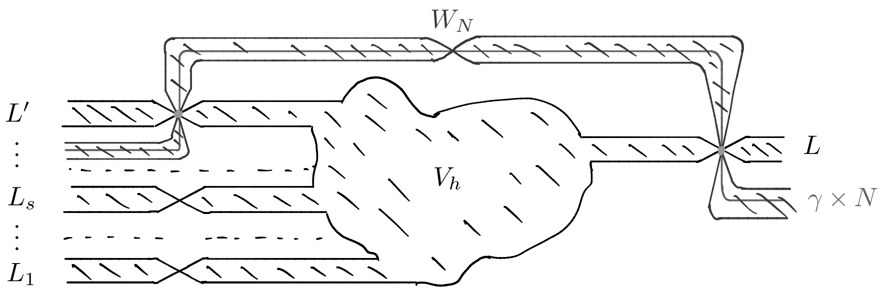

Figure 41. (Color online) The perturbed cobordism V_h and the small enough deformation W_N of $\gamma \times N$.

$\mathrm{mod}_{\mathcal{F}uk_i^*(M)}$. We denote this extension by $\bar{\phi}_V$. Because the category $\mathcal{F}uk_i^*(M)$ is homologically unital, it follows that there exists an element a_V^W as in the statement and this element is associated to $W = W_L$. In fact, it follows that formula (20) is true over $\mathcal{F}uk_i^*(M)$. □

REMARK 4.2.2

It is useful to notice that the extension $\bar{\phi}_V$ of ϕ_V depends on our choices of perturbations W_N and, implicitly, also depends on the choice of the curve γ. However, the properties of ϕ_V can be deduced from properties of any extension $\bar{\phi}_V$ because, up to a boundary, the morphisms ϕ_V are robust: they do not depend on the curve γ (assuming the correct behavior at the ends) or on perturbations.

We next consider a different curve γ' with the same ends at infinity as γ in Figure 41 and that intersects the bottlenecks of V_h in the same pattern as γ.

LEMMA 4.2.3

In the setting above, there is a sufficiently small deformation W' of $\gamma' \times L$ such that the associated element $a_V^{W'}$ is defined and $[a_V^W] = [a_V^{W'}]$.

The argument to prove this is to use a horizontal Hamiltonian deformation from γ to γ'. To use such a transformation to prove the statement of the lemma, we need to also move one of the bottlenecks along γ (because, in essence, the areas enclosed by the two curves γ and γ' between the bottlenecks are not the same). We also need to ensure unobstructedness of the Lagrangians involved. For this purpose, the perturbation associated to γ' (and along the transformation) might need to be taken very small compared to the perturbation associated to γ.

REMARK 4.2.4

A somewhat delicate point is that the size of the perturbation germs of L, as it appears at the two ends of $\gamma' \times L$, depends potentially on γ'. At the same time, the category $\mathcal{F}uk_i^*(M)$ does not "see" perturbation germs associated to its objects and, as mentioned in Remark 4.2.2, it follows that whenever we use various auxiliary curves γ, γ', sufficiently small deformations of associated trivial cobordisms are sufficient for our arguments even if they depend on the curves γ, γ'.

Returning to the properties of Θ, it follows from the definition of sufficiently small deformations of trivial cobordisms that the class $a_{V_h}^{V'_h}$ is the unit, and thus Θ applied to a trivial cobordism of the form $\gamma \times N$ is the identity. Showing that Θ respects the composition of cobordisms requires a slightly more involved argument.

LEMMA 4.2.5

Consider $V, V' \in \mathcal{L}ag^*(M)$ such that $V'' = V' \circ V$ is defined with the gluing over the end L'. Then $\Theta(V'') = \Theta(V') \circ \Theta(V)$.

Proof

We sketch V, V' and their composition V'' after deformation in Figure 42.

To fix ideas, we start under the assumption that neither V nor V' carry (nonmarked) teardrops. We also consider two curves γ and γ' (γ and γ' are homotopic in the complement of the marked points $P_i, P'_j \in \mathbb{C}$ associated to V, V' in case these carry nonmarked teardrops) as in the picture. We consider a sufficiently small deformation W of $\gamma \times L$ with regular bottlenecks at the points P, Q, R, T and with inverted bottlenecks at X, Y, Z. We assume that γ' is Hamiltonian isotopic to γ with the ends P and T fixed, and we also consider a sufficiently small deformation W' of $\gamma' \times L$ which is obtained by the associated Hamiltonian deformation of W. We claim that there exists such a γ and W (in particular, W is unobstructed) with the following additional properties: there is a complex C_Q with generators the intersection points $W \cap V'' \cap \pi^{-1}(U_Q)$, with U_Q as in the picture, and whose differential counts curves with an image that does not get out of U_Q. Under the translation bringing Q over the bottleneck H (that

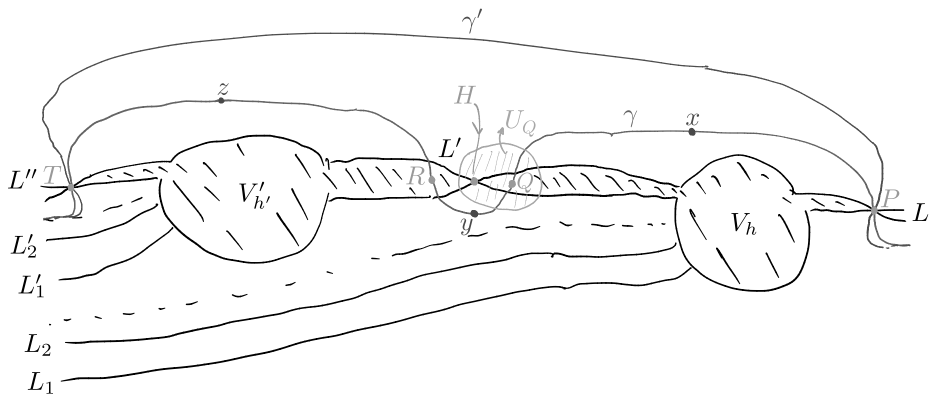

Figure 42. (Color online) The composition $V'' = V' \circ V$ after deformation and the curves γ and γ'. A sufficiently small deformation W of $\gamma \times L$ has regular bottlenecks at P, Q, R, T and inverted ones at X, Y, Z.

corresponds to the gluing of the end of V to the corresponding end of V') of V'' (by an isotopy coming from the plane), this complex is identified with $CF(L, L')$. There is a similar complex C_R associated to R with similar properties. The strips going from the elements in C_R to those in C_Q give a morphism homotopic to the identity. The existence of such a W results from a Gromov compactness argument applied by diminishing simultaneously the "loop" that γ does around H and the perturbation giving W. Using the identifications of the complexes C_R and C_Q with the complex $CF(L, L')$ over the bottleneck H, the complex $CF(W, V'')$ can be rewritten as a sum of four complexes corresponding to the four regular bottlenecks of W involving C_R, C_Q and complexes C_P, C_T associated in a similar way to the bottlenecks P and T. In this writing, the differential of $CF(W, V'')$ involves ϕ_V^W defined by strips going from C_P to C_Q, $\phi_{V'}^W$ defined by strips going from C_R to C_T, and a morphism homotopic to the identity from C_R to C_Q. As W' is Hamiltonian isotopic to W, we finally deduce that $\phi_{V''}^{W'} \simeq \phi_{V''}^W \simeq \phi_{V'}^W \circ \phi_V^W$ and, in view of Lemma 4.2.1, this ends the proof in case V and V' do not carry teardrops.

If V and V' carry teardrops, we notice that the curve γ is away from the pivots of both V and V', that γ and γ' are homotopic in the complement of the union of the pivots. Moreover, for the Floer complexes and maps that appear in this proof, plane-simple coherence is easy to check. Thus, it follows that the argument remains true in this case too, and this concludes the proof. □

REMARK 4.2.6

By using methods similar to those above, it is also possible to show, following [10], that the iterated cone-decomposition in (18) remains valid also in terms of modules over the immersed Fukaya category $\mathcal{F}uk_i^*(M)$ at least for $V \in \mathcal{L}ag_0^*(\mathbb{C} \times M)$.

To summarize what we have proved up to this point in the subsection, we have shown that $\Theta(V)$, as defined in (19), is characterized by the class $[a_V^W] \in HF(L, L')$ through equation (20). Moreover, it also follows that the definition of Θ is independent of the choice of the curves γ and that Θ is an actual functor, as it associates the identity to any small enough perturbation of a curve $\gamma \times L$. Finally, it follows that Θ does not "see" perturbation germs (as long as they are small enough).

We end the section with three other important features of the functor Θ.

Before going on, we simplify our drawing conventions to avoid drawing explicitly the projections onto the plane of the perturbations of the cylindrical ends of cobordisms.

Instead, from now on we will represent these configurations schematically by using the following conventions that apply to the cylindrical ends (or parts) of immersed cobordisms: we indicate the regular bottlenecks by a dot (such as P, Q, R, and T in Figure 42); we will no longer indicate the inverted bottlenecks (such as X, Y, Z in Figure 42) but we assume that a unique inverted bottleneck

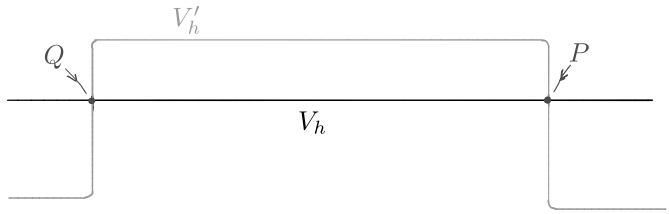

Figure 43. (Color online) The schematic version of Figure 39. The bottlenecks P and Q are regular for both V'_h and V_h.

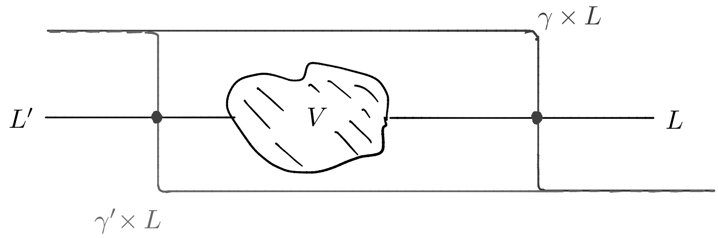

Figure 44. (Color online) The Lagrangians V, $\gamma \times L$ and $\gamma' \times L$ drawn without perturbations. The intersections at the bottlenecks drawn in red are as in Figure 39 (at the point P).

always exists between two regular ones; when two cylindrical components intersect at bottlenecks of both components, then the two bottlenecks are regular and positioned as in Figure 39; there could also be pivots in $\mathbb{C}$, away from the projection of the cobordisms considered, such as P_1, P_2 in Figure 34, that will be used to "stabilize" teardrops. Of course, more details will be given if needed. To give an example of our conventions, Figure 43 is the schematic version of Figure 39.

The next result provides an alternative description for $\Theta(V)$ when V is simple.

LEMMA 4.2.7

Assume that $V \in \mathcal{L}ag_0^*(\mathbb{C} \times M)$ (which means that V does not carry any J_0-teardrops) and is simple (which means that it has just two ends, one positive and one negative). Then $\Theta(V)$ is induced by the Hamiltonian translation moving a small enough perturbation of $\gamma \times L$ to $\gamma' \times L$ in Figure 44.

REMARK 4.2.8

There is a more general version of this statement when V is not simple, but we will not need it here.

Proof

The proof makes use of Figure 45 below. We consider the two cobordisms $W = \eta \times L$ and $W' = \eta' \times L$, where η, η' are the curves as in Figure 45. There is an obvious horizontal Hamiltonian diffeomorphism φ (induced from the plane

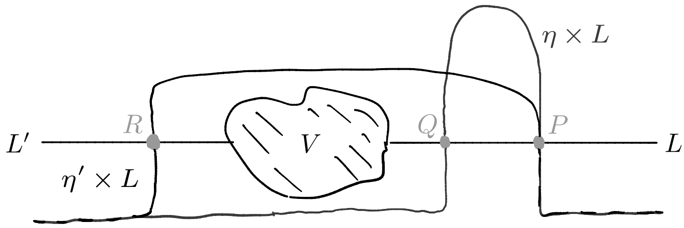

Figure 45. (Color online) Schematic representation of V, $W = \eta \times L$ and $W' = \eta' \times L$ obtained from W by the obvious Hamiltonian isotopy carrying the bottleneck P to Q (and maintaining P fixed).

$\mathbb{C}$) that carries W to W'. Put $C = CF(W, V)$ and $C' = CF(W', V)$, and deduce that there is a quasi-isomorphism $\psi : C \to C'$ induced by this Hamiltonian diffeomorphism. Of course, in the definition of these Floer complexes we need to use appropriate perturbations for W and W' (and we also consider a sufficiently small perturbation of V between Q and P). In view of this, C can be written as a cone over the identity morphism identifying the complex C_P over P to the complex over Q, C_Q, both being identified to $CF(L, L)$. The complex C' is the cone over the map $\phi_V^{W'}$ relating the complex C_P and $C_R = CF(L, L')$. The chain morphism ψ, induced by the Hamiltonian isotopy φ moving η to η' (that moves Q to R and keeps P fixed), restricts to the chain map $\psi_V^W : C_Q \to C_R$ that is easily seen to be homotopic to the chain map induced by the translation in the statement (to see this it is enough to notice that φ can be taken to be an appropriate translation moving Q to R). As ψ is a chain map, it follows from a simple algebraic calculation that $\phi_V^{W'}$ is chain homotopic to ψ_V^W, and this concludes the proof. □

Given a simple cobordism V, recall that $\bar{V}$ is the inverse cobordism, as defined in Section 2.1, by a 180° rotation in the plane. We can also transform the data associated to V in such a way as to get corresponding data for $\bar{V}$, thus allowing us to interpret $\bar{V}$ as an element of $\mathcal{L}ag^*(\mathbb{C} \times M)$. It is also possible to adjust the deformation of the ends, say, $h_{L'}$, so that $\bar{V} \circ V$ is defined. Using Lemma 4.2.7 and the functoriality of Θ, it is a simple exercise to deduce the following statement.

COROLLARY 4.2.9

If V is a simple cobordism $\in \mathcal{L}ag_0^*(\mathbb{C} \times M)$ and $\bar{V}$ is its inverse, then $\Theta(V) \circ \Theta(\bar{V}) = \mathrm{id}$.

REMARK 4.2.10

For simple cobordisms $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$ that might carry (nonmarked) teardrops, the same result is true, but the proof is more involved and is postponed to Lemma 4.3.2.

We now fix an object $\in \mathcal{L}ag^*(M)$. Thus, we are fixing a quadruple $(L, c, \mathcal{D}_L, h_L)$. Recall that L does not carry any J_0 (nonmarked) teardrops. By using Seidel's iterative process, it is easy to construct coherent perturbations with J_0 as a base almost complex structure, that make all moduli spaces with action-negative inputs (and output) regular (as described in Definition 3.1.2 and in Assumption I in Section 3.3.2). Let $\mathcal{D}'_L$ be such a choice of coherent perturbation data (the initial $\mathcal{D}_L$ is another such choice).

LEMMA 4.2.11

There exist a marking c', a perturbation h'_L, and a simple unobstructed cobordism $V \in \mathcal{L}ag_0^*(\mathbb{C} \times M)$ with ends $(L, c, \mathcal{D}_L, h_L)$ and $(L, c', \mathcal{D}'_L, h'_L)$. Geometrically, V is a small deformation of the trivial cobordism $\mathbb{R} \times L$ with a projection such as V_h in Figure 39.

Proof

We will make use of Figure 46 below. In short, we consider a deformation W of the Lagrangian immersion $\mathbb{R} \times L$ with three bottlenecks, P and Q regular and one inverted, X, between P and Q. The element V in the proof will be of the form $V = (\mathbb{R} \times L, \mathbf{c}, \overline{\mathcal{D}}, h_V)$, and we take $W = V_{h_V}$.

We first discuss the data $\overline{\mathcal{D}}$ associated to W. It is a coherent choice of perturbations such that $\overline{\mathcal{D}}$ restricts to $\mathcal{D}_L$ at P and to $\mathcal{D}'_L$ at both X and Q. Moreover, over the portion of the cobordism going from X to Q, $(W, \overline{\mathcal{D}})$ is a small deformation of $(\mathbb{R} \times L, i \times \mathcal{D}'_L)$. Such $\overline{\mathcal{D}}$ exists by the usual construction of coherent perturbations because the base almost-complex structure is the same J_0 for both $\mathcal{D}_L$ and $\mathcal{D}'_L$ and there are no J_0 (nonmarked) teardrops. The marking of W coincides with that of $(L, c, \mathcal{D}_L, h_L)$ over P and there are no markings over the bottleneck X. The purpose of the proof is to show that there is a marking c' over Q that consists only of self-intersection points that are action-negative and such that the marking $\mathbf{c} = (\{P\} \times c) \cup (\{Q\} \times c')$ makes V unobstructed.

We denote by $J_X = \{X\} \times I_L$ the self-intersection points of L viewed in the fiber over the bottleneck X, $J_X \subset \{X\} \times L \subset \{X\} \times M$. We let $\mathbb{Z}/2 < J_X >$ be the $\mathbb{Z}/2$-vector space of basis the elements of J_X and consider the element $z \in \mathbb{Z}/2\langle J_X \rangle$ given by $z = \sum_{y \in I_X} \#_2 \overline{\mathcal{M}}_c(\emptyset; y) y$. Here $\overline{\mathcal{M}}_c(\emptyset; y)$ is the moduli space of c-marked teardrops with boundary on W, defined with respect to the data $\overline{\mathcal{D}}$ and with output at y. As c is the marking over P, all these teardrops have

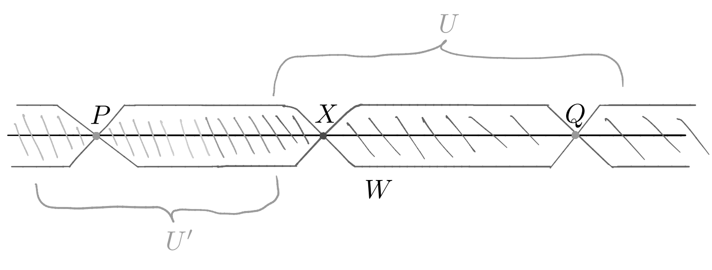

Figure 46. (Color online) The Lagrangian W. The data $\overline{\mathcal{D}}$ is a small deformation of the identity in the region U and is an interpolation between the data $\mathcal{D}_L$ and $\mathcal{D}'_L$ in the region U'.

marked inputs over P. Of course, only the elements in those moduli spaces that are 0-dimensional are counted (with coefficients in $\mathbb{Z}_2$). By taking the deformation W sufficiently close to $\mathbb{R} \times L$ and considering the action inequalities as in Remark 3.3.3(c), we see that whenever $\overline{\mathcal{M}}_c(\emptyset; y)$ is not void the point y is action-negative. We now define c' to be the union of all the self-intersection points $y \in I_L$ that appear in z and $\mathbf{c} = (\{P\} \times c) \cup (\{Q\} \times c')$. Because W is a sufficiently small deformation of the identity over the segment from X to Q, we deduce that at each point $y \in J_X$ the number of $\mathbf{c}$-marked teardrops is now 0, as the number of teardrops coming from P agrees with the number (mod2) of teardrops coming from Q (see also the proof of Lemma 3.4.2). It follows that V is unobstructed and this concludes the proof. □

4.2.2. Cabling and counting teardrops

In this section, we revisit the cabling construction from Section 2.1.4 with the aim of showing the following.

PROPOSITION 4.2.12

Let $V, V' \in \mathrm{Mor}_{\mathcal{C}ob^*(M)}(L, L')$. The two morphisms V and V' are cabling-equivalent (see Section 2.1.4 and Definition 2.1.9) if and only if $\Theta(V) = \Theta(V')$.

Along the way to proving this result, we will also need to adjust the cabling construction to our setting in a way that takes into account the deformations of the immersed cobordisms V and V'.

Proof

Using Lemma 4.2.1, we see that it is sufficient to show that V and V' are cabling-equivalent if and only if $[a_V^W] = [a_{V'}^W]$, where W is a small enough deformation of $\gamma \times L$ for an appropriate curve γ. We show this in three steps.

Step 1. Cabling taking into account markings and perturbations

We now consider two cobordisms $V, V' \in \mathrm{Mor}_{\mathcal{C}ob^*(M)}(L, L')$. We define the cabling of V and V', $V'' = \mathcal{C}(V, V'; c)$, by revisiting the construction in Section 2.1.4 and, in particular, Figure 12. We consider the Lagrangian represented in Figure 47 below (the components, V and V', appear in Figure 11). To fix ideas, we will first assume that $V, V' \in \mathcal{L}ag_0^*(M)$ and that L and L' intersect transversely.

We now list details of the construction. The data for V is $(V, h, c, \mathcal{D}_{V_h})$ and for V' it is $(V', h', c', \mathcal{D}_{V'_{h'}})$, and we have the two primitives f_V and $f_{V'}$ (that are viewed as defined on V_h and $V'_{h'}$).

(i) The bottlenecks R, S and R', S' are associated to the respective ends of V and V', respectively, with the structure induced from the deformed cobordisms V_h and $V'_{h'}$.

(ii) Postponing for the moment discussion of the relevant perturbation data, the segments SP, RP (resp., $S'P$, $R'P$) represent sufficiently small deformations

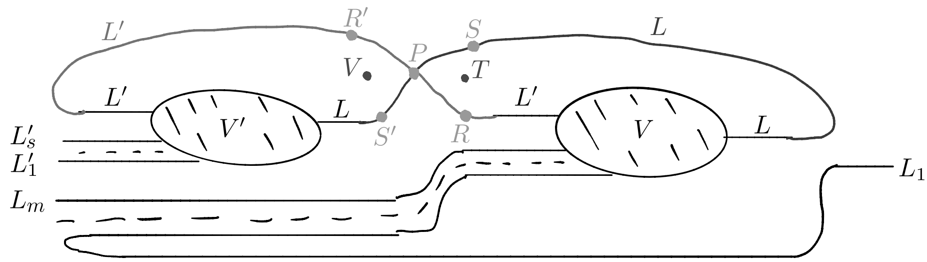

Figure 47. (Color online) Schematic representation of cabling, in the immersed perturbed, marked, exact setting. The points R, P, S, S', R' are regular bottlenecks, U and T are points in $\mathbb{C}$ used to stabilize teardrops.

of products of the form $A \times L$ (resp., $A \times L'$) with A the respective segment along the curves in green and blue in Figure 47.

(iii) The profile of the gluing at P for each of the two branches—one associated to L' with a restriction at P written as $\{P_+\} \times L'$ and the other associated to L with a restriction at P written as $\{P_-\} \times L$—is as in Figure 37.

(iv) The points U and T are pivots associated to V''.

(v) The primitives f_V and $f_{V'}$ are extended to a primitive $f_{V''}$. This is defined on the domain of the immersion and the "sizes" of the blue and green curves in Figure 47 might need to be adjusted to ensure the existence of the primitive $f_{V''}$ (by varying the areas in $\mathbb{C}\backslash\pi(V'')$ containing U and T, respectively). The restriction of $f_{V''}$ to the submanifold $\{P_+\} \times L'$ is denoted by $f_{V''}|_{L',P}$ and similarly for the restriction to $\{P_-\} \times L$ which is denoted by $f_{V''}|_{L,P}$. The curves in blue and green in Figure 47 are picked in such a way that $\min[f_{V''}|_{L,P}] \geq \max[f_{V''}|_{L',P}]$.

We further discuss the markings c'' and data $\mathcal{D}_{V''}$ associated to the cabling V''. To proceed, we first fix the perturbation data associated to the ends we are interested in. Thus, the data for L is $(L, h_L, c_1, \mathcal{D}_L)$ and that for L' is $(L', h_{L'}, c'_1, \mathcal{D}_{L'})$ with c_1, c'_1 and the rest of the data choices constrained as described below. A key point here is that for the perturbation data for the cabling to make sense it has to be chosen in a coherent way for both L and for L' when viewed as Lagrangians $\subset \{P\} \times M$. To achieve this, we will make use of Lemma 4.2.11. We first pick perturbation data $\mathcal{D}'$ that is coherent (and unobstructed) for both L and L'. By Lemma 4.2.11, there are unobstructed cobordisms V_1 from $(L, h_L, c_1, \mathcal{D}_L)$ to $(L, h_L, c_2, \mathcal{D}')$ and similarly V'_1 from $(L', h_{L'}, c'_1, \mathcal{D}_{L'})$ to $(L', h_{L'}, c'_2, \mathcal{D}')$. Here c_2, c'_2 are markings determined as in Lemma 4.2.11. We pick the perturbations in question here such that the cobordisms $\overline{V_1} \circ V_1$ and $\overline{V'_1} \circ V'_1$ are defined (see Corollary 4.2.9).

With these cobordisms V_1, V'_1 given with the associated perturbation data, the additional conditions on the cabling V'' are as follows:

(vi) The data $\mathcal{D}_{\mathcal{V}''}$ of V'' restricted to V', viewed from a neighborhood of S' and to a neighborhood of R' coincides with the data of V', including the markings. Similarly, the data of V'' restricted to V viewed from a neighborhood of S to a neighborhood of R coincides with the data of V.

(vii) The cobordism V'' together with its perturbation data restricts to V_1' along the segment RP and to the cobordism $\overline{V_1'}$ along PR'. Similarly, V'' restricts to V_1 along SP and to $\overline{V_1}$ along RS'.

(viii) The marking c'' of V'' contains the markings c and c' for V and V' as described in point (vi), the markings $\{P\} \times c_2$ and $\{P\} \times c_2'$ over P together with a set (possibly void) of new points of the form $((\{P_-\} \times \{x\}), (\{P_+\} \times \{y\}))$, where $(x, y) \in L \times L'$ such that $j_L(x) = j_{L'}(y)$. (We postpone the precise choice of these additional new points until Step 3 of the proof.)

(ix) The pivots of V'' are the union of the pivots of V and V' and of U and T. Finally, the data for V'' is coherent and regular also with respect to the pivots U, T.

Two cobordisms $V, V' \in \mathcal{L}ag^*(\mathbb{C} \times M)$ are, by definition, cabling-equivalent if there exists a cobordism V'', with a choice of coherent perturbation data as described above, such that V'' is unobstructed.

Step 2. Teardrops and the classes a_V^W

The purpose of this paragraph is to show that the homology class $[a_V^W]$ (from Lemma 4.2.1) can be determined by counting the teardrops in Figure 47 with output at P and with one interior marked point that is sent to T. More precisely, let $\mathcal{M}^T_{V'',c''}(\emptyset; y)$ be the moduli space of marked teardrops with one interior marked point being sent to T and output at the point $y \in \{P\} \times (L \cap L')$. Notice that the condition on the interior marked point implies that the marking of the potential entries of such a teardrop are only at points in $c \subset c''$. Consider the element $w = \sum_y \#_2 \mathcal{M}^T_{V'',c''}(\emptyset; y) y \in CF(L, L'; \mathcal{D}')$ (again the count is only for those moduli spaces of dimension 0). We claim that w is a cycle and that $[w]$ is identified (in a precise sense to be made explicit) with $[a_V^W]$. We will make use of Figure 48 below.

In this figure appears a sufficiently small perturbation W of $\gamma \times L$. This has again three regular bottlenecks at the points R', S, and R (as well as two inverted bottlenecks that are no longer pictured). By considering the 1-dimensional moduli

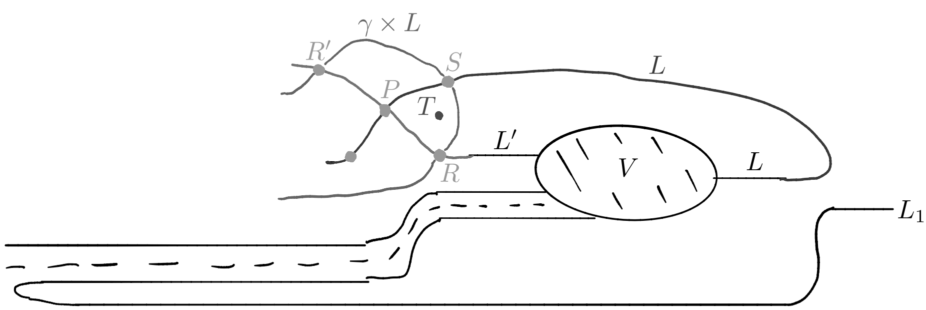

Figure 48. (Color online) Comparing teardrops and $\Theta(V)$ by using a small perturbation W of $\gamma \times L$.

spaces of the form $\mathcal{M}^{T}_{V'',c''}(\emptyset; y)$ and their boundary, it is easy to see that w is indeed a cycle. We then consider the 1-dimensional moduli space of marked Floer strips with one interior marked point (again sent to T) that have an input at the point S and with output at R', and with the $\{0\} \times \mathbb{R}$ part of the boundary along W and the $\{1\} \times \mathbb{R}$ part of their boundary along V''. From the structure of the boundary of this moduli space we see that there are two maps $\Psi_1 = \phi_1 \circ \phi^{W}_{V}$ and $\Psi_2 = \mu_2(-, w)$ that are chain homotopic. Here ϕ_1 is the chain map given by the marked Floer strips (with one interior marked point that is sent to T) that exit the point R and enter in R'. The operation μ_2 is given by (marked) triangles with entries at S and at P and output at R'. As W is a small deformation of $\gamma \times L$ and V'' is the composition of two inverse cobordisms over the interval RR', it follows that ϕ_1 is chain homotopic to the identity. We now apply Ψ_1 and Ψ_2 to the unit $[L] \in CF(L, L)$ viewed in the (marked) Floer complex in the fiber over S. We obtain $[\mu^2([L], w)] = [a^{W}_{V}]$. This means that $\eta_{V''}([a^{W}_{V}]) = [w]$, where η is the isomorphism between the complexes $CF(L, L')_R$ and $CF(L, L')_P$ induced by the cobordism V'' restricted to the segment RP, changing the data $\mathcal{D}_L$ to $\mathcal{D}'$ (as described in Lemma 4.2.11).

Denote now by $w' = \sum_y \#_2 \mathcal{M}^{U}_{V'',c''}(\emptyset; y) y \in CF(L, L'; \mathcal{D}')$ the cycle defined by counting marked teardrops with an interior marked point being sent to U (see Figure 47). It follows, from an argument similar to the one above, that $\eta'_{V''}([a^{W}_{V'}]) = [w']$, where $\eta'_{V''}$ is the identification using (the restriction of V'' to) the interval $S'P$.

Step 3. Conclusion of the proof

We start by remarking that the construction of the cabling V'' at Step 1, together with the coherent, regular perturbation data $\mathcal{D}_{V''}$ and the marking c'' is possible by the usual scheme. It is useful to note here that, as we are only using teardrops with interior marked points, we may pick coherent perturbations inductively, following the total valence $\|u\|$ of the curves u involved. In particular, the cycles w, w' are well defined as in Step 2. Notice that these cycles do not depend on the part of the marking c'' given by self-intersection points of V'' that belong to $\{P\} \times (L \cap L')$. We denote this part of the marking c'' by c''_P.

We now remark that there is a choice of c''_P making V'' unobstructed if and only if $[w] = [w']$. This shows the statement, as it implies by Step 2 that $\Theta(V) = \Theta(V')$ if and only if V and V' are cabling-equivalent.

In turn, the cobordism V'' (together with the data defined before) is unobstructed if and only if for any self-intersection point y the number (mod2) of marked teardrops with output at y is 0. Given that both V and V' are unobstructed, and in view of the presence of the bottleneck at the point P, the only points y where the number of teardrops might not vanish are the points $y \in \{P\} \times (L \cap L')$. We now consider such a point y and the moduli spaces of teardrops that have y as an output. A marked teardrop u with output at y has a projection onto the plane that is holomorphic in a neighborhood of P. As a result, this teardrop could be either completely contained in $\{P\} \times M$, or if its projection $v = \pi \circ u$ onto the plane is not constant, then the image of v is either to the left or to the right of P since otherwise the image of the curve v would enter one of the unbounded quadrants at the point P. Moreover, v is simple in a neighborhood of P. This is seen as follows. First, because u is a marked teardrop and in view of the action condition on the points belonging to the marking, none of the entry puncture points of u can belong to $\{P\} \times M$ as otherwise the image of v would again intersect one of the unbounded quadrants at P (if an output lies over P, then the curve v will "get out" of P through one of the unbounded quadrants). Second, due to asymptotic convergence at the exit boundary point, the restriction of v to some neighborhood of the exit puncture is simple and, again because the image of v cannot intersect one of the unbounded quadrants at P, it follows that v is simple over a possibly smaller neighborhood of the exit. Thus, we deduce that u is simple in a neighborhood of y.

We conclude that either $\pi \circ u$ is constant equal to P or u belongs to one of $\mathcal{M}^U_{V'',c''}(\emptyset; y)$ or $\mathcal{M}^T_{V'',c''}(\emptyset; y)$. Denote by $\mathcal{M}_{L\cup L',c'''}(\emptyset; y)$ the set of teardrops u with a constant projection. Here c''' is the marking $c''' = c_L \cup c_{L'} \cup c''_P$. Notice that we can write $\#_2\mathcal{M}_{L\cup L',c'''}(\emptyset; y) = \langle d_{L,L'}(c''_P), y\rangle$, where $d_{L,L'}$ is the differential of the Floer complex $CF(L, L')$. Of course, we work here—as before—with the Floer complex with a differential counting marked strips, with the marking c_L for L and $c_{L'}$ for L'. We conclude, therefore, that the total number $(\bmod 2)$ of marked teardrops at the point y is $n_y = \langle w, y\rangle - \langle w', y\rangle - \langle d_{L,L'}(c''_P), y\rangle$. Unobstructedness of V'' is equivalent to the vanishing of all these numbers and thus to the equality $w - w' = d_{L,L'}c''_P$. This means that if V'' is unobstructed, then $[w] = [w']$. Conversely, if $[w] = [w']$, then we can pick η to be such that $d_{L,L'}\eta = w - w'$ and then the cabling V'' defined as above and with $c''_P = \eta$ is unobstructed. □

REMARK 4.2.13

The fact that we work over $\mathbb{Z}/2$ is important in these calculations. It is clear that to work over $\mathbb{Z}$ one needs not only to fix spin structures (so that the various moduli spaces are oriented), but also to allow 0-size surgery with weights (essentially allowing integral multiples of the same intersection point) as well as redefine the marked Floer complex to take into account these weights.

Proposition 4.2.12 and the other properties of the functor Θ from Section 4.2.1 show that cabling is indeed an equivalence relation and that Axioms 1, 2, and 3 from Section 2.2 are satisfied.

4.3. From algebra to geometry

4.3.1. Surgery over cycles and geometrization of module morphisms

The main purpose of this subsection is to show that Axiom 4 from Section 2.2 is satisfied in our context. In other words, we need to show that for any morphism $V \in \mathrm{Mor}_{\mathsf{C}ob^*}(L, L')$ there exists a surgery morphism $S_V \in \mathcal{L}ag^*(\mathbb{C} \times M)$ that is cabling equivalent to V. By Proposition 4.2.12, this is equivalent to finding S_V such that $\Theta(S_V) = \Theta(V)$. This property is a consequence of the next result.

PROPOSITION 4.3.1

Let $L, L' \in \mathcal{L}ag^*(M)$, and let $a \in CF(L,L')$ be a cycle. Then there exists a marked cobordism $S_{L,L';a} : L \rightsquigarrow L'$, $S_{L,L';a} \in \mathcal{L}ag_0^*(\mathbb{C}\times M)$ such that

$$\Theta(S_{L,L';a}) = [a] \in HF(L,L').$$

Proof

It is not hard to construct a cobordism with double points containing 1-dimensional clean intersections, that represents the trace of a 0-size surgery: we can simply put $S_{L,L';a} = \gamma \times L \cup \gamma' \times L'$, where γ, γ' are the curves from Figure 49 (see also Section 2.1.2). This cobordism has three ends (indicated here with their markings): (L,c), (L',c'), and $(L\cup L', c\cup c'\cup a)$. The difficulty is to construct a deformation S_a of this $S_{L,L';a}$ together with the associated marking of S_a such that this marking restricts to the markings of the ends (as in Definition 3.2.2) and such that S_a is unobstructed, so that, together with the relevant perturbation data, we obtain an element in $\mathcal{L}ag_0^*(\mathbb{C}\times M)$.

We start with the two Lagrangians L and L' viewed as elements in $\mathcal{L}ag^*(M)$ and thus endowed with the additional data $L = (L, c, \mathcal{D}_L, h_L)$, $L' = (L', c', \mathcal{D}_{L'}, h_{L'})$. Additionally, the primitives f_L and $f_{L'}$ are also fixed. In view of Corollary 4.2.9 and Lemma 4.2.11, we may assume that L and L' are intersecting transversely at points different from the self-intersection points of both L and L' and that there is a choice of coherent data $\mathcal{D}''$ associated to the immersed Lagrangian $L\cup L'$ that extends both $\mathcal{D}_L$ and $\mathcal{D}_{L'}$. To fix ideas, we indicate from the start that the cobordism that we will construct will consist of the union of two sufficiently small deformations V of $\gamma\times L$ and V' of $\gamma'\times L'$, with γ and γ' two curves to be described below in Figure 49, together with additional marking points associated to a. In other words, S_a is a Lagrangian immersion with domain the disjoint union $V \coprod V'$ and $j_{S_a} = j_V \cup j_{V'}$. The primitive on S_a restricts to f_V on the V factor and to $f_{V'}$ on the V' factor.

Notice that the curves γ, γ' themselves are tangent at the point P and they coincide along the bottom leg in the picture, to the left of P. The points Q, P, R, and T are regular bottlenecks and there are inverted bottlenecks X, Y, Z. The primitives f_V and $f_{V'}$ restrict to the primitives f_L at Q and $f_{L'}$ at R and, by possibly having γ make an "upper bump" between Q and P, we may assume that the primitive f_V restricts at P along $L\subset \{P\}\times M$ to a primitive on L, $f_V|_{L,P}$, such that $\min(f_V|_{L,P}) > \max(f_{V'}|_{L',P})$, where $f_{V'}|_{L',P}$ is the corresponding restriction of $f_{V'}$.

There are four "regimes" governing the perturbations providing S_a and we describe them now. Let $D_i \subset L$ be small disks around the self-intersection points of L, and let $D'_i \subset L'$ be small disks around the self-intersections points of L' (in each case we view an intersection point as a pair $(P_-, P_+) \in I_L \subset L\times L$ and similarly for L'). Also, let K_i, K'_i be small disks around the intersection points of $x_i \in j_L(L)\cap j_{L'}(L')$ with $K_i \subset L$, $K'_i \subset L'$. The first regime is unperturbed: one portion of S_a, $(S_a)_1$, is given as

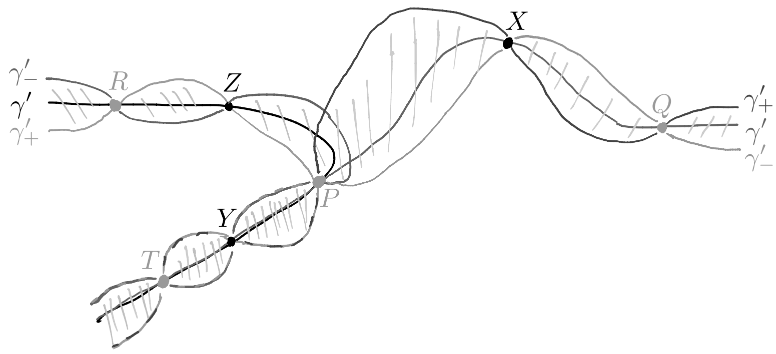

Figure 49. (Color online) The two cobordisms V—in blue—and V'—in gray—together with the (regular) bottlenecks P, Q, T, R.

$$(S_a)_1 = \Big(\gamma \times \Big(L \setminus \Big\{\bigcup_i K_i \cup \bigcup_j D_j\Big\}\Big)\Big) \cup \Big(\gamma' \times \Big(L' \setminus \Big\{\bigcup_i K'_i \cup \bigcup_k D'_k\Big\}\Big)\Big).$$

The second portion $(S_a)_2$ contains a perturbation around the self-intersection points of L in the sense of the small perturbation of the cobordism $\gamma \times L$ as in Section 3.4.4. In other words, if D_i, D_j are two disks with $P_- \in D_i$, $P_+ \in D_j$, $(P_-, P_+) \in I_L^{<0}$, then

$$(S_a)_2 = \gamma_- \times D_i \cup \gamma_+ \times D_j.$$

Similarly, there is another portion $(S_a)'_2$ that corresponds to a similar small perturbation of $\gamma' \times L'$.

Finally, there is yet one other portion, $(S_a)_3$, which is given in terms of the disks K_i, K'_i by

$$(S_a)_3 = \gamma_- \times K_i \cup \gamma'_+ \times K'_i.$$

These four parts of S_a are glued together by using perturbations on slighter larger disks $\tilde{D}_i$ such that $D_i \subset \tilde{D}_i$ (and similarly for D'_j, K_r, and K'_r). This is possible because the curves γ_-, γ_+ are Hamiltonian isotopic in the plane to γ (and similarly for the $(-)'$ curves). This Hamiltonian isotopy can be lifted to $\mathbb{C} \times M$ and the gluing uses an interpolation between this Hamiltonian defined on $\gamma \times D_i$ and the trivial one defined on the exterior of $\gamma \times \tilde{D}_i$.

The next point is to discuss the marking of S_a and the relevant data $\mathcal{D}_{S_a}$. The markings appear at the points R, Q, P, and T and are as follows. At R they are given by c', at Q by c, at P by $c \cup c' \cup a$, and at T by $c \cup c' \cup a$. Notice that this makes sense as a consists of pairs $(P_-, P_+) \in L \times L'$ such that $j_L(P_-) = j_{L'}(P_+)$. It follows that for the pairs as above we have $f_{S_a}(P_-) > f_{S_a}(P_+)$ (when all the relevant perturbations are sufficiently small). The data $\mathcal{D}_{S_a}$ restricts to the data $\mathcal{D}''$ in neighborhoods of the four regular bottlenecks Q, R, P, T. Away from the bottlenecks the data is a small deformation of $i \times \mathcal{D}''$. The deformations giving V, V' as well as $\mathcal{D}_{S_a}$ are taken sufficiently small such that with this data both

V and V' are unobstructed. Notice that S_a is a deformation of the immersed cobordism $S_{L,L';a}$, as desired.

We now want to show that, after possibly diminishing these perturbations even more, $L'' = (L \cup L', c \cup c' \cup a, \mathcal{D}'')$ (which is viewed as the end over the bottleneck T) is unobstructed and that the cobordism S_a is also unobstructed. We start with L''. Because both L and L' are unobstructed, it is immediate to see that the only marked teardrops with boundary on L'' and that are not purely on L or on L', have to have some marked inputs $\in a$. Moreover, because the markings are action-negative and because $\min(f_V|_{P,L}) > \max(f_{V'}|_{P,L'})$, we also have the same type of inequality at T. This implies that a marked teardrop with boundary on L'' can have at most one marked point $\in a$ as an input (to have more such points, the boundary of the teardrop should cross back from L' to L at some other input, but that input cannot belong to a due to the action inequality). In this case, obviously, the output is some intersection point $y \in L \cap L'$. We now consider the moduli space of such teardrops at y, $\mathcal{M}_{c\cup c'\cup a;L\cup L'},(\emptyset, y)$. It is immediate to see that this moduli space can be written as a union $\bigcup_{x\in a} \mathcal{M}_{\mathbf{c};L,L'}(x, y)$, where $\mathbf{c}$ are precisely the markings c for L and c' for L'. In other words, $\mathcal{M}_{\mathbf{c};L,L'}(x, y)$ is the moduli space of marked Floer strips with boundaries on L and L', originating at x and ending at y. Considering now only the 0-dimensional such moduli spaces and recalling that the data $\mathcal{D}''$ ensures regularity, we have

$$\#_2\mathcal{M}_{c\cup c'\cup a;L\cup L'}(\emptyset, y) = \sum_{x\in a} \#_2\mathcal{M}_{\mathbf{c};L,L'}(x, y) = \langle da, y\rangle = 0$$

because, by hypothesis, a is a cycle (here $d : CF(L, L') \to CF(L, L')$ is the (marked) Floer differential). Therefore, we conclude that L'' is unobstructed.

We proceed by analyzing S_a. We again start by using the fact that V and V' are each unobstructed. Therefore, the only type of teardrops with boundary on S_a that we need to worry about are those having a marked input $\in \{P\} \times a$ or $\in \{T\} \times a$ and output in $\{Y\} \times L''$. But we may assume that the cobordism S_a restricted to the region between P and T is in fact a sufficiently small deformation of $\gamma_{P,T} \times L''$, where $\gamma_{P,T}$ is the part of γ between P and T (it agrees there with γ'). This implies that this region of S_a is also unobstructed (by Lemma 3.4.2) and shows that S_a itself is unobstructed.

To conclude the proof of the proposition, we still need to show that $\Theta(S_a) = [a]$. For this argument we will use Figure 50 and notice that, by considering a sufficiently small deformation of $\gamma'' \times L$ (as in the figure) and the corresponding data, the morphism $\phi^W_{S_a}$ is given by counting marked strips from Q to R in Figure 50. There has to be only one such marking passing from V to V' and the only possibility for such a marking is at the point P. Thus this marking has to belong to $\{P\} \times \{a\}$. In other words, we have $\phi^W_{S_a}([L]) = \mu_2([L], a)$, where μ_2 is the multiplication for the following three Lagrangian cobordisms (W, V, V') given by counting triangles such as u in the figure. It is not difficult to see that $[L]$ is the unit (in homology) for this multiplication and thus $\Theta(S_a) = [a]$, which concludes the proof. $\square$

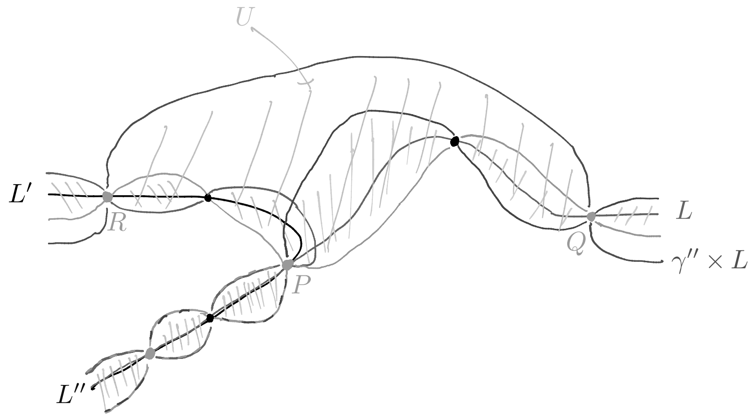

Figure 50. (Color online) The cobordisms S_a together with the schematic representation of a sufficiently small deformation W of $\gamma'' \times L$ and a triangle u.

4.3.2. $\mathsf{C}ob^*(M)$ *has surgery models*

In this subsection, we intend to finish showing that the category $\mathsf{C}ob^*(M)$ satisfies the axioms from Section 2.2. Given that we already constructed the functor Θ from Theorem 4.1.1 and that we showed in Proposition 4.2.12 that V and V' are cabling equivalent if and only if $\Theta(V) = \Theta(V')$, we already know that the quotient functor $\widehat{\Theta}$ is well defined and injective. It follows that Axioms 1 and 2 are satisfied. We showed in Section 4.3.1 that Axiom 4 is also satisfied.

It remains to show that Axiom 3, which was verified for elements in $\mathcal{L}ag_0^*(\mathbb{C}\times M)$ in Corollary 4.2.9, is also true for general V's $\in \mathcal{L}ag^*(\mathbb{C}\times M)$. Moreover, we need to show that the functor $\widehat{\Theta}$ is triangulated. Indeed, if this is so, Axiom 5—in the form made explicit in Remark 2.3.3(b)—is also satisfied because the relevant naturality properties are already satisfied in the triangulated category $H(\mathrm{mod}(\mathcal{F}uk^*(M)))$ which is the image of the injective functor $\widehat{\Theta}$.

Therefore, taken together, these two claims imply that $\mathsf{C}ob^*(M)$ has surgery models.

LEMMA 4.3.2

If V is a simple cobordism $V \in \mathcal{L}ag^*(\mathbb{C}\times M)$, possibly carrying (nonmarked) teardrops, then $\Theta(V)\circ\Theta(\bar{V}) = \mathrm{id}$.

Proof

Given that V may carry teardrops, we need to deal with the possible presence of pivots. The difficulty that these create is that the argument used in proving Corollary 4.2.9, which is based on a Hamiltonian isotopy induced by a translation in the plane, is no longer directly applicable as the image of the moving cobordism intersects the pivots. The present argument makes use of Figure 51 below. The ends of the cobordism V are L and L'. We consider a Lagrangian $W \subset \mathbb{C}\times M$, $W = \gamma \times N$, where γ is a circle intersecting the cylindrical region of V at the point P and Q and N is in $\mathcal{L}ag^*(M)$ in general position relative to L and

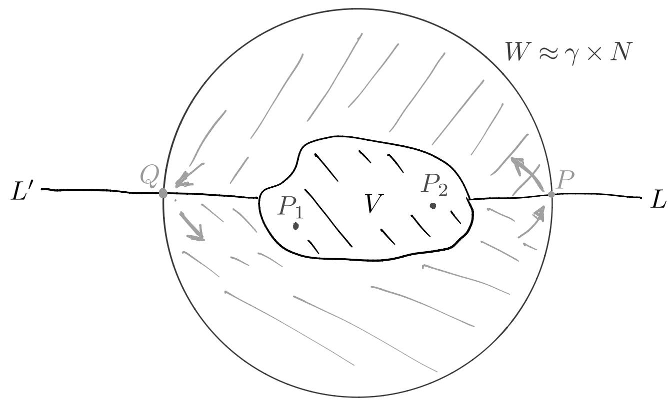

Figure 51. (Color online) Schematic representation of the cobordism V of the Lagrangian $W = \gamma \times N$; P_1 and P_2 are pivots for V.

L'. Assume without loss of generality that N is embedded. The points P and Q are regular bottlenecks for V. Following the definition of ϕ_V^W (see (19)), we see that this morphism is given by counting (marked strips) going from points $\{P\} \times \{x\}$ to points $\{Q\} \times \{y\}$, where $x \in N \cap L$, $y \in N \cap L'$. Similarly, $\phi_{\bar{V}}^W$ is given by counting the strips going from Q to P. We now consider the composition $\phi_V \circ \phi_{\bar{V}}$ and view it as providing a count of the points in a part of the boundary of the moduli space of marked strips going from Q to Q. Another such part consists of disks of Maslov class 2 with boundary on W, going through the points in $\{Q\} \times (N \cap L')$. There is precisely one such disk for each point in $\{Q\} \times (N \cap L')$, and we conclude that $\phi_V \circ \phi_{\bar{V}}$ is chain homotopic to the identity. The same argument actually shows that the composition $\phi_V \circ \phi_{\bar{V}}$ is chain homotopic to the identity as an A_∞-module map and concludes the proof. Indeed, recall that the construction of the morphism ϕ_V associated to the cobordism $V : L \rightsquigarrow L'$, as constructed in [10] (and discussed in Section 4.2.1), provides not only a chain morphism $CF(N, L) \to CF(N, L')$ for each appropriate N, but actually extends to an A_∞-module morphism $\mathcal{Y}_e(L) \to \mathcal{Y}_e(L')$. Similarly, the morphism $\phi_{\bar{V}}$ defines an A_∞-module morphism $\mathcal{Y}_e(L') \to \mathcal{Y}_e(L)$. Taking this into account, the arguments above are easily seen to show that the composition of these two morphisms, as A_∞-modules, is homotopic to the identity. $\square$

LEMMA 4.3.3
The functor $\widehat{\Theta}$ is triangulated.

Proof
Recall from Section 2.1.1 and Figure 5 the definition of the distinguished triangles in $\mathsf{C}ob^*(M)$. It is easy to see that the distinguished triangle associated to a surgery cobordism S_a, such as constructed in Section 4.3.1, is mapped by Θ to an exact triangle in the homological category of A_∞-modules over $\mathcal{F}uk^*(M)$ (a simple proof in this case is perfectly similar to the embedded case in [10]). However,

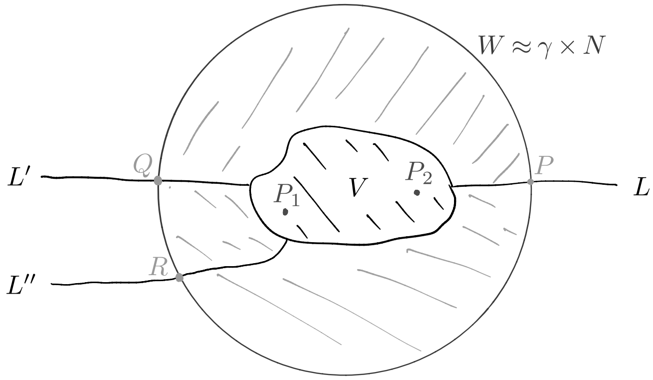

Figure 52. (Color online) Schematic representation of the cobordism V with three ends L, L', L'', and $W = \gamma \times N$; P_1 and P_2 are pivots for V.

we need to deal with the case of $V : L \rightsquigarrow (L'', L')$, $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$, when V possibly carries (nonmarked) teardrops. To do this we provide a different proof, along an approach similar to the one in Lemma 4.3.2. We consider Figure 52 below (cf. Figure 51). There are now three regular bottlenecks P, Q, R, each corresponding to one of the ends of V. The curve γ is a circle, as in Figure 51. We consider a Lagrangian $W = \gamma \times N$, where N is embedded, in general position relative to all three of L, L', L''.

We again consider the (marked) Floer complex C_P generated by the points $\in \{P\} \times (N \cap L)$ and similarly C_Q, C_R for the other ends of V. There is a chain map that we denote by $\psi^P_Q : C_P \to C_Q$ (and, similarly, ψ^Q_R, ψ^R_P) defined by counting Floer (marked) strips going from intersection points over P to intersection points over Q. These three maps are respectively identified with the three maps associated through the construction of Θ to the three maps in the distinguished triangle corresponding to V, as given in Figure 5. Thus, we need to show that these three maps fit into an exact triangle of chain complexes (and, by extension, of A_∞-modules). For this purpose, denote by $C_{P,Q}$ the cone of the map ψ^P_Q: it is given as a vector space by $C_P \oplus C_Q$ and has the differential $D_{P,Q} = (d_P + \psi^P_Q, d_Q)$. There is a chain map $\psi^{P,Q}_R : C_{P,Q} \to C_R$ defined by counting strips from intersection points over either P or Q to intersection points over R. Moreover, there is a chain map $\psi^R_{P,Q} : C_R \to C_{P,Q}$ given by counting strips from intersection points over R to intersection points that are either over P or Q. As in the proof of Lemma 4.3.2, these two maps are homotopy inverses. We notice that the projection $C_{P,Q} \to C_P$ composed with $\psi^R_{P,Q}$ agrees with ψ^R_P and that $\psi^{P,Q}_R$ composed with the inclusion $C_Q \to C_{P,Q}$ agrees with ψ^Q_R, which shows that the three maps ψ^P_Q, ψ^Q_R, ψ^R_P fit into an exact triangle in the homotopy category. The final step, as in Lemma 4.3.2, is to move from chain complexes to A_∞-modules over the category $\mathcal{F}uk^*_e(M)$. This argument is again a variant of the constructions in [10]. $\square$

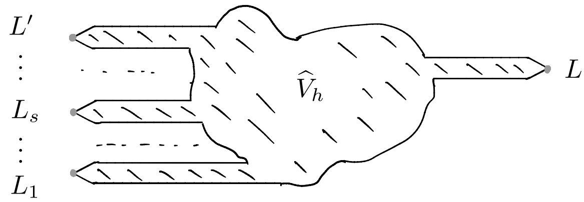

Figure 53. (Color online) The "truncated" cobordism $\hat{V}_h$ used to define the shadow of $(V, c, \mathcal{D}_V, h)$.

4.4. Proof of Theorem 4.1.1

Most of Theorem 4.1.1 has been already established. We only need to discuss the metric statements in this theorem.

We start by defining the shadow of an element $(V, c, \mathcal{D}_V, h) \in \mathcal{L}ag^*(\mathbb{C} \times M)$. Recall that V_h has regular bottlenecks along each of its ends, and let $\hat{V}_h$ be the part of V_h that is left after removing the regions to the left of the positive bottlenecks and to the right of the negative bottlenecks, as in Figure 53 (cf. the cobordism V_h from Figure 41). The shadow of $(V, c, \mathcal{D}_V, h)$ is by definition the shadow of $\hat{V}_h$, in the sense of Section 2.4.

Notice that the band width parameter, ϵ, appearing in the definition of the perturbation germs h_L associated to an element $L \in \mathcal{L}ag^*(M)$, as discussed in Section 3.4.1, can be taken arbitrarily small. Making use of this and noticing also that in the construction of the surgery morphisms in Section 4.3 we may take the various relevant deformations with an arbitrary small area covered by their shadows—this applies to h_L, $h_{L'}$ as well as the deformations giving the cobordisms V and V'—we deduce that the category $\mathsf{C}ob^*(M)$ has small surgery models in the sense of Section 2.4.

REMARK 4.4.1

The surgeries that occur in our framework are 0-sized, hence the shadow of surgery cobordisms is always 0. However, due to the way we have implemented the category $\mathsf{C}ob^*(M)$, we always perturb cobordisms V into V_h. These already have positive shadow, which can be taken to be arbitrarily small (by adjusting h). Thus the fact that $\mathsf{C}ob^*(M)$ has small surgery models is a result of these perturbations, and not of performing surgeries with small handle sizes. If we could realize this category without the perturbations V_h, then the surgery models would not just be small but actually of size 0.

The next step is to see that $\mathsf{C}ob^*(M)$ satisfies the rigidity Axiom 6 from Section 2.4. We will prove a stronger statement, namely, that $\mathsf{C}ob^*(M)$ is strongly rigid.

LEMMA 4.4.2

The category $\mathsf{C}ob^*(M)$ is strongly rigid (see Definition 2.4.2).

Proof
The first part of the proof is to show that $\mathsf{C}ob^*(M)$ is rigid. This means that we need to show that for any two Lagrangians $L, L' \in \mathcal{L}ag^*(M)$ with $j_L(L) \neq j_{L'}(L')$ there exists a constant $\delta(L, L') > 0$ such that for any simple cobordism $V : L \rightsquigarrow L'$, $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$, we have $\mathcal{S}(V) \geq \delta(L, L')$.

Recall that for two Lagrangians L and L', the Gromov width of L relative to L' is defined by

$$w(L; L') = \sup\left\{ \frac{\pi r^2}{2} \;\middle|\; \exists e : (B_r, \omega_o) \hookrightarrow (M, \omega), e(B_r) \cap L' = \emptyset, e^{-1}(L) = \mathbb{R}B_r \right\},$$

where e is a symplectic embedding, (B_r, ω_0) is the standard ball of radius r, and $\mathbb{R}B_r = B_r \cap (\mathbb{R}^n \times \{0\})$ (this definition is generally used in the embedded case but adjusts to the immersed situation without difficulty).

Assume now that $V : L \rightsquigarrow L'$, $V \in \mathcal{L}ag_0^*(\mathbb{C} \times M)$ is a simple cobordism. Proceeding as in [23], we deduce that through the center of a ball $e(B_r)$, as in the definition of $w(L; L')$, passes a perturbed marked J_0-holomorphic strip u with boundary along L and on L' and of energy not bigger than $\mathcal{S}(V)$.

The energy of u is bounded from below by the energy of $\mathrm{Image}(u) \cap e(B_r)$. By adjusting the data associated to the pair L, L', there exists a constant $c_{L,L'} > 0$ depending on the perturbations $\mathcal{D}_L$ and $\mathcal{D}_{L'}$ (and thus also of the base almost complex structure J_0) such that the energy of $\mathrm{Image}(u) \cap e(B_r)$ is bounded from below by $c_{L,L'} \frac{\pi r^2}{2}$ (if the almost-complex structure used inside B_r could be assumed to be the standard one and the perturbations vanish there, the energy in question would agree with the symplectic area and we could take $c_{L,L'} = 1$).

In summary, we have

$$\mathcal{S}(V) \geq c_{L,L'} w(L; L') \tag{21}$$

which implies that Axiom 6 is satisfied at least for simple cobordisms that do not carry nonmarked teardrops. In case $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$, the argument in [23] no longer applies because the argument there is based on a Hamiltonian isotopy such as the one appearing in Lemma 4.2.7 and this isotopy potentially intersects the poles associated to V. However, a quantitative adjustment of the argument in Lemma 4.3.2 does imply the statement even in this case.

We now turn to proving the nondegeneracy of the metric $d^{\mathcal{F},\mathcal{F}'}$, as described in Definition 2.4.2. In principle, this can be done following the arguments from [12] and adapting them to the immersed framework. Below, however, we will use the cabling construction to reduce the problem to the case of simple cobordisms and then use what we have proved in the preceding part of the proof.

Consider a cobordism $V : L \rightsquigarrow (F_1, \ldots, F_i, L', F_{i+1}, F_m)$, $V \in \mathcal{L}ag^*(\mathbb{C} \times M)$. Using the braiding in Remark 2.3.2, we may replace it by a cobordism $V' : L \rightsquigarrow L''$ where L'' is obtained by iterated 0-size surgeries involving the F_i's and L'. Because $\mathsf{C}ob^*(M)$ has small surgery models, the cabling used in the braiding can be assumed to add to the shadow of the initial cobordism V as little as wanted. In short, for any $\epsilon > 0$ we can find such a V' such that $\mathcal{S}(V') \leq \mathcal{S}(V) + \epsilon$. Assume now that there are families $\mathcal{F}$ and $\mathcal{F}'$ with $\overline{\bigcup_{F \in \mathcal{F}} F} \cap \overline{\bigcup_{F' \in \mathcal{F}'} F'}$ totally discrete and $L \neq L'$ such that $d^{\mathcal{F},\mathcal{F}'}(L, L') = 0$. Focusing for now just on the family $\mathcal{F}$, it follows from the arguments above that, for any ϵ, there is a cobordism $V'_\epsilon : L \rightsquigarrow L''_\epsilon$ of shadow less than ϵ with $L''_\epsilon = F_m \#_{c_m} \cdots \# L' \# \cdots \# F_1$. Consider now a point $x \in L \backslash L'$. As there exists a sufficiently small standard symplectic ball around x (as in the definition of width) that also does not intersect L', we deduce from (21) that $x \in \overline{\bigcup_{F\in\mathcal{F}} F}$. Take a small neighborhood $U \subset (L \backslash L')$ of x in L. By what we have just proved, we obtain that $U \subset \overline{\bigcup_{F\in\mathcal{F}} F}$. We apply the same argument to $\mathcal{F}'$ and we arrive at a contradiction because of our assumption on $\mathcal{F}$ and $\mathcal{F}'$. $\square$

To conclude the proof of Theorem 4.1.1 there are two more statements to prove. The most important part is that the restriction $\widehat{\Theta}_e$ of the functor $\widehat{\Theta}$ to the triangulated subcategory $\widehat{\mathsf{C}}ob^*_e(M)$ generated by the embedded Lagrangians gives a triangulated equivalence to $D\mathcal{F}uk^*(M)$. The functor $\widehat{\Theta}$ is triangulated, and all modules in $D\mathcal{F}uk^*(M)$ are iterated cones of Yoneda modules of embedded objects. Therefore, by using Proposition 4.3.1, recurrence on the number of cones, and equation (16), the result follows. Notice that it also follows from Proposition 4.3.1 and equation (16) that the category $\widehat{\mathsf{C}}ob^*(M)$ itself is the Donaldson category $\mathcal{D}on(\mathcal{L}ag^*(M))$ associated to the Lagrangians in $\mathcal{L}ag^*(M)$.

REMARK 4.4.3

Notice that in this construction the base almost-complex structure that one starts with can be any regular almost-complex structure J_0 on M because exact, embedded Lagrangians are unobstructed with respect to any such structure. Moreover, because we use in Proposition 4.3.1 0-size surgeries, all modules in $D\mathcal{F}uk^*(M)$ are represented by marked immersed Lagrangians $j_L : L \to M$ such that the immersion j_L restricts to embeddings on the connected components of L. In other words, in order to represent $D\mathcal{F}uk^*(M)$ the only immersions required are unions of embedded Lagrangians together with appropriate choices of markings that are used to deform the usual Floer differential as well as the higher A_∞ multiplications.

The last part of Theorem 4.1.1 claims that $\widehat{\Theta}_e$ is nondecreasing with respect to the pseudometrics mentioned there. This follows through the same arguments as in [12], but we will not give more details here.

5. Some further remarks

5.1. Related results

As mentioned before, the literature concerning Lagrangian cobordism starts with the work of Arnold [5], [6] who introduced the notion. Rigid behavior first appears in this context in a remarkable article due to Chekanov [21]. All the consequences and particular statements that result from the existence of the functor Θ in Section 4.2.1, when applied only to *embedded Lagrangians and cobordisms*, were established mainly in [9] and [10]. A variety of additional facts are known and, as they also contribute to various other aspects of the geometry-algebra dictionary discussed here, we will rapidly review the relevant literature here.

5.1.1. Simple cobordisms

As discussed before, unobstructed simple cobordisms induce isomorphisms between the Floer-theoretic invariants of the two ends. More is likely to be true: for instance, simple exact cobordisms are conjectured to coincide with Lagrangian suspensions. Steps toward proving this conjecture are due to Suárez and Barraud (see [8], [50]). They prove, under constraints on the behavior of the fundamental group inclusions of the ends, that an exact simple cobordism is a pseudoisotopy (i.e., the cobordism is smoothly trivial).

5.1.2. Additional algebraic structures

There are a variety of results of this sort. Among them: relations with Calabi–Yau structures due to Campling [17]; to stability conditions, discovered by Hensel [32]; Θ can be used (see [19]) to provide a categorification of Seidel's representation that views this representation as an algebraic translation of Lagrangian suspension; cobordisms can be studied also in the total space of Lefschetz fibrations over the complex plane (see [11]); they also can be used to approach a Seidel type exact sequence for generalized Dehn twists from a different angle (see [36]). Some connections to mirror symmetry appear in Sheridan–Smith [49] as well as in work of Hicks [33]

5.1.3. Cobordism groups

Following earlier results due to Audin [7] and Eliashberg [24] that compute in the flexible case cobordism groups in Euclidean spaces by means of algebraic topology, there are essentially two recent calculations in the rigid case. Both apply to surfaces: Haug [31] for the torus and Perrier [41] for surfaces of genus at least 2 (this last published article contains a number of gaps and errors; a corrected version is in preparation by Rathel-Fournier [45]).

5.1.4. Variants of the "surgery models" formalism

It is likely that the concept of a cobordism category with surgery models, or a variant, can be applied to some other contexts. There is some work in progress by Fontaine [26] which focuses on a related formalism for Morse functions. In a different direction, it is not clear for now whether the five axioms governing categories with surgery models are in any way optimal.

5.1.5. Measurements and shadows

As mentioned earlier, shadows of cobordisms first appeared in [23] for the case of simple cobordisms and in the wider setting of general cobordisms in [12], where fragmentation metrics were also introduced. A variety of interesting related results are due to Bisgaard [13], [14]. In a different direction, an interesting new approach is due to Chassé [20] who is exploring the behavior of shadow-type metrics for classes of Lagrangians under uniform curvature bounds.

5.1.6. Other approaches/points of view
Lagrangian cobordism has also been considered from a different perspective by Nadler and Tanaka (see [39], [52]). Another direction of study is concerned with Lagrangian cobordism in Weinstein manifolds in the context of the wrapped Fukaya category (see, e.g., [15], [51]). There are also connections, only superficially explored at the moment, to microlocal sheaf-theoretic techniques as evidenced in work of Viterbo [53, p. 55] (see also more recent contributions of Ganatra–Pardon–Shende [30], Nadler [37], and Nadler–Shende [38]).

Finally, there is considerable work today on cobordisms between Lagrangians through Legendrians. This is certainly a different point of view relative to the one that is central to us here, but there are also a variety of common aspects. For remarkable applications of this point of view, see Chantraine, Rizell, Ghiggini, and Golovko [18].

5.2. Some further questions

5.2.1. Nonmarked Lagrangians
As mentioned above, in general, rigid cobordism categories with surgery models have to contain immersed Lagrangians among their objects. A natural question is whether one can have objects that are nonmarked (or, in other words, with empty markings).

In some framework the answer to this question is expected to be in the affirmative (thus providing a nontautological solution to Kontsevich's conjecture). However, there are several difficulties in achieving this, some conceptual and some technical.

Here are two examples of the problems on the conceptual side.

(a) The ϵ-surgery $L\#_{\mathbf{x},\epsilon}L'$ at some intersection points $\mathbf{x}=\{x_i\}_{i\in I}\subset L\cap L'$ with $\epsilon>0$ and L, L' exact is only exact (with suitable choices of handles) if $f_L(x_i)-f_{L'}(x_i)$ is the same for all $i\in I$. As a result, the morphisms that can be obtained through an analogue of Proposition 4.3.1 are associated to cycles of intersection points that have the same action. On the other hand, there is no such control of the action of the points representing $\Theta(V)$ with Θ the functor in Section 4.2.1, and thus it is not clear whether all the image of $\Theta(V)$, at least as defined in Section 4.2.1, can be represented geometrically.

(b) We have not discussed here changing the base almost-complex structure J_0, but it is certain that such changes in any case lead to a setup involving bounding chains.

The technical difficulties involved in removing marked Lagrangians from the picture and replacing them with immersed but nonmarked ones already appear at the following point. We need to compare $L\#_{c,\epsilon}L'$ for varying ϵ and fixed c and for

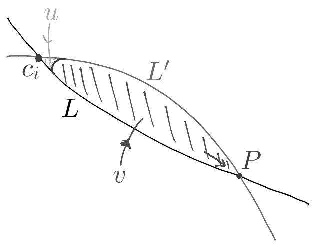

Figure 54. (Color online) Comparing teardrops to marked teardrops (or Floer strips) after rounding a type corner.

L and L' embedded. Assuming the marked Lagrangian $L\#_{c,0}L'$ is unobstructed (in the marked sense), the basic questions are whether it is true that

(i) the nonmarked Lagrangian $L\#_{c,\epsilon}L'$ is also unobstructed,
(ii) the surgery cobordism from $L'' = L\#_{c,0}L'$ to $L''' = L\#_{c,\epsilon}L'$ for small enough ϵ is also unobstructed and the relevant Yoneda modules of L'' and L''' are quasi-isomorphic.

A key step is to compare moduli spaces $\mathcal{N}_{L'',c}$ of marked teardrops with boundary on L'' to the moduli spaces $\mathcal{N}_{L'''}$ of nonmarked teardrops with boundary on L''', after performing surgery at the marking c, as pictured in Figure 54 (assuming the dimension of these moduli spaces is 0). The two moduli spaces need to be identified for ϵ small. While intuitively clear, this identification is not obvious. It depends on a gluing result that appears in [29] but, beyond that, there are other subtleties. For instance, to ensure regularity of $\mathcal{N}_{L'',c}$ it is natural to use nonautonomous almost-complex structures. On the other hand, $\mathcal{N}_{L'''}$ consists, by default, of autonomous complex structures, thus these moduli spaces are not so easy to render regular. Additionally, the gluing result in [29] applies to pseudoholomorphic curves with respect to autonomous almost-complex structures. Regularity of curves defined with respect to autonomous structures and without perturbations can potentially be approached as in [35], as attempted in [42] (the latter work is is progress and not complete at time of publication).

Further, the same methods also need to be applied to deal with moduli spaces of J-holomorphic disks and with more complicated curves. This requires a more general gluing statement such as of the type described in [29], that applies to polygons with more than a single corner having boundary on a generically immersed Lagrangian submanifold. Moreover, one also needs a version of the Gromov compactness theorem that applies to curves with boundary on a sequence of Lagrangians that converges to an immersed Lagrangian through a process that corresponds to reducing to zero the size of a handle attached through surgery.

An additional issue, related to the application of the Gromov compactness argument in this setting, is that to achieve the comparison between the moduli spaces $\mathcal{N}_{L'',c}$ and $\mathcal{N}_{L'''}$ one needs uniform energy bounds that apply to all ϵ-surgeries for ϵ small enough. Otherwise, the moduli spaces $\mathcal{N}_{L'''}$ might contain curves with energy going to ∞ when $\epsilon \to 0$ and these curves obviously are not "seen" by the spaces $\mathcal{N}_{L'',c}$.

The gluing and the compactness ingredients mentioned are basically "within reach" today. The gluing statement seems very close to [29] and the compactness one is likely to follow from some version of symplectic field theory compactness (see [16]). For both, see also the recent work [40].

Another question in the context of cobordism categories with nonmarked Lagrangian immersions has to do with the smallness of the surgery models. As explained in Remark 4.4.1, the category $\mathsf{C}ob^*(M)$ addressed in Theorem 4.1.1 has *small* surgery models precisely because its objects consist of marked immersions. Assuming that an implementation of such a category with *nonmarked* Lagrangian immersions is possible, it is not clear whether it would have *small* surgery models.

5.2.2. Bounding chains and invariance properties

As mentioned in Section 1.3.3, bounding chains and the markings of Lagrangian immersions used here are closely related. Recalling from [2] the Akaho–Joyce A_∞-algebra associated to an immersed Lagrangian L, it is easy to see that a marking associated to L provides a bounding chain for this algebra. Moreover, the marked Floer homology of two such marked Lagrangians L and L' is the deformed Floer homology associated to the respective bounding chains. Conversely, possibly under some additional conditions, it is expected that a bounding chain can be reduced to a marking.

Also of interest is a more detailed exploration of the relation between markings and bounding chains because the notion of unobstructedness that we use here is somewhat unsatisfying since it depends on a specific choice of a base almost-complex structure J_0. Certainly, it remains true that the resulting category $\widehat{\mathsf{C}ob}^*_e(M)$ is a model for the derived Fukaya category of the embedded objects and thus it does not depend on J_0 up to triangulated equivalence. Nonetheless, it is expected that a more invariant construction would make explicit use of bounding chains to adjust the markings, data, and so on when transitioning from one base almost-complex structure to another.

In this context, an interesting related question is to understand the behavior of bounding chains with respect to surgery. A first step in this direction appears in [40].

Other related questions are concerned with invariance properties of the marked Floer theory with respect to exact homotopies of Lagrangian immersions. By this we mean homotopies $\Psi : L \times [0,1] \longrightarrow M$ such that $\psi_t := \Psi(-,t) : L \longrightarrow M$ is an exact Lagrangian immersion for all t. (Of course, one can consider here also more restrictive types of homotopies.) Can an analogue of the Lagrangian suspension cobordism be associated with such a homotopy? Can it be endowed with a meaningful marking (resp., bounding chain) that extends a given marking (resp., bounding chain) on its negative end (i.e., the one corresponding to $t = 0$)? Will the marked (resp., deformed) Floer homology remain invariant under such homotopies?

5.2.3. Novikov coefficients

Going beyond $\mathbb{Z}/2$-coefficients seems possible, as is, in particular, using coefficients in the universal Novikov field Λ. This extension is, however, not completely immediate as it requires one to incorporate into the surgery construction information coming from the coefficients. In other words, where in the body of the article we viewed surgery as an operation associated to a cycle $c = \sum c_i$, where $c_i \in L \cap L'$, now we need to consider cycles of the form $c = \sum \alpha_i c_i$ with $\alpha_i \in \Lambda$ and the surgery needs to take into account the size of the α_i's. The sizes appearing in the exponents of the α_i's would be related to the size of handles in the surgery. The other coefficients in the α_i's (coming from the base field over which Λ is defined) can be taken into account by endowing the surgered Lagrangian with a suitable local system.

5.2.4. Idempotents

With motivation in questions coming from mirror symmetry it is natural to try to understand how to take into account idempotents in the global picture of the cobordism categories considered here. This turns out not to be difficult (even if we leave the details for a subsequent publication): one can view idempotents as an additional structure given by a class q_L in the quantum homology $\mathrm{QH}(L)$ defined for each one of the ends L of a cobordism and such that $q_L^2 = q_L$. A similar structure also makes sense for a cobordism V, where in this case the relevant class belongs to $\mathrm{QH}(V, \partial V)$. There is a connectant $\mathrm{QH}(V, \partial V) \to \mathrm{QH}(\partial V)$ and the requirement for a cobordism respecting the idempotents is for the classes on the boundary to descend from that on V. The case of standard cobordisms is recovered when all these classes are the units. For embedded Lagrangians, the additional structure provided by idempotents is not hard to control (compared to what was discussed before one also needs to show that the various morphisms induced by cobordisms also relate in the appropriate way the idempotents of the ends). However, we expect that the whole construction will be of increased interest when cobordisms and their ends are allowed to be immersed. The second author thanks Mohammed Abouzaid for a useful discussion on this topic.

5.2.5. Small scale metric structure

For an embedded Lagrangian $L \in \mathcal{L}ag_e^*(M)$, a basic question is whether there is some $\epsilon > 0$ such that, if $d^{\mathcal{F},\mathcal{F}'}(L, L') \leq \epsilon$ for some $L' \in \mathcal{L}ag^*(M)$, then L' is also embedded and, further, L' is Hamiltonian isotopic to L. Could it be that even more is true for ϵ possibly even smaller: there is a constant k depending only on $\mathcal{F}$, $\mathcal{F}'$ such that $d^{\mathcal{F},\mathcal{F}'}(L, L') \geq k d_H(L, L')$ (where d_H is the Hofer distance)?

5.2.6. Embedded versus immersed

There are two natural questions related to the formalism discussed here. Are there criteria that identify objects of $D\mathcal{F}uk^*(M)$ that are not representable by embedded Lagrangians? Similarly, are there criteria that identify unobstructed immersed, marked, Lagrangians that are not isomorphic to objects of $D\mathcal{F}uk^*(M)$?

Acknowledgments. This work started during the second author's stay at the Institute for Advanced Study (IAS) for the academic year 2015–2016. The first author was also hosted by the IAS during two weeks in the fall of 2015. Both authors thank the IAS and, in particular, Helmut Hofer for the hospitality extended to them. The second author also thanks the Forschungsinstitut für Mathematik (FIM) in Zurich for its hospitality during frequent research visits. We thank the referees for careful readings and very useful suggestions.

Cornea's work was partially supported by an individual Natural Sciences and Engineering Research Council (NSERC) Discovery Grant, a Simons Fellowship, and an IAS fellowship grant.

References

[1] M. Akaho, *Intersection theory for Lagrangian immersions*, Math. Res. Lett. **12** (2005), no. 4, 543–550. MR 2155230. DOI 10.4310/MRL.2005.v12.n4.a8.

[2] M. Akaho and D. Joyce, *Immersed Lagrangian Floer theory*, J. Differential Geom. **86** (2010), no. 3, 381–500. MR 2785840. DOI 10.4310/JDG/1303219427.

[3] G. Alston and E. Bao, *Exact, graded, immersed Lagrangians and Floer theory*, J. Symplectic Geom. **16** (2018), no. 2, 357–438. MR 3834425. DOI 10.4310/JSG.2018.v16.n2.a2.

[4] ———, *Immersed Lagrangian Floer cohomology via pearly trajectories*, preprint, arXiv:1907.03072v2 [math.SG].

[5] V. I. Arnold, *Lagrange and Legendre cobordisms, I*, Funktsional. Anal. i Prilozhen. **14** (1980), no. 3, 1–13. MR 0583797.

[6] ———, *Lagrange and Legendre cobordisms, II*, Funktsional. Anal. i Prilozhen. **14** (1980), no. 4, 8–17. MR 0595724.

[7] M. Audin, *Quelques calculs en cobordisme lagrangien*, Ann. Inst. Fourier (Grenoble) **35** (1985), no. 3, 159–194. MR 0810672. DOI 10.5802/AIF.1023.

[8] J.-F. Barraud and L. S. Suárez, *The fundamental group of a rigid Lagrangian cobordism*, Ann. Math. Què. **43** (2019), no. 1, 125–144. MR 3925140. DOI 10.1007/s40316-018-0109-2.

[9] P. Biran and O. Cornea, *Lagrangian cobordism, I*, J. Amer. Math. Soc. **26** (2013), no. 2, 295–340. MR 3011416. DOI 10.1090/S0894-0347-2012-00756-5.

[10] ———, *Lagrangian cobordism and Fukaya categories*, Geom. Funct. Anal. **24** (2014), no. 6, 1731–1830. MR 3283928. DOI 10.1007/s00039-014-0305-4.

[11] ———, *Cone-decompositions of Lagrangian cobordisms in Lefschetz fibrations*, Selecta Math. (N.S.) **23** (2017), no. 4, 2635–2704. MR 3703462. DOI 10.1007/s00029-017-0318-6.

[12] P. Biran, O. Cornea, and E. Shelukhin, *Lagrangian shadows and triangulated categories*, preprint, arXiv:1806.06630v1 [math.SG].

[13] M. R. Bisgaard, *Invariants of Lagrangian cobordisms via spectral numbers*, J. Topol. Anal. **11** (2019), no. 1, 205–231. MR 3918067. DOI 10.1142/S1793525319500092.

[14] ———, *A distance expanding flow on exact Lagrangian cobordism classes*, preprint, arXiv:1608.05821v1 [math.SG].

[15] V. Bosshard, *On Lagrangian cobordisms in Liouville manifolds*, Ph.D. dissertation, ETH Zürich, Zürich, in preparation.

[16] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888. MR 2026549. DOI 10.2140/gt.2003.7.799.

[17] E. Campling, *Fukaya categories of Lagrangian cobordisms and duality*, preprint, arXiv:1902.00930v1 [math.SG].

[18] B. Chantraine, G. D. Rizell, P. Ghiggini, and R. Golovko, *Geometric generation of the wrapped Fukaya category of Weinstein manifolds and sectors*, preprint, arXiv:1712.09126v3 [math.SG].

[19] F. Charette and O. Cornea, *Categorification of Seidel's representation*, Israel J. Math. **211** (2016), no. 1, 67–104. MR 3474957. DOI 10.1007/s11856-015-1261-x.

[20] J.-P. Chassé, Ph.D. dissertation, University of Montreal, Montreal, in preparation.

[21] Y. V. Chekanov, "Lagrangian embeddings and Lagrangian cobordism" in *Topics in Singularity Theory*, Amer. Math. Soc. Transl. Ser. 2 **180**, Amer. Math. Soc., Providence, 1997, 13–23. MR 1767110. DOI 10.1090/trans2/180/02.

[22] O. Cornea and F. Lalonde, *Cluster homology: An overview of the construction and results*, Electron. Res. Announc. Amer. Math. Soc. **12** (2006), 1–12. MR 2200949. DOI 10.1090/S1079-6762-06-00154-5.

[23] O. Cornea and E. Shelukhin, *Lagrangian cobordism and metric invariants*, J. Diff. Geom. **112** (2019), no. 1, 1–45. MR 3948226. DOI 10.4310/jdg/1557281005.

[24] Y. Eliashberg, "Cobordisme des solutions de relations différentielles" in *South Rhone Seminar on Geometry, I (Lyon, 1983)*, Travaux en Cours, Hermann, Paris, 1984, 17–31. MR 0753850.

[25] A. Floer, *Morse theory for Lagrangian intersections*, J. Differential Geom. **28** (1988), no. 3, 513–547. MR 0965228. DOI 10.4310/jdg/1214442477.

[26] P. Fontaine, Ph.D. dissertation, University of Montreal, Montreal, in preparation.

[27] K. Fukaya, *Unobstructed immersed Lagrangian correspondence and filtered A infinity functor*, preprint, arXiv:1706.02131v3 [math.SG].

[28] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian Intersection Floer Theory: Anomaly and Obstruction, I*, AMS/IP Stud. Adv. Math. **46**, Amer. Math. Soc., Providence, 2009. MR 2553465. DOI 10.1090/crmp/049/07.

[29] ———, *Lagrangian intersection Floer theory: Anomaly and obstruction*, Chapter 10, preprint, 2007, http://www.math.kyoto-u.ac.jp/~fukaya/Chapter10071117.pdf.

[30] S. Ganatra, J. Pardon, and V. Shende, *Microlocal Morse theory of wrapped Fukaya categories*, preprint, arXiv:1809.08807v2 [math.SG].

[31] L. Haug, *The Lagrangian cobordism group of T^2*, Selecta Math. (N.S.) **21** (2015), no. 3, 1021–1069. MR 3366925. DOI 10.1007/s00029-014-0173-7.

[32] F. Hensel, *Stability conditions and Lagrangian cobordisms*, J. Symplectic Geom. **18** (2020), no. 2, 463–536. MR 4118147. DOI 10.4310/JSG.2020.v18.n2.a4.

[33] J. Hicks, *Tropical Lagrangian hypersurfaces are unobstructed*, J. Topol. **13** (2020), no. 4, 1409–1454. MR 4125753. DOI 10.1112/topo.12165.

[34] F. Lalonde and J.-C. Sikorav, *Sous-variétés lagrangiennes et lagrangiennes exactes des fibrés cotangents*, Comment. Math. Helv. **66** (1991), no. 1, 18–33. MR 1090163. DOI 10.1007/BF02566634.

[35] L. Lazzarini, *Relative frames on J-holomorphic curves*, J. Fixed Point Theory Appl. **9** (2011), no. 2, 213–256. MR 2821362. DOI 10.1007/s11784-010-0004-1.

[36] C.-Y. Mak and W. Wu, *Dehn twist exact sequences through Lagrangian cobordism*, Compos. Math. **154** (2018), no. 12, 2485–2533. MR 3873526. DOI 10.1112/s0010437x18007479.

[37] D. Nadler, *Wrapped microlocal sheaves on pairs of pants*, preprint, arXiv:1604.00114v1 [math.SG].

[38] D. Nadler and V. Shende, *Sheaf quantization in Weinstein symplectic manifolds*, preprint, arXiv:2007.10154v2 [math.SG].

[39] D. Nadler and H. L. Tanaka, *A stable ∞-category of Lagrangian cobordisms*, Adv. Math. **366** (2020), art. ID 107026. MR 4070298. DOI 10.1016/j.aim.2020.107026.

[40] J. Palmer and C. Woodward, *Invariance of immersed Floer cohomology under Lagrangian surgery*, preprint, arXiv:1903.01943v4 [math.SG].

[41] A. Perrier, *Lagrangian cobordism groups of higher genus surfaces*, preprint, arXiv:1901.06002v1 [math.SG].

[42] ———, *Structure of J-holomorphic disks with immersed Lagrangian boundary conditions*, preprint, arXiv:1808.01849v1 [math.SG].

[43] L. Polterovich, *The surgery of Lagrange submanifolds*, Geom. Funct. Anal. **1** (1991), no. 2, 198–210. MR 1097259. DOI 10.1007/BF01896378.

[44] M. Pozniak, "Floer homology, Novikov rings and clean intersections" in *Northern California Symplectic Geometry Seminar*, Amer. Math. Soc. Transl. Ser. 2 **196**, Amer. Math. Soc., Providence, 1999, 119–181. MR 1736217. DOI 10.1090/trans2/196/08.

[45] D. Rathel-Fournier, Ph.D. dissertation, University of Montreal, Montreal, in preparation.

[46] F. Schmäschke, *Floer homology of Lagrangians in clean intersection*, preprint, arXiv:1606.05327v1 [math.SG].

[47] P. Seidel, *Graded Lagrangian submanifolds*, Bull. Soc. Math. France **128** (2000), no. 1, 103–149. MR 1765826.

[48] ———, *Fukaya Categories and Picard-Lefschetz Theory*, Zur. Lect. Adv. Math., Eur. Math. Soc. (EMS), Zürich, 2008. MR 2441780. DOI 10.4171/063.

[49] N. Sheridan and I. Smith, *Rational equivalence and Lagrangian tori on K3 surfaces*, Comment. Math. Helv. **95** (2020), no. 2, 301–337. MR 4115285. DOI 10.4171/CMH/489.

[50] L. S. Suárez, *Exact Lagrangian cobordism and pseudo-isotopy*, Internat. J. Math. **28** (2017), no. 8, art. ID 1750059. MR 3681121. DOI 10.1142/S0129167X17500598.

[51] H. Tanaka, *Generation for Lagrangian cobordisms in Weinstein manifolds*, preprint, arXiv:1810.10605v5 [math.SG].

[52] ———, *Cyclic structures and broken cycles*, preprint, arXiv:1907.03301v1 [math.AT].

[53] C. Viterbo, *Sheaf quantization of Lagrangians and Floer cohomology*, preprint, arXiv:1901.09440v1 [math.SG].

Biran: Department of Mathematics, ETH Zurich, Zurich, Switzerland; biran@math.ethz.ch

Cornea: Department of Mathematics and Statistics, University of Montreal, Montreal, Canada; cornea@dms.umontreal.ca