

Lagrangian matching invariants for fibred four-manifolds: II

TIM PERUTZ

In the second of a pair of papers, we complete our geometric construction of “Lagrangian matching invariants” for smooth four-manifolds equipped with broken fibrations. We prove an index formula, a vanishing theorem for connected sums and an analogue of the Meng–Taubes formula. These results lend support to the conjecture that the invariants coincide with Seiberg–Witten invariants of the underlying four-manifold, and are in particular independent of the broken fibration.

53D40, 57R57; 57R15

1 Introduction

This paper is the sequel to [24], which we shall refer to as Part I. Our overall aims and main theorems were stated in the introduction to that paper, but briefly, the goal is to understand Seiberg–Witten invariants of a four-manifold equipped with a “broken fibration”—that is, a singular Lefschetz fibration in the sense of Auroux–Donaldson–Katzarkov [1]—using constructions in symplectic topology directly reflecting the geometry of the fibration. We shall not achieve quite as much as that: what we shall manage is to construct “Lagrangian matching invariants”, as invariants of the fibration, and demonstrate, in a number of ways, their resemblance to Seiberg–Witten invariants of the four-manifold. However, equality of the Lagrangian matching invariants with the Seiberg–Witten invariants remains conjectural, as does their independence of the fibration.

1.1 Fibred coisotropic hypersurfaces

An important role in the programme is taken by certain Lagrangian correspondences between symmetric products of surfaces. In Part I we constructed these correspondences as vanishing cycles for certain degenerations and studied their properties. They arise from *fibred coisotropic hypersurfaces* in symmetric products. One way to describe their significance, which we did not do explicitly in Part I, is by means of a theorem which we now proceed to state.

Let Σ be a closed Riemann surface equipped with an area form α . Fix $n \geq 2$ and a Kähler form ω on $\text{Sym}^n(\Sigma)$ representing a cohomology class (in $H^2(\text{Sym}^n(\Sigma); \mathbb{R})$) invariant under action of the mapping class group of Σ . When $g(\Sigma) > 0$ the space of such invariant classes is two-dimensional, and there is a standard basis (η, θ) for it:

- η is dual to any of the smooth divisors $\text{Sym}^{n-1}(\Sigma) \subset \text{Sym}^n(\Sigma)$, embedded by the map $D \mapsto x + D$ for a fixed $x \in \Sigma$.
- θ corresponds to the intersection form on $H_1(\Sigma; \mathbb{R})$ under the natural map $\Lambda^2 H^1(\Sigma) \rightarrow H^2(\text{Sym}^n(\Sigma))$ (that is, the usual isomorphism $\Lambda^2 H^1(\Sigma) \rightarrow H^2(\text{Jac}(\Sigma))$ followed by AJ^* , where $\text{AJ}: \text{Sym}^n(\Sigma) \rightarrow \text{Jac}(\Sigma)$ is the Abel–Jacobi map).

When $g(\Sigma) = 0$, the class θ vanishes, but η spans the invariant part of the second cohomology. The classes η and θ will come up repeatedly in this paper, as they did in Part I.

After rescaling ω , we may write its cohomology class as $[\omega] = \eta + \lambda\theta$ where λ lies in (some bounded-below interval in) \mathbb{R} .

One obtains a homomorphism between symplectic mapping class groups

$$\kappa_{n,\lambda}: \frac{\text{Aut}(\Sigma, \alpha)}{\text{Ham}(\Sigma, \alpha)} \rightarrow \frac{\text{Aut}(\text{Sym}^n(\Sigma), \omega)}{\text{Ham}(\text{Sym}^n(\Sigma), \omega)}$$

by the following procedure. Given $\phi \in \text{Aut}(\Sigma, \alpha)$, its mapping torus $Y = \text{T}(\phi) \rightarrow S^1$ is a bundle with fibre Σ carrying a natural cohomology class $\alpha_\phi \in H^2(\text{T}(\phi); \mathbb{R})$ extending α . A choice of complex structure on its vertical tangent bundle makes its relative symmetric products $\text{Sym}_{S^1}^n(Y) = \{(t, D) : t \in S^1, D \in \text{Sym}^n(Y_t)\}$ into smooth manifolds. They fibre over S^1 with complex fibres. A point which arose in Part I is that there are natural maps

$$H^2(Y) \rightarrow H^2(\text{Sym}_{S^1}^n(Y)), \quad c \mapsto c^{[1]}$$

and
$$H^0(Y) \rightarrow H^2(\text{Sym}_{S^1}^n(Y)), \quad c \mapsto c^{[2]},$$

defined by means of the universal divisor

$$\Delta = \{(x, D) : x \in \text{Supp}(D)\} \subset Y \times_{S^1} \text{Sym}_{S^1}^n(Y)$$

and its Poincaré dual cohomology class δ . The second projection

$$\text{pr}_2: Y \times_{S^1} \text{Sym}_{S^1}^n(Y) \rightarrow \text{Sym}_{S^1}^n(\Sigma)$$

is a Σ -bundle. Using also the first projection pr_1 , one puts

$$c^{[1]} = (\text{pr}_2)_!(\text{pr}_1^*(c) \smile \delta), \quad c^{[2]} = (\text{pr}_2)_!(\text{pr}_1^*(c) \smile \delta \smile \delta).$$

There is a nonempty, convex set of closed two-forms on $\text{Sym}_{S^1}^n(Y)$ which are Kähler on the fibres and which globally represent the unique cohomology class in the linear span of $(\alpha_\phi)^{[1]}$ and $1^{[2]}$ which extends the class $\eta + \lambda\theta$ on a reference fibre. One verifies that $1^{[2]}$ restricts to the fibre $\text{Sym}^n(\Sigma)$ as $2(n\eta - \theta)$; this, in tandem with MacDonald's formula $c_1(\text{Sym}^n(\Sigma)) = (n + 1 - g)\eta - \theta$, shows that the class extending $\eta + \lambda\theta$ is

$$(1) \quad \frac{1 + \lambda n}{\langle \Sigma, \alpha \rangle} (\alpha_\phi)^{[1]} - \frac{\lambda}{2} 1^{[2]}.$$

We then define $\kappa_{n,\lambda}([\phi])$ to be the Hamiltonian-isotopy class of the symplectic monodromy of $\text{Sym}_{S^1}^n(Y)$. One checks that this procedure defines a homomorphism.

In Section 2 of Part I, we defined spherically fibred coisotropic submanifolds and Dehn twists along them.

Theorem 1.1 *Let τ_γ denote the positive Dehn twist along the embedded circle $\gamma \subset \Sigma$. Then, assuming $\lambda > 0$, $\kappa_{n,\lambda}(\tau_\gamma)$ coincides with the Hamiltonian isotopy class represented by a fibred Dehn twist along an S^1 -fibred coisotropic submanifold $V_\gamma \subset \text{Sym}^n(\Sigma)$, itself determined by γ up to Hamiltonian isotopy. The reduced space of V_γ is diffeomorphic to $\text{Sym}^{n-1}(\bar{\Sigma})$, where $\bar{\Sigma}$ is the result of surgery along γ .*

This theorem is essentially due to I Smith [37, Proposition 3.7], drawing on ideas of Seidel, though to obtain the monodromy symplectically (up to Hamiltonian isotopy), and not merely smoothly, we invoke the monodromy theorem from Part I, Section 2.

The theorem follows readily from Part I, in which the construction of V_γ was given in detail (Theorem A). Beware that whilst it is easy to think of a hypersurface in $\text{Sym}^n(\Sigma)$ determined by γ , namely the image of the natural map $\gamma \times \text{Sym}^{n-1}(\Sigma) \rightarrow \text{Sym}^n(\Sigma)$, this is not diffeomorphic, let alone isotopic, to V_γ (it is, however, homologous to V_γ). The significance of V_γ is more subtle: it is the vanishing cycle of a relative Hilbert scheme of points on the fibres of an elementary Lefschetz fibration. The monodromy theorem then shows that the monodromy of an elementary symplectic Morse–Bott fibration, such as this Hilbert scheme, is the fibred Dehn twist about its vanishing cycle (modulo Hamiltonian isotopies).

Corollary 1.2 *When $\lambda > 0$, $\text{im}(\kappa_{n,\lambda})$ is generated by fibred Dehn twists.*

Proof We claim that $\text{Aut}(\Sigma, \alpha)/\text{Ham}(\Sigma, \alpha)$ is generated (as a group, not a monoid) by Dehn twists; the corollary then follows from the theorem.

Moser [22] proved that $\text{Aut}(\Sigma, \alpha)$ has the same homotopy type as $\text{Diff}(\Sigma)$. Thus $\pi_0 \text{Aut}(\Sigma, \alpha)$ is the ordinary mapping class group $\pi_0 \text{Diff}(\Sigma)$, which is generated by

(area-preserving) Dehn twists. Hence it suffices to show that the kernel of the quotient map $\text{Aut}(\Sigma, \alpha)/\text{Ham}(\Sigma, \alpha) \rightarrow \pi_0 \text{Aut}(\Sigma, \alpha)$ is generated by Dehn twists. This kernel is isomorphic, via the flux homomorphism, to $H^1(\Sigma; \mathbb{R})/\Gamma$, where Γ is the flux group. By the Earle–Eells theorem [7], Γ is zero when $g(\Sigma) > 1$ and $\Gamma = H^1(\Sigma; \mathbb{Z})$ when $g(\Sigma) = 1$. By considering the Dehn twist along a curve γ , followed by the inverse twist along a parallel curve γ' , one sees that elements of $H^1(\Sigma; \mathbb{R})/\Gamma$ are indeed generated by Dehn twists. \square

It is unknown whether the corollary holds when $\lambda \leq 0$.

1.2 Plan of the paper

In Section 4 of Part I we showed how families of Lagrangian correspondences, parametrised by S^1 , could be cast as Lagrangian boundary conditions associated with elementary broken fibrations (Theorem B). In Part II we continue where we left off. We begin, in Section 2, by restating Theorem B, or rather, stating a corollary of it which is our jumping-off point for this part. In Section 3 we set out what we need from pseudo-holomorphic curve theory. This is then applied to relative Hilbert schemes of points arising from broken fibrations, with the Lagrangian boundary conditions we have constructed.

The definition of the Lagrangian matching invariants is given in Section 4.2. It invokes a particular case of the theory of pseudo-holomorphic sections already described. However, rather than applying the theory in its “raw” form, we show that the resulting invariants $\mathcal{L}_{(X, \pi)}$ for a broken fibration π on a four-manifold X can be organised in a way which is strongly reminiscent of Seiberg–Witten theory:

- A “topological sector” for our theory determines (and is in many cases determined by) a Spin^c -structure for X (Section 4).
- The local expected dimension of the moduli space of pseudo-holomorphic sections defining the invariants is equal to that of the Seiberg–Witten moduli space (Theorem D).
- There are relative invariants in Floer homology groups $HF_*(Y, \mathfrak{t})$, which one can take to be finitely generated over \mathbb{Z} for each Spin^c -structure. These are constructed as summands (corresponding to components of the twisted loop space) in the fixed-point Floer homology groups $HF_*(\kappa_{n, \lambda}(\phi))$, where $Y = T(\phi)$, and λ is chosen so as to avoid periods. The group $HF_*(Y, \mathfrak{t})$ is graded by the \mathbb{Z} -set of homotopy classes of oriented two-plane fields underlying \mathfrak{t} (Section 5). Such geometric gradings appear in Kronheimer and Mrowka’s monopole Floer homology theory for three-manifolds [16].

Computations are carried out in Section 6, including the proofs of Theorem E (“Meng–Taubes formula”) and Theorem F (“vanishing for connected sums”), both of which were stated in the introduction to Part I.

Acknowledgements I thank Simon Donaldson, Paul Seidel, Ivan Smith and Michael Usher for their useful input. This work was supported by the EPSRC via Research Grant EP/C535995/1.

2 Matching conditions

In this section we state some of our conclusions from Part I in a language we shall find convenient for the pseudo-holomorphic curve theory still to come.

Definition 2.1 Let S be a compact surface with an even number of boundary components. A *matching* for S is a decomposition

$$\partial S = \partial S_+ \amalg \partial S_-$$

(defined by a continuous function $\partial S \rightarrow \{\pm 1\}$) together with an orientation-reversing diffeomorphism $\tau: \partial S_+ \rightarrow \partial S_-$.

In Part I, Definition 2.1, a *symplectic Morse–Bott fibration* was defined in terms of data $(E, \pi, \Omega, J_0, j_0)$, where E is a manifold with boundary, $\pi: E \rightarrow S$ a smooth proper map to a surface, mapping ∂E submersively to ∂S ; Ω is a closed two-form, nondegenerate on the tangent distribution $\ker(D\pi)$; J_0 (resp. j_0) is the germ of an almost complex structure near the critical set X^{crit} (resp. near $\pi(X^{\text{crit}})$), making $D\pi$ holomorphic. X^{crit} is assumed to be a submanifold, and the complex Hessian form on its normal bundle nondegenerate. We additionally made certain integrability assumptions near X^{crit} , which we do not repeat here.

Definition 2.2 Let $(E^{2n+2}, \pi, \Omega, J_0, j_0)$ be a symplectic Morse–Bott fibration over a surface S equipped with a matching. Let $(\partial E_{\pm}, \pi_{\pm}, \Omega_{\pm})$ be the restriction of E to ∂S_{\pm} . A *Lagrangian matching condition* for (E, π) is a submanifold $Q^{n+1} \subset \partial E_+ \times_{\partial S_+} \tau^*(\partial E_-)$ such that (i) the projection $Q \rightarrow \partial S_+$ is a proper, surjective submersion, and (ii) $((-\Omega_+) \oplus \tau^*\Omega_-)|_Q = 0$.

Our main examples of matched surfaces arise from broken fibrations (Part I, Definition 1.1). Let (X, π) be a broken fibration over a closed, oriented surface S . Let Z be the one-dimensional part of X^{crit} , and suppose that $\pi|_{X^{\text{crit}}}$ is injective.

Let N be a narrow, closed tubular neighbourhood of $\pi(Z) \subset S$, with retraction $r: N \rightarrow Z$. Then r gives rise to a matching for $S' := S \setminus \text{int}(N)$. Indeed, if Γ is a component of $\pi(Z)$, there are two components of $\partial S'$ which retract to it under r . The one for which the fibres have lower Euler characteristic (higher genus, in the connected case) is designated as belonging to $\partial S'_+$, the other as belonging to $\partial S'_-$. Clearly the matching τ is an invariant of (X, π) , up to isotopy.

Let $X' = X|_{S'}$ be the restricted fibration.

Now let $\nu: S' \setminus S^{\text{crit}} \rightarrow \mathbb{Z}_{\geq 0}$ be a locally constant function with the property that

$$2\nu(s) + \chi(X_s) = \text{const.}$$

For example, ν could be the intersection number $s \mapsto \beta \cdot [X_s]$ for some relative homology class $\beta \in H_2(X, Z; \mathbb{Z})$ with $\delta(\beta) = [Z] \in H_1(Z; \mathbb{Z})$ (see Lemma 4.9).

Choose a positively oriented complex structure J on the tangent distribution $T^\vee X'$. Let

$$(2) \quad X^{[\nu]} = \text{Hilb}_{S'}^{\nu}(X')$$

be the relative Hilbert scheme, ie, the disjoint union, over components $C \subset S'$, of the relative Hilbert schemes $\text{Hilb}_C^{\nu(C)}(X|_C)$ (see Part I, Section 3). It comes with a map $\pi^{[\nu]}: X^{[\nu]} \rightarrow S'$, and a complex structure $J^{[\nu]}$ on its vertical tangent distribution $\ker D\pi^{[\nu]}$ (this distribution is not of constant rank).

Suppose that $W \in H^2(X; \mathbb{R})$ is a class such that $\langle W, F \rangle > 0$ for every class $F \in H_2(X; \mathbb{Z})$ represented by a component of a fibre of π . Such a W exists by definition of broken fibrations.

Definition 2.3 A closed two-form Ω on $X^{[\nu]}$ is W -admissible if it meets the following conditions:

- $(X^{[\nu]}, \pi^{[\nu]}, \Omega, J_0, j_0)$ is a symplectic Morse–Bott fibration for some (j_0, J_0) , where J_0 agrees with $J^{[\nu]}$ on tangent vectors where both are defined.
- Ω is Kähler on each regular fibre $\text{Sym}^{\nu(s)}(X_s)$, and also on the components of $\text{crit}(\pi^{[\nu]})$.
- There is a $\lambda > 0$ such that the restriction of Ω to each regular fibre $\text{Sym}^{\nu(s)}(X_s)$ represents the class $\eta_{X_s} + \lambda\theta_{X_s}$.
- The restriction of Ω to $\partial X^{[\nu]}$ represents the cohomology class

$$\frac{1 + \lambda n}{\langle W, \text{fibre} \rangle} W^{[1]} - \frac{\lambda}{2} 1^{[2]}$$

occurring in the construction of $\kappa_{n,\lambda}$ (cf equation (1)).

It is *admissible* if it is W -admissible for some W .

Lemma 2.4 *W -admissible two-forms exist. Moreover, the set of admissible forms is convex, hence they are unique up to deformation.*

Proof For existence, one uses the Thurston–Gompf method: see Gompf [13, Theorem 2.2] and also Donaldson–Smith [6, Lemma 4.4]. One can use Equation (17) to rewrite the class

$$\frac{1 + \lambda n}{\langle W, \text{fibre} \rangle} (W|_{\partial X'})^{[1]} - \frac{\lambda}{2} 1^{[2]}$$

in the cohomology of the boundary as

$$\left(\frac{1 + \lambda n}{\langle W, \text{fibre} \rangle} (W|_{\partial X'}) + \frac{\lambda}{2} c_1(T^\vee X)|_{\partial X'} \right)^{[1]} - \lambda c_1(T^\vee \partial X^{[v]}).$$

One then looks for a closed two-form globally representing the class

$$\left(\frac{1 + \lambda n}{\langle W, \text{fibre} \rangle} (W|_{\partial X'}) + \frac{\lambda}{2} c_1(T^\vee X') \right)^{[1]} - \lambda c_1(TX^{[v]}).$$

This is possible locally in the base: this is a nontrivial statement near a critical value, but it was established in Part I, Section 3. Hence, by Thurston–Gompf patching, it is also possible globally. The convexity statement is obvious. \square

Admissibility depends on the choice of J , but the set of J_t -admissible forms varies continuously in a path J_t .

Our rephrasing of Theorem B from Part I is the following:

Theorem 2.5 *For any choice of J -admissible two-form Ω , the symplectic Morse–Bott fibration $(X^{[v]}, \pi^{[v]}, \Omega, J_0, j_0)$ admits a canonical Lagrangian matching condition \mathcal{Q} , up to isotopy. In a smooth path of complex structures J_t , and of J_t -admissible two-forms Ω_t , there are smoothly varying matching conditions \mathcal{Q}_t .*

3 Pseudo-holomorphic sections with Lagrangian matching conditions

In this section we set out a general framework for “Lagrangian matching invariants”, not specifically tied to the Lagrangian matching conditions arising from broken fibrations on four-manifolds.

3.1 Almost complex structures adapted to symplectic Morse–Bott fibrations

Let M be a smooth manifold and $V \rightarrow M$ an oriented real vector bundle. Denote by $\mathcal{J}(V)$ the space of all orientation-compatible, C^∞ complex structures on V , with its C^∞ topology. If (V, σ) is a symplectic vector bundle then the subspace $\mathcal{J}(V, \sigma) \subset \mathcal{J}(V)$ of compatible complex structures—those $J \in \mathcal{J}(V)$ such that the formula $g_x(v_1, v_2) := \sigma(v_1, Jv_2)$ defines a metric on the vector bundle V —is contractible.

Definition 3.1 Let $(E, \pi, \Omega, J_0, j_0)$ be a symplectic Morse–Bott fibration over S , and let $E^* = E \setminus \text{crit}(\pi)$. A pair $(J, j) \in \mathcal{J}(TE) \times \mathcal{J}(TS)$ is *adapted* to the fibration if

- (1) π is (J, j) –holomorphic, ie $D\pi \circ J = j \circ D\pi$;
- (2) $J|_{T^\vee E^*} \in \mathcal{J}(T^\vee E^*, \Omega|_{T^\vee E^*})$, that is, J is compatible with Ω on the smooth points of the fibres;
- (3) J extends the germ J_0 , and j extends j_0 .

Trying not to make the notation *too* cumbersome, we write $\mathcal{J}(E, \pi)$ for the space of adapted pairs, and $\mathcal{J}(E, \pi, j) = \{J : (J, j) \in \mathcal{J}(E, \pi)\}$. With respect to the splitting $TE^* = T^\vee E^* \oplus \pi^* TS^*$ defined by the symplectic connection (or, in fact, any other connection), a pair $(J, j) \in \mathcal{J}(E, \pi)$ has a block decomposition over E^* of shape

$$(3) \quad J = \begin{pmatrix} J^{vv} & J^{vh} \\ 0 & j \end{pmatrix}, \quad J^{vv} \circ J^{vh} + J^{vh} \circ j = 0.$$

With J^{vv} and j fixed, J^{vh} is just a \mathbb{C} –antilinear homomorphism:

$$J^{vh} \in \Gamma \text{Hom}^{0,1}(T^h E^*, T^\vee E^*).$$

Lemma 3.2 *The spaces $\mathcal{J}(E, \pi, j)$ are contractible.*

Proof To see that $\mathcal{J}(E, \pi, j) \neq \emptyset$, fix a closed neighbourhood U of $\text{crit}(\pi)$ on which J_0 can be defined; let $U^* = U \setminus E^{\text{crit}}$. One can extend $J_0|_{T^\vee U^*}$ to an element $J^{vv} \in \mathcal{J}(T^\vee E^*, \Omega|_{T^\vee E^*})$. Over U^* , $J^{vv} \oplus j$ differs from J_0 by an antilinear homomorphism, which can be extended to one defined over E^* . This gives rise to a complex structure J of the right sort.

Now consider the restriction map $r: \mathcal{J}(E, \pi, j) \rightarrow \mathcal{J}(T^\vee E^*, \Omega|_{T^\vee E^*})$. Its fibres are affine spaces modelled on the vector spaces $\Gamma_c \text{Hom}^{0,1}(T^h E^*, T^\vee E^*)$, where Γ_c means sections supported outside X^{crit} . The map r admits a section s , and $\mathcal{J}(E, \pi, j)$ deformation-retracts to the contractible space $\text{im}(s)$, hence is contractible. \square

An easy matrix calculation, linearising (3), shows that $\mathcal{J}(E, \pi)$ has formal tangent spaces $T_J \mathcal{J}(E, \pi, j)$ which fit into short exact sequences

$$0 \rightarrow \text{Hom}^{0,1}((\pi^* TS, j), (T^\vee E, J^{vv})) \rightarrow T_J \mathcal{J}(E, \pi, j) \rightarrow T_{J^{vv}} \mathcal{J}(T^\vee E) \rightarrow 0.$$

Moreover, $\mathcal{J}(E, \pi, j)$ can be made into a smooth Fréchet manifold which fibres smoothly over $\mathcal{J}(T^\vee E)$.

The drawback of working with spaces of almost complex structures of class C^∞ is that they are not complete, which means that in establishing genericity results, one cannot directly invoke the implicit function and Baire category theorems. This is not a serious problem: one fix, due to Floer, is to observe that for any rapidly-decreasing sequence $\epsilon = (\epsilon_n)_{n \geq 1}$ of positive reals there is a dense Banach submanifold $\mathcal{J}(E, \pi, j)_\epsilon \subset \mathcal{J}(E, \pi, j)$ whose tangent vectors are bundle maps with finite “ C_ϵ^∞ –norm”. We refer to Schwarz [29, Section 4.2.1] for the details.

3.2 The moduli space

Now suppose that $(E, \pi, \Omega, J_0, j_0)$ is a symplectic Morse–Bott fibration over a matched surface, with Lagrangian matching condition Q (Definition 2.1 and Definition 2.2).

- Write $\text{sect}(E)$ for the space of C^∞ sections of π , and $\text{sect}(E, Q)$ for the subspace of sections u which “map ∂S into Q .” This is an abbreviated way of saying that the image of the section $(u|_{\partial S_+}, \tau^* u|_{\partial S_-})$ of the fibre product $\partial E_+ \times_{\partial S_+} \tau^* \partial E_-$ lies in Q .
- The tangent space of $\text{sect}(E, Q)$ at u is the subspace $C^\infty(u^* T^\vee E; T^\vee Q)$ of sections $\xi \in C^\infty(u^* T^\vee E)$ such that $(\xi|_{\partial S_+}, \tau^* \xi|_{\partial S_-})$ is a section of $T^\vee Q$.
- We need a Sobolev completion $\text{sect}_1^p(E, Q)$ of $\text{sect}(E, Q)$, defined using a Riemannian metric and a fixed number $p > 2$. It is the smallest subset of the continuous sections $\text{sect}_{C^0}(E, Q)$ which contains $\text{sect}(E, Q)$ and which also contains the section $x \mapsto \exp_{u(x)}(\xi(x))$ for any $u \in \text{sect}(E, Q)$, $\xi \in L_1^p(u^* T^\vee E; T^\vee Q)$. The space $\text{sect}_1^p(E, Q)$ is a smooth Banach manifold with tangent spaces $L_1^p(u^* T^\vee E; T^\vee Q)$.
- Fix $(J, j) \in \mathcal{J}(E, \pi, \Omega)$. The moduli space of (j, J) –pseudo-holomorphic sections with boundary in Q is the space

$$\mathcal{M}_{J,j}(E, Q) = \{u \in \text{sect}_1^p(E, Q) : J \circ (Du) = (Du) \circ j\} \subset \text{sect}_1^p(E, Q),$$

and it is this that we wish to analyse.

- Let $\mathcal{E}_{J,j}^Q \rightarrow \text{sect}(E, Q)$ be the natural infinite-rank vector bundle with fibres $(\mathcal{E}_{J,j}^Q)_u = L^p(\text{Hom}_Q^{0,1}(TS, u^*T^\vee E))$. Here $\text{Hom}_Q^{0,1}(TS, u^*T^\vee E)$ is the vector space of (j, J^{vv}) -antilinear homomorphisms $TS \rightarrow u^*T^\vee E$ which (in the obvious sense) carry $T(\partial S)$ to $u^*T^\vee Q$.

The moduli space $\mathcal{M}_{J,j}(E, Q)$ is the zero-set of the section

$$(4) \quad \bar{\partial}_{J,j} = \frac{1}{2}(D + J \circ D \circ j) \in \Gamma(\mathcal{E}_{J,j}^Q).$$

Allowing J to vary within $\mathcal{J}(E, \pi, j)$ one gets a “universal” vector bundle

$$\mathcal{E}_j^Q \rightarrow \text{sect}(E) \times \mathcal{J}(E, \pi; j).$$

This has a section $\bar{\partial}_j$, whose zero-set $\mathcal{M}_j(E, Q) = \bigcup_J \mathcal{M}_{J,j}(E)$ is the “universal moduli space”.

When $u \in \mathcal{M}_{J,j}(E, Q)$, there is an intrinsically defined linearised operator

$$D_u: L_1^p(u^*T^\vee E; Q) \rightarrow (\mathcal{E}_{J,j}^Q)_u.$$

When $u \in \text{sect}_1^p(E, Q)$ is not holomorphic, the linearisation is not intrinsic, but can be defined by choosing a connection on $\mathcal{E}_{J,j}^Q$. Then for connections ∇_u on the complex vector bundles $u^*T^\vee E$ and bundle maps $a_u: u^*T^\vee E \rightarrow \text{Hom}^{0,1}(TS, E)$, one has

$$D_u v = \nabla_u^{0,1} v + a_u v.$$

3.3 Transversality

Apart from minor differences of context, this is a review of standard theory. We follow Seidel [34] closely, since our set-up is almost the same as his.

The first point to make is that a C^1 section of π cannot intersect E^{crit} , because, along the image of the section, its derivative provides a right inverse to $D\pi$. Consequently the presence of critical points makes no difference to transversality theory for moduli spaces of sections. It does, however, affect their compactifications.

Introduce a torsion-free connection ∇ on $T(E \setminus E^{\text{crit}})$ which restricts to a connection on the vertical subbundle. The linearisation of $\bar{\partial}_{J,j}$ (4) at $u \in \mathcal{M}_{J,j}(E)$ is a linear map

$$(5) \quad D_u: \Gamma(u^*T^\vee E) \rightarrow (\mathcal{E}_{J,j}^Q)_u = \Gamma(\text{Hom}^{0,1}(TS, u^*T^\vee E)).$$

Explicitly,

$$D_u(v) = (u^*\nabla)^{0,1}(v) + \frac{1}{2}(\nabla_v J) \circ Du \circ j.$$

The linearisation of $\bar{\partial}_j$ is $D_{u,J}^{\text{univ}}$, where

$$D_{u,J}^{\text{univ}}: (v, Y) \mapsto D_u v + \frac{1}{2} Y \circ D_u \circ j.$$

The linearised operators extend continuously to maps between the Banach completions of their domains and targets. Thus we have

$$D_u: L_1^p(u^* T^\vee E, T^\vee Q) \rightarrow L^p(\text{Hom}_Q^{0,1}(TS, u^* T^\vee E),$$

$$D_{u,J}^{\text{univ}}: L_1^p(u^* T^\vee E, T^\vee Q) \oplus C_\epsilon^\infty(T_J \mathcal{J}(E, \pi, j)) \rightarrow L^p(\text{Hom}_Q^{0,1}(TS, u^* T^\vee E).$$

The crucial point is that D_u is Fredholm (we will come back to this point shortly). It follows that so too is $D_{u,J}^{\text{univ}}$. The latter is also *surjective*, whence $\mathcal{M}_j(E, Q)$ is a smooth Banach submanifold of $\text{sect}_1^p(E, Q) \times \mathcal{J}(E, \pi; j)_\epsilon$. This surjectivity statement is nonstandard only in one respect, namely, that we are considering almost complex structures on $\partial E_+ \times_{\partial S_+} \tau^* \partial E_-$ which respect the fibre product decomposition. However, the standard argument, involving unique continuation, is unaffected by this point (if $D_u^* \eta = 0$, where η is an L^q section of the dual bundle to $\text{Hom}_Q^{0,1}(TS, u^* T^\vee E)$, $q^{-1} + p^{-1} = 1$, then η must be supported in ∂S ; but $L^q \cap \ker D_u^*$ contains a dense subspace of continuous sections, hence $\eta = 0$).

Write $\mathcal{J}^{\text{reg}}(X, \pi; j)$ for the space of *regular* almost complex structures: those with the property that D_u is onto for every $u \in \mathcal{M}_{J,j}(X, Q)$. Their importance is that when J is regular, $\mathcal{M}_{J,j}(E, Q)$ is a smooth manifold of local dimension $\text{ind } D_u$. The key transversality statement, following from surjectivity of $D_{u,J}^{\text{univ}}$, is that $\mathcal{J}^{\text{reg}}(E, \pi, j) \cap \mathcal{J}(E, \pi, J)_\epsilon$ is a (dense) Baire subset of the complete space $\mathcal{J}(E, \pi, j)_\epsilon$ (as trailed earlier, we use Floer’s C_ϵ^∞ spaces here).

3.3.1 Intersection with cycles in fibres Mark a finite set of points $\{s_i\}_{i \in I}$ in $\text{int}(S)$, manifolds Z_i , and smooth maps $\zeta_i: Z_i \rightarrow X_{s_i}$. There is an evaluation map

$$(6) \quad \text{ev}_I^J = \prod_{i \in I} \text{ev}_i: \mathcal{M}_{J,j}(X, Q) \rightarrow \prod_i X_{s_i}.$$

The arguments that establish the surjectivity of $D_{u,J}^{\text{univ}}$ extend easily to give a “transversality of evaluation” lemma (see Seidel [34, Lemma 2.5]) which says that, for any $J \in \mathcal{J}(X, \pi; j)$, and any neighbourhood U of $\{s_i : i \in I\}$, there exist arbitrarily small perturbations J_t of J , supported in $\pi^{-1}(U)$, such that $J_t \in \mathcal{J}^{\text{reg}}(X, \pi; j)$ and $\text{ev}_I^{J_t}$ is transverse to $\zeta = \prod_i \zeta_i$. Similarly, given marked points $s'_k \in \partial S_+$, and maps $\zeta'_k: Z'_k \rightarrow Q_{s'_k}$, there exist small perturbations of J supported in a chosen neighbourhood of $\{s'_k\} \cup \{\tau^{-1}(s'_k)\}$ making ζ'_k transverse to the natural evaluation

map. (Again, the argument is unaffected by our use of almost complex structures which respect the fibre product structure.)

3.3.2 Fredholm theory It is a simple matter to prove that pseudo-holomorphic sections with Lagrangian matching conditions have a Fredholm deformation theory, granted the standard Fredholm theory for Lagrangian boundary conditions. The argument is most direct in the case where $S = S_+ \amalg S_-$ and the matching diffeomorphism τ of the boundary components extends to a diffeomorphism $\tau': S_+ \rightarrow S_-$: here, a Lagrangian matching condition is simply a Lagrangian boundary condition in a fibre product.

In general, the argument goes as follows. Near ∂S , we can extend τ to an diffeomorphism $\tilde{\tau}$ between collar neighbourhoods C_+ of ∂S_+ and C_- of ∂S_- . Take $u \in \text{sect}(E, Q)$, and let $u_{\pm} = u|_{C_{\pm}}$. Then we have a section $(u_+, \tilde{\tau}^*u_-)$ of the fibre bundle $E|_{C_+} \times_{C_+} \tilde{\tau}^*(E|_{C_-}) \rightarrow C_+$.

We claim that the map (5),

$$D_u: L_1^p(u^*T^v E; Q) \rightarrow (\mathcal{E}_{J,j}^Q)_u,$$

is Fredholm. Over compact subsets of $\text{int}(S)$, one has the usual elliptic estimate for D_u . Near the boundary, one can equivalently work with D_u over $C_+ \cup C_-$ or with the corresponding operator $D_{(u_+, \tilde{\tau}^*u_-)}$ over C_+ . The theory developed by Floer—see [9, Lemma 2.3]—gives an elliptic estimate for the latter. Together, these estimates imply that $\ker(D_u)$ is finite-dimensional. On small open sets in $\text{int}(S)$ or in C_+ , the operator D_u has a bounded right inverse, by elliptic theory, and these may be patched together to form a global parametrix (as in Donaldson [5, Proposition 3.6]), which shows that D_u has closed range and finite-dimensional cokernel.

To state the index formula for D_u , we need to define a “Maslov index” map

$$\mu_Q: \text{sect}(E, Q) \rightarrow \mathbb{Z}.$$

For a component $\partial S_+^{(i)}$ of ∂S_+ , let $v_+^{(i)} = u|_{\partial S_+^{(i)}}$. Let $v_-^{(i)} = u|_{\tau(\partial S_+^{(i)})}$. Choose symplectic trivialisations $t_+^{(i)}$ for $(v_+^{(i)})^*T^v E$ and $t_-^{(i)}$ for $\tau^*(v_-^{(i)})^*T^v E$. Then the pullback $(v_+^{(i)}, \tau \circ v_-^{(i)})^*T^v Q$ defines (via $t_+^{(i)} \oplus t_-^{(i)}$) a loop in a Lagrangian Grassmannian, and so has a classical Maslov index l_i . We put

$$\mu_Q(u) = 2c_1^{\text{rel}}(u^*T^v E) + \sum_i l_i,$$

where $c_1^{\text{rel}}(u^*T^v E)$ is the Chern number of $u^*T^v E \rightarrow S$ relative to the trivialisations $t_{\pm}^{(i)}$ over ∂S .

Lemma 3.3 *The index of D_u is*

$$\text{ind}(D_u) = \mu_Q(u) + \sum_{S_i \in \pi_0(S)} \text{rk}_{\mathbb{C}}(T^\vee E|_{S_i}) \chi(S_i).$$

Proof We use the observation from Eliashberg–Givental–Hofer [8] that the index of Cauchy–Riemann operators is invariant under “cutting and pasting”. For a positive boundary component $\partial S_+^{(i)}$, let $\gamma_+^{(i)}$ be a loop in S' , parallel to $\partial S_+^{(i)}$, and let $V_+^{(i)} = (u^* T^\vee E)|_{\gamma_+^{(i)}}$. Let $\delta \subset S^2$ be the equator, and let $U \rightarrow S^2$ be a complex vector bundle with $c_1(U) = 0$ and $U|_\delta = V_+^{(i)}$. Cut S' along $\gamma_+^{(i)}$, and S^2 along δ to obtain a Riemann surface with boundary, then re-glue the four boundary components so as to obtain a new Riemann surface (the result of surgery on $\gamma_+^{(i)}$) which inherits a vector bundle. By [8, 1.8, Axiom C3], this vector bundle also inherits a Cauchy–Riemann operator $D_{u,+}^{(i)}$, and $\text{ind}(D_{u,+}^{(i)}) = \text{ind}(D_u) + 2 \text{rank}_{\mathbb{C}}(V_+^{(i)})$. (According to [8] one should consider real analytic data here; however, ultimately this will make no difference.)

Perform this surgery operation for each i ; also perform similar surgeries on a loop $\gamma_-^{(i)}$ parallel to $\tau(\partial S_+^{(i)})$, for each i . The result is (i) a matched Riemann surface C which is the union of a closed Riemann surface C' and k pairs of discs $(D_+^{(i)}, D_-^{(i)})$ with matched boundaries; (ii) a Cauchy–Riemann operator D'_u over it. The ordinary Riemann–Roch theorem computes the index of D'_u over C' . Riemann–Roch for surfaces with boundary computes the index over $D_+^{(i)} \cup D_-^{(i)}$ as $l_i + \text{rank}(T^\vee E)|_{D_+^{(i)}} + \text{rank}(T^\vee E)|_{D_-^{(i)}}$, where l_i is a Maslov index as in the definition of $\mu_Q(u)$. Putting these things together gives the result. \square

3.4 Compactness

The *action* of a section $u \in \text{sect}(E, Q)$ is

$$\mathcal{A}(u) = \int_S u^* \Omega.$$

Though Ω itself need not be symplectic, $\Omega + c\pi^*\beta$ is symplectic (tamed by J) for any positive area form $\beta \in \Omega_S^2$ and $c \gg 0$. The symplectic area for sections then differs from the action by a constant, $c \int_S \beta$. Gromov’s compactness theorem (the relevant version is that of Ye [40]) says that any sequence u_1, u_2, \dots in $\mathcal{M}_{J,j}(X, Q)$ of bounded symplectic area—equivalently, bounded action—has a subsequence which converges, in Gromov’s topology, to a pseudo-holomorphic curve v with bubbles and boundary bubbles.

The principal component of a Gromov-limit is a differentiable map $S \rightarrow E$ which is a section over the complement of a discrete set in S (the roots of the bubble-trees),

hence is globally a section. Hence to compactify $\{u \in \mathcal{M}_{J,j}(E, Q) : \mathcal{A}(u) \leq \lambda\}$ one need only consider curves whose components are (i) sections; (ii) bubbles in regular fibres; (iii) bubbles in singular fibres; (iv) boundary bubbles.

3.4.1 Fibred monotonicity One does not expect moduli spaces of pseudo-holomorphic sections with Lagrangian boundary conditions to resemble closed manifolds in general, even to the extent of carrying fundamental homology classes. The reason is that there is typically a codimension 1 boundary, corresponding to bubbling off of discs. Various hypotheses can be imposed to make sure that there is in fact no such boundary in codimension 1. Ours will be fibred weak monotonicity; indeed, we will mostly be concerned with fibred monotonicity. In our applications, which will involve symmetric products $\text{Sym}^n(\Sigma)$, fibred monotonicity holds when $n \geq g(\Sigma)$.

Definition 3.4 (a) A closed Lagrangian submanifold L in a compact symplectic manifold (M, ω) is called *monotone* if there exists $c > 0$ such that

$$\int_{\bar{D}} u^* \omega = c \mu(u)$$

for every smooth map $u: (\bar{D}, \partial \bar{D}) \rightarrow (M, L)$.

(b) A symplectic Morse–Bott fibration (E, π, Ω) over a matched surface, with Lagrangian matching condition Q , is *fibred-monotone* if, for each $s \in \partial S_+$, Q_s is monotone in $(\partial E_+ \times_{\partial S_+} \tau^* \partial E_-)_s$.

For a Lagrangian submanifold $L \subset M$, let $\mu_L: \pi_2(M, L) \rightarrow \mathbb{Z}$ be the Maslov index homomorphism and $\mu_{\min}(L)$ the minimal Maslov index:

$$\text{im}(\mu_L) = \mu_{\min}(L)\mathbb{Z}, \quad \mu_{\min}(L) \geq 0.$$

Let $c_{\min}(M)$ be the minimal Chern number of M :

$$\text{im}(c_1: \pi_2(M) \rightarrow \mathbb{Z}) = c_{\min}(M)\mathbb{Z}, \quad c_{\min} \geq 0.$$

Since the Maslov index of a disc whose boundary is mapped to a point is twice its Chern number, $\mu_{\min}(L) \leq 2c_{\min}(M)$. For a Lagrangian boundary condition Q , let $\mu_{\min}(Q) = \gcd_{s \in \partial S} \mu_{\min}(Q_s)$.

In the following proposition, we suppose that $E^{\text{crit}} = \emptyset$, so that (E, π, Ω) is actually a locally Hamiltonian fibration (LHF).

Proposition 3.5 *Let Q be a Lagrangian matching condition for the LHF (E^{2n+2}, π, Ω) over a matched surface S , and fix $h \in \pi_0 \text{sect}(E, Q)$. Fix also finite sets $I \subset \text{int}(S)$,*

$J \subset \partial S_+$; and smooth cycles $\zeta_s: Z_s \rightarrow E_s$ for each $s \in S$, and $\zeta': Z_{s'} \rightarrow Q_{s'}$ for $s' \in J$. Let $Z = \prod Z_s$, $\zeta = \prod \zeta_s: Z \rightarrow \prod E_s$; and $Z' = \prod Z_{s'}$, $\zeta' = \prod \zeta'_{s'}: Z \rightarrow \prod Q_{s'}$. Suppose that

- (1) both E and Q are fibre-monotone;
- (2) the virtual dimension $d(h) = \text{ind}(D_u)$ ($u \in h$) satisfies

$$(7) \quad d(h) - \sum_{s \in I} (2n - \dim(Z_s)) - \sum_{s' \in J} (n - \dim(Z'_{s'})) - \mu_{\min}(Q) < 0.$$

Then for any $j \in \mathcal{J}(TS)$ and for a dense set of $J \in \mathcal{J}^{\text{reg}}(X, \pi, j)$, the fibre product

$$\mathcal{M}_{J,j}(E, Q) \times_{(\text{ev}, \zeta, \times \zeta')} (Z \times Z')$$

is a compact, smooth manifold of dimension

$$d(h) - \sum_I (2n - \dim(Z_s)) - \sum_J (n - \dim(Z'_{s'}))$$

when this number is nonnegative, and empty when it is negative.

Proof By the remarks on transversality of evaluation (6), there is a dense subset of $\mathcal{J}^{\text{reg}}(X, \pi, j)$ whose members J have the property that the evaluation map

$$\text{ev} = \text{ev}_A \times \text{ev}'_B: \mathcal{M}_{J,j}(X, Q) \rightarrow \prod_a E_{s_a} \times \prod_b Q_{s'_b}$$

is transverse to $\zeta \times \zeta'$. This means that there is a smooth moduli space

$$\mathcal{M} = \mathcal{M}_{J,j}(E, Q) \times_{(\text{ev}, \zeta, \times \zeta')} (Z \times Z').$$

Any sequence in \mathcal{M} , for which the sections all lie in h , has a subsequence with a Gromov limit comprising a principal component u_{prin} , nonconstant bubbles $\{u_l\}_{l \in L}$, and nonconstant boundary bubbles $\{u'_k\}_{k \in K}$. These satisfy

$$d(h) = \text{ind}(u_{\text{prin}}) + 2 \sum_{l \in L} c_1(u_l) + \sum_{k \in K} \mu(u'_k).$$

By fibred monotonicity, the Maslov indices $\mu(u'_k)$ are positive multiples of $\mu_{\min}(Q)$, hence $\geq \mu_{\min}(Q)$. The contribution of each term $c_1(\beta_l)$ from a bubble in a fibre is also positive, at least $\mu_{\min}(Q)$. Then if $L \cup K \neq \emptyset$ we will have $\text{ind}(u_{\text{prin}}) - \sum_a (2n - \dim(Z_a)) - \sum_b (n - \dim(Z'_b)) < 0$. But for regular almost complex structures, this number gives the dimension near u_{prin} of \mathcal{M} , which cannot be negative. \square

There are many possible hypotheses to make singularities in the fibres permissible, and we do not attempt an axiomatisation. In the case of relative Hilbert schemes of points on Lefschetz fibrations, they can be admitted because of the following observation of Donaldson and Smith:

Lemma 3.6 [6, Lemma A.11] *Let $\text{Hilb}_S^n(E)$ be the relative Hilbert scheme of n points on a Lefschetz fibration $E \rightarrow S$ with irreducible fibres, and $\text{Hilb}^n(E_s)$ a singular fibre. Then any holomorphic sphere in $\text{Hilb}^n(E_s)$ which arises as a bubble-component in a Gromov limit of pseudo-holomorphic sections is homotopic to a sphere in a regular fibre.*

3.4.2 Invariance We now give the parametric version of Proposition 3.5. All the data involved in the construction can be allowed to move in smooth families. To begin with, fix the symplectic Morse–Bott fibration $(E, \pi, \Omega, J_0, j_0)$, Lagrangian boundary condition Q , and homotopy class h .

Each j and each $J \in \mathcal{J}^{\text{reg}}(E, \pi; j)$ gives a moduli space $\mathcal{M}_{J,j}(E, Q;)_h$ which we know to be a smooth, oriented manifold; after a small perturbation of J , the fibre product $\mathcal{M}_{J,j}(E, Q;)_h \times_{\text{ev}, \zeta \times \zeta'} Z \times Z'$ with a family of smooth oriented cycles in the fibres (as considered above) is also smooth and oriented. It carries a smooth boundary-evaluation map

$$\text{ev}(J, j): \mathcal{M}_{J,j}(E, Q)_h \rightarrow \text{sect}(Q)$$

to the space of sections of Q .

Proposition 3.7 *Assume that the hypotheses of Proposition 3.5 hold, and that the left-hand side of (7) is $< -(k+1)$, for some $k \geq 0$. Then the oriented bordism class of the map $\text{ev}(J, j)$ is independent of (J, j) . Moreover, any smooth map $S^k \rightarrow \mathcal{J}^{\text{reg}}(X, \pi)$ induces a bundle \mathcal{M}_{S^k} over S^k and a map $\text{ev}: \mathcal{M}_{S^k} \rightarrow \text{sect}(Q)$ which extends to a map to $\text{sect}(Q)$ from a bundle over B^{k+1} .*

Proof One considers parametrised moduli spaces. The space $\mathcal{J}(E, \pi)$ is contractible; thus any map $f: S^k \rightarrow \mathcal{J}^{\text{reg}}(X, \pi)$ extends to a map $F: B^{k+1} \rightarrow \mathcal{J}(X, \pi)$. The transversality theory extends in a straightforward way, showing that after a small homotopy of F , fixing f , one gets a smooth moduli space of the appropriate dimension which inherits an orientation from that on B^{k+1} . The numerical conditions imply that compactness goes through as before, and this gives the result. \square

Likewise, the bordism class of the evaluation map is unchanged under deformations of Ω , J_0 , j_0 , the points s_a , s'_b and the cycles ζ_a and ζ'_b .

3.5 Fibred k -negativity

In the framework of fibred monotonicity, bubbling in low-dimensional moduli spaces (with regularity assumptions in force) is ruled out on the grounds that the principal component of a limit curve must have nonnegative index. A more refined method incorporates transversality for pseudo-holomorphic spheres in the fibres and discs in the boundary fibres. One can then determine numerical conditions under which moduli spaces of sections are generically disjoint from those of fibrewise spheres and discs. In the context of Hamiltonian Floer homology, these are the “weak monotonicity” conditions of Hofer and Salamon [15] (a symplectic manifold (M^{2n}, ω) is weakly monotone if it is monotone, or $c_{\min} = 0$, or $c_{\min} \geq n - 2$). To handle Lagrangian boundary conditions effectively, somewhat more conservative assumptions are required. We have seen that monotone Lagrangians in monotone symplectic manifolds are acceptable, and we will consider one other hypothesis, drawing on technical results of Lazzarini.¹

Definition 3.8 For any $k \geq 0$, a compact symplectic manifold (M^{2n}, ω) is k -negative if, for any sphere $f: S^2 \rightarrow M$ with $\int_{S^2} f^* \omega > 0$, one has $\langle f^* c_1(TM), [S^2] \rangle \leq 3 - n - k$. A Lagrangian $L \subset M$ is called k -negative if any disc $f: (D, \partial D) \rightarrow (M, L)$ with $\int_D f^* \omega > 0$ has Maslov index $\leq 2 - n - k$.

We shall need to consider 3-negative manifolds and 2-negative Lagrangians.

3.5.1 Transversality in the fibres For k -negative manifolds ($k \geq 0$) generic almost complex structures $J \in \mathcal{J}(M, \omega)$ are regular in the sense that for every element $\beta \in \mathcal{M}_J^{\text{si}}(M)$ of the moduli space of parametrised, simple J -holomorphic spheres, the relevant linearised operator D_β is surjective (a pseudo-holomorphic curve $u: C \rightarrow M$ is simple if it does not factor through a branched covering $C \rightarrow C'$ of degree > 1 ; by a lemma of McDuff (see McDuff and Salamon [21, Proposition 2.51]) this is equivalent to the set of injective points, $\{z \in S^2 : \beta^{-1}(\beta(z)) = \{z\}\}$, being nonempty, or indeed open and dense). For such J , $\mathcal{M}_J^{\text{si}}(M)$ is a manifold of local dimension $2(c_1(\beta) + n)$, and its quotient by the free action of $\text{Aut}(S^2) = \text{PSL}(2, \mathbb{C})$ has dimension $2(c_1(\beta) + n - 3)$. Thus there can be no simple holomorphic spheres. Since every pseudo-holomorphic sphere factors through a simple one, there are no nonconstant spheres at all. By the same

¹There is, unfortunately, no straightforward general theory for Lagrangians with $\mu_{\min} = 0$ in symplectic manifolds with $c_{\min} = 0$ (these are the characteristics of vanishing cycles in $\text{Sym}^{g-1}(\Sigma) \times \text{Sym}^{g-2}(\bar{\Sigma})$). This is a borderline situation in which there is just enough leeway to obtain compact zero-dimensional moduli spaces, but not enough to handle one-parameter families. We shall not venture into this hazardous territory.

argument, one sees that *in a k -negative symplectic manifold, there are no nonconstant spheres in a generic k -parameter family of almost complex structures.*

As with spheres, transversality theory for pseudo-holomorphic discs in (M, L) works in a straightforward way for *simple* discs—those for which the set of injective points is nonempty, or equivalently open and dense—but the complication is that not every nonconstant disc factors through a simple one. One can, however, use Lazzarini’s lemma [18]: *if δ is a nonconstant J -disc then there is a simple J -disc δ' such that $\delta'(\partial\bar{D}) \subset \delta(\partial\bar{D})$.* The index formula says that the moduli space of parametrised simple discs $\mathcal{M}_J^{\text{si}}(M, L)$ has virtual dimension $\mu_L + n$. Since $\text{Aut}(\bar{D})$ is 3-dimensional and acts freely, the unparametrised moduli space has dimension $\mu_L + n - 3$. Thus, for generic k -parameter families of almost complex structures, the moduli space of simple discs is empty. Lazzarini’s lemma allows us to remove “simple” from this conclusion.

Since we are interested in fibrations over surfaces, with Lagrangian boundary conditions, we should consider 3-negative fibres and 2-negative Lagrangians in them. (Remember that we need a spare parameter to prove invariance!)

3.5.2 Singular fibres Symplectic Morse–Bott fibrations can be handled using the method developed for Lefschetz fibrations in Seidel’s thesis [32]. One needs to make some hypothesis about the normalisation \tilde{M}_0 of a singular fibre M_0 . The almost complex structure $J_0 = J|_{M_0}$ on M_0 lifts to one \tilde{J}_0 on \tilde{M}_0 , integrable near the preimage of the normal crossing divisor; it is sufficient to assume that there \tilde{J}_0 is compatible with a 1-negative symplectic form on \tilde{M}_0 . The point is that any J_0 -sphere in M_0 lifts uniquely to a \tilde{J}_0 -sphere in \tilde{M}_0 . One cannot expect \tilde{J}_0 to be regular; however, for generic J , lifts of nonconstant spheres in M_0 are regular, hence have nonnegative Chern number: here we use a standard argument which shows that one can achieve regularity for all spheres passing through an open set U by making a perturbation of the almost complex structure supported in U .

Let us summarise our conclusion:

Proposition 3.9 *Let Q be a Lagrangian matching condition for a symplectic Morse–Bott fibration (E^{2n+2}, π, Ω) over a matched surface S , and fix $h \in \pi_0 \text{sect}(E, Q)$. Fix also finite sets of regular values $I \subset \text{int}(S)$ and $J \subset \partial S_+$; and smooth cycles $\zeta_s: Z_s \rightarrow E_s$ for each $s \in S$, and $\zeta'_{s'}: Z_{s'} \rightarrow Q_{s'}$ for $s' \in J$. Let $Z = \prod Z_s$, $\zeta = \prod \zeta_s: Z \rightarrow \prod E_s$; and $Z' = \prod Z'_{s'}$, $\zeta' = \prod \zeta'_{s'}: Z' \rightarrow \prod Q_{s'}$. Suppose that*

- (1) *the fibres of E are 3-negative; those of Q are 2-negative;*
- (2) *the components of E^{crit} are 1-negative;*

(3) the number

$$D(h; Z_i, Z_j) := n\chi(S) + \mu_Q(h) - \sum_{s \in I} (2n - d_s(Z)) - \sum_{s' \in J} (n - d'_{s'}(Z'))$$

is nonpositive.

Here $d_s(Z) = \dim Z_s$, $d'_{s'}(Z') = \dim Z'_{s'}$. Then, for any $j \in \mathcal{J}(TS)$ and for a dense set of $J \in \mathcal{J}^{\text{reg}}(E, \pi, j)$, the fibre product

$$\mathcal{M}_{J,j}(E, Q) \times_{(\text{ev}, \zeta \times \zeta')} (Z \times Z')$$

is a compact, smooth manifold of dimension $D(h; Z, Z')$ when this number is negative, and empty when it is zero. Moreover, for generic paths in $\mathcal{J}^{\text{reg}}(E, \pi, j)$, the one-parameter moduli space is a compact 1-manifold with boundary when $D(h; Z_i, Z'_j) = 0$, and empty when $D(h; Z, Z') < 0$.

Remark 3.10 The same transversality arguments also apply when the fibres, Lagrangians, and critical manifolds are all monotone; this is a more complicated approach to the one discussed above, but gives a slightly sharper result: one need only assume that $D(h) \leq 0$. Note also that this technique handles the singular fibres of $\text{Hilb}_S^n(E)$ for $n \geq g$ in the case previously neglected, when E has reducible singular fibres.

3.6 Orientations

Since the first version of this paper was written, a draft of Seidel's book on Fukaya categories [35] has become available. It includes a thorough account of orientations in Lagrangian Floer theory incorporating nonorientable Lagrangians. It seems churlish to ignore this useful reference, particularly as the draft of this section dealt unconvincingly with one or two points. The account we now present should be regarded, for the most part, as expository; besides Seidel's book, its sources are de Silva's thesis [36] and the book of Fukaya–Oh–Ohta–Ono [12].

Consider a space \mathcal{M} defined as the zero-set of a nonlinear Fredholm map f between Banach spaces. Suppose, moreover, that the linear Fredholm maps F_x linearising f are all surjective ("regularity"). Then \mathcal{M} is a smooth manifold with tangent spaces $\ker(F_x)$. Orientability of this manifold means triviality of $\Lambda^{\max} T\mathcal{M}$, the line bundle with fibres $\Lambda^{\max} \ker(F_x)$. In practice, \mathcal{M} sits inside a larger space \mathcal{B} parametrising a family of Fredholm operators $\mathcal{F} = \{F_x\}_{x \in \mathcal{B}}$, and $\Lambda^{\max} T\mathcal{M}$ extends to the larger space as the determinant index bundle $\text{Det}(\mathcal{F}) \rightarrow \mathcal{B}$, the natural line bundle with fibres $\Lambda^{\max} \ker(F_x) \otimes \Lambda^{\max}(\text{coker}(F_x))^*$. One then studies w_1 of this bundle.

The Fredholm operators we wish to consider are Cauchy–Riemann operators over a Riemann surface with boundary, subject to a Lagrangian boundary condition. Let S be a compact Riemann surface with one boundary component. Let $E \rightarrow S$ be a (trivial) symplectic vector bundle of rank $2n$. We consider the space $\text{Lag}(S, E)$ of Lagrangian subbundles $F \subset E|_{\partial S}$. By fixing a diffeomorphism $\partial S \cong S^1$ and a trivialisation of $E|_{\partial S}$, we can identify $\text{Lag}(S, E)$ with the space of unbased loops $\mathcal{L}\text{Gr}_n$ in Gr_n , the Grassmannian of Lagrangian subspaces of the standard symplectic vector space $(\mathbb{C}^n, \omega_{\mathbb{C}^n})$. Notice that the Maslov index decomposes $\text{Lag}(S, E)$ (or equally $\mathcal{L}\text{Gr}_n$) into connected components $\text{Lag}_k(S, E)$ (or $\mathcal{L}_k\text{Gr}_n$) indexed by $k \in \mathbb{Z}$.

By choosing a conjugate-linear bundle map $a \in V := \text{Hom}^{0,1}(TS, E)$, we get an \mathbb{R} -linear, Fredholm Cauchy–Riemann problem for (E, F) : that of solving $(\bar{\partial} + a)(u) = 0$, where $u \in L_1^p(E)$, for some fixed $p > 2$, and $u(z) \in F_z$ for $z \in \partial S$. Our object of study is the determinant index bundle $\text{Det}(\mathcal{F}) \rightarrow V \times \mathcal{L}\text{Gr}_n$.

The space $V \times \mathcal{L}\text{Gr}_n$ obviously deformation-retracts to the subspace $\{0\} \times \mathcal{L}\text{Gr}_n$. This in turn deformation-retracts to $\text{U}(n)/\text{O}(n)$, via a deformation retraction of $\text{Sp}(2n, \mathbb{R})$ to $\text{U}(n)$. Hence $\text{Lag}(S, E)$ is homotopy-equivalent to $\mathcal{L}(\text{U}(n)/\text{O}(n))$.

Remark 3.11 It is worth noting that if one considers only the based loops, and passes to the limit $\text{U}/\text{O} = \varinjlim_n \text{U}(n)/\text{O}(n)$, there is a homotopy equivalence

$$\Omega(\text{U}/\text{O}) \simeq \mathbb{Z} \times B\text{O}$$

with the classifying space of KO -theory. Bott [3] constructed such an equivalence using Morse theory, as the final step in the periodic octagon starting with $\mathbb{Z} \times B\text{O}$; there is also an equivalence $\Omega^7(\mathbb{Z} \times B\text{O}) \simeq \text{U}/\text{O}$. We shall not need Bott periodicity, but we shall be concerned with the functorial map

$$[X, \mathcal{L}(\text{U}/\text{O})] \rightarrow KO(X)$$

defined by stabilising the virtual index bundle, and with its restriction to $\Omega(\text{U}/\text{O})$.

Following [35] (which itself extends [36]), we now give the formula for $w_1(\text{Det}(\mathcal{F})) \in H^1(\text{Lag}(S, E); \mathbb{Z}/2)$. Evaluation $e: S^1 \times \mathcal{L}_k\text{Gr}_n \rightarrow \text{Gr}_n$ induces a map

$$e^*: H^*(\text{Gr}_n; \mathbb{Z}/2) \rightarrow H^*(S^1 \times \mathcal{L}_k\text{Gr}_n; \mathbb{Z}/2).$$

For $c \in H^i(\text{Gr}_n; \mathbb{Z}/2)$, let $T(c) = \int_{S^1} e^*c \in H^{i-1}(\text{Gr}_n; \mathbb{Z}/2)$, and let $U(c) = (e^*c)|_{(\{z\} \times \text{Gr}_n)} \in H^i(\text{Gr}_n; \mathbb{Z}/2)$. It is easy to check that

$$H^1(\mathcal{L}_k\text{Gr}_n; \mathbb{Z}/2) = \mathbb{Z}/2 \oplus \mathbb{Z}/2, \quad n \geq 3,$$

with generators $T(w_2)$ and $U(\mu)$, $\mu \in H^1(\mathcal{L}_k \text{Gr}_n; \mathbb{Z}/2)$ being the Maslov index reduced mod 2. The first Stiefel–Whitney class for the determinant index bundle over the component $\mathcal{L}_k \text{Gr}_n$ is then

$$(8) \quad w_1(\text{Det}(\mathcal{F})) = T(w_2) + (k + 1)U(\mu).$$

Given a loop $\gamma: S^1 \rightarrow \text{Lag}(S, E)$, lying in the component where the Maslov index is k , one can take the torus of boundary values $\partial\gamma: S^1 \times \partial S \rightarrow L$. What the formula says is that, if k is odd, one has

$$\langle w_1(\gamma^* \text{Det}(\mathcal{F})), [S^1] \rangle = \langle (\partial\gamma)^* w_2(TL), [S^1 \times \partial S] \rangle.$$

If k is even one must add a correction term $\langle w_1 l^* TL, [S^1] \rangle$, where l is the loop $(\partial\gamma)|_{S^1 \times \{x\}}: S^1 \rightarrow L$. In particular, $w_1(\gamma^* \text{Det})$ vanishes when L is spin.

Fukaya et al [12] show that one gets an actual orientation for the moduli space $\mathcal{M}_S(M, L)$ of J -holomorphic maps $(S, \partial S) \rightarrow (M, L)$ by giving a “relative spin structure”, that is, an oriented vector bundle $\xi \rightarrow M$, with $w_2(\xi|L) = w_2(TL)$, and a spin structure on $\xi|L \oplus TL$. This is not quite enough for our purposes, because it precludes nonorientable Lagrangian boundary conditions. We can get around this by working not with spin but with *pin* (again, cf [35]).

The Lie group $\text{Pin}^+(n)$ is one of the central extensions of $\text{O}(n)$ by $\mathbb{Z}/2$. Such central extensions of topological groups are classified by $H^2(\text{BO}(n); \mathbb{Z}/2)$, which is the direct sum of two copies of $\mathbb{Z}/2$ generated by the universal characteristic classes w_2 and w_1^2 . The central extension associated with w_2 (resp. $w_1^2 + w_2$) is $\text{Pin}^+(n)$ (resp. $\text{Pin}^-(n)$). A well-known algebraic construction assigns to any quadratic vector space (V, Q) a group $\text{Pin}(V, Q)$ as the subgroup of the unit group of the Clifford algebra $\text{Cl}(V)$ (in which the relation $v^2 = Q(v)1$ holds) generated by $Q^{-1}\{\pm 1\} \subset V \subset \text{Cl}(V)$. In these terms, $\text{Pin}^\pm(n) = \text{Pin}(\mathbb{R}^n, \pm \|\cdot\|^2)$.

A Pin^+ -structure for a real n -plane bundle $\zeta \rightarrow Z$ is a homotopy class of homotopy-liftings $Z \rightarrow B \text{Pin}^+(n)$ of the classifying map $Z \rightarrow B \text{O}(n)$. The obstruction to such a lift is $w_2(\zeta)$; when this is zero, the Pin^+ -structures form an affine space modelled on $H^1(Z; \mathbb{Z}/2)$. The inclusions $\text{O}(n) \rightarrow \text{O}(n + 1)$ lift to the Pin^+ -groups, and this gives the meaning of “stable”.

Example 3.12 When n is odd, one has an isomorphism

$$\text{O}(n) \xrightarrow{\cong} \mathbb{Z}/2 \times \text{SO}(n), \quad A \mapsto (\det(A), \det(A)^{-1} \cdot A).$$

The four central extensions of $\mathbb{Z}/2 \times \text{SO}(n)$ by $\mathbb{Z}/2$ are

$$(1) \quad \mathbb{Z}/2 \times \mathbb{Z}/2 \times \text{SO}(n),$$

- (2) $\mathbb{Z}/4 \times \mathrm{SO}(n)$,
- (3) $\mathbb{Z}/2 \times \mathrm{Spin}(n)$, and
- (4) $(\mathbb{Z}/4 \times \mathrm{Spin}(n))/(\mathbb{Z}/2)$ (quotient by the “diagonal” involution).

Which of these is $\mathrm{Pin}^+(n)$? Certainly one of the last two, since the covering $\mathrm{Pin}^+(n) \rightarrow \mathrm{O}(n)$ must restrict to the spin covering $\mathrm{Spin}(n) \rightarrow \mathrm{SO}(n)$. The algebraic fact [35, Lemma 11.14] which allows us to decide between them is that the two preimages of $-\mathrm{id} \in \mathrm{O}(n)$ in $\mathrm{Pin}^+(n)$ both have order 2 when $n \equiv 1 \pmod{4}$, and 4 when $n \equiv 3 \pmod{4}$. Hence $\mathrm{Pin}^+(n)$ is $\mathbb{Z}/2 \times \mathrm{Spin}(n)$ or $(\mathbb{Z}/4 \times \mathrm{Spin}(n))/(\mathbb{Z}/2)$ according to whether n is 1 or 3 mod 4.

Now suppose that (X^{2n+2}, π, Ω) is a LHF over a compact matched surface S , and $Q \subset \partial X$ a Lagrangian matching condition. Thus Q is a subbundle of $Y_+ \times_{\partial S_-} Y_-$, where $Y_- = \pi^{-1}(\partial S_-)$ and Y_+ is the pullback to ∂S_- of $\pi^{-1}(\partial S_-)$ by the matching $\tau: \partial S_- \rightarrow \partial S_+$.

Definition 3.13 A relative pin structure for (X, Q) is

- a stable oriented vector bundle $\xi \rightarrow X$, with $w_2(\xi|_Q) = w_2(T^\vee Q)$;
- a stable Pin^+ -structure on $\xi' \oplus T^\vee Q$. Here ξ' denotes the restriction to Q of the bundle $\mathrm{pr}_1^*(\tau^*\xi|_{Y_+}) \oplus \mathrm{pr}_2^*(\xi|_{Y_-})$ over $Y_+ \times_{\partial S_-} Y_-$.

Lemma 3.14 Suppose that the first projection embeds Q in Y_+ , and that the second projection is an S^1 -bundle over Y_- . There is then a canonical relative pin structure for (X, Q) .

Proof The boundaries ∂S_- and ∂S_+ lie on different components of S , since the respective fibres of π have different dimensions, $2n$ and $2(n-1)$ respectively. Let ξ be the bundle $T^\vee X$ over the components containing ∂S_- , and the trivial bundle over those containing ∂S_+ . Thus $\xi' \rightarrow Y_+ \times_{\partial S_-} Y_-$ is $\mathrm{pr}_2^* T^\vee Y_-$. We exhibit a pin structure on $\xi'|_Q \oplus T^\vee Q$.

$T^\vee Q$ contains a distinguished line-subbundle λ , the kernel of the projection to $T^\vee Y_-$. Choose a Euclidean metric on $T^\vee Q$, and notice that $\lambda^\perp = \mathrm{pr}_2^* T^\vee Y_-$. The splitting $T^\vee Q = \lambda \oplus \mathrm{pr}_2^* T^\vee Y_-$ reduces the structure group of $\xi'|_Q \oplus T^\vee Q$ to the subgroup

$$\mathrm{O}(1) \times \mathrm{SO}(2n-2) \hookrightarrow \mathrm{O}(4n-3), \quad (A, B) \mapsto \mathrm{diag}(A, B, B).$$

We claim that this homomorphism factors through $\mathrm{Pin}^+(4n-3)$. This will give a canonical relative pin structure.

The restriction to identity components, $\{1\} \times \mathrm{SO}(2n-2) \rightarrow \mathrm{SO}(4n-3)$, lifts to $\mathrm{Spin}(4n-3)$ because it kills $\pi_1 \mathrm{SO}(2n-2)$. Write this lift as l . Let $\sigma \in \mathrm{Pin}^+(4n-3)$ be an element of order 2 which maps to $(-1, I_{2n-2}, I_{2n-2})$; this exists by Example 3.12. The map $(A, B) \mapsto \sigma(A)l(A)$ gives the claimed lift. \square

Proposition 3.15 *Suppose J is a compatible complex structure for (E, Ω) such that $D\pi \circ J = i \circ D\pi$. Consider the moduli space of J -holomorphic sections $\mathcal{M}_J(X, Q)$; assume it is transversely cut out. A relative pin structure for (X, Q) induces an orientation for any component of $\mathcal{M}_J(X, Q)$ on which the Maslov index is odd. A relative pin structure together with an orientation for one fibre Q_x induces an orientation for components of even Maslov index.*

Proof We assume, up until the last moment, that the base is a disc D .

Step (i) The regularity assumption means that the cokernels of the relevant linear Cauchy–Riemann operators are trivial, and hence that the moduli space is smooth. The tangent space to $\mathcal{M}_J(X, Q)$ at u is then the index of a Fredholm operator associated with the Cauchy–Riemann problem $(u^*T^\vee X, \partial u^*T^\vee Q)$. Thus the top exterior power of the tangent space is the determinant line of the corresponding Fredholm operator. The usefulness of working over a disc is that $u^*T^\vee X$ can be trivialised in a canonical way, up to homotopy. Hence the top exterior power of the tangent bundle to $\mathcal{M}_J(X, Q)$ is isomorphic to the pullback of $\mathrm{Det}(\mathcal{F})$ (the determinant bundle of the universal family of Cauchy–Riemann operators over $\mathcal{L}_k \mathrm{Gr}_n$) by a classifying map

$$\tau: \mathcal{M}_J(X, Q) \rightarrow \mathcal{L}_k \mathrm{Gr}_n.$$

Step (ii) On the other hand, $\mathrm{Det}(\mathcal{F})$ gives rise to a double covering $p: \tilde{\mathcal{L}} \rightarrow \mathcal{L}_k \mathrm{Gr}_n$ as the S^0 -bundle sitting inside the line bundle. The pullback line bundle $p^* \mathrm{Det}(\mathcal{F})$ has a *tautological* trivialisation. It will therefore suffice to show that a stable pin structure determines a lift of τ to a map $\tilde{\tau}: \mathcal{M}_J(X, Q) \rightarrow \tilde{\mathcal{L}}$, so that the pullback by $\tilde{\tau}$ of the tautological trivialisation of $p^* \mathrm{Det}(\mathcal{F}) \rightarrow \tilde{\mathcal{L}}$ gives the sought orientation.

Step (iii) Now consider the pullback of the tautological \mathbb{R}^n -bundle $U \rightarrow \mathrm{Gr}_n$ (whose fibre at Λ is Λ itself) by the evaluation map $e: S^1 \times \mathcal{L} \mathrm{Gr}_n \rightarrow \mathrm{Gr}_n$. There is a double covering of $\mathcal{L} \mathrm{Gr}_n$ whose fibre at F is a stable pin structure in the bundle $(e^*U)|_{S^1 \times \{F\}}$ over S^1 . This double covering is the S^0 -bundle of a unique line bundle ν . By construction, $w_1(\nu) = T(w_2)$.

We are given a stable pin structure on $\xi|_Q \oplus T^\vee Q$. This pulls back to a stable pin structure on $(\partial u)^*(\xi|_Q \oplus T^\vee Q)$, for any $u \in \mathcal{M}_J(X, Q)$. But $(\partial u)^*(\xi|_Q)$ is canonically trivial, since it extends to an oriented vector bundle $u^*\xi$ over the disc.

Hence we obtain a stable pin structure on $(\partial u)^* T^\vee Q$. In other words, our relative pin structure gives rise to a section of the covering $\tau^* \nu$.

Step (iv) Our task is now to compare the information found in steps (ii) and (iii). Suppose that k is odd. According to Equation (8), one then has $w_1 \text{Det}(\mathcal{F}) = w_1(\nu)$; thus, $\text{Det}(\mathcal{F})$ and ν are isomorphic line bundles. To specify an isomorphism, one has only to do so for some fibre. The particular choice will not really matter for our purposes; notionally, however, we follow the convention of [35, Lemma 11.17]. *A fortiori*, one then has an isomorphism $\tau^* \text{Det}(\mathcal{F}) \cong \tau^* \nu$. As in (iii), the pin structure gives rise to a section σ of the S^0 -bundle in $\tau^* \nu$; composed with the isomorphism with $\tau^* \text{Det}(\mathcal{F})$, one then gets a lift $\tilde{\tau}$ as demanded in step (ii).

When k is even, one proceeds along similar lines, but because of Equation (8), the result is only an isomorphism $\text{Det}(\mathcal{F}) \cong \nu \otimes \lambda$, where λ is the pullback to $\mathcal{L}_n \text{Gr}_n$ of the line bundle $\Lambda^n U \rightarrow \text{Gr}_n$. The pin structure then trivialises $\tau^* \nu$, but one also needs the given orientation of Q_x to trivialise $\tau^* \lambda$.

Finally, in the general case where S is not a disc, one proceeds by a trick of “pinching off the boundary”, as explained in both [12; 35]. \square

Remark 3.16 Suppose k is even, that (X, Q) admits a relative pin structure, and that the fibres Q_z of $Q \rightarrow S^1$ are orientable but Q itself is not. Choosing a relative pin structure and an orientation for Q_z , we determine an orientation for $\mathcal{M}_S(X, Q)$. On the other hand, sliding z around the circle, we eventually come back to the *opposite* orientation of Q_z , and hence to the opposite orientation of the moduli space, which is absurd. Hence the moduli space is empty in this case—an observation which one can also prove in more elementary fashion.

3.6.1 Invariants We can now define our prototypical Lagrangian matching invariants, using a familiar procedure. Suppose that $(E, \pi, \Omega, J_0, j_0)$ is a symplectic Morse–Bott fibration, Q a Lagrangian matching condition, $h \in \pi_0 \text{sect}(E, Q)$. Let $I \subset \text{int}(S)$ and $J \subset \partial S$ be finite sets. Suppose *either* that the fibres, critical manifolds and Lagrangians are monotone, *or* that they satisfy the negativity assumptions of the last proposition.

We write $(E^\#, Q^\#)$ for (E, Q) with the enhancements of a relative pin structure for (E, Q) and, if Q is orientable, an orientation.

The invariant takes the form of homomorphism

$$(9) \quad \Phi_{E, Q; I, J}(h): \bigotimes_{s \in I} H_*(E_s; \mathbb{Z}) \otimes \bigotimes_{s' \in J} H_*(Q_{s'}; \mathbb{Z}) \rightarrow \mathbb{Z}.$$

Thinking of the products of cycles Z and Z' as defining monomials $[Z]$, $[Z']$ in the tensor products, we put

$$\Phi_{E,Q;I,J}(h)([Z] \otimes [Z']) = \begin{cases} \#\mathcal{M}_{J,j}(E, Q) \times_{(ev, \xi, \times \xi')} (Z \times Z'), & D(h; Z, Z') = 0; \\ 0 & D(h; Z, Z') \neq 0. \end{cases}$$

Here $\#$ is the signed count of points in a compact, oriented zero-manifold. This is independent of choices because the one-parametric moduli spaces are compact, oriented one-manifolds with boundary.

3.6.2 Floer homology Here we review the definition of Floer homology for symplectic automorphisms. This is a well-established theory, and we shall be brief (Fredholm and gluing theory are as in Schwartz [29]; transversality and compactness as in Hofer-Salamon [15] and Seidel [32]).

Preliminaries

- Recall that we have defined a locally Hamiltonian fibration (LHF) to be a triple (E, π, Ω) , where E is a smooth manifold with boundary, π a smooth proper submersion $\pi: E \rightarrow B$ to another manifold with boundary, mapping ∂E to ∂B , and $\Omega \in \Omega^2(E)$ a closed two-form which is nondegenerate on the fibres.
- For a commutative ring R , the universal Novikov ring Λ_R is the ring of formal series $\sum_{c \in \mathbb{R}} a(c)t^c$, where $a: \mathbb{R} \rightarrow R$ has the property that $\text{Supp}(a) \cap (-\infty, C)$ is a finite set for each $C \in \mathbb{R}$. When R is a field, so is Λ_R .

Floer homology theory assigns to each LHF (Y, π, σ) over a closed, oriented 1-manifold Z , a $\mathbb{Z}/2$ -graded Λ_R -module

$$HF_*(Y, \sigma) = HF_0(Y, \sigma) \oplus HF_1(Y, \sigma),$$

though we should give two caveats: (i) the fibres Y_z of π should be weakly monotone symplectic manifolds in the sense of [15] (the definition was given in the discussion preceding Definition 3.8); (ii) the definition is simpler when R has characteristic 2.

There is a splitting (of $\mathbb{Z}/2$ -graded modules) into “topological sectors”, ie components of the space of sections, $\text{sect}(Y)$:

$$HF_*(Y, \sigma) = \bigoplus_{\gamma \in \pi_0 \text{sect}(Y)} HF_*(Y, \sigma)_\gamma.$$

There is a continuation isomorphism $HF_*(Y, \sigma) \rightarrow HF_*(Y, \sigma')$ when there is an $\alpha \in \Omega^1(Y)$ with

$$\sigma - \sigma' = d\alpha, \quad \alpha(T^\vee Y) = 0.$$

$HF_*(Y, \sigma)$ is, formally, the Morse–Novikov homology of the action 1–form $\mathcal{A}_{Y, \sigma} \in \Omega^1(\text{sect}(Y))$:

$$\mathcal{A}_{Y, \sigma}(\gamma; \xi) = \int_Z \sigma(\dot{\gamma}, \xi), \quad \xi \in C_Z^\infty(\gamma^* T^\vee Y).$$

The set of zeros of $\mathcal{A}_{Y, \sigma}$ coincides with the set $\mathcal{H}(Y, \sigma)$ of horizontal sections defined by the Hamiltonian connection:

$$\mathcal{H}(Y, \sigma) = \{v: Z \rightarrow Y : \pi \circ v = \text{id}_Z, \text{im } D_z v \subset T^h X_{v(z)}\}.$$

One should again enlarge the space of sections $\text{sect}(Y)$ to a Banach manifold, namely $\text{sect}_1^2(Y)$, which has tangent spaces $L_1^2(\gamma^* T^\vee Y)$. The transversality condition for a zero v of $\mathcal{A}_{Y, \sigma}$ is surjectivity of the operator

$$\xi \mapsto \nabla_{\partial_t} \xi; \quad L_1^2(v^* T^\vee Y) \rightarrow L^2(v^* T^\vee Y),$$

where ∇ is the intrinsic connection along v . After fixing basepoints $z_i \in Z$, one in each component, we can consider the linear holonomy maps $L_{v, i} \in \text{End}(T^\vee Y_{v(z_i)})$. Transversality of v is equivalent in turn to the invertibility of the linear maps $\text{id} - L_{v, i} \in \text{End}(T_{v(z_i)}^\vee Y)$. This is the condition that $v(z)$ is a nondegenerate fixed point of the monodromy $\phi \in \text{Aut}(Y_z, \sigma|_{Y_z})$. It follows that when Y is *nondegenerate*, meaning that every $v \in \mathcal{H}(Y, \sigma)$ is transverse, the set $\mathcal{H}(Y, \sigma)$ is finite. In the next section, we shall see how to assign to each v a free R –module of rank 1, denoted by $|o(v)|$. When R has characteristic 2, $|o(v)| = R$. Now, for a nondegenerate LHF (Y, π, σ) , one sets

$$CF(Y, \sigma) = \bigoplus_{v \in \mathcal{H}(Y, \sigma)} |o(v)| \otimes_R \Lambda_R,$$

This is the module underlying Floer’s (co)chain complex. The definition does not yet use the orientation of Z . For a general LHF, one perturbs it to a nondegenerate one and proceeds as before.

Each $v \in \mathcal{H}(Y, \sigma)$ has a Lefschetz number $l_v \in \mathbb{Z}/2$, defined when Z is connected to be 0 if $\det(\text{id} - L_v) > 0$ and 1 if $\det(\text{id} - L_v) < 0$, and in general by summing the Lefschetz numbers over components of Z . These give the $\mathbb{Z}/2$ –grading of the complex, $CF = CF_0 \oplus CF_1$, with CF_i generated by $\{v : l_v = i\}$. It follows that when $Z = \coprod Z_i$ there are canonical $\mathbb{Z}/2$ –graded isomorphisms $CF(Y) \cong \bigotimes_i CF(Y|Z_i)$.

A compatible vertical almost complex structure $J^v \in \mathcal{J}(T^\vee Y, \sigma)$ extends uniquely to an almost complex structure J_0 on $TY \oplus \varepsilon^1$ (where ε^1 is the trivial real line bundle over Y) such that $J(\tilde{\partial}_t) = 1 \in \mathbb{R}$ (here ∂_t is the unit length vector field on $S^1 = \mathbb{R}/\mathbb{Z}$ and $\tilde{\partial}_t$ its horizontal lift).

Since $T(Y \times \mathbb{R}) = TY \oplus \varepsilon^1$, J_0 induces a translation-invariant complex structure on $T(Y \times \mathbb{R})$. It is still denoted J_0 , and called *cylindrical*. If we regard $Y \times \mathbb{R}$ as a

locally Hamiltonian fibration over the cylinder $S^1 \times \mathbb{R}$, with two-form $\text{pr}_1^* \sigma$, then J_0 preserves both vertical and horizontal subbundles, and the projection to $S^1 \times \mathbb{R} = \mathbb{C}/i\mathbb{Z}$ is holomorphic.

The differential on $CF_*(Y, \sigma)$ is defined via the moduli space $\mathcal{M}_{J,j}(Y, \pi, \sigma)$ of finite-action pseudo-holomorphic sections of

$$\pi \times \text{id}: Y \times \mathbb{R} \rightarrow Z \times \mathbb{R},$$

ie sections u satisfying $J \circ (Du) = (Du) \circ j$ and $\int_{Z \times \mathbb{R}} u^* \sigma < \infty$. Under the assumption of nondegeneracy, any $u \in \mathcal{M}_{J,j}(Y \times \mathbb{R}, \pi \times \text{id}, \text{pr}_1^* \Omega)$ has the following asymptotic behaviour:

- As $s \rightarrow \pm\infty$, the loops $u(s, \cdot)$ converges pointwise towards a horizontal section $v^\pm \in \mathcal{H}(Y, \sigma)$.
- The convergence is exponentially fast with respect to the Riemannian metric g given by $\sigma(\cdot, J\cdot)$ on $T^v Y$ and for which $\tilde{\partial}_t$ is a unit-length vector field.

When Y has weakly monotone fibres, there is a good transversality theory for gradient trajectories [11; 15]. This shows that a dense set of cylindrical almost complex structures are “regular” in the sense that (i) the deformation operators are cylindrical for all trajectories $u \in \mathcal{M}_{J,j}^{\leq 2}(Y, \pi, \sigma)$ (the superscript ≤ 2 refers to the index); (ii) these trajectories do not hit any point of Y which lies on a pseudo-holomorphic sphere $S \subset Y_t$ with $c_1(S) \leq 0$. For regular almost complex structures, the moduli spaces $\mathcal{M}_{J,C}^{\leq 2}(v_-|Y|v_+)$ of trajectories of index ≤ 1 , asymptotic limits v^\pm , and action $\leq C$, have compactness (as well as regularity) properties: any sequence has a subsequence converging in the Gromov–Floer topology to a broken trajectory.

We shall show in the next section how to assign to an arbitrary $u \in \mathcal{M}_{J,j}^1(v_-|Y|v_+)/\mathbb{R}$ an isomorphism $i_u: |o(v_-)| \rightarrow |o(v_+)|$. To define the differential in the Floer complex, we formally sum these isomorphisms, weighting them by their action. Thus, for $x \in |o(v_-)|$, we set

$$\partial x = \sum_{v_+ \in \mathcal{H}(Y, \sigma)} \sum_{u \in \mathcal{M}_{J,j}^1(v_-|Y|v_+)/\mathbb{R}} i_u(x) \otimes t^{A(u)},$$

and extend ∂ to a Λ_R -linear endomorphism of $CF_*(Y, \sigma)$. There are chain homotopy equivalences (“continuation maps”) between the Floer complexes for $(Y, \sigma; J)$ and (Y, σ', J') when there exists $\alpha \in \Omega^1(Y)$ with $\sigma - \sigma' = d\alpha$ and $\alpha(T^v Y) = 0$. The argument by which one proves that $\partial^2 = 0$ and that continuation maps are chain maps is famous enough that we will not even adumbrate it. However, we do still need to say what $|o(v)|$ and i_u are, and we shall do that in the next section.

3.7 Coherent orientations in fixed point Floer homology

Coherent orientations for Floer-theoretic moduli spaces were introduced by Floer and Hofer [10]. The essential technical ingredient is their (linear) gluing construction for cylindrical Cauchy–Riemann operators, together with an index theorem which expresses the determinant of the glued operator as the tensor product of those of the factors. Besides [10], we draw on Seidel’s treatment of the Lagrangian case [35].

In our sketch of Floer homology, we have so far avoided discussing the relevant linearised operators. These are *cylindrical Cauchy–Riemann operators*. Let S be a Riemann surface with cylindrical ends (and compact outside the ends). We compactify it to a Riemann surface \bar{S} with one boundary component for each end: this is done by identifying an outgoing end $S^1 \times (0, \infty) \subset S$ with $S^1 \times (0, 1)$ via the diffeomorphism $(0, \infty) \rightarrow (0, 1)$, $s \mapsto s(1 + s^2)^{-1/2}$, then adding a circle $S^1 \times \{1\}$; similarly for incoming ends. Take a hermitian vector bundle $(E, \|\cdot\|)$ over \bar{S} , and a unitary connection ∇ in it. The Cauchy–Riemann operator $\bar{\partial}_\nabla$ is the operator $C^\infty(S, E) \rightarrow \Omega^1(S, E)$ defined by

$$\bar{\partial}_\nabla \xi = (\nabla \xi)^{0,1}.$$

The behaviour of ∇ over the boundary of \bar{S} is of importance. Namely, for each component T of the boundary, one wants the linear map $\nabla_{\partial/\partial t}$, operating on sections of $E|_T$, to have trivial kernel. In this case, $\bar{\partial}_\nabla$ is called a *nondegenerate* Cauchy–Riemann operator, and it extends to a Fredholm operator $L_1^p(S, E) \rightarrow L^p(S, T^*S \otimes E)$.

Take $\nu \in \mathcal{H}(Y, \sigma)$, and form the pullback $\nu^*T^\nu Y \rightarrow S^1$ as a symplectic vector bundle. It carries a canonical symplectic connection $\nabla_{\nu, \sigma}$, the resulting of linearising the connection in $T^\nu Y$ defined by σ , and the nondegeneracy assumption is that the monodromy $L = L_\nu$ of $\nabla_{\nu, \sigma}$ does not have 1 as an eigenvalue.

Let $x = \nu([0]) \in M$, and consider the space \mathcal{P}_L of paths $\gamma: [0, 1] \rightarrow U(T_x M)$ with $\gamma(0) = 1$ and $\gamma(1) = L$. The component-set $\pi_0(\mathcal{P}_L)$ is an affine copy of \mathbb{Z} ; indeed, the homotopy classes of such paths γ make up the fibre over L of the universal covering $\tilde{U}(T_x M) \rightarrow U(T_x M)$. We denote by $L^\# [n]$ the result of acting on $L^\# \in \pi_0(\mathcal{P}_m)$ by $n \in \mathbb{Z}$.

Associated with any $\gamma \in \mathcal{P}_L$ is a Cauchy–Riemann operator $\bar{\partial}_\gamma$ in the trivial vector bundle $T_x M \times \mathbb{C} \rightarrow \mathbb{C}$. The assumption that L does not have 1 in its spectrum guarantees its nondegeneracy. It has a determinant line $\delta_\gamma = \det \bar{\partial}_\gamma$, and these lines make up a bundle $\delta \rightarrow \mathcal{P}$.

Now, each component of \mathcal{P} is simply connected, since its fundamental group is isomorphic to $\pi_2(U(T_x M)) = 0$; hence δ is a trivial line bundle. We can therefore associate

with a lift $L^\# \in \widetilde{U}(T_x M)$ of L a line $o(L^\#)$, unique up to canonical isomorphism, by putting $o(L^\#) = \delta_\gamma$ for any γ representing the component $L^\#$.

The orientation group $|o(L^\#)|$ is the abelian group generated by the two orientations ω_1 and ω_2 of o_γ , modulo the relation that $\omega_1 + \omega_2 = 0$. Thus a choice of isomorphism determines an isomorphism $|o(L^\#)| \cong \mathbb{Z}$.

How does $o(L^\#[n])$ compare to $o(L^\#)$? The answer is pleasantly simple: there is a canonical homotopy-class of isomorphisms between them. This is proved using the Floer–Hofer gluing theorem: $o(L^\#[n])^* \otimes o(L^\#)$ is the determinant line of a Cauchy–Riemann operator over S^2 , obtained by gluing. Such determinant lines are canonically oriented since they form a simply connected space which contains determinant lines of complex-linear operators.

The last observation shows that the orientation groups $|o(L^\#[n])|$ are canonically isomorphic to a single group which we may denote by $|o(v)|$. Formally, $|o(v)|$ is the group of families $(a_n)_{n \in \mathbb{Z}}$ where $a_n \in |o(L^\#[n])|$ and a_n maps to a_0 under the canonical isomorphism $|o(L^\#[n])| \rightarrow |o(L^\#)|$.

We now have the promised orientation lines $|o(v)|$. The next step is to show how $u \in \mathcal{M}_J(v_-, v_+)$ determines an isomorphism $|o(v_-)| \rightarrow |o(v_+)|$. For this we invoke, once again, the linear gluing theorem of Floer–Hofer. The tangent space to $\mathcal{M}_J(v_-, v_+)$ is the kernel of a surjective Fredholm operator D_u —namely, a Cauchy–Riemann operator $\bar{\partial}_\nabla$ in the bundle $u^* T^\vee Y$ over the cylinder, where ∇ extends the canonical symplectic connection over the ends $\{\pm\infty\} \times S^1$. We cap off the two ends of the cylinder to make S^2 . The bundle $v^* T^\vee Y$, being trivial, extends to a hermitian vector bundle over S^2 , and the connection ∇ also extends. The resulting Cauchy–Riemann operator L over S^2 is homotopic to a complex-linear operator, hence its real determinant line is canonically oriented. But it is also homotopic to an operator obtained by gluing $\bar{\partial}_\nabla$ to two cylindrical Cauchy–Riemann operators over discs, and this gives an isomorphism $\det L \cong \det(D_u) \otimes o(L_-^\#) \otimes o(L_+^\#)^*$, where L_\pm is the monodromy of $\nabla_{v_\pm, \sigma}$, $L_-^\#$ is an arbitrary lift of L_- , and $L_+^\#$ the resulting lift of L_+ (obtained from it by path-lifting). The upshot is that u determines an isomorphism $\Lambda^{\text{top}} T_u \mathcal{M}_J(v_- | Y | v_+) \cong o(L_-^\#)^* \otimes o(L_+^\#)$ (up to a positive factor).

Let us abbreviate $\mathcal{M}_J(v_- | Y | v_+)$ to \mathcal{M} , and $\mathcal{M}_J(v_- | Y | v_+)/\mathbb{R}$ to \mathcal{M}^* . If u is a nonconstant trajectory, translation in the \mathbb{R} -direction determines a canonical tangent vector $\eta \in T_u \mathcal{M}$, and hence an isomorphism $\iota(\eta): \Lambda^{\text{top}} T_u \mathcal{M} \rightarrow \Lambda^{\text{top}} T_{[u]} \mathcal{M}^*$. Thus there is a composite isomorphism

$$\Lambda^{\text{top}} T_{[u]} \mathcal{M}^* \cong o(L_-^\#)^* \otimes o(L_+^\#).$$

Now when $T_{[u]}\mathcal{M}^*$ is 0-dimensional, its top exterior power is simply \mathbb{R} , which has a canonical generator 1. Hence, plugging this generator into our isomorphism, we obtain a generator for $o(L_-^\#)^* \otimes o(L_+^\#)$. This, finally, induces an isomorphism

$$i_u: |o(v_-)| \rightarrow |o(v_+)|.$$

It is not difficult to see that these isomorphisms are compatible with gluing, in the sense that when $u_1 \# u_2$ is a glued trajectory, $i_{u_1 \# u_2}$ is the composite of i_{u_1} and i_{u_2} (cf the discussion in Seidel [35, Section 12b]). This, of course, is what is needed for the proof that $\partial^2 = 0$.

3.8 A field theory

As observed by Seidel in [33], the Floer homology theory just described has an extension to an open-closed topological field theory in $(1 + 1)$ dimensions, coupled to singular fibrations. This unifies it with the Lagrangian matching theory. An ‘‘object’’ in the field theory is an LHF over a disjoint union of oriented circles. ‘‘Cobordisms’’ between objects are symplectic Morse–Bott fibrations over punctured matched surfaces, equipped with Lagrangian boundary conditions. To be precise, a cobordism from $\bigcup_{i=1}^p (Y_i, \pi_i, \sigma_i)$ to $\bigcup_{j=1}^q (Y'_j, \pi'_j, \sigma'_j)$ is defined by the following data:

- A compact surface S with matched boundary, together with a finite set $F = \{z_1, \dots, z_{p+q}\}$ of interior marked points, of which p are labelled as ‘‘incoming’’ and q as ‘‘outgoing’’
- Small disjoint coordinate discs $\xi_i: (\Delta, 0) \rightarrow (S, z_i)$ contained in $\text{int}(S)$
- A symplectic Morse Bott fibration $(E, \pi, \Omega, J_0, j_0)$ over $S \setminus F$, equipped with a Lagrangian boundary condition Q
- Diffeomorphisms ϕ_j fitting into commutative diagrams

$$\begin{array}{ccc} \xi_j^* E & \xrightarrow{\phi_j} & Y'_j \times [0, \infty) \\ \downarrow & & \downarrow \\ \Delta^* & \xrightarrow{r e^{i\theta} \mapsto (e^{i\theta}, -\log r)} & S^1 \times [0, \infty) \end{array}$$

at each of the outgoing punctures, where $\Phi_j^* \Omega = \text{pr}_1^* \sigma'_j$. At the incoming punctures, one has rather

$$\begin{array}{ccc} \xi_i^* E & \xrightarrow{\psi_i} & Y_i \times (-\infty, 0] \\ \downarrow & & \downarrow \\ \Delta^* & \xrightarrow{r e^{i\theta} \mapsto (e^{-i\theta}, \log r)} & S^1 \times (-\infty, 0]. \end{array}$$

There is then an obvious notion of composition of cobordisms.

Remark 3.17 In formulating the field theory, it is convenient to work over the Novikov ring Λ_k of an arbitrary base field k , rather than over $\Lambda_{\mathbb{Z}}$. This is itself a field, so taking homology commutes with \otimes and Hom .

Assuming that (E, Q) is either a fibrewise monotone fibration with a fibrewise monotone boundary condition, or a fibrewise 3–negative fibration with a fibrewise 2–negative boundary condition, there are now relative invariants

$$\Phi_{E,Q} \in \text{Hom}_{\Lambda_k} \left(\bigotimes_{i=1}^p HF_*(Y_i, \sigma_i), \bigotimes_{i=1}^q HF_*(Y'_i, \sigma'_i) \right),$$

subject to a gluing law (see Seidel [32; 33]; the gluing theory is done in [29]).²

Another general property of Floer homology is the Poincaré duality isomorphism:

$$HF_*(-Y, \sigma) \cong HF_*(Y, \sigma)^*,$$

where $(-Y, \sigma)$ refers to the fibration obtained by switching the orientation of the base circle. A change in the incoming/outgoing label of a point ζ_i dualises the corresponding group.

3.9 Quantum cap product

If one chooses further finite sets $I \subset S$ and $J \subset \partial S_+$, disjoint from one another and from the set of punctures F , the relative invariant $\Phi_{E,Q} \in \text{Hom}(HF_*(Y, \sigma), HF_*(Y', \sigma'))$ generalises to a homomorphism

$$(10) \quad \Phi_{E,Q;I,J}: \bigotimes_{s \in I} H_*(E_s; k) \otimes \bigotimes_{s' \in J} H_*(Q_{s'}; k) \rightarrow \text{Hom}_*(HF_*(Y, \sigma), HF_*(Y', \sigma')).$$

The homomorphism is the map on homology associated with a chain-level map $C\Phi$. The latter is defined using moduli spaces of index 0, finite action pseudo-holomorphic sections with boundary on Q , hitting chosen cycles in the marked fibres E_{s_a} and $Q_{s'_b}$. When S is a twice-punctured sphere, $\Phi_{E;\{s\},\emptyset}: H_*(X_s) \rightarrow \text{End}_* HF_*(X)$ is, by definition, the quantum cap product on Floer homology.

For simplicity, we state the next lemma in a version which concerns a single marked point (the reader will have no difficulty in generalising it to allow more points).

²Note that this is a *topological* field theory; the invariants are defined by choosing a complex structure on the base, and are then independent of it. We do not consider the more subtle question of constructing a *conformal* field theory.

Lemma 3.18 *Let $\Gamma: H_*(E_{s_0}) \rightarrow H_*(E_{s_1})$ be the isomorphism obtained by parallel transport of cycles over a path $\{s_t\}_{t \in [0,1]}$ in S . Then*

$$\Gamma \circ \Phi_{E, Q; \{s_0\}, \emptyset}(c) = \Phi_{E, Q; \{s_1\}, \emptyset}(\Gamma(c)).$$

Hence $\Phi_{E, Q; I, J}(c)$ depends only upon the image of c in $H_*(E)$. When $s \in \partial S_+$, one has

$$(11) \quad \Phi_{E, Q, \{s\}, \emptyset}(c) = \Phi_{E, Q, \emptyset, \{s\}}(i_* \circ \text{pr}_1^!(c)),$$

where $i: Q_s \hookrightarrow E_s \times E_{\tau(s)}$ is the inclusion. One also has

$$\Phi_{E, Q, \{\tau(s)\}, \emptyset}(c) = \Phi_{E, Q, \emptyset, \{s\}}(i_* \circ \text{pr}_2^!(c)).$$

Proof The first assertion is proved by considering the ends of 1–dimensional moduli spaces associated of sections passing through the smooth singular chain obtained by parallel transporting a cycle from one fibre to another.

For the assertions about cycles in boundary fibres, one should consider a cycle b in the product $E_s \times E_{\tau(s)}$. The fibre product of b with a moduli space of sections M —taken via the evaluation map of M into $E_s \times E_{\tau(s)}$ —is identified with the fibre product of $b \cap Q_s$, taken via the evaluation map of M into Q . Applying this to cycles b of the form $c \times E_{\tau(s)}$ or $E_s \times c$ gives the two formulae. \square

3.9.1 A reduction to homology Take two symplectic manifolds (M, ω) and $(\bar{M}, \bar{\omega})$ with a Lagrangian submanifold $L \subset (M, -\omega) \times (\bar{M}, \bar{\omega})$. Let $Y = S^1 \times M$, σ the pullback of ω . Let $S = \Delta \cup \Delta$ be the union of two discs (with their boundaries labelled as $+$ then $-$) and equip it with the obvious matching τ . Take the union of trivial fibrations, $E = (M \times \Delta) \cup (\bar{M} \times \Delta) \rightarrow S$. Puncture S at interior points ζ_+ in the first disc and ζ_- in the second. One obtains, by restriction, an LHF $E^* \rightarrow S \setminus \{\zeta_+, \zeta_-\}$ with Lagrangian matching condition Q .

On the other hand, one has canonical (Piunikhin–Salamon–Schwarz, or “PSS”) isomorphisms³

$$\begin{aligned} \Pi_M: H_*(M; \Lambda_k) &\rightarrow HF_*(M \times S^1, \text{pr}_1^* \omega), \\ \Pi_{\bar{M}}: H_*(\bar{M}; \Lambda_k) &\rightarrow HF_*(\bar{M} \times S^1, \text{pr}_1^* \omega). \end{aligned}$$

³See Piunikhin–Salamon–Schwarz [25], or alternatively, McDuff–Salamon [21, Chapter 12]. The formal argument has to be backed up by transversality and compactness arguments (unproblematic in the monotone or 3–negative cases), by Floer’s gluing theorem for pseudo-holomorphic trajectories, and by a Gromov–Witten style gluing theorem for holomorphic spheres [21, Chapter 10].

Note that these are valid with our orientation convention (the reason is almost the same as that which permitted the construction of the maps i_u in our construction of Floer homology: the PSS map and its inverse are defined by means of holomorphic planes asymptotic to horizontal sections and subject to incidence conditions; the moduli space of such planes, when it is 0-dimensional, is “canonically oriented relative to the cylindrical end”).

Proposition 3.19 *Suppose that either (i) M and \bar{M} are monotone and L a monotone Lagrangian, or (ii) M and \bar{M} are 3-negative and L 2-negative. Then one has*

$$(\Pi_{\bar{M}})^{-1} \circ \Phi_{E^*, Q} \circ \Pi_M(a) = \pm \text{pr}_{2*}([L] \cap \text{pr}_1^! \cdot) + \text{h.o.t.}$$

The letters *h.o.t.*, abbreviating “higher order terms”, designate a map which carries the subgroup $H_*(M; \Lambda^{\geq 0}) \subset H_*(M; \Lambda_k)$ into $H_*(\bar{M}; \Lambda^{> 0})$. Here $\Lambda^{> 0} \subset \Lambda^{\geq 0} \subset \Lambda_k$ are the additive subgroups consisting of series $\sum a(r)t^r$ with $r > 0$ or ≥ 0 .

In words: the leading order part of the map on Floer homology (with respect to the filtration by the action functional, recorded by the Novikov coefficients) is the classical map between homology groups associated with the correspondence, up to a sign which we do not check.

Proof The relative invariant $\Phi_{E, Q; \{\zeta_+, \zeta_-\}, \{s\}}$ for the “completed” fibration E may be formulated as a map

$$\Phi_{E, Q; \{\zeta_+, \zeta_-\}, \{s\}}: H_*(\bar{M}; \Lambda_k) \otimes H^*(L; \Lambda_k) \rightarrow H^*(\bar{M}; \Lambda_k).$$

We can decompose this map as $\Phi_{E, Q; \{\zeta_+, \zeta_-\}, \{s\}} = \sum_{\lambda \in \mathbb{R}} \Phi_\lambda t^\lambda$, where Φ_λ takes values in $H_*(\bar{M}; k)$. We have $\Phi_\lambda = 0$ for $\lambda > 0$. In fact, in the negative (rather than monotone) case, $\Phi_\lambda = 0$ for $\lambda \neq 0$ for dimension reasons.

The vertical isomorphisms in the diagram are obtained from the relative invariants of trivial fibrations over the once-punctured two-sphere, and commutativity of the diagram follows from the gluing property for the relative invariants. It therefore suffices to prove that the leading term Φ_0 of $\Phi(E, Q; \{\zeta_+, \zeta_-\}, \{s\})$ is $a \otimes c \mapsto \text{pr}_{2*}(c \cap \text{pr}_1^! a)$.

Trivialising the fibrations, $\text{sect}(E, Q)$ is identified with the space of maps

$$(\Delta, \partial\Delta) \rightarrow (M \times \bar{M}, L).$$

We use a complex structure of the form $-J \oplus \bar{J}$. The relative invariant is then defined using two smooth singular cycles, Z_\pm , lying over ζ_\pm , and a third cycle Z' in L . We then consider discs $\delta(\Delta, \partial\Delta) \rightarrow (M \times \bar{M}, L)$ such that $\delta(0)$ lies in the image of $Z_+ \times Z_-$, and $\delta(1)$ lies in the image of L . The leading term Φ_0 is computed using

discs δ of index zero. Because of our monotonicity/negativity assumptions, index zero implies area zero, hence these discs are constant. Moreover, regularity for constant discs is the same thing as transversality for the cycles $Z_+ \times Z_-$ and Z' . This reduces us to classical intersection theory, and so gives the result. \square

Remark 3.20 The proposition says that the relative invariant gives a homomorphism $H^*(L) \otimes QH^*(M) \rightarrow QH^*(\bar{M})$, where QH^* is quantum cohomology, deforming the homomorphism on classical cohomology groups induced by the correspondence L . These homomorphisms are little explored, even when $\bar{M} = \{\text{pt}\}$.

3.9.2 A speculative interlude Lagrangian matching invariants should have computational consequences in Floer homology. The following conjecture (which is the subject of work in progress) explains how this might work.

Suppose that (M, ω) and $(\bar{M}, \bar{\omega})$ are symplectic, $\rho: V \rightarrow \bar{M}$ an S^k -bundle with structure group $SO(k + 1)$, and $i: V \hookrightarrow M$ an embedding such that $i^*\omega = \rho^*\bar{\omega}$. Suppose that $\phi \in \text{Aut}(M, \omega)$ satisfies $\phi(V) = V$ and that $\phi|_V$ covers $\bar{\phi} \in \text{Aut}(\bar{M}, \bar{\omega})$. The cohomology of \bar{M} acts on $HF_*(\bar{\phi})$ by quantum cap product; in particular, the Euler class of V defines an endomorphism $e = e(V) \cap \cdot: HF_*(\bar{\phi}) \rightarrow HF_*(\bar{\phi})[-k - 1]$. We form its mapping cone—a chain complex $\text{cone}(e)$.

There is a natural cobordism, in the sense of our field theory, from $T(\phi)$ to $T(\bar{\phi})$, and hence (assuming \hat{V} monotone in $-M \times \bar{M}$, say) a Lagrangian matching invariant

$$\Phi: HF_*(\bar{\phi}) \rightarrow HF_*(\phi).$$

It turns out that, when \hat{V} has high enough minimal Maslov index (at least $k + 2$), the composite $\Phi \circ e$ is nullhomotopic, in an essentially canonical way. Using the nullhomotopy, one builds a map $a: H_* \text{cone}(e) \rightarrow HF_*(\phi)$. On the other hand, there is the fibred Dehn twist $\tau_V \in \text{Aut}(M, \omega)$ studied in I.2. By considering a symplectic Morse–Bott fibration over an annulus, with one critical value, one gets a map $b: HF_*(\phi) \rightarrow HF_*(\phi \circ \tau_V)$.

Conjecture 3.21 Under conditions which render a and b well-defined, they fit into a long exact sequence:

$$\begin{array}{ccccc}
 H_* \text{cone}(e) & \xrightarrow{a} & HF_*(\phi) & \xrightarrow{b} & HF_*(\phi \circ \tau_V) \\
 & & & \searrow & \swarrow \\
 & & & & [-1]
 \end{array}$$

In the case of our Lagrangian matching conditions for relative Hilbert schemes, this can be translated into Seiberg–Witten theory, where it appears to be a somewhat unorthodox variant of the surgery triangle in monopole Floer homology [17] in which all the three-manifolds involved are fibred. Exactness of this triangle can in fact be deduced from that of the usual one together with a connected sum formula.

4 Lagrangian matching invariants for broken fibrations

The previous section was devoid of examples. This section is devoted to a single class of examples: the relative Hilbert schemes associated with broken fibrations, and the Lagrangian matching conditions of Theorem 2.5.

4.1 Monotonicity for symmetric products and their vanishing cycles

4.1.1 The nonseparating case We consider $\text{Sym}^n(\Sigma)$, where Σ is connected of genus g , and $n \geq 2$.

- (1) The Hurewicz map $\pi_2(\text{Sym}^n(\Sigma)) \rightarrow H_2(\text{Sym}^n(\Sigma))$ has rank one. Here is a different proof to those I have seen in the literature in Bertram–Thaddeus [2] and Donaldson–Smith [6]: we deduce the result from Hopf’s classic theorem that, in general, $\text{coker}(\pi_2(X) \rightarrow H_2(X)) \cong H_2(\pi_1(X))$.⁴ Since $\pi_1(\text{Sym}^n(\Sigma)) = H^1(\Sigma)$, and the second integral homology of a free abelian group is its exterior square, we have $H_2(\pi_1) \cong \Lambda^2 H^1(\Sigma)$. But $H_2(\text{Sym}^n(\Sigma)) \cong \mathbb{Z} \oplus \Lambda^2 H^1(\Sigma)$. This tells us $\text{im}(\pi_2 \rightarrow H_2)$ must be isomorphic to \mathbb{Z} , and that its generator must be a primitive class (ie, not a nonunit multiple of another integer class).
- (2) One can easily identify a spherical homology class $h \in H_2(\text{Sym}^n(\Sigma))$: take any pencil of linearly equivalent effective divisors. One has $\langle \eta, h \rangle = 1$ (which shows that h is primitive) and $\langle \theta, h \rangle = 0$. (The notation was explained in the introduction to this paper.) Since $c_1(T \text{Sym}^n(\Sigma)) = (n+1-g)\eta - \theta$, the Chern number of h is $n+1-g$. Thus $\text{Sym}^n(\Sigma)$ is monotone if $n \geq g$, regardless of the choice of symplectic form.
- (3) Let $\bar{\Sigma}$ be the result of surgery on a nonseparating loop $L \subset \Sigma$. Consider $\text{Sym}^n(\Sigma) \times \text{Sym}^{n-1}(\bar{\Sigma})$ with a symplectic form $-\omega \oplus \bar{\omega}$ that it is a sum of Kähler forms, with a sign reversal on the first factor, in cohomology classes $[\omega] = \eta_\Sigma + \lambda\theta_\Sigma$ and $[\bar{\omega}] = \eta_{\bar{\Sigma}} + \lambda\theta_{\bar{\Sigma}}$. This symplectic manifold is also monotone. The two generators for $\text{im}(\pi_2 \rightarrow H_2)$ (chosen with appropriate signs) each have Chern number $n+1-g$ and area 1.

⁴An isomorphism is induced by the map $X \rightarrow B\pi_1(X)$ classifying the universal cover. When $X = \text{Sym}^n(\Sigma)$ that map can be taken to be the Abel–Jacobi map to the Jacobian torus.

- (4) Now consider the Lagrangian vanishing cycle

$$\widehat{V}_L \subset (\text{Sym}^n(\Sigma) \times \text{Sym}^{n-1}(\overline{\Sigma}), -\omega \oplus \overline{\omega})$$

associated with the loop $L \subset \Sigma$ (Part I: Theorem A). It was shown in Part I that every disc in $\pi_2(\text{Sym}^n(\Sigma) \times \text{Sym}^{n-1}(\overline{\Sigma}), \widehat{V}_L)$ lifts to a sphere in the product (Lemma 3.21). Hence the positive area discs D for which $|\mu_{\widehat{V}_L}(D)|$ is least all have

$$\mu_{\widehat{V}_L}(D) = 2(n+1-g) = 2((n-1)+1-(g-1)),$$

and area 1. Thus \widehat{V}_L is monotone as a Lagrangian when $n \geq g$.

- (5) In the same situation, one checks that \widehat{V}_L is 2-negative when $n \leq (2g(\Sigma)-1)/4$. This condition does not imply 3-negativity of $\text{Sym}^n(\Sigma) \times \text{Sym}^{n-1}(\overline{\Sigma})$, but it does imply 3-negativity of the two factors.

4.1.2 The separating case The case where Σ is connected but $\overline{\Sigma}$ disconnected, with components Σ_1 and Σ_2 , is more delicate. One has

$$\text{Sym}^{n-1}(\overline{\Sigma}) = \bigcup_{i=0}^{n-1} \text{Sym}^i(\Sigma_1) \times \text{Sym}^{n-1-i}(\Sigma_2).$$

For brevity, we write M for $\text{Sym}^n(\Sigma)$, N for $\text{Sym}^{n-1}(\overline{\Sigma})$ and N_i for the component $\text{Sym}^i(\Sigma_1) \times \text{Sym}^{n-1-i}(\Sigma_2)$ of N . We put a symplectic form on $M \times N$ which, as before, is of shape $\omega \oplus -\overline{\omega}$ for Kähler forms ω on M and $-\overline{\omega}$ on N . Since we want to find Lagrangian correspondences, we should suppose that ω represents a class $\eta_\Sigma + \lambda\theta_\Sigma$ for some $\lambda > 0$, and that $\overline{\omega}$ represents $e + \lambda\theta_{\overline{\Sigma}}$, where e restricts to the component N_i as

$$\mu_i\eta_{\Sigma_1} + (1-\mu_i)\eta_{\Sigma_2}$$

for some $\mu_i \in (0, 1)$. When this is true, Theorem A from Part I gives us a correspondence $\widehat{V}_{L,k} \subset M \times N_i$ for each $i \in \{0, \dots, n-1\}$.

The factors $\text{Sym}^i(\Sigma_1)$ and $\text{Sym}^{n-1-i}(\Sigma_2)$ are all 3-negative when

$$n-1 \leq \min_{j=1,2} (g(\Sigma_j) - 1)/2;$$

under this hypothesis, we will have adequate control so far as bubbling of holomorphic spheres is concerned.

The awkward point is that it is no longer true in general that every disc with boundary in \widehat{V}_L lifts to a sphere: when Σ_1 and Σ_2 have positive genus, there is a cyclic group's worth of nonspherical discs (Part I: Lemma 3.21) for each component of \widehat{V}_L . We have to determine the Maslov indices for the generators for these cyclic groups.

Lemma 4.1 Assume that $g(\Sigma_1) > 0$, $g(\Sigma_2) > 0$ and $n > 1$. Then the group $\pi_2(M \times N_k, \widehat{V}_{L,k})$ contains primitive, nonspherical classes u_k such that

$$\mu_{\widehat{V}_{L,k}}(u_k) = k + 1 - 2g(\Sigma_1).$$

The class u_k generates $\pi_2(M \times N_k, \widehat{V}_{L,k})$ modulo the image of $\pi_2(M \times N_k)$.

Before we can prove this, we need to obtain a slightly more precise picture of $\widehat{V}_{L,k}$. We will achieve this by considering the model considered in Part I, Section 3.2.1, in which Σ , Σ_1 and Σ_2 are all 2-spheres. By interpreting the n -th symmetric product of \mathbb{P}^1 as \mathbb{P}^n , we obtained an explicit equation for $\widehat{V}_{L,k}$ (the explicit model was denoted by $\mathcal{L}_{n,k}$).

Lemma 4.2 One of the S^1 -fibres of $\mathcal{L}_{n,k} \subset \text{Sym}^n(\mathbb{P}^1)$ consists of n -tuples of the form $[e^{i\theta}, e^{i\theta}\eta, \dots, e^{i\theta}\eta^{k-1}; \infty, \dots, \infty]$, where $\eta = e^{2\pi i/k}$. Indeed, under the quotient map $\mathcal{L}_{n,k} \rightarrow \text{Sym}^{k-1}(\mathbb{P}^1) \times \text{Sym}^{n-k}(\mathbb{P}^1)$, this circle maps to the point

$$((k-1)\infty, (n-k)0).$$

Proof Our notation is from Part I, Section 3.2.1. The S^1 -action $a_{n,k}$ preserves this circle, so it suffices to take $\theta = 0$. Our identification $\text{Sym}^n(\mathbb{P}^1) = \mathbb{P}^n$ sends an n -tuple $\mathbf{x} = [x_1, \dots, x_k; \infty, \dots, \infty]$, where all the x_j lie in $\mathbb{C} \subset \mathbb{P}^1$, to the point $(\sigma_k(\mathbf{x}) : \dots : \sigma_1(\mathbf{x}) : 1 : 0 : \dots : 0)$ ($n-k$ zeros). Here σ_j stands for the j -th elementary symmetric polynomial, and our convention is such that $\sigma_1(\mathbf{x}) = -\sum x_i$ and $\sigma_k(\mathbf{x}) = (-1)^k \prod x_i$. The n -tuple $\mathbf{x} = [1, \eta, \dots, \eta^{k-1}; \infty, \dots, \infty]$ is mapped to $(1 : 0 : \dots : 0 : 1 : 0 : \dots : 0)$, ie to the point $(z_0 : \dots : z_n)$ where $z_0 = z_k = 1$ and $z_j = 0$ otherwise. To see this, note that $\sigma_k(1, \eta, \dots, \eta^{k-1}) = 1$ (clear) but that $\sigma_j(1, \eta, \dots, \eta^{k-1}) = 0$ for $j < k$, since by homogeneity of σ_j ,

$$\eta^j \sigma_j(1, \eta, \dots, \eta^{k-1}) = \sigma_j(\eta, \eta^2, \dots, \eta^k) = \sigma_j(1, \eta, \dots, \eta^{k-1}).$$

This point $(z_0 : \dots : z_n)$ certainly satisfies the defining equation for $\mathcal{L}_{n,k}$, which is $\sum_{j=0}^{k-1} |z_j|^2 = \sum_{j=k}^n |z_j|^2$. The reduction map (dividing out the S^1 -action) carries it to $((1 : \dots : 0 : 1), (1 : 0 : \dots : 0)) \in \mathbb{P}^{k-1} \times \mathbb{P}^{n-k}$, which corresponds to the divisor $((k-1)(1:0), (n-k)(0:1)) \in \text{Sym}^{k-1}(\mathbb{P}^1) \times \text{Sym}^{n-k}(\mathbb{P}^{n-k})$. \square

Proof of Lemma 4.1 By exchanging the roles of Σ_1 and Σ_2 , and of k and $n-1-k$, we see that it is sufficient to assume that $k > 1$.

Compactify each of the two components of $\Sigma \setminus \gamma$ to a Riemann surface-with-boundary $\widetilde{\Sigma}_i$, $i = 1, 2$. Thus there is a map $\widetilde{\Sigma}_i \rightarrow \Sigma_i$ collapsing the boundary to a point. Consider a holomorphic branched covering $b: \widetilde{\Sigma}_1 \rightarrow D$ over the closed unit disc, of degree k ,

restricting to an unramified cyclic covering $\partial\tilde{\Sigma} \rightarrow \partial D$ (also of degree k). (We are free to adjust the complex structure on Σ to make sure that the covering exists.) Fixing a point $z \in \Sigma_2$, we obtain a holomorphic map

$$b_z: D \rightarrow \text{Sym}^n(\Sigma), \quad x \mapsto b^{-1}(x) + (n-k)z.$$

Now, the coisotropic hypersurface $V_L \subset \text{Sym}^n(\Sigma)$ depends on the choice of a Kähler form on a relative Hilbert scheme, so it does not make sense to ask whether the boundary of b_z lies on V_L . However, if we allow ourselves to move V_L by totally real isotopies by the method of Part I, Section 3.6, adjusting the (nonclosed) “good two-form” (Part I, Section 2.3.2) defining the vanishing cycle, we may assume that this is indeed so. Furthermore we may arrange that the boundary curve of this holomorphic curve circumnavigates the fibre of the projection $V_L \rightarrow N_k$ over a point (P, Q) . Indeed, by the last lemma, this behaviour is exactly what happens in the genus-zero model, which may be transplanted into higher-genus situations as in Part I, Section 3.6. (Because of the homotopical nature of Maslov indices, it is perfectly acceptable to sacrifice closedness of the nondegenerate two-forms.)

We can then promote b_z to a holomorphic map

$$B_z: D \rightarrow M \times N_k, \quad x \mapsto (b_x(z), P, Q).$$

with boundary on \hat{V}_L . We can see that $[B_z]$ is primitive in $\pi_2(M \times N_k, \hat{V}_{L,k})$ simply by observing that its boundary is primitive in $\pi_1(\hat{V}_L)$ (this requires $g(\Sigma_1) > 0$).

To calculate the Maslov index of B_z , we shall exhibit a trivialisation (Y_1, \dots, Y_{2n-1}) of $(B_z|_{\partial D})^*T\hat{V}_{L,i}$. We then obtain a trivialisation $(Y_1, \dots, Y_{2n-1}; JY_1, \dots, JY_{2n-1})$ for $(B_z|_{\partial D})^*T(M \times N_i)$ via a complex structure J , and hence an orientation o for this last bundle. We then extend this orientation to a section of $\Lambda^{\text{top}}B_z^*T(M \times N_k)$, and count its zeros with signs; the Maslov index is this count.

A warm-up exercise, left to the reader, is to derive the Chern number $n + 1 - g$ for the generating sphere in $\pi_2(\text{Sym}^n(\Sigma))$ by thinking of this sphere as a degree n branched cover of S^2 and applying Riemann–Hurwitz. This exercise is useful in getting the local contributions to the Maslov index correct.

Start by trivialising $T(\partial D)$ via an anticlockwise-pointing vector field Z . Let $a_z = b_z|_{\partial D}$ and $A_z = B_z|_{\partial D}$. Then $a_{z*}Z$ is a nonvanishing section of $a_z^*TV_L$. The required trivialisation for $A_z^*T\hat{V}_L$ is obtained by putting $Y_i = X_i$ for $i = 1, \dots, 2n-2$, where (X_1, \dots, X_{2n-2}) is a positively oriented trivialisation of $T_{(P,Q)}(N_k)$, and $Y_{2n-1} = a_{z*}Z$. Now, Z extends to a vector field \tilde{Z} on D with a single, nondegenerate zero s of sign $+1$, which we may assume does not coincide with any of the critical

values of the covering b . We may extend the fields Y_i to fields \tilde{Y}_i over D by putting $\tilde{Y}_i = X_i$ for $i < 2n - 1$ and $\tilde{Y}_{2n-1} = b_{z*}\tilde{Z}$.

The field \tilde{Y}_{2n-1} has a k -fold zero at s , which contributes $2k$ to the Maslov index. The other contributions come from the critical points of b ; each adds -1 to the Maslov index, and there are $k - 1 + 2g(\Sigma_1)$ of them, by Riemann–Hurwitz. This gives us a Maslov index of

$$\mu_{\hat{V}_L}(B_z) = 2k - (k - 1 + 2g(\Sigma_1)) = k + 1 - 2g(\Sigma_1).$$

Putting $u_k = [B_z]$, we obtain the Maslov index claimed. \square

Now, we wish $\hat{V}_{L,k}$ to be 2-negative. We assume that $n \leq \frac{1}{2} \min(g(\Sigma_1), g(\Sigma_2))$. It is not clear to the author whether 2-negativity can be achieved simultaneously for all $k \in \{0, \dots, n - 1\}$; rather, we will adjust the available parameters so that it is 2-negative for any particular k . We first choose the coefficients μ_k so as to ensure that the product manifolds $M \times N_k$ are all negatively monotone, ie, that the areas of the two generators for π_2 are proportional to their respective Chern numbers (with a negative constant of proportionality, λ). We have to arrange that the area of the disc u_k is in the ratio 2λ to its Maslov index. Now, $\hat{V}_{L,k}$ depends on the Hamiltonian isotopy class of $L \subset \Sigma$. The areas of the disks that it bounds are sensitive to this Hamiltonian isotopy class: in fact, under a *Lagrangian* isotopy of L , the area of the resulting u_k will change proportionally to the flux of the isotopy. Hence by moving L we can ensure that the area of u_k is 2λ times its Maslov index.

Once we have arranged this negative monotonicity, we can apply Lemma 4.1 to see that 2-negativity holds provided that $n \leq \frac{2}{3} \min(g(\Sigma_1), g(\Sigma_2))$, which is certainly true when $n \leq \frac{1}{2} \min(g(\Sigma_1), g(\Sigma_2))$.

Our conclusion is expressed by the following lemma.

Lemma 4.3 *When $n \leq \frac{1}{2} \min(g(\Sigma_1), g(\Sigma_2))$, all the symplectic manifolds M and N_k are 3-negative. For each $k \in \{0, \dots, n - 1\}$, after an isotopy of L inside Σ , the corresponding Lagrangian $\hat{V}_{L,k} \subset M \times N_k$ is 2-negative.*

4.2 Defining the invariants for broken fibrations on closed four-manifolds

We are, at last, ready to give the definition of the Lagrangian matching invariants for broken fibrations on closed four-manifolds.

The invariants are indexed by Spin^c -structures. For now we explain in schematic form the connection between moduli spaces of sections and Spin^c -structures. The details will be given afterwards.

Let (X, π) be a broken fibration over a closed surface (X connected) and let $Z \subset X^{\text{crit}}$ be the 1–dimensional part of the set of critical points.

(i) The “Taubes map” τ_X is a bijection

$$\text{Spin}^c(X) \rightarrow \delta^{-1}([Z]) \subset H_2(X, Z),$$

where $\delta: H_2(X, Z) \rightarrow H_1(Z)$ is the boundary homomorphism, and Z is oriented by a vector field v such that $i\pi_*(v)$ points into the side of $\pi(Z) \subset S$ on which the Euler characteristic is higher (here i is the complex structure on S). The map τ_X arises from the canonical Spin^c –structure $\mathfrak{s}_{\text{can}}$ on the almost complex manifold $X \setminus Z$. It is characterised by the relation

$$(12) \quad \tau_X(\mathfrak{s}) = \beta \quad \text{if} \quad \mathfrak{s}|_{(X \setminus Z)} = \text{PD}(\beta) \cdot \mathfrak{s}_{\text{can}}.$$

In the context of near-symplectic manifolds, the map τ_X plays an important role in Taubes’ programme. In [38] it is proved that $\text{im}(\tau) \subset \delta^{-1}([Z])$; it is then easy to see that τ_X is bijective.

(ii) Construct $X^{[v]}$ as in (2), where $v: S^{\text{reg}} \rightarrow \mathbb{Z}_{\geq 0}$ satisfies $2v(s) + \chi(X_s) = 2d$, with d constant; and its Lagrangian matching condition \mathcal{Q} . We shall see presently how to build a map

$$\alpha_v: \pi_0 \text{sect}(X^{[v]}, \mathcal{Q}) \rightarrow \delta^{-1}([Z]) \subset H_2(X, Z; \mathbb{Z}).$$

If $\beta \in \text{im}(\alpha_v)$ then one has $\langle \beta, X_s \rangle = v(s)$. Also, $\langle \beta, [F] \rangle \geq 0$ for every homology class $[F] \in H_2(X \setminus Z; \mathbb{Z})$ represented by a component of a regular or nodal fibre.

Define the map

$$(13) \quad \sigma_v = \tau_X^{-1} \circ \alpha_v: \pi_0 \text{sect}(X^{[v]}, \mathcal{Q}) \rightarrow \text{Spin}^c(X).$$

We will prove that σ_v is injective, and, when $\inf(v) \geq 2$, describe its image.

Now take $\beta = \tau_X^{-1}(\mathfrak{s}) \in H_2(X, Z; \mathbb{Z})$. Let us recall a definition from I.1:

Definition 4.4 Suppose that the fibres of π are all connected. The Spin^c –structure \mathfrak{s} is then *admissible* if (i) for every homology class $[F] \in H_2(X \setminus Z; \mathbb{Z})$ represented by F , a regular fibre or an irreducible component of a nodal fibre, one has $\langle c_1(\mathfrak{s}), F \rangle \geq \chi(F)$ (equivalently, $\langle \beta, [F] \rangle \geq 0$); and (ii) one of the following two conditions holds: for each regular fibre X_s one has *either*

- (1) $\langle c_1(\mathfrak{s}), [X_s] \rangle \leq \chi(X_s)/2$ (equivalently, $\langle \beta, [X_s] \rangle \leq -\chi(X_s)/4$); or
- (2) $\langle c_1(\mathfrak{s}), [X_s] \rangle > 0$ (equivalently, $\langle \beta, X_s \rangle \geq g(X_s)$).

By the results of Section 4.1.1, condition (2) implies that the fibres $\text{Sym}^{\nu(s)}(X_s)$ are monotone, and that \mathcal{Q} is a fibrewise-monotone matching condition. Condition (1) implies that the fibres $\text{Sym}^{\nu(s)}(X_s)$ are 3–negative, and that the fibres of the Lagrangian matching condition are 2–negative. These are the conditions under which we are able to construct invariants.

Definition 4.5 In a broken fibration some of whose fibres have two connected components (but none more than two), we say that \mathfrak{s} is admissible when

$$\chi(F) \leq \langle c_1(\mathfrak{s}), [F] \rangle \leq \chi(F)/2$$

for each fibre-component F .

By the results of Section 4.1.2 (more particularly, Lemma 4.3) admissibility guarantees 3–negativity of the relevant symmetric products, and—after suitable adjustments are made—2–negativity of the Lagrangian correspondences.

Let $I_\pi \subset H_1(X; \mathbb{Z})$ be the subgroup of classes supported on a fibre of π . Let $I_\pi^* = \text{Hom}(I_\pi, \mathbb{Z})$. Write

$$\mathbb{A}(X, \pi) = \mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* I_\pi^*,$$

and make it a graded ring by declaring U to have degree 2. An element l of $\mathbb{A}(X, \pi)$ is a finite sum $\sum U^n \otimes l_n$. It can be thought of equally as a homomorphism $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* I_\pi^* \rightarrow \mathbb{Z}$ (sending $U^m \otimes \lambda$ to $l_m(\lambda)$) and we write $l(m)$ to denote its value on $m \in \mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* I_\pi^*$.

When $I_\pi = H_1(X; \mathbb{Z})$, we write $\mathbb{A}(X)$ instead of $\mathbb{A}(X, \pi)$. The reader can verify that this is always true when the base is S^2 .

Remark 4.6 At this point I should confess to a technical inconsistency with the announcements made in the introduction to Part I, arising because of a careless oversight. (This refers to the published version; the ArXiv version has been corrected.) In Part I, Section 1, $\mathbb{A}(X, \pi)$ was defined using the dual group of $K_\pi = \ker(\pi_*: H_1(X) \rightarrow H_1(S))$ in place of I_π^* . In general, K_π will be larger than I_π , and I have not established how to work with the larger group. I apologise for this change.

Definition 4.7 Fix an orientation for each orientable attaching surface Q_i for the broken fibration (X, π) . The *Lagrangian matching invariant*

$$\mathcal{L}_{(X, \pi)}(\mathfrak{s}) \in \mathbb{A}(X, \pi)$$

is a homogeneous element of degree $d(\mathfrak{s})$, where

$$(14) \quad d(\mathfrak{s}) = \frac{1}{4} \left(c_1(\mathfrak{s})^2 - 2\chi(X) - 3\sigma(X) \right),$$

characterised as follows. Let $m = U^a \otimes l_1 \wedge \cdots \wedge l_b$ be a monomial of degree $d(\mathfrak{s})$. Represent l_i by a loop $\gamma_i \subset X_{s_i}$ on a regular fibre. The points s_i should all be distinct. We associate with each γ_i a smooth codimension–1 cycle in a symmetric product:

$$\delta_{\gamma_i} = \gamma_i + \text{Sym}^{\nu(s)-1}(X_s) \subset \text{Sym}^{\nu(s)}(X_s).$$

Choose a points $x_j \in X$, lying in distinct regular fibres of π , disjoint from the fibres X_{s_i} . With each of these we associate the codimension–2 cycle

$$\delta_{x_j} = x_j + \text{Sym}^{\nu(s)-1}(X_{\pi(x_j)}) \subset \text{Sym}^{\nu(s)}(X_{\pi(s)}).$$

The symplectic Morse–Bott fibration $X^{[v]} \rightarrow S'$, and its Lagrangian matching condition \mathcal{Q} , give rise to a homomorphism

$$\Phi = \Phi_{X^{[v], \mathcal{Q}; \{s_i\} \cup \{\pi(x_j)\}, \emptyset}: \bigotimes_{i=1}^{a+b} H^*(\text{Sym}^{\nu(s_i)}(X_{s_i}); \mathbb{Z}) \rightarrow \mathbb{Z},$$

as in (9). Set

$$\mathcal{L}_{(X, \pi)}(\mathfrak{s})(U^a \otimes [\gamma_1] \wedge \cdots \wedge [\gamma_b]) = \Phi([\delta_{x_1}] \otimes \cdots \otimes [\delta_{x_a}] \otimes [\delta_{\gamma_1}] \otimes \cdots \otimes [\delta_{\gamma_b}])$$

(the cohomology classes here have degrees 2 and 1, and are Poincaré dual to the cycles $\delta_{x_j}, \delta_{\gamma_i}$).

The definition involves an orientation for the moduli space of pseudo-holomorphic sections. This is defined via the relative pin structure for \mathcal{Q} constructed by the general mechanism of Lemma 3.14.

We also need to orient the components \mathcal{Q}_i of \mathcal{Q} for which an orientation exists. This is equivalent to orienting the attaching tori $Q_i \subset X$ (the tori which degenerate to critical circles): indeed, \mathcal{Q}_i is an S^1 –family of circle-bundles over symplectic manifolds. One orients these circle-bundles by orienting a fibre then orienting the base. The base has its symplectic orientation, and one of the fibres may be taken to be a copy $\gamma \times (n-1)_X$ of the attaching circle $\gamma \subset \Sigma$. Finally, one orients \mathcal{Q}_i as

$$\begin{aligned} & \text{(canonical orientation of } S^1) \\ & \otimes \text{(orientation of fibre of } \widehat{V}_\gamma \rightarrow \text{Sym}^{n-1}(\overline{\Sigma})) \\ & \otimes \text{(symplectic orientation of } \text{Sym}^{n-1}(\overline{\Sigma})). \end{aligned}$$

That the degree of $\mathcal{L}_{(X, \pi)}(\mathfrak{s})$ is $d(\mathfrak{s})$ is a nontrivial matter: it is a direct consequence of our index formula, Theorem D. Other than that, the well-definedness of $\mathcal{L}_{(X, \pi)}(\mathfrak{s})$ is an easy consequence of the general theory we have set up. The following points are worth remarking:

- We must check independence from the choices of cycles δ_{x_j} and δ_{γ_i} . One can certainly move the point x_j inside X' : a path $x_j(t)$ gives rise to a homotopy of cycles $\text{Sym}^{\nu(x_j)-1}(X_{\pi(x_j(t))}) \subset \text{Sym}^{n(x_j)}(X_{\pi(x_j(t))})$, and the parametrised moduli space is a compact one-manifold with boundary. It is also legitimate to move x_j between different components of X' , because of (10). Similarly, one can replace γ_j by a homotopic loop in a nearby fibre, or by a homologous point loop in the same fibre, and one can also pass between different components of X' .
- We must show that Φ is symmetric in the δ_{x_j} -factors and antisymmetric in the δ_{γ_i} -factors. This holds because a pair of cycles in distinct fibres can be replaced by the intersection of two cycles in a single fibre; thus symmetry and antisymmetry hold because the relevant cycles have, respectively, even and odd codimension.

4.3 Sections of \mathcal{Q} versus Spin^c -structures

We now explain the map α_ν involved in the definition of σ_ν given in equation (13), starting with a prototypical case.

We use the same notation as in Part I, Section 4. So, $\pi: Y \rightarrow S^1$ and $\bar{\pi}: \bar{Y} \rightarrow S^1$ are fibred three-manifolds, and $(X_{\text{br}}, \pi_{\text{br}})$ an elementary broken fibration over an annulus $A = \{z \in \mathbb{C} : 1/2 \leq |z| \leq 2\}$ realising a cobordism between them:

$$Y = \pi_{\text{br}}^{-1}(\{|z| = 1/2\}), \quad \bar{Y} = \pi_{\text{br}}^{-1}(\{|z| = 2\}).$$

The fibre $\Sigma = \pi_{\text{br}}^{-1}(1/2)$ is connected of genus g , whilst $\bar{\Sigma} := \pi_{\text{br}}^{-1}(2)$ is either connected of genus $g - 1$, or else disconnected with components of genera g_1 and $g - g_1$.

We write $(Y^{[n]}, \pi^{[n]})$ for the relative symmetric product $\text{Sym}_{S^1}^n(Y) \rightarrow S^1$, and $(\bar{Y}^{[n-1]}, \bar{\pi}^{[n-1]})$ for $\text{Sym}_{S^1}^{n-1}(\bar{Y}) \rightarrow S^1$. As usual, after choosing complex structures on $T^\vee Y$ and $T^\vee \bar{Y}$ these become differentiable families of complex manifolds.

The first step is to look more carefully at the Taubes map $\tau_{X_{\text{br}}}$. Let $X = X_{\text{br}}$ and let $Z \subset X$ be the circle of critical points. Then we have a commutative diagram

$$\begin{array}{ccc} \text{Spin}^c(X) & \xrightarrow{\text{restr.}} & \text{Spin}^c(\bar{Y}) \\ \tau_X \downarrow & & \cong \downarrow \tau_{\bar{Y}} \\ H_2(X, \partial X \cup Z; \mathbb{Z}) & \xrightarrow{\delta} & H_1(\bar{Y}; \mathbb{Z}) \end{array}$$

where the following holds:

- The map δ is the obvious boundary homomorphism. A straightforward computation verifies that it is surjective.
- τ_X is the Taubes map, defined by the condition (12). In similar fashion, $\tau_{\bar{Y}}$ is defined by the relation

$$\mathfrak{t} = \text{PD}(\tau_{\bar{Y}}(\mathfrak{t})) \cdot \mathfrak{t}_{\text{can}},$$

where $\mathfrak{t}_{\text{can}} \in \text{Spin}^c(\bar{Y})$ is the canonical Spin^c -structure associated with the vertical two-plane field on \bar{Y} (it is the restriction to \bar{Y} of $\mathfrak{s}_{\text{can}}$).

For each integer d , let

$$\text{Spin}^c(\bar{Y})_d = \{\mathfrak{t} \in \text{Spin}^c(\bar{Y}) : \langle c_1(\mathfrak{t}), [\bar{\Sigma}] \rangle = 2d\},$$

and let $\text{Spin}^c(X)_d$ be the preimage of $\text{Spin}^c(\bar{Y})_d$ under the restriction map.

What is the image of $\text{Spin}^c(X)_d$ (resp. $\text{Spin}^c(\bar{Y})_d$) under $\tau(X)$ (resp. $\tau_{\bar{Y}}$)? Write

$$2d = \chi(\Sigma) + 2n = \chi(\bar{\Sigma}) + 2(n-1).$$

Then $\tau_X(\text{Spin}^c(X)_d)$ is the affine subgroup

$$H_2^{(n)} = \{\beta \in H_2(X, Z \cup \partial X) : \partial\beta = [Z] \bmod H_1(\partial X; \mathbb{Z}) \text{ and } \beta \cdot [\Sigma] = n\}.$$

This follows immediately from the fact that $\mathfrak{s}_{\text{can}}$ has degree $\chi(\bar{\Sigma})$ over $\bar{\Sigma}$. Similarly, $\tau_{\bar{Y}}(\text{Spin}^c(\bar{Y})_d)$ is the affine subgroup

$$H_1(\bar{Y}; \mathbb{Z})_n = \{\gamma \in H_1(\bar{Y}; \mathbb{Z}) : \gamma \cdot [\bar{\Sigma}] = n\}.$$

If $\bar{\Sigma}$ is disconnected, with components $\bar{\Sigma}_1$ and $\bar{\Sigma}_2$, it becomes necessary to refine the set $\text{Spin}^c(X)_d$. Write

$$\text{Spin}^c(\bar{Y})_d = \bigcup_{d_1+d_2=d} \text{Spin}^c(\bar{Y})_{d_1, d_2},$$

where $\mathfrak{t} \in \text{Spin}^c(\bar{Y})_{d_1, d_2}$ if $\langle c_1(\mathfrak{t}), [\bar{\Sigma}_i] \rangle = 2d_i$. Let $\text{Spin}^c(X)_{d_1, d_2}$ be the preimage of $\text{Spin}^c(\bar{Y})_{d_1, d_2}$ under restriction. Then $\tau_X(\text{Spin}^c(X)_{d_1, d_2}) = H_2^{(n_1, n_2)}$, where $n_i = d_i - \chi(\Sigma)/2$, and $H_2^{(n_1, n_2)}$ is the obvious refinement of $H_2^{(n)}$.

Now we bring in the space of sections of the Lagrangian matching condition \mathcal{Q} .

Proposition 4.8 *Assume $\bar{\Sigma}$ is connected. Then there is a unique map*

$$\alpha: \pi_0 \text{sect}(\mathcal{Q}) \rightarrow H_2^{(n)}$$

which makes the diagram

$$\begin{array}{ccc}
 \pi_0 \text{sect}(\mathcal{Q}) & \overset{\alpha}{\dashrightarrow} & H_2^{(n)} \\
 \downarrow & & \downarrow \partial \\
 \pi_0 \text{sect}(Y^{[n]}) \times \pi_0 \text{sect}(\bar{Y}^{[n-1]}) & \longrightarrow & H_1(Y; \mathbb{Z})_n \oplus H_1(\bar{Y}; \mathbb{Z})_{n-1}
 \end{array}$$

commute. Here the lower horizontal arrow is the natural “cycle map” sending homotopy classes of sections of symmetric product bundles to the homology classes of the cycles swept out by the points in their support. The map α is surjective, and if $n - 1 \geq 2$ also injective. Hence the composite

$$\sigma := \tau_X^{-1} \circ \alpha: \pi_0 \text{sect}(\mathcal{Q}) \rightarrow \text{Spin}^c(X)_d$$

is surjective, and bijective if $n - 1 \geq 2$.

Proof The elements of $H_2^{(n)}$ are represented by surfaces with boundary, of form $C_1 \cup \dots \cup C_{r-1} \cup C'$, where the C_i are cylinders which are sections of the broken fibration $X \rightarrow A$, and C' is a cylinder with $\partial C' = Z \cup \gamma$ for some circle $\gamma \subset Y$. Recall the attaching surface $Q \subset Y$ (torus of Klein bottle) which collapses to Z . The set $\pi_0 \text{sect}(\mathcal{Q})$ maps injectively to $H_1(Q; \mathbb{Z})$, and in this way it may be considered as an affine abelian group, isomorphic to \mathbb{Z} or $\mathbb{Z}/2$ according to whether or not Q is orientable. There is an “exact sequence”

$$0 \rightarrow \pi_0 \text{sect}(\mathcal{Q}) \rightarrow H_2^{(n)} \rightarrow H_1(\bar{Y}; \mathbb{Z})_{n-1} \rightarrow 0$$

(we use quotation marks because these are only affine abelian groups).

Notice that $H_2(X, \partial X \cup Z; \mathbb{Z})$ injects into $H_1(\partial X; \mathbb{Z})$ (and even into $H_1(Y; \mathbb{Z})$). Thus $H_2^{(n)}$ defines a correspondence between $H_1(Y; \mathbb{Z})_n$ and $H_1(\bar{Y}; \mathbb{Z})_{n-1}$. Because of this, α —if it exists—is fully determined by the commutative square in the statement of the theorem. We must show that the image of any section of \mathcal{Q} in $H_1(Y; \mathbb{Z})_n \oplus H_1(\bar{Y}; \mathbb{Z})_{n-1}$ lies in $\partial(H_2^{(n)})$.

To do so, we study $\pi_0 \text{sect}(\bar{Y}^{[n-1]})$, and relate it to $H_1(\bar{Y}; \mathbb{Z})_{n-1}$. In general, given a self-diffeomorphism $\psi \in \text{Diff}(M)$ of a connected manifold M , once one chooses a reference section of the mapping torus $T(\psi)$, the set of homotopy classes of sections, $\pi_0 \text{sect} T(\psi)$, is described by a natural commutative diagram

$$\begin{array}{ccc}
 \pi_0 \text{sect} T(\psi) & \xrightarrow{\cong} & \pi_1(M) / \sim \\
 \downarrow & & \downarrow \\
 H_1(T(\psi); \mathbb{Z}) & \longleftarrow & H_1(M) / \text{im}(1 - \psi_*)
 \end{array}$$

where \sim is the relation of “ ψ -twisted conjugacy” in $\pi_1(M)$:

$$\gamma \sim \beta \cdot \gamma \cdot (\psi \circ \beta)^{-1}$$

(concatenation of paths is read from left to right). This applies to \bar{Y} , the mapping torus of $\bar{\phi} \in \text{Diff}^+(\bar{\Sigma})$. In that case, the vertical map on the right (the abelianisation map) is always surjective, but when the genus of the fibre is > 1 it is not injective.

When we consider the symmetric products $\bar{Y}^{[n-1]} = \text{Sym}_{S^1}^{n-1}(\bar{Y})$, we can slightly modify the targets of the vertical arrows in the commutative diagram:

$$\begin{array}{ccc} \pi_0 \text{sect}(\bar{Y}^{[n-1]}) & \xrightarrow{\cong} & \pi_1(\text{Sym}^{n-1}(\bar{\Sigma})) / \sim \\ \downarrow & & \downarrow \\ H_1(\bar{Y}) & \longleftarrow & H_1(M) / \text{im}(1 - \bar{\phi}_*) \end{array}$$

Here we are considering sections of $\bar{Y}^{[n-1]}$ as $(n-1)$ -fold sections of \bar{Y} . If $n-1 \geq 2$, $\pi_1(\text{Sym}^{n-1}(\bar{\Sigma}))$ is abelian (isomorphic to $H_1(\bar{\Sigma}; \mathbb{Z})$), and the vertical map on the right is a bijection.

We conclude that $\pi_0 \text{sect}(\bar{Y}^{[n-1]})$ maps surjectively to $H_1(\bar{Y}; \mathbb{Z})_{n-1}$, and also injectively if $n-1 \geq 2$.

For any section of $\gamma \in \text{sect}(\bar{Y}^{[n-1]})$, γ^*Q is an S^1 -bundle over S^1 , and thus γ lifts to a section of $Q \rightarrow S^1$. If the surface Q is a torus, the pullback bundle γ^*Q is always orientable, and the available lifts (up to homotopy) are parametrised by \mathbb{Z} . If Q is a Klein bottle, it is never orientable, and the lifts are parametrised by $\mathbb{Z}/2$ (cf Part I, section 4). In either case, the lifts are in bijection with $\pi_0 \text{sect}(Q)$. Now \bar{Y} contains a distinguished braid B (which degenerates inside X to the critical circle Z). The section γ is homotopic to a section γ' whose image in \bar{Y} is well away from B . We noted in Part I (Remark 4.2) that, after moving Q by an isotopy, and restricting to an open subset $U \subset \bar{Y}^{[n-1]}$ of points $[x_1, \dots, x_{n-1}]$ where none of the x_i is close to B , we can suppose that the S^1 -bundle $Q \rightarrow \bar{Y}^{[n-1]}$ is identified with $\bar{Y}^{[n-1]} \times_{S^1} Q$. That is to say, to lift γ' to a section $\tilde{\gamma}'$ of Q , one just has to add a section of Q . But that means that $\alpha(\tilde{\gamma}') \in H_2^{(n)}$ exists.

Hence the map α exists, and we obtain a commutative diagram

$$\begin{array}{ccccc} \pi_0 \text{sect}(Q) & \longrightarrow & \pi_0 \text{sect}(Q) & \longrightarrow & \pi_0 \text{sect}(\bar{Y}^{[n-1]}) \\ & & \alpha \downarrow & & \downarrow \\ \pi_0 \text{sect}(Q) & \longrightarrow & H_2^{(d)} & \longrightarrow & H_1(\bar{Y}; \mathbb{Z})_{n-1} \end{array}$$

where the upper row is a “short exact sequence” of sets, the lower row a short exact sequence of affine abelian groups. We have already established the surjectivity (and when $n \geq 3$, injectivity) of the right-hand vertical arrow. This gives the result. \square

If $\bar{\Sigma}$ is disconnected, one should beware that $\text{sect}(\bar{Y}^{[n-1]})$ (and hence $\text{sect}(\mathcal{Q})$) may be empty, because $\bar{\phi}$ permutes the components of $\bar{\Sigma}$. However, we can still set up a commutative diagram

$$\begin{array}{ccc} \pi_0 \text{sect}(\mathcal{Q}) & \overset{\alpha}{\dashrightarrow} & \bigcup_{0 \leq d_1 \leq d} H_2^{(d_1, d-d_1)} \\ \downarrow & & \downarrow \delta \\ \pi_0 \text{sect}(Y^{[n]}) \times \pi_0 \text{sect}(\bar{Y}^{[n-1]}) & \longrightarrow & H_1(Y; \mathbb{Z})_n \oplus H_1(\bar{Y}; \mathbb{Z})_{n-1} \end{array}$$

which uniquely characterises α . This again leads to composite maps

$$\sigma = \tau_X^{-1} \circ \alpha: \pi_0 \text{sect}(\mathcal{Q}) \rightarrow \bigcup_{0 \leq d_1 \leq d} \text{Spin}^c(X)_{(d_1, d_2)}.$$

If $n - 1 \geq 2$, α (and hence σ) is injective. In general, however, it is not surjective.

4.3.1 Sections of relative Hilbert schemes versus Spin^c -structures Now take any broken fibration (X, π) over a surface S . Consider the boundary map

$$\delta: H_2(X, \partial X \cup Z; \mathbb{Z}) \rightarrow H_1(Z; \mathbb{Z}).$$

Let $\beta \in \delta^{-1}([Z])$. The intersection number of β with a regular fibre X_s is not constant, and nor is the Euler characteristic of the fibre, but one has the following:

Lemma 4.9 *The number $2\beta \cdot X_s + \chi(X_s)$ is independent of $s \in S \setminus S^{\text{crit}}$.*

Proof The Lefschetz dual $D(\beta)$ of β lies in $H^2(X \setminus Z; \mathbb{Z})$. The submanifold $X \setminus Z$ supports an almost complex structure compatible with ω and preserving the fibres of π , and whilst $c_1(X \setminus Z)$ does not extend to X , $c_1(X \setminus Z) + 2D(\beta)$ does (cf Taubes [38]). The result follows. \square

Note that one can define the Taubes map

$$\tau: \text{Spin}^c(X) \rightarrow H_2(X, \partial X \cup Z; \mathbb{Z})$$

precisely as in (12). As before, let $\nu: S \setminus S^{\text{crit}} \rightarrow \mathbb{Z}_{\geq 0}$ be a function satisfying $2\nu(s) + \chi(X_s) = d$, and form the relative Hilbert scheme $X^{[\nu]}$ and its matching condition \mathcal{Q} . Then any section $u \in \text{sect}(X^{[\nu]}, \mathcal{Q})$ defines a relative homology class

$$\alpha(u) \in \delta^{-1}([Z]) \subset H_2(X, \partial X \cup Z; \mathbb{Z}).$$

Indeed, over $S' = S \setminus \text{nd}(\pi(Z))$, u tautologically defines a cycle u' representing a class in $H_2(X, \pi^{-1}(\partial S); \mathbb{Z})$. This extends to a class in $H_2(X, Z; \mathbb{Z})$ as in Proposition 4.8.

The composite of τ^{-1} and α defines the map

$$\sigma_\nu = \tau^{-1} \circ \alpha: \pi_0 \text{sect}(X^{[\nu]}, \mathcal{Q}) \rightarrow \text{Spin}^c(X)_d$$

introduced in (13).

Proposition 4.10 *Assume $\text{inf}(\nu) \geq 2$. Then the map σ_ν is injective; if all fibres are connected, it is also surjective.*

Proof We claim that, if $S = S_1 \cup_T S_2$ is a splitting along a circle of regular values T , and $X = X_1 \cup_Y X_2$ is the corresponding splitting, then injectivity of σ_ν for X is implied by injectivity for X_1 and X_2 . Indeed, suppose sections $s_1, s_2 \in \text{sect}(X^\nu, \mathcal{Q})$ satisfy $\sigma_\nu(s_1) = \sigma_\nu(s_2)$. Then, by injectivity of σ_ν for the X_i , s_1 is homotopic to a new section s'_1 which differs from s_2 only over a small neighbourhood of T . The obstruction to extending the homotopy over T is a class $\delta \in \pi_1(Y^{[n]}, \gamma)$, where $n = \nu(s)$, $s \in T$, and $\gamma = s_2|_T$. But there is a natural injection $i: \pi_1(Y^{[n]}, \gamma) \rightarrow H_2(Y; \mathbb{Z}) = H^1(Y; \mathbb{Z})$, and $i(\delta)$ is the lift of $\sigma_\nu(s_1) - \sigma_\nu(s_2)$ to $H^1(Y; \mathbb{Z})$. Hence $\delta = 0$.

Similarly, surjectivity of σ_ν for X_1 and X_2 implies surjectivity for X , providing that $\text{inf } \nu \geq 2$. The point is that $\pi_0 \text{sect}(Y^{[n]})$ maps surjectively to $H_1(Y; \mathbb{Z})_n$ if $n \geq 2$ ($\pi_0 \text{sect}(Y^{[n]})$ is an affine space modelled on the abelianisation of $\pi_1(\text{Sym}^n(\Sigma))$, where Σ is the fibre of $Y \rightarrow S^1$).

With this patching procedure to hand, it suffices to check the result in three cases: (i) S is a connected surface with nonempty boundary and $X^{\text{crit}} = \emptyset$; (ii) S is a disc and $X \rightarrow S$ is an elementary Lefschetz fibration; (iii) S is an annulus and $X \rightarrow S$ an elementary broken fibration. Case (i) is easy once one observes that S retracts to a bouquet of circles; (ii) is also straightforward and left to the interested reader. Proposition 4.8 and the remarks following its proof deal with (iii). The result follows. \square

4.4 Computing the index

The map $\sigma_\nu: \pi_0 \text{sect}(\mathcal{Q}) \rightarrow \text{Spin}^c(X)_d$ gives a tidy way of partitioning the sections of \mathcal{Q} , but its importance goes beyond book-keeping: we can use it to understand the index problem for pseudo-holomorphic sections with boundary on \mathcal{Q} . Whilst setting up invariants from such sections depends critically on compactness and transversality theorems, the index computation can be done in a “soft” topological context. Underlying

this and virtually all such computations is the invariance of the Fredholm index under compact perturbations.

The proof of the index formula is rather complicated. As a warm-up, and so as to make clear which aspects are tied specifically to our matching conditions, we first consider the index problem for the relative Hilbert scheme $X^{[n]} = \text{Hilb}_S^n(X)$ of a *Lefschetz* fibration $\pi: X \rightarrow S$ over a closed surface. In the case $S = S^2$, this was done by Smith [37, Proposition 4.3] using an analytic form of Grothendieck–Riemann–Roch. It can be done in more elementary fashion as follows.

Lemma 4.11 *Let $u \in \text{sect}(X^{[n]})$, and let $A \in H_2(X; \mathbb{Z})$ be its underlying homology class. Then the index of u (which, according to Riemann–Roch, is equal to $2\langle c_1(T^\vee X^{[n]}), [u] \rangle + n\chi(S)$) is*

$$\text{ind}(u) = A \cdot A + \langle c_1(TX), A \rangle.$$

The proof is preceded by an algebraic digression. Suppose that X is actually a complex surface (not necessarily compact) and $\pi: X \rightarrow S$ a holomorphic submersion (not necessarily proper). To give a holomorphic section $u: S \rightarrow X^{[n]}$ of the relative symmetric product $\text{Sym}_S^n(X)$ is to give a commutative triangle

$$\begin{array}{ccc} T_u & \xrightarrow{\tilde{u}} & X \\ \downarrow t_u & \searrow \pi & \\ S & & \end{array}$$

where T_u is a Riemann surface, t_u a proper holomorphic map of degree d , and \tilde{u} a holomorphic mapping into X . One constructs T_u , as a branched cover of S , as the space of germs of sections of π mapping each point s into $\text{Supp}(u(s))$.

Lemma 4.12 *For any holomorphic section $u: S \rightarrow X$, there is an isomorphism of holomorphic vector bundles over S ,*

$$u^*T^\vee \text{Sym}_S^n(X) \cong t_{u*}\tilde{u}^*(T^\vee X).$$

Proof This is essentially the infinitesimal version of the correspondence just described. The standard algebro-geometric description of $T \text{Sym}^n(X_s)$ at an effective divisor D is as the vector space of “Laurent tails” $H^0(X_s; \mathcal{O}_{X_s, D}(D))$ (here $\mathcal{O}_{X_s, D}(D)$ is the cokernel of the natural inclusion of sheaves $\mathcal{O}_{X_s} \rightarrow \mathcal{O}_{X_s}(D)$). Thus the fibre of $u^*T^\vee \text{Sym}_S^n(X)$ at x is $H^0(X_s; \mathcal{O}_{X_s, D}(D))$, where D is $u(x)$, counted with multiplicities. On the other hand, the fibre of $t_{u*}\tilde{u}^*(T^\vee X)$ at x is the space of D -jets of holomorphic maps into X_s at D , which is also naturally identified with $H^0(\mathcal{O}_{X_s, D}(D))$. □

Proof of Lemma 4.11 Because u is a smooth section, it does not map to any critical points of $\pi^{[n]}$. The index only depends on u only through its homotopy class; hence, by making small perturbations to u , we may arrange that u has generic properties. Namely, we can assume that it is transverse to each stratum of the diagonal in $X^{[n]}$, and hence only intersects the diagonal in its top stratum (ie, there are no triple self-intersections). Moreover, we may suppose that, for any such intersection point $u(s)$, there is a small disc $D_s \subset S$ centred at s over which u is either holomorphic or antiholomorphic, according to the sign of the intersection: having identified the self-intersection points $u(s)$, we adjust the almost complex structure on X so that it becomes integrable in a neighbourhood of X_s .

In this situation, we can again construct a “parametrisation” of u as a commutative diagram

$$\begin{array}{ccc} T_u & \xrightarrow{\tilde{u}} & X \\ & \searrow t_u & \swarrow \pi \\ & & S. \end{array}$$

Here $t_u: T_u \rightarrow S$ is a degree n map of smooth surfaces whose fibre over x is $\text{Supp}(u(x))$, *except* when $u(x)$ hits the diagonal. In the latter case, u is locally (anti)holomorphic, and T_u can be constructed in the way discussed in the preamble to this proof. Thus t_u has simple ramification. We still have an isomorphism of complex vector bundles $u^* T^\vee \text{Sym}_S^n(X) \cong t_{u*} \tilde{u}^*(T^\vee X)$: the previous lemma gives such an isomorphism near the branch values, and there is then an obvious (and canonical) extension to the rest of S .

The image of \tilde{u} is an embedded surface U in X ; the branch points of t_u (in T_u) map to tangencies with the fibres of X . Of these tangencies, some (N_+ , say), have coincident orientations with the fibre of X , and some (say N_-) have opposite orientation. The two local models for t_u are respectively $z \mapsto z^2$ and $z \mapsto \bar{z}^2$. In this context, the Riemann–Hurwitz formula reads

$$\chi(T_u) = n\chi(S) - (N_+ + N_-).$$

On the other hand, for a complex line bundle $L \rightarrow T_u$,

$$(15) \quad \langle c_1(t_{u*}L), [S] \rangle = \langle c_1(L), [T_u] \rangle + \frac{1}{2}(N_+ - N_-).$$

We now calculate

$$\begin{aligned}
 \text{ind}(u) &= 2\langle c_1(u^* T^\vee X^{[n]}), [S] \rangle + n\chi(S) && \text{(Riemann–Roch)} \\
 &= 2\langle \tilde{u}^* c_1(T^\vee X), [T_u] \rangle + N_+ - N_- + n\chi(S) && \text{(by Equation (15))} \\
 &= 2\langle \tilde{u}^* c_1(TX) - \tilde{u}^* \pi^* c_1(TS), [T_u] \rangle \\
 &\quad + N_+ - N_- + n\chi(S) && \text{(since } TX = T^\vee X \oplus \pi^* TS) \\
 &= 2\langle \tilde{u}^* c_1(TX), [T_u] \rangle - n\chi(S) + N_+ - N_- && \text{(since } \deg(t_u) = n) \\
 &= 2\langle \tilde{u}^* c_1(TX), [T_u] \rangle - \chi(T_u) - 2N_- && \text{(Riemann–Hurwitz).}
 \end{aligned}$$

We can now make the link with the homology class A , since

$$\langle \tilde{u}^* c_1(TX), [T_u] \rangle = \langle c_1(TX), A \rangle.$$

If $N_- = 0$, we can homotope the almost complex structure on X so that it preserves TU , whereupon the adjunction formula gives

$$-\chi(T_u) = A \cdot A - \langle c_1(TX), A \rangle.$$

If $N_- > 0$ then we may take U to be almost complex except at the N_- orientation-reversing vertical tangencies. There is a relation

$$-\chi(T_u) + kN_- = A \cdot A - \langle c_1(TX), A \rangle$$

for a *universal* integer k ; indeed, k is the first Chern class of the trivial \mathbb{C}^2 -bundle over the disc, Δ , relative to a certain trivialisation over its boundary. Hence

$$\text{ind}(u) = \langle c_1(TX), A \rangle + A \cdot A - (k + 2)N_-.$$

The cheapest way to deduce that $k = -2$ is to observe that (there is some example in which) one can deform an initial $u_0 \in \text{sect}(X^{[n]})$ which has $N_- = 0$ to a homotopic section u_1 which is generic and has $N_- > 0$. We have $\text{ind}(u_0) = \text{ind}(u_1)$, but the only way this can occur is if $k = -2$. The result now follows. \square

The Spin^c -structure $\mathfrak{s} = \tau_X^{-1}(A)$ corresponding to A is $L_A \otimes \mathfrak{s}_{\text{can}}$, where L_A is the line bundle with a section with transverse zero-locus A ; thus $c_1(\mathfrak{s}) = c_1(TX) + 2\text{PD}(A)$. Since $c_1(TX)^2 = 2e(X) + 3\sigma(X)$, one has $A^2 + \langle c_1(TX), A \rangle = d(\mathfrak{s})$, where $d(\mathfrak{s})$ is as in Equation (14).

We now turn to broken fibrations, and to one of the results previewed in Part I of this pair of papers.

Theorem D Suppose that $u \in \text{sect}(X^{[v]}, \mathcal{Q})$ represents the Spin^c -structure \mathfrak{s} , ie that $\sigma_v(u) = \mathfrak{s}$. Then the index of u is

$$\text{ind}(u) = d(\mathfrak{s}).$$

A part of the proof is to verify the result by hand for each of two broken fibrations (and at least one section u in each case). The attaching surface in one of these fibrations is a torus; in the other, a pair of Klein bottles. We do this now.

Example 4.13 This example is a broken fibration $\pi: X \rightarrow S^2$. The total space is $X = X_+ \cup X_0 \cup X_-$ where X_+ is a trivial T^2 -bundle over a north-polar disc D_+ ; X_- is a trivial S^2 -bundle over the south-polar disc D_- ; and these are joined in the trivial, “untwisted”, way by a broken fibration X_0 over the equatorial annulus. The vanishing surface $\mathcal{Q} \subset \partial X_+$ is a torus, and there is a section u of $(D_+, \partial D_+) \rightarrow (X_+, \mathcal{Q})$ of Maslov index 0. Thus $\text{ind}(u) = 1$. In fact, for a standard almost complex structure, the moduli space of pseudo-holomorphic sections is S^1 .

It is easy to see that $e(X) = 2$ and $b_1(X) = 1$, so $b_2(X) = 2$ and, since there is a section, $\sigma(X) = 0$ (in fact, as shown in the last section of [1], $X \cong (S^1 \times S^3) \# (S^2 \times S^2)$). The section u can be completed to a section v over S^2 , which we may take to be pseudo-holomorphic (adjusting the almost complex structure on $X \setminus Z$). We have $[v] \cdot [v] = 0$, so by the adjunction formula, $\langle c_1(X \setminus Z), [v] \rangle = 2$.

The Spin^c -structure \mathfrak{s} associated with u has $c_1(\mathfrak{s}) = 2[\text{fibre}] + 2[v]$. Note here that, on $X \setminus Z$, $c_1(\mathfrak{s}) = c_1(X \setminus Z) + 2\text{PD}[u]$, and that one can calculate the intersection of $c_1(X \setminus Z) + 2\text{PD}[u]$ with a surface in $X \setminus Z$ as a sum of two intersection numbers. Thus $c_1(\mathfrak{s})^2 = 8$, and $d(\mathfrak{s}) = 2 - 1 = 1$.

Example 4.14 In this example of a broken fibration $X \rightarrow S^2$ (a homotopy- $S^2 \times S^2$), the vanishing surface \mathcal{Q} is the union of two Klein bottles. We have $X = X_+ \cup X_0 \cup X_-$, where $X_+ \rightarrow D_+$ and $X_- \rightarrow D_-$ are now trivial S^2 -bundles. The equatorial part X_0 is a broken fibration with two circles of critical points, mapping to parallel circles (the “tropics”). The fibre over the equator is $S^2 \cup S^2$, and the monodromy around the equator interchanges the two components. The vanishing Klein bottle $\mathcal{Q}_- \subset S^2 \times \partial D_-$ is $\bigcup_{z \in S^1} l_z \times \{z\}$, where l_z is the circle $z^{1/2}\mathbb{R} \cup \{\infty\} \subset \mathbb{C} \cup \{\infty\} = S^2$. The fibration is symmetric under reflection in the equator in the base S^2 ; in particular, the other Klein bottle \mathcal{Q}_+ is the reflection of the first.

We have $\pi_1(X) = \{1\}$ (as one verifies with a little care) and $e(X) = 4$, so $b_2(X) = 2$. There exists a two-fold section, so the fibres are nontrivial in homology and the signature is zero.

There is a section u_- over D_- with boundary on Q_- , given by $u_-(z) = 0 \in \mathbb{C} \subset S^2$. This has Maslov index $+1$. There is a similar section u_+ over D_+ . Their union, $u = u_+ \cup u_-$, has index $\text{ind}(u) = 2(1 + 1) = 4$.

To calculate $c_1(\mathfrak{s})^2$ for the associated Spin^c -structure \mathfrak{s} , we look more closely at $H_2(X; \mathbb{Z})$. Observe that there is a two-fold section s , given on D_- by $s(z) = \{0, \infty\}$, and similarly on D_+ . In the equatorial annulus, $s(z)$ is given by one point on each of the two components of the fibre X_z . Moreover, we can arrange firstly that $S := \text{im}(s)$ is pseudo-holomorphic for some almost complex structure compatible with the broken fibration, and secondly that S is an embedded sphere with $[S] \cdot [S] = 4$. By adjunction, $\langle c_1(X \setminus Z), [S] \rangle = 6$. Notice that S is homologous inside $X \setminus Z$ to another surface S' transverse to u and meeting $\text{im}(u)$ with (multiplicity $+1$) at $0 \in D_+$ and at $0 \in D_-$. Thus $\langle c_1(\mathfrak{s}), [S] \rangle = 6 + 2 \cdot 2 = 10$. We also have $\langle c_1(\mathfrak{s}), F \rangle = 2$, where F is an equatorial fibre. Hence $c_1(\mathfrak{s}) = 2[S] + 2F$. Then $c_1(\mathfrak{s})^2 = 24$, and $d(\mathfrak{s}) = 4 = \text{ind}(u)$.

Proof of Theorem D We will show that $\text{ind}(u) - d(\mathfrak{s})$ is some universal multiple of the number of tori, plus a universal multiple of the number of Klein bottles. The foregoing computations, showing that $\text{ind}(u) - d(\mathfrak{s}) = 0$ in each of the two examples above, then complete the proof of the theorem.

The argument is a little less intricate when $S = S^2$ and there is just one critical circle Z , so we will deal with that case then indicate how it is generalised. So, $S' = S_+ \amalg S_-$, the union of two discs, and there is an orientation-reversing diffeomorphism $\tilde{\tau}: S_+ \rightarrow S_-$.

One can think of u as a section (u_+, u_-) of $X_+^{[n]} \times_{S_+} \tilde{\tau}^* X_-^{[-1]}$. The Riemann–Roch theorem for surfaces with boundary gives the index of u as

$$\text{ind}(u) = \mu_{\mathcal{Q}}(u) + n + (n - 1),$$

where $\mu_{\mathcal{Q}}(u)$ is the Maslov index for u relative to \mathcal{Q} , when u is considered as a section of $X_+^{[n]} \times \tilde{\tau}^* X_-^{[n-1]}$. We compute this index in a number of steps.

Step 1 Consider the matching condition \mathcal{Q} lying over ∂S_+ . At a homotopical level, we can make a sequence of simplifications. Let $Y = \pi^{-1}(\partial S_+)$ and $\bar{Y} = \pi^{-1}(\tau(\partial S_+))$.

- (1) \bar{Y} contains a distinguished braid, B a two-fold section of $Y \rightarrow S^1$, associated with the broken fibration: see the discussion at the beginning of Section 4.1 in Part I. Using Remark 4.2 from Part I, we can isotope \mathcal{Q} through totally real subbundles so that, over an open set $U \subset \bar{Y}^{[n-1]}$ of points $[x_1, \dots, x_{n-1}]$ with $x_i \neq \text{nd}(B)$, we have $\mathcal{Q} \cong \bar{Y}^{[n-1]} \times_{S^1} \mathcal{Q}$ as an S^1 -bundle over $\bar{Y}^{[n-1]}$. Furthermore, the embedding $\mathcal{Q} \rightarrow Y$ sends $([x_1, \dots, x_{n-1}]; q)$ to $[x_1, \dots, x_{n-1}, q] \in Y^{[n]}$. Such a totally real isotopy leaves Fredholm indices unaltered.

- (2) Suppose $\Gamma = (\gamma, \bar{\gamma}) \in \text{sect}(\mathcal{Q})$. We can perturb $\bar{\gamma}$ to another section of \bar{Y}^{n-1} which avoids U , and lift the homotopy to \mathcal{Q} , so obtaining a homotopic section of \mathcal{Q} .
- (3) After a further (generic) perturbation, γ does not intersect the diagonal in $Y^{[n]}$, and $\bar{\gamma}$ does not intersect the diagonal in $\bar{Y}^{[n-1]}$.

When Γ satisfies point (3), we can canonically associate with it an n -fold covering $t_\gamma: T_\gamma \rightarrow S^1$; an $(n-1)$ -fold covering $t_{\bar{\gamma}}: T_{\bar{\gamma}} \rightarrow S^1$; and maps

$$\tilde{\gamma}: T_\gamma \rightarrow Y, \quad \tilde{\bar{\gamma}}: T_{\bar{\gamma}} \rightarrow \bar{Y}$$

such that $\gamma = \tilde{\gamma} \circ t_\gamma^{-1}$ and $\bar{\gamma} = \tilde{\bar{\gamma}} \circ t_{\bar{\gamma}}^{-1}$. We then have

$$\gamma^* T^\vee Y^{[n]} = (t_\gamma)_* \tilde{\gamma}^* T^\vee Y, \quad \bar{\gamma}^* T^\vee \bar{Y}^{[n-1]} = (t_{\bar{\gamma}})_* \tilde{\bar{\gamma}}^* T^\vee \bar{Y}.$$

By point 2, T_γ is isomorphic to a disjoint union $S^1 \cup T_{\bar{\gamma}} \rightarrow S^1$. Moreover, $\gamma^* T^\vee Y^{[n]} \cong q^* T^\vee Y \oplus \bar{\gamma}^* T^\vee \bar{Y}^{[n-1]}$, where $q: S^1 \rightarrow Q_i$ is the restriction of γ to $S^1 \subset T_\gamma$. To put this more plainly: γ consists of a section q of the vanishing surface $Q_i \subset Y$, together with an $(n-1)$ -fold section disjoint from Q_i which matches with the $(n-1)$ -fold section $\bar{\gamma}$ of \bar{Y} .

With a generic $u \in \text{sect}(X^{[v]}, \mathcal{Q})$ we can associate, as in the proof of Lemma 4.11, a surface with boundary T_u , a branched covering $t_u: T_u \rightarrow S^1$, and a smooth map $\tilde{u}: T_u \rightarrow X'$, such that $\tilde{u} \circ t_u^{-1} = u$ and

$$u^* T^\vee X^{[v]} = (t_u)_* \tilde{u}^* T^\vee X'.$$

$T^\vee \mathcal{Q}$ defines a (linearised) boundary condition for deformations of u , and this translates into a boundary condition for deformations of (T_u, \tilde{u}) . We understand the boundary condition using point (1): T_u has a distinguished component C over ∂S_+ , on which the map \tilde{u} is q . On this boundary component, the boundary condition is $q^* T^\vee \mathcal{Q} \subset q^* T^\vee Y$. The remaining components of T_u over ∂S_+ , say $\gamma_1, \dots, \gamma_k$, are matched with the components of T_u over ∂S_- . On each pair γ_i , the boundary condition is the diagonal Lagrangian matching condition $\text{diag}_{\gamma_i^* T^\vee \bar{Y}} \subset \gamma_i^* T^\vee \bar{Y} \oplus \gamma_i^* T^\vee \bar{Y}$.

Step 2 We can go further by joining up the matching ends of T_u , eliminating the diagonal boundary condition. We have $S = S' \cup N$, where N is an annulus. We can find a covering $T_{\text{tunnel}} \rightarrow N$ which joins up with T_u to give a surface with boundary $T_{\text{joined}} = T \cup T_{\text{tunnel}}$ with a map $t_{\text{joined}}: T_{\text{joined}} \rightarrow S$. Moreover, there is a map $v = u_{\text{joined}}: T_{\text{joined}} \rightarrow X$, lifting t_{joined} and extending $\tilde{u}: T_u \rightarrow X'$. The boundary $\partial T_{\text{joined}}$ maps diffeomorphically to ∂S_+ under t_{joined} , while $u_{\text{joined}}(\partial T_{\text{joined}}) \subset \mathcal{Q}$. Now, one has

$$\mu_{\mathcal{Q}}(u) = \mu_{\mathcal{Q}}(v) + (N_+ - N_-),$$

where N_+ and N_- are numbers of branch points of $T_{\text{joined}} \rightarrow S$, as in the proof of Lemma 4.11. (This is a small modification of (15).)

Step 3 The next step is to cap off T_{joined} —not in X , but in a manifold $X_{Z,f}$ obtained by surgery on X along Z . That is, we choose a framing f for Z , excise a small neighbourhood $Z \times D^3$ of Z from X , and glue in $D^2 \times S^2$ in its place. Then $\partial T_{\text{joined}}$ bounds a standard embedded disc D in X_Z .

Notice that the framings can be chosen according to a standard recipe. If Q is a torus, we may parametrise a neighbourhood of $Z_i \cong S^1$ as $S^1 \times D^3$ in such a way that $Q = S^1 \times \{(x, y, 0) : x^2 + y^2 = 1\} \subset \partial(S^1 \times D^3)$ and such that the curve $\text{im}(q) \subset Q$ is $S^1 \times \{(1, 0, 0)\}$. If Q is a Klein bottle, we can choose fix some other model in which Q and q appear in a standard way in $\partial(S^1 \times D^3)$.

The almost complex structure J on $X \setminus Z$ then extends to an almost complex structure J' on $X_{Z,f}^* := X_{Z,f} \setminus \{p\}$, where p is a point in the added $D^2 \times S^2$. There is no obstruction to choosing J' so that the disc D is pseudo-holomorphic.

Now D has complex normal bundle, and tangents to Q form a real subbundle over ∂D . There is therefore a Maslov index $\mu_Q(D)$, which is actually zero when Q is a torus. We can pin down our universal model in the Klein bottle case by stipulating that the Maslov index (which can be any odd integer) should be 1.

Let $W = T_{\text{joined}} \cup \bigcup_i D_i$ be the capped surface, and $w: W \rightarrow X_{Z,f}$ the natural map extending v . The complex line bundles $v^*T^v X$ and $N_{D/X_{Z,f}}$ are isomorphic over the common boundary of T_u and D , so join together to give a line bundle $L \rightarrow W$. We then have

$$(16) \quad \mu_Q(u_{\text{joined}}) + \mu_Q(D) = 2\langle c_1(L), [W] \rangle.$$

Notice that there is a degree n map $p: W \rightarrow S$ extending $t_u: T_u \rightarrow S'$. We now compute as in Lemma 4.11:

$$\begin{aligned} \text{ind}(u) &= \mu_Q(u) + (2n - 1) && \text{(Riemann–Roch)} \\ &= \mu_Q(u_{\text{joined}}) + (2n - 1) + N_+ - N_- && \text{(by Step 2)} \\ &= 2\langle c_1(L), W \rangle - \mu_Q(D) + (2n - 1) + N_+ - N_- && \text{(by Equation (16))} \\ &= 2\langle w^*c_1(TX_{Z,f}^*) - w^*p^*c_1(TS), [W] \rangle - \mu_Q(D) \\ &\quad + (2n - 1) + N_+ - N_- && \text{(splitting } w^*TX_{Z,f}^*) \\ &= 2\langle w^*c_1(TX_{Z,f}^*), [W] \rangle - \mu_Q(D) - 2n - 1 + N_+ - N_- && \text{(deg}(p) = n) \\ &= 2\langle w^*c_1(TX_{Z,f}^*), [W] \rangle - \mu_Q(D) - \chi(W) - 2N_- - 1 && \text{(Riemann–Hurwitz)} \\ &= \langle w^*c_1(TX_{Z,f}^*), [W] \rangle + w_*[W] \cdot w_*[W] - \mu_Q(D) - 1. && \text{(adjunction)} \end{aligned}$$

The last line uses the variant of the adjunction formula discussed in the proof of Lemma 4.11.

Thus $\text{ind}(u)$ differs from $[W]^2 + \langle c_1(TX_{Z,f}^*), W \rangle$ by $\mu_Q(D) + 1$, which is 1 or 2, depending on the topology of Q . But $c_1(TX_{Z,f}^*)^2$ differs from $2\chi(X_{Z,f}) + 3\sigma(X_{Z,f})$ just by contributions (“Hopf invariants”) from the singularities of p [14]. These depend only on the topology of Q . It follows that

$$\text{ind}(u) - \frac{1}{4}[c_1(\mathfrak{s}_W)^2 - 2e(X_{Z,f}) - 3\sigma(X_{Z,f})]$$

depends only on the topology of Q , where \mathfrak{s}_W is the J' -canonical Spin^c -structure on $X_{Z,f}^*$, twisted by $\text{PD}(W)$. Hence the same is true of $\text{ind}(u) - d(\mathfrak{s})$. This establishes our claim.

We now drop the assumptions that $S = S^2$ and that Z is a single circle. By the index formula of Lemma 3.3, we now have

$$\text{ind}(u) = \mu_Q(u) + \sum_{S_i \in \pi_0(S')} v(S_i)\chi(S_i)$$

where $\mu_Q(u)$ is a generalised Maslov index. One can form T_{joined} and v as before, and $\mu_Q(u) = \mu_Q(v) + (N_+ + N_-)$. The computation now proceeds as before. It shows, as claimed, that $\text{ind}(u) - d(\mathfrak{s})$ is a universal multiple of the number of torus components of Q , plus a universal multiple of the number of Klein bottle components. \square

4.5 The relative theory: Floer homology

We now begin to lay out the relative version of the theory—the version for broken fibrations over surfaces with boundary. We begin by setting out the special features of Floer homology for relative symmetric products of fibred three-manifolds. These groups first appear, in connection with Seiberg–Witten theory, in the work of D Salamon [28]. They are studied in the author’s thesis [23] and also by Usher [39], who noticed how to achieve the monotonicity condition explained below.

4.5.1 Topological sectors Let (Y, π) be a surface bundle over S^1 , $\Sigma = \pi^{-1}([0])$ its fibre, $j \in \mathcal{J}(T^*Y)$ a vertical complex structure, and $Y^{[n]} = \text{Sym}_{S^1}^n(Y)$ the relative symmetric product. Let $H_1(Y; \mathbb{Z})_n \subset H_1(Y; \mathbb{Z})$ be the affine subgroup of classes γ with $\gamma \cdot [\Sigma] = n$. As discussed in the last section, there is a “tautological” map

$$a: \pi_0 \text{sect}(Y^{[n]}) \rightarrow H_1(Y; \mathbb{Z})_n.$$

Assuming Σ connected, this map is surjective and, if $n > 1$, also injective (see the proof of Proposition 4.8).

The map $\tau: H_1(Y; \mathbb{Z}) \rightarrow \text{Spin}^c(Y)$, sending γ to $\text{PD}(\gamma) \cdot \mathfrak{t}_{\text{can}}$, maps $H_1(Y; \mathbb{Z})$ bijectively to the subset $\text{Spin}^c(Y)_{n+\chi(\Sigma)/2}$ of Spin^c -structures \mathfrak{t} with $\langle c_1(\mathfrak{t}), [\Sigma] \rangle = 2n + \chi(\Sigma)$. We put

$$HF_*(Y, \mathfrak{t}) = \bigoplus_{\gamma \in (\tau \circ a)^{-1}(\mathfrak{t})} HF_*(Y^{[n]}, \gamma).$$

The group $HF_*(Y, \mathfrak{t})$ should really be notated as $HF_*(Y^{[n]}, \pi^{[n]}, \mathfrak{t}; \Omega)$, since it depends upon a choice of closed two-form Ω . The matter of choosing Ω takes us back to the homomorphism $\kappa_{n,\lambda}$ constructed at the beginning of this paper. The monodromy m of $Y \rightarrow S^1$ is only defined modulo isotopies, but by choosing a suitable two-form on Y we can find a lift \tilde{m} to the symplectic mapping class group of Σ . We can then apply $\kappa_{n,\lambda}$ to it. Thus we choose Ω to be a fibrewise-Kähler form whose monodromy represents $\kappa_{n,\lambda}(m)$.

The adjustable parameters in this procedure are the choice of lift \tilde{m} of the monodromy m , and the value of λ . There is no obligation to make the same choices for different \mathfrak{t} , but we can pick out a canonical, \mathfrak{t} -dependent choice by insisting that a monotonicity condition holds. This is the topic of the following paragraphs.

4.5.2 Periods and monotonicity Consider the mapping torus $T(\phi)$ of a self-diffeomorphism $\phi \in \text{Diff}^+(M)$ of a connected manifold M . Assume that $\phi(x) = x$; we then use x as a basepoint of M , and c_x (the “constant section at x ”) as a basepoint for the twisted free loop space $\text{sect}(T(\phi))$.

Lemma 4.15 *The map $e: \text{sect}(T(\phi)) \rightarrow M$ given by evaluation at $[0] \in S^1$ gives rise to an exact sequence*

$$\pi_2(M) / \text{im}(1 - \phi_*) \rightarrow \pi_1(\text{sect}(T(\phi)), c_x) \xrightarrow{e_*} \pi_1(M)^\phi \rightarrow 1$$

where $\pi_1(M)^\phi \subset \pi_1(M)$ is the group of invariants.

A very similar assertion appears in M Poźniak’s thesis [27].

Proof The evaluation map e is a Serre fibration—just like the evaluation map $\mathcal{L}M \rightarrow M$ on the free loop space—and the sequence can be obtained from the associated exact sequence of homotopy groups. However, it can also be derived directly.

Think of loops in $\text{sect}(T(\phi))$ as maps $\Gamma: S^1 \times [0, 1] \rightarrow M$ satisfying $\Gamma(t, 1) = \phi \circ \Gamma(t, 0)$. Given a based loop $\gamma: S^1 \rightarrow M$ representing a ϕ -invariant class $[\gamma]$, a choice of homotopy $\phi_*\gamma \simeq \gamma$ gives rise to a loop in $\pi_1(\text{sect}(T(\phi)), c_x)$, hence e_* maps onto $\pi_1(M)^\phi$. One obtains a map $d: \pi_2(M) \rightarrow \pi_1(\text{sect}(T(\phi)), \ell)$ by modifying the

constant loop at l by a map $(D^2, \partial D^2) \rightarrow (M, \{x\})$, where D^2 is embedded as a small disc inside $S^1 \times [0, 1]$. Elements of $\pi_2(M)$ shape $s - \phi_*s$ then lie in the kernel of d , and $\text{im}(d) \subset \ker(e_*)$. Moreover, if $[\Gamma] \in \ker(e_*)$ then, after a homotopy, Γ is constant along $S^1 \times \{0\}$, and so after a further homotopy differs from l by a map supported in a disc contained interior of $S^1 \times [0, 1]$. Hence $[\Gamma] \in \text{im}(d)$. If $[r] \in \pi_2(M)$ lies in $\ker(d)$ then one choose a nullhomotopy $\{\Gamma_t\}_{t \in [0,1]}$ of Γ_0 (the image of r). Then $e \circ \Gamma_t$ defines a 2-sphere s in M , and one has $r \simeq s - \phi_*s$. \square

Proposition 4.16 *If $n > 2$, the natural map*

$$p: \pi_1 \text{sect}(Y^{[n]}, \gamma) \rightarrow H_2(Y; \mathbb{Z}) = H^1(Y; \mathbb{Z}),$$

is injective, with cokernel isomorphic to \mathbb{Z}/n .

Proof When $n > 2$ one has $\pi_2(\text{Sym}^n(\Sigma)) = \mathbb{Z}$: this is true when $n \gg 0$ because $\text{Sym}^n(\Sigma)$ is then a projective space bundle over a torus. It then follows for all $n \geq 3$ by descending induction, using the Lefschetz hyperplane theorem. The action of ϕ_* on π_2 is trivial. After homotoping the monodromy of $Y^{[n]}$ so that it has a fixed point, the previous lemma then gives the upper row in the following commutative diagram:

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathbb{Z} & \longrightarrow & \pi_1 \text{sect}(Y^{[n]}, \gamma) & \longrightarrow & H_1(\text{Sym}^n(\Sigma))^\phi & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & H_2(\Sigma) & \longrightarrow & H_2(Y) & \longrightarrow & H_1(\Sigma)^\phi & \longrightarrow & 0 \end{array}$$

The lower row is the short exact sequence describing homology for a mapping torus; the maps in it are the ones induced by inclusion of a fibre and by intersection with a fibre. The vertical maps are the “tautological” ones. The one on the right is an isomorphism; that on the left is injective, with cokernel \mathbb{Z}/n . By the snake lemma, $\pi_1 \text{sect}(Y^{[n]}, \gamma) \rightarrow H_2(Y; \mathbb{Z})$ is also injective with cokernel \mathbb{Z}/n . \square

The Chern class of $T^\vee Y^{[n]}$ is given, in the notation of the introduction, by

$$(17) \quad c_1(T^\vee Y^{[n]}) = (c_1(T^\vee Y)^{[1]} + 1^{[2]})/2.$$

The two sides agree on the fibres $\text{Sym}^n(\Sigma)$ since the restrictions of $c_1(T^\vee Y)^{[1]}$ and $1^{[2]}$ are respectively $(2 - 2g)\eta_\Sigma$ and $2n\eta_\Sigma - 2\theta_\Sigma$. One way to prove this formula is to use Grothendieck–Riemann–Roch (the calculation is very similar to that of the index in [37]). As pointed out to the author by Michael Usher (private communication), it has a more concrete expression:

$$\langle c_1(T^\vee Y^{[n]}), [\Gamma] \rangle = \langle c_1(T^\vee Y) + 2\text{PD}(\gamma), p(\Gamma) \rangle / 2 = \langle c_1(t_\gamma), p(\Gamma) \rangle / 2,$$

where Γ is a loop in $\text{sect}(Y)$ based at γ and $t_\gamma = \tau(\gamma)$. This can also be verified directly, using the method of Lemma 4.11.

Now, Floer homology for an LHF $(T \rightarrow S^1, \Omega)$ is at its simplest when the homomorphism

$$[\Omega]: \pi_1(\text{sect}(T), \gamma) \rightarrow \mathbb{Z}$$

is a multiple of the index homomorphism defined by $c_1(T^\vee T)$. This is the ‘‘monotonicity’’ condition, and has been exploited by a number of authors (most relevantly, by Seidel [33] in connection with the mapping class group). In our set-up, monotonicity can be achieved provided that $n \neq -\chi(\Sigma)/2$.

Lemma 4.17 *Fix a Spin^c -structure $t = t_\gamma \in \text{Spin}^c(Y)_{2n+\chi(\Sigma)}$, where $2n + \chi(\Sigma) \neq 0$. Put*

$$W_\lambda = (1 + \lambda n) (w_\lambda)^{[1]} - \frac{\lambda}{2} 1^{[2]} \in H^2(Y^{[n]}; \mathbb{R}),$$

where
$$w_\lambda = \frac{\lambda \text{PD}(\gamma) + (\chi(\Sigma) + 2n)^{-1} c_1(t_\gamma)}{1 + \lambda n}, \quad \lambda > -\frac{1}{n}.$$

Then the following monotonicity relation holds: $W_\lambda - (2/(\chi(\Sigma) + 2n))c_1(T^\vee Y^{[n]})$ vanishes on $\pi_1(\text{sect}(Y^{[n]}, \gamma))$.

Proof Take $w \in H^2(Y; \mathbb{R})$ with $\langle w, [\Sigma] \rangle = 1$. On $Y^{[n]}$ we consider the class

$$W = (1 + \lambda n) w^{[1]} - \frac{\lambda}{2} 1^{[2]}, \quad \lambda > 0$$

(cf Equation (1)). One has

$$\langle W, [\Gamma] \rangle = \langle (1 + \lambda n) w - \lambda \text{PD}(\gamma), p(\Gamma) \rangle.$$

We have learned that

$$c_1(T^\vee Y^{[n]}, [\Gamma]) = \frac{1}{2} \langle c_1(t_\gamma), p(\Gamma) \rangle.$$

If we put $w = w_\lambda = (1 + \lambda n)^{-1}(\lambda \text{PD}(\gamma) + (\chi(\Sigma) + 2n)^{-1} c_1(t_\gamma))$ then we will have, $W = W_\lambda$, $\langle w, [\Sigma] \rangle = 1$, and

$$\left\langle W_\lambda - \frac{2}{\chi(\Sigma) + 2n} c_1(T^\vee Y^{[n]}, [\Gamma]) \right\rangle = 0. \quad \square$$

We can express this lemma in terms of monodromy. Fix an area-form α on Σ of total area 1, and let $m \in \pi_0 \text{Diff}^+(\Sigma)$ be the monodromy of $Y \rightarrow S^1$. Lift it to $\tilde{m} \in \text{Aut}(\Sigma, \alpha) / \text{Ham}(\Sigma, \alpha)$ by specifying that the class $[\alpha_{\tilde{m}}]$ of the closed two-form

$\alpha_{\tilde{m}}$ induced by α on $T(\tilde{m})$ should equal w , as in the lemma. The monodromy of the form Ω constructed in the lemma is then $\kappa_{n,\lambda}(\tilde{m})$.

In this paper we have usually insisted that λ be strictly positive, so it is worth noting that for the last lemma that would be an unnecessary stricture.

We now go back to the Floer homology group $HF_*(Y, \mathfrak{t})$ which we incompletely defined above. We can complete the definition by specifying that the closed, fibrewise-symplectic two form on $Y^{[n]}$ should be drawn from the convex set of closed, *fibrewise-Kähler* two-forms, representing the class W_λ for some $\lambda \geq 0$.

The monotonicity property established in the lemma tells us that we can define Floer homology over \mathbb{Z} , not $\Lambda_{\mathbb{Z}}$. Indeed, two index-zero trajectories with the same asymptotic limits must also have the same area (the argument is given in [33], for example).

Thus, adopting these \mathbb{Z} coefficients, $HF_*(Y, \mathfrak{t})$ is a *finitely generated abelian group*. The convexity of the set of allowed two-forms (and the contractibility of the space of allowed vertical almost complex structures) implies that $HF_*(Y, \mathfrak{t})$ is a well-defined group, up to canonical isomorphism.

We have not included λ in the notation, and in fact $HF_*(Y, \mathfrak{t})$ is independent of $\lambda \geq 0$, up to isomorphism. Indeed, it is a general principle (proved by Y-J Lee) that fixed-point Floer homology $HF_*(T, \Omega_0)$ does not change under a deformation $\{\Omega_t\}_{t \in [0,1]}$ provided that (i) the groups are well defined for all t , over the same coefficient ring, and (ii) the periods do not change either. Condition (ii) means, more precisely, that the energy (or action) homomorphism $[\Omega_t]: K \rightarrow \mathbb{R}$ on the subgroup $\ker(c_1(T^*T)) \subset \pi_1(\text{sect}(T), \gamma)$ (where γ is a reference section, acting as a basepoint), should be constant along the deformation. When this is satisfied, the subgroup $HF_*(T, \Omega_t)_\gamma$, corresponding to the component of $\text{sect}(T)$ that contains γ , is independent of t . The proof uses Lee's delicate bifurcation analysis [20]; see also the discussions in [19; 39].

We can make the following conclusion.

Proposition 4.18 *Given Y and $\mathfrak{t} = \mathfrak{t}_\gamma \in \text{Spin}^c(Y)_{2n+\chi(\Sigma)}$, where either $n \geq g$ or $n \leq (g-1)/2$, there is a Floer homology group $HF_*(Y, \mathfrak{t})$ which is a finitely generated abelian group, and is an invariant of (Y, π, \mathfrak{t}) .*

For $Y = S^1 \times \Sigma$, with its trivial fibration over S^1 , the PSS isomorphism relating Floer homology to ordinary homology gives

$$HF_*(Y, \mathfrak{t}_n) \cong H_*(\text{Sym}^n(\Sigma); \mathbb{Z}),$$

as $\mathbb{Z}/2$ -graded abelian groups, where \mathfrak{t}_n is given by $c_1(\mathfrak{t}_n) = (2n + \chi(\Sigma))\text{pr}_2^*[o_\Sigma]$.

One has $HF_*(Y, \mathfrak{t}') = 0$ when $c_1(\mathfrak{t}')$ is not the pullback of a class on Σ .

4.5.3 Extending closed two-forms defined on the boundary There are relative invariants which take the form of homomorphisms with values in the groups $HF_*(Y, \mathfrak{t})$, defined as in (10).

Suppose (X, π) is a broken fibration over a surface-with-boundary S . Label the components of ∂S as incoming or outgoing, and decompose $Y = \partial X$ as $Y_{\text{in}} \cup Y_{\text{out}}$ accordingly.

Recall the procedure discussed in Section 2 as a preamble to Theorem 2.5. As in that discussion, we let S' be the complement of a neighbourhood of $\pi(Z) \subset S$ (where Z is the one-dimensional part of the critical set), and $X' = X|_{S'}$. Let $\mathfrak{s} \in \text{Spin}^c(X)$ be a Spin^c -structure such that $\langle c_1(\mathfrak{s}), [\text{fibre}] \rangle = 2d$. There is then a locally constant function $\nu: S' \rightarrow \mathbb{Z}_{\geq 0}$ such that $\nu(s) + \chi(X_s)/2 = d$ for all regular values s , and we can form the relative Hilbert scheme $X^{[\nu]} \rightarrow S'$.

Recall that the Floer homology $HF_*(Y_{\text{out}}, \mathfrak{t}_{Y_{\text{out}}})$ is defined using a two-form Ω_{out} on $Y_{\text{out}}^{[\nu]}$ with

$$[\Omega_{\text{out}}] = (1 + \lambda\nu) (w_\lambda)^{[1]} - \frac{\lambda}{2} 1^{[2]}$$

where $w_\lambda = \frac{\lambda \text{PD}(\gamma_{\text{in}}) + (2d)^{-1} c_1(\mathfrak{t}_{\gamma_{\text{in}}})}{1 + \lambda\nu}$, $\lambda \geq -\frac{1}{\nu}$.

The same goes for the Floer homology associated with the incoming boundary. We can then extend $[\Omega_{\text{out}}]$ (and its incoming version) to $X^{[\nu]}$ as the class

$$(18) \quad (1 + \lambda\nu) (W_\lambda)^{[1]} - \frac{\lambda}{2} 1^{[2]} \in H^2(X^{[\nu]}; \mathbb{R}).$$

The notation $^{[1]}$ and $^{[2]}$ make sense on the relative Hilbert scheme, as distinct from the relative symmetric product: the definition is as at the beginning of this paper, except that the universal divisor is replaced by the universal sheaf. We refer to Part I, Section 3 for further details.

By Lemma 2.4, this class is represented by an admissible two-form Ω , in the sense of Definition 2.3, provided that $\lambda > 0$. We can moreover arrange that Ω extends both Ω_{out} and its incoming version Ω_{in} .

Finally, Theorem 2.5 applies to give us a Lagrangian matching condition \mathcal{Q} for $(X^{[\nu]}, \pi^{[\nu]}, \Omega)$ associated with \mathfrak{s} . Notice that, according to this prescription, we choose a different two form Ω , and hence are obliged to use a different \mathcal{Q} , for each Spin^c -structure (which, however, we decline to notate). This is the price of working with the canonical, finitely generated Floer homology groups.

4.5.4 Relative Lagrangian matching invariants Using the last paragraph, we conclude that one gets functorial relative invariants

$$\mathcal{L}_{(X,\pi)}(\mathfrak{s}): HF_*(Y_{\text{in}}, \mathfrak{t}_{\text{in}}) \rightarrow HF_*(Y_{\text{out}}, \mathfrak{t}_{\text{out}}),$$

for any admissible Spin^c -structure (cf Part I, Section 1). These are defined as an instance of (10), with no marked points, save for a caveat we shall come to momentarily. The fibration is $X^{[v]} \rightarrow S'$, equipped with a form Ω representing the class defined by formula (18), as discussed in the last paragraph. The Lagrangian matching condition Ω then arises as in Theorem 2.5.

The caveat is that *we only count sections representing the Spin^c -structure \mathfrak{s}* . To be precise, consider the space $\text{sect}(X^{[v]}, \Omega; x_{\text{in}}, x_{\text{out}})$ of sections of the cylindrical-end completion of $(X^{[v]}, \Omega)$, asymptotic to fixed sections x_{in} and x_{out} over the ends. There is a natural map $\sigma_v: \pi_0 \text{sect}(X^{[v]}, \Omega; x_{\text{in}}, x_{\text{out}}) \rightarrow \text{Spin}^c(X)$, defined in just the same way as in (13).

Thus, to define $\mathcal{L}_{(X,\pi)}(\mathfrak{s})$ one counts finite-action, index 0 pseudo-holomorphic sections v of the cylindrical-ends completion of $(X^{[v]}, \Omega)$ such that $\sigma_v(v) = \mathfrak{s}$.

Since we are not using Novikov coefficients, it is important to check that the count is finite. When $\inf(v) \geq 2$, $\sigma_v(v)$ determines the free homotopy class of v over $X^{[v]}$. By “free” we mean that we allow homotopies through sections which do not end on x_{in} and x_{out} . This claim is a straightforward variant of Proposition 4.8. Thus any two sections (v and v' , say) on which σ_v agree differ by an element (u_1, u_2) of $\pi_1(\text{sect}(Y_{\text{in}}, x_{\text{in}})) \times \pi_1(\text{sect}(Y_{\text{out}}, x_{\text{out}}))$, encoding the difference between free and constrained homotopy. We are only interested in rigid (ie, index zero) pseudo-holomorphic sections. For these, the indices of u_1 and u_2 must sum to zero. The monotonicity property (Lemma 4.17) then implies that the energies of v and v' are equal. Gromov–Floer compactness then implies that the count is finite.

4.5.5 Quantum module structure By applying (10) to cylinders $Y \times [0, 1] \rightarrow S^1 \times [0, 1]$, with one marked point in the base, one finds that $HF_*(Y, \mathfrak{t})$ is a module over $\mathbb{Z}[U]$. Here U acts by quantum cap product with a codimension-two cycle:

$$U \cdot c = \delta_x \cap c.$$

It decreases degree by 2 (the degree comes into play when one imposes relative gradings on $HF_*(Y, \mathfrak{t})$). Elements $l \in H_1(\Sigma; \mathbb{Z})$ also act (decreasing degree by 1):

$$l \cdot c = \delta_\gamma \cap c,$$

where γ is a loop in a fibre, representing l . Hence $HF_*(Y, \mathfrak{t})$ becomes a module over the ring

$$\mathbb{Z}[U] \otimes \Lambda^* H_1(Y; \mathbb{Z}).$$

At this point it is easy to check that the groups $HF_*(Y, \mathfrak{t})$ and the elements $\mathcal{L}_{(X, \pi)}(\mathfrak{s})$ give a field theory having the familiar properties set out in the introduction to Part I.

The only point there which we have not yet addressed is the somewhat unorthodox grading properties of the groups; we deal with that next.

5 Geometric gradings for Floer homology groups

Our treatment of the gradings of the Floer homology groups, and of the degrees of the cobordism-maps, is a variation of the conventional method in symplectic Floer theory. It was motivated by treatment of gradings and degrees in Kronheimer–Mrowka’s monopole Floer homology theory (see Kronheimer–Mrowka [16] and Kronheimer–Mrowka–Osváth–Szabó [17]), and it highlights the similarity to that theory.

If Y is an oriented three-manifold, the set $J(Y)$ of homotopy classes of oriented two-plane fields has a convenient algebro-topological description which we now summarise. Spin^c -structures provide a convenient formulation (which seems to originate in [16]; see also the introduction to [17]), but the substance is in Pontrjagin’s homotopy-classification of maps $Y \rightarrow S^2$ [26].

An oriented two-plane field effects a reduction of structure group of $TY \oplus \varepsilon^1$ (ε^1 the trivial real line bundle) from $\text{SO}(3)$ to $\text{U}(1) \subset \text{U}(2)$, and hence determines a Spin^c -structure. We write $J(Y, \mathfrak{t})$ for the set of homotopy classes of oriented two-plane fields underlying $\mathfrak{t} \in \text{Spin}^c(Y)$. One obtains a transitive \mathbb{Z} -action on $J(T, \mathfrak{t})$ by modifying two-plane fields via automorphisms of TY supported in a ball $B^3 \subset Y$: $n \in \mathbb{Z}$ sends $[\xi]$ to $[\alpha_n^* \xi]$, where $\alpha_n: (B^3, \partial B^3) \rightarrow (\text{SO}(3), \{1\})$ is a smooth map of degree $2n$, acting on TY via a trivialisation of TB^3 . When $c_1(\mathfrak{t}_\gamma)$ is torsion, the \mathbb{Z} -action is free; otherwise, its stabiliser is the divisibility $\text{div}(c_1(\mathfrak{t}))$ of $c_1(\mathfrak{t})$ in $H^2(Y)/\text{torsion}$. We declare that $\text{div}(c) = 0$ when c is a torsion class, so that $J(T, \mathfrak{t})$ is always identified, up to a shift, with $\mathbb{Z}/\text{div}(c_1(\mathfrak{t}))$.

Now suppose that X is a cobordism from Y_0 to Y_1 , $j_i \in J(Y_i, \mathfrak{t}_i)$ for $i = 0, 1$, and $\mathfrak{s} \in \text{Spin}^c(X)$ a Spin^c -structure which restricts to $\mathfrak{t}_i \in \text{Spin}^c(Y_i)$. Write

$$j_0 \overset{\mathfrak{s}}{\sim} j_1$$

if there is an almost complex structure I on X which induces the Spin^c -structure \mathfrak{s} and preserves representatives for j_0 and j_1 . Given j_0 and \mathfrak{s} , there is a unique j_1 such that $j_0 \overset{\mathfrak{s}}{\sim} j_1$, so this relation defines a map of \mathbb{Z} -sets $J(X, \mathfrak{s}): J(Y_0, \mathfrak{t}_0) \rightarrow J(Y_1, \mathfrak{t}_1)$.

Kronheimer–Mrowka’s monopole Floer homology group $HM_*(Y, \mathfrak{t})$ is graded by $J(Y, \mathfrak{t})$.⁵ Thus $HM_*(Y, \mathfrak{t})$ is the direct sum of subgroups $HM_j(Y, \mathfrak{t})$ indexed by $j \in J(Y, \mathfrak{t})$; any two homogeneous elements have a relative degree in $\mathbb{Z}/\text{div}(c_1(\mathfrak{t}))$. The cobordism map for (X, \mathfrak{s}) sends HM_{j_0} to $HM_{J(X, \mathfrak{s})(j_0)}$. The following theorem asserts that the same is true in the symplectic Floer homology of symmetric products of fibred three-manifolds.

Theorem 5.1 *Let (Y_0, π_0, J_0) be a bundle of Riemann surfaces over an oriented 1–manifold. Let $\mathfrak{t}_0 \in \text{Spin}^c(Y)$ be an admissible Spin^c –structure. Then $HF_*(Y_0, \mathfrak{t}_0)$ carries a functorial grading by $J(Y, \mathfrak{t}_0)$. Given a broken fibration (X, π) realising a cobordism from Y_0 to Y_1 , and an admissible Spin^c –structure \mathfrak{s} with $\mathfrak{s}|_{Y_i} = \mathfrak{t}_i$, the cobordism map*

$$\mathcal{L}_{(X, \pi)}(\mathfrak{s}): HF_*(Y_0, \mathfrak{t}_0) \rightarrow HF_*(Y_1, \mathfrak{t}_1)$$

sends HF_{j_0} to HF_{j_1} , where $j_1 = J(X, \mathfrak{s})(j_0)$.

To prove this result, we attach to each generator γ for the Floer complex a member θ of the \mathbb{Z} –set of reductions of structure group of the vector bundle $\gamma^*T^vY^{[n]}$ from $\text{Sp}(2n, \mathbb{R})$ to $\text{O}(n)$. These “graded sections” (γ, θ) are related to the two-plane fields involved in the monopole grading through a variant of the Pontrjagin–Thom construction.

Definition 5.2 Let (T^{2n+1}, π, σ) be a locally Hamiltonian fibration (LHF) over S^1 .

- A *grading* for a section $\gamma \in \text{sect}(T)$ is a reduction θ of the structure group of the vector bundle $\gamma^*(T^vT)$ from $\text{Sp}(2n, \mathbb{R})$ to the diagonally embedded subgroup $\text{O}(n)$. Thus θ is the isomorphism class of a pair consisting of a principal $\text{O}(n)$ –bundle $P \rightarrow S^1$ and an isomorphism $P \times_{\text{O}(n)} \mathbb{C}^n \cong \gamma^*T^vT$. The pair (γ, θ) is called a *graded section*.
- A grading for γ determines a Lagrangian subbundle $\Lambda_\theta \subset \gamma^*T^vT$, up to homotopy: $\Lambda_\theta = P \times_{\text{O}(n)} \mathbb{R}^n \subset P \times_{\text{O}(n)} \mathbb{C}^n = \gamma^*T^vT$. We say that (γ, θ) is *even* if Λ_θ is an orientable vector bundle; otherwise it is *odd*.

The space $\text{sect}_{gr}(T)$ of graded sections is a covering space of $\text{sect}(T)$. The fibre is $\pi_1(\text{Sp}(2n, \mathbb{R})/\text{O}(n)) = \mathbb{Z}$. Given a graded section (γ_0, θ_0) and a path $\Gamma = \{\gamma_t\}_{t \in [0, 1]}$ in $\text{sect}(T)$ starting at γ_0 , path lifting gives a grading θ_1 for γ_1 . For any other grading θ'_1 , we write $\delta(\theta_0, \theta'_1)$ for the difference $\theta'_1 - \theta_1$. If θ_0 and θ'_1 are both even, one has

$$\delta(\theta_0, \theta'_1) = 2c_1(\Gamma; \theta_0, \theta'_1),$$

⁵Monopole Floer homology has three basic versions, signified by various decorations, as well as perturbed and local-coefficient variants. The grading properties are common to all of these; we do not, however, consider the completed groups $HM_\bullet(Y, \mathfrak{t})$.

where $c_1(\Gamma; \theta_0, \theta'_1)$ is the relative Chern number—the Euler number of $\Gamma^*T^\vee T$ relative to a trivialisation over the boundary which lifts θ_0 and θ'_1 .

The connection between graded sections and index problems in Floer theory begins to emerge when one looks at the set $\hat{\mathcal{J}}(T, [\gamma])$ of homotopy classes of graded sections of T representing the same class as γ in $\pi_0 \text{sect}(T)$. Recall that the indeterminacy in the relative grading of two homotopic sections γ, γ' which are horizontal (and so represent generators for the Floer complex) is $2N_\gamma\mathbb{Z}$, where

$$N_\gamma\mathbb{Z} = \langle c_1(T^\vee T), \pi_1(\text{sect}(T), \gamma) \rangle.$$

But gradings θ_0, θ_1 for the same section γ are homotopic if and only if

$$\theta_1 - \theta_0 \equiv 0 \pmod{2N_\gamma}.$$

Thus the action of \mathbb{Z} upon $\hat{\mathcal{J}}(Y, [\gamma])$ is transitive with stabiliser $2N_\gamma$.

5.1 Lifting to $\hat{\mathcal{J}}(T, [\gamma])$

We now explain how to assign a grading to each generator for the Floer complex.

A generator is a section γ of $T \rightarrow S^1$ which is horizontal with respect to the connection determined by σ , ie, $\sigma(\dot{\gamma}(t), v) = 0$ for all $t \in S^1$ and all $v \in T_{\gamma(t)}^\vee T$. Because γ is horizontal, there is a canonical symplectic connection on $\gamma^*T^\vee T$. One says that γ is nondegenerate if the monodromy of this connection has no fixed points besides zero. We wish to assign a grading to a nondegenerate horizontal section.

What a nondegenerate horizontal section gives us is

- a symplectic vector bundle $\xi \rightarrow S^1$;
- a symplectic connection ∇ on ξ such that the monodromy L_∇ does not have 1 as an eigenvalue.

∇ determines an $O(n)$ -reduction of ξ as follows: Choose a symplectic vector bundle $\eta \rightarrow \bar{D}$ and an isomorphism $\iota: \eta|_{\partial\bar{D}} \cong \xi$. By trivialising η one obtains, *a fortiori*, a reduction θ_η of ξ to the orthogonal group. However, there is also a Conley–Zehnder index $\text{CZ}(\eta, \nabla) \in \mathbb{Z}$.⁶ We assign to γ the unique grading θ such that

$$\delta(\theta_\eta, \theta) = \text{CZ}(\eta, \nabla).$$

⁶This number is derived from the usual Conley–Zehnder index for paths in the symplectic group as follows. One chooses an extension of $\iota^*\nabla$ to a symplectic connection on ξ and a trivialisation of η along a radial arc, whereupon one obtains a path in the symplectic group from the identity to the monodromy of ∇ . This has a CZ index.

Notice that the parity of θ coincides with the usual parity (or mod 2 Conley–Zehnder index) of (ξ, ∇) , derived from the sign of $\det(\text{id} - L_\nabla)$.

Proposition 5.3 *The grading θ does not depend on the extension η .*

Proof Suppose η_0 and η_1 are two such extensions, inducing gradings θ_0, θ_1 . Glue the two copies of the disc together along their common boundary, reversing the orientation on the first copy; likewise, glue the two vector bundles together to obtain $\eta \rightarrow S^2$. We can evaluate the Chern number $2c_1(\eta)$ in two ways. First, by the additivity properties of the Conley–Zehnder index,

$$2c_1(\eta) = \text{CZ}(\eta_1, \nabla) - \text{CZ}(\eta_0, \nabla).$$

Second, let f_0, f_1 be symplectic trivialisations of ξ which lift the gradings θ_0, θ_1 (if both are even) or $\theta_0 + 1, \theta_1 + 1$ (if both are odd). Then

$$\begin{aligned} 2c_1(\eta) &= 2c_1(\eta_1; f_1) - 2c_1(\eta_0; f_1) \\ &= 2c_1(\eta_1; f_1) - 2c_1(\eta_0; f_0) + (\theta_0 - \theta_1) \\ &= (\theta_1 - \theta_{\eta_1}) - (\theta_0 - \theta_{\eta_0}) + (\theta_0 - \theta_1) \\ &= \text{CZ}(\eta_1, \nabla) - \text{CZ}(\eta_0, \nabla) + (\theta_0 - \theta_1). \end{aligned}$$

Hence $\theta_1 - \theta_0 = 0$. □

We have therefore established a grading map $\gamma \mapsto \hat{\gamma} = (\gamma, \theta)$. It is evidently compatible with the Floer-theoretic index: Suppose one has an LHF (E, π, Ω) over the cylinder $S^1 \times [0, 1]$, realising a cobordism between nondegenerate locally Hamiltonian fibre bundles (Y_0, π_0, σ_0) and (Y_1, π_1, σ_1) . Let Γ be a smooth section such that $\partial\Gamma = \gamma_1 - \gamma_0$ is horizontal. Then one has

$$\delta(\hat{\gamma}_0, \hat{\gamma}_1) = \text{CZ}(\Gamma).$$

5.1.1 Relation to two-plane fields We now specialise to surface-bundles $Y \rightarrow S^1$ and their relative symmetric products $Y^{[n]} = \text{Sym}_{S^1}^n(Y)$. A section $\gamma \in \text{sect}(Y^{[n]})$ determines a homology class $[\gamma] \in H_1(Y; \mathbb{Z})$, and therefore a Spin^c -structure $\mathfrak{t}_\gamma = \text{PD}(\gamma) \cdot \mathfrak{t}_{\text{can}}$.

The object of this section is to compare the \mathbb{Z} -sets $\hat{J}(Y, [\gamma])$ and $J(Y, \mathfrak{t}_\gamma)$. A first task is to compare the stabilisers $2N_\gamma\mathbb{Z}$ and $\text{div}(c_1(\mathfrak{t}_\gamma))\mathbb{Z}$.

Lemma 5.4 *$\text{div}(c_1(\mathfrak{t}_\gamma))$ divides $2N_\gamma$, which in turn divides $2(n + 1 - g)$.*

Proof For a class $\Gamma \in H_2(Y^{[n]}; \mathbb{Z})$ represented by a loop $\{\gamma_t\}_{t \in S^1}$ of sections $\gamma_t \in \text{sect}(Y^{[n]})$ one has, by Equation (17),

$$\langle c_1(T^\vee Y^{[n]}), \Gamma \rangle = \frac{1}{2} \langle c_1(T^\vee Y) + 2h, p(\Gamma) \rangle$$

where $p(\Gamma) \in H_2(Y; \mathbb{Z})$ is the class tautologically determined by Γ , and $h \in H^2(Y; \mathbb{Z}) = H_1(Y; \mathbb{Z})$ that determined by γ_0 . Thus $\langle c_1(T^\vee Y^{[n]}), \Gamma \rangle = \langle c_1(\mathfrak{t}_\gamma), p(\Gamma) \rangle / 2$, and $2N_\gamma$ is a multiple of $\text{div}(\mathfrak{t}_\gamma)$.

That N_γ divides $n + 1 - g$ is due to the fact that $c_{\min}(\text{Sym}^n(\Sigma)) = |n + 1 - g|$. \square

We define a map

$$t: \widehat{J}(Y, [\gamma]) \rightarrow J(Y, \mathfrak{t}_\gamma)$$

in three stages:

- (1) To define $t(\gamma, \theta)$, we first perturb γ so that it does not meet the diagonal. It then defines an embedded 1-manifold in Y , and the grading is represented by a line subbundle l_θ of its normal bundle.
- (2) The $[-\epsilon, \epsilon]$ -subbundle of l_θ exponentiates to a surface in Y . The boundary of this surface is a 1-submanifold γ' with a natural (outward-pointing) framing. We call γ' the framed double of (γ, θ) .
- (3) We can twist $T^\vee Y$ along γ' (as we can for any framed 1-submanifold transverse to $T^\vee Y$) using the inverse of the Pontrjagin–Thom construction. The framing determines a tubular neighbourhood $\gamma' \times D^2 \hookrightarrow Y$. Projection onto D^2 , followed by the collapsing map $D^2 \rightarrow D^2 / \partial D^2 = S^2$ gives a map of the neighbourhood to S^2 which extends to a map $\alpha: Y \rightarrow S^2$. In the presence of a trivialisation of TY , any two plane-fields ξ, ξ' differ by a map $\alpha: Y \rightarrow \text{SO}(3) / \text{SO}(2) = S^2$. Conversely, given a map $\alpha: Y \rightarrow S^2$, one can twist the plane-field ξ to a new plane-field $\xi' = \alpha \cdot \xi$. Applying this twisting to the neighbourhood of γ' , with the trivialisation determined by the framing, and the map $\alpha: Y \rightarrow S^2$, we obtain a plane-field $\alpha \cdot T^\vee Y$. This field is $t(\gamma, \theta)$.

The difference in Euler classes, $e(t(\gamma, \theta)) - e(T^\vee Y)$, is Poincaré dual to $[\gamma'] = 2[\gamma]$, hence $t(\gamma, \theta)$ does represent \mathfrak{t}_γ .

Lemma 5.5 t is a map of \mathbb{Z} -sets.

Proof It is evident that there is a universal constant c such that $t(\gamma, \theta + 1) = t(\gamma, \theta) + c$. The Pontrjagin–Thom map is a map of \mathbb{Z} -sets: when one adds a twist to the framing of a submanifold, the resulting two-plane field changes by $+1$. But adding two to the grading θ results in adding two twists to the framing of γ' (one twist for each component, if θ is even), hence $c = 1$. \square

5.1.2 Cobordism maps Suppose (X, π, Ω) is a Lefschetz fibration over a surface with boundary, realising a cobordism between (Y_0, π_0, σ_0) and (Y_1, π_1, σ_1) .

Fix sections $\alpha_i \in \text{sect}(Y_i)$; $r \in \mathbb{N}$; and a component β of the space of sections Γ of $X^{[n]}$ with $\partial\Gamma = \alpha_1 - \alpha_0$. A grading of α_0 extends to a unique trivialisation of $\Gamma^*T^\vee X^{[n]}$, up to homotopy, and this defines a map

$$\hat{J}(X, \beta): \hat{J}(Y_0, \alpha_0) \rightarrow \hat{J}(Y_1, \alpha_1).$$

This map encodes the degree of the cobordism map $HF(X, \beta)$ on Floer homology, in the sense that $\mathcal{L}_{(X, \pi)}(\mathfrak{s}_\beta)$ maps $HF_\theta(Y_0, \mathfrak{t}_{\alpha_0})$ into $HF_{\hat{J}_{X, \beta}(\theta)}(Y_1, \mathfrak{t}_{\alpha_1})$.

Lemma 5.6 *The diagram*

$$\begin{array}{ccc} \hat{J}(Y_0, [\gamma_0]) & \xrightarrow{\hat{J}_{X, \beta}} & \hat{J}(Y_1, [\gamma_0]) \\ t \downarrow & & \downarrow t \\ J(Y_0, \mathfrak{t}_{\gamma_0}) & \xrightarrow{J_{X, \mathfrak{s}_\beta}} & J(Y_1, \mathfrak{t}_{\gamma_1}) \end{array}$$

commutes.

Proof It suffices to construct an almost complex structure, representing \mathfrak{s}_β and preserving the two-plane fields $t(\gamma_i, \theta_i)$, where θ_0 and θ_1 extend to a grading $\{\theta_t\}_{t \in [0, 1]}$ for a homotopy $\Gamma = \{\gamma_t\}_{t \in [0, 1]}$.

We may assume that γ_0 and γ_1 are disjoint from the diagonals in $Y_i^{[r]}$, and that Γ intersects the strata of the diagonal $\Delta \subset \text{Hilb}_S^n(X)$ transversely. Then Γ hits only the top stratum of Δ , and the intersections are isolated points. The subset $\Gamma' \subset X$ determined by Γ is an embedded surface, and the intersections with Δ appear as points $\{p_1, \dots, p_k\}$ where Γ' is tangent to the fibres of X .

Suppose, for the time being, that Γ is disjoint from the diagonal. We may then take Γ' to be pseudo-holomorphic, with respect to an almost complex structure I on X preserving the vertical distribution $\ker(D\pi)$. The idea is to define an new almost complex structure by “twisting along the double of Γ' ”. The procedure is similar to the one used to define the map t . One obtains a new almost complex structure $\alpha \cdot I := \alpha I \alpha^{-1}$ from I from any map $\alpha: X \rightarrow \text{SO}(TX)/\text{U}(TX, I)$. But along Γ' we can choose a unitary trivialisation of TX (by trivialising $T\Gamma'$ and $T^\vee X|_{\Gamma'}$). Then, near Γ' , $\text{SO}(TX)/\text{U}(TX, I)$ is identified with $\text{SO}(4)/\text{U}(2) = S^2$. In a tubular neighbourhood $D^2 \times \Gamma'$ of Γ' , we take α to be the composite $(D^2 \times \Gamma', \partial D^2 \times \Gamma') \rightarrow (D^2, \partial D^2) \rightarrow (S^2, *) \rightarrow (\text{SO}(4)/\text{U}(2), [\text{id}])$. This defines a twisted almost complex structure $I' = \alpha \cdot I$, compatible with the twisted tangent distribution on ∂X .

It remains to deal with the vertical tangencies. The failure of the square to commute is given $N\Gamma \cdot \Delta$, where N is some constant. It follows that, for Lefschetz fibrations on closed four-manifolds, one has $\text{ind}(\Gamma) - d(\mathfrak{s}_\Gamma) = N\Gamma \cdot \Delta$. Hence $N = 0$ by Lemma 4.11 (the Lefschetz fibration case of Theorem D). \square

The proof of Theorem 5.1 is now almost complete; all that remains is to check functoriality for broken fibrations. If (X, π) is a broken fibration, realising a cobordism from Y_0 to Y_1 , and $\beta \in \pi_0 \text{sect}(X^{[v]}, \mathcal{Q})$, we define $\hat{J}(X, \beta)$ by sending (γ_0, θ_0) to the (γ_1, θ_1) , where $\gamma_1 - \gamma_0 = \partial\Gamma$ for a representative Γ of β , and there is a trivialisation of $\Gamma^*T^\nu X^{[v]}$ compatible with γ_0, γ_1 and $T^\nu\mathcal{Q}$. Then $\hat{J}(X, \beta)$ defines the degree of the map $\Phi(X, \beta)$ on Floer homology, and we must compare $t \circ \hat{J}(X, \beta)$ with $J_{X, \mathfrak{s}_\beta} \circ t$. Again, the fact (Theorem D) that $\text{ind}(\beta)$ equals $d(\mathfrak{s}_\beta)$ when X is closed implies that the discrepancy must be zero.

6 Calculations

We now compute the Lagrangian matching invariants in some simple cases, and show that they coincide with Seiberg–Witten invariants. The calculations presented here are limited to situations where (i) the monodromies around either side of each circle of critical values are trivial, or (ii) one considers Sym^n only for $n \leq 1$.

6.1 First examples

6.1.1 Trivial bundle over a disc Let $\Sigma \times \Delta \rightarrow \Delta$ be the trivial bundle over the unit disc. Let \mathfrak{s}_d be the Spin^c -structure with $\langle c_1(\mathfrak{s}), [\Sigma] \rangle = d = \chi(\Sigma) + 2n$, and \mathfrak{t}_d its restriction to $\Sigma \times S^1$. When \mathfrak{s}_d is admissible, the relative invariant is an element

$$\mathcal{L}(\mathfrak{s}_d) \in HF_*(\Sigma \times S^1, \mathfrak{t}_d),$$

the “fundamental class”. The ring $\mathbb{Z}[U] \otimes \Lambda^* H_1(\Sigma; \mathbb{Z})$ operates on $HF_*(\Sigma \times S^1, \mathfrak{t}_d)$. Indeed, $l \in H_1(\Sigma; \mathbb{Z})$ acts by $l \cdot x = \mu(l) \cap x$, where $\mu: H_1(\Sigma; \mathbb{Z}) \rightarrow H^1(\text{Sym}^n(\Sigma); \mathbb{Z})$ is the usual μ -isomorphism (see, eg, Part I, Section 3.6), and \cap is quantum cap product. The degree two element U acts by $U \cdot x = \eta \cap x$ (η as in the introduction to this paper). By the PSS isomorphism, the elements $m \cdot \mathcal{L}(\mathfrak{s}_d)$ give a basis for $HF_*(\Sigma \times S^1)$ as m runs over monomials

$$m = U^i \otimes l_1 \wedge \cdots \wedge l_{n-i}, \quad 0 \leq i \leq n.$$

So we have an isomorphism

$$\begin{aligned} \Phi: \mathbb{Z}[U] \otimes \Lambda^* H_1(\Sigma; \mathbb{Z}) &\rightarrow HF_*(\Sigma \times S^1, \mathfrak{t}_d), \\ \eta^i \otimes l_1 \wedge \cdots \wedge l_{n-i} &\mapsto U^i \cdot (l_1 \cdots l_{n-1} \cdot \mathcal{L}(\mathfrak{s}_d)). \end{aligned}$$

If we reverse the orientation of the disc, we obtain a map $\mathcal{L}^\vee: HF_*(\Sigma \times S^1, \mathfrak{t}_d) \rightarrow \mathbb{Z}$. This sends $U^n \cdot \mathcal{L}$ to 1, and all the other basis elements $m \cdot \mathcal{L}$ to 0.

We can describe the action of $\mathbb{Z}[U] \otimes \Lambda^* H_1(\Sigma; \mathbb{Z})$ on $HF_*(\Sigma \times S^1, \mathfrak{t}_d)$ by quoting Bertram and Thaddeus' formulae [2] for the quantum cup product on $\text{Sym}^n(\Sigma)$. We have

$$l \cdot \Phi(c) = \Phi(l \cdot c), \quad l \in H_1(\Sigma; \mathbb{Z})$$

(the quantum product with classes of odd degree is undeformed, as always). We also have

$$U \cdot \Phi(c) = \Phi(\eta \cdot c), \quad n \leq (g - 1)/2$$

for dimension reasons. However,

$$U \cdot \Phi(\eta^i) = \Phi(\eta^{i+1} + \theta_{g-n+i} - \theta_{g-n} \eta^i), \quad n \geq g > 0,$$

where $\theta_m = \theta^m/m!$ for $m \geq 0$, and $\theta_m = 0$ for $m < 0$ (recall that $\theta \in \Lambda^2 H_1(\Sigma) = \Lambda^2 H^1(\Sigma)^*$ is the cup-product form).

When $g = 0$, one has $U \cdot \eta^i = \eta^{i+1} + \eta^{i-n}$ (negative powers are read as zero). In particular,

$$U^{n+1} \cdot = \text{id}.$$

6.1.2 The trivial S^2 -bundle over S^2 This example, the fibre bundle

$$\text{pr}_2: S^2 \times S^2 \rightarrow S^2,$$

is included just to show how to use the field theory before it gets mixed up with other ingredients. We use the basis $([S^2 \times \{\text{pt}\}], [\{\text{pt}\} \times S^2])$ for homology. Consider the homology class (m, n) , $n \geq 1$ and its associated Spin^c -structure $\mathfrak{s}(m, n)$. It has $c_1(\mathfrak{s}) = \text{PD}(2 + 2m, 2 + 2n)$ and $d(\mathfrak{s}(m, n)) = 2[mn + m + n]$.

We compute $\mathcal{L}_{(S^2 \times S^2, \text{pr}_2)}(\mathfrak{s}(m, n))$ by splitting the two-sphere into the southern hemisphere D_- and northern hemisphere D_+ . We have

$$\mathcal{L}_{S^2 \times S^2, \text{pr}_2}(\mathfrak{s}(m, n)) = \mathcal{L}_{D_+}(U^{mn+m+n} \cdot \mathcal{L}_{D_-}) = \mathcal{L}_{D_+}(U^n \cdot U^{m(n+1)} \cdot \mathcal{L}_{D_-}).$$

Here \mathcal{L}_{D_-} is the invariant associated with the southern hemisphere D_- ,

$$\mathcal{L}_{D_-} \in HF_*(S^2 \times S^1, \mathfrak{t}_{2+2n}),$$

where $\langle c_1(\mathfrak{t}_k), [S^2 \times \{\text{pt}\}] \rangle = k$. By the PSS isomorphism,

$$\{\mathcal{L}_{D_-}, U \cdot \mathcal{L}_{D_-}, \dots, U^n \cdot \mathcal{L}_{D_-}\}$$

is a \mathbb{Z} -basis for $HF_*(S^2 \times S^1, \mathfrak{t}_{2+2n})$. Because of the structure of the quantum cohomology just quoted, $U^{n+1} \cdot \mathcal{L}_{D_-} = \mathcal{L}_{D_-}$, and $U^n \cdot U^{m(n+1)} \cdot \mathcal{L}_{D_-} = U^n \cdot \mathcal{L}_{D_-} =$

$\Phi([\eta^n])$ is a primitive element of Floer homology. Moreover, this element maps to $1 \in \mathbb{Z}$ under $\mathcal{L}_{D_+}: HF_*(S^2 \times S^1, t_{2+2n}) \rightarrow \mathbb{Z}$.

Hence

$$\mathcal{L}_{S^2 \times S^2, \text{pr}_2}(\mathfrak{s}(m, n)) = \begin{cases} U^{(m+1)(n+1)-1}, & n \geq 1, m \geq 0 \\ 0, & n \geq 1, m < 0. \end{cases}$$

We can compare this result with the Seiberg–Witten invariants. By the wall-crossing formula, we have

$$SW_{S^2 \times S^2}^\tau(\mathfrak{s}(m, n)) = \begin{cases} U^{(m+1)(n+1)-1}, & n \geq 0, m \geq 0 \\ 0, & n \geq 1, m < 0. \end{cases}$$

The superscript τ designates Taubes’ chamber for a symplectic form $\omega_{S^2} \oplus \omega_{S^2}$.

6.1.3 A broken fibration on $(S^1 \times S^3) \# (S^2 \times S^2)$ This broken fibration, $\pi: X \rightarrow S^2$, was introduced in Example 4.13. Recall its structure:

- Over the southern hemisphere, $X_- \rightarrow D_-$ is a trivial torus-bundle. The vanishing torus $Q \subset \partial X_-$ is $l \times \partial D_-$, where $l \subset T^2$ is the projection of a line in \mathbb{R}^2 to $\mathbb{R}^2/\mathbb{Z}^2$, so Q is Lagrangian with respect to the standard, product symplectic form.
- Over an equatorial annulus, $X_0 \rightarrow A$ is the elementary broken fibration constructed from Q .
- Over the northern hemisphere, $X_+ \rightarrow D_+$ is a trivial S^2 -bundle. There are two possible ways to glue ∂X_+ to ∂X_0 , corresponding to the two elements of $\pi_1(\text{SO}(3))$. We choose the “untwisted” gluing, for which the union $X = X_+ \cup X_0 \cup X_-$ is diffeomorphic to $(S^1 \times S^3) \# (S^2 \times S^2)$ (the other gluing gives $(S^1 \times S^3) \# \mathbb{P}^2 \# \overline{\mathbb{P}^2}$; see Auroux–Donaldson–Katzarkov [1, Section 8.2]).

For the standard complex structure J_0 on $T^2 \times D_-$, the only holomorphic sections of $X_- \rightarrow D_-$ with boundary on Q are the constant sections c_x , $x \in l$, and since $\text{ind}(c_x) = 1$, J_0 is regular. These constant sections represent a class $\beta \in H_2(X, \mathbb{Z}; \mathbb{Z})$ with $\beta \cdot [T^2] = 1$. The corresponding Spin^c -structure \mathfrak{s}_β has $c_1(\mathfrak{s}_\beta) = (2, 2)$ in the basis for $H_2(X; \mathbb{Z})$ given by (F, S) , where F is the fibre class and S is represented by a square-zero section. Hence

$$\mathcal{L}_{X, \pi}(\mathfrak{s}_\beta) = \pm 1 \otimes \lambda \in \mathbb{A}(X) = \mathbb{Z}[U] \otimes \Lambda^* H^1(X; \mathbb{Z}),$$

where λ is the generator for $H^1(X; \mathbb{Z}) = \mathbb{Z}$. This is in agreement with the Seiberg–Witten invariant $\text{SW}_X(\mathfrak{s}_\beta)$, calculated in the Taubes chamber of a compatible near-symplectic form.

We can compute the invariants for more general Spin^c -structures using our field theory. Note that the Lagrangian matching invariants for n -fold sections of X_- and $(n-1)$ -fold sections of X_+ are well defined for any $n > 0$.

Write $\mathfrak{s}_{p,q}$ for the unique Spin^c -structure on X with $c_1(\mathfrak{s}_{p,q}) = (2p, 2q)$. The relative homology class $\beta + (n-1)S + mF \in H_2(X, Z; \mathbb{Z})$ has associated Spin^c -structure $\mathfrak{s}_{\beta+(n-1)S+mF} = \mathfrak{s}_{m+1,n}$. The virtual dimension of the moduli space is $d(\mathfrak{s}_{m+1,n}) = 2n(m+1) - 1$.

We shall calculate the invariant of $\mathfrak{s}_{m+1,n}$ ($m \geq 0$) by splitting S^2 into its three parts D_- , A and D_+ . By the gluing rule,

$$\mathcal{L}_{X,\pi}(\mathfrak{s}_{m+1,n}) = \mathcal{L}_{D_+} \circ (U^{n(m+1)-1} \cdot \mathcal{L}_A) \circ (l \cdot \mathcal{L}_{D_-}).$$

The relative invariant of D_- is the element $\mathcal{L}_{D_-} = \Phi(1) \in HF_*(T^2 \times S^1, \mathfrak{t}_{2n})$, and $l \cdot \mathcal{L}_{D_-} = \Phi(l)$. The relative invariant \mathcal{L}_A for $X_0 \rightarrow A$ is a map

$$\begin{aligned} H_*(\text{Sym}^n(T^2); \mathbb{Z}) &= HF_*(T^2 \times S^1, \mathfrak{t}_{2n}) \\ &\rightarrow HF_*(S^2 \times S^1, \mathfrak{t}_{2n}) \cong H_*(\text{Sym}^{n-1}(S^2); \mathbb{Z}). \end{aligned}$$

The map \mathcal{L}_A is identified with the map on homology of symmetric products induced by the fundamental class of the correspondence \widehat{V} . Monotonicity ensures that there are no ‘‘quantum corrections’’: two holomorphic sections with the same asymptotic limits and the same index must also have the same energy, and the horizontal sections are exactly the ones with zero energy. Thus \mathcal{L}_A sends $\Phi(l)$ to $\Phi(1)$.

Next, $U^{nm+n-1} \cdot \Phi(1) = U^{n-1} \cdot \Phi(1) = \Phi(\eta^{n-1})$. Finally, \mathcal{L}_{D_+} maps $\Phi(\eta^{n-1})$ to 1.

We deduce that

$$\mathcal{L}_{X,\pi}(\mathfrak{s}_{m+1,n}) = \begin{cases} \pm U^{n-1} \otimes \lambda, & m \geq 0, n \geq 0; \\ 0, & m < 0, n \geq 0. \end{cases}$$

This is again in agreement with the Seiberg–Witten invariants, calculated in the Taubes chamber of a compatible near-symplectic form.

The same method applies also to the broken fibration on $(S^1 \times S^3) \# \mathbb{P}^2 \# \overline{\mathbb{P}^2}$ mentioned above.

6.1.4 Connected sum with \mathbb{P}^2 This is one of the topological examples from [1]. Start with a broken fibration $\pi: X \rightarrow S$ with an isolated critical point c . Let $D \subset S$ be a small disc containing $\pi(c)$. There is a torus $Q \subset \pi^{-1}(\partial D)$, the union of the vanishing cycles for radial paths in D . Then there is a broken fibration $\pi': X \rightarrow S'$

which coincides with π outside D , and which contains a circle Z of critical points mapping to a small circle in D . The manifold X' is diffeomorphic to $X \# \mathbb{P}^2$.

Let $\beta \in H_2(X', Z; \mathbb{Z})$ be a class with $\beta \cdot [\Sigma] = 1$, where Σ is a fibre over a point in ∂D (so $\beta \cdot [\bar{\Sigma}] = 0$ for fibres $\bar{\Sigma}$ over the centre of D).

Proposition 6.1 *We have $\mathcal{L}_{X', \pi'}(\mathfrak{s}_\beta) = 0$.*

Proof Choose a near-symplectic form ω on X' such that Q is Lagrangian. The monodromy of π' around ∂D is a Dehn twist $\tau_L \in \text{Aut}(\Sigma)$ about the circle $L = Q \cap \Sigma$. We analyse the element $\mathcal{L} \in HF_*(Y, \mathfrak{t})$, where Y is the mapping torus of the Dehn twist τ_L on Σ and \mathfrak{t} the restriction of \mathfrak{s}_β .

Notice first that $\mathcal{L} \in HF_{\text{odd}}(Y, \mathfrak{t})$. It is known by Seidel's work [31] that $HF_{\text{odd}}(Y, \mathfrak{t}) \cong H_1(\Sigma, L; \mathbb{Z})$, by an isomorphism which is equivariant under the action by symplectic automorphisms of Σ on the two sides. If $\delta \in \text{Aut}(\Sigma, \omega|_\Sigma)$ acts as the identity near L , then $\delta_*(\mathcal{L}) = \mathcal{L}$, because δ induces an automorphism of the LHF $Y \rightarrow S^1$ which acts trivially on Q . But we can obtain any element of $\text{Sp}(H_1(\Sigma); \mathbb{Z})$ which fixes $[L]$ from such a δ , so there are no nonzero elements of $H_1(\Sigma, L; \Lambda_{\mathbb{Z}})$ fixed by all the automorphisms δ . The result follows. \square

6.1.5 A broken fibration on $(S^1 \times S^3) \# (S^2 \times S^2) \# \overline{\mathbb{P}^2}$ In this example we take a step in the direction of an ‘‘isotropic blow-up’’ formula for Lagrangian matching invariants. This involves a model for the differentiable blow-up of a broken fibration described in [1]. This model is unusual in that there is a compatible near-symplectic form for which the exceptional sphere is *isotropic*. The sphere lies over an arc in the base connecting an isolated critical value and a point on a circle of critical values.⁷

Rather than studying the isotropic blow-up process in general, we apply it to our earlier example on $(S^1 \times S^3) \# (S^2 \times S^2)$ and to a variant of that example on $(S^1 \times S^3) \# \mathbb{P}^2 \# \overline{\mathbb{P}^2}$. Let $\hat{\pi}: \hat{X} \rightarrow S^2$ be either of the following two broken fibrations:

- The set of critical values is the union of the equator and the point $0 \in D_-$.
- The regular fibres over D_- have genus 1. In the fibre $T^2 = \hat{\pi}^{-1}(1)$, the vanishing cycle for the path $[0, 1]$ is a nonseparating circle L . The vanishing surface $Q \subset \hat{\pi}^{-1}(\partial D_-)$ has $Q \cap T^2 = L$.
- The fibres over D_+ have genus 0, so $\hat{X}_+ := \hat{\pi}^{-1}(D_+)$ is a trivial bundle $S^2 \times D_+ \rightarrow D_+$.

⁷In four-dimensional symplectic geometry, Lagrangian two-spheres have self-intersection -2 , but in near-symplectic geometry there are other possibilities, including isotropic (-1) -spheres.

- There are two possible ways of gluing $\partial\hat{X}_+$ to \hat{X}_0 . One of these has total space $\hat{X} \cong (S^1 \times S^3) \# (S^2 \times S^2) \# \overline{\mathbb{P}^2}$, the other $\hat{X} \cong (S^1 \times S^3) \# \mathbb{P}^2 \# \overline{\mathbb{P}^2} \# \overline{\mathbb{P}^2}$.

We will only consider those homology classes $\beta \in H_2(\hat{X}, \mathbb{Z}; \mathbb{Z})$ with $\beta \cdot [T^2] = 1$, so the fibration over D_+ is irrelevant. A useful local model is the fibration

$$q: (z, w) \mapsto z^2 + w^2$$

on $E^r = \{(z, w) \in \mathbb{C}^2 : |q(z, w)| \leq r, (|z|^2 + |w|^2)^2 - |q(z, w)|^2 \leq \epsilon\}$, with its standard symplectic form $\omega_{\mathbb{C}^2}$ and Lagrangian boundary condition

$$Q^r = \bigcup_{s \in \partial\bar{D}(0; r)} Q_s^r, \quad Q_s^r = s^{1/2} S^1$$

(where S^1 is the unit circle in $\mathbb{R}^2 \subset \mathbb{C}^2$). This model was carefully analysed by Seidel [34, Section 2.3]. For the standard almost complex structure, the holomorphic sections with boundary on Q^r are given by

$$u_{a, \pm}^r(s) = (as + \bar{a}, \pm i(as - \bar{a})), \quad |a|^2 = r/2.$$

Moreover, all the sections $u_{a, \pm}^r$ are regular.

We can embed E^1 as a subfibration of $\hat{X}_- \rightarrow D_-$ (its complement is then a trivial fibration). We are then free to shrink the base from $D_- = \bar{D}(0; 1)$ to $\bar{D}(0; r)$, using the Lagrangian boundary condition Q^r . By the degeneration argument of [34, Lemmas 2.14 and 2.15] when r is sufficiently small, every pseudo-holomorphic section of $\hat{X}_-|_{\bar{D}(0; r)} \rightarrow \bar{D}(0; r)$ lies inside E^r , and is therefore one of the $u_{a, \pm}^r$.

It is important to note that the circles in Q^r traced out by $u_{a, +}^r$ and $u_{a, -}^r$ are *not* mutually homologous, since they intersect transversely at a single point. Hence when we embed $E(r)$ into X_- , $u_{a, +}^r$ and $u_{a, -}^r$ represent *distinct* homology classes β_+ , β_- . Thus we have moduli spaces of sections $\mathcal{M}_{\beta_+}^r = \{u_{a, +}^r : |a|^2 = r/2\}$ and $\mathcal{M}_{\beta_-}^r = \{u_{a, -}^r : |a|^2 = r/2\}$. Each is diffeomorphic to S^1 (by $u_{a, \pm}^r \mapsto \sqrt{2}r^{-1/2}a$) but precisely one of these diffeomorphisms is orientation-preserving (another consequence of Seidel's degeneration argument, which shows that $\mathcal{M}_{\beta_-}^r \cup \mathcal{M}_{\beta_+}^r$ is the boundary of a moduli space $\mathcal{M}' \cong S^1 \times [-1, 1]$). Let $\lambda \subset T^2$ be a loop representing the generator of $\pi_1(\hat{X})$, and intersecting L transversely at a point. Then the evaluation maps $\text{ev}_+ : \mathcal{M}_{\beta_+}^r \rightarrow L$ and $\text{ev}_- : \mathcal{M}_{\beta_-}^r \rightarrow L$ are both transverse to λ . The fibre product moduli spaces, $\mathcal{M}_{\beta_{\pm}}^r \times_{\text{ev}_{\pm}} \lambda$, are singletons, each consisting of one point, counted with the same signs. Hence

$$\mathcal{L}_{\hat{X}, \hat{\pi}}(\mathfrak{s}_{\beta_{\pm}}) = \varepsilon[\lambda],$$

where $\varepsilon = \pm 1$ is a common sign. Moreover, $\mathcal{L}_{\widehat{X}, \widehat{\pi}}(\mathfrak{s}\beta') = 0$ for every other class β' with $\beta' \cdot T^2 = 1$.

To compare with Seiberg–Witten theory, let us take $\widehat{X} \cong (S^1 \times S^3) \# (S^2 \times S^2) \# \overline{\mathbb{P}^2}$. We use the basis (S, S', E) for $H_2(\widehat{X}; \mathbb{Z})$, where $S = [S^2 \times \{\text{pt}\}]$, $S' = [\{\text{pt}\} \times S^2]$, E the class of a (-1) -sphere in $\overline{\mathbb{P}^2}$. For the Spin^c -structure $\mathfrak{s}_{p,q,k}$ with $c_1(\mathfrak{s}_{p,q,k}) = 2pS + 2qS' + (2k+1)E$, one has $d(\mathfrak{s}_{p,q,k}) = 2pq - k(k+1) - 1$. For the classes β_{\pm} one has $p = q = 1$ and $d(\mathfrak{s}_{\beta_{\pm}}) = 1$, and it follows that $\{\mathfrak{s}_{\beta_+}, \mathfrak{s}_{\beta_-}\} = \{\mathfrak{s}_{1,1,0}, \mathfrak{s}_{1,1,1}\}$. By the blow-up formula for Seiberg–Witten invariants, one has

$$\text{SW}_{\widehat{X}}^{\tau}(\mathfrak{s}_{1,1,k}) = \begin{cases} \pm[\lambda], & k \in \{0, 1\}, \\ 0, & k \notin \{0, 1\}. \end{cases}$$

6.2 Separating model for $X_1 \# X_2$

There is a simple realisation of the connected sum (*not* fibre sum!) of broken fibrations.

Suppose $\pi_i: X_i \rightarrow S$, $i = 1, 2$, are broken fibrations over the same base. Take a simultaneous regular value $s \in S$, and a small coordinate disc $(\Delta, 0) \hookrightarrow (S, s)$. There is then a broken fibration $\pi: X \rightarrow S$ such that

- $\pi^{-1}(S \setminus \Delta) = \pi_1^{-1}(S \setminus \Delta) \sqcup \pi_2^{-1}(S \setminus \Delta)$, and on this region π is the disjoint union of π_1 and π_2 ;
- $X^{\text{crit}} \cap \pi^{-1}(\Delta)$ is a circle;
- s is a regular value and the fibre $X_s = \pi^{-1}(s)$ is the connected sum of fibres $\pi_1^{-1}(s) \# \pi_2^{-1}(s)$.

The construction is best understood pictorially (see Figure 1).

Lemma 6.2 $X \cong X_1 \# X_2$.

Proof We need to find a separating three-sphere. We work over the unit disc Δ , taking the circle of critical values to be $\{|z| = 1/2\}$. As the connected sum of two surfaces, the central fibre $\pi^{-1}(0)$ contains a separating circle L (drawn in the figure). Over each ray $\{re^{i\theta} : 0 \leq r \leq 1/2\} \subset \Delta$, there is a disc in X bounded by L and containing exactly one critical point of π . Let Y be the union (over $\theta \in [0, 2\pi]$) of these discs. It is the union of a copy of $D^2 \times S^1$ (lying over $\{|z| \leq 1/4\}$) and a copy of $S^1 \times D^2$ (over $1/4 \leq |z| \leq 1/2$), joined in the obvious way along their common boundary $S^1 \times S^1$. Hence Y is diffeomorphic to S^3 . The complement $X \setminus Y$ breaks into two components $X'_i \cong X_i \setminus B^4$, so $X \cong X_1 \# X_2$. \square




Figure 1: Connected sum of broken fibrations. The first diagram shows the disjoint union of two surface-bundles over the disc. One converts them the disjoint union into a broken fibration as in the second diagram.

6.2.1 Vanishing of the invariant The Lagrangian matching invariant of the connected sum of broken fibration vanishes:

Theorem F $\mathcal{L}_{(X_1\#X_2, \pi_1\#\pi_2)}(\mathfrak{s}) = 0$ for each admissible Spin^c -structure \mathfrak{s} .

The admissible ones are essentially those corresponding to n_1 points on Σ_1 and n_2 points on Σ_2 , where $n_i \leq \frac{1}{2}g(\Sigma_i) - 1$ for $i = 1, 2$ (though one could do a little better than this). The precise definition was given in Section 4.2.

Proof $\mathcal{L}_{(X_1\#X_2, \pi_1\#\pi_2)}(\mathfrak{s})$ factors through the map on homology induced by the fundamental class $[\widehat{V}_L]$:

$$[\widehat{V}_L]: H_*(\text{Sym}^n(\Sigma_1 \cup \Sigma_2); \mathbb{Z}) \rightarrow H_*(\text{Sym}^{n-1}(\Sigma_1\#\Sigma_2); \mathbb{Z}).$$

But here $L \subset \Sigma_1\#\Sigma_2$ is the loop separating the two connect summands, so the fundamental class map is zero (Part I, Lemma 3.18). \square

6.3 Dimensional reduction

The cobordism category of broken fibrations contains a much simpler category $\mathcal{C}_{\text{Morse}}$ in which all the structures have S^1 -symmetry.

An object of $\mathcal{C}_{\text{Morse}}$ is a triple (Σ, c, r) consisting of a closed, oriented surface with a locally constant function $c: \Sigma \rightarrow S^1$ and $r \in H_0(\Sigma; \mathbb{Z})$. A morphism (Y, f, a) from (Σ_1, c_1, r_1) to (Σ_2, c_2, r_2) is a cobordism Y , equipped with a circle-valued Morse function $f: Y \rightarrow S^1$ with no interior maximum or minimum, and a homology class $a \in H_1(Y, \partial Y \cup \text{crit}(f); \mathbb{Z})$. There are two requirements on a : its boundary must be

$$\partial a = r_2 - r_1 + \sum_{x \in \text{crit}(f)} (-1)^{\text{ind}(x)+1} [x] \in H_0(\partial Y \cup \text{crit}(f); \mathbb{Z}),$$

and the number $a \cap [F] - \dim(H_1(F))/2$ must be constant as F ranges over regular fibres F of f .

One obtains a broken fibration from a morphism in $\mathcal{C}_{\text{Morse}}$ by crossing with the identity map $\text{id}: S^1 \rightarrow S^1$.

6.3.1 Elementary cobordisms A basic case to consider is one where Y is an elementary cobordism C_{elem}^- from Σ to $\bar{\Sigma}$; the Morse function f is real valued, and has a single critical point, of index $+1$.

This gives rise to an invariant

$$\mathcal{L}(\mathfrak{s}_d): HF_*(\Sigma \times S^1, \mathfrak{t}_d) \rightarrow HF_*(\bar{\Sigma} \times S^1, \mathfrak{t}_d),$$

where \mathfrak{s}_d is the unique Spin^c -structure which is pulled back from C_{elem}^- and which satisfies $\langle \mathfrak{s}_d, [\Sigma] \rangle = d := \chi(\Sigma) + 2n$. Under the PSS isomorphisms of Floer and singular homology, this is the map

$$[\hat{V}_L]: H_*(\text{Sym}^n(\Sigma); \mathbb{Z}) \rightarrow H_{*+1}(\text{Sym}^{n-1}(\bar{\Sigma}); \mathbb{Z}).$$

We can also consider an elementary cobordism C_{elem}^+ running in the opposite direction (so the Morse function has one critical point, of index 2). Its invariant is the adjoint homomorphism (with respect to Poincaré duality)

$$H_*(\text{Sym}^{n-1}(\bar{\Sigma}); \mathbb{Z}) \rightarrow H_{*+1}(\text{Sym}^n(\Sigma); \mathbb{Z})$$

It is not difficult to give completely explicit formulae for these maps, but what is more illuminating is to observe that they are the same as the maps that arise in a well-known $(2+1)$ -dimensional TQFT, described in [30; 4], which we shall call the Segal–Donaldson model. This can be characterised as follows.

Definition 6.3 The *admissible $(2+1)$ -dimensional cobordism category* of degree d is the category in which an object is an oriented 2 -manifold, and a morphism is an “admissible cobordism” Y equipped with a line bundle which has degree d over the incoming and outgoing ends. Admissible means that either

$$\begin{aligned} \partial Y = \emptyset, \quad H_1(Y; \mathbb{Z}) \cong \mathbb{Z}, \text{ or} \\ \partial Y \neq \emptyset, \quad H_1(\partial Y; \mathbb{Z}) \rightarrow H_1(Y; \mathbb{Z}) \text{ is onto.} \end{aligned}$$

The Segal–Donaldson model is a TQFT on the category of admissible $(2+1)$ -dimensional cobordisms, ie a functor from this category to super \mathbb{Q} -vector spaces, preserving the multiplicative structures (disjoint union; super tensor product) and dualities. For our purposes, the morphisms are defined only up to a sign ± 1 .

Theorem 6.4 (Donaldson [4]) *For any sequence of nonnegative integers n_0, n_1, \dots , there is a unique TQFT V_d on the degree- d admissible cobordism category, defined over \mathbb{Q} , such that*

$$V_d(\Sigma) = \bigoplus_{i \geq 0} n_i \Lambda_i \Sigma$$

for connected surfaces Σ , as representations of the mapping class group $\pi_0 \text{Diff}^+(\Sigma)$. Here $\Lambda_i \Sigma$ denotes $\Lambda^{g(\Sigma)-i} H_1(\Sigma; \mathbb{Q})$. If Y^3 is closed, so that there is a unique line bundle L of degree d , then

$$V_d(Y, L) = \sum_i n_i a_i,$$

where $A(Y) = \pm(a_0 + \sum_i a_i(t^i + t^{-i}))$ is the normalised Alexander polynomial of Y .

(“Unique” means up to natural isomorphism.) We shall take V_d to be the unique TQFT with

$$V_d(\Sigma) = \bigoplus_{j \geq 1} j \Lambda_{d+j} \Sigma.$$

This is of interest because then

$$V_d(\Sigma) \cong H_*(\text{Sym}^n(\Sigma); \mathbb{Q})$$

as representations of $\pi_0 \text{Diff}^+(\Sigma)$, where $d = \chi(\Sigma) + 2n$.

We do not know that Lagrangian matching invariants give a TQFT; they depend, in principle, on the fibrations. However, if we show that $\mathcal{L}(\mathfrak{s}_d)$ coincides with $V_d(C_{\text{elem}}^-)$ then our field theory must coincide with V_d .

Lemma 6.5 *Considered as a map $H_*(\text{Sym}^n(\Sigma); \mathbb{Z}) \rightarrow H_*(\text{Sym}^{n-1}(\bar{\Sigma}))$, $\mathcal{L} = \mathcal{L}(\mathfrak{s}_d)$ satisfies the following properties:*

- (1) *Let $(\alpha, \bar{\alpha}) \in H_1(\Sigma; \mathbb{Z}) \times H_1(\bar{\Sigma}; \mathbb{Z})$ be such that the images of α and $\bar{\alpha}$ in $H_1(C_{\text{elem}}^-; \mathbb{Z})$ are equal. Then $\bar{\alpha} \wedge \mathcal{L}(x) = \mathcal{L}(\alpha \wedge x)$.*
- (2) *$U \cdot \mathcal{L}(x) = \mathcal{L}(U \cdot x)$ and $\theta_{\bar{\Sigma}} \cdot \mathcal{L}(x) = \mathcal{L}(\theta_{\Sigma} \cdot x)$.*
- (3) *On the degree 1 part $H_1(\Sigma; \mathbb{Z}) = H^1(\text{Sym}^n(\Sigma); \mathbb{Z})$, \mathcal{L} is given, up to sign, by $x \mapsto x \cap [L]$.*

Moreover, \mathcal{L} is uniquely characterised (up to sign) by these properties.

Proof The first two properties are instances of (11). The third is another way of writing the equation $\text{pr}_{1*}[\widehat{V}_L] = [\delta_L]$ for the fundamental class of the vanishing cycle \widehat{V}_L (see Part I, Lemma 3.18).

For the uniqueness, we note that by (2) it suffices to show that \mathcal{L} is uniquely determined on the images of the maps $\Lambda^k H_1(\Sigma; \mathbb{Z}) \rightarrow H^k \text{Sym}^n(\Sigma; \mathbb{Z})$ sending $x_1 \wedge \cdots \wedge x_k$ to $\mu(x_1) \cup \cdots \cup \mu(x_k)$; this is a simple consequence of (1) and (3). \square

One can easily verify from Donaldson's model that the Donaldson–Segal map $V_d(C_{\text{elem}}^-)$ satisfies these same properties.

Corollary 6.6 $V_d(C_{\text{elem}}^-) = \pm \mathcal{L}$, and $V_d(C_{\text{elem}}^+)$ is its adjoint $\pm \mathcal{L}^\vee$ (the Lagrangian matching invariant of $C_{\text{elem}}^+ \times S^1$).

This has the following consequence.

Theorem E Let Y be a closed 3-manifold with $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}$, and $f: Y \rightarrow S^1$ a Morse function with connected fibres and no extrema. Let

$$\text{id} \times f: S^1 \times Y \rightarrow S^1 \times S^1$$

be the corresponding broken fibration, and let \mathfrak{s}_d be an admissible Spin^c -structure on $S^1 \times Y$ which is pulled back from Y . Then

$$\mathcal{L}_{S^1 \times Y, \text{id} \times f}(\mathfrak{s}_d) = \pm \sum_{i \geq 0} i a_{d+i},$$

where $A(Y) = \pm(a_0 + \sum_{i \geq 0} a_i(t^i + t^{-i}))$ is the Alexander polynomial.

The assumption $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}$ implies that such Morse functions exist. It may be possible to push the argument through for a general Y with $b_1 > 0$, replacing the Alexander polynomial by the Milnor–Turaev torsion.

References

- [1] **D Auroux, S K Donaldson, L Katzarkov**, *Singular Lefschetz pencils*, *Geom. Topol.* 9 (2005) 1043–1114 MR2140998
- [2] **A Bertram, M Thaddeus**, *On the quantum cohomology of a symmetric product of an algebraic curve*, *Duke Math. J.* 108 (2001) 329–362 MR1833394
- [3] **R Bott**, *The stable homotopy of the classical groups*, *Ann. of Math.* (2) 70 (1959) 313–337 MR0110104

- [4] **S K Donaldson**, *Topological field theories and formulae of Casson and Meng–Taubes*, from: “Proceedings of the Kirbyfest (Berkeley, CA, 1998)”, *Geom. Topol. Monogr.* 2, *Geom. Topol. Publ.*, Coventry (1999) 87–102 MR1734402
- [5] **S K Donaldson**, *Floer homology groups in Yang–Mills theory*, *Cambridge Tracts in Mathematics* 147, Cambridge University Press (2002) MR1883043 With the assistance of M Furuta and D Kotschick
- [6] **S Donaldson, I Smith**, *Lefschetz pencils and the canonical class for symplectic four-manifolds*, *Topology* 42 (2003) 743–785 MR1958528
- [7] **C J Earle, J Eells**, *A fibre bundle description of Teichmüller theory*, *J. Differential Geometry* 3 (1969) 19–43 MR0276999
- [8] **Y Eliashberg, A Givental, H Hofer**, *Introduction to symplectic field theory*, *Geom. Funct. Anal.* (2000) 560–673 MR1826267 GAFA 2000 (Tel Aviv, 1999)
- [9] **A Floer**, *The unregularized gradient flow of the symplectic action*, *Comm. Pure Appl. Math.* 41 (1988) 775–813 MR948771
- [10] **A Floer, H Hofer**, *Coherent orientations for periodic orbit problems in symplectic geometry*, *Math. Z.* 212 (1993) 13–38 MR1200162
- [11] **A Floer, H Hofer, D Salamon**, *Transversality in elliptic Morse theory for the symplectic action*, *Duke Math. J.* 80 (1995) 251–292 MR1360618
- [12] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Lagrangian intersection Floer theory: anomaly and obstruction*, book manuscript (2000)
- [13] **R E Gompf**, *The topology of symplectic manifolds*, *Turkish J. Math.* 25 (2001) 43–59 MR1829078
- [14] **F Hirzebruch, H Hopf**, *Felder von Flächenelementen in 4–dimensionalen Mannigfaltigkeiten*, *Math. Ann.* 136 (1958) 156–172 MR0100844
- [15] **H Hofer, D A Salamon**, *Floer homology and Novikov rings*, from: “The Floer memorial volume”, *Progr. Math.* 133, Birkhäuser, Basel (1995) 483–524 MR1362838
- [16] **P Kronheimer, T Mrowka**, *Monopoles and contact structures*, *Invent. Math.* 130 (1997) 209–255 MR1474156
- [17] **P Kronheimer, T Mrowka, P Ozsváth, Z Szabó**, *Monopoles and lens space surgeries*, *Ann. of Math.* (2) 165 (2007) 457–546 MR2299739
- [18] **L Lazzarini**, *Existence of a somewhere injective pseudo-holomorphic disc*, *Geom. Funct. Anal.* 10 (2000) 829–862 MR1791142
- [19] **Y-J Lee**, *Heegaard splittings and Seiberg–Witten monopoles*, from: “Geometry and topology of manifolds”, *Fields Inst. Commun.* 47, Amer. Math. Soc. (2005) 173–202 MR2189932
- [20] **Y-J Lee**, *Reidemeister torsion in Floer–Novikov theory and counting pseudo-holomorphic tori. I*, *J. Symplectic Geom.* 3 (2005) 221–311 MR2199540

- [21] **D McDuff, D Salamon**, *J-holomorphic curves and symplectic topology*, Amer. Math. Soc. Colloquium Publ. 52, Amer. Math. Soc. (2004) MR2045629
- [22] **J Moser**, *On the volume elements on a manifold*, Trans. Amer. Math. Soc. 120 (1965) 286–294 MR0182927
- [23] **T Perutz**, *Surface-fibrations, four-manifolds, and symplectic Floer homology*, PhD thesis, Imperial College London (2005)
- [24] **T Perutz**, *Lagrangian matching invariants for fibred four-manifolds. I*, Geom. Topol. 11 (2007) 759–828 MR2302502
- [25] **S Piunikhin, D Salamon, M Schwarz**, *Symplectic Floer–Donaldson theory and quantum cohomology*, from: “Contact and symplectic geometry (Cambridge, 1994)”, Publ. Newton Inst. 8, Cambridge Univ. Press (1996) 171–200 MR1432464
- [26] **L Pontrjagin**, *A classification of mappings of the three-dimensional complex into the two-dimensional sphere*, Rec. Math. [Mat. Sbornik] N. S. 9 (51) (1941) 331–363 MR0004780
- [27] **M Poźniak**, *Floer homology, Novikov rings and clean intersections*, PhD thesis, University of Warwick (1994)
- [28] **D A Salamon**, *Seiberg-Witten invariants of mapping tori, symplectic fixed points, and Lefschetz numbers*, from: “Proceedings of 6th Gökova Geometry-Topology Conference”, Turkish J. Math. 23 (1999) 117–143 MR1701642
- [29] **M Schwarz**, *Cohomology operations from S^1 -cobordisms in Floer homology*, PhD thesis, ETH-Zürich (1995)
- [30] **G Segal**, *The definition of conformal field theory*, from: “Topology, geometry and quantum field theory”, London Math. Soc. Lecture Note Ser. 308, Cambridge Univ. Press (2004) 421–577 MR2079383 preprint circa 1990
- [31] **P Seidel**, *The symplectic Floer homology of a Dehn twist*, Math. Res. Lett. 3 (1996) 829–834 MR1426539
- [32] **P Seidel**, *Floer homology and the symplectic isotopy problem*, PhD thesis, University of Oxford (1997)
- [33] **P Seidel**, *Symplectic Floer homology and the mapping class group*, Pacific J. Math. 206 (2002) 219–229 MR1924827
- [34] **P Seidel**, *A long exact sequence for symplectic Floer cohomology*, Topology 42 (2003) 1003–1063 MR1978046
- [35] **P Seidel**, *Fukaya categories and Picard–Lefschetz theory*, book manuscript available at <http://www-math.mit.edu/~seidel> (2006)
- [36] **V de Silva**, *Products in the symplectic Floer homology of Lagrangian intersections*, PhD thesis, University of Oxford (1999)

- [37] **I Smith**, *Serre–Taubes duality for pseudoholomorphic curves*, *Topology* 42 (2003) 931–979 MR1978044
- [38] **C H Taubes**, *Seiberg–Witten invariants and pseudo-holomorphic subvarieties for self-dual, harmonic 2–forms*, *Geom. Topol.* 3 (1999) 167–210 MR1697181
- [39] **M Usher**, *Vortices and a TQFT for Lefschetz fibrations on 4–manifolds*, *Algebr. Geom. Topol.* 6 (2006) 1677–1743 MR2263047
- [40] **R Ye**, *Gromov’s compactness theorem for pseudo holomorphic curves*, *Trans. Amer. Math. Soc.* 342 (1994) 671–694 MR1176088

*DPMMS, Centre for Mathematical Sciences, University of Cambridge
Wilberforce Road, Cambridge CB3 0WB, UK*

*Department of Mathematics, Columbia University
2990 Broadway, New York, NY 10027, USA*

`tim.perutz@cantab.net`

Proposed: Peter Ozsváth
Seconded: Simon Donaldson, Ron Stern

Received: 7 June 2006
Revised: 14 November 2007