

CANONICAL ORIENTATIONS IN HEEGAARD FLOER THEORY

MOHAMMED ABOUZAID AND CIPRIAN MANOLESCU

ABSTRACT. We set up Heegaard Floer theory over the integers, using canonical orientations coming from coupled Spin structures on the Lagrangian tori. We prove naturality of Heegaard Floer homology, sutured Floer homology, and link Floer homology over \mathbb{Z} . We give a new proof of the surgery exact triangle in this context, as well as a definition of involutive Heegaard Floer homology over \mathbb{Z} .

1. INTRODUCTION

In a series of papers, such as [30], [29], [31], [26], Ozsváth and Szabó developed Heegaard Floer theory: a collection of invariants of 3-manifolds, 4-manifolds, and knots. Since then, the theory has become an important tool in low-dimensional topology, leading to numerous applications.

In the original papers, the 3-manifold and knot invariants were defined as homology groups with coefficients in \mathbb{Z} , and the 4-manifold invariants were integers defined up to sign. However, at some point the community started ignoring signs, and it became customary to work over the coefficient field \mathbb{F}_2 . Notably, naturality of the Heegaard Floer invariants [16] and their full functoriality properties under cobordisms [46, 48] were only established over \mathbb{F}_2 . This is sufficient for many applications, but it does constitute a limitation for others. By contrast, monopole Floer homology [19] was defined over \mathbb{Z} , and the Seiberg-Witten invariants [44] are integers (with the sign determined by a homology orientation on the 4-manifold).

In Lagrangian Floer homology, there are several ways to choose signs for the counts of moduli spaces. One way is by *coherent orientations* [6], where one trivializes the determinant line bundles over sufficiently many homotopy classes of disks arbitrarily, and then trivializes them over the rest of the homotopy classes in the unique compatible way. Another way, which is more popular in the recent literature, is through *canonical orientations* induced from Spin or Pin structures on the Lagrangians; see [41, 10, 40]. Yet another way is by using twisted derived local systems on the Lagrangians [37].

In their papers on Heegaard Floer homology, Ozsváth and Szabó used coherent orientations to pin down the signs. The purpose of our paper is to offer a different perspective on Heegaard Floer theory over \mathbb{Z} , by using canonical orientations instead. These are better suited for questions about naturality and functoriality. We establish the naturality properties here and thus place the theory over \mathbb{Z} on a solid footing.

The canonical orientations on moduli spaces are constructed as follows. Ideally, we would like to fix Spin structures on the Lagrangian tori \mathbb{T}_α and \mathbb{T}_β that appear in Heegaard Floer homology. However, a Spin structure requires in particular an orientation, and in our case the Lagrangians do not have natural orientations. Rather, they have a *coupled orientation* (that is, an orientation on their product), which has the same origin as the absolute $\mathbb{Z}/2$ -grading in Heegaard Floer homology defined in [29, Section 10.4].

Instead of Spin structures, one can settle for Pin structures on the Lagrangians; these do not need a background orientation. We choose Lie group Pin structures, that is, structures

invariant under the torus action. Up to homotopy, there is only one Lie group Pin structure on the torus, and this suffices to determine the isomorphism class of the Heegaard Floer groups over \mathbb{Z} . Nevertheless, specifying the choice only up to homotopy is not enough for discussing naturality. The space of Lie group Pin structures has the homotopy type of $\mathbb{R}P^\infty$, with $\pi_1 = \mathbb{Z}/2$, and the monodromy on Floer homology around a non-trivial loop in this space is multiplication by -1 .

To rectify this problem, note that a Lie group Pin structure on the Lagrangian torus \mathbb{T}_α is specified by the Pin structure in the tangent space at each point. In turn, this tangent space can be identified with $\mathbb{A} = H^1(U_\alpha; \mathbb{R})$, where U_α is the alpha handlebody in the Heegaard splitting. Similarly, the tangent space to \mathbb{T}_β can be identified with $\mathbb{B} = H^1(U_\beta; \mathbb{R})$, where U_β is the beta handlebody. While there are no canonical Pin structures on \mathbb{A} and \mathbb{B} , it turns out that there is a canonical *coupled Spin structure* on the pair (\mathbb{A}, \mathbb{B}) . Coupled Spin structures are a new concept that we introduce in this paper. They make sense for a pair of parallelizable Lagrangians (such as our tori), and make use of an underlying coupled orientation (which exists in our setting). We will see that coupled Spin structures produce canonical orientations on the moduli spaces of J -holomorphic strips, and thus give Lagrangian Floer groups that are well-defined up to canonical isomorphism. In our context, we use them to define Heegaard Floer complexes over \mathbb{Z} .

With these choices, we establish the invariance of Heegaard Floer homology under the usual Heegaard moves from [30], as well as its naturality, following the proof scheme in [16]:

Theorem 1.1. *Let Y be a closed, oriented 3-manifold equipped with a basepoint $z \in Y$ and a $Spin^c$ structure \mathfrak{s} . The Heegaard Floer homologies \widehat{HF} , HF^+ , HF^- and HF^∞ (defined using the canonical coupled Spin structure) are invariants of the triple (Y, \mathfrak{s}, z) , well-defined up to natural isomorphism in the category of $\mathbb{Z}[U]$ -modules.*

The canonical orientations on moduli spaces that we construct produce in particular a coherent orientation system as in [30, Definition 3.11]. Thus, our Heegaard Floer groups are isomorphic to those constructed by Ozsváth-Szabó in [30], for some (unspecified) orientation system. Up to equivalence, there are $2^{b_1(Y)}$ orientation systems. In [29, Theorem 10.12], Ozsváth and Szabó further select a special orientation system, using their calculation of a twisted version \underline{HF}^∞ . We expect that their system is the same as the one arising from our work, but do not prove this here.

We mention one important way in which the proof of naturality in Theorem 1.1 departs from the arguments in [16]. Juhász, Thurston, and Zemke used maps associated to equivalences between the curve systems on the Heegaard diagram. (An equivalence is replacing such a system with any other that represents the same handlebody. We can think of it as a composition of isotopies and handleslides.) Over \mathbb{F}_2 , it is relatively easy to see that the maps induced by equivalences commute. Over \mathbb{Z} , it is harder; one can prove this up to a sign, as in [12]. What we do instead is to use maps associated to isotopies and handleslides. We then have to check certain relations associated to loops of handleslides, using the description of those loops in [16, Appendix A].

With our framework in place, we give a new proof of the surgery exact triangle from [29, Section 9], with \mathbb{Z} coefficients. (See Theorem 6.1.) Interestingly, the natural maps induced by counting holomorphic triangles (using the coupled Spin structures) are not the ones that appear in the triangle. One needs to use twisted coefficients for the maps, and we show that one can do so without changing the three Heegaard Floer groups themselves.

Another application of Theorem 1.1 is a definition of involutive Heegaard Floer homology over \mathbb{Z} . Over \mathbb{F}_2 , this theory was defined in [14], by taking the homology of the mapping

cone of an involution ι on Heegaard Floer complexes. The map ι involves moves on Heegaard diagrams and thus relies on naturality. Since we now have this property over \mathbb{Z} , the same construction can be done with integer coefficients. (See Section 7.4 below.)

In a different direction, we prove analogues of Theorem 1.1 for sutured Floer homology and for link Floer homology. Sutured Floer homology is an invariant of balanced sutured manifolds introduced by Juhász [15]. Roughly, a balanced sutured manifold is a compact 3-manifold M with boundary, whose boundary ∂M is split along a collection of sutures γ into two parts R_+ and R_- . The balanced condition says that R_+ and R_- have the same Euler characteristic. The construction of sutured Floer homology proceeds along similar lines to that of Heegaard Floer homology, but we no longer have a canonical coupled Spin structure. Instead, we are forced to choose one, and this is equivalent to choosing a coupled Spin structure S on the pair of vector spaces $(H_1(M, R_-; \mathbb{R}), H_1(M, R_+; \mathbb{R}))$. We call this data a *homological coupled Spin structure*. From this we can define a sutured Floer homology group $SFH(M, \gamma, \mathfrak{s}, S)$ over \mathbb{Z} .

Theorem 1.2. *Let (M, γ) be a sutured manifold equipped with a $Spin^c$ structure \mathfrak{s} and a homological coupled Spin structure S . Then, the sutured Floer homology $SFH(M, \gamma, \mathfrak{s}, S)$ is a natural invariant of the quadruple $(M, \gamma, \mathfrak{s}, S)$.*

Finally, we consider link Floer homology, an invariant of links in three-manifolds introduced by Ozsváth and Szabó in [26], [28]. The original construction was over \mathbb{Z} for knots (see [26], [36], but only over \mathbb{F}_2 for links, due to the presence of disk bubbles. For links in S^3 , a refinement over \mathbb{Z} was constructed using grid homology in [25]. For links in arbitrary 3-manifolds, Sarkar [38] explained the existence of several possible theories over \mathbb{Z} (depending on some choices). In our set-up, we find a canonical coupled Spin structure on the Lagrangian tori, and we use it to identify a single preferred theory over \mathbb{Z} . This theory comes in several flavors, such as \widehat{HFL} or HFL^- .

Theorem 1.3. *Let $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ be a multi-based oriented link in a three-manifold Y , and $\underline{\mathfrak{s}}$ a $Spin^c$ structure on Y relative to L . Then, the link Floer homologies $\widehat{HFL}(Y, \mathbb{L}, \underline{\mathfrak{s}})$ and $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ are natural invariants of the triple $(Y, \mathbb{L}, \underline{\mathfrak{s}})$.*

In the proofs of the naturality statements in Theorems 1.2 and 1.3, we again need to consider loops of handleslides. We will make use of the work of Qin [35], who recently generalized the results of [16, Appendix A] to these contexts.

In future work, we will use canonical orientations to prove that Heegaard Floer homology and link Floer homology are functorial over \mathbb{Z} under cobordisms, refining the results of Ozsváth-Szabó [31] and Zemke [46], [48]. The cobordism maps are determined by homology orientations, in a manner reminiscent of what happens in monopole Floer homology [19].

Organization of the paper. In Section 2 we discuss Spin, Pin, and coupled Spin structures on vector bundles. In Section 3 we explain the general construction of canonical orientations from Pin structures, in the context of Lagrangian Floer theory. In Section 4 we define Heegaard Floer homology over \mathbb{Z} and prove that its isomorphism class is a 3-manifold invariant. In Section 5 we prove Theorem 1.1 in its full strength, by establishing naturality. In Section 6 we prove the surgery exact triangle. In Section 7 we discuss signs in a few other settings, and in particular prove Theorems 1.2 and Theorem 1.3.

Acknowledgements. We thank Irving Dai, András Juhász, Tye Lidman, Robert Lipshitz, Maggie Miller, Tom Mrowka, Qianhe Qin, Sucharit Sarkar, Eha Srivastava, and Ian Zemke for helpful conversations. MA was partially supported by NSF award DMS-2103805.

CM was partially supported by a Simons Investigator Award and the Simons Collaboration Grant on New Structures in Low-Dimensional Topology.

2. SPIN, PIN, AND COUPLED SPIN

In this section we review some generalities about Spin and Pin structures, and then introduce the new notion of a coupled Spin structure.

2.1. Pin structures. We start by listing some relevant facts about Pin structures, comparing them to the better-known Spin structures. For more details, we refer to [18] and [40].

Throughout the paper let $\text{Pin}(n)$ denote the *positive* Pin group in dimension n . This is the central extension of $\text{O}(n)$ by $\mathbb{Z}/2$ determined (as a group extension) by the element $w_2 \in H^2(\text{BO}(n); \mathbb{Z}/2)$. It can be constructed explicitly as follows. Let $\mathcal{C}\ell(n)$ be the real Clifford algebra on n generators e_1, \dots, e_n , subject to the relations $e_i^2 = 1$ for all i , and $e_i e_j = -e_j e_i$ for all $i \neq j$. Then, $\text{Pin}(n)$ is the subset of $\mathcal{C}\ell(n)$ consisting of elements of the form $v_1 v_2 \cdots v_k$ where v_i are unit vectors in $\mathbb{R}^n = \text{Span}(e_1, \dots, e_n)$. The subgroup $\text{Spin}(n) \subset \text{Pin}(n)$ consists of those $v_1 v_2 \cdots v_k$ with k even, and the double cover map $\text{Pin}(n) \rightarrow \text{O}(n)$ is given by composing reflections across the hyperplanes v_i^\perp .

We obtain a commutative diagram of Lie group homomorphisms

$$\begin{array}{ccccccc} 1 & \longrightarrow & \mathbb{Z}/2 & \longrightarrow & \text{Spin}(n) & \longrightarrow & \text{SO}(n) \longrightarrow 1 \\ & & \downarrow \text{id} & & \downarrow & & \downarrow \\ 1 & \longrightarrow & \mathbb{Z}/2 & \longrightarrow & \text{Pin}(n) & \longrightarrow & \text{O}(n) \longrightarrow 1 \end{array}$$

Remark 2.1. The group $\text{Pin}(n)$ is denoted Pin_n^+ in [18]. It is to be distinguished from the *negative* Pin group, denoted Pin_n^- in [18], which is the central extension of $\text{O}(n)$ by $\mathbb{Z}/2$ classified by $w_2 + w_1^2 \in H^2(\text{BO}(n); \mathbb{Z}/2)$. For example, for $n = 1$ we have

$$\text{Pin}_1^+ \cong \mathbb{Z}/2 \times \mathbb{Z}/2, \quad \text{Pin}_1^- \cong \mathbb{Z}/4,$$

so that the surjection $\text{Pin}(1) \rightarrow \text{O}(1)$ splits.

For $n = 2$, we have $\text{Pin}_2^+ \cong \text{O}(2)$, whereas Pin_2^- is the group appearing in gauge theory from the symmetries of the Seiberg-Witten equations [11, 23].

From now on, whenever we talk about a base B for a vector bundle, we will assume that B is paracompact.

Definition 2.2. Let $E \rightarrow B$ a real vector bundle of rank n with an inner product, and $\text{Fr}(E) \rightarrow B$ the principal $\text{O}(n)$ -bundle of orthonormal frames in E . A *Pin structure* on E is defined to be a lift of $\text{Fr}(E)$ to a principal $\text{Pin}(n)$ -bundle. In particular, if M is a smooth manifold, a Pin structure on M is a Pin structure on the tangent bundle TM .

Definition 2.3. An *isomorphism* between two Pin structures on a bundle $E \rightarrow B$ is an isomorphism of principal $\text{Pin}(n)$ -bundles commuting with the projections to $\text{Fr}(E)$.

If E is oriented, then the choice of a Pin structure on E is equivalent to that of a Spin structure. However, Pin structures can be defined without choosing orientations, and they can exist on non-orientable bundles.

Proposition 2.4. ([18, p.185-186]) *A vector bundle $E \rightarrow B$ admits a Pin structure if and only if $w_2(E) = 0$. If one exists, the set of Pin structures on E up to isomorphism forms an affine space over $H^1(B; \mathbb{Z}/2)$.*

Thus, given a Pin structure $P^\#$ on E and an element $\eta \in H^1(B; \mathbb{Z}/2)$, we can twist $P^\#$ by η and obtain a new Pin structure on E (up to isomorphism), denoted $P^\# \otimes \eta$.

A concept closely related to isomorphism is that of homotopy.

Definition 2.5. A *homotopy* between two Pin structures $P_0^\#, P_1^\#$ on a bundle $E \rightarrow B$ consists of a Pin structure on the pullback of E to $B \times [0, 1]$, whose restriction to the two ends agrees with $P_0^\#$ and $P_1^\#$, respectively. Two homotopies are *equivalent* if the corresponding Pin structure over the product of B with the boundary of $[0, 1] \times [0, 1]$ extends to a Pin structure on the interior.

Remark 2.6. Alternatively, one may formulate these notions as follows: A Pin structure on a vector bundle E over B is a lift of a representative classifying map $B \rightarrow \mathrm{BO}(n)$ to $\mathrm{BPin}(n)$. A homotopy of Pin structures is given by a homotopy between the lifts (through other such lifts), and the notion of equivalence corresponds to a homotopy of homotopies.

The classification of Pin structures up to homotopy is the same as up to isomorphism.

Lemma 2.7. *Two Pin structures on a vector bundle $E \rightarrow B$ are isomorphic if and only if they are homotopic.*

Proof. If the two Pin structures are homotopic, then after choosing a Pin connection in the bundle over $B \times [0, 1]$ we get an isomorphism by parallel transport. Conversely, if we have an isomorphism, we can construct a Pin bundle over $B \times [0, 1]$ by gluing two trivial bundles over $B \times [0, 2/3]$ and $B \times (1/3, 1]$ using that isomorphism. \square

The advantage of talking about homotopies is that we can iterate them. We can talk about homotopies of homotopies (what we called equivalences), and also about higher homotopies. Using this data we can construct a simplicial set whose vertices are Pin structures on E , and whose n -simplices correspond to n -homotopies. We call this the *space of Pin structures on E* .

Remark 2.8. Up to homotopy equivalence, the space of Pin structures on a vector bundle E does not depend on the inner product on E . For simplicity, we will suppress the inner products from our future discussions.

Proposition 2.9. *Let $E \rightarrow B$ be a vector bundle that admits a Pin structure. Then the space of such structures is (non-canonically) homotopy equivalent to the mapping space $\mathrm{Map}(B, \mathbb{RP}^\infty)$.*

Proof. We think of Pin structures in terms of maps to the classifying space, as in Remark 2.6. The fibration

$$\mathrm{Map}(B, \mathrm{B}\mathbb{Z}/2) \rightarrow \mathrm{Map}(B, \mathrm{BPin}(n)) \rightarrow \mathrm{Map}(B, \mathrm{BO}(n))$$

gives the desired result, using the fact that $\mathrm{B}\mathbb{Z}/2 = \mathbb{RP}^\infty$. \square

The above result has various generalizations one of which we will use: given a point $b \in B$, the space of Pin structures on E which are fixed at b is homotopy equivalent to the space of based maps from B to $\mathrm{BO}(n)$.

Example 2.10. Let M be the circle S^1 , equipped with either orientation. It is well-known that M admits two isomorphism classes of Spin structures:

- the *bounding* one, which is obtained by restricting a Spin structure on D^2 , and hence represents the zero element in the Spin bordism group $\Omega_{\mathrm{Spin}}^1 \cong \mathbb{Z}/2$. In this case the oriented frame bundle of TS^1 is just S^1 itself, and this Spin structure corresponds to its nontrivial (connected) double cover;

- the *Lie group* one, which is obtained by choosing a Spin structure on the tangent space to S^1 at a single point, and extending it to the whole circle in an S^1 -equivariant way. This represents the nontrivial element in Ω_{Spin}^1 , and corresponds to the trivial double cover of S^1 .

See for example [17, p.35-36].

We will use the same terminology (*bounding* and *Lie group*) for the corresponding isomorphism classes of Pin structures on the circle, which are independent of the orientation. Note, however, that this is a slight misnomer: the Pin bordism group Ω_{Pin}^1 is trivial, because a Lie group Pin structure also bounds (a Möbius band equipped with its own Pin structure).

Definition 2.11. Fix a splitting of the extension $\text{Pin}(1) \rightarrow \text{O}(1)$. This determines a canonical Pin structure on any one-dimensional vector space, and hence (by equivariant extension) a canonical representative of the Lie group Pin structures on the circle. When we are interested in a Pin structure on S^1 on the nose (not just up to isomorphism), this is what we call the *Lie group Pin structure*.

By contrast, one can show that the action of rotation on the space of Pin structures on D^2 generates a non-trivial loop, so that there is no natural choice of bounding Pin structure on S^1 .

Remark 2.12. In Seidel's book [40], a Lie group Pin structure on the circle is called trivial, and a bounding one nontrivial. We will not use this terminology here, to prevent confusion—since either kind of structure can be considered trivial from a different perspective: that of bordism, or that of equivariant trivializations.

Example 2.13. Let V be a vector space, viewed as a vector bundle over a point. Then, there are non-equivalent homotopies between Pin structures on V . The simplest case to consider is when the two Pin structures that we are comparing are the same; in that case, such isomorphisms correspond bijectively to Pin structures on the circle which are fixed at a point, as can be seen by gluing the endpoints of the interval. In fact, by Proposition 2.9, the space of Pin structures on V is (non-canonically) homotopy equivalent to $\text{B}\mathbb{Z}/2 = \mathbb{R}\mathbb{P}^\infty$, and $\pi_1(\mathbb{R}\mathbb{P}^\infty) = \mathbb{Z}/2$.

Example 2.14. On a torus T^n we will be interested in the *Lie group Pin structures*, those that are equivariant with respect to the Lie group multiplication on T^n . The space of such structures is the same as that of Pin structures on the tangent space at any point, and thus homotopy equivalent to $\mathbb{R}\mathbb{P}^\infty$. Unlike in the case $n = 1$ considered in Example 2.10 and Definition 2.11, for $n \geq 2$ we do not have a canonical representative anymore.

Remark 2.15. One disadvantage of Pin (as opposed to Spin) structures is that they do not behave well with respect to direct sums. Indeed, suppose that E and E' are two vector bundles over the same base B . Then

$$w_2(E \oplus E') = w_2(E) + w_2(E') + w_1(E)w_1(E').$$

Thus, if E and E' are Spin (have $w_1 = w_2 = 0$) then so is $E \oplus E'$, and in fact we get an induced Spin structure on $E \oplus E'$. By contrast, if E and E' are Pin (have $w_2 = 0$), this does not guarantee that $E \oplus E'$ is Pin.

The source of this issue is the lack of a natural multiplication map $\text{Pin}(n) \times \text{Pin}(n') \rightarrow \text{Pin}(n + n')$. If we use the explicit description of the Pin groups as subsets of the Clifford algebras, we could try to define the map by

$$(v_1 \dots v_k, v'_1 \dots v'_{k'}) \mapsto v_1 \dots v_k v'_1 \dots v'_{k'}.$$

However, the left hand side is equal to

$$(1, v'_1 \dots v'_{k'}) \cdot (v_1 \dots v_k, 1) \in \text{Pin}(n) \times \text{Pin}(n')$$

which should be mapped to $v'_1 \dots v'_{k'} v_1 \dots v_k \in \text{Pin}(n + n')$. Thus, we run into the problem that the elements $v_1 \dots v_k$ and $v'_1 \dots v'_{k'}$ commute in $\text{Pin}(n + n')$ only up to a sign $(-1)^{kk'}$. On the other hand, there is a natural multiplication $\text{Spin}(n) \times \text{Spin}(n') \rightarrow \text{Spin}(n + n')$, since in that case we only work with even values of k and k' .

2.2. Coupled orientations. The notion of coupled orientation that we introduce here will play a role in our take on Lagrangian Floer homology in Section 3.2. We will then come back to it in Section 4.1.

Definition 2.16. For $n, m \geq 0$, we define the *coupled special orthogonal group*¹ $\text{SO}(n; m)$ to be the group of pairs of orthogonal matrices with the same determinant:

$$\text{SO}(n; m) = \{(A, B) \in \text{O}(n) \times \text{O}(m) \mid \det(A) = \det(B)\}.$$

Definition 2.17. Let E, F be two real vector bundles (with inner products) over the same base B , of ranks n and m respectively. A *coupled orientation* on (E, F) is a lift of $\text{Fr}(E) \times \text{Fr}(F) \rightarrow B$ from a principal $\text{O}(n) \times \text{O}(m)$ bundle to a principal $\text{SO}(n, m)$ bundle.

Lemma 2.18. *A coupled orientation on (E, F) is equivalent to an orientation on the direct sum $E \oplus F$. It exists if and only if $w_1(E) = w_1(F)$.*

Proof. For the first part, observe that there is a pull-back diagram

$$\begin{array}{ccc} \text{SO}(n; m) & \hookrightarrow & \text{SO}(n + m) \\ \downarrow & & \downarrow \\ \text{O}(n) \times \text{O}(m) & \hookrightarrow & \text{O}(n + m) \end{array}$$

Thus, lifting a principal $\text{O}(n) \times \text{O}(m)$ bundle to $\text{SO}(n; m)$ is the same as lifting the induced $\text{O}(n + m)$ bundle to $\text{SO}(n + m)$. The latter is the data of an orientation on the direct sum.

The second part follows from the relation $w_1(E \oplus F) = w_1(E) + w_1(F)$. \square

Definition 2.19. Let E, F, E', F' be four real vector bundles over a base B . Given coupled orientations on (E, F) and (E', F') , we can form their *direct sum*. This is the coupled orientation on $(E \oplus E', F \oplus F')$ induced by the map

$$(1) \quad \text{SO}(n; m) \times \text{SO}(n', m') \rightarrow \text{SO}(n + n'; m + m'),$$

$$((A, B), (C, D)) \mapsto \left(\left(\begin{array}{c|c} A & 0 \\ \hline 0 & C \end{array} \right), \left(\begin{array}{c|c} B & 0 \\ \hline 0 & D \end{array} \right) \right).$$

Definition 2.20. For any vector bundle E , there is a *canonical coupled orientation* on (E, E) , induced from the natural map

$$(2) \quad \text{O}(n) \rightarrow \text{SO}(n; n), \quad A \mapsto (A, A).$$

Lemma 2.21. *Let E and E' be two vector bundles over the same base B . Equip (E, E) and (E', E') with their canonical coupled orientations, then take their direct sum. The result is the canonical coupled orientation on $(E \oplus E', E \oplus E')$.*

¹Warning: the similar notation $\text{SO}(n, m)$ is often used in the literature to denote a different group, that of transformations that preserve an indefinite bilinear form of signature (n, m) and have determinant one.

Proof. This is a consequence of the commutativity of the diagram

$$\begin{array}{ccc} \mathrm{O}(n) \times \mathrm{O}(n') & \longrightarrow & \mathrm{SO}(n; n) \times \mathrm{SO}(n'; n') \\ \downarrow & & \downarrow \\ \mathrm{O}(n + n') & \longrightarrow & \mathrm{SO}(n + n'; n + n') \end{array}$$

which involves maps of the form (1) and (2). \square

Observe that the equivalence in Lemma 2.18 is based on the map

$$(3) \quad \mathrm{O}(n) \times \mathrm{O}(m) \rightarrow \mathrm{O}(n + m), \quad (A, B) \mapsto \left(\begin{array}{c|c} A & 0 \\ \hline 0 & B \end{array} \right).$$

Under this equivalence, the canonical coupled orientation on (E, E) corresponds to the orientation on $E \oplus E$ given by an ordered basis of the form

$$((v_1, 0), \dots, (v_n, 0), (0, v_1), \dots, (0, v_n)),$$

where $\{v_1, \dots, v_n\}$ is any basis of E . Let us call this the *concatenated orientation* on $E \oplus E$.

There is another natural choice of orientation on $E \oplus E$, from bases of the form

$$((v_1, 0), (0, v_1), \dots, (v_n, 0), (0, v_n)).$$

We call this the *shuffled orientation*. It differs from the concatenated one by $(-1)^{n(n-1)/2}$. The shuffled orientation behaves better with respect to direct sums, in the following sense: Given shuffled orientations on $E \oplus E$ and $E' \oplus E'$, their direct sum is an orientation on $(E \oplus E) \oplus (E' \oplus E')$ which corresponds to the shuffled orientation on $(E \oplus E') \oplus (E \oplus E')$ under the isomorphism that interchanges the second and third summands.

To make everything be compatible with direct sums, it is desirable to have the canonical coupled orientation on (E, E) correspond to the shuffled orientation on $E \oplus E$. We can achieve this by adjusting the map (3) in the particular case where $n = m$. In that case we consider the map

$$(4) \quad \mathrm{O}(n) \times \mathrm{O}(n) \rightarrow \mathrm{O}(2n), \quad (A, B) \mapsto C^{-1} \left(\begin{array}{c|c} A & 0 \\ \hline 0 & B \end{array} \right) C,$$

where C is the permutation matrix taking the standard basis $(e_1, e_2, \dots, e_{2n-1}, e_{2n})$ to the basis $(e_1, e_{n+1}, e_2, e_{n+2}, \dots, e_n, e_{2n})$.

Convention 2.22. In this paper we will only use coupled orientations on pairs (E, F) , where E and F have the same rank. When doing so, to go from a coupled orientation on (E, F) to an orientation on $E \oplus F$ we will use the map (4) instead of (3). Therefore, the canonical coupled orientation on a pair (E, E) will correspond to the shuffled orientation on $E \oplus E$.

So far we have discussed coupled orientations on vector bundles. Let us turn to the analogous notion for manifolds.

Definition 2.23. Let M_0 and M_1 be two smooth manifolds. For $i = 0, 1$, let $\pi_i : M_0 \times M_1 \rightarrow M_i$ be the projection. A *coupled orientation* on (M_0, M_1) is a coupled orientation on the pair of vector bundles $(\pi_0^*TM_0, \pi_1^*TM_1)$ on $M_0 \times M_1$.

Remark 2.24. By Lemma 2.18, a coupled orientation on (M_0, M_1) corresponds to an orientation on the bundle

$$T(M_0 \times M_1) = \pi_0^*TM_0 \oplus \pi_1^*TM_1.$$

In other words, we can think of a coupled orientation on (M_0, M_1) as an orientation of the product $M_0 \times M_1$. We will only use this notion when $\dim M_0 = \dim M_1$; then, the correspondence to the orientation on $M_0 \times M_1$ is defined using the map (4), as in Convention 2.22.

2.3. Coupled Spin structures. We now define some notions that will be used in Section 4 and later.

Definition 2.25. We introduce the group $\widetilde{\text{Spin}}(n; m)$ to be the $(\mathbb{Z}/2 \times \mathbb{Z}/2)$ -cover of $\text{SO}(n; m)$ given as the pull-back

$$(5) \quad \begin{array}{ccc} \widetilde{\text{Spin}}(n; m) & \hookrightarrow & \text{Pin}(n) \times \text{Pin}(m) \\ \downarrow & & \downarrow \\ \text{SO}(n, m) & \hookrightarrow & \text{O}(n) \times \text{O}(m) \end{array}$$

We let the *coupled Spin group* $\text{Spin}(n; m)$ be the quotient of $\widetilde{\text{Spin}}(n; m)$ by the diagonal subgroup $\mathbb{Z}/2 \subset \mathbb{Z}/2 \times \mathbb{Z}/2$.

Explicitly, $\widetilde{\text{Spin}}(n; m)$ is the subgroup of $\text{Pin}(n) \times \text{Pin}(m)$ consisting of pairs of the form $(v_1 \dots v_k, u_1 \dots u_l)$ with $k + l$ even, and $\text{Spin}(n; m)$ is its quotient by the equivalence relation

$$(v_1 \dots v_k, u_1 \dots u_l) \sim (-v_1 \dots v_k, -u_1 \dots u_l).$$

Definition 2.26. Let E, F be two real vector bundles (with inner products) over the same base B , of ranks n and m respectively. A *Spin structure* on (E, F) is a lift of $\text{Fr}(E) \times \text{Fr}(F) \rightarrow B$ to a principal $\widetilde{\text{Spin}}(n, m)$ bundle. A *coupled Spin structure* on (E, F) is a lift of $\text{Fr}(E) \times \text{Fr}(F) \rightarrow B$ to principal $\text{Spin}(n, m)$ bundle.

Lemma 2.27. *A coupled Spin structure on a pair (E, F) exists if and only if $w_1(E) = w_1(F)$ and $w_2(E) = w_2(F)$. If one exists, then the space of such coupled Spin structures is (non-canonically) homotopy equivalent to $\text{Map}(B, \mathbb{R}\mathbb{P}^\infty)$. In particular, isomorphism classes of coupled Spin structures form a torsor over $H^1(B; \mathbb{Z}/2)$.*

Proof. The pull-back diagram (5) shows that a $\widetilde{\text{Spin}}$ structure is the same as the data of a coupled orientation together with Pin structures on the two bundles. From Lemma 2.18 and the definition of the Pin group, a $\widetilde{\text{Spin}}$ structure exists if and only if $w_1(E) = w_1(F)$ and $w_2(E) = w_2(F) = 0$.

In fact, the same diagram shows that we can characterize $\widetilde{\text{Spin}}(n; m)$ as the central extension of $\text{SO}(n; m)$ by $\mathbb{Z}/2 \times \mathbb{Z}/2$ determined by the pull-back of the element $\pi_1^* w_2 \oplus \pi_2^* w_2$ from

$$H^2(\text{BO}(n) \times \text{BO}(m); \mathbb{Z}/2 \times \mathbb{Z}/2) \cong H^2(\text{BO}(n) \times \text{BO}(m); \mathbb{Z}/2) \oplus H^2(\text{BO}(n) \times \text{BO}(m); \mathbb{Z}/2)$$

to $H^2(\text{BSO}(n; m); \mathbb{Z}/2 \times \mathbb{Z}/2)$, where π_1 and π_2 are the projections from $\text{BO}(n) \times \text{BO}(m)$ to the two factors. Dividing $\widetilde{\text{Spin}}(n; m)$ by the diagonal $\mathbb{Z}/2$ action gives $\text{Spin}(n; m)$ as the central extension of $\text{SO}(n; m)$ by $\mathbb{Z}/2$ determined by the pull-back of

$$\pi_1^* w_2 + \pi_2^* w_2 \in H^2(\text{BO}(n) \times \text{BO}(m); \mathbb{Z}/2)$$

to $H^2(\text{BSO}(n; m); \mathbb{Z}/2)$. This says that, once we have a coupled orientation (which exists iff $w_1(E) = w_1(F)$), the remaining obstruction to the existence of a coupled Spin structure is $w_2(E) + w_2(F)$.

The classification statement about $\text{Map}(B, \mathbb{R}\mathbb{P}^\infty)$ follows from the same argument as in the proof of Proposition 2.9. Furthermore, isomorphism is equivalent to homotopy (by the same argument as in Lemma 2.7), and we have $\pi_0(\text{Map}(B, \mathbb{R}\mathbb{P}^\infty)) = H^1(B; \mathbb{Z}/2)$. \square

Remark 2.28. We will mostly be interested in coupled Spin structures on a pair (V, W) of vector spaces, viewed as vector bundles over a point. Suppose that the pair (V, W) is equipped with a coupled orientation. Then, Pin structures on the two vector spaces determine a $\widetilde{\text{Spin}}$ structure and hence (in a non-injective way) a coupled Spin structure on (V, W) . If we fix the isomorphism class of the Pin structure on V , the space of its representatives is homotopy equivalent to $\mathbb{R}\mathbb{P}^\infty$; the same goes for W . The map to coupled Spin structures is modeled on the multiplication $\mathbb{R}\mathbb{P}^\infty \times \mathbb{R}\mathbb{P}^\infty \rightarrow \mathbb{R}\mathbb{P}^\infty$ which comes from viewing $\mathbb{R}\mathbb{P}^\infty$ as $B\mathbb{Z}/2$. In particular, suppose we have a loop of Pin structures on (V, W) that gives the standard generator of $\pi_1(\mathbb{R}\mathbb{P}^\infty) \cong \mathbb{Z}/2$ for both V and W . Then, the induced loop in the space of coupled Spin structures is trivial, representing $1 + 1 = 0$ in $\pi_1(\mathbb{R}\mathbb{P}^\infty) \cong \mathbb{Z}/2$.

Remark 2.29. We also have a commutative diagram

$$(6) \quad \begin{array}{ccc} \text{Spin}(n) \times \text{Spin}(m) & \longrightarrow & \widetilde{\text{Spin}}(n; m) \\ \downarrow & & \downarrow \\ \text{SO}(n) \times \text{SO}(m) & \longrightarrow & \text{SO}(n, m). \end{array}$$

This implies that Spin structures on two bundles E and F determine a $\widetilde{\text{Spin}}$ and hence a coupled Spin structure of (E, F) .

Given a coupled orientation, Lemma 2.27 and Remark 2.28 show that coupled Spin structures are a weaker condition than having Pin structures on each bundle. However, one advantage they have over Pin structures is that they behave well with respect to direct sums:

Lemma 2.30. *Let E, F, E', F' be four real vector bundles over a base B . Given coupled Spin structures on (E, F) and (E', F') , there is an induced coupled Spin structure on $(E \oplus E', F \oplus F')$.*

Proof. This follows from the existence of a natural multiplication map

$$\text{Spin}(n, m) \times \text{Spin}(n', m') \rightarrow \text{Spin}(n + n', m + m')$$

as follows. If we denote the elements of $\text{Pin}(n)$ by $v_1 \dots v_k$, and those of $\text{Pin}(m)$ by $u_1 \dots u_l$, then the elements of $\text{Spin}(n, m)$ are equivalence classes of pairs $[(v_1 \dots v_k, u_1 \dots u_l)]$ with $k + l$ even. Similarly, the elements of $\text{Spin}(n', m')$ are equivalence classes $[(v'_1 \dots v'_{k'}, u'_1 \dots u'_{l'})]$ with $k' + l'$ even. We let the multiplication be

$$([(v_1 \dots v_k, u_1 \dots u_l)], [(v'_1 \dots v'_{k'}, u'_1 \dots u'_{l'})]) \mapsto [(v_1 \dots v_k v'_1 \dots v'_{k'}, u_1 \dots u_l u'_1 \dots u'_{l'})].$$

This lands in the right place because $k + k' + l + l'$ is even. Furthermore, we do not encounter the same commutativity problem as the one for Pin discussed at the end of Section 2.1. In the current setting, when we commute the two factors, the result is

$$[(v'_1 \dots v'_{k'} v_1 \dots v_k, u'_1 \dots u'_{l'} u_1 \dots u_l)] = [((-1)^{kk'} v_1 \dots v_k v'_1 \dots v'_{k'}, (-1)^{ll'} u_1 \dots u_l u'_1 \dots u'_{l'})].$$

The right hand side is same equivalence class as $[(v_1 \dots v_k v'_1 \dots v'_{k'}, u_1 \dots u_l u'_1 \dots u'_{l'})]$. Indeed, we have $(-1)^{kk'} = (-1)^{ll'}$ because $k \equiv k' \pmod{2}$ and $l \equiv l' \pmod{2}$. \square

We end with a few more properties of coupled Spin structures.

Lemma 2.31. *Given any vector bundle E , the pair (E, E) has a canonical coupled Spin structure.*

Proof. This comes from the existence of a natural map $O(n) \rightarrow \text{Spin}(n, n)$ defined as follows. We view $O(n)$ as $\text{Pin}(n)/\pm 1$. Given an element $v_1 \dots v_k \in \text{Pin}(n)$, we map its image in $O(n)$ to

$$[(v_1 \dots v_k, v_1 \dots v_k)] \in \text{Spin}(n, n).$$

Observe that changing $v_1 \dots v_k$ by a sign produces the same equivalence class in $\text{Spin}(n, n)$. \square

Lemma 2.32. *Given vector bundles E, F and S over the same base B , a coupled Spin structure on (E, F) naturally induces one on $(E \oplus S, F \oplus S)$.*

Proof. This follows from Lemmas 2.30 and 2.32. \square

Lemma 2.33. *Let $f : E \rightarrow F$ be a bundle map between vector bundles over the same base B (covering the identity on B). Suppose f has fixed rank, so that $K = \ker(f)$ and $C = \text{coker}(f)$ form vector bundles over B . Then, a coupled Spin structure on (C, K) induces one on (F, E) .*

Proof. Recall that our vector bundles are implicitly equipped with inner products; see Remark 2.8. Therefore, if we let $I = \text{im}(f)$ then by taking orthogonal complements we get isomorphisms $E \cong K \oplus I$ and $F \cong C \oplus I$. The conclusion follows from Lemma 2.32. \square

Lemma 2.34. *Let (E, ω) be a symplectic vector bundle, and $L_0, L_1 \subset E$ be Lagrangian sub-bundles such that $L_0 \cap L_1$ is also a sub-bundle (i.e., has fixed rank). Then:*

- (1) *An inner product on E induces a canonical bundle isomorphism $\tau_{L_0, L_1} : L_0 \rightarrow L_1$;*
- (2) *There is a canonical coupled Spin structure on (L_0, L_1) .*

Proof. (1) Let $(L_0 \cap L_1)^\omega$ be the symplectic complement to $L_0 \cap L_1$ in E . Then $E' = (L_0 \cap L_1)^\omega / (L_0 \cap L_1)$ is symplectic, and inside it we have Lagrangian sub-bundles $L'_0 = L_0 / (L_0 \cap L_1)$ and $L'_1 = L_1 / (L_0 \cap L_1)$. Using the inner product we can identify each quotient bundle with the corresponding orthogonal complement. Furthermore, since L'_0 and L'_1 intersect in the zero section, the symplectic form gives an identification between L'_0 and the dual to L'_1 . The inner product then gives an identification between L'_0 and L'_1 . Taking the direct sum of this with the identity on $L_0 \cap L_1$ we get the desired identification between L_0 and L_1 .

(2) Apply part (a) and Lemma 2.31. \square

Remark 2.35. In the situation of Lemma 2.34 (1), note that τ_{L_1, L_0} is not the inverse to τ_{L_0, L_1} . For example, when L_0 and L_1 are transverse lines in \mathbb{R}^2 , the isomorphism τ_{L_0, L_1} is counterclockwise rotation, and its composition with τ_{L_1, L_0} is $-\text{id}$ on L_0 .

Remark 2.36. On the other hand, the isomorphisms from Lemma 2.34 (1) behave well with respect to direct sums. If $L_0, L_1 \subset E$ are as in the lemma, and $L'_0, L'_1 \subset E'$ is another pair of the same kind, then

$$\tau_{L_0 \oplus L'_0, L_1 \oplus L'_1} = \tau_{L_0, L_1} \oplus \tau_{L'_0, L'_1}.$$

This implies that if we take the direct sum of the canonical coupled Spin structures on (L_0, L_1) and (L'_0, L'_1) , we get the canonical coupled Spin structure on $(L_0 \oplus L'_0, L_1 \oplus L'_1)$.

3. CANONICAL ORIENTATIONS IN LAGRANGIAN FLOER HOMOLOGY

In this section we explain how Pin structures produce canonical orientations on the moduli spaces appearing in Lagrangian Floer homology.

Convention 3.1. Our exposition is inspired from Seidel's book [40], but we will use the usual conventions in Heegaard Floer theory [30]. Specifically, we work with Floer homology instead of Floer cohomology, and when we talk about polygon maps, the Lagrangian boundary conditions are ordered clockwise rather than counterclockwise. Our Floer chain groups $CF_*(L_0, L_1)$ correspond to Seidel's Floer cochain groups $CF^{n-*}(L_1, L_0)$.

3.1. Orientation spaces. We review here the definition of orientation spaces associated to linear Lagrangian branes, following [40, Section 11]. We describe a slight modification that is convenient for our purposes.

Given a $(2n)$ -dimensional symplectic vector space V , we let $\text{Gr}(V)$ be the Lagrangian Grassmannian, whose points are the Lagrangian subspaces of V . In [40, Section 11h], Seidel considers a natural map to a product of Eilenberg-MacLane spaces

$$(7) \quad (\mu, w_2) : \text{Gr}(V) \rightarrow K(\mathbb{Z}, 1) \times K(\mathbb{Z}/2, 2),$$

and pulls back the product of the universal fibrations to obtain a $\mathbb{Z} \times \mathbb{R}\mathbb{P}^\infty$ bundle over $\text{Gr}(V)$, denoted $\text{Gr}^\#(V)$. The points $\Lambda^\# \in \text{Gr}^\#(V)$ are called *abstract linear Lagrangian branes*.

The first component of (7),

$$\mu : \text{Gr}(V) \rightarrow K(\mathbb{Z}, 1) \cong S^1$$

can be described explicitly as the squared phase map associated to a compatible complex structure I_V and a quadratic complex volume form on (V, I_V) ; see [40, Section 11j].

In our setting we shall mostly focus on the second component

$$w_2 : \text{Gr}(V) \rightarrow K(\mathbb{Z}/2, 2).$$

We let $\text{Gr}^\dagger(V)$ be the pullback of the universal fibration, which is an $\mathbb{R}\mathbb{P}^\infty$ bundle over $\text{Gr}(V)$. Furthermore, $\text{Gr}^\#(V)$ is a \mathbb{Z} -cover of $\text{Gr}^\dagger(V)$.

An alternate description of $\text{Gr}^\dagger(V)$ is as follows. Let

$$i_V : \text{Gr}(V) \rightarrow \text{BO}(n)$$

be the composition of the forgetful map from $\text{Gr}(V)$ to the ordinary Grassmannian of n -planes in V , with the map to $\text{BO}(n)$ induced by some inclusion $V \hookrightarrow \mathbb{R}^\infty$. Then, $\text{Gr}^\dagger(V)$ is the pullback of the bundle $\text{BPin}(n) \rightarrow \text{BO}(n)$ under i_V ; compare [40, p.60]. Thus, we can think of a point $\Lambda^\dagger \in \text{Gr}^\dagger(V)$ as a Pin structure on the underlying Lagrangian subspace $\Lambda \in \text{Gr}(V)$; i.e., a principal homogeneous $\text{Pin}(n)$ -space equipped with an isomorphism $P^\# \times_{\text{Pin}(n)} \mathbb{R}^n \cong \Lambda$. This interpretation is particularly useful in families: a family of Lagrangian subspaces in V parametrized by a base space B (i.e., a Lagrangian bundle over B) is described by a map $B \rightarrow \text{Gr}(V)$, and a lift of this map to $\text{Gr}^\dagger(V)$ is the same as a Pin structure on the Lagrangian bundle.

In the notation of [40, Section 11j], an element of $\text{Gr}^\#(V)$ is a triple

$$\Lambda^\# = (\Lambda, \alpha, P^\#),$$

where $\Lambda \in \text{Gr}(V)$, $\alpha \in \mathbb{R}$ is such that $e^{2\pi i \alpha} = \mu(\Lambda) \in S^1$, and $P^\#$ is a Pin structure on Λ . The pair $\Lambda^\dagger = (\Lambda, P^\#)$ is the projection of $\Lambda^\#$ to $\text{Gr}^\dagger(V)$.

Next, consider two points $\Lambda_0^\#, \Lambda_1^\# \in \text{Gr}^\#(V)$, such that the underlying Lagrangians $\Lambda_0, \Lambda_1 \subset V$ intersect transversely. Choose a path


$$\rho^\# = (\Lambda_t^\#)_{t \in [0,1]}$$

connecting the two points, and let $\rho = (\Lambda_t)_{t \in [0,1]}$ be its projection to $\text{Gr}(V)$. There is a technical condition we need to impose on ρ , that it has negative definite crossing form at

$t = 1$ with the constant path Λ_1 ; see [40, Section 11g]. Then, the Maslov index of the path ρ depends only on $\Lambda_0^\#$ and $\Lambda_1^\#$, and is denoted

$$i(\Lambda_0^\#, \Lambda_1^\#) \in \mathbb{Z}.$$

Furthermore, one can consider a Fredholm problem on a capped half-infinite strip H , with moving Lagrangian boundary conditions $\{\Lambda_t\}$ that are constant near the infinite end, as in [40, Section 11g] or [9, p.143], but reflected to be in line with our Conventions 3.1.:



We call H a *cap*. The $\bar{\partial}$ operator on H associated to this set-up is denoted $\bar{\partial}_\rho$. Consider its determinant index bundle

$$\det(\bar{\partial}_\rho) = \lambda^{\text{top}}((\text{coker } \bar{\partial}_\rho)^*) \otimes \lambda^{\text{top}}(\ker \bar{\partial}_\rho).$$

The *orientation space*

$$o(\Lambda_0^\#, \Lambda_1^\#)$$

is the abelian group generated by the two orientations $\omega, \bar{\omega}$ on $\det(\bar{\partial}_\rho)$, modulo the relation $\bar{\omega} = -\omega$. Since this is a free abelian group of rank-1, it is non-canonically isomorphic to \mathbb{Z} .

Remark 3.2. Rotating our cap by 180° , we see that it is equivalent to the cap in Seidel's book with boundary conditions given by the reflected path Λ_{1-t} from Λ_1 to Λ_0 . By Equation (11.27) in [40], we have

$$(8) \quad \text{index}(\det(\bar{\partial}_\rho)) = i(\Lambda_1^\#, \Lambda_0^\#) = n - i(\Lambda_0^\#, \Lambda_1^\#).$$

Also, the orientation space $o(\Lambda_0^\#, \Lambda_1^\#)$ is canonically isomorphic to the one in Seidel's book tensored with $\lambda^{\text{top}}\Lambda_t$ (for any t). Note that, compared with [40], we have not changed the definition of the Maslov index $i(\Lambda_0^\#, \Lambda_1^\#)$; but we have changed the definition of $o(\Lambda_0^\#, \Lambda_1^\#)$. We have also changed the meaning of the notation H : in [40], our cap is denoted \bar{H} , and its reflection across a vertical axis is called H .

Lemma 3.3. *The orientation space*

$$o(\Lambda_0^\#, \Lambda_1^\#)$$

depends, up to canonical isomorphism, only on $\Lambda_0^\#$ and $\Lambda_1^\#$ (in particular, not on the choice of path $\rho^\#$ with these endpoints).

Proof. This is discussed in [40, p.164]. The essential point in the following: the space of paths in $\text{Gr}(V)$, with endpoints Λ_0 and Λ_1 , has \mathbb{Z} -many components, each of which has fundamental group $\mathbb{Z}/2$. The orientation space of the operator $\det(\bar{\partial}_\rho)$ associated to paths ρ defines a non-trivial local system over this space, as shown by de Silva in [41]. Imposing the grading condition distinguishes a homotopy class of paths. More importantly, considering paths of Lagrangians equipped with Pin structures yields a simply connected space of paths, which can be thought of as a 2-fold cover of the space of paths in $\text{Gr}(V)$. Its components over a fixed path in $\text{Gr}(V)$ are distinguished by the homotopy class of the path of Pin structures relative the endpoints. Thus, parallel transport through paths of paths in $\text{Gr}^\#(V)$ yields isomorphisms of determinant lines that are independent of all choices. \square

Lemma 3.4. *The local system formed by the orientation spaces $o(\Lambda_0^\#, \Lambda_1^\#)$ over the space of all Pin structures on Λ_1 is non-trivial (and similarly for Λ_0).*

Proof. Recall that the space of Pin structures on Λ_1 has the homotopy type of $\mathbb{R}\mathbb{P}^\infty$, and thus $\pi_1 = \mathbb{Z}/2$; see Example 2.13. Consider a homotopically non-trivial loop $\{P_{1,s}^\#\}_{s \in [0,1]}$ of Pin structures on Λ_1 , restricting to $P_1^\#$ at $t \in \{0,1\}$. Concatenating this with the path (Λ_t) gives a family of boundary conditions for the cap, parametrized by $s \in [0,1]$. There is a corresponding family of Cauchy-Riemann operators. Appealing once again to the non-triviality of $\det(\bar{\partial}_\rho)$ over the space of paths in $\text{Gr}(V)$ [41], we conclude that the monodromy on the orientation space is -1 . Compare [40, Remark 11.19]. \square

The following result is ultimately derived from the fact that every even index loop in the Grassmanian of Lagrangians is homotopic to a loop in the image of the action by the unitary group; that equips the orientation operator with a trivialization coming from deforming the corresponding Fredholm problem to one which is complex linear. This will not be apparent in our proof, which uses instead formal properties established in [40].

Proposition 3.5. *Suppose we are given $\Lambda_0^\dagger, \Lambda_1^\dagger \in \text{Gr}^\dagger(V)$, and a value $\epsilon \in \mathbb{Z}/2$. Pick lifts $\Lambda_0^\#, \Lambda_1^\# \in \text{Gr}^\#(V)$ of $\Lambda_0^\dagger, \Lambda_1^\dagger$ so that*

$$(9) \quad \epsilon = i(\Lambda_0^\#, \Lambda_1^\#) \pmod{2}.$$

Then, the orientation spaces $o(\Lambda_0^\#, \Lambda_1^\#)$ are canonically isomorphic, for all choices of $\Lambda_0^\#, \Lambda_1^\#$ satisfying (9).

Proof. A shift operation S on the elements $\Lambda^\# = (\Lambda, \alpha, P^\#) \in \text{Gr}^\#(V)$ is introduced in [40, Section 11k]:

$$S\Lambda^\# = (\Lambda, \alpha - 1, P^\# \otimes \lambda^{\text{top}}(\Lambda)).$$

Lemma 11.21 in [40] says that

$$(10) \quad i(\Lambda_0^\#, S\Lambda_1^\#) = i(\Lambda_0^\#, \Lambda_1^\#) - 1$$

and

$$(11) \quad o(\Lambda_0^\#, S\Lambda_1^\#) \cong o(\Lambda_0^\#, \Lambda_1^\#).$$

For $n = 1$, the isomorphism (11) is determined by clockwise rotation by π in the plane. In higher dimensions, it is determined by asking for it to be compatible with direct sums.

Similar arguments to those in the proof of Lemma 11.21 in [40] show that

$$(12) \quad i(S\Lambda_0^\#, \Lambda_1^\#) = i(\Lambda_0^\#, \Lambda_1^\#) + 1$$

and

$$(13) \quad o(S\Lambda_0^\#, \Lambda_1^\#) \cong o(\Lambda_0^\#, \Lambda_1^\#).$$

For us, a more relevant shift operation Σ is the one given by

$$\Sigma\Lambda^\# = (\Lambda, \alpha - 1, P^\#),$$

since this preserves the underlying $\Lambda^\dagger \in \text{Gr}^\dagger(V)$. Observe that the corresponding double shift is the same:

$$(14) \quad \Sigma \circ \Sigma = S \circ S.$$

The index does not depend on the Pin structure, so from (10) and (12) we deduce that

$$i(\Lambda_0^\#, \Lambda_1^\#) = i(\Lambda_0^\#, \Sigma\Lambda_1^\#) + 1 = i(\Sigma\Lambda_0^\#, \Lambda_1^\#) - 1.$$

Therefore, to establish what we need for the proposition, it suffices to find canonical isomorphisms

$$o(\Lambda_0^\#, \Lambda_1^\#) \cong o(\Lambda_0^\#, \Sigma^2 \Lambda_1^\#) \cong o(\Sigma^2 \Lambda_0^\#, \Lambda_1^\#) \cong o(\Sigma \Lambda_0^\#, \Sigma \Lambda_1^\#).$$

The first two isomorphisms follow from (11) and (13), together with (14). The final isomorphism, $o(\Lambda_0^\#, \Lambda_1^\#) \cong o(\Sigma \Lambda_0^\#, \Sigma \Lambda_1^\#)$, is clear because by using the family $\{\Sigma \Lambda_t^\#\}$ for the index problem on the cap H , the index problem does not change: the Lagrangian paths $\{\Lambda_t\}$ are the same. \square

In view of Proposition 3.5, we introduce the notation

$$(15) \quad o(\Lambda_0^\dagger, \Lambda_1^\dagger, \epsilon) := o(\Lambda_0^\#, \Lambda_1^\#)$$

for $\Lambda_0^\dagger, \Lambda_1^\dagger \in \text{Gr}^\dagger(V)$, and $\epsilon \in \mathbb{Z}/2$. Here, $\Lambda_0^\#$ and $\Lambda_1^\#$ are any lifts satisfying (9).

3.2. Lagrangian Floer homology. Lagrangian Floer homology over \mathbb{Z} is developed in [9], [10] using Spin structures on the Lagrangians. In Seidel's book [40], it was noted that only Pin structures are necessary, but the exposition there was under the assumption that the symplectic manifold M satisfies $2c_1(TM) = 0$ (so that the Floer complex is \mathbb{Z} -graded). In the context of Heegaard Floer theory, it will be convenient to use Pin structures, but the manifold does not have $2c_1(TM) = 0$. Nevertheless, since the only Lagrangians considered in Heegaard Floer theory are orientable, this setting admits a canonical relative $\mathbb{Z}/2$ grading. In fact, one can specify an absolute $\mathbb{Z}/2$ grading, which then suffices for the sign discussion in [40] to go through, using what we set up in Section 3.1.

We present here an adaptation of the techniques in [10, Chapter 8] and [40, Section 11] to the setting of interest to us.

Let (M, ω) be a closed symplectic manifold of dimension $2n$. Let $L_0, L_1 \subset M$ be two connected, transversely intersecting Lagrangians, such that L_0 and L_1 are orientable, and equipped with Pin structures $P_0^\#$ and $P_1^\#$. We do not choose orientations on the Lagrangians, but we will assume that they are coupled oriented, in the sense of Definition 2.23; i.e., that we have an orientation on $L_0 \times L_1$. (See Remark 2.24.)

Given a coupled orientation on (L_0, L_1) , every intersection point $\mathbf{x} \in L_0 \cap L_1$ admits an absolute mod 2 grading

$$\text{gr}(\mathbf{x}) \in \mathbb{Z}/2,$$

obtained as follows. We compare the orientation of $T_{\mathbf{x}}L_0 \oplus T_{\mathbf{x}}L_1 \cong T_{\mathbf{x},\mathbf{x}}(L_0 \times L_1)$ with the orientation of $T_{\mathbf{x}}M$ coming from ω^n . If they are the same, we set $\text{gr}(\mathbf{x}) = n \pmod{2}$. If they are different, we set $\text{gr}(\mathbf{x}) = n + 1 \pmod{2}$.

Given $\mathbf{x} \in L_0 \cap L_1$, in the notation of Section 3.1, we have elements

$$\Lambda_i^\dagger = (T_{\mathbf{x}}L_i, P_i^\#|_{T_{\mathbf{x}}L_i}) \in \text{Gr}^\dagger(T_{\mathbf{x}}M), \quad i = 0, 1.$$

We consider the orientation space

$$(16) \quad o(\mathbf{x}) := o(\Lambda_0^\dagger, \Lambda_1^\dagger, \text{gr}(\mathbf{x})).$$

We denote a generator of $o(\mathbf{x})$, i.e. an orientation of the corresponding determinant line, by $\omega_{\mathbf{x}}$.

The Lagrangian Floer complex is

$$CF_*(L_0, L_1) = \bigoplus_{\mathbf{x} \in L_0 \cap L_1} o(\mathbf{x}),$$

with the differential

$$(17) \quad \partial\omega_{\mathbf{x}} = \sum_{\mathbf{y} \in L_0 \cap L_1} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} (\#\widehat{\mathcal{M}}(\phi)) \cdot \omega_{\mathbf{y}}.$$

Here, $\pi_2(\mathbf{x}, \mathbf{y})$ denotes the space of relative homotopy classes between \mathbf{x} and \mathbf{y} (with boundary on the two Lagrangians), and $\mu(\phi)$ denotes the Maslov index of such a class. Furthermore, $\widehat{\mathcal{M}}(\phi) = \mathcal{M}(\phi)/\mathbb{R}$ is the count of J -holomorphic strips (solutions to Floer's equation) in the class ϕ , after dividing by translation by \mathbb{R} . It remains to explain what we mean by the signed count $\#\mathcal{M}(\phi) \in \mathbb{Z}$, which depends on the orientations $\omega_{\mathbf{x}} \in o(\mathbf{x})$ and $\omega_{\mathbf{y}} \in o(\mathbf{y})$.

Remark 3.6. In Heegaard Floer theory, one considers several variants of the Floer complex, keeping track of the intersections of the holomorphic disks with a (real codimension two) symplectic hypersurface $D \subset M$ which is disjoint from the Lagrangians. This additional complication does not affect our discussion of signs.

Remark 3.7. The coupled orientation on (L_0, L_1) makes the Lagrangian Floer complex absolutely $\mathbb{Z}/2$ -graded. In certain situations, the grading can be lifted to a relative or absolute $\mathbb{Z}/2N$ - or \mathbb{Z} -grading. Again, this issue is independent from our discussion of signs.

To define the signed count $\#\widehat{\mathcal{M}}(\phi)$, we need to orient the moduli space. Let

$$\mathbf{P}(L_0, L_1) = \{\gamma \in C^\infty([0, 1], M) \mid \gamma(0) \in L_0, \gamma(1) \in L_1\}$$

and let $\Omega(\phi)$ be the space of (non necessarily holomorphic) Whitney disks from \mathbf{x} to \mathbf{y} in the class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$; that is, the space of smooth paths in $\mathbf{P}(L_0, L_1)$ from the constant path $c_{\mathbf{x}}$ at \mathbf{x} to the constant path $c_{\mathbf{y}}$ at \mathbf{y} , in the class ϕ . By picking a homeomorphism from $[0, 1]$ to \mathbb{R} , we can view the moduli space $\mathcal{M}(\phi)$ as a subset of $\Omega(\phi)$.

The orientation on $\mathcal{M}(\phi)$ is governed by orienting the determinant index bundle $\det(D\bar{\partial})$ over $\Omega(\phi)$. Recall that the orientation spaces $o(\mathbf{x})$ and $o(\mathbf{y})$ also come from determinant index bundles, for operators on the cap H associated to paths of linear Lagrangian branes. We will denote these by $\bar{\partial}_{\mathbf{x}}$ and $\bar{\partial}_{\mathbf{y}}$, for convenience (instead of using the path of branes in the subscript, as we did in Section 3.1). Once we have an orientation on $\mathcal{M}(\phi)$, we get one on $\widehat{\mathcal{M}}(\phi) = \mathcal{M}(\phi)/\mathbb{R}$ using the ordered convention $\mathcal{M}(\phi) \cong \mathbb{R} \times \widehat{\mathcal{M}}(\phi)$.

To see that orientations $\omega_{\mathbf{x}} \in o(\mathbf{x})$ and $\omega_{\mathbf{y}} \in o(\mathbf{y})$ induce one on $\det(D\bar{\partial})$, it suffices to prove the following lemma.


Lemma 3.8. *The Pin structures and coupled orientation on the Lagrangians (L_0, L_1) induce a canonical isomorphism:*

$$(18) \quad \det(\bar{\partial}_{\mathbf{x}}) \otimes \det(D\bar{\partial}) \cong \det(\bar{\partial}_{\mathbf{y}}).$$

Proof. The arguments are as in Proposition 11.13 and Section (12b) in [40], with minor modifications. Given $u \in \Omega(\phi)$, we view it as a strip $u : [0, 1] \times \mathbb{R} \rightarrow M$, and trivialize u^*TM over its domain. Over each of the boundaries $\{0\} \times \mathbb{R}$ and $\{1\} \times \mathbb{R}$ we get a family of linear Lagrangians of the same symplectic vector space $V = u^*TM$. These families interpolate between $T_{\mathbf{x}}L_i$ and $T_{\mathbf{y}}L_i$, and are equipped with Pin structures.

After some deformations of Fredholm operators, we can glue together the operators $\bar{\partial}_{\mathbf{x}}$ and $D\bar{\partial}_u$ to obtain a new operator $\bar{\partial}_{\mathbf{y}}^{\text{new}}$ on the cap H , with boundary conditions interpolating

between $T_{\mathbf{y}}L_0$ and $T_{\mathbf{y}}L_1$:



Under the gluing operation, the indices are added and the determinant lines are tensored:

$$(19) \quad \text{index}(\bar{\partial}_{\mathbf{x}}) + \text{index}(D\bar{\partial}) = \text{index}(\bar{\partial}_{\mathbf{y}}^{\text{new}}),$$

$$(20) \quad \det(\bar{\partial}_{\mathbf{x}}) \otimes \det(D\bar{\partial}) \cong \det(\bar{\partial}_{\mathbf{y}}^{\text{new}}).$$

From the construction of the orientation spaces in Proposition 3.5, together with Equation (8), we know we must have

$$\text{index}(\bar{\partial}_{\mathbf{x}}) \equiv n - \text{gr}(\mathbf{x}) \pmod{2}.$$

Equation (19) implies that

$$\text{index}(\bar{\partial}_{\mathbf{y}}^{\text{new}}) \equiv (n - \text{gr}(\mathbf{x})) + (\text{gr}(\mathbf{x}) - \text{gr}(\mathbf{y})) \equiv n - \text{gr}(\mathbf{y}) \pmod{2}.$$

This means that the boundary conditions for $\bar{\partial}_{\mathbf{y}}^{\text{new}}$ are of the kind needed to define the orientation space $o(\mathbf{y})$, so the operator $\bar{\partial}_{\mathbf{y}}^{\text{new}}$ can play the role of $\bar{\partial}_{\mathbf{y}}$. Equation (20) gives the desired isomorphism (18). \square

This completes the description of the differential ∂ on $CF_*(L_0, L_1)$. To make sure that the sum in (17) is finite, and that $\partial^2 = 0$, we need extra assumptions. The usual requirement is that the Lagrangians are monotone with Maslov numbers greater than 2. However, this is not the case in Heegaard Floer theory; instead, some *ad hoc* arguments are used there. In particular, one shows that the contributions from disk and sphere bubbles to ∂^2 are zero. (The signs of these contributions are discussed in the next section.) Assuming we are in a situation where the bubble counts are zero, the only other place in the proof that $\partial^2 = 0$ where signs play a role is in the fact that the orientations on $\widehat{\mathcal{M}}(\phi)$ are compatible with gluing. For orientations coming from Pin structures as above, the proof of compatibility with gluing is as in [40, Section 12f]. Thus, we obtain Lagrangian Floer homology groups, denoted $HF_*(L_0, L_1)$.

Proposition 3.9. *Given Lagrangians $L_0, L_1 \subset M$ equipped with Pin structures and a coupled orientation (and satisfying the assumptions above), the Floer homology groups $HF_*(L_0, L_1)$ are well-defined up to canonical isomorphism.*

Proof. Proposition 3.5 ensures that the orientation spaces $o(\mathbf{x})$ are well-defined (up to canonical isomorphisms), and these isomorphisms commute with the differentials. \square

3.3. Bubbles. When showing $\partial^2 = 0$ in the Lagrangian Floer complex, we may encounter moduli spaces of stable disk bubbles, consisting of sphere bubbles attached to disks. For future reference, let us discuss their orientations.

For sphere bubbles, the linearization of the $\bar{\partial}$ operator is complex, so those moduli spaces are canonically oriented.

For disk bubbles, the ones relevant to $\partial^2 = 0$ are those in relative homotopy classes $\phi \in \pi_2(M, L_i)$ ($i = 0$ or 1) with one fixed boundary point:

$$(21) \quad u : (D^2, \partial D^2) \rightarrow (M, L_i), \quad u(1) = \mathbf{x} \in L_0 \cap L_1,$$

and with Maslov index $\mu(\phi) \leq 2$. In Heegaard-Floer theory, topological constraints preclude all disks of non-positive Maslov index, so only the case of Maslov index 2 disks is relevant. Following the notation in [30, Section 3.7], we let $\mathcal{N}(\phi)$ be the moduli space of such J -holomorphic disks, and we let $\widehat{\mathcal{N}}(\phi)$ be its quotient by the (two-dimensional) automorphism group $\text{Aut}(D^2, 1)$. To define a signed count $\#\widehat{\mathcal{N}}(\phi)$, we need the following fact.

Proposition 3.10. *Fix $i \in \{0, 1\}$. The Pin structure on the Lagrangian L_i induces a canonical orientation of $\widehat{\mathcal{N}}(\phi)$.*

Proof. To orient $\mathcal{N}(\phi)$, we need to trivialize the determinant index bundle $\det(D\bar{\partial})$ over the space of disks u satisfying (21). Given such a disk, we trivialize u^*TM over the disk D^2 . The pullback u^*TL_i gives a loop $\rho : S^1 \rightarrow \text{Gr}(\mathbb{C}^n)$ of Lagrangian subspaces along ∂D^2 . Let us denote the Cauchy-Riemann operator on D^2 , with these boundary conditions, by $\bar{\partial}_\rho$. This is studied in [40], where it is denoted $D_{D,\rho}$. Since the Maslov index of the path ρ is even (equal to 2), Lemma 11.17 in [40] says that the Pin structure induces an isomorphism

$$\det(\bar{\partial}_\rho) \cong \lambda^{\text{top}}(\rho(1)).$$

The actual operator we are interested in, $D\bar{\partial}_u$, differs from $\bar{\partial}_\rho$ in the fact that its value at 1 is fixed to be zero (because $u(1) = \mathbf{x}$ is fixed). Thus, we have

$$\text{ind}(D\bar{\partial}_u) = \text{ind}(\bar{\partial}_\rho) - \rho(1)$$

as virtual vector spaces. It follows that

$$\det(D\bar{\partial}_u) \cong \det(\bar{\partial}_\rho) \otimes \lambda^{\text{top}}(\rho(1))^{-1} \cong \mathbb{R},$$

so the bundle $\det(D\bar{\partial})$ is trivialized, as claimed.

The orientation on $\mathcal{N}(\phi)$ induces one on the quotient $\widehat{\mathcal{N}}(\phi)$ after fixing an orientation on the automorphism group $\text{Aut}(D^2, 1)$, which we do as follows. We identify D^2 with the upper half-space \mathbb{H} (preserving their orientations as subsets of \mathbb{C}), such that $1 \in \partial D^2$ gets mapped to infinity. Then $\text{Aut}(\mathbb{H}, \infty)$ is generated by translations by $t \in \mathbb{R}$ and dilations by $e^s, s \in \mathbb{R}$. We choose the orientation on its tangent space to be given by the ordered basis $(\partial/\partial s, \partial/\partial t)$. Finally, to relate the orientations on $\widehat{\mathcal{N}}(\phi)$ and $\mathcal{N}(\phi)$, we use the convention $\mathcal{N}(\phi) \cong \text{Aut}(D^2, 1) \times \widehat{\mathcal{N}}(\phi)$. \square

Example 3.11. Let L be the unit circle in the plane $M = \mathbb{R}^2$, and let ϕ be the class of the unit disk. This has Maslov index two, and the moduli space $\widehat{\mathcal{N}}(\phi)$ consists of a single point. Proposition 3.10 assigns to this point the positive orientation if L is equipped with the Lie group Pin structure, and the negative one if it is equipped with the bounding Pin structure. This can be read from the proof of Lemma 11.17 in [40], in which the study of orientations on $\det(\bar{\partial}_\rho)$ is reduced to the case where the Lagrangian loop ρ is constant. In that case, the trivialization rule is spelled explicitly, depending on the Pin structure as noted above.

Given a class $\phi \in \pi_2(M, L_i)$, there is a corresponding homotopy class of strips (still denoted ϕ) in $\pi_2(\mathbf{x}, \mathbf{x})$. If $\mu(\phi) = 2$, then the moduli space of strips $\widehat{\mathcal{M}}(\phi)$ is one-dimensional, and the moduli space of disk bubbles $\widehat{\mathcal{N}}(\phi)$ is part of its (zero-dimensional) boundary.

In this case, observe that Lemma 3.8 provides a canonical orientation on $\widehat{\mathcal{M}}(\phi)$, because when $\mathbf{x} = \mathbf{y}$, the tensor product $\det(\bar{\partial}_{\mathbf{x}}) \otimes \det(\bar{\partial}_{\mathbf{y}})$ is trivial.




FIGURE 1. Orientations for gluing a disk bubble to a constant trajectory.

Proposition 3.12. *For disk bubbles with boundary on L_0 , the orientation on $\widehat{\mathcal{N}}(\phi)$ from Proposition 3.10 agrees with the boundary orientation of $\widehat{\mathcal{M}}(\phi)$ from Lemma 3.8. For disk bubbles with boundary on L_1 , the two orientations disagree.*

Proof. Since $\mu(\phi) = 2$, the moduli space $\mathcal{M}(\phi)$ is one-dimensional, with some part of its boundary coming from disk bubbles in $\mathcal{N}(\phi) \cong \text{Aut}(D^2, 1) \times \widehat{\mathcal{N}}(\phi)$. Near that part, the strips in $\mathcal{M}(\phi)$ are obtained by gluing a disk bubble to the constant trajectory at \mathbf{x} . The gluing involves a parameter $T \rightarrow \infty$. Given a disk u in $\mathcal{N}(\phi)$, gluing it with the parameter T is equivalent to gluing some dilation of u by e^s with parameter 1, where $s \rightarrow -\infty$ as $T \rightarrow \infty$. Thus, the dilation coordinate $-s$ on $\text{Aut}(D^2, 1) \cong \text{Aut}(\mathbb{H}, \infty)$ corresponds to the gluing parameter. On the other hand, the translation coordinate t becomes (after gluing) either the \mathbb{R} -translation coordinate on $\mathcal{M}(\phi) \cong \mathbb{R} \times \widehat{\mathcal{M}}(\phi)$, or its opposite, depending on whether we consider bubbles with boundary on L_0 or L_1 . See Figure 1.

Since the \mathbb{R} factor comes first in the identifications $\mathcal{M}(\phi) \cong \mathbb{R} \times \widehat{\mathcal{M}}(\phi)$ and $\mathcal{N}(\phi) \cong \text{Aut}(D^2, 1) \times \widehat{\mathcal{N}}(\phi)$, the claim follows from the way we chose the orientation on $\text{Aut}(D^2, 1)$: we used $(\partial/\partial s, \partial/\partial t)$, which has the same orientation as $(\partial/\partial t, -\partial/\partial s)$. \square

The count of disk bubbles with boundary on a Lagrangian L appears in the calculation of the Floer homology of L with itself, which can be formulated using the pearl complex [3]. We state here the result when the count is zero.

Proposition 3.13. *Under the assumptions from Section 3.2, suppose $L_0 = L_1$ (with the same Pin structure), and denote this Lagrangian by L . Pick the coupled orientation on (L, L) to be the product of either orientation on L with itself. Assume the count of stable disk bubbles vanishes. Then, we have a canonical isomorphism*

$$HF_*(L, L) \cong H_*(L; |\lambda^{\text{top}}TL|),$$

where $|\lambda^{\text{top}}TL|$ denotes the orientation local system on L .

Proof. This follows from the PSS isomorphism [34], which is also discussed in [40, Section (12e)] and [45]. The PSS isomorphism involves counts of configurations consisting of holomorphic caps joined to Morse trajectories, and these can be graded using Pin structures by an analogue of Lemma 3.8. The isomorphism is usually stated in terms of Floer cohomology, saying that

$$HF^*(L, L) \cong H^*(L; \mathbb{Z}).$$

The statement for Floer homology can be deduced from this in view of our Conventions 3.1 and Poincaré duality for L . \square

Remark 3.14. Since L is assumed to be orientable, the local system $|\lambda^{\text{top}}TL|$ is trivial, so Proposition 3.13 implies the existence of an isomorphism $HF_*(L, L) \cong H_*(L; \mathbb{Z})$. However, this isomorphism is not canonical, because it depends on the choice of an orientation on L .

3.4. Decompositions. Inside the path space $\mathbf{P}(L_0, L_1)$, we have the constant paths $c_{\mathbf{x}}$ for every $\mathbf{x} \in L_0 \cap L_1$. Recall that $\pi_2(\mathbf{x}, \mathbf{y})$ denotes the set of relative homotopy classes of Whitney disks from \mathbf{x} to \mathbf{y} , i.e., relative homotopy classes of paths from $c_{\mathbf{x}}$ to $c_{\mathbf{y}}$ in $\mathbf{P}(L_0, L_1)$. There is a concatenation operation:

$$* : \pi_2(\mathbf{x}, \mathbf{y}) \times \pi_2(\mathbf{y}, \mathbf{z}) \rightarrow \pi_2(\mathbf{x}, \mathbf{z}).$$

In case $\mathbf{P}(L_0, L_1)$ is disconnected, some of the sets $\pi_2(\mathbf{x}, \mathbf{y})$ could be empty. We define an equivalence relation on $L_0 \cap L_1$ by

$$(22) \quad \mathbf{x} \sim \mathbf{y} \iff \pi_2(\mathbf{x}, \mathbf{y}) \neq \emptyset.$$

We let \mathbf{S} be the set of equivalence classes. Then, the Lagrangian Floer complex splits into a direct sum

$$CF_*(L_0, L_1) = \bigoplus_{\mathfrak{s} \in \mathbf{S}} CF_*(L_0, L_1, \mathfrak{s}),$$

where

$$CF_*(L_0, L_1, \mathfrak{s}) = \bigoplus_{\mathbf{x} \in \mathfrak{s}} o(\mathbf{x})$$

with the differential as in (17).

3.5. Twisted coefficients. There is a well-known variant of Lagrangian Floer homology, which uses twisted coefficients. In the context of Heegaard Floer theory, this was introduced by Ozsváth and Szabó in [29, Section 8]. We present it here in a more general framework.

We keep the set-up from Sections 3.2 and 3.4. Fix $\mathfrak{s} \in \mathbf{S}$ and a base intersection point $\mathbf{x}_0 \in \mathfrak{s}$. Let

$$G := \pi_1(\mathbf{P}(L_0, L_1), c_{\mathbf{x}_0}) = \pi_2(\mathbf{x}_0, \mathbf{x}_0).$$

We denote the elements of the group ring $\mathbb{Z}[G]$ by e^ψ , where $\psi \in G$.

Definition 3.15. A *complete set of paths for \mathfrak{s}* is a choice of relative homotopy classes $\theta_{\mathbf{x}} \in \pi_2(\mathbf{x}_0, \mathbf{x})$, one for each $\mathbf{x} \in \mathfrak{s}$.

Let A be a $\mathbb{Z}[G]$ -module, and fix a complete set of paths $\{\theta_{\mathbf{x}}\}$ for \mathfrak{s} . Then, for any $\mathbf{x}, \mathbf{y} \in \mathfrak{s}$, we get an identification

$$(23) \quad \pi_2(\mathbf{x}, \mathbf{y}) \xrightarrow{\cong} G, \quad \phi \mapsto \theta_{\mathbf{x}} * \phi * \theta_{\mathbf{y}}^{-1}.$$

Using this identification, for every $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we can make sense of the action of an element $e^\phi \in \mathbb{Z}[\pi_2(\mathbf{x}, \mathbf{y})]$ on A .

We define the Lagrangian Floer complex with coefficients in A by

$$\underline{CF}_*(L_0, L_1, \mathfrak{s}; A) = \bigoplus_{\mathbf{x} \in \mathfrak{s}} o(\mathbf{x}) \otimes_{\mathbb{Z}} A,$$

with the differential

$$\partial(\omega_{\mathbf{x}} \otimes a) = \sum_{\mathbf{y} \in \mathfrak{s}} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} (\#\widehat{\mathcal{M}}(\phi)) \cdot \omega_{\mathbf{y}} \otimes (e^\phi \cdot a).$$

The moduli spaces $\widehat{\mathcal{M}}(\phi)$ are oriented just as in Section 3.2. The resulting Floer homology with twisted coefficients is denoted $\underline{HF}_*(L_0, L_1, \mathfrak{s}; A)$.

Remark 3.16. A more canonical way of defining Floer complexes with twisted coefficients in A , without choosing a complete set of paths (but still choosing the basepoint \mathbf{x}_0), is to set

$$\underline{CF}_*^{\text{can}}(L_0, L_1, \mathfrak{s}; A) = \bigoplus_{\mathbf{x} \in \mathfrak{s}} o(\mathbf{x}) \otimes_{\mathbb{Z}} (\mathbb{Z}[\pi_2(\mathbf{x}_0, \mathbf{x})] \otimes_{\mathbb{Z}[G]} A),$$

where $\mathbb{Z}[\pi_2(\mathbf{x}_0, \mathbf{x})]$ is the free abelian group generated by $\pi_2(\mathbf{x}_0, \mathbf{x})$. The group ring $\mathbb{Z}[G] = \mathbb{Z}[\pi_2(\mathbf{x}_0, \mathbf{x}_0)]$ acts on $\mathbb{Z}[\pi_2(\mathbf{x}_0, \mathbf{x})]$ via concatenation of paths. We denote a typical generator of $\underline{CF}_*^{\text{can}}(L_0, L_1, \mathfrak{s}; A)$ by $\omega_{\mathbf{x}} \otimes (e^{\psi} \otimes a)$, where $\omega_{\mathbf{x}} \in o(\mathbf{x})$, $\psi \in \pi_2(\mathbf{x}_0, \mathbf{x})$, and $a \in A$. Then, the differential on $\underline{CF}_*(L_0, L_1, \mathfrak{s}; A)$ is given by

$$\partial(\omega_{\mathbf{x}} \otimes (e^{\psi} \otimes a)) = \sum_{\mathbf{y} \in \mathfrak{s}} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} (\#\widehat{\mathcal{M}}(\phi)) \cdot \omega_{\mathbf{y}} \otimes (e^{\psi * \phi} \otimes a).$$

The complete set of paths determines an isomorphism between $\mathbb{Z}[\pi_2(\mathbf{x}_0, \mathbf{x})]$ and $\mathbb{Z}[G]$, and hence between $\underline{CF}_*^{\text{can}}(L_0, L_1, \mathfrak{s}; A)$ and $\underline{CF}_*(L_0, L_1, \mathfrak{s}; A)$. In particular, this shows that the complexes $\underline{CF}_*(L_0, L_1, \mathfrak{s}; A)$, for different complete sets of paths, are all isomorphic.

Remark 3.17. When $A = \mathbb{Z}$ is the trivial $\mathbb{Z}[G]$ -module, the complex $\underline{CF}_*(L_0, L_1, \mathfrak{s}; \mathbb{Z})$ becomes the usual Lagrangian Floer complex. On the other hand, $\underline{CF}_*^{\text{can}}(L_0, L_1, \mathfrak{s}; \mathbb{Z})$ is only non-canonically isomorphic to $CF_*(L_0, L_1, \mathfrak{s}; \mathbb{Z})$; the isomorphism depends on the complete sets of paths. For our purposes, it is clearer to just choose such a complete set of paths from the very beginning, and work with the non-canonical complexes $\underline{CF}_*(L_0, L_1, \mathfrak{s}; A)$.

We will mostly be interested in the following kind of modules A . Let

$$(p_0, p_1) : \mathbf{P}(L_0, L_1) \rightarrow L_0 \times L_1$$

be the map taking a path to its endpoints. Given $\eta \in H^1(L_0; \mathbb{Z}/2)$ and $\zeta \in H^1(L_1; \mathbb{Z}/2)$, the pull-backs $p_0^* \eta$ and $p_1^* \zeta$ are elements of $H^1(\mathbf{P}(L_0, L_1); \mathbb{Z}/2)$ or, equivalently, morphisms $G \rightarrow \mathbb{Z}/2$. We let $A_{\eta, \zeta}$ be the $\mathbb{Z}[G]$ -module which is \mathbb{Z} as an abelian group, and such that elements $e^{\psi} \in \mathbb{Z}[G]$ act on $A_{\eta, \zeta}$ by multiplication with

$$(24) \quad (-1)^{p_0^* \eta(\psi)} \cdot (-1)^{p_1^* \zeta(\psi)} \in \{\pm 1\}.$$

We can define these modules for any $\mathfrak{s} \in \mathbf{S}$, so we could combine the resulting Floer complexes into one:

$$\underline{CF}_*(L_0, L_1; A_{\eta, \zeta}) = \bigoplus_{\mathfrak{s} \in \mathbf{S}} \underline{CF}_*(L_0, L_1, \mathfrak{s}; A_{\eta, \zeta}).$$

These twisted coefficients govern the changes in Pin structures on the Lagrangians.

Proposition 3.18. *Fix $\mathfrak{s} \in \mathbf{S}$. For $\eta \in H^1(L_0; \mathbb{Z}/2)$, let L_0^η denote the Lagrangian L_0 with its Pin structure changed from $P_0^\#$ to $P_0^\# \otimes \eta$. Similarly, for $\zeta \in H^1(L_1; \mathbb{Z}/2)$, let L_1^ζ denote L_1 with its Pin structure changed from $P_1^\#$ to $P_1^\# \otimes \zeta$. Then, we have a canonical isomorphism of Floer complexes*

$$CF_*(L_0^\eta, L_1^\zeta, \mathfrak{s}) \cong \underline{CF}_*(L_0, L_1, \mathfrak{s}; A_{\eta, \zeta}).$$

Proof. For simplicity, let us just consider a change in the Pin structure on L_0 ; i.e., we assume $\zeta = 0$, and we denote $A_{\eta, 0}$ by A_η . (The effect of a change on the Pin structure on L_1 can be analyzed in a similar way.)

Recall that we have a base point $\mathbf{x}_0 \in \mathfrak{s}$. We fix an isomorphism between the restrictions of the Pin structures $P_0^\#$ and $P_0^\# \otimes \eta$ to $T_{\mathbf{x}_0} L_0$.

Recall that

$$(25) \quad \underline{CF}_*(L_0, L_1, \mathfrak{s}) = \bigoplus_{\mathbf{x} \in L_0 \cap L_1} o(\mathbf{x}) \otimes_{\mathbb{Z}} A_\eta.$$

On the other hand,

$$(26) \quad CF_*(L_0^\eta, L_1, \mathfrak{s}) = \bigoplus_{\mathbf{x} \in L_0 \cap L_1} o_\eta(\mathbf{x}),$$

where $o_\eta(\mathbf{x})$ is the orientation space at \mathbf{x} coming from the determinant index bundle $\det(\bar{\partial}_\mathbf{x}^\eta)$ associated to the new Pin structure $P_0^\# \otimes \eta$. To construct an isomorphism from the chain complex (25) to (26), we need isomorphisms

$$(27) \quad \Psi_\mathbf{x} : o(\mathbf{x}) \otimes_{\mathbb{Z}} A_\eta \rightarrow o_\eta(\mathbf{x})$$

for all $\mathbf{x} \in \mathbf{S}$.

Let $\omega_\mathbf{x} \in o(\mathbf{x})$, $a \in A_\eta$, and $\omega_\mathbf{x}^\eta \in o_\eta(\mathbf{x})$. Note that a is just an integer, although the elements of $\mathbb{Z}[G]$ act on it nontrivially. We set

$$(28) \quad \Psi_\mathbf{x}(\omega_\mathbf{x} \otimes a) = \sigma_\mathbf{x} \cdot a \cdot \omega_\mathbf{x}^\eta,$$

with the sign $\sigma_\mathbf{x} \in \{\pm 1\}$ is determined as follows. As part of the complete set of paths, we have a class $\theta_\mathbf{x} \in \pi_2(\mathbf{x}_0, \mathbf{x})$. Let $D\bar{\partial}(\theta_\mathbf{x})$ be the linearized Cauchy-Riemann operator for a path in that class. By Lemma 3.8, we have isomorphisms:

$$(29) \quad \begin{aligned} \det(\bar{\partial}_{\mathbf{x}_0}) \otimes \det(D\bar{\partial}(\theta_\mathbf{x})) &\cong \det(\bar{\partial}_\mathbf{x}), \\ \det(\bar{\partial}_{\mathbf{x}_0}^\eta) \otimes \det(D\bar{\partial}(\theta_\mathbf{x})) &\cong \det(\bar{\partial}_\mathbf{x}^\eta). \end{aligned}$$

Our identification of the Pin structures on $T_{\mathbf{x}_0}L_0$ gives an isomorphism between $\det(\bar{\partial}_{\mathbf{x}_0})$ and $\det(\bar{\partial}_{\mathbf{x}_0}^\eta)$. We obtain an isomorphism

$$(30) \quad \det(\bar{\partial}_\mathbf{x}) \cong \det(\bar{\partial}_\mathbf{x}^\eta).$$

We let the sign $\sigma_\mathbf{x}$ in (28) be $+1$ if $\omega_\mathbf{x}$ gets taken to $\omega_\mathbf{x}^\eta$ under this isomorphism, and we let it be -1 otherwise.

Let us check that the isomorphisms $\Psi_\mathbf{x}$ commute with the chain complex differentials. Consider the contribution of a generator $\omega_\mathbf{y}$ to $\partial\omega_\mathbf{x}$ in $\underline{CF}_*(L_0, L_1, \mathfrak{s})$ coming from moduli spaces in a class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, and compare it to the contribution of some $\omega_\mathbf{y}^\eta$ to $\partial\omega_\mathbf{x}^\eta$ in $CF_*(L_0^\eta, L_1, \mathfrak{s})$ from the same ϕ . Let $\psi = \theta_\mathbf{x} * \phi * \theta_\mathbf{y}^{-1} \in \pi_2(\mathbf{x}_0, \mathbf{x}_0)$ be the class that corresponds to ϕ under the identification (23). Because the differential on $\underline{CF}_*(L_0, L_1, \mathfrak{s})$ involves twisted coefficients, it picks up a sign of $(-1)^{p_0^* \eta(\psi)}$. Thus, what we need to check is that the diagram

$$\begin{array}{ccc} \det(\bar{\partial}_\mathbf{x}) \otimes \det(D\bar{\partial}(\phi)) & \xrightarrow{\cong} & \det(\bar{\partial}_\mathbf{x}^\eta) \otimes \det(D\bar{\partial}(\phi)) \\ \cong \downarrow & & \downarrow \cong \\ \det(\bar{\partial}_\mathbf{y}) & \xrightarrow{\cong} & \det(\bar{\partial}_\mathbf{y}^\eta) \end{array}$$

commutes up to the sign $(-1)^{p_0^* \eta(\psi)}$. In this diagram, the top horizontal isomorphism is the one in (30), using the class $\theta_\mathbf{x}$; the bottom horizontal isomorphism is its analogue using the class $\theta_\mathbf{y}$, tensored with the identity on the $\det(D\bar{\partial}(\phi))$ factors; the vertical isomorphisms are from Lemma 3.8 and use the class ϕ .

Going back to the definition of the horizontal isomorphisms from (29), we see that we can consider instead the diagram

$$(31) \quad \begin{array}{ccc} \det(\bar{\partial}_{\mathbf{x}_0}) \otimes \det(D\bar{\partial}(\psi)) & \xrightarrow{\cong} & \det(\bar{\partial}_{\mathbf{x}_0}^\eta) \otimes \det(D\bar{\partial}(\psi)) \\ \cong \downarrow & & \downarrow \cong \\ \det(\bar{\partial}_{\mathbf{x}_0}) & \xrightarrow{\cong} & \det(\bar{\partial}_{\mathbf{x}_0}^\eta) \end{array}$$

where the horizontal maps come from the identification of the Pin structures at \mathbf{x}_0 , and the vertical isomorphisms use Lemma 3.8 for the class ψ . We are left to show that (31) commutes up to the sign $(-1)^{p_0^* \eta(\psi)}$.

The left vertical isomorphism in (31) is induced by the original Pin structures on the Lagrangians, and the right one by the new Pin structures (where the structure on L_0 is tensored with η). The construction of the isomorphisms in Lemma 3.8 is based on identifying the orientation spaces $o(\mathbf{y})$ (where in our case $\mathbf{y} = \mathbf{x}_0$) coming from $\bar{\partial}_{\mathbf{y}}^{\text{new}}$ and $\bar{\partial}_{\mathbf{y}}$. This relies on Proposition 3.5, which identifies orientation spaces associated to paths in $\text{Gr}^\dagger(V)$ with fixed endpoints and fixed index. In our setting, we compare two such paths from $T_{\mathbf{x}_0}L_0$ to $T_{\mathbf{x}_0}L_1$. These differ by concatenating with a loop from $T_{\mathbf{x}_0}L_0$ to itself, and whether this loop is nontrivial in $\pi_1(\text{Gr}^\dagger(T_{\mathbf{x}_0}L_0))$ is determined by the sign $(-1)^{p_0^* \eta(\psi)}$. If the sign is positive, the two paths are homotopic and produce the same isomorphism, so the diagram (31) commutes. If the sign is negative, Lemma 3.4 shows that the isomorphisms differ by -1 , so the diagram (31) anti-commutes.

We have now shown that the formula (27) produces an isomorphism of chain complexes. Recall that we started our construction by choosing an isomorphism between the restrictions of the Pin structures $P_0^\#$ and $P_0^\# \otimes \eta$ to $T_{\mathbf{x}_0}L_0$. There is a connected space of such isomorphisms, so this choice will not affect the canonical nature of the final isomorphism between chain complexes, which is a discrete object. \square

Corollary 3.19. *Let $\mathfrak{s} \in \mathbf{S}$ and $\alpha \in H^1(M; \mathbb{Z}/2)$. Consider the inclusions $\iota_0 : L_0 \hookrightarrow M$ and $\iota_1 : L_1 \hookrightarrow M$. For $i \in \{0, 1\}$, let L_i^α denote the Lagrangian L_i with its Pin structure changed by tensoring with $\iota_i^* \alpha$. Then, there is a canonical isomorphism:*

$$CF_*(L_0, L_1, \mathfrak{s}) \cong CF_*(L_0^\alpha, L_1^\alpha, \mathfrak{s}).$$

Proof. This is a consequence of Proposition 3.18, by taking $\eta = \iota_0^* \alpha$ and $\zeta = \iota_1^* \alpha$. Indeed, note that in the formula (24) we have $\pi_0^* \eta(\psi) = \pi_1^* \zeta(\psi)$, so the module $A_{\eta, \zeta}$ is \mathbb{Z} with the trivial $\mathbb{Z}[G]$ -action. \square

Corollary 3.19 shows that orientations of the moduli spaces in Floer theory are induced by Pin structures on the Lagrangians modulo those coming from the ambient space; that is, by an affine space on the cokernel of the map

$$H^1(M; \mathbb{Z}/2) \rightarrow H^1(L_0; \mathbb{Z}/2) \oplus H^1(L_1; \mathbb{Z}/2).$$

3.6. Polygon maps. Our discussion of orientations on moduli spaces extends to holomorphic polygons, much as in Seidel's book [40]. We state the results here, and leave the proofs to the reader.

For $m \geq 2$, suppose we have monotone Lagrangians $L_0, L_1, \dots, L_m \subset M$, equipped with Pin structures. We also assume that we are given coupled orientations on the pairs (L_0, L_1) , (L_1, L_2) , \dots , (L_{m-1}, L_m) . These induce a coupled orientation on (L_0, L_m) . We define a map

$$(32) \quad f_{L_0, \dots, L_m} : CF_*(L_0, L_1) \otimes CF_*(L_1, L_2) \otimes \cdots \otimes CF_*(L_{m-1}, L_m) \rightarrow CF_*(L_0, L_m)$$

by

$$f_{L_0, \dots, L_m}(\omega_{\mathbf{x}_1} \otimes \omega_{\mathbf{x}_2} \otimes \cdots \otimes \omega_{\mathbf{x}_m}) = \sum_{\mathbf{y} \in L_0 \cap L_m} \sum_{\substack{\phi \in \pi_2(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m, \mathbf{y}) \\ \mu(\phi)=0}} (\#\mathcal{M}(\phi)) \cdot \omega_{\mathbf{y}}.$$

Here, $\mathbf{x}_i \in L_{i-1} \cap L_i$ for $i = 1, \dots, m$, and $\mathbf{y} \in L_0 \cap L_m$. Also, $\omega_{\mathbf{x}_i} \in o(\mathbf{x}_i)$ and $\omega_{\mathbf{y}} \in o(\mathbf{y})$ are orientations for the respective determinant lines. The notation $\pi_2(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m, \mathbf{y})$ indicates the set of (relative) homotopy classes of $(k+1)$ -gons in M with boundaries on L_0, L_1, \dots, L_m and vertices at $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m, \mathbf{y}$, in clockwise order. (See Conventions 3.1.) The moduli space $\mathcal{M}(\phi)$ consists of holomorphic polygon maps in the class ϕ , of index $\mu(\phi) = 0$. Its orientation is induced from the Pin structures and coupled orientations, using a straightforward generalization of Lemma 3.8.

When $m = 1$, we let f_{L_0, L_1} denote the differential ∂ on the complex $CF_*(L_0, L_1)$.

Let us adjust the signs to define

$$\tilde{f}_{L_0, \dots, L_m}(\omega_{\mathbf{x}_1} \otimes \cdots \otimes \omega_{\mathbf{x}_m}) = (-1)^{\text{gr}(\mathbf{x}_m) + 2\text{gr}(\mathbf{x}_{m-1}) + \cdots + m\text{gr}(\mathbf{x}_1)} f_{L_0, \dots, L_m}(\omega_{\mathbf{x}_1} \otimes \cdots \otimes \omega_{\mathbf{x}_m}).$$

Then, these new polygon maps satisfy the A_∞ relations:

$$\sum_{0 \leq i < j \leq m} (-1)^* \tilde{f}_{L_0, \dots, L_i, L_j, \dots, L_m}(\omega_{\mathbf{x}_1} \otimes \cdots \otimes \omega_{\mathbf{x}_i} \otimes \tilde{f}_{L_i, L_{i+1}, \dots, L_j}(\omega_{\mathbf{x}_{i+1}} \otimes \cdots \otimes \omega_{\mathbf{x}_j}) \otimes \omega_{\mathbf{x}_{j+1}} \otimes \cdots \otimes \omega_{\mathbf{x}_m}) = 0$$

where $*$ = $\text{gr}(\mathbf{x}_{j+1}) + \cdots + \text{gr}(\mathbf{x}_m) + (m - j)$. Note that the signs are slightly different from those in [40], because we use homological rather than cohomological conventions.

When $m = 1$, the A_∞ relation says that $\partial^2 = 0$ on $CF_*(L_0, L_1)$. When $m = 2$, it says that

$$f_{L_0, L_1, L_2} : CF_*(L_0, L_1) \otimes CF_*(L_1, L_2) \rightarrow CF_*(L_0, L_2)$$

is a chain map. Hence, it induces a product map on Floer homology:

$$(33) \quad F_{L_0, L_1, L_2} : HF_*(L_0, L_1) \otimes HF_*(L_1, L_2) \rightarrow HF_*(L_0, L_2).$$

In Section 3.4 we decomposed the Floer complexes according to equivalence classes of intersection points. Pick such equivalence classes $\mathfrak{s}_i \subset L_{i-1} \cap L_i$ for $i = 1, \dots, m$, and $\mathfrak{t} \subset L_0 \cap L_m$. We then define an equivalence relation on tuples $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \phi)$ with $\mathbf{x}_i \in \mathfrak{s}_i$, $\mathbf{y} \in \mathfrak{t}$, and $\phi \in \pi_2(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$, where

$$(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \phi) \sim (\mathbf{x}'_1, \dots, \mathbf{x}'_m, \mathbf{y}', \phi')$$

if ϕ' differs from ϕ by concatenation with classes in $\pi_2(\mathbf{x}_1, \mathbf{x}'_1), \dots, \pi_2(\mathbf{x}_m, \mathbf{x}'_m)$ and $\pi_2(\mathbf{y}, \mathbf{y}')$. We let \mathbf{S} be the set of such equivalence classes. For every $\mathfrak{s} \in \mathbf{S}$, we have a polygon map

$$f_{L_0, \dots, L_m}^{\mathfrak{s}} : CF_*(L_0, L_1, \mathfrak{s}_1) \otimes \cdots \otimes CF_*(L_{m-1}, L_m, \mathfrak{s}_m) \rightarrow CF_*(L_0, L_m, \mathfrak{t})$$

which counts only polygons in classes ϕ such that $[(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \phi)] = \mathfrak{s} \in \mathbf{S}$. The map (32) splits as the sum of these maps, over all $\mathfrak{s} \in \mathbf{S}$.

There is also a variant of the polygon maps with twisted coefficients; compare [29, Section 8.2]. It can be described using bimodule tensor products in the formalism of Remark 3.16, but we present instead the explicit definition: Fix base intersection points $\mathbf{x}_{i,0} \in \mathfrak{s}_i$ for $i = 1, \dots, m$, as well as $\mathbf{y}_0 \in \mathfrak{t}$. Let

$$G_i = \pi_1(\mathbf{P}(L_{i-1}, L_i), c_{\mathbf{x}_{i,0}}) = \pi_2(\mathbf{x}_{i,0}, \mathbf{x}_{i,0}), \quad i = 1, \dots, m$$

and let also

$$G = \pi_1(\mathbf{P}(L_0, L_m), c_{\mathbf{y}_0}) = \pi_2(\mathbf{y}_0, \mathbf{y}_0).$$

Denote by $\underline{\mathbf{S}} = \pi_2(\mathbf{x}_{1,0}, \dots, \mathbf{x}_{m,0}, \mathbf{y}_0)$. Observe that \mathbf{S} is the quotient of $\underline{\mathbf{S}}$ by the actions of all the groups G_i and G ; for $\mathfrak{s} \in \mathbf{S}$, we let $\underline{\mathbf{S}}(\mathfrak{s}) \subset \underline{\mathbf{S}}$ be its equivalence class.

Fix also complete sets of paths for all \mathfrak{s}_i and \mathfrak{t} , as in Definition 3.15. These induce identifications between $\underline{\mathbf{S}}$ and $\pi_2(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$, for all $\mathbf{x}_i \in \mathfrak{s}_i$ and $\mathbf{y} \in \mathfrak{t}$. Given $\phi \in \pi_2(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$, we let $\underline{\mathfrak{s}}(\phi)$ denote the corresponding element in $\underline{\mathbf{S}}$.

Fix $\mathfrak{s} \in \mathbf{S}$. Suppose we have modules A_i over $\mathbb{Z}[G_i]$. These induce a module $A^\mathfrak{s}$ over $\mathbb{Z}[G]$ given by

$$(34) \quad A^\mathfrak{s} = \frac{\{(a_1, \dots, a_m, \underline{\mathfrak{s}}) \in A_1 \times \dots \times A_m \times \underline{\mathbf{S}}(\mathfrak{s})\}}{(a_1, \dots, a_m, \underline{\mathfrak{s}}) \sim (e^{\psi_1} a_1, \dots, e^{\psi_m} a_m, (\psi_1 \times \dots \times \psi_m) \underline{\mathfrak{s}})}$$

where ψ_i are arbitrary elements of G_i for $i = 1, \dots, m$. For $\psi \in G$, the element $e^\psi \in \mathbb{Z}[G]$ acts on $[(a_1, \dots, a_m, \underline{\mathfrak{s}})]$ by taking it to $[(a_1, \dots, a_m, \psi \underline{\mathfrak{s}})]$.

We get a polygon map

$$(35) \quad \underline{f}_{L_0, \dots, L_m}^\mathfrak{s} : \underline{CF}_*(L_0, L_1, \mathfrak{s}_1; A_1) \otimes \dots \otimes \underline{CF}_*(L_{m-1}, L_m, \mathfrak{s}_m; A_m) \rightarrow \underline{CF}_*(L_0, L_m, \mathfrak{t}; A^\mathfrak{s})$$

given by

$$\begin{aligned} \underline{f}_{L_0, \dots, L_m}^\mathfrak{s} ((\omega_{\mathbf{x}_1} \otimes a_1) \otimes \dots \otimes (\omega_{\mathbf{x}_m} \otimes a_m)) = \\ \sum_{\mathbf{y} \in \mathfrak{t}} \sum_{\underline{\mathfrak{s}} \in \underline{\mathbf{S}}(\mathfrak{s})} \sum_{\substack{\phi \in \pi_2(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m, \mathbf{y}) \\ [(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \phi)] = \mathfrak{s} \\ \mu(\phi) = 0}} (\#\mathcal{M}(\phi)) \cdot (\omega_{\mathbf{y}} \otimes [(a_1, \dots, a_m, \underline{\mathfrak{s}}(\phi))]). \end{aligned}$$

These polygon maps satisfy A_∞ relations just like the untwisted ones. In particular, the triangle maps give chain maps and hence maps on the Floer homology with twisted coefficients.

As in Section 3.5, we are particularly interested in modules that govern the change in Pin structures. Pick elements $\eta_i \in H^1(L_i; \mathbb{Z}/2)$ for $i = 0, \dots, m$. From these we obtain modules $A_i := A_{\eta_{i-1}, \eta_i}$ over $\mathbb{Z}[G_i]$ for $i = 1, \dots, m$, as well as A_{η_0, η_m} over $\mathbb{Z}[G]$. For $\mathfrak{s} \in \mathbf{S}$, let $A^\mathfrak{s}$ be as in (34).

Pick any (not necessarily holomorphic) polygon in a class in $\underline{\mathbf{S}}(\mathfrak{s}) \subset \pi_2(\mathbf{x}_{1,0}, \dots, \mathbf{x}_{m,0}, \mathbf{y}_0)$, and let Γ be its boundary. This consists of paths γ_i on L_i from $\mathbf{x}_{i,0}$ to $\mathbf{x}_{i+1,0}$ for $i = 1, \dots, m-1$, as well as γ_0 on L_0 from \mathbf{y}_0 to $\mathbf{x}_{1,0}$ and γ_m on L_m from $\mathbf{x}_{m,0}$ to \mathbf{y}_0 . Since the class \mathfrak{s} is fixed, the homotopy class of Γ (relative to its vertices) is fixed up to concatenation with boundaries of disks in $G_i = \pi_2(\mathbf{x}_{i,0}, \mathbf{x}_{i,0})$ and $G = \pi_2(\mathbf{y}_0, \mathbf{y}_0)$.

For $i = 0, \dots, m$, restriction to the part of the boundary in L_i produces maps q_i from $\underline{\mathbf{S}} = \pi_2(\mathbf{x}_{1,0}, \dots, \mathbf{x}_{m,0}, \mathbf{y}_0)$ to the set of relative homotopy classes of paths between the two relevant basepoints on L_i . The latter space can be identified with $\pi_1(L_i)$ by concatenating with the path γ_i . Viewing the class η_i as a morphism $\pi_1(L_i) \rightarrow \mathbb{Z}/2$, we obtain a map

$$\eta_i \circ q_i : \underline{\mathbf{S}} \rightarrow \mathbb{Z}/2.$$

We now define a map

$$(36) \quad A^\mathfrak{s} \rightarrow A_{\eta_0, \eta_m}, \quad [(a_1, a_2, \dots, a_m, \underline{\mathfrak{s}})] \rightarrow a_1 a_2 \dots a_m \cdot (-1)^{\sum_{i=0}^m (\eta_i \circ q_i)(\underline{\mathfrak{s}})}.$$

It is straightforward to check that this is well-defined, and a morphism of $\mathbb{Z}[G]$ -modules.

The morphism (36) induces a map $\underline{CF}_*(L_0, L_m, \mathfrak{t}; A^\mathfrak{s}) \rightarrow \underline{CF}_*(L_0, L_m, \mathfrak{t}; A_{\eta_0, \eta_m})$. Composing the polygon maps in (35) with this yields a map

$$\underline{f}_{L_0, \dots, L_m}^\mathfrak{s} : \underline{CF}_*(L_0, L_1, \mathfrak{s}_1; A_{\eta_0, \eta_1}) \otimes \dots \otimes \underline{CF}_*(L_{m-1}, L_m, \mathfrak{s}_m; A_{\eta_{m-1}, \eta_m}) \rightarrow \underline{CF}_*(L_0, L_m, \mathfrak{t}; A_{\eta_0, \eta_m}).$$

Proposition 3.20. *For $i = 0, \dots, m$, let $L_i^{\eta_i}$ denote L_i with its Pin structure changed by tensoring with η_i . Under the identifications from Proposition 3.18, the map $\underline{f}_{L_0, \dots, L_m}^\mathfrak{s}$ corresponds*

to the ordinary polygon map

$$f_{L_0^{\eta_0}, \dots, L_m^{\eta_m}} : CF_*(L_0^{\eta_0}, L_1^{\eta_1}, \mathfrak{s}_1) \otimes \cdots \otimes CF_*(L_{m-1}^{\eta_{m-1}}, L_m^{\eta_m}, \mathfrak{s}_m) \rightarrow CF_*(L_0^{\eta_0}, L_m^{\eta_m}, \mathfrak{t}).$$

The proof is routine (along the same lines as that of Proposition 3.18), so we omit it.

3.7. Examples. We give a few simple examples of sign counts of polygons, starting with the case when M is a surface, and the Lagrangians are curves. Recall from Definition 2.11 that in this case, there are canonical Lie group Pin structures on the Lagrangians (not just up to isomorphism).

If V is a one-dimensional vector space, we denote by $|V|$ the abelian group generated by orientations ω on V , modulo the relation $\bar{\omega} = -\omega$.

Lemma 3.21. *Let M be a two-dimensional symplectic manifold, and $L_0, L_1 \subset M$ be curves equipped with their Lie group Pin structures and a coupled orientation. Let $\mathbf{x} \in L_0 \cap L_1$ be a transverse intersection point.*

(a) *If $\text{gr}(\mathbf{x}) = 1 \in \mathbb{Z}/2$, then the orientation space $o(\mathbf{x})$ from (16) has a canonical trivialization.*

(b) *If $\text{gr}(\mathbf{x}) = 0$, then $o(\mathbf{x})$ is canonically identified with $|T_{\mathbf{x}}L_0|$.*

Proof. (a) The condition $\text{gr}(\mathbf{x}) = 1$ means that the orientation on $T_{\mathbf{x}}L_0 \oplus T_{\mathbf{x}}L_1$ agrees with the one on $T_{\mathbf{x}}M$. Up to isotopy, there is a unique linear isomorphism from $T_{\mathbf{x}}M$ to \mathbb{C} taking $T_{\mathbf{x}}L_0$ to \mathbb{R} and $T_{\mathbf{x}}L_1$ to $i\mathbb{R}$ (respecting the coupled orientation). In view of Remark 3.2, the $\bar{\partial}$ operator on the cap H with boundary conditions interpolating from \mathbb{R} to $i\mathbb{R}$ via $e^{i\theta}\mathbb{R}$, $\theta \in [0, \pi/2]$ is equivalent to the one discussed in [40, proof of Lemma 11.11], where it is shown to be bijective. Hence, its determinant index bundle can be trivialized by picking the positive generator.

(b) If $\text{gr}(\mathbf{x}) = 0$, then the orientation on $T_{\mathbf{x}}L_0 \oplus T_{\mathbf{x}}L_1$ disagrees with the one on $T_{\mathbf{x}}M$. We can again identify $o(\mathbf{x})$ with a standard orientation space, this time using Lagrangians interpolating from $i\mathbb{R}$ to \mathbb{R} . This is the reverse of the path considered in (a). By Equation (11.27) in [40], the orientation space differs from the one in part (a) by tensoring with $|T_{\mathbf{x}}L_0|$. \square

Remark 3.22. In view of Lemma 3.21, when $\dim(M) = 2$ and the Lagrangians have the Lie groups Pin structures, choosing an orientation on L_0 (or, equivalently, on L_1 , in a way compatible with the coupled orientation) allows us to trivialize $o(\mathbf{x})$ for all intersection points $\mathbf{x} \in L_0 \cap L_1$. We can then view the Lagrangian Floer complex $CF_*(L_0, L_1)$ as generated by the intersection points: we identify each \mathbf{x} with the positive generator of $o(\mathbf{x}) \cong \mathbb{R}$.

Example 3.23. We review Example 13.5 in [40], using our Convention 3.1. Let $M \cong S^1 \times [0, 1]$ be an open annulus, which we view as the complement of a small disk $D^2(\epsilon)$ in \mathbb{C} . Let α be the unit circle, and β a Hamiltonian translate of α intersecting it in two points, as in Figure 2. Equip α and β with the Lie group Pin structures. Pick also counterclockwise orientations on α and β as in the figure. The resulting coupled orientation on (α, β) gives $\text{gr}(\mathbf{x}) = 1$ and $\text{gr}(\mathbf{y}) = 0$.

Following Remark 3.22, the Floer complex $CF_*(\alpha, \beta)$ is generated by the intersection points \mathbf{x} and \mathbf{y} . We have

$$\partial \mathbf{x} = \pm \mathbf{y} \pm \mathbf{y},$$

where the two signs come from the bigons A and B in Figure 2. By Proposition 3.13, we know that $HF_*(\alpha, \beta) \cong H_*(S^1)$, so we must have $\partial \mathbf{x} = 0$. Thus, the two bigons A and B must come with opposite signs. In fact, it is computed in [40, Example 13.5] that A contributes +1 and B contributes -1.




FIGURE 2. Two circles in an annulus.




FIGURE 3. (a) A trivial triangle at a triple intersection. (b), (c): Its perturbations.

Example 3.24. Let M be any surface, and consider three curves $\alpha, \beta, \gamma \subset M$, having a triple intersection point \mathbf{x} . Suppose α, β, γ appear in this cyclic order around \mathbf{x} , counterclockwise. We equip them with the Lie group Pin structures. There is a trivial α - β - γ triangle of index 0 that contributes to the polygon map

$$f_{\alpha,\beta,\gamma} : CF_*(\alpha, \beta) \otimes CF_*(\beta, \gamma) \rightarrow CF_*(\alpha, \gamma).$$

Suppose we choose orientations on α, β, γ either as in Figure 3 (a) or all three opposite to those shown there. This is equivalent to asking that the induced coupled orientations on each pair (α, β) , (β, γ) and (α, γ) are such that $\text{gr}(\mathbf{x}) = 1 \in \mathbb{Z}/2$ in all three Floer complexes. Then, the triangle map takes the form

$$f_{\alpha,\beta,\gamma}(\mathbf{x} \otimes \mathbf{x}) = \pm \mathbf{x} + \dots$$

where $\pm \mathbf{x}$ is the contribution of the trivial triangle. We claim that its sign is positive. This can be checked explicitly by following the definitions. Alternatively, we can determine the sign from a specific example: embed the picture into one with three circles in the annulus, all Hamiltonian translates of the other. Then, under the PSS isomorphism, the generators \mathbf{x} become identified with $[S^1] \in H_*(S^1)$, and the triangle map becomes the intersection product on homology:

$$H_*(S^1) \otimes H_*(S^1) \rightarrow H_*(S^1).$$

See for example [1] (where this is stated in terms of cohomology, with the triangle map becoming the cup product) or [45]. The intersection of $[S^1]$ and $[S^1]$ is $[S^1]$, with a positive sign.

Example 3.25. If α, β, γ are as in Example 3.24, we can slightly perturb them to get a triple α', β', γ' as in Figure 3 (b) or (c). The small α' - β' - γ' triangle between the curves is a deformation of the trivial triangle from (a), so it also contributes positively to the polygon

map. Note that the two triangles in Figure 3(b) and (c) are, in fact, isomorphic. They differ by a 180° rotation, which reverses the orientations on the curves; but what matter are only the coupled orientations, which stay the same.

Remark 3.26. If α, β, γ are as in Example 3.24, except in the cyclic order α, γ, β , then the trivial α - β - γ triangle is obstructed, and thus does not contribute to the polygon map. Indeed, a slight perturbation of one of the curves produces a small α - γ - β triangle instead of an α - β - γ one.

Example 3.27. Suppose we have two symplectic manifolds M and M' (of dimensions $2n$ and $2n'$) and two pairs of Lagrangians as in Section 3.2: (L_0, L_1) in M , and (L'_0, L'_1) in M' . We obtain a product pair $(L_0 \times L'_0, L_1 \times L'_1)$ in $M \times M'$.

Unfortunately, Pin structures do not behave well with respect to products; see Remark 2.15. We will use Spin structures instead. Assume L_0, L_1, L'_0, L'_1 come equipped with Spin structures (and, in particular, are oriented). We consider the product Spin structures on $L_0 \times L'_0$ and $L_1 \times L'_1$. Note that Spin structures on a pair of Lagrangians produce both Pin structures and orientations, hence coupled orientations, so the Floer chain complex is well-defined.

With this in mind, we consider the Floer chain complexes $CF_*(L_0, L_1)$, $CF_*(L'_0, L'_1)$ and $CF_*(L_0 \times L'_0, L_1 \times L'_1)$. Given intersection points $\mathbf{x} \in L_0 \cap L_1$ and $\mathbf{x}' \in L'_0 \cap L'_1$, there is an intersection point $\mathbf{x} \times \mathbf{x}' \in (L_0 \times L'_0) \cap (L_1 \times L'_1)$. Given index problems on the factors, their direct sum is the index problem on the product. Thus, we have identifications between the orientation spaces $o(\mathbf{x} \times \mathbf{x}') \cong o(\mathbf{x}) \otimes o(\mathbf{x}')$. We also have an identification between the moduli spaces $\mathcal{M}(\phi \times \phi') \cong \mathcal{M}(\phi) \times \mathcal{M}(\phi')$. It follows that $CF_*(L_0 \times L'_0, L_1 \times L'_1)$ is the tensor product of the chain complexes $CF_*(L_0, L_1)$ and $CF_*(L'_0, L'_1)$. (See [42, 2].)

More generally, if we have an $(m+1)$ -tuple of Spin Lagrangians on M and an $(m+1)$ -tuple of Spin Lagrangians on M' , the moduli space of holomorphic polygons on the product is the fiber product of the moduli spaces of holomorphic polygons on the factors over the abstract moduli spaces of $(m+1)$ -polygons. In particular, for $m = 2$, it is simply a product.

4. THREE-MANIFOLD INVARIANTS

We present here a construction of Heegaard Floer homology, following [30], but with two departures from the original. The first (which is our main point) is that we define the signs in the differential using canonical orientations instead of coherent orientations. The second is that we apply Perutz's work [32] to equip the symmetric product with a symplectic form so that the Heegaard tori are Lagrangians, rather than only totally real as in [30]. This allows us to be in the setting of Section 3. An alternative would have been to extend the results from that section to the totally real case; this can indeed be done, because the totally real Grassmannian deformation retracts onto the Lagrangian one. However, Perutz's perspective will also be useful to us when we discuss handleslides in Section 4.8.

4.1. Heegaard Floer homology: a preliminary version. Let Y be a closed, connected, oriented 3-manifold and $z \in Y$ a basepoint. A (pointed) *Heegaard diagram* for Y consists of the data

$$\mathcal{H} = (\Sigma, \boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_g\}, \boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}, z)$$

where Σ is a surface of genus g containing z , and splitting Y into two handlebodies U_α and U_β . Further, $\boldsymbol{\alpha}$ is a collection of g disjoint simple closed curves on Σ that are attaching curves for U_α , in the sense that we obtain U_α by attaching disks with boundaries on the curves α_i , and then a 3-ball. Similarly, $\boldsymbol{\beta}$ is a collection of attaching curves for U_β . We assume that the

alpha and beta curves intersect transversely, and are disjoint from z . Note that U_α and U_β inherit orientations from Y . We orient Σ as $\partial U_\alpha = -\partial U_\beta$.

We consider the symmetric product

$$M = \text{Sym}^g(\Sigma) = \Sigma^{\times g}/S_g,$$

where S_g is the symmetric group. Inside $H^2(\text{Sym}^g(\Sigma); \mathbb{Z})$ we have a class η Poincaré dual to $\{z\} \times \text{Sym}^{g-1}(\Sigma)$. For each oriented simple closed curve γ on Σ , we also have a class $\langle \gamma \rangle$ in $H^1(\text{Sym}^g(\Sigma); \mathbb{Z})$, dual to $\gamma \times \text{Sym}^{g-1}(\Sigma)$. If we pick curves $\gamma_1, \dots, \gamma_g$ and duals $\gamma_1^*, \dots, \gamma_g^*$ so that they produce a symplectic basis of $H_1(\Sigma)$, the class

$$\theta := \sum_{i=1}^g \langle \gamma_i \rangle \langle \gamma_i^* \rangle \in H^2(\text{Sym}^{g-1}(\Sigma); \mathbb{Z})$$

is invariant under the action of the mapping class group of Σ .

For $\epsilon > 0$ sufficiently small, Perutz [32] constructs a symplectic form ω on Σ in the class $\eta + \epsilon\theta$, such that the tori

$$\mathbb{T}_\alpha = \alpha_1 \times \dots \times \alpha_g, \quad \mathbb{T}_\beta = \beta_1 \times \dots \times \beta_g$$

are Lagrangians with respect to ω . As an aside, we note that $c_1(TM) = \eta - \theta$.

In [29, Section 10.4], Ozsváth and Szabó define an absolute grading on the Heegaard Floer groups and hence on the generators $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. We will denote their grading by $\text{gr}_{\text{HF}}(\mathbf{x}) \in \mathbb{Z}/2$.

Let \mathbb{A} and \mathbb{B} be the subspaces of $H_1(\Sigma; \mathbb{R})$ spanned by the alpha and beta curves respectively. We have canonical isomorphisms

$$(37) \quad \mathbb{A} = \ker(H_1(\Sigma; \mathbb{R}) \rightarrow H_1(U_\alpha; \mathbb{R})) \cong H_1(\mathbb{T}_\alpha; \mathbb{R}),$$

$$(38) \quad \mathbb{B} = \ker(H_1(\Sigma; \mathbb{R}) \rightarrow H_1(U_\beta; \mathbb{R})) \cong H_1(\mathbb{T}_\beta; \mathbb{R}).$$

Further, \mathbb{A} can be canonically identified with the tangent space to \mathbb{T}_α at any point, and \mathbb{B} to the tangent space to \mathbb{T}_β at any point.

We define orientations on the moduli spaces $\widehat{\mathcal{M}}(\phi)$ as in Section 3.2, by making the following choices:

- We choose Lie group Pin structures on the tori \mathbb{T}_α and \mathbb{T}_β , as in Example 2.14. To specify these, it suffices to choose them on the tangent spaces to \mathbb{T}_α and \mathbb{T}_β at arbitrary points; i.e., we need Pin structures $P_\alpha^\#$ on \mathbb{A} and $P_\beta^\#$ on \mathbb{B} . For now, we let $P_\alpha^\#$ and $P_\beta^\#$ be arbitrary;
- For the coupled orientation on $(\mathbb{T}_\alpha, \mathbb{T}_\beta)$, we choose it so that the resulting $\mathbb{Z}/2$ -grading gr on the Heegaard Floer complex (constructed as in Section 3.2) is given by the formula:

$$(39) \quad \text{gr}(\mathbf{x}) = \text{gr}_{\text{HF}}(\mathbf{x}) + g + b_1(Y) \pmod{2}.$$

The intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ decompose according to Spin^c structures on the 3-manifold Y . A subset of those, $\mathbf{S} \subseteq \text{Spin}^c(Y)$, is the set of equivalence classes described in Section 3.4. (Not all the Spin^c structures have to be represented in a given Heegaard diagram.)

Fix $\mathfrak{s} \in \text{Spin}^c(Y)$. We assume that the Heegaard diagram is strongly \mathfrak{s} -admissible in the sense of [30, Definition 4.10]. There are several versions of the Heegaard Floer complex, differing in how we keep track of the basepoint z . A preliminary infinity version $CF_{\text{prel}}^\infty(\mathcal{H}, \mathfrak{s}) =$

$CF_{\text{prel}}^\infty(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s})$ is generated by pairs $[\omega_{\mathbf{x}}, i]$, where $\omega_{\mathbf{x}}$ is an orientation of $\det(\bar{\partial}_{\mathbf{x}})$ for $\mathbf{x} \in \mathfrak{s} \subset \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $i \in \mathbb{Z}$. The differential is given by

$$\partial[\omega_{\mathbf{x}}, i] = \sum_{\mathbf{y} \in \mathfrak{s}} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} (\#\widehat{\mathcal{M}}(\phi)) \cdot [\omega_{\mathbf{y}}, i - n_z(\phi)].$$

Here, $n_z(\phi)$ denotes the intersection product of ϕ with the class of the divisor $\{z\} \times \text{Sym}^{g-1}(\Sigma)$. There is a $\mathbb{Z}[U]$ -action on the complex, given by $U \cdot [\omega_{\mathbf{x}}, i] = [\omega_{\mathbf{x}}, i - 1]$.

Other Heegaard Floer complexes are:

- $CF_{\text{prel}}^-(\mathcal{H}, \mathfrak{s})$, the subcomplex of $CF_{\text{prel}}^\infty(\mathcal{H}, \mathfrak{s})$ generated by $[\omega_{\mathbf{x}}, i]$ with $i < 0$;
- $CF_{\text{prel}}^+(\mathcal{H}, \mathfrak{s})$, the corresponding quotient complex, generated by $[\omega_{\mathbf{x}}, i]$ with $i \geq 0$;
- $\widehat{CF}_{\text{prel}}(\mathcal{H}, \mathfrak{s})$, generated directly by $\omega_{\mathbf{x}}$ and counting only disks in classes ϕ with $n_z(\phi) = 0$.

To ensure that the differentials are finite, we need to impose a strong admissibility condition on the Heegaard diagram, as in [30, Section 4.2]. If we only care about $CF_{\text{prel}}^+(\mathcal{H}, \mathfrak{s})$ and $\widehat{CF}_{\text{prel}}(\mathcal{H}, \mathfrak{s})$, a weaker admissibility condition is sufficient; see [30, Theorem 4.15].

To show that the differentials square to zero, Ozsváth and Szabó prove that the contributions from disk and sphere bubbles are zero; see [30, Section 3.7]. Their proof uses degeneration arguments and the fact that the counts of 0-dimensional moduli spaces of bubbles are constant in families. As such, the proof works just as well for any signed counts with this property; in particular, for those induced by Pin structures as in Section 3.3.

From here, we obtain preliminary Heegaard Floer homology groups:

$$HF_{\text{prel}}^\infty(\mathcal{H}, \mathfrak{s}), \quad HF_{\text{prel}}^-(\mathcal{H}, \mathfrak{s}), \quad HF_{\text{prel}}^+(\mathcal{H}, \mathfrak{s}), \quad \widehat{HF}_{\text{prel}}(\mathcal{H}, \mathfrak{s}).$$

4.2. The absolute $\mathbb{Z}/2$ -grading. Before moving forward, let us make some comments about the absolute $\mathbb{Z}/2$ -grading gr_{HF} used in (39). In [29, Section 10.4], this grading is defined by asking that a version of Heegaard Floer homology with twisted coefficients, \underline{HF}^∞ , is supported in even degrees.

To see that the argument is not circular, note that it suffices to first define and compute \underline{HF}^∞ without signs (i.e., using \mathbb{F}_2 instead of \mathbb{Z}). This still defines the absolute $\mathbb{Z}/2$ -grading, and hence a sign on the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and a coupled orientation on $(\mathbb{T}_\alpha, \mathbb{T}_\beta)$. We then use it to upgrade the Floer complex to \mathbb{Z} coefficients.

For the purpose of studying naturality, we shall however need a more concrete definition of the absolute $\mathbb{Z}/2$ grading gr_{HF} , which makes no reference to the calculation of \underline{HF}^∞ ; see [8, Section 2.4] and [33, Section 3]. An orientation of \mathbb{T}_α is equivalent to one of \mathbb{A} , and one of \mathbb{T}_β to one of \mathbb{B} . Thus, to define gr_{HF} it suffices to specify a coupled orientation on (\mathbb{A}, \mathbb{B}) .

By associating to each alpha or beta curve the disk it bounds in the respective handlebody, and combining this with Lefschetz duality, we get canonical isomorphisms

$$(40) \quad \mathbb{A} \cong H_2(U_\alpha, \partial U_\alpha; \mathbb{R}) \cong H^1(U_\alpha; \mathbb{R}),$$

$$(41) \quad \mathbb{B} \cong H_2(U_\beta, \partial U_\beta; \mathbb{R}) \cong H^1(U_\beta; \mathbb{R}).$$

Therefore, \mathbb{A} is the dual vector space to $H_1(U_\alpha; \mathbb{R})$ and \mathbb{B} to $H_1(U_\beta; \mathbb{R})$. We have a Mayer-Vietoris sequence

$$(42) \quad 0 \rightarrow H_2(Y; \mathbb{R}) \rightarrow H_1(\Sigma; \mathbb{R}) \rightarrow H_1(U_\alpha; \mathbb{R}) \oplus H_1(U_\beta; \mathbb{R}) \rightarrow H_1(Y; \mathbb{R}) \rightarrow 0.$$

An orientation on a vector space determines one on its dual. Since $H_2(Y; \mathbb{R})$ and $H_1(Y; \mathbb{R})$ are related by Poincaré duality, there is a coupled orientation on them. Furthermore, $H_1(\Sigma; \mathbb{R})$

is a symplectic vector space and has a canonical orientation. Combining these, we obtain an orientation on $H_1(U_\alpha; \mathbb{R}) \oplus H_1(U_\beta; \mathbb{R})$ and hence on $\mathbb{A} \oplus \mathbb{B}$. This produces a coupled orientation on (\mathbb{A}, \mathbb{B}) using Lemma 2.18 and Conventions 2.22. (In [8, Section 2.4], at the end they multiply the orientation on $\mathbb{A} \oplus \mathbb{B}$ by $(-1)^{g(g-1)/2}$. This final sign change is the same as that needed to go from the concatenated to the shuffled orientation; see Section 2.2. In our context, the sign change is automatically implemented by the transition to a coupled orientation, by Convention 2.22.)

It is not hard to see that this definition of gr_{HF} can be rephrased more simply as follows. Since \mathbb{A} and \mathbb{B} are Lagrangian subspaces in the vector space $H_1(\Sigma; \mathbb{R})$, equipping the latter with an inner product produces an isomorphism

$$\tau_{\mathbb{A}, \mathbb{B}} : \mathbb{A} \rightarrow \mathbb{B}$$

as in Lemma 2.34 (1). The coupled orientation on (\mathbb{A}, \mathbb{B}) is obtained by choosing any orientation on \mathbb{A} and pairing it with its image under $\tau_{\mathbb{A}, \mathbb{B}}$. We will take this as the working definition of gr_{HF} in our paper.

Remark 4.1. There is also an absolute $\mathbb{Z}/2$ grading on monopole Floer homology, defined by Kronheimer and Mrowka in [19, Subsection 22.4]. This corresponds to $\text{gr}_{HF} + b_1(Y)$, which is an expression that will appear naturally in the statement of Proposition 4.7 below.

4.3. Adjusting the definition. The isomorphism classes of the preliminary Heegaard Floer groups turn out to be three-manifold invariants. However, in the definition we had to choose Pin structures $P_\alpha^\#$ on \mathbb{A} and $P_\beta^\#$ on \mathbb{B} . To make the Heegaard Floer groups themselves into *canonical* invariants, we need to get rid of their dependence on this data. We will do this by tensoring with certain lines (rank one free abelian groups).

Let us make use of the concept of a coupled Spin structure defined in Section 2.3. We view the vector spaces \mathbb{A} and \mathbb{B} as vector bundles over a point. Since they are Lagrangian subspaces of $H_1(\Sigma; \mathbb{R})$, Lemma 2.34 (2) gives a canonical coupled Spin structure on (\mathbb{A}, \mathbb{B}) ; we denote this structure by $S_{\text{can}}^\#$.

In Section 4.1 we also introduced a coupled orientation on (\mathbb{A}, \mathbb{B}) : the one defining the grading gr_{HF} . Given the description in Section 4.2, we see that the coupled Spin structure $S_{\text{can}}^\#$ produces this coupled orientation on (\mathbb{A}, \mathbb{B}) , under the natural map $\text{Spin}(g, g) \rightarrow \text{SO}(g, g)$.

As noted in Remark 2.28, once we fix the coupled orientation, a pair of Pin structures $P_\alpha^\#$ on \mathbb{A} and $P_\beta^\#$ on \mathbb{B} gives rise to a coupled Spin structure on (\mathbb{A}, \mathbb{B}) , which we denote simply by $(P_\alpha^\#, P_\beta^\#)$. The space of coupled Spin structures on (\mathbb{A}, \mathbb{B}) compatible with the given coupled orientation is homotopy equivalent to $\mathbb{R}P^\infty$. Since $\pi_1(\mathbb{R}P^\infty) = \mathbb{Z}/2$, there are exactly two homotopy classes of paths from $(P_\alpha^\#, P_\beta^\#)$ to $S_{\text{can}}^\#$. If these classes are denoted γ_1 and γ_2 , we let $\ell(P_\alpha^\#, P_\beta^\#)$ be the rank one abelian group generated by γ_1 and γ_2 modulo the relation $\gamma_1 = -\gamma_2$.

For $\circ \in \{\widehat{}, +, -, \infty\}$, we define the chain complex

$$CF^\circ(\mathcal{H}, \mathfrak{s}) = CF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s}) \otimes \ell(P_\alpha^\#, P_\beta^\#)$$

with homology

$$HF^\circ(\mathcal{H}, \mathfrak{s}) = HF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s}) \otimes \ell(P_\alpha^\#, P_\beta^\#).$$

As we shall soon see, these groups are natural invariants of the three-manifold and the Spin^c structure. We will often denote them by $HF^\circ(Y, \mathfrak{s})$.

For now, let us prove the first step towards naturality.

Proposition 4.2. *For $\circ \in \{\hat{\cdot}, +, -, \infty\}$, the homology groups $HF^\circ(\mathcal{H}, \mathfrak{s})$ are independent of the choice of Pin structures $P_\alpha^\#$ and $P_\beta^\#$, up to canonical isomorphism.*

Proof. We will prove this at the level of chain complexes. Let $(R_\alpha^\#, R_\beta^\#)$ be another choice for the pair of Pin structures on (\mathbb{A}, \mathbb{B}) . We claim that the complexes $CF^\circ(\mathcal{H}, \mathfrak{s})$ defined from $(P_\alpha^\#, P_\beta^\#)$ are canonically isomorphic to those from $(R_\alpha^\#, R_\beta^\#)$.

Pick paths of Pin structures γ from $P_\alpha^\#$ to $R_\alpha^\#$ and η from $P_\beta^\#$ to $R_\beta^\#$. They induce an isomorphism of orientation spaces $o(\mathbf{x})$ from those defined using $P_\alpha^\#$ and $P_\beta^\#$ to those defined using $R_\alpha^\#$ and $R_\beta^\#$. The isomorphism is compatible with the differentials, so we get an isomorphism between the corresponding preliminary Heegaard Floer chain groups $CF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s})$. Moreover, combining γ and η we obtain a path (γ, η) of coupled Spin structures from $(P_\alpha^\#, P_\beta^\#)$ to $(R_\alpha^\#, R_\beta^\#)$. Using concatenation with the paths to $S_{\text{can}}^\#$, this gives an isomorphism of lines $\ell(P_\alpha^\#, P_\beta^\#) \cong \ell(R_\alpha^\#, R_\beta^\#)$. Altogether, we get the desired isomorphism between the respective $CF^\circ(\mathcal{H}, \mathfrak{s})$.

We need to show that this last isomorphism does not depend on the choices of γ and η . Since the space of Pin structures has $\pi_1 = \mathbb{Z}/2$ (cf. Remark 2.13), up to homotopy there are only two possible choices for γ and two for η . Changing such a choice results in a sign change for the isomorphism on orientation spaces and hence on $CF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s})$ (see Lemma 3.4), but this is cancelled by another sign change for the isomorphism between $\ell(P_\alpha^\#, P_\beta^\#)$ and $\ell(R_\alpha^\#, R_\beta^\#)$. \square

Convention 4.3. We will sometimes drop the Spin^c structure \mathfrak{s} from notation and consider preliminary chain groups $CF_{\text{prel}}^\circ(\mathcal{H}) = CF_{\text{prel}}^\circ(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ with homology $HF_{\text{prel}}^\circ(\mathcal{H}) = HF_{\text{prel}}^\circ(\mathbb{T}_\alpha, \mathbb{T}_\beta)$, as well as adjusted chain groups $CF^\circ(\mathcal{H}) = CF^\circ(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ with homology $HF^\circ(\mathcal{H}) = HF^\circ(\mathbb{T}_\alpha, \mathbb{T}_\beta)$. By these we always mean the direct sum over the set \mathbf{S} of Spin^c structures that are represented by intersection points in the diagram \mathcal{H} . Thus, $HF^\circ(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ may be different from $HF^\circ(Y)$, the latter being the direct sum of $HF^\circ(Y, \mathfrak{s})$ over all Spin^c structures \mathfrak{s} .

4.4. More on Pin structures. Our definition of Heegaard Floer homology relies on Pin structures $P_\alpha^\#$ and $P_\beta^\#$. Although in Proposition 4.2 we showed that in the end different choices produce canonically isomorphic groups, in practice it is useful to work with concrete Pin structures. One way to produce them is as follows.

Definition 4.4. Choose some additional data a , consisting of an ordering of the alpha curves together with an orientation of each. This trivializes the tangent space $T_{\mathbf{x}}\mathbb{T}_\alpha$, and lets us define a trivial Spin structure $S_a^\#$ on \mathbb{T}_α . We can think of it as the product of the trivial Lie group Spin structures on each α_i (with the product taken in the given order). We let $P_a^\#$ be the Pin structure on \mathbb{A} induced from $S_a^\#$.

Similarly, we let $P_b^\#$ be the Pin structure on \mathbb{B} obtained from data b that consists of an ordering of the beta curves together with an orientation of each.

Remark 4.5. Data a and b as above also specifies orientations of \mathbb{A} and \mathbb{B} , and hence a coupled orientation. It is convenient to choose a and b so that this agrees with the coupled orientation described in Section 4.2.

Going back to the case of general Pin structures, it is worth emphasizing the following fact.

Proposition 4.6. *Let $(P_\alpha^\#, P_\beta^\#)$ and $(R_\alpha^\#, R_\beta^\#)$ be two pairs of Pin structures that represent the same coupled Spin structure on (\mathbb{A}, \mathbb{B}) . Then, for each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the orientation lines $o(\mathbf{x})$ constructed from $(P_\alpha^\#, P_\beta^\#)$ and $(R_\alpha^\#, R_\beta^\#)$ are canonically isomorphic.*

Proof. In the proof of Proposition 4.2 we argued that the groups $o(\mathbf{x}) \otimes \ell(P_\alpha^\#, P_\beta^\#)$ are canonically isomorphic, for any choices of Pin structures. In the particular case when the Pin structures produce the same coupled Spin structure, the lines $\ell(P_\alpha^\#, P_\beta^\#)$ are canonically isomorphic by definition. Thus, the same holds for $o(\mathbf{x})$. \square

Since the isomorphisms in Proposition 4.6 are compatible with the differentials, it follows that the preliminary groups $CF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s})$ constructed from those pairs are canonically isomorphic. Thus, what matters in the construction of $CF_{\text{prel}}^\circ(\mathcal{H}, \mathfrak{s})$ is only the coupled Spin structure, not its Pin representatives.

4.5. Describing some generators. Observe that the generators of $CF^\circ(\mathcal{H}, \mathfrak{s})$ are elements of the lines $o(\mathbf{x}) \otimes \ell(P_\alpha^\#, P_\beta^\#)$ for $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The following result will be useful.

Proposition 4.7. *If $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ has $\text{gr}(\mathbf{x}) = g \pmod{2}$, i.e. $\text{gr}_{\text{HF}}(\mathbf{x}) = b_1(Y) \pmod{2}$, then the line $o(\mathbf{x}) \otimes \ell(P_\alpha^\#, P_\beta^\#)$ admits a canonical trivialization.*

Proof. In view of the definition of $\text{gr}(\mathbf{x})$ in Section 3.2, the condition $\text{gr}(\mathbf{x}) = g \pmod{2}$ means that the orientation on $T_{\mathbf{x}}\mathbb{T}_\alpha \oplus T_{\mathbf{x}}\mathbb{T}_\beta$ (coming from the coupled orientation) agrees with the one on $T_{\mathbf{x}}M$. Suppose $\mathbf{x} = \{x_1, \dots, x_g\}$ with $x_i \in \alpha_i \cap \beta_{\sigma(i)}$, for some permutation σ . We view each pair $(\alpha_i, \beta_{\sigma(i)})$ as Lagrangians in Σ and we choose the coupled orientation on them so that it agrees with the one on Σ ; that is, so that $\text{gr}(x_i) = 1$ for all i . The product of these coupled orientations gives the coupled orientation on $(\mathbb{T}_\alpha, \mathbb{T}_\beta)$.

Let us now choose some additional data a as in Definition 4.4, consisting of an ordering of the alpha curves together with an orientation of each. With our \mathbf{x} fixed, the data a produces a similar data b on the beta curves, consisting of an ordering of them together with orientations. Indeed, for the ordering, we simply transplant the ordering of the alpha curves using the permutation σ (determined by the generator \mathbf{x}). For the orientations, we choose them so that we get the coupled orientation on each $(\alpha_i, \beta_{\sigma(i)})$ that we already specified. As in Definition 4.4, let $P_a^\#$ and $P_b^\#$ be the Pin structures induced from a and b .

Let us first prove the lemma in the case where the groups are defined from Pin structures $P_\alpha^\# = P_a^\#$ and $P_\beta^\# = P_b^\#$ for some data a (and the induced data b). In that case, recall that the orientation space $o(\mathbf{x})$ is associated to an index problem on the cap H , with boundary conditions interpolating between $T_{\mathbf{x}}\mathbb{T}_\alpha$ and $T_{\mathbf{x}}\mathbb{T}_\beta$. This is the direct sum of the index problems corresponding to each x_i , and therefore $o(\mathbf{x})$ can be canonically identified with the tensor product of the orientation spaces $o(x_i)$ for $i = 1, \dots, g$; compare Example 3.27. (Since we trivialized the tangent spaces, we have Spin and not just Pin structures, so Example 3.27 is applicable here.) Because $\text{gr}(x_i) = 1$, Lemma 3.21 gives a canonical trivialization for each of $o(x_i)$. Tensoring them together we get a trivialization of $o(\mathbf{x})$. Furthermore, using Remark 2.36 about the behavior of canonical Spin structures under direct sum, we deduce that $(P_a^\#, P_b^\#)$ represents the canonical coupled Spin structure on (\mathbb{A}, \mathbb{B}) . Thus, by choosing the constant path we trivialize the line $\ell(P_a^\#, P_b^\#)$.

Let us now consider the case where $P_\alpha^\#$ and $P_\beta^\#$ are arbitrary. From the proof of Proposition 4.2 we know that there is a canonical isomorphism between the line $o(\mathbf{x}) \otimes \ell(P_\alpha^\#, P_\beta^\#)$ and

the similar line defined using $(P_a^\#, P_b^\#)$. Since the latter line is trivialized, so is the former. However, we need to make sure that this trivialization does not depend on the chosen data a ; that is, if we have two sets of data a and a' , their canonical trivializations (defined in the previous paragraph) are related by the isomorphism from the proof of Proposition 4.2.

If the change in a comes from a change in the orientation of one of the curves α_i , note that there must be a corresponding change in the orientation of the curve $\beta_{\sigma(i)}$. Overall, the coupled orientation and the coupled Spin structure on the pair $(\alpha_i, \beta_{\sigma(i)})$ are preserved. It follows that we are in the situation of Proposition 4.6: The coupled Spin structure being fixed, the spaces $o(x_i)$ (before and after we made the orientation changes) are canonically identified. The origin of these identifications is the same as that in the proof of Proposition 4.2, so everything is compatible.

It remains to study the case of a change in the ordering of the alpha curves. It suffices to consider a transposition, say changing the ordering (α_1, α_2) to (α_2, α_1) , which must come in tandem with changing $(\beta_{\sigma(1)}, \beta_{\sigma(2)})$ to $(\beta_{\sigma(2)}, \beta_{\sigma(1)})$. This produces a genuine change in the Pin structures. We obtain a local problem near \mathbf{x} , and we can restrict attention to small neighborhoods of each $x_i = \alpha_i \cap \beta_{\sigma(i)}$ in Σ , whose product is a neighborhood of \mathbf{x} in $\text{Sym}^g(\Sigma)$, away from the diagonal. Furthermore, without loss of generality, we can assume that σ is the identity, and that we are in two dimensions (that is, we focus on the neighborhoods of x_1 and x_2 , as the rest is not affected by the ordering change).

The problem we are left with is the following. For $i = 1, 2$, we are given curves α_i and β_i in $\mathbb{C} \cong \mathbb{R}^2$, intersecting transversely at x_i . We may as well assume α_i and β_i are lines through $x_i = 0$. We denote $\mathbf{x} = (0, 0) \in \mathbb{C}^2$, $\mathbb{T}_\alpha = \alpha_1 \times \alpha_2$, $\mathbb{T}_\beta = \beta_1 \times \beta_2$. Because the coupled orientations are fixed, and we are otherwise free to choose orientations on α_i , we can assume that (α_1, β_1) and (α_2, β_2) are both positive bases of \mathbb{R}^2 . Let a and b be the data consisting of the chosen orientations of the curves, together with the ordering (α_1, α_2) ; let a' be the similar data with the ordering (α_2, α_1) . These produce Pin structures $P_a^\#$ and $P_{a'}^\#$ on $T_{\mathbf{x}}\mathbb{T}_\alpha$. Similarly, the orderings (β_1, β_2) and (β_2, β_1) give data b and b' and produce Pin structures $P_b^\#$ and $P_{b'}^\#$ on $T_{\mathbf{x}}\mathbb{T}_\beta$. The lines $o(\mathbf{x}) \otimes \ell(P_a^\#, P_b^\#)$ defined using $(P_a^\#, P_b^\#)$ and $(P_{a'}^\#, P_{b'}^\#)$ are related by two isomorphisms, and our problem is to show that these isomorphisms are the same.

Specifically, one isomorphism comes from viewing $o(\mathbf{x})$ as a tensor product in two different ways: $o(x_1) \otimes o(x_2) \cong o(x_2) \otimes o(x_1)$. We combine this with the identity on $\ell(P_a^\#, P_b^\#)$, as the coupled Spin structure is the same for both orderings.

The other isomorphism is described in the proof of Proposition 4.2, where we had to choose paths γ from $P_a^\#$ to $P_{a'}^\#$ and η from $P_b^\#$ to $P_{b'}^\#$. We let γ be given by the Pin structures coming from the trivializations of

$$T_{\mathbf{x}}\mathbb{T}_\alpha \cong T_{x_1}\alpha_1 \oplus T_{x_2}\alpha_2 \cong \mathbb{R}^2$$

using the bases

$$((\cos t, \sin t), (-\sin t, \cos t)), \quad t \in [0, \pi/2].$$

Similarly, we let η come from the trivializations of

$$T_{\mathbf{x}}\mathbb{T}_\beta \cong T_{x_1}\beta_1 \oplus T_{x_2}\beta_2 \cong \mathbb{R}^2$$

using the same bases as above. The coupled Spin structure is kept constant throughout this process, so $\ell(P_a^\#, P_b^\#)$ is trivialized with the constant path. Furthermore, following these paths we see decompositions of $o(\mathbf{x})$ as products of two lines, starting with $o(x_1) \otimes o(x_2)$ and

ending with $o(x_2) \otimes o(x_1)$. In other words, we get the same isomorphism as the first one we described. \square

4.6. The homology action. Following [30, Section 4.2.5], we can equip our Heegaard Floer homology groups $HF^\circ(Y, \mathfrak{s})$ with an action of the exterior algebra $\Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors})$. We have an isomorphism

$$H^1(\mathbf{P}(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z}) \cong \mathbb{Z} \oplus (H^1(Y; \mathbb{Z})/\text{Tors}).$$

Given a cocycle $\zeta \in Z^1(\mathbf{P}(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z})$, we define its action on the preliminary Floer complex $CF_{\text{prel}}^\infty(\mathcal{H}, \mathfrak{s})$ by

$$\partial[\omega_{\mathbf{x}}, i] = \sum_{\mathbf{y} \in \mathfrak{s}} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} \zeta(\phi) \cdot (\#\widehat{\mathcal{M}}(\phi)) \cdot [\omega_{\mathbf{y}}, i - n_z(\phi)].$$

The same proof as in [30, Proposition 4.17] shows that A_γ is a chain map, depends only on the homology class of γ , and satisfies $A_\gamma^2 = 0$ on homology. Restricting to $[\gamma]$ in the $H^1(Y; \mathbb{Z})/\text{Tors}$ yields an action of $\Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors})$ on $HF_{\text{prel}}^\infty(\mathcal{H}, \mathfrak{s})$, which we then tensor with the identity on $\ell(P_\alpha^\#, P_\beta^\#)$ to get one on $HF^\infty(Y, \mathfrak{s})$. The actions on the other versions $HF^\circ(Y, \mathfrak{s})$ are constructed similarly.

4.7. Twisted coefficients. There are yet more variants of Heegaard Floer homology that can be considered; e.g., with twisted coefficients, as in Section 3.5. In the Heegaard Floer setting, the group $G = \pi_2(\mathbf{x}_0, \mathbf{x}_0)$ maps into $\mathbb{Z} \oplus H^1(Y; \mathbb{Z})$, and this map is an isomorphism for $g > 1$; see [30, Proposition 2.15]. The \mathbb{Z} summand captures the quantity $n_z(\phi)$, and is not relevant, because in the differential we already keep track of this quantity. The more relevant group is $\widehat{\pi}_2(\mathbf{x}_0, \mathbf{x}_0)$, which only includes classes ϕ such that $n_z(\phi) = 0$. This group maps to $H^1(Y; \mathbb{Z})$, and produces genuine twisted coefficients. Thus, we obtain groups

$$\underline{HF}^\infty(Y, \mathfrak{s}; A), \quad \underline{HF}^-(Y, \mathfrak{s}; A), \quad \underline{HF}^+(Y, \mathfrak{s}; A), \quad \widehat{HF}(Y, \mathfrak{s}; A)$$

for any module A over $\mathbb{Z}[H^1(Y; \mathbb{Z})]$. Compare [29, Section 8].

In particular, we can consider changes in the Pin structures on the Lagrangians, which correspond to twisting coefficients by modules of the form $A_{\eta, \zeta}$, as in Proposition 3.18. As explained at the end of Section 3.5, different systems of orientations on the moduli spaces are parametrized by an affine space on

$$(43) \quad \text{coker}(H^1(M; \mathbb{Z}/2) \rightarrow H^1(\mathbb{T}_\alpha; \mathbb{Z}/2) \oplus H^1(\mathbb{T}_\beta; \mathbb{Z}/2)).$$

Let $M' = \text{Sym}^g(\Sigma - \{z\})$. We have $\pi_1(\Omega M') = \pi_2(M') = H_2(M') = 0$, and $\pi_1(M') = \pi_1(M) = H_1(\Sigma) = \mathbb{Z}^{2g}$. Let also $\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ be the space of homotopy classes of Whitney disks inside M' . This has $\pi_1(\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta)) \cong H^1(Y; \mathbb{Z})$ for $g > 1$.

We have the short exact sequence of a fibration

$$0 = \pi_1(\Omega M') \rightarrow \pi_1(\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta)) \rightarrow \pi_1(\mathbb{T}_\alpha \times \mathbb{T}_\beta) \rightarrow \pi_1(M') \rightarrow 0.$$

After we abelianize, the exact sequence splits, because $H_1(M) = \mathbb{Z}^{2g}$ is free. Therefore we can dualize it and get a short exact sequence

$$0 \rightarrow H^1(M; \mathbb{Z}/2) \rightarrow H^1(\mathbb{T}_\alpha \times \mathbb{T}_\beta; \mathbb{Z}/2) \rightarrow H^1(\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z}/2) \rightarrow 0$$

so the cokernel in (43) is (at least for $g > 1$)

$$H^1(\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z}/2) \cong \text{Hom}(\pi_1(\widehat{\mathbf{P}}(\mathbb{T}_\alpha, \mathbb{T}_\beta)), \mathbb{Z}/2) \cong \text{Hom}(H^1(Y; \mathbb{Z}), \mathbb{Z}/2).$$




FIGURE 4. A handleslide. We draw the β circles in blue and the γ circles in green. The gray disks are feet of handles that can be connected to other parts of the diagram.

This exactly corresponds to the different *coherent orientation systems* considered by Ozsváth and Szabó in [30, Section 4.2.4]. It was observed there that such systems form an affine space over $\text{Hom}(H^1(Y; \mathbb{Z}), \mathbb{Z}/2)$.

Thus, on a 3-manifold Y , there are $2^{b_1(Y)}$ orientation systems. Our choice of coupled Spin structures picks up a particular one. In [29, Theorem 10.12], Ozsváth and Szabó also choose a particular one, but in a different way, using their study of \widehat{HF}^∞ as a $\mathbb{Z}[H^1(Y; \mathbb{Z})]$ -module. We do not attempt to show that our orientation system corresponds to theirs in general. Rather, let us focus on the case $Y = \#^g(S^1 \times S^2)$, which will become relevant to us in Section 5.

For $\#^g(S^1 \times S^2)$, we have a Heegaard diagram with the β curves isotopic to the α curves; see [30, Lemma 9.1]. There, they choose the orientation system so that

$$(44) \quad \widehat{HF}(\#^g(S^1 \times S^2)) \cong H_*(T^g; \mathbb{Z}).$$

We check that our construction produces the same answer.

Lemma 4.8. *Equation (44) holds for the Heegaard Floer groups defined over \mathbb{Z} in our setting, with $\widehat{HF}(\#^g(S^1 \times S^2))$ equipped with the $\text{gr}_{\widehat{HF}}$ grading.*

Proof. In the genus g diagram where the α_i curve is Hamiltonian isotopic to β_i , the Lagrangian tori \mathbb{T}_α is Hamiltonian isotopic to \mathbb{T}_β . Further, this isotopy preserves the Lie group Pin structures. Thus, we are simply computing $\widehat{HF}(\mathbb{T}_\alpha, \mathbb{T}_\alpha)$, which is isomorphic to $H_*(\mathbb{T}_\alpha; \mathbb{Z})$ by Proposition 3.13. This proves the statement as relatively graded groups. The absolutely graded version also holds, because we know from [29] that it does so when we work with $\mathbb{Z}/2$ coefficients. \square

4.8. Handleslides. In proving the invariance of Heegaard Floer homology, a key role will be played by the maps induced by handleslides, which we discuss here.

There are two kinds of handleslides: among the alpha curves, and among the beta curves. We will describe the maps associated to beta handleslides; those associated to alpha handleslides are similar. Furthermore, we focus on the minus version of HF for concreteness.

Following [30, Section 9], a beta handleslide consists of replacing a curve collection $\beta = \{\beta_1, \dots, \beta_g\}$ by another collection $\gamma = \{\gamma_1, \dots, \gamma_g\}$ such that:

- The curve γ_1 is obtained from β_1 by sliding it over β_2 ;
- For $i > 1$, the curves γ_i is obtained from β_i by a small Hamiltonian isotopy;
- For all i , the curve γ_i intersects β_i in two points, and does not intersect any other beta curve.

See Figure 4. We let $\mathbb{T}_\gamma = \gamma_1 \times \dots \times \gamma_g$ denote the resulting Lagrangian torus.

Recall that in the construction of the preliminary groups $HF_{\text{prel}}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ we used Lie group Pin structures on \mathbb{T}_α and \mathbb{T}_β , coming from Pin structures $P_\alpha^\#$ and $P_\beta^\#$ on the vector spaces \mathbb{A} and \mathbb{B} from (37) and (38). Since the handleslide does not change the Heegaard splitting (only the diagram), the same vector space \mathbb{B} can be identified with the tangent space to \mathbb{T}_γ , and we will use the same Pin structure $P_\beta^\#$ to get one on \mathbb{T}_γ .

In [32], Perutz shows that a handleslide induces a Hamiltonian isotopy $\psi : M \rightarrow M$ taking \mathbb{T}_β to \mathbb{T}_γ . The following lemma shows that our chosen Pin structures on the Lagrangian tori are preserved by ψ .

Lemma 4.9. *The Hamiltonian isotopy ψ induced by a handleslide takes the Lie group Pin structure on \mathbb{T}_β from $P_\beta^\#$ on \mathbb{B} to the Lie group Pin structure on \mathbb{T}_γ coming from the same $P_\beta^\#$ (up to canonical homotopy of Pin structures).*

Proof. Pick a point p_i on each curve β_i , so that $p = p_1 \times \cdots \times p_g$ is a basepoint on \mathbb{T}_β . From the description of ψ in [32], we can arrange so that in a neighborhood of each p_i , the curve γ_i is a small translate of the corresponding β_i (including for $i = 1$); and the isotopy ψ is locally given by the product of these translations. Thus, if we identify $T_p\mathbb{T}_\beta$ and $T_{\psi(p)}\mathbb{T}_\gamma$ to \mathbb{B} , we see that the Pin structures on these two tangent spaces both come from $P_\beta^\#$. To obtain the Lie group Pin structures on the whole Lagrangians \mathbb{T}_β and \mathbb{T}_γ , we translate the ones at the given points.

Let us identify each β_i to a curve (still denoted β_i) on \mathbb{T}_β , namely the product of β_i and the basepoints $\{z_j\}$ for $j \neq i$. For $i = 1, \dots, g$, let $\beta'_i \subset \mathbb{T}_\gamma$ be the image of $\beta_i \subset \mathbb{T}_\beta$ under the Hamiltonian isotopy ψ . The collection of curves $\{\beta'_i\}$ differ from $\{\gamma_i\}$ by a linear shear L on \mathbb{T}_γ . The transformation L takes translations on the torus to other translations, and the image of the Pin structure on \mathbb{T}_β under ψ is invariant under these translations, just like the given Pin structure on \mathbb{T}_γ (the one coming from $P_\beta^\#$). Since these two Pin structures on \mathbb{T}_γ agree at a point, they agree everywhere. \square

Remark 4.10. The fact that we chose Lie group Pin structures on the alpha and beta curves was key in the proof of handleslide invariance. Consider for example the alternative of equipping the Lagrangian tori with the Pin structure induced by the product of bounding Spin structures on each curve. (Recall that we can take products of Spin but not Pin structures; see Remark 2.15.) This “bounding Pin structure” would not have been invariant under reparametrizations of the torus by linear maps $L \in \text{GL}(n, \mathbb{Z})$. For example, up to isomorphism, the bounding Pin structure on T^2 differs from the Lie group Pin structure (which we know is independent of reparametrization) by the class $(1, 1) \in H^1(T^2; \mathbb{Z}/2) = \mathbb{Z}/2 \oplus \mathbb{Z}/2$. Clearly, this class is not preserved by the action of $\text{GL}(n, \mathbb{Z})$; e.g. the shear $(x, y) \mapsto (x, x + y)$ takes it to $(1, 0)$.

Since handleslides give Hamiltonian isotopies of the Lagrangian tori, we can associate continuation maps to them, which are isomorphisms on Floer homology. In our case, *a priori* we get isomorphisms on the preliminary Heegaard Floer homology groups from Section 4.1. To get the adjusted groups from Section 4.3, we tensor the preliminary groups with $\ell(P_\alpha^\#, P_\beta^\#)$. Since this line is the same after doing the handleslide (the space \mathbb{B} stays fixed), we can simply tensor the preliminary map with the identity on $\ell(P_\alpha^\#, P_\beta^\#)$ and obtain an isomorphism

$$\Gamma_{\beta \rightarrow \gamma}^\alpha : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s}) \xrightarrow{\cong} HF^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathfrak{s}).$$

Nevertheless, for the proofs in the rest of the paper, it will be helpful to make a closer connection to the current Heegaard Floer literature. Therefore, we will work with triangle maps instead of continuation maps. Triangle maps are used in the original proof of handleslide invariance, in [30, Section 9].

The treatment of handleslides in [30, Section 9] starts by showing that

$$HF^-(\mathbb{T}_\beta, \mathbb{T}_\gamma, \mathfrak{s}_0) \cong \mathbb{Z}[U] \otimes H_*(T^g; \mathbb{Z}),$$

and picking a generator $\Theta_{\beta, \gamma}$ of the top degree part of this group. (Here, \mathfrak{s}_0 is the torsion Spin^c structure. This is the only Spin^c supported by the β - γ diagram, so we can drop it from the notation, following Convention 4.3.) Then, we consider the product map

$$F_{\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma} : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s}) \otimes HF^-(\mathbb{T}_\beta, \mathbb{T}_\gamma) \rightarrow HF^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathfrak{s})$$

as in (33). We define the following map associated to the handleslide:

$$(45) \quad \Psi_{\beta \rightarrow \gamma}^\alpha : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s}) \rightarrow HF^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathfrak{s}), \quad \mathbf{x} \mapsto F_{\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma}(\mathbf{x} \otimes \Theta_{\beta, \gamma}).$$

It is shown in [30, Section 9] that this map is an isomorphism (for suitable coherent orientation systems).

Let us construct $\Psi_{\beta \rightarrow \gamma}^\alpha$ in our setting, with the signs coming from Lie group Pin structures. The work of Perutz [32] shows that \mathbb{T}_γ is Hamiltonian isotopic to \mathbb{T}_β . Therefore, by Proposition 3.13 (keeping track of the basepoint as in HF^-), we have canonical isomorphisms

$$(46) \quad HF_{\text{prel}}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma) \cong HF_{\text{prel}}^-(\mathbb{T}_\beta, \mathbb{T}_\beta) \cong \mathbb{Z}[U] \otimes H_*(\mathbb{T}_\beta; |\lambda^{\text{top}} T\mathbb{T}_\beta|) \cong \mathbb{Z}[U] \otimes H_*(T^g).$$

To go from the preliminary to the adjusted Heegaard Floer groups, we tensor them with $\ell(P_\beta^\#, P_\beta^\#)$. Note that the pair $(P_\beta^\#, P_\beta^\#)$ gives exactly the canonical coupled Spin structure $S_{\text{can}}^\#$ on (\mathbb{B}, \mathbb{B}) . It follows that the line $\ell(P_\beta^\#, P_\beta^\#)$ is canonically trivialized (by the constant path), so we also have canonical isomorphisms

$$(47) \quad HF^-(\mathbb{T}_\beta, \mathbb{T}_\gamma) \cong HF^-(\mathbb{T}_\beta, \mathbb{T}_\beta) \cong \mathbb{Z}[U] \otimes H_*(\mathbb{T}_\beta; |\lambda^{\text{top}} T\mathbb{T}_\beta|) \cong \mathbb{Z}[U] \otimes H_*(T^g).$$

Further, note that (Σ, β, γ) is a Heegaard diagram for the manifold $\#^g(S^1 \times S^2)$ with $b_1 = g$, and the absolute mod 2 grading gr_{HF} on its Heegaard Floer homology coincides with that on $\mathbb{Z}[U] \otimes H_*(T^g)$ under the above isomorphisms. (This follows from the definition of the mod 2 grading in terms of \underline{HF}^∞ in [29, Section 10.4].) Moreover, in this case formula (39) implies $\text{gr} = \text{gr}_{\text{HF}}$.

The chain complex $CF_*^-(\mathbb{T}_\beta, \mathbb{T}_\gamma)$ has 2^g generators, the same as its homology, so the differential must be zero. Consider the intersection point in the top homological degree:

$$\Theta_{\beta, \gamma} = \{\theta_1, \dots, \theta_g\}$$

with $\theta_i \in \beta_i \cap \gamma_i$ as in Figure 4. Since $\text{gr}(\Theta_{\beta, \gamma}) = g$, Proposition 4.7 tells us that the line

$$o(\Theta_{\beta, \gamma}) \otimes \ell(P_\beta^\#, P_\beta^\#)$$

is canonically trivialized. Since we already know the same about $\ell(P_\beta^\#, P_\beta^\#)$, we get that $o(\Theta_{\beta, \gamma})$ is canonically trivialized. Therefore, we can identify $\Theta_{\beta, \gamma}$ with the positive generator of this orientation space, and thus view it as a generator of the Floer homology $HF_{\text{prel}}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma)$. We pick this $\Theta_{\beta, \gamma}$ to define the handleslide map on preliminary groups as a triangle map, as in (45). We then tensor it with the identity on $\ell(P_\alpha^\#, P_\beta^\#)$ to get the true handleslide map

$$(48) \quad \Psi_{\beta \rightarrow \gamma}^\alpha : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s}) \rightarrow HF^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathfrak{s}).$$

Remark 4.11. In general, triangle maps are defined only on preliminary Heegaard Floer groups (following the recipe in Section 3.6). To define them on the adjusted groups HF^- , we would need maps

$$\ell(P_\alpha^\#, P_\beta^\#) \otimes \ell(P_\beta^\#, P_\gamma^\#) \rightarrow \ell(P_\alpha^\#, P_\gamma^\#).$$

A natural definition of these maps only exists in certain cases, for example when $P_\beta^\# = P_\gamma^\#$ as in our situation.

Convention 4.12. By a slight abuse of notation, when discussing handleslides, from now on we will not focus on the distinction between preliminary and adjusted Heegaard Floer groups. Since $\ell(P_\beta^\#, P_\beta^\#)$ is canonically trivialized we will write $\Theta_{\beta,\gamma}$ for the generator of $HF^-(\mathbb{T}_\beta, \mathbb{T}_\gamma)$ as well, and describe the map in (48) as

$$\Psi_{\beta \rightarrow \gamma}^\alpha = F_{\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma}(\cdot \otimes \Theta_{\beta,\gamma}).$$

The fact that we define it like this on the preliminary groups and then tensor with the identity is implicit.

Proposition 4.13. *The map $\Psi_{\beta \rightarrow \gamma}^\alpha$ from (48) is an isomorphism.*

Proof. This appears as Theorem 9.5 in [30]. We do not repeat the proof, but rather explain what needs to be modified when we use canonical orientations instead of coherent orientations. To show that $\Psi_{\beta \rightarrow \gamma}^\alpha$ is an isomorphism, Ozsváth and Szabó construct an inverse

$$\Psi_{\gamma \rightarrow \delta}^\alpha : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathfrak{s}) \rightarrow HF^-(\mathbb{T}_\alpha, \mathbb{T}_\delta, \mathfrak{s}), \quad \mathbf{x} \mapsto F_{\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\mathbf{x} \otimes \Theta_{\gamma,\delta}),$$

where δ is a collection of curves Hamiltonian isotopic to β (as in Figure 5), and $HF^-(\mathbb{T}_\alpha, \mathbb{T}_\delta)$ is identified with $HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ using this isotopy. Showing that $\Psi_{\gamma \rightarrow \delta}^\alpha \circ \Psi_{\beta \rightarrow \gamma}^\alpha = \text{id}$ boils down to proving the relation

$$(49) \quad F_{\mathbb{T}_\beta, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\Theta_{\beta,\gamma} \otimes \Theta_{\gamma,\delta}) = \Theta_{\beta,\delta}$$

which involves counting index zero holomorphic triangles in the diagram in Figure 5. There is only one such triangle (in the symmetric product), which is the product of the g darkly shaded triangles on the Heegaard surface.

Let us choose counterclockwise orientations on each of the β , γ , and δ curves in Figure 5. We also order the curves in each set with those indexed by 1 and 2 being the ones involved in the handleslide, again as shown in Figure 5. This equips the Lagrangians with orientations, in a way compatible with their existing coupled orientations coming from the grading. In fact, we now have data that specifies a Pin structure (and even a Spin structure) on each Lagrangian, as in Definition 4.4. Furthermore, each curve becomes an oriented Lagrangian on the Heegaard surface, which we also endow with its Lie group Spin structure. Using the given ordering of the curves, we get product Spin structures on the tori. Then, all the vertices of the darkly shaded triangles have $\text{gr} = 1$. In the setting of coherent orientations, we know from Example 3.25 that each of the darkly shaded triangles gets a sign of $+1$; see Figure 3 (b), with α' , β' and γ' replaced by β_i , γ_i , δ_i , in this order. Since the three triangles come with a positive sign, so does their product; see Example 3.27. Relation (49) follows from here. \square

This concludes the description of the triangle map induced by a beta handleslide on HF^- . The case of an alpha handleslide (changing α into a new curve collection γ) is similar, with the triangle map

$$\Psi_\beta^{\alpha \rightarrow \gamma} : HF^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathfrak{s}) \rightarrow HF^-(\mathbb{T}_\gamma, \mathbb{T}_\beta, \mathfrak{s}), \quad \mathbf{x} \mapsto F_{\mathbb{T}_\gamma, \mathbb{T}_\alpha, \mathbb{T}_\beta}(\Theta_{\gamma,\alpha} \otimes \mathbf{x})$$

involving the top degree generator $\Theta_{\gamma,\alpha} \in HF^-(\mathbb{T}_\gamma, \mathbb{T}_\alpha)$.




FIGURE 5. An index zero holomorphic triangle contributing +1.

The maps on the other versions of HF are defined similarly. In the case of HF^∞ and \widehat{HF} , in Equation 47 we need to replace $\mathbb{Z}[U]$ with $\mathbb{Z}[U, U^{-1}]$ and \mathbb{Z} , respectively, and use the corresponding Θ elements. In the case of HF^+ , we still use the Θ element on HF^- to construct the triangle map; see Theorem 8.12 and Section 9 in [30].

4.9. Stabilizations. We now turn to maps induced by stabilizations. The stabilization move involves replacing a Heegaard diagram $\mathcal{H} = (\Sigma, \alpha, \beta)$ with $\mathcal{H}' = (\Sigma', \alpha', \beta')$, where Σ' is the connected sum of Σ with a (two-dimensional) torus E , and the curve collections α', β' are obtained from α resp. β by adding two new curves on E : α_E and β_E , intersecting in a single point \mathbf{x}_E . (Since here we work with embedded Heegaard diagrams, E is a torus inside Y , disjoint from Σ , and we take their connect sum inside Y .) We equip the new curves with the coupled orientation induced from the one on $E \cong \alpha_E \times \beta_E$. We also equip them with the Lie group Pin structures $P_{\alpha, E}^\#$ and $P_{\beta, E}^\#$ (recall that canonical choices of Lie group Pin structures exist in dimension 1). Together, the coupled orientation and Pin structures produce the canonical coupled Spin structure $S_{\text{can}, E}^\#$ on (α_E, β_E) . It follows that the line $\ell(P_{\alpha, E}^\#, P_{\beta, E}^\#)$ is canonically trivialized.

Starting from Pin structures $P_\alpha^\#$ and $P_\beta^\#$ in the setting before stabilization, we can take the product of the coupled Spin structure $(P_\alpha^\#, P_\beta^\#)$ with $S_{\text{can}, E}^\#$ and obtain a coupled Spin structure $(P_{\alpha'}^\#, P_{\beta'}^\#)$ after stabilization. Since canonical coupled Spin structures are preserved by products (cf. Remark 2.36), we have canonical isomorphisms

$$(50) \quad \ell(P_{\alpha'}^\#, P_{\beta'}^\#) \cong \ell(P_\alpha^\#, P_\beta^\#) \otimes \ell(P_{\alpha, E}^\#, P_{\beta, E}^\#) \cong \ell(P_\alpha^\#, P_\beta^\#).$$

Recall from Proposition 4.6 that the orientation spaces $o(\mathbf{x})$ depend on Pin structures only through their combined coupled Spin structure. Coupled Spin structures behave well with regard to products, and in fact we have isomorphisms

$$o(\mathbf{x}) \otimes o(\mathbf{x}_E) \cong o(\mathbf{x} \times \mathbf{x}_E),$$

for all $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. (Compare Example 3.27.)

The proof of stabilization invariance in [30, Section 10] is based on a neck stretching argument. Ultimately, we identify the J -holomorphic strips on $\text{Sym}^{g+1}(\Sigma')$ (contributing to the Heegaard Floer group defined from \mathcal{H}') with products of those on $\text{Sym}^g(\Sigma)$ (contributing to the Heegaard Floer group defined from \mathcal{H}) and on E (with boundaries on α_E and β_E). We end up in the situation of Example 3.27 and obtain an isomorphism

$$(51) \quad HF_{\text{prel}}^-(\mathbb{T}_{\alpha'}, \mathbb{T}_{\beta'}) \cong HF_{\text{prel}}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta) \otimes_{\mathbb{Z}[U]} HF_{\text{prel}}^-(\alpha_E, \beta_E).$$

Strictly speaking, to be in the setting of Example 3.27 one needs to lift the coupled Spin to Spin structures; but the resulting isomorphism does not depend on this lift.

Combining (51) with (50), we also get an isomorphism

$$HF^-(\mathbb{T}_{\alpha'}, \mathbb{T}_{\beta'}) \cong HF^-(\mathbb{T}_{\alpha}, \mathbb{T}_{\beta}) \otimes_{\mathbb{Z}[U]} HF^-(\alpha_E, \beta_E).$$

Observe that (E, α_E, β_E) is a Heegaard diagram for S^3 . The group $HF^-(\alpha_E, \beta_E) \cong HF^-(S^3)$ is a rank one free module over $\mathbb{Z}[U]$ supported in grading $\text{gr}_{\text{HF}} = 0$, which by formula (39) corresponds to $\text{gr} = 1$. Lemma 3.21 (a) says that the line $o(\mathbf{x}_E) \otimes \ell(P_{\alpha, E}^{\#}, P_{\beta, E}^{\#})$ is canonically trivialized; so we can view $HF^-(\alpha_E, \beta_E)$ as being generated by \mathbf{x}_E . (Compare Proposition 4.7.) We obtain the desired stabilization isomorphism:

$$(52) \quad S = \cdot \otimes \mathbf{x}_E : HF^-(\mathbb{T}_{\alpha}, \mathbb{T}_{\beta}) \xrightarrow{\cong} HF^-(\mathbb{T}_{\alpha'}, \mathbb{T}_{\beta'}).$$

The stabilization isomorphisms for the other variants of HF are constructed similarly.

4.10. Invariance. We are now ready to prove a weak version of Theorem 1.1: we establish the existence of isomorphisms between Heegaard Floer homologies associated to different Heegaard diagrams for the same 3-manifold. We leave the proof that this isomorphisms are canonical for the next section.

For simplicity, we will denote by HF° any of the variants of Heegaard Floer homology, $\circ \in \{\hat{\cdot}, +, -, \infty\}$.

Proposition 4.14. *Let Y be a closed, oriented 3-manifold equipped with a basepoint $z \in Y$ and a Spin^c structure \mathfrak{s} . Then, the isomorphism class of the Heegaard Floer homology HF° (defined using Pin structures on the Lagrangian tori) is an invariant of the pair (Y, \mathfrak{s}) .*

Proof. Following [30], we have to prove invariance under three kinds of moves on Heegaard diagrams: Hamiltonian isotopies of the alpha and beta curves, handleslides of the alpha and beta curves, and stabilizations.

Hamiltonian isotopies of the alpha and beta curves induce Hamiltonian isotopies of the Lagrangian tori, and these preserve the Lie group Pin structures. Then, as usual in Floer theory, continuation maps associated to Hamiltonian isotopies give rise to isomorphisms between the Floer homologies.

Handleslides were discussed in Section 4.8. We associate to them triangle maps of the form $\Psi_{\beta \rightarrow \gamma}^{\alpha}$, and Proposition 4.13 shows that these are isomorphisms. Similarly, to stabilizations we associate the isomorphisms S defined in Section 4.9.

So far we have only discussed pointed Heegaard moves, i.e. those supported away from the basepoint z . It is shown in [30, Proposition 7.1] that any two Heegaard diagrams representing the same Y (possibly with different basepoints) become diffeomorphic after a finite sequence of pointed Heegaard moves. It follows that the isomorphism class of HF° is independent of z . \square

Remark 4.15. If we were only interested in Proposition 4.14 (and not in the naturality results discussed in the next section), there would be no need of coupled Spin structures or of adjusting the preliminary Heegaard Floer groups as in Section 4.3. We could have chosen any Lie group Pin structures on the Lagrangian tori, and worked with the preliminary Heegaard Floer groups. Indeed, up to (non-canonical) isomorphism, these are the same as the adjusted groups.

5. NATURALITY

To complete the proof of Theorem 1.1, it remains to show that the isomorphisms induced by (pointed) Heegaard moves on HF° are natural. In other words, if we go from one Heegaard

diagram \mathcal{H} to another diagram \mathcal{H}' by a sequence of moves, the resulting isomorphism between the Heegaard Floer homologies of \mathcal{H} and \mathcal{H}' does not depend on the sequence of moves we chose.

In [16], Juhász, Thurston and Zemke gave a list of conditions that need to be checked to ensure that naturality holds for a 3-manifold invariant defined from Heegaard diagrams. We review here some of their set-up. While they work with sutured manifolds, we will restrict here to based 3-manifolds, for simplicity; this is the class \mathcal{S}_{man} in their notation.

Remark 5.1. When talking about naturality, it is important to fix the basepoint z . Even over \mathbb{F}_2 , when one moves z to another basepoint z' by following a path on Y , there is an isomorphism between the Floer homologies that depends on this path. See [49], [46].

5.1. Strong Heegaard invariants. We define an *isotopy diagram* to be an equivalence class of Heegaard diagrams (for based 3-manifolds), where two diagrams are equivalent if the underlying Heegaard surfaces are the same, their alpha curves differ by isotopies, and their beta curves differ by isotopies. Further, we say that two isotopy diagrams are α -equivalent if they differ by a sequence of α -handleslides. We define β -equivalence similarly. (Of course, all isotopies and handleslides here are supposed to avoid the basepoint.)

We let \mathcal{G} be the oriented (multi-)graph whose vertices are all isotopy diagrams, and whose edges are associated to diagram moves: α -equivalences, β -equivalences, stabilizations, destabilizations, and diffeomorphisms. If we restrict to only one kind of these moves, we denote the respective subgraphs (with the same vertices as \mathcal{G}) by \mathcal{G}_α , \mathcal{G}_β , $\mathcal{G}_{\text{stab}}$ and $\mathcal{G}_{\text{diff}}$.

Definition 5.2. A *distinguished rectangle* in \mathcal{G} is a subgraph

$$\begin{array}{ccc} \mathcal{H}_1 & \xrightarrow{e} & \mathcal{H}_2 \\ \downarrow f & & \downarrow g \\ \mathcal{H}_3 & \xrightarrow{h} & \mathcal{H}_4 \end{array}$$

such that one of the following holds:

- (1) Both e and h are α -equivalences, whereas f and g are β -equivalences;
- (2) Both e and h are equivalences of the same type (α or β), whereas f and g are stabilizations;
- (3) Both e and h are equivalences of the same type, while $f = g$ is a diffeomorphism.
- (4) All of e , f , g and h are stabilizations, with e and h both consisting of replacing the same disk D_1 in the Heegaard surface with the same punctured torus T_1 , and similarly f and g both consisting of replacing a disk D_2 with a punctured torus T_2 . We require that $D_1 \cap D_2 = \emptyset$ and $T_1 \cap T_2 = \emptyset$;
- (5) Both e and h are stabilizations, whereas f and g are diffeomorphisms, taking one stabilization to the other.

We also need the notion of a *simple handleswap*, which is a triangle in \mathcal{G} consisting of an α -equivalence (a handleslide), a β -equivalence (another handleslide), and a diffeomorphism, all being the identity except on a punctured genus two surface, where they look like in Figure 6.

Definition 5.3. Let \mathcal{C} be a category. A *strong Heegaard invariant* is an assignment $F : \mathcal{G} \rightarrow \mathcal{C}$ that takes the vertices of \mathcal{G} (isotopy diagrams) to objects in \mathcal{C} , and edges of \mathcal{G} to isomorphisms in \mathcal{C} , with the following properties:

- (1) **Functoriality:** The restrictions of F to \mathcal{G}_α , \mathcal{G}_β , and $\mathcal{G}_{\text{diff}}$ are functors. Further, if e is a stabilization and e' the corresponding destabilization, then $F(e') = F(e)^{-1}$.




FIGURE 6. A simple handleswap. We reproduce here Figure 4 in [16]. The dashed curve is the boundary of a genus two surface obtained by identifying the boundaries of the grayed circles in pairs. (In each diagram, the circles at the same height are identified via reflection in a vertical axis). The α circles are in red and the β circles in blue.

- (2) **Commutativity:** For every distinguished rectangle in \mathcal{G} , applying F yields a commutative diagram in \mathcal{C} ;
- (3) **Continuity:** If e is an edge in \mathcal{G} associated to a diffeomorphism ψ from the same diagram to itself, and ψ is isotopic to the identity, then $F(e)$ is the identity;
- (4) **Handleswap invariance:** For every simple handleswap

$$\begin{array}{ccc}
 \mathcal{H}_1 & & \\
 \uparrow g & \searrow e & \\
 \mathcal{H}_3 & \xleftarrow{f} & \mathcal{H}_2
 \end{array}$$

the composition $F(g) \circ F(f) \circ F(e)$ is the identity.

Juhász, Thurston and Zemke proved that strong Heegaard invariants are natural:

Theorem 5.4 (Theorem 2.38 in [16]). *Let $F : \mathcal{G} \rightarrow \mathcal{C}$ be a strong Heegaard invariant. If two isotopy diagrams \mathcal{H}_1 and \mathcal{H}_2 are related by a path of arrows in \mathcal{G} , then the isomorphism from $F(\mathcal{H}_1)$ to $F(\mathcal{H}_2)$ induced by composing the morphisms along the path is independent of the choice of this path.*

In [16, Theorem 2.33 (2)], they further proved that the HF^- for $\circ \in \{\hat{}, +, -, \infty\}$, defined without signs, are strong Heegaard invariants into the category of $\mathbb{F}_2[U]$ -modules. We need to show the same thing with signs, i.e., that they are strong Heegaard invariants into the category of $\mathbb{Z}[U]$ -modules. (A weaker version of this claim, in the projective category of $\mathbb{Z}[U]$ -modules where morphisms are defined up to a sign, was proved by Gartner in [12].)

5.2. Loops of handleslides. In our context it is helpful to decompose each α - and β -equivalence f into a sequence of handleslides, and define $F(f)$ as the composition of maps

associated to those handleslides. In [16, Appendix A], Juhász, Thurston and Zemke describe a set of conditions that produce a strong Heegaard invariant, using maps induced by handleslides instead of equivalences. More details for their proofs appear in the work of Qin [35].

First, if f is an equivalence, to prove that the map $F(f)$ is well-defined (that is, it does not depend on how we present f as a composition of handleslides), we would need to show that a loop of handleslides produces the identity. Definition A.3 and Proposition A.5 in [16] give a finite set of types of such loops that generate all others. To describe them, let us first define an *attaching set* on a based surface (Σ, z) to be an isotopy class of an unordered collection of g disjoint simple closed curves β_1, \dots, β_g on $\Sigma \setminus \{z\}$ that are linearly independent in $H_1(\Sigma)$. For example, the collections α and β in a Heegaard diagram are attaching sets. (Unlike in [16], here we pick β instead of α as the notation for a typical attaching set. This is because β -handleslides, where we fix the α curves, are more commonly considered in Heegaard Floer theory; e.g. in [30] and in Section 4.8 of this paper. Of course, this convention is of little importance.)

Definition 5.5. A *handleslide loop* is one of the following sequences of attaching sets on Σ connected by handleslides:

- (1) A *slide triangle*, formed of three attaching sets of the form

$$\begin{array}{ccc} \{\beta_1, \beta_2\} \cup \vec{\beta} & \text{-----} & \{\beta_2, \beta_3\} \cup \vec{\beta} \\ & \searrow \quad \swarrow & \\ & \{\beta_1, \beta_3\} \cup \vec{\beta} & \end{array}$$

where $\beta_1, \beta_2, \beta_3$ bound a pair-of-pants, and $\vec{\beta}$ is a fixed collection of $g - 2$ curves.

- (2) A *commuting slide square*, formed of four attaching sets of the form


$$\begin{array}{ccc} \{\beta_1, \beta_2, \beta_3, \beta_4\} \cup \vec{\beta} & \text{-----} & \{\beta'_1, \beta_2, \beta_3, \beta_4\} \cup \vec{\beta} \\ | & & | \\ \{\beta_1, \beta_2, \beta'_3, \beta_4\} \cup \vec{\beta} & \text{-----} & \{\beta'_1, \beta_2, \beta'_3, \beta_4\} \cup \vec{\beta} \end{array}$$

where β'_1 is obtained by sliding β_1 over β_2 , and β'_3 is obtained by sliding β_3 over β_4 . Here, $\vec{\beta}$ is a fixed collection of $g - 4$ curves.

- (3) A square of the form


$$\begin{array}{ccc} \{\beta_1, \beta_2, \beta_3\} \cup \vec{\beta} & \text{-----} & \{\beta'_1, \beta_2, \beta_3\} \cup \vec{\beta} \\ | & & | \\ \{\beta''_1, \beta_2, \beta_3\} \cup \vec{\beta} & \text{-----} & \{\beta'''_1, \beta_2, \beta'_3, \beta_4\} \cup \vec{\beta} \end{array}$$

where β'_1 is obtained by sliding β_1 over β_2 , whereas β''_1 is obtained by sliding β_1 over β_3 , and β'''_1 is obtained by sliding β'_1 over β_3 (or, equivalently, sliding β''_1 over β_2). Here, $\vec{\beta}$ is a fixed collection of $g - 3$ curves.




(1)

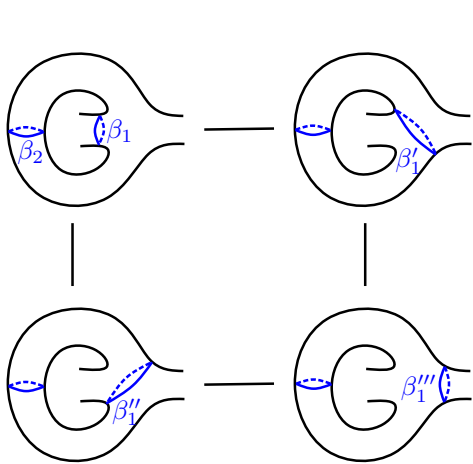
(2)




(3)



(4)



(5)



(6)

FIGURE 7. The six types of handleslide loops.

(4) A square of the form

$$\begin{array}{ccc} \{\beta_1, \beta_2, \beta_3\} \cup \vec{\beta} & \text{---} & \{\beta'_1, \beta_2, \beta_3\} \cup \vec{\beta} \\ | & & | \\ \{\beta_1, \beta_2, \beta'_3\} \cup \vec{\beta} & \text{---} & \{\beta'_1, \beta_2, \beta'_3\} \cup \vec{\beta} \end{array}$$

where β'_1 is obtained by sliding β_1 over β_2 , and β'_3 is obtained by sliding β_3 over β_2 , so that the two slides are happening along arcs that reach β_2 from opposite sides. Here, $\vec{\beta}$ is a fixed collection of $g - 3$ curves.

(5) A square of the form

$$\begin{array}{ccc} \{\beta_1, \beta_2\} \cup \vec{\beta} & \text{---} & \{\beta'_1, \beta_2\} \cup \vec{\beta} \\ | & & | \\ \{\beta''_1, \beta_2\} \cup \vec{\beta} & \text{---} & \{\beta'''_1, \beta_2\} \cup \vec{\beta} \end{array}$$

where β'_1 and β''_1 are both obtained by sliding β_1 over β_2 , but from opposite sides; and β'''_1 is obtained from β'_1 by doing both of these slides over β_2 . Here, $\vec{\beta}$ is a fixed collection of $g - 2$ curves.

(6) A pentagon of the form

$$\begin{array}{ccc} & \{\beta_1, \beta_2, \beta_3\} \cup \vec{\beta} & \\ & / \quad \backslash & \\ \{\beta'_1, \beta_2, \beta_3\} \cup \vec{\beta} & & \{\beta_1, \beta'_2, \beta_3\} \cup \vec{\beta} \\ | & & | \\ \{\beta''_1, \beta_2, \beta_3\} \cup \vec{\beta} & \text{---} & \{\beta'_1, \beta'_2, \beta_3\} \cup \vec{\beta} \end{array}$$

where β'_1 is obtained from β_1 by sliding over β_2 , whereas β'_2 is obtained from β_2 by sliding it over β_3 , and β''_1 is obtained from β'_1 by sliding it over β_3 (or, equivalently, from β_1 by sliding it over β'_2). Here, $\vec{\beta}$ is a fixed collection of $g - 3$ curves.

See Figure 7, which is based on Figures 21 and 22 in [16].

We also need one other loop, which involves stabilizations in addition to handleslides. This is the stabilization slide from [16, Definition 7.7]. It was described there as a triangle where two edges are stabilizations, and one is an equivalence. Since that equivalence is the composition of two handleslides, in our context the triangle becomes a square.

Definition 5.6. A *stabilization β -slide* is a square composed of four Heegaard diagrams that differ locally as in Figure 8. A *stabilization α -slide* is similar, with the alpha and beta curves reversed.

We let \mathcal{G}' be the oriented (multi-)graph whose vertices are all isotopy diagrams, and whose edges are associated to diagram moves: α -handleslides, β -handleslides, stabilizations, destablizations, and diffeomorphisms. Note that \mathcal{G}' is a subgraph of \mathcal{G} , with the same vertices but fewer edges.

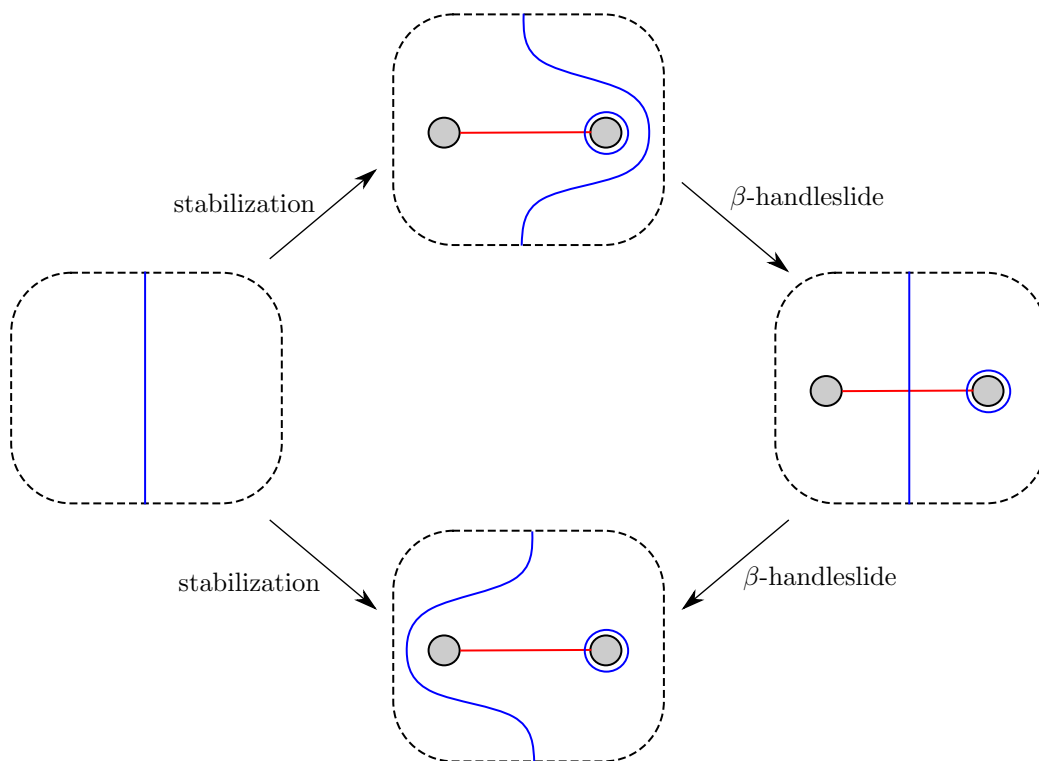


FIGURE 8. A stabilization slide. The two handleslides are performed along the same β -curve, but the second is done through the handle.

Theorem 5.7 (Theorem A.6 in [16]; Theorem 1.3 in [35]). *Let \mathcal{C} be a category. Suppose we are given an assignment $F : \mathcal{G}' \rightarrow \mathcal{C}$ that takes the vertices of \mathcal{G} (isotopy diagrams) to objects in \mathcal{C} , and edges of \mathcal{G}' to isomorphisms in \mathcal{C} , with the following properties:*

- (1) **Functoriality:** *The restriction of F to $\mathcal{G}_{\text{diff}}$ (which is a subgraph of \mathcal{G}') is a functor. Further, if e is a stabilization and e' the corresponding destabilization, then $F(e') = F(e)^{-1}$.*
- (2) **Commutativity along distinguished rectangles:** *For every distinguished rectangle in \mathcal{G}' (i.e., such that the only equivalences involved are handleslides), applying F yields a commutative diagram in \mathcal{C} ;*
- (3) **Continuity:** *condition (3) in Definition 5.3;*
- (4) **Handleswap invariance:** *condition (4) in Definition 5.3;*
- (5) **Commutativity along handleslide loops:** *F' commutes along any of the handleslide loops in Definition 5.5;*
- (6) **Commutativity along stabilization slides:** *F' commutes along any stabilization slide (see Definition 5.6).*

Then, F' extends uniquely to a strong Heegaard invariant $F : \mathcal{G} \rightarrow \mathcal{C}$.

Remark 5.8. The original stabilization slide from [16, Definition 7.7] is an example of a distinguished rectangle in \mathcal{G} , where one horizontal edge is trivial, the other is the composition of the two handleslides, and the vertical edges are the stabilizations. Since one edge is a composition of two handleslides rather than a single handleslide, the stabilization slide does

not count as a distinguished rectangle in \mathcal{G}' . This is the reason for listing it separately from condition (2).

5.3. Heegaard Floer homology as a strong Heegaard invariant.

Proof of Theorem 1.1. We focus on the minus version. In view of Theorem 5.4, it suffices to show that HF^- is a strong Heegaard invariant. Juhász, Thurston and Zemke did this in [16, Theorem 2.33 (2)] for HF^- as an $\mathbb{F}_2[U]$ -module, by directly checking the conditions in Definition 5.3. Here we work with $\mathbb{Z}[U]$ -modules, and we will check the conditions in Theorem 5.7 instead.

We view HF^- as an assignment from \mathcal{G}' to $\mathbb{Z}[U]$ -modules. Notice that in Section 4.1, HF^- is defined on (admissible) Heegaard diagrams, rather than isotopy diagrams. We define it on an isotopy diagram as the colimit of the transitive system of Heegaard Floer homologies over all admissible Heegaard diagrams that produce the given isotopy diagram; see [16, Definition 9.19]. To construct the colimit we use the continuation maps defined in [30, Section 7.3], which we denote by $\Gamma_{\beta \rightarrow \beta'}^\alpha$ or $\Gamma_\beta^{\alpha \rightarrow \alpha'}$.

To have a well-defined colimit, we further need to ensure that the continuation maps relating Heegaard Floer homologies give rise to a transitive system in the sense of [16, Definition 1.1]; i.e., that loops of isotopies induce the identity on Heegaard Floer homology. We can decompose such loops into smaller (triangular) loops such that the quadruple Heegaard diagram for each loop is admissible; this can be done as in the proofs of Lemmas 9.7, 9.10, 9.12 in [16], which in turn are based on Lemma 9.5 in [16]. A typical such triangular loop is

$$\begin{array}{ccc} (\Sigma, \alpha, \beta, z) & \xrightarrow{\quad\quad\quad} & (\Sigma, \alpha, \beta', z) \\ & \swarrow \quad \quad \quad \searrow & \\ & (\Sigma, \alpha, \beta'', z) & \end{array}$$

where the quadruple diagram $(\Sigma, \alpha, \beta, \beta', \beta'', z)$ is admissible. We claim that

$$(53) \quad \Gamma_{\beta'' \rightarrow \beta}^\alpha \circ \Gamma_{\beta' \rightarrow \beta''}^\alpha \circ \Gamma_{\beta \rightarrow \beta'}^\alpha = \text{id}.$$

For a punctured surface (Σ, z) of genus $g > 0$, the identity component of its diffeomorphism group is contractible; see [4], [5]. It follows that the composition of isotopies taking $\beta \rightarrow \beta' \rightarrow \beta'' \rightarrow \beta$ is homotopic to the identity. Admissibility ensures that we can associate a continuation map to this homotopy, which induces a chain homotopy between the maps on Heegaard Floer complexes and the identity; passing to homology, we obtain (53).

The other kind of triangular loop, where the alpha curves vary, is entirely analogous. This completes the proof that we have a transitive system, and therefore HF^- is well-defined on isotopy diagrams.

Remark 5.9. In [16, Section 9.1], to define HF^- over $\mathbb{F}_2[U]$ for isotopy diagrams, the authors proceeded differently: they assigned triangle maps to isotopies, and they also showed that these triangle maps are the same as continuation maps. They proved that the triangle maps form a transitive system by making use of the uniqueness of top degree generators of the form $\Theta_{\beta, \beta'}$ or $\Theta_{\alpha, \alpha'}$ over $\mathbb{F}_2[U]$. Over $\mathbb{Z}[U]$, for arbitrary isotopies, it is unclear which triangle maps to consider, because there are two choices (differing by a sign) for the top degree generators.

Next, we need to assign maps to the edges in \mathcal{G}' (i.e. to moves between isotopy diagrams). For an α - or β -handleslide, we use the maps $\Psi_\beta^{\alpha \rightarrow \gamma}$ and $\Psi_{\beta \rightarrow \gamma}^\alpha$, constructed in Section 4.8. To a diffeomorphism between diagrams we assign the obvious identification of Floer homologies.

To a stabilization we assign the map S from Equation (52), and to a destabilization its inverse.

Convention 5.10. Throughout this proof, when discussing Heegaard moves appearing in a multiple Heegaard diagram, we will draw pictures where each curve is oriented, and there is an ordering of the curves in each attaching set. The orientation choices will be made carefully, so as to be compatible with how the coupled orientations for handleslides and stabilizations were defined in Sections 4.8 and 4.9. (Reversing all the orientations at the same time would produce an equally valid picture.) We will then have data to specify Spin structures on each Lagrangian, as in Definition 4.4. This will enable us to appeal to the product principle from Example 3.27, just as we did in the proof of Proposition 4.13.

We now proceed to check the conditions in Theorem 5.7.

(1) **Functoriality:** This is immediate from the definitions.

(2) **Commutativity along distinguished rectangles:** There are five types of such rectangles to check; see Definition 5.2. Type (1) is commutation between α - and β -handleslides, i.e., a relation of the form

$$\Psi_{\beta'}^{\alpha \rightarrow \alpha'} \circ \Psi_{\beta \rightarrow \beta'}^{\alpha} = \Psi_{\beta \rightarrow \beta'}^{\alpha'} \circ \Psi_{\beta}^{\alpha \rightarrow \alpha'}.$$

This follows from the A_{∞} relations for polygon maps; compare [16, Proposition 9.10 (3)]. The proofs of commutativity along the other four types of distinguished rectangles are just as in [16, Section 9.2]. The arguments are based on choosing specific almost complex structures suitable for diffeomorphisms and stabilizations, and do not involve the signs in an essential way. The only thing of note is for type (4), which is the commutation of two stabilizations; there, we use the fact that coupled Spin structures behave well with regard to direct sums (see Remark 2.36), and in particular that the result of two direct sums does not depend on their order.

(3) **Continuity:** The proof of this over $\mathbb{F}_2[U]$ in [16, Proposition 9.27] is based on relating continuation maps to changes in the almost complex structure, and works just as well over $\mathbb{Z}[U]$.

(4) **Handleswap invariance:** The proof of invariance under a simple handleswap in [16, Section 9.3] (over $\mathbb{F}_2[U]$) is based on degenerations (neck-stretching) to reduce it to calculations in the genus 2 picture in Figure 6. The degenerations involve passing to Lipshitz's cylindrical reformulation of Heegaard Floer homology [21]. These degeneration arguments can be extended to the signed case (over $\mathbb{Z}[U]$); for an explanation, see Section 5.4 below.

This reduces the problem to several curve counts in the genus 2 picture. They involve understanding the effect of the two handleslides in Figure 6. Let us discuss the α -handleslide. (The β -handleslide is treated similarly.) The α -handleslide is shown in Figure 9, which is based on Figure 59 in [16]. We have a triple pointed Heegaard diagram $(\Sigma_0, \alpha'_0, \alpha_0, \beta_0, p_0)$, where Σ_0 is a surface of genus two with basepoint p_0 and three attaching sets:


$$\alpha'_0 = \{\alpha'_1, \alpha'_2\}, \quad \alpha_0 = \{\alpha_1, \alpha_2\}, \quad \beta_0 = \{\beta_1, \beta_2\}.$$

Let

$$\mathbf{a} = \{a_1, a_2\} = \mathbb{T}_{\alpha_0} \cap \mathbb{T}_{\beta_0}, \quad \mathbf{b} = \{b_1, b_2\} = \mathbb{T}_{\alpha'_0} \cap \mathbb{T}_{\beta_0}.$$

Let also $\Theta = \{\theta_1, \theta_2\} \in \mathbb{T}_{\alpha'_0} \cap \mathbb{T}_{\alpha_0}$ be the maximal degree generator, as shown in Figure 9.

We equip the curves on Σ_0 with orientations and orderings as shown in the figure, following Convention 5.10.

FIGURE 9. The α -handleslide from Figure 6.

The triangle map for the α -handleslide is computed (mod 2) in Proposition 9.31 of [16]. We claim that the same calculation holds over \mathbb{Z} . In the \mathbb{F}_2 calculation, there are three points at which explicit moduli spaces are counted (mod 2). We discuss them in turn, explaining how to refine the counts to get the answer in \mathbb{Z} .

The first count appears in Lemma 9.52 in [16]. The lemma states that the differential on $\widehat{CF}(\Sigma_0, \alpha'_0, \alpha_0, p_0)$ vanishes. This is true in the signed case as well, because the rank of the Floer homology is 4, according to Lemma 4.8.

Another explicit count is in Lemma 9.53 in [16], which states that the map

$$\Psi_{\beta_0}^{\alpha_0 \rightarrow \alpha'_0} : \widehat{CF}(\Sigma_0, \alpha_0, \beta_0, p_0) \rightarrow \widehat{CF}(\Sigma_0, \alpha'_0, \beta_0, p_0)$$

satisfies $\Psi_{\beta_0}^{\alpha_0 \rightarrow \alpha'_0}(\mathbf{a}) = \mathbf{b}$. There is a unique holomorphic triangle (in the symmetric product) connecting Θ , \mathbf{a} and \mathbf{b} , shown in Figure 9 as the product of the two darkly shaded triangles. Each of these two triangles has edges oriented as in Figure 3 (b) (with $\alpha'_i, \alpha_i, \beta_i$ playing the roles of α', β', γ'), and therefore (according to Example 3.25) its sign is +1. It follows from Example 3.27 that the product of the triangles has sign +1 as well, so we indeed have $\Psi_{\beta_0}^{\alpha_0 \rightarrow \alpha'_0}(\mathbf{a}) = \mathbf{b}$.

One last count is at the very end of the proof of Lemma 9.58 in [16]. There, the authors compute that

$$(54) \quad \#\mathcal{M}_{(\mathbf{a}, \mathbf{a})}(d) \equiv 1 \pmod{2},$$

where $\mathcal{M}_{(\mathbf{a}, \mathbf{a})}(d)$ is the moduli space of Maslov index 2 holomorphic curves u on $(\Sigma_0, \alpha_0, \beta_0)$ satisfying $u(d) = p_0$, with $d \in [0, 1] \times \mathbb{R}$ being any fixed point. In our setting, we claim that

$$(55) \quad \#\mathcal{M}_{(\mathbf{a}, \mathbf{a})}(d) = 1.$$

The proof of (54) in [16] is based on an argument from [21, Appendix A]. We adapt it here to our purposes. We consider the Heegaard diagrams for $S^1 \times S^2$ shown in Figure 10. On the left we have the class ϕ of a bigon, with $\widehat{\mathcal{M}}(\phi)$ consisting of a single point; this can be positive or negative, depending on the coupled orientations and Pin structures on the curves there. On the right we have the double stabilization of ϕ , which we denote by ϕ'' . Stabilization invariance (as discussed in the proof of Proposition 4.14) implies that

$$\#\widehat{\mathcal{M}}(\phi'') = \#\widehat{\mathcal{M}}(\phi).$$




FIGURE 10. (a) A bigon on a diagram for $S^1 \times S^2$. (b) Its double stabilization.




FIGURE 11. A triple diagram appearing in a handleslide loop of type (2).

On the other hand, by stretching the neck in the cylindrical reformulation, we see that for large neck length we must have

$$\widehat{\mathcal{M}}(\phi'') = \mathcal{M}_{(\mathbf{a}, \mathbf{a})}(d) \times \widehat{\mathcal{M}}(\phi).$$

This proves (55).

The rest of the arguments in the proof of handleswap invariance go through just as in [16, Section 9.3].

(5) **Commutativity along handleslide loops:** We need to check each of the six types of loops from Figure 7. The first type is the most complicated, so we discuss the others first.

Consider a loop of type (2). We change the four attaching sets appearing in the commuting square by small isotopies, so that each curve intersects the curves obtained from it by isotopies or handleslides transversely at two points. After this change we re-label the attaching sets as

$$(56) \quad \begin{array}{ccc} \beta & \text{---} & \gamma \\ | & & | \\ \epsilon & \text{---} & \delta \end{array}$$

so that

$$\begin{aligned} \beta &= \{\beta_1, \beta_2, \beta_3, \beta_4\} \cup \vec{\beta}, & \gamma &\approx \{\beta'_1, \beta_2, \beta_3, \beta_4\} \cup \vec{\beta} \\ \epsilon &\approx \{\beta_1, \beta_2, \beta'_3, \beta_4\} \cup \vec{\beta}, & \delta &\approx \{\beta'_1, \beta_2, \beta'_3, \beta_4\} \cup \vec{\beta}, \end{aligned}$$

where \approx denotes isotopy. Figure 11 shows the relevant part of the Heegaard surface with the curves in the systems β , γ , and δ . (We do not show ϵ .)

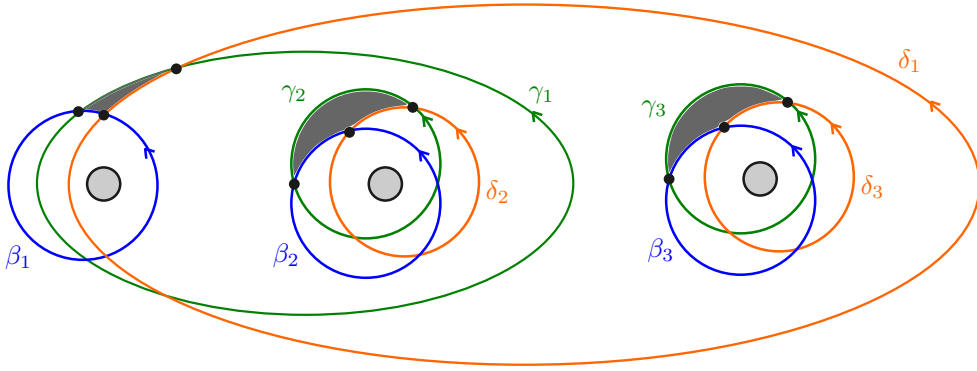


FIGURE 12. A triple diagram appearing in a handleslide loop of type (3).

There are top degree intersection points $\Theta_{\beta,\gamma}$, $\Theta_{\beta,\epsilon}$, $\Theta_{\gamma,\delta}$ and $\Theta_{\epsilon,\delta}$, which are used to define the respective handleslide maps. The commutativity of the handleslide loop can be written as

$$\Psi_{\gamma \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \gamma}^{\alpha} = \Psi_{\epsilon \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \epsilon}^{\alpha}.$$

Using the A_{∞} relations for polygon maps, this would follow if we can prove that

$$(57) \quad F_{\mathbb{T}_{\beta}, \mathbb{T}_{\gamma}, \mathbb{T}_{\delta}}(\Theta_{\beta,\gamma} \otimes \Theta_{\gamma,\delta}) = F_{\mathbb{T}_{\beta}, \mathbb{T}_{\epsilon}, \mathbb{T}_{\delta}}(\Theta_{\beta,\epsilon} \otimes \Theta_{\epsilon,\delta}).$$

Even though the attaching sets β and δ do not differ by a single handleslide, observe that we still have a unique top degree intersection point

$$\Theta_{\beta,\delta} \in \mathbb{T}_{\beta} \cap \mathbb{T}_{\delta}.$$

We claim that Equation (57) holds true because both sides are equal to $\Theta_{\beta,\delta}$. We will check this for the left hand side; the right hand side is similar.

To show that

$$(58) \quad F_{\mathbb{T}_{\beta}, \mathbb{T}_{\gamma}, \mathbb{T}_{\delta}}(\Theta_{\beta,\gamma} \otimes \Theta_{\gamma,\delta}) = \Theta_{\beta,\delta},$$

observe that there is a unique holomorphic triangle of index zero connecting the three intersection points. This is the product of g triangles on the Heegaard surface, four of which are shown darkly shaded in Figure 11. If we orient all the curves in the figure counterclockwise, we obtain the required coupled orientations in each handleslide. It follows from Examples 3.25 and 3.27 that the triangle comes with a positive sign; see Figure 3(c), with α' , β' , γ' replaced by β_i , γ_i , δ_i . Therefore, the relation (58) holds.

The handleslide loops of types (3), (4) and (5) are very similar to type (2). We isotope and then re-label each of the four attaching sets as in (56), in the order given in Definition 5.5 for the respective type of loop. In each case we claim that the relation (57) holds, and we prove it by establishing (58). The latter relation follows by exhibiting a unique holomorphic triangle, which appears with sign +1 according to the analysis in Example 3.25. Figures 12, 13 and 14 show the relevant triangles. Each darkly shaded triangle is oriented as in Figure 3, either (b) or (c).




FIGURE 13. A triple diagram appearing in a handleslide loop of type (4).





FIGURE 14. A triple diagram appearing in a handleslide loop of type (5). The two disks labeled A co-bound a handle.

The handleslide loop of type (6) is also somewhat similar. In this case we have a pentagon. We isotope and re-label the attaching sets from Definition 5.5 as



We want to show that

$$(59) \quad \Psi_{\gamma \rightarrow \delta}^{\alpha} \circ \Psi_{\epsilon \rightarrow \gamma}^{\alpha} \circ \Psi_{\beta \rightarrow \epsilon}^{\alpha} = \Psi_{\zeta \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \zeta}^{\alpha}.$$

Even though the attaching set β differs from γ by two handleslides (rather than one), there is still a unique top degree intersection point $\Theta_{\beta, \gamma} \in \mathbb{T}_{\beta} \cap \mathbb{T}_{\gamma}$, which defines a map $\Psi_{\beta \rightarrow \gamma}^{\alpha}$. Similarly, there is an intersection point $\Theta_{\beta, \delta} \in \mathbb{T}_{\beta} \cap \mathbb{T}_{\delta}$ defining a map $\Psi_{\beta \rightarrow \delta}^{\alpha}$. To prove (59), we will show that

$$\Psi_{\epsilon \rightarrow \gamma}^{\alpha} \circ \Psi_{\beta \rightarrow \epsilon}^{\alpha} = \Psi_{\beta \rightarrow \gamma}^{\alpha}, \quad \Psi_{\zeta \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \zeta}^{\alpha} = \Psi_{\beta \rightarrow \delta}^{\alpha}, \quad \Psi_{\gamma \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \gamma}^{\alpha} = \Psi_{\beta \rightarrow \delta}^{\alpha}.$$




FIGURE 15. A triple diagram appearing in a handleslide loop of type (6).




FIGURE 16. Another triple diagram appearing in a handleslide loop of type (6).

which in turn follow from the relations:

$$(60) \quad F_{\mathbb{T}_\beta, \mathbb{T}_\epsilon, \mathbb{T}_\gamma}(\Theta_{\beta, \epsilon} \otimes \Theta_{\epsilon, \gamma}) = \Theta_{\beta, \gamma},$$

$$(61) \quad F_{\mathbb{T}_\beta, \mathbb{T}_\zeta, \mathbb{T}_\delta}(\Theta_{\beta, \zeta} \otimes \Theta_{\zeta, \delta}) = \Theta_{\beta, \delta},$$

$$(62) \quad F_{\mathbb{T}_\beta, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\Theta_{\beta, \gamma} \otimes \Theta_{\gamma, \delta}) = \Theta_{\beta, \delta}.$$

Relation (60) is the same as the one proved in the study of a handleslide of type (3) (see Figure 12), because it comes from a curve handlesliding over two other curves. Relations (61) and (62) follow from investigating the unique holomorphic triangles in Figures 15 and 16, which both come with positive signs.

Finally, let us discuss handleslide loops of type (1), i.e., slide triangles. After an isotopy, we denote the three attaching sets appearing in the triangle by




Figure 17 shows the resulting triple Heegaard diagram. Note that, because of admissibility issues, we cannot arrange for all three attaching sets to be in the standard position for handleslides from Figure 4. We did arrange this for the (β, γ) and (β, δ) pairs, but not for the (γ, δ) pair. Nevertheless, there is a unique top degree intersection point in $\mathbb{T}_\gamma \cap \mathbb{T}_\delta$, which




FIGURE 17. The triple diagram from a handleslide loop of type (1).




FIGURE 18. An isotopy to get a standard handlebody diagram.

we denote by $\Theta_{\gamma,\delta}$; this induces a map

$$\Psi_{\gamma \rightarrow \delta}^{\alpha} : HF^{-}(\mathbb{T}_{\alpha}, \mathbb{T}_{\gamma}) \rightarrow HF^{-}(\mathbb{T}_{\alpha}, \mathbb{T}_{\delta}).$$

The holomorphic triangle darkly shaded in Figure 17 comes with positive sign, showing that

$$F_{\mathbb{T}_{\beta}, \mathbb{T}_{\gamma}, \mathbb{T}_{\delta}}(\Theta_{\beta,\gamma} \otimes \Theta_{\gamma,\delta}) = \Theta_{\beta,\delta}$$

and therefore

$$(63) \quad \Psi_{\gamma \rightarrow \delta}^{\alpha} \circ \Psi_{\beta \rightarrow \gamma}^{\alpha} = \Psi_{\beta \rightarrow \delta}^{\alpha}.$$

However, we are not done! The map $\Psi_{\gamma \rightarrow \delta}^{\alpha}$ is not a handleslide map like those in Section 4.8, where maps were defined from diagrams of the kind shown in Figure 4. We can obtain a diagram of that kind by isotoping the curves γ_2 and δ_1 as in Figure 18, so that they intersect each other and no longer intersect δ_2 and γ_1 . We let γ' and δ' denote the new attaching sets, where:

- For $i \neq 2$, the curve γ'_i is obtained from γ_i by a small Hamiltonian isotopy so that γ'_i intersects γ_i in two points, and does not intersect γ_j for any $j \neq i$;
- For $i \neq 1$, the curve δ'_i is obtained from δ_i by a small Hamiltonian isotopy so that δ'_i intersects δ_i in two points, and does not intersect δ_j for any $j \neq i$;
- The curves γ'_2 and δ'_1 are as in Figure 18.

Even so, the orderings and orientations of the curves γ' in Figure 18 differ from those chosen for a handleslide in Figure 5. To get to the standard picture, we should switch the ordering of γ'_1 and γ'_2 , and also the orientation on γ'_2 (i.e., on the new γ'_1). The two changes have the combined effect of preserving orientation on the torus $\mathbb{T}_{\gamma'}$, and hence preserving the coupled orientation. They do change the Pin structure, since we are using two different sets of data in Definition 4.4. The two Pin structures come from trivializations of $T\mathbb{T}_{\gamma'}$ that differ




FIGURE 19. Figure 18 with new orderings and orientations.

by replacing the ordered pair (γ'_1, γ'_2) with $(\gamma'_2, -\gamma'_1)$. Let us relate them by the homotopy given by 90° rotation in the (γ'_1, γ'_2) plane. This gives a homotopy between the Pin structures, which allows us to identify the respective Θ generators. Therefore, let us implement these changes for γ' , and also make the corresponding changes for the set γ . The result consists in replacing Figure 18 with Figure 19.

In the new picture (just as in the old), the attaching set γ' represents the same isotopy diagram as γ , so there is a continuation map $\Gamma_{\gamma \rightarrow \gamma'}^\alpha$. Similarly, δ' represents the same isotopy diagram as δ , and we have a continuation map $\Gamma_{\delta' \rightarrow \delta}^\alpha$. These continuation maps are used in the colimit that defines HF^- for an isotopy diagram. Moreover, the handleslide map between our two isotopy diagrams was defined as $\Psi_{\gamma' \rightarrow \delta'}^\alpha$. Thus, to deduce the commutativity of the slide triangle from (63), it suffices to prove the following lemma.

Lemma 5.11. *We have $\Psi_{\gamma \rightarrow \delta}^\alpha = \Gamma_{\delta' \rightarrow \delta}^\alpha \circ \Psi_{\gamma' \rightarrow \delta'}^\alpha \circ \Gamma_{\gamma \rightarrow \gamma'}^\alpha$.*

Proof. In the Heegaard diagram $(\Sigma, \gamma, \gamma', z)$, there is a unique top degree intersection point $\Theta_{\gamma, \gamma'}$, which produces a triangle map $\Psi_{\gamma \rightarrow \gamma'}^\alpha$. This map is the same as the continuation map $\Gamma_{\gamma \rightarrow \gamma'}^\alpha$, because at the chain level we can interpolate between the two maps by counting monogons as in [21, Proposition 11.4].

Similarly, there is a top degree intersection point $\Theta_{\delta', \delta}$ that defines a triangle map $\Psi_{\delta' \rightarrow \delta}^\alpha = \Gamma_{\delta' \rightarrow \delta}^\alpha$. We are left to show that

$$(64) \quad \Psi_{\gamma \rightarrow \delta}^\alpha = \Psi_{\delta' \rightarrow \delta}^\alpha \circ \Psi_{\gamma' \rightarrow \delta'}^\alpha \circ \Psi_{\gamma \rightarrow \gamma'}^\alpha.$$

We can further find a unique top degree intersection point $\Theta_{\gamma, \delta'}$, defining a map $\Psi_{\gamma \rightarrow \delta'}^\alpha$. We will deduce (64) from the relations

$$\Psi_{\gamma \rightarrow \delta'}^\alpha = \Psi_{\gamma' \rightarrow \delta'}^\alpha \circ \Psi_{\gamma \rightarrow \gamma'}^\alpha, \quad \Psi_{\gamma \rightarrow \delta}^\alpha = \Psi_{\delta' \rightarrow \delta}^\alpha \circ \Psi_{\gamma \rightarrow \delta'}^\alpha.$$

In turn, to prove these it suffices to show that

$$(65) \quad F_{\mathbb{T}_\gamma, \mathbb{T}_{\gamma'}, \mathbb{T}_{\delta'}}(\Theta_{\gamma, \gamma'} \otimes \Theta_{\gamma', \delta'}) = \Theta_{\gamma, \delta'}$$

and

$$(66) \quad F_{\mathbb{T}_\gamma, \mathbb{T}_{\delta'}, \mathbb{T}_\delta}(\Theta_{\gamma, \delta'} \otimes \Theta_{\delta', \delta}) = \Theta_{\gamma, \delta}.$$

Equation (65) follows from investigating the unique holomorphic triangle between the top degree generators in Figure 20; this comes with a positive sign.

Equation (66) follows from a more involved analysis, based on Figure 21. There, the homotopy class of triangles between theta elements that we need to consider is not made of two bigons, but rather is the darkly shaded hexagon. This is still combined with $g - 2$ other standard triangles in the rest of the triple diagram, not shown in the picture. The standard




FIGURE 20. A triple diagram leading to Equation (65).




FIGURE 21. A triple diagram with a darkly shaded hexagon leading to Equation (66).

triangles come with a positive sign, so to deduce (66), it remains to prove that the count of holomorphic triangles in the hexagon class (in the second symmetric product) is +1.

Let us re-draw the hexagon by itself in Figure 22 (a). To count its holomorphic representatives, we use Lipshitz’s cylindrical picture from [21] and stretch the almost complex structure along the dashed line in Figure 22. This kind of deformation was analyzed in [21, Appendix A] and used at various places in the literature; see for example [28, Theorem 5.1], [49, Section 14], or [47, Proposition 4.3]. In our case, the result of the neck stretching is shown in Figure 21. (See Section 5.4 below for more about signs in the cylindrical picture.)

Let T be the standard 2-simplex (triangle), where we remove neighborhoods of the vertices and attach three infinite half-strips instead. In the cylindrical picture, we are counting holomorphic maps $u : S \rightarrow \Sigma \times T$, where S is a branched cover of Δ (in our case, a hexagon). Let

$$\pi_T : \Sigma \times T \rightarrow T, \quad \pi_\Sigma : \Sigma \times T \rightarrow \Sigma$$

be the projections. We impose some boundary conditions on u : the edges of S should map under $\pi_T \circ u$ to the edges of the triangle, and under $\pi_\Sigma \circ u$ to the edges of the hexagon shown in Figure 22; further, vertices should map to the corresponding vertices. Let $\mathcal{M}(\psi)$ be the moduli space of holomorphic maps of this form, in the hexagon class ψ from Figure 22.




FIGURE 22. (a) The hexagon from Figure 21. (b) Its degeneration.

Let us extend the dashed line in Figure 22 to a separating circle on Σ . After the degeneration, the surface Σ turns into the wedge sum of two surfaces Σ_1 and Σ_2 , glued to each other at points $p \in \Sigma_1$ and $q \in \Sigma_2$. The class ψ gets split into ψ_1 and ψ_2 . For $i = 1, 2$, we have moduli spaces $\mathcal{M}(\psi_i)$, consisting of holomorphic maps $u_i : T \rightarrow \Sigma_i \times T$ in the class ψ_i (with the obvious boundary conditions). Consider the maps

$$\rho_1 : \mathcal{M}(\psi_1) \rightarrow \mathbb{R}, \quad \rho_1 = \pi_T((\pi_\Sigma \circ u)^{-1}(p)),$$

$$\rho_2 : \mathcal{M}(\psi_2) \rightarrow \mathbb{R}, \quad \rho_2 = \pi_T((\pi_\Sigma \circ u)^{-1}(q)).$$

(In principle, these maps should have image in T , but they land on the green edge of ∂T , which we identify with \mathbb{R} by orienting it using the boundary orientation from T .)

For large enough neck length, the space $\mathcal{M}(\psi)$ is identified with the fiber product

$$(67) \quad \mathcal{M}(\psi_1) \times_{\mathbb{R}} \mathcal{M}(\psi_2) = \{(u_1, u_2) \in \mathcal{M}(\psi_1) \times \mathcal{M}(\psi_2) \mid \rho_1(u_1) + \rho_2(u_2) = 0\},$$

and this identification preserves orientations. Since a triangle with three marked points on the boundary has a unique automorphism, we see that $\mathcal{M}(\psi_1)$ is a point, so its image under ρ_1 is also a fixed point in \mathbb{R} . On the other hand, $\mathcal{M}(\psi_2)$ is one-dimensional, because when we apply the Riemann mapping theorem to construct a holomorphic representative, we can vary the length of a slit along either the orange or the brown curve.

Thus, the fiber product in (67) is simply the preimage $\rho_2^{-1}(t)$ for some $t \in \mathbb{R}$. To count its points (with sign), we can vary t at will. In the limit $t \rightarrow +\infty$, the class ψ_2 splits into a triangle and a bigon, as in Figure 23. The respective moduli spaces both consist of a point with positive sign; see Examples 3.25 and 3.23. Therefore, the count of points in $\mathcal{M}(\psi) \cong \rho_2^{-1}(t)$ is $+1$. \square

Remark 5.12. If we reverse the orientations of all the curves in Figures 21, 22 and 23, the coupled Spin structures stay the same (cf. Convention 5.10), and the two triangles we consider still come with positive sign. However, the bigon on the right of Figure 23 looks like the bigon labeled B (rather than A) in Figure 2, so it comes with a negative sign. This sign is cancelled by the fact that the orientation of the green curve is reversed, which changes the orientation of the line \mathbb{R} over which we take the fiber product in (67). Thus, the overall sign of the hexagon is still $+1$.

An alternate way to think about this in Lipshitz's cylindrical picture is that we orient the moduli space of hexagons S that are double covers of the triangle T branched at a single point. This moduli space is just the triangle T itself, because the hexagon is determined




FIGURE 23. The class ψ_2 from Figure 22(b), in the limit when it splits along the δ_2 curve. (Note that, after breaking the orange δ_2 segment in two, the orientation on one half of the segment is reversed. This is as it should be: We want the same coupled orientation between the green and orange curves at the two sides of the breaking point.)

by the position of the branch point $x \in T$. The degeneration in Figure 22 corresponds to sending x to an edge, and the one in Figure 23 corresponds to further sending it to a vertex. In our discussion, we have conveniently chosen orientations of the moduli spaces compatible with these degenerations: the triangle T is oriented as a subset of the plane, the green curve is oriented as part of its boundary, and $t \rightarrow +\infty$ is the positive end of \mathbb{R} .

Now that Lemma 5.11 is proved, we have finished checking invariance under handleslide loops. We continue with the last step in the proof of Theorem 1.1.

(6) **Commutativity along stabilization slides:** Let us denote the Heegaard diagrams that appear in Figure 8 by

$$\begin{array}{ccccc}
 & & (\Sigma, \alpha, \beta) & & \\
 & \nearrow S & & \searrow \Psi_{\beta \rightarrow \gamma}^\alpha & \\
 (\Sigma^u, \alpha^u, \beta^u) & & & & (\Sigma, \alpha, \gamma) \\
 & \searrow S' & & \swarrow \Psi_{\gamma \rightarrow \delta}^\alpha & \\
 & & (\Sigma, \alpha, \delta) & &
 \end{array}$$

with Σ^u be the original (un-stabilized) surface, and Σ its stabilization. The two handleslide maps on the right hand side are induced by top degree generators $\Theta_{\beta, \gamma}$ and $\Theta_{\gamma, \delta}$.

Even though they differ by two handleslides rather than one, the attaching sets β and δ still produce a unique top degree generator $\Theta_{\beta, \delta}$, and a triangle map $\Psi_{\beta \rightarrow \delta}^\alpha$. The holomorphic triangle darkly shaded in Figure 24 comes with positive sign, which implies that

$$F_{\mathbb{T}_\beta, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\Theta_{\beta, \gamma} \otimes \Theta_{\gamma, \delta}) = \Theta_{\beta, \delta}$$

and therefore $\Psi_{\gamma \rightarrow \delta}^\alpha \circ \Psi_{\beta \rightarrow \gamma}^\alpha = \Psi_{\beta \rightarrow \delta}^\alpha$.

It remains to show that

$$\Psi_{\beta \rightarrow \delta}^\alpha \circ S = S'.$$

This is a consequence of Theorem 10.4 in [29], which says that triangle maps commute with stabilizations. In the un-stabilized diagram, the remaining β and δ curves are just small Hamiltonian perturbations of the β^u curves, so the respective triangle map is just the nearest point map (which we can think of as the identity, once we identify the corresponding curves by the small isotopy).

By applying Theorems 5.7 and 5.4, we deduce that HF^- is a natural invariant of based 3-manifolds. The proof for the other versions is similar. \square




FIGURE 24. A triple diagram appearing in the stabilization slide.

5.4. Signs in the cylindrical picture. In [21], Lipshitz reinterprets Heegaard Floer homology as follows. Instead of strips $u : \mathbb{R} \times [0, 1] \rightarrow \text{Sym}^g(\Sigma)$ with boundaries on \mathbb{T}_α and \mathbb{T}_β , he considers holomorphic curves

$$u : S \rightarrow \Sigma \times \mathbb{R} \times [0, 1]$$

where S is a smooth surface with boundary and infinite ends (of any genus), and the boundary components of ∂S are mapped to Lagrangians of the form $\alpha_i \times \{1\} \times \mathbb{R}$ and $\beta_i \times \{1\} \times \mathbb{R}$. (For the exact conditions on these curves, we refer to [21, Section 1].)

There is a tautological correspondence between the curves counted in Lipshitz's setup and the strips in the symmetric product that are counted in Ozsváth and Szabó's original construction. This correspondence is compatible with the determinant index bundles of $D\bar{\partial}$ operators, which allows Lipshitz to show that his theory is isomorphic to the original one over \mathbb{Z} . (See Section 13 in [21], particularly Proposition 13.7.) The cylindrical picture and the tautological correspondence can also be extended to polygon maps.

In [21], the signs for the differential come from coherent orientations, just as in the original [30]. In our context, the coupled Spin structure induces signs in the Ozsváth-Szabó picture, and hence (via the tautological correspondence) in Lipshitz's cylindrical picture.

In Heegaard Floer theory, it is common to move to the cylindrical picture when one studies various degenerations. See for example [28, Section 5], [16, Section 9.3], [46, Section 6], or what we did in the proofs of handleswap invariance and in Lemma 5.11 above. In such cases, we need to know that the orientations of the moduli spaces of curves are compatible with these degenerations. Transplanting orientations from the symmetric product picture is not fully satisfactory, because the degeneration arguments are less clear there.

Instead, it suffices to note that one can also construct canonical orientations in Lipshitz's picture, and that these agree with those from the symmetric product. In all the situations we encounter (such as in Lemma 5.11), we fix Spin structures as well as orderings on the alpha curves, as well as on the beta curves. By taking the product of these with the canonical Lie group Spin structure on \mathbb{R} , we get Spin structures on Lagrangians such as $\alpha_i \times \{1\} \times \mathbb{R}$ and $\beta_i \times \{1\} \times \mathbb{R}$, which give the boundary conditions in Lipshitz's picture. Thus, we can apply Seidel's work [40] and get orientations on the moduli spaces of curves there. In particular, we can construct orientation spaces $o(x)$ for every $x \in \alpha_i \cap \beta_j$, and set

$$o(\mathbf{x}) = o(x_1) \otimes \cdots \otimes o(x_g)$$

for each $\mathbf{x} = \{x_1, \dots, x_g\} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Since the Spin structures on \mathbb{T}_α and \mathbb{T}_β are products of those on the alpha and beta curves, respectively, we find that the orientation spaces $o(\mathbf{x})$ which generate the Floer complexes in the two pictures can be identified. Furthermore, following the identification of the index bundles in [21, Section 13], we see that the signs of differentials in the two pictures agree. A similar argument applies to the signs of higher polygons.

Whenever we consider a degeneration of polygons in Lipshitz's picture, the index bundles are compatible with gluing, and therefore so are the orientations induced by Spin structures. This justifies extending the degeneration arguments to \mathbb{Z} , provided the Spin structures before and after the degeneration are compatible. In fact, we always choose Lie group Spin structures on the curves, which are canonical once an orientation is fixed. Therefore, in each situation it suffices to describe orientations on the curves that are compatible with the degeneration. (An example is in Figure 22.)

6. THE SURGERY EXACT TRIANGLE

The surgery exact triangle is one of the most useful tools in Heegaard Floer theory. It was originally proved in [29, Theorem 9.12], using \mathbb{Z} coefficients from coherent orientations. A different proof, using the triangle detection lemma from homological algebra, was given in [27], with \mathbb{F}_2 coefficients. In this section we prove it over \mathbb{Z} using our set-up with canonical orientations.

A similar exact triangle exists in Lagrangian Floer homology, due to Seidel [39]; see also [40, Theorem 17.16]. This applies to the case where one Lagrangian is obtained from the other by a Dehn twist or, more generally, a fibered Dehn twist ([22], [43], [42]). Neither of these is quite the situation in Heegaard Floer homology but, nevertheless, one could imagine an extension of Seidel's exact triangle that applies to Heegaard Floer theory. With regard to orientations, an interesting point to note is that for Seidel's exact triangle, one needs to choose a bounding rather than a Lie group Pin structure on the circle; see [40, Example 17.15]. (In his terminology, nontrivial means bounding.) In our situation, it is also the case that the simplest exact triangle that can be constructed would involve bounding Pin structures. Nevertheless, by using twisted coefficients, we can arrive at a triangle with orientations coming from Lie group Pin structures.

We state the triangle here in the same generality as in monopole theory; compare [20, Theorem 2.4].

Theorem 6.1. *Let Y be a compact oriented 3-manifold with torus boundary, and let β, γ, δ be three oriented simple closed curves on ∂Y , pairwise intersecting in a single point, with intersection numbers*

$$\beta \cdot \gamma = \gamma \cdot \delta = \delta \cdot \beta = -1.$$

Let Y_β, Y_γ , and Y_δ be the closed three-manifolds obtained from Y by filling it with a solid torus $S^1 \times D^2$ such that $\{1\} \times \partial D^2$ is attached to the respective curve (β, γ or δ). Then, there is an exact sequence

$$\dots \rightarrow HF^+(Y_\beta) \rightarrow HF^+(Y_\gamma) \rightarrow HF^+(Y_\delta) \rightarrow \dots$$

Before proving the theorem, we recall the setup for the corresponding result in [29, Section 9], and explain why the naive strategy fails. We can choose an admissible pointed Heegaard multi-diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, z)$ such that:

- The Heegaard diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta})$ represent Y_β, Y_γ , and Y_δ respectively;




FIGURE 25. A pair of triangles on the torus. We use the point p to define polygon maps with twisted coefficients.

- For $i = 1, \dots, g - 1$, the curves β_i , γ_i and δ_i are small isotopic translates of each other, each pairwise intersecting in a pair of transverse intersection points (with the isotopies being supported away from z);
- We have $\beta_g = \beta$, $\gamma_g = \gamma$, and $\delta_g = \delta$.

Thus, we can re-write the desired exact triangle as

$$(68) \quad \cdots \rightarrow HF^+(\mathbb{T}_\alpha, \mathbb{T}_\beta) \xrightarrow{f_1} HF^+(\mathbb{T}_\alpha, \mathbb{T}_\gamma) \xrightarrow{f_2} HF^+(\mathbb{T}_\alpha, \mathbb{T}_\delta) \xrightarrow{f_3} \cdots$$

Note that, for the purposes of proving it, we do not need to worry about naturality issues. We choose Lie group Pin structures $P_\alpha^\#, P_\beta^\#, P_\gamma^\#, P_\delta^\#$ on $\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma, \mathbb{T}_\delta$, and also choose arbitrary trivializations of the lines $\ell(P_\alpha^\#, P_\beta^\#)$, $\ell(P_\alpha^\#, P_\gamma^\#)$, $\ell(P_\alpha^\#, P_\delta^\#)$. This allows us to identify the Heegaard Floer groups with the preliminary ones from Section 4.1, for which we have well-defined polygon maps. (Compare Remark 4.11.)

As a first guess, we could try to define the maps in (68) by counting triangles:

$$f_1(\cdot) = F_{\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma}(\cdot \otimes \Theta_{\beta, \gamma}), \quad f_2(\cdot) = F_{\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\cdot \otimes \Theta_{\gamma, \delta}), \quad f_3(\cdot) = F_{\mathbb{T}_\alpha, \mathbb{T}_\delta, \mathbb{T}_\beta}(\cdot \otimes \Theta_{\delta, \beta}).$$

Here, $\Theta_{\beta, \gamma}$ can be any of the two top-degree generators of the group $HF^+(\mathbb{T}_\beta, \mathbb{T}_\gamma) \cong H_*(T^{g-1}) \otimes \mathbb{Z}[U^{-1}]$; and similarly for $\Theta_{\gamma, \delta}$ and $\Theta_{\delta, \beta}$. It does not matter which generator we choose, because they differ by a sign, and changing one of the maps in an exact triangle by -1 keeps it exact.

However, defining the maps f_i in this way does not make the double compositions be zero. Indeed, for example, the proof that $f_2 \circ f_1 = 0$ in [29, Section 9] reduces to showing that

$$(69) \quad F_{\mathbb{T}_\beta, \mathbb{T}_\gamma, \mathbb{T}_\delta}(\Theta_{\beta, \gamma} \otimes \Theta_{\gamma, \delta}) = 0.$$

Proposition 9.5 in [29] shows that the relevant (index zero) holomorphic triangles between $\mathbb{T}_\beta, \mathbb{T}_\gamma$ and \mathbb{T}_δ come in pairs, with the triangles in each pair having basepoint multiplicity $n_z = k(k-1)/2$ for $k \geq 1$. This can be established by a degeneration argument to reduce it to the genus 1 case, where triangles can be counted explicitly on the diagram. The two triangles with $k = 1$ (so $n_z = 0$) are shown in Figure 25. The triangles for higher k wrap around the torus more times.

Observe that rotation by 180° around z takes the two triangles in a pair to each other. Furthermore, the Lie group Pin structure on each circle is uniquely characterized as being

preserved by translations, and this property is preserved by the 180° rotation. It follows that the two triangles in a pair are related by a symplectomorphism of the torus that preserves all the underlying structures, and therefore are counted with the same sign. This means that they do not cancel out, so we cannot deduce Equation (69). (They would cancel out if we had the bounding Pin structure on either one or three of the curves β , γ , δ , and the Lie group Pin structure on the rest.)

We shall remedy this problem while keeping Lie group Pin structures on all Lagrangians, by introducing twisted coefficients:

Proof of Theorem 6.1. Let $\iota : \partial Y \rightarrow Y$ be the inclusion. By the “half lives, half dies” lemma, the kernel of the map

$$\iota_* : H_1(\partial Y; \mathbb{Z}/2) \rightarrow H_1(Y; \mathbb{Z}/2)$$

is one-dimensional. Since β , γ and δ are simple closed curves on the torus ∂Y , they represent non-trivial elements in $H_1(\partial Y; \mathbb{Z}/2) \cong (\mathbb{Z}/2)^2$. By our hypotheses, their sum is zero (mod 2) in homology. Hence, the three curves represent the three distinct non-zero elements in $H_1(\partial Y; \mathbb{Z}/2)$. Exactly one of them must be in the kernel of ι_* ; without loss of generality, assume this is δ .

We pick a point p on δ (away from the other curves), and twist the Heegaard Floer groups and maps based on boundary intersections with p , similarly to [29, Section 9.3] but only using signs. Specifically, whenever we count some polygons in a class ϕ where one of the edges of the polygon is on \mathbb{T}_δ , we introduce a sign of

$$(-1)^{\#(V \cap \partial_\delta \phi)},$$

where $\partial_\delta \phi$ is the part of the boundary of ϕ on \mathbb{T}_δ , and $V = \delta_1 \times \cdots \times \delta_{g-1} \times \{p\} \subset \text{Sym}^g(\mathbb{T}_\delta)$.

At the level of the Heegaard Floer groups, this changes $HF^+(\mathbb{T}_\alpha, \mathbb{T}_\delta)$ into $\underline{HF}^+(\mathbb{T}_\alpha, \mathbb{T}_\delta; A)$ where A is a module over $\mathbb{Z}[H^1(Y; \mathbb{Z})]$ as in Section 4.7. Precisely, A is \mathbb{Z} as an abelian group, and the action of an element $c \in H^1(Y; \mathbb{Z})$ is given by $(-1)^{c([\delta])}$. Since $[\delta] = 0 \in H_1(Y; \mathbb{Z})$, the action is actually trivial, and we obtain the same group $HF^+(\mathbb{T}_\alpha, \mathbb{T}_\delta)$. The same logic applies to the groups $HF^+(\mathbb{T}_\gamma, \mathbb{T}_\delta)$ and $HF^+(\mathbb{T}_\delta, \mathbb{T}_\beta)$, so we can still keep the same elements $\Theta_{\gamma, \delta}$, $\Theta_{\delta, \beta}$ as before.

The groups $HF^+(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ and $HF^+(\mathbb{T}_\beta, \mathbb{T}_\gamma)$ are also unchanged, as they do not involve the Lagrangian \mathbb{T}_δ . The only difference is with regard to the maps $F_{\mathbb{T}_\alpha, \mathbb{T}_\gamma, \mathbb{T}_\delta}$ and $F_{\mathbb{T}_\alpha, \mathbb{T}_\delta, \mathbb{T}_\beta}$, which are now twisted, and produce new maps f_2 , f_3 in (68). As in Section 3.5, the twisted maps still satisfy the A_∞ polygon relations. Checking that double compositions such as $f_2 \circ f_1$ are zero in (68) reduces to counting triangles between \mathbb{T}_β , \mathbb{T}_γ and \mathbb{T}_δ , but now the two triangles in a pair with the same n_2 come with the opposite signs, because one of the triangles has p with odd multiplicity on its boundary.

It follows that with the new definitions, the sequence (68) is a chain complex. To check exactness, either of the two proofs in the literature applies with no significant change. The proof in [29, Section 9] uses energy filtrations, while the proof in [27] (stated there for \widehat{HF} , but applicable to HF^+) uses that a count of holomorphic quadrilaterals is ± 1 ; whether this count is $+1$ or -1 is not relevant for the argument. \square

Remark 6.2. In monopole Floer homology, the analogous surgery exact triangle was proved over \mathbb{F}_2 in [20]. An extension to coefficients in $\mathbb{Z}[i]$ is developed in [7], which is based on twisted coefficients. It is argued there that the same class of twistings (more limited than in Heegaard Floer theory) cannot produce an extension to \mathbb{Z} .

7. OTHER INVARIANTS

We now explain how to generalize the canonical orientations in Heegaard Floer homology to other related Floer homologies.

7.1. Sutured Floer homology. Juhász [15] developed a Floer homology theory for balanced sutured manifolds. Let us recall a few concepts.

Definition 7.1 (Definition 2.2 in [15]). A *balanced sutured manifold* (M, γ) is a compact oriented 3-manifold M together with a set $\gamma \subset \partial M$ of pairwise disjoint annuli, and a decomposition of its complement $\partial M \setminus \text{Int}(\gamma)$ into two pieces R_+ and R_- , with the following properties:

- Each annulus in γ has one boundary component in R_+ and one in R_- ;
- The middle curve $S^1 \times \{0\} \subset S^1 \times [-1, 1]$ in each annulus in γ is called a *suture*, and is oriented as follows. We equip R_+ with its orientation induced from M , and R_- with the opposite of the orientation induced from M . Then, we ask that any component of ∂R_+ or ∂R_- , with its induced boundary orientation, represents the same homology class in the respective annulus as the suture;
- The manifold M has no closed components, and the map $\pi_0(\gamma) \rightarrow \pi_0(\partial M)$ is onto;
- The Euler characteristics of R_+ and R_- coincide.

Definition 7.2 (Definition 2.13 in [16]). A *sutured Heegaard diagram* (Σ, α, β) for the balanced sutured manifold (M, γ) consists of an oriented surface $\Sigma \subset M$ whose (oriented) boundary is the union of sutures, and sets of attaching closed curves

$$\alpha = \{\alpha_1, \dots, \alpha_n\}, \quad \beta = \{\beta_1, \dots, \beta_n\},$$

such that:

- The components of α bound disks on the negative side of Σ , and compressing these disks yields a surface isotopic to R_- relative to γ ;
- The components of β bound curves on the positive side of Σ , and compressing these disks yields a surface isotopic to R_+ relative to γ .

Given a sutured Heegaard diagram, one considers the Lagrangians $\mathbb{T}_\alpha = \alpha_1 \times \dots \times \alpha_n$ and $\mathbb{T}_\beta = \beta_1 \times \dots \times \beta_n \subset \text{Sym}^n(\Sigma)$. Their Lagrangian Floer homology (taken over \mathbb{F}_2) is taken as the definition of the sutured Floer homology $SFH(M, \gamma)$, which Juhász proved to be an invariant of the balanced sutured manifold [15].

This theory can be upgraded to \mathbb{Z} coefficients along the same lines as we did with HF° . We only point out the relevant differences.

First, the surface Σ is no longer closed, and its genus g does not have to be equal to the number n of alpha (or beta) curves. Perutz's work [32] was originally phrased for closed surfaces and $g = n$. Nevertheless, we can fill in the boundary of Σ with disks to make it closed, and only consider holomorphic disks that avoid the new disks. The construction of the symplectic form on $\text{Sym}^n(\Sigma)$ in [32, Section 7] is for any n . The argument about handleslides corresponding to Hamiltonian isotopies can be adapted to any n .

Secondly, the surface Σ splits M not into handlebodies, but into two sutured compression bodies U_α and U_β . We can still define \mathbb{A} and \mathbb{B} as before, and we have the canonical isomorphisms (37) and (38). However, \mathbb{A} and \mathbb{B} are not Lagrangian subspaces of $H_1(\Sigma; \mathbb{R})$.

In [8], Friedl, Juhász and Rasmussen construct an absolute $\mathbb{Z}/2$ grading on sutured Floer homology, which is determined by a *homology orientation* ω for the pair (M, R_-) : that is,

by an orientation of the vector space

$$H_*(M, R_-; \mathbb{R}) = H_1(M, R_-; \mathbb{R}) \oplus H_2(M, R_-; \mathbb{R}).$$

Observe that the alpha and beta curves correspond to 1- and 2-cells, respectively, in a relative cell decomposition for (M, R_-) . Thus, there is an exact sequence

$$(70) \quad 0 \rightarrow H_2(M, R_-; \mathbb{R}) \rightarrow \mathbb{B} \rightarrow \mathbb{A} \rightarrow H_1(M, R_-; \mathbb{R}) \rightarrow 0.$$

It follows that fixing a homology orientation gives a coupled orientation on (\mathbb{A}, \mathbb{B}) , and this is how the absolute $\mathbb{Z}/2$ -grading is defined in [8]. For our purposes, we need more: a coupled Spin structure on the same pair (\mathbb{A}, \mathbb{B}) . To find one, apply Lemma 2.33 to the same exact sequence (70). We deduce that it suffices to specify a coupled Spin structure on $(H_1(M, R_-; \mathbb{R}), H_2(M, R_-; \mathbb{R}))$. By Lefschetz duality for triples, the space $H_2(M, R_-; \mathbb{R})$ is dual to $H_1(M, R_+; \mathbb{R})$, so we can identify them using an inner product. This suggests the following definition.

Definition 7.3. A *homological coupled Spin structure* on the sutured manifold (M, γ) is a coupled Spin structure on the pair of vector spaces $(H_1(M, R_-; \mathbb{R}), H_1(M, R_+; \mathbb{R}))$.

In general, there is a $\mathbb{Z}/2 \times \mathbb{R}\mathbb{P}^\infty$ worth of homological coupled Spin structures, where the $\mathbb{Z}/2$ comes from the two choices of homology orientation.

We have just seen that a homological coupled Spin structure S determines a coupled Spin structure on the pair (\mathbb{A}, \mathbb{B}) . Given a Spin^c structure on (M, γ) as in [15, Section 4], this allows us to define sutured Floer homology $SFH(M, \gamma, \mathfrak{s}, S)$ over \mathbb{Z} , in a manner similar to how we defined Heegaard Floer homology in Section 4: We choose Pin structures $P_\alpha^\#$ and $P_\beta^\#$ on the Lagrangians, define preliminary Floer complexes, and then tensor them with lines $\ell(P_\alpha^\#, P_\beta^\#)$.

Proof of Theorem 1.2. The arguments are entirely similar to those for Heegaard Floer homology in Sections 4.10 and 5. With regard to Theorem 5.7, the fact that the same loops of handleslides are sufficient in the sutured case follows from the work of Qin [35]. \square

7.2. Heegaard Floer homology with multiple basepoints. Let us go back to a closed, connected, oriented three-manifold Y . In [28, Section 4], Ozsváth and Szabó generalized the definition of Heegaard Floer homology to allow for multiple basepoints instead of just one. We follow the notation from that paper, letting $\mathbf{w} = \{w_1, \dots, w_\ell\}$ be the collection of basepoints on Y .

There is a notion of an ℓ -pointed balanced Heegaard diagram for (Y, \mathbf{w}) . This consists of a surface $\Sigma \subset Y$ of genus g containing \mathbf{w} , together with $g + \ell - 1$ alpha curves and $g + \ell - 1$ beta curves. We ask that:

- the surface Σ splits Y into two handlebodies U_α and U_β ;
- the alpha curves bound disks in U_α , and $\Sigma \setminus \boldsymbol{\alpha}$ consists of ℓ planar surfaces, each containing a basepoint;
- the beta curves bound disks in U_β , and $\Sigma \setminus \boldsymbol{\beta}$ consists of ℓ planar surfaces, each containing a basepoint.

In a certain sense, this set-up is a particular case of what we had in Section 7.1. Indeed, to (Y, \mathbf{w}) we can associate a sutured manifold (M, γ) as follows. We pick small disjoint balls $B_1, \dots, B_\ell \subset Y$, with B_i centered at the basepoint w_i . We let γ_i be an annulus around the equator for ∂B_i , and let $R_{+,i}$ and $R_{-,i}$ be the two components of $B_i \setminus \gamma_i$. We let

$$M = Y \setminus (B_1 \cup \dots \cup B_\ell),$$

with γ the union of γ_i 's, and R_+ resp. R_- the union of all $R_{+,i}$ resp. $R_{-,i}$. (See Example 2.3 in [15].) Furthermore, from an ℓ -pointed balanced Heegaard diagram for Y we get a sutured diagram for (M, γ) by deleting disks around each w_i .

Let \mathfrak{s} be a Spin^c structure on Y . In [28, Section 4], Ozsváth and Szabó define a version of Heegaard Floer homology $HF^-(Y, \mathbf{w}, \mathfrak{s})$ with coefficients in the ring $\mathbb{F}_2[U_1, \dots, U_\ell]$. By setting $U_1 = \dots = U_\ell = 0$ at the chain level, and then taking homology, they also get a hat version $\widehat{HF}(Y, \mathbf{w}, \mathfrak{s})$ over \mathbb{F}_2 . Let us upgrade their constructions to \mathbb{Z} instead of \mathbb{F}_2 .

In the case of $\widehat{HF}(Y, \mathbf{w}, \mathfrak{s})$, this is just the sutured Floer homology $SFH(M, \gamma, \mathfrak{s}, S)$ which we constructed in Section 7.1, with a dependence on the homological coupled Spin structure S . In the case at hand, observe that both $H_1(M, R_-; \mathbb{R})$ and $H_1(M, R_+; \mathbb{R})$ are canonically identified with $H_1(Y, \mathbf{w}; \mathbb{R})$, by collapsing a hemisphere onto the center w_i of the ball B_i . By Lemma 2.31, there is a canonical coupled Spin structure on the pair

$$(H_1(Y, \mathbf{w}; \mathbb{R}), H_1(Y, \mathbf{w}; \mathbb{R})).$$

This gives a canonical homological coupled Spin structure, and we use it to define $\widehat{HF}(Y, \mathbf{w}, \mathfrak{s})$ without any additional dependence.

The homological coupled Spin structure gives a coupled Spin structure on the pair (\mathbb{A}, \mathbb{B}) , which we can also use to define the minus Floer homology HF^- , much as we did in the singly based case in Section 4. When counting holomorphic strips in a class ϕ , we keep track of their intersection $n_{w_i}(\phi)$ with $\{w_i\} \times \text{Sym}^{g+\ell-2}(\Sigma)$ by a factor of $U_i^{n_{w_i}(\phi)}$. The one main difference is that we now have boundary degenerations (disk bubbles), of the kind discussed in Section 3.3. In [28, Theorem 5.5], Ozsváth and Szabó prove that for $\ell > 1$ the count of such bubbles with boundary on \mathbb{T}_α is 1 (mod 2); and the same is true for those with boundary on \mathbb{T}_β . Here is the refinement of that result with \mathbb{Z} coefficients.

Proposition 7.4. *Let Σ be a surface of genus g equipped with a set α of $g + \ell - 1$ attaching circles for a handlebody, where $\ell > 1$. Let $M = \text{Sym}^{g+\ell-1}(\Sigma)$ and $\mathbb{T}_\alpha \subset M$ be the product of the alpha curves, equipped with a Lie group Pin structure. Suppose $\mathbf{x} \in \mathbb{T}_\alpha$ and $\phi \in \pi_2(M, \mathbb{T}_\alpha)$ is a relative homotopy class of disks as in (21), with index $\mu(\phi) = 2$ and such that the domain of ϕ on Σ has only non-negative coefficients. Then the signed count of holomorphic disks in the class ϕ is $\#\widehat{\mathcal{N}}(\phi) = 1$.*

Proof. As noted in the proof of [28, Theorem 5.5], the conditions in the theorem imply that ϕ must correspond to a domain D on Σ which is one of the connected components of $\Sigma \setminus \alpha$. This is a planar surface with boundary, and by de-stabilization and gluing arguments one can reduce the problem to the case where D is a disk with one marked point. In our setting, we can choose orientations and orderings on the alpha curves (as in Definition 4.4) to view \mathbb{T}_α as the product of $g + \ell - 1$ circles, each of which equipped with the Lie group Spin structure. The arguments in the proof over \mathbb{F}_2 go through over \mathbb{Z} , and we are left with the moduli space for a disk D whose boundary has the Lie group Spin structure. (In fact, once we reduced to a single circle we only care about Pin.) The moduli space $\mathcal{N}(\phi)$ is just $\text{Aut}(D^2, 1)$, and the quotient $\widehat{\mathcal{N}}(\phi)$ is a point. The fact that this point is counted with a positive sign can be read off the proof of Lemma 11.17 in [40], using the fact that the Pin structure on ∂D is the Lie group one (or “trivial”, in the terminology of [40]). \square

Of course, Proposition 7.4 applies equally well to the beta curves. When considering ∂^2 for the minus Floer complex, we get contributions from both the alpha and the beta degenerations. By Proposition 3.12, the beta degenerations actually come with negative sign. Overall, we get ℓ contributions of +1 from the components of $\Sigma \setminus \alpha$, and ℓ contributions




FIGURE 26. An index zero/three stabilization.

of -1 from the components of $\Sigma \setminus \beta$. Each component contains a single basepoint w_i , so it gets multiplied by a factor U_i . The two contributions from the domains containing w_i cancel out, implying that $\partial^2 = 0$. We obtain a well-defined homology $HF^-(Y, \mathbf{w}, \mathfrak{s})$, with coefficients in $\mathbb{Z}[U_1, \dots, U_\ell]$.

Proposition 7.5. *Let Y be a closed, connected, oriented 3-manifold equipped with a set of basepoints $\mathbf{w} \in Y$ and a $Spin^c$ structure \mathfrak{s} . Then, the Heegaard Floer homologies \widehat{HF} and HF^- are natural invariants of the triple $(Y, \mathbf{w}, \mathfrak{s})$.*

Proof. Similar to that of Theorems 1.1 and 1.2. \square

In [28, Theorem 4.4], Ozsváth and Szabó also investigate how the Heegaard Floer homology depends on the basepoints, up to (non-canonical) isomorphism. Here is the analogue of their result for \mathbb{Z} coefficients.

Proposition 7.6. *Let $\mathbf{w} = \{w_1, \dots, w_\ell\} \subset Y$. Then, the actions of multiplication by U_i on $HF^-(Y, \mathbf{w}, \mathfrak{s})$ for $i = 1, \dots, \ell$ are the same, giving it the structure of a $\mathbb{Z}[U]$ -module (where U is any U_i). Moreover, we have an isomorphism of $\mathbb{Z}[U]$ -modules*

$$HF^-(Y, \mathbf{w}, \mathfrak{s}) \cong HF^-(Y, \mathfrak{s}).$$

We also have an isomorphism of abelian groups

$$\widehat{HF}(Y, \mathbf{w}, \mathfrak{s}) \cong \widehat{HF}(Y, \mathfrak{s}) \otimes H_*(T^{\ell-1}).$$

Proof. The proof of [28, Theorem 4.4] proceeds by using Heegaard moves to reduce to the case of a simple index zero/three stabilization. There, we have a Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w})$ and its connected sum

$$(\Sigma', \alpha', \beta', \mathbf{w}') = (\Sigma, \alpha, \beta, \mathbf{w}) \# (S, \alpha, \beta, v_1, v_2),$$

where S is a sphere with two great circles α and β intersecting at two points x and y , as in Figure 26. The connected sum is taken at a basepoint $w_1 \in \mathbf{w}$ on Σ , which gets identified with v_1 on S . The point v_2 becomes a new basepoint $w_{\ell+1}$, so that $\mathbf{w}' = \mathbf{w} \cup \{w_{\ell+1}\}$. Degeneration arguments show that we have an isomorphism of chain complexes

$$CF^-(\Sigma', \alpha', \beta', \mathbf{w}') \cong CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \otimes CF^-(S, \alpha, \beta, v_1, v_2)$$

where the tensor product is over $\mathbb{F}[U_1]$ and the second factor is the mapping cone of $U_1 - U_{\ell+1}$. In our setting, after we equip each curve with the Lie group $Spin$ structure, the same arguments work over \mathbb{Z} . Indeed, $CF^-(S, \alpha, \beta, v_1, v_2)$ is generated by x and y , and in its differential the two empty bigons from y to x come with opposite signs (as we have seen in Example 3.23). The two bigons from x to y also come with opposite signs, and contribute $U_1 - U_{\ell+1}$. Thus, $CF^-(\Sigma', \alpha', \beta', \mathbf{w}')$ is isomorphic to the mapping cone

$$CF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}] \xrightarrow{U_1 - U_{\ell+1}} CF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}].$$

From here we deduce that the $\mathbb{F}[U_1]$ -modules HF^- for the two diagrams are isomorphic, whereas the hat versions \widehat{HF} differ by tensoring with a $H_*(S^1)$ factor. This suffices to establish the desired results. (To see that the U_i actions on HF^- are the same, we relate them to a singly based diagram by Heegaard moves, using any of the U_i as the variable in the coefficient ring.) \square

7.3. Link Floer homology. Link Floer homology is an invariant of links in three-manifolds developed by Ozsváth and Szabó in [28], with coefficients in \mathbb{F}_2 . It generalizes the construction for knots in [26], [36]. There are several versions of link Floer homology. We start by reviewing \widehat{HFL} and HFL^- , borrowing some terminology and notation from [28], [24], and [48].

Definition 7.7. Let Y be a closed oriented 3-manifold. A *multi-based link* in Y is a triple $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ where $L \subset Y$ is an embedded, oriented link, and \mathbf{w} and \mathbf{z} are disjoint collections of basepoints on L with the following properties:

- every component of L has at least two basepoints, and
- as we traverse any component, the basepoints alternate between those in \mathbf{w} and those in \mathbf{z} .

Definition 7.8. Fix (Y, \mathbb{L}) as above. A *multi-based Heegaard diagram* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ representing (Y, \mathbb{L}) is a Heegaard diagram for Y with $\mathbf{w}, \mathbf{z} \subset \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$, such that after we attach disks to the alpha and beta curves in their respective handlebody, the handlebodies are split into balls; each ball should have exactly one w and one z basepoint on its boundary, and as we join these two basepoints by an arc inside the ball, the union of these arcs should be the link L . The orientation of L should be such that the arcs go from the w to the z basepoints inside the alpha handlebody.

Let (Y, \mathbb{L}) be a 3-manifold with a multi-based link, and \mathfrak{s} be a Spin^c structure on $Y \setminus \text{nbhd}(L)$ relative its boundary, as in [28, Section 3.2]. Suppose we have m basepoints of type w and m of type z . Given a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ representing (Y, \mathbb{L}) , we define $HFL^-(Y, \mathbb{L}, \mathfrak{s})$ as the homology of a Floer complex over $\mathbb{F}_2[U_1, \dots, U_m]$, where the Lagrangians are \mathbb{T}_α and \mathbb{T}_β as usual, and we only count holomorphic strips in classes ϕ such that

$$n_{z_1}(\phi) = n_{z_2}(\phi) = \dots = n_{z_m}(\phi) = 0.$$

Further, when counting the strips we include a factor of

$$U_1^{n_{w_1}(\phi)} U_2^{n_{w_2}(\phi)} \dots U_m^{n_{w_m}(\phi)},$$

just as we did in Section 7.2.

For $\widehat{HFL}(Y, \mathbb{L}, \mathfrak{s})$, we set all U_i to be zero in the complex above and then take homology; that is, we only count strips that do not go over any basepoints.

Let us upgrade these constructions to \mathbb{Z} instead of \mathbb{F}_2 . To imitate what we did for HF° in Section 4, we need a coupled Spin structure on (\mathbb{A}, \mathbb{B}) , where \mathbb{A} and \mathbb{B} are the spans of the alpha and beta curves, respectively. Observe that to (Y, \mathbb{L}) we can associate a sutured manifold (M, γ) as follows. We let M be the complement $Y \setminus \text{nbhd}(L)$. Take meridians $\mu(w_i), \mu(z_i) \subset \partial M$ around the basepoints. Let the sutures γ be disjoint annuli around each of these meridians in ∂M , and let R_- and R_+ be the remaining parts of ∂M that are in U_α and U_β , respectively. (Compare Example 2.4 in [15].)

We can now appeal to the results from Section 7.1. In fact, notice that $\widehat{HFL}(Y, \mathbb{L}, \mathfrak{s})$ is nothing else than the sutured Floer homology of (M, γ) . A coupled Spin structure on (\mathbb{A}, \mathbb{B}) is determined by a homological coupled Spin structure on (M, γ) . In our setting, where

the sutured manifold comes from a multi-based link, the pairs (M, R_-) and (M, R_+) are homotopy equivalent, by an isotopy supported in a neighborhood of ∂M which slides each annulus from R_- into the next one from R_+ by following the orientation of the link. Hence, the spaces $H_1(M, R_-; \mathbb{R})$ and $H_1(M, R_+; \mathbb{R})$ are canonically identified, and by Lemma 2.31 we have a preferred homological coupled Spin structure.

Thus, we can define $\widehat{HFL}(Y, \mathbb{L}, \underline{\mathfrak{s}})$ without making any additional choices. The same goes for $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$, because the spaces \mathbb{A} and \mathbb{B} are the same. Note that for these versions of link Floer homology we do not have any contributions from disk bubbles, because their domains would go through the z_i basepoints and hence they are not counted. The resulting theories are natural invariants of the multi-based link.

Proof of Theorem 1.3. Similar to those of Theorem 1.1, Theorem 1.2 and Proposition 7.5. \square

We also have an analogue of Proposition 7.6, with a similar proof.

Proposition 7.9. *Let $(Y, \mathbb{L}, \underline{\mathfrak{s}})$ be a 3-manifold with a multi-based link and a relative Spin^c structure. Then, the actions of multiplication by U_i on $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ are the same for all variables U_i corresponding to basepoints on the same component of the link. This gives $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ the structure of a module over a polynomial ring with one variable for each link component. As such, $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ is independent of the number and position of the base points, up to (non-canonical) isomorphism.*

In the case of $HFL^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$, if $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ and $\mathbb{L}' = (L, \mathbf{w}', \mathbf{z}')$ are such that $|\mathbf{w}| = |\mathbf{z}| = m$ and $|\mathbf{w}'| = |\mathbf{z}'| = m' \leq m$, then we have a (non-canonical) isomorphism

$$\widehat{HFL}(Y, \mathbb{L}, \underline{\mathfrak{s}}) \cong \widehat{HFL}(Y, \mathbb{L}', \underline{\mathfrak{s}}) \otimes H_*(T^{m-m'}).$$

One can also consider more general versions of link Floer complexes, where we allow strips that go over the z basepoints. For example, in [48], Zemke works with a curved chain complex $\mathcal{CFL}^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ for $\underline{\mathfrak{s}} \in \text{Spin}^c(Y)$. This is a module over $\mathbb{F}_2[U_1, \dots, U_m, V_1, \dots, V_m]$, and its differential counts holomorphic strips in a class ϕ with a coefficient of

$$U_1^{n_{w_1}(\phi)} \dots U_m^{n_{w_m}(\phi)} \cdot V_1^{n_{z_1}(\phi)} \dots V_m^{n_{z_m}(\phi)}.$$

To do this over $\mathbb{Z}[U_1, \dots, U_m, V_1, \dots, V_m]$, we use our preferred homological coupled Spin structure. We now have disk bubbles as in Section 7.2, and they no longer cancel. Rather, there are positive contributions to ∂^2 from disks on \mathbb{T}_α , and negative ones from disks on \mathbb{T}_β . We get

$$\partial^2 = \left(\sum_{i=1}^m (U_i V_{\sigma(i)} - V_i U_{\tau(i)}) \right) \cdot \text{id},$$

where $z_{\sigma(i)}$ is the basepoint following w_i as we go around the respective component of L (with the given orientation), and $w_{\tau(i)}$ is the basepoint following $\tau(i)$. The complex $\mathcal{CFL}^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ is curved, so we cannot take its homology (unless we have exactly two basepoints per component, in which case $\partial^2 = 0$). Nevertheless, a variant of Theorem 1.3 still applies, saying that $\mathcal{CFL}^-(Y, \mathbb{L}, \underline{\mathfrak{s}})$ form a transitive system of curved chain complexes in the sense of [48, Definition 2.15].

Remark 7.10. In [38], Sarkar constructs $2^{\ell-1}$ versions of link Floer homology over \mathbb{Z} , where the different choices correspond to whether, for each of the ℓ components of the link, it is the alpha or the beta degenerations with basepoints on that component which get counted with +1 (versus -1). Our canonical version corresponds to his where all the alpha degenerations are counted with +1. This agrees with the signs from grid homology [25, Definition 4.1]. To

get the other choices in [38], one can use twisted coefficients, or equip one or both of the Lagrangians with a Pin structure different from the Lie group one. (Compare Section 4.7.)

7.4. Involutive Heegaard Floer homology. In [14], Hendricks and the second author define an invariant of 3-manifolds called involutive Heegaard Floer homology. Let us review that construction. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a based Heegaard diagram representing a 3-manifold Y . For simplicity, we focus on Spin^c structures \mathfrak{s} that are self-conjugate, i.e., $\mathfrak{s} = \bar{\mathfrak{s}}$. (This is the only interesting case.) Given such an \mathfrak{s} , there is a canonical conjugation isomorphism between Heegaard Floer chain complexes:

$$\eta : CF^\circ(\mathcal{H}, \mathfrak{s}) \xrightarrow{\cong} CF^\circ(\bar{\mathcal{H}}, \mathfrak{s})$$

where $\bar{\mathcal{H}}$ is the Heegaard diagram $(-\Sigma, \boldsymbol{\beta}, \boldsymbol{\alpha}, z)$ and $\circ \in \{\wedge, +, -, \infty\}$. Furthermore, since \mathcal{H} and $\bar{\mathcal{H}}$ represent the same 3-manifold, there is a chain homotopy equivalence induced by Heegaard moves

$$\Phi(\bar{\mathcal{H}}, \mathcal{H}) : CF^\circ(\bar{\mathcal{H}}, \mathfrak{s}) \xrightarrow{\sim} CF^\circ(\mathcal{H}, \mathfrak{s}).$$

We let

$$\iota = \Phi(\bar{\mathcal{H}}, \mathcal{H}) \circ \eta : CF^\circ(\mathcal{H}, \mathfrak{s}) \rightarrow CF^\circ(\mathcal{H}, \mathfrak{s}).$$

The involutive Heegaard Floer homology $HFI^\circ(Y, \mathfrak{s})$ is defined as the homology of the mapping cone

$$CF^\circ(\mathcal{H}, \mathfrak{s}) \xrightarrow{Q(1+\iota)} Q \cdot CF^\circ(\mathcal{H}, \mathfrak{s})[-1],$$

where Q is a formal variable and $[-1]$ is a shift in degree.

We can repeat the same construction over \mathbb{Z} , using the Floer complexes and maps from Section 4.

Proposition 7.11. *The isomorphism class of the involutive Heegaard Floer homology $HFI^\circ(Y, \mathfrak{s})$, as a graded $\mathbb{Z}[Q, U]/(Q^2)$ -module, is an invariant of the pair (Y, \mathfrak{s}) .*

Proof. This is similar to the proof of the analogous Theorem 1.1 in [14]. The key input is that the map $\Phi(\bar{\mathcal{H}}, \mathcal{H})$ is invariant up to chain homotopy equivalence. This is a version of naturality at the chain level, which in our context follows from the proofs in Section 5. \square

There is also a naturality result for involutive Heegaard Floer homology (with \mathbb{F}_2 coefficients), proved in [13]: It says that HFI° is a natural invariant of $(Y, \mathfrak{s}, z, \xi)$, where ξ is a framing of $T_z Y$. While we expect the same statement with \mathbb{Z} coefficients, proving it is beyond the scope of the current paper.

REFERENCES

- [1] M. Abouzaid, *A topological model for the Fukaya categories of plumbings*, J. Differential Geom., **87**(2011), no. 1, 1–80.
- [2] L. Amorim, *The Künneth theorem for the Fukaya algebra of a product of Lagrangians*, Internat. J. Math., **28**(2017), no. 4, 1750026, 38.
- [3] P. Biran and O. Cornea, *Quantum structures for Lagrangian submanifolds*, preprint, arXiv:0708.4221, 2007.
- [4] C. J. Earle and J. Eells, *A fibre bundle description of Teichmüller theory*, J. Differential Geometry, **3**(1969), 19–43.
- [5] C. J. Earle and A. Schatz, *Teichmüller theory for surfaces with boundary*, J. Differential Geometry, **4**(1970), 169–185.
- [6] A. Floer and H. Hofer, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z., **212**(1993), no. 1, 13–38.
- [7] J. Freeman, *The Surgery Exact Triangle in Monopole Floer Homology with $\mathbb{Z}[i]$ Coefficients*, ProQuest LLC, Ann Arbor, MI, thesis (Ph.D.)–Massachusetts Institute of Technology, 2021.

- [8] S. Friedl, A. Juhász, and J. Rasmussen, *The decategorification of sutured Floer homology*, J. Topol., **4**(2011), no. 2, 431–478.
- [9] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*, volume 46 of *AMS/IP Studies in Advanced Mathematics*, American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.
- [10] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, volume 46 of *AMS/IP Studies in Advanced Mathematics*, American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.
- [11] M. Furuta, *Monopole equation and the $\frac{11}{8}$ -conjecture*, Math. Res. Lett., **8**(2001), no. 3, 279–291.
- [12] M. Gartner, *Projective naturality in Heegaard Floer homology*, Algebr. Geom. Topol., **23**(2023), no. 3, 963–1054.
- [13] K. Hendricks, J. Hom, M. Stoffregen, and I. Zemke, *Naturality and functoriality in involutive Heegaard Floer homology*, preprint, arXiv:2201.12906, 2022.
- [14] K. Hendricks and C. Manolescu, *Involutive Heegaard Floer homology*, Duke Math. J., **166**(2017), no. 7, 1211–1299.
- [15] A. Juhász, *Holomorphic discs and sutured manifolds*, Algebr. Geom. Topol., **6**(2006), 1429–1457.
- [16] A. Juhász, D. Thurston, and I. Zemke, *Naturality and mapping class groups in Heegaard Floer homology*, Mem. Amer. Math. Soc., **273**(2021), no. 1338.
- [17] R. C. Kirby, *The topology of 4-manifolds*, volume 1374 of *Lecture Notes in Mathematics*, Springer-Verlag, Berlin, 1989.
- [18] R. C. Kirby and L. R. Taylor, *Pin structures on low-dimensional manifolds*, in *Geometry of low-dimensional manifolds, 2 (Durham, 1989)*, Cambridge Univ. Press, Cambridge, volume 151 of *London Math. Soc. Lecture Note Ser.*, pp. 177–242, 1990.
- [19] P. Kronheimer and T. Mrowka, *Monopoles and three-manifolds*, volume 10 of *New Mathematical Monographs*, Cambridge University Press, Cambridge, 2007.
- [20] P. Kronheimer, T. Mrowka, P. Ozsváth, and Z. Szabó, *Monopoles and lens space surgeries*, Ann. of Math. (2), **165**(2007), no. 2, 457–546.
- [21] R. Lipshitz, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol., **10**(2006), 955–1096.
- [22] C. Y. Mak and W. Wu, *Dehn twist exact sequences through Lagrangian cobordism*, Compos. Math., **154**(2018), no. 12, 2485–2533.
- [23] C. Manolescu, *Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture*, J. Amer. Math. Soc., **29**(2016), no. 1, 147–176.
- [24] C. Manolescu, P. Ozsváth, and S. Sarkar, *A combinatorial description of knot Floer homology*, Ann. of Math. (2), **169**(2009), no. 2, 633–660.
- [25] C. Manolescu, P. Ozsváth, Z. Szabó, and D. Thurston, *On combinatorial link Floer homology*, Geom. Topol., **11**(2007), 2339–2412.
- [26] P. Ozsváth and Z. Szabó, *Holomorphic disks and knot invariants*, Adv. Math., **186**(2004), no. 1, 58–116.
- [27] P. Ozsváth and Z. Szabó, *On the Heegaard Floer homology of branched double-covers*, Adv. Math., **194**(2005), no. 1, 1–33.
- [28] P. Ozsváth and Z. Szabó, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, Algebr. Geom. Topol., **8**(2008), no. 2, 615–692.
- [29] P. S. Ozsváth and Z. Szabó, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. (2), **159**(2004), no. 3, 1159–1245.
- [30] P. S. Ozsváth and Z. Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2), **159**(2004), no. 3, 1027–1158.
- [31] P. S. Ozsváth and Z. Szabó, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math., **202**(2006), no. 2, 326–400.
- [32] T. Perutz, *Hamiltonian handleslides for Heegaard Floer homology*, in *Proceedings of Gökova Geometry-Topology Conference 2007*, Gökova Geometry/Topology Conference (GGT), Gökova, pp. 15–35, 2008.
- [33] I. Petkova, *An absolute $\mathbb{Z}/2$ grading on bordered Heegaard Floer homology*, New York J. Math., **28**(2022), 90–116.
- [34] S. Piunikhin, D. Salamon, and M. Schwarz, *Symplectic Floer-Donaldson theory and quantum cohomology*, in *Contact and symplectic geometry (Cambridge, 1994)*, Cambridge Univ. Press, Cambridge, volume 8 of *Publ. Newton Inst.*, pp. 171–200, 1996.
- [35] Q. Qin, *Loops of handleslides for sutured diagrams*, preprint, arXiv:2508.04042, 2025.

- [36] J. A. Rasmussen, *Floer homology and knot complements*, ProQuest LLC, Ann Arbor, MI, thesis (Ph.D.)–Harvard University, 2003.
- [37] S. Rezchikov, *Floer homology via twisted loop spaces*, preprint, arXiv:1909.01325, 2019.
- [38] S. Sarkar, *A note on sign conventions in link Floer homology*, *Quantum Topol.*, **2**(2011), no. 3, 217–239.
- [39] P. Seidel, *A long exact sequence for symplectic Floer cohomology*, *Topology*, **42**(2003), no. 5, 1003–1063.
- [40] P. Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008.
- [41] V. D. Silva, *Products in the symplectic Floer homology of Lagrangian intersections*, PhD thesis, University of Oxford, 1998.
- [42] K. Wehrheim and C. Woodward, *Orientations for pseudoholomorphic quilts*, preprint, arXiv:1503.07803, 2015.
- [43] K. Wehrheim and C. T. Woodward, *Exact triangle for fibered Dehn twists*, *Res. Math. Sci.*, **3**(2016), Paper No. 17, 75.
- [44] E. Witten, *Monopoles and four-manifolds*, *Math. Res. Lett.*, **1**(1994), no. 6, 769–796.
- [45] F. Zapolsky, *The Lagrangian Floer-quantum-PSS package and canonical orientations in Floer theory*, preprint, arXiv:1507.02253, 2015.
- [46] I. Zemke, *Graph cobordisms and Heegaard Floer homology*, preprint, arXiv:1512.01184, 2015.
- [47] I. Zemke, *Quasistabilization and basepoint moving maps in link Floer homology*, *Algebr. Geom. Topol.*, **17**(2017), no. 6, 3461–3518.
- [48] I. Zemke, *Link cobordisms and functoriality in link Floer homology*, *J. Topol.*, **12**(2019), no. 1, 94–220.
- [49] I. Zemke, *A graph TQFT for hat Heegaard Floer homology*, *Quantum Topol.*, **12**(2021), no. 3, 439–460.

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY, 450 JANE STANFORD WAY, BUILDING 380, STANFORD, CA 94305-2125, USA.

Email address: abouzaid@stanford.edu

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY, 450 JANE STANFORD WAY, BUILDING 380, STANFORD, CA 94305-2125, USA.

Email address: cm5@stanford.edu