

Proof of the Arnold chord conjecture in three dimensions, II

MICHAEL HUTCHINGS
CLIFFORD HENRY TAUBES

In “Proof of the Arnold chord conjecture in three dimensions, I” [12], we deduced the Arnold chord conjecture in three dimensions from another result, which asserts that an exact symplectic cobordism between contact three-manifolds induces a map on (filtered) embedded contact homology satisfying certain axioms. The present paper proves the latter result, thus completing the proof of the three-dimensional chord conjecture. We also prove that filtered embedded contact homology does not depend on the choice of almost complex structure used to define it.

53D40, 57R58

1 Introduction

The main goal of this paper is to prove that an exact symplectic cobordism between contact 3-manifolds induces a map on (filtered) embedded contact homology (ECH) satisfying certain axioms. This result appears here as [Theorem 1.9](#), and was previously stated in Hutchings and Taubes [12, Theorem 2.4], where it was used to prove the Arnold Chord Conjecture in three dimensions. This result also has additional applications, for example it gives rise to new obstructions to symplectic embeddings in four dimensions; see Hutchings [8]. Along the way to proving [Theorem 1.9](#), we will also prove that filtered ECH does not depend on the choice of almost complex structure used to define it ([Theorem 1.3](#) below). Although this paper is a sequel to [12], we will not use anything from the latter paper except for some basic definitions. We begin by briefly reviewing these definitions. For more about ECH, see Hutchings [7; 4] and the references therein.

1.1 Embedded contact homology

Let Y be a closed oriented 3-manifold. (For simplicity, all 3-manifolds in this paper are assumed connected except where otherwise stated.) Let λ be a contact form on Y , let R denote the associated Reeb vector field, and let $\xi = \text{Ker}(\lambda)$ denote the associated contact structure. Assume that λ is nondegenerate, ie, all Reeb orbits are nondegenerate.

Let J be an almost complex structure on $\mathbb{R} \times Y$ such that J is \mathbb{R} -invariant, $J(\partial_s) = R$ where s denotes the \mathbb{R} coordinate, and J sends ξ to itself, rotating ξ positively with respect to the orientation on ξ given by $d\lambda$. We call such an almost complex structure *symplectization-admissible*. The reason for the terminology is that the noncompact symplectic manifold $(\mathbb{R} \times Y, d(e^s\lambda))$ is called the “symplectization” of (Y, λ) . Note that a symplectization-admissible almost complex structure is equivalent to an almost complex structure J on ξ that rotates positively with respect to $d\lambda$. In particular, the space of symplectization-admissible almost complex structures is contractible.

Given a generic symplectization-admissible structure J , and given $\Gamma \in H_1(Y)$, the *embedded contact homology* $\text{ECH}_*(Y, \lambda, \Gamma; J)$ is the homology of a chain complex $\text{ECC}_*(Y, \lambda, \Gamma; J)$ defined as follows. Recall that an *orbit set* is a finite set of pairs $\Theta = \{(\Theta_i, m_i)\}$, where the Θ_i are distinct embedded Reeb orbits, and the m_i are positive integers. The homology class of the orbit set Θ is defined by

$$[\Theta] := \sum_i m_i [\Theta_i] \in H_1(Y).$$

The orbit set $\Theta = \{(\Theta_i, m_i)\}$ is called *admissible* if $m_i = 1$ whenever Θ_i is hyperbolic, ie, the linearized Reeb flow around Θ_i has real eigenvalues. Define $\text{ECC}_*(Y, \lambda, \Gamma; J)$ to be the free $\mathbb{Z}/2$ -module generated by admissible orbit sets Θ with $[\Theta] = \Gamma$. Although ECH can also be defined over \mathbb{Z} , see Hutchings and Taubes [11, Section 9], in this paper we always use $\mathbb{Z}/2$ coefficients for simplicity.

To specify the differential ∂ on the chain complex, we need the following:

Definition 1.1 Given a symplectization-admissible J , and given orbit sets $\Theta = \{(\Theta_i, m_i)\}$ and $\Theta' = \{(\Theta'_j, m'_j)\}$, define a “ J -holomorphic curve from Θ to Θ' ” to be a J -holomorphic curve in $\mathbb{R} \times Y$ (whose domain is a possibly disconnected punctured compact Riemann surface) with positive ends at covers of Θ_i with total multiplicity m_i , negative ends at covers of Θ'_j with total multiplicity m'_j , and no other ends. Here a *positive end* of a holomorphic curve at a (not necessarily embedded) Reeb orbit γ is an end that is asymptotic to the cylinder $\mathbb{R} \times \gamma$ as the \mathbb{R} coordinate $s \rightarrow +\infty$. A *negative end* is defined analogously with $s \rightarrow -\infty$. Let $\mathcal{M}^J(\Theta, \Theta')$ denote the moduli space of J -holomorphic curves from Θ to Θ' , where two such curves are considered equivalent if they represent the same current in $\mathbb{R} \times Y$, up to translation of the \mathbb{R} coordinate.

Given admissible orbit sets Θ and Θ' with $[\Theta] = [\Theta'] = \Gamma$, the differential coefficient $\langle \partial\Theta, \Theta' \rangle \in \mathbb{Z}/2$ is defined to be the mod 2 count of J -holomorphic curves in $\mathcal{M}^J(\Theta, \Theta')$ with “ECH index” equal to 1. For the definition of the ECH index, see Hutchings [5; 6]. If J is generic, then ∂ is well-defined and $\partial^2 = 0$, as shown in

Hutchings and Taubes [10, Section 7]. A symplectization-admissible almost complex structure that is generic in this sense will be called *ECH-generic* here.

The ECH index defines a relative $\mathbb{Z}/d(c_1(\xi) + 2 \text{PD}(\Gamma))$ grading on the chain complex, where d denotes divisibility in $H^2(Y; \mathbb{Z})/\text{Torsion}$. However the grading will not play a major role in this paper.

It is shown in Taubes [21; 22; 23; 24] that ECH is isomorphic to a version of Seiberg–Witten Floer cohomology as defined by Kronheimer and Mrowka [13]. The precise statement is that there is a canonical isomorphism of relatively graded $\mathbb{Z}/2$ -modules¹

$$(1) \quad \text{ECH}_*(Y, \lambda, \Gamma; J) \simeq \widehat{\text{HM}}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}).$$

Here $\widehat{\text{HM}}^*$ denotes Seiberg–Witten Floer cohomology with $\mathbb{Z}/2$ coefficients, and $\mathfrak{s}_{\xi, \Gamma}$ denotes the spin-c structure $\mathfrak{s}_{\xi} + \text{PD}(\Gamma)$ on Y , where \mathfrak{s}_{ξ} denotes the spin-c structure determined by oriented 2-plane field ξ ; see Example 2.1.

1.2 Filtered ECH

If $\Theta = \{(\Theta_i, m_i)\}$ is an orbit set, its *symplectic action* or *length* is defined by

$$(2) \quad \mathcal{A}(\Theta) := \sum_i m_i \int_{\Theta_i} \gamma.$$

Since J is symplectization-admissible, it follows that the ECH differential decreases the action, ie, if $\langle \partial\Theta, \Theta' \rangle \neq 0$ then $\mathcal{A}(\Theta) > \mathcal{A}(\Theta')$. Thus for any real number L , it makes sense to define the *filtered ECH*, denoted by $\text{ECH}_*^L(Y, \lambda, \Gamma; J)$, to be the homology of the subcomplex $\text{ECC}_*^L(Y, \lambda, \Gamma; J)$ of the ECH chain complex spanned by ECH generators with action less than L .

There are various natural maps defined on filtered ECH. First, if $L < L'$ then there is a map

$$(3) \quad \iota_J^{L, L'} : \text{ECH}_*^L(Y, \lambda, \Gamma; J) \longrightarrow \text{ECH}_*^{L'}(Y, \lambda, \Gamma; J)$$

induced by the inclusion of chain complexes. The usual ECH is recovered as the direct limit

$$(4) \quad \text{ECH}_*(Y, \lambda, \Gamma; J) = \lim_{L \rightarrow \infty} \text{ECH}_*^L(Y, \lambda, \Gamma; J).$$

In addition, if c is a positive constant, then there is a canonical “scaling” isomorphism

$$(5) \quad s_J : \text{ECH}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \text{ECH}_*^{cL}(Y, c\lambda, \Gamma; J^c),$$

¹Ordinarily $\widehat{\text{HM}}^*$ is defined over \mathbb{Z} – see [13] – and it is shown in [23] that one can lift the isomorphism (1) to \mathbb{Z} coefficients.

where J^c is defined to agree with J when restricted to the contact planes ξ . This is because the chain complexes on both sides have the same generators, and the self-diffeomorphism of $\mathbb{R} \times Y$ sending $(s, y) \mapsto (cs, y)$ induces a bijection between J -holomorphic curves and J^c -holomorphic curves.

Note that to define $\text{ECH}_*^L(Y, \lambda, \Gamma; J)$, one does not need the full assumption that λ is nondegenerate and J is ECH-generic, but only the following conditions:

Definition 1.2 The contact form λ is L -nondegenerate if all Reeb orbits of length less than L are nondegenerate, and if there is no orbit set² of action exactly L . Given an L -nondegenerate contact form λ , a symplectization-admissible almost complex structure J for λ is ECH^L -generic if the genericity conditions from [11] hold for orbit sets of action less than L so that the ECH differential ∂ is well-defined on admissible orbit sets of action less than L and satisfies $\partial^2 = 0$.

1.3 J -independence of filtered ECH (statement)

We now state a theorem asserting that filtered ECH and the various maps on it do not depend on J . Before stating the result, let us recall precisely what it means to say that objects or maps between them are independent of choices.

Let $\{G_i \mid i \in I\}$ be a collection of groups indexed by some index set I . We say that “the groups G_i are canonically isomorphic to each other”, or “ G_i does not depend on i ”, if for every pair $i_1, i_2 \in I$ there is a canonical isomorphism $\phi_{i_1, i_2}: G_{i_1} \xrightarrow{\cong} G_{i_2}$, such that $\phi_{i_2, i_3} \circ \phi_{i_1, i_2} = \phi_{i_1, i_3}$ for every triple $i_1, i_2, i_3 \in I$. In this case all the groups G_i are canonically isomorphic to a single group G . Specifically one can define G to be the disjoint union of the groups G_i , modulo the equivalence relation that $g \in G_{i_1}$ is equivalent to $\phi_{i_1, i_2}(g) \in G_{i_2}$, with group operation induced by the operations on the groups G_i .

Now let $\{H_j \mid j \in J\}$ be another such collection of groups that are canonically isomorphic to a single group H via isomorphisms $\psi_{j_1, j_2}: H_{j_1} \xrightarrow{\cong} H_{j_2}$ as above. Then a collection of maps $\{f_{i, j}: G_i \rightarrow H_j \mid i \in I, j \in J\}$ induces a well-defined map

²The condition that there is no orbit set of action exactly L is not needed to define filtered ECH, but it will be convenient to choose L this way when we relate filtered ECH to Seiberg–Witten Floer cohomology, starting in Lemma 2.3.

$f: G \rightarrow H$ provided that the diagram

$$\begin{array}{ccc}
 G_{i_1} & \xrightarrow{f_{i_1, j_1}} & H_{j_1} \\
 \phi_{i_1, i_2} \downarrow \simeq & & \psi_{j_1, j_2} \downarrow \simeq \\
 G_{i_2} & \xrightarrow{f_{i_2, j_2}} & H_{j_2}
 \end{array}$$

commutes for all $i_1, i_2 \in I$ and $j_1, j_2 \in J$.

With these conventions, we now have:

Theorem 1.3 *Let Y be a closed oriented connected 3-manifold, and let $\Gamma \in H_1(Y)$.*

- (a) *If λ is an L -nondegenerate contact form on Y , then $\text{ECH}_*^L(Y, \lambda, \Gamma; J)$ does not depend on the choice of ECH^L -generic J , so we can denote it by $\text{ECH}_*^L(Y, \lambda, \Gamma)$.*
- (b) *If $L < L'$ and if λ is L' -nondegenerate, then the maps $i_J^{L, L'}$ in (3) induce a well-defined map*
- (6)
$$i^{L, L'}: \text{ECH}_*^L(Y, \lambda, \Gamma) \longrightarrow \text{ECH}_*^{L'}(Y, \lambda, \Gamma).$$
- (c) *If λ is a nondegenerate contact form on Y , then $\text{ECH}_*(Y, \lambda, \Gamma; J)$ does not depend on the choice of ECH -generic J , so we can denote it by $\text{ECH}_*(Y, \lambda, \Gamma)$.*
- (d) *If $c > 0$, then the scaling isomorphisms s_J in (5) induce a well-defined isomorphism*
- (7)
$$s: \text{ECH}_*^L(Y, \lambda, \Gamma) \xrightarrow{\simeq} \text{ECH}_*^{cL}(Y, c\lambda, \Gamma).$$
- (e) *The isomorphism (1) does not depend on J and so determines a canonical isomorphism*
- (8)
$$\text{ECH}_*(Y, \lambda, \Gamma) \simeq \widehat{\text{HM}}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}).$$

The proof of [Theorem 1.3](#) uses Seiberg–Witten theory, similarly to parts of the proof of the isomorphism (1), and is given in [Section 3](#).

Remark 1.4 Although this is not necessary for the proof of the chord conjecture, the proof of [Theorem 1.3](#) works just as well with \mathbb{Z} coefficients by [\[23\]](#). Parts (a)–(d) of [Theorem 1.3](#) also hold for disconnected three-manifolds, by a straightforward modification of the proof.

At times it is convenient to ignore the homology class Γ in the definition of ECH, and simply define

$$(9) \quad \text{ECH}_*(Y, \lambda) := \bigoplus_{\Gamma \in H_1(Y)} \text{ECH}_*(Y, \lambda, \Gamma).$$

This is the homology of a chain complex $\text{ECC}_*(Y, \lambda; J)$ generated by all admissible orbit sets, and by (8) this homology is canonically isomorphic (as a relatively graded $\mathbb{Z}/2$ -module) to

$$\widehat{\text{HM}}^{-*}(Y) := \bigoplus_{\mathfrak{s} \in \text{Spin}^c(Y)} \widehat{\text{HM}}^{-*}(Y, \mathfrak{s}).$$

Note that while $\text{ECH}_*(Y, \lambda)$ is a topological invariant of Y , the filtered version $\text{ECH}_*^L(Y, \lambda)$ depends strongly on λ and L .

1.4 Exact symplectic cobordisms

If Y_+ and Y_- are closed oriented (connected) 3-manifolds, our convention is that a “cobordism from Y_+ to Y_- ” is a compact oriented smooth 4-manifold with $\partial X = Y_+ - Y_-$. Such a cobordism induces a map of ungraded $\mathbb{Z}/2$ -modules

$$(10) \quad \widehat{\text{HM}}^*(X): \widehat{\text{HM}}^*(Y_+) \longrightarrow \widehat{\text{HM}}^*(Y_-).$$

If λ_{\pm} are nondegenerate contact forms on Y_{\pm} , we define

$$(11) \quad \Phi(X): \text{ECH}_*(Y_+, \lambda_+) \longrightarrow \text{ECH}_*(Y_-, \lambda_-)$$

to be the map on ECH obtained by composing the map (10) on Seiberg–Witten Floer cohomology with the canonical isomorphism (8) on both sides.

If (Y_{\pm}, λ_{\pm}) are as above, an *exact symplectic cobordism* from (Y_+, λ_+) to (Y_-, λ_-) is a pair (X, ω) , where X is a cobordism from Y_+ to Y_- , and ω is a symplectic form on X , such that there exists a 1-form λ on X with $d\lambda = \omega$ and $\lambda|_{Y_{\pm}} = \lambda_{\pm}$. A 1-form with these properties is called a *Liouville form* for (X, ω) . When we wish to specify a Liouville form, we denote the exact symplectic cobordism by (X, λ) , and we continue to write $\omega = d\lambda$.

When (X, ω) is an exact symplectic cobordism as above, we would like to relate the map (11) to holomorphic curves. To prepare for this, let λ be a Liouville form. This determines a Liouville vector field V characterized by $\iota_V \omega = \lambda$. If $\varepsilon > 0$ is sufficiently small, then the flow of V starting on Y_- for times in $[0, \varepsilon]$ defines a diffeomorphism

$$(12) \quad N_- \simeq [0, \varepsilon] \times Y_-,$$

where N_- is (the closure of) a neighborhood of Y_- . If s denotes the $[0, \varepsilon]$ coordinate in (12), then $\lambda = e^s \lambda_-$ on N_- . Likewise we obtain a neighborhood

$$(13) \quad N_+ \simeq [-\varepsilon, 0] \times Y_+$$

of Y_+ in which $\lambda = e^s \lambda_+$. Using the identifications (12) and (13), one can then glue symplectization ends to X to obtain the “completion”

$$(14) \quad \bar{X} := ((-\infty, 0] \times Y_-) \cup_{Y_-} X \cup_{Y_+} ([0, \infty) \times Y_+),$$

which is a noncompact symplectic 4-manifold.

Note that the completion (14) depends on the Liouville form in the following sense: If λ' is another Liouville form for ω , then the obvious identification between the completions (14) for λ and λ' is a homeomorphism, and will be a diffeomorphism if λ and λ' agree near ∂X .

Definition 1.5 An almost complex structure J on \bar{X} is *cobordism-admissible* if it is ω -compatible on X , and if it agrees with symplectization-admissible almost complex structures J_+ for λ_+ on $[0, \infty) \times Y_+$ and J_- for λ_- on $(-\infty, 0] \times Y_-$.

Given a cobordism-admissible J , and given (not necessarily admissible) orbit sets $\Theta^+ = \{(\Theta_j^+, m_j^+)\}$ in Y_+ and $\Theta^- = \{(\Theta_j^-, m_j^-)\}$ in Y_- , we define a “ J -holomorphic curve in \bar{X} from Θ^+ to Θ^- ” analogously to Definition 1.1, and denote the moduli space of such curves by $\mathcal{M}^J(\Theta^+, \Theta^-)$, where two such curves are considered equivalent if they represent the same current in \bar{X} . More generally, we make the following definition:

Definition 1.6 Let J, J_\pm be as in Definition 1.5. A *broken J -holomorphic curve from Θ^+ to Θ^-* is a collection of holomorphic curves $\{C_k\}_{1 \leq k \leq N}$ called “levels”, and (not necessarily admissible) orbit sets Θ^{k+} and Θ^{k-} for each k , such that there exists $k_0 \in \{1, \dots, N\}$ such that:

- Θ^{k+} is an orbit set in (Y_+, λ_+) for each $k \geq k_0$, Θ^{k-} is an orbit set in (Y_-, λ_-) for each $k \leq k_0$, $\Theta^{N+} = \Theta^+$; $\Theta^{1-} = \Theta^-$ and $\Theta^{k-} = \Theta^{k-1,+}$ for each $k > 1$.
- If $k > k_0$ then $C_k \in \mathcal{M}^{J+}(\Theta^{k+}, \Theta^{k-})$, if $k < k_0$ then $C_k \in \mathcal{M}^{J-}(\Theta^{k+}, \Theta^{k-})$, and $C_{k_0} \in \mathcal{M}^J(\Theta^{k_0,+}, \Theta^{k_0,-})$.
- If $k \neq k_0$ then C_k is not \mathbb{R} -invariant (as a current).

Let $\overline{\mathcal{M}^J(\Theta^+, \Theta^-)}$ denote the moduli space of broken J -holomorphic curves from Θ^+ to Θ^- as above.

Note that $\mathcal{M}^J(\Theta^+, \Theta^-)$ is a subset of $\overline{\mathcal{M}^J(\Theta^+, \Theta^-)}$ corresponding to broken curves as above in which the number of levels $N = 1$. (It is perhaps a misnomer to use the term “broken” when there is just one level.)

We would now like to relate the map (11) to broken J -holomorphic curves in \overline{X} , where J is cobordism-admissible.

1.5 Statement of the main theorem

Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) , and assume that the contact forms λ_{\pm} are nondegenerate. Fix a cobordism-admissible almost complex structure J on \overline{X} that restricts to symplectization-admissible almost complex structures J_+ on $[0, \infty) \times Y_+$ and J_- on $(-\infty, 0] \times Y_-$, as in Definition 1.5. We now recall some definitions from [12].

Product cylinders If the cobordism (X, λ) and the almost complex structure J on \overline{X} are very special, then X may contain regions that look like pieces of a symplectization, in the following sense:

Definition 1.7 A *product region* in X is the image of an embedding $[s_-, s_+] \times Z \rightarrow X$, where $s_- < s_+$ and Z is an open 3-manifold, such that:

- $\{s_{\pm}\} \times Z$ maps to Y_{\pm} , and $(s_-, s_+) \times Z$ maps to the interior of X .
- The pullback of the Liouville form λ on X to $[s_-, s_+] \times Z$ has the form $e^s \lambda_0$, where s denotes the $[s_-, s_+]$ coordinate, and λ_0 is a contact form on Z .
- The pullback of the almost complex structure J on X to $[s_-, s_+] \times Z$ has the following two properties: First, the restriction of J to $\text{Ker}(\lambda_0)$ is independent of s . Second, $J(\partial/\partial_s) = f(s)R_0$, where f is a positive function of s and R_0 denotes the Reeb vector field for λ_0 .

Given a product region as above, the embedded Reeb orbits of λ_{\pm} in $\{s_{\pm}\} \times Z$ are identified with the embedded Reeb orbits of λ_0 in Z . If γ is such a Reeb orbit, then we can form a J -holomorphic cylinder in \overline{X} by taking the union of $[s_-, s_+] \times \gamma$ in $[s_-, s_+] \times Z$ with $(-\infty, 0] \times \gamma$ in $(-\infty, 0] \times Y_-$ and $[0, \infty) \times \gamma$ in $[0, \infty) \times Y_+$.

Definition 1.8 We call a J -holomorphic cylinder as above a *product cylinder*.

Composition of cobordisms If (X^+, λ^+) is an exact symplectic cobordism from (Y_+, λ_+) to (Y_0, λ_0) , and if (X^-, λ^-) is an exact symplectic cobordism from (Y_0, λ_0) to (Y_-, λ_-) , then we can compose them to obtain an exact symplectic cobordism

$(X^- \circ X^+, \lambda)$ from (Y_+, λ_+) to (Y_-, λ_-) . Here $X^- \circ X^+$ is obtained by gluing X^- and X^+ along Y_0 analogously to (14), and $\lambda|_{X^\pm} = \lambda^\pm$.

Homotopy of cobordisms Two exact symplectic cobordisms (X, ω_0) and (X, ω_1) from (Y_+, λ_+) to (Y_-, λ_-) with the same underlying four-manifold X are *homotopic* if there is a smooth one-parameter family of symplectic forms $\{\omega_t \mid t \in [0, 1]\}$ on X such that (X, ω_t) is an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) for each $t \in [0, 1]$.

Theorem 1.9 *Let (Y_+, λ_+) and (Y_-, λ_-) be closed oriented connected 3-manifolds with nondegenerate contact forms. Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) . Then there exist maps of ungraded $\mathbb{Z}/2$ -modules*

$$(15) \quad \Phi^L(X, \lambda): \text{ECH}_*^L(Y_+, \lambda_+) \longrightarrow \text{ECH}_*^L(Y_-, \lambda_-)$$

for each real number L , such that:

(Homotopy invariance) *The map $\Phi^L(X, \lambda)$ depends only on L and the homotopy class of (X, ω) .*

(Inclusion) *If $L < L'$ then the following diagram commutes:*

$$\begin{CD} \text{ECH}_*^L(Y_+, \lambda_+) @>\Phi^L(X, \lambda)>> \text{ECH}_*^L(Y_-, \lambda_-) \\ @VVi^{L, L'}V @VVi^{L, L'}V \\ \text{ECH}_*^{L'}(Y_+, \lambda_+) @>\Phi^{L'}(X, \lambda)>> \text{ECH}_*^{L'}(Y_-, \lambda_-) \end{CD}$$

(Direct limit)

$$\lim_{L \rightarrow \infty} \Phi^L(X, \lambda) = \Phi(X): \text{ECH}_*(Y_+, \lambda_+) \longrightarrow \text{ECH}_*(Y_-, \lambda_-),$$

where $\Phi(X)$ is as in (11).

(Composition) *If (X, λ) is the composition of (X^-, λ^-) and (X^+, λ^+) as above with λ_0 nondegenerate, then*

$$\Phi^L(X^- \circ X^+, \lambda) = \Phi^L(X^-, \lambda^-) \circ \Phi^L(X^+, \lambda^+).$$

(Scaling) *If c is a positive constant then the following diagram commutes:*

$$\begin{CD} \text{ECH}_*^L(Y_+, \lambda_+) @>\Phi^L(X, \lambda)>> \text{ECH}_*^L(Y_-, \lambda_-) \\ @V s \downarrow \simeq VV @V s \downarrow \simeq VV \\ \text{ECH}_*^{cL}(Y_+, c\lambda_+) @>\Phi^{cL}(X, c\lambda)>> \text{ECH}_*^{cL}(Y_-, c\lambda_-) \end{CD}$$

(Holomorphic curves) *Let J be a cobordism-admissible almost complex structure on \bar{X} such that J_+ and J_- are ECH^L -generic. Then there exists a (noncanonical) chain map*

$$\hat{\Phi}: \text{ECC}_*^L(Y_+, \lambda_+, J_+) \longrightarrow \text{ECC}_*^L(Y_-, \lambda_-, J_-)$$

inducing $\Phi^L(X, \lambda)$, such that if Θ^+ and Θ^- are admissible orbit sets for (Y_+, λ_+) and (Y_-, λ_-) respectively with action less than L , then:

- (i) *If there are no broken J -holomorphic curves in \bar{X} from Θ^+ to Θ^- , then $\langle \hat{\Phi}\Theta^+, \Theta^- \rangle = 0$.*
- (ii) *If the only broken J -holomorphic curve in \bar{X} from Θ^+ to Θ^- is a union of covers of product cylinders, then $\langle \hat{\Phi}\Theta^+, \Theta^- \rangle = 1$.*

Our proof of [Theorem 1.9](#) uses Seiberg–Witten theory. While it would be natural to try to define the maps $\Phi^L(X, \lambda)$ more directly by counting (broken) holomorphic curves in \bar{X} with ECH index 0, there are substantial technical difficulties with this approach; see the discussion in [[4](#), Section 5.5].

Remark 1.10 The maps $\Phi^L(X, \lambda)$ respect the decomposition (9) in the this sense: The image of $\text{ECH}_*(Y_+, \lambda_+, \Gamma_+)$ has a nonzero component in $\text{ECH}_*(Y_-, \lambda_-, \Gamma_-)$ only if $\Gamma_+ \in H_1(Y_+)$ and $\Gamma_- \in H_1(Y_-)$ map to the same class in $H_1(X)$. This follows from part (i) of the holomorphic curves axiom (or more simply by keeping track of the spin-c structures in the construction of $\Phi^L(X, \lambda)$).

Remark 1.11 Part (ii) of the holomorphic curves axiom includes the case where Θ^\pm and the product region are empty, in which case there is a unique holomorphic curve, namely the empty set. It then follows that $\Phi^L(X, \lambda)$ sends the ECH contact invariant for (Y_+, λ_+) (the class in ECH represented by the empty set of Reeb orbits) to the ECH contact invariant for (Y_-, λ_-) .

Remark 1.12 If we allow Y_+ and Y_- to be disconnected, then all of [Theorem 1.9](#) except for the direct limit axiom still holds, by a straightforward modification of the proof. (The statement of the direct limit axiom does not make sense in this case because the relevant Seiberg–Witten Floer cohomology needed to define the map $\Phi(X)$ has not been defined for disconnected three-manifolds).

Remark 1.13 We expect that [Theorem 1.9](#) also holds with \mathbb{Z} coefficients. Note that the cobordism maps on Seiberg–Witten Floer homology defined in [[13](#)] depend on a choice of “homology orientation” of the cobordism. However we expect to be able to define cobordism maps on ECH without choosing a homology orientation, as this works

in those cases where ECH cobordism maps can be defined by counting holomorphic curves; cf Latschev and Wendl [14, Lemma A.14]. The direct limit axiom should then hold for a suitable homology orientation.

The rest of this paper is organized as follows. Section 2 reviews the definition of Seiberg–Witten Floer cohomology and discusses how to perturb the relevant Seiberg–Witten equations using a contact form. Section 3 relates filtered ECH to a filtered version of Seiberg–Witten Floer cohomology and proves Theorem 1.3. Section 4 reviews the construction of cobordism maps on Seiberg–Witten Floer cohomology and introduces a perturbation of the relevant Seiberg–Witten equations on an exact symplectic cobordism using the symplectic form. Section 5 proves all of Theorem 1.9, except for the holomorphic curves axiom, which is proved in Section 6, and using some results relating Seiberg–Witten solutions on exact symplectic cobordisms to holomorphic curves, which are proved in Section 7.

Acknowledgments

The first author was partially supported by NSF grant DMS-0806037. The second author was partially supported by the Clay Mathematics Institute, the Mathematical Sciences Research Institute and the NSF. Both authors thank MSRI for its hospitality.

2 Seiberg–Witten Floer cohomology and contact forms

We now review how to define Seiberg–Witten Floer cohomology, with the Seiberg–Witten equations perturbed by a contact form.

2.1 Seiberg–Witten Floer cohomology

We begin by briefly reviewing the relevant parts of the definition of Seiberg–Witten Floer cohomology. We follow the conventions in the book by Kronheimer and Mrowka [13], which explains the full details of this theory.

Let Y be a closed oriented (connected) 3–manifold, and let g be a Riemannian metric on Y . A *spin- c structure* on Y consists of a rank 2 Hermitian vector bundle \mathbb{S} over Y , whose sections are called *spinors*, together with a bundle map $\text{cl}: TY \rightarrow \text{End}(\mathbb{S})$, called *Clifford multiplication*, such that

$$(16) \quad \text{cl}(a) \text{cl}(b) + \text{cl}(b) \text{cl}(a) = -2\langle a, b \rangle$$

for $a, b \in T_y Y$, and

$$\text{cl}(e_1) \text{cl}(e_2) \text{cl}(e_3) = 1$$

when (e_1, e_2, e_3) is an oriented orthonormal basis for $T_y Y$. We denote the spin-c structure by $\mathfrak{s} = (\mathbb{S}, \text{cl})$. Two spin-c structures (\mathbb{S}, cl) and $(\mathbb{S}', \text{cl}')$ are considered equivalent if there is a bundle isomorphism $\phi: \mathbb{S} \xrightarrow{\sim} \mathbb{S}'$ respecting the Clifford multiplications, meaning that $\text{cl}'(v)\phi(\psi) = \phi(\text{cl}(v)\psi)$ for $v \in T_y Y$ and $\psi \in \mathbb{S}_y$. The set of spin-c structures is then an affine space over $H^2(Y; \mathbb{Z})$. The definition of the action is that if $e \in H^2(Y; \mathbb{Z})$, then

$$(17) \quad (\mathbb{S}, \text{cl}) + e := (\mathbb{S} \otimes L_e, \text{cl} \otimes 1),$$

where L_e denotes the complex line bundle with $c_1(L_e) = e$. If $\mathfrak{s} = (\mathbb{S}, \text{cl})$ is a spin-c structure, we define $c_1(\mathfrak{s}) := c_1(\mathbb{S}) \in H^2(Y; \mathbb{Z})$.

A spin-c structure is also equivalent to a lift of the frame bundle of TY from a principal $\text{SO}(3)$ -bundle to a principal $U(2)$ -bundle. The set of spin-c structures on Y does not depend on the metric g .

Example 2.1 An oriented 2-plane field ξ on Y determines a spin-c structure \mathfrak{s}_ξ as follows. The spinor bundle is given by

$$\mathbb{S} = \underline{\mathbb{C}} \oplus \xi,$$

where $\underline{\mathbb{C}}$ denotes the trivial complex line bundle on Y , and ξ is regarded as a Hermitian line bundle using its orientation and the metric on Y . Clifford multiplication is defined as follows: if (e_1, e_2, e_3) is an oriented orthonormal basis for $T_y Y$ such that (e_2, e_3) is an oriented orthonormal basis for ξ_y , then in terms of the basis $(1, e_2)$ for \mathbb{S} ,

$$\text{cl}(e_1) = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, \quad \text{cl}(e_2) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad \text{cl}(e_3) = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}.$$

Now fix a spin-c structure (\mathbb{S}, cl) . A *spin-c connection* is a connection $\mathbb{A}_\mathbb{S}$ on \mathbb{S} which is compatible with Clifford multiplication in the following sense: If v is a section of TY and ψ is a spinor, then

$$(18) \quad \nabla_{\mathbb{A}_\mathbb{S}}(\text{cl}(v)\psi) = \text{cl}(\nabla v)\psi + \text{cl}(v)\nabla_{\mathbb{A}_\mathbb{S}}\psi,$$

where ∇v denotes the covariant derivative of v with respect to the Levi-Civita connection. A spin-c connection $\mathbb{A}_\mathbb{S}$ is equivalent to a (Hermitian) connection \mathbb{A} on the determinant line bundle $\det(\mathbb{S})$. Note that adding an imaginary-valued 1-form a to \mathbb{A} has the effect of adding $a/2$ to $\mathbb{A}_\mathbb{S}$. A spin-c connection $\mathbb{A}_\mathbb{S}$, identified in this way with a connection \mathbb{A} on $\det(\mathbb{S})$, determines a *Dirac operator* $D_\mathbb{A}$, which is defined to be the composition

$$(19) \quad C^\infty(Y; \mathbb{S}) \xrightarrow{\nabla_{\mathbb{A}_\mathbb{S}}} C^\infty(Y; T^*Y \otimes \mathbb{S}) \xrightarrow{\text{cl}} C^\infty(Y; \mathbb{S}).$$

Here Clifford multiplication is extended to cotangent vectors by using the metric on Y to identify T^*Y with TY .

Now fix an exact 2-form η on Y . The *Seiberg–Witten equations* with perturbation η concern a pair (\mathbb{A}, Ψ) , where \mathbb{A} is a connection on $\det(\mathbb{S})$ and Ψ is a spinor. The equations are

$$(20) \quad \begin{aligned} D_{\mathbb{A}}\Psi &= 0, \\ *F_{\mathbb{A}} &= \tau(\Psi) + i*\eta. \end{aligned}$$

Here $*$ denotes the Hodge star, $F_{\mathbb{A}}$ denotes the curvature of \mathbb{A} , and $\tau: \mathbb{S} \rightarrow iT^*Y$ is a quadratic bundle map defined by

$$\tau(\Psi)(v) = \langle \text{cl}(v)\Psi, \Psi \rangle$$

for $\Psi \in \mathbb{S}_y$ and $v \in T_yY$. A pair (\mathbb{A}, Ψ) solves the Seiberg–Witten equations (20) if and only if it is a critical point of the functional α_η on the set of pairs (\mathbb{A}, Ψ) defined by

$$(21) \quad \alpha_\eta(\mathbb{A}, \Psi) := -\frac{1}{8} \int_Y (\mathbb{A} - \mathbb{A}_0) \wedge (F_{\mathbb{A}} + F_{\mathbb{A}_0} - 2i\eta) + \frac{1}{2} \int_Y \langle D_{\mathbb{A}}\Psi, \Psi \rangle.$$

Here \mathbb{A}_0 is any reference connection on $\det(\mathbb{S})$; changing this reference connection will add a constant to the functional (21).

The *gauge group* $\mathcal{G} := C^\infty(Y; S^1)$ acts on the set of pairs (\mathbb{A}, Ψ) by

$$(22) \quad u \cdot (\mathbb{A}, \Psi) := (\mathbb{A} - 2u^{-1}du, u\Psi),$$

and this action preserves the set of solutions to the Seiberg–Witten equations. Two solutions are considered equivalent if one is obtained from the other by the action of the gauge group. A solution (\mathbb{A}, Ψ) is called *reducible* if $\Psi \equiv 0$, and *irreducible* otherwise. If the exact 2-form η is suitably generic, then there are only finitely many irreducible solutions to (20) (modulo gauge equivalence), each of which is cut out transversely in an appropriate sense. Fix such a 2-form η .

The chain complex for defining Seiberg–Witten Floer cohomology (with $\mathbb{Z}/2$ coefficients) can be decomposed into submodules (not subcomplexes)

$$\widehat{\text{CM}}^* = \widehat{\text{CM}}_{\text{irr}}^* \oplus \widehat{\text{CM}}_{\text{red}}^*.$$

Here $\widehat{\text{CM}}_{\text{irr}}^*$ is the free $\mathbb{Z}/2$ -module generated by the irreducible solutions, while $\widehat{\text{CM}}_{\text{red}}^*$ is a more complicated term arising from the reducibles. The details of the reducible part $\widehat{\text{CM}}_{\text{red}}^*$ do not concern us here, because soon we will be restricting attention to a certain subcomplex of $\widehat{\text{CM}}^*$, for a particular perturbation η , which lives entirely within $\widehat{\text{CM}}_{\text{irr}}^*$ as explained in the proof of [Lemma 2.3](#) below.

For the same reason, our primary interest is in the part of the chain complex differential that maps $\widehat{\text{CM}}_{\text{irr}}^*$ to itself. To describe this, let (\mathbb{A}_+, Ψ_+) and (\mathbb{A}_-, Ψ_-) be two solutions to the equations (20). An *instanton* from (\mathbb{A}_-, Ψ_-) to (\mathbb{A}_+, Ψ_+) is a smooth one-parameter family of pairs $(\mathbb{A}(s), \Psi(s))$ parametrized by $s \in \mathbb{R}$, where $\mathbb{A}(s)$ is a connection on $\det(\mathbb{S})$ and $\Psi(s)$ is a spinor, satisfying the equations

$$(23) \quad \begin{aligned} \frac{\partial}{\partial s} \Psi(s) &= -D_{\mathbb{A}(s)} \Psi(s), \\ \frac{\partial}{\partial s} \mathbb{A}(s) &= -*F_{\mathbb{A}(s)} + \tau(\Psi(s)) + i*\eta, \\ \lim_{s \rightarrow \pm\infty} (\mathbb{A}(s), \Psi(s)) &= (\mathbb{A}_{\pm}, \Psi_{\pm}). \end{aligned}$$

A solution to these equations is a downward gradient flow line of the functional (21) from (\mathbb{A}_-, ψ_-) to (\mathbb{A}_+, ψ_+) . Here the metric on the space of pairs (\mathbb{A}, Ψ) is induced by the Hermitian inner product on \mathbb{S} together with $1/4$ of the L^2 inner product on $\Omega^1(Y; i\mathbb{R})$. The gauge group $C^\infty(Y; S^1)$ again acts on the space of such instantons. Also \mathbb{R} acts on the space of instantons by translating the s coordinate. If $(\mathbb{A}_{\pm}, \Psi_{\pm})$ are irreducible, then the coefficient of (\mathbb{A}_-, Ψ_-) in the differential of (\mathbb{A}_+, Ψ_+) counts index 1 instantons from (\mathbb{A}_-, Ψ_-) to (\mathbb{A}_+, Ψ_+) , modulo gauge equivalence and translation of s . Here the “index” is the local expected dimension of the moduli space of instantons modulo gauge equivalence. The index defines a relative $\mathbb{Z}/d(c_1(\mathfrak{s}))$ -grading on the chain complex, such that the differential increases the grading by 1.

All we need to know about the rest of the differential is that if (\mathbb{A}_+, Ψ_+) is irreducible, and if there is no index one instanton to (\mathbb{A}_+, Ψ_+) from a reducible solution (\mathbb{A}_-, Ψ_-) , then the differential sends (\mathbb{A}_+, Ψ_+) to an element of $\widehat{\text{CM}}_{\text{irr}}^*$. Here when (\mathbb{A}_-, Ψ_-) is reducible, the index is defined to be the local expected dimension of the moduli space of instantons modulo gauge equivalence that have the same asymptotic decay rate as $s \rightarrow -\infty$.

In general, to obtain transversality of the moduli spaces of instantons as needed to define the differential, some “abstract” perturbations of equations (20) and (23) are required. These are described in [13, Chapter 11], where a Banach space \mathcal{P} of such perturbations is defined. Below, an *abstract* perturbation is one from \mathcal{P} , a *small* abstract perturbation is one with small \mathcal{P} -norm, and a *generic* abstract perturbation is one from a residual subset (depending on context) of \mathcal{P} . As noted previously, if the exact 2-form η in (20) is suitably generic, then there are only finitely many irreducible solutions to (20), and these are all cut out transversely. For such a choice of η , the generic abstract perturbation needed to define the differential can be chosen to vanish to any given order on the irreducible solutions to (20), and in particular so that the generators of $\widehat{\text{CM}}_{\text{irr}}^*$

are unchanged, ie, every solution to the perturbed version of (20) is a solution to the unperturbed version and vice-versa; see [21, Section 3h, Part 5]. When η is generic in this sense, we always assume that the abstract perturbations needed to define the differential (and also the cobordism maps reviewed in Section 4.1) are chosen this way. The abstract perturbations then have little conceptual role in the arguments below – see Proposition 3.1(c) and also [21, Theorem 4.4] – so we usually suppress them from the notation.

We denote the homology of this chain complex by $\widehat{HM}^*(Y, \mathfrak{s}; g, \eta)$. The homologies for different choices of (g, η) (and abstract perturbations) are canonically isomorphic to each other. The isomorphisms between the homologies for different choices are a special case of the cobordism maps reviewed in Section 4.1. Thus the homologies for different choices are canonically isomorphic to a single $\mathbb{Z}/2$ -module, which is denoted by $\widehat{HM}^*(Y, \mathfrak{s})$.

2.2 Perturbing the equations using a contact form

Now suppose λ is a contact form on Y . Choose an almost complex structure J on the contact planes ξ as needed to define a symplectization-admissible almost complex structure on Y ; see Section 1.1. The choices of λ and J determine a metric g on Y such that Reeb vector field R has length 1 and is orthogonal to the contact plane field ξ , on which the metric is given by

$$(24) \quad g(v, w) = \frac{1}{2}d\lambda(v, Jw).$$

In this metric one has

$$(25) \quad |\lambda| = 1, \quad d\lambda = 2*\lambda.$$

Remark 2.2 The factor of 1/2 in (24) and the factor of 2 in (25) could probably be dropped, but we have included these factors for consistency with the papers Taubes [19; 21] and their sequels.

With these choices made, if $\mathfrak{s} = (\mathbb{S}, \text{cl})$ is any spin-c structure, then there is a canonical decomposition

$$(26) \quad \mathbb{S} = E \oplus K^{-1}E$$

into eigenbundles of $\text{cl}(\lambda)$, where E is the $+i$ eigenbundle, and K^{-1} denotes the contact structure ξ , regarded as a Hermitian line bundle via J . When E is the trivial line bundle \mathbb{C} , one recovers Example 2.1. In this case it turns out that there is a

distinguished connection $A_{K^{-1}}$ on K^{-1} such that $D_{A_{K^{-1}}}(1, 0) = 0$. In the general case, a connection \mathbb{A} on $\det(\mathbb{S}) = K^{-1}E^2$ can be written as

$$(27) \quad \mathbb{A} = A_{K^{-1}} + 2A,$$

where A is a connection on E . Using (27), we henceforth identify a spin-c connection with a Hermitian connection A on E (instead of with a connection \mathbb{A} on $\det(\mathbb{S})$ as in Section 2.1), and denote its corresponding Dirac operator by D_A .

As in [21], given a spin-c structure \mathfrak{s} as above, we consider, for a connection A on E and a section ψ of \mathbb{S} , the following version of the Seiberg–Witten equations:

$$(28) \quad \begin{aligned} D_A \psi &= 0, \\ *F_A &= r(\tau(\psi) - i\lambda) - \frac{1}{2} *F_{A_{K^{-1}}} + i*\mu. \end{aligned}$$

Here r is a positive real number (which below we will take to be very large), and μ is an exact 2–form satisfying certain conditions described in the next paragraph. Under the identification (27), the equations (28) are equivalent to the Seiberg–Witten equations (20) with perturbation

$$(29) \quad \eta = -r d\lambda + 2\mu,$$

if we rescale the spinor by

$$(30) \quad \Psi = \sqrt{2r} \psi.$$

The 2–form μ above is a suitably generic exact smooth 2–form from a certain Banach space Ω of such forms defined in [21]. The precise details of Ω are not relevant here; we just need to recall the following facts. First, Ω is dense in the space of smooth exact 2–forms. Also Ω is a Banach subspace of \mathcal{P} , meaning that if $\mu \in \Omega$, then the equations (28) without further perturbation, together with the corresponding version of (23), namely (31) below, constitute one of the “abstract perturbations” from \mathcal{P} . The \mathcal{P} –norm of an element of Ω controls its derivatives to all orders. We always assume that the form μ in (29) has \mathcal{P} –norm less than 1 and C^0 –norm less than 1/100. Also the space Ω can be chosen so as to contain $d\lambda$, and this is assumed below.³ Finally, the spaces Ω and \mathcal{P} depend on the metric, and thus on the pair (λ, J) . However Ω and \mathcal{P} can be chosen so as to define smooth Banach space bundles over the space of metrics. We do not indicate this dependence of Ω and \mathcal{P} on the metric in the notation below.

The version of the Seiberg–Witten Floer chain complex obtained from solutions to (28) for a given data set (λ, J, r, μ) and abstract perturbation from \mathcal{P} (if necessary

³The fact that $d\lambda \in \Omega$ will be used in the proof of Lemma 5.7.

to obtain suitable transversality) is denoted below by $\widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda, J, r)$. Here we are suppressing μ and the abstract perturbation from the notation. The corresponding Seiberg–Witten Floer cohomology is denoted by $\widehat{\text{HM}}^*(Y, \mathfrak{s}; \lambda, J, r)$. The irreducible component $\widehat{\text{CM}}_{\text{irr}}^*$ of the chain complex is generated by irreducible solutions to (28). If (A_+, ψ_+) and (A_-, ψ_-) are two such irreducible solutions, then the component of (A_-, ψ_-) in the differential of (A_+, ψ_+) counts index 1 solutions to a correspondingly perturbed version of the equations

$$\begin{aligned}
 (31) \quad & \frac{\partial}{\partial s} \psi(s) = -D_{A(s)} \psi(s), \\
 & \frac{\partial}{\partial s} A(s) = -*F_{A(s)} + r(\tau(\psi(s)) - i\lambda) - \frac{1}{2} *F_{A_{K-1}} + i*\mu, \\
 & \lim_{s \rightarrow \pm\infty} (A(s), \psi(s)) = (A_{\pm}, \psi_{\pm}),
 \end{aligned}$$

modulo gauge equivalence and s -translation.

2.3 The energy filtration

When r above is large, the chain complex $\widehat{\text{CM}}^*$ has (up to some level) a filtration analogous to the symplectic action filtration on ECH. This works as follows. If (A, ψ) is a solution to (28), define the *energy*

$$(32) \quad E(A) := i \int_Y \lambda \wedge F_A.$$

The idea is that given an ECH generator Θ , if r is sufficiently large then there is a corresponding irreducible solution (A, ψ) to (28) for which the zero set of the E component of ψ (see equation (26)) is close to the Reeb orbits in Θ , the curvature F_A is concentrated in a radius $O(r^{-1/2})$ neighborhood of the Reeb orbits in Θ , and the energy $E(A)$ is approximately $2\pi\mathcal{A}(\Theta)$.

This motivates defining the following analogue of the filtered ECH chain complex ECC_*^L from Section 1.2: If L is a real number, define $\widehat{\text{CM}}_L^*$ to be the submodule of $\widehat{\text{CM}}_{\text{irr}}^*$ generated by irreducible solutions (A, ψ) to (28) with $E(A) < 2\pi L$.

Lemma 2.3 *Fix Y, λ, J as above and $L \in \mathbb{R}$. Suppose that λ has no orbit set of action exactly L . Fix r sufficiently large, and a 2-form $\mu \in \Omega$ with \mathcal{P} -norm less than 1 so that all irreducible solutions to (28) are cut out transversely. Then for every \mathfrak{s} and for every sufficiently small generic abstract perturbation:*

- (a) $\widehat{\text{CM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$ is a subcomplex of $\widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda, J, r)$.

(b) If $L' < L$ and if there is no orbit set with action in the interval $[L', L]$, then the inclusion

$$\widehat{\text{CM}}_{L'}^*(Y, \mathfrak{s}; \lambda, J, r) \longrightarrow \widehat{\text{CM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$$

is an isomorphism of chain complexes.

Proof First observe that if r is sufficiently large then all solutions (A, ψ) to the perturbed Seiberg–Witten equations (28) with $E(A) < 2\pi L$ are irreducible, because it follows from (28) that the energy of a reducible solution $(A, 0)$ to (28) is a linear, increasing function of r .

Now if we fix the spin-c structure \mathfrak{s} , then part (a) of the lemma follows from the first bullet in [21, Theorem 4.4], and part (b) of the lemma follows from [21, Proposition 4.12]. Lemma 2.3 (without the spin-c structure fixed) then follows from Lemma 2.4 below. \square

Lemma 2.4 *Given a real number \mathcal{E} , there exists an integer κ such that if r is sufficiently large, then at most κ spin-c structures admit solutions (A, ψ) to (28) with $E(A) < \mathcal{E}$.*

Proof Let (A, ψ) be a solution to (28) for some spin-c structure \mathfrak{s} . It follows from the curvature equation in (28) and the a priori estimates on ψ in [24, Lemma 2.3] that if r is sufficiently large, then the L^1 -norm of F_A over Y is bounded by $c_0 + c_1 E(A)$, where c_0 and c_1 are independent of r and \mathfrak{s} . This implies the lemma, because a bound on the L^1 -norm of F_A gives an upper bound on the absolute values of the pairings of $c_1(\mathfrak{s})$ with a set of generators for $H_2(Y)$. \square

When Lemma 2.3 is applicable, we denote the homology of the subcomplex $\widehat{\text{CM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$ of $\widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda, J, r)$ by $\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$. If r is larger than some (λ, J) -dependent constant, then this homology does not depend on the 2-form μ , the small abstract perturbation, or r . This follows from [21, Lemma 4.6], and a generalization is proved in Lemma 3.4 below. We always assume that r is sufficiently large as above so that $\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$ is well-defined and independent of r , although we keep r in the notation. We will see in Section 3 that this homology is isomorphic to $\text{ECH}_*^L(Y, \lambda, \text{PD}(\mathfrak{s} - \mathfrak{s}_\xi); J)$.

3 SWF cohomology and (filtered) ECH

We now explain the relation between filtered ECH and the filtered version of Seiberg–Witten Floer cohomology defined in Section 2.3. Along the way we review the construction of the isomorphism (1) between ECH and $\widehat{\text{HM}}^*$ and prove Theorem 1.3 on the J -independence of filtered ECH.

3.1 L -flat approximations

In order to define a chain map from the filtered ECH chain complex to the Seiberg–Witten Floer chain complex, it is convenient to modify the pair (λ, J) so that it has a certain nice form in a tubular neighborhood of each Reeb orbit of symplectic action less than L . Specifically, we say that the pair (λ, J) is L -flat if near each Reeb orbit of length less than L it satisfies the conditions in [21, Equation (4.1)]. (We do not need to recall these conditions in detail here.) The reasons for introducing this condition are discussed in [21, Section 5c, Part 2]. In particular, we have the following key fact:

Proposition 3.1 *Fix Y, λ, J and $L \in \mathbb{R}$. Suppose that λ is L -nondegenerate; see Definition 1.2. Then for all r sufficiently large, and for all $\Gamma \in H_1(Y)$, the following hold:*

- (a) *There is a canonical map from the set of generators of $\widehat{\text{CM}}_L^*(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda, J, r)$ to the set of orbit sets in the homology class Γ of length less than L .*
- (b) *If (λ, J) is L -flat, then the generators of $\widehat{\text{CM}}_L^*(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda, J, r)$ are cut out transversely, and the map in part (a) is a bijection from the set of these generators to the set of admissible orbit sets in the homology class Γ of length less than L .*
- (c) *Suppose (λ, J) is L -flat and J is ECH^L -generic. Fix a 2-form μ from Ω with \mathcal{P} -norm less than 1, and fix a small generic abstract perturbation. Then the bijection in part (b) induces an isomorphism of relatively graded chain complexes*

$$(33) \quad \text{ECC}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \widehat{\text{CM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda, J, r).$$

Proof Part (a) follows directly from [19, Section 6]. Part (b) follows from [21, Theorem 4.2], and part (c) follows from the second bullet of [21, Theorem 4.4]. \square

The basic picture for part (a) is that when r is large, generators (A, ψ) of $\widehat{\text{CM}}_L^*$ have F_A concentrated near a collection of Reeb orbits as described preceding Lemma 2.3, and this defines the map from generators of $\widehat{\text{CM}}_L^*$ to orbit sets with symplectic action less than L . The idea for part (c) is then that the instantons that define the differential on the right hand side of (33) correspond in a similar manner to the holomorphic curves that define the differential on the left hand side of (33).

To make use of the above proposition, we need to suitably approximate an arbitrary pair (λ, J) by an L -flat pair.

Definition 3.2 Let λ be an L -nondegenerate contact form, and let J be an ECH^L -generic symplectization-admissible almost complex structure for λ . An L -flat approximation to (λ, J) is an L -flat pair (λ_1, J_1) , which is the endpoint of a smooth homotopy $\{(\lambda_t, J_t) \mid t \in [0, 1]\}$ with $(\lambda_0, J_0) = (\lambda, J)$ such that:

- (i) For each $t \in [0, 1]$, λ_t is an L -nondegenerate contact form, and J_t is an ECH^L -generic symplectization-admissible almost complex structure for λ_t .
- (ii) The Reeb orbits of λ_t with length less than L , and their lengths, do not depend on t .

We will see in Lemma 3.6 below that L -flat approximations always exist. Note that if $\{(\lambda_t, J_t)\}$ is a homotopy as in Definition 3.2, then by condition (i) there is a canonical isomorphism of chain complexes

$$(34) \quad \text{ECC}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \text{ECC}_*^L(Y, \lambda_1, \Gamma; J_1),$$

induced by the canonical identification of generators from condition (ii). Combining this with the isomorphism (33) for (λ_1, J_1) , we conclude that if (λ_1, J_1) is an L -flat approximation to (λ, J) , and if r is sufficiently large, then there is a canonical isomorphism of chain complexes

$$(35) \quad \text{ECC}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \widehat{\text{CM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda_1, J_1, r).$$

3.2 Deforming λ and J

We now state and prove a key lemma regarding the behavior of $\widehat{\text{HM}}_L^*$ under certain special deformations of λ and J .

Definition 3.3 An *admissible deformation* is a smooth 1-parameter family $\rho = \{(\lambda_t, L_t, J_t, r_t) \mid t \in [0, 1]\}$ such that for all $t \in [0, 1]$:

- λ_t is an L_t -nondegenerate contact form on Y .
- J_t is a symplectization-admissible almost complex structure for λ_t .
- r_t is a positive real number.

The following is a slight generalization of [21, Lemmas 4.6 and 4.16].

Lemma 3.4 Let $\rho = \{(\lambda_t, L_t, J_t, r_t) \mid t \in [0, 1]\}$ be an admissible deformation. If the real numbers $\{r_t\}$ are sufficiently large, then for any $\mathfrak{s} \in \text{Spin}^c(Y)$, the admissible deformation ρ induces an isomorphism

$$(36) \quad \Phi_\rho: \widehat{\text{HM}}_{L_0}^*(Y, \mathfrak{s}; \lambda_0, J_0, r_0) \xrightarrow{\cong} \widehat{\text{HM}}_{L_1}^*(Y, \mathfrak{s}; \lambda_1, J_1, r_1)$$

with the following properties:

- (a) Φ_ρ is invariant under homotopy of admissible deformations.
- (b) If ρ_1 and ρ_2 are composable admissible deformations, then $\Phi_{\rho_1 \circ \rho_2} = \Phi_{\rho_1} \circ \Phi_{\rho_2}$.

(c) *The diagram*

$$(37) \quad \begin{array}{ccc} \widehat{\text{HM}}_{L_0}^*(Y, \mathfrak{s}; \lambda_0, J_0, r) & \xrightarrow{\Phi_\rho} & \widehat{\text{HM}}_{L_1}^*(Y, \mathfrak{s}; \lambda_1, J_1, r) \\ \downarrow & & \downarrow \\ \widehat{\text{HM}}^*(Y, \mathfrak{s}; \lambda_0, J_0, r) & \longrightarrow & \widehat{\text{HM}}^*(Y, \mathfrak{s}; \lambda_1, J_1, r) \end{array}$$

commutes, where the vertical arrows are induced by the inclusions of chain complexes, and the bottom arrow is the canonical isomorphism on Seiberg–Witten Floer cohomology.

(d) *If for all $t \in [0, 1]$, the pair (λ_t, J_t) is L_t –flat and J_t is ECH^{L_t} –generic, then under the isomorphism (35), the map Φ_ρ is induced by the isomorphism of chain complexes*

$$\text{ECC}_{-*}^{L_0}(Y, \lambda_0, PD(\mathfrak{s} - \mathfrak{s}_\xi); J_0) \longrightarrow \text{ECC}_{-*}^{L_1}(Y, \lambda_1, PD(\mathfrak{s} - \mathfrak{s}_\xi); J_1)$$

determined by the canonical bijection on generators.

Proof As explained in [13], the canonical isomorphism on Seiberg–Witten Floer cohomology at the bottom of (37) is induced by a chain map which is defined from a suitable 1–parameter family of data sets that interpolates between those used to define the two chain complexes. Various relevant aspects of this are summarized in [21, Section 3h]. In the case at hand, the relevant 1–parameter family of data sets has the form

$$\{D_t = (\lambda_t, J_t, r_t, \mu_t, \mathfrak{p}_t) \mid t \in [0, 1]\}.$$

Here $\{\mu_t \mid t \in [0, 1]\}$ is a smooth family of 2–forms in Ω with \mathcal{P} –norm less than 1, and $\{\mathfrak{p}_t \mid t \in [0, 1]\}$ is a generic smooth family of abstract perturbations with small \mathcal{P} –norm. More precisely, recall from Section 2.2 that Ω and \mathcal{P} are smooth Banach space bundles over the space of metrics on Y ; the families $\{\mu_t\}$ and $\{\mathfrak{p}_t\}$ are sections of the restrictions of these bundles to the path of metrics determined by $\{(\lambda_t, J_t)\}$. The family $\{\mathfrak{p}_t\}$ can and should be chosen so that for generic $t \in [0, 1]$, the necessary transversality holds so that the Seiberg–Witten Floer chain complex $\widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda_t, J_t, r_t)$ is defined.

To prove parts (a)–(c), let N be a large positive integer, and choose numbers $0 = t_0 < t_1 < \dots < t_N = 1$ with $t_i - t_{i-1} < 2/N$ for each $i = 1, \dots, N$, such that the chain complex $\widehat{\text{CM}}^*$ is defined for each data set D_{t_i} . As explained in [21, Section 3h Part 3], if $\{\mathfrak{p}_t \mid t \in [0, 1]\}$ is generic then for each $i = 1, \dots, N$, the family of data sets parametrized by $t \in [t_{i-1}, t_i]$ can be used to define a chain map

$$(38) \quad \widehat{I}_i: \widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda_{t_{i-1}}, J_{t_{i-1}}, r_{t_{i-1}}) \longrightarrow \widehat{\text{CM}}^*(Y, \mathfrak{s}; \lambda_{t_i}, J_{t_i}, r_{t_i}).$$

Let I_i denote the map on $\widehat{\text{HM}}^*$ induced by \widehat{I}_i . The canonical isomorphism on the bottom of (37) is then given by the composition $I_N \circ \dots \circ I_1$.

Since L_t varies continuously with t , it follows from a compactness argument that there exists $\varepsilon > 0$ such that for each $t \in [0, 1]$, the contact form λ_t has no orbit set with action in the interval $[L_t - \varepsilon, L_t + \varepsilon]$. If N is sufficiently large, then for each i and for each $t \in [t_{i-1}, t_i]$, we have $|L_t - L_{t_{i-1}}| < \varepsilon$, and in particular the contact form λ_t has no orbit set of action exactly $L_{t_{i-1}}$. It then follows from [21, Lemma 4.6] that if the numbers $\{r_t\}$ are sufficiently large, then \widehat{I}_i restricts to a chain map

$$\widehat{\text{CM}}^*_{L_{t_{i-1}}} \rightarrow \widehat{\text{CM}}^*_{L_{t_i}}$$

which induces an isomorphism

$$\widehat{\text{HM}}^*_{L_{t_{i-1}}}(Y, \mathfrak{s}; \lambda_{t_{i-1}}, J_{t_{i-1}}, r_{t_{i-1}}) \xrightarrow{\cong} \widehat{\text{HM}}^*_{L_{t_i}}(Y, \mathfrak{s}; \lambda_{t_i}, J_{t_i}, r_{t_i}).$$

Finally, it follows from Lemma 2.3(b) that, again if the numbers $\{r_t\}$ are sufficiently large, then there is an isomorphism

$$\widehat{\text{HM}}^*_{L_{t_{i-1}}}(Y, \mathfrak{s}; \lambda_{t_i}, J_{t_i}, r_{t_i}) \xrightarrow{\cong} \widehat{\text{HM}}^*_{L_{t_i}}(Y, \mathfrak{s}; \lambda_{t_i}, J_{t_i}, r_{t_i})$$

induced by the inclusion of one chain complex into the other, depending on which of $L_{t_{i-1}}$ and L_{t_i} is larger. We now define

$$\Phi_{\rho|_{[t_{i-1}, t_i]}}: \widehat{\text{HM}}^*_{L_{t_{i-1}}}(Y, \mathfrak{s}; \lambda_{t_{i-1}}, J_{t_{i-1}}, r_{t_{i-1}}) \xrightarrow{\cong} \widehat{\text{HM}}^*_{L_{t_i}}(Y, \mathfrak{s}; \lambda_{t_i}, J_{t_i}, r_{t_i})$$

to be the composition of the above two isomorphisms, and

$$\Phi_{\rho} := \Phi_{\rho|_{[t_{N-1}, t_N]}} \circ \dots \circ \Phi_{\rho|_{[t_0, t_1]}}.$$

A two-parameter version of the above subdivision construction, again using [21, Lemma 4.6] and assuming that the numbers $\{r_t\}$ are sufficiently large, shows that the map Φ_{ρ} on homology is independent of the choices made above and satisfies the homotopy invariance property (a). Properties (b) and (c) are then immediate from the construction.

We now show that property (d) holds for a given $\{(\lambda_t, J_t, L_t)\}$ provided that $\{r_t\}$ is sufficiently large. By Lemma 2.4, we can fix the spin-c structure \mathfrak{s} . Arguing by contradiction, suppose that for each positive integer j we have a path $\{r_{j,t} \mid t \in [0, 1]\}$ for which property (d) fails, with $\lim_{j \rightarrow \infty} \min_{t \in [0, 1]} r_{j,t} = +\infty$.

For each j , for each positive integer k , choose a path $\{p_{j,k,t} \mid t \in [0, 1]\}$ of abstract perturbations suitable for defining the map Φ_{ρ} , such that the following hold for each j, k, t :

- (i) $\mathfrak{p}_{j,k,t}$ has \mathcal{P} -norm less than k^{-1} .
- (ii) There are no $\mathfrak{p}_{j,k,t}$ -instantons that have negative index between generators of $\widehat{\text{CM}}_{L_t}^*(Y, \mathfrak{s}; \lambda_t, J_t, r_{j,t})$. (This can be arranged by the Sard–Smale Theorem as in [19, Section 7].)

Now fix j and k . Since property (d) fails for $\{r_{j,t}\}$, it follows that if we construct the corresponding map Φ_ρ using $\{\mathfrak{p}_{j,k,t}\}$, then for each N in the construction of Φ_ρ , there exists $i \in \{1, \dots, N\}$ such that the corresponding chain map \widehat{I}_i as in (38) is not the canonical bijection of generators. Taking $N \rightarrow \infty$, a compactness argument using (ii) then finds $t_{j,k} \in [0, 1]$ and an index zero, non- \mathbb{R} -invariant $\mathfrak{p}_{j,k,t_{j,k}}$ -instanton $\mathfrak{d}_{j,k}$ between two generators of

$$\widehat{\text{CM}}_{L_{t_{j,k}}}^*(Y, \mathfrak{s}; \lambda_{t_{j,k}}, J_{t_{j,k}}, r_{j,t_{j,k}}).$$

For each j , pass to a subsequence of the k 's such that the sequence $\{t_{j,k}\}$ converges to $t_j \in [0, 1]$. Next pass to a subsequence of the j 's such that t_j converges to $t_* \in [0, 1]$.

Given the doubly indexed sequence $\{\mathfrak{d}_{j,k}\}$ of $\mathfrak{p}_{j,k,t_{j,k}}$ -instantons constructed above, the argument in [24, Section 8(b)] can now be repeated almost verbatim⁴ to conclude the following: There exists a broken J_{t_*} -holomorphic curve in $\mathbb{R} \times Y$ between two generators of $\text{ECC}_{t_*}^{L_{t_*}}(Y, \lambda_{t_*}; J_{t_*})$, with each level non- \mathbb{R} -invariant as in Definition 1.6, and with total ECH index zero. But this contradicts the fact that J_{t_*} is $\text{ECH}^{L_{t_*}}$ -generic; see Hutchings and Sullivan [9, Corollary 11.5] or Hutchings [4, Proposition 3.7]. \square

We can now deduce that $\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$ does not depend on J or r .

Corollary 3.5 *Suppose λ is an L -nondegenerate contact form and \mathfrak{s} is a spin-c structure on Y . Then the relatively graded $\mathbb{Z}/2$ -modules $\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$ for different r and J (where r is sufficiently large with respect to λ, L, J) are canonically isomorphic to a single relatively graded $\mathbb{Z}/2$ -module $\widehat{\text{HM}}_L^*(Y, \lambda, \mathfrak{s})$, with the following properties:*

- (a) *Inclusion of chain complexes induces a well-defined map*

$$\widehat{\text{HM}}_L^*(Y, \lambda, \mathfrak{s}) \longrightarrow \widehat{\text{HM}}^*(Y, \mathfrak{s}).$$

- (b) *If $L < L'$ and if λ is also L' -nondegenerate, then inclusion of chain complexes induces a well-defined map*

$$\widehat{\text{HM}}_L^*(Y, \lambda, \mathfrak{s}) \longrightarrow \widehat{\text{HM}}_{L'}^*(Y, \lambda, \mathfrak{s}).$$

⁴Here one uses the stability condition in Remark 7.2 below to deal with the fact that t_j depends on j .

(c) If $c > 0$ then there is a canonical “scaling” isomorphism

$$(39) \quad s: \widehat{\text{HM}}_L^*(Y, \lambda, \mathfrak{s}) \xrightarrow{\simeq} \widehat{\text{HM}}_{cL}^*(Y, c\lambda, \mathfrak{s}).$$

Proof Since the space of symplectization-admissible almost complex structures for λ is contractible, it follows that if r_i is sufficiently large with respect to J_i for $i = 0, 1$, then Lemma 3.4(a) provides a well-defined isomorphism

$$\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J_0, r_0) \xrightarrow{\simeq} \widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J_1, r_1),$$

induced by an admissible deformation of the form $\rho = \{(\lambda, L, J_t, r_t)\}$. By Lemma 3.4(b), these isomorphisms satisfy the necessary composition property to identify the modules

$$\widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r)$$

for different J, r with a single relatively graded $\mathbb{Z}/2$ -module $\widehat{\text{HM}}_L^*(Y, \lambda, \mathfrak{s})$.

Property (a) now follows immediately from Lemma 3.4(c). Property (b) follows similarly from the construction of the maps Φ_ρ .

To prove property (c), fix J and fix r sufficiently large with respect to J . Consider the admissible deformation

$$(40) \quad \rho_c := \{((1 - t + ct)\lambda, (1 - t + ct)L, J, r)\}.$$

Here we are regarding J as an almost complex structure on ξ , so that the same J can be used for any positive multiple of λ . By Lemma 3.4, the admissible deformation (40) induces an isomorphism

$$\Phi_{\rho_c}: \widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r) \xrightarrow{\simeq} \widehat{\text{HM}}_{cL}^*(Y, \mathfrak{s}; c\lambda, J, r).$$

We claim that this isomorphism induces a well-defined isomorphism as in (39). To prove this, we need to check that given another pair (J', r') , if ρ'_c is the primed analogue of (40), then the diagram

$$\begin{array}{ccc} \widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J, r) & \xrightarrow{\Phi_{\rho_c}} & \widehat{\text{HM}}_{cL}^*(Y, \mathfrak{s}; c\lambda, J, r) \\ \downarrow \Phi_{\rho_1} & & \downarrow \Phi_{\rho_2} \\ \widehat{\text{HM}}_L^*(Y, \mathfrak{s}; \lambda, J', r') & \xrightarrow{\Phi_{\rho'_c}} & \widehat{\text{HM}}_{cL}^*(Y, \mathfrak{s}; c\lambda, J', r') \end{array}$$

commutes. Here $\rho_1 = \{(\lambda, L, J_t, r_t)\}$ and $\rho_2 = \{(c\lambda, cL, J_t, r_t)\}$, where $\{(J_t, r_t)\}$ is a homotopy from (J, r) to (J', r') . We now observe that both $\rho'_c \circ \rho_1$ and $\rho_2 \circ \rho_c$ are homotopic through admissible deformations to

$$\{((1 - t + ct)\lambda, (1 - t + ct)L, J_t, r_t)\},$$

and so commutativity of the above diagram follows from Lemma 3.4(a), (b). □

Below, when we are not concerned with the spin-c structure, we write

$$\widehat{HM}_L^*(Y, \lambda) := \bigoplus_{\mathfrak{s} \in \text{Spin}^c(Y)} \widehat{HM}_L^*(Y, \lambda, \mathfrak{s}).$$

3.3 The filtered isomorphism

We now define an isomorphism from filtered embedded contact homology to filtered Seiberg–Witten Floer cohomology, and describe how it behaves under scaling and inclusion of chain complexes. To obtain a *canonical* isomorphism, we will need the following lemma:

Lemma 3.6 [21, Proposition B.1] *If λ is L -nondegenerate and if J is ECH^L -generic, then there exist “preferred” L -flat approximations to (λ, J) , and for each preferred L -flat approximation (λ_1, J_1) there exist “preferred” homotopies $\{(\lambda_t, J_t) \mid t \in [0, 1]\}$ as in Definition 3.2, such that:*

- (a) *If (λ_1, J_1) is a preferred L -flat approximation, then any two preferred homotopies for (λ_1, J_1) are homotopic through admissible deformations.*
- (b) *If (λ_1^0, J_1^0) and (λ_1^1, J_1^1) are two preferred L -flat approximations, then they are connected by a homotopy of L -flat pairs $\{(\lambda_1^\nu, J_1^\nu) \mid \nu \in [0, 1]\}$ with the following properties:*
 - (i) *The Reeb orbits of λ_1^ν do not depend on ν .*
 - (ii) *$\{(\lambda_1^\nu, J_1^\nu) \mid \nu \in [0, 1]\}$ is homotopic through admissible deformations to the composition of a preferred homotopy for (λ_1^1, J_1^1) with the inverse of a preferred homotopy for (λ_1^0, J_1^0) .*
- (c) *For every $\varepsilon > 0$, there exists a preferred L -flat approximation (λ_1, J_1) with a preferred homotopy $\{(\lambda_t, J_t)\}$ such that each (λ_t, J_t) agrees with (λ, J) except within distance ε of the Reeb orbits of action less than L .*

Part (c) of the above lemma will be used in Section 6.3.

We can now relate filtered ECH to filtered Seiberg–Witten Floer cohomology:

Lemma 3.7 *Suppose that λ is L -nondegenerate and J is ECH^L -generic. Then for all $\Gamma \in H_1(Y)$, there is a canonical isomorphism of relatively graded $\mathbb{Z}/2$ -modules*

$$(41) \quad \Psi^L: \text{ECH}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \widehat{HM}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma})$$

with the following properties:

- (a) If $L < L'$, if λ is L' -nondegenerate, and if J is $\text{ECH}^{L'}$ -generic, then the diagram

$$\begin{CD} \text{ECH}_*^L(Y, \lambda, \Gamma; J) @>\Psi^L>> \widehat{\text{HM}}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}) \\ @V\iota_J^{L, L'}VV @VVV \\ \text{ECH}_*^{L'}(Y, \lambda, \Gamma; J) @>\Psi^{L'}>> \widehat{\text{HM}}_{L'}^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}) \end{CD}$$

commutes, where $\iota_J^{L, L'}$ is the inclusion-induced map (3), and the right arrow is the inclusion-induced map in Corollary 3.5(b).

- (b) If $c > 0$, then the diagram

$$\begin{CD} \text{ECH}_*^L(Y, \lambda, \Gamma; J) @>\Psi^L>> \widehat{\text{HM}}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}) \\ @V s_J VV @VV s V \\ \text{ECH}_*^{cL}(Y, c\lambda, \Gamma; J) @>\Psi^{cL}>> \widehat{\text{HM}}_{cL}^{-*}(Y, c\lambda, \mathfrak{s}_{\xi, \Gamma}) \end{CD}$$

commutes, where s_J is the scaling isomorphism (7), and s is the scaling isomorphism in Corollary 3.5(c).

Proof Let (λ_1, J_1) be a preferred L -flat approximation to (λ, J) , and let $\{(\lambda_t, J_t) \mid t \in [0, 1]\}$ be a preferred homotopy from (λ, J) to (λ_1, J_1) . If r is sufficiently large, then by (35) we have a canonical isomorphism

$$\text{ECH}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda_1, J_1, r).$$

By Lemma 3.4, the admissible deformation

$$(42) \quad \rho_1 = \{(\lambda_{1-t}, L, J_{1-t}, r) \mid t \in [0, 1]\}$$

determines an isomorphism

$$(43) \quad \Phi_{\rho_1}: \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda_1, J_1, r) \xrightarrow{\cong} \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda, J, r).$$

By Lemmas 3.6(a) and 3.4(a), the map (43) does not depend on the choice of preferred homotopy. Let

$$\widetilde{\Psi}^L: \text{ECH}_*^L(Y, \lambda, \Gamma; J) \xrightarrow{\cong} \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda, J, r)$$

denote the composition of the previous two isomorphisms. We claim that $\widetilde{\Psi}^L$ induces a well-defined map Ψ^L as in (41).

We first show that $\widetilde{\Psi}^L$ does not depend on the choice of preferred L -flat approximation. Given two preferred L -flat approximations (λ_1^0, J_1^0) and (λ_1^1, J_1^1) , let $\{(\lambda_1^v, J_1^v) \mid v \in [0, 1]\}$ be a homotopy of L -flat pairs provided by Lemma 3.6(b). By Lemma 3.4(a), (b), the isomorphisms (43) for the two preferred L -flat approximations differ by the isomorphism

$$\Phi_{\rho_2}: \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda_1^0, J_1^0, r) \xrightarrow{\cong} \widehat{\text{HM}}_L^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}; \lambda_1^1, J_1^1, r)$$

induced by the admissible deformation

$$\rho_2 = \{(\lambda_1^v, L, J_1^v, r) \mid v \in [0, 1]\}.$$

Applying Lemma 3.4(d) to the latter path then shows that the two versions of $\widetilde{\Psi}^L$ defined using the two preferred L -flat approximations agree.

We now show that Ψ^L does not depend on the choice of r . Suppose that r, r' are both sufficiently large to define the isomorphism $\widetilde{\Psi}_L$. To prove that the versions of Ψ_L defined using r and r' agree, it is enough to show that the following diagram commutes:

$$\begin{array}{ccccc} \text{ECH}_*^L(\lambda, J) & \xrightarrow{\cong} & \widehat{\text{HM}}_L^{-*}(\lambda_1, J_1, r) & \xrightarrow{\Phi_{\rho_1}} & \widehat{\text{HM}}_L^{-*}(\lambda, J, r) \\ \parallel & & \Phi_{\rho_3} \downarrow & & \Phi_{\rho_4} \downarrow \\ \text{ECH}_*^L(\lambda, J) & \xrightarrow{\cong} & \widehat{\text{HM}}_L^{-*}(\lambda_1, J_1, r') & \xrightarrow{\Phi_{\rho'_1}} & \widehat{\text{HM}}_L^{-*}(\lambda, J, r') \end{array}$$

Here we have dropped Y and Γ from the notation, the horizontal isomorphisms on the left are given by (35), the admissible deformation ρ'_1 is defined as in (42) but with r replaced by r' , and

$$\begin{aligned} \rho_3 &= \{(\lambda_1, L, J_1, (1-t)r + tr') \mid t \in [0, 1]\}, \\ \rho_4 &= \{(\lambda, L, J, (1-t)r + tr') \mid t \in [0, 1]\}. \end{aligned}$$

The left square commutes by Lemma 3.4(d). The right square commutes by Lemma 3.4(a), (b), because both $\rho_4 \circ \rho_1$ and $\rho'_1 \circ \rho_3$ are homotopic through admissible deformations to

$$\{(\lambda_{1-t}, L, J_{1-t}, (1-t)r + tr') \mid t \in [0, 1]\}.$$

This completes the proof that Ψ^L is well-defined.

To prove that Ψ^L satisfies property (a), choose a preferred L' -flat approximation (λ_1, J_1) to define $\Psi^{L'}$. Then this is also a preferred L -flat approximation that can be

used to define Ψ^L . It now suffices to show that the diagram

$$\begin{array}{ccccc}
 \text{ECH}_*^L(\lambda, J) & \xrightarrow{\cong} & \widehat{\text{HM}}_L^{-*}(\lambda_1, J_1, r) & \xrightarrow{\Phi_{\rho_1}} & \widehat{\text{HM}}_L^{-*}(\lambda, J, r) \\
 \downarrow \iota_J^{L, L'} & & \downarrow & & \downarrow \\
 \text{ECH}_*^{L'}(\lambda, J) & \xrightarrow{\cong} & \widehat{\text{HM}}_{L'}^{-*}(\lambda_1, J_1, r) & \xrightarrow{\Phi_{\rho_1''}} & \widehat{\text{HM}}_{L'}^{-*}(\lambda, J, r)
 \end{array}$$

commutes. Here ρ_1'' is defined as in (42) but with L replaced by L' , and the vertical arrows in the diagram are induced by inclusions of chain complexes. Now the left square commutes by the definition of the isomorphism (33), while the right square commutes by a straightforward analogue of Lemma 3.4(c).

To prove property (b), let us further drop r from the notation and consider the diagram:

$$\begin{array}{ccccccc}
 \text{ECH}_*^L(\lambda, J) & \xrightarrow{\cong} & \text{ECH}_*^L(\lambda_1, J_1) & \xrightarrow{\cong} & \widehat{\text{HM}}_L^{-*}(\lambda_1, J_1) & \xrightarrow{\Phi_{\rho_1}} & \widehat{\text{HM}}_L^{-*}(\lambda, J) \\
 \downarrow s_J & & \downarrow s_{J_1} & & \downarrow \Phi_{\rho_c^1} & & \downarrow \Phi_{\rho_c} \\
 \text{ECH}_*^{cL}(c\lambda, J) & \xrightarrow{\cong} & \text{ECH}_*^{cL}(c\lambda_1, J_1) & \xrightarrow{\cong} & \widehat{\text{HM}}_{cL}^{-*}(c\lambda_1, J_1) & \xrightarrow{\Phi_{\rho_1^c}} & \widehat{\text{HM}}_{cL}^{-*}(c\lambda, J)
 \end{array}$$

Here ρ_c was defined in (40), ρ_c^1 denotes the analogue of (40) for (λ_1, J_1) , and ρ_1^c is obtained from (42) by multiplying the contact forms and L by c . Also the horizontal isomorphisms on the left are induced by (34), and the horizontal isomorphisms in the middle are induced by (33). By definition, the composition of the horizontal arrows in the top row of the above diagram is Ψ^L , and the composition of the horizontal arrows in the bottom row is Ψ^{cL} . So to prove property (b) it is enough to show that the above diagram commutes. The left square commutes at the chain level because each map in the left square sends each admissible orbit set to itself. The middle square commutes by Lemma 3.4(d). The right square commutes by Lemma 3.4(a), (b), because both $\rho_c \circ \rho_1$ and $\rho_1^c \circ \rho_c^1$ are homotopic to

$$\{(1-t+ct)\lambda_{1-t}, (1-t+ct)L, J_{1-t}, r) \mid t \in [0, 1]\}$$

through admissible deformations. □

3.4 J -independence of filtered ECH (proof)

We now have enough machinery in place to prove Theorem 1.3, asserting that ECH and ECH^L do not depend on the choice of almost complex structure used to define them.

Proof of Theorem 1.3 We may assume, by slightly decreasing L if necessary, that there is no orbit set of action exactly L . Part (a) then follows from the canonical isomorphism (41) given by Lemma 3.7. Part (b) follows from Lemma 3.7(a). Part (c) follows from part (b) by taking direct limits. Part (d) follows from Lemma 3.7(b). Part (e) follows from the definition of the isomorphism between ECH and \widehat{HM}^* reviewed in Section 3.5 below. \square

Remark 3.8 ECH has various additional structures on it that we are not using in this paper, for example a degree -2 map U . It is shown in [25] that these agree with analogous structures on Seiberg–Witten Floer cohomology under the isomorphism determined by (41) (see Section 3.5 below). Consequently the proof of Theorem 1.3 shows that these additional structures are also independent of J .

3.5 The full isomorphism

We are now in a position to write down the full isomorphism from embedded contact homology to Seiberg–Witten Floer cohomology.

Let Y be a closed oriented connected 3–manifold with a nondegenerate contact form λ , and fix $\Gamma \in H_1(Y)$. By Lemma 3.7, if λ has no orbit set of action L , then for each $\Gamma \in H_1(Y)$ there is a well-defined isomorphism

$$(44) \quad \text{ECH}_*^L(Y, \lambda, \Gamma) \xrightarrow{\cong} \widehat{HM}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}).$$

By Corollary 3.5(a), there is a well-defined map

$$(45) \quad \widehat{HM}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}) \longrightarrow \widehat{HM}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma}).$$

We now define

$$(46) \quad T^L: \text{ECH}_*^L(Y, \lambda, \Gamma) \longrightarrow \widehat{HM}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma})$$

to be the composition of the maps (44) and (45) above.

If $L < L'$, then it follows from Lemma 3.7(a) that

$$T^L = T^{L'} \circ \iota^{L, L'},$$

where $\iota^{L, L'}$ is the inclusion-induced map (6). This means that it makes sense to define

$$(47) \quad T: \text{ECH}_*(Y, \lambda, \Gamma) \longrightarrow \widehat{HM}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma})$$

to be the direct limit over L of the maps T^L in (46). The main theorem of [21] (after passing to $\mathbb{Z}/2$ coefficients) can now be stated as follows:

Theorem 3.9 [21] *The map T in (47) is an isomorphism of relatively graded $\mathbb{Z}/2$ -modules.*

Knowing that (44) is an isomorphism, the rest of the proof of Theorem 3.9 amounts to showing that the maps (45) induce an isomorphism

$$\lim_{\rightarrow} \widehat{\text{HM}}_L^{-*}(Y, \lambda, \mathfrak{s}_{\xi, \Gamma}) \xrightarrow{\cong} \widehat{\text{HM}}^{-*}(Y, \mathfrak{s}_{\xi, \Gamma});$$

see [21, Theorem 4.5]. (This is not immediately obvious because one has to increase r as one increases L in order to define the left hand side; see Section 2.3.)

4 Seiberg–Witten Floer cobordism maps and symplectic forms

We now review from [13, Chapter 24] the maps on Seiberg–Witten Floer cohomology induced by a (smooth) cobordism. We then introduce a perturbation of the relevant Seiberg–Witten equations on an exact symplectic cobordism using the symplectic form.

4.1 Smooth cobordisms

Let Y_+ and Y_- be closed oriented (connected) three-manifolds. Let X be a cobordism from Y_+ to Y_- as in Section 1.4.

Given some metric on X , a *spin-c structure* on X is a lift of the frame bundle of TX from $\text{SO}(4)$ to

$$\text{Spin}^c(4) = \text{Spin}(4) \times_{\mathbb{Z}/2} U(1).$$

This is equivalent to a Hermitian vector bundle $\mathbb{S} = \mathbb{S}_+ \oplus \mathbb{S}_-$, where \mathbb{S}_+ and \mathbb{S}_- each have rank 2, together with a Clifford multiplication $\text{cl}: TX \rightarrow \text{End}(\mathbb{S})$ satisfying (16), such that $\text{cl}(v)$ exchanges \mathbb{S}_+ and \mathbb{S}_- for each $v \in TX$, and

$$\text{cl}(e_1) \text{cl}(e_2) \text{cl}(e_3) \text{cl}(e_4) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$$

on $\mathbb{S}_+ \oplus \mathbb{S}_-$ whenever (e_1, e_2, e_3, e_4) is an oriented orthonormal basis for $T_x X$. The set $\text{Spin}^c(X)$ of isomorphism classes of spin-c structures on X is an affine space over $H^2(X; \mathbb{Z})$, with the action as in (17), which does not depend on the choice of metric. Given a spin-c structure on X , a spin-c connection is defined as in (18). A spin-c connection $\mathbb{A}_{\mathbb{S}}$ is equivalent to a Hermitian connection \mathbb{A} on $\det(\mathbb{S}_+) = \det(\mathbb{S}_-)$, and adding an imaginary-valued 1-form a to \mathbb{A} adds $a/2$ to $\mathbb{A}_{\mathbb{S}}$. As in (19), the connection \mathbb{A} defines a Dirac operator

$$D_{\mathbb{A}}: C^\infty(X; \mathbb{S}_{\pm}) \longrightarrow C^\infty(X; \mathbb{S}_{\mp}).$$

A spin-c structure \mathfrak{s} on X restricts to a spin-c structure $\mathfrak{s}|_{Y_{\pm}}$ on Y_{\pm} as follows. Let v denote the outward pointing unit normal vector to Y_+ , and the inward pointing unit normal vector to Y_- . If $\mathbb{S} = \mathbb{S}_+ \oplus \mathbb{S}_-$ is the spin bundle for \mathfrak{s} with Clifford multiplication cl , then we define the spin bundle $\mathbb{S}_{Y_{\pm}}$ for $\mathfrak{s}|_{Y_{\pm}}$ to be

$$(48) \quad \mathbb{S}_{Y_{\pm}} := (\mathbb{S}_+)|_{Y_{\pm}}$$

with the Clifford action $TY \rightarrow \text{End}(\mathbb{S}_{Y_{\pm}})$ given by $\text{cl}(v)^{-1} \text{cl}(\cdot)$.

If \mathfrak{s} is a spin-c structure on X with $\mathfrak{s}_{\pm} := \mathfrak{s}|_{Y_{\pm}}$, then there is a cobordism map (of ungraded $\mathbb{Z}/2$ -modules)

$$(49) \quad \widehat{\text{HM}}^*(X, \mathfrak{s}): \widehat{\text{HM}}^*(Y_+, \mathfrak{s}_+) \longrightarrow \widehat{\text{HM}}^*(Y_-, \mathfrak{s}_-).$$

We now review the basic formalism of the definition of this map; the details are explained in [13]. Choose a metric g_{\pm} , exact 2-form η_{\pm} , and abstract perturbation \mathfrak{p}_{\pm} as needed to define the chain complex $\widehat{\text{CM}}^*(Y_{\pm}, \mathfrak{s}_{\pm}; g_{\pm}, \eta_{\pm})$. One defines a chain map

$$(50) \quad \widehat{\text{CM}}^*(Y_+, \mathfrak{s}_+; g_+, \eta_+) \longrightarrow \widehat{\text{CM}}^*(Y_-, \mathfrak{s}_-; g_-, \eta_-)$$

as follows. Attach cylindrical ends to X to obtain

$$\bar{X} := ((-\infty, 0] \times Y_-) \cup_{Y_-} X \cup_{Y_+} ([0, \infty) \times Y_+).$$

Choose a metric g on \bar{X} that on the ends agrees with the product of the standard metric on $(-\infty, 0]$ or $[0, \infty)$ with the chosen metric g_{\pm} on Y_{\pm} . Choose a self-dual 2-form η on \bar{X} that on each end agrees with the self-dual part of (the pullback of) η_{\pm} , namely $\frac{1}{2}(\eta_{\pm} + *\eta_{\pm})$, where $*$ denotes the Hodge star on \bar{X} . The spin-c structure on X has a canonical extension over \bar{X} , so that on each end, \mathbb{S}_+ and \mathbb{S}_- are both identified with the boundary spinor bundle, and if s denotes the $(-\infty, 0]$ or $[0, \infty)$ coordinate, then $\text{cl}(\partial_s): \mathbb{S}_+ \xrightarrow{\cong} \mathbb{S}_-$ preserves the identifications with the boundary spinor bundle.

We now consider solutions to the Seiberg–Witten equations on \bar{X} . These equations concern a pair (\mathbb{A}, Ψ) , where \mathbb{A} is a connection on $\det(\mathbb{S}_+)$ and Ψ is a section of \mathbb{S}_+ . Without abstract perturbation terms (which we will describe shortly), the equations are

$$(51) \quad \begin{aligned} D_{\mathbb{A}} \Psi &= 0, \\ F_{\mathbb{A}}^+ &= \frac{1}{2} \rho(\Psi) + i\eta. \end{aligned}$$

Here $F_{\mathbb{A}}^+$ denotes the self-dual part of the curvature $F_{\mathbb{A}}$, and $\rho: \mathbb{S}_+ \rightarrow \bigwedge_+^2 T^*X$ is a quadratic bundle map defined by

$$\rho(\Psi)(v, w) = -\frac{1}{2} \langle [\text{cl}(v), \text{cl}(w)] \Psi, \Psi \rangle$$

for $\Psi \in (\mathbb{S}_+)_x$ and $v, w \in T_x X$. The gauge group $C^\infty(X; S^1)$ acts on the set of solutions, again by (22).

A connection \mathbb{A} on $\det(\mathbb{S}_+)$ is in *temporal gauge* on the ends if on $(-\infty, 0] \times Y_-$ and $[0, \infty) \times Y_+$ one has

$$(52) \quad \nabla_{\mathbb{A}} = \frac{\partial}{\partial s} + \nabla_{\mathbb{A}(s)},$$

where $\mathbb{A}(s)$ is a connection on the bundle $\det(\mathbb{S}_{Y_\pm})$ over the 3-manifold Y_\pm , depending on s . Any connection can be placed into temporal gauge by an appropriate gauge transformation. After this has been done, the equations (51) on the ends are equivalent to the instanton equations (23).

To define cobordism maps, we also need to consider abstract perturbations of the equations (51). Suppose that p_+ and p_- are abstract perturbations for use in defining the perturbations of the equations (20) and (23) on Y_+ and Y_- . It is explained in [13, Chapter 11] how these are extended as an abstract perturbation p over all of \bar{X} . The resulting perturbation of (51) agrees on $(-\infty, 0] \times Y_-$ or $[0, \infty) \times Y_+$ with the corresponding perturbation of (23) via p_- or p_+ . Any such extension must be suitably generic in order to use the solutions of the perturbed version of (51) to define the chain map (50). In particular, a nonzero extension may be necessary even when p_- and p_+ are both zero.

Let $(\mathbb{A}_\pm, \Psi_\pm)$ be solutions to the three-dimensional Seiberg–Witten equations (20) for $(Y_\pm, \mathfrak{s}_\pm; g_\pm, \eta_\pm)$. We are interested in solutions to the abstract perturbation of the four-dimensional Seiberg–Witten equations (51) that on the ends are in temporal gauge and satisfy the convergence conditions

$$(53) \quad \begin{aligned} \lim_{s \rightarrow \infty} (\mathbb{A}(s), \Psi(s)) &= (\mathbb{A}_+, \Psi_+) \quad \text{as } s \rightarrow +\infty \text{ on } [0, \infty) \times Y_+, \\ \lim_{s \rightarrow -\infty} (\mathbb{A}(s), \Psi(s)) &= (\mathbb{A}_-, \Psi_-) \quad \text{as } s \rightarrow -\infty \text{ on } (-\infty, 0] \times Y_-. \end{aligned}$$

A solution to the perturbed equations (51) satisfying (53) will be called an “instanton from (\mathbb{A}_-, Ψ_-) to (\mathbb{A}_+, Ψ_+) ”. We often denote an instanton as above by \mathfrak{d} and write $\mathfrak{d}|_s := (\mathbb{A}(s), \Psi(s))$ and $\mathfrak{c}_\pm := (\mathbb{A}_\pm, \Psi_\pm)$. Every instanton has an *index*, which is the expected dimension of the corresponding component of the moduli space of instantons (with the same asymptotic decay rate as $s \rightarrow +\infty$ or $s \rightarrow -\infty$ if (\mathbb{A}_+, Ψ_+) or (\mathbb{A}_-, Ψ_-) , respectively, is reducible) modulo gauge equivalence. The component of the chain map (50) from an irreducible generator (\mathbb{A}_+, Ψ_+) to an irreducible generator (\mathbb{A}_-, Ψ_-) counts index zero instantons from (\mathbb{A}_-, Ψ_-) to (\mathbb{A}_+, Ψ_+) modulo gauge equivalence. All we need to know about the remaining components of (50) is the following: if (\mathbb{A}_+, Ψ_+) is irreducible, and if there are no index zero instantons to

(\mathbb{A}_+, Ψ_+) from a reducible (\mathbb{A}_-, Ψ_-) , then the chain map (50) sends (\mathbb{A}_+, Ψ_+) to an element of $\widehat{\text{CM}}_{\text{irr}}^*(Y_-, \dots)$.

Although the chain map (50) may depend on the abstract perturbations, the induced map on homology

$$(54) \quad \widehat{\text{HM}}^*(X, \mathfrak{s}; g, \eta): \widehat{\text{HM}}^*(Y_+, \mathfrak{s}_+; g_+, \eta_+) \longrightarrow \widehat{\text{HM}}^*(Y_-, \mathfrak{s}_-; g_-, \eta_-)$$

does not. To show that this map does not depend on the extension \mathfrak{p} of \mathfrak{p}_+ and \mathfrak{p}_- , given a homotopy of extensions \mathfrak{p} one defines a chain homotopy between the corresponding chain maps by counting index -1 instantons. The proof that the map (54) does not depend on \mathfrak{p}_+ or \mathfrak{p}_- either is a special case of a more general argument, which we will outline shortly.

In the special case when X is a product cobordism $[0, 1] \times Y$, the maps (54) define the canonical isomorphisms that prove that $\widehat{\text{HM}}^*(Y, \mathfrak{s}; g, \eta)$ does not depend⁵ on g or η (or the abstract perturbations that we are suppressing from the notation). The necessary composition property for these isomorphisms follows from the following more general composition property. Let X^+ be a cobordism from Y_+ to Y_0 , let X^- be a cobordism from Y_0 to Y_- , and let $X = X^- \cup_{Y_0} X^+$ be the composite cobordism from Y_+ to Y_- . If $\mathfrak{s}^\pm \in \text{Spin}^c(X^\pm)$, if (g^\pm, η^\pm) are choices to define the cobordism map on X^\pm , and if (g, η) are choices to define the cobordism map on X , then

$$(55) \quad \widehat{\text{HM}}^*(X^-, \mathfrak{s}^-; g^-, \eta^-) \circ \widehat{\text{HM}}^*(X^+, \mathfrak{s}^+; g^+, \eta^+) = \sum_{\{\mathfrak{s} \in \text{Spin}^c(X) \mid \mathfrak{s}|_{X^\pm} = \mathfrak{s}^\pm\}} \widehat{\text{HM}}^*(X, \mathfrak{s}; g, \eta).$$

Note here that the sum on the right is well-defined, because by [13, Proposition 24.6.6] the cobordism map (54) is nonzero for only finitely many spin-c structures on X . Equation (55) is proved by “stretching the neck” along Y_0 and counting index -1 instantons to define a chain homotopy between the corresponding chain maps.

The special case of (55) when X^+ and X^- are both product cobordisms gives the composition property needed to show that $\widehat{\text{HM}}^*(Y, \mathfrak{s})$ is well-defined. The special case of (55) when just one of X^+ or X^- is a product cobordism then implies that the map (54) induces a well-defined map (49). With these identifications, (55) now translates to

$$(56) \quad \widehat{\text{HM}}^*(X^-, \mathfrak{s}^-) \circ \widehat{\text{HM}}^*(X^+, \mathfrak{s}^+) = \sum_{\{\mathfrak{s} \in \text{Spin}^c(X) \mid \mathfrak{s}|_{X^\pm} = \mathfrak{s}^\pm\}} \widehat{\text{HM}}^*(X, \mathfrak{s}).$$

⁵[21, Section 3h] says more about this in the case when g and η are determined by a contact form as in (28) and (31).

One can also combine the cobordism maps (49) into a single cobordism map

$$(57) \quad \widehat{\text{HM}}^*(X) := \sum_{\mathfrak{s} \in \text{Spin}^c(X)} \widehat{\text{HM}}^*(X, \mathfrak{s}): \widehat{\text{HM}}^*(Y_+) \longrightarrow \widehat{\text{HM}}^*(Y_-).$$

The composition property (56) implies that

$$\widehat{\text{HM}}^*(X) = \widehat{\text{HM}}^*(X_-) \circ \widehat{\text{HM}}^*(X_+).$$

Note that when X is not a product, the cobordism map (57) generally does not preserve the relative gradings, although there is a weaker relation between the gradings of the inputs and outputs of this map explained in [13, Section 3.4]. We will simply regard (57) as a map of ungraded $\mathbb{Z}/2$ -modules.

4.2 Perturbing the equations on an exact symplectic cobordism

We now introduce a useful perturbation of the four-dimensional Seiberg–Witten equations on an exact symplectic cobordism. This is closely related to the perturbation of the three-dimensional Seiberg–Witten equations defined in Section 2.2.

Let (Y_+, λ_+) and (Y_-, λ_-) be closed oriented (connected) 3-manifolds with contact forms. Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) . Recall the notion of “cobordism-admissible almost complex structure” from Definition 1.5. Below it will be convenient to work with a slightly stronger notion. Note that if $\varepsilon > 0$ is as in (12) and (13), then the completion \overline{X} in (14) contains subsets identified with $(-\infty, \varepsilon] \times Y_-$ and $[-\varepsilon, \infty) \times Y_+$.

Definition 4.1 An almost complex structure J on \overline{X} is *strongly cobordism-admissible* if it is ω -compatible on X , and if it agrees with symplectization-admissible almost complex structures J_+ for λ_+ on $[-\varepsilon, \infty) \times Y_+$ and J_- for λ_- on $(-\infty, \varepsilon] \times Y_-$, for some $\varepsilon > 0$ as in (12) and (13).

Given ε and J as above, we define a 1-form $\tilde{\lambda}$ on \overline{X} as follows. Fix a smooth increasing function $\phi_-: (-\infty, \varepsilon] \rightarrow (-\infty, \varepsilon]$ with $\phi_-(s) = 2s$ for $s \leq \varepsilon/10$ and $\phi_-(s) = s$ for $s > \varepsilon/2$. Likewise fix a smooth increasing function $\phi_+: [-\varepsilon, \infty) \rightarrow [-\varepsilon, \infty)$ with $\phi_+(s) = s$ for $s \leq -\varepsilon/2$ and $\phi_+(s) = 2s$ for $s \geq -\varepsilon/10$. Now define

$$(58) \quad \tilde{\lambda} := \begin{cases} e^{\phi_-} \lambda_- & \text{on } (-\infty, \varepsilon] \times Y_-, \\ \lambda & \text{on } X \setminus (([0, \varepsilon] \times Y_-) \cup ([-\varepsilon, 0] \times Y_+)), \\ e^{\phi_+} \lambda_+ & \text{on } [-\varepsilon, \infty) \times Y_+. \end{cases}$$

Write $\tilde{\omega} = d\tilde{\lambda}$; this is a symplectic form on all of \overline{X} . Also, J is $\tilde{\omega}$ -compatible on all of \overline{X} .

Remark 4.2 It would be more usual to define $\tilde{\lambda}$ by extending the 1-form λ on all of X to agree with $e^s \lambda_+$ on $[0, \infty) \times Y_+$, and with $e^s \lambda_-$ on $(-\infty, 0] \times Y_-$. We are using the more nonstandard 1-form (58) because of the factors of 2 discussed in Remark 2.2.

We next define a metric g on \bar{X} as follows. Let g_{\pm} denote the metric on Y_{\pm} determined by λ_{\pm} and J_{\pm} as in Section 2.2. Fix a smooth positive function σ_- on $(-\infty, \varepsilon]$ such that $\sigma_-(s) = 2e^{2s}$ for $s \leq \varepsilon/10$ and $\sigma_-(s) = 2$ for $s \geq \varepsilon/2$. Likewise fix a smooth positive function σ_+ on $[-\varepsilon, \infty)$ such that $\sigma_+(s) = 2$ for $s \leq -\varepsilon/2$ and $\sigma_+(s) = 2e^{2s}$ for $s \geq -\varepsilon/10$. Also require that⁶

$$(59) \quad \sigma_{\pm}(s) \in [3/2, 5/2] \quad \text{for } \pm s \in [0, \varepsilon].$$

Define a positive function σ on \bar{X} to equal σ_- on $(-\infty, \varepsilon] \times Y_-$, to equal σ_+ on $[-\varepsilon, \infty) \times Y_+$, and to equal 2 on the rest of \bar{X} . Define a metric g on \bar{X} by

$$(60) \quad g(\cdot, \cdot) = \sigma^{-1} \tilde{\omega}(\cdot, J(\cdot)).$$

Note that g agrees with the product metric with g_{\pm} on the ends $[0, \infty) \times Y_+$ and $(-\infty, 0] \times Y_-$. Also, $\tilde{\omega}$ is self-dual with respect to g and has norm $|\tilde{\omega}| = \sqrt{2} \sigma$. Define $\hat{\omega} = \sqrt{2} \tilde{\omega} / |\tilde{\omega}| = \sigma^{-1} \tilde{\omega}$.

Let \mathfrak{s} be a spin-c structure on X with spinor bundle $\mathbb{S} = \mathbb{S}_+ \oplus \mathbb{S}_-$. There is a canonical decomposition

$$(61) \quad \mathbb{S}_+ = E \oplus K^{-1} E$$

into eigenbundles of $\text{cl}(\hat{\omega})$, where E is the $-2i$ eigenbundle, and K denotes the canonical bundle of (X, J) . Note that on $[0, \infty) \times Y_+$ or $(-\infty, 0] \times Y_-$, under the identification (48), this splitting agrees with the splitting determined by Clifford multiplication by λ_+ or λ_- as in (26). When E is the trivial line bundle $\underline{\mathbb{C}}$, there is a distinguished connection $A_{K^{-1}}$ on K^{-1} such that $D_{A_{K^{-1}}}(1, 0) = 0$. As in the three-dimensional case (27), this allows us to identify a spin-c connection for a general spin-c structure with a Hermitian connection A on the corresponding line bundle E .

Now choose exact 2-forms μ_{\pm} on Y_{\pm} as in Section 2.2, and let μ be an exact 2-form on \bar{X} that agrees with μ_{\pm} on the ends. For the arguments later in this paper we need to choose μ so that its derivatives up to some sufficiently large (but constant) order have absolute value less than $1/100$. Let μ_* denote the self-dual part of μ . We now consider, for a connection A on E and a section ψ of \mathbb{S}_+ , the following version of

⁶The condition (59) will be used in Lemma 7.6.

the four-dimensional Seiberg–Witten equations on \overline{X} :

$$(62) \quad \begin{aligned} D_A \psi &= 0, \\ F_A^+ &= \frac{r}{2} (\rho(\psi) - i\hat{\omega}) - \frac{1}{2} F_{A_{K-1}}^+ + i\mu_*. \end{aligned}$$

Here r is a positive real number, which will be taken to be very large below. The equations (62) are equivalent to the equations (51) with perturbation

$$(63) \quad \eta = -r\hat{\omega} + 2\mu_*$$

after rescaling the spinor as in (30). On $[\varepsilon, \infty) \times Y_+$ and $(-\infty, -\varepsilon] \times Y_-$, if A is in temporal gauge, then the equations (62) are equivalent to the perturbed instanton equations (31) (with a \pm subscript on μ). Thus we can use the equations (51) (with appropriate small abstract perturbations) to define a chain map

$$(64) \quad \widehat{\text{CM}}^*(X, \mathfrak{s}; \lambda, J, r): \widehat{\text{CM}}^*(Y_+, \mathfrak{s}_+; \lambda_+, J_+, r) \longrightarrow \widehat{\text{CM}}^*(Y_-, \mathfrak{s}_-; \lambda_-, J_-, r).$$

Here \mathfrak{s}_\pm denotes the restriction of $\mathfrak{s} \in \text{Spin}^c(X)$ to Y_\pm . In general we expect the chain map (64) to depend on the choice of μ (and on the choice of abstract perturbations), although as explained in Section 4.1 the induced map on homology does not.

5 ECH cobordism maps

The goal of this section is to define the maps on (filtered) ECH induced by an exact symplectic cobordism, and to prove that they satisfy all of the axioms in Theorem 1.9, except for the holomorphic curves axiom, which will be proved in Section 6.

5.1 Cobordism maps and holomorphic curves (statements)

We now state some key properties of the map on Seiberg–Witten Floer cohomology induced by an exact symplectic cobordism with the Seiberg–Witten equations perturbed as in Section 4.2. To simplify notation we henceforth ignore the decomposition via spin-c structures, as in (57), although it is straightforward to insert spin-c structures into the discussion below.

The following proposition asserts that the instantons that are used to define the chain map (64) give rise to broken holomorphic curves, and in particular respect the symplectic action filtration. It also proves similar statements for certain chain homotopies, for which we need the following strengthening of the notion of “homotopy of exact symplectic cobordisms” defined in Section 1.5:

Definition 5.1 Two exact symplectic cobordisms (X, λ_0) and (X, λ_1) from (Y_+, λ_+) to (Y_-, λ_-) with the same underlying four-manifold X are *strongly homotopic* if there is a smooth one-parameter family of 1-forms $\{\lambda_t \mid t \in [0, 1]\}$ on X such that (X, λ_t) is an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) for each $t \in [0, 1]$, and there exists $\varepsilon > 0$ such that the identifications (12) and (13) for λ_t do not depend on t .

Note that the last condition in the above definition ensures that the completions (14) of (X, λ_t) for different t are diffeomorphic via the obvious identification.

Proposition 5.2 Fix $L \in \mathbb{R}$, closed connected contact 3-manifolds (Y_+, λ_+) and (Y_-, λ_-) such that λ_{\pm} are L -nondegenerate, symplectization-admissible almost complex structures J_{\pm} for λ_{\pm} , 2-forms μ_{\pm} on Y_{\pm} with \mathcal{P} -norm less than 1, and generic perturbations \mathfrak{p}_{\pm} on Y_{\pm} as needed to define the chain complexes $\widehat{\text{CM}}^*(Y_{\pm}; \lambda_{\pm}, J_{\pm}, r)$.

(a) Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) . Suppose J is a strongly cobordism-admissible almost complex structure on \overline{X} that restricts to J_+ on $[0, \infty) \times Y_+$ and to J_- on $(-\infty, 0] \times Y_-$. Let μ be a small exact 2-form on \overline{X} extending μ_{\pm} , and let \mathfrak{p} be a generic extension of \mathfrak{p}_{\pm} over \overline{X} . Assume that r is sufficiently large, and that \mathfrak{p}_{\pm} and \mathfrak{p} are sufficiently small for the given r . Let \mathfrak{d} be a solution to the corresponding perturbed version of (62) with index 0 and with $E(\mathfrak{c}_+) < 2\pi L$. Then:

- (i) $E(\mathfrak{c}_-) < 2\pi L$.
- (ii) There exists a broken J -holomorphic curve from Θ^+ to Θ^- , where Θ^{\pm} is the orbit set determined by \mathfrak{c}_{\pm} via Proposition 3.1(a).

(b) Let $\{(X, \lambda_t) \mid t \in [0, 1]\}$ be a strong homotopy of exact symplectic cobordisms from (Y_+, λ_+) to (Y_-, λ_-) . Let $\{(J_t, \mu_t, \mathfrak{p}_t) \mid t \in [0, 1]\}$ be a one-parameter family of choices as in part (a) with $\{\mathfrak{p}_t\}$ generic. Suppose that r is sufficiently large and that \mathfrak{p}_{\pm} and each \mathfrak{p}_t are sufficiently small for the given r . Let $t \in [0, 1]$ and let \mathfrak{d} be a solution to the corresponding perturbed version of (62) with index -1 and with $E(\mathfrak{c}_+) < 2\pi L$. Then $E(\mathfrak{c}_-) < 2\pi L$.

Proposition 5.2 is proved in Section 7 below. We can now define cobordism maps on filtered Seiberg–Witten Floer cohomology:

Corollary 5.3 Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) , where λ_{\pm} is L -nondegenerate. Let J_{\pm} be a symplectization-admissible almost complex structure for λ_{\pm} . Suppose r is sufficiently large. Fix 2-forms μ_{\pm}

with \mathcal{P} -norm less than 1 and fix sufficiently small abstract perturbations \mathfrak{p}_\pm as needed to define the chain complexes $\widehat{\text{CM}}^*(Y_\pm; \lambda_\pm, J_\pm, r)$. Then there is a well-defined map

$$(65) \quad \widehat{\text{HM}}_L^*(X, \lambda): \widehat{\text{HM}}_L^*(Y_+; \lambda_+, J_+, r) \longrightarrow \widehat{\text{HM}}_L^*(Y_-; \lambda_-, J_-, r),$$

depending only on $X, \lambda, L, r, J_\pm, \mu_\pm, \mathfrak{p}_\pm$, with the following properties:

- (a) If $L' < L$ and if λ_\pm is also L' -nondegenerate, then the diagram

$$\begin{array}{ccc} \widehat{\text{HM}}_{L'}^*(Y_+; \lambda_+, J_+, r) & \xrightarrow{\widehat{\text{HM}}_{L'}^*(X, \lambda)} & \widehat{\text{HM}}_{L'}^*(Y_-; \lambda_-, J_-, r) \\ \downarrow & & \downarrow \\ \widehat{\text{HM}}_L^*(Y_+; \lambda_+, J_+, r) & \xrightarrow{\widehat{\text{HM}}_L^*(X, \lambda)} & \widehat{\text{HM}}_L^*(Y_-; \lambda_-, J_-, r) \end{array}$$

commutes, where the vertical arrows are induced by inclusions of chain complexes.

- (b) Likewise the diagram

$$(66) \quad \begin{array}{ccc} \widehat{\text{HM}}_L^*(Y_+; \lambda_+, J_+, r) & \xrightarrow{\widehat{\text{HM}}_L^*(X, \lambda)} & \widehat{\text{HM}}_L^*(Y_-; \lambda_-, J_-, r) \\ \downarrow & & \downarrow \\ \widehat{\text{HM}}^*(Y_+; \lambda_+, J_+, r) & \longrightarrow & \widehat{\text{HM}}^*(Y_-; \lambda_-, J_-, r) \end{array}$$

commutes, where the bottom arrow is the Seiberg–Witten Floer cobordism map induced by (64).

- (c) If $\{(X_t, \lambda_t) \mid t \in [0, 1]\}$ is a strong homotopy of exact symplectic cobordisms from (Y_+, λ_+) to (Y_-, λ_-) , and if r is sufficiently large, then

$$\widehat{\text{HM}}_L^*(X, \lambda_0) = \widehat{\text{HM}}_L^*(X, \lambda_1).$$

Note that for now the map $\widehat{\text{HM}}_L^*(X, \lambda)$ may depend on $r, J_\pm, \mu_\pm, \mathfrak{p}_\pm$, although these choices are not indicated in the notation. Proposition 5.5 below will show that in fact $\widehat{\text{HM}}_L^*(X, \lambda)$ is independent of these choices.

Proof Choose a strongly cobordism-admissible almost complex structure J on \overline{X} extending J_+ and J_- , and choose small perturbations μ and \mathfrak{p} extending μ_\pm and \mathfrak{p}_\pm , as needed to define the chain map (64). Summing over spin-c structures, we then have a chain map

$$(67) \quad \widehat{\text{CM}}^*(Y_+; \lambda_+, J_+, r) \longrightarrow \widehat{\text{CM}}^*(Y_-; \lambda_-, J_-, r),$$

whose induced map on homology is the bottom arrow in (66). It follows from Proposition 5.2(a) that if r is sufficiently large, and if the perturbations are sufficiently small, then the chain map (67) restricts to a chain map

$$(68) \quad \widehat{\text{CM}}_L^*(Y_+; \lambda_+, J_+, r) \longrightarrow \widehat{\text{CM}}_L^*(Y_-; \lambda_-, J_-, r).$$

We define $\widehat{\text{HM}}_L^*(X, \lambda)$ to be the map on homology induced by (68).

We now show that $\widehat{\text{HM}}_L^*(X, \lambda)$ does not depend on the choice of extensions J, μ, \mathfrak{p} of $J_\pm, \mu_\pm, \mathfrak{p}_\pm$ over X . Given two choices of extensions, we can choose a homotopy between them. This homotopy induces a chain homotopy between the corresponding maps (67), which counts index -1 instantons that appear during the homotopy. It follows from Proposition 5.2(b) that if r is sufficiently large, then this chain homotopy maps $\widehat{\text{CM}}_L^*$ to $\widehat{\text{CM}}_L^*$ (here we are ignoring the gradings as usual), and hence restricts to a chain homotopy between the corresponding maps (68).

Properties (a) and (b) above now hold by construction. One proves property (c) by using Proposition 5.2(b) to define a chain homotopy. □

The induced maps on $\widehat{\text{HM}}_L^*$ constructed above behave nicely under composition of exact symplectic cobordisms:

Proposition 5.4 *Suppose (X, λ) is the composition of an exact symplectic cobordism (X^-, λ^-) from (Y_0, λ_0) to (Y_-, λ_-) with an exact symplectic cobordism (X^+, λ^+) from (Y_+, λ_+) to (Y_0, λ_0) , where λ_\pm and λ_0 are L -nondegenerate. Let J_\pm and J_0 be symplectization-admissible almost complex structures for λ_\pm and λ_0 . Let μ_\pm, μ_0 be 2-forms from the Y_\pm, Y_0 versions of Ω with \mathcal{P} -norm less than 1, fix r sufficiently large, and let $\mathfrak{p}_\pm, \mathfrak{p}_0$ be sufficiently small generic abstract perturbations as needed to define the chain complexes $\widehat{\text{CM}}^*$. Then the maps in Corollary 5.3 for these data satisfy*

$$\widehat{\text{HM}}_L^*(X, \lambda) = \widehat{\text{HM}}_L^*(X^-, \lambda^-) \circ \widehat{\text{HM}}_L^*(X^+, \lambda^+).$$

Proposition 5.4 is proved in Section 7 using a neck stretching argument.

5.2 Invariance of cobordism maps

The goal of this subsection is to prove the following proposition, asserting that the map $\widehat{\text{HM}}_L^*(X, \lambda)$ defined in Corollary 5.3 depends only on X, λ, L , and not on the additional choices made in its definition.

Proposition 5.5 *Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) where λ_{\pm} is L -nondegenerate. Then the map defined in [Corollary 5.3](#) induces a well-defined map*

$$\widehat{HM}_L^*(X, \lambda): \widehat{HM}_L^*(Y_+, \lambda_+) \longrightarrow \widehat{HM}_L^*(Y_-, \lambda_-),$$

where $\widehat{HM}_L^*(Y_{\pm}, \lambda_{\pm})$ is defined as in [Corollary 3.5](#).

To prepare for the proof of this proposition, we need the following lemma, which relates the maps on \widehat{HM}_L^* induced by exact product symplectic cobordisms to the canonical isomorphisms between different versions of \widehat{HM}_L^* .

Lemma 5.6 *Let $\rho = \{(\lambda_t, L_t, J_t, r) \mid t \in [0, 1]\}$ be an admissible deformation as in [Definition 3.3](#). Assume further that:*

- $\lambda_t = f_t \lambda_0$, where $f: [0, 1] \times Y \rightarrow \mathbb{R}^{>0}$ satisfies $\partial f / \partial t < 0$ everywhere.
- $dL_t / dt \leq 0$.

Let (X, λ) be the product exact cobordism $([-1, 0] \times Y, \lambda_{-s})$ from (Y, λ_0) to (Y, λ_1) . Suppose r is sufficiently large. Fix small perturbations μ_i, p_i for $i = 0, 1$ as needed to define the chain complexes \widehat{CM}^* for $t = 0, 1$. Then the cobordism map $\widehat{HM}_{L_0}^*(X, \lambda)$ in [Corollary 5.3](#) is the composition

$$\widehat{HM}_{L_0}^*(Y; \lambda_0, J_0, r) \xrightarrow{\Phi_\rho} \widehat{HM}_{L_1}^*(Y; \lambda_1, J_1, r) \longrightarrow \widehat{HM}_{L_0}^*(Y; \lambda_1, J_1, r),$$

where Φ_ρ is the isomorphism from [\(36\)](#), and the map on the right is induced by the inclusion of chain complexes.

Proof The proof has two steps.

Step 1 We start by making choices as in the proof of [Lemma 3.4](#) to define Φ_ρ . Choose a path of data $\{(\mu_t, p_t) \mid t \in [0, 1]\}$ from (μ_0, p_0) to (μ_1, p_1) , where μ_t is a 2-form from Ω with \mathcal{P} -norm less than 1, and p_t is a small abstract perturbation such that the data $D_t = (\lambda_t, J_t, r, \mu_t, p_t)$ is suitable for defining the chain complex $\widehat{CM}^*(Y; \lambda_t, J_t, r)$ for generic $t \in [0, 1]$. Let N be a large positive integer, and choose numbers $0 = t_0 < t_1 < \dots < t_N = 1$ such that $t_i - t_{i-1} < 2/N$ and such that $\widehat{CM}^*(Y, s; \lambda_{t_i}, J_{t_i}, r)$ is defined for each $i = 1, \dots, N$.

To shorten the notation below, write $\widehat{HM}_L^*(t)$ to denote $\widehat{HM}_L^*(Y; \lambda_t, J_t, r)$. Also, for $t < t'$ let $\widehat{HM}_L^*([t, t'])$ denote the cobordism map induced by the portion of the cobordism parametrized by $[-t', -t] \times Y$, and let $\rho_{[t, t']}$ denote the portion of the

admissible deformation parametrized by the interval $[t, t']$, reparametrized by the interval $[0, 1]$.

Choose N sufficiently large that λ_{t_i} has no orbit sets of action in the interval $[L_{t_i}, L_{t_{i-1}}]$ for each $i = 1, \dots, N$. Then for r sufficiently large, the lemma holds for the portion of the cobordism parametrized by $[-t_i, -t_{i-1}] \times Y$. That is, the cobordism map

$$(69) \quad \widehat{\text{HM}}_{L_{t_{i-1}}}^*([t_{i-1}, t_i])$$

equals the composition

$$\widehat{\text{HM}}_{L_{t_{i-1}}}^*(t_{i-1}) \xrightarrow{\Phi_{\rho_{[t_{i-1}, t_i]}}} \widehat{\text{HM}}_{L_{t_i}}^*(t_i) \longrightarrow \widehat{\text{HM}}_{L_{t_{i-1}}}^*(t_i),$$

where the map on the right is induced by the inclusion of chain complexes. The reason is that the map on the right is an isomorphism on the chain level by Lemma 2.3(b), so that the cobordism map (69) actually maps to $\widehat{\text{HM}}_{L_{t_i}}^*(t_i)$. Then the cobordism map (69), regarded as a map to $\widehat{\text{HM}}_{L_{t_i}}^*(t_i)$, agrees with Φ_{ρ_i} by the definition of the latter.

Step 2 We now show by induction on i that the lemma holds for the portion of the cobordism parametrized by $[-t_i, 0] \times Y$. The case $i = 1$ follows from Step 1. Now let $i > 1$ and suppose the claim is true for $i - 1$. We need to show that the cobordism map

$$\widehat{\text{HM}}_{L_0}^*([0, t_i])$$

agrees with the composition

$$\widehat{\text{HM}}_{L_0}^*(0) \xrightarrow{\Phi_{\rho_{[0, t_i]}}} \widehat{\text{HM}}_{L_{t_i}}^*(t_i) \longrightarrow \widehat{\text{HM}}_{L_0}^*(t_i),$$

where the arrow on the right is induced by inclusion.

By Proposition 5.4 we have

$$\widehat{\text{HM}}_{L_0}^*([0, t_i]) = \widehat{\text{HM}}_{L_0}^*([t_{i-1}, t_i]) \circ \widehat{\text{HM}}_{L_0}^*([0, t_{i-1}]).$$

And by Lemma 3.4(b) we have

$$\Phi_{\rho_{[0, t_i]}} = \Phi_{\rho_{[t_{i-1}, t_i]}} \circ \Phi_{\rho_{[0, t_{i-1}]}}.$$

So by the inductive hypothesis and Step 1, we just need to show that the diagram

$$\begin{array}{ccc} \widehat{\text{HM}}_{L_0}^*(t_{i-1}) & \xrightarrow{\widehat{\text{HM}}_{L_0}^*([t_{i-1}, t_i])} & \widehat{\text{HM}}_{L_0}^*(t_i) \\ \uparrow & & \uparrow \\ \widehat{\text{HM}}_{L_{t_{i-1}}}^*(t_{i-1}) & \xrightarrow{\widehat{\text{HM}}_{L_{t_{i-1}}}^*([t_{i-1}, t_i])} & \widehat{\text{HM}}_{L_{t_{i-1}}}^*(t_i) \end{array}$$

commutes, where the vertical arrows are induced by inclusion. But this holds by Corollary 5.3(a). \square

We can now prove Proposition 5.5. The latter is an immediate consequence of the following lemma:

Lemma 5.7 *Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) , where λ_{\pm} is L -nondegenerate. Let $\{J_{\pm}^t \mid t \in [0, 1]\}$ be a one-parameter family of symplectization-admissible almost complex structures for λ_{\pm} . Suppose that $\{r_t \mid t \in [0, 1]\}$ is a sufficiently large one-parameter family of real numbers. Let*

$$(\mu_{\pm}^0, \mathfrak{p}_{\pm}^0) \quad \text{and} \quad (\mu_{\pm}^1, \mathfrak{p}_{\pm}^1)$$

be small perturbations as needed to define the chain complexes $\widehat{\text{CM}}^*$ for $t = 0, 1$. Then the versions of $\widehat{\text{HM}}_L^*(X, \lambda)$ for $t = 0, 1$ fit into a commutative diagram

$$(70) \quad \begin{array}{ccc} \widehat{\text{HM}}_L^*(Y_+, \lambda_+, J_+^0, r_0) & \xrightarrow[\simeq]{\Phi_{\rho_+}} & \widehat{\text{HM}}_L^*(Y_+, \lambda_+, J_+^1, r_1) \\ \widehat{\text{HM}}_L^*(X, \lambda)_{t=0} \downarrow & & \downarrow \widehat{\text{HM}}_L^*(X, \lambda)_{t=1} \\ \widehat{\text{HM}}_L^*(Y_-, \lambda_-, J_-^0, r_0) & \xrightarrow[\simeq]{\Phi_{\rho_-}} & \widehat{\text{HM}}_L^*(Y_-, \lambda_-, J_-^1, r_1), \end{array}$$

where $\Phi_{\rho_{\pm}}$ is the isomorphism from Lemma 3.4, and

$$\rho_{\pm} = \{(\lambda_{\pm}, L, J_{\pm}^t, r_t) \mid t \in [0, 1]\}.$$

Proof First note that because $d\lambda_{\pm}$ is in the space Ω for Y_{\pm} (see Section 2.2), a small change in r can be effected by a change in μ_+ and μ_- . Thus, by the homotopy properties in Lemma 3.4(a), (b), it is enough to prove the lemma in the case when r_t does not depend on t ; let us write $r_t = r$.

To prove the lemma for constant r , it is enough to do so in the special case when $(J_-^t, \mu_-^t, \mathfrak{p}_-^t)$ do not depend on t ; let us denote these by $(J_-, \mu_-, \mathfrak{p}_-)$. (The case when $(J_-^t, \mu_-^t, \mathfrak{p}_-^t)$ do depend on t , but $(J_+^t, \mu_+^t, \mathfrak{p}_+^t)$ do not, is proved by a symmetrical argument, and these two cases together imply the general case.)

Let $\varepsilon > 0$ be as in (12) and (13), so that a neighborhood of Y_+ in (X, λ) is identified with $([-\varepsilon, 0] \times Y_+, e^s \lambda_+)$. Choose ε sufficiently small so that λ_+ has no orbit sets with action in the closed interval $[e^{-\varepsilon} L, L]$. We can decompose the exact cobordism $X = X^0 \circ X^+$, where $X^+ = [-\varepsilon, 0] \times Y_+$, and X^0 is the closure of $X \setminus X^+$. Make

choices $(J_0, \mu_0, \mathfrak{p}_0)$ as needed to define the chain complex $\widehat{\text{CM}}^*(Y_+, e^{-\varepsilon}\lambda_+, J_0, r)$. We now have a commutative diagram:

$$\begin{array}{ccc}
 \widehat{\text{HM}}_L^*(Y_+, \lambda_+, J_+^0, r) & \xrightarrow{\Phi_{\rho_+}} & \widehat{\text{HM}}_L^*(Y_+, \lambda_+, J_+^1, r) \\
 \downarrow \Phi_{\rho_0} & & \downarrow \Phi_{\rho_1} \\
 \widehat{\text{HM}}_{e^{-\varepsilon}L}^*(Y_+, e^{-\varepsilon}\lambda_+, J_0, r) & \xlongequal{\quad} & \widehat{\text{HM}}_{e^{-\varepsilon}L}^*(Y_+, e^{-\varepsilon}\lambda_+, J_0, r) \\
 \downarrow & & \downarrow \\
 \widehat{\text{HM}}_L^*(Y_+, e^{-\varepsilon}\lambda_+, J_0, r) & \xlongequal{\quad} & \widehat{\text{HM}}_L^*(Y_+, e^{-\varepsilon}\lambda_+, J_0, r) \\
 \downarrow \widehat{\text{HM}}_L^*(X^0, \lambda) & & \downarrow \widehat{\text{HM}}_L^*(X^0, \lambda) \\
 \widehat{\text{HM}}_L^*(Y_-, \lambda_-, J_-, r) & \xlongequal{\quad} & \widehat{\text{HM}}_L^*(Y_-, \lambda_-, J_-, r)
 \end{array}$$

Here

$$\rho_0 = \{(e^{-\varepsilon t}\lambda_+, e^{-\varepsilon t}L, \widehat{J}_t, r) \mid t \in [0, 1]\},$$

where $\{\widehat{J}_t \mid t \in [0, 1]\}$ is a path of almost complex structures from J_+^0 to J_0 . The admissible deformation ρ_1 is defined analogously. The top square in the diagram commutes by Lemma 3.4(a), (b). The vertical arrows in the middle of the diagram are induced by the inclusion of chain complexes. By Lemma 5.6, the composition of the two vertical arrows on the upper left is the cobordism map $\widehat{\text{HM}}_L^*(X^+, e^s\lambda_+)$ defined by Corollary 5.3 from the choices $(J_+^0, \mu_+, \mathfrak{p}_+, r)$ and $(J_0, \mu_0, \mathfrak{p}_0, r)$. Then by Proposition 5.4, the composition of the three vertical arrows on the left is $\widehat{\text{HM}}_L^*(X, \lambda)_{t=0}$. Likewise, the composition of the three vertical arrows on the right is $\widehat{\text{HM}}_L^*(X, \lambda)_{t=1}$. Thus the above diagram gives the desired commutative diagram (70). \square

This completes the proof of Proposition 5.5.

We also note the following special case of Lemma 5.7, which is needed in Section 5.3:

Corollary 5.8 *Suppose X is a product cobordism $([-\varepsilon, 0] \times Y, e^s\lambda)$ where $\varepsilon > 0$, the variable s denotes the $[-\varepsilon, 0]$ coordinate, and λ is an L -nondegenerate contact form on Y . Then $\widehat{\text{HM}}_L^*(X, e^s\lambda)$ is the composition*

$$\widehat{\text{HM}}_L^*(Y, \lambda) \xrightarrow{s} \widehat{\text{HM}}_{e^{-\varepsilon}L}^*(Y, e^{-\varepsilon}\lambda) \longrightarrow \widehat{\text{HM}}_L^*(Y, e^{-\varepsilon}\lambda),$$

where s is the scaling isomorphism from Corollary 3.5(c), and the right arrow is the inclusion-induced map from Corollary 3.5(b).

Proof Choose a symplectization-admissible almost complex structure J for λ , and let r be large. The claim then follows by applying [Lemma 5.6](#) to the admissible deformation $\rho = \{(e^{-\varepsilon t} \lambda, e^{-\varepsilon t} L, J, r) \mid t \in [0, 1]\}$, because Φ_ρ agrees with the scaling isomorphism s by the definition of the latter in the proof of [Corollary 3.5\(c\)](#). \square

5.3 Construction of ECH cobordism maps

We now begin the proof of [Theorem 1.9](#) by defining the map (15) on filtered ECH induced by an exact symplectic cobordism.

Let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) , and assume that λ_+ and λ_- are nondegenerate. Fix a real number L . Without loss of generality, we can assume (by slightly decreasing L if necessary) that λ_+ and λ_- do not have any orbit sets of action exactly L . By [Proposition 5.5](#), we have a well-defined map

$$(71) \quad \widehat{\text{HM}}_L^*(X, \lambda): \widehat{\text{HM}}_L^*(Y_+, \lambda_+) \longrightarrow \widehat{\text{HM}}_L^*(Y_-, \lambda_-).$$

On the other hand, by [Lemma 3.7](#) we have canonical isomorphisms

$$(72) \quad \text{ECH}_*^L(Y_\pm, \lambda_\pm) \xrightarrow{\cong} \widehat{\text{HM}}_L^{-*}(Y_\pm, \lambda_\pm).$$

Definition 5.9 Define a map of ungraded $\mathbb{Z}/2$ -modules

$$(73) \quad \Phi^L(X, \lambda): \text{ECH}_*^L(Y_+, \lambda_+) \longrightarrow \text{ECH}_*^L(Y_-, \lambda_-)$$

to be the composition of the map (71) with the isomorphisms (72).

We now prove all of [Theorem 1.9](#) except for the holomorphic curves axiom:

Proposition 5.10 *The map (73) satisfies the homotopy invariance, inclusion, direct limit, composition and scaling axioms in [Theorem 1.9](#).*

Proof The inclusion axiom follows from [Lemma 3.7\(a\)](#) and [Corollary 5.3\(a\)](#). The direct limit axiom follows from [Corollary 5.3\(b\)](#). The composition axiom follows from [Proposition 5.4](#).

To prove the homotopy invariance axiom, let $\{(X, \omega_t) \mid t \in [0, 1]\}$ be a homotopy of exact symplectic cobordisms, let λ_0 be a Liouville form for ω_0 , and let λ_1 be a Liouville form for ω_1 . We need to show that $\Phi^L(X, \lambda_0) = \Phi^L(X, \lambda_1)$. Since the space of Liouville forms for a given exact symplectic cobordism (X, ω) is affine linear, there is no obstruction to connecting λ_0 and λ_1 by a smooth one-parameter family $\{\lambda_t \mid t \in [0, 1]\}$ of 1-forms on X such that λ_t is a Liouville form for ω_t . Next, fix $\varepsilon > 0$ such that there are disjoint neighborhoods (12) and (13) for each λ_t . We can

then find a smooth one-parameter family $\{\varphi_t \mid t \in [0, 1]\}$ of diffeomorphisms of X with $\varphi_0 = \text{id}_X$ and $\varphi_t|_{\partial X} = \text{id}_{\partial X}$ such that φ_t pulls back the neighborhoods (12) and (13) for λ_t to those for λ_0 . Then $\{(X, \varphi_t^* \lambda_t) \mid t \in [0, 1]\}$ is a strong homotopy from (X, λ_0) to $(X, \varphi_1^* \lambda_1)$ as in Definition 5.1. By Corollary 5.3(c), $\Phi^L(X, \lambda_0) = \Phi^L(X, \varphi_1^* \lambda_1)$. Now the diffeomorphism φ_1 extends to a symplectomorphism between the completions (14) of (X, λ_1) and $(X, \varphi_1^* \lambda_1)$, and so by construction $\Phi^L(X, \varphi_1^* \lambda_1) = \Phi^L(X, \lambda_1)$.

To prove the scaling axiom, let (X, λ) be an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) , where the contact forms λ_{\pm} are nondegenerate and have no orbit sets of action L . Write $c = e^{-\varepsilon}$, and assume without loss of generality that $\varepsilon > 0$. We need to show that the upper square in the diagram

$$\begin{array}{ccc}
 \text{ECH}_*^L(Y_+, \lambda_+) & \xrightarrow{\Phi^L(X, \lambda)} & \text{ECH}_*^L(Y_-, \lambda_-) \\
 \downarrow s & & \downarrow s \\
 \text{ECH}_*^{e^{-\varepsilon}L}(Y_+, e^{-\varepsilon}\lambda_+) & \xrightarrow{\Phi^{e^{-\varepsilon}L}(X, e^{-\varepsilon}\lambda)} & \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, e^{-\varepsilon}\lambda_-) \\
 \downarrow i^{e^{-\varepsilon}L, L} & & \downarrow i^{e^{-\varepsilon}L, L} \\
 \text{ECH}_*^L(Y_+, e^{-\varepsilon}\lambda_+) & \xrightarrow{\Phi^L(X, e^{-\varepsilon}\lambda)} & \text{ECH}_*^L(Y_-, e^{-\varepsilon}\lambda_-)
 \end{array}$$

commutes, where s denotes the scaling isomorphism (7) for $c = e^{-\varepsilon}$.

Since the composition of two scaling isomorphisms is a scaling isomorphism, we may assume without loss of generality that ε is sufficiently small so that λ_{\pm} has no orbit set of action in the interval $[L, e^\varepsilon L]$. Then the lower vertical arrows in the above diagram are isomorphisms by Lemma 2.3(b). Also, by the inclusion axiom, the lower square commutes. So to prove that the upper square commutes, it is enough to show that the outer rectangle commutes.

For this purpose consider the product exact cobordisms $(X^+ = [-\varepsilon, 0] \times Y_+, e^s \lambda_+)$ and $(X^- = [-\varepsilon, 0] \times Y_-, e^s \lambda_-)$. By Corollary 5.8 and Lemma 3.7(a), (b), the compositions of the vertical arrows in the above diagram are given by

$$(74) \quad \Phi^L(X^\pm, e^s \lambda_\pm) = i^{e^{-\varepsilon}L, L} \circ s: \text{ECH}_*^L(Y_\pm, \lambda_\pm) \longrightarrow \text{ECH}_*^L(Y_\pm, e^{-\varepsilon}\lambda_\pm).$$

So to prove that the outer rectangle in the above diagram commutes, it is enough to prove that the square

$$\begin{array}{ccc}
 \text{ECH}_*^L(Y_+, \lambda_+) & \xrightarrow{\Phi^L(X, \lambda)} & \text{ECH}_*^L(Y_-, \lambda_-) \\
 \downarrow \Phi^L(X^+, e^s \lambda_+) & & \downarrow \Phi^L(X^-, e^s \lambda_-) \\
 \text{ECH}_*^L(Y_+, e^{-\varepsilon}\lambda_+) & \xrightarrow{\Phi^L(X, e^{-\varepsilon}\lambda)} & \text{ECH}_*^L(Y_-, e^{-\varepsilon}\lambda_-)
 \end{array}$$

commutes. By the composition axiom, this is equivalent to the assertion that

$$(75) \quad \Phi^L((X, e^{-\varepsilon}\lambda) \circ (X^+, e^s\lambda_+)) = \Phi^L((X^-, e^s\lambda_-) \circ (X, \lambda)).$$

But these two compositions of exact symplectic cobordisms are homotopic through exact symplectic cobordisms from (Y_+, λ_+) to $(Y_-, e^{-\varepsilon}\lambda_-)$ if ε is sufficiently small as in (12) and (13). Thus equation (75) follows from the homotopy axiom. \square

6 Proof of the holomorphic curves axiom

Let (X, λ) be an exact symplectic cobordism as in the statement of Theorem 1.9. To complete the proof of Theorem 1.9, we now prove that the maps $\Phi^L(X, \lambda)$ defined in Section 5.3 satisfy the holomorphic curves axiom. For this purpose fix a cobordism-admissible almost complex structure J on \bar{X} as in the statement of the holomorphic curves axiom. Let J_{\pm} denote the symplectization-admissible almost complex structure that J determines on $\mathbb{R} \times Y_{\pm}$, and assume that this is ECH^L-admissible.

In the analysis in this section, we adopt the convention that c denotes a positive constant whose value may increase from one appearance to the next.

6.1 Reduction to the strongly cobordism-admissible case

The first step in the proof of the holomorphic curves axiom is to reduce to the case where J is strongly cobordism-admissible; see Definition 4.1. The latter condition ensures that the chain map (68) is defined, and will also be convenient in Section 6.3.

Lemma 6.1 *To prove the holomorphic curves axiom, it suffices to prove it in the special case when J is strongly cobordism-admissible.*

Proof Assume that the holomorphic curves axiom is true in the strongly cobordism-admissible case, and let J be any (not necessarily strongly) cobordism-admissible almost complex structure. Fix L such that λ_+ and λ_- have no ECH generators of action exactly L . Choose $\varepsilon > 0$ sufficiently small that λ_+ has no ECH generators with action in the interval $[L, e^{\varepsilon}L]$, and λ_- has no ECH generators with action in the interval $[e^{-\varepsilon}L, L]$. Define an exact cobordism

$$X' = ([-\varepsilon, 0] \times Y_-, e^s\lambda_-) \circ X \circ ([0, \varepsilon] \times Y_+, e^s\lambda_+)$$

from $(Y_+, e^{\varepsilon}\lambda_+)$ to $(Y_-, e^{-\varepsilon}\lambda_-)$.

We use the cobordism-admissible almost complex structure J on \bar{X} to define a strongly cobordism-admissible almost complex structure J' on \bar{X}' as follows. Note that there

is a natural identification $\overline{X'} = \overline{X}$, sending $(-\infty, 0] \times Y_-$ and $[0, \infty) \times Y_+$ in $\overline{X'}$ to $(-\infty, -\varepsilon] \times Y_-$ and $[\varepsilon, \infty) \times Y_+$ in \overline{X} . Under this identification, the almost complex structure J on \overline{X} is not quite cobordism-admissible for $\overline{X'}$, because on the ends $J(\partial_s)$ is not the Reeb vector field, but rather the Reeb vector field times $e^{\pm\varepsilon}$. To repair this defect, choose a diffeomorphism $\varphi_+ : [0, \infty) \rightarrow [0, \infty)$ such that $\varphi_+(s) = s$ for s close to 0, and $d\varphi_+(s)/ds = e^{-\varepsilon}$ for $s \geq \varepsilon/2$. Likewise choose a diffeomorphism $\varphi_- : (-\infty, 0] \rightarrow (-\infty, 0]$ such that $\varphi_-(s) = s$ for s close to 0, and $d\varphi_-(s)/ds = e^\varepsilon$ for $s \leq -\varepsilon/2$. Define a diffeomorphism $\phi : \overline{X'} \rightarrow \overline{X}$ by setting $\phi|_X = \text{id}_X$ and defining $\phi(s, y) = (\varphi_+(s), y)$ on $[0, \infty) \times Y_+$ and $\phi(s, y) = (\varphi_-(s), y)$ on $(-\infty, 0] \times Y_-$. Now $J' := \phi^*J$ is a strongly cobordism-admissible almost complex structure on $\overline{X'}$. Furthermore, product regions for X' with respect to J' correspond to product regions for X with respect to J .

By hypothesis, the holomorphic curves axiom holds for (X', J') , so there is a chain map

$$\widehat{\Phi}' : \text{ECC}_*^L(Y_+, e^\varepsilon \lambda_+, J_+) \longrightarrow \text{ECC}_*^L(Y_-, e^{-\varepsilon} \lambda_-, J_-)$$

which induces $\Phi^L(X')$ and fulfills conditions (i) and (ii) in the holomorphic curves axiom. To deduce the holomorphic curves axiom for (X, J) , define a chain map $\widehat{\Phi}$ by composing $\widehat{\Phi}'$ with the composition

$$(76) \quad \text{ECC}_*^L(Y_+, \lambda_+, J_+) \longrightarrow \text{ECC}_*^{e^\varepsilon L}(Y_+, e^\varepsilon \lambda_+, J_+) \longrightarrow \text{ECC}_*^L(Y_+, e^\varepsilon \lambda_+, J_+)$$

on the left, and the composition

$$(77) \quad \text{ECC}_*^L(Y_-, e^{-\varepsilon} \lambda_-, J_-) \longrightarrow \text{ECC}_*^{e^\varepsilon L}(Y_-, \lambda_-, J_-) \longrightarrow \text{ECC}_*^L(Y_-, \lambda_-, J_-)$$

on the right. In each of (76) and (77), the left arrow is the scaling isomorphism, and the right arrow is the inverse of the map induced by inclusion of chain complexes (which is an isomorphism since we chose ε sufficiently small). Each of the compositions (76) and (77) is the obvious geometric identification of generators, and so since $\widehat{\Phi}'$ satisfies conditions (i) and (ii) in the holomorphic curves axiom, it follows that $\widehat{\Phi}$ satisfies these conditions as well. Finally, it follows from (74) and the composition axiom that $\widehat{\Phi}$ induces the map $\Phi^L(X)$, as required. □

Assume henceforth that J is strongly cobordism-admissible.

6.2 The L -flat case

We now prove the holomorphic curves axiom in the special case when (λ_+, J_+) and (λ_-, J_-) are L -flat. In this case, define a chain map

$$(78) \quad \widehat{\Phi} : \text{ECC}_*^L(Y_+, \lambda_+; J_+) \longrightarrow \text{ECC}_*^L(Y_-, \lambda_-; J_-)$$

by composing the chain map (68) for r large (and some choice of small 2-form μ and small abstract perturbation \mathfrak{p}) with the isomorphisms of chain complexes (33) on both sides.

Proposition 6.2 *If (λ_+, J_+) and (λ_-, J_-) are L -flat, then there exists an abstract perturbation \mathfrak{p} such that the chain map $\widehat{\Phi}$ in (78) fulfills the holomorphic curves axiom.*

Proof We need to show that conditions (i) and (ii) in the holomorphic curves axiom hold for this $\widehat{\Phi}$. It follows immediately from Proposition 5.2(a) that condition (i) holds.

To prove (ii), suppose that Θ^+ is a generator of $\text{ECC}_*^L(Y_+, \lambda_+; J_+)$ in a product region, and let Θ^- denote the corresponding generator of $\text{ECC}_*^L(Y_-, \lambda_-; J_-)$. Let $C_\Theta \in \mathcal{M}^J(\Theta^+, \Theta^-)$ denote the corresponding union of product cylinders. We need the following proposition:

Proposition 6.3 *Suppose that (λ_+, J_+) and (λ_-, J_-) are L -flat, and let Θ^\pm, C_Θ be as above. If r is sufficiently large then:*

- (a) *There exists a solution $\mathfrak{d}_\Theta = (A_\Theta, \psi_\Theta)$ to the equations (62) (without abstract perturbation) such that on the ends, $\lim_{s \rightarrow \pm\infty} \mathfrak{d}_\Theta$ corresponds to Θ^\pm via (33).*
- (b) *The operator $D_{\mathfrak{d}_\Theta}$ obtained from linearizing the equations (62) at \mathfrak{d}_Θ (this is the $\mathfrak{p} = 0$ case of the operator in [21, Equation (3.9)]) has index 0 and trivial cokernel.*
- (c) *If C_Θ is the only broken J -holomorphic curve from Θ^+ to Θ^- , then the instanton \mathfrak{d}_Θ in (a) is unique up to gauge equivalence.*

Proof (a) If Θ^\pm is the empty set, then this is proved similarly to [20, Proposition 4.3], giving a solution $\mathfrak{d}_\emptyset = (A_\emptyset, \psi_\emptyset)$ such that $\lim_{s \rightarrow \pm\infty} \mathfrak{d}_\emptyset$ corresponds to the empty set via (33), and $|\psi_0| \geq 1 - \kappa r^{-1}$ and $|F_{A_0}| \leq \kappa$ everywhere for some r -independent constant κ .

In the general case, this is proved by repeating the construction in [22, Section 4–7] with cosmetic changes. We now briefly summarize this construction.

One starts as in [22, Section 5a] by building a complex line bundle E over \overline{X} and a pair (A^*, ψ^*) consisting of a connection on E and a section of $\mathbb{S}_+ = E \oplus K^{-1}E$ (see (61)) that are close to solving (62). The bundle E is such that there is a section of E whose zero set with multiplicity is given by C_Θ . On the complement of a small radius neighborhood of C_Θ , the bundle E is identified with the trivial line bundle, and (A^*, ψ^*) is close to the instanton $(A_\emptyset, \psi_\emptyset)$ constructed above. Near a product cylinder $\mathbb{R} \times \gamma$, where (γ, m) is an element of the orbit set Θ^\pm , the pair (A^*, ψ^*) is

determined by a map $\mathfrak{v}: \mathbb{R} \times \gamma \rightarrow \mathfrak{C}_m$. Here \mathfrak{C}_m denotes the moduli space of degree m vortices on \mathbb{C} . The space \mathfrak{C}_m is naturally diffeomorphic to \mathbb{C}^m with coordinates $(\sigma_1, \dots, \sigma_m)$; see [22, Section 2] for details. The map \mathfrak{v} is required to be asymptotic to $0 \in \mathbb{C}^m$ when the \mathbb{R} coordinate s goes to $\pm\infty$. It is also required to satisfy a certain nonlinear Cauchy–Riemann equation. For each collection of maps \mathfrak{v} satisfying these conditions, a gluing construction in [22, Section 5] then perturbs the corresponding pair (A^*, ψ^*) to an instanton.⁷

When $m = 1$, it turns out that the unique solution \mathfrak{v} for (γ, m) is given by the constant map $\mathbb{R} \times \gamma \rightarrow 0 \in \mathbb{C}$. If $m > 1$, and thus γ is elliptic, let T denote the symplectic action of γ . Then the L -flatness and “product region” assumptions imply that a neighborhood of $\mathbb{R} \times \gamma$ can be identified with $\mathbb{R} \times (\mathbb{R}/T\mathbb{Z}) \times \mathbb{C}$ with coordinates s, t, z , so that the Reeb vector field is given by $R = \partial_t$, and the Liouville form λ is given by

$$\lambda = e^s \left(\left(1 - \frac{2\pi R}{T} |z|^2 \right) dt + \frac{i}{2} (z d\bar{z} - \bar{z} dz) \right).$$

Here R is a constant, the “rotation number”, which is irrational by the nondegeneracy assumptions. Meanwhile, $J\partial_s = f(s)\partial_t$ and $J\partial_z = i\partial_z$, where the function $f(s)$ is positive and equal to constants when $s \gg 0$ and $s \ll 0$. (This is only slightly different from the symplectization context of [22], where one would have $f \equiv 1$.) Now the key point is that in this case, similarly to [22], the nonlinear Cauchy–Riemann equation that $\mathfrak{v} = (\sigma_1, \dots, \sigma_m)$ has to satisfy is linear, namely the equation

$$(f^{-1}\partial_s + i\partial_t)\sigma_q + \frac{2\pi Rq}{T}\sigma_q = 0.$$

Thus this equation has a (unique) solution \mathfrak{v} with the required asymptotics

$$\lim_{s \rightarrow \pm\infty} \mathfrak{v}(s, \cdot) = 0,$$

namely $\mathfrak{v} \equiv 0$. In conclusion, we obtain a (unique) solution \mathfrak{v} for each product cylinder $(\mathbb{R} \times \gamma, m)$ in C_Θ , and this gives rise to the desired instanton.

- (b) This follows similarly to the proof of nondegeneracy in [23, Theorem 1.2].
- (c) Suppose that C_Θ is the only broken J -holomorphic curve from Θ^+ to Θ^- . We need to show that if r is sufficiently large then ∂_Θ is the unique (up to gauge equivalence) solution to (62) such that $\lim_{s \rightarrow \pm\infty} \partial_\Theta$ corresponds to Θ^\pm via (33).

⁷The argument in [22] is complicated in order to handle non- \mathbb{R} -invariant holomorphic curves having multiple ends converging to (covers of) the same Reeb orbit, or ends converging to multiple covers of a Reeb orbit. For the union of product cylinders C_Θ , one can avoid these complications and instead use (with appropriate cosmetic changes) the simpler construction in Taubes [17], which constructs a Seiberg–Witten solution from a holomorphic torus with self-intersection number zero in a closed symplectic manifold.

Suppose $\vartheta' = (A', \psi')$ is another such instanton. First observe that for any $\delta > 0$, if r is sufficiently large then we must have $1 - |\psi'| < \delta$ on the complement of the radius δ neighborhood of C_Θ . Otherwise Proposition 7.1 would imply that there is a “generalized broken J -holomorphic curve” (see Section 7.1) from Θ^+ to Θ^- containing a point not on C_Θ , contradicting our hypothesis.

Using the above observation, the arguments in [24, Section 6] carry over⁸ to show that ϑ' is gauge equivalent to ϑ_Θ . The idea is that ϑ' must be obtained from the gluing construction in (a), and uniqueness for the instanton then follows because the solutions \mathfrak{v} to their respective nonlinear Cauchy–Riemann equations are unique and cut out transversely. □

To complete the proof of Proposition 6.2, we need to show that if C_Θ is the only broken J -holomorphic curve from Θ^+ to Θ^- , then

$$(79) \quad \langle \hat{\Phi}\Theta^+, \Theta^- \rangle = 1 \in \mathbb{Z}/2.$$

Let c_\pm denote the Seiberg–Witten Floer generator corresponding to Θ^\pm via (33). Recall that to define the chain map (78), one fixes small abstract perturbations p_\pm as necessary to define the Seiberg–Witten chain complexes for Y_\pm , and extends these perturbations over \bar{X} as in [13, Chapter 24] to obtain a small perturbation p as needed to perturb the equations (62). Recall from Section 2.1 that we choose the perturbations p_\pm so that c_\pm are still solutions to the perturbed version of the Seiberg–Witten equations (20). Likewise the perturbation p can be chosen to vanish to second order on the instantons ϑ_Θ given by Proposition 6.3, so that these are transverse solutions to the perturbed version of the instanton equations (62). A limiting argument similar to Step 2 of the proof of Proposition 5.2 in Section 7.6 now shows that if r is sufficiently large and if p_\pm and p are sufficiently small, then any other solution ϑ' to the corresponding perturbed version of (62) with $\lim_{s \rightarrow \pm} \vartheta' = c_\pm$ is gauge equivalent to ϑ_Θ . It follows that (79) holds as desired. □

6.3 The non- L -flat case

To prove the holomorphic curves axiom in the non- L -flat case, we reduce to the L -flat case by defining a sequence of modified exact symplectic cobordisms $\{(X, \lambda_n)\}_{n=1,2,\dots}$ between L -flat pairs, equipped with strongly cobordism-admissible almost complex structures J_n , such that (λ_n, J_n) converges to (λ, J) in an appropriate sense. Fix $\varepsilon > 0$ as in Definition 4.1. We can then write $X = X^- \circ X^0 \circ X^+$, where $X^- = [0, \varepsilon] \times Y_-$ and

⁸As in the proof of part (a), the argument needed here can be simplified from that in [24] and differs only cosmetically from analogous arguments in [17].

$X^\pm = [-\varepsilon, 0] \times Y_\pm$ and $\lambda|_{X^\pm} = e^s \lambda_\pm$. Here s denotes the $[-\varepsilon, 0]$ or $[0, \varepsilon]$ coordinate as usual. The idea of the construction is to define (λ_n, J_n) by suitably modifying (λ, J) on X^\pm , and in neighborhoods of product regions. The construction has four steps.

Step 1 To begin the construction, fix a positive integer n . Let U_\pm denote the set of points in Y_\pm that are within distance $1/n$ of a Reeb orbit with action less than L , as measured using some arbitrary metrics on Y_\pm . By Lemma 3.6(c), there exists a preferred homotopy $\{(\lambda_\pm^t, J_\pm^t) \mid t \in [0, 1]\}$ on Y_\pm where $(\lambda_\pm^0, J_\pm^0) = (\lambda_\pm, J_\pm)$, the pair (λ_\pm^1, J_\pm^1) is L -flat, and (λ_\pm^t, J_\pm^t) agrees with (λ_\pm, J_\pm) on $Y_\pm \setminus U_\pm$. To ensure smooth gluing below, let us reparametrize the homotopy so that $(\lambda_\pm^t, J_\pm^t) = (\lambda_\pm, J_\pm)$ for t in a neighborhood of 0, and $(\lambda_\pm^t, J_\pm^t) = (\lambda_\pm^1, J_\pm^1)$ for t in a neighborhood of 1. Also, we can assume that if a component of U_\pm is contained in Z where $[s_-, s_+] \times Z$ is a product region, then

$$e^{-s_+} \lambda_+^t = e^{-s_-} \lambda_-^t, \quad J_+^t|_{\text{Ker}(\lambda_0)} = J_-^t|_{\text{Ker}(\lambda_0)}$$

on this component.

Keep in mind that (λ_\pm^t, J_\pm^t) depends on n , although we do not indicate this in the notation. We will need the following estimates on this n -dependence:

Lemma 6.4 *There exists an n -independent constant $c > 0$ such that the homotopy $\{(\lambda_\pm^t, J_\pm^t)\}$ above can be chosen so that*

$$(80) \quad \left\| \frac{\partial \lambda_\pm^t}{\partial t} \right\|_{C^1}, \left\| \frac{\partial J_\pm^t}{\partial t} \right\|_{C^0} \leq cn^{-1},$$

$$(81) \quad \left\| \frac{\partial \lambda_\pm^t}{\partial t} \right\|_{C^2}, \left\| \frac{\partial J_\pm^t}{\partial t} \right\|_{C^1} \leq c.$$

Proof Let $\gamma: \mathbb{R}/T\mathbb{Z} \rightarrow Y_\pm$ be a Reeb orbit of action less than L . Recall from [21, Equation (2-1)] that there exists a disk $D \subset \mathbb{C}$ about the origin and an extension of γ to an embedding $\varphi: (\mathbb{R}/T\mathbb{Z}) \times D \rightarrow Y_\pm$ such that:

- If t denotes the $\mathbb{R}/T\mathbb{Z}$ coordinate and z denotes the \mathbb{C} coordinate, then

$$(82) \quad \varphi^* \lambda_\pm^0 = (1 - 2\nu|z|^2 - \mu\bar{z}^2 - \bar{\mu}z^2) dt + \frac{i}{2}(z d\bar{z} - \bar{z} dz) + \dots,$$

where ν and μ are respectively real and complex valued functions on $\mathbb{R}/T\mathbb{Z}$, and the unwritten terms are $O(|z|^3)$.

- At $z = 0$, the restriction of J_\pm to ξ is the standard almost complex structure on \mathbb{C} .

By [21, Equation (2-11)], λ_{\pm}^1 differs from λ_{\pm}^0 only in the μ terms and higher order terms in (82), and these differences occur only where $|z| \leq c/n$. It follows that $\lambda_{\pm}^1 - \lambda_{\pm}^0$ satisfies the C^1 and C^2 bounds in (80) and (81), and because of the way a preferred homotopy is constructed in [21, Appendix A], $\partial\lambda_{\pm}^t/\partial t$ also satisfies these bounds.

It also follows from [21, Equation (2-11)] and the second bullet point above that J_{\pm}^1 and J_{\pm}^0 agree along γ , and therefore their difference is $O(|z|)$. Since their difference is supported where $|z| \leq c/n$, it follows from this and the cutoff construction of J_{\pm}^1 in [21, Appendix A] that $J_{\pm}^1 - J_{\pm}^0$ satisfies the C^0 and C^1 bounds in (80) and (81). It then follows from the construction of a preferred homotopy that $\partial J_{\pm}^t/\partial t$ also satisfies these bounds. □

As a first step to defining λ_n , define a 1-form λ'_n on X by

$$\lambda'_n := \begin{cases} e^s \lambda_+^{1+\varepsilon^{-1}s} & \text{on } X^+ = [-\varepsilon, 0] \times Y_+, \\ \lambda & \text{on } X^0, \\ e^s \lambda_-^{1-\varepsilon^{-1}s} & \text{on } X^- = [0, \varepsilon] \times Y_-. \end{cases}$$

It follows from (80) that if n is sufficiently large (which we assume that it is), then (X, λ'_n) is an exact symplectic cobordism from (Y_+, λ_+^1) to (Y_-, λ_-^1) .

Step 2 We now relate the maps on ECH induced by (X, λ) to those induced by (X, λ'_n) .

Lemma 6.5 *The following diagram commutes:*

$$(83) \quad \begin{array}{ccc} \text{ECH}_*^L(Y_+, \lambda_+) & \xrightarrow{\cong} & \text{ECH}_*^L(Y_+, \lambda_+^1) \\ \downarrow \Phi^L(X, \lambda) & & \downarrow \Phi^L(X, \lambda'_n) \\ \text{ECH}_*^L(Y_-, \lambda_-) & \xrightarrow{\cong} & \text{ECH}_*^L(Y_-, \lambda_-^1) \end{array}$$

Here the horizontal arrows are induced by the canonical isomorphism of chain complexes (35).

Proof Let Ψ_+ and Ψ_- denote the top and bottom arrows in (83). By the composition axiom we have

$$\begin{aligned} \Phi^L(X, \lambda'_n) &= \Phi^L(X^-, \lambda'_n) \circ \Phi^L(X^0, \lambda'_n) \circ \Phi^L(X^+, \lambda'_n), \\ \Phi^L(X, \lambda) &= \Phi^L(X^-, \lambda) \circ \Phi^L(X^0, \lambda) \circ \Phi^L(X^+, \lambda). \end{aligned}$$

Since λ'_n agrees with λ on X^0 , it then suffices to show that

$$(84) \quad \Phi^L(X^+, \lambda) = \Phi^L(X^+, \lambda'_n) \circ \Psi_+,$$

$$(85) \quad \Phi^L(X^-, \lambda'_n) = \Psi_- \circ \Phi^L(X^-, \lambda).$$

To prove (84), observe that by Lemmas 5.6, 3.4(a), (b) and 3.7, we have a commutative diagram:

$$\begin{array}{ccc} \text{ECH}_*^L(Y_+, \lambda_+) & \xrightarrow{\Psi_+} & \text{ECH}^L(Y_+, \lambda_+^1) \\ \downarrow s & & \downarrow \Phi^L(X^+, \lambda'_n) \\ \text{ECH}_*^{e^{-\varepsilon}L}(Y_+, e^{-\varepsilon}\lambda_+) & \xrightarrow{\iota} & \text{ECH}_*^L(Y_+, e^{-\varepsilon}\lambda_+) \end{array}$$

By (74), the composition $\iota \circ s$ in the above square is equal to $\Phi^L(X^+, \lambda)$.

To prove (85), by Lemmas 5.6, 3.4(a), (b) and 3.7 again, we have a commutative diagram:

$$\begin{array}{ccc} \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, \lambda_-) & \xrightarrow{s} & \text{ECH}_*^L(Y_-, e^\varepsilon\lambda_-) \\ \downarrow \simeq & & \downarrow \Phi^L(X^-, \lambda'_n) \\ \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, \lambda_-^1) & \xrightarrow{\iota} & \text{ECH}_*^L(Y_-, \lambda_-^1) \end{array}$$

Here the left vertical arrow is induced by (35). Similarly to Lemma 3.4(c), the latter map fits into a commutative diagram:

$$\begin{array}{ccc} \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, \lambda_-) & \xrightarrow{\iota} & \text{ECH}_*^L(Y_-, \lambda_-) \\ \downarrow \simeq & & \downarrow \Psi_- \\ \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, \lambda_-^1) & \xrightarrow{\iota} & \text{ECH}_*^L(Y_-, \lambda_-^1) \end{array}$$

Combining the above two diagrams gives a commutative diagram:

$$\begin{array}{ccc} \text{ECH}_*^{e^{-\varepsilon}L}(Y_-, \lambda_-) & \xrightarrow{s} & \text{ECH}_*^L(Y_-, e^\varepsilon\lambda_-) \\ \downarrow \iota & & \downarrow \Phi^L(X^-, \lambda'_n) \\ \text{ECH}_*^L(Y_-, \lambda_-) & \xrightarrow{\Psi_-} & \text{ECH}_*^L(Y_-, \lambda_-^1) \end{array}$$

By (74) again, the composition $\iota \circ s^{-1}$ in the above square is equal to $\Phi^L(X^-, \lambda)$. \square

Step 3 We now construct a strongly cobordism-admissible almost complex structure J'_n for (X, λ'_n) . On X^0 we take $J'_n = J$. To define J'_n on X^\pm , write $t = 1 \pm \varepsilon^{-1}s$,

and let R_{\pm}^t denote the Reeb vector field associated to λ_{\pm}^t . As a step towards defining J'_n , define an almost complex structure J''_n on X^{\pm} by

$$(86) \quad J''_n \frac{\partial}{\partial s} = R_{\pm}^t, \quad J''_n|_{\text{Ker}(\lambda_{\pm}^t)} = J_{\pm}^t|_{\text{Ker}(\lambda_{\pm}^t)}.$$

It follows from (80) and (81) that

$$(87) \quad \|J''_n - J\|_{C^0} \leq cn^{-1}, \quad \|J''_n\|_{C^1} \leq c.$$

It also follows from (80) that if n is sufficiently large, then J''_n is $d\lambda'_n$ -tame. However J''_n is not necessarily $d\lambda'_n$ -compatible, except near $s = 0, \mp \varepsilon$. We can measure the failure of compatibility by a 2-form Ω on X^{\pm} defined by

$$\Omega(v_1, v_2) := d\lambda'_n(v_1, J''_n v_2) - d\lambda'_n(v_2, J''_n v_1).$$

By (80) and (81), we have

$$\|\Omega\|_{C^1} \leq cn^{-1}.$$

Now Ω , regarded as a bundle map from the space of almost complex structures on X^{\pm} to the space of real $(1, 1)$ -forms, is transverse to 0 at each fiber. It then follows from the inverse function theorem and (87) that if n is sufficiently large, then we can find a $d\lambda'_n$ -compatible almost complex structure J'_n that agrees with J''_n near $s = 0, \mp \varepsilon$, and that satisfies

$$(88) \quad \|J'_n - J\|_{C^0} \leq cn^{-1}, \quad \|J'_n\|_{C^1} \leq c.$$

Step 4 The last step in the construction is to replace (λ'_n, J'_n) by a pair (λ_n, J_n) , which is better behaved with respect to product regions. Let us call an embedded Reeb orbit γ in Y_+ a “product Reeb orbit (with respect to (X, λ, J))” if $\gamma \subset \{s_+\} \times Z$ where $[s_-, s_+] \times Z$ is a product region in X (with respect to λ and J). Fix $\delta > 0$ such that if γ is a product Reeb orbit with action less than L , then:

- (i) If Z denotes the radius δ neighborhood of γ , then $[s_-, s_+] \times Z$ is a product region in X for some s_-, s_+ .
- (ii) γ has distance at least 2δ from all other Reeb orbits in Y_+ with action less than L .

Lemma 6.6 *If $n > \delta^{-1}$, then there is a 1-form λ_n on X such that (X, λ_n) is an exact symplectic cobordism from (Y_+, λ_+^1) to (Y_-, λ_-^1) , and a strongly cobordism-admissible almost complex structure J_n on \bar{X} for λ_n , with the following properties:*

- (a) *The exact symplectic cobordisms (X, λ_n) and (X, λ'_n) from (Y_+, λ_+^1) to (Y_-, λ_-^1) are homotopic in the sense of Section 1.5.*

- (b) If γ is a product Reeb orbit in Y_+ of action less than L with respect to (X, λ, J) , then the radius δ neighborhood of γ is contained in a product region for (X, λ_n, J_n) .
- (c) (λ_n, J_n) agrees with (λ'_n, J'_n) on $\bar{X} \setminus X$, and on the complement in X of the product regions $[s_-, s_+] \times Z$ where Z is the radius $1/n$ neighborhood of a product Reeb orbit in Y_+ of action less than L .
- (d) $\|J_n - J\|_{C^0} \leq cn^{-1}$ and $\|J_n\|_{C^1} \leq c$.

Proof Let γ be a product Reeb orbit in Y_+ with action less than L , and let $[s_-, s_+] \times Z$ be the corresponding product region as in (i) above. By the construction of λ_{\pm}^t , the 1-forms $e^{-s} + \lambda_+^t$ and $e^{-s} - \lambda_-^t$ agree on Z , so let us denote this 1-form simply by λ_0^t . Now on $[s_-, s_+] \times Z$, replace λ_n^t by

$$\lambda_n := e^s \lambda_0^1.$$

To construct J_n on $[s_-, s_+] \times Z$, recall from the construction of J_{\pm}^1 that the restrictions of J_+^1 and J_-^1 to $\text{Ker}(\lambda_0^1)$ agree. Let R_0^t denote the Reeb vector field associated to λ_0^t , and recall from the definition of “product region” that on this region, $J(\partial/\partial s) = fR_0^0$ where f is some function of s , which, by the definition of “strongly cobordism-admissible”, equals e^s near $s = s_{\pm}$. Now define J_n on this region by

$$J_n \frac{\partial}{\partial s} = fR_0^1, \quad J_n|_{\text{Ker}(\lambda_0^1)} = J_{\pm}^1|_{\text{Ker}(\lambda_0^1)}.$$

Let (λ_n, J_n) be obtained by modifying (λ'_n, J'_n) as above for each product Reeb orbit of action less than L . These satisfy properties (a), (b) and (c) by construction, and property (d) follows from (80), (81) and (88). □

We now state a lemma implying that if the hypothesis of (i) or (ii) in the holomorphic curves axiom holds for (X, λ, J) , then it also holds for (X, λ_n, J_n) when n is sufficiently large. Consider pairs (Θ_+, Θ_-) , where Θ_{\pm} is an ECH generator for λ_{\pm} of action less than L . Recall from Definition 3.2 that Θ_{\pm} corresponds to an ECH generator for λ_{\pm}^1 of action less than L , and we denote this also by Θ_{\pm} . Let A denote the set of pairs (Θ_+, Θ_-) for which there exists no broken J -holomorphic curve from Θ_+ to Θ_- . Let A_n denote the set of pairs (Θ_+, Θ_-) for which there exists no broken J_n -holomorphic curve from Θ_+ to Θ_- . Let B denote the set of pairs (Θ_+, Θ_-) for which the only broken J -holomorphic curve from Θ_+ to Θ_- is a union of covers of product cylinders. Let B_n denote the set of pairs (Θ_+, Θ_-) for which the only broken J_n -holomorphic curve from Θ_+ to Θ_- is a union of covers of product cylinders.

Lemma 6.7 *If n is sufficiently large, then $A \subset A_n$ and $B \subset B_n$.*

Lemma 6.7 is proved by a Gromov compactness argument in Section 6.4 below. Assuming this, we can now give:

Proof of the holomorphic curves axiom (strongly cobordism-admissible case) Let n be sufficiently large as in Lemmas 6.6 and 6.7. Define a chain map

$$\widehat{\Phi}: \text{ECC}_*^L(Y_+, \lambda_+; J_+) \longrightarrow \text{ECC}_*^L(Y_-, \lambda_-; J_-)$$

as the composition

$$\begin{aligned} \text{ECC}_*^L(Y_+, \lambda_+; J_+) &\rightarrow \text{ECC}_*^L(Y_+, \lambda_+^1; J_+^1) \\ &\rightarrow \text{ECC}_*^L(Y_-, \lambda_-^1; J_-^1) \rightarrow \text{ECC}_*^L(Y_-, \lambda_-; J_-). \end{aligned}$$

Here the first map is the canonical isomorphism of chain complexes (35) for Y_+ , the second map is the chain map (78) for the cobordism (X, λ_n, J_n) , and the third map is the inverse of the canonical isomorphism of chain complexes (35) for Y_- . By Lemmas 6.5 and 6.6(a) and the Homotopy Invariance axiom, the chain map $\widehat{\Phi}$ induces the map $\Phi^L(X, \lambda)$ on homology.

To prove that $\widehat{\Phi}$ fulfills conditions (i) and (ii) in the holomorphic curves axiom, we must show that if $(\Theta_+, \Theta_-) \in A$ then $\langle \widehat{\Phi}\Theta_+, \Theta_- \rangle = 0$, and if $(\Theta_+, \Theta_-) \in B$ then $\langle \widehat{\Phi}\Theta_+, \Theta_- \rangle = 1$. If $(\Theta_+, \Theta_-) \in A$ (resp. B), then by Lemma 6.7 we have $(\Theta_+, \Theta_-) \in A_n$ (resp. B_n), and by Proposition 6.2 applied to (X, λ_n, J_n) we have $\langle \widehat{\Phi}\Theta_+, \Theta_- \rangle = 0$ (resp. 1). □

6.4 Gromov compactness

We now prove Lemma 6.7. Continuing with the setting of Section 6.3, it is enough to show the following:

Lemma 6.8 *Let Θ_{\pm} be ECH generators for λ_{\pm} of action less than L . Suppose that (n_1, n_2, \dots) is an increasing infinite sequence of positive integers such that for each $n \in \{n_1, n_2, \dots\}$ there exists a broken J_n -holomorphic curve $u_n \in \mathcal{M}^{J_n}(\Theta_+, \Theta_-)$. Then:*

- (a) *After passing to a subsequence, the broken J_n -holomorphic curves u_n converge (in the sense of [5, Section 9], using currents instead of maps) to a broken J -holomorphic curve $u \in \mathcal{M}^J(\Theta_+, \Theta_-)$.*
- (b) *If u is a union of covers of product cylinders, then so is u_n for all sufficiently large n .*

To clarify assertion (a), note that by construction, the Liouville forms λ_n and λ on X have the same Liouville vector field near ∂X , and so there is a canonical diffeomorphism between the completions (14) of (X, λ_n) and (X, λ) , which is the identity on each of the three subsets in (14).

To prove Lemma 6.8, note first that part (b) follows quickly from part (a). The reason is that if u is a union of covers of product cylinders, then by Lemma 6.6(b), if n is sufficiently large then each level of u_n is either (i) a J_n -holomorphic curve in \bar{X} from Θ_+ to Θ_- contained in a product region for (X, λ_n, J_n) , or (ii) a J_{\pm}^1 -holomorphic curve in $\mathbb{R} \times Y_{\pm}$ from Θ_{\pm} to itself. In case (ii), since $d\lambda_{\pm}^1$ is pointwise nonnegative on any J_{\pm}^1 -holomorphic curve, and zero only where the holomorphic curve is tangent to \mathbb{R} cross the Reeb flow, it follows by Stokes' Theorem that any level of type (ii) maps to a union of \mathbb{R} -invariant cylinders, and in particular does not exist by the nontriviality condition in our definition of "broken holomorphic curve". So there is only a level of type (i), and the same argument shows that this maps to a union of product cylinders.

To prove Lemma 6.8(a), first note that the arguments for [5, Lemma 9.8] can be used with only minor notational changes to see that it is enough to prove the following assertion about unbroken holomorphic curves:

Lemma 6.9 *Let Θ_{\pm} be ECH generators for λ_{\pm} of action less than L . Suppose that (n_1, n_2, \dots) is an increasing sequence of positive integers such that for each $n \in \{n_1, n_2, \dots\}$ there is a J_n -holomorphic curve $C_n \in \mathcal{M}^{J_n}(\Theta_+, \Theta_-)$. Then:*

- (a) *After passing to a subsequence, the J_n -holomorphic curves C_n converge as currents on \bar{X} to a J -holomorphic curve $C \in \mathcal{M}^J(\Theta'_+, \Theta'_-)$ for some orbit sets Θ'_{\pm} for λ_{\pm} .*
- (b) *Let s_n be a sequence of positive real numbers with $\lim_{n \rightarrow \infty} s_n = \infty$. Let $C'_n \subset [-s_n, s_n] \times Y_+$ denote the translate by $-s_n$ of the intersection of C_n with $[0, 2s_n] \times Y_+ \subset \bar{X}$. Then after passing to a subsequence, the curves C'_n converge as a current to a J_+ -holomorphic curve in $\mathbb{R} \times Y_+$ between some orbit sets for λ_+ .*
- (c) *Likewise, let s_n be a sequence of negative real numbers with $\lim_{n \rightarrow \infty} s_n = -\infty$. Let $C'_n \subset [s_n, -s_n] \times Y_-$ denote the translate by $-s_n$ of the intersection of C_n with $[2s_n, 0] \times Y_- \subset \bar{X}$. Then after passing to a subsequence, the curves C'_n converge as a current to a J_- -holomorphic curve in $\mathbb{R} \times Y_-$ between some orbit sets for λ_- .*

Note that this lemma does not directly follow from standard Gromov compactness results, because the sequence $\{J_n\}$ does not converge to J in C^1 ; we just have C^0 convergence and a C^1 bound from Lemma 6.6(d).

Proof of Lemma 6.9 We will just prove part (a), as the proofs of parts (b) and (c) are essentially the same. The argument has three steps.

Step 1 We first obtain convergence to some current (which we will later show is J -holomorphic).

Let $\Sigma \subset \bar{X}$ denote the union of the product cylinders $\mathbb{R} \times \gamma$, where γ is a product Reeb orbit of length less than L , the half-cylinders $[-\varepsilon, \infty) \times \gamma_+$, where γ_+ is a Reeb orbit of λ_+ of length less than L , and the half-cylinders $(-\infty, \varepsilon] \times \gamma_-$, where γ_- is a Reeb orbit of λ_- of action less than L . Let $\Sigma_{1/n} \subset \bar{X}$ denote the radius $1/n$ neighborhood of Σ . By construction, (λ_n, J_n) agrees with (λ, J) on $\bar{X} \setminus \Sigma_{1/n}$.

Observe that by Stokes' Theorem,

$$\int_{C_n \cap ((-\infty, 0] \times Y_-)} d\lambda_-^1 + \int_{C_n \cap X} d\lambda_n + \int_{C_n \cap ([0, \infty) \times Y_+)} d\lambda_+^1 = \int_{\Theta_+} \lambda_+^1 - \int_{\Theta_-} \lambda_-^1 \leq L.$$

It follows from this that for any compact set $K \subset \bar{X}$, the area of $C_n \cap K$ has an n -independent upper bound. It now follows from the compactness theorem for currents – see Federer [2, 4.2.17] or Morgan [15, Theorem 5.5] – that we can pass to a subsequence so that $\{C_n\}$ converges weakly as a current to an integral rectifiable current C with locally finite 2-dimensional Hausdorff measure.

Lemma 6.10 *The convergence to C is pointwise in the sense that*

$$(89) \quad \lim_{n \rightarrow \infty} \left(\sup_{x \in C_n \cap K} \text{dist}(x, C_n) + \sup_{x \in C_n \cap K} \text{dist}(x, C) \right) = 0$$

for every compact set $K \subset \bar{X}$.

Proof This is proved by copying the arguments in Taubes [18, Section 5c] and using Lemma 6.11 below. □

Given $\rho > 0$ and $x \in C_n$, let $a_n(x, \rho)$ denote the integral of $d\lambda$ over the subset of C_n with distance less than or equal to ρ from x .

Lemma 6.11 *There exists a constant $\kappa > 1$ such that for all $n > \kappa$ and $x \in C_n$, if $\kappa^{-1} > \rho > \rho' > 0$, then*

$$a_n(x, \rho) > \kappa^{-1} (\rho/\rho')^2 a_n(x, \rho').$$

Proof It follows from Lemma 6.6(d) that J_n is tamed by $d\lambda$ for all sufficiently large n . Moreover, if $|\cdot|$ denotes the metric determined by $d\lambda$ and J , then there exists a constant $\delta > 0$ such that if n is sufficiently large then $d\lambda(v, J_n v) \geq \delta|v|^2$. One can then apply Ye [27, Theorem 2.1]. □

Step 2 We now recall a criterion for C to be J -holomorphic.

Let D denote the closed unit disk. Call a smooth map $\sigma: D \rightarrow \bar{X}$ *admissible* if $\sigma(\partial D) \subset \bar{X} \setminus C$.

Definition 6.12 (cf [18, Section 6a]) A *positive cohomology assignment* is an assignment, to each admissible map σ , of an integer $I(\sigma)$, satisfying the following conditions:

- (a) $I(\sigma) = 0$ if the image of σ is disjoint from C .
- (b) If σ_0 and σ_1 are admissible maps that are homotopic through admissible maps, then $I(\sigma_0) = I(\sigma_1)$.
- (c) If σ is admissible and if $\phi: D \rightarrow D$ is a smooth map so that $\phi: \partial D \rightarrow \partial D$ is a degree k covering, then $I(\sigma \circ \phi) = kI(\sigma)$.
- (d) Suppose that σ is admissible and that $\sigma^{-1}(C)$ is contained in the interior of a finite disjoint union $\bigsqcup_i D_i$ where each D_i is the image of an orientation-preserving embedding $\theta_i: D \rightarrow D$. Then $I(\sigma) = \sum_i I(\sigma \circ \theta_i)$.
- (e) If σ is a J -holomorphic embedding whose image intersects C , then $I(\sigma) > 0$.

If there exists a positive cohomology assignment, then it follows as in [24, Lemma 4.4] that C is a J -holomorphic subvariety of \bar{X} . The arguments in [5, Lemma 9.8] then show that C is an element of $\mathcal{M}^J(\Theta'_+, \Theta'_-)$ for some Θ'_\pm .

Step 3 To complete the proof of Lemma 6.9, we define a positive cohomology assignment I as follows. If $\sigma: D \rightarrow C$ is an admissible map, then it follows from the pointwise convergence (89) that $\sigma(\partial D)$ is disjoint from C_n whenever n is sufficiently large. It then follows from the convergence of currents that the intersection number of D with C_n is independent of n when n is sufficiently large. Define $I(\sigma)$ to be this intersection number.

Conditions (a)–(d) in Definition 6.12 follow directly from the definition of I , together with the fact that C_n converges to C both as a current and pointwise in the sense of (89). Condition (e) is immediate in the special case when σ maps to $\bar{X} \setminus \Sigma$, because then C_n is J -holomorphic in a neighborhood of $\sigma(D)$ for all sufficiently large n . In particular, it follows from [24, Lemma 4.4] that $C \cap (\bar{X} \setminus \Sigma)$ is a J -holomorphic submanifold on the complement of a discrete set. This last fact can also be deduced from standard Gromov compactness theorems – see eg Hummel [3], Wolfson [26], Ye [27] – since the intersection of C with any compact subset of $\bar{X} \setminus \Sigma$ is a pointwise limit of J -holomorphic subvarieties.

It remains to prove condition (e) when $\sigma(D)$ is allowed to intersect Σ . By [18, Lemma 5.5], any holomorphic disk (without boundary constraint) can be perturbed to a holomorphic disk that is transverse to Σ . So by conditions (a) and (d), we can reduce to the case where $\sigma(D)$ has small radius and intersects Σ only at its center point, transversely, which is also in C . To prove property (e) in this case, we use the following lemma, which allows us to perturb a family of J -holomorphic disks to a family of J_n -holomorphic disks.

Lemma 6.13 *Let D_1, D_2 be disks centered at the origin in \mathbb{C} , and let $\phi: D_1 \times D_2 \rightarrow \bar{X}$ be a map such that $\phi|_{D_1 \times \{z_2\}}$ is a J -holomorphic embedding for each $z_2 \in D_2$, and $\phi^{-1}(\Sigma) = \{0\} \times D_2$. After replacing D_1 by a sufficiently small radius subdisk, given $\epsilon > 0$, if n is sufficiently large, then there exists a smooth map $\varphi_n: [0, 1] \times D_1 \times D_2 \rightarrow \bar{X}$ with the following properties:*

- $\varphi_n(0, \cdot, \cdot) = \phi$.
- For each $z_2 \in D_2$, the map $\varphi_n(1, \cdot, z_2)$ is an embedding with J_n -holomorphic image.
- $\sup_{t \in [0, 1], z_1 \in D_1, z_2 \in D_2} \text{dist}(\phi(z_1, z_2), \varphi_n(t, z_1, z_2)) < \epsilon$.

Granted Lemma 6.13, the proof of property (e) is completed as follows. Let $\sigma: D_1 \rightarrow \bar{X}$ be an admissible map that intersects Σ only at its center point, transversely, which is also in C . By [18, Lemma 5.5], we can then find $\phi: D_1 \times D_2 \rightarrow \bar{X}$ as in Lemma 6.13 such that ϕ restricts to a diffeomorphism from a neighborhood of $(0, 0)$ to an open set U in \bar{X} . We can shrink D_1 as in Lemma 6.13, and also shrink D_2 , so that $\phi|_{D_1 \times \{z_2\}}$ is admissible for all $z_2 \in D_2$. By the pointwise convergence (89), if n is sufficiently large, then C_n intersects U . It follows that if ϵ in Lemma 6.13 is chosen sufficiently small, and if n is sufficiently large, then $\varphi_n(1, \cdot, z_2)$ intersects C_n for some $z_2 \in D_2$. Moreover, it follows from the pointwise convergence (89) that if ϵ is sufficiently small and n is sufficiently large then $\varphi_n(1, \cdot, z_2)$ is homotopic to σ through disks whose boundaries do not intersect C_n . Therefore $I(\sigma_1)$ equals the intersection number of C_n with $\varphi_n(1, \cdot, z_2)$ when ϵ is sufficiently small and n is sufficiently large. Since the latter disk is J_n -holomorphic and intersects C_n , we conclude that $I(\sigma_1) > 0$ as desired. \square

Proof of Lemma 6.13 To simplify notation we will just prove the lemma in the case when D_2 is a point, and we will drop z_2 from the notation and write $D = D_1$. The lemma in the general case then follows by noting that the estimates used to prove the lemma when D_2 is a point vary continuously with a smooth family of holomorphic disks. So let $\phi: D \rightarrow \bar{X}$ be a holomorphic map such that $\phi^{-1}(\Sigma) = \{0\}$; we need to show that after replacing D by a smaller radius disk, given $\epsilon > 0$, if n is sufficiently

large then there exists $\varphi_n: [0, 1] \times D \rightarrow \bar{X}$ such that $\varphi_n(0, \cdot) = \phi$, the map $\varphi_n(1, \cdot)$ is an embedding with J_n -holomorphic image, and $\sup_{t \in [0, 1]} \sup_{z \in D} \text{dist}(\phi(z), \varphi_n(t, z)) < \epsilon$. We do so in five steps.

Step 1 We first write down the equations we need to solve in a convenient coordinate system.

We can choose complex coordinates (z, w) for a neighborhood of $\phi(0)$ in \bar{X} with the following properties: First, the intersection of D with this neighborhood is given by $w = 0$. Second, each constant z slice is J -holomorphic. Third, the J version of $T^{1,0}\bar{X}$ is spanned by

$$(90) \quad dz + \sigma d\bar{z}, \quad dw + \gamma d\bar{w},$$

where σ and γ are smooth functions that obey $|\sigma(\cdot, w)| + |\gamma(\cdot, w)| \leq c|w|$. Such coordinates can be found in a neighborhood of any point on a J -holomorphic curve in an almost complex 4-manifold, as explained in [18, Section 5d]. Similarly to (90), the J_n version of $T^{1,0}\bar{X}$ is spanned by

$$(91) \quad dz + \sigma_n d\bar{z} + \mu_n d\bar{w}, \quad dw + \gamma_n d\bar{z} + \nu_n d\bar{w},$$

where $\sigma_n, \mu_n, \gamma_n, \nu_n$ are smooth functions. By Lemma 6.6(d), these satisfy $|\sigma_n - \sigma| + |\mu_n| + |\gamma_n - \gamma| + |\nu_n| \leq cn^{-1}$, and the first derivatives of $\sigma_n - \sigma, \mu_n, \gamma_n - \gamma$ and ν_n are bounded in absolute value by c .

Now fix $r > 0$ such that the coordinates z and w are defined where both have norm less than $2r$, and replace D with the disk $(w = 0, |z| \leq r)$. Let $\eta: D \rightarrow \mathbb{C}$ be a smooth function with $|\eta| < r$. It follows from (91) that the graph $w = \eta(z)$ is J_n -holomorphic if and only if

$$\frac{\partial \eta}{\partial \bar{z}} + \gamma_n - \sigma_n \frac{\partial \eta}{\partial z} + \nu_n \frac{\partial \bar{\eta}}{\partial \bar{z}} + (\mu_n \gamma_n - \sigma_n \nu_n) \frac{\partial \bar{\eta}}{\partial z} + \mu_n \left(\frac{\partial \bar{\eta}}{\partial z} \frac{\partial \eta}{\partial \bar{z}} - \frac{\partial \bar{\eta}}{\partial \bar{z}} \frac{\partial \eta}{\partial z} \right) = 0.$$

It proves useful to rewrite the above equation in the schematic form

$$(92) \quad \frac{\partial \eta}{\partial \bar{z}} + \gamma - \sigma \frac{\partial \eta}{\partial z} + \tau_0 + \tau_1(\eta, \nabla \eta) + \tau_2(\eta, \nabla \eta).$$

Here $\tau_0 = (\gamma_n - \gamma)|_{w=0}$ is a function of z with $|\tau_0| \leq cn^{-1}$ and with first derivatives that are bounded in absolute value by c . Meanwhile $\tau_1(a, \cdot)$ for fixed a is a z -dependent affine linear function that obeys $|\tau_1(a, b)| \leq cn^{-1}(|a| + |b|)$. The first derivatives of $\tau_1(a, \cdot)$ are bounded in absolute value by c . Finally, $\tau_2(a, \cdot)$ for fixed a is a quadratic function of its second entry with $|\tau_2(a, b)| \leq cn^{-1}|b|^2$. The first derivatives of τ_2 with respect to both z and a are bounded in absolute value by c . Also observe that since $\phi^{-1}(\Sigma) = \{0\}$, it follows that for any $\delta > 0$, if n is sufficiently large then $\tau_0 = 0$ where $|z| > \delta$.

To prove [Lemma 6.13](#), it now suffices to show that for every $\varepsilon > 0$, if n is sufficiently large then there exists a solution η_n to the equation (92) with $|\eta_n| < \varepsilon$. One can then define $\varphi_n(t, z) = (z, w = t\eta_n(z))$.

Step 2 We will solve (92) using a fixed point construction in a certain Banach space \mathcal{H} of C^1 functions.

To define the Banach space \mathcal{H} , fix once and for all a number $\nu \in (0, 1/16)$. If \mathbb{V} is any finite-dimensional normed vector space over \mathbb{C} , define a norm $\|\cdot\|_\diamond$ on the space of bounded smooth functions $f: \mathbb{C} \rightarrow \mathbb{V}$ by

$$\|f\|_\diamond^2 := \sup_{z \in \mathbb{C}} \sup_{\rho \in [0, 1]} \rho^{-\nu} \int_{|z'-z| < \rho} |f(z')|^2.$$

Now let \mathcal{C} denote the space of smooth functions $\eta: \mathbb{C} \rightarrow \mathbb{C}$ that are holomorphic on the complement of the unit disk and that satisfy $\lim_{|z| \rightarrow \infty} \eta(z) = 0$. Define a norm $\|\cdot\|_*$ on \mathcal{C} by

$$\|\eta\|_* := \|\nabla\eta\|_2 + \|\nabla\eta\|_\diamond + \|\nabla\nabla\eta\|_\diamond.$$

Finally, define \mathcal{H} to be the completion of \mathcal{C} with respect to the norm $\|\cdot\|_*$. The following lemma about \mathcal{H} will be needed below:

Lemma 6.14 \mathcal{H} is a subset of the Hölder space $C^{1, \nu/2}$, and the inclusion $\mathcal{H} \rightarrow C^{1, \nu/2}$ is a bounded linear map of Banach spaces.

Proof By Morrey [16, Theorem 3.5.2], there exists a constant c (depending on ν) such that

$$(93) \quad |\eta| \leq c\|\nabla\eta\|_\diamond, \quad |\nabla\eta| \leq c\|\nabla\nabla\eta\|_\diamond,$$

and the exponent $\nu/2$ Hölder norm of $|\nabla\eta|$ is also bounded by $c\|\nabla\nabla\eta\|_\diamond$. □

Step 3 (of the proof of [Lemma 6.13](#)) Fix a smooth function $\chi: \mathbb{C} \rightarrow [0, 1]$ that is equal to 1 on the disk of radius $r/4$ and equal to 0 outside of the disk of radius $r/2$. Given $\eta \in \mathcal{C}$, a standard use of the Green’s function for $\bar{\partial}$ on \mathbb{C} finds a unique solution $T = T(\eta) \in \mathcal{C}$ of the equation

$$(94) \quad \frac{\partial T}{\partial \bar{z}} = -\chi \left(\gamma - \sigma \frac{\partial \eta}{\partial z} + \tau_0 + \tau_1(\eta, \nabla\eta) + \tau_2(\eta, \nabla\eta) \right).$$

Here τ_1 and τ_2 should be extended arbitrarily for $|w| > r$ so that they still satisfy the estimates from Step 1. It follows from (94), using (93) and [16, Theorems 3.5.2 and 5.4.1], that

$$(95) \quad \|T\|_* \leq c(\|\tau_0\|_\infty + \|\nabla\tau_0\|_\diamond + n^{-1}\|\eta\|_* + \|\eta\|_*^2).$$

Step 4 Fix $\varepsilon > 0$ and let $\mathcal{H}_\varepsilon \subset \mathcal{H}$ denote the ball of radius ε centered at the origin. We claim that if n is sufficiently large, then the map $\eta \mapsto T(\eta)$ maps $\mathcal{H}_\varepsilon \cap \mathcal{C}$ to itself. By (95), it is enough to show that

$$(96) \quad \|\nabla \tau_0\|_\diamond < \frac{1}{2}c^{-1}\varepsilon$$

if n is sufficiently large, where c here denotes the same constant as in (95). To do so, recall that for any $\delta > 0$, if n is large enough then τ_0 is supported in the disk of radius δ . Then the bound $|\nabla \tau_0| \leq c$ implies that for each z we have

$$\int_{|z'-z|<\rho} |\nabla \tau_0(z')|^2 \leq c \min(\rho^2, \delta^2).$$

It follows that $\|\nabla \tau_0\|_\diamond \leq c\delta^{1-\nu/2}$. By taking δ sufficiently small, we conclude that the desired inequality (96) holds if n is sufficiently large.

Step 5 By Step 4, for any $\varepsilon > 0$, if n is sufficiently large then $T^k(0) \in \mathcal{H}_\varepsilon$ for all $k \geq 0$. By Lemma 6.14 and the Arzelà–Ascoli Theorem, the sequence $\{T^k(0)\}_{k=0,1,\dots}$ then converges uniformly in the C^1 topology to a C^1 function η . Since the convergence is in C^1 , the limit function η obeys (92). Also, elliptic bootstrapping shows that η is in fact C^∞ . Finally, by (93) we have $|\eta| < c\varepsilon$, where c does not depend on ε . As explained at the end of Step 1, this completes the proof of Lemma 6.13. \square

7 Cobordism maps and holomorphic curves (proofs)

To complete the unfinished business, this section proves Propositions 5.2 and 5.4, which were used in Section 5 to define the map on $\widehat{\text{HM}}_L^*$ induced by an exact symplectic cobordism.

7.1 Statement of Proposition 7.1

Propositions 5.2 and 5.4 will be deduced from Proposition 7.1 below, which describes how Seiberg–Witten solutions in a cobordism give rise to holomorphic curves. The statement of Proposition 7.1 requires the following preliminaries.

The Seiberg–Witten action functional Let Y be a closed oriented 3–manifold with a contact form λ , and let J be a symplectization-admissible almost complex structure on $\mathbb{R} \times Y$. These determine a metric on Y according to the conventions in Section 2.2. Fix a spin-c structure and recall the splitting (26).

As noted in Section 2.2, solutions to our perturbed Seiberg–Witten equations (31) on $\mathbb{R} \times Y$ correspond to gradient flow lines of the functional (21), under the identifications

(27), (29) and (30). However it will be convenient below to regard these solutions as gradient flow lines of a different functional α on connections on E and sections of S defined by

$$(97) \quad \alpha(A, \psi) := \frac{1}{2}(cs(A) - rE(A)) + \epsilon_\mu(A) + r \int_Y \langle D_A \psi, \psi \rangle,$$

where the terms in (97) are defined as follows.

Choose a reference (Hermitian) connection A_E on the line bundle E . An arbitrary connection A on E differs from A_E by an imaginary-valued 1-form. We define the Chern–Simons functional

$$cs(A) := - \int_Y (A - A_E) \wedge d(A - A_E) - 2 \int_Y (A - A_E) \wedge (F_{A_E} + \frac{1}{2}F_{A_{K^{-1}}}).$$

Here $A_{K^{-1}}$ is the distinguished connection on K^{-1} defined in Section 2.2. Also, $E(A)$ in (97) is the energy defined in (32), and

$$\epsilon_\mu(A) := i \int_Y (A - A_E) \wedge \mu.$$

The functionals (21) and (97) differ by a constant as follows: If we make the identifications (27), (29) and (30), and choose $\mathbb{A}_0 = A_{K^{-1}} + 2A_E$, then

$$(98) \quad \alpha_\eta(\mathbb{A}, \Psi) = \alpha(A, \psi) + \frac{ir}{2} \int_Y F_{A_E} \wedge \lambda.$$

Geometric setup Proposition 7.1 is applicable to two geometric setups:

Case 1 The first geometric setup, which is needed for Proposition 5.2, is where (X, λ) is an exact symplectic cobordism from (Y_+, λ_+) to (Y_-, λ_-) . In this case let \bar{X} denote the completion of X as in (14). Let us denote the ends of \bar{X} by $\mathcal{E}_- := (-\infty, 0] \times Y_-$ and $\mathcal{E}_+ := [0, \infty) \times Y_+$. Also let $s_*: \bar{X} \rightarrow \mathbb{R}$ denote the piecewise smooth function that agrees with the $(-\infty, 0]$ coordinate on \mathcal{E}_- , agrees with the $[0, \infty)$ coordinate on \mathcal{E}_+ , and equals 0 on X .

Recall from Section 4.2 that to write down the Seiberg–Witten equations (62) on \bar{X} , we need to choose a strongly cobordism-admissible almost complex structure J on \bar{X} – see Definition 4.1 – that restricts to symplectization-admissible almost complex structures J_\pm for λ_\pm on \mathcal{E}_\pm . Then λ and J determine a metric g on \bar{X} , as well as the 2-form $\hat{\omega}$ that appears in (62). We also need to choose small exact 2-forms μ_\pm on Y_\pm , and a small exact 2-form μ on \bar{X} that restricts to μ_\pm on \mathcal{E}_\pm .

Case 2 The second geometric setup, which is needed for Proposition 5.4, considers the composition (X, λ) of exact symplectic cobordisms (X^+, λ^+) from (Y_+, λ_+)

to (Y_0, λ_0) and (X^-, λ^-) from (Y_0, λ_0) to (Y_-, λ_-) . For the purposes of “neck stretching”, given $R \geq 0$ consider the diffeomorphic manifold

$$(99) \quad X_R = X^- \cup \bigcup_{\{-R\} \times Y_0} ([-R, R] \times Y_0) \cup \bigcup_{\{R\} \times Y_0} X^+.$$

Define the completion \overline{X}_R as usual by attaching ends $\mathcal{E}_- = (-\infty, 0] \times Y_-$ and $\mathcal{E}_+ = [0, \infty) \times Y_+$ to X_R . We now specify how to write down a version of the Seiberg–Witten equations (62) on \overline{X}_R .

To start, define $s_*: \overline{X}_R \rightarrow \mathbb{R}$ as follows. Let $s_*^\pm: \overline{X}^\pm \rightarrow \mathbb{R}$ denote the function defined in Case 1 above. Then define s_* to agree with $s_*^- - R$ on $\mathcal{E}_- \cup X^-$, to agree with the $[-R, R]$ coordinate on $[-R, R] \times Y_0$, and to agree with $s_*^+ + R$ on $X^+ \cup \mathcal{E}_+$.

Let $\tilde{\lambda}^\pm$ denote the 1-form on the completion \overline{X}^\pm defined in (58). Define a 1-form $\tilde{\lambda}_R$ on \overline{X}_R by

$$(100) \quad \tilde{\lambda}_R = \begin{cases} e^{-2R} \tilde{\lambda}^- & \text{on } \mathcal{E}_- \cup X^-, \\ e^{2s_*} \lambda_0 & \text{on } [-R, R] \times Y_0, \\ e^{2R} \tilde{\lambda}^+ & \text{on } X^+ \cup \mathcal{E}_+. \end{cases}$$

When R is fixed, we usually denote $\tilde{\lambda}_R$ simply by $\tilde{\lambda}$. Define $\tilde{\omega} = d\tilde{\lambda}$ as before. Note that $(X_R, \tilde{\lambda}|_{X_R})$ is an exact symplectic cobordism from $(Y_+, e^{2R}\lambda_+)$ to $(Y_-, e^{-2R}\lambda_-)$. However, below, references to the “length” of Reeb orbits on Y_\pm refer to the length as defined by λ_\pm , which does not depend on R . We denote this length as usual by A .

Let J_\pm and J_0 be symplectization-admissible almost complex structures for λ_\pm and λ_0 respectively. Let J^\pm be strongly cobordism-admissible almost complex structures on X^\pm restricting to J_\pm and J_0 on the ends. These determine a strongly cobordism-admissible almost complex structure J on \overline{X}_R , which agrees with J^\pm on $\mathcal{E}_\pm \cup X^\pm$, and which agrees with J_0 on $[-R, R] \times Y_0$.

Let g^\pm be the metric on \overline{X}^\pm determined by λ^\pm and J^\pm as in Section 4.2. These extend to a metric g on \overline{X}_R that agrees with g^\pm on $\mathcal{E}_\pm \cup X^\pm$, and which on $[-R, R] \times Y_0$ agrees with the \mathbb{R} -invariant metric on $\mathbb{R} \times Y_0$ determined by λ_0 and J_0 according to the conventions in Section 2.2. Using the metric g , define $\hat{\omega} := \sqrt{2}\tilde{\omega}/|\tilde{\omega}|$ as before.

Finally, let μ_\pm and μ_0 be small exact 2-forms on Y_\pm and Y_0 . Let μ^\pm be small exact 2-forms on X^\pm as in Case 1 that restrict to μ_\pm and μ_0 on the ends. These determine an exact 2-form μ on \overline{X}_R , which restricts to μ^\pm on $\mathcal{E}_\pm \cup X^\pm$, and which restricts to μ_0 on $[-R, R] \times Y_0$.

Below, when we wish to consider both geometric setups simultaneously, we let X_* denote X in Case 1 and X_R in Case 2. Likewise, \overline{X}_* denotes \overline{X} or \overline{X}_R as appropriate.

Variations in the data Proposition 7.1 considers variations in the given data (λ, J, μ) . To clarify, fix $\varepsilon > 0$ for use in defining neighborhoods as in (12) and (13) of the positive and negative boundaries of X in Case 1 or X^\pm in Case 2, and for defining the data on their completions as in Section 4.2. A “variation” then consists of data (λ', J', μ') , which are constrained to be usable above for given data $(\lambda_\pm, J_\pm, \mu_\pm)$ (and (λ_0, J_0, μ_0) in Case 2), with the further requirement that λ' agree with λ on the above boundary neighborhoods. The proposition refers to a “neighborhood” of (λ, J, μ) ; this consists of data (λ', J', μ') as above in a C^∞ –Fréchet neighborhood of (λ, J, μ) .

Index and action difference Let \mathfrak{d} be a instanton solution to (62) on \overline{X}_* . We now introduce two numbers associated to \mathfrak{d} , which will be needed below.

First, let $i_\mathfrak{d}$ denote the index of the instanton \mathfrak{d} . This is the Fredholm index of the operator $D_\mathfrak{d}$ obtained from linearizing the equations (62) at \mathfrak{d} .

Second, recall that the solutions to the perturbed Seiberg–Witten equations in (28) are the critical points of the “Seiberg–Witten action” functional (97) on the space of pairs (A, ψ) . As in Section 4.1, let c_\pm denote the $s_* \rightarrow \pm\infty$ limit of \mathfrak{d} . Let α_\pm denote the Y_\pm version of the action functional. We then define

$$A_\mathfrak{d} := \alpha_-(c_-) - \alpha_+(c_+).$$

Note that while the functionals α_\pm are generally not gauge invariant, the quantity $A_\mathfrak{d}$ is still gauge invariant.

Spinor decomposition If ψ is a section of S_+ , we write $\psi = (\alpha, \beta)$, where α and β respectively denote the E and $K^{-1}E$ components of ψ in the decomposition (61).

Generalized broken J –holomorphic curves If Θ_\pm are orbit sets in Y_\pm , we define a *generalized broken J –holomorphic curve* from Θ_+ to Θ_- to be a collection of holomorphic curves $\{C_k\}_{1 \leq k \leq N}$ as in Definition 1.6, but with one difference: Recall that in Definition 1.6 the curves C_k for $k > k_0$ are in $\mathbb{R} \times Y_+$, the curve C_{k_0} is in \overline{X}_* , and the curves C_k for $k < k_0$ are in $\mathbb{R} \times Y_-$. The difference is that now we do *not* mod out by \mathbb{R} –translation of the curves C_k in $\mathbb{R} \times Y_\pm$ for $k \neq k_0$. Note that if $k > k_0$ we can then identify $C_k \cap ([0, \infty) \times Y_+)$ with a subset of \overline{X}_* , and if $k < k_0$ we can likewise identify $C_k \cap ((-\infty, 0] \times Y_-)$ with a subset of \overline{X}_* .

Proposition 7.1 Fix a data set consisting of (λ, J, μ) . Let $\mathcal{K} \geq 1$ be given, and assume that all Reeb orbits of λ_\pm (and λ_0 in Case 2) of length less than or equal to $(2\pi)^{-1}\mathcal{K}$ are nondegenerate. Then there exist:

- (i) $\kappa \geq 1$,
- (ii) a neighborhood of the given data set, and

(iii) given $\delta > 0$, a number $\kappa_\delta \geq 1$,

such that the following holds: Take $r \geq \kappa_\delta$ and a data set from the given neighborhood (and take any R in Case 2) so as to define (62) on \bar{X}_* . Let $\mathfrak{d} = (A, \psi = (\alpha, \beta))$ denote an instanton solution to this version of (62) with $A_\mathfrak{d} \leq \mathcal{K}r$ or $i_\mathfrak{d} > -\mathcal{K}r$. Assume also that $E(c_+) \leq \mathcal{K}$. Then:

- $E(c_-) \leq E(c_+) + \delta$.
- Each point in \bar{X}_* where $|\alpha| \leq 1 - \delta$ has distance less than $\kappa r^{-1/2}$ from $\alpha^{-1}(0)$.
- There exist
 - (a) a positive integer $N \leq \kappa$ and a partition of \mathbb{R} into intervals $I_1 < \dots < I_N$, each of length at least $2\delta^{-1}$, with $[-1, 1] \subset I_{k_0}$, and
 - (b) a generalized broken J -holomorphic curve $\{C_k\}_{1 \leq k \leq N}$ in \bar{X}_* from an orbit set Θ^+ in Y_+ to an orbit set Θ^- in Y_-

such that for each $k = 1, \dots, N$, with the above identifications of subsets of C_k with subsets of \bar{X}_* , we have

$$\sup_{z \in C_k \cap \mathfrak{d}_*^{-1}(I_k)} \text{dist}(z, \alpha^{-1}(0)) + \sup_{z \in \alpha^{-1}(0) \cap \mathfrak{d}_*^{-1}(I_k)} \text{dist}(C_k, z) < \delta.$$

In particular, Θ_\pm is the orbit set determined by c_\pm under the map in Proposition 3.1(a).

Remark 7.2 The constants κ and κ_δ in Proposition 7.1 can be chosen to be *stable*, by which we mean that they satisfy the conclusions of the proposition for data in some neighborhood of the given data (λ, J, μ) . Various lemmas in the proof of Proposition 7.1 below also refer to constants that are stable in this sense. In general, we omit proofs of stability, as these follow from the proofs below with only cosmetic changes.

7.2 Preliminaries to the proof of Proposition 7.1

An analogue of Proposition 7.1 for a symplectization $\mathbb{R} \times Y$ with \mathbb{R} -invariant (J, μ) was proved in [24, Proposition 5.5]. A slight difference is that [24, Proposition 5.5] applies only to a single data set (λ, J, μ) , while Proposition 7.1 applies to every suitable data set (λ', J', μ') in some neighborhood of a given (λ, J, μ) and to every R where applicable. The proof of Proposition 7.1 below mostly follows the proof of [24, Proposition 5.5], indicating the necessary modifications for our situation. Before starting the proof, we need to make a few more definitions.

The spectral flow function

Returning to the setting of the beginning of Section 7.1, given $r \geq 1$, a pair $c = (A, \psi)$ of a connection on E and a section of \mathbb{S} determines a self-adjoint operator $\mathcal{L}_{r,c}$ defined in [22, Equation (3.8)]. Roughly speaking this operator is the Hessian of α at c (after modding out by gauge transformations). Let us call a pair (r, c) *nondegenerate* if the corresponding operator $\mathcal{L}_{r,c}$ has trivial kernel.

Now fix a reference pair c_* such that the pair $(1, c_*)$ is nondegenerate, and fix $r \geq 1$. If c is such that the pair (r, c) is nondegenerate, then we define the *spectral flow function* $f(c)$ to be the spectral flow from \mathcal{L}_{1,c_*} to $\mathcal{L}_{r,c}$.

If the spin-c structure has non-torsion first Chern class, ie, if $c_1(\det(\mathbb{S}))$ is not torsion in $H^2(Y; \mathbb{Z})$, then the functional α is not invariant under the action of the gauge group $C^\infty(Y; S^1)$, and neither is the spectral flow function f . However the combination

$$(101) \quad \alpha^f(\cdot) := \alpha(\cdot) - 2\pi^2 f(\cdot)$$

is always gauge invariant.

Index and spectral flow

Returning to the setting of Proposition 7.1, we now relate the index of an instanton to the spectral flow functions on Y_\pm . Fix a spin-c structure \mathbb{S} on \bar{X}_* and let E be defined by the splitting (61). Fix a reference pair $\mathfrak{d}_* = (A_*, \psi_*)$ of a connection on E and a section of \mathbb{S}_+ with the following properties: First, the restriction to the $\pm s > 1$ portion of \bar{X}_* is pulled back from a configuration $c_{\pm*}$ on Y_\pm . Second, require that the pair $(r = 1, c_{\pm*})$ is nondegenerate in the sense described above. This guarantees that the operator $D_{\mathfrak{d}_*}$ is Fredholm. (Note that this operator is defined regardless of whether \mathfrak{d}_* solves the Seiberg–Witten equations (62).) Let i_* denote the index of $D_{\mathfrak{d}_*}$. Let f_\pm denote the spectral flow function on Y_\pm defined using $c_{\pm*}$ as the reference pair.

If \mathfrak{d} is an instanton solution to (62) with nondegenerate $s_* \rightarrow \pm\infty$ limits c_\pm , then it follows from [1] that its index is given by⁹

$$(102) \quad i_{\mathfrak{d}} = i_* + f_+(c_+) - f_-(c_-).$$

7.3 Estimates on instantons

To begin the proof of Proposition 7.1, we now establish various estimates for instanton solutions to (62) on \bar{X}_* , parallel to [24, Section 3], where analogous estimates are derived for instantons on a symplectization. Assume in what follows that $(\lambda_\pm, J_\pm, \mu_\pm)$

⁹In a symplectization with \mathbb{R} -invariant (J, μ) , one can take \mathfrak{d}_* to be independent of the \mathbb{R} factor, so that $i_* = 0$. In this case $i_{\mathfrak{d}}$ agrees with the quantity $f_{\mathfrak{d}}$ in [24].

(and $(\lambda_0, J_0, \mu_0, R)$ in Case 2) are given. Fix data (λ, J, μ) as in Section 7.1. Below, c_0 denotes a number that is greater than 1, that is stable in the sense of Remark 7.2, and that does not depend on any given solution to (28) or (62) or on the value of r used to define these equations. The value of c_0 can increase from one appearance to the next.

Lemma 7.3 (cf [24, Lemma 3.1]) *There exists a stable $\kappa \geq 1$ such that if $r \geq \kappa$ and if $(A, \psi = (\alpha, \beta))$ is an instanton solution to (62) on \bar{X}_* , then*

$$|\alpha| \leq 1 + \kappa r^{-1},$$

$$|\beta|^2 \leq \kappa r^{-1}(1 - |\alpha|^2) + \kappa^2 r^{-2}.$$

Proof This follows from the maximum principle as in [24, Lemma 3.1], using the corresponding inequalities in the 3–dimensional case [19, Lemma 2.2] to obtain the necessary bounds as $s \rightarrow \pm\infty$. □

Lemma 7.4 (cf [24, Lemma 3.2]) *There exists a stable $\kappa \geq 1$ with the following property: Suppose that $r \geq \kappa$ and that $\mathfrak{d} = (A, \psi)$ is an instanton solution to (62) on \bar{X}_* with $A_{\mathfrak{d}} \leq r^2$ or $i_{\mathfrak{d}} \geq -r^2$. Then $|F_A| \leq \kappa r$.*

Proof Copy the proof of [24, Lemma 3.2], replacing [24, Lemma 3.3] in that argument with Lemma 7.5 below. □

To state the next lemma, let ∇_A denote the covariant derivative on $S_+ = E \oplus K^{-1}E$ determined by the connection A on E together with the distinguished connection $A_{K^{-1}}$ on K^{-1} from Section 4.2. Note that under the identification (27), the difference $\nabla_A - \nabla_{\mathbb{A}} \in \Omega^1(\bar{X}_*; \text{End}(S_+))$ is bounded in C^0 and does not depend on A .

Lemma 7.5 (cf [24, Lemma 3.3]) *There exists a stable $\kappa \geq 1$ with the following property: Suppose that $r \geq \kappa$ and that $\mathfrak{d} = (A, \psi)$ is an instanton solution to (62) on \bar{X}_* with $A_{\mathfrak{d}} \leq r^2$ or $i_{\mathfrak{d}} \geq -r^2$. Let $I \subset \mathbb{R}$ denote an interval of length 2. Then*

$$\int_{s_*^{-1}(I)} (|F_A|^2 + r|\nabla_A \psi|^2) \leq \kappa r^2.$$

The proof of Lemma 7.5 requires two additional lemmas. To state these, let $I_0 = \{0\}$ in Case 1 and let $I_0 = \{-R, R\}$ in Case 2. On $s_*^{-1}(\mathbb{R} \setminus I_0)$, define two 1–forms by

$$(103) \quad \mathcal{B}_{(A, \psi)} := *F_A - r(\tau(\psi) - ia) - i*\mu + \frac{1}{2}*F_{A_{K^{-1}}},$$

$$E_A := F_A \left(\frac{\partial}{\partial s}, \cdot \right).$$

Here $*$ denotes the three-dimensional Hodge star, a denotes the relevant contact form (λ_{\pm} or λ_0), and s denotes the \mathbb{R} coordinate on $\mathbb{R} \times Y_{\pm}$ or $\mathbb{R} \times Y_0$. Also let $\nabla_{A,s}$ denote the covariant derivative with respect to the connection ∇_A on S_+ in the direction $\partial/\partial s$. We then have:

Lemma 7.6 (cf [24, Lemma 3.4]) *There exists a stable $\kappa \geq 1$ such that if $r \geq \kappa$ and if $\bar{\mathfrak{d}} = (A, \psi)$ is an instanton solution to (62), then the following hold:*

(a) *Suppose that $s_+ \geq s_-$ are in the same component of $\mathbb{R} \setminus I_0$. Then*

$$(104) \quad \alpha(\bar{\mathfrak{d}}|_{s_* = s_-}) - \alpha(\bar{\mathfrak{d}}|_{s_* = s_+}) = \frac{1}{2} \int_{s_* \in [s_-, s_+]} (|E_A|^2 + |\mathcal{B}_{(A, \psi)}|^2 + 2r(|\nabla_{A,s} \psi|^2 + |D_{A(s_*)} \psi|^2)).$$

Here α denotes the functional (97) on Y_+ , Y_- or Y_0 as appropriate.

(b) *In Case 1,*

$$(105) \quad \kappa^{-1} \int_X (|F_A|^2 + r|\nabla_A \psi|^2) \leq -\alpha(\bar{\mathfrak{d}}|_{\partial X}) + \kappa r,$$

where $\alpha(\bar{\mathfrak{d}}|_{\partial X}) := \alpha(\bar{\mathfrak{d}}|_{\{0\} \times Y_+}) - \alpha(\bar{\mathfrak{d}}|_{\{0\} \times Y_-})$. In Case 2, analogous inequalities hold with X replaced by X^- or X^+ .

(c) *If $s_+, s_- \in \mathbb{R} \setminus I_0$ and $s_+ > s_-$ then:*

$$\kappa^{-1} \int_{s_*^{-1}[s_-, s_+]} (|F_A|^2 + 2r|\nabla_A \psi|^2) \leq \alpha(\bar{\mathfrak{d}}|_{s_* = s_-}) - \alpha(\bar{\mathfrak{d}}|_{s_* = s_+}) + \kappa(s_+ - s_-)r^2 + \kappa r.$$

$$(d) \quad \frac{1}{2} \int_{s_*^{-1}(\mathbb{R} \setminus I_0)} (|E_A|^2 + |\mathcal{B}_{(A, \psi)}|^2 + 2r(|\nabla_{A,s} \psi|^2 + |D_A \psi|^2)) \leq \alpha(c_-) - \alpha(c_+) + \kappa r.$$

Proof (a) We can apply a gauge transformation to put the connection A into temporal gauge (52) on $s_*^{-1}[s_-, s_+]$. Equation (104) then becomes

$$(106) \quad \alpha(\bar{\mathfrak{d}}|_{s_* = s_-}) - \alpha(\bar{\mathfrak{d}}|_{s_* = s_+}) = \frac{1}{2} \int_{s_*^{-1}[s_-, s_+]} \left(\left| \frac{\partial A}{\partial s_*} \right|^2 + |\mathcal{B}_{(A, \psi)}|^2 + 2r \left(\left| \frac{\partial \psi}{\partial s_*} \right|^2 + |D_{A(s_*)} \psi|^2 \right) \right).$$

This is equivalent to the first equation in [24, Lemma 3.4]. An alternate way to understand this equation is to recall that $(A(s_*), \psi(s_*))$ is a downward gradient flow line of the functional α in (97). In particular, the L^2 gradient of α at (A, ψ) is

$(\mathcal{B}_{(A,\psi)}, \sqrt{2r} D_A \psi)$. Equation (106) then follows from the fact that if $\gamma(s)$ is a downward gradient flow line of a function f then

$$(107) \quad f(s_-) - f(s_+) = \frac{1}{2} \int_{s \in [s_-, s_+]} (\|\nabla f\|^2 + \|\partial\gamma/\partial s\|^2).$$

(b) We just consider Case 1 since the proof in Case 2 is the same. Recall that our solution (A, ψ) to (62) corresponds to a solution (\mathbb{A}, Ψ) of (51) via (27), (30) and (63). Identify Y_{\pm} with $\{0\} \times Y_{\pm}$ in \bar{X} . By (98), $\alpha(\partial_{\partial X})$ differs by an $O(r)$ constant from $\alpha_{\eta}(\partial_{\partial X}) := \alpha_{\eta_+}(\partial_{Y_+}) - \alpha_{\eta_-}(\partial_{Y_-})$, so it is enough to prove the claim with $\alpha(\partial_{\partial X})$ replaced by $\alpha_{\eta}(\partial_{\partial X})$.

Recall from Section 4.2 that $\hat{\omega} = \sigma^{-1} d\tilde{\lambda}$ where $\sigma: X \rightarrow [3/2, 5/2]$ is a smooth function with $\sigma|_{\partial X} = 2$. Now start with the Bochner–Weitzenböck formula

$$D_{\mathbb{A}}^* D_{\mathbb{A}} \Psi = \nabla_{\mathbb{A}}^* \nabla_{\mathbb{A}} \Psi + \frac{1}{2} \text{cl}(F_{\mathbb{A}}^+) \Psi + \frac{s}{4} \Psi,$$

where s denotes the scalar curvature of X . Putting in the Dirac equation $D_{\mathbb{A}} \Psi = 0$ from (51), multiplying the resulting equation by σ , taking the inner product with Ψ and integrating by parts gives

$$(108) \quad 0 = \int_X \sigma |\nabla_{\mathbb{A}} \Psi|^2 + \frac{1}{2} \int_X \sigma \langle \text{cl}(F_{\mathbb{A}}^+) \Psi, \Psi \rangle + \frac{1}{4} \int_X \sigma s |\Psi|^2 + \int_X \langle d\sigma \otimes \Psi, \nabla_{\mathbb{A}} \Psi \rangle + 2 \int_{Y_+} \langle D_{\mathbb{A}|_{Y_+}} \Psi, \Psi \rangle - 2 \int_{Y_-} \langle D_{\mathbb{A}|_{Y_-}} \Psi, \Psi \rangle.$$

Second, taking the norm square of the curvature equation in (51), multiplying by σ , and integrating over X gives

$$(109) \quad 0 = \int_X \sigma |F_{\mathbb{A}}^+|^2 + \frac{1}{2} \int_X \sigma |\Psi|^4 + \int_X \sigma |\eta|^2 - \int_X \sigma \langle \text{cl}(F_{\mathbb{A}}^+) \Psi, \Psi \rangle - 2 \int_X \sigma \langle F_{\mathbb{A}}^+, i\eta \rangle + \int_X \sigma \langle \text{cl}(i\eta) \Psi, \Psi \rangle.$$

Third, by (21) and Stokes' Theorem we have

$$(110) \quad \alpha_{\eta}(\partial|_{\partial X}) = \frac{1}{8} \int_X (|F_{\mathbb{A}}^+|^2 - |F_{\mathbb{A}}^-|^2 + F_{\mathbb{A}_0} \wedge F_{\mathbb{A}_0}) + \frac{1}{2} \int_{Y_+} \langle D_{\mathbb{A}|_{Y_+}} \Psi, \Psi \rangle - \frac{1}{2} \int_{Y_-} \langle D_{\mathbb{A}|_{Y_-}} \Psi, \Psi \rangle + \frac{1}{4} \int_{Y_+} (\mathbb{A} - \mathbb{A}_0) \wedge i\eta_+ - \frac{1}{4} \int_{Y_-} (\mathbb{A} - \mathbb{A}_0) \wedge i\eta_-.$$

Here \mathbb{A}_0 is any reference connection on $\det(\mathbb{S})$ over X extending the chosen reference connections over Y_+ and Y_- , and η_{\pm} denotes the perturbation (29) for Y_{\pm} .

Adding two times equation (108) to equation (109) and subtracting eight times equation (110) gives

$$\begin{aligned}
 (111) \quad & -8a_{\eta}(\mathfrak{D}|_{\partial X}) \\
 & = 2 \int_X \sigma |\nabla_{\mathbb{A}} \Psi|^2 + \int_X ((\sigma - 1)|F_{\mathbb{A}}^+|^2 + |F_{\mathbb{A}}^-|^2) + \int_X \langle d\sigma \otimes \Psi, \nabla_{\mathbb{A}} \Psi \rangle \\
 & \quad + \frac{1}{2} \int_X \sigma s |\Psi|^2 - \int_X F_{\mathbb{A}_0} \wedge F_{\mathbb{A}_0} \\
 & \quad + \frac{1}{2} \int_X \sigma |\Psi|^4 + \int_X \sigma \langle \text{cl}(i\eta)\Psi, \Psi \rangle + \int_X \sigma |\eta|^2 \\
 & \quad - 2 \int_X \sigma \langle F_{\mathbb{A}}^+, i\eta \rangle - 2 \int_{Y_+} (\mathbb{A} - \mathbb{A}_0) \wedge i\eta_+ + 2 \int_{Y_-} (\mathbb{A} - \mathbb{A}_0) \wedge i\eta_-.
 \end{aligned}$$

On the right side of (111), in the first term we have $|\nabla_{\mathbb{A}} \Psi|^2 = 2r|\nabla_A \psi|^2 + O(r)$, since $(2r)^{-1/2}|\Psi| = |\psi| = O(1)$ by Lemma 7.3; in the second term we have $F_{\mathbb{A}} = 2F_A + O(1)$; and in the third term we have

$$\langle d\sigma \otimes \Psi, \nabla_{\mathbb{A}} \Psi \rangle \geq -\frac{1}{100} |\nabla_{\mathbb{A}} \Psi|^2 - c_0 |\Psi|^2,$$

where $|\Psi|^2$ is $O(r)$ by Lemma 7.3. The second line on the right side of (111) is $O(r)$ by Lemma 7.3 again. Using $\Psi = \sqrt{2r}(\alpha, \beta)$ and (63), we can expand the sum of the integrands in the third line of the right side of (111) as σ times

$$\begin{aligned}
 \frac{1}{2} |\Psi|^4 + \langle \text{cl}(i\eta)\Psi, \Psi \rangle + |\eta|^2 & = 2r^2((|\alpha|^2 - 1)^2 + 2|\alpha|^2|\beta|^2 + 2|\beta|^2 + |\beta|^4) \\
 & \quad + \langle \text{cl}(2i\mu_*)\Psi, \Psi \rangle - 4r\langle \widehat{\omega}, \mu_* \rangle + 4|\mu_*|^2,
 \end{aligned}$$

which is $O(r)$ by Lemma 7.3. Since the 2-forms η_{\pm} on Y_{\pm} extend over X to the exact 2-form $-r d\tilde{\lambda} + 4\mu$, the fourth line on the right side of (111) can be rewritten using Stokes' Theorem as

$$(112) \quad 2 \int_X F_{\mathbb{A}_0} \wedge i(r d\tilde{\lambda} - 4\mu) + \int_X (\langle F_{\mathbb{A}}^+, 4i\sigma\mu \rangle + \langle F_{\mathbb{A}}^-, 8i\mu \rangle).$$

The first term in (112) is $O(r)$. Since we assumed in Section 4.2 that $|\mu| \leq 1/100$, the second term in (112) is bounded from below by $\frac{-1}{10} \int_X (|F_{\mathbb{A}}^+|^2 + |F_{\mathbb{A}}^-|^2)$, so we can combine this with the second term in the first line on the right hand side of (111) to obtain the desired inequality (105).

(c) By part (b), it is enough to show that the stated inequality holds when s_+ and s_- are in the same component of $\mathbb{R} \setminus I_0$. We can further replace the functional α by α_{η_0} , where η_0 denotes the perturbation (29) for Y_+ , Y_- or Y_0 as appropriate.

As in (107), we have

$$\alpha_{\eta_0}(\partial_{s_* = s_-}) - \alpha_{\eta_0}(\partial_{s_* = s_+}) = \frac{1}{2} \int_{s_*^{-1}[s_-, s_+]} \left(\frac{1}{4} |-*F_{\mathbb{A}(s_*)} + \tau(\Psi(s_*)) + i*\eta_0|^2 + |D_{\mathbb{A}(s_*)}\Psi|^2 + \frac{1}{4} |\partial_s \mathbb{A}|^2 + |\partial_s \Psi|^2 \right).$$

Expanding the first term in the integrand, and using the Bochner–Weitzenböck formula for the three-dimensional Dirac operator $D_{\mathbb{A}(s_*)}$ on constant s_* slices to expand the second term in the integrand, the right hand side becomes

$$\frac{1}{2} \int_{s_*^{-1}[s_-, s_+]} \left(\frac{1}{4} |F_{\mathbb{A}}|^2 + |\nabla_{\mathbb{A}} \Psi|^2 + \frac{s}{4} |\Psi|^2 + \frac{1}{4} |\tau(\Psi(s_*))|^2 + \frac{1}{4} |\eta_0|^2 + \frac{1}{2} \langle \tau(\Psi(s_*)), i*\eta_0 \rangle - \frac{1}{2} \langle *F_{\mathbb{A}(s_*)}, i*\eta_0 \rangle \right).$$

The sum of the first two terms in the integrand is $|F_A|^2 + 2r|\nabla_A \psi|^2 + O(r)$. The third and fourth terms are $O(r)$ by Lemma 7.3, the fifth term is $O(r^2)$ by (29), and likewise the sixth term is $O(r^{3/2})$. The last term is $O(r^2)$ because $|*F_{\mathbb{A}(s_*)}| \leq 2|F_{\mathbb{A}}^+|$, which is $O(r)$ as noted in the proof of (b).

(d) This follows immediately from (a) and (b). □

Continuing with the proof of Lemma 7.5, note that the case $A_{\mathfrak{d}} \leq r^2$ follows immediately from Lemma 7.6(a)–(c). To deal with the remaining cases we need:

Lemma 7.7 (cf [22, Lemma 3.5]) *There exists a stable constant $\kappa \geq 1$ such that if $\mathfrak{d} = (A, \psi)$ is an instanton solution to (62), then*

$$\alpha(c_-) - \alpha(c_+) \leq -2\pi^2 i_{\mathfrak{d}} + \frac{r}{2} (E(A_+) - E(A_-)) + \kappa r^{2/3} (\ln r)^\kappa (1 + |E(A_+)|^{4/3} + |E(A_-)|^{4/3}).$$

Proof By equations (101) and (102) we have

$$\alpha(c_-) - \alpha(c_+) = \alpha^f(c_-) - \alpha^f(c_+) - 2\pi^2 (i_{\mathfrak{d}} - i_*) .$$

The lemma then follows from [21, Proposition 4.10]. □

Proof of Lemma 7.5 Using Lemmas 7.6(a), (d) and 7.7, the arguments in the proof of [24, Lemma 3.3] establish the assertions of Lemma 7.5 if $\text{dist}(I, I_0) \geq T = c_0(\ln r)^{c_0}$.

To deal with the remaining cases, we will restrict to Case 1, as the proof in Case 2 is very similar. By what was just said, there exist points $s_- \in [-T - 2, -T]$ and $s_+ \in [T, T + 2]$ such that

$$(113) \quad \int_{s_*=s_{\pm}} (|F_A|^2 + r|\nabla_A \psi|^2) \leq c_0 r^2.$$

Now let $A_{E_{\pm}}$ denote the reference connection on $E|_{Y_{\pm}}$ used to define the functional α in (97) for Y_{\pm} . It is convenient below to choose the reference connection $A_{E_{\pm}}$ so that $F_{A_{E_{\pm}}} + \frac{1}{2}F_{A_{K-1}}$ is harmonic on Y_{\pm} . Choose identifications of $E_-|_{(-\infty, 0] \times Y_-}$ and $E_+|_{[0, \infty) \times Y_+}$ with the pullbacks of E_- and E_+ respectively. Extend $A_{E_{\pm}}$ to a reference connection A_E on E over \bar{X} such that on $(-\infty, 0] \times Y_-$ and $[0, \infty) \times Y_+$, with respect to the above identifications, we have

$$(114) \quad \nabla_{A_E} = \partial_s + \nabla_{A_{E_{\pm}}}.$$

Let

$$(115) \quad \nu_E := \frac{i}{\pi} \left(F_{A_E} + \frac{1}{2} F_{A_{K-1}} \right).$$

This is a closed form that represents the cohomology class $\frac{1}{2}c_1(\mathfrak{s})$. Write $A = A_E + \hat{a}$. We claim that

$$(116) \quad \left| \alpha(\partial|_{s_*=s_-}) - \alpha(\partial|_{s_*=s_+}) \right| \leq c_0 \left| \int_{s_*=s_-} \hat{a} \wedge \nu_E - \int_{s_*=s_+} \hat{a} \wedge \nu_E \right| + c_0 r^2.$$

To prove (116), note that the functional α on Y_{\pm} is invariant under nullhomotopic gauge transformations. Thus to evaluate $\alpha(\partial_{s_*=s_+})$ or $\alpha(\partial_{s_*=s_-})$, we may assume that $\hat{a}|_{s_*=s_{\pm}} = \sigma + \hat{a}'$ where σ is harmonic, \hat{a}' is co-closed, and \hat{a}' is orthogonal to the space of harmonic forms on Y_{\pm} , so that

$$\int_{s_*=s_{\pm}} \nu_E \wedge \hat{a} = \int_{s_*=s_{\pm}} \nu_E \wedge \sigma \quad \text{and} \quad \|\hat{a}'\|_{L^2_1} \leq c_0 \|d\hat{a}'\|_{L^2}.$$

By (113), the last inequality implies that $\|\hat{a}'\|_{L^2_1} \leq c_0 r$, and (113) also implies that $\|\nabla_A \psi|_{s_*=s_{\pm}}\|_{L^2} \leq c_0 r^{1/2}$. Putting these last two inequalities and Lemma 7.3 into (97), we obtain

$$\alpha(\partial|_{s_*=s_{\pm}}) = i\pi \int_{s_*=s_{\pm}} \hat{a} \wedge \nu_E + O(r^2),$$

from which (116) follows.

Applying Stokes' Theorem to the right hand side of (116), and using the fact that $|v_E|$ and $|F_{AE}|$ enjoy r -independent upper bounds, we obtain

$$\alpha(\partial|_{s_*=s_-}) - \alpha(\partial|_{s_*=s_+}) \leq c_0 T^{1/2} \left(\int_{s_*^{-1}[s_-, s_+]} |F_A|^2 \right)^{1/2} + c_0 T + c_0 r^2.$$

It follows from this and Lemma 7.6(c) that if r is larger than some stable constant, then

$$\int_{s_*^{-1}[s_-, s_+]} |F_A|^2 \leq c_0 T r^2.$$

Putting this inequality back into the previous one, we obtain

$$\alpha(\partial|_{s_*=s_-}) - \alpha(\partial|_{s_*=s_+}) \leq c_0 r^2.$$

It follows from this and Lemma 7.6(a), (b) that

$$(117) \quad \alpha(\partial|_{s_*=s'_-}) - \alpha(\partial|_{s_*=s'_+}) \leq c_0 r^2$$

whenever $s_- \leq s'_- \leq s'_+ \leq s_+$. When $s'_+ - s'_- = 2$, using Lemma 7.6(c) with (117) proves the remaining cases of Lemma 7.5. □

We will also need counterparts of the estimates in [24, Lemmas 3.6–3.10, 4.3]. Since these lemmas and their proofs carry over almost verbatim to our setting, we will not repeat them here, but just note the following modifications: First, the constants κ provided by these lemmas are stable. The constant κ_q provided by [24, Lemma 3.6] is also stable, although the neighborhood of stability may depend on q . Second, $[x, y] \times M$ is to be replaced by $s_*^{-1}[x, y]$, and $\mathbb{R} \times M$ is to be replaced by \bar{X}_* . Third, f_∂ is to be replaced by i_∂ . Finally, $\frac{\partial}{\partial s} A \pm B_A$ is to be replaced by F_A^\pm , the (anti-)self-dual part of the curvature F_A .

7.4 Instantons and holomorphic curves

We now establish counterparts of results from [24, Section 4]. The latter explains how instantons can be used to define parts of holomorphic curves.

We need to introduce the following notation. Continue with the geometric setup from Section 7.1. If $(A, \psi = (\alpha, \beta))$ is an instanton solution to the perturbed Seiberg–Witten equations (62) on \bar{X}_* , define a function $\underline{M}: \mathbb{R} \rightarrow \mathbb{R}$ by

$$\underline{M}(s) := r \int_{s_*^{-1}[s-1, s+1]} (1 - |\alpha|^2).$$

The idea of this function is that an r -independent upper bound on \underline{M} will allow us to find, for large r , a holomorphic curve near the zero locus of α whose area in

$s_*^{-1}[s - 1, s + 1]$ is approximately $\frac{1}{2\pi} \underline{M}(s)$. In particular, the propositions to follow assume certain upper bounds on \underline{M} ; we will establish upper bounds on \underline{M} later in Section 7.5.

Meanwhile, continuing with the notation, define a connection \widehat{A} on E by

$$\widehat{A} := A - \frac{1}{2}(\bar{\alpha} \nabla_A \alpha - \alpha \nabla_A \bar{\alpha}).$$

Note that this has curvature

$$F_{\widehat{A}} = (1 - |\alpha|^2) F_A - \nabla_A \bar{\alpha} \wedge \nabla_A \alpha.$$

Below, on the subsets of \bar{X}_* that are identified with $(-\infty, 0) \times Y_-$ or $[0, \infty) \times Y_+$, or $[-R, R] \times Y_0$ in Case 2, let λ_Y denote the pullback of the relevant contact form λ_{\pm} or λ_0 on Y_{\pm} or Y_0 .

Proposition 7.8 (cf [24, Proposition 4.1]) *Given $\delta > 0$ and $\mathcal{K} \geq 1$, there exist a stable $\kappa \geq 1$ and a neighborhood of the given data set (λ, J, μ) such that the following holds: Let $r \geq \kappa$, and let $\mathfrak{d} = (A, \psi = (\alpha, \beta))$ be an instanton solution to (62) as defined with a data set from this neighborhood of (λ, J, μ) (and any R in Case 2). Assume that $A_{\mathfrak{d}} \leq r^2$ or $i_{\mathfrak{d}} \geq -r^2$. Let \mathbb{I} be a connected subset of \mathbb{R} of length at least $2\delta^{-1} + 16$ such that $\underline{M}(\cdot) \leq \mathcal{K}$ on \mathbb{I} . Let $I \subset \mathbb{I}$ be a connected set of points with distance at least 7 from the boundary of \mathbb{I} and length $2\delta^{-1}$. Then:*

- Each point in $s_*^{-1}(I)$ where $|\alpha| \leq 1 - \delta$ has distance at most $\kappa r^{-1/2}$ from $\alpha^{-1}(0)$.
- There exists a finite set \mathcal{C} of pairs of the form (C, m) where C is a closed, irreducible J -holomorphic subvariety in a neighborhood of the closure of $s^{-1}(I)$, m is a positive integer, and the subvarieties C for different elements of \mathcal{C} are distinct, such that:

(a)
$$\sup_{z \in \bigcup_{(C,m) \in \mathcal{C}} C \cap s_*^{-1}(I)} \text{dist}(z, \alpha^{-1}(0)) + \sup_{z \in \alpha^{-1}(0) \cap s_*^{-1}(I)} \text{dist}\left(z, \bigcup_{(C,m) \in \mathcal{C}} C\right) < \delta.$$

- (b) Let v be a 2-form on \bar{X}_* with support in $s_*^{-1}(I)$ such that $|v| \leq 1$ and $|\nabla v| \leq \delta^{-1}$. Then

$$\left| \frac{i}{2\pi} \int_{s_* \in I} v \wedge F_{\widehat{A}} - \sum_{(C,m) \in \mathcal{C}} m \int_C v \right| \leq \delta.$$

- (c) Let $I_Y \subset I$ denote a component of the subset of I where the distance to I_0 is at least 2. Then

$$\sum_{(C,m) \in \mathcal{C}} m \int_{C \cap s_*^{-1}(I_Y)} d\lambda_Y \leq \kappa.$$

Proof The proof of the first bullet differs only in notation from the proof of the first bullet of [24, Proposition 4.1], except for the following change: Violation of the first bullet requires sequences $\{(\lambda_n, J_n, \mu_n)\}_{n=1,2,\dots}$ and $\{r_n, A_n, \psi_n\}_{n=1,2,\dots}$, as well as a sequence $\{R_n\}_{n=1,2,\dots}$ in Case 2, where $\{(\lambda_n, J_n, \mu_n)\}_{n=1,2,\dots}$ converges to (λ, J, μ) , the pair (A_n, ψ_n) solves the version of (62) defined using $(\lambda_n, J_n, \mu_n, r_n, R_n)$, and the rest of the assumptions on the sequence are the same as in [24]. Note that the *stability* of the constants κ provided by the lemmas in Section 7.3 must be used to obtain the contradiction that proves the first bullet.

The proof of the second bullet is obtained by similarly modifying the proof of the second bullet of [24, Proposition 4.1], using the stability of the constants κ , and making the usual notational changes to replace $\mathbb{R} \times M$ in [24] by \bar{X}_* here. In particular, the form $ds \wedge a + \frac{1}{2} *a$ in [24, Equation (4.5)] is to be replaced by \hat{w} here. \square

The following proposition is similar to Proposition 7.8, but with the interval I expanded.

Proposition 7.9 (cf [24, Proposition 4.5]) *Given $\mathcal{K} \geq 1$, suppose that each Reeb orbit with length $\leq \mathcal{K}/2\pi$ of λ_{\pm} , and of λ_0 in Case 2, is nondegenerate. Then there exists $\kappa \geq 1$, and given $\delta > 0$ there exist $\kappa_{\delta} \geq 1$ and a neighborhood of the given data set (λ, J, μ) such that the following holds: Suppose that $r \geq \kappa_{\delta}$ and that $\mathfrak{d} = (A, \psi = (\alpha, \beta))$ is an instanton solution to (62) as defined with a data set from this neighborhood of (λ, J, μ) (and any R in Case 2). Assume that $A_{\mathfrak{d}} \leq r^2$ or $i_{\mathfrak{d}} \geq -r^2$. Let $\mathbb{I} \subset \mathbb{R}$ be a connected set of length at least $4\delta^{-1} + 16$ such that $\underline{M}(\cdot) \leq \mathcal{K}$ on \mathbb{I} . Assume in addition that $I_0 \cap \mathbb{I}$ has distance at least $\frac{4}{3}\delta^{-1}$ from $\partial\mathbb{I}$. Let $I \subset \mathbb{I}$ denote the set of points with distance at least 7 from the boundary of \mathbb{I} . Then:*

- Each point in $s_*^{-1}(I)$ where $|\alpha| \leq 1 - \delta$ has distance less than $\kappa r^{-1/2}$ from $\alpha^{-1}(0)$.
- There exist:
 - (1) A positive integer $N \leq \kappa$ and a cover $I = \bigcup_{k=1}^N I_k$ where each I_k is a connected open set of length at least $2\delta^{-1}$, such that $I_k \cap I_{k'} = \emptyset$ when $|k - k'| > 1$. If $|k - k'| = 1$, then $I_k \cap I_{k'}$ has length between $\frac{1}{128}\delta^{-1}$ and $\frac{1}{64}\delta^{-1}$. Finally, each boundary point of each I_k has distance at least δ^{-1} from $I_0 \cap \mathbb{I}$.
 - (2) For each $k \in \{1, \dots, N\}$, a finite set \mathcal{C}_k of pairs (C, m) where m is a positive integer and C is a closed irreducible J -holomorphic subvariety in a neighborhood of $s_*^{-1}(I_k)$. The subvarieties C for different elements of \mathcal{C}_k are distinct.

These are such that:

$$(a) \quad \sup_{\substack{z \in \bigcup_{(C,m) \in \mathcal{C}_k} C \cap s_*^{-1}(I_k)}} \text{dist}(z, \alpha^{-1}(0)) + \sup_{z \in \alpha^{-1}(0) \cap s_*^{-1}(I_k)} \text{dist}\left(z, \bigcup_{(C,m) \in \mathcal{C}_k} C\right) < \delta.$$

(b) Let $I' \subset I_k$ be an interval of length 1 and let ν be a 2-form on $s_*^{-1}(I')$ with $|\nu| \leq 1$ and $|\nabla \nu| \leq \delta^{-1}$. Then

$$\left| \frac{i}{2\pi} \int_{s_*^{-1} I'} \nu \wedge F_{\hat{A}} - \sum_{(C,m) \in \mathcal{C}_k} m \int_{C \cap s_*^{-1}(I')} \nu \right| \leq \delta.$$

(c)
$$\sum_{(C,m) \in \mathcal{C}_k} m \int_{C \cap (\bar{X}_* - s_*^{-1}(I_0))} d\lambda_Y < \kappa.$$

• Suppose that \mathbb{I} is unbounded from above. Fix $E_+ \leq \mathcal{K}$, and require nondegenerate Reeb orbits only for length at most $\frac{1}{2\pi} E_+$. Assume also that $E(c_+) \leq E_+$. Then the preceding conclusions hold with κ depending on \mathcal{K} and E_+ , and with κ_δ depending only on \mathcal{K} , E_+ and δ . Moreover, if $\mathbb{I} = \mathbb{R}$ then $E(c_-) \leq E_+ + \delta$.

Proof The first bullet follows from the first bullet in Proposition 7.8. The proof of the rest of Proposition 7.9 is a slight modification of the proof of [24, Proposition 4.5]. The latter proof has five parts. The first two parts establish [24, Lemma 4.6, Corollary 4.7, Lemma 4.8], which are applicable here with the contact manifold M in [24] replaced by Y_\pm or Y_0 here. The third part of the proof of [24, Proposition 4.5] has the following analogue here:

Lemma 7.10 (cf [24, Lemma 4.9]) *Given $\mathcal{K} \geq 1$, suppose that each Reeb orbit with length at most $\mathcal{K}/2\pi$ of λ_\pm , and λ_0 in Case 2, is nondegenerate. Given also $\varepsilon > 0$, there exists $\kappa \geq 1$ and a neighborhood of the given data set (λ, J, μ) such that the following holds: Suppose that $r \geq \kappa$ and that $\mathfrak{d} = (A, \psi = (\alpha, \beta))$ is an instanton solution to (62) as defined with a data set in this neighborhood of (λ, J, μ) (and any R in Case 2), with $A_\mathfrak{d} \leq r^2$ or $i_\mathfrak{d} \geq -r^2$. Let $\mathbb{I} \subset \mathbb{R} \setminus I_0$ denote a connected subset of length at least 16 such that $\underline{M}(\cdot) \leq \mathcal{K}$ on \mathbb{I} . Let \mathcal{I} denote the set of integers k such that $[k, k + 1] \in \mathbb{I}$ and*

$$\frac{i}{2\pi} \int_{s_*^{-1}[k, k+1]} d\lambda_Y \wedge F_{\hat{A}} \geq \varepsilon.$$

Let I' be a component of $\mathbb{I} \setminus \bigcup_{k \in \mathcal{I}} [k, k + 1]$. Then

$$\frac{i}{2\pi} \int_{s_*^{-1}(I')} d\lambda_Y \wedge F_{\hat{A}} \geq -\varepsilon^2.$$

Proof Copy the proof of [24, Lemma 4.9]. Wherever the latter proof invokes lemmas from [24, Section 3], replace these as indicated in Section 7.3 above. □

The fourth part of the proof of [24, Proposition 4.5] carries over with only notational changes to deduce the second bullet in Proposition 7.9 from Proposition 7.8.

The fifth part of the proof of [24, Proposition 4.5] carries over to prove the third bullet in Proposition 7.9, with the following modification: The key step is to show that given $\varepsilon_0 > 0$, if r is sufficiently large, then if $k_- < k_+$ are integers in \mathbb{I} with $k_+ - k_- < \varepsilon_0^{-1}$, then

$$(118) \quad \frac{i}{2\pi} \int_{s_*^{-1}(k_+, k_+ + 1)} ds \wedge \lambda_Y \wedge F_{\hat{A}} - \frac{i}{2\pi} \int_{s_*^{-1}(k_-, k_- + 1)} ds \wedge \lambda_Y \wedge F_{\hat{A}} > -\varepsilon_0.$$

If the intervals $(k_-, k_- + 1)$ and $(k_+, k_+ + 1)$ are in the same component of $\mathbb{R} \setminus I_0$, then the inequality (118) follows from an integration by parts in [24, Section 4d, Part 5]. So to complete the proof, we just need to prove (118) when $k_- + 1 = k_+ \in I_0$. To simplify notation, restrict to Case 1, so that $k_- + 1 = k_+ = 0$. The aforementioned integration by parts can be used to show that the integrals on the left hand side of (118) satisfy

$$\left| \frac{i}{2\pi} \int_{s_*^{-1}(k_{\pm}, k_{\pm} + 1)} ds \wedge \lambda_Y \wedge F_{\hat{A}} - \frac{i}{2\pi} \int_{\{0\} \times Y_{\pm}} \lambda_{\pm} \wedge F_{\hat{A}} \right| < \frac{\varepsilon_0}{3}$$

if r is sufficiently large. So to prove (118), it is enough to show that

$$\frac{i}{2\pi} \int_X d\lambda \wedge F_{\hat{A}} > -\frac{\varepsilon_0}{3}$$

if r is sufficiently large. This last inequality follows from the a priori estimates in Lemma 7.3 and [24, Lemma 3.8]. □

7.5 Proof of Proposition 7.1

We now carry over material from [24, Section 5] to our setting and prove Proposition 7.1. The proof of Proposition 7.1 uses the following proposition, which is similar to the $\mathbb{I} = \mathbb{R}$ case of Proposition 7.9, but with the assumption on \underline{M} replaced by an assumption on $E(c_+)$.

Proposition 7.11 (cf [24, Proposition 5.1]) *Fix $\mathcal{K} \geq 1$ and $E_+ \leq \mathcal{K}$. Assume all Reeb orbits of λ_{\pm} , and λ_0 in Case 2, of length at most $\frac{1}{2\pi}E_+$ are nondegenerate. Then there exists $\kappa \geq 1$, and given $\delta > 0$ there exist $\kappa_{\delta} \geq 1$ and a neighborhood of the given data set (λ, J, μ) such that the following holds: Suppose that $r \geq \kappa_{\delta}$ and that $\mathfrak{v} = (A, \psi = (\alpha, \beta))$ is an instanton solution to (62) as defined with a data set in this neighborhood of (λ, J, μ) (and any R in Case 2). Assume that $A_{\mathfrak{v}} \leq \mathcal{K}r$ or $i_{\mathfrak{v}} \geq -\mathcal{K}r$. Assume also that $E(c_+) \leq E_+$. Then:*

- $E(c_-) \leq E_+ + \delta$.
- The first two bullets of Proposition 7.9 hold with $I = \mathbb{R}$.

Proof This follows from Proposition 7.9 if we can show that given an instanton solution \mathfrak{d} to (62) with $A_{\mathfrak{d}} \leq \mathcal{K}r$ or $i_{\mathfrak{d}} \geq -\mathcal{K}r$ and $E(c_+) < E_+$, there exists an r -independent upper bound on $\underline{M}(\cdot)$ when r is sufficiently large. We now explain how to obtain such a bound on \underline{M} by modifying the arguments in [24, Section 5], which obtain an analogous bound on \underline{M} when $\overline{X}_* = \mathbb{R} \times M$.

First note that our assumptions imply that

$$(119) \quad A_{\mathfrak{d}} < c_0(\mathcal{K} + 1)r.$$

To prove (119), we can assume that $i_{\mathfrak{d}} > -\mathcal{K}r$, and this implies that

$$a(c_-) - a(c_+) < a^f(c_-) - a^f(c_+) + c_0\mathcal{K}r.$$

By [21, Proposition 4.11], the assumption $E(c_+) < \mathcal{K}$ implies that $a^f(c_+) > -c_0\mathcal{K}r$; see (134) below. Meanwhile, an almost verbatim version of an argument from [24, Section 5d] proves that $a^f(c_-) \leq c_0$ when $i_{\mathfrak{d}} > -\mathcal{K}r$. The inequality (119) follows.

We now explain how to obtain a bound on $\underline{M}(s)$ when $s \geq R + 2$. If the interval $[s, s + 1]$ does not intersect I_0 , define

$$\underline{E}(s) := i \int_{s_*^{-1}[s, s+1]} ds \wedge \lambda_Y \wedge F_A.$$

When $s \geq R + 2$, we will first obtain a bound $\underline{E}(s)$, and then use this to bound $\underline{M}(s)$.

To obtain bounds on \underline{E} , we need three inequalities. To state these, recall the reference connection A_E from (114) and the 2-form ν_E defined from its curvature in (115). Let $u_{\pm}: Y_{\pm} \rightarrow S^1$, and also $u_0: Y_0 \rightarrow S^1$ in Case 2, be gauge transformations. If $s > R$, write the connection component of $u_+ \cdot \mathfrak{d}$ as $A_E + \hat{a}_+$, and define

$$p_+(s) := -i \int_{s_*^{-1}(s)} \hat{a}_+ \wedge \nu_E.$$

Here and below, our convention is that $R = 0$ in Case 1. Define $p_-(s)$ analogously if $s < -R$, and define p_0 analogously in Case 2 if $-R < s < R$.

The first inequality asserts that if $s > R + 3$, then

$$(120) \quad r \underline{E}(s - 1) \leq -c_0 a(u_+ c_+) + c_0(r + a(\mathfrak{d})|_{s_* = R+3} - a(c_+)) + c_0 r^{2/3} \sup_{x \geq s} |\underline{E}(x)|^{4/3} + c_0 \sup_{[s, s+1]} |p_+|.$$

The second inequality asserts that if $s < s' < -R - 3$, then for suitable $s_- \in [s', s' + 1]$ we have

$$(121) \quad r\underline{E}(s - 1) \leq -c_0\alpha(u_- \partial|_{s_* = s_-}) + c_0(r + \alpha(c_-) - \alpha(\partial|_{s_* = s_-})) \\ + c_0 r^{2/3} \sup_{x \in [s, s_-]} |\underline{E}(x)|^{4/3} + c_0 \sup_{[s, s+1]} |p_-|.$$

Here s_- is “suitable” if $O(s_-) \leq \int_{s'}^{s'+1} O(s_*) ds_*$, where $O(s)$ is defined in (126) below. The third inequality asserts that in Case 2, if $-R + 3 < s < s' < R - 3$, and if $s_0 \in [s', s' + 1]$ satisfies $O(s_0) \leq \int_{s'}^{s'+1} O(s_*) ds_*$, then

$$(122) \quad r\underline{E}(s - 1) \leq -c_0\alpha(u_0 \partial|_{s_* = s_0}) + c_0(r + \alpha(\partial|_{s_* = -R+3}) - \alpha(\partial|_{s_* = s_0})) \\ + c_0 r^{2/3} \sup_{x \in [s, s_0]} |\underline{E}(x)|^{4/3} + c_0 \sup_{[s, s+1]} |p_0|.$$

The inequalities (120)–(122) are all proved analogously to [24, (5-18)].

To exploit the inequalities (120)–(122), we need appropriate bounds on the terms that do not involve $\underline{E}(\cdot)$. We first observe that the action differences in (120)–(122) are bounded by

$$(123) \quad \alpha(\partial|_{s_* = R+3}) - \alpha(c_+) \leq c_0(\mathcal{K} + 1)r, \\ \alpha(c_-) - \alpha(\partial|_{s_* = s_-}) \leq c_0(\mathcal{K} + 1)r, \\ \alpha(\partial|_{s_* = -R+3}) - \alpha(\partial|_{s_* = s_0}) \leq c_0(\mathcal{K} + 1)r.$$

To prove (123), first use Lemma 7.6(a), (b) to see that each action difference is bounded from above by $A_0 + c_0 r$, and then use (119).

To bound the remaining terms in (120), the discussion in [24, Section 5d] finds a gauge transformation $u_+ : Y_+ \rightarrow S^1$ such that $\alpha(u_+ c_+) > -c_0 E_+ r$ and $\lim_{s \rightarrow \infty} |p_+| \leq c_0$. The first of these conditions allows us to replace (120) by

$$(124) \quad \underline{E}(s - 1) \leq c_0(1 + \mathcal{K}) + c_0 r^{-1/3} \sup_{x \geq s} |\underline{E}(x)|^{4/3} + c_0 r^{-1} \sup_{[s, s+1]} |p_+|$$

for $s \geq R + 3$. The arguments in [24, Section 5d] can be applied verbatim using (124) to give the bound $\underline{E}(s) \leq c_0 \mathcal{K}$ for $s \geq R + 2$. The arguments in [24, Section 5d] also explain why this last bound implies that $\underline{M}(s) \leq c_0 \mathcal{K}$ for $s \geq R + 2$.

It remains to bound $\underline{M}(s)$ for $s \leq R + 2$. Let $t_* \in \{-R, R\}$ and suppose that \underline{E} (where defined) and \underline{M} have been bounded above by c_0 on $[t_* + 2, \infty)$. Let $t_{**} = -\infty$ if $R < 10$ or $t_* = -R$, and let $t_{**} = -R + 2$ otherwise. We now explain how to extend a bound of this sort on \underline{M} and \underline{E} to the interval (t_{**}, ∞) , in two steps. Applying this procedure once if $R < 10$, and twice if $R \geq 10$, will give the desired bound on $\underline{M}(s)$ for all $s \in \mathbb{R}$.

Step 1 This step bounds \underline{E} (where defined) and \underline{M} on $[t_* - 100, t_* + 2]$.

For $s \in \mathbb{R} \setminus I_0$, define

$$E(s) := i \int_{s_*^{-1}(s)} \lambda_Y \wedge F_A.$$

Recall that $d\tilde{\lambda} = \sigma\hat{\omega}$ where $\sigma: \bar{X}_R \rightarrow \mathbb{R}$ agrees with $2e^{2s_*}$ on $s_*^{-1}(\mathbb{R} \setminus I_0)$. Now use Stokes' Theorem and (62) to see that

$$\begin{aligned} (125) \quad E(s) &= ie^{-2s} \int_{s_*^{-1}(s)} \tilde{\lambda} \wedge F_A = ie^{-2s} \int_{s_*^{-1}(-\infty, s]} \sigma\hat{\omega} \wedge F_A \\ &= re^{-2s} \int_{s_*^{-1}(-\infty, s]} \sigma(1 - |\alpha|^2 + |\beta|^2) + O(1). \end{aligned}$$

Integrating this equation over $s \in [t_* + 2, t_* + 3]$ (or a slight upward translation of this interval as needed to avoid I_0) and using the a priori bounds in Lemma 7.3 shows that the bound on $\underline{E}(s)$ for $s \geq t_* + 2$ gives rise to a bound on \underline{M} on $[t_* - 100, t_* + 2]$.

Similarly to (125), if $s_- < s_+$ are in $\mathbb{R} \setminus I_0$ then

$$E(s_+) - e^{-2(s_+ - s_-)}E(s_-) = re^{-2s_+} \int_{s_*^{-1}[s_-, s_+]} \sigma(1 - |\alpha|^2 + |\beta|^2) + O(1).$$

Using this equation and Lemma 7.3 shows that the bound on $\underline{E}(s)$ for $s \geq t_* + 2$ gives rise to a bound on \underline{E} (where defined) on $[t_* - 100, t_* + 2]$.

Step 2 We now extend the bounds on \underline{E} and \underline{M} over $[t_{**}, t_* - 100]$. We assume below that $t_{**} < t_* - 100$.

For $s \in \mathbb{R} \setminus I_0$ define

$$(126) \quad O(s) := \int_{s_*^{-1}(s)} (|\mathcal{B}_{(A, \psi)}|^2 + r|D_A(s)\psi|^2),$$

where $\mathcal{B}_{(A, \psi)}$ was defined in (103). Also, define $\underline{O}(s) = \int_s^{s+1} O(s_*) ds_*$ when $[s, s + 1]$ does not intersect I_0 . Write $Y = Y_-$ when $t_{**} = -\infty$ and $Y = Y_0$ when t_{**} is finite. There exists $s_Y \in [t_* - 100, t_* - 99]$ such that $O(s_Y) \leq \underline{O}(t_* - 100)$. Then

$$(127) \quad O(s_Y) \leq \underline{O}(t_* - 100) \leq 2A_0 + c_0r \leq c_0(\mathcal{K} + 1)r$$

by Lemma 7.6(d) and the inequality (119).

Next, note that there is a map $u_Y: Y \rightarrow S^1$ such that the connection component of $u_Y \cdot \partial|_{s_Y}$ can be written as $A_E + \hat{a}_Y$ where \hat{a}_Y is a co-closed, i -valued 1-form on Y whose L^2 orthogonal projection to the space of harmonic 1-forms is bounded by c_0 . Combining this with (121) or (122) as appropriate with $s' = t_* - 100$, and using the

bound (127) on $O(s_Y)$ and the bound on $\underline{E}(s)$ for $s \in [t_* - 100, t_* - 2]$, the arguments leading to [24, Equation (5.14)] can be used to obtain a lower bound

$$(128) \quad \alpha(A_E + \widehat{a}_Y) \geq -c_0 r.$$

To continue, extend the map u_Y to all of $[t_{**}, t_* - 99]$ to be independent of s_* , and replace \mathfrak{d} on this portion of \overline{X}_* by $u_Y \cdot \mathfrak{d}$. It follows from (123), (128), and the relevant inequality (121) or (122) with s_- or s_0 set equal to s_Y , that for $s \in [t_{**}, t_* - 100]$ we have

$$(129) \quad \underline{E}(s - 1) \leq c_0(1 + \mathcal{K}) + c_0 r^{-1/3} \sup_{x \in [s, s_Y]} |\underline{E}(x)|^{1/3} + c_0 r^{-1} \sup_{[s, s+1]} |\mathfrak{p}_Y|.$$

Here \mathfrak{p}_Y denotes \mathfrak{p}_- or \mathfrak{p}_0 as appropriate. Moreover, we have

$$(130) \quad |\mathfrak{p}_Y| \leq c_0 \mathcal{K} \quad \text{on } [t_* - 100, t_* - 99].$$

To see why (130) is true, note that by our choice of u_Y we have $|\mathfrak{p}_Y(s_Y)| \leq c_0$. Meanwhile, [24, Lemma 3.9] bounds the derivative of the function $s \mapsto |\mathfrak{p}_Y(s)|$ by $c_0(1 + |M(s)|)$. Integrating this derivative bound and applying the conclusions from Step 1 gives (130).

Granted (129) and (130), arguments from [24, Section 5d] can be used in an almost verbatim fashion to bound \underline{E} on $[t_{**}, t_* - 100]$ by $c_0 \mathcal{K}$. One just needs to replace all references to the $s \rightarrow \infty$ limit of \mathfrak{d} by $\mathfrak{d}_{s_* = t_* - 100}$. As noted previously, arguments from [24, Section 5b] can be used with this bound on \underline{E} to bound \underline{M} by $c_0 \mathcal{K}$ on $[t_{**}, t_* - 100]$. □

Proof of Proposition 7.1 The first two bullets of Proposition 7.1 follow immediately from Proposition 7.11. The third bullet of Proposition 7.1 is deduced from Proposition 7.11 in the same way that the third bullet of [24, Proposition 5.5] is deduced from Proposition 5.1 in [24, Section 5e]. □

7.6 Proof of Propositions 5.2 and 5.4

Proof of Proposition 5.2 (a) We consider Case 1 of the geometric setup in Section 7.1. If the perturbations \mathfrak{p}_\pm and \mathfrak{p} are zero, then assertion (a) follows immediately from Case 1 of Proposition 7.1. For the case when the perturbations \mathfrak{p}_\pm and \mathfrak{p} are not zero, the proof has two steps.

Step 1 We claim that if $r > c_0$ and if the \mathcal{P} -norm of \mathfrak{p} is bounded by c_0^{-1} , then an instanton \mathfrak{d} as in (a) satisfies

$$(131) \quad \alpha(c_-) - \alpha(c_+) < (c_0 + 2\pi L)r.$$

Here and below, $\alpha(\mathfrak{c}_\pm)$ denotes the sum of the action functional (97) for Y_\pm and the abstract perturbation \mathfrak{p}_\pm . To prove (131), first note that the same integration by parts that proves Lemma 7.6(a), (b) implies that

$$(132) \quad \alpha(\mathfrak{c}_-) > \alpha(\mathfrak{c}_+) - c_0 r$$

if the \mathcal{P} -norm of \mathfrak{p} is bounded by c_0^{-1} . (See also the remark after [13, Proposition 24.6.4].) Since \mathfrak{d} has index 0, it follows from (101), (102) and (132) that

$$(133) \quad \alpha^f(\mathfrak{c}_-) > \alpha^f(\mathfrak{c}_+) - c_0 r.$$

Here f denotes the spectral flow function on Y_\pm defined using \mathfrak{p}_\pm . Meanwhile, by [21, Proposition 4.11] we have

$$(134) \quad \alpha^f(\mathfrak{c}_+) = -\frac{1}{2} r E(\mathfrak{c}_+) (1 + o(1)).$$

By this and (133) we have $\alpha^f(\mathfrak{c}_-) > -(c_0 + 2\pi L)r$. Consequently [21, Proposition 4.11] can be invoked a second time to give

$$(135) \quad \alpha^f(\mathfrak{c}_-) = -\frac{1}{2} r E(\mathfrak{c}_-) (1 + o(1)).$$

On the other hand, [24, Lemma 2.3] implies that $E(\mathfrak{c}_-) > -c_0$. This together with (135) implies that $\alpha^f(\mathfrak{c}_-) < c_0 r$. Since \mathfrak{d} has index 0, it follows from this last inequality and (134) that (131) holds.

Step 2 Now let r be large, and assume to get a contradiction that the conclusion of Proposition 5.2(a) is false. Then there exist data (J, μ) and a sequence of perturbations $\{\mathfrak{p}_k\}_{k=1,2,\dots}$ with $\lim_{k \rightarrow \infty} \mathfrak{p}_k = 0$, for which there is an instanton \mathfrak{d}_k with index 0 and $E(\mathfrak{c}_{k+}) < 2\pi L$ such that assertion (i) or (ii) in Proposition 5.2(a) fails. Here $\mathfrak{c}_{k\pm}$ denotes the $s \rightarrow \pm\infty$ limit of \mathfrak{c}_k . By passing to a subsequence we may assume that \mathfrak{c}_{k+} does not depend on k , so that we can denote it by \mathfrak{c}_+ .

Now (131) applies to each \mathfrak{d}_k to give

$$(136) \quad \alpha(\mathfrak{c}_{k-}) - \alpha(\mathfrak{c}_+) < (c_0 + 2\pi L)r.$$

Consequently, [13, Proposition 24.6.4] implies that the sequence of instantons $\{\mathfrak{d}_k\}_{k=1,2,\dots}$ has a subsequence that converges in the sense of [13, Section 26] to a broken trajectory, from \mathfrak{c}_+ to some generator \mathfrak{c}_- , that is defined using the equations (62) on \bar{X} and (31) on $\mathbb{R} \times Y_\pm$, without abstract perturbations. In particular, we can pass to a further subsequence so that $\mathfrak{c}_{k-} = \mathfrak{c}_-$ for all k . Let $\{\mathfrak{d}^n\}_{n=1,\dots,N}$ denote the ordered set of instantons that comprise the limiting broken trajectory. Let \mathfrak{c}_\pm^n denote the $s \rightarrow \pm\infty$ limit of \mathfrak{d}^n . These limits satisfy $\mathfrak{c}_-^1 = \mathfrak{c}_-$, $\mathfrak{c}_+^N = \mathfrak{c}_+$ and $\mathfrak{c}_+^n = \mathfrak{c}_-^{n+1}$ for $1 \leq n < N$.

There is a unique n_0 such that \mathfrak{d}^{n_0} is an instanton on \overline{X} . By Lemma 7.6(a), (b), we have $\alpha(\mathfrak{c}_-^{n_0}) > \alpha(\mathfrak{c}_+^{n_0}) - c_0 r$, and by [24, Lemma 3.4] we have $\alpha(\mathfrak{c}_-^n) > \alpha(\mathfrak{c}_+^n)$ for all $n \neq n_0$. These inequalities together with (136) imply that $\alpha(\mathfrak{c}_-^n) - \alpha(\mathfrak{c}_+^n) \leq c_0 r$ for each n . Consequently, if r is sufficiently large, then Proposition 7.1 applies to \mathfrak{d}^{n_0} , and [24, Proposition 5.5] applies to \mathfrak{d}^n for each $n \neq n_0$, to produce a broken J -holomorphic curve. These propositions (or the existence of these broken J -holomorphic curves) also imply that $E(\mathfrak{c}_-) < 2\pi L$ if r is sufficiently large. The concatenation of these N broken J -holomorphic curves is a broken J -holomorphic curve from Θ_+ to Θ_- , where Θ_\pm is determined by \mathfrak{c}_\pm via Proposition 3.1(a). It follows that if r is sufficiently large, then assertions (i) and (ii) in Proposition 5.2(a) are true for each \mathfrak{d}_k . This is the desired contradiction.

(b) This is essentially the same as the proof of (a), the only difference being that in Step 2, one now considers a sequence $\{\mathfrak{d}_k\}$ of instantons that solves the perturbed equations (62) for the data corresponding to some $t_k \in [0, 1]$. One then passes to a subsequence such that $\lim_{k \rightarrow \infty} t_k = t_*$. The arguments in [13, Section 24] can be used to show that the sequence of instantons \mathfrak{d}_k has a subsequence that converges to a broken trajectory for the data corresponding to $t = t_*$. Now the constants in Lemma 7.6 and Proposition 7.1, because they are stable, can be chosen to work for the data corresponding to all $t \in [0, 1]$. Thus if r is sufficiently large (independently of the value of t_*), then the rest of the proof of (a) can be repeated verbatim to prove part (b). □

Proof of Proposition 5.4 We now consider Case 2 of the geometric setup in Section 7.1. Recall the 1-form $\tilde{\lambda}_R$ on \overline{X}_R defined in (100). Define a 1-form λ' on X to agree with $\tilde{\lambda}_0$ on $[-\varepsilon, \varepsilon] \times Y_0$, where ε was fixed in Section 7.1, and to agree with λ on the rest of X . Note that the exact symplectic cobordism (X, λ') from (Y_+, λ_+) to (Y_-, λ_-) is strongly homotopic to (X, λ) . So by Corollary 5.3(c), if r is sufficiently large then $\widehat{HM}_L^*(X, \lambda) = \widehat{HM}_L^*(X, \lambda')$. Thus to prove Proposition 5.4, it is enough to show that if r is sufficiently large then

$$(137) \quad \widehat{HM}_L^*(X, \lambda') = \widehat{HM}_L^*(X^-, \lambda^-) \circ \widehat{HM}_L^*(X^+, \lambda^+).$$

To prove (137), we fix r large and vary R in Case 2. Let $\mathfrak{p}_\pm, \mathfrak{p}_0$ be abstract perturbations as needed to define the respective Seiberg–Witten Floer chain complexes on Y_\pm, Y_0 . Extend these to abstract perturbations \mathfrak{p}^\pm on \overline{X}^\pm as needed to define chain maps that induce $\widehat{HM}_L^*(X^\pm, \lambda^\pm)$; denote these chain maps by Φ_\pm . The perturbations \mathfrak{p}^\pm , with suitable cutoff functions, then determine an abstract perturbation \mathfrak{p}_R on \overline{X}_R , as explained in [13, Section 11, Section 24.1] (see [13, Equation (24.1)]). Let

$$\Phi_R: \widehat{CM}_L^*(Y_+; \lambda_+, J_+, r) \longrightarrow \widehat{CM}_L^*(Y_-; \lambda_-, J_-, r)$$

denote the chain map defined by counting index 0 instantons on \overline{X}_R . (One may need to perturb \mathfrak{p}_R to obtain transversality, in which case the chain map will depend on this perturbation.) It follows as in the proof of [Proposition 5.2\(a\)](#) that if r is sufficiently large, then for any R , if the abstract perturbations are sufficiently small then Φ_R is defined.

When $R = 0$, the induced map on homology is

$$(\Phi_0)_* = \widehat{\text{HM}}_L^*(X, \lambda'),$$

because by construction $(\overline{X}_0, \tilde{\lambda}_0) = (\overline{X}, \tilde{\lambda})$. On the other hand, for $R > 0$ the manifold $(\overline{X}_R, \tilde{\lambda}_R)$ is obtained by gluing $(\overline{X}^-, \tilde{\lambda}^-)$ with the $s \geq R$ part of the positive end removed to $(\overline{X}^+, \tilde{\lambda}^+)$ with the $s \leq -R$ part of the negative end removed. It then follows from [\[13, Proposition 26.1.6\]](#) that there exists R_0 such that if $R \geq R_0$, then Φ_R is defined without any further perturbation of \mathfrak{p}_R , and

$$\Phi_R = \Phi_- \circ \Phi_+.$$

So to complete the proof of [\(137\)](#), it is enough to show that the chain maps

$$\Phi_0, \Phi_{R_0}: \widehat{\text{CM}}_L^*(Y_+; \lambda_+, J_+, r) \longrightarrow \widehat{\text{CM}}_L^*(Y_-; \lambda_-, J_-, r)$$

are chain homotopic. To construct the desired chain homotopy one counts index -1 instantons in the family $\{\overline{X}_R \mid R \in [0, R_0]\}$ with a generic small family of abstract perturbations. If r is sufficiently large, then this chain homotopy will be well defined as in the proof of [Proposition 5.2\(b\)](#). \square

References

- [1] **MF Atiyah, V K Patodi, IM Singer**, *Spectral asymmetry and Riemannian geometry, I*, Math. Proc. Cambridge Philos. Soc. 77 (1975) 43–69 [MR0397797](#)
- [2] **H Federer**, *Geometric measure theory*, Grundle. Math. Wissen. 153, Springer, New York (1969) [MR0257325](#)
- [3] **C Hummel**, *Gromov’s compactness theorem for pseudo-holomorphic curves*, Progress in Mathematics 151, Birkhäuser, Basel (1997) [MR1451624](#)
- [4] **M Hutchings**, *Lecture notes on embedded contact homology* [arXiv:1303.5789](#)
- [5] **M Hutchings**, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. 4 (2002) 313–361 [MR1941088](#)
- [6] **M Hutchings**, *The embedded contact homology index revisited*, from: “New perspectives and challenges in symplectic field theory”, (M Abreu, F Lalonde, L Polterovich, editors), CRM Proc. Lecture Notes 49, Amer. Math. Soc. (2009) 263–297 [MR2555941](#)

- [7] **M Hutchings**, *Embedded contact homology and its applications*, from: “Proc. ICM, Vol. II”, (R Bhatia, A Pal, G Rangarajan, V Srinivas, M Vanninathan, editors), Hindustan Book Agency, New Delhi (2010) 1022–1041 [MR2827830](#)
- [8] **M Hutchings**, *Quantitative embedded contact homology*, *J. Differential Geom.* 88 (2011) 231–266 [MR2838266](#)
- [9] **M Hutchings**, **M Sullivan**, *Rounding corners of polygons and the embedded contact homology of T^3* , *Geom. Topol.* 10 (2006) 169–266 [MR2207793](#)
- [10] **M Hutchings**, **C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders, I*, *J. Symplectic Geom.* 5 (2007) 43–137 [MR2371184](#)
- [11] **M Hutchings**, **C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders, II*, *J. Symplectic Geom.* 7 (2009) 29–133 [MR2491716](#)
- [12] **M Hutchings**, **C H Taubes**, *Proof of the Arnold chord conjecture in three dimensions, I*, *Math. Res. Lett.* 18 (2011) 295–313 [MR2784673](#)
- [13] **P Kronheimer**, **T Mrowka**, *Monopoles and three-manifolds*, New Mathematical Monographs 10, Cambridge Univ. Press (2007) [MR2388043](#)
- [14] **J Latschev**, **C Wendl**, *Algebraic torsion in contact manifolds*, *Geom. Funct. Anal.* 21 (2011) 1144–1195 [MR2846386](#)
- [15] **F Morgan**, *Geometric measure theory: A beginner’s guide*, 4th edition, Elsevier, Amsterdam (2009) [MR2455580](#)
- [16] **C B Morrey, Jr**, *Multiple integrals in the calculus of variations*, Grundle Math. Wissen. 130, Springer, New York (1966) [MR0202511](#)
- [17] **C H Taubes**, $Gr \Rightarrow SW$: *From pseudo-holomorphic curves to Seiberg–Witten solutions*, from: “Seiberg–Witten and Gromov invariants for symplectic 4–manifolds”, (R Wentworth, editor), First Int. Press Lect. Ser. 2, International Press (2000) 163–273
- [18] **C H Taubes**, $SW \Rightarrow Gr$: *From the Seiberg–Witten equations to pseudo-holomorphic curves*, from: “Seiberg–Witten and Gromov invariants for symplectic 4–manifolds”, (R Wentworth, editor), First Int. Press Lect. Ser. 2, International Press (2000) 1–97
- [19] **C H Taubes**, *The Seiberg–Witten equations and the Weinstein conjecture*, *Geom. Topol.* 11 (2007) 2117–2202 [MR2350473](#)
- [20] **C H Taubes**, *The Seiberg–Witten equations and the Weinstein conjecture, II: More closed integral curves of the Reeb vector field*, *Geom. Topol.* 13 (2009) 1337–1417 [MR2496048](#)
- [21] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, I*, *Geom. Topol.* 14 (2010) 2497–2581 [MR2746723](#)
- [22] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, II*, *Geom. Topol.* 14 (2010) 2583–2720 [MR2746724](#)

- [23] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, III*, *Geom. Topol.* 14 (2010) 2721–2817 [MR2746725](#)
- [24] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, IV*, *Geom. Topol.* 14 (2010) 2819–2960 [MR2746726](#)
- [25] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, V*, *Geom. Topol.* 14 (2010) 2961–3000 [MR2746727](#)
- [26] **J G Wolfson**, *Gromov’s compactness of pseudo-holomorphic curves and symplectic geometry*, *J. Differential Geom.* 28 (1988) 383–405 [MR965221](#)
- [27] **R Ye**, *Gromov’s compactness theorem for pseudo holomorphic curves*, *Trans. Amer. Math. Soc.* 342 (1994) 671–694 [MR1176088](#)

Mathematics Department, University of California, Berkeley
970 Evans Hall, Berkeley, CA 94720, USA

Department of Mathematics, Harvard University
Cambridge, MA 02138, USA

hutchings@math.berkeley.edu, chtaubes@math.harvard.edu

<http://math.berkeley.edu/~hutchings>

Proposed: Tom Mrowka

Received: 15 November 2011

Seconded: Peter Ozsváth, Leonid Polterovich

Accepted: 8 February 2013