Embedded contact homology and Seiberg–Witten Floer cohomology I

CLIFFORD HENRY TAUBES

This is the first of five papers that construct an isomorphism between the embedded contact homology and Seiberg–Witten Floer cohomology of a compact 3–manifold with a given contact 1–form. This paper describes what is involved in the construction.

57R17; 57R57

1 Introduction

The purpose of this article is to describe an isomorphism between the Seiberg–Witten Floer cohomology of a compact, oriented 3–manifold and the embedded contact homology as defined by a given contact 1–form on the 3–manifold. What follows momentarily is a very brief description of how these cohomology/homology groups are defined. A more detailed description is provided for both later in the article.

Consider first the Seiberg–Witten side of the story. Let M denote the 3–manifold in question. The Seiberg–Witten Floer homology/cohomology is defined with the choice of a $\mathrm{Spin}_{\mathbb{C}}$ structure on M. With a Riemannian metric chosen, the latter is an equivalence class of lifts of the oriented, orthonormal frame bundle to a principal $\mathrm{Spin}_{\mathbb{C}}$ bundle. Each $\mathrm{Spin}_{\mathbb{C}}$ structure has an associated cohomology class in $H^2(M;\mathbb{Z})$; this is its first Chern class. Let $p \in \{2, 4, \ldots\}$ denote the divisibility of this class. (It is always divisible by 2.) Each $\mathrm{Spin}_{\mathbb{C}}$ structure with nontorsion first Chern class defines $\mathbb{Z}/p\mathbb{Z}$ graded homology and cohomology groups. These are the associated Seiberg–Witten Floer homology and cohomology. These groups are always finitely generated. When the first Chern class is torsion, the associated Seiberg–Witten Floer homology and cohomology is $\mathbb{Z}$ graded. The Seiberg–Witten Floer homology in this case is finitely generated in each degree. As explained by Kronheimer and Mrowka [16], the degrees in which the Seiberg–Witten Floer homology is nonzero are bounded from above, but never from below. This group is designated by $\widehat{HM}_*$ [16]. There is a corresponding Seiberg–Witten Floer cohomology group as well; this is designated by $\mathcal{H}^{\mathrm{SW}}$, or when reference to the chain complex is relevant, by $\mathcal{H}(\mathcal{C}^{\mathrm{SW}})$. In all cases, the generators of

Published: 15 December 2010 DOI: 10.2140/gt.2010.14.2497

these homology and cohomology groups are the solutions to certain versions of the Seiberg–Witten equations on M; and the differentials are defined via a weighted count of certain sorts of solutions to the Seiberg–Witten equations on $\mathbb{R} \times M$. (The book [16] is taken here to be the reference bible for the Seiberg–Witten side of the story.)

Consider next the contact homology story. Embedded contact homology was invented by Michael Hutchings (see Hutchings and Sullivan [12] and Hutchings and Taubes [13]). The definition requires first the choice of a contact 1–form a on M that is compatible with the orientation. Thus, the form a is such that $a \wedge da$ is nowhere zero and orients the 3–manifold. A (suitably generic) contact form of this sort defines a version of embedded contact homology and its associated cohomology for each $\mathrm{Spin}_{\mathbb{C}}$ structure on M. These groups have $\mathbb{Z}/p\mathbb{Z}$ grading when the first Chern class is not torsion, and they are $\mathbb{Z}$ graded otherwise. The generators for these groups consist of finite sets where any given element is a pair that consists of a closed integral curve of the vector field that generates the kernel of da and a positive integer. Note that closed orbits with hyperbolic return map are paired only with the integer 1. The differential in each case is defined via a weighted count of certain embedded, pseudoholomorphic curves in the symplectization, $\mathbb{R} \times M$, of M.

The theorem that follows states formally what is said in the opening paragraph.

Theorem 1 *Let M denote a compact, oriented three dimensional manifold and let a denote a suitably generic contact 1–form on M that gives the chosen orientation. Fix a $\mathrm{Spin}_{\mathbb{C}}$ structure on M. Then there is an isomorphism between the associated embedded contact homology and the Seiberg–Witten Floer cohomology that reverses the sign of the relative $\mathbb{Z}/p\mathbb{Z}$ grading.*

Note that this theorem implies two conjectures by Michael Hutchings: The embedded contact homology does not depend on the contact 1–form; and the embedded contact homology is finitely generated in each degree.

There are circumstances where both the Seiberg–Witten Floer cohomology and embedded contact homology have additional structure. This additional structure is not discussed further until the fifth paper [29] in this series, except for the remark that the isomorphism that is described here is compatible with these additional structures.

Theorem 1 can be viewed as the 3–manifold analog of the equivalence proved by the author in [21; 22] between the Seiberg–Witten invariants of a compact symplectic manifold and certain of the Gromov invariants that are computed by counting its pseudoholomorphic curves. Theorem 1 can also be viewed as a generalization of the author's papers [24; 25], which use the existence of certain nontrivial Seiberg–Witten Floer homology classes to find closed orbits of the vector field that generates the kernel of da.

The proof of Theorem 1 occupies Section 4 of this article plus its three immediate sequels [26; 27; 28]. The proof uses many of the constructions and observations that are used in [22]. These parts of the argument are summarized by Theorems 4.2 and 4.3 to come. The subsequent papers in this series contain the proofs of Theorems 4.2 and 4.3. Ideas from [24; 25] also play a central role to the proof of Theorem 1. Theorems 4.4 and 4.5 contain most of the input to Theorem 1 from [24; 25]. These last two theorems are proved below in Section 4.h.

What follows is a table of contents for the remaining parts of this article.

Section 2 gives the definition of embedded contact homology. Also stated here is Proposition 2.5; this very useful proposition asserts that any given contact 1–form has a suitable deformation to one with properties that very much simplify the subsequent analysis.

Section 3 introduces the Seiberg–Witten equations and then the Seiberg–Witten Floer cohomology. It also describes the special versions of the Seiberg–Witten equations that can be defined with the help of a contact 1–form.

Section 4 proves Theorem 1 modulo two technical results, Theorems 4.2 and 4.3. The various parts of Theorems 4.2 and 4.3 are proved in the sequels [26; 27; 28].

Section 5 is meant to give a very rough picture of two key maps that are supplied by Theorems 4.2 and 4.3. This section also indicates how the proof of Theorem 1 would proceed without the approximation result from Proposition 2.5.

The Appendix contains the proof of Proposition 2.5.

Acknowledgements Before continuing, the author hereby acknowledges the immense debt owed to Michael Hutchings, Peter Kronheimer and Tom Mrowka for sharing their thoughts and knowledge about the subject matter in this paper. In particular, Michael Hutchings made many cogent suggestions for improving and clarifying an earlier version of what follows.

This work was supported in part by the National Science Foundation.

2 Embedded contact homology

The purpose of this section is to give the definition of embedded contact homology. As noted in the introduction, this homology theory was introduced by Michael Hutchings. Most of what follows here paraphrases parts of the accounts in [12] and [13].

2.a Reeb orbits

Use v in what follows to denote the vector field on M that generates the kernel of da and pairs with a so as to equal 1. It is traditional to call v the *Reeb* vector field. A *Reeb orbit* denotes here an embedded circle with tangent v, thus a closed integral curve of v. A Reeb orbit is implicitly oriented by v.

Let γ denote a Reeb orbit. The integral of the contact 1–form along γ is denoted by ℓ_γ. This integral, a positive number, is called the *symplectic action* of γ. The set of Reeb orbits whose symplectic action is bounded by any given positive number is a compact subset in $C^\infty(S^1; M)$.

Fix an almost complex structure, J, on the kernel of a so that $da(\cdot, J(\cdot))$ defines a Riemannian metric on the kernel of a. A Reeb orbit γ has a neighborhood that is parametrized by the product of S^1 and a disk $D \subset \mathbb{C}$ about the origin by an embedding $\varphi\colon S^1 \times D \to M$ which makes a, da, and the Reeb vector field v appear as

$$\begin{aligned}
&\frac{2\pi}{\ell_\gamma}\varphi^* a = (1 - 2\nu|z|^2 - \mu\bar{z}^2 - \bar{\mu}z^2)dt + \frac{i}{2}(z\,d\bar{z} - \bar{z}\,dz) + \cdots, \\
&\frac{2\pi}{\ell_\gamma}\varphi^* da = i\,dz \wedge d\bar{z} - 2(\nu z + \mu\bar{z})\,d\bar{z}\wedge dt - 2(\nu\bar{z} + \bar{\mu}z)\,dz\wedge dt + \cdots, \\
&\frac{\ell_\gamma}{2\pi}(\varphi^{-1})_* v = \frac{\partial}{\partial t} + 2i(\nu z + \mu\bar{z})\frac{\partial}{\partial z} - 2i(\nu\bar{z} + \bar{\mu}z)\frac{\partial}{\partial\bar{z}} + \cdots.
\end{aligned} \tag{2-1}$$

Here, ν and μ are respectively real and complex valued functions on S^1. The unwritten terms in the top equation are $\mathcal{O}(|z|^3)$ and those in the lower two equations are $\mathcal{O}(|z|^2)$. In (2-1) and in what follows, the circle S^1 is implicitly identified with $\mathbb{R}/(2\pi\mathbb{Z})$ and $t \in \mathbb{R}/(2\pi\mathbb{Z})$ is used to denote its affine coordinate. These coordinates are such that the vector field $\partial/\partial z$ at $z = 0$ pushes forward via φ so as to generate the $+i$ eigenspace of J on kernel(a).

It follows as a consequence of (2-1) that the integral curves of v appear in this coordinate chart as the graphs of maps from an interval in $\mathbb{R}$ to D that obey an equation of the form

$$\frac{i}{2}\frac{d}{dt}z + \nu z + \mu\bar{z} = \mathfrak{r}, \tag{2-2}$$

where $\mathfrak{r}$ is a smooth function of t and z with $|\mathfrak{r}| \le c_0|z|^2$ and $|d\mathfrak{r}| \le c_0|z|$.

The left hand side of (2-2) defines a first order, $\mathbb{R}$–linear symmetric operator on $C^\infty(\mathbb{R}; \mathbb{C})$, this is the operator that takes a function $t \to z(t)$ to

$$Lz = \frac{i}{2}\frac{d}{dt}z + \nu z + \mu\bar{z}. \tag{2-3}$$

Such an operator is defined given any pair $(\nu, \mu) \in C^\infty(S^1; \mathbb{R} \oplus \mathbb{C})$. When z is written in terms of real functions x and y as $z = x + iy$ and then any function in the kernel of (2-3) can be written as

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = U \begin{pmatrix} x(0) \\ y(0) \end{pmatrix}, \text{ where } U|_t \in \mathrm{SL}(2; \mathbb{R}) \text{ for each } t \in \mathbb{R}. \tag{2-4}$$

As t varies in $[0, 2\pi]$, the map $t \to U|_t$ defines a path in $\mathrm{SL}(2 : \mathbb{R})$ from the identity. (The matrix $U|_{2\pi}$ is the linearization of the Reeb flow on $\mathrm{kernel}(a)$ along the Reeb orbit.)

A pair of functions (ν, μ) is said to be *nondegenerate* when the corresponding matrix U has $\mathrm{trace}(U|_{2\pi}) \neq 2$. The pair is deemed to be *elliptic* when $|\mathrm{trace}(U|_{2\pi})| < 2$ and *hyperbolic* when $|\mathrm{trace}(U|_{2\pi})| > 2$. Note that when (ν, μ) is hyperbolic, then the k–th power of $U|_{2\pi}$ does not have eigenvalue 1 for any k. Such is the case because $U|_{2\pi}$ in this case has two real eigenvalues, one with absolute value greater than 1 and the other of the same sign with absolute value less than 1. When elliptic, the pair (ν, μ) is said to be *n–elliptic* when the k–th power of $U|_{2\pi}$ does not have eigenvalue 1 for all $k \leq n$. Note that a matrix in $\mathrm{SL}(2; \mathbb{R})$ whose trace has absolute value less than 2 has two complex eigenvalues; these are on the unit circle and one is the conjugate of the other. A Reeb orbit γ is said to be respectively *nondegenerate*, *hyperbolic*, or *n–elliptic* when such is the case for the functions (ν, μ) that come from (2-2). Note that the labeling of γ as either hyperbolic or n–elliptic is an intrinsic property of γ; it does not depend on the choice of φ or the almost complex structure on the kernel of a.

The notion of hyperbolic or n–elliptic can be viewed as a condition on the operator L in (2-3). In particular, L has trivial kernel on the space of maps from S^1 to $\mathbb{C}$ if and only if (ν, μ) is nondegenerate. If (ν, μ) is hyperbolic, then L has trivial kernel on the space of $2\pi k$–periodic maps from $\mathbb{R}$ to $\mathbb{C}$ for any positive integer k. Meanwhile, if (ν, μ) is n–elliptic, then L has trivial kernel on the space of $2\pi k$ periodic maps from $\mathbb{R}$ to $\mathbb{C}$ for all $k \in \{1, 2, \ldots, n\}$.

A complex valued function on $\mathbb{R}$ is said to be an eigenvector of L if L sends the function to a constant, real multiple of itself. The constant in question is called the eigenvalue. An eigenvector is said to be $2\pi k$–periodic for a given integer $k > 1$ if it is $2\pi k$–periodic but not $2\pi k'$ periodic for any positive integer $k' < k$. Two nontrivial eigenvectors can have the same eigenvalue only if they have the same periodicity and the same degree as a map from S^1 to $\mathbb{C} - \{0\}$. (A nontrivial eigenvector is nowhere zero.)

The definition of embedded contact homology requires the Reeb orbits to be nondegenerate, and the elliptic ones to be n–elliptic for all n. The following lemma asserts a well known fact that there exist such contact forms.

Lemma 2.1 *There exists a residual set of contact forms in $C^\infty(M;T^*M)$ whose associated Reeb orbits are either hyperbolic or else are n–elliptic for all positive integers n. In fact, given a positive integer n, and $L \geq 1$, there is a dense open subset of contact forms in $C^\infty(M;T^*M)$ with the following property: If γ is an associated Reeb orbit with $\ell_\gamma \leq L$, then γ is hyperbolic or n–elliptic.*

A nondegenerate Reeb orbit is isolated in the following sense: There is an open, concentric disk $D' \subset D$ such that there are no Reeb orbits in $\varphi(S^1 \times D')$ except γ that generate the homology of $\varphi(S^1 \times D')$. Hyperbolic Reeb orbits have a stronger isolation property: There are no Reeb orbits except γ in $\varphi(S^1 \times D')$. In the case when γ is n–elliptic, then the following is true: But for multiple covers of γ, there are no Reeb orbits that generate the class of $k\gamma$ in $H_1(\varphi(S^1 \times D');\mathbb{Z})$ for any $k \in \{1,\ldots,n\}$.

The residual set of contact forms given in Lemma 2.1 is denoted by $\mathcal{N}_M$.

2.b Pseudoholomorphic subvarieties

The manifold $\mathbb{R} \times M$ has a family of almost complex structures that it inherits from the contact geometry of M. These almost complex structures are characterized by the following properties: They are invariant with respect to translations along the $\mathbb{R}$ factor of $\mathbb{R} \times M$; they map the generator, $\partial/\partial s$, of these translations to v; and they preserve the kernel of a. Such an almost complex structure endows M with a metric that sets the Hodge star of da equal to $2a$ and gives v norm 1. Almost complex structures of this sort are said to be *compatible with a*. These are the only ones used in this article.

Let J now denote an a–compatible, almost complex structure. An irreducible, pseudoholomorphic subvariety in $\mathbb{R}\times M$ is defined to be a closed subset with the following two properties: First, the complement of a finite set of points is a connected, 2-dimensional submanifold whose tangent space is J–invariant. Second, the integral of da over this submanifold is finite. A pseudoholomorphic subvariety is defined to be a finite union of irreducible, pseudoholomorphic subvarieties. What follows describes some of the salient properties of irreducible, pseudoholomorphic subvarieties $\mathbb{R} \times M$. The basic story on such curves is presented in a series of seminal papers of Hofer [4; 5; 6; 7] and Hofer, Wysocki and Zehnder [9; 8; 10].

Agree to use s in what follows to denote the Euclidean coordinate on the $\mathbb{R}$ factor of $\mathbb{R} \times M$. The first point to note is that these subvarieties are well behaved where $|s|$ is large. To say more, suppose for the moment that the contact form comes from the residual subset that is described in Lemma 2.1. With an almost complex structure fixed, let Σ denote a given pseudoholomorphic subvariety. Then there exists $s_0 > 1$ such that the $|s| \geq s_0$ portion of Σ is a disjoint union of properly embedded cylinders to which the function s restricts without critical points. Each such cylinder is called an *end* of Σ. The ends on which s $\geq s_0$ are called the positive side ends, and those where s $\leq -s_0$ are called the negative side ends.

The subvariety Σ may contain irreducible components of the form $\mathbb{R}\times\gamma$ with γ a Reeb orbit. Such cylinders are the only $\mathbb{R}$–invariant pseudoholomorphic subvarieties where $\mathbb{R}$ is understood to act on $\mathbb{R}\times M$ as the constant translations along the $\mathbb{R}$ factor.

Let $\mathcal{E}$ denote an end of Σ that is not in an $\mathbb{R}$–invariant cylinder. Then there is a Reeb orbit, $\gamma = \gamma_{\mathcal{E}}$, and a positive integer $q_{\mathcal{E}}$ such that the following is true: Each constant s slice of $\mathcal{E}$, thus $\mathcal{E}|_s \subset M$, is a braid in the $S^1\times D$ tubular neighborhood of γ that projects as a $q_{\mathcal{E}}$–to-1 covering map to the central circle. Moreover, as $|s|\to\infty$, these braids converge pointwise to γ.

To say more about this, note that a tubular neighborhood map $\varphi\colon S^1\times D\to M$ for γ can be chosen so as to have the following additional properties: First, the vector field $\partial/\partial z$ along D pushes forward to define a type $(1,0)$ tangent vector along $\mathbb{R}\times\gamma$ in $\mathbb{R}\times M$ with length $2^{-1/2}$. Second, the $\mathbb{C}$–valued 1–form

$$dz - 2i(\nu z + \mu\bar{z})\,dt \tag{2-5}$$

differs from the φ–pull back of a 1–form in $T^{1,0}(\mathbb{R}\times M)$ by $\mathcal{O}(|z|^2)\,dt, \mathcal{O}(|z|^2)\,dz$ and $\mathcal{O}(|z|)\,d\bar{z}$. Note that it follows from the second equation of (2-1) that this form has length $(4\pi/\ell_\gamma)^{1/2}$ to $\mathcal{O}(|z|^2)$. Henceforth, all tubular neighborhood maps are assumed to be of this sort.

Now, suppose that $\mathcal{E}$ denotes an $s \ll -1$ end of Σ that is not part of an $\mathbb{R}$–invariant cylinder. Let $\gamma = \gamma_{\mathcal{E}}$ and $q_{\mathcal{E}}$ be as described above. The end $\mathcal{E}$ can be viewed using the tubular neighborhood map as a subvariety in $\mathbb{R}\times(S^1\times D)$, this the image of a map from $(-\infty,-s_0]\times\mathbb{R}/(2\pi q_{\mathcal{E}}\mathbb{Z})$ into $\mathbb{R}\times(S^1\times D)$ that sends any given (s,t) to the point $(s,t,z(s,t))$ where z is a certain $\mathbb{C}$–valued function. To say more about $z(\cdot)$, introduce $\lambda_{q_{\mathcal{E}}}$ to denote the least negative of the eigenvalues of the set of $2\pi q_{\mathcal{E}}$ periodic eigenvectors of (2-3). Next, denote by $\mathrm{div}_{\mathcal{E}}\subset\{1,2,\dots,q_{\mathcal{E}}\}$ the subset with the following two properties: First, any given $q'\in\mathrm{div}_{\mathcal{E}}$ is a divisor of q_E. Second, there is a $2\pi q'$ periodic eigenvector of (2-3) with eigenvalue $\lambda_{q'}$ such that $0>\lambda_{q'}\geq\lambda_{q_{\mathcal{E}}}$. The function $z(\cdot)$ is given in terms of this data as

$$z(s,t) = \sum_{q'\in\mathrm{div}_{\mathcal{E}}} (\varsigma_{q'}(t)+\mathfrak{r}_{q'})e^{-2\lambda_{q'}s}, \tag{2-6}$$

where $\varsigma_{q'}$ is a $2\pi q'$ periodic eigenvector of (2-3) with eigenvalue $\lambda_{q'}$; and where $\mathfrak{r}_{q'}$ is $2\pi q'$ periodic and its norms and that of its derivative are bounded by $e^{-\varepsilon|s|}$ with ε a positive constant. Note in this regard that $\varsigma_{q'}=0$ is allowed, but if $\varsigma_{q_{\mathcal{E}}}=0$, then $\mathfrak{r}_{q_{\mathcal{E}}} \neq 0$. A positive end of Σ is described by (2-6) but with $\lambda_{q_{\mathcal{E}}} > 0$ the least positive eigenvalue of the set of $2\pi q_{\mathcal{E}}$ eigenvectors; and with the eigenvalue condition on membership in $\mathrm{div}_{\mathcal{E}}$ now reading $\lambda_{q_{\mathcal{E}}} \geq \lambda_{q'} > 0$. Equation (2-6) can be derived from what is done, for example, in Siefring [19].

The set, $\mathcal{M}$, of irreducible, pseudoholomorphic subvarieties has nice properties also. To say more, endow $\mathcal{M}$ with the topology whereby the neighborhoods of a given element $\Sigma \in \mathcal{M}$ are generated by sets of the following sort: A given basis set is labeled by $\varepsilon > 0$ and it consists of the subvarieties $\Sigma' \subset \mathcal{M}$ with the following two properties:

- $\sup_{z\in\Sigma} \mathrm{dist}(z, \Sigma') + \sup_{z'\in\Sigma'} \mathrm{dist}(\Sigma, z') < \varepsilon$.

(2-7)
- If ϖ is a compactly supported 2–form on $\mathbb{R}\times M$, then $\left|\int_\Sigma \varpi - \int_{\Sigma'} \varpi\right| \leq \varepsilon \sup_{\mathbb{R}\times M} |\varpi|$.

What follows is the basic structure theorem for $\mathcal{M}$.

Lemma 2.2 *Fix $\Sigma \subset \mathcal{M}$ with the following property: Let $\mathcal{E}$ be any given end of Σ and let γ denote the Reeb orbit that is approached by the $|s| \to \infty$ limit of the constant s slices of $\mathcal{E}$. Then γ is either hyperbolic or $q_{\mathcal{E}}$–elliptic. Assuming that Σ has these properties, there exists a Fredholm operator, $\mathcal{D}_\Sigma$, a ball $B \subset \mathrm{kernel}(\mathcal{D}_\Sigma)$, a smooth map $f\colon B \to \mathrm{cokernel}(\mathcal{D}_\Sigma)$ and a homeomorphism from $f^{-1}(0)$ to a neighborhood of Σ in $\mathcal{M}$. Here, $f(0) = 0$ and the homeomorphism sends 0 to Σ. Furthermore:*

- *Let $\mathcal{M}_{\mathrm{reg}} \subset \mathcal{M}$ denote the set that consists of those Σ with $\mathrm{cokernel}(\mathcal{D}_\Sigma) = 0$. This set $\mathcal{M}_{\mathrm{reg}}$ is open and it has the structure of a smooth manifold.*
- *Let $\Sigma \in \mathcal{M}_{\mathrm{reg}}$. Then the just described homeomorphism from $B \subset \mathrm{kernel}(\mathcal{D}_\Sigma)$ into $\mathcal{M}$ gives a smooth coordinate chart for a neighborhood of Σ.*
- *If the contact form comes from Lemma 2.1's residual set, then there is a residual set of compatible almost complex structures for which $\mathcal{M}_{\mathrm{reg}} = \mathcal{M}$.*

To say more about $\mathcal{D}_\Sigma$, remember that an irreducible, pseudoholomorphic subvariety, C, has a *model curve*; this a complex curve, C_0, together with an almost everywhere 1–1 pseudoholomorphic map $\phi\colon C_0 \to \mathbb{R}\times M$ whose image is C. Assuming that C has only immersion singularities, there is a well-defined pullback normal bundle over C_0. This is the bundle, $N \to C_0$, whose fiber at any given point is the normal 2–plane in $T(\mathbb{R}\times M)$ at the point's ϕ–image to an embedded disk in C. The composition of an exponential map with a section of a suitable disk sub-bundle of N defines a deformation of C in $\mathbb{R}\times M$. Here, an exponential map is a smooth map from a uniform radius disk subbundle in N to $\mathbb{R}\times M$ that restricts as ϕ to the zero section and has surjective differential along the zero section.

The almost complex structure gives N a complex bundle structure, and the induced metric from $\mathbb{R}\times M$ gives N a hermitian structure and thus the structure of a holomorphic line bundle. As such, there is an associated d-bar operator that maps sections of N to those of $N\otimes T^{1,0}C_0$. This d-bar operator enters the story in the following manner: A deformation of C that preserves to first order the pseudoholomorphic condition is the image via an exponential map of a section of N that is annihilated by an operator that can be viewed as the d-bar operator with an extra, $\mathbb{R}$–linear, zero–th order term. It can be identified as the operator that sends a section ζ to

$$\mathcal{D}_C\zeta = \bar{\partial}\zeta + \nu_C\zeta + \mu_C\bar{\zeta}, \tag{2-8}$$

where ν_C is a section of $T^{0,1}C_0$ and where μ_C is one of $N^2\otimes T^{0,1}C_0$. Note that the parametrization given in (2-6) for any given end of C induces a trivialization of N and TC_0 on such an end with the following property: When written using this trivialization, the pair (ν_C,μ_C) converges as $|s|\to\infty$ on the end to the pair (ν,μ) that appears in the associated version of (2-1).

The operator $\mathcal{D}_C$ defines a bounded, $\mathbb{R}$–linear Fredholm map from the Sobolev space $L^2_1(C_0;N)$ to $L^2(C_0;N\otimes T^{0,1}C_0)$ if the following is true: The constant s slices of each end of C limit as $|s|\to\infty$ as some integer q–fold cover of a Reeb orbit that is either hyperbolic or q–elliptic.

A more complicated version of this $\mathbb{R}$–linear operator defines $\mathcal{D}_C$ in the cases when C has nonimmersion singularities. The operator $\mathcal{D}_\Sigma$ is $\mathcal{D}_C$ if $\Sigma = C$. If each irreducible component of Σ is an immersed curve of the sort just described, then $\mathcal{D}_\Sigma$ is the direct sum of the corresponding operators with domain and range the direct sum of the corresponding Sobolev spaces.

2.c Embedded contact homology

The following definition of Hutchings' embedded contact homology is taken from Section 11 in [12]. To set the stage for the definition, note that each class in $H_1(M;\mathbb{Z})$ labels a version of this homology. This understood, fix a class Γ. Assume that the contact form a is from Lemma 2.1's residual set $\mathcal{N}_M$.

The chain complex The chain complex for Γ's version of embedded contact homology is the free $\mathbb{Z}$ module that is generated by equivalence classes of pairs $(\Theta,\mathfrak{o})$ where Θ and $\mathfrak{o}$ are as follows. First Θ is a finite set of pairs of the form (γ,m) where γ is a Reeb orbit and m is a positive integer subject to three constraints. First, no two pairs have the same Reeb orbit. Second, $m=1$ when γ is hyperbolic. Third, the formal sum $[\Theta]=\sum_{(\gamma,m)\in\Theta} m\gamma$ should define a closed cycle that generates the class Γ in

$H_1(M;\mathbb{Z})$. Note that the empty set $\Theta = \varnothing$ defines a generator in the case when Γ is the trivial class.

Meanwhile, $\mathfrak{o}$ is an ordering of the pairs in Θ whose Reeb orbit component is hyperbolic and whose version of (2-4) has matrix U with $\operatorname{trace}(U|_{2\pi}) > 2$. The equivalence relation identifies $(\Theta, \mathfrak{o})$ with $\pm(\Theta, \mathfrak{o}')$ where the $\pm$ factor is the image in $\{\pm 1\}$ of the permutation that takes $\mathfrak{o}$ to $\mathfrak{o}'$.

The free $\mathbb{Z}$ module so generated is denoted by $\mathcal{C}_{\mathrm{ech}}$ in what follows. Keep in mind that it depends on Γ. In most of what follows, the pair $(\Theta, \mathfrak{o})$ will be denoted as Θ with the presence of the ordering $\mathfrak{o}$ implicit.

The grading Let $K^{-1} \subset TM$ denote the 2–plane bundle given by the kernel of a. Orient this bundle using da, and let $-\hat{e}_K \in H^2(M;\mathbb{Z})$ denote its Euler class. Let $P(\Gamma) \in H^2(M;\mathbb{Z})$ denote the Poincaré dual of Γ and let p denote the divisibility of $-\hat{e}_K + 2P(\Gamma)$. The $\mathbb{Z}$ module $\mathcal{C}_{\mathrm{ech}}$ has a relative $\mathbb{Z}/p\mathbb{Z}$ grading whose definition is given in the five steps that follow.

Step 1 The path $U\colon [0, 2\pi] \to \mathrm{SL}(2;\mathbb{R})$ can be used to assign a *rotation number* to any hyperbolic or n–elliptic path (ν, μ). This rotation number is defined as follows: When the pair (ν, μ) is hyperbolic, there is a homotopy of U through a 1–parameter family of paths such that the $t = 2\pi$ element of each path on this family has $|\operatorname{trace}(\cdot)| \geq 2$, and so that the end member path is a path of pure rotations. As such, the end member path rotates $\mathbb{C}$ by a total of πk radians for some $k \in \mathbb{Z}$. This integer k is the rotation number for (ν, μ). In the case when (ν, μ) is n–elliptic, there is a similar homotopy of U, now through a family of paths such that each path in this family has its $t = 2\pi$ point conjugate to $U(2\pi)$. The end-member of this homotopy is again a path of rotations, this time rotating $\mathbb{C}$ by an angle $2\pi\mathrm{R}$. This angle R is the rotation number. Note in this regard that the n–elliptic condition means that $2k\mathrm{R}$ is not in $\mathbb{Z}$ when $k \in \{1, \dots, n\}$. The numbers k or R depend on φ, but not k mod (2) or R mod $(\mathbb{Z})$.

Suppose that the pair (ν, μ) is defined from a given Reeb orbit γ and coordinate map φ. Use $z_{\gamma,1}$ to denote the rotation number $k \in \mathbb{Z}$ when γ is hyperbolic. When γ is m–elliptic and $q \in \{1, \dots, m\}$, use $z_{\gamma,q}$ to denote 1 plus twice the greatest integer less than $q\mathrm{R}$.

Step 2 Let Θ_- and Θ_+ define generators in $\mathcal{C}_{\mathrm{ech}}$. Fix a tubular neighborhood embedding as described above for each Reeb orbit from the pairs that comprise Θ_- and Θ_+. As Θ_- and Θ_+ define the class Γ in the manner just described, there is a smooth, oriented, properly immersed surface $Z \subset \mathbb{R} \times M$ with transversal self-intersections that has the following properties: The $|s| \gg 1$ portion of this surface is a disjoint union of embedded cylinders on which s restricts as a function with no critical points. The cylinders that sit where $s \gg 1$ are distinguished in part by the elements in Θ_+. In particular, a given pair (γ, m) labels m such cylinders. Given γ's tubular neighborhood map φ, then each of the m cylinders sits in $\mathbb{R} \times \varphi(S^1 \times D)$ as the image of the graph over $\mathbb{R} \times S^1$ of the function that sends (s,t) to $e^{-2\lambda s + ix} \in \mathbb{C}$ where $\lambda > 0$ and $x \in \mathbb{R}/(2\pi\mathbb{Z})$. However, if C and C' are two such cylinders, then the corresponding points x and x' must define distinct points in the circle. There is an analogous correspondence between the cylinders that sit where $s \ll -1$ with the elements in Θ_-. The only difference is that λ is now required to be negative.

Step 3 Let Z denote a surface as described in Step 2. Then Z has a relative self-intersection number defined to be its intersection number with a deformation, Z', whose restriction to any given $|s| \gg 1$ cylinder deforms the latter so as to change the parameter x to $x + \varepsilon$ with $\varepsilon > 0$ but very small. With Z denoting the surface in question, Q_Z is used to denote this self-intersection number.

Step 4 The surface Z also has a well-defined pairing with the Euler class of the bundle K^{-1}. This pairing is defined by the usual count of the zeros of a section of this bundle over the surface with the proviso that the section should restrict to each $|s| \gg 1$ cylinder so as to be nonzero, and to be constant with respect to the trivialization of K on $\varphi(S^1 \times D)$ that is given by the coordinate vector field $\partial/\partial z$ on D. Note in this regard that K and the tangent space to D agree along $S^1 \times \{0\}$. This pairing is denoted here by $-\langle c_1, Z\rangle$.

Step 5 With Σ chosen as in Step 2, introduce the integer

$$I(\Theta_-, \Theta_+; Z) = -\langle c_1, Z\rangle + Q_Z + \sum_{(\gamma,m)\in\Theta_+} \sum_{1\le q\le m} z_{\gamma,q} - \sum_{(\gamma,m)\in\Theta_-} \sum_{1\le q\le m} z_{\gamma,q}. \tag{2-9}$$

Although the various tubular neighborhood embeddings of $S^1 \times D$ are needed to make sense of the terms in (2-9), the value of $I(\Theta_-, \Theta_+; Z)$ does not, in fact, depend on them. Moreover, the image of I in $\mathbb{Z}/p\mathbb{Z}$ depends only on the ordered pair (Θ_-, Θ_+). This is proved by Hutchings in [11]; see also Hutchings and Sullivan [12]. Hutchings also proves that this image in $\mathbb{Z}/p\mathbb{Z}$ obeys the sum rule $I(\Theta_1, \Theta_2) + I(\Theta_2, \Theta_3) = I(\Theta_1, \Theta_3)$. The relative $\mathbb{Z}/p\mathbb{Z}$ degree assignments to the generators of $\mathcal{C}_{\mathrm{ech}}$ are made so that $I(\Theta_-, \Theta_+) = \mathrm{degree}(\Theta_-) - \mathrm{degree}(\Theta_+)$.

The differential The differential that is used by Hutchings decreases the $\mathbb{Z}/p\mathbb{Z}$ grading by 1. Its definition is given in the three steps that follow.

Step 1 The almost complex structure should be chosen so as to be generic in the sense given in Lemma 2.2. Suppose that Θ_- and Θ_+ are generators of $\mathcal{C}_{\mathrm{ech}}$. Introduce $\mathcal{M}_1(\Theta_-, \Theta_+)$ to denote the set whose elements are finite sets of pairs of the form

(C, m) where C is a pseudoholomorphic subvariety and m is a positive integer. The elements in this set are constrained as follows: First, $m = 1$ unless $C = \mathbb{R} \times \gamma$ with γ a closed Reeb orbit. Second, if (C, m) and (C', m') are distinct pairs, then C is not a translate of C' along the $\mathbb{R}$ factor of $\mathbb{R} \times M$. To state the third constraint, let π_M denote the projection from $\mathbb{R} \times M$ to M. Here is the third constraint: A given element $\Sigma \in \mathcal{M}_1(\Theta_-, \Theta_+)$ defines the formal sum $\sum_{(C,m) \in \Sigma} m \pi_M(C)$, here viewed as a 2–cycle. The boundary of this 2–cycle must be $\sum_{(\gamma,m) \in \Theta_+} m\gamma - \sum_{(\gamma,m) \in \Theta_-} m\gamma$. To state the fourth constraint, let $H_2(M, \Theta_-, \Theta_+)$ denote the set that consists of the relative homology classes of 2–chains $z \subset M$ with $\partial z = \sum_{(\gamma,m) \in \Theta_+} m\gamma - \sum_{(\gamma,m) \in \Theta_-} m\gamma$. To be more explicit, chains z and z' define the same class in $H_2(M, \Theta_-, \Theta_+)$ when the closed cycle $z - z'$ is the boundary of a 3–cycle in M. Thus, $H_2(M, \Theta_-, \Theta_+)$ is an affine space modeled on $H_2(M; \mathbb{Z})$. Let $Z \subset \mathbb{R} \times M$ denote a surface as described above that can be used to define the invariant $I(\Theta_-, \Theta_+, \cdot)$, but one such that $\pi_M(Z)$ and $\sum_{(C,m) \in \Sigma} m \pi_M(C)$ define the same element in $H_2(M, \Theta_-, \Theta_+)$. Here is the fourth constraint: $I(\Theta_-, \Theta_+, Z) = 1$. Note in this regard that $I(\Theta_-, \Theta_+, \cdot)$ takes identical values on surfaces Z and Z' that give the same element in $H_2(M, \Theta_-, \Theta_+)$. Take $\mathcal{M}_1(\Theta_-, \Theta_+) = \varnothing$ when $I(\Theta_-, \Theta_+) \neq 1$.

The set $\mathcal{M}_1(\Theta_-, \Theta_+)$ inherits a topology and a local structure of the sort described by Lemma 2.2 using its tautological embedding into a disjoint union of products of $\mathcal{M}$.

Define $\mathcal{M}_{\leq 0}(\Theta_-, \Theta_+)$ by copying the preceding definition but with the condition that $Z = \sum_{(C,m) \in \Sigma} m \pi_M(C)$ is such that $I(\Theta_-, \Theta_+, Z) \leq 0$.

Step 2 As noted by Hutchings, there exists a residual set of almost complex structures for which the resulting version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ has the following properties:

(1) If $\Sigma \in \mathcal{M}_1(\Theta_-, \Theta_+)$, then $\bigcup_{(C,m) \in \Sigma} C$ is an embedded, pseudoholomorphic subvariety. Meanwhile, $\mathcal{M}_{\leq 0}(\Theta_-, \Theta_+) = \varnothing$ unless $\Theta_- = \Theta_+$, in which case it contains precisely one element, and each subvariety from this element is an $\mathbb{R}$–invariant cylinder.

(2) The space $\mathcal{M}_1(\Theta_-, \Theta_+)$ has a finite set of components and each component is a smooth, 1–dimensional manifold.

(2-10) (3) The $\mathbb{R}$ action on $\mathbb{R} \times M$ induces a free $\mathbb{R}$ action on each component. As a consequence, each element in $\mathcal{M}_1(\Theta_-, \Theta_+)$ consists of a disjoint union of $\mathbb{R}$–invariant cylinders with integer weights and one pseudoholomorphic submanifold that is not $\mathbb{R}$–invariant.

(4) Let $\Sigma \subset \mathcal{M}_1(\Theta_-, \Theta_+)$ and let $\mathcal{E} \subset \Sigma$ denote an end. Let $q_{\mathcal{E}}, \mathrm{div}_{\mathcal{E}}$ and $\varsigma_{q_{\mathcal{E}}}$ denote the data that appear in $\mathcal{E}$'s version of (2-6). Then $\mathrm{div}_{\mathcal{E}} = \{q_{\mathcal{E}}\}$ and $\varsigma_{q_{\mathcal{E}}} \neq 0$.

(5) Let $\Sigma \subset \mathcal{M}_1(\Theta_-, \Theta_+)$ and let $\mathcal{E}$ and $\mathcal{E}'$ denote distinct pairs of either positive or negative ends of Σ with $\gamma_{\mathcal{E}} = \gamma_{\mathcal{E}'}$ and $q_{\mathcal{E}} = q_{\mathcal{E}'}$. Let $\varsigma_{q_{\mathcal{E}}}$ and $\varsigma_{q_{\mathcal{E}'}}$ denote the $2\pi q$–periodic eigenvector that appears in the respective $\mathcal{E}$ and $\mathcal{E}'$ versions of (2-6). Then $\varsigma_{q_{\mathcal{E}}}|_t \neq \varsigma_{q_{\mathcal{E}'}}|_{t+2\pi k}$ for any $t \in S^1$ and $k \in \mathbb{Z}$.

Properties (1)–(3) are proved by Hutchings [11]; see also Hutchings and Sullivan [12]. Properties (4) and (5) follow from what is done in Section 3 of Hutchings and Taubes [14]. Property (4) also needs some facts that can be derived from what is said in Section 11 of [11] and its Remark 1.

Let $\mathcal{J}_a'$ denote the residual set of a–compatible almost complex structures that have the properties listed in (2-10) and lie in Lemma 2.2's residual set. Fix an almost complex structure from $\mathcal{J}_a'$. Let Θ_- and Θ_+ denote generators of the embedded contact homology chain complex. Hutchings uses constructions of Bourgeois and Mohnke [1] to associate a sign, ± 1, to each component of $\mathcal{M}_1(\Theta_-, \Theta_+)$. The full details of this are given in Section 9 and especially Section 9.5 of [14]. Let $\sigma(\Theta_-, \Theta_+)$ denote the sum of these signs when $\mathcal{M}_1(\Theta_-, \Theta_+)$ is nonempty, and 0 otherwise.

To say a wee bit more about these signs, note that the sign that is associated to any given component of $\mathcal{M}_1(\Theta_-, \Theta_+)$ is obtained by comparing two natural orientations. The first is that induced by the generator of the $\mathbb{R}$ action. The second is defined using ideas of Quillen [18] about determinant line bundles of parametrized families of Fredholm operators. As noted by Hutchings and explained in Section 9.5 of [14], these ideas of Quillen can be used along lines explained in [1] so as to define a second orientation to each component of $\mathcal{M}_1(\Theta_-, \Theta_+)$. (The respective parts of Θ_- and Θ_+ that involve the ordering of their even rotation number hyperbolic Reeb orbits is needed solely to define this second orientation of $\mathcal{M}_1(\Theta_-, \Theta_+)$.)

Step 3 The differential, δ, on $\mathcal{C}_{\mathrm{ech}}$ is defined on any given generator Θ by the formula $\delta\Theta = \sum_{\Theta' \in \mathcal{C}_{\mathrm{ech}}} \sigma(\Theta', \Theta)\ \Theta'$. The proof that $\delta^2 = 0$ appears in [13; 14]. The proof requires that the almost complex structure come from a certain residual subset of $\mathcal{J}_a'$. This last residual set is denoted by $\mathcal{J}_a$ in what follows. Almost complex structures will always be chosen from $\mathcal{J}_a$ unless explicitly noted otherwise.

The embedded contact homology for the class Γ is defined to be the homology of δ on $\mathcal{C}_{\mathrm{ech}}$. As defined, δ decreases the $\mathbb{Z}/p\mathbb{Z}$ degree by 1 so this homology is $\mathbb{Z}/p\mathbb{Z}$ graded. The homology defined here is denoted in what follows by $\mathcal{H}_{\mathrm{ech}}$, where it is understood that the class Γ is fixed in advance and not subsequently changed.

The filtration As noted in Section 2.a, the set of Reeb orbits with an a priori symplectic action bound is compact. If all such orbits are nondegenerate, then there are but a finite set with symplectic action less than any given amount. Granted this last point, fix $L > 0$ and let $\mathcal{C}_{\mathrm{ech}}^L \subset \mathcal{C}_{\mathrm{ech}}$ denote the submodule that is generated by elements Θ that obey $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma < L$. This is a finitely generated chain complex.

If $\Sigma \subset \mathcal{M}_1(\Theta,\Theta')$, then $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma \leq \sum_{(\gamma,m)\in\Theta'} m\ell_\gamma$. As a result, the differential on $\mathcal{C}_{\mathrm{ech}}$ maps $\mathcal{C}_{\mathrm{ech}}^L$ to itself. This understood, let $\mathcal{H}_{\mathrm{ech}}^L$ denote the homology that is defined by δ on $\mathcal{C}_{\mathrm{ech}}^L$. Then $\mathcal{H}_{\mathrm{ech}} = \operatorname{dir\,lim}_{L\to\infty} \mathcal{H}_{\mathrm{ech}}^L$, where the homomorphisms for this direct limit are induced by the L and $L' > L$ versions of the submodule inclusion homomorphism from $\mathcal{C}_{\mathrm{ech}}^L$ into $\mathcal{C}_{\mathrm{ech}}^{L'}$.

2.d Changing the contact structure

As it turns out, the proof of Theorem 1 is considerably shorter when the contact structure is approximated by one which has a canonical form near some of its Reeb orbits. The following lemma describes these canonical forms.

Lemma 2.3 *Suppose that $(\nu,\mu) \in C^\infty(S^1;\mathbb{R}\oplus\mathbb{C})$.*

The elliptic case *Suppose that (ν,μ) is elliptic with rotation angle $\mathrm{R} \in \mathbb{R}$. There is a homotopy of (ν,μ) through elliptic pairs with rotation angle R to the pair $(\frac{1}{2}\mathrm{R}, 0)$.*

The hyperbolic case *Suppose that (ν,μ) is hyperbolic with rotation number k. If $\varepsilon > 0$ is small, there is a homotopy of (ν,μ) through hyperbolic pairs to the pair $(\frac{1}{4}k, i\varepsilon e^{ikt})$.*

Proof of Lemma 2.3 The statement in the elliptic case is straightforward; it follows readily from the geometry of $\mathrm{SL}(2;\mathbb{R})$ that any two elliptic pairs with the same rotation number are homotopic through a family of constant rotation number elliptic pairs. In the hyperbolic case, remark first that any two hyperbolic pairs with the same rotation number are homotopic through hyperbolic pairs. As a consequence, it is enough to verify that the pair $(\frac{1}{2}k, i\varepsilon e^{-ikt})$ is hyperbolic with rotation number k when ε is sufficiently small. The calculation is straightforward and left to the reader. □

Assume that the contact structure a is from Lemma 2.1's residual set and the almost complex structure, J, is from $\mathcal{J}_a$. Fix $L \geq 1$ and $\delta > 0$. A pair $(\hat{a}, \hat{J})$ of contact structure on M and compatible almost complex structure on $\mathbb{R}\times M$ is said to be a *(δ, L)–approximation* for (a, J) when the following is true: There is a smooth, 1–parameter family $\{(a_\tau, J_\tau)\}_{\tau\in[0,1]}$ of pairs of contact structure and compatible almost complex structure with $(a_0, J_0) = (a, J)$ and $(a_1, J_1) = (\hat{a}, \hat{J})$; and such that

(1) For each $\tau \in [0, 1]$, the respective sets of a and a_τ Reeb orbits with symplectic action less than L are identical.

(2) Let γ denote a Reeb orbit for a with $\ell_\gamma < L$. If γ is elliptic or hyperbolic as defined using a, then it is respectively elliptic or hyperbolic as defined using any $\tau \in [0, 1]$ version of a_τ and the rotation number is independent of τ.

(3) Let Θ_- and Θ_+ denote generators of $\mathcal{C}^L_{\mathrm{ech}}$. For each $\tau \in [0, 1]$, there is a 1–1 correspondence between the components of the respective J and J_τ versions of $\mathcal{M}_1(\Theta_-, \Theta_+)$. The correspondence between the respective versions of $\mathcal{M}_1(\Theta_-, \Theta_+)$ is such that partnered components contribute the same sign to the respective J and J_τ versions of $\sigma(\Theta_-, \Theta_+)$. Meanwhile, the space $\mathcal{M}_{\leq 0}(\Theta_-, \Theta_+)$ is empty unless $\Theta_- \neq \Theta_+$, and in this case, it contains but one element, and each subvariety from the
(2-11) latter is an $\mathbb{R}$–invariant cylinder.

(4) Let γ denote a Reeb orbit with $\ell_\gamma < L$. There is a coordinate embedding $\varphi\colon S^1 \times D \to M$ of the sort described in the preceding with the following property: If γ is hyperbolic with rotation number k, then the $\hat{a}$–version of the pair (ν, μ) is equal to $(\frac{1}{2}k, i\varepsilon e^{-ikt})$ for some $\varepsilon \in (0, \delta)$. If γ is elliptic with rotation number R, then

   (i) $$(2\pi/\ell_\gamma)\varphi^*\hat{a} = (1 - \mathrm{R}|z|^2)\,dt + \frac{i}{2}(z\,d\bar{z} - \bar{z}\,dz).$$

   (ii) The φ^*–pullback of the $\hat{J}$–version of $T^{1,0}(\mathbb{R} \times M)$ is spanned by the forms
   $$ds + i\hat{a} \quad \text{and} \quad \frac{\ell_\gamma}{2\pi}(dz - i\mathrm{R}z\,dt).$$

(5) The contact structure $\hat{a}$ comes from Lemma 2.1's residual set and the almost complex structure $\hat{J}$ comes from the set $\mathcal{J}_{\hat{a}}$.

The next subsection gives a first indication as to why pairs $(\hat{a}, \hat{J})$ as just described are easy to work with.

The proposition that follows asserts that the homology of $\mathcal{C}^L_*$ as defined by a pair (a, J) is isomorphic to that defined by a (δ, L) approximation.

Proposition 2.4 *Let a denote a contact 1–form from the residual set given in Lemma 2.1 and let J denote a complex structure from $\mathcal{J}_a$. Fix $L \geq 1$ such that there is no generator Θ of $\mathcal{C}_{\mathrm{ech}}$ with $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma = L$. Fix also $\delta > 0$. Let $(\hat{a}, \hat{J})$ denote a (δ, L) approximation to the given pair (a, J). Then the identification provided by the* *first item in (2-11) between the Reeb orbits with symplectic action less than L induces a degree preserving isomorphism between the a and $\hat{a}$ versions of $\mathcal{C}_{\mathrm{ech}}^{L}$ that intertwines the respective differentials. Thus, it induces an isomorphism between the respective (a, J) and $(\hat{a}, \hat{J})$ versions of $\mathcal{H}_{\mathrm{ech}}^{L}$.*

Proof of Proposition 2.4 The fact that the isomorphism preserves degree follows from the third item in (2-11). The fact that it intertwines the differential follows from the fourth item in (2-11). □

The final proposition asserts that there are in all cases (δ, L) approximations.

Proposition 2.5 *Let a denote a contact form from Lemma 2.1's residual set and let $J \in \mathcal{J}_a$. Fix $L \geq 1$ such that there is no generator $\Theta \in \mathcal{C}_{\mathrm{ech}}$ with $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma = L$. Given $\delta > 0$, there exist (δ, L) approximations to (a, J). In fact, there exists $\kappa > 1$ that depends only on (a, J) and has the following significance: Fix $\rho > 0$ and there exists a (δ, L) approximation to (a, J) that is the end member of a family $\{(a_\tau, J_\tau)\}_{\tau\in[0,1]}$ which has initial member (a, J), and is such that each $\tau \in [0, 1]$ member obeys*

- *$(a_\tau, J_\tau) = (a, J)$ on the complement of the radius ρ tubular neighborhoods of the Reeb orbits with length less than L.*
- *$a_\tau - a$ has C^1–norm less than ρ and C^2–norm less than κ.*
- *$J_\tau - J$ has C^0–norm less than ρ and C^1–norm less than κ.*

This proposition is proved in the Appendix to this article.

2.e Pseudoholomorphic subvarieties for (δ, L) approximating pairs

This last subsection is an aside of sorts whose purpose is to say something about the pseudoholomorphic curves for a pair (a, J) of contact form on $S^1 \times \mathbb{C}$ and compatible almost complex structure on $\mathbb{R} \times S^1 \times \mathbb{C}$ where

$$a = \frac{\ell}{2\pi}\left((1 - \mathrm{R}|z|^2)\, dt + \frac{i}{2}(z\, d\bar{z} - \bar{z}\, dz)\right), \tag{2-12}$$

and where $T^{1,0}(\mathbb{R} \times S^1 \times \mathbb{C})$ is spanned by $ds + ia$ and $dz - i\mathrm{R}z\, dt$. Here, $\ell > 0$ and R are constant. A straightforward calculation verifies the following: Let $w = (2\pi/\ell)s - (1/2)|z|^2$. Then

- $dw + i\,dt \in T^{1,0}(\mathbb{R} \times M)$; thus the constant (w, t) planes are pseudoholomorphic.
- (2-13) Subvarieties $z = f(w, t)$ are pseudoholomorphic if and only if

$$\left(\frac{\partial}{\partial w} + i\frac{\partial}{\partial t} + \mathrm{R}\right) f = 0.$$

Note that the complex structure in this case is integrable. Local holomorphic coordinates are $u = w + it$ and $x = e^{\mathrm{R}w} z$.

3 Seiberg–Witten Floer (co)homology

The purpose of this section is to say more about the relevant versions of Seiberg–Witten Floer homology and cohomology. As a complete treatment of the subject is given by Kronheimer and Mrowka in [16], what follows focuses for the most part on those aspects of the story that are relevant to the case when M comes with a contact 1–form. In any event, much of what is said below paraphrases the definitions and discussion in Kronheimer and Mrowka's book [16].

3.a The Seiberg–Witten equations on M and $\mathbb{R} \times M$

Fix a Riemannian metric on M so as to define the bundle of oriented, orthonormal frames for TM. Let $\mathrm{Fr} \to M$ denote this principal SO(3) bundle. A $\mathrm{Spin}_{\mathbb{C}}$ lift of this bundle denotes here a principal $U(2)$ bundle, $F \to M$ such that $F/U(1) = \mathrm{Fr}$. Such a lift is called a $\mathrm{Spin}_{\mathbb{C}}$ structure. Two lifts, F and F', are deemed equivalent if there is a bundle isomorphism from F to F' that covers the projections to Fr. The set of equivalence classes of lifts can be put in 1–1 correspondence with elements in $H^2(M;\mathbb{Z})$.

Let $F \to M$ denote now a $\mathrm{Spin}_{\mathbb{C}}$ structure. Use $\mathbb{S}$ to denote the associated $\mathbb{C}^2$ bundle $F \times_{U(2)} \mathbb{C}^2$. Use $\det(\mathbb{S})$ in what follows to denote the complex hermitian line bundle $F \times_{U(1)} \mathbb{C}$. Having fixed a $\mathrm{Spin}_{\mathbb{C}}$ structure, the associated Seiberg–Witten equations constitute a system of equations for a pair (A, ψ) where A here denotes a connection on $\det(\mathbb{S})$ and ψ denotes a section of $\mathbb{S}$.

To say more about these equations, introduce the Clifford multiplication homomorphism $\mathrm{cl}\colon T^*M \to \mathrm{End}(\mathbb{S})$. This homomorphism is such that $\mathrm{cl}(b)^\dagger = -\mathrm{cl}(b)$ and $\mathrm{cl}(b)\,\mathrm{cl}(b') = -\mathrm{cl}(*(b \wedge b')) - \langle b, b'\rangle$. Here, $\langle\ ,\ \rangle$ denotes the metric inner product and $*$ denotes the associated Hodge star. The Seiberg–Witten equations involve two related homomorphisms. The first, $\hat{c}\colon \mathbb{S} \otimes T^*M \to \mathbb{S}$, is defined so as to send any given decomposable element $\eta \otimes b$ to $\mathrm{cl}(b)\eta$. The second is a quadratic, bundle preserving map from $\mathbb{S}$ to iT^*M. The image of any given $\eta \in \mathbb{S}$ under the latter map is written in what follows as $\eta^\dagger \tau \eta$. It is defined by the rule $\langle b, \eta^\dagger \tau \eta\rangle = \eta^\dagger\,\mathrm{cl}(b)\eta$.

Let A now denote a connection on $\det(\mathbb{S})$. In what follows, the Hodge star of its curvature 2–form is denoted by B_A, this being a section of iT^*M. The connection A

and the Levi-Civita connection on TM define a Hermitian connection on $\mathbb{S}$. The associated covariant derivative is denoted in what follows by ∇_A. This covariant derivative is used to define the Dirac operator, $D^A = \hat{c}(\nabla_A) : C^\infty(M;\mathbb{S}) \to C^\infty(M;\mathbb{S})$.

A pair (A, ψ) of connection on $\det(\mathbb{S})$ and section ψ of $\mathbb{S}$ obeys the simplest version of the Seiberg–Witten equations when

$$B_A - \psi^\dagger \tau \psi = 0 \quad \text{and} \quad D^A \psi = 0. \tag{3-1}$$

A rigorous definition of the Seiberg–Witten Floer homology involves solutions to perturbed versions of the equations in (3-1). The description of these perturbed equations requires a brief digression to set the stage.

To start the digression, remark that the equations in (3-1) are gauge invariant in the following sense: If u is a smooth map from M to $U(1)$, then the pair $(A - 2u^{-1}du, u\psi)$ solves (3-1) if and only if (A, ψ) does. A function, $\mathfrak{g}$, of pairs consisting of a connection on $\det(\mathbb{S})$ and a section of $\mathbb{S}$ is deemed gauge invariant when $\mathfrak{g}(A - 2u^{-1}du, u\psi) = \mathfrak{g}(A, \psi)$ for all $u \in C^\infty(M; U(1))$. The allowed sorts of functions form what Kronheimer and Mrowka call a large, separable Banach space of tame perturbations. Such a Banach space is described in Chapter 11 of [16]. Somewhat more is said below about this. This Banach space of tame perturbations that is used here is denoted by $\mathcal{P}$. If $\mathfrak{g} \in \mathcal{P}$, then the differential of $\mathfrak{g}$ at any given (A, ψ) defines section $(\mathfrak{T}|_{(A,\psi)}, \mathfrak{S}|_{(A,\psi)})$ of $iT^*M \oplus \mathbb{S}$ by writing $\frac{d}{dt}\mathfrak{g}(A + tb, \psi + t\eta)|_{t=0}$ as $\int_M (b \wedge *\mathfrak{T} - \frac{1}{2}(\eta^\dagger \mathfrak{S} + \mathfrak{S}^\dagger \eta))$. Each $\mathfrak{g} \in \mathcal{P}$ gives the equation

$$B_A - \psi^\dagger \tau \psi - \mathfrak{T}_{(A,\psi)} = 0 \quad \text{and} \quad D^A \psi - \mathfrak{S}|_{(A,\psi)} = 0. \tag{3-2}$$

Note that if u is a smooth map from M to $U(1)$, then $(A - 2u^{-1}du, u\psi)$ solves (3-2) if and only if (A, ψ) does. Pairs of connection and section that are related in this way are said to be *gauge equivalent*.

There are corresponding Seiberg–Witten equations on $\mathbb{R} \times M$ that constitute a system of equations for a pair $\mathfrak{d} = (A, \psi)$, where A now denotes a map from $\mathbb{R}$ into the space of Hermitian connections on $\det(\mathbb{S})$ and ψ denotes a map from $\mathbb{R}$ into the space of sections of $\mathbb{S} \to M$. With $s \in \mathbb{R}$ denoting the Euclidean coordinate, these equations read

(3-3)
- $\frac{\partial}{\partial s} A + B_A - \psi^\dagger \tau^k \psi - \mathfrak{T}(A, \psi) = 0.$
- $\frac{\partial}{\partial s} \psi + D^A \psi - \mathfrak{S}(A, \psi) = 0.$

Of particular interest here are *instanton* solutions. An instanton is a solution to (3-3) with $s \to +\infty$ limit and $s \to -\infty$ limit, each a solution to (3-1). If u is a smooth map from M to $U(1)$ and (A, ψ) is a solution to (3-3), then so is $(A - 2u^{-1}du, u\psi)$.

3.b An overview of Seiberg–Witten Floer homology/cohomology

This subsection very briefly summarizes the story from [16]. To start, Kronheimer and Mrowka prove that (3-2) has but a finite set of solutions up to gauge equivalence if $\mathfrak{g}$ is chosen from a certain residual subset in the Banach space $\mathcal{P}$. With one caveat, these equivalence classes form a basis for the chain complex that defines the Seiberg–Witten Floer homology. The caveat concerns the case when the first Chern class of $\det(\mathbb{S})$ is a torsion class. The situation here is more complicated by virtue of the fact that (3-2) admits solutions with ψ identically zero when $c_1(\det(\mathbb{S}))$ is torsion. These $\psi = 0$ solutions are deemed to be *reducible*, and those with ψ somewhere nonzero are deemed to be *irreducible*. Here is the salient distinction: The group $C^\infty(M; U(1))$ acts with trivial stabilizer on any pair (A, ψ) with ψ somewhere nonzero, but it acts with stabilizer $U(1)$ on any $(A, 0)$. Here, $U(1) \subset C^\infty(M; U(1))$ is identified with the constant maps. This distinction makes for a chain complex when $c_1(\det(\mathbb{S}))$ is torsion wtih one generator for each gauge equivalence class of irreducible solution to (3-2), and a countable set of generators for each gauge equivalence class of reducible solution to (3-2). The chain complex for the Seiberg–Witten Floer homology is denoted in what follows by $\mathcal{C}_{\mathrm{SW}}$. This $\mathbb{Z}$ module is finite when $c_1(\det(\mathbb{S}))$ is not torsion, but not finitely generated otherwise.

The complex $\mathcal{C}_{\mathrm{SW}}$ has a natural, relative $\mathbb{Z}/p\mathbb{Z}$ grading, where p here denotes the divisibility of the class $c_1(\det(\mathbb{S}))$ in $H^2(M; \mathbb{Z})$/torsion. The complex is $\mathbb{Z}$ graded when $c_1(\det(\mathbb{S}))$ is torsion. This grading is described in some detail momentarily. Suffice it to say for now that the relative grading between two irreducible generators is defined to be minus the spectral flow for a certain 1–parameter family of unbounded, self-adjoint, operators (with compact resolvent) on $L^2(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$. This family is constructed from a path, parametrized by $[0, 1]$, of pairs (A, ψ) with A a connection on $\det(\mathbb{S})$ and ψ a section of $\mathbb{S}$. This path starts at the first irreducible solution, and ends at the second. Meanwhile, the operator that is parametrized by any such pair (A, ψ) is a self map of $C^\infty(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$; it is defined from the linearization of (3-2) at (A, ψ).

In the case where $c_1(\det(\mathbb{S}))$ is torsion, the countable set of cycles that correspond to any given reducible solution can be labeled by a set of the form $\{k, k-2, \ldots\}$, where $k \in \mathbb{Z}$. The relative grading between cycles $k - 2j$ and $k - 2j'$ is $2(j - j')$. The integer k can be fixed once a fiducial, irreducible configuration is chosen to define the zero point for the grading. Given such a choice, k is then minus the spectral flow for a family of self-adjoint differential operators that starts at a certain operator that is parametrized by the fiducial configuration and ends at one parametrized by a suitable irreducible configuration lying very near the given reducible solution. In the case when $c_1(\det(\mathbb{S}))$ is torsion, the $\mathbb{Z}$–module $\mathcal{C}_{\mathrm{SW}}$ is finitely generated in each degree.

The differential that defines the Seiberg–Witten Floer homology is defined using an algebraic count of instanton solutions to (3-3) as defined by any $\mathfrak{g}$ from a certain residual subset in $\mathcal{P}$. The differential decreases that $\mathbb{Z}/p\mathbb{Z}$ degree by 1. To say more about the differential, note that Kronheimer and Mrowka prove the following: Let $\mathfrak{c}_-$ and $\mathfrak{c}_+$ denote pairs of irreducible solutions to (3-2) and introduce $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$ to denote the set of instanton solutions to (3-3) with $s \to -\infty$ limit equal to $\mathfrak{c}_-$ and with $s \to \infty$ limit equal to $u \cdot \mathfrak{c}_+$ with $u \in C^\infty(M; S^1)$. This set depends only on the gauge equivalence classes of $\mathfrak{c}_-$ and $\mathfrak{c}_+$ in the following sense: Suppose that $u \in C^\infty(M; U(1))$. Let $\mathfrak{d} = (A, \psi) \in \mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$; then $u \cdot \mathfrak{d} \in \mathfrak{M}(u \cdot \mathfrak{c}_-, \mathfrak{c}_+)$ where $u \cdot (A, \psi)$ is shorthand for $(A - 2u^{-1}du, u\psi)$.

Given that $\mathfrak{g}$ comes from a certain residual subset of $\mathcal{P}$, this $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$ has the structure of a smooth, finite dimensional manifold. There is one zero dimensional component if $\mathfrak{c}_-$ is gauge equivalent to $\mathfrak{c}_+$; and in this case, $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$ consists of the constant map $s \to \mathfrak{c}_-$. There are no zero dimensional components otherwise. Meanwhile, there is a finite set of 1–dimensional components of $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$; and each component is an orbit of the $\mathbb{R}$ action that is induced by translation along the $\mathbb{R}$ factor of $\mathbb{R} \times M$. Such 1–dimensional components exist only in the case where the degree of $\mathfrak{c}_+$ is one less than that of $\mathfrak{c}_-$. Use $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ in what follows to denote the space of 1–dimensional components of $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$.

Each component of $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ has a corresponding sign. This sign is obtained by comparing the orientation given by the generator of the $\mathbb{R}$ action with an orientation that is defined using Quillen's ideas about determinant line bundles. Somewhat more is said about this below, but in any event, the full story is given in [16]. Let $\sigma(\mathfrak{c}_-, \mathfrak{c}_+)$ denote the sum of these signs when $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ is nonempty, or zero when it is. In the case when $c_1(\det(\mathbb{S}))$ is not torsion, the differential that defines the Seiberg–Witten Floer homology acts on any given generator $\mathfrak{c}$ as $\delta\mathfrak{c} = \sum_{\mathfrak{c}' \in \mathcal{C}_{\mathrm{SW}}} \sigma(\mathfrak{c}, \mathfrak{c}')\mathfrak{c}'$. In the case when $c_1(\det(\mathbb{S}))$ is torsion, what is written here defines the part of the differential that involves the irreducible generators. Only this part is needed for the proofs of the theorems in the introduction. This being the case, the reader can consult [16] to see how the rest of the differential is defined.

The homology of this differential on $\mathcal{C}_{\mathrm{SW}}$ is denoted by $\widehat{\mathcal{H}}_*$ in [16], and so denoted by $\widehat{\mathcal{H}}_*$ here. This homology is finitely generated in the case when $c_1(\det(\mathbb{S}))$ is not torsion. In the case when this class is torsion, $\widehat{\mathcal{H}}_*$ is finitely generated in each degree; and the set of degrees where it is nonzero is bounded from above but unbounded from below.

The Seiberg–Witten Floer cohomology is defined by the dual differential on the $\mathbb{Z}$–module $\mathcal{C}^{\mathrm{SW}} = \mathrm{Hom}(\mathcal{C}_{\mathrm{SW}}, \mathbb{Z})$. This differential, δ^*, acts on any given cocycle C by

$(\delta^* C)(\cdot) = C(\delta(\cdot))$. Note that the basis described above for $\mathcal{C}_{\mathrm{SW}}$ supplies a canonical dual basis for $\mathcal{C}^{\mathrm{SW}}$ and a $\mathbb{Z}/p\mathbb{Z}$ grading. This differential sends any given basis element $\mathfrak{c}$ of $\mathcal{C}^{\mathrm{SW}}$ to

$$\delta^* \mathfrak{c} = \sum_{\mathfrak{c}' \in \mathcal{C}^{\mathrm{SW}}} \sigma(\mathfrak{c}', \mathfrak{c})\mathfrak{c}'. \tag{3-4}$$

Note that it increases the $\mathbb{Z}/p\mathbb{Z}$ degree by 1. The resulting cohomology groups are denoted in what follows by $\mathcal{H}^{\mathrm{SW}}$.

Keep in mind that the definition of these groups requires the choice of a function from a certain residual subset in $\mathcal{P}$. However, two such functions give isomorphic versions of Seiberg–Witten homology and cohomology. Section 3.d says more about the criteria for admission in this residual set. Section 3.h, an appendix to Section 3, says more about these isomorphisms.

3.c Contact forms and Seiberg–Witten equations

Suppose now that a is a given contact 1–form on M. Fix a metric on M for which $*da = 2a$ and $|a| = 1$. Note that such a metric on TM is neither more nor less than an almost complex structure, J, on kernel(a) such $da(\cdot, J(\cdot))$ is a metric on the kernel of a. In particular, a pair (a, J) of contact form and almost complex structure in $\mathcal{J}_a$ supplies M with a canonical metric.

With the metric fixed, let $F \to M$ denote a $\mathrm{Spin}_{\mathbb{C}}$ structure. The endomorphism $\mathrm{cl}(a)$ on $\mathbb{S}$ has square -1 and so its eigenspaces in each fiber define a splitting of $\mathbb{S}$ as the orthogonal, direct sum of two complex, Hermitian line bundles. This direct sum is written in what follows as $E \oplus EK^{-1}$ where $E \to \mathbb{S}$ and $K \to \mathbb{S}$ are complex line bundles. The convention has $\mathrm{cl}(a)$ act as i on the first summand and $-i$ on the second. The bundle $K^{-1} \to \mathbb{S}$ is isomorphic as an $\mathrm{SO}(2)$ bundle to the kernel of a in TM with the orientation defined by da. Note that any given equivalence class of complex line bundles can arise in this manner from some $\mathrm{Spin}_{\mathbb{C}}$ structure on M. Moreover, two $\mathrm{Spin}_{\mathbb{C}}$ structures have isomorphic versions of E if and only if they are equivalent.

The contact form a determines a *canonical* $\mathrm{Spin}_{\mathbb{C}}$ *structure*, the $\mathrm{Spin}_{\mathbb{C}}$ structure for which the spinor bundle decomposes as $\mathbb{S} = \mathbb{S}_I = I_{\mathbb{C}} \oplus K^{-1}$, where $I_{\mathbb{C}} \to M$ denotes the trivial complex line bundle. Fix a unit norm section $1_{\mathbb{C}}$ of $I_{\mathbb{C}}$. Such a section defines a *canonical connection* on $K^{-1} = \det(\mathbb{S}_I)$. This is the unique connection for which the section $\psi_I = (1_{\mathbb{C}}, 0)$ of $\mathbb{S}_I$ is annihilated by the corresponding Dirac operator. This canonical connection is written as A_K.

Let $\mathbb{S} = E \oplus EK^{-1}$ now denote the spinor bundle for some other $\mathrm{Spin}_{\mathbb{C}}$ structure. Any given connection on $\det(\mathbb{S}) = E^2 K^{-1}$ can be written as $A_K + 2A$ where A is a connection on E. With A a connection on E, the symbol D_A is used to denote the operator D^{A_K+2A} that appears in (3-1)–(3-3). $\mathrm{Conn}(E)$ is used in what follows to denote the Fréchet space of smooth, Hermitian connections on E.

With the splitting $\mathbb{S} = E \oplus EK^{-1}$ given, the corresponding components of any given section ψ of $\mathbb{S}$ are written as (α, β). Thus, α is a section of E and β one of EK^{-1}.

The contact form is also used to define a certain family of perturbations for use in (3-2) and (3-3). This family is parametrized by $[1, \infty)$. To set the stage, view $\mathcal{P}$ now as a Banach space of functions on $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$. Any given $r \in [1, \infty)$ version of these equations requires the choice of a function $\mathfrak{g}$ from $\mathcal{P}$. These equations, viewed now as equations for a pair (A, ψ) from $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$, read

(3-5)
- $B_A - r(\psi^\dagger\tau\psi - ia) - \mathfrak{T}|_{(A,\psi)} + \frac{1}{2}B_{A_K} = 0.$
- $D_A\psi - \mathfrak{S}|_{(A,\psi)} = 0.$

Here, $\mathfrak{T}$ and $\mathfrak{S}$ are defined from $\mathfrak{g}$ as before. Meanwhile, $*B_{A_K}$ is the curvature 2–form for the connection A_K. The associated version of (3-3) for a map $s \to (A, \psi)$ from $\mathbb{R}$ to $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ is

(3-6)
- $\frac{\partial}{\partial s}A + B_A - r(\psi^\dagger\tau\psi - ia) - \mathfrak{T}|_{(A,\psi)} + \frac{1}{2}B_{A_K} = 0.$
- $\frac{\partial}{\partial s}\psi + D_A\psi - \mathfrak{S}|_{(A,\psi)} = 0.$

Equations (3-5) and (3-6) can be made to look like (3-2) and (3-3) by replacing ψ in the latter by $r^{1/2}\psi$. It is left to the reader to derive the relation between the respective versions of what is denoted by $\mathfrak{g}$.

3.d The Banach space $\mathcal{P}$

As noted in Section 3.b, the Seiberg–Witten Floer homology and cohomology can be defined only after choosing a function $\mathfrak{g}$ from a certain residual subset of $\mathcal{P}$. There are two criteria for membership in this set when the first Chern class of E is not torsion and three criteria when it is. The first concerns the linearized version of (3-5). To say more, fix any pair $\mathfrak{c} = (A, \psi) \in \mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$. Define an operator $\mathfrak{L}_{\mathfrak{c}}$ with domain and range $C^\infty(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$ as follows: It sends any given triple (b, η, ϕ) in its domain to the section of $iT^*M \oplus \mathbb{S} \oplus i\mathbb{R}$ whose three components are

(3-7)
- $*db - d\phi - 2^{-1/2}r^{1/2}(\psi^\dagger\tau\eta + \eta^\dagger\tau\psi) - \mathfrak{t}_{(A,\psi)}(b, \eta),$
- $D_A\eta + 2^{1/2}r^{1/2}(cl(b)\psi + \phi\psi) - \mathfrak{s}_{(A,\psi)}(b, \eta),$
- $*d{*b} - 2^{-1/2}r^{1/2}(\eta^\dagger\psi - \psi^\dagger\eta),$

where the pair $(\mathfrak{t}_{(A,\psi)}, \mathfrak{s}_{(A,\psi)})$ denotes the operator on $C^\infty(M; iT^*M \oplus \mathbb{S})$ that sends a given section (b,η) to $(\frac{d}{dt}\mathfrak{T}(A+tb, \psi+t\eta),\ \frac{d}{dt}\mathfrak{S}(A+tb, \psi+t\eta))|_{t=0}$. The operator $\mathfrak{L}_c$ is symmetric and extends to $L^2(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$ as an unbounded, self-adjoint operator with dense domain $L^2_1(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$. As such, it has pure point spectrum and each eigenvalue has finite multiplicity. Moreover, the spectrum is unbounded from above and from below. Finally, every eigenvector is smooth.

A function $\mathfrak{g}$ from $\mathcal{P}$ can be used to define the Seiberg–Witten Floer homology only when the following criterion is met:

(3-8) **Criterion 1** If $\mathfrak{c}$ is an irreducible solution to (3-5), then the operator $\mathfrak{L}_\mathfrak{c}$ has trivial kernel. If $\mathfrak{c}$ is a reducible solution to (3-5), then kernel$(\mathfrak{L}_\mathfrak{c})$ consists of the constant sections of the $i\mathbb{R}$ summand of $iT^*M \oplus \mathbb{S} \oplus i\mathbb{R}$.

A solution that satisfies this criterion is said to be *nondegenerate*.

The second required property for $\mathfrak{g}$ involves the operator on $\mathbb{R}\times M$ that arises from the linearized version of (3-6). To elaborate, suppose that $s \to \mathfrak{d}(s) = (A,\psi)$ is a smooth map from $\mathbb{R}$ into $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ that has $s \to \pm\infty$ limits. Let $\mathfrak{c}_\pm$ denote the latter. Now define the operator $\mathfrak{D}_\mathfrak{d}$ from $C^\infty(\mathbb{R}\times M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$ to itself as follows: It sends a triple (b,η,ϕ) to the section whose respective three components are

(3-9)
- $\frac{\partial}{\partial s} b + {*db} - d\phi - 2^{-1/2} r^{1/2}(\psi^\dagger \tau\eta + \eta^\dagger\tau\psi) - \mathfrak{t}_{(A,\psi)}(b,\eta)$,
- $\frac{\partial}{\partial s}\eta + D_A\eta + 2^{1/2} r^{1/2}(\mathrm{cl}(b)\psi + \phi\psi) - \mathfrak{s}_{(A,\psi)}(b,\eta)$,
- $\frac{\partial}{\partial s}\phi + {*d{*b}} - 2^{-1/2} r^{1/2}(\eta^\dagger\psi - \psi^\dagger\eta)$.

This operator extends to define a bounded operator from $L^2_1(\mathbb{R}\times M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$ to $L^2(\mathbb{R}\times M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$. If both $\mathfrak{c}_+$ and $\mathfrak{c}_-$ are irreducible and if both the $\mathfrak{c} = \mathfrak{c}_+$ and $\mathfrak{c} = \mathfrak{c}_-$ versions of $\mathfrak{L}_\mathfrak{c}$ have trivial kernel, then this extended version of $\mathfrak{D}_\mathfrak{d}$ is a Fredholm operator. This understood, what follows is the second requirement on $\mathfrak{g}$ for its use to define the Seiberg–Witten Floer homology and cohomology.

(3-10) **Criterion 2** Let $s \to \mathfrak{d}(s)$ denote an instanton solution to (3-6) such that both $|s| \to \infty$ limits are irreducible and such that their corresponding versions of $\mathfrak{L}_\mathfrak{c}$ have trivial kernel. Then $\mathfrak{D}_\mathfrak{d}$ has trivial kernel.

An instanton that satisfies this criterion is also said to be *nondegenerate*.

The third requirement on $\mathfrak{g}$ for its use in defining the Seiberg–Witten Floer homology and cohomology concerns the operator $\mathfrak{D}_\mathfrak{d}$ when $s \to \mathfrak{d}(s)$ is an instanton solution to

(3-6) with at least one $|s| \to \infty$ limit reducible. As these solutions play no essential role in what follows, this third criterion will not be stated explicitly. The reader can refer instead to [16].

A perturbation that satisfies the aforementioned Criteria (3-8) and (3-10) is said to be *suitable*. The upcoming Section 3.h says more about perturbations.

The space $\mathcal{P}$ has the following property: Let $\mathfrak{g} \in \mathcal{P}$. Just as the first derivatives of $\mathfrak{g}$ at any given $(A,\psi) \in \mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ define the smooth section $(\mathfrak{T}|_{(A,\psi)}, \mathfrak{S}|_{(A,\psi)})$ of the bundle $iT^*M \oplus \mathbb{S}$, so the derivatives of $\mathfrak{g}$ to order $k \geq 1$ at (A,ψ) define a smooth section of $\bigotimes_k (iT^*M \oplus \mathbb{S})$. Let $\mathfrak{g}_k|_{(A,\psi)}$ denote the latter. The derivatives of this section to any given order are bounded by an appropriate (A,ψ)–dependent multiple of $\|\mathfrak{g}\|_{\mathcal{P}}$. Here, $\|\cdot\|_{\mathcal{P}}$ denotes the Banach space norm on $\mathcal{P}$. For example, bounds of this sort appear in Theorem 11.6.1 of [16]. See also Proposition 2.5 in [24].

What follows describes some of the simplest nonconstant functions in $\mathcal{P}$. To start, let μ denote a smooth, coclosed 1–form on M. Thus, $d{*}\mu = 0$. Use $\mathfrak{e}_\mu$ to denote the function on $\mathrm{Conn}(E)$ whose value on any given connection A is

$$\mathfrak{e}_\mu(A) = i \int_M \mu \wedge *B_A. \tag{3-11}$$

View $\mathfrak{e}_\mu$ as a function on $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ that is independent of the second factor. Viewed in this light, $\mathfrak{e}_\mu$ is a candidate for a function from $\mathcal{P}$; and this is the case if μ comes from a certain Banach space of coclosed 1–forms. To say more, let Ω_0 denote the vector space of finite linear combinations of coclosed eigenfunctions of the operator $*d$, here sending $C^\infty(M;T^*M)$ to itself. Then $\mathcal{P}$ contains the linear space $\{\mathfrak{e}_\mu : \mu \in \Omega_0\}$. Moreover, the following is true: For each $k \geq 0$, there is a constant c_k such that the C^k norm of any given $\mu \in \Omega_0$ is bounded by $c_k \|\mathfrak{e}_\mu\|_{\mathcal{P}}$.

Note that these norm bounds imply that the function $\mu \to \|\mathfrak{e}_\mu\|_{\mathcal{P}}$ defines a norm on Ω_0; and they imply that the completion of Ω_0 with respect to this norm is a subspace of $C^\infty(M;T^*M)$. Use Ω in what follows to denote this completion. This norm on Ω is called the "$\mathcal{P}$ norm" in what follows.

The perturbations of particular interest in what follows have the from $\mathfrak{e}_\mu + \mathfrak{p}$ where $\mu \in \Omega$ and $\mathfrak{p} \in \mathcal{P}$ with $\|\mathfrak{p}\|_{\mathcal{P}}$ very much smaller than $\|\mathfrak{e}_\mu\|_{\mathcal{P}}$ and with $\|\mathfrak{e}_\mu\|_{\mathcal{P}} \ll 1$. Note that with $\mathfrak{g} = \mathfrak{e}_\mu$, the pair $(\mathfrak{T},\mathfrak{S})$ in (3-6) is $(\mathfrak{T} = i{*}d\mu, \mathfrak{S} = 0)$. In this case, the terms $\mathfrak{t}$ and $\mathfrak{s}$ are absent in (3-7).

3.e Degrees and signs

This subsection has two parts. The first elaborates on the $\mathbb{Z}/p\mathbb{Z}$ degree assignments to the generators of $\mathcal{C}_{\mathrm{SW}}$, and the second elaborates on the signs that are used to define the differential on $\mathcal{C}_{\mathrm{SW}}$.

Part 1 As mentioned above, the relative $\mathbb{Z}/p\mathbb{Z}$ degree between two generators of the Seiberg–Witten Floer complex is defined using the spectral flow for a 1–parameter family of unbounded, self-adjoint operators on $L^2(M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$. To elaborate on the relevant case, suppose that $\mathfrak{c}_-$ and $\mathfrak{c}_+$ are irreducible solutions to some r and $\mathfrak{g}$ version of (3-5) and are such that the respective $\mathfrak{c} = \mathfrak{c}_-$ and $\mathfrak{c} = c_+$ versions of $\mathfrak{L}_{\mathfrak{c}}$ have trivial cokernel. Fix a path $s \to \mathfrak{d}(s) \in \mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ parametrized by $[0, 1]$ such that $\mathfrak{d}(0) = \mathfrak{c}_-$ and $\mathfrak{d}(1) = \mathfrak{c}_+$. If the chosen path $s \to \mathfrak{d}(s)$ is sufficiently generic, then there will be at most one eigenvalue very near 0 at any $s \in [0, 1]$. Such an eigenvalue will have multiplicity 1 and vary smoothly as the parameter s is changed as long as the eigenvalue is sufficiently close to 0. Moreover, if it changes sign as s varies, the eigenvalue crosses zero with nonzero derivative. This understood, the spectral flow for the path is equal to the number of points in $(0, 1)$ where an eigenvalue crosses zero with positive derivative, minus the number where it crosses zero with negative derivative. (See, for example, the author's paper [23].) This spectral flow is denoted in what follows by $f(\mathfrak{c}_-, \mathfrak{c}_+)$. Although a particular path must be chosen to compute this number, the number itself does not depend on the path. However, only the $\mathbb{Z}/p\mathbb{Z}$ reduction of $f(\mathfrak{c}_-, \mathfrak{c}_+)$ is gauge invariant. Granted the preceding, now view $\mathfrak{c}_-$ and $\mathfrak{c}_+$ as generators of $\mathcal{C}_{\mathrm{SW}}$. Then

$$\mathrm{degree}(\mathfrak{c}_+) - \mathrm{degree}(\mathfrak{c}_-) = -f(\mathfrak{c}_-, \mathfrak{c}_+) \mod (p). \tag{3-12}$$

When $c_1(\det(\mathbb{S}))$ is torsion, there will be reducible solutions to (3-5). As noted above, the countable set of cycles that correspond to any given reducible generator can be labeled by a set of the form $\{k, k-2, \ldots\}$ where $k \in \mathbb{Z}$. The relative grading between cycles $k - 2j$ and $k - 2j'$ is $2(j - j')$. Let $\mathfrak{c} \in \mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ denote a pair where $\mathfrak{L}_{\mathfrak{c}}$ has trivial kernel. Then the relative degree difference, $k - \mathrm{degree}(\mathfrak{c})$, is defined to be minus the spectral flow for the family $\{\mathfrak{L}_{\mathfrak{d}(s)}\}_{s \in [0,1]}$ where $\mathfrak{d}\colon [0, 1] \to \mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ is a path that starts at $\mathfrak{c}$ and ends at an irreducible configuration that is very near the reducible one.

Part 2 The signs that appear in (3-4) are also defined using families of operators; in this case the operators that appear in (3-9). A digression is needed first to say more about how this is done (See Chapter 20 of [16]). To start the digression, fix r and $\mathfrak{g}$. Let $\mathfrak{c}_-$ and $\mathfrak{c}_+$ denote elements in $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ where $\mathfrak{L}_{(\cdot)}$ has trivial kernel. Introduce $\mathfrak{P} = \mathfrak{P}(\mathfrak{c}_-, \mathfrak{c}_+)$ to denote the space of piecewise differentiable maps from $\mathbb{R}$ to $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ that have $s \to -\infty$ limit which is gauge equivalent to $\mathfrak{c}_-$ and $s \to \infty$ limit which is gauge equivalent to $\mathfrak{c}_+$. Each $\mathfrak{d} \in \mathfrak{P}$ has its corresponding version of $\mathfrak{D}_\mathfrak{d}$ as given in (3-9); here viewed as a Fredholm operator from $L^2_1(\mathbb{R} \times M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$ to $L^2(\mathbb{R} \times M; iT^*M \oplus \mathbb{S} \oplus i\mathbb{R})$. Quillen [18] explained how such a family of operators can be used to construct a real line bundle, $\det(\mathfrak{D}) \to \mathfrak{P}$. The fiber of $\det(\mathfrak{D})$ at a given $\mathfrak{d} \in \mathfrak{P}$ is canonically identified with $\bigwedge^{\max}(\mathrm{kernel}(\mathfrak{D}_\mathfrak{d})) \otimes_\mathbb{R} (\bigwedge^{\max}\mathrm{cokernel}(\mathfrak{D}_\mathfrak{d}))^*$ if either the kernel or cokernel of $\mathfrak{D}_\mathfrak{d}$ is nontrivial. As explained in [16], this real line bundle is suitably gauge invariant and has gauge invariant orientation. Use $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ to denote the 2–element set of orientations for $\det(\mathfrak{D}) \to \mathfrak{P}$; viewed here as a nontrivial $\mathbb{Z}/2\mathbb{Z}$ module.

Collectively, the modules $\Lambda(\cdot,\cdot)$ have the following three important features: To state the first, let $\mathfrak{c}_-$, $\mathfrak{c}_0$ and $\mathfrak{c}_+$ denote elements in $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ where $\mathfrak{L}_{(\cdot)}$ has trivial kernel. There is in this case the composition law $\Lambda(\mathfrak{c}_-, \mathfrak{c}_0) \otimes_{\mathbb{Z}/2\mathbb{Z}} \Lambda(\mathfrak{c}_0, \mathfrak{c}_+) = \Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$. Second, $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)^* = \Lambda(\mathfrak{c}_+, \mathfrak{c}_-)$. Note that these last two properties imply that $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ can be written as $\Lambda(\mathfrak{c}_-) \otimes_{\mathbb{Z}/2\mathbb{Z}} \Lambda(\mathfrak{c}_+)^*$ where $\Lambda(\cdot)$ is a $\mathbb{Z}/2\mathbb{Z}$ module that is assigned to each gauge equivalence class of pairs $\mathfrak{c} \in \mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ where $\mathfrak{L}_\mathfrak{c}$ has trivial kernel.

To state the final salient property, assume now that both $\mathfrak{c}_-$ and $\mathfrak{c}_+$ are irreducible solutions to (3-5) and that both the $\mathfrak{c} = \mathfrak{c}_-$ and $\mathfrak{c} = \mathfrak{c}_+$ versions of $\mathfrak{L}_\mathfrak{c}$ have trivial kernel. Assume that each $\mathfrak{d} \in \mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$ version of $\mathfrak{D}_\mathfrak{d}$ has trivial cokernel. Then the restriction of $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ to $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$ is canonically isomorphic to the latter's orientation sheaf.

With the digression now over, assume now that r and $\mathfrak{g}$ obey Criteria (3-8) and (3-10). What follows explains how the signs for the differential on $\mathcal{C}_{\mathrm{SW}}$ are determined. Assign to each gauge equivalence class of irreducible solutions to (3-5) an element, $\mathfrak{o}(\cdot)$, in the corresponding version of $\Lambda(\cdot)$. Suppose next that $\mathfrak{c}_-$ and $\mathfrak{c}_+$ are irreducible solutions to (3-5). Then $\mathfrak{o}(\mathfrak{c}_-)\mathfrak{o}(\mathfrak{c}_+) \in \Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ and so defines an orientation for each component of $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+)$. Meanwhile, each 1–dimensional component of this space is oriented by the generator of the $\mathbb{R}$ action that is induced by translation along the $\mathbb{R}$ factor of $\mathbb{R} \times M$. This understood, a given 1–dimensional component contributes $+1$ to $\sigma(\mathfrak{c}_-, \mathfrak{c}_+)$ when these two orientations agree; and it contributes -1 to $\sigma(\mathfrak{c}_-, \mathfrak{c}_+)$ when these two orientations disagree.

3.f Other functions on $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$

Various functions on $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ play a central role in subsequent parts of the story. The first of these is the gauge invariant function, E, on $\mathrm{Conn}(E)$ with value on $A \in \mathrm{Conn}(E)$ given by

$$\mathrm{E}(A) = i \int_M a \wedge *B_A. \tag{3-13}$$

The second function, the *Chern–Simons function*, is a function on $\mathrm{Conn}(E)$. Its definition requires first a choice, A_E, of a fiducial connection on E. It proves useful to choose the latter to be a connection whose curvature 2–form is harmonic. With A_E chosen once and for all, the value of the Chern–Simons function on $A \in \mathrm{Conn}(E)$ is given as follows: Write $A = A_E + \hat{a}_A$. Then

$$\mathfrak{cs}(A) = -\int_M \hat{a}_A \wedge *d\hat{a}_A - 2\int_M \hat{a}_A \wedge *\big(B_E + \tfrac{1}{2}B_{A_K}\big), \tag{3-14}$$

where $*B_E$ is the curvature of A_E and $*B_{A_K}$ is that of A_K. Note that $\mathfrak{cs}$ is fully gauge invariant only in the case where $c_1(\det(\mathbb{S}))$ is a torsion class.

The third of the four functions is denoted by $\mathfrak{a}$. Its critical points are the solutions to (3-5) and the maps that solve (3-6) parametrize the integral curves of its gradient vector field. This function is given by

$$\mathfrak{a} = \tfrac{1}{2}(\mathfrak{cs} - r\mathrm{E}) + \mathfrak{g} + r\int_M \psi^\dagger D_A\psi. \tag{3-15}$$

The definition of the fourth of these functions requires first the choice of a section $\psi_E \in C^\infty(M;\mathbb{S})$ such that the $r=1$, $\mathfrak{g}=0$ and $\mathfrak{c}=(A_E,\psi_E)$ version of $\mathfrak{L}_\mathfrak{c}$ has trivial kernel. With this done, the fourth function, f, is a locally constant function defined off of a codimension 1 subvariety in $\mathrm{Conn}(E)\times C^\infty(M;\mathbb{S})$. Its value on any given (A,ψ) is the spectral flow for the path of operators $\{\mathfrak{L}^s\}_{s\in[0,1]}$ with $\mathfrak{L}^s$ denoting the version of (3-7) that has r^s in lieu of r, $s\mathfrak{g}$ in lieu of $\mathfrak{g}$, and $(A_E + s\hat{a}_A, \psi_E + s(\psi-\psi_E))$ in lieu of (A,ψ).

The respective values of $\mathfrak{cs}$, $\mathfrak{a}$, and f on any given pair (A,ψ) are identical to those on $(A - u^{-1}du, u\psi)$ if $u \in C^\infty(M;S^1$ is homotopically trivial or if $c_1(\det(\mathbb{S}))$ is a torsion class. However, the functions

$$\mathfrak{cs}^f = \mathfrak{cs} - 4\pi^2 f \quad \text{and} \quad \mathfrak{a}^f = \mathfrak{a} - 2\pi^2 f \tag{3-16}$$

are fully gauge invariant. This is to say that their values on any given (A,ψ) are identical to those on $(A-u^{-1}du, u\psi)$ for all $u \in C^\infty(M;S^1)$. Note, however, that $\mathfrak{cs}^f$ and $\mathfrak{a}^f$ are only defined on the complement of the codimension 1 subvariety that consists of the elements $\mathfrak{c} \in \mathrm{Conn}(E)\times C^\infty(M;S^1)$ where $\mathfrak{L}_c$ has a nontrivial kernel.

3.g Special cases

Of principal interest in what follows are the cases of (3-5) and (3-6) where the metric is such that $|a| = 1$ and $*da = 2a$. Thus, the metric comes via an almost complex structure J on the kernel of a such that $da(\cdot, J(\cdot))$ is a metric on kernel(a). Assume as well that the function $\mathfrak{g}$ has the form $\mathfrak{e}_\mu$ with $\mathfrak{e}_\mu$ as in (3-11) as defined using a given 1–form $\mu \in \Omega$. In this case, the equations in (3-5) read

(3-17)
- $B_A - r(\psi^\dagger\tau\psi - ia) - i*d\mu + \frac{1}{2}B_{A_K} = 0.$
- $D_A\psi = 0.$

With $\mu \in \Omega$ fixed, and $r \geq 1$ chosen, $\mathcal{M}^r$ henceforth denotes the space of gauge equivalence classes of solutions to (3-17).

In this case, the equations in (3-6) for instantons are

(3-18)
- $\frac{\partial}{\partial s}A + B_A - r(\psi^\dagger\tau\psi - ia) - i*d\mu + \frac{1}{2}B_{A_K} = 0.$
- $\frac{\partial}{\partial s}\psi + D_A\psi = 0.$

As noted above, the equations in (3-18) assert that (A, ψ) is a critical point of the function

$$\mathfrak{a} = \frac{1}{2}(\mathfrak{cs} - rE) + \mathfrak{e}_\mu + r\int_M \psi^\dagger D_A\psi; \tag{3-19}$$

and the equations in (3-19) assert that the path $s \to (A, \psi)|_s$ is an integral curve of minus the gradient of $\mathfrak{a}$.

3.h Isomorphisms and perturbations

What follows here is an appendix to Section 3 whose purpose is to compare the respective Seiberg–Witten Floer cochain complexes as defined by distinct pairs of metric on M and perturbation. The discussion that follows has five parts.

Part 1 Any given metric on M and suitable perturbation from $\mathcal{P}$ can be used to define a chain complex using solutions to (3-2) and differential using solutions to (3-3). Both the differential and chain complex depend on the chosen metric and perturbation. To say that the resulting cohomology group does not depend on these choices means the following: Suppose that $(g_0, \mathfrak{g}_0)$ and $(g_1, \mathfrak{g}_1)$ are the relevant pairs of metric on TM and suitable perturbation. Let $\mathcal{H}_0$ and $\mathcal{H}_1$ denote the corresponding versions of the cohomology that is defined by the respective $(g_0, \mathfrak{g}_0)$ and $(g_1, \mathfrak{g}_1)$ versions of (3-2) and (3-3). As explained in [16], there is a canonical isomorphism from $\mathcal{H}_1$ to $\mathcal{H}_0$.

This isomorphism is denoted in what follows by $I_{0,1}$. Note that $I_{0,1}$ preserves the $\mathbb{Z}/p\mathbb{Z}$ gradings of these $\mathbb{Z}$–modules. More is said about $I_{0,1}$ in Part 2 immediately below. The $\mathbb{Z}$–modules $\mathcal{H}_0$ and $\mathcal{H}_1$ are said in what follows to be *realizations* of the Seiberg–Witten Floer cohomology.

Let $\mathcal{H}_0$ denote a given realization of the Seiberg–Witten Floer cohomology. Suppose that h is a given $\mathbb{Z}$–module and let $\phi_0\colon h \to \mathcal{H}_0$ denote a homomorphism. Given that any two realizations of the Seiberg–Witten Floer cohomology are canonically isomorphic, this homomorphism ϕ_0 defines a homomorphism from h to $\mathcal{H}^{\mathrm{SW}}$. To say more about what this means, let $\mathcal{H}_1$ denote a second realization of the Seiberg–Witten Floer cohomology and let $\phi_1\colon h \to \mathcal{H}_1$ denote some homomorphism. Then ϕ_1 also defines a homomorphism from h_1 into $\mathcal{H}^{\mathrm{SW}}$. These two homomorphisms from h into $\mathcal{H}^{\mathrm{SW}}$ agree if and only if $\phi_0 = I_{0,1}\phi_1$.

Part 2 The isomorphism $I_{0,1}$ between realizations $\mathcal{H}_0$ and $\mathcal{H}_1$ is obtained via a homomorphism, $\hat{I}_{0,1}$, between the corresponding cochain complexes that intertwines the respective differentials. To construct $\hat{I}_{0,1}$, first fix a smooth, 1–parameter family of metrics $x \to g_x$ parametrized by $x \in [0,1]$ such that $g_x = g_0$ for $x \le \frac{1}{8}$, and $g_x = g_1$ for $x \ge \frac{7}{8}$. Define a family of metrics $\{g_s\}_{s\in\mathbb{R}}$ so that $g_s = g_0$ for $s < 0$, $g_s = g_x$ for $x \in [0,1]$ and $g_s = g_1$ for $s > 1$. Each g_x has a corresponding Banach space of perturbations for use in defining (3-2) and (3-3). These can be viewed as defining a smooth Banach space bundle over the space of metrics. The bundle is denoted by $\mathcal{P}$. A perturbation chosen for a given metric is understood to come from its fiber over $\mathcal{P}$. This understood, the fiber over any given metric is also denoted by $\mathcal{P}$. Granted this sloppy notation, fix an analogous, but generic path $x \to \mathfrak{g}_x \in \mathcal{P}$ with $\mathfrak{g}_x = \mathfrak{g}_0$ for $x \le \frac{1}{8}$ and $\mathfrak{g}_x = \mathfrak{g}_1$ for $x \ge \frac{7}{8}$, and likewise extend this to a family $\{\mathfrak{g}_s\}_{s\in\mathbb{R}}$. Be forewarned that a path $\{\mathfrak{g}_x\}_{x\in[0,1]}$ can be used to construct the desired isomorphism only if certain criteria are met. The latter are much like the ones in Section 3.d. In any event, a reasonably generic path will suffice. A path with the desired properties is said to be *suitable*.

Use the data $\{(g_s, \mathfrak{g}_s)\}_{s\in R}$ to define an s–dependent version of (3-3) where the metric on any given constant s slice of $\mathbb{R} \times M$ is given by g_s, and where the perturbation terms have s–dependence; thus $(\mathfrak{T}, \mathfrak{S})$ at $s \in \mathbb{R}$ is obtained from $\mathfrak{g}_s$. The instanton solution to the resulting version of (3-3) with s–dependent parameters is said to be a *cobordism instanton*. The homomorphism $\hat{I}_{0,1}$ is defined by using an algebraic count of the equivalence classes of cobordism instantons whose corresponding version of (3-9) has Fredholm index equal to zero. For example, in the case when $g_0 = g_1$ and $\mathfrak{g}_0 = \mathfrak{g}_1$, one can take the constant family $x \to (g_0, \mathfrak{g}_0)$. The set of contributing instantons in this case consists of the set of $\mathbb{R}$–invariant instantons, and the equivalence class of each such instanton contributes $+1$ to the count.

Kronheimer and Mrowka, in Chapters 23 and 24 of [16], give the details of the construction of this homomorphism $\widehat{I}_{0,1}$. They explain why it intertwines the corresponding differentials, why the induced homomorphism $I_{0,1}$ between the respective cohomology modules is an isomorphism, and why this $I_{0,1}$ does not depend on the choice of the interpolating data set.

Part 3 A point to note with regards to these cobordism isomorphisms concerns their behavior with respect to composition. To set the context, suppose that $(g_0, \mathfrak{g}_0)$ and $(g_1, \mathfrak{g}_1)$ are two data sets that are suitable for defining the Seiberg–Witten Floer cohomology. Suppose, in addition that $x \to (g_x, \mathfrak{g}_x)$ is a suitable path of data sets that is parametrized by $x \in [0, 1]$, and in particular is a path that can be used to construct the canonical identification $I_{0,1}$ between the respective $(g_0, \mathfrak{g}_0)$ and $(g_1, \mathfrak{g}_1)$ versions of Seiberg–Witten Floer cohomology via a cobordism homomorphism between their cochain complexes. Now consider a finite set of points $0 = x_0 < x_1 < \cdots < x_{N-1} < x_N = 1$ such that each $x = x_N$ version of $(g_x, \mathfrak{g}_x)$ is suitable for defining the Seiberg–Witten Floer cochain complex and its differential. Note that if the family $\{x \to (g_x, \mathfrak{g}_x)\}_{x\in[0,1]}$ is sufficiently generic, there is a residual set of points in $[0, 1]$ with this property. In any event, given this set, there exists for each $k \in \{1, \dots, N\}$ a cobordism isomorphism between the respective Seiberg–Witten Floer cohomology as defined by the $x = x_k$ and $x = x_{k-1}$ versions of $(g_x, \mathfrak{g}_x)$. Let I_k denote the latter. Then the canonical isomorphism $I_{0,1}$ is obtained by composing in the reverse order the isomorphisms in the set $\{I_k\}_{1\le k\le N}$.

Part 4 This part of the digression concerns small changes in the data set. Suppose that $(g, \mathfrak{g})$ are suitable for defining the cochain complex and differential using the solutions to the corresponding versions of (3-2) and (3-3). Suppose in addition that an integer k has been given in the case when $c_1(\det(\mathbb{S}))$ is a torsion class, and that all solutions to (3-2) with degree k or greater are irreducible. The three points that follow should be kept in mind. To set the stage, remark that there is a natural way to identify respective spinor bundles for different metrics by viewing any given metric's orthonormal frame bundles as a submanifold inside the $\mathrm{Gl}(3;\mathbb{R})$ bundle of oriented frames in TM. This leads to a natural way to discuss the manner in which solutions to versions of (3-17) (and also (3-18)) vary as the metric on M is varied.

To state the first point, suppose that $(g', \mathfrak{g}')$ is chosen from a suitable neighborhood of $(g, \mathfrak{g})$. Then there is a 1–1 correspondence between the respective $(g, \mathfrak{g})$ and $(g', \mathfrak{g}')$ solutions to (3-2), at least between those with degree k or greater if $c_1(\det(S))$ is torsion. This correspondence pairs solutions with the same spectral flow. There is also a 1–1 correspondence between the respective sets of $(g,\mathfrak{g})$ and $(g',\mathfrak{g}')$ instantons that are used to define the differential on the corresponding sets of generators (in degree k or greater if $c_1(\det(\mathbb{S}))$ is torsion). The latter correspondence pairs instantons with the same sign contribution to the respective differentials.

The second point elaborates on the first: The first point alludes to correspondences with the following property: Corresponding solutions vary smoothly over the given neighborhood of $(g,\mathfrak{g})$.

To set the stage for the third point, fix $(g',\mathfrak{g}')$ from this neighborhood of $(g,\mathfrak{g})$. The correspondences between solutions of (3-2) and (3-3) extends by linearity to give an isomorphism between the respective cochain complexes that intertwines the respective differentials (in degree k or greater if $c_1(\det(S))$ is torsion.) With this understood, what follows is the third point: The latter isomorphism is realized by a cobordism isomorphism of the sort described in Part 2 of this digression.

Part 5 The relevant metrics for the proof of Theorem 1 are defined by pairs (a,J) of contact 1–form and suitable almost complex structure. The perturbations in question are those of the sort to give (3-5) and (3-6). (In this regard, it is assumed in what follows that the function E from (3-13) is always an element in $\mathcal{P}$.) In general, it is not possible to make do only with the simpler (3-17) and (3-18), at least with regards to the differential. Even so, (3-17) and (3-18) play the prominent role. The three points that follow indicate why (3-17) is sufficient for most purposes.

The first point concerns a partial version of what is said in Part 4. To set the stage, fix $r>0$ and a 1–form μ for use in (3-17). Suppose that $\mathcal{Z}$ is a finite set of distinct, gauge equivalence classes of solution to the (a,J), r, and μ version of (3-17). Suppose in addition that each element in $\mathcal{Z}$ is represented by an irreducible, nondegenerate solution. Let g denote the metric defined by the pair (a,J). Then there exists a neighborhood of the pair $(g,\mathfrak{g}=\mathfrak{e}_\mu)$ with the following significance: Suppose that $(g',\mathfrak{g}')$ comes from this neighborhood. Then there is a 1–1 correspondence between the equivalence classes that comprise $\mathcal{Z}$ and a set of equivalence classes of irreducible, nondegenerate solution to the $(g',\mathfrak{g}')$ version of (3-5). This equivalence lifts to an equivalence between corresponding sets of solutions, and the latter pairs solutions with the same spectral flow. This last correspondence is such that $(g',\mathfrak{g}')$ solutions vary smoothly as $(g',\mathfrak{g}')$ is varied.

The second point concerns $(g',\mathfrak{g}')$ instantons when the $(g,\mathfrak{g})$ instantons that interpolate between elements in $\mathcal{Z}$ are nondegenerate. Make the following assumption: Let $\mathfrak{d}$ denote a $(g,\mathfrak{g})$ instanton whose version of (3-9) has index 1 and has $s\to\pm\infty$ limits in $\mathcal{Z}$. Then $\mathfrak{d}$ is nondegenerate. Assume also that there exist but a finite set of such instantons. If this is the case, there exists perturbations very near g with the following property: First, the perturbation, with g, is suitable for defining the generators of the Seiberg–Witten Floer cochain complex and its differential. Second, the perturbation is zero to any given order at each element in $\mathcal{Z}$; and it is zero to any given order on the path in $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ that is defined by any $(g, \mathfrak{g})$ instanton of the sort just described.

The third point says more about instantons: Suppose that $(g', \mathfrak{g}')$ is sufficiently close to $(g, \mathfrak{g})$. Then, there is a 1–1 correspondence between the set of instantons described in the preceding paragraph and a subset of the $(g', \mathfrak{g}')$ instantons. Each instanton in this subset defines a path between $(g', \mathfrak{g}')$ solutions to (3-5) that correspond to the $(g, \mathfrak{g})$ solutions from $\mathcal{Z}$. The correspondence here has the same properties as its analog in Part 4.

4 Proof of Theorem 1

The purpose of this section is to explain how Theorem 1 follows from a collection of theorems about the large r versions of (3-17) and (3-18). The proof itself is given in the final subsection; the intervening subsections supply the necessary tools.

To set the stage, note that the definition of embedded contact homology requires the choice of a suitable pair (a, J) of contact 1–form from Lemma 2.1's residual set and almost complex structure from $\mathcal{J}_a$. Two different choices can, in principle, define different complexes and/or different differentials. (As noted above, Mike Hutchings conjectured that the resulting homology groups are isomorphic.) Likewise, the definition of the Seiberg–Witten Floer cohomology or homology requires choices. Such choices in this case consist of a 4–tuple $(\hat{a}, \hat{J}, r, \mathfrak{g})$ where $\hat{a}$ is a suitably chosen contact form and $\hat{J}$ is an almost complex structure from $\mathcal{J}_{\hat{a}}$. These are used to define the metric on M; and then $\hat{a}$ is used as the contact 1–form in the corresponding versions of (3-17) and (3-18). Meanwhile $r \geq 1$ is a real number and $\mathfrak{g} \in \mathcal{P}$ is a suitable perturbation term. The choice of this data determines the chain complex and the differential. These can and will differ with differing choices of $(\hat{a}, \hat{J}), r, \mathfrak{g}$.

As noted in Section 3.h, any two such choices give isomorphic realizations of the Seiberg–Witten Floer cohomology. In any event, the set of generators can be defined by using only solutions to (3-17). Given what is said in Part 5 of Section 3.h, this follows from:

Proposition 4.1 *Suppose that (a, J) is a given pair of contact 1–form and almost complex structure. Use the latter to define the metric, g, on M. If $c_1(\det(\mathbb{S}))$ is torsion, then fix an integer k. There exists $r_k > 0$ with the following significance:*

- *Fix $r > r_k$ and a 1–form μ with $\mathcal{P}$–norm less than 1. Then all solutions to the corresponding version of (3-17) are irreducible if $c_1(\det(\mathbb{S}))$ is not torsion. If $c_1(\det(\mathbb{S}))$ is torsion, then all solutions to (3-17) with degree k or less are irreducible.*
- *Fix $r > r_k$. If a form μ with $\mathcal{P}$–norm less than 1 is chosen from an appropriate residual set, then all of the aforementioned solutions to (3-17) are nondegenerate.*
- *The form μ can be chosen with $\mathcal{P}$ norm less than 1 from an appropriate residual set so that the following is true: There is a nonaccumulating, discrete set $\mathcal{U} \in [r_k, \infty)Q$ such that if $r > r_k$ is not in $\mathcal{U}$, then all irreducible solutions to the r and μ version of (3-17) are nondegenerate.*

Proof of Proposition 4.1 This summarizes results from Section 3 in [24] and Section 4 in [25]. □

Granted this proposition, it follows from the observations of Part 5 in Section 3.h that there are suitable perturbations of the form $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ where $\mathfrak{p}$ has the following properties: First, it comes from any given neighborhood of 0 in $\mathcal{P}$. Second, it vanishes to any desired order on the irreducible solutions to the (a, J), r and μ version of (3-17).

4.a From Reeb orbits to Seiberg–Witten solutions on M

Suppose that $(\hat{a}, \hat{J})$ is a pair of contact structure from Lemma 2.1's residual set and almost complex structure in $\mathcal{J}_{\hat{a}}$. Fix L. The chain complex $\mathcal{C}_{\mathrm{ech}}^L$ has generators that are equivalence classes of pairs of the form $(\Theta, \mathfrak{o})$ where Θ is a set whose typical element is a pair of Reeb orbit and positive integer subject to various constraints. It proves useful to introduce now $\mathcal{Z}_{\mathrm{ech}}^L$ to denote the set of such Θ. Assume now the following:

Suppose that $\hat{a}$ is a contact structure on M that defines the given orientation. Given $L \geq 1$, say that $\hat{a}$ is L–nondegenerate when two requirements are met. These requirements refer to finite sets of the following sort: Let Θ denote the set in question. This is a finite collection of pairs whose elements have the form (γ, m) where γ is a Reeb orbit and m is a positive integer. Moreover, no two pairs share the same Reeb orbit. The first requirement on Θ demands that $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma \neq L$. To state the second requirement, assume that this last sum is less than L. Require that all Reeb orbits that appear in a pair from Θ are nondegenerate; and that if $(\gamma, m) \in \Theta$ with γ elliptic, then γ is m–elliptic.

It proves useful to introduce $\mathcal{Z}_{\text{ech}}^L$ to denote the collection of sets Θ of the sort just described, and with $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma < L$. With $L \geq 1$ chosen, suppose that $(\hat{a}, \hat{J})$ is such that $\hat{a}$ is an L–nondegenerate contact form and that $\hat{J}$ is an almost complex structure on the kernel of $\hat{a}$ such that $d\hat{a}(\cdot, \hat{J}\cdot)$ is positive definite. Assume now the following:

(4-1)

- There is no element $\Theta \in \mathcal{Z}_{\text{ech}}^L$ with $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma = L$.
- Suppose that γ is a Reeb orbit with $\ell_\gamma < L$. Then γ has a tubular neighborhood map $\varphi\colon S^1 \times D \to M$ as described in Section 2.a such that if γ is hyperbolic with rotation number k, then $(\nu, \mu) = (\frac{1}{4}k, i\varepsilon e^{ikt})$ with $\varepsilon > 0$ but very small. Meanwhile, if γ is elliptic, then its rotation number R is irrational. Furthermore:
  - (i) The pair $(\nu, \mu) = (\frac{1}{2}\mathrm{R}, 0)$.
  - (ii) The φ^*–pull back of $T^{1,0}(\mathbb{R} \times M)$ is spanned by $ds + ia$ and $\frac{\ell_\gamma}{2\pi}(dz - i\mathrm{R}z\, dt)$.

  Moreover, these two forms are orthogonal and have norm $\sqrt{2}$.

Fix a large value for r and a 1–form μ from the $(\hat{a}, \hat{J})$ version of Ω with $\mathcal{P}$ norm less than 1. Use $\mathcal{M}^r$ to denote the set of equivalence classes of solutions to the corresponding version of (3-17).

The theorem that follows asserts the existence of a map from $\mathcal{Z}_{\text{ech}}^L$ into the version of $\mathcal{M}^r$ that is defined using $(\hat{a}, \hat{J})$, μ, and a sufficiently large r, This map is used to define the isomorphism for Theorem 1.

Theorem 4.2 *Fix $L \geq 1$ and a pair $(\hat{a}, \hat{J})$ as described above that obeys (4-1). There exists $\kappa \geq 1$ with the following significance: Define $\mathcal{M}^r$ using $r \geq \kappa$ and a 1–form $\mu \in \Omega$ with $\mathcal{P}$ norm bounded by 1. There exists a map $\Phi^r\colon \mathcal{Z}_{\text{ech}}^L \to \mathcal{M}^r$ with the three properties listed below.*

- *Φ^r is a bijection onto the subset in $\mathcal{M}^r$ of elements $\mathfrak{c} = (A, \psi)$ with $\mathrm{E}(A) < 2\pi L$.*
- *If $\mathfrak{c} \in \mathcal{M}^r$ is in the image of Φ^r, then the operator $\mathfrak{L}_{\mathfrak{c}}$ has trivial kernel.*
- *Let Θ and Θ' denote any two elements in $\mathcal{Z}_{\text{ech}}^L$ and let z and z' denote their respective $\mathbb{Z}/p\mathbb{Z}$ degrees. Meanwhile, let x and x' denote the respective $\mathbb{Z}/p\mathbb{Z}$ degrees of $\Phi^r(\Theta)$ and $\Phi^r(\Theta')$. Then $x - x' = -(z - z')$ modulo $p\mathbb{Z}$.*

The image of any given element $\Theta \in \mathcal{Z}_{\mathrm{ech}}^L$ *via the map* Φ^r *can be characterized in part as follows: Let* $\gamma \in M$ *denote a Reeb orbit with* $\ell_\gamma < L$ *and let* $D \subset M$ *denote a transverse disk of the sort described by (4-1). Let* (A, ψ) *denote a solution to (3-17) that defines the equivalence class* $\Phi^r(\Theta)$*. Then the integral of* $(i/2\pi) * B_A$ *over the concentric disk in* D *with radius* $r^{-1/4}$ *is bounded in absolute value by* $\kappa r^{-1/2}$ *unless* γ *comes from a pair* $(\gamma, m) \in \Theta$*. In this case, the integral differs from* m *by at most* $\kappa r^{-1/2}$*.*

The upcoming Section 5 says more about what Φ^r looks like. The actual construction of Φ^r is in Paper II of this series [26]. The assertion that it defines a bijection as described by the Theorem's first bullet is proved in Paper IV of the series [28]. The assertions of the second and third bullets are proved in Paper III of the series [23].

A rather more complicated version of Theorem 4.2 holds when the second item in (4-1) is not invoked. In the latter case, each $\Theta \in \mathcal{Z}_{\mathrm{ech}}^L$ parametrizes a subset of $\mathcal{M}^r$ such that the collection of these subsets elements $\mathfrak{c} = (A, \psi)$ with $E(A) < 2\pi L$. More is said in Section 5 about this more general version of Theorem 4.2.

4.b From pseudoholomorphic curves to Seiberg–Witten solutions on $\mathbb{R} \times M$

Fix $L \geq 1$. Suppose that $(\hat{a}, \hat{J})$ and $\mu \in \Omega$ are suitable input for Theorem 4.2. Any given large r version of Theorem 4.2's map Φ^r identifies $\mathcal{Z}_{\mathrm{ech}}^L$ with a subset in $\mathcal{M}^r$. Define $\mathcal{C}_{\mathrm{ech}}^L$ from the elements $\mathcal{Z}_{\mathrm{ech}}^L$ as in Section 2.c. The map Φ^r can be used to define a monomorphism from $\mathcal{C}_{\mathrm{ech}}^L$ into the Seiberg–Witten cochain complex. This is done as follows: Let $\Theta \in \mathcal{Z}_{\mathrm{ech}}^L$. Order the subset of pairs $(\gamma, 1) \in \Theta$ for which γ is hyperbolic with even rotation number. Doing so for all such Θ identifies $\mathcal{Z}_{\mathrm{ech}}^L$ with a set of generators of $\mathcal{C}_{\mathrm{ech}}^L$. The image of $\mathcal{Z}_{\mathrm{ech}}^L$ via Φ^r defines a set of generators of $\mathcal{C}^{\mathrm{SW}}$. Extend this map of generators in a $\mathbb{Z}$–linear fashion. The monomorphism so constructed is canonical up to precomposing by an isomorphism from $\mathcal{C}_{\mathrm{ech}}^L$ to itself that changes the sign of some of the generators. A monomorphism that is obtained from Φ^r in this way is denoted by T_Φ.

The next theorem asserts in part that there is a choice for T_Φ that intertwines the embedded contact homology differential with the Seiberg–Witten Floer cohomology differential. This theorem reintroduces the space $\mathcal{M}_1(\Theta_-, \Theta_+)$ from Section 2.c and the space $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ from Section 3.b. Note that both spaces admit a canonical $\mathbb{R}$–action, this induced by the action of $\mathbb{R}$ on $\mathbb{R} \times M$ as the group of constant translations of the $\mathbb{R}$ factor in $\mathbb{R} \times M$.

The upcoming Theorem 4.3 is used subsequently to prove that there are suitable pairs $(\hat{a}, \hat{J})$ with a version of T_Φ that intertwines the embedded contact homology differential with the Seiberg–Witten Floer cohomology differential. What follows directly sets the stage.

To start, the theorem assumes implicitly that $\hat{a}$ is suitably generic, which is to say $\hat{a} \in \mathcal{N}_M$. It also assumes implicitly that $\hat{J} \in \mathcal{J}_a$. Needless to say, it assumes (4-1). Granted the preceding assumptions, introduce the space $\mathcal{M}_1(\Theta_-, \Theta_+)$ from Section 2.c as defined for pairs $(\Theta_-, \Theta_+) \in \mathcal{Z}_{\mathrm{ech}}^L$.

Theorem 4.3 uses $(\hat{a}, \hat{J})$, a 1–form $\mu \in \Omega$ with $\mathcal{P}$–norm less than 1, and a specified value for r to define the space $\mathcal{M}^r$ and Theorem 4.2's map Φ^r. Theorem 4.3 then reintroduces the space $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ from Section 3.b as defined for solutions $\mathfrak{c}_-$ and $\mathfrak{c}_+$ of (3-17) that define the respective equivalence classes $\Phi^r(\Theta_-)$ and $\Phi^r(\Theta_+)$. Recall that this is the space of instanton solutions to $((\hat{a}, \hat{J}), r, \mu)$ version of (3-18) with $s \to -\infty$ limit equal to $\mathfrak{c}_-$, with $s \to \infty$ limit $\mathfrak{c}_+$, and with (3-9) defining an index 1 Fredholm operator.

The theorem also discusses certain perturbed versions of $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$. The perturbed version is defined using a suitable, small element $\mathfrak{p} \in \mathcal{P}$ that vanishes to second order on $\mathfrak{c}_-$ and $\mathfrak{c}_+$. This perturbed version is denoted by $\mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+)$. An element in $\mathfrak{d} \in \mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+)$ is an instanton solution to the $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ version of (3-6) with $s \to -\infty$ limit $\mathfrak{c}_-$, with $s \to \infty$ limit gauge equivalent to $\mathfrak{c}_+$, and with (3-9) again defining a Fredholm operator with index 1. As per what is said in Part 5 of Section 3.h, there exists a residual set of such perturbations that make the data $(\hat{a}, \hat{J})$ with r, and $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ suitable for defining the Seiberg–Witten Floer cochain complex.

Keep in mind that $\mathcal{M}_1(\Theta_-, \Theta_+)$ and any $\mathfrak{p} \in \mathcal{P}$ version of $\mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+)$ admits a canonical $\mathbb{R}$–action, which is induced by the action of $\mathbb{R}$ on $\mathbb{R} \times M$ as the group of constant translations of the $\mathbb{R}$ factor in $\mathbb{R} \times M$.

Theorem 4.3 *Fix $L \geq 1$ and a pair $(\hat{a}, \hat{J})$ as described above that obeys (4-1). There exists $\kappa \geq 1$ with the following significance: Define the space $\mathcal{M}^r$ using the pair $(\hat{a}, \hat{J})$, a 1–form $\mu \in \Omega$ with $\mathcal{P}$ norm bounded by 1, and $r \geq \kappa$.*

- *Let Θ_- and Θ_+ denote any two elements in $\mathcal{Z}_{\mathrm{ech}}^L$. Use $\mathfrak{c}_-$ and $\mathfrak{c}_+$ to denote solutions to (3-17) whose gauge equivalence classes are the respective images in $\mathcal{M}^r$ of Θ_- and Θ_+ via Theorem 4.2's map Φ^r.*
  - (i) *The space $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ has a finite set of components, and each component is an orbit of the canonical $\mathbb{R}$ action. In addition, if $\mathfrak{d} \in \mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$, then the corresponding $\mathfrak{g} = \mathfrak{e}_\mu$ version of (3-9) has trivial cokernel.*
  - (ii) *There is an $\mathbb{R}$–equivariant diffeomorphism, Ψ^r, from $\mathcal{M}_1(\Theta_-, \Theta_+)$ to $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$.*

(iii) *Let $\mathfrak{p} \in \mathcal{P}$ denote a sufficiently small element that vanishes to second order on the image of Φ^r. Define $\mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+)$ as above. There is also an $\mathbb{R}$–equivariant diffeomorphism from $\mathcal{M}_1(\Theta_-, \Theta_+)$ to $\mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+)$.*

- *Let $\mathfrak{p} \in \mathcal{P}$ be as described in item* (iii) *above, and such that the data $(\hat{a}, \hat{J})$, r and $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ is suitable for defining the Seiberg–Witten Floer cochain complex. There is a choice for T_Φ such that if Θ_- and Θ_+ now denote any given pair of generators of $\mathcal{C}^L_{\mathrm{ech}}$, then the contribution, $+1$ or -1, of any given component in $\mathcal{M}_1(\Theta_-, \Theta_+)$ to the embedded contact homology differential is the same as the contribution of its image via item* (ii)*'s diffeomorphism to the Seiberg–Witten Floer cohomology differential.*

A rough picture of Ψ^r is given in Section 5. The full construction is in Paper II of this series [26]. The proof that Ψ^r is an $\mathbb{R}$–equivariant embedding onto an open set is given in Paper III [24]. The assertion that Ψ^r is a surjection is given as Theorem 1.2 in Paper IV [28] and proved in the latter's Sections 3–7. The proof of item (iii) is in Section 8 of [28]. The proof of the assertion in the theorem's second bullet is in Paper III. With regards to the second bullet, Paper III discusses only the case for Ψ^r. Given what is said in Part 4 of Section 3.h and given the first bullet of the theorem, the assertion for the general case follows directly from the assertion for the special case of Ψ^r. Moreover, given items (i) and (ii) of the first bullet, it follows from what is said in Part 5 of Section 3.h that it is sufficient with regards to the second bullet to consider elements $\mathfrak{p} \in \mathcal{P}$ that vanish to second order along the image of all paths in $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ of the form $\{s \to \mathfrak{d}(s)\}_{s\in\mathbb{R}}$ where $\mathfrak{d}$ is an instanton solution to (3-18) from the set $\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$. If this is assumed, then $\mathfrak{M}_{1,\mathfrak{p}}(\mathfrak{c}_-, \mathfrak{c}_+) = \mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+)$ and item (iii)'s diffeomorphism is item (ii)'s map Ψ^r.

Theorem 4.3 has a replacement of sorts when the second item in (4-1) is not present. The latter is vastly more complicated to state, let alone prove. More is said on this score in Section 5.

The next theorem asserts that T_Φ can be chosen so that any sufficiently large r version of $T_\Phi(\mathcal{C}^L_{\mathrm{ech}})$ is mapped to itself by the Seiberg–Witten Floer cohomology differential. Of course, the definition of this differential may require a perturbation term. The latter is taken to have the form $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ with $\mathfrak{p} \in \mathcal{P}$ very small. As noted in Part 5 of Section 3.h, it is permissible to choose $\mathfrak{p}$ so as to vanish to second order on the image of any given large r version of Φ^r. As noted in the preceding paragraph, it is likewise permissible to choose $\mathfrak{p}$ to vanish along the image of all paths in $\mathrm{Conn}(E) \times C^\infty(M; \mathbb{S})$ of the form $\{s \to \mathfrak{d}(s)\}_{s\in\mathbb{R}}$ where $\mathfrak{d}$ is an instanton solution to (3-18) taken from the set $\{\mathfrak{M}_1(\mathfrak{c}_-, \mathfrak{c}_+) : \mathfrak{c}_+$ and $\mathfrak{c}_-$ are in the image of $\Phi^r\}$. This constraint is not strictly necessary; it is imposed implicitly in what follows only to simplify the presentation.

Theorem 4.4 *Fix $L \geq 1$ and a pair $(\hat{a}, \hat{J})$ as above for use in Theorem 4.2. This is to say that $\hat{a}$ is L–nondegenerate and the pair obeys (4-1). There exists $\kappa \geq 1$ with the following significance: Define $\mathcal{M}^r$ using $r \geq \kappa$ and a 1–form $\mu \in \Omega$ with $\mathcal{P}$ norm bounded by 1. Fix a very small normed element $\mathfrak{p} \in \mathcal{P}$ with the properties described above. Assume that $((\hat{a}, \hat{J}), r, \mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p})$ is suitable for defining the Seiberg–Witten Floer cochain complex and its differential. Let $\mathcal{C}^{\mathrm{SW}}$ denote this cochain complex.*

- *Let $\mathfrak{d}$ denote a $((\hat{a}, \hat{J}), r, \mathfrak{e}_\mu + \mathfrak{p})$ instanton with $s \to \infty$ limit in the image of Φ^r. Then the $s \to -\infty$ limit of $\mathfrak{d}$ is also in Φ^r.*
- *Assume in addition that $\hat{a} \in \mathcal{N}_M$ and that $\hat{J} \in \mathcal{J}_a$ so that the conditions for Theorem 4.3 are met. Then the monomorphism T_Φ can be chosen so as to intertwine the action of the boundary operator on $\mathcal{C}^L_{\mathrm{ech}}$ with the action of the differential on $\mathcal{C}^{\mathrm{SW}}$.*

Theorem 4.4 implies that when $(\hat{a}, \hat{J})$ is suitable for Theorem 4.3 and r is large, then there is a version of the monomorphism T_Φ that intertwines the embedded contact homology differential with the differential that defines the Seiberg–Witten Floer cohomology. As such, this T_Φ identifies $\mathcal{C}^L_{\mathrm{ech}}$ as a subcomplex in the Seiberg–Witten Floer cochain complex.

Theorem 4.4 is proved in the upcoming Section 4.h.

4.c The image of T_Φ in $\mathcal{H}^{\mathrm{SW}}$

Fix a pair (a, J) of contact 1–form from Lemma 2.1's residual set and almost complex structure from $\mathcal{J}_a$. Use this data to define the embedded contact homology chain $\mathbb{Z}$–module $\mathcal{C}_{\mathrm{ech}}$ and its differential. The latter has a filtration $\{\mathcal{C}^L_{\mathrm{ech}}\}_{L\geq 1}$ with the corresponding homology groups. Fix $L \geq 1$ such that the top item in (4-1) holds for the pair (a, J). Fix small $\delta > 0$ and Proposition 2.5 supplies a (δ, L) approximation, $(\hat{a}, \hat{J})$, to (a, J). The latter defines the analogous set of $\mathbb{Z}$–modules $\{\hat{\mathcal{C}}^{L'}_{\mathrm{ech}}\}_{L'\geq 1}$ and corresponding homology groups. Let $\hat{\mathcal{Z}}^L_{\mathrm{ech}}$ denote the $\hat{\mathcal{C}}^L_{\mathrm{ech}}$ analog of $\mathcal{Z}^L_{\mathrm{ech}}$. In this regard, the elements in $\hat{\mathcal{Z}}^L_{\mathrm{ech}}$ are geometrically identical to those in $\mathcal{Z}^L_{\mathrm{ech}}$. As noted in Proposition 2.4, the corresponding versions of $\mathcal{H}^L_{\mathrm{ech}}$ are canonically isomorphic.

The upcoming Theorem 4.5 refers to classes in the realization of the Seiberg–Witten Floer cohomology as defined $(\hat{a}, \hat{J})$, a 1–form μ with $\mathcal{P}$–norm bounded by 1, a very large r, and a suitable perturbation of the form $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p})$, with $\mathfrak{p}$ having very small norm. Theorem 4.5 assumes implicitly that r is large enough to invoke Theorems 4.2–4.4. As in Theorem 4.4, it proves convenient to require $\mathfrak{p}$ to vanish to second order on the image of the approximation, $(\hat{a}, \hat{J})$, r and μ version of Φ^r. Given what is said in Theorem 4.4 and Part 5 of Section 3.h, it is also permissible to choose $\mathfrak{p}$ to vanish along the image of all paths in $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ of the form $\{s \to \mathfrak{d}(s)\}_{s\in\mathbb{R}}$ where $\mathfrak{d}$ is an instanton solution to (3-18) from the set $\{\mathfrak{M}_1(\cdot, \mathfrak{c}_+) : \mathfrak{c}_+$ is in the image of $\Phi^r\}$. This constraint is also assumed implicitly, again to simplify the presentation.

Theorem 4.5 *Fix a pair (a, J) where a is a contact 1–form from Lemma 2.1's residual set, and where $J \in \mathcal{J}_a$. Also, fix $k \in \mathbb{Z}$ when $c_1(\det(\mathbb{S}))$ is a torsion class. There exist constants $\kappa \geq \kappa_* \geq 1$ with the following significance: Fix $L_* \geq \kappa_*$ and there exists $L \geq \kappa$ such that the (a, J) and L version of the first item in (4-1) holds. Fix $\delta \in (0, \kappa^{-1})$; then choose a (δ, L) approximation, $(\hat{a}, \hat{J})$, to the pair (a, J). Choose a 1–form $\mu \in \Omega$ with $\mathcal{P}$ norm bounded by 1 as described in Proposition 4.1. Take r very large and fix a suitable perturbation of the form described above with $\mathfrak{p}$ very small, and so that the data $(\hat{a}, \hat{J})$, r and $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ are suitable for defining the Seiberg–Witten Floer cochain complex and differential. Use $\mathcal{C}^{\mathrm{SW}}$ to denote this cochain complex, and use $T_\Phi\colon \hat{\mathcal{C}}^L_{\mathrm{ech}} \to \mathcal{C}^{\mathrm{SW}}$ to denote the monomorphism given by the third bullet in Theorem 4.4.*

- *Let θ denote any Seiberg–Witten Floer cohomology class, but of degree k or more if $c_1(\det(\mathbb{S}))$ is a torsion class. The class θ has a representative cocycle in $T_\Phi(\hat{\mathcal{C}}^{L_*}_{\mathrm{ech}})$.*
- *Suppose that $\upsilon \in \hat{\mathcal{C}}^{L_*}_{\mathrm{ech}}$ and that $T_\Phi(\upsilon)$ is a coboundary. If $c_1(\det(\mathbb{S}))$ is a torsion class, also assume that υ has degree k or less. Then υ is a boundary in $\hat{\mathcal{C}}^{L}_{\mathrm{ech}}$.*

This theorem is also proved in Section 4.h.

4.d Filtrations of the Seiberg–Witten Floer cohomology

This section and Sections 4.e–4.g supply the background material that is used subsequently to prove Theorems 4.4 and 4.5. To begin, let (a, J) denote a pair consisting of a contact 1–form and compatible almost complex structure. Fix $\mu \in \Omega$ with $\mathcal{P}$ norm less than 1, fix r very large and choose a very small normed element $\mathfrak{p} \in \mathcal{P}$ so that the data $(\hat{a}, \hat{J})$, r and $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$ are suitable for defining the Seiberg–Witten Floer cochain complex. Use $\mathcal{C}^{\mathrm{SW}}$ in what follows to denote this cochain complex, and use δ^* to denote the differential. Use $\mathfrak{g}$ also to define the function $\mathfrak{a}$ in (3-15) and the function $\mathfrak{a}^f$ in (3-16). Suppose that $\mathfrak{A} \in \mathbb{R}$ is such that the following is true:

- If $c_1(\det(\mathbb{S}))$ is a torsion class, then (3-5) has no solutions with $\mathfrak{a} = r\mathfrak{A}$.

(4-2)
- If $c_1(\det(\mathbb{S}))$ is not torsion, then (3-5) has no solutions with $\mathfrak{a}^f \in [r\mathfrak{A} - 1, r\mathfrak{A}]$.

Given that δ^* increases $\mathfrak{a}$ and $\mathfrak{a}^f$, this cochain complex has as a subcomplex the set $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ that is generated by the elements in $\mathcal{M}^r$ with $\mathfrak{a}^f > r\mathfrak{A}$ or $\mathfrak{a} > r\mathfrak{A}$ as the case may be. The short exact sequence

$$0 \to \mathcal{C}^{\mathrm{SW},\mathfrak{A}} \to \mathcal{C}^{\mathrm{SW}} \to \mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}} \to 0 \tag{4-3}$$

induces a corresponding long exact sequence in cohomology, thus

$$\cdots \to \mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}}) \to \mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}) \to \mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}}) \to \mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}}) \to \cdots . \tag{4-4}$$

The exact sequence in (4-4) is also stable with respect to certain relatively large changes in the data set. To elaborate, suppose that $\{(a_x, J_x)\}_{x\in[0,1]}$ is a smoothly parametrized family with each term consisting of a contact 1–form and compatible almost contact structure. Fix a smoothly parametrized family $\{r_x\}_{x\in[0,1]}$ with all members very large. Let $\{\mu_x\}_{x\in[0,1]}$ denote a smoothly parametrized family of forms in Ω with $\mathcal{P}$ norm less than 1. Assume that (4-2) holds for the version of (3-17) that is defined by each of the data sets from the family $\{((a_x, J_x), r_x, \mathfrak{g}_x = \mathfrak{e}_{\mu_x})\}_{x\in[0,1]}$. Now let $\{\mathfrak{p}_x\}_{x\in[0,1]}$ denote a family in $\mathcal{P}$ with very small norm. Note that if each $x \in [0,1]$ version of $\mathfrak{p}_x$ has sufficiently small norm, then (4-2) will also hold when (3-5) is defined by any $x \in [0,1]$ version of $((a_x, J_x), r_x, \mathfrak{g}_x = \mathfrak{e}_{\mu_x} + \mathfrak{p}_x)$. This constraint on the norms is assumed implicitly in what follows. Assume that $((a_0, J_0), r_0, \mathfrak{g}_0 = \mathfrak{e}_{\mu_0} + \mathfrak{p}_0)$ and $((a_1, J_1), r_1, \mathfrak{g}_1 = \mathfrak{e}_{\mu_1} + \mathfrak{p}_1)$ are suitable for defining the Seiberg–Witten Floer cochain complex and differential.

Granted these assumptions, there are corresponding versions of (4-3) and also (4-4) for the data sets $((a_0, J_0), r_0, \mathfrak{g}_0 = \mathfrak{e}_{\mu_0} + \mathfrak{p}_0)$ and $((a_1, J_1), r_1, \mathfrak{g}_1 = \mathfrak{e}_{\mu_1} + \mathfrak{p}_1)$.

Lemma 4.6 *Suppose that the family of data sets*

$$\{\mathrm{D}_x = ((a_x, J_x), r_x, \mathfrak{g}_x = \mathfrak{e}_{\mu_x}) + \mathfrak{p}_x\}_{x\in[0,1]}$$

has all of the properties just described.

- *If each r_x is sufficiently large, and each $\mathfrak{p}_x$ is sufficiently small and suitably generic, then the resulting cobordism homomorphism between the respective $x = 0$ and $x = 1$ realizations of the Seiberg–Witten Floer cochain complex preserves (4-3) and induces an isomorphism between the respective $x = 0$ and $x = 1$ terms in (4-4) that intertwines the arrows.*
- *The isomorphisms described in the preceding bullet between the respective terms in (4-4) are invariant under homotopies in the following sense: Assume that*
  (i) *The families*

$$\{\mathrm{D}_x^0\}_{x\in[0,1]} \quad \text{and} \quad \{\mathrm{D}_x^1\}_{x\in[0,1]}$$

are 1–parameter families of data sets of the sort considered above.

(ii) *The family*

$$\big\{y \to \{\mathrm{D}_{(x,y)} = ((a_{(x,y)}, J_{(x,y)}), r_{(x,y)}, \mathfrak{g}_{(x,y)} = \mathfrak{e}_{\mu_{(x,y)}} + \mathfrak{p}_{(x,y)})\}_{x\in[0,1]}\big\}_{y\in[0,1]}$$

is a two parameter family of data sets with respective $y = 0$ and $y = 1$ members $\{\mathrm{D}^0_x\}_{x\in[0,1]}$ and $\{\mathrm{D}^1_x\}_{x\in[0,1]}$, and such that both $\{\mathrm{D}_{(0,y)}\}_{y\in[0,1]}$ and $\{\mathrm{D}_{(1,y)}\}_{y\in[0,1]}$ are constant. Require that (4-2) holds when (3-5) is defined by each $(x, y) \in \times_2[0, 1]$ version of $D_{(x,y)}$.

Then the respective isomorphisms between the terms in (4-4) that are defined by $\{\mathrm{D}^0_x\}_{x\in[0,1]}$ and $\{\mathrm{D}^1_x\}_{x\in[0,1]}$ agree.

Remark that the assertion in the second bullet is not used in what follows.

Proof of Lemma 4.6 The proof has nine steps. Steps 1–6 prove the first bullet and Steps 7–9 prove the second bullet.

Step 1 Given what is said in Part 2 of Section 3.h, the family $\{\mathrm{D}_x\}_{x\in[0,1]}$ defines a version of the data needed to define a cochain homomorphism, $\hat{I}_{0,1}$, from the $x = 1$ realization of the Seiberg–Witten cochain complex to the $x = 0$ realization. With what is said in Part 3 of Section 3.h in mind, choose $\{\mathfrak{p}_x\}_{x\in(0,1)}$ in a suitably generic fashion so as to make the following construction: Fix some very large integer, N, and break up the interval $[0, 1]$ into N intervals of length at most $2/N$ to factor the canonical isomorphism as the concatenation of a sequence of isomorphisms $\{I_k\}_{1\le k\le N}$, this as described in Part 3 of Section 3.h. By way of reminder, the interval is broken into segments $\{[x_{k-1}, x_k]\}_{1\le k\le N}$ with $x_0 = 0$ and $x_1 = 1$, and then any given $k \in \{1, \dots, N\}$ version of I_k is constructed as described in Part 2 of Section 3.h from a corresponding cochain homomorphism, $\hat{I}_k$, from the $x = x_k$ realization of the Seiberg–Witten Floer cochain complex to the $x = x_{k-1}$ realization. Note that the points $\{x_k\}_{1\le k\le N}$ are chosen so that each $x \in \{x_k\}_{1\le k\le N}$ version of D_x is suitable for defining the Seiberg–Witten Floer cochain complex. In particular, each defines a version of (4-4).

Step 2 Recall that each $\hat{I}_k$ is defined using the cobordism instantons for the $x \in [x_{k-1}, x_k]$ portion of $[0, 1]$. In particular, the definition uses only cobordism instantons whose version of (3-9) has Fredholm index zero. If N is large, then the following is true: Let $\mathfrak{d}$ denote a cobordism instanton that contributes to $\hat{I}_k$. Let $\mathfrak{c}_-$ and $\mathfrak{c}_+$ denote the respective $s \to -\infty$ and $s \to \infty$ limits of $\mathfrak{d}$. If the $x = x_{k-1}$ version of $\mathfrak{a}$ on $\mathfrak{c}_-$ is less than that of the $x = x_k$ version of $\mathfrak{a}$ on $\mathfrak{c}_+$, then their difference will be very small in absolute value. That this is so follows by considering the alternative, and constructing a suitable limit as $N \to \infty$ from a sequence of cobordism instantons on shorter and shorter intervals with a uniformly negative change in $\mathfrak{a}$. The latter will converge in a suitable fashion to what Kronheimer and Mrowka refer to as a broken trajectory for some fixed $x \in [0,1]$ data set. Each segment of such a broken trajectory corresponds to an instanton for this data set, and instantons do not decrease $\mathfrak{a}$. However the sum of the decreases in $\mathfrak{a}$ over the broken trajectory will be bounded away from zero due to the initial assumption. This contradiction proves the point in question.

Step 3 As noted above, $\hat{I}_k$ is defined using cobordism instantons that define index zero versions of (3-9). Given this fact, and given what is said in Step 2, it follows that $\hat{I}_k$ can only decrease the corresponding versions of $\mathfrak{a}^f$ a small amount if N is large. This and (4-2) have the following consequence: Each $\hat{I}_k$ intertwines the respective k and $k-1$ versions of (4-3) and so defines a homomorphism between the respective versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ and between respective versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$. This last observation implies that $I_{0,1}$ defines a homomorphism between the respective $x=0$ and $x=1$ versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$, and between the respective versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$.

Step 4 This step explains why the homomorphism that $I_{0,1}$ defines between the $x=0$ and $x=1$ versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ is an isomorphism. Lemma 4.6 follows from this fact given that $I_{0,1}$ defines an isomorphism between the respective versions of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}})$. Use $I_{0,1*}$ to denote this restriction of $I_{0,1}$ to $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$.

Given that $I_{0,1}$ can be factored as in Steps 2 and 3, it is sufficient to consider $I_{0,1*}$ in the case when D_0 and D_1 are nearly identical. This assumption enters momentarily. In any case, define now a 2–parameter family of data sets $\{\mathrm{D}_{(x,y)}\}_{(x,y)\in[0,1]\times[0,1]}$ as follows: For $x \in [0,\frac{1}{2}]$, set $\mathrm{D}_{(x,y)} = \mathrm{D}_{2yx}$. For $x \in [\frac{1}{2},1]$, set $\mathrm{D}_{(x,y)} = \mathrm{D}_{2y(1-x)}$. Each $y \in [0,1]$ version of $\mathrm{D}_{(x,y)}$ starts and ends at D_0 as x increases in $[0,1]$. Meanwhile, the $y=0$ member is D_0 for all x, and the $y=1$ starts at D_0 when $x=0$, runs to D_1 at $x=\frac{1}{2}$, and returns to D_0 at $x=1$.

If D_1 is close to D_0, then each $y \in [0,1]$ family $\{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}$ defines a version of $\hat{I}$ that preserves the D_0 version of (4-3). Thus, each defines an endomorphism of the D_0 version of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$. Denote the latter by I^y. The $y=0$ version is the identity endomorphism. Meanwhile, the $y=1$ family $\{\mathrm{D}_{(x,1)}\}_{x\in[0,1]}$ is such that $I^{(1)}$ factors on $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ as $I' \circ I_{0,1*}$. It follows that $I_{0,1*}$ is an isomorphism if $I^{(1)}$ is. Meanwhile, the latter follows if there is a chain homotopy between $\hat{I}^{(1)}$ and $\hat{I}^{(0)}$. As explained in Step 5, such a chain homotopy is obtained from a certain chain homotopy between the identity endomorphism of $\mathcal{C}^{\mathrm{SW}}$ and the endomorphism that is defined by the data set $\{\mathrm{D}_{(x,1)}\}_{x\in[0,1]}$. Let $\hat{o}\colon \mathcal{C}^{\mathrm{SW}} \to \mathcal{C}^{\mathrm{SW}}$ denote the desired chain homotopy. Then $\hat{o}$ defines a chain homotopy between $\hat{I}^{(1)}$ and $\hat{I}^{(0)}$ if $\hat{o}$ maps $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ to itself.

Step 5 This step provides some background for Steps 6 and 7. To set the stage, suppose that $\{y \to \{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}\}_{y\in[0,1]}$ is a 2–parameter family of data sets as described by the second bullet of the lemma. The family from Step 4 gives one example. By assumption, both the $y = 0$ and $y = 1$ members are suitable for defining a homomorphism between the respective realizations of the Seiberg–Witten Floer cochain complex that induces the canonical isomorphism between their cohomologies. Let $\hat{I}_0$ and $\hat{I}_1$ these homomorphisms. Chapters 23 and 24 of [16] explain how to construct a chain homotopy $\hat{I}_0$ and $\hat{I}_1$ from the 2–parameter family $\{y \to \{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}\}_{y\in[0,1]}$.

The construction requires first choosing a very small normed and sufficiently generic map $\mathfrak{q}\colon \times_2[0,1] \to \mathcal{P}$ that vanishes on the boundary of the square $\times_2[0,1]$. If $\mathfrak{q}$ is sufficiently generic, then a chain homotopy is defined using an algebraic count of the number of cobordism instantons of the following sort: For each $(x,y) \in \times_2[0,1]$, let $\mathrm{D}^{\mathfrak{q}}_{(x,y)}$ denote the data set that is obtained from $\mathrm{D}_{(x,y)}$ by replacing $\mathfrak{p}_{(x,y)}$ by $\mathfrak{p}_{(x,y)} + \mathfrak{q}_{(x,y)}$. The instantons in question are defined by a $y \in [0,1]$ version of the data set $\{\mathrm{D}^{\mathfrak{q}}_{(x,y)}\}_{x\in[0,1]}$, and they are such that the corresponding version of (3-9) has Fredholm index equal to -1. Let $\mathfrak{c}$ denote a given irreducible generator of the $x = 0$ version of the Seiberg–Witten chain complex. If $\mathfrak{q}$ is sufficiently generic, there will be but a finite number of such instantons whose $s \to \infty$ limit gives $\mathfrak{c}$.

Let $\hat{o}$ denote the chain homotopy. To say a sentence more about its definition, let $\mathfrak{c}$ again denote a given irreducible generator of the $x = 0$ version of the Seiberg–Witten Floer cochain complex. Then $\hat{o}(\mathfrak{c})$ is a suitably weighted sum of those generators of the $x = 1$ cochain complex that are the $s \to -\infty$ limit of some cobordism instanton of the sort just described whose $s \to \infty$ limit gives $\mathfrak{c}$.

Step 6 Consider what is said in Step 5 in the context of the family $\{\mathrm{D}_{(x,y)}\}_{x,y\in[0,1]}$ from Step 4. The respective $x = 0$ and $x = 1$ realizations of the Seiberg–Witten Floer cochain complexes are the same, which is to say $\mathcal{C}^{\mathrm{SW}}$. To say more, let $\hat{o}$ denote the chain homotopy. Let $\mathfrak{c}$ denote an irreducible generator of $\mathcal{C}^{\mathrm{SW}}$, and let $\mathfrak{d}$ denote a cobordism instanton that contributes to $\hat{o}(\mathfrak{c})$. Let $\mathfrak{c}'$ denote the $s \to -\infty$ limit of $\mathfrak{d}$. Given that D_1 is very close to D_0, then $\mathfrak{a}(\mathfrak{c}') - \mathfrak{a}(\mathfrak{c})$ can be negative only by a very small amount. Meanwhile, $\mathfrak{d}$ decreases the spectral flow function by 1. This being the case, it follows that $\mathfrak{a}^f(\mathfrak{c}') - \mathfrak{a}^f(\mathfrak{c}) + 1$ can be negative only by a very small amount. In particular if $\mathfrak{a}^f(\mathfrak{c}) > r\mathfrak{A}$, $\mathfrak{a}^f(\mathfrak{c}') > r\mathfrak{A} - 1$, and so (4-2) requires that $\mathfrak{a}^f(\mathfrak{c}') > r\mathfrak{A}$. As a consequence, $\hat{o}$ maps $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ to itself.

Step 7 To prove the second bullet, consider first the case where the family of $\{\{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}\}_{y\in[0,1]}$ is such that for any fixed $y \in [0,1]$, the change in the data set $\{\mathrm{D}_{(x,y)}\}_{x\in(0,1)}$ is very small. To be precise about this, suppose that $\delta > 0$ is given such that the following strong form of (4-2) holds: If $x \in [0,1]$, then there is no solution to the $\mathrm{D}_{(x,0)}$ version of (3-5) with $\mathfrak{a}^f \in [r\mathfrak{A} - \delta, r\mathfrak{A} + \delta]$ or $\mathfrak{a}^f \in [r\mathfrak{A} - 1 - \delta, r\mathfrak{A} + \delta]$. Let $\mathfrak{c}$ denote a generator of the $x = 1$ version of $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ and let $\mathfrak{d}$ denote a cobordism instanton that is defined using some $y \in [0,1]$ version of the data set $\{\mathrm{D}^{\mathfrak{q}}_{(x,y)}\}_{x\in[0,1]}$ whose $s \to \infty$ limit gives $\mathfrak{c}$ and whose version of (3-7) has Fredholm index -1. Let $\mathfrak{c}'$ denote the $s \to -\infty$ limit of $\mathfrak{d}$. With δ fixed, if the data that defines each $x \in [0,1]$ version of $\mathrm{D}_{(x,y)}$ is sufficiently close to the data defining $\mathrm{D}_{(x,0)}$, and if each $x \in [0,1]$ version of $\mathfrak{q}_{(x,y)}$ has sufficiently small norm, then $smash fa^f(\mathfrak{c}') - \mathfrak{a}^f(\mathfrak{c}) + 1$ can be no less than $-\frac{1}{2}\delta$. As $\mathfrak{d}$ has Fredholm index -1, this implies $\mathfrak{c}'$ is a generator of the $x = 0$ version of $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$. It follows that the chain map $\hat{o}$ restricts to the $x = 1$ version of $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ so as to map the latter to the $x = 0$ version of $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$. This implies that the respective $y = 0$ and $y = 1$ versions of the isomorphism from the $x = 1$ realization of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ to the $x = 0$ version of $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ are identical.

Step 8 What is said in Step 7 is exploited by breaking Lemma 4.6's version of the family $\big\{\{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}\big\}_{y\in[0,1]}$ into a set of "small" homotopies of the sort considered in Step 7. This can be done as follows: Fix a very large, odd integer N, and a partition of the interval $[0,1]$ as $0 = x_1 < x_2 < \cdots < x_N = 1$ with $|x_k - x_{k-1}| < 2N^{-1}$ for all k. For odd $k \in \{1,\dots,N-1\}$, let $\{\mathrm{A}^k_\sigma\}_{\sigma\in[0,1]}$ denote the path in the square $\times_2[0,1]$ that starts at $(x_k, 0)$, increases the y–coordinate to 1 with $x = x_k$, then keeps $y = 1$ and increases x to $x_k + 1$. For even $k \in \{1,\dots,N-1\}$, let $\{\mathrm{A}^k_\sigma\}_{\sigma\in[0,1]}$ denote the path in the square that starts at $(x_k, 1)$, increases the x–coordinate to x_{k+1} with y fixed at 1, then fixes x at x_{k+1} and decreases the y coordinate so as to end at $(x_{k+1}, 0)$. If the points $\{\mathfrak{p}^0_x\}_{x\in[0,1]}$ and $\{\mathfrak{p}^1_x\}_{x\in[0,1]}$ are suitably generic, and if the points $\{x_k\}_{k=1,\dots,N}$ are suitably generic, and if $\mathfrak{q}$ is suitably generic, then each path $\{\mathrm{A}^k\}_{k=1,2,\dots,N-1}$ is suitable for defining a cobordism homomorphism between the respective Seiberg–Witten Floer cochain complexes that are defined by its endpoints. Let $\{\hat{I}^k_{\mathrm{A}}\}_{k=1,\dots,N-1}$ denote these homomorphisms, and let $\{I^k_{\mathrm{A}}\}_{k=1,\dots,N-1}$ denote the resulting canonical isomorphisms between the respective realizations of the Seiberg–Witten Floer cohomologies. Let I^1 denote the isomorphism that is defined by the data set $\{\mathrm{D}^1_x\}_{x\in[0,1]}$. The factorization property of the canonical isomorphism with respect to compositions (as explained in Part 3 of Section 3.h) implies that $I^1_{\mathrm{A}} \circ \cdots \circ I^{N-1}_{\mathrm{A}} = I^1$.

Define next a second set of paths $\big\{\{\mathrm{B}^k_\sigma\}_{\sigma\in[0,1]}\big\}_{k=1,\dots,N-1}$ as follows: For odd k, the path B^k starts at $(x_k, 0)$, then increases the x coordinate with $y = 0$ to reach $(x_{k+1}, 0)$, then keeps the x coordinate fixed and increases y so as to end at $(x_{k+1}, 1)$. For even k, the path starts at $(x_k, 1)$, decreases the y coordinate with x fixed to reach $(x_k, 0)$, and then keeps y at 0 as x increases so as to end at $(x_{k+1}, 0)$. Given suitable genericity assumptions, this set of paths also defines a corresponding set of cobordism homomorphisms $\{\hat{I}^k_{\mathrm{B}}\}_{k=1,\dots,N-1}$ and a resulting set of canonical isomorphisms $\{I^k_{\mathrm{B}}\}_{k=1,\dots,N-1}$. Let I_0 denote the isomorphism that is defined by the data set $\{\mathrm{D}^0_x\}_{x\in[0,1]}$. The composition of the latter is such that $I^1_{\mathrm{B}} \circ \cdots \circ I^{N-1}_{\mathrm{B}} = I^0$.

Step 9 For any given $k \in \{1, \dots, N-1\}$, the paths A^k and B^k have the same starting and ending points. Thus, the corresponding $\hat{I}^k_{\mathrm{A}}$ and $\hat{I}^k_{\mathrm{B}}$ are homomorphisms between the same pair of cochain complexes. This understood, it follows from the composition identities given above that it is enough to prove that the conclusions of the second bullet hold with roles of $\{\mathrm{D}^0_x\}_{x\in[0,1]}$ and $\{\mathrm{D}^1_x\}_{x\in[0,1]}$ played respectively by a suitable parametrization of each $k \in \{1, \dots, N-1\}$ version of A^k and B^k. The role of the family $\{\{\mathrm{D}_{(x,y)}\}_{x\in[0,1]}\}_{y\in[0,1]}$ is played by a suitable parametrization of the following family of paths: Let $z \in [0,1]$ denote this parameter. If k is odd, then the $z = 0$ path is A^k and the $z = 1$ path is B^k. In general, any given $z \in [0,1]$ path starts at $(x_k, 0)$, increases x_k to equal $x_k + z(x_{k+1} - x_k)$ with y held at 0, then holds x fixed and increases y to $y = 1$, then with y fixed, it increases x to x_{k+1}. If k is even, the $z = 0$ path is B^k and the $z = 1$ path is A^k. Any given $z \in [0,1]$ path starts at $(x_k, 1)$, increases x to $x_k + z(x_{k+1} - x_k)$ with y held at 1, then decreases y to 0 with x fixed, and finally increases x to x_{k+1} with y held at 0. If N is large, a suitable reparametrization of such a family as $\{\{\Delta_{(w,z)}\}_{w\in[0,1]}\}_{z\in[0,1]}$ is such that for any fixed z, the data that defines the path $\{\Delta_{(w,z)}\}_{w\in[0,1]}$ is very close to the data that defines the $\{\Delta_{(w,0)}\}_{w\in[0,1]}$. In particular, the observations in Step 6 apply if N is large to this family and so prove what is needed to complete the proof of Lemma 4.6. $\square$

4.e Min-max

This section introduces various additional notions that were used in [24; 25]. Minor modifications are made on these in preparation for their use in the upcoming proof of Theorem 4.5. The upcoming Proposition 4.7 and Proposition 4.8 supply these notions.

To set the stage, fix a pair (a, J) of contact 1–form and compatible almost complex structure. If $c_1(\det(\mathbb{S}))$ is torsion, also fix an integer k. Let $\mu \in \Omega$ denote a 1–form with $\mathcal{P}$ norm less than 1. Suppose that a small element $\mathfrak{p} \in \mathcal{P}$ has been chosen so that the data $((a, J), r, \mathfrak{e}_\mu + \mathfrak{p})$ is suitable for defining the Seiberg–Witten Floer cochain complex.

Let θ be a nonzero Seiberg–Witten Floer cohomology class in a given degree, this k or greater if $c_1(\det(\mathbb{S}))$ is torsion. Let $\mathfrak{n} = \sum_{\mathfrak{c}} z_{\mathfrak{c}}\mathfrak{c}$ denote a representative of this class in the Seiberg–Witten Floer cochain complex. Here, $z_{\mathfrak{c}} \in \mathbb{Z}$ and $\mathfrak{c} \in \mathcal{M}^r$. Define $\mathfrak{a}^f[\mathfrak{n}, r]$ to be the minimum of the values of $\mathfrak{a}^f$ on the set $\{\mathfrak{c} : z_{\mathfrak{c}} \neq 0\}$. Set $\mathfrak{a}^f_\theta[r]$ to denote the maximum of $\{\mathfrak{a}^f[\mathfrak{n}, r] : \mathfrak{n}$ represents $\theta\}$. Note that $\mathfrak{a}^f_\theta[\cdot]$ is defined by first taking a minimum and then a maximum; and that this is opposite to the order used in [24; 25]. The order is switched because θ is a *cohomology* class rather than a *homology* class. If $\mathfrak{c}$ is a given generator from $\mathcal{M}^r$, then a generator $\mathfrak{c}'$ appears on the right hand side of (3-4)'s definition of $\delta^*\mathfrak{c}$ only if $\mathfrak{M}_1(\mathfrak{c}', \mathfrak{c}) \neq$. As such, $\delta^*\mathfrak{c}$ is a sum of generators on which $\mathfrak{a}^f(\cdot) > \mathfrak{a}^f(\mathfrak{c})$.

Proposition 4.7 *Fix a pair (a, J), and fix $k \in \mathbb{Z}$ if $c_1(\det(\mathbb{S}))$ is a torsion class. There is a residual set in Ω such that if μ is from this set, then there exists $r_k > 1$ and a nonaccumulating, discrete set $\mathfrak{U} \in [r_k, \infty)$ such that the following is true:*

- *Take $r \in [r_k, \infty) - \mathfrak{U}$. All solutions to the corresponding version of (3-17) are irreducible and nondegenerate if $c_1(\det(\mathbb{S}))$ is not torsion. If $c_1(\det(\mathbb{S}))$ is torsion, all solutions to (3-17) with degree k or more are irreducible and nondegenerate.*
- *For $r \in [r_k, \infty) - \mathfrak{U}$, use the data (a, J), μ and r to define the generators for the Seiberg–Witten Floer cochain complex, but in degrees k or greater if $c_1(\det(\mathbb{S}))$ is torsion. Fix a decreasing map $\sigma\colon [r_k, \infty) \to (0, 1)$ with limit zero. Let $r \to \mathfrak{p}_r$ denote a map from $[r_k, \infty) - \mathfrak{U}$ to $\mathcal{P}$ such that*
  - (i) *Any given version of $\mathfrak{p}_r$ has norm less than $\sigma(r)$.*
  - (ii) *$\mathfrak{p}_r$ vanishes to order two on all solutions (with degree k or greater if $c_1(\det(\mathbb{S}))$ is torsion) to the (a, J), μ and r version of (3-17). If the map $r \to \mathfrak{p}_r$ is suitably generic subject to these last constraints, then*
    - (a) *The data $((a, J), r, \mathfrak{e}_\mu + \mathfrak{p}_r)$ is suitable for defining the differential for the Seiberg–Witten Floer cochain complex if $r \in [r_k, \infty) - \mathfrak{U}$ is chosen from a discrete set, $\mathfrak{V}$, that accumulates only on the points in $\mathfrak{U}$.*
    - (b) *The various $r \in [r_k, \infty) - (\mathfrak{U} \cup \mathfrak{V})$ versions of the Seiberg–Witten Floer cohomology groups (in degrees k or greater when $c_1(\det(\mathbb{S}))$ is torsion) can be identified so that the following is true: If θ is any given nonzero cohomology class, then the assignment $r \to \mathfrak{a}^f_\theta[r]$ is the restriction of a continuous, piecewise differentiable function on the half line $[r_k, \infty)$.*

Proof of Proposition 4.7 Except for the ordering change with regards to "min" and "max", the argument is the same as that used to prove Proposition 4.2 in [24] when $c_1(\det(\mathbb{S}))$ is torsion. But for this same ordering change, the argument is essentially that used for Proposition 2.5 in [25] when $c_1(\det(\mathbb{S}))$ is not torsion.

Now fix $\mathfrak{A} \in \mathbb{R}$ and assume that (4-2) holds. A similar min-max construction can be done for classes in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ and also for those in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$. To say more, suppose again that $\mu \in \Omega$ has $\mathcal{P}$ norm less than 1 and that $\mathfrak{p} \in \mathcal{P}$ is small and such that $((a, J), r, \mathfrak{e}_\mu + \mathfrak{p})$ is suitable for defining the Seiberg–Witten Floer cochain complex. Assume that $\mathfrak{p}$ is such that (4-2) also holds for the corresponding solutions to (3-5). Let θ denote a nonzero class in either $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$ or $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$. Assume that θ has degree k or greater in the case when $c_1(\det(\mathbb{S}))$ is torsion. Given a cocycle $\mathfrak{n}$ that represents θ, define $\mathfrak{a}^f[\mathfrak{n}, r]$ to be the minimum of $\mathfrak{a}^f(\cdot)$ on the generators that represent θ. Then define $\mathfrak{a}^f_\theta[r]$ to be the maximum of the elements in $\{\mathfrak{a}^f[\mathfrak{n}, r] : \mathfrak{n} \text{ represents } \theta\}$. The analog of Proposition 4.7 in this new context is given by the next proposition. Its proof is essentially the same as that for Proposition 4.7 so the details are left to the reader. □

Proposition 4.8 *The conclusions of Proposition 4.7 can be augmented with the following: As r varies in $[r_k, \infty) - (\mathfrak{U} \cup \mathfrak{V})$, the corresponding versions of the $\mathbb{Z}$–modules in (4-4) can be identified (in degrees k or greater when $c_1(\det(\mathbb{S}))$ is torsion) so that if θ denotes any given cohomology class in any of the three cohomologies, then the assignment $r \to \mathfrak{a}^f_\theta[r]$ is the restriction of a continuous, piecewise differentiable function on the half line $[r_k, \infty)$.*

The continuity and piecewise differentiability of $\mathfrak{a}^f_\theta[\cdot]$ is exploited in the coming subsections.

4.f Bounds on E from $\mathfrak{a}^f$ and vice versa

The next proposition plays one of the key roles in the proof of Theorem 4.5. To set the terminology, suppose that (a, J) is given to define the metric on M. Fix $\mu \in \Omega$ with $\mathcal{P}$ norm less than 1 and a map $r \to \mathfrak{p}_r$ as described in Proposition 4.8. Suppose in addition that $\mathfrak{A} \in \mathbb{R}$ has been specified and that (4-2) holds. In what follows, the various $r \in [r_k, \infty) - (\mathfrak{U} \cup \mathfrak{V})$ versions of (4-4) are implicitly identified using one of the identifications provided by Proposition 4.8. Such an identification should be understood when reference is made to a class in one of the groups in (4-4) with no reference to the precise value of r.

Proposition 4.9 *Fix a pair (a, J), and if $c_1(\det(\mathbb{S}))$ is torsion, fix an integer k. Choose μ as described in Proposition 4.7 and Proposition 4.8. Fix $\mathfrak{A} \in \mathbb{R}$ and suppose that (4-2) holds. There exists $\mathcal{K} \geq 1$ with the following significance: Fix a map $r \to \mathfrak{p}_r$ as described in Proposition 4.8 and such that each $\mathfrak{p}_r$ has very small norm. Suppose that θ is a nonzero cohomology class of fixed degree (k or greater if $c_1(\det(\mathbb{S}))$ is torsion) in any of the cohomology groups that appear in (4-4). There exists an increasing, unbounded sequence $\{r_i\}_{i=1,2,\ldots}$ in $[r_k, \infty) - (\mathfrak{U} \cup \mathfrak{V})$ such that the $r = r_i$ version of the class θ has a representative cocycle with $\mathrm{E}(\cdot) \leq 2\pi\mathcal{K}$ on each generator that appears with a nonzero coefficient.*

The proof of this last proposition makes use of the following restatement of results from [24; 25].

Proposition 4.10 *Let (a, J) denote a pair consisting of a contact 1–form and compatible almost complex structure. There exists $\kappa \geq 1$ such that the following is true: Fix $\mu \in \Omega$ with $\mathcal{P}$ norm bounded by 1. Suppose that $\mathfrak{c} = (A, \psi)$ is a solution to the (a, J), μ and r version of (3-17). Then*

- $|\mathfrak{cs}^f| < \kappa r^{31/16}$.
- *If* $c_1(\det(\mathbb{S}))$ *is torsion, then* $|\mathfrak{cs}| \leq \kappa r^{2/3}(1+|\mathrm{E}|^{4/3})$.
- *If* $c_1(\det(\mathbb{S}))$ *is not torsion, then* $|\mathfrak{cs}^f| \leq \kappa r^{2/3}(\ln r)^{\kappa}(1+|\mathrm{E}|^{4/3})$.

Proof of Proposition 4.9 But for notation, the arguments in the case $c_1(\det(\mathbb{S}))$ is torsion are those used in Propositions 4.6 and Corollary 4.7 in [24] with Propositions 4.8 and 4.10 added. In the case that $c_1(\det(\mathbb{S}))$ is not torsion, the arguments are the same but for notation as those used in Section 2.3 of [25] but with Propositions 4.8 and 4.10 added. □

Proof of Proposition 4.10 The bound in the first bullet follows from Proposition 5.1 in [24] in the case when $c_1(\det(\mathbb{S}))$ is torsion, and Proposition 1.10 of [25] when $c_1(\det(\mathbb{S}))$ is not torsion. The bound in the second bullet follows from (4-2) in [24] and Lemma 2.4 in [24]. The bound in the third bullet restates Proposition 1.9 of [25]. □

Proposition 4.10 has an additional very important corollary:

Proposition 4.11 *Let* (a, J) *denote a pair consisting of a contact* 1*–form and compatible almost complex structure. There exists* $\kappa \geq 1$ *with following significance: Fix* $\mu \in \Omega$ *with* $\mathcal{P}$ *norm bounded by 1. Take* $r \geq \kappa$ *and suppose that* $\mathfrak{c}$ *is a solution to the* (a, J), μ *and* r *version of (3-17) with* $\mathfrak{a}^f(\mathfrak{c}) > -\kappa r^{31/16}$. *Then*

- $|2r^{-1}\mathfrak{a}^f(\mathfrak{c}) + \mathrm{E}(\mathfrak{c})| \leq r^{-1/50}(1+|\mathrm{E}(\mathfrak{c})|)$.
- *In addition, if* $c_1(\det(\mathbb{S}))$ *is torsion,* $|2r^{-1}\mathfrak{a}(\mathfrak{c}) + \mathrm{E}(\mathfrak{c})| \leq r^{-1/50}(1+|\mathrm{E}(\mathfrak{c})|)$.

The proof of this proposition introduces a convention that is used throughout this paper and its sequels: In all appearances, c_0 denotes a constant greater than 1 whose value is independent of r, μ, and any given (A, ψ). Subsequent appearances of c_0 are allowed to have different values, but these can be assumed to increase from one appearance to the next.

Proof of Proposition 4.11 Since $|\mathfrak{cs}^f| \leq c_0 r^{31/16}$, and since it is assumed that $\mathfrak{a}^f$ is greater than $-c_0 r^{31/16}$, it follows that $\mathrm{E} \leq c_0 r^{15/16}$. Hold this last bound for the moment. Use the third bullet in Proposition 4.10 to see that $|\mathfrak{cs}^f| \leq c_0 r^{2/3}\mathrm{E}^{4/3}(\ln r)^{c_0}$. These bounds and the bound just derived for E imply that $|\mathfrak{cs}|$ or $|\mathfrak{cs}^f|$ is no greater than $c_0 r^{-1/50}(r\mathrm{E})$. The latter bound implies the assertion of the first bullet. The argument for the second bullet is identical but for an appeal now to the second bullet in Proposition 4.10. □

4.g Bounds on E and $\mathfrak{a}^f$ for families

The next result asserts a parametrized version of what is proved in Section 6d of [24]. To set things up, suppose that $\{(a_\tau, J_\tau)\}_{\tau\in[0,1]}$ is a smoothly parametrized family of pairs consisting of a contact 1–form and compatible almost complex structure. This family is assumed now to have one additional attribute. Suppose that $L \geq 1$ has been specified such that the following condition holds for each $\tau \in [0,1]$:

> Let Θ denote a set of pairs of the form (γ, m) with γ a Reeb orbit as defined by a_τ and m a positive integer. Assume that distinct pairs have distinct Reeb orbit component. Then

$$\sum_{(\gamma,m)\in\Theta} m\ell_\gamma \neq L. \tag{4-5}$$

Let $\{\mu_\tau\}_{\tau\in[0,1]}$ denote a corresponding family of 1–forms, each in Ω and each with $\mathcal{P}$ norm bounded by 1.

Proposition 4.12 *Given the data $\{(a_\tau, J_\tau), \mu_\tau\}_{\tau\in[0,1]}$ and L, there exists κ with the following significance: Fix $\tau \in [0,1]$. Suppose that $r \geq \kappa$ and that $\mathfrak{c} = (A, \psi)$ is a solution to the version of (3-17) that is defined by the data (a_τ, J_τ), μ_τ and r. Assume in addition that $\mathrm{E}(A) \leq 2\pi L + \kappa^{-1}$. Then $\mathrm{E}(A) < 2\pi L - \kappa^{-1}$.*

This last proposition leads to the following important observation.

Proposition 4.13 *Fix the data $\{(a_\tau, J_\tau),\ \mu_\tau\}_{\tau\in[0,1]}$ and L. There exists $\kappa \geq 100$ with the following significance: Fix $\tau \in [0,1]$ and $r \geq \kappa$. Suppose that $\mathfrak{c} = (A, \psi)$ is a solution to the version of (3-17) that is defined by (a_τ, J_τ), μ_τ and r. Then*

$$\mathfrak{a}^f(\mathfrak{c}) \notin \left[-\pi L(1-\kappa^{-1})r,\ -\pi L(1+\kappa^{-1})r\right].$$

Proof of Proposition 4.12 In what follows, the spinor bundle $\mathbb{S}$ is written as $E \oplus EK^{-1}$, and corresponding components of a given section ψ are denoted by (α, β).

To start the proof, suppose that the proposition is not true. Then there exists an unbounded sequence $\{r_n\}_{n=1,2,\ldots} \subset [1,\infty)$, a convergent sequence $\{\tau_n\}_{n=1,2,\ldots} \subset [0,1]$, and a corresponding sequence $\{(A_n, \psi_n)\}_{n=1,2,\ldots}$ where any given (A_n, ψ_n) obeys (3-17) as defined using $r = r_n$ and the $\tau = \tau_n$ version of (a_τ, J_τ) and μ_τ. Let $\tau \in [0,1]$ denote the limit point of the sequence $\{\tau_n\}_{n=1,2,\ldots}$. The four steps that follow explain why the existence of such a sequence leads to nonsense and so proves the lemma.

Step 1 The arguments used in Section 6d of [24] can be repeated with only cosmetic changes to find a possibly empty set, Θ, of the following sort: First, the typical element in Θ is a pair (γ, m) where γ is an a_τ Reeb orbit and m is a positive integer. Moreover, distinct pairs have distinct Reeb orbit components. Finally, there exists some subsequence of $\{(A_n, \psi_n = (\alpha_n, \beta_n))\}_{n=1,2,\dots}$, hence renumbered consecutively from 1, and there is a sequence $\{\varepsilon_n\}_{n=1,2,\dots} \subset (0, \frac{1}{100}]$ with limit zero such that

(4-6)

- $|\alpha_n|^2 \geq 1 - \varepsilon_n$ at all points with distance ε_n or greater from $\bigcup_{(\gamma,m)\in\Theta} \gamma$.
- Fix $(\gamma, m) \in \Theta$ and let $\varphi\colon S^1 \times D \to M$ denote a coordinate chart for a tubular neighborhood of γ of the sort depicted in (2-1). Assume that the closure of the image of φ is disjoint from all other Reeb orbits from Θ. Then α_n vanishes on any given constant $t \in S^1$ slice of $\varphi(S^1 \times D)$ with degree m.

Step 2 Fix a smooth function $\chi\colon \mathbb{R} \to [0, 1]$ such that $\chi = 1$ on $(-\infty, \frac{5}{16}]$ and $\chi = 0$ on $[\frac{7}{16}, \infty)$. Given a pair (A, α) of connection on E and section of E, introduce the connection

$$\hat{A} = A - \tfrac{1}{2}(1 - \chi(|\alpha|^2))|\alpha|^{-2}(\bar{\alpha}\nabla_A\alpha - \alpha\nabla_A\bar{\alpha}). \tag{4-7}$$

Note that $\hat{A}$ is flat where $|\alpha|^2 > \frac{1}{2}$; and here $\alpha/|\alpha|$ is covariantly constant. Use $\hat{A}_n$ in what follows to denote the (A_n, α_n) version of $\hat{A}$. It follows from the first item in (4-6) that $\hat{A}_n$ is flat at all points in M with distance ε_n or greater from $\bigcup_{(\gamma,m)\in\Theta} \gamma$. Let $(\gamma, m) \in \Theta$. Let $\varphi\colon S^1 \times D \to M$ denote the coordinate chart map from the second item of (4-6). It is a consequence of this second item that

$$i \int_{S^1\times D} dt \wedge \varphi^*(*B_{\hat{A}_n}) = 2\pi m. \tag{4-8}$$

What with (2-1) and Lemma 2.2 in [24], the equality in (4-8) implies that

$$i \int_M a_\tau \wedge *B_{\hat{A}_n} = 2\pi \sum_{(\gamma,m)\in\Theta} m\ell_\gamma + \mathfrak{e}, \tag{4-9}$$

where $|\mathfrak{e}| \leq c_0\varepsilon_n(1 + |\mathrm{E}(A_n)|)$. Given that the sequence $\{\tau_n\}_{n=1,2,\dots}$ converges to τ, a second appeal to Lemma 2.2 in [24] now applied to (4-9) finds

$$i \int_M a_{\tau_n} \wedge *B_{\hat{A}_n} = 2\pi \sum_{(\gamma,m)\in\Theta} m\ell_\gamma + \mathfrak{e}_n, \tag{4-10}$$

where $|\mathfrak{e}_n| \leq c_0\delta_n(1 + L)$ with $\{\delta_n\} \in (0, 1)$ a sequence with limit zero as $n \to \infty$.

Step 3 Integrate by parts to see that

$$\left|\int_M a_{\tau_n} \wedge (*B_{\hat{A}_n} - *B_{A_n})\right| = \left|\int_M da_{\tau_n} \wedge (1-\chi(|\alpha_n|^2))|\alpha_n|^{-2}(\bar{\alpha}_n \nabla_{A_n}\alpha_n - \alpha_n \nabla_{A_n}\bar{\alpha}_n)\right|. \tag{4-11}$$

To bound the right hand side of (4-11), first define the function g on the domain $[0,\infty)$ by setting the rule

$$g(t) = -\int_t^2 (1-\chi(s))s^{-1}\,ds. \tag{4-12}$$

Since

$$(1-\chi(|\alpha|^2))|\alpha|^{-2}(\bar{\alpha}\nabla_A\alpha + \alpha\nabla_A\bar{\alpha}) = d(g(|\alpha|^2),$$

an integration by parts on the right hand side of (4-1) identifies the latter with

$$2\left|\int_M da_{\tau_n} \wedge (1-\chi(|\alpha_n|^2))|\alpha_n|^{-2}\alpha_n \nabla_{A_n}\bar{\alpha}_n\right|. \tag{4-13}$$

Next, use the Dirac equation to identify the covariant derivative of α_n along the Reeb vector field for a_{τ_n} with derivatives of β_n. Make this identification, and then use Hölder's inequality to bound (4-13) by $\|\nabla_{A_n}\beta_n\|_2$, where the subscript indicates the L^2 norm.

To complete the bound on the right hand side of (4-11), integrate both sides of what is written in the second line of Equation (6.4) in [24] over M and use the latter's Lemma 2.2 again to bound the L^2 norm of $\nabla_{A_n}\beta_n$ by $c_0 r_n^{-1/2}$.

Step 4 Given that the right hand side of (4-11) is bounded by $c_0 r_n^{-1/2}$, it thus follows from (4-10) that $\mathrm{E}(A_n) \le 2\pi\Sigma_{(\gamma,m)\in\Theta} m\ell_\gamma + c_0\delta_n{}'$ where $\{\delta_n{}'\}_{n=1,2,\dots}$ is a sequence with limit zero as $n\to\infty$. However, this then implies that $\mathrm{E}(A_n) < 2\pi L$ for all n sufficiently large. Given the assumptions, this is nonsense. □

Proof of Proposition 4.13 Suppose for the moment that the constant κ that appears in Proposition 4.11 can be chosen so that Proposition 4.11's conclusions hold with the same constant κ for all $\tau\in[0,1]$ versions of (a_τ, J_τ). Suppose that there exists $\tau\in[0,1]$, $\epsilon>0$ and a solution to (3-17) as defined with the data (a_τ, J_τ), μ_τ and r such that $\mathfrak{a}^f$ lies in the interval $[-\pi L(1-\epsilon)r,\ -\pi L(1+\epsilon)r]$. Appeal to Proposition 4.11 to conclude that the value of E on this solution obeys $\mathrm{E} \le 2\pi(L+\epsilon) + c_0 r^{-1/50}$. If $\epsilon < c_0^{-1}$, then a further appeal to Proposition 4.12 finds that $\mathrm{E} < 2\pi L - c_0^{-1}$ if $r \ge c_0$. Another look at Proposition 4.11 finds that $\mathfrak{a}^f \ge -\pi(L-c_0^{-1})r$.

The constant κ that appears in Proposition 4.11 comes from the version that appears in Proposition 4.10. The value of κ that makes the first item of Proposition 4.10 true comes from Proposition 5.1 in [24] and Proposition 1.10 in [25]. A look at the proofs of Proposition 5.1 in [24] and Proposition 1.10 in [25] finds that the latter version of κ can be chosen so as to hold for every $\tau \in [0,1]$ version of (a_τ, J_τ) and μ_τ.

The value of κ that makes the second item in Proposition 4.10 true comes from Lemma 2.4 in [24]; and the value of κ that makes the third item true comes via Proposition 1.9 in [25]. A close look at the proofs of both of these propositions shows that their contributions to κ can be assumed to be τ independent. □

4.h Proofs of Theorems 4.4 and 4.5

The preceding three subsections supply all of the heavy machinery for the proof. It remains only to put the machinery to use.

Proof of Theorem 4.4 Define $\mathfrak{a}$ as in (3-15) and $\mathfrak{a}^f$ as in (3-16) using $\mathfrak{g} = \mathfrak{e}_\mu + \mathfrak{p}$. Note that all solutions to (3-5) are solutions to (3-17) and vice-versa. Meanwhile, if $\mathfrak{c}$ is a solution to (3-17), then both $\mathfrak{a}(\mathfrak{c})$ and $\mathfrak{a}^f(\mathfrak{c})$ are given by the version of (3-15) with $\mathfrak{g} = \mathfrak{e}_\mu$. Also keep in mind that $\mathrm{E} < 2\pi L$ on $\Phi^r(\mathcal{Z}^L_{\mathrm{ech}})$.

To prove the assertion in the first bullet, use the first item in Proposition 4.10 and Proposition 4.11 to see that $\mathfrak{a}^f \geq -(\pi L - c_0^{-1})r$ on $\Phi^r(\mathcal{Z}^L_{\mathrm{ech}})$ if $r \geq c_0$. Now let $\mathfrak{d}$ denote an instanton solution to the relevant version of (3-6). The assignment $s \to \mathfrak{a}(\mathfrak{d}(s))$ defines a decreasing function on $\mathbb{R}$, and so $\mathfrak{a}(\mathfrak{c}_-) > -\pi L r$. Since $\mathfrak{d}$'s version of (3-9) has Fredholm index 1, so $f(\mathfrak{c}_-) = f(\mathfrak{c}_+) - 1$. As a consequence, $\mathfrak{a}^f(\mathfrak{c}_-) > \mathfrak{a}^f(\mathfrak{c}_+) > -\pi L r$. This understood, Proposition 4.11 implies that $\mathrm{E}(\mathfrak{c}_-) < 2\pi L$ if $r \geq c_0$. According to Theorem 4.2, this implies that $\mathfrak{c}_- \in \Phi^r(\mathcal{Z}^L_{\mathrm{ech}})$.

The assertion in the second bullet is an immediate consequence of the first bullet and Theorem 4.3. □

Proof of Theorem 4.5 The proof has three parts.

Part 1 The lemma that follows summarizes the contribution from this part. The lemma refers to notation that is introduced in Propositions 4.7–4.9.

Lemma 4.14 *Let (a, J) be a pair of contact 1–form from Lemma 2.1's residual set and almost complex structure from $\mathcal{J}_a$. Fix $k \in \mathbb{Z}$ when $c_1(\det(\mathbb{S}))$ is torsion. Then there exists $\kappa > 1$ and $\mathcal{K} \geq 1$ with the following significance: Fix $\mu_0 \in \Omega$ with $\mathcal{P}$ norm less than 1 as described in Proposition 4.7–4.9 and likewise a map $r \to \mathfrak{p}_r$ from $[r_k, \infty) - \mathfrak{U}$ to $\mathcal{P}$ with very small norm. Use (a, J), μ and $r \in [r_k + \kappa, \infty) - (\mathfrak{U} \cup \mathfrak{V})$ to define $\mathcal{M}^r$ and the Seiberg–Witten Floer cochain complex.*

- *Let θ denote a nonzero class in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}})$, but with degree k or greater if $c_1(\det(\mathbb{S}))$ is torsion. Then θ has a representative cocycle with $\mathfrak{a}^f$ greater than $-\pi\mathcal{K}r$ and with $\mathrm{E} < 2\pi\mathcal{K}$ on each generator that appears with nonzero coefficient.*
- *Let $\mathfrak{A} = -\pi\mathcal{K}$. There exists $r_{\mathfrak{A}} \geq 1$ such that (4-2) holds when $r > r_{\mathfrak{A}}$.*
- *Let θ denote a nonzero class in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$, but with degree k or greater if $c_1(\det(\mathbb{S}))$ is torsion. Then θ has a representative cocycle with $\mathfrak{a}^f > -\pi\kappa\mathcal{K}r$ and with $\mathrm{E} < 2\pi\kappa\mathcal{K}$ on each generator that appears with nonzero coefficient.*
- *There exists $r_{\kappa\mathfrak{A}}$ such that (4-2) holds using $\kappa\mathfrak{A}$ in lieu of $\mathfrak{A}$ when $r > r_{\mathfrak{A}}$.*

Proof of Lemma 4.14 Let θ denote a nonzero class in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW}})$, but of degree k or greater when $c_1(\det(\mathbb{S}))$ is torsion. It follows from Proposition 4.9 that there exists $\mathcal{K} \geq 1$ and an unbounded set $\{r_i\}_{i=1,2,\ldots} \in [r_k, \infty) - (\mathcal{U} \cup \mathcal{V})$ such that when $r = r_i$, then θ has a cocycle representative, $\mathfrak{n}$, with $\mathrm{E} \leq 2\pi\mathcal{K}$ on each generator. Given that the cohomology with degree greater than any given integer is finitely generated, the constant $\mathcal{K}$ can be taken so as to be independent of the choice for θ. Note that $\mathcal{K}$ can be chosen so that (4-5) holds with $L = \mathcal{K}$.

Fix r_i and let $\mathfrak{n}$ denote the cocycle representative described above. It follows from Proposition 4.12 that $\mathrm{E} < 2\pi\mathcal{K} - c_0^{-1}$ on each generator that appears in $\mathfrak{n}$. The second and third items of Proposition 4.10 then find that $|\mathfrak{cs}^f| \leq c_0 r_i^{-1/50} r_i(1 + |\mathrm{E}|)$ on each generator that appears in $\mathfrak{n}$. This implies that $\mathfrak{a}^f_\theta > -2\pi(\mathcal{K} - c_0^{-1})r_i$. Given that $\mathfrak{a}^f_\theta[\cdot]$ is a continuous function, it follows from Proposition 4.13 that $\mathfrak{a}^f_\theta[r] > -\pi\mathcal{K}$ for $r > c_0$. It then follows from Propositions 4.11 and 4.12 that each $r > c_0$ version of θ has a cocycle representative with the property that $\mathfrak{a}^f > -\pi\mathcal{K}$ and $\mathrm{E} < 2\pi\mathcal{K} - c_0^{-1}$ on each generator.

Proposition 4.13 implies what is asserted by the second bullet of the lemma. To obtain the third bullet's assertion, let θ denote a nonzero class in $\mathcal{H}^*(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$, but of degree k or greater when $c_1(\det(\mathbb{S}))$ is torsion. Repeat the arguments for the first bullet using this class θ to find $c_0 > 1$ such that $\mathfrak{a}^f > -\pi c_0\mathcal{K}$ and $\mathrm{E} < 2\pi c_0\mathcal{K} - c_0^{-1}$. The final bullet again follows from Proposition 4.13. □

Part 2 Let $\mathcal{K}$ denote the constant given in Lemma 4.14 and set $\mathfrak{A} = -\pi\mathcal{K}$. It follows from Proposition 4.13 that (4-2) holds if $r > c_0$, and so (4-3) and (4-4) are well defined. Lemma 4.14 implies that the homomorphism $\mathcal{H}^j(\mathcal{C}^{\mathrm{SW},\mathfrak{A}}) \to \mathcal{H}^j(\mathcal{C}^{\mathrm{SW}})$ is surjective, at least if $j < k$ in the case when $c_1(\det(\mathbb{S}))$ is torsion.

Set $L = \mathcal{K}$ and fix some very small, but positive δ. Let $(\hat{a}, \hat{J})$ denote a (δ, L) approximation to (a, J). Fix $\mu \in \Omega$ with $\mathcal{P}$ norm less than 1 as described in the $(\hat{a}, \hat{J})$

version of Proposition 4.1. It follows from Proposition 4.13 that (4-2) is also obeyed by $(\hat{a}, \hat{J})$ and μ. Let $\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}}$ and $\hat{\mathcal{C}}^{\mathrm{SW}}$ denote the $\mathbb{Z}$ modules that appear in the latter's version of (4-3). It is a consequence of Lemma 4.6 that $\mathcal{H}^j(\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}}) \approx \mathcal{H}^j(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$, at least for $j > k$ when $c_1(\det(\mathbb{S}))$ is torsion. Meanwhile, the $(\hat{a}, \hat{J})$ and μ versions of Propositions 4.11 and 4.12 imply that $\mathrm{E} < 2\pi L - c_0^{-1}$ on all generators of $\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}}$ when $r > c_0$. Conversely, any generator of $\hat{\mathcal{C}}^{\mathrm{SW}}$ with $\mathrm{E} < 2\pi L + c_0^{-1}$ is a generator of $\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}}$. This follows from the second and third bullets in Proposition 4.10.

Granted the preceding, the $(\hat{a}, \hat{J})$ and μ versions of Theorems 4.2–4.4 identify $\mathcal{H}^j(\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}}) \approx H_{-j}(\hat{\mathcal{C}}^L_{\mathrm{ech}})$, at least for $j > k$ when $c_1(\det(\mathbb{S}))$ is torsion. This is what is claimed in the first bullet of Theorem 4.5.

Part 3 To address the second bullet of Theorem 4.5, let T_Φ denote the monomorphism given by the $(\hat{a}, \hat{J})$ version of Theorem 4.3. Suppose that $\upsilon \in \hat{\mathcal{C}}^{L_*}_{\mathrm{ech}}$ is such that $T_\Phi(\upsilon)$ is a coboundary. Let $\mathfrak{A}_* = -\pi L_*$. Given that T_Φ identifies $\mathcal{H}^j(\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}_*})$ with $H_{-j}(\hat{\mathcal{C}}^{L_*}_{\mathrm{ech}})$, it is sufficient to assume that the class $\hat{\lambda} \in \mathcal{H}^*(\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}_*})$ represented by $T_\Phi(\upsilon)$ is the image by the connecting homomorphism of (4-4) of a nonzero class in $\mathcal{H}^{j-1}(\hat{\mathcal{C}}^{\mathrm{SW}}/\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}_*})$. It follows from Lemma 4.6 that this class corresponds to one in $\mathcal{H}^{j-1}(\mathcal{C}^{\mathrm{SW}}/\mathcal{C}^{\mathrm{SW},\mathfrak{A}_*})$. Let θ denote the latter. If $r \geq c_0$, Propositions 4.12 and 4.13 find $\mathcal{K}_* \geq 1$ and a representative cocycle with $\mathrm{E} < 2\pi\mathcal{K}_*$ and $\mathfrak{a}^f > -\pi\mathcal{K}_*$ on each generator. If L is such that $L > \mathcal{K}_*$, then it follows that θ is represented by a cocycle in $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$. Meanwhile, $\hat{\lambda}$ corresponds via the isomorphism from Lemma 4.6 to a class $\lambda \in \mathcal{H}^j(\mathcal{C}^{\mathrm{SW},\mathfrak{A}_*})$. The fact that θ is represented by a cocycle in $\mathcal{C}^{\mathrm{SW},\mathfrak{A}}$ implies that λ is zero in $\mathcal{H}^j(\mathcal{C}^{\mathrm{SW},\mathfrak{A}})$. This implies (again via the isomorphism from Lemma 4.6) that $\hat{\lambda}$ is zero in $\mathcal{H}^j(\hat{\mathcal{C}}^{\mathrm{SW},\mathfrak{A}})$. Given the identification of the latter with $H_{-j}(\hat{\mathcal{C}}^L_{\mathrm{ech}})$, this means that υ is a boundary in $\hat{\mathcal{C}}^L_{\mathrm{ech}}$. □

4.i Proof of Theorem 1

The proof has three parts.

Part 1 This part returns to the notation and conventions of Section 3.h. To set the stage for what is to come, suppose g_0 and g_1 are given metrics on M, and that $\mathfrak{g}_0$ and $\mathfrak{g}_1$ are respective elements in the g_0 and g_1 versions of $\mathcal{P}$ and such that both $(g_0, \mathfrak{g}_0)$ and $(g_1, \mathfrak{g}_1)$ are suitable for defining realizations of the Seiberg–Witten Floer complex. Suppose that $\{(g_x, \mathfrak{g}_x)\}_{x\in[0,1]}$ is a given suitable path of data sets as described in Part 2 of Section 3.h; thus a path that can be used to construct the canonical isomorphism from the $(g_1, \mathfrak{g}_1)$ realization of the Seiberg–Witten Floer cohomology to the $(g_0, \mathfrak{g}_0)$ realization. Recall that this isomorphism, $I_{0,1}$, is induced by a homomorphism, $\hat{I}_{0,1}$ from the $(g_1, \mathfrak{g}_1)$ version of the Seiberg–Witten chain complex to that defined by

$(g_0, \mathfrak{g}_0)$. For each $x \in [0,1]$, let $\mathcal{Z}_x$ denote a finite set of generators for the $(g_x, \mathfrak{g}_x)$ version of the Seiberg–Witten Floer cochain complex. Suppose that each generator in $\mathcal{Z}_x$ is irreducible and nondegenerate. In addition, suppose the following: Let $\mathfrak{d}$ denote an instanton solution to the $(g_x, \mathfrak{g}_x)$ version of (3-3) whose $s \to \infty$ limit in $\mathbb{R} \times M$ is an element in $\mathcal{Z}_x$. Then

(4-14)
- The $s \to -\infty$ limit of $\mathfrak{d}$ is also in $\mathcal{Z}_x$.
- The Fredholm index of $\mathfrak{d}$'s version of (3-9) is nonnegative.

What follows is a final constraint, this concerning the variation with x of the family $\{\mathcal{Z}_x\}_{x\in[0,1]}$: Require that these sets be identified one with the other as x varies between 0 and 1. Here, the identification between $\mathcal{Z}_x$ and $\mathcal{Z}_{x'}$ for x near x' is that described in Part 5 of Section 3.h.

The following comment concerns the requirement in the second bullet of (4-14): The Smale–Sard theorem can be used (as, for example in [24, Section 7]) to find a residual set of suitable paths such that the second bullet is obeyed at each $x \in [0,1]$ by all instantons that interpolate between irreducible, nondegenerate solutions of the corresponding version of (3-6). This understood, the constraint given by the second bullet is easy to satisfy.

Introduce $\mathcal{C}\mathcal{Z}_x$ to denote the linear span of $\mathcal{Z}_x$. The assumptions about the family $\{\mathcal{Z}_x\}_{x\in[0,1]}$ have consequences with regards to $\{\mathcal{C}\mathcal{Z}_x\}_{x\in[0,1]}$, which are stated in the upcoming lemma. This lemma refers to the isomorphism $\iota\colon \mathcal{C}\mathcal{Z}_1 \to \mathcal{C}\mathcal{Z}_0$ that is induced by the identification between the respective generating sets $\mathcal{Z}_0$ and $\mathcal{Z}_1$ that comes by varying x from 0 to 1.

Lemma 4.15 *Given the preceding assumptions, the following is true:*

- *$\mathcal{C}\mathcal{Z}_x$ is mapped to itself by the Seiberg–Witten differential.*
- *The homomorphism $\hat{I}_{0,1}$ maps $\mathcal{C}\mathcal{Z}_1$ to $\mathcal{C}\mathcal{Z}_0$.*
- *The homomorphism $\hat{I}_{0,1}$ restricts to $\mathcal{C}\mathcal{Z}_1$ as $\iota + K$ where K has the following property: Introduce the function $\mathfrak{a}^f$ from (3-16). Let $\mathfrak{c} \in \mathcal{Z}_0$ denote any generator. Then $K(\mathfrak{c})$ can be written as a sum of elements in $\mathcal{C}\mathcal{Z}_1$ on which $\mathfrak{a}^f$ is strictly greater than $\mathfrak{a}^f(\mathfrak{c})$.*

Proof of Lemma 4.15 The first bullet follows from the assumption concerning the instantons with $s \to \infty$ limit in $\mathfrak{g}_x$. To prove the second, assume the contrary. Fix N large and break the interval $[0,1]$ as in Part 3 of Section 3.h into subintervals using break points $0 = x_0 < x_1 < \cdots < x_N = 1$ with $|x_k - x_{k-1}| < 2N^{-1}$ for each $k \in \{1, \ldots, N\}$ so as to factor the isomorphism $I_{0,1}$ as a suitable product of isomorphisms $\{I_k\}_{1\le k\le N}$. By way of reminder, the part of the path $\{(g_x,\mathfrak{g}_x)\}_{x\in[0,1]}$ with $x\in[x_k,x_{k-1}]$ supplies the data for defining the canonical isomorphism, I_k, from the realization of the Seiberg–Witten Floer cohomology defined by the $x=x_k$ member of the path to that defined by the $x=x_{k-1}$ member. Let $\hat{I}_k$ denote the corresponding homomorphism between the relevant cochain complexes. Given what was said about compositions in Part 3 of Section 3.h, there exists $k\in\{1,\dots,N\}$ such that $\hat{I}_k$ maps some element in $\mathcal{Z}_{x_k}$, the $x=x_k$ version of $\mathcal{Z}_x$, to some element in the $x=x_{k-1}$ Seiberg–Witten Floer cochain complex that is not $\mathcal{C}\mathcal{Z}_{x_{k-1}}$ in the latter's version of $\mathcal{C}\mathcal{Z}_x$. Taking a suitable limit of an $N\to\infty$ subsequence of the relevant instantons that define the cochain homomorphisms supplies $x\in[0,1]$ and a $(g_x,\mathfrak{g}_x)$ instanton whose $s\to\infty$ limit on $\mathbb{R}\times M$ is in $\mathcal{Z}_x$ and whose $x\to-\infty$ limit is not in $\mathcal{Z}_x$. That this is so follows from what is said about limits in Chapter 24.6 of [16]. This event is precluded by the assumptions.

The same sort of limiting argument proves the third assertion using two additional observations. What follows is the first. Let $\mathfrak{d}$ denote a given $(g_x,\mathfrak{g}_x)$ instanton and let $\mathfrak{c}_-$ and $\mathfrak{c}_+$ denote the respective $s\to-\infty$ and $s\to+\infty$ limits of $\mathfrak{d}(s)$. Then $\mathfrak{a}(\mathfrak{c}_-)\ge\mathfrak{a}(\mathfrak{c}_+)$ with equality if and only if $\mathfrak{c}_-=\mathfrak{c}_+$ and $\mathfrak{d}$ is $\mathbb{R}$–invariant. Meanwhile, the spectral flow function, f, is such that $f(\mathfrak{c}_+)-f(\mathfrak{c}_-)\ge 0$. Indeed, this difference is the Fredholm index of $\mathfrak{d}$'s version of (3-9), and this index is nonnegative by assumption. Thus, $\mathfrak{a}^f(\mathfrak{c}_-)\ge\mathfrak{a}^f(\mathfrak{c}_+)$ unless $\mathfrak{c}_-=\mathfrak{c}_+$ and $\mathfrak{d}$ is $\mathbb{R}$–invariant. The subset of $(g_x,\mathfrak{g}_x)$ instantons that increase $\mathfrak{a}^f$ account for what is denoted by K in the third bullet of (4-14).

What follows is the second observation. The appearance of ι in the third bullet of (4-14) can be seen using perturbation theory. To elaborate, perturbation theory with what is said in Chapter 24 of [16] can be used to prove the following: If N is sufficiently large and if $k\in\{1,\dots,N\}$, then any given $\mathbb{R}$–invariant instanton that is defined by an element in $\mathcal{Z}_{x_k}$ defines a cobordism instanton whose limit as the $\mathbb{R}$ parameter tends to ∞ in $\mathbb{R}\times M$ is the given element and whose limit as this parameter tends to $-\infty$ is the canonical partner in $\mathcal{Z}_{x_{k-1}}$. In addition, this cobordism instanton is the only one that interpolates between these two generators. Finally, the sign contribution from this cobordism instanton is $+1$. $\square$

Part 2 Fix a pair (a,J) of contact 1–form and almost complex structure suitable for defining $\{\mathcal{H}^L_{\mathrm{ech}}\}_L\ge 1$. The values of L that are realized as $\sum_{(m,\gamma)\in\Theta} m\ell_\gamma$ with Θ a generator of the embedded contact homology chain complex form a discrete set with no accumulation points. Let Λ_{ech} denote this set. Choose an increasing sequence $\{L_k\}_{k=1,2,\dots}\subset[1,\infty)$ so that no member is in Λ_{ech} and so that each interval of the form $[L_k,L_{k+1}]$ contains precisely one member of Λ_{ech}. Use ψ_k in what follows to denote the direct limit homomorphism from $\mathcal{H}^{L_k}_{\mathrm{ech}}$ to $\mathcal{H}_{\mathrm{ech}}$.

For each $k \in \{1, 2, \dots\}$, fix a preferred, (δ, L_k) approximation to (a, J). Denote the latter by $(\hat{a}_k, \hat{J}_k)$. Preferred (δ, L) approximations are described in Proposition B.1 of the Appendix. Fix an increasing sequence $\{r_k\}_{k=1,2,\dots} \subset [1, \infty)$, a sequence $\{\mu_k\}_{k=1,2,\dots}$ of 1–forms, and a sequence $\{\mathfrak{p}_k\}_{k=1,2,\dots} \subset \mathcal{P}$ so that the following is true for each k: The data $(\hat{a}_k, \hat{J}_k)$, μ_k and $\mathfrak{p}_k$ is suitable for defining the Seiberg–Witten cochain complex. Assume in addition that Theorems 4.2–4.4 can be invoked using r_k, with the data $(\hat{a}_k, \hat{J}_k)$, μ_k and $\mathfrak{p}_k$. Further constraints on these sequences will be described momentarily.

If $c_1(\det(\mathbb{S}))$ is not torsion, then $\mathcal{H}^{\mathrm{SW}}$ is a finitely generated $\mathbb{Z}$–module. If $c_1(\det(\mathbb{S}))$ is torsion, then any fixed degree summand in $\mathcal{H}^{\mathrm{SW}}$ is finitely generated. This understood, let $\mathcal{H}$ denote either the whole of $\mathcal{H}^{\mathrm{SW}}$ if $c_1(\det(\mathbb{S}))$ is torsion, or a fixed degree summand of $\mathcal{H}^{\mathrm{SW}}$ if $c_1(\det(\mathbb{S}))$ is a torsion class. Use $\mathcal{H}_k$ to denote the realization of $\mathcal{H}^{\mathrm{SW}}$ that is defined using the data $(\hat{a}_k, \hat{J}_k), r_k$, μ_k and $\mathfrak{p}_k$. Use $\mathcal{C}_k$ to denote the corresponding free $\mathbb{Z}$–module of cochains.

Fix $k \in \{1, 2, \dots\}$ and define T_k to be the $(\hat{a}_k, \hat{J}_k)$ version of the injective homomorphism T_Φ as defined in Section 4.b. Thus, T_k maps $\mathcal{C}_{\mathrm{ech}}^{L_k}$ to the $((\hat{a}_k, \hat{J}_k), r, \mu_k)$ version of the Seiberg–Witten cochain complex, $\mathcal{C}_k$. This homomorphism T_k induces a homomorphism, also denoted by T_k, from ${\mathcal{H}_{\mathrm{ech}}}^{L_k}$ to $\mathcal{H}_k$. Since $\mathcal{H}_k$ is finitely generated, it follows from what is said by the first bullet of Theorem 4.5 that T_k is surjective if k is sufficiently large. Fix such k and denote it by k_*. Fix $k_{**} > k_*$ so that the conclusions of the second bullet in Theorem 4.5 hold with $L_* = L_{k_*}$ and with $L = L_{k_{**}}$.

Note for future reference that any given $k \geq k_{**}$ and sufficiently large r version of $\mathcal{H}_k$ comes with a filtration,

$$\mathcal{H}_k{}^1 \subseteq \cdots \subseteq \mathcal{H}_k{}^{k_*} = \mathcal{H}_k, \tag{4-15}$$

where a given $k' \in \{1, \dots, k_*\}$ submodule consists of elements that are represented by cochains from the image via T_k of $\mathcal{H}_{\mathrm{ech}}^{L_{k'}}$. Here is another way to say this: An element in $\mathcal{H}_k{}^{k'}$ is represented by a sum of generators that are given by equivalence classes of the $((\hat{a}_k, \hat{J}_k), r_k, \mu_k)$ version of (3-17) on which the function E from (3-13) is less than $2\pi L_{k'}$.

Part 3 As explained momentarily, a homomorphism $Q\colon \mathcal{H}_{k_{**}} \to \mathcal{H}_{\mathrm{ech}}$ can be defined as follows: Given $\mathfrak{z} \in \mathcal{H}_{k_{**}}$, take an element $z \in \mathcal{H}_{\mathrm{ech}}^{L_{k_*}}$ that is mapped by $T_{k_{**}}$ to $\mathfrak{z}$. Set $Q(\mathfrak{z}) = \psi_{k_*}(z)$. To see that this is well defined, suppose that z' is an alternate choice. As the image via $T_{k_{**}}$ of $z - z'$ is zero, it follows from what is said by the second bullet of Theorem 4.5 that z and z' are mapped to zero in $\mathcal{H}_{\mathrm{ech}}^{L_{k_{**}}}$. This being the case, their difference is mapped to zero in $\mathcal{H}_{\mathrm{ech}}$ by ψ_{k_*}.

What follows next explains why Q is injective. To start, suppose that $\mathfrak{z} \in \mathcal{H}_{k_{**}}$ is mapped to zero by Q. Let $k \in \{1, \dots, k_*\}$ denote the minimal integer such that $\mathfrak{z} \in {\mathcal{H}_{k_{**}}}^k$. Let $z \in {\mathcal{H}_{\mathrm{ech}}}^{L_k}$ denote an element that is mapped via $T_{k_{**}}$ to $\mathfrak{z}$. To say that $Q(\mathfrak{z})$ is zero means that there exists some $k' \geq k_{**}$ such that z is zero in $\mathcal{H}_{\mathrm{ech}}^{L'_k}$. This implies that $T_{k'}(z)$ is zero in the realization $\mathcal{H}_{k'}$.

Lemma 4.16 *There exists a cobordism isomorphism, $\hat{I}$, that maps the Seiberg–Witten cochain complex $\mathcal{C}_{k_{**}}$ to the corresponding complex $\mathcal{C}_{k'}$. This homomorphism restricts to the image via the cochain version of $T_{k_{**}}$ so as to define a homomorphism from ${\mathcal{C}_{\mathrm{ech}}}^{L_k}$ into $\mathcal{C}_{k'}$ with the following properties:*

- *This homomorphism from ${\mathcal{C}_{\mathrm{ech}}}^{L_k}$ into $\mathcal{C}_{k'}$ maps into the image via $T_{k'}$ of ${\mathcal{C}_{\mathrm{ech}}}^{L_k}$ and so induces an endomorphism from ${\mathcal{C}_{\mathrm{ech}}}^{L_k}$ to itself.*
- *The induced endomorphism of ${\mathcal{C}_{\mathrm{ech}}}^{L_k}$ maps any given cycle $\hat{z}$ to one of the form $\hat{z} + \hat{u}$ where $\hat{u} \in \mathcal{C}_{\mathrm{ech}}^{L_{k-1}}$.*

Granted Lemma 4.16, suppose that $\hat{z}$ is a closed cycle that gives the class z. Use the second bullet in Lemma 4.16 to define $\hat{u}$. Note that the latter is also closed. This follows from Theorem 4.3 since $\hat{I}$ intertwines the differentials for the respective Seiberg–Witten Floer complexes. Let u denote the class in $\mathcal{H}_{\mathrm{ech}}^{L_{k-1}}$ defined by $\hat{u}$. The class z is mapped by I to the class $T_{k'}(z+u)$. As $T_{k'}(z) = 0$, this implies that $T_{k_{**}}(z) = -T_{k_{**}}(u)$. Given the definition of k, the latter conclusion implies that $\mathfrak{z} = 0$.

The homomorphism Q is also surjective onto the relevant degree summand in $\mathcal{H}_{\mathrm{ech}}$. To see that such is the case, fix a given class in $\mathcal{H}_{\mathrm{ech}}$ with the relevant degree. Let k denote the minimum integer in the range $\{k_*, k_*{+}1, \dots\}$ such this class is represented by an element in ${\mathcal{H}_{\mathrm{ech}}}^{L_k}$. Use z' to denote the latter element. Fix $k' > k$ so that the conclusions of Theorem 4.5 hold for $L = L_k$ and $L_* = L_{k'}$. It follows from Theorem 4.5 that $T_{k'}(z') \neq 0$ in $\mathcal{H}_{k'}$. Use I to again denote the cobordism isomorphism from $\mathcal{H}_{k_{**}}$ to $\mathcal{H}_{k'}$. Then $T_{k'}(z')$ can be written as $I(T_{k_{**}}(z))$ for some $z \in \mathcal{H}_{\mathrm{ech}}^{L_{k*}}$. It then follows from Lemma 4.16 that $T_{k'}(z') = T_{k'}(z'')$ where z'' is represented by a chain in ${\mathcal{C}_{\mathrm{ech}}}^{L_{k*}}$. Granted this last point, then the second bullet in Theorem 4.5 implies that the given class in $\mathcal{H}_{\mathrm{ech}}$ is also represented by z''. Therefore, the given class from $\mathcal{H}_{\mathrm{ech}}$ is in the image of Q.

Proof of Lemma 4.16 Keep in mind that each $k \in \{1, 2, \dots\}$ version of $(\hat{a}_k, \hat{J}_k)$ is a preferred $(L{=}L_k, \delta)$ approximation to (a, J). Fix such k and then $k' > k$. Proposition B.1 supplies a certain sort of family, $\{(\hat{a}^x, \hat{J}^x)\}_{x\in[0,1]}$ of contact 1–form and almost complex structure whose $x = 0$ member is $(\hat{a}_k, \hat{J}_k)$ and whose $x = 1$ member is $(\hat{a}_{k+1}, \hat{J}_{k+1})$. If r is sufficiently large, then r with any given $x \in [0, 1]$

version of $(\hat{a}^x, \hat{J}^x)$ can be used as input for Theorem 4.2 for the choice $L = L_k$. Keep this fact in mind.

For each $x \in [0,1]$, let g_x denote the metric on M that is defined by $(\hat{a}^x, \hat{J}^x)$. Fix a smoothly varying path $\{x \to \mu_x\}_{x \in [0,1]}$ of 1–forms that interpolates between μ_{k-1} and μ_k, and such that any given μ_x is in the g_x version of Ω. Fix a family $\{x \to r_x\}_{x\in[0,1]}$ with all members sufficiently large, and in particular such that the data $((\hat{a}^x, \hat{J}^x), r_x, \mu_x)$ for any $x \in [0,1]$ can be used as input for Theorem 4.2 and for the first bullet in Theorem 4.4. Let $\mathcal{Z}_x$ for $x \in [0,1]$ denote the image of $\mathcal{C}_{\mathrm{ech}}^{L_k}$ via the corresponding version of Φ^r. What follows explains why the 1–parameter family $\{x \to \mathcal{Z}_x\}_{x\in[0,1]}$ obeys the conditions that are needed to invoke Lemma 4.15 for the path $\{(g_x, \mathfrak{g}_x = \mathfrak{e}\mu_x + \mathfrak{p}_x)\}_{x\in[0,1]}$ for certain paths $\{\mathfrak{p}_x\}_{x\in[0,1]}$ of elements in $\mathcal{P}$. The conclusions of this lemma imply what is asserted by Lemma 4.16.

The conditions for Lemma 4.15 are discussed in order opposite to the order given. To start, note that the family $\{\mathcal{Z}_x\}_{x\in[0,1]}$ varies in the appropriate, smooth manner as x varies; this is a consequence of Theorem 4.2 if each $x \in [0,1]$ version of $\mathfrak{p}_x$ has very small norm. In fact, this path can be chosen so that each $\mathfrak{p}_x$ vanishes to second order on the elements in $\mathcal{Z}_x$.

To argue for the second bullet of (4-14), remark that the path $\{\mathfrak{p}_x\}_{x\in[0,1]}$ can be chosen so that the following is true for each $x \in [0,1]$:

(4-16) Every instanton solution to the $(g_x, r_x, \mathfrak{g}_x)$ version of (3-6) has a corresponding version of (3-9) with nonnegative Fredholm index.

Choose $\{\mathfrak{p}_x\}_{x\in[0,1]}$ so that (4-16) holds. As noted in Part 1 above, this is always possible. The second bullet in (4-14) is therefore satisfied for each $x \in [0,1]$. To see about the first bullet, of (4-14), let $\mathfrak{d}$ denote an instanton solution to a given $x \in [0,1]$ version of (3-6) with $s \to \infty$ limit given by $\mathfrak{c}_+ \in \mathcal{Z}_x$. It follows from Theorem 4.2 that $E(\mathfrak{c}_+) < 2\pi L_k$. If $\mathfrak{p}_x$ has sufficiently small norm then the first bullet of Theorem 4.4 guarantees that $\mathfrak{c}_- \in \mathcal{Z}_x$. □

5 Theorems 4.2 and 4.3

The constructions that lead to Φ^r and Ψ^r and the arguments for Theorems 4.2 and 4.3 are modifications of those used in [21; 22] to prove the equivalence between the Gromov and Seiberg–Witten invariants of compact, symplectic 4–manifolds. What follows in this section is a brief description of what is involved. The details are contained in Papers II, III and IV of this series [26]–[28].

5.a Vortices on $\mathbb{C}$

Both Φ^r and Ψ^r use the solutions to the vortex equations on $\mathbb{C}$ to construct pairs of connection on E and section of $\mathbb{S}$ from Reeb orbit or pseudoholomorphic curves as the case may be. Solutions to the vortex equations played a similar role in the constructions done in [22]; see its article Gr $\Rightarrow$ SW. What follows in this subsection provides a brief summary of the vortex part of the story.

The vortex moduli spaces are labeled by a nonnegative integer, with the integer n version of the vortex moduli space denoted by $\mathfrak{C}_n$. The latter consists of certain equivalence classes of pairs (A, α), where A is a hermitian connection on the trivial complex line bundle over $\mathbb{C}$, and where α is a section of this bundle. A pair $\mathfrak{c} = (A, \alpha)$ is in $\mathfrak{C}_n$ if and only if the curvature of A and the A–covariant derivative of α satisfy

(5-1)

- $*F_A = -i(1-|\alpha|^2)$.
- $\bar\partial_A \alpha = 0$.
- $|\alpha| \le 1$.
- The function $(1-|\alpha|^2)$ is integrable on $\mathbb{C}$ and $\int_{\mathbb{C}} (1-|\alpha|^2) = 2\pi n$.

Here, $\bar\partial_A$ denotes the d–bar operator that is defined by the connection A. The equivalence relation that defines a point in $\mathfrak{C}_n$ identifies pairs (A, α) and (A', α') when $A' = A - u^{-1}du$ and $\alpha' = u\alpha$ where u is a smooth map from $\mathbb{C}$ to S^1.

The space $\mathfrak{C}_0$ consists of a single element, the gauge equivalence class of the pair $(A{=}0,\ \alpha{=}1)$. When $n \ge 1$, the space $\mathfrak{C}_n$ has the structure of a smooth, complex manifold that is biholomorphic to $\mathbb{C}^n$. This holomorphic identification is realized as follows: As is explained by the author in Section 2 of the article Gr $\Rightarrow$ SW from [22], [20] and [15], if (A, α) solves (5-1), then α has precisely n zeros counting multiplicities. Let $\mathfrak{Z}_{\mathfrak{c}} = \{z_1, \ldots, z_n\}$ denote the resulting set in the n–th symmetric product of $\mathbb{C}^n$. A holomorphic diffeomorphism from $\mathfrak{C}_n$ to $\mathbb{C}^n$ sends $\mathfrak{c}$ to the point in $\mathbb{C}^n$ whose q–th coordinate is $\sum_{1\le j\le n} z_j^q$. As it turns out

$$\sum_{1\le j\le n} z_j^q = \frac{1}{2\pi}\int_{\mathbb{C}} z^q(1-|\alpha|^2). \tag{5-2}$$

With regards to the integral in (5-2), note that

$$0 < 1-|\alpha|^2 < c_0 \sum_{1\le j\le m} e^{-\sqrt{2}|z-z_j|} \quad \text{and} \quad |\nabla_A \alpha|^2 \le c_0 \sum_{1\le j\le m} e^{-\sqrt{2}|z-z_j|}, \tag{5-3}$$

where c_0 is a constant that is independent of n and $(A, \alpha) \in \mathfrak{C}_m$.

This holomorphic identification to $\mathbb{C}^n$ does *not* provide the natural Riemannian metric on $\mathfrak{C}_n$ when $n > 1$. The relevant metric is described momentarily. To set the stage for the story on the metric, remark that the $(1,0)$ tangent space to $\mathfrak{C}_n$ at a given $\mathfrak{c} = (A, \alpha)$ is naturally isomorphic to a certain vector space of pairs (x, ι) where x is a $\mathbb{C}$–valued function on $\mathbb{C}$ and ι is a section of the trivial bundle. To lie in $T_{1,0}\mathfrak{C}_n|_{\mathfrak{c}}$, both x and ι must be square integrable on $\mathbb{C}$ and obey the coupled system

$$\partial x + 2^{-1/2}\bar{\alpha}\iota = 0 \quad \text{and} \quad \bar{\partial}_A \iota + 2^{-1/2}\alpha x = 0. \tag{5-4}$$

In this equation, ∂ is shorthand for $\partial/\partial z$. The pair (x, ι) provides the first order change in (A, α), which adds $2^{-1/2}x$ to the $(0,1)$ part of A and adds ι to α. The relevant metric on $\mathfrak{C}_n$ is defined so that the metric norm of (x, ι) is $\pi^{-1/2}$ times its L^2 norm as defined by integration on $\mathbb{C}$. This metric is a Kähler metric with respect to the complex structure.

5.b Constructing the map Φ^r in Theorem 4.2

This subsection is meant to give an indication of what is involved in constructing the map Φ^r. To set the stage, write the spinor bundle $\mathbb{S}$ as $E \oplus EK^{-1}$. With $L > 1$ fixed, introduce as notation $\mathcal{Z}^L$ to denote the set whose typical element, Θ, consists of pairs of the form (γ, m) where γ is a Reeb orbit and m a positive integer. Require that distinct pairs from Θ have distinct Reeb orbit components. In addition, require $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma < L$; and require that the formal sum $\sum_{(\gamma,m)\in\Theta} m\gamma$ define a cycle whose class in $H_1(M;\mathbb{Z})$ is Poincaré dual to $c_1(E)$.

The subsequent description of Φ^r has four parts.

Part 1 Fix $\Theta \in \mathcal{Z}^L$. Assign to each $(\gamma, m) \in \Theta \in \mathcal{Z}^L$ a smooth map, $\mathfrak{c}_\gamma\colon S^1 \to \mathfrak{C}_m$. The first step to defining Φ^r is to construct a pair in $\mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$ from the data $\{\mathfrak{c}_\gamma\}_{(\gamma,m)\in\Theta}$. The map $\mathfrak{c}_\gamma\colon S^1 \to \mathfrak{C}_m$ is lifted so as to give a connection on, and a section of, the product complex line bundle over $S^1 \times \mathbb{C}$. The pullback of this pair to any given constant $t \in S^1$ slice of $S^1 \times \mathbb{C}$ is a solution to (5-1). Define $\hat{r}_\gamma\colon S^1 \times \mathbb{C} \to S^1 \times \mathbb{C}$ so that

$$\hat{r}_\gamma^*(t, z) = \left(t, \left(\frac{\ell_\gamma}{2\pi} r\right)^{1/2} z\right).$$

Use $(A^{(\gamma)}, \alpha^{(\gamma)})$ to denote the pullback via $\hat{r}_\gamma$ of the chosen lift of $\mathfrak{c}_\gamma$. As might be expected from (5-3), the connection $A^{(\gamma)}$ is nearly flat where $|z| \gg r^{-1/2}$ on $S^1 \times \mathbb{C}$. Meanwhile, $\alpha^{(\gamma)}$ is nearly $A^{(\gamma)}$–covariantly constant with norm 1 on this same part of $S^1 \times \mathbb{C}$.

Suppose now that a tubular neighborhood map has been chosen for each Reeb orbit from Θ of the sort that is described in Section 2.a. Use such a map to identify a tubular neighborhood of each Reeb orbit with $S^1 \times D \subset S^1 \times \mathbb{C}$. Then $(A^{(\gamma)}, \alpha^{(\gamma)})$ can be viewed as a pair of connection on and section of the product complex line bundle over a tubular neighborhood of γ in M with $A^{(\gamma)}$ nearly flat near the boundary of this tubular neighborhood, and with $\alpha^{(\gamma)}$ having norm 1 and nearly covariantly constant near this same boundary.

This understood, the collection $\{(A^{(\gamma)}, \alpha^{(\gamma)})\}_{(\gamma,m)\in\Theta}$ are pasted together using "bump" functions so as to obtain a pair (A^*, α^*) of connection on and section of a complex line bundle over M. The latter is isomorphic to E. Note that A^* is flat except very near the Reeb orbits from Θ; and likewise α^* is covariantly constant with norm 1 except very near these same Reeb orbits.

Write $\mathbb{S} = E \oplus EK^{-1}$ and define $(A^*, \psi^* = (\alpha^*, 0)) \in \mathrm{Conn}(E) \times C^\infty(M;\mathbb{S})$. A calculation shows that (A^*, ψ^*) comes reasonably close to solving (3-17) if r is large. Such is the case by virtue of the fact that the pairs in Θ involve *Reeb orbits*. Indeed, the construction just described can be applied to a set such as Θ whose typical element is a pair (γ, m) where γ is an embedded loop in M. If γ is not a Reeb orbit, then the resulting (A^*, ψ^*) will not come close to solving (3-17) when r is large.

The plan is to look for a solution to (3-17) near (A^*, ψ^*). Such a solution can be found when r is large if the collection $\{\mathfrak{c}_\gamma \colon S^1 \to \mathfrak{C}_m\}_{(\gamma,m)\in\Theta}$ are suitably constrained.

Part 2 What follows describes the constraints on $\{\mathfrak{c}_\gamma\}_{(\gamma,m)\in\Theta}$. To this end, return to the vortex moduli space $\mathfrak{C}_m$. Let (ν, μ) denote a pair consisting of a real number and a complex number. Any such pair defines a function, h, on $\mathfrak{C}_m$ given by

$$h = \frac{1}{4\pi}\int_{\mathbb{C}} (2\nu|z|^2 + (\mu\bar{z}^2 + \bar{\mu}z^2))(1-|\alpha|^2). \tag{5-5}$$

As with any function on $\mathfrak{C}_m$, this one defines a Hamiltonian vector field. Now suppose that ν and μ are respectively a real valued function on S^1 and a $\mathbb{C}$–valued function on S^1. Then (5-5) defines a 1–parameter family of Hamiltonian vector fields on $\mathfrak{C}_m$. Of interest are the closed, integral curves of the latter. These are maps $\mathfrak{c} \colon S^1 \to \mathfrak{C}_m$ that obey at each $t \in S^1$ the equation

$$\frac{i}{2}\mathfrak{c}' + \nabla^{(1,0)}h|_{\mathfrak{c}} = 0, \tag{5-6}$$

where $\mathfrak{c}'$ is shorthand for the $(1,0)$ part of $\mathfrak{c}_*(d/dt)$, and where $\nabla^{(1,0)}h$ denotes the $(1,0)$ part of the gradient of h.

Now suppose that γ is a Reeb orbit. Fix a tubular neighborhood map for γ of the sort described in Section 2.a. Then γ has an associated pair (ν, μ) for use in (5-5), namely the pair that appears in (2-1) and (2-3). With the preceding understood, what follows is the key observation. Suppose that the following is true:

(5-7) Each $(\gamma, m) \in \Theta$ version of $\mathfrak{c}_\gamma$ is a solution to the corresponding version of (5-6).
In addition, the linearized version of the left hand side of (5-6) at this $\mathfrak{c}_\gamma$ defines an operator with trivial kernel.

Note that the linearization of (5-6) at a given map $\mathfrak{c}\colon S^1 \to \mathfrak{C}_m$ defines a first order, elliptic and symmetric operator on $C^\infty(S^1, \mathfrak{c}^* T^{1,0}\mathfrak{C}_m)$.

Under the assumption in (5-7), perturbation theory can be employed to modify (A^*, ψ^*) when r is large so that the result solves the corresponding version of (3-17). To put this in a more formal way, introduce $\mathfrak{C}\Theta$ for $\Theta \in \mathcal{Z}^L$ to denote the set whose typical element assigns to each $(\gamma, m) \in \Theta$ a corresponding solution to (5-6). Say that a solution to (5-6) is *nondegenerate* when the linearization of the left hand side of (5-6) at the solution has trivial kernel. Use $\mathfrak{C}\Theta^*$ to denote the subset where each assigned solution is nondegenerate. Given $L \geq 1$, let $\mathfrak{C}\mathcal{Z}^L$ denote $\{\mathfrak{C}\Theta : \sum_{(\gamma,m)\in\Theta} m\ell_\gamma < L\}$. Use $\mathfrak{C}\mathcal{Z}^{L*}$ to denote the subset $\{\mathfrak{C}\Theta^* : \Theta \in \mathcal{Z}^L\} \subset \mathfrak{C}\mathcal{Z}^L$.

(5-8) If the contact form comes from Lemma 2.1's residual subset, if $\mathfrak{C}\mathcal{Z}^{L*} = \mathfrak{C}\mathcal{Z}^L$, and if r is sufficiently large, then perturbation theory defines an injective map $\Phi^r\colon \mathfrak{C}\mathcal{Z}^L \to \mathcal{M}^r$ whose image consists of the set of elements with $E < 2\pi L$.

This map is constructed in [26]; and [28] proves that it is injective and surjective onto the $E < 2\pi L$ subset in $\mathcal{M}^r$. It is fair to say that these parts of [26; 28] do little more than reinterpret parts of the respective Gr $\Rightarrow$ SW and Gr $=$ SW and articles in [22].

If $\mathfrak{C}\mathcal{Z}^L \neq \mathfrak{C}\mathcal{Z}^{L*}$, then perturbation theory constructs, for each sufficiently large r, an injective map from a certain subset of $\mathfrak{C}\mathcal{Z}^L$ into $\mathcal{M}^r$ whose image consists of the set of elements with $E < 2\pi L$. The latter is also denoted by Φ^r in what follows.

Part 3 What follows says some things about the space of solutions to (5-6). To start, note that the solution space to (5-6) is compact if (ν, μ) is either hyperbolic or m–elliptic. This is proved in [28]. In either case, there exists a unique solution for $m = 1$; this is the vortex with $\alpha^{-1}(0) = \{0\}$. The unique $m = 1$ solution to (5-6) is nondegenerate if (ν, μ) is nondegenerate. In this $m = 1$ case, the corresponding linear operator is the operator L that is depicted in (2-3).

It is not known whether the solution space to the $m>1$ versions of (5-6) consists of solely nondegenerate solutions, even if (ν,μ) are chosen in a generic fashion. However, if each Reeb orbit with $\ell_\gamma \le L$ is either hyperbolic or m–elliptic, and if each of the corresponding versions of (5-6) has solely nondegenerate solutions when $\ell_\gamma \le L$, then $\mathfrak{C}\mathcal{Z}^L$ is a finite set.

All solutions to (5-6) are known for some specific choices of ν and μ:

- $(\nu,\mu)=(\frac{1}{4}k, i\varepsilon e^{ikt})$ is hyperbolic with rotation number k. Here, $\varepsilon>0$ but very small. Then there are no $m>1$ solutions to (5-6).

(5-9)
- $(\nu,\mu)=(\frac{1}{2}\mathrm{R},0)$ with R irrational. Then there is a unique solution to (5-6) for each m, this being the vortex with $\alpha^{-1}(0)=\{0\}$. Moreover, the latter is nondegenerate.

These last facts are proved in [26]. If (5-9) holds for each Reeb orbit γ with $\ell_\gamma < L$, then the set $\mathfrak{C}\mathcal{Z}^L$ is precisely the set $\mathcal{Z}^L_{\mathrm{ech}}$ that gives the generators of the embedded contact homology subcomplex $\mathcal{C}^L_{\mathrm{ech}}$. In this case, Φ^r is a map from $\mathcal{Z}^L_{\mathrm{ech}}$ into $\mathcal{M}^r$; this is the map used in Theorem 4.2.

Part 4 As it turns out, the proof that embedded contact homology is isomorphic to Seiberg–Witten Floer cohomology does not require knowledge of all solutions to (5-6). Knowledge of the corresponding Hamiltonian Floer cohomology groups is sufficient.

To elaborate, Floer [2; 3] introduced his celebrated "Floer (co)homology" to resolve a famous conjecture of Arnold that concerned closed orbits of time-dependent Hamiltonian vector fields on symplectic manifolds. What is written in (5-6) is an example of just such a Hamiltonian dynamical system. In particular, if (ν,μ) is a nondegenerate pair, then there are well-defined, $\mathbb{Z}$–graded Floer homology and cohomology groups whose generators are solutions to a suitably generic compactly supported (on $\mathfrak{C}_m$) perturbation of (5-6).

There is one subtle point here, involving the instantons that define the differentials. In this context, an instanton is a smooth map from $\mathfrak{c}\colon \mathbb{R}\times S^1 \to \mathfrak{C}_m$ that obeys the equation

$$\bar\partial\mathfrak{c} + \nabla^{1,0}h|_{\mathfrak{c}} = 0, \tag{5-10}$$

where $\bar\partial$ is a suitably defined version of the d–bar operator on $\mathfrak{c}^*T^{1,0}\mathfrak{C}_m \to S^1\times\mathbb{R}$. The map $\mathfrak{c}$ must also limit as $s\to\pm\infty$ to a solution of (5-6). In order to obtain a well-defined differential for the Hamiltonian Floer (co)homology, it is necessary to prove that the moduli space of instanton solutions to (5-10) can be compactified by adding "broken trajectories". This can be done when (ν,μ) are either hyperbolic or m–elliptic.

In any event, it can be shown that there are well-defined Hamiltonian Floer (co)homology groups for (5-6) when (ν, μ) are either hyperbolic or m–elliptic. Furthermore, it can be shown (using (5-9)) that the Hamiltonian Floer cohomology groups are as follows:

$$\begin{cases} \mathbb{Z} & \text{if } m = 1. \\ 0 & \text{if } m > 1 \text{ and } (\nu, \mu) \text{ is hyperbolic.} \\ \mathbb{Z} & \text{if } m > 1 \text{ and } (\nu, \mu) \text{ is } m\text{–elliptic.} \end{cases} \tag{5-11}$$

Suppose that each $\Theta \in \mathcal{Z}^L$ contains only pairs (γ, m) such that γ is either hyperbolic or m–elliptic. Granted only this assumption, it is nonetheless the case that the points in $\mathfrak{C}\Theta$ that are mapped by Φ^r into $\mathcal{M}^r$ carry, in a suitable sense, the product of the Hamiltonian Floer cohomology groups as defined by the various $(\gamma, m) \in \Theta$ versions of (5-6). Somewhat more is said about this in the next subsection.

5.c Constructing the map Ψ^r in Theorem 4.3

This section is meant to give a rough indication of how Ψ^r is constructed. There are three parts to what follows. The final two parts say something about what is involved when (5-9) does not hold.

Part 1 The idea is to mimic as much as possible what is done in the Gr $\Rightarrow$ SW article in [22]. As done there, the first step constructs an approximate solution to (3-18) such that the connection is flat except very near to the given curve in $\mathcal{M}_1(\Theta_-, \Theta_+)$, the section of $\mathbb{S}$ lies only in the E summand, and this section of E is covariantly constant with norm 1 except very near the given curve. Step 2 uses perturbation theoretic techniques to find an honest solution to (3-18) that differs little from the approximate one. There are, however, serious new issues that do not arise in [22], relating to the behavior of the elements of $\mathcal{M}_1(\Theta_-, \Theta_+)$ where $|s|$ is large on $\mathbb{R} \times M$.

To elaborate on this last point, suppose that $\Sigma \subset \mathbb{R} \times M$ is an embedded, pseudo-holomorphic curve. As such, Σ has a well-defined normal bundle, $N \to \Sigma$, and a fixed radius disk bundle $N_1 \subset N$ with an exponential map $e_\Sigma \colon N_1 \to \mathbb{R} \times M$ that immerses N_1 and embeds a neighborhood of the zero section as a neighborhood of Σ. Even so, there need not exist a fixed radius disk subbundle of N that is everywhere embedded by e_Σ. The point being that the constant s slices of distinct ends of Σ can limit as $s \to \pm\infty$ to the same Reeb orbit. In addition, the constant s slices of any given end need not define a degree 1 braid in the tubular neighborhood of the nearby Reeb orbit.

These remarks about the fixed radius disk bundle are relevant because the constructions in the article Gr $\Rightarrow$ SW from [22] require an embedding of just such a bundle. When a fixed radius disk bundle is embedded by e_Σ, then the constructions in [22] can be copied with only minor changes to produce an approximate solution to (3-18) and then a deformation of the latter to an honest solution. In this regard, the approximate solution has the following appearance: Use the exponential map to identify the fixed radius subbundle of N_1 with a tubular neighborhood of Σ in $\mathbb{R}\times M$. The pullback of the connection and the section of E to any given fiber of N_1 differs little from the pullback of a solution to (5-1) via the map from $\mathbb{C}$ to $\mathbb{C}$ that sends any given $z\in\mathbb{C}$ to $r^{1/2}z$.

In general, only the following can be guaranteed: Given $R>1$, there exists $\rho>0$ such that e_Σ embeds the radius ρ disk bundle in N_1 where $|s|<R$. This understood, the constructions in the article Gr $\Rightarrow$ SW from [22] need modifications at large $|s|$ on $\mathbb{R}\times M$. The full details are given in [26; 28]; they account for the length of these papers. Said briefly, the approximate solution at moderate values of $|s|$ is constructed as in the article Gr $\Rightarrow$ SW from [22]. At points where $s\ll -1$, the curve Σ is very near the $s\ll -1$ part of a union of $\mathbb{R}$–invariant cylinders, each of the form $\mathbb{R}\times\gamma$ with $\gamma\subset M$ a Reeb orbit. The constructions in Gr $\Rightarrow$ SW from [22] are applied using these cylinders in lieu of Σ to obtain an approximate solution where $s\ll -1$. Likewise, the constructions in Gr $\Rightarrow$ SW from [22] are applied to the $s\gg 1$ part of another union of $\mathbb{R}$ invariant cylinders to obtain the approximate solutions on this same part of $\mathbb{R}\times M$. The approximate solutions on these three regions in $\mathbb{R}\times M$ are then glued together where the regions overlap so as to obtain an approximate solution on the whole of $\mathbb{R}\times M$. It is a consequence of (5-3) that a gluing of this sort will result in a pair (A^*,ψ^*) that nearly solves (3-18) when r is large.

With the approximate solution in hand, a perturbative construction finds a nearby (A,ψ) that obeys (3-18) on the nose. The latter construction is somewhat more complicated than that in Gr $\Rightarrow$ SW from [22].

Part 2 The assumption in (5-9) greatly simplifies matters. The analog of Theorem 4.3 when (5-9) is not assumed is very much more complicated. The complications are twofold: First, Φ^r now associates to each $\Theta\in\mathcal{Z}^L$ a set, Φ_Θ, of elements in $\mathcal{M}^r$. These elements do not all have the same degree and there will, in general, be instanton solutions to (3-19) with both $s\to-\infty$ limit and $s\to+\infty$ limit in Φ_Θ. With degrees and signs taken into account, these sorts of instantons compute the product of the Hamiltonian Floer cohomology given in (5-11) for the various pairs $(\gamma,m)\in\Theta$.

Meanwhile, if Θ_- and Θ_+ are distinct elements in $\mathcal{Z}^L$, there may be instanton solutions with $s\to-\infty$ and $s\to+\infty$ limits in respectively Φ_{Θ_-} and Φ_{Θ_+}. However, the set of such solutions is not necessarily in 1–1 correspondence with $\mathcal{M}_1(\Theta_-,\Theta_+)$. If (5-9) does not hold, then each instanton in $\mathcal{M}_1(\Theta_-,\Theta_+)$ can determine a number of instanton solutions to (3-19), even when both Φ_{Θ_-} and Φ_{Θ_+} consist of a single element. To say more about this last point, recall that an approximate solution to (3-19) for, say $s \ll -1$, is constructed using as template what is done in Gr $\Rightarrow$ SW from [22] with the pseudoholomorphic curve taken to be a product of cylinders. The template from the article Gr $\Rightarrow$ SW in [22] requires a solution to (5-10) for each such cylinder. In this regard, the $s \to -\infty$ limit of the solution for a cylinder $\mathbb{R} \times \gamma$ with $(\gamma, m) \in \Theta_-$ must be a solution $\mathfrak{c}_\gamma$ to (5-6). However, the $s \to \infty$ limit must be quite different since it has to match up with what is done at moderate values of $|s|$ using the template from Gr $\Rightarrow$ SW from [22] as applied to the given $\Sigma \subset \mathcal{M}_1(\Theta_-, \Theta_+)$. The precise behavior of $\alpha^{-1}(0)$ at large s is determined by the various versions of (2-6) that come from the ends of Σ whose constant s slices converge as $s \to -\infty$ to γ. In particular, the solution to (5-10) must be such that $\sup_{t \in S^1} \{\mathrm{dist}(\alpha^{-1}(0), 0)|_{(s,t)}\}$ diverges as $s \to \infty$. The story is even more complicated if there are two or more ends involved and they have distinct versions of what is denoted as $q_{\mathcal{E}}$ in Section 2.b.

In the case when (5-9) holds, there is but a single relevant solution to (5-10); and the story, though still long, is more or less straightforward. If (5-9) does not hold, then there may be many relevant solutions to (5-9), and then each will determine a distinct instanton solution to (3-19).

Part 3 The upshot of all of this is that when (5-9) does not hold, the proof that the embedded contact homology is isomorphic to Seiberg–Witten Floer cohomology requires much more work, both on the analytic side and on the algebraic side. What follows is meant to give a rough indication of what is involved on the algebraic side: Each element in $\mathcal{Z}^L$ determines some number of generators in $\mathcal{C}^{\mathrm{SW}}$. This is the case even for elements that are not in $\mathcal{Z}^L_{\mathrm{ech}}$ and so are not considered generators of $\mathcal{C}^L_{\mathrm{ech}}$. Note that an element in $\mathcal{Z}^L - \mathcal{Z}^L_{\mathrm{ech}}$ pairs one or more hyperbolic Reeb orbits with an integer greater than 1. In any event, with the differentials taken into account, each $\Theta \in \mathcal{Z}^L$ determines a submodule, $\mathbb{Z}\Phi_\Theta \subset \mathcal{C}^{\mathrm{SW}}$. Let $L_1 < L_2 < \cdots$ denote the ordered set of numbers that can be obtained as $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma$ with $\Theta \in \mathcal{Z}^L$. If r is sufficiently large, then the $\mathfrak{a}^f \geq -\pi L r$ subcomplex of the Seiberg–Witten Floer cochain complex can be filtered as

$$\cdots \subset \bigoplus_{\Theta \in Z^{L_k}} \mathbb{Z}\Phi_\Theta \subset \bigoplus_{\Theta \in Z^{L_{k+1}}} \mathbb{Z}\Phi_\Theta \subset \cdots . \tag{5-12}$$

The E_2 term of the corresponding spectral sequence is isomorphic to the free $\mathbb{Z}$ module generated by the generators of $\mathcal{C}^L_{\mathrm{ech}}$. This follows from the aforementioned fact that the cohomology of $\mathbb{Z}\Phi_\Theta$ is isomorphic to the product of the various (γ, m) versions of (5-11). In particular, this cohomology is isomorphic to either $\mathbb{Z}$ or 0 with $\mathbb{Z}$ arising if and only if Θ is in $\mathcal{Z}^L_{\mathrm{ech}}$ and so gives a generator of $\mathcal{C}^L_{\mathrm{ech}}$. The induced differential on the E_2 term of the spectral sequence corresponding to (5-12) should be identical to the differential on $\mathcal{C}^L_{\mathrm{ech}}$.

Appendix: (δ, L) approximations

This appendix consists of two subappendices A and B. The first gives a proof of Proposition 2.5; the second states some observations about the proof.

A Proof of Proposition 2.5

Let $\mathcal{R}_L$ denote the set of Reeb orbits for a with symplectic action less than L. To set the stage for the constructions that follow, agree to associate a tubular neighborhood map from $S^1 \times D$ as described in Sections 2.a and 2.b for each Reeb orbit in $\mathcal{R}_L$. Since there are but a finite number of such Reeb orbits, no generality is lost by assuming that these tubular neighborhood maps have pairwise disjoint image. When $\gamma \in \mathcal{R}_L$, the associated tubular neighborhood map is denoted by φ_γ. This map is used to identify its image and domain so as to view the 1–form a near γ as the 1–form on $S^1 \times D$ that is given by ℓ_γ times what is written on the right hand side of (2-1). Although not strictly necessary for what follows, it nonetheless proves convenient to choose φ_γ so that the following is true: If γ is elliptic, then the rotation number R is between 0 and 1. If γ is hyperbolic, then its rotation number is either 0 or 1.

Use φ_γ to define the functions (ν, μ) that appear in (2-1). This done, fix a homotopy $\{\tau \to (\nu_\tau, \mu_\tau)\}_{\tau\in[0,1]}$ as described in Lemma 2.3 with $(\nu_0, \mu_0) = (\nu, \mu)$. Choose this homotopy to be independent of τ near 0 and near 1.

What follows describes how the proof proceeds: Let Q denote a very large integer. The plan is to define a sequence $\{(a_k, J_k)\}_{k=0,1,\dots,Q}$ such that $(a_0, J_0) = (a, J)$ and such that for each $k \in \{0, \dots, Q\}$, all but the fourth item in (2-11) are satisfied if $\hat{a} = a_k$ and $\hat{J} = J_k$. For $k \ge 1$, the latter is replaced by

(A-1)

- $$\frac{2\pi}{\ell_\gamma}\varphi^* a_k = (1 - 2\nu_{\tau=k/Q}|z|^2 - \mu_{\tau=k/Q}\bar{z}^2 - \bar{\mu}_{\tau=k/Q}z^2)\,dt + \frac{i}{2}(z\,d\bar{z} - \bar{z}\,dz)$$
- $\varphi^* T^{1,0}(\mathbb{R} \times M)$ is spanned by
$$ds + i\varphi^* a_k \quad\text{and}\quad \frac{\ell_\gamma}{2\pi}(dz - 2i(\nu_{\tau=k/Q}z + \mu_{\tau=k/Q}\bar{z})\,dt).$$

An inductive argument is used to make these constructions. A lower bound for the integer Q is described below.

Three facts play a prime role in the construction of the sequence $\{(a_k, J_k)\}_{k=1,2,\dots}$. Here is the first:

(A-2) The assignment of $(\tau, t) \in [0, 1] \times S^1$ to any given Reeb orbit's version of the pair $(\nu_\tau(t), \mu_\tau(t))$ defines a smooth map to $\mathbb{R} \times \mathbb{C}$. The derivatives of this map to any given fixed order enjoy a uniform bound that is independent of $\gamma \in \mathcal{R}_L$.

To state the second key fact, introduce L_0 to denote the smallest of the lengths of all closed Reeb orbits.

(A-3) For any given positive integer $q \le L_0^{-1} L + 1$, there is a positive lower bound, independent of $\gamma \in \mathcal{R}_L$ and $\tau \in [0, 1]$, to the absolute value of any eigenvalue of the corresponding (ν_τ, μ_τ) version of (2-3) on the space of $2\pi q$–periodic functions.

The statement of the third fact requires a digression to set the notation. To start it, fix generators Θ_- and Θ_+ from $\mathcal{Z}^L_{\mathrm{ech}}$. An element from J's version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ consists of some number of $\mathbb{R}$–invariant cylinders with integer weights, and one non–$\mathbb{R}$–invariant, irreducible submanifold. Let C denote either one of these cylinders, or the non–$\mathbb{R}$–invariant submanifold. Deformations of C that preserve to first order the J invariance of its tangent space can be viewed with the help of a suitable exponential map as sections of the normal bundle $N \to C$ that obey a certain first order, $\mathbb{R}$–linear elliptic equation. The linear operator that defines this equation is denoted by $\mathcal{D}_C$. This is the operator in (2-8) that was briefly described in the paragraph that follows Lemma 2.2. As noted there, it defines a bounded, Fredholm map from $L^2_1(C; N)$ to $L^2(C; N \otimes T^{0,1}C)$. In this guise, its cokernel is trivial, and so it has an inverse that gives a bounded, linear map to the L^2 orthogonal complement of its kernel. This kernel is trivial if $C = \mathbb{R} \times \gamma$. The inverse of $\mathcal{D}_C$ is denoted by $\mathcal{D}_C^{-1}$.

With this notation set, here is the third point:

(A-4) There is a bound, independent of (Θ_-, Θ_+) and C from $\mathcal{M}_1(\Theta_-, \Theta_+)$, to the norm of $\mathcal{D}_C^{-1}$.

To explain, remember that there are but a finite number of pairs (Θ_-, Θ_+) to choose from, and for each, there are but a finite number of components in $\mathcal{M}_1(\Theta_-, \Theta_+)$. Thus, up to the action of $\mathbb{R}$, there are but a finite number of possible choices for C. The lower bound for the integer Q is determined by the bounds in (A-2), (A-3) and (A-4). To this end, let λ_0 denote the bound that is alluded to in (A-3) and let σ_0 denote the bound that is alluded to in (A-4).

To initiate the induction, note that (a_0, J_0) are such that $(\hat{a} = a_0, \hat{J} = J_0)$ satisfy (A-1) and all but the fourth item of (2-11). Suppose that $k \in \{1, 2, \dots, Q-1\}$ and that a pair (a_k, J_k) have been defined so that $(\hat{a} = a_k, \hat{J} = J_k)$ satisfies all but the fourth item of (2-11) and (A-1) if $k \geq 1$. The assertion in (A-4) also holds when $\mathcal{M}_1(\cdot,\cdot)$ is defined using J_k. This is because the set of subvarieties in question is finite.

The induction from k to $k+1$ requires one additional and crucial input: a constant, $\sigma_* \geq 1$, whose definition follows. To start, let $\gamma \in \mathcal{R}_L$ and let $C = \mathbb{R} \times \gamma$. For each $\tau \in [0, 1]$, use L_τ to denote the version of (2-3) that has γ's version of (ν, μ) replaced by (ν_τ, μ_τ). For each integer $q \in \{1, \dots, L_0/L + 1\}$, let $\mathcal{D}_{C,\tau,q}$ denote the operator $\partial/\partial s + L_\tau$ with domain the space of complex valued, L^2_1 functions on $\mathbb{R} \times (\mathbb{R}/2\pi q\mathbb{Z})$ and range the space of complex valued, L^2 functions on $\mathbb{R} \times (\mathbb{R}/2\pi q\mathbb{Z})$. By virtue of (A-3), this operator is invertible. Take σ_* to be a q, τ and $\gamma \in \mathcal{R}_L$ independent upper bound for the norm of this inverse. This constant σ_* is determined by λ_0 and a sup norm bound for all $\gamma \in \mathcal{R}_L$ versions of the pair (ν, μ).

The completion of the induction from k to $k+1$ is presented below in eight parts. Before starting, take note of the convention used here that c_0 denotes a constant that is independent of the relevant variables. Its value is greater than 1 and it can be assumed to increase between subsequent appearances.

Part 1 To construct a candidate for a_{k+1}, remark first that there exists, by assumption, some $\rho_k \ll \delta$ with the following significance: Let $\gamma \in \mathcal{R}_L$; and use φ_γ again to identify a tubular neighborhood of γ with $S^1 \times D$. The 1–form a_k on the $|z| < \rho_k$ part of $S^1 \times D$ is given by the version of (2-1) that has (ν, μ) replaced by $(\nu_{\tau=k/Q}, \mu_{\tau=k/Q})$.

The construction of a_{k+1} also requires a smooth, nonincreasing function, $\chi\colon [0, \infty) \to [0, 1]$ with value 1 on $[0, \frac{5}{16}]$ and value 0 on $[\frac{7}{16}, \infty)$. This function should be fixed once and for all. Given $\rho > 0$ with $\rho \ll \rho_k$, define a function τ_ρ on D to equal

$$\tau_\rho = \frac{k}{Q} + \frac{1}{Q}\chi\left(\frac{1}{\rho}|z|\right). \tag{A-5}$$

So $\tau_\rho = (k+1)/Q$ where $|z| \leq \frac{1}{4}\rho$ and $\tau_\rho = k/Q$ where $|z| > \rho$. Note as well that

$$|d\tau_\rho| \leq c_0 \frac{1}{Q}\rho^{-1} \quad \text{and} \quad |\nabla d\tau_\rho| \leq c_0 \frac{1}{Q}\rho^{-2}. \tag{A-6}$$

With ρ chosen as above, define the 1–form on $S^1 \times D^2$ by the formula

$$\begin{aligned}\frac{2\pi}{\ell_\gamma} a^\rho = (1 - 2\nu_{\tau_\rho}|z|^2 - \mu_{\tau_\rho}\bar{z}^2 - \bar{\mu}_{\tau_\rho} z^2)\, dt \\ + \frac{i}{2}(z\, d\bar{z} - \bar{z}\, dz) + \left(1 - \chi\left(\frac{1}{\rho}|z|\right)\right)(\cdots),\end{aligned} \tag{A-7}$$

where the terms indicated by the three dots on the right are identical to those that appear in (2-1). It is a consequence of (A-6) that what is written in (A-7) defines a contact

1–form when ρ is sufficiently small. To see this, remark that $a^\rho = a_k$ where $|z| > \rho$, and that where $|z| \le \rho$,

$$\text{(A-8)}\qquad \begin{aligned} \frac{2\pi}{\ell_\gamma} a^\rho &= dt + \frac{i}{2}(z\,d\bar z - \bar z\,dz) + \mathfrak{r}_0, \\ \frac{2\pi}{\ell_\gamma} da^\rho &= i\,dz \wedge d\bar z - 2(\nu_{\tau_\rho} z + \mu_{\tau_\rho}\bar z)\,d\bar z \wedge dt - 2(\nu_{\tau_\rho}\bar z + \bar\mu_{\tau_\rho} z)\,dz\wedge dt + \mathfrak{r}_1, \end{aligned}$$

where $|\mathfrak{r}_0| \le c_0|z|^2$ and $|\mathfrak{r}_1| \le c_0(1/Q)|z|$.

Define $a_{k+1,\rho}$ as follows: If $\gamma \in \mathcal{R}_L$, set $a_{k+1,\rho}$ to equal γ's version of a^ρ on the image of φ_γ. Meanwhile, set $a_{k+1,\rho} = a_k$ on the complement of the union of these tubular neighborhoods.

Lemma A.1 *There exists $\kappa > 1$ with the following significance: If $Q > \kappa$ and if ρ is sufficiently small, then $a_{k+1,\rho}$ satisfies the first and second items in (2-11) plus (A-1).*

Proof of Lemma A.1 The proof has four steps.

Step 1 Let v_k denote the Reeb vector field for the contact form a_k and let $v_{k+1,\rho}$ denote the Reeb vector field for $a_{k+1,\rho}$. The latter agrees with v_k except near a Reeb orbit from $\mathcal{R}_L$. Let γ denote such an orbit. As before, use φ_γ to view a neighborhood of γ as a neighborhood of $S^1 \times \{0\}$ in $S^1 \times D$. This done, then $v_{k+1,\rho}$ agrees with v_k except on the part of $S^1 \times D$ where $|z| < \rho$. Meanwhile, it follows from the second equation in (A-8) that

$$\text{(A-9)}\qquad \frac{\ell_\gamma}{2\pi} v_{k+1,\rho} = \frac{\partial}{\partial t} + 2i(\nu_{\tau_\rho} z + \mu_{\tau_\rho}\bar z)\frac{\partial}{\partial z} - 2i(\nu_{\tau_\rho}\bar z + \bar\mu_{\tau_\rho} z)\frac{\partial}{\partial \bar z} + \mathfrak{v},$$

where

$$\text{(A-10)}\qquad |\mathfrak{v}| \le c_0\frac{1}{Q}|z| \quad\text{and}\quad |\nabla\mathfrak{v}| \le c_0.$$

The formula for $v_{k+1,\rho}$ indicates first that γ is also a $v_{k+1,\rho}$ Reeb orbit, and that it is elliptic or hyperbolic as a $v_{k+1,\rho}$ Reeb orbit for all small ρ. It also indicates that γ's rotation number as a $v_{k+1,\rho}$ Reeb orbit is independent of ρ and identical to its rotation number as v_k and v orbit.

Step 2 This step proves that every $v_{k,\rho}$ Reeb orbit with symplectic action less than L lies in a tubular neighborhood of some Reeb orbit from $\mathcal{R}_L$. To this end, suppose that there exists a sequence $\{\rho_j\}_{j=0,1,\dots}$ with limit zero and a corresponding sequence $\{\gamma_j\}$ of Reeb orbits for the $\rho = \rho_j$ version of $v_{k+1,\rho}$, all with symplectic action as defined by the $\rho = \rho_j$ version of $a_{k+1,\rho}$ bounded by L. View these loops as images of maps from S^1 into M. The bounds in (A-10) on $v_{k+1,\rho}$ and its first derivative guarantee the existence of a convergent subsequence in $C^1(S^1; M)$ whose limit map has the following property: Its image is a closed integral curve of v_k with symplectic action less than L. Thus, each large j version of γ_j must lie in the image of φ_γ that is associated to some Reeb orbit $\gamma \in \mathcal{R}_L$.

The following is a direct consequence: There exists ρ_0 such that if ρ is less than ρ_0 and if γ' is a $v_{k+1,\rho}$ Reeb orbit with symplectic action less than L, then γ' lies in the image of the tubular neighborhood map φ_γ that is associated to some Reeb orbit $\gamma \in \mathcal{R}_L$.

Step 3 A virtual repeat of what is said in Step 2 strengthens Step 2's conclusions as follows: Given $\sigma > 0$, there exists ρ_σ such that if $\rho < \rho_\sigma$ and if γ' is a $v_{k+1,\rho}$ Reeb orbit with symplectic action less than L, then γ' lies in the image of the tubular neighborhood map φ_γ that is associated to a Reeb orbit $\gamma \in \mathcal{R}_L$. Moreover, if γ' is such a $v_{k+1,\rho}$ Reeb orbit, and if it lies in $\varphi_\gamma(S^1 \times D)$, then the coordinate z for D obeys $|z| \le \sigma$ on γ'.

Step 4 Let $\gamma \in \mathcal{R}_L$ and let γ' denote a $v_{k+1,\rho}$ Reeb orbit with symplectic action less than L that lies in $\varphi_\gamma(S^1 \times D)$. Suppose for the sake of argument that $\gamma' \ne \gamma$. Let q denote the winding number of γ' in $S^1 \times D$. It follows from (A-9) that γ' can be viewed as a $2\pi q$ periodic map from $\mathbb{R}$ to $S^1 \times D$ by parametrizing it so that the pullback of dt is the Euclidean 1–form on $\mathbb{R}$. This done, use $z'\colon \mathbb{R}/(2\pi q\mathbb{Z}) \to \mathbb{C}$ to denote the function that sends $t \to z(\gamma'(t)) \in \mathbb{C}$.

Let $\hat{\nu}\colon \mathbb{R}/(2\pi q\mathbb{Z}) \to \mathbb{R}$ denote the map whose value at any given point $t \in \mathbb{R}/(2\pi q\mathbb{Z})$ is that of ν_{τ_ρ} at $(t, z'(t))$. Define $\hat{\mu}$ in an similar fashion. It then follows that

$$|\hat{\nu} - \nu_{\tau=k/Q}| + |\hat{\mu} - \mu_{\tau=k/Q}| \le c_0 \frac{1}{Q}. \tag{A-11}$$

Here, c_0 depends solely on the first derivative bounds that are alluded to in (A-2).

The preceding inequality implies that the function z' is a $2\pi q$–periodic solution to an equation that has the schematic form

$$\frac{i}{2}\frac{d}{dt}z' + \nu_{\tau=k/Q}z' + \mu_{\tau=k/Q}\overline{z}' = \mathfrak{r} \quad \text{where } |r| \le c_0\left(\frac{1}{Q}|z'| + |z'|^2\right). \tag{A-12}$$

Step 5 Let $\lambda_0 > 0$ again denote the eigenvalue bound that is alluded to in (A-3). Let c_0 denote the specific constant that appears in (A-12). Then (A-12) requires $z' = 0$ if Q is chosen so that $Q^{-1} \le (100c_0)^{-1}\lambda_0$. With this choice of Q, each Reeb orbit for any $\rho \ll \rho_k$ version of $v_{k+1,\rho}$ with symplectic action less than L must be a loop from $\mathcal{R}_L$. □

Given Lemma A.1, each $\rho \ll \rho_k$ version of $a_{k+1,\rho}$ is a candidate for a_{k+1}.

Part 2 This part of the proof defines an almost complex structure that is compatible with each such small ρ version of $a_{k+1,\rho}$. This almost complex structure is denoted in what follows by $J_{k+1,\rho}$.

To start, set $J_{k+1,\rho}$ equal to J_k on the complement of images of the tubular neighborhood maps for the Reeb orbits in $\mathcal{R}_L$. Now let γ denote such a Reeb orbit, and let φ_γ denote its tubular neighborhood map. As before, use φ_γ to identify $S^1 \times D$ with a neighborhood of γ. Since $J_{k+1,\rho}(\partial/\partial s) = v_{k+1,\rho}$, the only ambiguity concerns the action of $J_{k+1,\rho}$ on the kernel of $a_{k+1,\rho}$. A look at (A-1) and (A-7) indicates that $J_{k+1,\rho}$ can be chosen so as to have the following properties: First, $J_{k+1,\rho} = J_k$ except at points where $|z| \leq \rho$. Second,

$$|J_{k+1,\rho} - J_k| \leq c_0 \frac{1}{Q} |z| \tag{A-13}$$

with $\quad |\nabla(J_{k+1,\rho} - J_k)| \leq c_0 \frac{1}{Q} \quad$ and $\quad |\nabla^2(J_{k+1,\rho} - J_k)| \leq c_0 \frac{1}{Q} \rho^{-1}$.

Third, the φ pullback of the $J_{k+1,\rho}$ version of $T^{1,0}(\mathbb{R} \times M)$ is spanned by $ds + i\varphi^* a_{k+1,\rho}$ and by $(\ell_\gamma/2\pi)(dz - 2i(\nu_{\tau_\rho} z + \mu_{\tau_\rho} \bar{z})\, dt)$ at points where $|z| \leq \frac{1}{4}\rho$. Fix $J_{k+1,\rho}$ with these properties.

Part 3 The next task is to construct a 1–1 map from the set of components of the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ to those of the $J_{k+1,\rho}$ version. To start, take Σ from the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$. Then each $\mathbb{R}$–invariant cylinder from Σ is $J_{k+1,\rho}$ pseudoholomorphic because each has the form $\mathbb{R} \times \gamma$ with $\gamma \in \mathcal{R}_L$. Let $C \subset \Sigma$ denote the component that is not $\mathbb{R}$–invariant. Then C is $J_{k+1,\rho}$ pseudoholomorphic except where it intersects the product of $\mathbb{R}$ with a tubular neighborhood of a Reeb orbit in $\mathcal{R}_L$. To say more on this, let γ denote such a Reeb orbit. Again use φ_γ to identify $S^1 \times D$ with its φ_γ image. Given $R > 1$, there exists $\rho_R \ll \rho_k$ such that if $\rho < \rho_R$, then the intersection of C with the $|z| < \rho$ part of $\mathbb{R} \times (S^1 \times D)$ can occur only in the following two ways:

(A-14)

- Intersection occurs in a disk of radius R^{-1} in C centered around each point from the finite set where C intersects $\mathbb{R} \times \gamma$.
- Intersection can occur where $|s| > R$ on those ends of C that are labeled by a pair from $\Theta_- \cup \Theta_+$ whose Reeb orbit component is γ.

Let $\chi_{C,\rho}$ denote the characteristic function of the support of $|J_{k+1,\rho} - J_k|$ on C.

In any event, (A-13) implies that C is nearly $J_{k+1,\rho}$ pseudoholomorphic in that each of its tangent planes is nearly $J_{k+1,\rho}$ invariant. Moreover, C is nearly $J_{k+1,\rho}$ pseudoholomorphic in an L^2 sense. To elaborate, reintroduce the normal bundle $N \to C$ and let $\pi\colon T(\mathbb{R}\times M)|_C \to N$ denote the orthogonal projection. It then follows from (2-6), (A-13) and (A-14) that

$$\int_C |\pi \circ J_{k+1,\rho}|^2 \le c_0 \frac{1}{Q^2}\rho^2. \tag{A-15}$$

The relatively small L^∞ and L^2 norms of $\pi \circ J_{k+1,\rho}$ suggest a perturbative construction of a 1–1 map from the set of components of the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ to the set of components of the $J_{k+1,\rho}$ version that pairs components so as to satisfy the third item in (2-11). Such a construction is given in the three steps that follow.

Step 1 This step sets up this perturbative construction. To start, fix a component of the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ and take Σ in this component. A partner for Σ in the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ is described next. This partner has the same set of $\mathbb{R}$–invariant cylinders with the same integer weights as does Σ. Let $C \subset \Sigma$ denote the component that is not $\mathbb{R}$–invariant. The analogous component of the partner to Σ is constructed as a deformation of C that comes via a section of C's normal bundle N by composing the section with a suitably chosen exponential map from a disk subbundle in N to $\mathbb{R}\times M$.

To say more, suppose that $N_1 \subset N$ is a constant radius disk subbundle and suppose that $e_C\colon N_1 \to \mathbb{R}\times M$ is an exponential map that embeds each fiber disk as a pseudoholomorphic disk. Such maps are constructed in Section 5d of SW $\Rightarrow$ Gr from [22]. Note that e_C cannot embed the whole of N_1 unless each pair in $\Theta_- \cup \Theta_+$ has its second component equal to 1. In any event, given e_C, let η denote a section of N_1 over Σ which has $|s| \to \infty$ limit equal to zero. Then $e_C \circ \eta$ is $J_{k+1,\rho}$–pseudoholomorphic if and only if it obeys an equation that has the schematic form

$$\mathcal{D}_C\eta + \mathfrak{p}_1\cdot\eta + (\mathfrak{R}_1(\eta) + \mathfrak{p}_2)\cdot\nabla_C\eta + \mathcal{R}_0(\eta) + \mathfrak{p}_0 = 0. \tag{A-16}$$

To elaborate for a moment on the notation, $\mathcal{D}_C$ denotes the operator that appears in (2-8). Meanwhile, $\mathfrak{p}_1$ is a zero–th order, $\mathbb{R}$–linear operator that obeys $|\mathfrak{p}_1| \le c_0(1/Q)$ and with support in the two regions that are listed in (A-14). What is called $\mathfrak{p}_2$ in (A-16) is a homomorphism with support where $J_{k+1,\rho} \neq J_k$. It has norm $|\mathfrak{p}_2| \le c_0(1/Q)\rho$, because it is bounded by $c_0|J_{k+1,\rho} - J_k|$. What is denoted by $\mathfrak{p}_0$ is obtained from $\pi \circ J_{k+1,\rho}$ by restricting the latter to the (0, 1) tangent space of C. Finally, $\mathfrak{R}_1$ denotes a fiber preserving map from N_1 to $\mathrm{Hom}(T^{1,0}C, T^{0,1}C)$ and $\mathfrak{R}_2$ denotes a fiber preserving map from N_1 to $N \otimes T^{0,1}C$. By virtue of (2-6) and (A-13), these two maps obey

(A-17)
- $|\mathfrak{R}_1(b)| \le c_0 \dfrac{1}{Q}|b|$ and $|\mathfrak{R}_0(b)| \le c_0|b|^2$.
- $|\nabla\mathfrak{R}_1| \le c_0 \dfrac{1}{Q}$ and $|\nabla\mathfrak{R}_0| \le c_0 \dfrac{1}{Q}|b|$.

Step 2 Granted what was just said, a contraction mapping argument can be used to find small normed solutions to (A-16) when ρ is small given that the linear term in (A-16) is invertible as a map between suitable Banach spaces, and given that $\mathfrak{p}_0$ has suitably small norm as an element in the range Banach space.

Lemma A.2 *There exists $\kappa > 1$ with the following significance: Suppose that $Q \ge \kappa$, that $k \in \{1, \dots, Q\}$ and that $\{(a_i, J_i)\}_{1\le i\le k}$ has been constructed. If ρ is sufficiently small, then $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$ has bounded inverse as a map from the L^2–orthogonal complement of the kernel of $\mathcal{D}_C$ in $L^2_1(C; N)$ to $L^2(C; N \otimes T^{0,1}C)$. Such is also the case for the operator $\mathcal{D}_C + \sigma(\mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1)$ for each $\sigma \in [0, 1]$.*

Note that the $\sigma \ne 1$ version of the lemma is needed only to compare respective ± 1 weights that are used to define the embedded contact homology differential. This lemma is proved below in Part 7. Assume it for now.

The L^2_1 norm does not dominate the L^∞ norm, and this makes L^2_1 unsuitable as the Banach space for the contraction mapping. However, a slightly stronger norm can be used to define a suitable Banach space. To say more, introduce a norm on the space of compactly supported sections of either N or $N \otimes T_{\mathbb{C}}C$ as follows: Its square assigns to a section, ζ, the number

$$\int_C |\zeta|^2 + \sup_{z\in C} \sup_{x\in(0,1)} x^{-1/100} \int_{\mathrm{dist}(z,\cdot)<x} |\zeta|^2. \tag{A-18}$$

The Banach space for the contraction mapping argument is the completion of the space of compactly supported sections of N using the norm whose square assigns to any given section η the sum of three terms. The first is the square of the L^2_1 norm, and the next two are the respective $\zeta = \eta$ and $\zeta = \nabla\eta$ versions of (A-18). This space is denoted by $\mathcal{B}_1$ and its norm is denoted by $\|\cdot\|_*$. An appeal to Theorem 3.5.2 in [17] finds a constant, c_C, such that $|\cdot| \le c_C\|\cdot\|_*$. Note that this constant c_C depends on the curve C. Let $\mathcal{B}_1^\perp \subset \mathcal{B}_1$ denote the subspace of elements that are L^2–orthogonal to the kernel of $\mathcal{D}_C$.

Use $\mathcal{B}_0$ to denote the completion of the space of compactly supported sections of $N \otimes T^{0,1}C$ using the norm whose square is depicted in (A-18). If $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$

has a bounded inverse mapping $L^2(C; N \otimes T^{0,1}C)$ to $L^2_1(N)$, then an argument using Theorem 5.4.1 of [17] finds that the inverse of $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$ restricts to $\mathcal{B}_0$ so as to define a bounded, linear map from $\mathcal{B}_0$ to $\mathcal{B}_1^\perp$.

With the preceding understood, fix $\sigma > 0$ so that elements of $\mathcal{B}_1$ with $\|\cdot\|_* < 2\sigma$ define sections of the disk bundle N_1. Let $\mathcal{U}_\sigma \subset \mathcal{B}_1^\perp$ denote the ball of radius σ centered on the origin and define the map $\mathcal{T}\colon \mathcal{U}_\sigma \to \mathcal{B}_1^\perp$ by setting

(A-19)
$$\mathcal{T}(\eta) = -(\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1)^{-1}(\mathfrak{R}_1(\eta) \cdot \nabla_C \eta + \mathcal{R}_0(\eta) + \mathfrak{p}_0).$$

Lemma A.3 *There exists $\sigma' \in (0, \sigma)$ such that if ρ is sufficiently small then the following is true: Suppose that C is a component of a submanifold in the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$. Then $\mathcal{T}$ defines a contraction mapping from $\mathcal{U}_{\sigma'}$ to itself. For such ρ, the map $\mathcal{T}$ has a unique fixed point in $\mathcal{U}_{\sigma'}$. Moreover, this fixed point has $\|\cdot\|_*$ norm bounded by $c_C\rho$, and it is a smooth section of N that obeys (A-16). Here, c_C is independent of ρ but depends on C. In the case that $C = \mathbb{R} \times \gamma$, this fixed point is $\eta = 0$.*

Proof of Lemma A.3 It follows from (A-13) and the first line in (A-17) that

(A-20)
$$\|\mathcal{T}(\eta)\|_* \le c_{C1}(\|\eta\|_*^2 + \rho),$$

where c_{C1} is a constant that is independent of ρ but dependent on C. This last bound implies that $\mathcal{T}$ maps the ball in $\mathcal{B}_1^\perp$ of radius $\frac{1}{4}c_{C1}^{-1}$ to itself when $\rho < \frac{1}{8}c_{C1}^{-2}$. Meanwhile, the second line in (A-17) implies that

(A-21)
$$\|\mathcal{T}(\eta) - \mathcal{T}(\eta')\|_* \le c_{C2}(\|\eta\|_* + \|\eta'\|_*)\|\eta - \eta'\|_*,$$

where c_{C2} is a second C dependent but ρ independent constant. This last bound implies that $\mathcal{T}$ maps the ball of radius σ' to itself as a contraction mapping if $\sigma' < \frac{1}{4}(c_{C1} + c_{C2})^{-1}$ and if ρ is sufficiently small. The remaining assertions of the lemma follow using standard elliptic regularity arguments as found in Chapter 5 of [17]. □

Step 3 Let C now denote the non–$\mathbb{R}$–invariant component of Σ. With ρ very small, let $C' \subset \mathbb{R} \times M$ denote the immersed subvariety that is obtained from C using the section η given by Lemma A.3. This subvariety is $J_{k+1,\rho}$ pseudoholomorphic by construction. Introduce Σ' to denote the union of C' and the other $\mathbb{R}$–invariant elements in Σ with their associated integer weights. As a parenthetical remark, note that Proposition 11.4 in [12] implies the following: The union of the subvarieties that comprise Σ' is an embedded subvariety in $\mathbb{R} \times M$. In any event, Σ' defines an element in the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$.

The association of Σ to Σ' is an injective map from the set of components of the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ into the set of components of the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$. This map is denoted in what follows by $\mathcal{F}$.

Part 4 This part verifies that if ρ is sufficiently small, then the components of the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ that lie in the image of $\mathcal{F}$ are smooth points in this version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ in the sense of Lemma 2.2. It also verifies that the correspondence that is defined by $\mathcal{F}$ satisfies the third item in (2-11) if $\hat{J}$ is replaced there by $J_{k+1,\rho}$.

Lemma A.4 *Suppose that ρ is very small. Let Σ' denote a subvariety from a component of the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ that is in the image of the map $\mathcal{F}$. Let C' denote a component of Σ'. The associated deformation operator $\mathcal{D}_{C'}$ has trivial cokernel. Thus, the component of Σ' in the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ is an orbit of the $\mathbb{R}$–action on this space. Moreover, the sign that this component would contribute to the embedded contact homology differential is the same as that of its $\mathcal{F}$–inverse image in the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$.*

Proof of Lemma A.4 The invertibility of $\mathcal{D}_{C'}$ when C' is an $\mathbb{R}$–invariant cylinder from Σ' is automatic since this operator doesn't change when (a_k, J_k) is replaced by $(a_{k+1,\rho}, J_{k+1,\rho})$. Let C' denote the non–$\mathbb{R}$–invariant component of Σ'. Let Σ denote the subvariety from the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ that gives rise to Σ', and let $C \subset \Sigma$ denote the non–$\mathbb{R}$–invariant component that is used to construct C' via Lemma A.3. Since C' is the image via the exponential map of a section of C's normal bundle, it follows that C' is immersed and so has a normal bundle, $N' \to C'$. Note for reference momentarily that Lemma A.2 asserts that $\mathcal{D}_C + \sigma(\mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1)$ is invertible for any constant $\sigma \in [0, 1]$.

Let $\pi\colon N \to C$ again denote the normal bundle to C. Use the exponential map e_C to view C' as the graph in N of the section η. This identifies the normal bundle of C' with the restriction of the bundle $\pi * N$ to this graph. The view of C' as the graph of η also supplies an $\mathbb{R}$–linear isomorphism between $T^{0,1}C$ and $T^{0,1}C'$. These identifications allow $\mathcal{D}_{C'}$ to be viewed as a bounded operator from $L^2_1(C; N)$ to $L^2(C; N \otimes T^{0,1}C)$. As such, it has the form $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1 + \mathfrak{r}$ where $\mathfrak{r}$ has operator norm as a map from $\mathcal{B}_1$ to $\mathcal{B}_0$ that is bounded by $c_C\varepsilon(\rho)$ where $\rho \to \varepsilon(\rho)$ is a decreasing function with limit 0 as $\rho \to 0$. Meanwhile, c_C is independent of ρ but not C. Indeed, the latter fact follows because the derivative terms in $\mathfrak{r}$ have coefficients bounded by $c_C|\eta|$ and the zero–th order terms in $\mathfrak{r}$ have coefficients bounded by $c_C|\nabla\eta|$. The operator norms of these terms can be bounded by $c_C\|\eta\|_*$ using, respectively, Theorem 3.5.2 and Lemma 5.4.1 in [17].

Granted that $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$ is invertible, it follows from what was just said about $\mathfrak{r}$ that any sufficiently small ρ version of $\mathcal{D}_{C'}$ has trivial cokernel when viewed as a map from $L^2_1(C'; N')$ to $L^2(C'; N' \otimes T^{0,1}C')$. The fact that $\mathcal{D}_{C'}$ has trivial cokernel implies that the Σ' is a smooth point of the $J_{k+1,\rho}$ version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ in the sense of Lemma 2.2. Thus, its component is isomorphic to $\mathbb{R}$ with tangent vector field the generator of the $\mathbb{R}$ action that is induced by the constant translations on the $\mathbb{R}$ factor of $\mathbb{R} \times M$.

The small norm of $\mathcal{D}_{C'} - (\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1)$ and the fact that $(\mathcal{D}_C + \sigma(\mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1))$ is invertible for all $\sigma \in [0, 1]$ imply that the sign contribution of the component of Σ' to the embedded contact homology differential is the same as that of the component of Σ. Here is why: These signs are defined using the determinant line bundles for the operators $\mathcal{D}_{C'}$ and $\mathcal{D}_C$. The linear interpolation between $\mathcal{D}_{C'}$ and $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$ provides a canonical isomorphism between the respective determinant lines, as does the linear interpolation between $\mathcal{D}_C + \mathfrak{p}_2 \cdot \nabla_C + \mathfrak{p}_1$ and $\mathcal{D}_C$. Meanwhile, the orientations that are defined by the $\mathbb{R}$ actions on the components of the respective J_k and $J_{k+1,\rho}$ versions of $\mathcal{M}_1(\Theta_-, \Theta_+)$ are compatible as the construction of C' from C is $\mathbb{R}$–equivariant. $\square$

Part 5 Make a very small perturbation $a_{k+1,\rho}$ and $J_{k+1,\rho}$ so that the resulting contact structure is in Lemma 2.1's residual set and so that the complex structure comes from the set $\mathcal{J}_{a_{k+1,\rho}}$. The difference between the perturbed and unperturbed pair should have support in the radius tubular neighborhood of the Reeb orbits in $\mathcal{R}_L$, but its support should be disjoint from these Reeb orbits, and from the pseudoholomorphic curves that appear in $\mathcal{F}$'s image. Note in this regard that the projection to M of the union of the curves that appear in elements from $\mathcal{F}$'s image defines a codimension 1 subvariety in M. The arguments to justify that such perturbations exist are very much like those used in Section 4 of [14] and will not be presented. These differences between the perturbed and unperturbed pair can be as small as desired as measured with respect to any $q \gg 1$ version of the C^q norm, but in any event, this difference should have C^3 norm less than ρ^3. Make a very small perturbation. Agree to use $(a_{k+1,\rho}, J_{k+1,\rho})$ henceforth to denote this slightly perturbed version of the contact form and almost complex structure given in Part 2.

Part 6 This part proves that the small ρ versions of the map $\mathcal{F}$ are onto. A very similar argument proves that $J_{k+1,\rho}$ version of $\mathcal{M}_0(\Theta_-, \Theta_+)$ is described by the third item in (2-11) when ρ is small. The latter argument is not given.

To start the proof that small versions of $\mathcal{F}$ are onto, suppose to the contrary that such is not the case so as to derive some nonsense. Under this assumption, there exists a decreasing sequence $\{\rho_\upsilon\}_{\upsilon=1,2,\ldots}$ with limit zero, and for each index υ, there would exist a point Σ_υ in the J_{k+1,ρ_υ} version of $\mathcal{M}_1(\Theta_-,\Theta_+)$ that does not arise as described from a point in the J_k version of $\mathcal{M}_1(\Theta_-,\Theta_+)$. What are now standard compactness arguments can be used to prove that the sequence $\{\Sigma_\upsilon\}$ has a subsequence that converges to what is often called a *broken trajectory*. Indeed, the limiting behavior is defined by a finite, ordered set $\Lambda = \{\Sigma_1,\ldots,\Sigma_p\}$ where any given Σ_j consists of a finite set of pairs of the form (S,m) with $S \subset \mathbb{R}\times M$ an irreducible, J_k pseudoholomorphic subvariety and m a positive integer. Moreover, there are constraints on the $|s|\to\infty$ limits of the pairs that comprise any given Σ_j. The digression that follows is needed to describe these constraints.

To start the digression, suppose that $(S,m)\in\Sigma_j$. The large $|s|$ slices of any end $E\subset S$ converge as $|s|\to\infty$ as a multiple cover of some Reeb orbit. If γ is such a limit Reeb orbit, define $m_{\gamma,S-}$ to denote the multiplicity of this covering. Set $m_{\gamma,S-}=0$ if γ is not multiply covered by the $s\to-\infty$ limit of the constant s slices of any negative end of S. Now associate to Σ_j the set Θ_{j-} whose elements are pairs of the form (γ,q) where γ is a Reeb orbit that is multiply covered by the $s\to-\infty$ limit of the constant s slices of some end in $\cup_{(S,m)\in\Sigma_j} S$, and where $q=\sum_{(S,m)\in\Sigma_j} m\, m_{\gamma,S-}$. Likewise define Θ_{j+}.

What follows are the constraints on the pairs that comprise the elements from Λ:

(A-22) $\Theta_{1-}=\Theta_-,\quad \Theta_{p+}=\Theta_+ \text{ and } \Theta_{j+}=\Theta_{j+1-} \text{ for each } j\in\{1,\ldots,p-1\}.$

Note that these constraint imply that each $\Theta_{j\pm}\in\mathcal{C}_*^L$. Given Equation (102) in [12] and Propositions 11.4 and Corollary 11.5 in [12], it follows that Λ has just one element and that this element is in the J_k version $\mathcal{M}_1(\Theta_-,\Theta_+)$. The argument here is essentially the same as that used to Theorem 1.8 in [11]. See also the proof of Lemma 7.19 in [13]. Keep in mind that this limit element in the J_k version of $\mathcal{M}_1(\Theta_-,\Theta_+)$ consists of a set of disjoint cylinders with weights, and one non–$\mathbb{R}$–invariant submanifold that is disjoint from the cylinders. Let Σ denote the element in question.

The manner of convergence of the subsequence of $\{\Sigma_\upsilon\}_{\upsilon=1,2,\ldots}$ to Σ is described next. There is a sequence $\{s_\upsilon\}_{\upsilon=1,2,\ldots}\in\mathbb{R}$ such that the translation $s\to s_\upsilon$ of Σ_υ along the $\mathbb{R}$ factor of $\mathbb{R}\times M$ gives a new subsequence, now renamed $\{\Sigma_\upsilon\}_{\upsilon=1,2,\ldots}$, such that

(A-23)

- $\lim_{\upsilon\to\infty}\big(\sup_{z\in(\bigcup_{(C,m)\in\Sigma_\upsilon} C)}\operatorname{dist}(z,\Sigma)+\sup_{z\in(\bigcup_{(C,m)\in\Sigma} C)}\operatorname{dist}(z,\Sigma_\upsilon)\big) = 0.$
- $\lim_{\upsilon\to\infty}\sum_{(C,m)\in\Sigma_\upsilon} m\int_C w = \sum_{(C,m)\in\Sigma} m\int_C w$ for any 2–form w on $\mathbb{R}\times M$ with compact support.

The sort of convergence that is dictated in (A-23) requires the existence of a 1–1 correspondence between the set of components of any sufficiently large υ version of Σ_υ and the set of components of Σ. This correspondence is such that if (S, m) is a component of Σ and $(S_\upsilon, m_\upsilon) \subset \Sigma_\upsilon$, then the convergence in (A-23) holds with (S_υ, m_υ) replacing Σ_υ and (S, m) replacing Σ.

Granted this last point, the contraction mapping theorem argument from Lemma A.3 proves that the partner from any sufficiently large υ version of Σ_υ to an $\mathbb{R}$–invariant cylinder in Σ must coincide with this cylinder and its respective Σ_υ and Σ integer weights must agree. Here is why: Let C denote a pseudoholomorphic subvariety for any given almost complex structure compatible to any given contact 1–form. Suppose that C is not $\mathbb{R}$ invariant. View $\pi(C)$ as a cycle in M and write the boundary of this cycle as $\partial_+ C - \partial_- C$, where $\partial_\pm C$ are the respective positive integer weighted sums of Reeb orbits that arise by taking the $s \to \pm\infty$ limits of the constant s slices of C. Then these two weighted sums cannot be equal because the integral over C of the exterior derivative of the contact form is strictly positive.

Meanwhile, any partner in a sufficiently large υ version Σ_υ to the non–$\mathbb{R}$–invariant component of Σ must be that given via Lemma A.3 and its contraction mapping. These conclusions contradict the assumptions made at the outset as they imply that the component of each large υ version of Σ_υ in the J_{k+1,ρ_υ} version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ is in the image of the 1–1 map from the set of components of the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$.

Part 7 Proof of Lemma A.2 The assertion follows by construction when C' is an $\mathbb{R}$–invariant cylinder. Consider the case when C' is not $\mathbb{R}$–invariant. Let Σ denote the subvariety in the J_k version of $\mathcal{M}_1(\Theta_-, \Theta_+)$ that is paired by $\mathcal{F}$ with Σ' and let $C \subset \Sigma$ denote the non–$\mathbb{R}$–invariant component that gives rise to C' via Lemma A.3's contraction mapping.

As noted previously, the operator norm of $\mathfrak{p}_2 \cdot \nabla_C$ is bounded by $c_{C1}\rho$ where c_{C1} depends on C but not on ρ. This understood, Lemma A.2 follows with a proof that $\mathcal{D}_C + \sigma\mathfrak{p}_1$ also obeys its conclusions if ρ is small.

To prove the latter assertion, note that $\mathfrak{p}_1$ is nonzero only on the domains that are described by (A-14). The contribution to the operator norm of $\mathfrak{p}_1$ from the disks that are described in the first item of (A-14) is bounded by $c_0 R^{-1}$ since $\mathfrak{p}_1$ has a ρ–independent point wise bound and the disks have area R^{-2}. Granted that such is the case, the Sobolev theorems in dimension 2 imply the following: Given $\upsilon \in (0, 2)$, there exists a constant $c_0(\upsilon) \geq 1$ such that the contribution of this part of $\mathfrak{p}_1$ to the norm is less than $c_0(\upsilon) R^{-\upsilon}$.

This understood, let $\mathfrak{p}_{1+}$ denote the part of $\mathfrak{p}_1$ with support far out on the ends of C. Granted what was just said about $\mathfrak{p}_1 - \mathfrak{p}_{1*}$, it is sufficient to prove that all sufficiently small ρ versions of $\mathcal{D}_C + \mathfrak{p}_{1+}$ obey the conclusions of Lemma A.2 when ρ is small.

To prove the latter assertion, fix $R \gg 1$ so that the $|s| \geq R$ portion of C is far out on C's ends. Take ρ small so as to guarantee that $\mathfrak{p}_{1+}$ has support only where $|s| > 100R$. Let u_R denote the function on $\mathbb{R}$ that equals 0 where $|s| < R$, equals $(|s|/R - 1)$ where $|s| \in [R, 2R]$ and equals 1 where $|s| > 2R$. This function is Lipschitz. Now write

$$\|(\mathcal{D}_C + \mathfrak{p}_{1+})\eta\|_2^2 = \|u_R(\mathcal{D}_C + \mathfrak{p}_{1+})\eta\|_2^2 + \|(1-u_R)\mathcal{D}_C\eta\|_2^2. \tag{A-24}$$

Note that $\mathfrak{p}_1$ is absent from the far right term in (A-24) by virtue of the fact that $\mathfrak{p}_1$ is zero whereas $1 - u_R$ is not. Commute the functions u_R and $(1 - u_R)$ past the derivatives to obtain

$$\begin{aligned}\|(\mathcal{D}_C + \mathfrak{p}_{1+})\eta\|_2^2 \geq (1 - c_0 R^{-1})&\|(\mathcal{D}_C + \mathfrak{p}_{1+})(u_R\eta)\|_2^2 \\ &+ \|\mathcal{D}_C((1-u_R)\eta)\|_2^2 - c_0 R^{-1}\|\eta\|_2^2.\end{aligned} \tag{A-25}$$

The next point is that $u_R\eta$ has support far out on the ends of C. This is to say that each component of the support of u_R in C sits where C is represented as a multi-valued graph over either the $s \geq R$ or $s \leq -R$ part of some $\mathbb{R}$–invariant cylinder; this as depicted in (2-6). To see what this implies, suppose that $\mathcal{E} \subset C$ is an end, and let γ denote the associated Reeb orbit. Represent $\mathcal{E}$ as in (2-6) where the eigenfunctions and eigenvalues are those of the operator L_k that is given by replacing γ's version of (ν, μ) in (2-3) with $(\nu_{\tau=k/Q}, \mu_{\tau=k/Q})$. The operator $\mathcal{D}_C$ on $\mathcal{E}$ differs from $\partial/\partial s + L_k$ by terms that are bounded by $c_0 e^{-2\lambda s}$ with λ an eigenvalue of L_k that is respectively positive or negative when $\mathcal{E}$ is positive or negative. Since $|\mathfrak{p}_1| \leq c_0(1/Q)$, this implies that

$$\|(\mathcal{D}_C + \mathfrak{p}_{1+})(u_R\eta)\|_2^2 \geq \left(1 - c_0\left(\frac{1}{Q} + R^{-\lambda}\right)\right)\sigma_*^{-2}\|u_R\eta\|_{2,1}^2. \tag{A-26}$$

Here, $\|\cdot\|_{2,1}$ denotes the L_1^2 norm. Meanwhile, with $\prod_C$ denoting the L^2 projection orthogonal to the kernel of $\mathcal{D}_C$,

$$\|\mathcal{D}_C((1-u_R)\eta)\|_2^2 \geq \sigma_k^{-2}\Big\|\prod_C((1-u_R)\eta)\Big\|_{2,1}^2. \tag{A-27}$$

Here, σ_k is a bound on the inverse of $\mathcal{D}_C$. These last three equations imply that any sufficiently small ρ version of $\mathcal{D}_C + \mathfrak{p}_1$ is invertible as a map from the L^2–orthogonal complement in $L_1^2(C; N)$ of the kernel of $\mathcal{D}_C$ to $L^2(C; N \otimes T^{0,1}C)$. $\square$

Part 8 Given what has been said in the preceding parts, all sufficiently small ρ versions of the pair $(a_{k+1,\rho}, J_{k+1,\rho})$ are such that all but the fourth item of (2-11) are obeyed with $\hat{a} = a_{k+1,\rho}$ and $\hat{J} = J_{k+1,\rho}$. As (A-1) is obeyed, the induction can proceed with a_{k+1} and J_{k+1} set equal to any very small ρ version of $a_{k+1,\rho}$, $J_{k+1,\rho}$.

B Preferred (δ, L) approximations

A (δ, L) approximation to (a, J) of the sort just constructed will be called a preferred (δ, L) approximation. The construction of such an approximation requires the following choices: For each $\gamma \in \mathcal{R}_L$, a suitable tubular neighborhood map φ_γ must be chosen. This being done, a suitable homotopy $\{\tau \to (\nu_\tau, \mu_\tau)\}_{\tau\in[0,1]}$ for the φ_γ version of the functions (ν, μ) must be selected. With the latter in hand, a very large integer, Q, is chosen next. Given Q, a choice must be made, for each $k \in \{1, \dots, Q-1\}$, of a positive, but sufficiently small number, ρ. In addition, a choice must be made of a very small perturbation of $(a_{k+1,\rho}, J_{k+1,\rho})$ as described in Part 5 above.

The following proposition can be used to compare preferred (δ, L) approximations that are defined by different choices.

Proposition B.1 *Suppose that $L_1 > L_0 > 1$ are such that there is no generator $\Theta \in \mathcal{C}_{\mathrm{ech}}$ with $\sum_{(\gamma,m)\in\Theta} m\ell_\gamma$ equal to either L_1 or L_0. Let $(\hat{a}^0, \hat{J}^0)$ and $(\hat{a}^1, \hat{J}^1)$ denote given preferred, (δ, L_0) and (δ, L_1) approximations to (a, J). There exists a smoothly parametrized family $\{x \to (\hat{a}^x, J^x)\}_{x\in[0,1]}$ of contact form and compatible almost complex structure with the following properties:*

- *The end members are the given $(\hat{a}^0, \hat{J}^0)$ and $(\hat{a}^1, \hat{J}^1)$.*
- *If $\rho > 0$ is such that $(\hat{a}^0, \hat{J}^0)$ and $(\hat{a}^1, \hat{J}^1)$ are described by the L–version of the three bullets in Proposition 2.5, then such is case for each $x \in [0, 1]$ version of $(\hat{a}^x, \hat{J}^x)$.*
- *Each $x \in [0, 1]$ member of the family obeys all but the fifth item in (2-11); and the latter is obeyed if x is from a certain residual subset of $[0, 1]$.*

Proof of Proposition B.1 What follows outlines the construction. The construction starts by smoothly modifying the respective constructions for $(\hat{a}^0, \hat{J}^0)$ and $(\hat{a}^1, \hat{J}^1)$ so as to smoothly decrease the various choices for the parameters $\{\rho_k\}$ that appear. These are decreased so that the largest, ρ', is very much smaller than the smaller of those originally chosen. In particular, straightforward modifications to the constructions from the preceding part of this appendix will create two such families; one is parametrized by $x \in [0, \frac{1}{8}]$ and starts at $(\hat{a}^0, \hat{J}^0)$, the other is parametrized by $x \in [\frac{7}{8}, 1]$ and ends at $(\hat{a}^1, \hat{J}^1)$. This is done so that both families obey the conclusions of the second and third bullets of Proposition B.1; and so that the end member of the first and the starting member of the second obey the conclusions of the ρ' and L version of the three bullets in Proposition B.1. If ρ' is sufficiently small, a second, also straightforward set of modifications to the constructions used in the preceding part of the appendix will construct the remaining, $x \in (\frac{1}{8}, \frac{7}{8})$, part of the desired family. Note in this regard that the set of homotopy choices $\{\tau(\nu_\tau, \mu_\tau)\}_{\tau\in[0,1]}$ that are needed for defining $(\hat{a}^0, \hat{J}^0)$ and $(\hat{a}^1, \hat{J}^1)$ is contractible; this because the universal cover of $\mathrm{Sl}(2;\mathbb{R})$ is contractible.

As remarked above, the needed modifications to the constructions in the preceding part of this appendix are straightforward; this being the case, the details are omitted. □

References

[1] **F Bourgeois**, **K Mohnke**, *Coherent orientations in symplectic field theory*, Math. Z. 248 (2004) 123–146 MR2092725

[2] **A Floer**, *Morse theory for fixed points of symplectic diffeomorphisms*, Bull. Amer. Math. Soc. (N.S.) 16 (1987) 279–281 MR876964

[3] **A Floer**, *Morse theory for Lagrangian intersections*, J. Differential Geom. 28 (1988) 513–547 MR965228

[4] **H Hofer**, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, Invent. Math. 114 (1993) 515–563 MR1244912

[5] **H Hofer**, *Dynamics, topology, and holomorphic curves*, from: "Proceedings of the International Congress of Mathematicians, Vol. I (Berlin, 1998)", Extra Vol. I (1998) 255–280 MR1648034

[6] **H Hofer**, *Holomorphic curves and dynamics in dimension three*, from: "Symplectic geometry and topology (Park City, UT, 1997)", (Y Eliashberg, L Traynor, editors), IAS/Park City Math. Ser. 7, Amer. Math. Soc. (1999) 35–101 MR1702942

[7] **H Hofer**, *Holomorphic curves and real three-dimensional dynamics*, Geom. Funct. Anal. (2000) 674–704 MR1826268 GAFA 2000 (Tel Aviv, 1999)

[8] **H Hofer**, **K Wysocki**, **E Zehnder**, *Properties of pseudo-holomorphic curves in symplectisations. II. Embedding controls and algebraic invariants*, Geom. Funct. Anal. 5 (1995) 270–328 MR1334869

[9] **H Hofer**, **K Wysocki**, **E Zehnder**, *Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics*, Ann. Inst. H. Poincaré Anal. Non Linéaire 13 (1996) 337–379 MR1395676

[10] **H Hofer**, **K Wysocki**, **E Zehnder**, *Properties of pseudoholomorphic curves in symplectizations. III. Fredholm theory*, from: "Topics in nonlinear analysis", (J Escher, G Simonett, editors), Progr. Nonlinear Differential Equations Appl. 35, Birkhäuser, Basel (1999) 381–475 MR1725579

[11] **M Hutchings**, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. (JEMS) 4 (2002) 313–361 MR1941088

[12] **M Hutchings**, **M Sullivan**, *Rounding corners of polygons and the embedded contact homology of T^3*, Geom. Topol. 10 (2006) 169–266 MR2207793

[13] **M Hutchings**, **C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders. I*, J. Symplectic Geom. 5 (2007) 43–137 MR2371184

[14] **M Hutchings**, **C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders. II*, J. Symplectic Geom. 7 (2009) 29–133 MR2491716

[15] **A Jaffe**, **C Taubes**, *Vortices and monopoles. Structure of static gauge theories*, Progress in Physics 2, Birkhäuser, Boston (1980) MR614447

[16] **P Kronheimer**, **T Mrowka**, *Monopoles and three-manifolds*, New Math. Monogr. 10, Cambridge Univ. Press (2007) MR2388043

[17] **C B Morrey, Jr**, *Multiple integrals in the calculus of variations*, Grund. der math. Wissenschaften 130, Springer, New York (1966) MR0202511

[18] **D Quillen**, *Determinants of Cauchy–Riemann operators on Riemann surfaces*, Funktsional. Anal. i Prilozhen. 19 (1985) 37–41, 96 MR783704

[19] **R Siefring**, *Relative asymptotic behavior of pseudoholomorphic half-cylinders*, Comm. Pure Appl. Math. 61 (2008) 1631–1684 MR2456182

[20] **C H Taubes**, *Arbitrary N–vortex solutions to the first order Ginzburg–Landau equations*, Comm. Math. Phys. 72 (1980) 277–292 MR573986

[21] **C H Taubes**, *The Seiberg–Witten and Gromov invariants*, Math. Res. Lett. 2 (1995) 221–238 MR1324704

[22] **C H Taubes**, *Seiberg–Witten and Gromov invariants for symplectic 4–manifolds*, (R Wentworth, editor), First Int. Press Lecture Ser. 2, Int. Press, Somerville, MA (2000) MR1798809

[23] **C H Taubes**, *Asymptotic spectral flow for Dirac operators*, Comm. Anal. Geom. 15 (2007) 569–587 MR2379805

[24] **C H Taubes**, *The Seiberg–Witten equations and the Weinstein conjecture*, Geom. Topol. 11 (2007) 2117–2202 MR2350473

[25] **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

[26] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology II*, Geom. Topol. 14 (2010) 2583–2720

[27] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology III*, Geom. Topol. 14 (2010) 2721–2817

[28] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology IV*, Geom. Topol. 14 (2010) 2819–2960

[29] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology V*, Geom. Topol. 14 (2010) 2961–3000

Department of Mathematics, Harvard University
Cambridge, MA 02138, USA

chtaubes@math.harvard.edu

Proposed: Rob Kirby, Tom Mrowka
Seconded: Danny Calegari, Peter Ozsváth

Received: 15 November 2008
Accepted: 21 August 2009