

HF = HM, II

Reeb orbits and holomorphic curves for the ech/Heegaard Floer correspondence

ÇAĞATAY KUTLUHAN
YI-JEN LEE
CLIFFORD HENRY TAUBES

This is the second of five papers that construct an isomorphism between the Seiberg–Witten Floer homology and the Heegaard Floer homology of a given compact, oriented 3–manifold. The isomorphism is given as a composition of three isomorphisms; the first of these relates a version of embedded contact homology on an auxiliary manifold to the Heegaard Floer homology on the original. This paper describes this auxiliary manifold, its geometry and the relationship between the generators of the embedded contact homology chain complex and those of the Heegaard Floer chain complex. The pseudoholomorphic curves that define the differential on the embedded contact homology chain complex are also described here as a first step to relate the differential on the latter complex with that on the Heegaard Floer complex.

53C07, 53C15

1. The geometry of M , the manifold Y and the geometry of Y	2858
2. The closed integral curves of ν	2883
3. Pseudoholomorphic subvarieties and the geometry of $\mathbb{R} \times Y$	2899
4. Ech-HF subvarieties	2913
5. Quantitative concerns	2928
6. Heegaard Floer curves	2960
7. Ech-HF subvarieties and Heegaard Floer curves	2977
Appendix	2990
List of symbols	3003
References	3010

This paper and its sequel [12] supply the geometric and analytic results to relate Peter Ozsváth and Zóltan Szabó's Heegaard Floer homology [17; 18] and a version of Michael Hutchings' embedded contact homology [8]. The precise relationship is reported in [11].

By way of background, the Heegaard Floer homology of a given compact and oriented three-manifold is computed using a suitably chosen Morse function and associated pseudogradient vector field. This manifold is denoted by M . As explained in what follows, this data is used to construct geometric data on the connect sum of M with a certain number of copies of $S^1 \times S^2$. Use Y in what follows to denote this connect sum but with orientation reversed from that on M . The geometric data on Y can be used to define a stable Hamiltonian version of Michael Hutchings' embedded contact homology.

The plan for this paper is to first describe how the data that is needed to compute the Heegaard Floer homology for M is used to construct Y and the stable Hamiltonian structure needed to compute the relevant version of embedded contact homology for Y . With the respective geometric structures in place, the generators for the relevant version of embedded contact homology for Y are determined. As is explained in Proposition 2.8 below, each generator of the Heegaard Floer chain complex on M determines a set of generators for the relevant embedded contact homology complex on Y . With the generators understood, the balance of this article explains how the pseudoholomorphic curves on Y that are used to compute the embedded contact homology differential determine the data that is used on M to compute the Heegaard Floer differential. This is done using Robert Lipshitz's reformulation of Heegaard Floer homology [13].

What follows directly is an overview of the paper's contents.

Section 1 The manifold Y and its stable Hamiltonian structure are constructed from the Heegaard Floer geometry on M . Section 1E summarizes the most relevant aspects of the stable Hamiltonian structure and the associated geometry. Section 1F gives brief overview of what is needed to define Heegaard Floer homology on M and the relevant version of embedded contact homology on Y .

Section 2 The relevant generators of the embedded contact homology chain on Y complex are described. These generators are defined from integral curves of a vector field that is annihilated by the 2-form from the stable Hamiltonian structure. (This is the vector field v .) Proposition 2.7 describes the integral curves of v and Proposition 2.8

describes the corresponding generators of the embedded contact homology chain complex.

Section 3 A constrained set of almost complex structures on $\mathbb{R} \times Y$ are introduced in Section 3A. They are then used to construct pseudoholomorphic foliations of parts of $\mathbb{R} \times Y$ that play a central role in the subsequent sections and in [12]. Propositions 3.1–3.4 and Lemmas 3.5 and 3.6 describe the relevant properties of the subvarieties that make up these foliations.

Section 4 This section introduces the pseudoholomorphic subvarieties that define the differential and various other endomorphisms of the embedded contact homology of Y . These are the ech-HF subvarieties. The various propositions and lemmas in this section describe the properties of these subvarieties and, in particular, how they look in the part of Y from M and in the 1–handle parts of Y .

Section 5 This section states and then proves a crucial a priori bound on the value of the integral of the stable Hamiltonian 2–form over an ech-HF subvariety. The statement of the bound is part of Proposition 5.1. The a priori bounds in Proposition 5.1 play a central role in all of the subsequent analysis in the article and in [12]. For example, the bounds from Proposition 5.1 are used in the appendix to this article and in [12] to study the geometric limits of certain infinite sequences of ech-HF subvarieties.

Section 6 This section constitutes a digression to summarize the aspects of Lipshitz’s formulation of Heegaard Floer homology on M that are relevant to the embedded contact homology on Y . In particular, the section describes the sorts of pseudoholomorphic subvarieties that are used in Lipshitz’s formulation of Heegaard Floer homology on M , and it states and proves various key lemmas about the geometry of these subvarieties.

Section 7 This section defines a correspondence that maps any given ech-HF subvariety to a subvariety of the sort that is used by Lipshitz to describe the Heegaard Floer homology of M . Proposition 7.2 makes an assertion to the effect that each ech-HF subvariety on the $\mathbb{R} \times M$ part of $\mathbb{R} \times Y$ is close to a subvariety that appears in Lipshitz’s reformulation of Heegaard Floer homology. This proposition (and a slightly weaker Proposition 7.1) are used in [12] to compare the differential on the Heegaard Floer complex for M to the differential that is used to define relevant version of embedded contact homology on Y .

Appendix The appendix proves that the relevant version of embedded contact homology on Y can be defined using the rules laid out by Hutchings. Theorem A.1 makes a precise assertion to this effect.

This article supplies the proofs of Theorems 2.1 and 2.2 in [11]; the former is restated as Proposition 2.8 and proved in Section 2 and the latter is restated in the appendix, where it is proved.

The following notational conventions are used: What is denoted by c_0 is in all cases is a constant in $(1, \infty)$ whose value is independent of all relevant parameters. The value of c_0 can increase between subsequent appearances. A second convention concerns a function that is denoted by χ . This is a fixed, nonincreasing function on \mathbb{R} with value 1 on $(-\infty, 0]$ and equal to 0 on $[1, \infty)$. A list of symbols is supplied after the appendix.

Acknowledgements

The authors owe much to Michael Hutchings for sharing his thoughts about embedded contact homology. Another debt is owed to Robert Lipshitz for sharing thoughts about his version of Heegaard Floer homology.

The authors also acknowledge the generous support of the Mathematical Sciences Research Institute, where the three authors did most of the work reported here. Lee and Taubes were also supported in part by the National Science Foundation. Taubes was also supported by a David Eisenbud Fellowship. Lee has also been supported by Hong Kong Research Grants Council grants GRF-401913, 14316516, 14305541.

1 The geometry of M , the manifold Y and the geometry of Y

Let M denote a compact, oriented 3-manifold with a self-indexing Morse function, $f: M \rightarrow [0, 3]$ with one index 0 critical point and one index 3 critical point. Let G denote the number of index 1 critical points; this is the same as the number of index 2 critical points. The function f is used in what follows to construct what is called a *stable Hamiltonian structure* on $Y = M \# (\#_{G+1}(S^1 \times S^2))$. This geometric data on Y consists of a pair (a, w) of 1-form and 2-form with $dw = 0$, da at each point is in $\text{span}(w)$ and $a \wedge w$ is nowhere zero. The orientation for Y is chosen so that $a \wedge w > 0$. Section 1E summarizes the most relevant aspects of the resulting geometry.

The constructions require first fixing a 1-to-1 correspondence between the set of index 2 critical points of f and index 1 critical points. Let Λ denote the resulting ordered set of paired critical points. The connect sum is defined by attaching 1-handles at paired index 2 and index 1 critical points, and then at the index 0 and index 3 critical points with the handle attached by removing small radius balls about the respective members of each pair.

1A Attaching the handles

This part of the section explains how the 1–handles are attached to M so as to obtain Y . There are three parts to the discussion. The first part describes the 1–handle that is labeled by a given pair from Λ . The second part describes the 1–handle that is labeled by the index 0 and 3 critical points. The final part introduces some notational conventions that are used implicitly throughout this article. The definitions that follow involve a chosen constant $\delta_* \in (0, c_0^{-1})$. An upper bound for δ_* is determined by M and f . Any choice below this upper bound will work.

Part 1 Let $p \in M$ denote an index 1 critical point. Take coordinates (x, y, z) near p so that f appears as

$$(1-1) \quad f = 1 + x^2 + y^2 - 2z^2.$$

(The function f is taken here to be a harmonic function with respect to the Euclidean coordinates. The point is to have $*df$ be a closed 2–form.) Fix $\delta_* \in (0, \frac{1}{100})$ such that these coordinates are defined for $(x^2 + y^2 + z^2)^{1/2} \leq 10\delta_*$. By way of notation, if $d \in (0, 10\delta_*)$, the ball in this coordinate chart where $(x^2 + y^2 + z^2)^{1/2} < d$ is said to be the *radius d coordinate ball* centered on p . Implicit in the term “coordinate ball” is the use of the special coordinate system where f appears as above. Introduce the spherical coordinates $(r_+, (\theta_+, \varphi_+))$ of the usual sort: the Euclidean coordinates (x, y, z) are given in terms of the spherical ones by the rule $x = r_+ \sin \theta_+ \cos \varphi_+$, $y = r_+ \sin \theta_+ \sin \varphi_+$ and $z = r_+ \cos \theta_+$. (These are such that the positive z axis points to the north pole on the sphere.) The function f in the spherical coordinates is given by the formula $f = 1 + r_+^2(1 - 3 \cos^2 \theta_+)$. It proves convenient in what follows to define a function $t_+ = \ln r_+$.

Let p' denote an index 2 critical point. Take coordinates (x, y, z) here so that the function f appears as

$$(1-2) \quad f = 2 - (x^2 + y^2 - 2z^2).$$

Introduce again radial coordinates $(r_-, (\theta_-, \varphi_-))$. The function f appears in these coordinates as $f = 2 - r_-^2(1 - 3 \cos^2 \theta_-)$. The constant δ_* can be chosen so that these coordinates are defined where the radial coordinate $r_- \leq 10\delta_*$. By analogy with the index 1 critical point case, given $d \in (0, 10\delta_*)$, the ball in this coordinate chart where $r_- < d$ is said to be the *radius d coordinate ball* centered at p' . Use t_- to denote $\ln r_-$.

Let $\mathfrak{p} = (p, p')$ denote a pair from Λ . The convention here and subsequently has the left entry, p , the index 1 critical point and the right entry, p' , the index 2 critical point. The pair \mathfrak{p} labels one of the 1–handles that is used to obtain Y from M . This handle is denoted by $\mathcal{H}_{\mathfrak{p}}$. The attaching of the 1–handles requires the choice of a parameter; this a positive number R . The number R is constrained so that $R \gg -100 \ln \delta_*$. There are no upper bound constraints. What follows explains how to attach the 1–handle $\mathcal{H}_{\mathfrak{p}}$.

A 1–handle is by definition diffeomorphic to $[-1, 1] \times S^2$ and it is attached by identifying the respective boundaries to constant radii spheres in the radius $8\delta_*$ coordinate balls centered on p and p' . To say more, introduce coordinates $(u, (\theta, \phi))$ for the 1–handle where $u \in [-R + \ln(7\delta_*), R - \ln(7\delta_*)]$ and where (θ, ϕ) are the spherical coordinates on S^2 . The 1–handle $\mathcal{H}_{\mathfrak{p}}$ is defined by making the following identifications:

- (1-3) • Identify the $t_+ \in [-2R - \ln(7\delta_*), \ln(7\delta_*)]$ part of the radius $7\delta_*$ coordinate ball centered on p with $\mathcal{H}_{\mathfrak{p}}$ by writing $(t_+ = u - R, \theta_+ = \theta, \varphi_+ = \phi)$.
- Identify the $t_- \in [-2R - \ln(7\delta_*), \ln(7\delta_*)]$ part of the radius $7\delta_*$ coordinate ball centered on p' with $\mathcal{H}_{\mathfrak{p}}$ by writing $(t_- = -u - R, \theta_- = \pi - \theta, \varphi_- = \phi)$.

Adding the G handles $\{\mathcal{H}_{\mathfrak{p}}\}_{\mathfrak{p} \in \Lambda}$ gives a new manifold that is diffeomorphic to the connect sum of M with G copies of $S^1 \times S^2$.

Part 2 Fix Euclidean coordinates on respective small radius balls about the index 0 and index 3 critical points of f so that f appears as

$$(1-4) \quad f = x^2 + y^2 + z^2 \quad \text{and} \quad f = 3 - x^2 - y^2 - z^2.$$

Take δ_* so that these coordinates are defined for $(x^2 + y^2 + z^2)^{1/2} \leq 10\delta_*$. Use B_0 and B_3 to denote the respective radius $10\delta_*$ coordinate balls about the index 0 and 3 critical points. If $d \in (0, 10\delta_*)$, the subset of either ball where $(x^2 + y^2 + z^2)^{1/2} < d$ is called the radius d coordinate ball about the critical point in question. Introduce the spherical coordinates (r_+, θ_+, ϕ_+) for B_0 and (r_-, θ_-, ϕ_-) for B_3 . Use t_+ and t_- to denote $\ln r_+$ and $\ln r_-$.

The attaching spheres for the relevant 1–handle are the radius $7\delta_*$ coordinate balls centered on these two critical points. The handle is again parametrized using coordinates (u, θ, ϕ) where $u \in [-R - \ln 7\delta_*, R + \ln 7\delta_*]$ and with (θ, ϕ) the spherical angles. The handle attachment here is defined by (1-3) with it understood that p is the index 0 critical point and p' is the index 1 critical point.

This 1–handle in Y is denoted in what follows by \mathcal{H}_0 .

Part 3 This part introduces three conventions. First, the various handles from the set $\{\mathcal{H}_0, \{\mathcal{H}_p\}_{p \in \Lambda}\}$ are identified with their images in Y . To state the second, suppose that $r \in [e^{-R}, 7\delta_*)$. The complement of the union of the $u \in (-R - \ln r, R + \ln r)$ part of each 1-handle defines a submanifold with boundary in Y . The latter is denoted by M_r . The identifications depicted in (1-3) identify M_r with the complement in M of the union of the radius r coordinate balls about each of the critical points of f . This part of M is also denoted by M_r and these two versions of M_r are, for the most part, identified implicitly in all that follows. The third convention involves the function f . The latter is defined on M , and so it is defined on the incarnation of M_r in M . As a consequence, the function f is defined also on the incarnation of M_r in Y . The following is a sketch of M_r for small r (leaving out balls centered on the index 0 and 3 critical points of f). Figure 1 indicates the boundaries of the balls around the index 1 and 2 critical points of f , some trajectories of the gradient of f between these critical points, the surface Σ and the basepoint z_0 in Σ .




Figure 1

1B Geometry of Y : the pair (a, ω) on $\bigcup_{p \in \Lambda} \mathcal{H}_p$

The construction of the desired pair (a, ω) of 1-form and 2-form on $\bigcup_{p \in \Lambda} \mathcal{H}_p$ requires specification of two additional parameters. The first is a positive number; this denoted by δ . This number is constrained only by the upper bound $\delta < c_0^{-1}\delta_*$. The second parameter is also a positive number; this denoted by χ_0 . With δ chosen, the constant χ_0 is constrained only by an upper bound $\chi_0 < \delta^3$. With (δ, χ_0) chosen, the lower bound for the parameter R must be revised upwards to $R > -c_0 \ln \chi_0$. The constant c_0 in both cases depends only on the particular choice for the function χ .

Keep in mind that the goal here is to define a stable Hamiltonian structure. A contact 1–form is, in particular, a stable Hamiltonian structure, and a as constructed below is a contact 1–form on most of any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$. With this understood, fix $\mathfrak{p} \in \Lambda$. The specification of (a, w) on $\mathcal{H}_{\mathfrak{p}}$ has three parts.

Part 1 The definitions require the introduction of three new functions of the coordinate u . All are defined using the chosen, nonincreasing function $\chi: \mathbb{R} \rightarrow [0, 1]$. Recall that the latter equals 1 on $(-\infty, 0]$ and it equals 0 on $[1, \infty)$. The first of the new functions is denoted by χ , and it is given by the rule

$$(1-5) \quad \chi(u) = \chi_0 \chi(|u| - R - \ln \delta + 12).$$

Thus, $\chi = 0$ where $|u| \geq R + \ln \delta - 11$ and $\chi = \chi_0$ where $|u| \leq R + \ln \delta - 12$. The second and third functions are denoted respectively by χ_- and χ_+ . These are defined by the rules

$$(1-6) \quad \chi_-(u) = \chi(u - \frac{1}{4}R) \quad \text{and} \quad \chi_+(u) = \chi(-u - \frac{1}{4}R).$$

Thus, $\chi_+ = 0$ where $u \leq -\frac{1}{4}R - 1$ and $\chi_+ = 1$ where $u \geq -\frac{1}{4}R$. Note that this function is nondecreasing. Meanwhile, $\chi_- = 0$ where $u > \frac{1}{4}R + 1$ and $\chi_- = 1$ where $u \leq \frac{1}{4}R$ and so χ_- is nonincreasing. Use χ'_+ and χ'_- in what follows to denote their respective u -derivatives. Looking ahead, the purpose of χ_+ is to modify terms where $u \ll -1$ that would be $\mathcal{O}(e^{-4R})$ near $u = -R$ so as to make them equal to zero where $u \ll -1$. These are typically terms that are $\mathcal{O}(1)$ near $u = R$. By the same token, χ_- is used to modify terms where $u \gg 1$ that would be $\mathcal{O}(e^{-4R})$ near $u = R$ to make them equal to zero where $u \gg 1$. These terms are typically $\mathcal{O}(1)$ near $u = -R$.

Part 2 Define the 1–form a by the rule

$$(1-7) \quad a = (\chi + (2\chi_+ + \chi'_+) e^{2(u-R)} + (2\chi_- - \chi'_-) e^{-2(u+R)})(1 - 3 \cos^2 \theta) du \\ - \sqrt{6} (\chi + 2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)})) \cos \theta \sin^2 \theta d\phi \\ + 6(\chi_+ e^{2(u-R)} - \chi_- e^{-2(u+R)}) \cos \theta \sin \theta d\theta,$$

and the 2–form w on $\mathcal{H}_{\mathfrak{p}}$ by the rule $w = da$. The admittedly complicated formula for a is motivated by various concerns with the first and second being that it should be a contact form on $\mathcal{H}_{\mathfrak{p}}$, and it should be extendable into the rest of Y as a stable Hamiltonian structure. The extension to Y is described in the subsequent subsections. For the present purposes, the important observations are that w is closed and that

$a \wedge w > 0$ provided that $\delta < c_0^{-1}$, $\chi_0 < \delta^3$ and $R > -c_0 \ln \chi_0$. This claim is proved in the next paragraphs.

To verify the preceding claim about $a \wedge w$, first introduce on \mathcal{H}_p the 1-form

$$(1-8) \quad a_* = (1 - 3 \cos^2 \theta) du - \sqrt{6} \cos \theta \sin^2 \theta d\phi.$$

The reader can verify that this is a contact 1-form. (This 1-form appears in a 4-manifold context; with u having values in $\mathbb{R}/2\pi\mathbb{Z}$, this is the contact form on $S^1 \times S^2$ that is used to describe the pseudoholomorphic geometry near the zero locus of a self-dual, harmonic 2 form on a Riemannian 4-manifold. See for example [21].) Note that the 1-form a_* looks very much like a without the terms proportional to χ_{\pm} and with χ replaced by 1. Having introduced a_* , define a function f_* by the rule

$$(1-9) \quad f_* = (\chi_+ e^{2(u-R)} - \chi_- e^{-2(u+R)})(1 - 3 \cos^2 \theta).$$

Note that this is the function f where $|u| \gg 1$. Use f_* and a_* to write

$$(1-10) \quad a = \chi a_* - 2\sqrt{6}(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)}) \cos \theta \sin^2 \theta d\phi + df_*.$$

As noted above, the 1-form a_* is contact 1-form. This observation with (1-7) imply that $a \wedge w > 0$ where $|u| \geq R + \ln \chi_0 - c_0$. Here is why this is: Note first $\chi = \chi_0$ on this part of \mathcal{H}_p and as a consequence a can be written as $a = \chi_0 a_* + \tau$ and w as $w = \chi_0 da_* + \epsilon$ where τ and ϵ are bounded by $c_0 e^{-2(|u|-R)}$. Thus, $a \wedge w$ can be written as $(\chi_0^2 + \mathfrak{k})a_* \wedge da_*$ with \mathfrak{k} , the contribution from τ and ϵ , obeying $|\mathfrak{k}| \leq c_0 e^{-2(|u|-R)}(\chi_0 + e^{-2(|u|-R)})$. If $|u| \leq R + \ln \chi_0 - c_0$, then this is less than $\frac{1}{100} \chi_0^2$ and so $a \wedge w$ can be written as $(1 + \mathfrak{z})\chi_0^2 a_* \wedge da_*$ with $|\mathfrak{z}| < \frac{1}{100}$.

To continue, introduce $f_- = e^{2(u-R)}(1 - 3 \cos^2 \theta)$. Then

$$(1-11) \quad \begin{aligned} \bullet \quad df_- \wedge d(-2\sqrt{6}e^{2(u-R)} \cos \theta \sin^2 \theta d\phi) \\ \quad \quad \quad = 4\sqrt{6}e^{4(u-R)}(1 + 3 \cos^4 \theta) du \sin \theta d\theta d\phi, \\ \bullet \quad df_- \wedge da_* = 2\sqrt{6}e^{2(u-R)}(1 - 3 \cos^2 \theta)^2 du \sin \theta d\theta d\phi. \end{aligned}$$

The top line in (1-11) implies that $a \wedge w > 0$ where $\chi = 0$ and $u < 0$, and it implies that $a \wedge w > 0$ where $d\chi \neq 0$ and $u < 0$ if $\delta < c_0^{-1}$ and $\chi_0 \leq \delta^3$. The two bullets together with the fact that a_* is a contact form imply that $a \wedge w > 0$ on the $u \leq 0$ part of \mathcal{H}_p . Much the same analysis with f_+ will prove that $a \wedge w > 0$ on the $u \geq 0$ part of \mathcal{H}_p .

1C Geometry of Y : the pair (a, w) on $f^{-1}([1, 2]) \cap M_\delta$

The definition of a and w on M_δ requires the choice of additional data. The first of these is a class $c_{1M} \in H^2(M; \mathbb{Z})/\text{torsion}$ which is divisible by 2. The second is a suitably constrained *pseudogradient* vector field on M for the function f . Having chosen such a vector field, let v denote the normalized version that pairs with df to give 1. The vector field v is defined on the complement of the critical points of f . There are various constraints to the choice for v ; these are described in what follows when they are needed. The construction has nine parts.

Part 1 There are constraints on v near the index 1 and index 2 critical points of f on M . These are described here. Let p denote an index 1 critical point. Reintroduce the coordinates $(t_+, (\theta_+, \varphi_+))$ for the radius $8\delta_*$ coordinate ball centered at p .

Constraint 1 The vector field v on the radius $8\delta_*$ coordinate ball centered at p is

$$v = \frac{1}{2(1 + 3 \cos^4 \theta_+)} e^{-2t_+} \left((1 - 3 \cos^2 \theta_+) \frac{\partial}{\partial t_+} + 2 \cos \theta_+ \sin \theta_+ \frac{\partial}{\partial \theta_+} \right).$$

The analogous formula for v on the radius $8\delta_*$ coordinate ball centered on a given index 2 is obtained from (1-12) by first replacing (t_+, θ_+) with (t_-, θ_-) and then multiplying the resulting expression by -1 .

The fact that the pairing of v with df is 1 can be seen by writing df on this part of M in the $(t_+, (\theta_+, \varphi_+))$ coordinates:

$$(1-12) \quad df = 2e^{2t_+} ((1 - 3 \cos^2 \theta_+) dt_+ + 3 \cos \theta_+ \sin \theta_+ d\theta_+).$$

The analogous equation on the radius $8\delta_*$ coordinate ball centered on an index 2 critical point of f is obtained from (1-12) by replacing t_+ with t_- and multiplying the resulting expression by -1 .

It follows from Constraint 1 that v spans the kernel of the 2-form w on the part of M_δ in the radius $7\delta_*$ coordinate ball centered on any given index 1 or index 2 critical point of f . To see why this is, let p denote an index 1 critical point. The claim follows via a straightforward computation from the fact that the 2-form w appears where $t_+ > \ln \delta$ in the radius $7\delta_*$ coordinate ball as

$$(1-13) \quad w = -2\sqrt{6} d(e^{2t_+} \cos \theta_+ \sin^2 \theta_+ d\varphi_+).$$

The corresponding formula in the case when p is an index 2 critical point verifies that v also generates the kernel of w in the $t_- > \ln \delta$ part of the radius $4\delta_*$ coordinate ball.

Part 2 Introduce $\Sigma = f^{-1}(\frac{3}{2})$ to denote the Heegaard surface. The construction of a and w involves what Ozsváth and Szabó call a Heegaard diagram of Σ . Such a diagram is determined by the choice of \mathfrak{v} . In particular, the construction of Heegaard Floer homology in [17] requires special sorts of Heegaard diagrams, and so constrains \mathfrak{v} . As is evident in what follows, the same sorts of constraints are needed here to construct a and w .

Constraint 2 *The ascending disks from the index 1 critical points of f should have transversal intersection with the descending disks from the index 2 critical points.*

Let p denote a given index 1 critical point of f . The ascending disk from p intersects Σ as an embedded circle, this denoted by C_{p+} . Let p' denote a given index 2 critical point. The descending disk from p' intersects Σ in an embedded circle, this denoted by $C_{p'}$. These two circles intersect transversally if and only if the corresponding ascending disk and descending disk intersect transversally. By way of notation, C_+ is used in what follows to denote the union of the index 1 critical point versions of C_{p+} and C_- is used to denote the union of the index 2 critical point versions of $C_{p'}$.

The Lie transport by \mathfrak{v} identifies the $\theta_+ = \frac{\pi}{2}$ equatorial circle in any given constant $t_+ \geq \ln \delta$ sphere in the radius $8\delta_*$ coordinate ball with C_{p+} , and it identifies the part of such a sphere where $1 - 3 \cos^2 \theta_+ > 0$ with an annular neighborhood of this same circle. Let T_{p+} for the moment denote the annular neighborhood of C_{p+} that corresponds to the annulus in the radius $7\delta_*$ coordinate ball where $1 - 3 \cos^2 \theta_+ > 0$. Choosing $\delta_* < c_0^{-1}$ will guarantee that distinct index 1 critical points define respective versions of T_{p+} with disjoint closures in Σ .

Part 3 Introduce the function $h_+ = 2e^{2t_+} \cos \theta_+ \sin^2 \theta_+$ on the $t_+ > \ln \delta$ part of the radius $8\delta_*$ coordinate ball centered on p . This function has no critical points, and it is such that $\mathfrak{v}(h_+) = 0$. Likewise, $\mathfrak{v}(\varphi_+) = 0$. This understood, Lie transport by \mathfrak{v} of the functions (φ_+, h_+) identifies the latter as coordinates on T_{p+} . Given that (1-13) writes the 2-form w as $\sqrt{6} d\varphi_+ \wedge dh_+$, this same Lie transport identifies w with an area form on T_{p+} .

Meanwhile, when p' is a given index 2 critical point of f , Lie transport by \mathfrak{v} defines a corresponding annular neighborhood, $T_{p'}$ of the circle $C_{p'}$. As with the index 1 critical point case, distinct index 2 critical points can be assumed to define respective versions of $T_{p'}$ with disjoint closures in Σ . Given an index 2 critical point p' , Lie transport by \mathfrak{v} endows $T_{p'}$ with the corresponding coordinates (φ_-, h_-) , and gives the latter the area form $-\sqrt{6} d\varphi_- \wedge dh_-$.

Let p and p' denote a respective pair of index 1 and index 2 critical points of f . If $\delta_* < c_0^{-1}$, then each component of the intersection between T_{p_+} and $T_{p'_-}$ will appear in suitable coordinates as the intersection in \mathbb{R}^2 of $(-1, 1) \times \mathbb{R}$ with $\mathbb{R} \times (1, -1)$. Make such a choice for δ_* .

Constraint 3 A suitable choice for \mathfrak{v} guarantees that $(d\varphi_+, dh_+) = \pm(dh_-, d\varphi_-)$ on the intersection of T_{p_+} with $T_{p'_-}$, with the $+$ sign taken when the pair of vectors $(\frac{\partial}{\partial\varphi_+}, \frac{\partial}{\partial\varphi_-})$ defines an oriented basis for $T\Sigma$ at the corresponding point in $C_{p_+} \cap C_{p'_-}$. This constraint guarantees that the area form $\sqrt{6} d\varphi_+ \wedge dh_+$ for T_{p_+} is equal to the area form $-\sqrt{6} d\varphi_- \wedge dh_-$ for T_{p_-} on $T_{p_+} \cap T_{p'_-}$.

There are no obstructions to choosing \mathfrak{v} so as to satisfy this constraint.

Part 4 Introduce T_+ to denote the union of the various index 1 critical point versions of T_{p_+} and introduce T_- to denote the union of the various index 2 critical point versions of T_{p_-} . Fix a point in the complement of the closure of $T_+ \cup T_-$. This point is denoted by z_0 .

The data consisting of $\{z_0, T_+, T_-\}$ defines what Ozsváth and Szabó call a *pointed Heegaard diagram*. A given class in $H^2(M; \mathbb{Z})/\text{torsion}$ determines what they call a *strongly admissible Heegaard diagram*. Required here is a certain sort of strongly admissible diagram for the chosen class c_{1M} . Diagrams of the desired sort are constructed in Lemma 5.4 of [17].

Some terminology from [17] is required to say more about how the needed Heegaard diagrams are used. A fundamental domain is a component of $\Sigma \setminus (T_+ \cup T_-)$. The fundamental domain with the basepoint is denoted by \mathcal{D}_0 . The set of fundamental domains is ordered and labeled as $\{\mathcal{D}_i\}_{i=0, \dots, N}$. Each fundamental domain is viewed in what follows as a 2-chain on Σ . A periodic domain $\mathcal{P} \subset \Sigma$ is a 2-chain whose boundary is an integer-weighted sum of components of T_+ and T_- with the property that it can be written as the formal sum $\mathcal{P} = \sum_{0 \leq i \leq N} z_i \mathcal{D}_i$, where $z_i \in \mathbb{Z}$ and where the coefficient z_0 of the distinguished domain \mathcal{D}_0 is zero. The \mathbb{Z} -module of periodic domains is isomorphic to $H_2(M; \mathbb{Z})$ with the isomorphism given by adding to the 2-chain \mathcal{P} a weighted union of ascending and descending disks of the respective index 1 and index 2 critical points to create a closed 2-chain in M . The weight of each disk is that of its corresponding component in T_+ and T_- from $\partial\mathcal{P}$.

The lemma that follows brings the notion of a strongly admissible Heegaard diagram into the story.

Lemma 1.1 *There exist strongly admissible Heegaard diagrams for the class c_{1M} such that Σ admits an area form (to be denoted by w_Σ) with the following properties:*

- *The area of Σ is 2.*
- *The integral of w_Σ over any given periodic domain in Σ is equal to the value of c_{1M} on the corresponding homology class in $H_2(M; \mathbb{Z})$.*

Proof When \mathcal{P} is a periodic domain, use $c_{1M}(\mathcal{P})$ to denote the value of the class c_{1M} on the class determined by \mathcal{P} in $H_2(M; \mathbb{Z})$. This is an even integer. Ozsváth and Szabó introduce in Definition 5.3 of [17] the notion of a renormalized, c_{1M} -periodic domain. This is a 2-chain of the form $\mathcal{Q} = \mathcal{P} - \frac{1}{2}c_{1M}(\mathcal{P})[\Sigma]$ with \mathcal{P} being a periodic domain. Such a domain defines a cycle of the form $\mathcal{Q} = \sum_{0 \leq i \leq N} z_i \mathcal{D}_i$ with $\partial \mathcal{Q} = \partial \mathcal{P}$. Their proof of Lemma 5.4 in [17] constructs a c_{1M} -strongly admissible, pointed Heegaard diagram for which each renormalized, c_{1M} -periodic domain has both positive and negative coefficients. They assert at the very start of their proof that this is what the proof will obtain.

Granted such a Heegaard diagram, the argument used for the proof of Lemma 4.12 in [17] can be repeated to find an area 2-form on Σ such that each renormalized, c_{1M} -periodic domain has total signed area equal to zero. With \mathcal{Q} written as above, this means that $\sum_{0 \leq i \leq N} z_i A_i = 0$, where A_i here denotes the area of the corresponding fundamental domain \mathcal{D}_i . Here is the argument: The vector space over \mathbb{R} spanned by the renormalized c_{1M} -periodic domains generates a vector subspace of \mathbb{R}^{N+1} with dimension equal to the dimension of $H_2(M; \mathbb{R})$. This vector subspace intersects the positive quadrant only at the origin. As a consequence, there is a vector in the interior of the positive quadrant that is orthogonal to all vectors in this span. This vector can be written as $\sum_i A_i \mathcal{D}_i$ with $A_i > 0$. The desired area form can be any area form for Σ such that the area of any given fundamental domain \mathcal{D}_i is the corresponding coefficient A_i .

Normalize this area form so that $\sum_i A_i = 2$. This implies that the area of Σ is 2. Let \mathcal{P} denote a periodic domain. As the signed area of $\mathcal{Q} = \mathcal{P} - \frac{1}{2}c_{1M}(\mathcal{P})[\Sigma]$ is zero, it follows that the signed area of \mathcal{P} is $c_{1M}(\mathcal{P})$. This being the case, the area form obeys the conditions asserted by the lemma. □

Part 5 With Lemma 1.1 in hand, here is the next constraint on \mathfrak{v} :

Constraint 4 *The pseudogradient vector field \mathfrak{v} gives a Heegaard diagram of the sort described in Lemma 1.1.*

Let w_Σ denote an area form for Σ of the sort described by Lemma 1.1. There is no obstruction to choosing w_Σ so that it agrees on each index 1 critical point version of T_{p_+} with the area form $\sqrt{6} d\varphi_+ \wedge dh_+$ if $\delta_* < c_0^{-1}$. It follows from what is said in Part 3 that there is no obstruction to also requiring that $w_\Sigma = -\sqrt{6} d\varphi_- \wedge dh_-$ on each index 2 critical point version of T_{p_-} . Fix such a form w_Σ . Note the following important point: the form w_Σ can be assumed independent of the data δ , χ_0 and R that are used to define the geometry of the various 1–handles.

Use the integral curves of \mathfrak{v} to define a diffeomorphism between $(1, 2) \times \Sigma$ and the domain $f^{-1}(1, 2) \subset M$. Let t denote the Euclidean coordinate on the interval $[1, 2]$. This diffeomorphism identifies \mathfrak{v} with the vector field $\frac{\partial}{\partial t}$ and f with the coordinate t . Pullback via the projection from $[1, 2] \times \Sigma$ to Σ defines w_Σ as a 2–form on the whole of $[1, 2] \times \Sigma$ and thus on the $f^{-1}([1, 2])$ part of M_δ . It follows from (1-13) that this form extends w from $\bigcup_{p \in \Lambda} \mathcal{H}_p$ to the latter’s union with the $f^{-1}([1, 2])$ part of M_δ . This extension is denoted by w also.

Part 6 With w extended to $f^{-1}([1, 2]) \cap M_\delta$, turn now to the task of extending the 1–form a . The next lemma is needed for this. The lemma uses b_1 to denote the dimension of $H^1(M; \mathbb{R})$, the first Betti number of M .

Lemma 1.2 *There exists a set of b_1 points in $\Sigma \setminus (T_+ \cup T_-)$ and a 1–form, a_Σ , that is defined on the complement of z_0 and this set, and which has the following properties:*

- $da_\Sigma = w_\Sigma$.
- *The integral of a_Σ is zero on each component of C_+ and on each component of C_- .*

Proof Let \mathcal{I} denote the free \mathbb{Z} –module of rank $2G$ that is generated by the components of C_+ and C_- . The Mayer–Vietoris exact sequence for the homology of M leads via Lefschetz duality to an exact sequence

$$(1-14) \quad 0 \rightarrow H^1(M; \mathbb{R}) \rightarrow H^1(\Sigma; \mathbb{R}) \rightarrow \text{Hom}(\mathcal{I}; \mathbb{R}) \xrightarrow{F} H^2(M; \mathbb{R}) \rightarrow 0.$$

Use the fact that $H^2_{\text{de Rham}}(\Sigma \setminus z_0) = 0$ to choose a smooth 1–form on $\Sigma \setminus z_0$ whose exterior derivative is w_Σ . Let a_0 denote the latter. Integration of a_0 over the various components of C_+ and C_- define an element $\underline{a}_0 \in \text{Hom}(\mathcal{I}; \mathbb{R})$. Let F denote the homomorphism in from $\text{Hom}(\mathcal{I}; \mathbb{R})$ to $H^2(M; \mathbb{R})$. If $F(\underline{a}_0) = 0$, then it follows from (1-14) that a_0 can be modified by adding a suitable closed 1–form on Σ so that its integral over the various components of C_+ and C_- is zero.

For each $i \in \{1, \dots, N\}$, fix a point in the domain \mathcal{D}_i that sits in the complement of the latter's intersection with the closure of $T_+ \cup T_-$. Let z_i denote this point. Any given point $R = (R_1, \dots, R_N) \in \mathbb{R}^N$ determines a homomorphism from $H_2(M; \mathbb{R})$ to \mathbb{R} as follows: Represent the given class as a periodic domain, \mathcal{P} . Write \mathcal{P} as $\sum_{1 \leq i \leq N} Z_i \mathcal{D}_i$. The homomorphism in question sends \mathcal{P} to $\sum_{1 \leq i \leq N} R_i Z_i$. This then defines a surjective linear map $\psi: \mathbb{R}^N \rightarrow H^2(M; \mathbb{R})$.

Fix b_1 points from $\{z_i\}_{1 \leq i \leq N}$ so that the restriction of ψ to the corresponding b_1 -dimensional subspace in \mathbb{R}^N is an isomorphism. Let \mathbb{V} denote this subspace. Meanwhile, the given a_0 can be modified on $\Sigma \setminus \{z_i\}_{0=1, \dots, N}$ so that its integral over the boundary of each $i \in \{1, \dots, N\}$ version of \mathcal{D}_i is R_i . This modification changes $F(a_0)$ to $F(a_0) + \psi(R)$. This understood, and given what was said in the previous paragraph, there exists $R \in \mathbb{V}$ such that the resulting version of a_0 has integral zero on each component of C_+ and on each component of C_- . □

Part 7 Let $\mathbb{Y} \subset \Sigma$ denote the set composed of z_0 and the b_1 points given by Lemma 1.2. The resulting 1-form a_Σ and the constant δ_* can be chosen so that a_Σ has a more-or-less canonical form near each point in \mathbb{Y} . In particular, a_Σ and δ can be chosen so that there is a set $\{D_z\}_{z \in \mathbb{Y}}$ of disjoint disks in $\Sigma \setminus (T_+ \cup T_-)$, each with area $\frac{\pi}{100(1+b_1)} \delta_*^2$, and with the following additional properties: Fix $z \in \mathbb{Y}$. Then $z \in D_z$ and there exists a real number R_z and radial coordinates (ρ, φ) for D_z that make w_Σ and a_Σ appear as

$$(1-15) \quad w_\Sigma = \rho \, d\rho \, d\varphi \quad \text{and} \quad a_\Sigma = \left(\frac{1}{\pi} R_z + \frac{1}{2} \rho^2\right) d\varphi.$$

To see that this can be arranged, note first that coordinates (ρ, φ) can be found on a disk D'_z centered at z to make $w_\Sigma = \rho \, d\rho \, d\varphi$ because w_Σ is an area form. To see about a_Σ , let $a_{\Sigma 0}$ denote an initial choice for this 1-form. Then $d(a_\Sigma - \frac{1}{2} \rho^2 d\varphi) = 0$ and so $a_\Sigma - \frac{1}{2} \rho^2 d\varphi$ is a closed 1-form on $D'_z \setminus z$. Since $d\varphi$ generates the de Rham dimension 1 cohomology of $D'_z \setminus z$, it follows that a_Σ can be written as $(\frac{1}{\pi} R_z + \frac{1}{2} \rho^2 d\varphi) + d\mathfrak{f}$ with \mathfrak{f} being a function on $D'_z \setminus z$. Let D_z denote the disk centered at z with half the radius of D'_z and suppose that β is a nonnegative function that is equal to 1 near $\partial D'_z$ and equal to 0 on the D_z . The 1-form $a_\Sigma = (\frac{1}{\pi} R_z + \frac{1}{2} \rho^2 d\varphi) + d(\beta \mathfrak{f})$ obeys $da_\Sigma = w_\Sigma$ and it is the form depicted in (1-15) on D_z .

Granted the preceding, define a smooth 1-form, \hat{a}_Σ , on Σ as follows: Set \hat{a}_Σ to equal a_Σ on $\Sigma \setminus (\bigcup_{z \in \mathbb{Y}} D_z)$. Fix $z \in \mathbb{Y}$. To define \hat{a}_Σ on D_z , first define the function $\chi_\#$

of the coordinate ρ by the rule

$$\chi_{\#}(\rho) = \chi\left(1 - \frac{1000(1 + b_1)}{\delta_{*}}\rho\right).$$

This function is zero near z , but 1 on the complement of a compact set in D_z . Set

$$(1-16) \quad \widehat{a}_{\Sigma} = \chi_{\#}\left(\frac{1}{\pi}R_z + \frac{1}{2}\rho^2\right) d\varphi.$$

Note for future reference that $d\widehat{a}_{\Sigma} = Hw_{\Sigma}$, where H is a smooth function that equals 1 on the complement of $\bigcup_{z \in \mathfrak{Y}} D_z$.

Part 8 Let p denote an index 1 or index 2 critical point of f . When $r \in (0, 7)$, use respectively $T_{p+}^r \subset T_{p+}$ or $T_{p-}^r \subset T_{p-}$ to denote the subannulus that comes from the radius $r\delta_{*}$ coordinate ball centered on p .

Suppose that p is an index 1 critical point of f . As can be seen from (1-12), the forms a and df restrict to the $r_{+} > \delta$ part of the radius $7\delta_{*}$ coordinate ball centered on p as a 1-form that is invariant with respect to the Lie transport defined by \mathfrak{v} . Lie transport by \mathfrak{v} of $a - df$ from this part of p 's radius $7\delta_{*}$ coordinate ball defines a 1-form, a_{p+} , on T_{p+}^7 that obeys $da_{p+} = w_{\Sigma}$. If p is an index 2 critical point of f , then the analogous construction defines from $a - df$ a 1-form, a_{p-} , on T_{p-}^7 . In general, $a_{p+} \neq a_{p-}$ where their domains of definition intersect, and, in general, neither is equal to a_{Σ} . This part of the construction addresses these issues.

Let p and p' now denote respective index 1 and index 2 critical points of f and suppose that $T \subset T_{p+} \cap T_{p'-}$ is a given component. As $T_{p+} \cap T_{p'-}$ is contractible and as both a_{p+} and $a_{p'-}$ are antiderivatives of w_{Σ} on T , their restriction to $T_{p+}^7 \cap T_{p'-}^7$ can be written as $a_{p+} - a_{p'-} = d\mathfrak{k}$ where \mathfrak{k} is a smooth function on $T_{p+}^7 \cap T_{p'-}^7$. Fix a compactly supported, nonnegative function, q , on $T_{p+}^7 \cap T_{p'-}^7$ that is equal to 1 on $T_{p+}^6 \cap T_{p'-}^6$. Define a 1-form A_{p+} on T_{p+}^6 by the rule $A_{p+} = a_{p+} + \frac{1}{2}d(q \cdot \mathfrak{k})$. Note that $dA_{p+} = w_{\Sigma}$. Meanwhile, define a 1-form $A_{p'-}$ on $T_{p'-}^6$ by the rule $A_{p'-} = a_{p'-} - \frac{1}{2}d(q \cdot \mathfrak{k})$. The forms A_{p+} and $A_{p'-}$ agree on $T_{p+}^6 \cap T_{p'-}^6$.

Make these modifications to all intersections between T_{+} and T_{-} . Given $r \in (0, 7)$, let $T_{+}^r \subset T_{+}$ and $T_{-}^r \subset T_{-}$ denote the subdomains given by the union of relevant versions of $T_{p\pm}^r$. The modified 1-forms define a smooth form, A , on $T_{+}^6 \cup T_{-}^6$ that obeys $dA = w_{\Sigma}$.

The 1-form A is a smooth 1-form defined on $T_{+}^6 \cup T_{-}^6$ with $dA = w_{\Sigma}$. Meanwhile, \widehat{a}_{Σ} on $T_{+}^6 \cup T_{-}^6$ also obeys $d\widehat{a}_{\Sigma} = w_{\Sigma}$. It follows from (1-7) that the 1-form A

has integral zero over each component of C_+ and over each component of C_- . By construction, this is also true for \hat{a}_Σ . As a consequence, $A - \hat{a}_\Sigma$ is exact and so $A - \hat{a}_\Sigma = d\eta$ with η a smooth function on $T_+^6 \cup T_-^6$.

Fix a smooth function, q_Σ , with compact support on $T_+^6 \cup T_-^6$ that is equal to 1 on $T_+^5 \cup T_-^5$. A smooth 1-form, \hat{a} , is defined on Σ by the rule

$$(1-17) \quad \hat{a} = \hat{a}_\Sigma + d(q_\Sigma \eta).$$

The 1-form \hat{a} obeys $d\hat{a} = H\omega_\Sigma$ on Σ , where H is smooth and equals 1 on $\Sigma \setminus (\bigcup_{z \in \mathbb{Y}} D_z)$.

Part 9 This final part of the construction requires the introduction of yet more “bump” functions, plus a constant $r \geq 0$ whose lower bound depends only on the chosen bump functions. There is no upper bound to the choice for r .

The first of the bump functions, q_* , is a smooth, compactly supported function on $T_+^5 \cap T_-^5$ that is equal to 1 on $T_+^4 \cap T_-^4$. The second, Q , is a smooth function on $[1, \frac{3}{2}]$ that equals 1 where $t > \frac{5}{4}$ and equals 0 where $t < \frac{9}{8}$. The third, K , is a smooth, nonnegative function on $[1, \frac{3}{2}]$, this equal to 1 where $t \geq \frac{9}{8}$ and equal to 0 where $t < \frac{17}{16}$. The key feature of K are that it equals 1 where $\frac{d}{dt}Q \neq 0$ and that it vanishes where $t < 1 + 8\delta_*^2$.

With such choices in hand, the extension of a to the $f \in [1, \frac{3}{2}]$ part of M_δ is defined by the 1-form

$$(1-18) \quad a = (1 + rK) dt + \hat{a} + \frac{1}{2}d(Qq_*\ell).$$

The extension of a to the $f \in [\frac{3}{2}, 2]$ part of M_δ is defined by the version of (1-18) that is obtained by replacing the function $t \rightarrow Q(t)$ by the function $t \rightarrow -Q(3 - t)$ and the function K by the function $t \rightarrow K(3 - t)$.

The extension just defined is such that $a \wedge w > 0$ if $r > c_0$. In addition

$$(1-19) \quad \begin{aligned} &\bullet \quad da = w \text{ on the complement of } [1, 2] \times (\bigcup_{z \in \mathbb{Y}} D_z); \\ &\bullet \quad da = Hw \text{ on } [1, 2] \times (\bigcup_{z \in \mathbb{Y}} D_z). \end{aligned}$$

Thus, the pair (a, w) defines a stable Hamiltonian structure on the union of the $f \in [1, 2]$ part of M_δ with $\bigcup_{p \in \Lambda} \mathcal{H}_p$.

1D Geometry of Y : the pair (a, w) on the $f \in [0, 1] \cup [2, 3]$ part of M_δ and \mathcal{H}_0

This section finishes the construction of the pair (a, w) by extending what was constructed in Sections 1B and 1C to the remainder of Y . Various constraints on \mathfrak{v} must be imposed here. All of these can be made where $f \in [0, 1] \cup [2, 3]$ and thus they can be imposed without compromising those listed previously. The construction has four parts.

Part 1 This part of the subsection extends a and w to a neighborhood in Y of the union the $f \in [1, 2]$ part of M_δ with $\bigcup_{p \in \Lambda} \mathcal{H}_p$. To do this, let $\Sigma_1 \subset M_\delta$ denote $f^{-1}(1) \cap M_\delta$. This is an embedded surface with boundary. It is diffeomorphic to the complement in S^2 of G disjoint disks. Let p denote an index 1 critical point of f . As can be seen from Constraint 1 in Section 1C, the vector field \mathfrak{v} is along the part of the boundary of Σ_1 near p is tangent to the boundary sphere of the radius δ coordinate ball centered at p . This understood, Lie transport of Σ_1 by \mathfrak{v} defines an embedding of $[1 - 2\delta_*^2, 1 + \delta_*^2] \times \Sigma_1$ into M_δ with three properties: First, the image of the $t = 0$ slice is the identity map on Σ_1 . Second, the image of the embedding intersects the boundary of M_δ only on the boundary of Σ_1 . Third, the Euclidean coordinate t on the $[1 - 2\delta_*^2, 1 + \delta_*^2]$ factor pulls back as f , and the corresponding coordinate vector field $\frac{\partial}{\partial t}$ pushes forward as \mathfrak{v} .

Lie transport by \mathfrak{v} of the image in M_δ of the $t = 1 + \delta_*^2$ boundary identifies Σ_1 with the complement in Σ of the union of the annular neighborhoods of the index 1 critical points of f where the coordinate h_+ obeys $|h_+| \leq \frac{3}{3\sqrt{3}}\delta_*^2$. Use this identification to view w_Σ as an area form on Σ_1 and thus a closed 2-form w_Σ on $[1 - 2\delta_*^2, 1 + \delta_*^2] \times \Sigma_1$. The embedding into M_δ identifies the latter on the $f \geq 1$ part of the image with w . This understood, the embedding extends w to the image of the whole embedding. This extension agrees on the intersection with any given $p \in \Lambda$ version of \mathcal{H}_p with the latter's version of w and so extends w to the union $\bigcup_{p \in \Lambda} \mathcal{H}_p$ with the $f \in [1 - 2\delta_*^2, 2]$ part of M_δ .

Meanwhile, the 1-form a depicted in (1-18) is invariant with respect to Lie transport by \mathfrak{v} on the $[1, 1 + \delta_*^2] \times \Sigma_1$ part of M_δ , and so Lie transport by \mathfrak{v} defines a smooth extension of (1-18) to the $[1 - 2\delta_*^2, 1 + \delta_*^2] \times \Sigma_1$ part of M_δ . It follows from (1-7) that this extension agrees with the 1-form a on the $r_+ > \delta$ part of the radius δ_* coordinate ball centered at each index 1 critical point of f . This extension of a to the union

of $\bigcup_{p \in \Lambda} \mathcal{H}_p$ with the $f \in [1 - 2\delta_*^2, 2]$ part of \mathcal{M}_δ is such that a and w define a stable Hamiltonian structure.

The analogous construction defined where $f \in [2 - \delta_*^2, 2 + 2\delta_*^2]$ extends a and w to the union of $\bigcup_{p \in \Lambda} \mathcal{H}_p$ with the $f \in [1 - 2\delta_*^2, 2 + 2\delta_*^2]$ part of M_δ as a stable Hamiltonian structure with \mathfrak{v} generating the kernel of w in M_δ .

Part 2 The extension of a and w to the remainder of M_δ and thus to the whole of Y requires more constraints on \mathfrak{v} . These constraints can be met by modifying any given pseudogradient for f in the radius $10\delta_*$ coordinate balls about the index 0 and index 3 critical points of f , and on the $f \in [100\delta_*^2, 1 - 3\delta_*^2]$ and $f \in [2 - 3\delta_*^2, 3 - 100\delta_*^2]$ parts of M_δ . What follows states the constraint on \mathfrak{v} in the coordinate balls.

Constraint 5 *The pseudogradient \mathfrak{v} appears in the radius $8\delta_*$ coordinate ball centered on the index 0 critical point of f as the vector field*

$$\mathfrak{v} = \frac{1}{2}e^{-2t_+} \frac{\partial}{\partial t_+}.$$

The constraint on \mathfrak{v} in the radius $8\delta_$ coordinate ball centered on the index 3 critical point of f is obtained from what is written above by first replacing t_+ by t_- on the right-hand side and then reversing the sign of the resulting expression.*

To set the background for the next constraint, keep in mind that the $f = 1 - \delta_*^2$ slice of M_δ is an embedded 2–sphere. The integral of w over this slice is 2. This is also the area given by the 2–form $\frac{1}{2\pi} \sin \theta_+ d\theta_+ d\varphi_+$ to the boundary of any coordinate ball of radius less than $8\delta_*$ centered on the index 0 critical point of f .

The gradient flow of \mathfrak{v} identifies the set of points $\mathbb{Y} \in \Sigma$ with a set of points in this slice, and it identifies the set of disks $\{D_z\}_{z \in \mathbb{Y}}$ with a corresponding set of disks in this slice. This understood, view \mathbb{Y} and the set $\{D_z\}_{z \in \mathbb{Y}}$ as subsets of the $f = 1 - \delta_*^2$ sphere in M_δ . Keep in mind that the integral of w over each such disk is $\frac{\pi}{100(1+b_1)}\delta_*^2$. Choose a set of $1 + b_1$ points in the boundary 2–sphere of the radius $7\delta_*$ coordinate ball about the index 0 critical point. These points should be equally spaced along the equator. Let \mathbb{Y}_0 denote this set. Fix a round disk about each point in \mathbb{Y}_0 with area $\frac{\pi}{100(1+b_1)}\delta_*^2$. Given $z \in \mathbb{Y}_0$, use D_{0z} to denote the corresponding disk. Fix a corresponding set of $1 + b_1$ points in the boundary 2–sphere of the radius $7\delta_*$ coordinate ball about the index 3 critical point of f . Denote the latter set by \mathbb{Y}_3 . Fix the analogous round disk about each point in \mathbb{Y}_3 , and use D_{3z} to denote this disk when $z \in \mathbb{Y}_3$.

Note next that the $f = 1 - \delta_*^2$ slice of M_δ is identified via the integral curves of any given pseudogradient vector field for f with the boundary of any radius δ or greater coordinate ball centered on the index 0 critical point of f .

With the last three paragraphs as background, use arguments from Moser's proof that area forms are determined up to diffeomorphism by their total area to impose the following constraint:

Constraint 6 *The pseudogradient \mathfrak{v} can be chosen subject to Constraint 5 to have the following properties:*

- *On the set where $f \leq 1 - 3\delta_*^2$, the resulting identification of the $f = 1 - \delta_*^2$ slice in M_δ with the boundary of any radius $r \in [\delta, 8\delta_*]$ coordinate ball centered on the index 0 critical point of f is such that:*
 - (1) *The identification makes ω equal to $\sin \theta_+ d\theta_+ d\varphi_+$.*
 - (2) *The identification sends \mathbb{Y} to \mathbb{Y}_0 and it identifies the sets $\{D_z\}_{z \in \mathbb{Y}}$ and $\{D_{0z}\}_{z \in \mathbb{Y}_0}$.*
- *On the set where $f \geq 2 + 3\delta_*^2$, the resulting identification of the $f = 2 + \delta_*^2$ slice in M_δ with the boundary of any radius $r \in [\delta, 8\delta_*]$ coordinate ball centered on the index 3 critical point of f is such that:*
 - (1) *The identification makes ω equal to $\sin \theta_- d\theta_- d\varphi_-$.*
 - (2) *The identification sends \mathbb{Y} to \mathbb{Y}_3 and it identifies the sets $\{D_z\}_{z \in \mathbb{Y}}$ and $\{D_{3z}\}_{z \in \mathbb{Y}_3}$.*

Assume henceforth that \mathfrak{v} obeys this constraint also.

The final constraint concerns identification between the partnered disks in the sets $\{D_z\}_{z \in \mathbb{Y}}$ and $\{D_{0z}\}_{z \in \mathbb{Y}_0}$. The statement of the constraint refers to radial coordinates (ρ_0, φ_0) for each $z \in \mathbb{Y}_0$ version of D_{0z} . These are chosen so that the form $\frac{1}{2\pi} \sin \theta_+ d\theta_+ d\varphi_+$ appears as the form $\rho_0 d\rho_0 d\varphi_0$. Meanwhile, the corresponding $z \in \mathbb{Y}$ version of D_z has its chosen coordinates (ρ, φ) . The statement refers to an analogous set (ρ_3, φ_3) of radial coordinates on each $z \in \mathbb{Y}_3$ version of D_3 .

Constraint 7 *The identification given by the pseudogradient \mathfrak{v} of each $z \in \mathbb{Y}$ version of D_z with its $z \in \mathbb{Y}_0$ partner D_{0z} identifies the coordinates (ρ, φ) with (ρ_0, φ_0) . The identification given by pseudogradient \mathfrak{v} of each $z \in \mathbb{Y}$ version of D_z with its $z \in \mathbb{Y}_3$ partner D_{3z} identifies the coordinates (ρ, φ) with (ρ_3, φ_3) .*

There is no obstruction to choosing \mathfrak{v} so as to obey this constraint.

Part 3 Use the integral curves of the vector field \mathfrak{v} to define a diffeomorphism between the $f \in [\delta^2, 1 - \delta_*^2]$ part of M_δ and $[\delta^2, 1 - \delta_*^2] \times S^2$. The latter identifies \mathfrak{v} with the Euclidean vector field on the first factor and it identifies the coordinate t on the first factor with f . Use this identification to extend w to the $f \in [\delta^2, 1 - \delta_*^2]$ part of M_δ .

The extension of a requires the choice of a certain smooth 1-form on S^2 . This form is denoted by \hat{a}_0 . It is chosen so that

$$(1-20) \quad \bullet \quad d\hat{a}_0 = \frac{1}{2\pi} \sin \theta_+ d\theta_+ d\varphi_+ \text{ on } S^2 \setminus \left(\bigcup_{z \in \mathbb{Y}_0} D_{0z}\right),$$

$$\bullet \quad \hat{a}_0 = \chi_\#(R_z + \frac{1}{2}\rho_0^2 d\varphi_0) \text{ on any given } z \in \mathbb{Y}_0 \text{ version of } D_{0z}.$$

The 1-form a on the $f \in [1 - 2\delta_*^2, 1 - \delta_*^2]$ part of M_δ is invariant with respect to Lie transport by \mathfrak{v} . Use Lie transport by \mathfrak{v} to extend a to the $f \in [\delta^2, 1 - \delta_*^2]$ part of M_δ . Let a' denote this extension. The identification given by the integral curves of \mathfrak{v} between the $f \in [\delta^2, 1 - \delta_*^2]$ part of M_δ and $[\delta^2, 1 - \delta_*^2] \times S^2$ writes a' as $dt + \hat{a}_1$, where \hat{a}_1 is a smooth, t -independent 1-form on the S^2 factor. It follows from (1-20) that $d\hat{a}_1 = d\hat{a}_0$. There is, as a consequence, a smooth function f_0 on S^2 such that $\hat{a}_1 = \hat{a}_0 + df_0$.

Fix a smooth function on $[0, 1]$ that equals 1 on $[\frac{3}{4}, 1]$ and equals 0 on $[0, \frac{1}{4}]$. Denote this function by Q_0 . Choose a second smooth function on this domain, this one denoted by K_0 . Require that $K_0 \geq 0$, that $K_0 = 1$ on $[\frac{3}{16}, \frac{13}{16}]$, and that $K_0 = 0$ on both $[0, \frac{1}{8}]$ and $[\frac{7}{8}, 1]$. Let r denote the constant used in (1-21). The desired extension of a to the $f \in [\delta^2, 1 - 2\delta_*^2]$ part of M_δ is obtained by first identifying this part of M_δ as done above as with $[\delta^2, 1 - 2\delta_*^2] \times S^2$; then set the extension equal to

$$(1-21) \quad a = (1 + rK_0) dt + \hat{a}_0 + d(Q_0f_0).$$

An analogous construction for the $f \in [2 + 2\delta_*^2, 3 - \delta_*^2]$ part of M_δ extends a now to the whole of $M_\delta \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$. The latter identifies w where $t_- \in [\delta, 7\delta_*]$ with the 2-form $-\frac{1}{2\pi} \sin \theta_- d\theta_- d\varphi_-$, and it writes a as in (1-21) but with \hat{a}_0 replaced by a 1-form \hat{a}_3 , with f_0 replaced by a function f_3 and with suitable replacements for Q_0 and K_0 . In particular, the 1-form \hat{a}_3 is the pullback of \hat{a}_0 by the involution of S^2 that sends θ to $\pi - \theta$.

If $r > c_0$, then the resulting extensions defines a 1-form a and a closed 2-form w on the whole of $M_\delta \cup (\bigcup_p \mathcal{H}_p)$ with the property that $a \wedge w > 0$ and $da \in \text{span}(w)$.

Part 4 This last part of the subsection extends a and w over the handle \mathcal{H}_0 . To this end, reintroduce the coordinates $(u, (\theta, \phi))$ for \mathcal{H}_0 as defined in Section 1A. Define

this extension by setting

$$(1-22) \quad \bullet \quad a = 2(\chi_+ e^{2(|u|-R)} + \chi_- e^{-2(|u|+R)}) du + \hat{a}_0,$$

$$\bullet \quad w = \sin \theta d\theta d\phi.$$

The definition is such that

$$(1-23) \quad dw = 0, \quad a \wedge w > 0 \quad \text{and} \quad da \in \text{span}(w)$$

on \mathcal{H}_0 . In particular, the pair (a, w) defines a stable Hamiltonian structure on \mathcal{H}_0 that extends a and w from $M_\delta \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$ to the whole of Y .

1E Properties of a and w

The pair a and w as just constructed define a stable Hamiltonian structure on Y . Five properties of this pair play key roles in what follows in this paper and in its sequel [12]. These properties are listed below. The notation uses v for the vector field on Y that generates the kernel of w and has pairing equal to 1 with a .

Here are the first two properties:

Property 1 On \mathcal{H}_0 ,

$$w = \sin \theta d\theta \wedge d\phi \quad \text{and} \quad v = \frac{1}{2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)})} \frac{\partial}{\partial u}.$$

Property 2 On M_δ , the 2-form w on M_δ restricts to each constant f slice as an area form. The vector field v on M_δ is the pseudogradient vector field v .

The third property introduces by way of notation functions f and g of the coordinate u which are given by

$$(1-24) \quad f = \chi + 2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)}), \quad g = (\chi_+ e^{2(u-R)} - \chi_- e^{-2(u+R)}).$$

Their derivatives are denoted by f' and g' . This upcoming third property uses α to denote a certain positive function of the coordinates u and θ .

Property 3 Fix $p \in \Lambda$. Then a, w and v on \mathcal{H}_p are

- $a = (\chi + g')(1 - 3 \cos^2 \theta) du - \sqrt{6} f \cos \theta \sin^2 \theta d\phi + 6g \cos \theta \sin \theta d\theta,$
- $w = 6x \cos \theta \sin \theta d\theta \wedge du - \sqrt{6} d\{f \cos \theta \sin^2 \theta d\phi\},$
- $v = \alpha^{-1} \{f(1 - 3 \cos^2 \theta) \partial_u - \sqrt{6} \chi \cos \theta \partial_\phi + f' \cos \theta \sin \theta \partial_\theta\}.$

The fourth property concerns a closed 1-form on Y that extends df from M_δ . This 1-form is denoted in what follows by ν_\diamond . The form ν_\diamond restricts to any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ as df_* with f_* given by (1-9). Meanwhile, ν_\diamond on \mathcal{H}_0 is given by

$$(1-25) \quad \nu_\diamond = 2(\chi_+ e^{2(|u|-R)} + \chi_- e^{-2(|u|+R)}) du.$$

Note that ν_\diamond is exact on $M_\delta \cup (\bigcup_{\mathfrak{p} \in \Lambda} \mathcal{H}_\mathfrak{p})$ as the function f_* extends f over any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$. Even so, ν_\diamond is not an exact form on the whole of Y .

Property 4 *The 1-form ν_\diamond is such that $\nu_\diamond \wedge w \geq 0$ with equality only on the locus in each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ where both $u = 0$ and $1 - 3 \cos^2 \theta = 0$.*

The fact that $\nu_\diamond \wedge w > 0$ on $Y \setminus (\bigcup_{\mathfrak{p} \in \Lambda} \mathcal{H}_\mathfrak{p})$ follows directly from the first and second properties. Meanwhile the formula for w in the third property implies the following on any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$:

$$(1-26) \quad \nu_\diamond \wedge w = \sqrt{6} (g' f (1 - 3 \cos^2 \theta)^2 + 6g f' \cos^2 \theta \sin^2 \theta) du \sin \theta d\theta d\phi.$$

The final property concerns the homology of Y . Note in this regard that the Mayer-Vietoris sequence defines a canonical isomorphism

$$(1-27) \quad H_2(Y; \mathbb{Z}) = H_2(M; \mathbb{Z}) \oplus H_2(\mathcal{H}_0; \mathbb{Z}) \oplus \left(\bigoplus_{\mathfrak{p} \in \Lambda} H_2(\mathcal{H}_\mathfrak{p}; \mathbb{Z}) \right).$$

The summands that correspond to the 1-handles are isomorphic to \mathbb{Z} , and any oriented, cross-sectional sphere is a generator. Orient these spheres with the 2-form $\sin \theta d\theta d\phi$.

Property 5 *Integration of the 2-form w defines the linear map from $H_2(Y; \mathbb{Z})$ to \mathbb{Z} that*

- *has value 2 on the generator of $H_2(\mathcal{H}_0; \mathbb{Z})$;*
- *has value zero on each $\mathfrak{p} \in \Lambda$ version of $H_2(\mathcal{H}_\mathfrak{p}; \mathbb{Z})$;*
- *acts on the $H_2(M; \mathbb{Z})$ summand in (1-27) as the pairing with the chosen class c_{1M} .*

Some additional, less-central properties of a , w and v are pointed out as needed in subsequent sections of this paper.

1F Heegaard Floer homology on M and embedded contact homology on Y

This subsection briefly describes the purely geometric aspects of Heegaard Floer homology and embedded contact homology that have bearing to what follows in this paper and its sequel [12]. This subsection has five parts.

Part 1 The definition of Heegaard Floer homology requires the Morse function f , the choice of a $\text{Spin}^{\mathbb{C}}$ structure on M , and then the choice of a Heegaard diagram that is strongly admissible for the $\text{Spin}^{\mathbb{C}}$ structure's first Chern class. These sorts of Heegaard diagrams are described by Definition 4.10 in [17]. The definition of strong admissibility involves only the image of the first Chern class in $H^2(M; \mathbb{Z})/\text{torsion}$. This understood, take c_{1M} in Lemma 1.1 to be this image in what follows.

There are three versions of Heegaard Floer homology, and the chain complex for these are constructed from suitably constrained and decorated sets of G points in Σ . The G points must lie in $C_- \cap C_+$ with no two in the same version of C_{p_+} nor in the same version of C_{p_-} . As explained in Section 2.6 of [17], this set of G points with the point z_0 defines a $\text{Spin}^{\mathbb{C}}$ structure on M . (More is said about this in a moment.) The set of G points is further constrained so this associated $\text{Spin}^{\mathbb{C}}$ structure is the $\text{Spin}^{\mathbb{C}}$ structure that was chosen a priori. Use \mathcal{Z}_{HF} to denote the collection of sets of G points that satisfy all of these constraints. The decoration, when present, adds an integer label to a given set of points from \mathcal{Z}_{HF} . The differential for the chain complex involves certain sorts of pseudoholomorphic disks in the G -fold symmetric product of Σ .

It proves more convenient for the purposes at hand to use Lipshitz's reformulation of Heegaard Floer homology [13] so as to view the set \mathcal{Z}_{HF} as a collection $\hat{\nu} = \{\nu_1, \dots, \nu_G\}$ of G integral curves of ν in $f^{-1}([1, 2])$ that pair the index 1 critical points of f to the index 2 critical points. The correspondence between this view and the original comes by writing $f^{-1}([1, 2])$ as $(1, 2) \times \Sigma$, this done by using the integral curves of the pseudogradient vector field as explained in Part 5 of Section 1C. By way of a reminder, this identification equates ν with the Euclidean vector field on the $(1, 2)$ factor, and so identifies the set $\hat{\nu}$ with a set of G points in $C_+ \cap C_-$. The pseudoholomorphic disks in the G -fold symmetric product of Σ appear in Lipshitz's formulation as pseudoholomorphic maps from certain sorts of surfaces with boundary into $\mathbb{R} \times [1, 2] \times \Sigma$. The upcoming Section 6 has a detailed description of the sorts of maps that arise and so no more will be said about them here.

The Lipschitz view of Heegaard Floer geometry as geometry on $[1, 2] \times \Sigma$ is used in what follows to interpret the Heegaard Floer geometry as geometry on $f^{-1}((1, 2)) \subset M_\delta$ and thus geometry on Y and $\mathbb{R} \times Y$. This is the bridge that connects the embedded contact homology geometry to the Heegaard Floer geometry.

More is said about this connection in Sections 2, 6 and 7 of this article.

Part 2 The definition of the embedded contact homology chain complex requires the choice of a $\text{Spin}^{\mathbb{C}}$ structure on Y . A $\text{Spin}^{\mathbb{C}}$ structure on the connect sum of two compact, oriented 3-manifolds determines a $\text{Spin}^{\mathbb{C}}$ structure on each of the summands, and conversely a choice of $\text{Spin}^{\mathbb{C}}$ structure on each of the summands determines a $\text{Spin}^{\mathbb{C}}$ structure on their connect sum. Therefore, because $Y = M \# (\#_{G+1}(S^1 \times S^2))$, a $\text{Spin}^{\mathbb{C}}$ structure on Y determines one on M and on each of the $G + 1$ copies of $S^1 \times S^2$, and vice versa. In particular, given a $\text{Spin}^{\mathbb{C}}$ structure on M and a labeled set, $\{n_0, \{n_p\}_{p \in \Lambda}\}$, of even integers, there is a unique $\text{Spin}^{\mathbb{C}}$ structure on Y that induces the given $\text{Spin}^{\mathbb{C}}$ structure on M and is such that $\{n_0, \{n_p\}_{p \in \Lambda}\}$ are the pairings of the first Chern class of the $\text{Spin}^{\mathbb{C}}$ structure on Y with the generators of the 1-handle summands in (1-27). (The $\text{Spin}^{\mathbb{C}}$ structures on $S^1 \times S^2$ are determined by the associated first Chern class, and this can be any even multiple of the generator of $H^2(S^1 \times S^2; \mathbb{Z})$.)

With the preceding understood, and with a $\text{Spin}^{\mathbb{C}}$ structure on M chosen, then the discussion henceforth concerns only the $\text{Spin}^{\mathbb{C}}$ structure on Y with the following properties:

- (1-28) • The $\text{Spin}^{\mathbb{C}}$ structure on Y determines the given $\text{Spin}^{\mathbb{C}}$ structure on M .
 - Its first Chern class has pairing two with the generator of $H_2(\mathcal{H}_0; \mathbb{Z})$.
 - Its first Chern class has pairing zero with the generators of $\bigoplus_{p \in \Lambda} H_2(\mathcal{H}_p; \mathbb{Z})$.

As explained in [8], a generator of embedded contact homology chain complex can be viewed as a finite set, Θ , with each element being a pair (γ, m) where γ is a closed, integral curve of v and where m is a positive integer. The chosen $\text{Spin}^{\mathbb{C}}$ structure further constrains the pairs that form Θ . For the present purposes, the following remarks about this suffice: As done in [25], the set of $\text{Spin}^{\mathbb{C}}$ structures on Y is first identified with the set of equivalence classes of nowhere-zero vector fields on Y with two vector fields being equivalent if they are homotopic through nowhere-zero vector fields on the complement of a ball in Y . Thus, the vector field v defines a $\text{Spin}^{\mathbb{C}}$ structure on Y . This is the *canonical $\text{Spin}^{\mathbb{C}}$ structure*. Meanwhile, the set of $\text{Spin}^{\mathbb{C}}$ structures can also be viewed in a canonical way as an affine space modeled on $H^2(Y; \mathbb{Z})$. Thus,

given the canonical $\text{Spin}^{\mathbb{C}}$ structure, any other $\text{Spin}^{\mathbb{C}}$ structure is determined by (and determines) a class in $H^2(Y; \mathbb{Z})$. This understood, suppose that Θ is a finite collection of pairs of the form (γ, m) with γ being a closed integral curve of v and m being a positive integer. The vector field v orients γ , and with this orientation understood, the loop γ defines a class in $H_1(Y; \mathbb{Z})$. Use $[\gamma]^{\text{Pd}}$ to denote the Poincaré dual class in $H^2(Y; \mathbb{Z})$. Then the $\text{Spin}^{\mathbb{C}}$ structure defined by Θ is obtained from the canonical $\text{Spin}^{\mathbb{C}}$ structure by adding the class $\sum_{(\gamma, m) \in \Theta} m[\gamma]^{\text{Pd}}$. This formula for the $\text{Spin}^{\mathbb{C}}$ structure constrains the elements that form the set Θ by requiring that Θ 's $\text{Spin}^{\mathbb{C}}$ structure be the one on Y that was given a priori. (And keep in mind that the latter must obey the constraints in (1-28).)

If the $\text{Spin}^{\mathbb{C}}$ structure is defined from a set Θ using the rules in the preceding paragraph, then its first Chern class (to be denoted by c_1) can be written as follows: The kernel of a defines a 2-plane subbundle in TY that is oriented by the 2-form ω . Let K^{-1} denote this oriented, real 2-plane bundle and use $e_{K^{-1}}$ to denote the Euler class of K^{-1} , which is a class in $H^2(Y; \mathbb{Z})$. (It is the first Chern class of the canonical $\text{Spin}^{\mathbb{C}}$ structure.) The class c_1 is given by the formula

$$(1-29) \quad c_1 = e_{K^{-1}} + 2 \sum_{(\gamma, m) \in \Theta} m[\gamma]^{\text{Pd}}.$$

Note for future reference that the class $e_{K^{-1}}$ has the following properties:

- (1-30) • It has pairing 2 with the generator of $H_2(\mathcal{H}_0; \mathbb{Z})$.
 • It has pairing -2 with the generators of $\bigcup_{p \in \Lambda} H_2(\mathcal{H}_p; \mathbb{Z})$.

The first bullet follows from what is said in Property 1 of Section 1E, and the second follows from what is said in Property 3 in Section 1E.

Part 3 Hutchings [8] puts an additional constraint on the integers that can appear in pairs from Θ . This part of the subsection describes this new constraint. It is important to keep in mind with regards to this new constraint that the generating set for embedded contact homology can be defined only if certain closed integral curves of v are suitably nondegenerate. More is said about this in a moment. Assume for now that such is the case. What follows directly introduces some background that is needed to define the integer constraint.

Let γ denote a closed integral curve of v . Lie transport by v on a circumnavigation of γ starting at p defines an automorphism $U_\gamma: K^{-1}|_p \rightarrow K^{-1}|_p$ which preserves the 2-form ω since ω is annihilated by v 's Lie derivative. Thus, U_γ acts symplectically

on $K^{-1}|_p$ and so it has determinant 1. The eigenvalues of U_γ are independent of p because a change in p changes U_γ by a conjugation.

The curve γ is said to be *nondegenerate* when U_γ has no eigenvalue equal to 1 or -1 . The curve is said to be *totally nondegenerate* when no power of U_γ has eigenvalue 1 or -1 . The curve γ is said to be *hyperbolic* when U_γ has real eigenvalues, and it is said to be *elliptic* otherwise. The eigenvalues are on the unit circle when γ is elliptic.

Use \mathcal{Z}_{ech} to denote the collection of sets of pairs that obey the first Chern class constraint. The corresponding embedded contact homology chain complex can be defined when all closed integral curves that appear in sets from \mathcal{Z}_{ech} are totally nondegenerate. The upcoming Section 2 explains why this is so for the $\text{Spin}^{\mathbb{C}}$ structure that is specified in (1-28). Granted that such is the case, the embedded contact homology chain complex uses only the elements in \mathcal{Z}_{ech} whose constituent pairs obey the following constraint: supposing that (γ, m) is a pair with γ being a hyperbolic orbit of \mathfrak{v} , the integer m is 1. The subset of elements in \mathcal{Z}_{ech} that obey this constraint is denoted in what follows by $\mathcal{Z}_{\text{ech}, M}$.

Part 4 The differential for the embedded contact homology chain complex and various other important endomorphisms require the choice of a suitable almost complex structure on $\mathbb{R} \times Y$. The almost complex structure is constrained to be invariant with respect to translations along the \mathbb{R} factor of $\mathbb{R} \times Y$. Use s in what follows to denote the Euclidean coordinate for this \mathbb{R} factor. The almost complex structure must map $\frac{\partial}{\partial s}$ to v and it must tame w in the sense that it defines with w a positive-definite, bilinear form on $K^{-1} \subset \mathbb{R} \times Y$. Let J denote such an almost complex structure.

A subset $C \subset \mathbb{R} \times Y$ is said here to be J -holomorphic subvariety in $\mathbb{R} \times Y$ when the following conditions are met:

- (1-31) • The complement in C of a finite set of points is a 2-dimensional submanifold with J -invariant tangent space. Moreover, C has no point components.
- The integral of w over C is finite.
 - There is an $s \in \mathbb{R}$ independent bound for the integral of $ds \wedge a$ over C 's intersection with $[s, s + 1] \times Y$.

With regards to the last two bullets, note that the constraint given in the second bullet allows J to orient the tangent space to the smooth part of C . With this orientation understood, w restricts as nonnegative 2-forms on this tangent space. The almost

complex structure J is not required to preserve the kernel of a and so there is no assumption that $ds \wedge a$ is positive everywhere on the tangent space of C . Even so, there exists $r_* \geq 1$ such that the 2-form

$$(1-32) \quad \omega = ds \wedge a + r_* w$$

tames J in the sense that bilinear form $\omega(\cdot, J\cdot)$ is positive definite on $T(\mathbb{R} \times Y)$. Indeed, this follows because $w(\cdot, J\cdot)$ is positive definite on the kernel of a and the pairing of a with v is 1.

The conditions in (1-32) with the observation that ω tames J imply the following: There exists $s_C \geq 1$ such that the $|s| > s_C$ part of C is a disjoint union of embedded cylinders on which s has no critical points. Moreover, the large $|s|$ slices of any such cylinder appear in Y as a braid in a fixed radius tubular neighborhood of some closed integral curve of v , and the resulting s -parametrized family of braids converges pointwise to this integral curve as $|s| \rightarrow \infty$. This is explained by Hofer [4] under some assumptions, and by Siefring [19] in general. See also Section 2 of [10] for a discussion of the cases that arise in the present context.

A cylinder of the sort just described is said to be an *end* of C . A given end $\mathcal{E} \subset C$ is said to be *positive* when s is unbounded from above on \mathcal{E} and negative when s is unbounded from below. Associated to the given end \mathcal{E} is a pair $(\gamma_{\mathcal{E}}, m_{\mathcal{E}})$ of the following sort: First, $\gamma_{\mathcal{E}}$ is the closed integral curve of v that gives the limit of the large $|s|$ slices of \mathcal{E} . Meanwhile, $m_{\mathcal{E}}$ is the positive integer that describes the degree of the associated projection from any given large $|s|$ slice of \mathcal{E} to $\gamma_{\mathcal{E}}$. Note that the 1-form a orients the large $|s|$ slices and so this degree is in all cases positive.

A J -holomorphic subvariety is said to be *irreducible* when the complement of any finite set of points is connected.

Part 5 The differential and some other relevant endomorphisms of the embedded contact homology chain complex are defined using sets of following sort: A suitable set ϑ consists of a finite collection of pairs, with each of the form (C, n) where C is an irreducible embedded, J -holomorphic submanifold and where n is a positive integer. The integer n is constrained to equal 1 unless C is \mathbb{R} -invariant. The collection ϑ is further constrained in various ways. What follows are the three constraints that are germane to what follows here and in [12]. First, no two pair share the same J -holomorphic subvariety. The second constraint involves the positive ends of the subvarieties from ϑ and the third involves the negative ends of the subvarieties from ϑ .

To say more about the second constraint, let Θ_{ϑ_+} denote the collection of pairs of the form (γ, m) where $\gamma = \gamma_{\mathcal{E}}$ with \mathcal{E} being a positive end of some subvariety from ϑ . Meanwhile, m is a sum of positive integers. A given term in the sum is labeled by a pair $((C, n), \mathcal{E})$ where $(C, n) \in \vartheta$ and \mathcal{E} is a positive end of C with $\gamma_{\mathcal{E}} = \gamma$. The corresponding contribution to m is $nm_{\mathcal{E}}$. The set Θ_{ϑ_+} is defined so that no two pairs share the same closed integral curve. The second constraint requires that Θ_{ϑ_+} define an element in $\mathcal{Z}_{\text{ech}, M}$.

There is a negative end analog of Θ_{ϑ_+} , this denoted by Θ_{ϑ_-} , and the third constraint requires that Θ_{ϑ_-} to likewise define an element in $\mathcal{Z}_{\text{ech}, M}$.

2 The closed integral curves of v

This section first describes various properties of the closed integral curves of v . These properties are used to characterize and then list the elements that form the generating set $\mathcal{Z}_{\text{ech}, M}$. Proposition 2.8 gives the list.

2A The closed integral curves that intersect \mathcal{H}_0

Write \mathcal{H}_0 as $[-R - \ln 7\delta_*, R + \ln 7\delta_*] \times S^2$ as in Section 1A. Property 1 in Section 1E asserts that the vector field v on \mathcal{H}_0 is proportional to the vector field ∂_u . This the case, all integral curves of v that intersect \mathcal{H}_0 do so as $[-R - \ln 7\delta_*, R + \ln 7\delta_*] \times \text{point}$. All such curves intersect any given constant $u \in [-R - \ln 7\delta_*, R + \ln 7\delta_*]$ cross-sectional sphere with the same intersection sign, this being positive when these 2-spheres are oriented by w . In particular, any closed integral curve of v that intersects \mathcal{H}_0 defines a nontorsion class in $H_1(Y; \mathbb{Z})$ with the following property: its Poincaré dual in $H^2(Y; \mathbb{Z})$ has positive pairing with the generator of the $H_2(\mathcal{H}_0; \mathbb{Z})$ summand in (1-27). This last fact with (1-28) and (1-29) imply the following:

(2-1) *The integral curves of v from an element in $\mathcal{Z}_{\text{ech}, M}$ do not intersect \mathcal{H}_0 .*

Granted this, no more will be said here about the integral curves of v that intersect \mathcal{H}_0 except for the following observation: The construction in Section 1D provides $1 + b_1$ of curves of this sort. These curves are labeled by the points in the set \mathbb{Y} that is defined in Parts 6 and 7 of Section 1C. If $z \in \mathbb{Y}$, the corresponding curve intersects the surface Σ at z . This curve is denoted by $\gamma^{(z)}$.

2B Closed integral curves that are disjoint from \mathcal{H}_0

An integral curve of v that intersects M_δ where $f \in (0, 1]$ or where $f \in [2, 3)$ must intersect \mathcal{H}_0 . This follows from Property 2 in Section 1E: the vector field v on this part of Y is a pseudogradient vector field for f . This last fact with (2-1) leads to the following:

(2-2) *The integral curves of v from an element in $\mathcal{Z}_{\text{ech}, M}$ sit entirely in*

$$f^{-1}(1, 2) \cup \left(\bigcup_{\mathfrak{p}} \mathcal{H}_{\mathfrak{p}} \right).$$

The six parts that follow describe the closed integral curves of v that obey (2-2).

Part 1 This part of the subsection describes the parts of the integral curves of v that sit entirely in any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$. The following lemma summarizes. It uses the coordinates $(u, (\theta, \phi))$ for $\mathcal{H}_{\mathfrak{p}}$ from Section 1A.

Lemma 2.1 *Fix $\mathfrak{p} \in \Lambda$. The two circles in $\mathcal{H}_{\mathfrak{p}}$ where $\{u = 0, \cos(\theta) = \pm \frac{1}{\sqrt{3}}\}$ are closed integral curves of v . Both closed integral curves are hyperbolic and such that the associated linear return map has positive eigenvalues. These are the only integral curves of v that lie entirely in $\mathcal{H}_{\mathfrak{p}}$.*

Proof The fact that the circles in question are integral curves of v follows from the formula for v given in Property 3 of Section 1E. Note in this regard that the function f is positive at $u = 0$, but f' is zero and so v is proportional to the vector field ∂_ϕ where both $u = 0$ and $1 - 3 \cos^2 \theta = 0$. The fact that these curves are hyperbolic follows from the form of the linearization of v along the $\{u = 0, \cos(\theta) = \pm \frac{1}{\sqrt{3}}\}$ locus. This is because the linearization is proportional to

$$\partial_\phi + \beta(\chi_0 + 4e^{-2R})((\theta - \theta_*)\partial_u + 16e^{-2R}u\partial_\theta),$$

where β is a suitable constant and where θ_* is the relevant angle.

What follows explains why these are the only two integral curves of v that lie entirely in $\mathcal{H}_{\mathfrak{p}}$. To start, note that the $\theta = \frac{\pi}{2}$ slice is foliated by the integral curves of v . Since the function f is positive, the function u increases along any integral curve of v in this slice, so there can be no such integral curve that stays entirely in $\mathcal{H}_{\mathfrak{p}}$. The $\theta = 0, \pi$ arcs in $\mathcal{H}_{\mathfrak{p}}$ are also integral curves of v , but neither is closed in $\mathcal{H}_{\mathfrak{p}}$. Consider next a point in $\mathcal{H}_{\mathfrak{p}}$ where $\theta \in (0, \frac{\pi}{2})$. Let θ_* denote the angle with $\cos(\theta_*) = \frac{1}{\sqrt{3}}$. If both $u > 0$

and $\theta \geq \theta_*$, then u increases as does θ along the integral curve of v from the point, and so the integral curve must exit \mathcal{H}_p at the $u = R + \ln 7\delta_*$ end. If $u > 0$ and $\theta < \theta_*$, then u decreases initially and θ increases. If the curve does not limit to the $(u = 0, \theta = \theta_*)$ circle, then either θ becomes greater than θ_* at some positive u value, in which case both u and θ increase from that point on and the curve exits the $u = R + \ln 7\delta_*$ end of \mathcal{H}_p . On the other hand, if u becomes negative with $\theta < \theta_*$, then θ begins to decrease, and u continues to decrease, so the curve exits \mathcal{H}_p at the $u = -R - \ln 7\delta_*$ end. A similar analysis applies when u is negative at the given point. \square

Part 2 The upcoming lemma states an important feature of the integral curves of v in a given $p \in \Lambda$ version of \mathcal{H}_p that cross from one boundary 2–sphere to the other.

Lemma 2.2 *There exists $\kappa \geq 100$ that depends only on the Heegaard Floer data and has the following significance: Assume that $\delta < \kappa^{-1}$ and that $\chi_0 \leq \kappa^{-1}\delta^2$. Fix $p \in \Lambda$. Suppose that γ is a connected component in the $|u| \leq \ln R + \delta_*$ part of \mathcal{H}_p of an integral curve of v . Let (θ_-, ϕ_-) and (θ_+, ϕ_+) denote the respective coordinates of the start and endpoints of γ on the $|u| = R + \ln \delta_*$ locus.*

- *Either $\cos(\theta) = 0$ or $|\cos \theta| > 0$ along the whole of γ .*
- *The following assertions are equivalent:*
 - (a) *The endpoints of γ are on different spheres.*
 - (b) *$|1 - 3 \cos^2 \theta| > 0$ along the whole of γ .*
 - (c) *$\theta_+ = \theta_-$.*
 - (d) *$|\cos \theta_{\pm}| \sin^2 \theta_{\pm} < \frac{2}{3\sqrt{3}} \delta_*^{-2} (\chi_0 + 4e^{-R})$.*
- *The start and endpoints of γ are on the respective sphere where $u = -R - \ln \delta_*$ and the sphere where $u = R + \ln \delta_*$ if and only if $1 - 3 \cos^2 \theta > 0$ on γ .*
- *Assume that $u = -R - \ln \delta_*$ on the start point of γ and that $u = R + \ln \delta_*$ on the endpoint. Lift the coordinate ϕ along γ to an \mathbb{R} –valued function and let $\Delta\phi$ denote the net change in this function from start point to endpoint along γ . Then*

$$\Delta\phi = -\sqrt{6} \int_{[-R-\ln \delta_*, R+\ln \delta_*]} \frac{\chi(u)}{f(u)} \frac{\cos(\theta(u))}{1 - 3 \cos^2(\theta(u))} du,$$

where the function $u \rightarrow \theta(u)$ is such that $\theta(u)$ is the unique solution to the equation

$$\cos \theta \sin^2 \theta = 2\delta_*^2 \cos \theta_{\pm} \sin^2 \theta_{\pm} \frac{1}{f(u)},$$

with $1 - 3 \cos^2 \theta > 0$.

- Assume that $u = -R - \ln \delta_*$ on the start point of γ and that $u = R + \ln \delta_*$ on the endpoint. The integral of the 1-form a over the $u \in [-R - \ln \delta_*, R + \ln \delta_*]$ portion of γ is equal to

$$2\chi_0(1 + \epsilon) \int_{[-R - \ln \delta_*, R + \ln \delta_*]} \frac{1 + 3 \cos^4 \theta(u)}{|1 - 3 \cos^2 \theta(u)|} du + \epsilon',$$

where ϵ and ϵ' are such that $|\epsilon| \leq \kappa^{-1}$ and $|\epsilon'| \leq \kappa$.

This lemma is proved in a moment. What follows directly is a lemma to elaborate on the assertion in the fourth bullet regarding the function $u \rightarrow \theta(u)$.

Lemma 2.3 Let $\mathcal{J} \subset [0, \pi]$ denote the open set in $[0, \pi]$ where $|\cos \theta| \sin^2 \theta < \frac{2}{3\sqrt{3}}$. The set \mathcal{J} has three components, these denoted respectively by \mathcal{J}_+ , \mathcal{J}_- and \mathcal{J}_0 . The components \mathcal{J}_+ and \mathcal{J}_- are the respective $\cos \theta > 0$ and $\cos \theta < 0$ parts of the set where $1 - 3 \cos^2 \theta < 0$. The component \mathcal{J}_0 is the region where $1 - 3 \cos^2 \theta > 0$. The maps from \mathcal{J}_+ to $[0, \frac{2}{3\sqrt{3}})$ from \mathcal{J}_- to $(-\frac{2}{3\sqrt{3}}, 0]$ and from \mathcal{J}_0 to $(-\frac{2}{3\sqrt{3}}, \frac{2}{3\sqrt{3}})$ given by $\theta \rightarrow \cos \theta \sin^2 \theta$ are homeomorphisms. Moreover, the map from \mathcal{J}_0 is a diffeomorphism, while those from \mathcal{J}_+ and \mathcal{J}_- are diffeomorphisms on their interiors.

Proof This follows from the inverse function theorem given that the differential of the map $\theta \rightarrow \cos \theta \sin^2 \theta$ is equal to $-(1 - 3 \cos^2 \theta) \sin \theta$. □

Proof of Lemma 2.2 The first bullet’s assertion follows from what is said in the proof of Lemma 2.1.

What is said in the proof of Lemma 2.1 also proves that item (a) of the second bullet implies (b) and that (c) implies (a). The fact that (b) implies (a) follows from the formula for v given in Property 3 of Section 1E as the latter implies that the coefficient in front of ∂_u is nowhere zero. To obtain (c) from (b), use the fact that the coefficient in front of ∂_u is nowhere zero when (b) holds. Then u can be used as an affine parameter along γ and the fact that $f'(-u) = -f'(u)$ implies (c). To obtain (d) from (b), view θ as a function of u along the integral curve. Then it follows from the third bullet of Property 3 in Section 1E that $\frac{d\theta}{du}$ is given by

$$(2-3) \quad \frac{d\theta}{du} = \left(\frac{f'(u)}{f(u)} \right) \frac{\cos \theta \sin \theta}{1 - 3 \cos^2 \theta}.$$

The u and θ variables can be separated by writing this equation as $d(\ln(f(u))) = d(\ln(\cos \theta \sin^2 \theta))$, which can be integrated to see that

$$(2-4) \quad \frac{\cos(\theta(u)) \sin^2(\theta(u))}{\cos(\theta_{\pm}) \sin^2(\theta_{\pm})} = \frac{f(R + \ln \delta_*)}{f(u)}.$$

To obtain the constraint on θ_{\pm} , note first that the function $\theta \rightarrow |\cos \theta| \sin^2 \theta$ is maximized where $\cos^2 \theta = \frac{1}{3}$. Meanwhile, $f(R + \ln \delta_*) = 2\delta_*^2$ and the minimum of the function $f(\cdot)$ is at $u = 0$, where $f(0) = \chi_0 + 4e^{-R}$. These observations with (2-4) give the constraint on $|\cos \theta_{\pm}| \sin^2 \theta_{\pm}$ in (d). Similar considerations prove that the other items follow from (d).

The assertion that is made by the third bullet follows from the fact that the coefficient in front of ∂_u in the formula for v in Property 3 of Section 1E is positive if and only if $1 - 3 \cos^2 \theta > 0$.

To prove the fourth bullet, view an \mathbb{R} -valued lift of the coordinate ϕ along γ as a function of u also. This done, the formula for v in Property 3 of Section 1E implies that the function $u \rightarrow \phi(u)$ obeys

$$(2-5) \quad \frac{d\phi}{du} = -\sqrt{6} \frac{\chi(u)}{f(u)} \frac{\cos \theta(u)}{1 - 3 \cos^2 \theta(u)}.$$

The formula given in the fourth bullet follows by integrating (2-5).

The last bullet follows by first parametrizing γ by the function u , and then using the formula for a in Property 3 of Section 1E with the equations in (2-3) and (2-5) for $\frac{d\theta}{du}$ and $\frac{d\phi}{du}$ to write the pullback of a to γ . The terms ϵ and ϵ' are the contributions from terms with g and f' and where $\chi \neq \chi_0$. □

Part 3 This part of the subsection constitutes a digression to relate certain integral curves of v with those of the pseudogradient vector field \mathfrak{v} for f . To this end, keep in mind that $v = \mathfrak{v}$ on M_{δ} . The next lemma states a consequence of this fact.

Lemma 2.4 *There exists $\kappa \geq 1$ with the following significance: Assume that $\delta < \kappa^{-1} \delta_*^3$. Let p and p' denote respective index 1 and index 2 critical points of f . Suppose that γ is a closed, connected segment of an integral curve of v that starts on the boundary of the radius δ coordinate ball centered at p and ends on the boundary of the radius δ coordinate ball centered at p' . Assume, in addition, that the interior of γ is disjoint from $\mathcal{H}_0 \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$. Then γ sits in the radius $\kappa \delta$ tubular neighborhood of an integral curve of \mathfrak{v} in M that starts at p and ends at p' .*

Proof The segment γ is an integral curve of \mathfrak{v} . Parametrize γ as an embedding $s \rightarrow \gamma(s)$ from $[0, 1]$ such that $\gamma(0)$ sits on the boundary of the radius δ coordinate ball centered at p . The vector field \mathfrak{v} in the radius δ_* ball about p is depicted in (1-14) and $v = \mathfrak{v}$ where the distance from p is at least δ . It follows as a consequence that $1 - 3 \cos^2 \theta \geq 0$ on $\gamma(0)$. This implies that $f \geq 1$ on γ . Note that f on γ cannot be greater than 2 for if so, then either the interior of γ must intersect some radius δ coordinate ball about an index 2 critical point, or else γ would intersect \mathcal{H}_0 . This understood, the whole of γ must be a segment of an integral curve of \mathfrak{v} lying in $f^{-1}[1, 2]$ and $\gamma(1)$ must lie on the $f \leq 2$ part of the boundary of the radius δ coordinate ball centered on p' . This is to say that there is an integral curve of \mathfrak{v} that connects the radius δ coordinate ball centered on p with that centered on p' . It then follows (by arguing to the contrary to obtain a contradiction) that there exists $c_0 \geq 1$ with the following significance: if $\delta \leq c_0^{-1}$, then γ must lie in the radius $c_0\delta$ tubular neighborhood of an integral curve of \mathfrak{v} that runs from p to p' . \square

Part 4 Lemma 2.4 has the following consequence:

Lemma 2.5 *There exists $\kappa \geq 1$ with the following significance: Assume that $\delta < \kappa^{-1}\delta_*^3$. Suppose that γ is a closed, integral curve of v that does not intersect \mathcal{H}_0 and does not lie entirely in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$. Then:*

- *The intersection of γ with M_δ must lie where $f \in [1, 2]$ and each component of this intersection lies in a radius $\kappa\delta$ tubular neighborhood of an integral curve of \mathfrak{v} that runs from an index 1 critical point of f to an index 2 critical point.*
- *A properly embedded, connected segment in γ must cross at least one $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$ starting on the boundary 2-sphere at the index 2 critical point end of $\mathcal{H}_{\mathfrak{p}}$ and ending on the boundary 2-sphere at the index 1 critical point end.*
- *Let $\mathfrak{p} = (p, p')$ denote a given point in Λ . If γ intersects the boundary of the radius δ coordinate ball about p' , then it does so where $f \leq 2$ and the continuation of γ from this intersection point crosses $\mathcal{H}_{\mathfrak{p}}$ to an intersection point with the boundary of the radius δ coordinate ball about p .*
- *Let $\mathfrak{p} = (p, p')$ again denote a given point in Λ . If γ intersects the boundary of the radius δ coordinate ball about p , then it does so where $f \geq 1$ and the continuation of γ backwards from this intersection point crosses $\mathcal{H}_{\mathfrak{p}}$ in reverse to an intersection point with the boundary of the radius δ coordinate ball p' .*

Proof Let $\gamma_1 \subset \gamma$ denote the closure of a connected component of the intersection of γ with M_δ . Since $v = \mathfrak{v}$ on γ_1 , this cannot be the whole of γ . As a consequence, γ_1 must have a boundary. One boundary point must lie on the $f \geq 1$ part of the boundary of the radius δ coordinate ball centered on some index 1 critical point, and the other on the $f \leq 2$ part of the boundary of the radius δ coordinate ball centered on an index 2 critical point. This follows from the fact that \mathfrak{v} near such critical points obeys Constraint 1 in Section 1C. This understood, the claim made by the first bullet follows directly from Lemma 2.4.

Let p' denote the endpoint of γ_1 on the boundary of the radius δ coordinate ball centered on the index 2 critical point. Let γ_2 denote the connected component of γ in the corresponding version of \mathcal{H}_p that has p' as its start point. The other endpoint of this segment must lie on the boundary 2–sphere at the index 1 critical point end of \mathcal{H}_p . Indeed, if this were not the case, then it follows from Constraint 1 in Section 1C that the other endpoint would lie on the $f > 2$ part of the boundary of the radius δ ball centered at p' , and so continuing along γ from this endpoint would lead to an intersection with \mathcal{H}_0 .

These last observations imply what is claimed by the second and third bullets of the lemma. A very much analogous argument proves the claim made by the fourth bullet. \square

The following is a direct consequence of Lemma 2.5:

Corollary 2.6 *There exists $\kappa \geq 1$ such that when $\delta < \kappa^{-1} \delta_*^3$, what follows is true. Suppose that γ is a closed, integral curve of v that does not intersect \mathcal{H}_0 and does not lie entirely in some $p \in \Lambda$ version of \mathcal{H}_p . Then there exists a positive integer N and a cyclically ordered set $\{p^1, \dots, p^N\}$ of not necessarily distinct elements from Λ with the following property: The curve γ is an end-to-end concatenation of $2N$ closed, connected segments $\{\gamma^{p^1}, \gamma_1, \gamma^{p^2}, \gamma_2, \dots, \gamma^{p^N}\}$ where:*

- *For each $k \in \{1, \dots, N\}$, the segment γ^{p^k} is a component of $\gamma \cap \mathcal{H}_{p^k}$. This segment starts on the $f \leq 2$ part of the boundary 2–sphere of the index 2 critical point end of the handle and ends on the $f \geq 1$ part of the boundary of the index 1 critical point end of the handle.*
- *For each $k \in \{1, \dots, N\}$, the segment γ_k lies in M_δ . It starts at the endpoint of γ^{p^k} and, for $k < N$, it ends at the start point of $\gamma^{p^{k+1}}$. The segment γ_N ends at the start point of γ^{p^1} .*

- For each $k \in \{1, \dots, N\}$, the segment γ_k lies in the radius $\kappa\delta$ tubular neighborhood of an integral curve of \mathfrak{v} that runs from the index 1 critical point component of \mathfrak{p}^k to the index 2 critical point component of \mathfrak{p}^{k+1} when $k < N$, and to the index 2 critical point component of \mathfrak{p}^1 when $k = N$.

This corollary motivates the following definition: An *index 1-2 cycle* is a nonempty, finite, cyclically ordered set $\{\nu_1, \dots, \nu_N\}$ whose components are the closures of integral curves of \mathfrak{v} that connect index 1 critical points of f to index 2 critical points of f and are such that the following is true: For $k < N$, the endpoint of ν_k and the start point of ν_{k+1} define a pair from Λ . This is also true for the endpoint of ν_N and the start point of ν_1 . The integer N is said to be the length of the index 1-2 cycle.

Part 5 The preceding corollary asserts that any closed integral curve of v that avoids \mathcal{H}_0 and does not lie entirely in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$ defines an index 1-2 cycle in the sense that the restriction of the curve to M_{δ} is a union of some number N segments that have distance $\mathcal{O}(\delta)$ to the segments that define an index 1-2 cycle. Conversely, each segment from this index 1-2 cycle has distance $\mathcal{O}(\delta)$ to a segment of the M_{δ} part of the integral curve. The upcoming Proposition 2.7 describes the set closed integral curves of v that correspond in this way to any given index 1-2 cycle.

The proposition refers to an *orientation sign*, either $+$ or $-$, that is associated to a given integral curve of \mathfrak{v} that runs from an index 1 critical point of f to an index 2 critical point. What follows directly explains how this sign is determined. To start, note that the pseudogradient \mathfrak{v} is such that the ascending disks from the index 1 critical points of f have transversal intersection with the descending disks from the index 2 critical points. These intersections are the integral curves of \mathfrak{v} that run from the index 1 to the index 2 critical points. Let p denote a given index 1 critical point. The intersection of the ascending disk from p with $\Sigma = f^{-1}(\frac{3}{2})$ is the circle C_{p+} from Part 2 of Section 1C. By way of a reminder, the integral curves of \mathfrak{v} identify C_{p+} with the $\theta_+ = \frac{\pi}{2}$ circle in any constant radius sphere in the radius $8\delta_*$ coordinate ball centered at p . Introduce from Part 3 of Section 1C the functions (φ_+, h_+) for the tubular neighborhood T_{p+} of C_{p+} and then orient C_{p+} using the 1-form $d\varphi_+$.

Let p' denote a given index 2 critical point of f and let $C_{p'}$ denote the intersection between the descending disk from p' with Σ . Introduce the coordinates (φ_-, h_-) from Part 3 of Section 1C and orient $C_{p'}$ using the 1-form $d\varphi_-$. Meanwhile, the level set Σ is oriented using the area form ω_{Σ} . Granted this orientation, any given

intersection point in Σ between C_{p+} and C_{p-} has an associated sign; this sign is obtained by comparing the orientation of Σ with that defined by the pair $(\frac{\partial}{\partial\varphi_+}, \frac{\partial}{\partial\varphi_-})$ of respective oriented tangent vectors to C_{p+} and C_{p-} at the intersection point. Each such intersection point corresponds to a unique integral curve of v from p to p' . This orientation sign is the sign associated to the given integral curve.

Proposition 2.7 *There exists $\kappa \geq 1$ and, given $\delta < \kappa^{-1}\delta_*^3$, there exists $\kappa_\delta > \kappa |\ln \delta|$ with the following significance: Take $R \geq \kappa |\ln \delta|$. Let $\{v_1, \dots, v_N\}$ denote an index 1-2 cycle. The set of closed orbits of v that correspond to this cycle enjoys a 1-to-1 correspondence with the set of N -tuples of integers. To elaborate, let $\{\xi_1, \dots, \xi_N\}$ denote a given N -tuple of integers. The corresponding closed integral curve of v has a decomposition into segments as described in Corollary 2.6 such that for each $k \in \{1, \dots, N\}$, the change in the angle ϕ for the segment γ^{p^k} can be written as $\sigma_k + 2\pi\xi_k$ with $\sigma_k \in [0, 2\pi)$. In all cases, the resulting closed integral curve is hyperbolic, and the sign of the eigenvalues is $(-1)^N$ times the product of the orientation signs of the integral curves $\{v_1, \dots, v_N\}$. The integral, ℓ , of the 1-form a over this integral curve obeys*

$$\kappa^{-1} \sum_{1 \leq k \leq N} (\chi_0 R + |\xi_k|) - \kappa \leq \ell \leq \kappa \sum_{1 \leq k \leq N} (\chi_0 R + |\xi_k|) + \kappa.$$

Two paired sets of length N index 1-2 cycles and N integers label the same closed integral curve of v if one is obtained from the other by the action of $\mathbb{Z}/(N\mathbb{Z})$ that relabels both the set $\{v_1, \dots, v_N\}$ and the set $\{\xi_1, \dots, \xi_N\}$ by a simultaneous cyclic permutation of the labels $\{1, \dots, N\}$.

To elaborate on this business of $\mathbb{Z}/(N\mathbb{Z})$ permutations, suppose for the moment that $\{(v_1, \dots, v_N), (\xi_1, \dots, \xi_N)\}$ is a data set as described in the proposition such that $\mathbb{Z}/(N\mathbb{Z})$ acts with nontrivial stabilizer. Let $n \in \{2, \dots, N - 1\}$ denote the smallest integer such that the cycle of length n acts trivially on this set. The number n is a divisor of N . A suitable cyclic permutation of (v_1, \dots, v_N) writes this set as N/n consecutive copies of (v_1, \dots, v_n) and the same cyclic permutation simultaneously writes the set (ξ_1, \dots, ξ_N) as N/n consecutive copies of (ξ_1, \dots, ξ_n) . The former is an index 1-2 cycle also, and the data $\{(v_1, \dots, v_n), (\xi_1, \dots, \xi_n)\}$ defines the same integral curve of v as does the original, larger set. (The larger set can be viewed as labeling an N/n times multiple cover of the smaller set.) This understood, no generality is lost by restricting attention to sets where $\mathbb{Z}/(N\mathbb{Z})$ acts freely. This restriction is implicit in what follows.

Proof of Proposition 2.7 The proof is given in four steps.

Step 1 Keep in mind for what follows that the pseudogradient \mathfrak{v} that gives v on M_δ was chosen in part to guarantee that the ascending disks from the index 1 critical points of f intersect the descending disks from the index 2 critical points in a transversal fashion on Σ as described in Constraint 3 of Section 1C.

Let p denote an index 1 critical point of f and let ν denote an integral curve of \mathfrak{v} that runs from p to some index 2 critical point, p' . Reintroduce the coordinates (φ_+, \hat{h}_+) for the annular neighborhood $T_{p_+} \subset \Sigma$ of C_{p_+} . Keep in mind that φ_+ is $\mathbb{R}/2\pi\mathbb{Z}$ -valued and that $|\hat{h}_+| < \frac{128}{3\sqrt{3}}\delta_*^2$. The functions (φ_+, \hat{h}_+) are annihilated by \mathfrak{v} , and this has the following implication: an integral curve of \mathfrak{v} that intersects T_{p_+} with coordinates $\varphi_+ = \varphi$ and $\hat{h}_+ = \hat{h}$ intersects the radius $8\delta_*$ coordinate ball with the coordinates $(r_+, \theta_+, \varphi_+)$ with r_+ and θ_+ such that $2r_+^2 \cos \theta_+ + \sin^2 \theta_+ = \hat{h}_+$ and with $\varphi_+ = \varphi$.

Reintroduce the corresponding coordinates (φ_-, \hat{h}_-) for the annular neighborhood $T_{p'_-}$ of $C_{p'_-}$. As noted in Constraint 3 of Section 1C, the respective \pm pairs of coordinate differentials on ν 's component of $T_{p_+} \cap T_{p'_-}$ are such that $(d\varphi_+, d\hat{h}_+) = \varepsilon_\nu(d\hat{h}_-, d\varphi_-)$ where $\varepsilon_\nu \in \{-1, 1\}$ is the orientation sign for ν . The intersection point of Σ with ν is a point in $C_{p_+} \cap C_{p'_-}$. This point occurs where $\hat{h}_+ = 0$ and $\hat{h}_- = 0$. Write the respective φ_+ and φ_- coordinates of this point as ϕ_{ν_+} and ϕ_{ν_-} . The aforementioned identity between the \pm pairs of coordinate differentials implies that the functions (φ_+, \hat{h}_+) and (φ_-, \hat{h}_-) are related on ν 's component of $T_{p_+} \cap T_{p'_-}$ as follows: Lift ϕ_{ν_+} to $[0, 2\pi)$ and likewise lift ϕ_{ν_-} . These lifts determine lifts of φ_+ and φ_- to \mathbb{R} -valued functions. So interpreted, these \mathbb{R} -valued versions with \hat{h}_+ and \hat{h}_- obey

$$(2-6) \quad \varphi_- - \phi_{\nu_-} = \varepsilon_\nu \hat{h}_+ \quad \text{and} \quad \hat{h}_- = \varepsilon_\nu(\varphi_+ - \phi_{\nu_+})$$

on ν 's component of $T_{p_+} \cap T_{p'_-}$.

Step 2 Let γ denote a given integral curve of the vector field v that leaves the radius δ_* coordinate ball centered p_- with given values for φ_+ and \hat{h}_+ on the boundary 2-sphere. Write the \hat{h}_+ value as \hat{h}_{+0} . Let φ_{+0} denote a lift to \mathbb{R} of the φ_+ value. Assume that this lift and \hat{h}_+ are such that

$$(2-7) \quad |\varphi_{+0} - \phi_{\nu_+}| + |\hat{h}_{+0}| < \frac{4}{3\sqrt{3}}\delta_*^2.$$

This guarantees that γ intersects ν 's component of $T_{p_+} \cap T_{p'_-}$ and thus the boundary of the radius δ_* coordinate ball centered on p' .

Let \mathfrak{p} denote the pair from Λ with index 2 critical point p' and let p_1 denote the index 1 critical point that is paired with p' in \mathfrak{p} . Given (2-7), Lemma 2.2 asserts the following: The integral curve γ passes through the 1-handle $\mathcal{H}_{\mathfrak{p}}$ so as to exit the radius δ_* ball centered on p_1 if and only if this intersection point occurs where (φ_-, h_-) are such that

$$(2-8) \quad |h_-| < \frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-R}).$$

Assume that (2-8) also holds for the chosen integral curve. What with (2-6), this requires that

$$(2-9) \quad |\varphi_{+0} - \phi_{v_+}| < \frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-R}).$$

Let (φ_{+1}, h_{+1}) denote the corresponding pair that are defined by the values of (φ_+, h_+) where the integral curve γ intersects the boundary of the radius δ_* ball centered on p_1 . It follows from Lemma 2.2 and (2-6) that (φ_{+1}, h_{+1}) are given in terms of (φ_{+0}, h_{+0}) by the rule

$$(2-10) \quad \begin{aligned} &\bullet \quad \varphi_{+1} - \phi_{v_-} = \varepsilon_v h_{+0} + \Delta\phi_{\mathfrak{p}}, \\ &\bullet \quad h_{+1} = -\varepsilon_v(\varepsilon_{+0} - \phi_{v_+}), \end{aligned}$$

where $\Delta\phi_{\mathfrak{p}}$ denotes the integral that is written in the fourth bullet of Lemma 2.2.

Step 3 This step proves Proposition 2.7 for the simplest case, this when the index 1-2 cycle has the single element $\mathfrak{p} = \mathfrak{p}^{(1)}$. Write $\mathfrak{p} = (p, p')$. Apply what is said in the previous steps with $p_1 = p$. The segment, γ , is a closed integral curve of v if and only if (2-10) holds with

$$(2-11) \quad h_{+1} = h_{+0} \quad \text{and} \quad \varphi_{+1} = \varphi_{+0} + 2\pi\mathfrak{k} \quad \text{for some } \mathfrak{k} \in \mathbb{Z}_-.$$

To determine the solutions, use the bottom equation in (2-10) to write $\varphi_{+0} - \phi_{v_+} = -\varepsilon_v h_{+1}$. Having done so, equation (2-11) follows from the top equation in (2-10) if and only if

$$(2-12) \quad 2\pi\mathfrak{k} = \phi_{v_-} - \phi_{v_+} + 2\varepsilon_v h_{+1} + \Delta\phi_{\mathfrak{p}}.$$

To see about solving (2-12), first use Lemma 2.3 to find $c_0 > 1$ such that if $\delta_*^{-2}\chi_0 \leq c_0^{-1}$, then the set in $[0, \pi]$ where $|\cos \theta| \sin^2 \theta < \frac{2}{3\sqrt{3}}\delta_*^{-2}(\chi_0 + 4e^{-2R})$ has three components with one an interval centered on $\frac{\pi}{2}$. The function $h = \delta_*^2 \cos \theta \sin^2 \theta$ maps this component diffeomorphically to the open interval $(-\frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-2R}), \frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-2R}))$.

Let \mathcal{I} denote this interval. As is explained in the next paragraph, there is a unique $h_{+1} \in \mathcal{I}$ that obeys (2-12) for any given integer m .

Write h_{+1} as h so as to view the right-hand side of (2-12) as a map, $h \rightarrow F(h)$, from \mathcal{I} to \mathbb{R} . It follows from the fourth bullet of Lemma 2.2 that F is a proper map from \mathcal{I} to \mathbb{R} if $R\delta_*^2\chi_0^{-1} \geq c_0$. Indeed, the chain rule with the definition of $\theta(u)$ in the fourth bullet of Lemma 2.2 can be used to see that the derivative of F on \mathcal{I} obeys

$$(2-13) \quad \frac{dF}{dh} = -\sqrt{6} \int_{[-R-\ln\delta, R+\ln\delta]} \frac{\chi(u)}{f(u)^2} \frac{1 + 3 \cos^2(\theta(u))}{(1 - 3 \cos^2(\theta(u)))^3} du + 2\varepsilon_\nu.$$

The integral that appears on the right-hand side of (2-13) is no less than $c_0^{-1}R\chi_0^{-1}$. This fact about (2-13) implies that F is a diffeomorphism from \mathcal{I} to \mathbb{R} if $R\chi_0^{-1} > c_0$.

Assuming the bound $R\chi_0^{-1} > c_0$, it then follows that for each $m \in \mathbb{Z}$, there is a unique $h_{1+} = h_m \in \mathcal{I}$ such that solves (2-12). Note that (2-9) follows now automatically from the top bullet in (2-10), and thus (2-7) as well.

To continue with the proof of Proposition 2.7 when $N = 1$, let γ denote the closed integral curve of ν that is defined as above for a given $\mathfrak{k} \in \mathbb{Z}$. It follows from (2-6), (2-10), (2-12) and (2-13) that the linearized return map defines a matrix in $SL(2; \mathbb{R})$ that has the form

$$(2-14) \quad \begin{pmatrix} -\varepsilon_\nu \sigma R & \varepsilon_\nu \\ -\varepsilon_\nu & 0 \end{pmatrix},$$

with $\sigma > c_0^{-1}$. The eigenvalues of this matrix are real, and one has distance $\mathcal{O}(R)$ from the unit circle when R is large. The sign of the large eigenvalue is that of $-\varepsilon_\nu$, thus -1 times the orientation sign given to ν .

The final assertion of Proposition 2.7 gives a bound from above and below for the integral of a along γ . This bound follows from the last bullet of Lemma 2.2.

Step 4 This step proves Proposition 2.7 for the $N > 1$ cases. To set notation, the index k in what follows is from the set $\{1, \dots, N\}$, and if $k = N$, then $k + 1$ denotes the integer 1.

Write the pairs in Λ that are associated to the index 1 start points of the curves from the set $\{\nu_1, \dots, \nu_N\}$ as $\{p^1, \dots, p^N\}$ and write any given $k \in \{1, \dots, N\}$ version of p^k as $(p^{(k)}, p'^{(k)})$. The closed integral curve ν_k from the cycle runs from $p^{(k)}$ to $p'^{(k+1)}$. For each $k \in \{1, \dots, N\}$, let ϕ_{k+} denote the $[0, 2\pi)$ lift of the φ_+ coordinate of the intersection point between ν_k and the boundary sphere of the radius δ_* coordinate

ball centered on $p^{(k)}$, and use ϕ_{k-} to denote the $[0, 2\pi)$ lift of the φ_- coordinate of the intersection point between ν_k and the boundary sphere in the radius δ_* coordinate ball centered on $p^{(k)}$. Each $k \in \{1, \dots, N\}$ version of ν_k intersects Σ in a distinct component of $T_- \cap T_+$. Each such component has the two sets of coordinates (φ_+, \hat{h}_+) and (φ_-, \hat{h}_-) , and these are related by (2-6). Use ε_{ν_k} for the sign ε_ν that appears in ν_k 's version of (2-6).

Suppose that $\{\gamma_{1*}, \dots, \gamma_{N*}\}$ are segments of integral curves of v in M_δ such that each $k \in \{1, \dots, N\}$ version of γ_k intersects ν_k 's component of $T_+ \cap T_-$. Use $(\varphi_{+k}, \hat{h}_{+k})$ to denote the (φ_+, \hat{h}_+) coordinates of this intersection point with φ_{+k} a lift to \mathbb{R} of φ_+ such that the pair $(\varphi_{+k} - \phi_{k+}, \hat{h}_{+k})$ obey (2-7). Meanwhile, use $(\varphi_{-k}, \hat{h}_{-k})$ to denote the (φ_-, \hat{h}_-) coordinates with φ_{-k} denoting a lift to \mathbb{R} of φ_- such that $(\varphi_{-k} - \phi_{k-}, \hat{h}_{-k})$ also obey (2-7). Assume that \hat{h}_{-k} obeys (2-8) and so $\varphi_{+k} - \phi_{k+}$ obeys (2-9). This condition is necessary and sufficient to conclude that γ_k extends across $\mathcal{H}_{p^{k+1}}$ as a segment of an integral curve of v .

The segment γ_k enters the handle $\mathcal{H}_{p^{k+1}}$ with the value $-\hat{h}_{-k}$ for $\hat{h} = f \cos \theta \sin^2 \theta$ and with ϕ coordinate that lifts to \mathbb{R} as φ_{-k} . The extended segment exits this handle with the lifted ϕ coordinate and \hat{h} equal to $\varphi_{-k} + \Delta\phi_{p^{k+1}}$ and $-\hat{h}_{-k}$. Here, $\Delta\phi_{p^{(k+1)}}$ is the relevant version of the integral on the right-hand side of the third bullet in Lemma 2.2. This understood, the extended γ_k concatenates with γ_{k+1} to form a segment of an integral curve of v if and only if there exists $\mathfrak{k}_{k+1} \in \mathbb{Z}$ such that

$$(2-15) \quad \varphi_{-k} + \Delta\phi_{p^{(k+1)}} = \varphi_{+(k+1)} + 2\pi\mathfrak{k}_{k+1} \quad \text{and} \quad -\hat{h}_{-k} = \hat{h}_{k+1}.$$

Meanwhile, $\varphi_{-k} - \phi_{k-} = \varepsilon_{\nu_k} \hat{h}_{+k}$ and $\hat{h}_{-k} = \varepsilon_{\nu_k} (\varphi_{+k} - \phi_{k+})$. Thus, the N extended segments concatenate to define a closed integral curve of v if and only if the following condition is met for each $k \in \{1, \dots, N\}$: there exists an integer \mathfrak{k}_{k+1} such that

$$(2-16) \quad 2\pi\mathfrak{k}_{k+1} = \phi_{k-} - \phi_{(k+1)+} + \varepsilon_k \hat{h}_{+k} + \Delta\phi_{p^{k+1}}.$$

What follows explains why there is a unique solution to the system of equations in (2-16) for any given set $\{\mathfrak{k}_1, \dots, \mathfrak{k}_N\} \in \times_N \mathbb{Z}$ if $\delta_*^{-2} \chi_0 < c_0^{-1}$.

Reintroduce from Step 3 the interval $\mathcal{I} = (-\frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-2R}), \frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-2R}))$. The right-hand side of (2-16) defines a proper, smooth map from $\times_N \mathcal{I}$ to $\times_N \mathbb{R}$. It follows from (2-13) that this map defines a diffeomorphism if $R\delta_*^2 \chi_0^{-1} > c_0$. Assume this. Then there exists, for each N -tuple $\{\mathfrak{k}_1, \dots, \mathfrak{k}_N\} \in \times_N \mathbb{Z}$, a unique N -tuple $(\hat{h}_{+1}, \dots, \hat{h}_{+N}) \in \times_N \mathcal{I}$ that solves (2-16). As in the $N = 1$ case, (2-9) and (2-7) follow automatically.

The preceding gives the desired 1-to-1 correspondence between $\times_N \mathbb{Z}$ and the set of closed trajectories of ν that define the given index 1-2 cycle. The fact that each such closed trajectory is hyperbolic can be seen as follows: View the right-hand side of (2-15) as defining a map from \mathbb{R}^2 to itself. The differential of this map at $(\varphi_{-k}, \hat{h}_{-k})$ is a matrix in $\mathrm{SL}(2; \mathbb{R})$. The linearized return map is the appropriately ordered product of these N matrices. It follows from (2-13) that this ordered product has trace that can be written as

$$(2-17) \quad \sigma_N (-1)^{N_+} R^N \prod_{1 \leq k \leq n} \varepsilon_{\nu_k},$$

where $\sigma_N > c_N^{-1}$ with $c_N \geq 1$ determined by N . Here, N_+ is the number of segments of the index 1-2 cycle with positive orientation sign. Granted (2-17), it follows that the closed integral curve in question is hyperbolic, and its eigenvalues have the sign asserted by Proposition 2.7. The assertion about the integral of a over this closed integral curve follows directly using the fifth bullet of Lemma 2.2.

The final assertion about the $\mathbb{Z}/(N\mathbb{Z})$ action is a consequence of the fact that the map from $\times_N \mathcal{I}_N$ to $\times_N \mathbb{R}$ defined by the right-hand side of (2-16) is a diffeomorphism when $R\delta_*^2 \chi_0^{-1}$ is large. \square

2C The set $\mathcal{Z}_{\mathrm{ech}, \mathcal{M}}$

What is said in Sections 2A and 2B allows for a complete description of the set $\mathcal{Z}_{\mathrm{ech}, \mathcal{M}}$. This is provided by the upcoming Proposition 2.8. The four parts that follow directly supply the necessary background.

Part 1 The correspondence from Proposition 2.7 is canonical with the choice of the following data: Choose a lift to $[0, 2\pi)$ for the φ_+ coordinate function on the radius $8\delta_*$ coordinate ball about each index 1 critical point of f . Assume that such a lift has been chosen. Note that these G lifts supply via the handle identifications lifts to $[0, 2\pi)$ of the φ_- coordinate function on the radius $8\delta_*$ ball about any given index 2 critical point of f . The correspondence also requires a labeling of the elements in Λ as $\{\mathfrak{p}^1, \dots, \mathfrak{p}^G\}$.

Part 2 Reintroduce from Part 1 of Section 1F the set $\mathcal{Z}_{\mathrm{HF}}$. As noted in Section 1F, this is the set that is used to define the generators for the Heegaard Floer chain complex for the given $\mathrm{Spin}^{\mathbb{C}}$ structure. The set $\mathcal{Z}_{\mathrm{HF}}$ is viewed in what follows as the set of

unordered G -tuples whose elements are integral curves of v that start at an index 1 critical point of f and end at an index 2 critical point with the G -tuple constrained as follows: no two elements start at the same index 1 critical point and no two end at the same index 2 critical point. This collection of G arcs is also constrained so that these arcs and the integral curve of v through z_0 define the given $\text{Spin}^{\mathbb{C}}$ structure on M using the rules from Section 2.6 of [17]. To say more about the $\text{Spin}^{\mathbb{C}}$ structure that is defined by an element, \hat{v} , in \mathcal{Z}_{HF} note the G arcs that form \hat{v} and the closure of the integral curve through z_0 have pairwise disjoint neighborhoods that are embedded balls in M with smooth, disjoint boundaries. Since each ball contains two critical points of f , one with an even index and one with an odd index, the pseudogradient vector field v can be extended over these balls to define a nowhere-zero vector field on M . As explained in [25], a nowhere-zero vector field defines a $\text{Spin}^{\mathbb{C}}$ structure. This is the $\text{Spin}^{\mathbb{C}}$ structure that is associated to the element \hat{v} . It follows from what is said in [25] that this $\text{Spin}^{\mathbb{C}}$ structure does not depend on the extension of v over the $G + 1$ balls.

Part 3 Fix $\hat{v} = \{v_1, \dots, v_G\} \in \mathcal{Z}_{\text{HF}}$. The set $\{v_1, \dots, v_G\}$ has a unique partition modulo ordering into a set of at most G irreducible index 1-2 cycles. Let \hat{O} denote this unordered set of index 1-2 cycles. A set of this sort is called an *HF cycle*. Any given HF cycle defines a corresponding partition of Λ . Let \hat{o} denote a given index 1-2 cycle from \hat{O} , and let $\Lambda_{\hat{o}} \subset \Lambda$ denote the corresponding subset. Let $\{\mathfrak{k}_p\}_{p \in \Lambda_{\hat{o}}}$ denote a given set of integers. Proposition 2.7 associates to the data $(\hat{o}, \{\mathfrak{k}_p\}_{p \in \Lambda_{\hat{o}}})$ a unique, closed integral curve of v . Use $\gamma_{\hat{o}}$ to denote the latter curve.

Part 4 Fix $p \in \Lambda$. Use $\hat{\gamma}_p^+$ denote the integral curve of v in \mathcal{H}_p that sits in the locus where $u = 0$ and $\cos \theta = \frac{1}{\sqrt{3}}$. Use $\hat{\gamma}_p^-$ to denote the integral curve of v in \mathcal{H}_p that sits where $u = 0$ and $\cos \theta = -\frac{1}{\sqrt{3}}$.

Introduce O to denote the 4-element set $\{0, 1, -1, (1, -1)\}$. Let O_p denote a given element in O . What follows associates an element in \mathcal{Z} to O_p . The latter is denoted by Θ_p .

- (2-18) • If $O_p = 0$, then $\Theta_p = \emptyset$.
- If $O_p = 1$ or -1 , then $\Theta_p = (\hat{\gamma}_p^+, 1)$ or $(\hat{\gamma}_p^-, 1)$ as the case may be.
 - If $O_p = (1, -1)$, then $\Theta_p = \{((\hat{\gamma}_p^+, 1), (\hat{\gamma}_p^-, 1))\}$.

What follows is the promised description of $\mathcal{Z}_{\text{ech}, M}$.

Proposition 2.8 *There exists $\kappa > 1$ such that if $\delta < \kappa^{-1}$ and $\chi_0 < \kappa^{-1}\delta^2$, then the following is true: Fix a $\text{Spin}^{\mathbb{C}}$ structure on M to define \mathcal{Z}_{HF} . Then the elements in $\mathcal{Z}_{\text{ech},M}$ with the $\text{Spin}^{\mathbb{C}}$ structure from (1-28) are in 1-to-1 correspondence with $\mathcal{Z}_{\text{HF}} \times \prod_{p \in \Lambda} (\mathbb{Z} \times \mathbb{O})$. This correspondence is such that a given element $(\hat{v}, \{\mathfrak{k}_p, \mathbb{O}_p\}_{p \in \Lambda}) \in \mathcal{Z}_{\text{HF}} \times \prod_{p \in \Lambda} (\mathbb{Z} \times \mathbb{O})$ corresponds to the element*

$$\Theta = \{(\gamma_{\hat{\sigma}}, 1)\}_{\hat{\sigma} \in \hat{\mathcal{O}}} \cup \left\{ \bigcup_{p \in \Lambda} \Theta_p \right\} \in \mathcal{Z}_{\text{ech},M}.$$

Proof The proof must establish that any collection Θ of the sort described in the proposition has the $\text{Spin}^{\mathbb{C}}$ structure that is described in (1-28), and it must establish that the proposition accounts for all elements in $\mathcal{Z}_{\text{ech},M}$.

Assume for the moment that the elements define the desired $\text{Spin}^{\mathbb{C}}$ structure to focus on the second aspect of the proof. According to (2-2), each closed integral curve of v from a given element in $\mathcal{Z}_{\text{ech},M}$ sits entirely in $f^{-1}(1, 2) \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$. Let $\Theta \in \mathcal{Z}_{\text{ech},M}$. Fix $p \in \Lambda$. It follows from Lemma 2.1 that Θ can contain either 0, 1 or 2 pairs of the form $(\gamma, 1)$ where γ is an integral curve of v that lies entirely in \mathcal{H}_p . Meanwhile, Lemma 2.2 with (1-29) and (1-30) has the following implication: there is one and only one segment of an integral curve of v that crosses \mathcal{H}_p from the $u = -R - \ln 7\delta_*$ end to the $u = R + \ln 7\delta_*$ end. This last observation, Proposition 2.7 and what was just said about the integral curves entirely inside $\bigcup_{p \in \Lambda} \mathcal{H}_p$ lead directly to Proposition 2.8.

Consider now the $\text{Spin}^{\mathbb{C}}$ structure issue. Given what is said in (1-30), the construction guarantees that the second and third bullets of (1-28) are obeyed. Note in this regard that the integral curves $\{\hat{\gamma}_p^+, \hat{\gamma}_p^-\}_{p \in \Lambda}$ are null-homologous as each lies in the $u = 0$ sphere of its labeling handle.

The verification of the top bullet of (1-28) has three steps.

Step 1 Fix an element \hat{v} from \mathcal{Z}_{HF} . The arcs from the set \hat{v} and the closure in M of the flow line of v through the point z_0 have pairwise disjoint neighborhoods with each being an embedded ball with 3-sphere boundary. Fix a set of balls of this sort so that the complement in M of their union is in M_δ . Label these balls as $\{B_\alpha\}_{\alpha=0,1,\dots,G}$, with B_0 containing the flow line of v through the point z_0 .

Let $Q_0 \subset Y$ denote the union of the ball B_0 with the handle \mathcal{H}_0 . This is diffeomorphic as a manifold with boundary to the complement of a ball in $S^2 \times S^1$. Meanwhile, let $Q \subset Y$ denote the union of the balls from the set $\{B_\alpha\}_{\alpha=1,\dots,G}$ and the handles from

the set $\{\mathcal{H}_p\}_{p \in \Lambda}$. This is a smooth manifold with boundary, the boundary being the G boundary spheres of the balls $\{B_\alpha\}_{\alpha=1, \dots, G}$. The number of components of Q is the number of irreducible index 1-2 cycles that are defined by \hat{v} . An irreducible, index 1-2 cycle from \hat{v} with length $N \geq 1$ gives a component of Q that is diffeomorphic as a manifold with boundary to the complement of N disjoint balls in $S^2 \times S^1$.

The $G + 1$ balls $\{B_\alpha\}_{\alpha=0, 1, \dots, G+1}$ in M can be used to depict Y as the surgery $(M \setminus (\bigcup_{\alpha=0, 1, \dots, G} B_\alpha)) \#_{\partial Q} (Q_0 \cup Q)$. Supposing that Θ is an element from $\mathcal{Z}_{\text{ech}, M}$ with \hat{v} as its \mathcal{Z}_{HF} component, then all of the integral curves of v that form Θ are in the Q part of Y .

Step 2 As explained in Part 2 of Section 1F, the $\text{Spin}^{\mathbb{C}}$ structure on Y that is defined by Θ is obtained from the canonical $\text{Spin}^{\mathbb{C}}$ structure by adding the cohomology class $\sum_{(\gamma, m) \in \Theta} m[\gamma]^{\text{Pd}}$. A $\text{Spin}^{\mathbb{C}}$ structure is also an equivalence of nowhere-zero vector fields on Y . Since the canonical $\text{Spin}^{\mathbb{C}}$ structure is defined by the vector field v (generating the kernel of the 2-form w), the nowhere-zero vector field for the $\text{Spin}^{\mathbb{C}}$ structure defined by Θ can be obtained by suitably modifying v in a small tubular neighborhood of the curves that form the set Θ . Since all of these curves are in the Q of the connect sum, it follows that the $\text{Spin}^{\mathbb{C}}$ structure on the $(M \setminus (\bigcup_{\alpha=0, 1, \dots, G} B_\alpha))$ part of the connect sum that defines Y is defined by the nowhere-zero vector field v . Since $M \setminus (\bigcup_{\alpha=0, 1, \dots, G} B_\alpha)$ is contained in M_δ , the vector field v on the $M \setminus (\bigcup_{\alpha=0, 1, \dots, G} B_\alpha)$ part of Y is the pseudogradient vector field \mathfrak{v} . This implies what is said by the top bullet in (1-28) because the $\text{Spin}^{\mathbb{C}}$ structure that is defined by the Heegaard Floer element \hat{v} from \mathcal{Z}_{HF} is also defined by \mathfrak{v} on $M \setminus (\bigcup_{\alpha=0, 1, \dots, G} B_\alpha)$.

Step 3 The second bullet in (1-28) follows from (1-29) and (1-30) because the Q_0 part of Y has no curves from Θ . The third bullet of (1-28) also follows from (1-29) and (1-30) since each component of Q has a single closed integral curve of v which has intersection number $+1$ with any cross-sectional 2-sphere in each of the relevant handles from the set $\{\mathcal{H}_p\}_{p \in \Lambda}$. □

3 Pseudoholomorphic subvarieties and the geometry of $\mathbb{R} \times Y$

This section takes a preliminary step towards describing the pseudoholomorphic subvarieties on $\mathbb{R} \times Y$ that are used to compute the differential and certain other endomorphisms of the $\mathcal{Z}_{\text{ech}, M}$ -generated embedded contact homology chain complex. The first subsection describes the sorts of almost complex structures that are considered here. The

remaining subsections describe a foliation of most of $\mathbb{R} \times Y$ by pseudoholomorphic submanifolds. The missing parts are the closed orbits of v that sit where $u = 0$ and $1 - 3 \cos^2 \theta = 0$ in the various $\mathfrak{p} \in \Lambda$ versions of $\mathcal{H}_{\mathfrak{p}}$. This foliation plays an absolutely central role in the subsequent analysis of the curves that are relevant to $\mathcal{Z}_{\text{ech}, M}$.

3A Almost complex structures and holomorphic subvarieties

The notion of pseudoholomorphic requires first the specification of an almost complex structure on $\mathbb{R} \times Y$. The almost complex structures that are considered in this section are constrained in certain ways. The constraints are described in Part 1 of this subsection. Additional constraints are imposed in Section 7 so as to make contact with the Heegaard Floer story. Part 2 of this subsection describes two particularly important features of the almost complex structures that obey the constraints of Part 1. Part 3 describes a topology for the set of pseudoholomorphic subvarieties.

Part 1 Let J denote a given almost complex structure. What follows lists the constraints that are imposed in this section on J .

Constraint 1 J maps the Euclidean tangent vector ∂_s to the \mathbb{R} factor of $\mathbb{R} \times Y$ to the vector field v .

Constraint 2 J is not changed by constant translations of the coordinate s on the \mathbb{R} factor of $\mathbb{R} \times Y$.

The statement of the third constraint refers to a certain 1-form, \hat{a} , on Y . The definition of \hat{a} involves the 1-form ν_{\diamond} that is defined in Section 1E. The definition also requires a certain function of the coordinate u on any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$. This function is denoted by χ_{δ} and it is defined by the rule

$$(3-1) \quad u \rightarrow \chi_{\delta}(u) = \chi(|u| - R - \ln \delta + 10).$$

The function χ_{δ} equals 0 where $|u| \geq R + \ln \delta - 9$ and it equals 1 where $|u| < R + \ln \delta - 10$. In particular, $\chi_{\delta} = 1$ where the function χ is nonzero. The 1-form \hat{a} is given by the rule

$$(3-2) \quad \hat{a} = \nu_{\diamond} \quad \text{on } M_{\delta} \cup \mathcal{H}_0 \quad \text{and} \quad \hat{a} = \chi_{\delta} a + (1 - \chi_{\delta}) \nu_{\diamond} \quad \text{on } \bigcup_{\mathfrak{p} \in \Lambda} \mathcal{H}_{\mathfrak{p}}.$$

What was said above about χ_{δ} guarantees that \hat{a} is smooth, that $\hat{a} \wedge w > 0$ and that \hat{a} has pairing 1 with v . That $\hat{a} \wedge w > 0$ implies that the kernel of \hat{a} defines a 2-dimensional subbundle in TY and that w restricts to this subbundle as a symplectic form.

Constraint 3 J preserves the kernel of the 1-form \hat{a} and as such, it defines the orientation given by w .

To give some indication of what this constraint does, note that since $\hat{a} = v_\diamond$ on M_δ and $v_\delta = df$ on M_δ , the almost complex structure J preserves the level sets of f on $\mathbb{R} \times M_\delta$. Since \hat{a} is proportional to du on \mathcal{H}_0 , this J also preserves the level sets of u on $\mathbb{R} \times \mathcal{H}_0$. Meanwhile, as $\hat{a} = a$ on the $|u| < R + \ln \delta - 10$ part if any given $p \in \Lambda$ version of \mathcal{H}_p , the almost complex structure J on $\mathbb{R} \times \mathcal{H}_p$ preserves the contact distribution given by the kernel of a .

The remaining constraints concern only the behavior of J on $\{\mathbb{R} \times \mathcal{H}_p\}_{p \in \Lambda}$.

Constraint 4 J on $\mathbb{R} \times \mathcal{H}_p$ is unchanged by $\mathbb{R}/2\pi\mathbb{Z}$ translations of the coordinate ϕ .

To set the background for the next constraint, use Properties 3 and 4 of Section 1E to see that the kernel of \hat{a} is spanned where $u \sin \theta$ and $(1 - 3 \cos^2 \theta) \sin \theta$ are not both zero by the vector fields

$$(3-3) \quad \begin{aligned} \bullet \quad e_1 &= -6g \cos \theta \sin \theta \partial_u + (\chi + g')(1 - 3 \cos^2 \theta) \partial_\theta, \\ \bullet \quad e_2 &= \partial_\phi + \sqrt{6} \chi_\delta f \cos \theta \sin^2 \theta v. \end{aligned}$$

(The function g is defined in (1-24) by the rule $g = (\chi_+ e^{2(u-R)} - \chi_- e^{-2(u+R)})$.) With the preceding understood, the last of the constraints in the section follows directly.

Constraint 5 J on $\mathbb{R} \times \mathcal{H}_p$ is such that $Je_1 = \sigma^{-1}e_2$ with σ a positive function of u and θ .

Note that the function σ is bounded near the locus where $u = 0$ and $1 - \cos^2 \theta = 0$, and is further constrained there so as to ensure that J is smooth. (This last constraint is not crucial to what follows here and in [12], but it does simplify some arguments.)

Assume henceforth that J obeys these constraints.

Part 2 The almost complex structure J has the following two properties:

- (3-4) \bullet The bilinear form $w(\cdot, J\cdot)$ is positive definite on the kernel of the 1-form a .
- \bullet There exists $r \geq 1$ such that the 2-form $\omega = ds \wedge a + rw$ tames J .

To see about the top bullet, note first that Constraint 3 in Part 1 means that the bilinear form $w(\cdot, J\cdot)$ restricts to the kernel of \hat{a} as a positive-definite inner product. With

this in mind, write a given vector $e \in \text{kernel}(a)$ as $e = xv + \hat{e}$ where $x \in \mathbb{R}$ and $\hat{e} \in \text{kernel}(\hat{a})$. The fact that $a \wedge w > 0$ implies that $|x| \geq c_0 \omega(\hat{e}, J\hat{e})^{1/2}$ and so $e = 0$ if $\hat{e} = 0$. Meanwhile, $\omega(e, Je)$ is equal to $\omega(\hat{e}, J\hat{e})$ because v generates the kernel of w . To see about the second bullet, write a given tangent vector as $q = x_1 \partial_s + x_2 v + \hat{e}$ where $\hat{e} \in \text{kernel}(\hat{e})$. Then

$$(3-5) \quad \omega(q, Jq) = x_1^2 + x_2^2 + x_1 a(J\hat{e}) - x_2 a(\hat{e}) + r\omega(\hat{e}, J\hat{e}).$$

This is positive if and only if $q \neq 0$ when $r \geq c_0$ because $\omega(\cdot, J\cdot)$ is positive definite on the kernel of \hat{a} .

Note for reference below that J is compatible with the 2-form

$$(3-6) \quad \hat{\omega} = ds \wedge \hat{a} + w,$$

in the sense that the bilinear form $\hat{\omega}(\cdot, J\cdot)$ defines a Riemannian metric on $T(\mathbb{R} \times Y)$. This metric is used implicitly in what follows to measure distances and to define the norms of vectors, covectors and other tensors.

Part 3 Given that J maps ∂_s to v , the properties in (3-4) imply that the almost complex structure J obeys the conditions that are assumed in Part 4 of Section 1F. This being the case, (1-31) defines what is meant by a J -holomorphic subvariety. The set of J -holomorphic subvarieties is given the topology whereby a basis of open neighborhoods of a given subvariety C consists of sets that are indexed by a pair of positive number and smooth, compactly supported 2-form on $\mathbb{R} \times Y$. If (ε, ϖ) is such a pair, then a J -holomorphic subvariety C' is in the corresponding open neighborhood when

$$(3-7) \quad \begin{aligned} & \bullet \sup_{z \in C} \text{dist}(z, C') + \sup_{z \in C'} \text{dist}(z, C) < \varepsilon, \\ & \bullet \left| \int_C \varpi - \int_{C'} \varpi \right| < \varepsilon. \end{aligned}$$

The resulting topological space is called the *moduli space* of J -holomorphic subvarieties.

The group \mathbb{R} acts as a continuous group of homeomorphisms on the moduli space of J -holomorphic subvarieties. The action of the generator $1 \in \mathbb{R}$ on a given subvariety C translates C by 1 along the \mathbb{R} factor in $\mathbb{R} \times Y$. Let Z denote a given space with an action of \mathbb{R} . A given map from a moduli space component to Z or from Z to a component is said to be \mathbb{R} -equivariant when it intertwines the given \mathbb{R} action on Z with this particular \mathbb{R} action on the moduli space of J -holomorphic subvarieties.

3B Level sets

The almost complex structure J preserves the kernel of the 1-form ν_\diamond on $M_\delta \cup \mathcal{H}_0$. Meanwhile, the kernel of ν_δ is tangent to a surface either with or without boundary. These surface are the level sets in M_δ and the level sets of the coordinate u in \mathcal{H}_0 . Those with boundary are the level sets of f with $f \in [1 - 2\delta_*^2, 1 + \delta_*^2]$ or $f \in [2 - \delta_*^2, 2 + \delta_*^2]$. The surfaces without boundary are J -holomorphic submanifolds as defined by (1-31). As noted in the propositions that follow, those with boundary are subsets of J -holomorphic submanifolds that are described by (1-31). These propositions also indicate that there are important distinctions to be made depending on the following three alternatives: either the level set in question lies in the union of \mathcal{H}_0 with the $f \notin [1, 2]$ part of M_δ ; or it is a level set in M_δ where $f \in (1, 2)$; or it is either the $f = 1$ or $f = 2$ level set in M_δ . The proofs of the propositions are given at the end of the subsection.

The first proposition talks about the $f \notin [1, 2]$ case.

Proposition 3.1 *There exists a component, \mathcal{M}_0 , of the moduli space of J -holomorphic subvarieties which has the properties listed below:*

- *Each subvariety from \mathcal{M}_0 is an embedded 2-sphere in $\mathbb{R} \times Y$.*
- *\mathcal{M}_0 has the structure of a smooth manifold which is \mathbb{R} -equivariantly diffeomorphic to $\mathbb{R} \times (-1, 1)$ with the action on the latter space given by the constant translations along the \mathbb{R} factor. Such a diffeomorphism can be chosen so that:*
 - (a) *Suppose $(x, y) \in \mathbb{R} \times [-\frac{1}{2}, \frac{1}{2}]$. The corresponding submanifold is $\{x\} \times S_y$, where $S_y \subset \mathcal{H}_0$ is the level set $u = 2(R + \ln \delta)y$.*
 - (b) *Suppose $(x, y) \in \mathbb{R} \times [\frac{1}{2}, 1)$. The corresponding submanifold intersects the subset $\mathbb{R} \times M_\delta$ as $\{x\} \times S_y$, where S_y is the $f = \delta_2 + (1 - \delta^2)(2y - 1)$ level set.*
 - (c) *Suppose $(x, y) \in \mathbb{R} \times (-1, -\frac{1}{2}]$. The corresponding submanifold intersects the subset $\mathbb{R} \times M_\delta$ as $\{x\} \times S_y$, where S_y is the $f = 3 - \delta^2 - (1 - \delta^2)(2y - 1)$ level set.*

The next proposition talks about the $f \in (1, 2)$ level sets.

Proposition 3.2 *There exists a component, \mathcal{M}_Σ , of the moduli space of J -holomorphic subvarieties which has the properties listed below:*

- *Each subvariety from \mathcal{M}_Σ is an embedded surface of genus G in $\mathbb{R} \times Y$.*

- \mathcal{M}_Σ has the structure of a smooth manifold which is \mathbb{R} -equivariantly diffeomorphic to $\mathbb{R} \times (1, 2)$ with the action on the latter space defined to be the translations along the \mathbb{R} factor. Such a diffeomorphism can be chosen so that if $(x, y) \in \mathbb{R} \times (1, 2)$, then the corresponding manifold intersects $\mathbb{R} \times M_\delta$ as $\{x\} \times S_y$, where S_y is the $f = y$ level set.

The final proposition in this subsection concerns the $f = 1$ and $f = 2$ level sets. These subvarieties have ends. When talking about the ends of these subvarieties, the proposition refers to the integral curve of v and positive integer that is associated to a given end of a given J -holomorphic subvariety.

Proposition 3.3 *There exist components, \mathcal{M}_1 and \mathcal{M}_2 , of the moduli space of J -holomorphic subvarieties which have the properties listed below:*

- Each subvariety from either \mathcal{M}_1 or \mathcal{M}_2 is a properly embedded, $2G$ -times punctured sphere in $\mathbb{R} \times Y$.
- Each such submanifold from \mathcal{M}_1 or \mathcal{M}_2 has $2G$ negative ends and no positive ends.
- The ends of each such subvariety are labeled by the set $\{(p, \varepsilon)\}_{p \in \Lambda, \varepsilon \in \{-, +\}}$. This labeling is such that the integral curve of v that is associated to the (p, ε) labeled end is the curve $\hat{\gamma}_p^\varepsilon$, this the locus in \mathcal{H}_p where $u = 0$ and $\cos \theta = \varepsilon \frac{1}{\sqrt{3}}$. The associated integer in each case is equal to 1.
- \mathcal{M}_1 and \mathcal{M}_2 have the structure of a smooth manifold that is \mathbb{R} -equivariantly diffeomorphic to the group \mathbb{R} . These diffeomorphisms can be chosen so that:
 - (a) The $x \in \mathbb{R}$ labeled submanifold in \mathcal{M}_1 intersects $\mathbb{R} \times M_\delta$ as $\{x\} \times f^{-1}(1)$.
 - (b) The $x \in \mathbb{R}$ labeled submanifold in \mathcal{M}_2 intersects $\mathbb{R} \times M_\delta$ as $\{x\} \times f^{-1}(2)$.

The remainder of this subsection is occupied with the:

Proof of Propositions 3.1–3.2 The proof of these propositions is given in six steps. The first five steps describe the J -holomorphic subvarieties that restrict to $\mathbb{R} \times M_\delta$ as $\{x\} \times S$ with $S \subset M_\delta$ a level set of f with f in $[1 - 2\delta^2, 1 + \delta^2]$ or with f in $[2 - \delta^2, 2 + \delta^2]$. Note that the arguments are very similar to those used in Section 4a of [21].

Step 1 Fix $p \in \Lambda$. Let $S \subset \mathcal{H}_p$ denote a closed, connected submanifold which is invariant under rotations of the angle ϕ . Suppose in addition that the tangent space

of S is mapped to itself by J . The fact that S is ϕ -invariant implies that ∂_ϕ is tangent to S . Use (3-3) to write

$$(3-8) \quad \partial_\phi = e_2 - \sqrt{6} \chi_\delta f \cos \theta \sin^2 \theta v,$$

and then use the fifth constraint in Part 1 to see that

$$(3-9) \quad J \cdot \partial_\phi = -\sigma e_1 + \sqrt{6} \chi_\delta f \cos \theta \sin^2 \theta \partial_s.$$

As J maps TS to itself, the vector field $J \cdot \partial_\phi$ is also tangent to S . Use (3-3) and (3-9) to restate this fact as follows: the vector field $J \cdot \partial_\phi$ is $-\sigma$ times

$$(3-10) \quad 6g \cos \theta \sin \theta \partial_u - (\chi + g')(1 - 3 \cos^2 \theta) \partial_\theta + \sigma^{-1} \sqrt{6} \chi_\delta f \cos \theta \sin^2 \theta \partial_s,$$

with g as defined in (1-24). As such, the integral curves of this vector field foliate the complement in S of the fixed points of the S^1 action that rotates the angle ϕ , and any two such curves are related by a rotation of the angle ϕ .

Step 2 Let S be as described in Step 1 and let $\nu \subset S$ denote an integral curve of the vector field in (3-10). The function $u \rightarrow g(u)$ vanishes where $u = 0$ and this implies that u is either strictly positive, strictly negative or identically zero on S . The case where $u = 0$ on S is relevant to what is said in the next subsection and is not discussed further here.

The coefficient of ∂_θ in (3-10) vanishes where $1 - 3 \cos^2 \theta = 0$. This implies that the function $1 - 3 \cos^2 \theta$ is either strictly negative, strictly positive, or identically zero on S . These different cases are considered respectively in Steps 3, 4 and 5. What is said in these steps explains why these cases respectively concern Propositions 3.1, 3.2 and 3.3.

Step 3 This step assumes that $1 - 3 \cos^2 \theta < 0$ on S . It follows as a consequence that $\cos \theta$ is nowhere zero on S . Assume here that $\cos \theta > 0$. Suppose that $g(u)$ and thus u are positive on S . If this is the case, then θ , u and s all decrease when traveling backwards along the curve ν . As a consequence, the curve ν must end where $\cos \theta = 1$, which is by necessity a fixed point of ϕ on S . This fixed point is the point on S where both u and s take on their minimum values.

Moving in the oriented direction on ν out of this point increases u , θ and s . The function u increases faster and faster relative to θ as $1 - 3 \cos^2 \theta$ increases towards zero. As a consequence, u is eventually greater than $R + \ln \delta$ on ν and this part of ν is in the radius $7\delta_*$ coordinate ball centered on p 's index 1 critical point. Noting that S

is an orbit of ν with respect to the S^1 action generated by constant translations of ϕ , it follows that S is an embedded disk in $\mathbb{R} \times \mathcal{H}_p$.

The function χ_δ is zero where $u > R + \ln \delta$ on ν , and therefore s is constant on this part of ν . As $J \cdot \partial_s = \nu$, such a slice must be a level set of f and as $1 - 3 \cos^2 \theta < 0$, this level set is one where $f < 1$. This implies that S seamlessly extends some submanifold in $\mathbb{R} \times M_\delta$ of the form $\{x\} \times S_y$ as described in Proposition 3.1(b) by attaching a disk to its $\cos \theta > 0$ boundary component on the boundary of the radius δ coordinate ball centered on p 's index 1 critical point. The case where $\cos \theta < 0$ on S has the analogous conclusion and is proved by essentially identical arguments.

Let $\{x\} \times S_y \subset \mathbb{R} \times M_\delta$ denote a surface with S_y being an $f \in [1 - 2\delta^2, 1)$ level set. Integrate (3-10) from any point on its boundary components on the radius δ coordinate ball centered on p 's index 1 critical point to obtain two disks as just described that extends $\{x\} \times S_y$ into $\mathbb{R} \times \mathcal{H}_p$.

Much the same argument applies to the case where $u < 0$ on S . In this case, S extends some submanifold in $\mathbb{R} \times M_\delta$ of the form $\{x\} \times S_y$ with S_y as described in Proposition 3.1(c).

Apply the preceding construction for each $p \in \Lambda$ and for a given surface in $\mathbb{R} \times Y$ of the form $\{x\} \times S_y$ with $S_y \subset M_\delta$ an $f \in [1 - 2\delta^2, 1)$ or $f \in (2, 2 - 2\delta^2]$ level set. The result is a compact, J -holomorphic submanifold in $\mathbb{R} \times Y$ that is diffeomorphic to S^2 .

Step 4 This step assumes that $1 - 3 \cos^2 \theta > 0$ on S . Consider first the case where $u > 0$ and so $g > 0$. Start on ν at a point where $\cos \theta < 0$. The coefficient that multiplies ∂_θ in (3-10) is negative and so θ decreases along ν . Meanwhile, both s and u also decrease. Eventually, the $\theta = \frac{\pi}{2}$ point is reached, and then $\cos \theta$ turns positive. Both u and s start to increase along ν and do so from then on. Thus, the minimum of both u and s occur at the $\theta = \frac{\pi}{2}$ point on ν . As $\cos \theta$ approaches $\frac{1}{\sqrt{3}}$, the rate of increase of u relative to θ becomes very large, and so eventually $u \geq R + \ln \delta$. Here χ_δ is zero, so s is constant, as is f . Rotating ν by changing ϕ shows that the $u \geq R + \ln \delta$ part of S where $\cos \theta > 0$ is an annular neighborhood of a boundary component of a surface of the form $\{x\} \times S_y \subset \mathbb{R} \times M_\delta$ where S_y is some $y \in (1, 1 + \delta^2)$ level set of f . Going in the opposite direction on ν from the $\cos \theta < 0$ start point has u increasing and s also increasing, so the same story applies to the part of S where $u \geq R + \ln \delta$ and $\cos \theta < 0$. Note that the same surface S_y appears here because changing θ to $\pi - \theta$ changes the sign of all terms in (3-10). Granted this, then attaching the $u < R + \ln \delta$

part of S to $\{x\} \times S_y$ seamlessly extends the latter by attaching a cylinder to its two boundary components on the radius δ coordinate ball centered on the index 1 critical point from \mathfrak{p} .

Let $(x, y) \in \mathbb{R} \times (1, 1 + \delta^2]$ and let $\{x\} \times S_y$ denote the surface in $\mathbb{R} \times M_\delta$ with S_y the level set of $f = y$. Integrate (3-10) from any point on its $\cos \theta < 0$ boundary on the radius δ coordinate ball centered on \mathfrak{p} 's index 1 critical point to obtain a cylinder S as just described that extends $\{x\} \times S_y$ into $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$.

Much the same argument applies in the case where $u < 0$ on ν . Here, the cylinder S extends a surface of the form $\{x\} \times S_y$ with $y \in [2 - \delta^2, 2)$ and with S_y again denoting the $f = y$ level set.

Apply the preceding construction for each $\mathfrak{p} \in \Lambda$ and for a given surface in $\mathbb{R} \times Y$ of the form $\{x\} \times S_y$ with $S_y \subset M_\delta$ an $f \in (1, 1 + \delta^2]$ or $f \in [2 - \delta^2, 2)$ level set. The result is a compact, J -holomorphic submanifold in $\mathbb{R} \times Y$ that is diffeomorphic to Σ .

Step 5 Consider here the case where $1 - 3 \cos^2 \theta = 0$ on S . Suppose that $u > 0$ and that $\cos \theta > 0$. Both u and s decrease as ν is traveled in reverse. The rate of decrease of u as it nears zero is proportional to u , while the rate of change of s remains bounded away from zero. As a consequence, s diverges towards $-\infty$ as $u \rightarrow 0$ on ν . Moving in the forward direction on ν increases u and eventually $u \geq R + \ln \delta$ on ν . Here, s is constant as is f . As S is invariant under translations of the coordinate ϕ , this implies that the part of S where $u \geq R + \ln \delta$ is an annular neighborhood of a boundary component of a surface in $\mathbb{R} \times M_\delta$ of the form $\{x\} \times S_y$ where S_y is the $f = 1$ level set in M_δ . In particular, the part of S where $u \leq R + \ln \delta$ seamlessly attaches to this boundary component to extend the surface $\{x\} \times S_y$. Integrating (3-10) will extend any surface of this sort into the $\cos \theta > 0$ part of $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$. A very similar construction extends the $\cos \theta = -\frac{1}{\sqrt{3}}$ part of the intersection of such a surface with the radius δ coordinate ball centered on the index 1 critical point of f from \mathfrak{p} .

But for cosmetics, the same construction applies in the case where $u < 0$ on ν . Here, the cylinder S extends a surface of the form $\{x\} \times S_y$ with S_y the $f = 2$ level set, and any surface of this sort can be so extended by integrating (3-10).

Apply the preceding construction for each $\mathfrak{p} \in \Lambda$ and for a given surface in $\mathbb{R} \times Y$ of the form $\{x\} \times S_y$ with $S_y \subset M_\delta$ an $f = 1$ or $f = 2$ level set. The result is a compact, J -holomorphic submanifold in $\mathbb{R} \times Y$ that is diffeomorphic to a sphere with $2G$ punctures. Such a submanifold has $2G$ ends, all negative ends, and with the $s \ll -1$ asymptotics as described by Proposition 3.3.

Note that the construction just completed supplies all of the J -holomorphic subvarieties that intersect $\mathbb{R} \times M_\delta$ in an $f = 1$ or $f = 2$ level set. The argument for this is as follows: Suppose that C is a J -holomorphic subvariety in $\mathbb{R} \times Y$ such that $C \cap (\mathbb{R} \times M_\delta)$ consists only of points where $f = 1$ or $f = 2$. Suppose that p is such a point and suppose for the sake of argument that $f = 1$ at p . Let x_p denote the \mathbb{R} coordinate of p . If C does not coincide with the subvariety labeled by x_p in \mathcal{M}_1 , then p will be an isolated intersection point between this subvariety and C and so there is a disk in the x_p member of \mathcal{M}_1 with center p that intersects C only at p and its such that its boundary has some positive distance δ from C . Denote this disk by D . Since local intersections between J -holomorphic varieties are positive, any J -holomorphic deformation of D that is sufficiently close to D will also intersect C . This implies in particular that C must intersect elements from the space \mathcal{M}_Σ of Proposition 3.2 that are labeled by $x \in \mathbb{R}$ near x_p with f values near 1. This contradicts the initial assumption about C intersecting $\mathbb{R} \times M_\delta$ only where $f = 1$ or $f = 2$.

Step 6 This step completes the proof of Proposition 3.1. But for notation, the arguments given below complete the proof of Propositions 3.2 and 3.3. This being the case, the final arguments for Propositions 3.2 and 3.3 are not supplied.

Proposition 3.1 parametrizes its subvarieties by $\mathbb{R} \times (-1, 1)$. It is a straightforward task to verify that this parametrization is continuous with respect to the topology defined in Part 3 of Section 3A. The parametrization is \mathbb{R} -equivariant because the coefficients in (3-10) do not depend on the \mathbb{R} factor of $\mathbb{R} \times Y$. As explained in a moment, this homeomorphism of $\mathbb{R} \times (-1, 1)$ into the moduli space of J -holomorphic subvarieties is onto a component of this moduli space. Granted that such is the case, the smooth structure on this component is defined by that on $\mathbb{R} \times (-1, 1)$.

To see about completeness, suppose that $\{C_k\}_{k=1,2,\dots}$ is a sequence of J -holomorphic subvarieties in $\mathbb{R} \times Y$ that converges to a subvariety C in \mathcal{M}_0 . Since all submanifolds from \mathcal{M}_0 are compact, this must also be the case for all large k versions of C_k . If C_k is not in \mathcal{M}_0 , it must either intersect some subvariety on \mathcal{M}_0 , or else it lies entirely in the union of the $f \in [1, 2]$ part of $\mathbb{R} \times M_\delta$ with the $1 - 3 \cos^2 \theta \leq 0$ part of each $p \in \Lambda$ version of \mathcal{H}_p . If it intersects some subvariety in \mathcal{M}_0 , then it must intersect all because J -holomorphic subvarieties intersect with strictly positive local intersection number (see eg [3] or [14]). This can't happen if C_k is compact because there are spheres in \mathcal{M}_0 with any given \mathbb{R} coordinate. If it intersects the $f \in (1, 2)$ part of $\mathbb{R} \times M_\delta$, then it must intersect all of the subvarieties that are constructed in Step 4, and so it

again cannot be compact. If it intersects $\mathbb{R} \times M_\delta$ as the $f = 1$ or $f = 2$ level set, then it is one of the subvarieties constructed in Step 5 and so it cannot be compact.

By way of a converse, suppose that $\{C_k\}_{k=1,2,\dots} \subset \mathcal{M}_0$ is a sequence that converges to some J -holomorphic subvariety C . If C is not in \mathcal{M}_0 , then C must intersect $\mathbb{R} \times M_\delta$ as either the $f = 1$ or $f = 2$ level set. Thus, C is one of the subvarieties that are constructed in Step 5. Since C is noncompact, the sort of convergence required by the first bullet in (3-7) is not possible. \square

3C Subvarieties that lie entirely in $\mathbb{R} \times \mathcal{H}_p$

The proposition that follows describes the J -holomorphic subvarieties that sit entirely in some $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$.

Proposition 3.4 *Fix $p \in \Lambda$. A J -holomorphic subvariety that lies entirely in $\mathbb{R} \times \mathcal{H}_p$ sits in the $u = 0$ slice of $\mathbb{R} \times \mathcal{H}_p$. These subvarieties comprise three components of the moduli space of J -holomorphic subvarieties. The latter are denoted by \mathcal{M}_{p+} , \mathcal{M}_{p-} and \mathcal{M}_{p_0} . These components and their constituent subvarieties are described below:*

- *The component \mathcal{M}_{p+} : this moduli space component is smooth manifold that is equivariantly diffeomorphic to \mathbb{R} .*
 - (1) *Each element in \mathcal{M}_{p+} is an embedded, open disk that lies where $\cos \theta > \frac{1}{\sqrt{3}}$.*
 - (2) *The end of each element is positive. The associated closed integral curve of v is the curve $\hat{\gamma}_p^+$, this the $(u = 0, \cos \theta = \frac{1}{\sqrt{3}})$ locus in \mathcal{H}_p . The associated integer is 1.*
 - (3) *The function s has a single critical point on each element; this is a nondegenerate critical point where s is smallest. It occurs where $\theta = 0$.*
 - (4) *The parametrization by \mathbb{R} can be chosen so that the parameter gives the minimum value of s .*

- *The component \mathcal{M}_{p-} : this moduli space component is smooth manifold that is equivariantly diffeomorphic to \mathbb{R} .*
 - (1) *Each element in \mathcal{M}_{p+} is an embedded, open disk that lies where $\cos \theta < -\frac{1}{\sqrt{3}}$.*
 - (2) *The end of each element is positive. The associated closed integral curve of v is the curve $\hat{\gamma}_p^-$, this the $(u = 0, \cos \theta = -\frac{1}{\sqrt{3}})$ locus in \mathcal{H}_p . The associated integer is 1.*

- (3) The function s has a single critical point on each element; this is a nondegenerate critical point where s is smallest. It occurs where $\theta = \pi$.
- (4) The parametrization by \mathbb{R} can be chosen so that the parameter gives the minimum value of s .
- The component \mathcal{M}_{p_0} : this moduli space component is a smooth manifold that is equivariantly diffeomorphic to \mathbb{R} .
 - (1) Each element is an embedded cylinder that lies where $1 - 3 \cos^2 \theta > 0$.
 - (2) The ends of each element are positive. The curves $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ are their respective associated closed integral curves of v . The associated integer is 1 for both ends.
 - (3) The function s has a single critical value on each element, the circle where $\theta = \frac{\pi}{2}$.
 - (4) The parametrization by \mathbb{R} can be chosen so that the parameter gives the minimum value of s .

Proof Except for some minor changes, the arguments are the same as those used in the previous subsection to prove Propositions 3.1–3.3. In fact, the analysis of the J -holomorphic subvarieties that sit where $u = 0$ is simpler by virtue of the fact that the vector field in (3-10) can be replaced where $u = 0$ by the vector field

$$(3-11) \quad -(1 - 3 \cos^2 \theta) \partial_\theta + \sigma^{-1} \sqrt{6} \cos \theta \sin^2 \theta \partial_s.$$

The reader can also see a very similar argument used in [21, Section 4a, Example 2]. \square

3D Compact J -holomorphic subvarieties

The following lemma summarizes the contents of this last subsection.

Lemma 3.5 *The moduli space components \mathcal{M}_0 and \mathcal{M}_Σ contain the only irreducible, compact J -holomorphic subvarieties in $\mathbb{R} \times Y$.*

Proof To see why this is, suppose to the contrary that C is an irreducible, compact J -holomorphic subvariety that is not in \mathcal{M}_0 nor \mathcal{M}_Σ . This subvariety cannot be in \mathcal{M}_1 nor can it be in \mathcal{M}_0 , nor any $p \in \Lambda$ version $\mathcal{M}_{p\pm,0}$ as all of the latter contain only noncompact subvarieties. This being the case, C must intersect some submanifold from \mathcal{M}_0 or from \mathcal{M}_Σ . Given that J -holomorphic subvarieties intersect with strictly positive local intersection number, and given that all manifolds involved are compact,

C must intersect all subvarieties from \mathcal{M}_0 if it intersects just one. Likewise, it must intersect all subvarieties from M_Σ if it intersects just one. This can happen only if $|s|$ is unbounded on C . □

3E Fredholm and ech indices

This section makes some parenthetical remarks (for now) about the subvarieties that are described by Propositions 3.1–3.4. These remarks concern two integers that Hutchings associates to a J –holomorphic subvariety in $\mathbb{R} \times Y$. The first is the *ech index*; it is from Definition 2.14 of [7]. As explained in the appendix, subvarieties with ech index 1 are used to define the differential for embedded contact homology. The second integer is the *Fredholm index*. This is given in Definition 4.12 of [7]. The Fredholm index is the difference between the kernel dimension and the cokernel dimension of a certain differential operator that is canonically associated to the J –holomorphic subvariety in $\mathbb{R} \times Y$. If the subvariety is embedded, the operator in question is obtained from the $\bar{\partial}$ operator acting on sections of the normal bundle by adding an \mathbb{R} –linear, zeroth-order term. To say more, note that J along C gives C the structure of a holomorphic surface and thus a $\bar{\partial}$ operator. Meanwhile, J with the 2–form $ds \wedge a + \omega$ gives C ’s normal bundle the structure of a holomorphic, hermitian line bundle. (The normal bundle is denoted by N in what follows.) This data associated to C and N supplies a $\bar{\partial}$ operator that maps sections of N to sections of $N \otimes T^{0,1}C$. The differential operator of interest is defined jointly by $\bar{\partial}$ and a pair, v and μ , of section of $T^{0,1}C$ and section of $\bar{N}^2 \otimes T^{-0,1}C$. The desired operator sends any given section \mathfrak{s} of N to the section

$$(3-12) \quad \bar{\partial}\mathfrak{s} + v\mathfrak{s} + \mu\bar{\mathfrak{s}}.$$

The pair (v, μ) are defined by the 1–jet of J in the directions normal to C (see Part 2 of Section 6C in this regard.) The operator in (3-12) extends to map Sobolev class L^2_1 sections of N (the section and its covariant derivatives to first order are square integrable) to L^2 sections of N . It is Fredholm in the latter incarnation and Hutchings’ Fredholm index is the difference between the dimension of the kernel and the dimension of the cokernel. (The cokernel of this Fredholm version of (3-12) are the smooth sections of $N \otimes T^{0,1}C$ that are annihilated by the formal L^2 adjoint operator.) The subvariety is said to be *obstructed* when the cokernel is not 0. If the cokernel is 0, then the subvariety is said to be unobstructed.

Lemma 3.6 *The following lists the ech index and Fredholm indices of the subvarieties in Propositions 3.1–3.4. The list also says whether or not the subvarieties are obstructed.*

- Each subvariety in the space \mathcal{M}_0 of Proposition 3.1 has ech index equal to 2 and Fredholm index equal to 2. Each is unobstructed.
- Each subvariety in the space \mathcal{M}_Σ of Proposition 3.2 has Fredholm index $2 - 2G$ and each is obstructed.
- Each subvariety in the spaces \mathcal{M}_1 and \mathcal{M}_2 of Proposition 3.3 has Fredholm index and ech index equal to $2 - 2G$ and each is obstructed.
- Each subvariety from any $\mathfrak{p} \in \Lambda$ version of the spaces $\mathcal{M}_{\mathfrak{p}_+}$ and $\mathcal{M}_{\mathfrak{p}_-}$ of Proposition 3.4 has ech index 1 and Fredholm index 1. It is also unobstructed.
- Each subvariety from any $\mathfrak{p} \in \Lambda$ version of the space $\mathcal{M}_{\mathfrak{p}_0}$ of Proposition 3.4 has ech index 0 and Fredholm index 0. It is obstructed.

Proof The ech and Fredholm indices are computed using the formulas in Definitions 2.14 and 4.12 of [7]. (See also Theorem 4.15 in [7].) If the Fredholm index is negative (as is the case with \mathcal{M}_Σ , \mathcal{M}_1 and \mathcal{M}_2), then there must be nontrivial cokernel and so the subvariety is obstructed by definition. If the Fredholm index is zero, as is the case with the subvarieties from $\{\mathcal{M}_{\mathfrak{p}_0}\}_{\mathfrak{p} \in \Lambda}$, then the existence of a nontrivial element in the kernel requires a nontrivial element in the cokernel, so the subvariety is obstructed. The kernel in this case is nontrivial for each subvariety because it is in a 1-dimensional space of subvarieties.

The proof that the subvarieties in \mathcal{M}_0 are unobstructed can be had along the following lines: Since the Fredholm index is 2, it suffices to prove that the dimension of the kernel of the operator in (3-12) is also 2. The fact that \mathcal{M}_0 is 2-dimensional implies that the kernel of the operator in (3-12) has dimension at least 2. If the kernel had dimension greater than 2, then at least one nontrivial element in the kernel would have a zero at some point in the subvariety. The fact that this section is annihilated by the operator in (3-12) implies that all of its zeros have positive local degree. (The zeros of an element in the kernel of the operator in (3-12) can be shown to be isolated, and to have finite degree. Granted this, Taylor's expansion with remainder and the vanishing of (3-12) can be used to prove the following: The section near a zero looks like a nonzero multiple of $z^p + \mathcal{O}(|z|^{p+1})$ where $p \in \{1, 2, \dots\}$ and where z is a local holomorphic parameter for the curve.) The sum of the local degrees of a section is equal to the Euler class of the normal bundle, but this can't be positive in the case at hand because the normal bundle is isomorphic to the product bundle.

A proof that the subvarieties in each $\mathfrak{p} \in \Lambda$ version of either $\mathcal{M}_{\mathfrak{p}_+}$ or $\mathcal{M}_{\mathfrak{p}_-}$ can be made along the following lines: Fix $\mathfrak{p} \in \Lambda$ and let C denote one of the subvarieties in

\mathcal{M}_{p_+} or \mathcal{M}_{p_-} . The subvariety C is fixed by the constant rotations of the φ coordinate on the $u = 0$ slice of \mathcal{H}_p . Since J and w and a are also invariant, it follows that the operator in (3-12) is invariant under this $SO(2)$ action, and thus so is its formal L^2 adjoint. This implies that the cokernel can be written as a direct sum of character eigenspaces for the $SO(2)$ action. Meanwhile, the assertion that the formal L^2 adjoint of the operator in (3-12) annihilates an element in an $SO(2)$ character subspace in $C^\infty(C; N \otimes T^{1,0}C)$ implies that the element in question obeys a first-order ordinary differential equation. This equation can be integrated to see that there are no square integrable solutions. \square

4 Ech-HF subvarieties

Let ϑ denote as in Part 5 of Section 1F a finite collection of sets of the form (C, n) such that C is an irreducible, J -holomorphic subvariety and n is a positive integer. Part 5 of Section 1F explains how to associate to ϑ two elements, Θ_{ϑ_+} and Θ_{ϑ_-} , in \mathcal{Z}_{ech} . Of primary interest in what follows are those ϑ where Θ_{ϑ_+} and Θ_{ϑ_-} are in $\mathcal{Z}_{ech, M}$. Note that all pairs in such a set ϑ whose subvariety component is noncompact have integer component equal to 1. This is because all pairs in $\mathcal{Z}_{ech, M}$ have integer component 1. This is also why distinct closed integral curves of v are associated to distinct ends from the union of the subvarieties in ϑ , and why each end from this union has associated integer 1.

If all integer components of ϑ are 1, then ϑ can be viewed as the set of irreducible components of $\bigcup_{(C,1) \in \vartheta} C$. This view is used in what follows. This view also motivates the following definition: A J -holomorphic subvariety whose irreducible components define a set ϑ with Θ_{ϑ_+} and Θ_{ϑ_-} in $\mathcal{Z}_{ech, M}$ is said to be an *ech subvariety*. When C denotes an ech subvariety, the corresponding sets in $\mathcal{Z}_{ech, M}$ are henceforth labeled as Θ_{C_+} and Θ_{C_-} . An ech subvariety is said to be an *ech-HF subvariety* if it lacks irreducible components that are described by Propositions 3.1–3.4.

This section begins the task of describing the ech-HF subvarieties.

4A Projections to \mathcal{H}_0 , the $f \notin (1, 2)$ part of M_δ and the $1 - 3 \cos^2 \theta \leq 0$ part of \mathcal{H}_p

The following proposition summarizes the conclusions of this subsection:

Proposition 4.1 *Let $C \subset \mathbb{R} \times Y$ denote an ech-HF subvariety. Then:*

- C has no intersection with $\mathbb{R} \times \mathcal{H}_0$ and with the parts of $\mathbb{R} \times M_\delta$ where $f \leq 1$ or $f \geq 2$.
- C has no intersection with the $1 - 3 \cos^2 \theta < 0$ part of any $\mathfrak{p} \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$.
- An irreducible component of C that intersects the locus in $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$ where $1 - 3 \cos^2 \theta = 0$ does so at $u = 0$. This component is the cylinder $\mathbb{R} \times \gamma$ where $\gamma \subset \mathcal{H}_\mathfrak{p}$ is the relevant $u = 0$ and $1 - 3 \cos^2 \theta = 0$ closed integral curve of v ; thus either $\widehat{\gamma}_\mathfrak{p}^+$ or $\widehat{\gamma}_\mathfrak{p}^-$.

Proof If the projection of C to Y contains a point in \mathcal{H}_0 or in M_δ where $f < 1$ or $f > 2$, then C must intersect a J -holomorphic subvariety in \mathcal{M}_0 . Such an intersection has locally positive intersection number, and this means that C must intersect every subvariety in \mathcal{M}_0 . This is true despite the fact that C is noncompact because the closed integral curves of v that are associated to C 's ends are disjoint from \mathcal{H}_0 , from the part of M_δ where $f \notin (1, 2)$, and from the part $1 - 3 \cos^2 \theta < 0$ part of any $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$. However, this same observation about C 's ends implies the following: for any $\varepsilon > 0$, there exists $x_\varepsilon > 1$ such that C has no intersections with the subvarieties from \mathcal{M}_0 that are parametrized as in Proposition 3.1 by values of (x, y) in $\mathbb{R} \times (-1 + \varepsilon, 1 - \varepsilon)$ with $|x| \gg x_\varepsilon$.

Suppose that C 's projection to Y intersects M_δ where $f = 1$ or where $f = 2$, or that it intersects some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ where $1 - 3 \cos^2 \theta = 0$. Two cases can be distinguished. The first occurs when such an intersection occurs at some point that is not on the locus in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ where both $u = 0$ and $1 - 3 \cos^2 \theta = 0$. Since C has no component in either \mathcal{M}_0 or \mathcal{M}_1 , so C must intersect some submanifold from \mathcal{M}_0 also. This is ruled out by what is said in the previous paragraph.

In the second case, the projection of C to Y intersects the relevant set only in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ where $u = 0$ and $1 - 3 \cos^2 \theta = 0$. For the sake of argument, suppose this occurs where $\cos \theta = \frac{1}{\sqrt{3}}$. The subvariety C therefore intersects the J -holomorphic subvariety $\mathbb{R} \times \widehat{\gamma}_\mathfrak{p}^+$. Suppose that the irreducible component involved is not this same $\mathbb{R} \times \widehat{\gamma}_\mathfrak{p}^+$. Fix $\varepsilon \in (0, 1 - \frac{1}{\sqrt{3}}]$ and an angle $\phi_* \in \mathbb{R}/2\pi\mathbb{Z}$. It follows from (2-4) and Lemma 2.2 that there exists a segment of an integral curve of v in $\mathcal{H}_\mathfrak{p}$ with the following properties: it enters $\mathcal{H}_\mathfrak{p}$ from the index 1 critical point end; it exits from the index 2 critical point end; it lies where the function $1 - 3 \cos^2 \theta$ is negative; and it intersects the $u = 0$ locus at the point where $\cos \theta = \frac{1}{\sqrt{3}} + \varepsilon$ and $\phi = \phi_*$. If ε

is sufficiently small, this curve must intersect C . In particular, there are points on C where $1 - 3 \cos^2 \theta$ is negative, and so the previous analysis applies. \square

4B Projection to the $f \in (1, 2)$ part of M_δ

This subsection makes some first observations about the intersection between a given ech-HF subvariety and the $f \in (1, 2)$ part of M_δ . The discussion here has two parts.

Part 1 The lemma that follows states a fact that underpins all of the subsequent analysis.

Lemma 4.2 *Let C denote an ech-HF subvariety. Then C has intersection number G with each subvariety from \mathcal{M}_Σ .*

Proof The positive ends of C determine a set $\Theta_{C_+} \in \mathcal{Z}_{\text{ech}, M}$. The integral curves of v that form Θ_{C_+} are pairwise disjoint in Y because their restrictions to M_δ are pairwise disjoint. Let $Z_+ \subset Y$ denote the union of the integral curves of v from Θ_{C_+} . The submanifold $\mathbb{R} \times Z_+ \subset \mathbb{R} \times Y$ is a J -holomorphic subvariety.

The intersection of Z_+ with M_δ is a disjoint union of segments with each starting on the boundary of the radius δ coordinate ball centered on an index 1 critical point of f and ending on the boundary of the radius δ coordinate ball centered on an index 2 critical point of f . Let $\{\gamma_1, \dots, \gamma_G\}$ denote these segments. The submanifold with boundary $\mathbb{R} \times (\bigcup_{1 \leq k \leq G} \gamma_k) \subset \mathbb{R} \times Z$ is a properly embedded cylinder with boundary in $\mathbb{R} \times M_\delta$.

Fix a given point $y \in (1 + \delta^2, 2 - \delta^2)$ and let $S_y \subset M_\delta$ denote the corresponding level set of f . If $x \in \mathbb{R}$, then $\{x\} \times S_y \in \mathcal{M}_\Sigma$. This subvariety has intersection number G with $\mathbb{R} \times (\bigcup_{1 \leq k \leq G} \gamma_k)$ and thus with $\mathbb{R} \times Z_+$ since S_y is disjoint from $Z_+ \cap (\bigcup_{p \in \Delta} \mathcal{H}_p)$. This implies that C has intersection number G with all $x \gg 1$ versions of $\{x\} \times S_y$.

Given that intersection numbers are locally constant under small perturbations, and given that $|s|$ is bounded on any given submanifold from \mathcal{M}_Σ , what is said in the previous paragraph implies that C has intersection number G with every submanifold from \mathcal{M}_Σ . \square

Part 2 This part uses Lemma 4.2 with Proposition 4.1 to say something about the intersections between a given ech-HF subvariety and the surfaces in \mathcal{M}_Σ that are

parametrized by pairs (x, y) with y near 1 or with y near 2. The relevant remarks are supplied by the next lemma.

Lemma 4.3 *Let C denote an ech-HF subvariety. Fix $\varepsilon \in (\delta, \delta_*)$; then there exists $z > 0$ with the following significance:*

- *Fix $(x, y) \in \mathbb{R} \times (1, 1+z]$ and let $S \in \mathcal{M}_\Sigma$ denote the corresponding submanifold. Then projection to Y of $C \cap S$ has precisely one point in the $|u| < R + \ln \varepsilon$ part of each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$.*
- *Fix $(x, y) \in \mathbb{R} \times [2-z, 2)$ and let $S \in \mathcal{M}_\Sigma$ denote the corresponding submanifold. Then projection to Y of $C \cap S$ has precisely one point in the $|u| < R + \ln \varepsilon$ part of each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$.*

Proof The submanifold C has G intersections with $\mathbb{R} \times Z_+$. Suppose that there exists a sequence $\{(x_k, z_k)\}_{k=1,2,\dots} \subset \mathbb{R} \times (0, \delta_*^2]$ with no convergent subsequence with the following property: for each $k \in \{1, 2, \dots\}$, there exists at least one intersection point between C and the $(x = x_k, z = z_k)$ version of S whose projection to Y lies in M_ε . The sequence $\{x_k\}$ is bounded from above. To see why this is, note first that any constant $s \gg 1$ slice of C is very close to Z_+ . Meanwhile, the curves that form Z_+ do not intersect an $f = 1 + z$ level set of f in M_δ if $z < c_0^{-1}$ since this is true for the integral curves of \mathfrak{v} that run from the index 1 critical points of f to the index 2 critical points of f . A similar argument using the integral curves of \mathfrak{v} from Θ_{C_-} proves that $\{x_k\}_{k=1,2,\dots}$ is bounded from below. As a consequence, no generality is lost by assuming that the latter sequence converges. Assuming that this is the case, and given that $\{(x_k, z_k)\}_{k=1,2,\dots}$ has no convergent subsequences, it follows that the sequence $\{z_k\}_{k=1,2,\dots}$ must have limit zero. But then C must intersect some subvariety from \mathcal{M}_1 since M_δ is compact. Such an intersection is ruled out by what is said in Proposition 4.1.

Suppose now that there is a sequence $\{(x_k, z_k)\}_{k=1,2,\dots} \subset \mathbb{R} \times (0, \delta_*^2]$ which does not converge and is such that for each $k \in \{1, 2, \dots\}$, there exists at least two intersection points between C and the $(x = x_k, z = z_k)$ version of S whose projection to Y lies in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ where $|u| < R + \ln \varepsilon$. An argument much like that given in the preceding paragraph proves that $\{x_k\}_{k=1,2,\dots}$ must be bounded and so can be assumed convergent. Let x_* denote the limit point. As before, the sequence $\{z_k\}_{k=1,2,\dots}$ can be assumed to converge to zero. Fix $x_0 \gg 1$ such that the very large s part of C is very close to the correspondingly large s part of $\mathbb{R} \times Z_+$. If $x > x_0$

and k is large, then C has 1 intersection point with each $(x > x_0, z = z_k)$ version of S where $|u| < R + \ln \varepsilon$ in \mathcal{H}_p . As x is now decreased from x_0 towards x_* , there must exist, for each large x , some point x'_k between x_0 and x_k with the property that C intersects the $(x = x'_k, z = z_k)$ version of S where $|u| > R + 2 \ln \varepsilon$ in \mathcal{H}_p . But this is ruled out by what is said in the previous paragraph.

These last two lemmas provide a picture of C 's intersection with $f^{-1}(1, 2)$ that plays a central role in Section 7 where the Heegaard Floer geometry enters the story. To set the stage now, use the integral curves of v as in Part 5 of Section 1C to define a diffeomorphism from $f^{-1}(1, 2) \subset M$ to $(1, 2) \times \Sigma$ which identifies v with the Euclidean vector field ∂_t along the $(1, 2)$ factor and which identifies f with the coordinate t on this same factor. The next lemma uses this diffeomorphism to view the $\mathbb{R} \times M_\delta$ part of C as a submanifold in $\mathbb{R} \times (1, 2) \times \Sigma$. The lemma also refers to the respective sets C_+ and C_- in Σ and their annular neighborhoods T_+ and T_- . These are introduced in Parts 2 and 4 of Section 1C. □

Lemma 4.4 *There exists $\kappa > 1$ that is independent of the parameters (δ, χ_0, R) and which has the following significance: Suppose that C is an ech-HF subvariety. View the $\mathbb{R} \times M_\delta$ part of C as a subset of $\mathbb{R} \times (1, 2) \times \Sigma$ as explained above. There exists $\varepsilon > 0$ that depends on C and is such that:*

- *The $\mathbb{R} \times M_\delta$ part of C intersects $\mathbb{R} \times (1, 1 + \varepsilon) \times \Sigma$ in $\mathbb{R} \times (1, 1 + \varepsilon) \times T_+$.*
- *The $\mathbb{R} \times M_\delta$ part of C intersects $\mathbb{R} \times (2 - \varepsilon, 2) \times \Sigma$ in $\mathbb{R} \times (2, 2 - \varepsilon) \times T_-$.*
- *The projection to Σ of the $|s| > \varepsilon^{-1}$ part of C 's intersection with $\mathbb{R} \times M_\delta$ lies in the union of the radius $\kappa\delta$ disks centered at the points in $C_- \cap C_+$.*

Proof Fix $\varepsilon > 0$ and use ε to get z from Lemma 4.3. Supposing that $y \in (1, 1 + z)$ or $y \in (2, 2 - z)$, it follows from Lemma 4.2 and 4.3 that C 's intersection with the locus $\mathbb{R} \times f^{-1}(y)$ is in $\bigcup_{p \in \Lambda} (\mathbb{R} \times \mathcal{H}_p)$. Meanwhile, it follows from Proposition 4.1 that $1 - 3 \cos^2 \theta > 0$ on the part of C in $\mathbb{R} \times (M_\delta \cap (\bigcup_{p \in \Lambda} \mathcal{H}_p))$. This implies what is said by the first two bullets. The third bullet follows from Corollary 2.6 since the $s \rightarrow \pm\infty$ limit of the constant $s \in \mathbb{R}$ slices of C converge to the closed orbits of v from Corollary 2.6. □

4C The projection to $\bigcup_{p \in \Lambda} \mathcal{H}_p$

Fix $p \in \Lambda$. This subsection talks about the intersection of a given ech-HF subvariety with $\mathbb{R} \times \mathcal{H}_p$. The section has five parts. Proposition 4.5 in Part 2 summarizes what is

said. Part 1 supplies the background for the statement of this proposition. Parts 3 and 4 contain the proof of Proposition 4.5. Part 5 concerns the values of the function θ on the $\mathbb{R} \times \mathcal{H}_p$ part of an ech-HF subvariety.

This and subsequent sections use \mathcal{H}_p^+ to denote the part of \mathcal{H}_p here the coordinates u and θ obey $|u| < R + \ln \delta_*$ and $1 - 3 \cos^2 \theta > 0$.

Part 1 What follows directly describes an extremely useful parametrization of \mathcal{H}_p^+ . This parametrization involves the function \hat{h} on \mathcal{H}_p given by

$$(4-1) \quad \hat{h}(u, \theta) = (\chi + 2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)})) \cos \theta \sin^2 \theta.$$

In the notation of (1-24), the function $\hat{h} = f \cos \theta \sin^2 \theta$. This function has no critical points on \mathcal{H}_p^+ . In fact, $d\hat{h}$ has strictly negative pairing with the vector field e_1 in (3-3). More is true: it follows from (3-10) that $d\hat{h}$ pulls back without critical points to the \mathcal{H}_p^+ part of any submanifold from \mathcal{M}_Σ , and it follows from (3-11) that $d\hat{h}$ pulls back without critical points to any submanifold from \mathcal{M}_0 . As the vector field ∂_ϕ annihilates \hat{h} , what was just said implies the following: the pair (\hat{h}, ϕ) define coordinates on the \mathcal{H}_p^+ part of any submanifold from \mathcal{M}_Σ and on any submanifold from \mathcal{H}_0 .

Note for future reference that the vector field v also annihilates \hat{h} . (Use the formula for v in Property 3 from Section 1E to see this.) It follows as a consequence that \hat{h} is constant on the segments of the integral curves of v in \mathcal{H}_p . In particular, the level sets of \hat{h} are foliated by the integral curves of v , and this implies that \hat{h} , if not constant, lacks local extrema on any submanifold in \mathcal{H}_p^+ whose tangent space is preserved by J .

Define an embedding from $\mathbb{R} \times \mathcal{H}_p^+$ into

$$(4-2) \quad \mathbb{R} \times (-R - \ln \delta_*, R + \ln \delta_*) \times (\mathbb{R}/2\pi\mathbb{Z}) \times \left(-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2\right),$$

as follows: Let q denote a given point in $\mathbb{R} \times \mathcal{H}_p^+$. The image of q via this embedding has coordinates $(x, \hat{u}, \hat{\phi}, h)$ where $\hat{\phi} = \phi(q)$ and $h = \hat{h}(q)$, and where (x, \hat{u}) are the respective values of $s \in \mathbb{R}$ and u at the $\theta = \frac{\pi}{2}$ point on the integral curve of the vector field in (3-10) that contains the point q . (The vector field in (3-10) is $-\sigma^{-1} J \cdot \partial_\phi$.)

Let Ψ_p denote the inverse of the map just defined. The domain of Ψ_p sits in (4-2) as a product $\mathbb{R} \times \mathcal{X}$ with $\mathcal{X} \subset (-R - \ln \delta_*, R + \ln \delta_*) \times (\mathbb{R}/2\pi\mathbb{Z}) \times \left(-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2\right)$ a certain p -independent open set.

To obtain a sense of this parametrization, let $(x, \hat{u}, \hat{\phi}, h)$ with $\hat{u} \neq 0$ denote a point in $\mathbb{R} \times \mathcal{X}$. Then $\Psi_p(x, \hat{u}, \hat{\phi}, h)$ is the point where $(\hat{h} = h, \phi = \hat{\phi})$ on the surface

from \mathcal{M}_Σ that also contains the point in $\mathbb{R} \times \mathcal{H}_p$ where $s = x$, $u = \hat{u}$, $\theta = \frac{\pi}{2}$ and $\phi = \hat{\phi}$. Meanwhile, a point $(x, 0, \hat{\phi}, h)$ is sent by Ψ_p to the point in the surface from \mathcal{M}_{p_0} that also contains the point where $s = x$, $u = 0$, $\theta = \frac{\pi}{2}$ and $\phi = \hat{\phi}$.

The map Ψ_p has the following very important properties:

- (4-3) • The constant (x, \hat{u}) surfaces in $\mathbb{R} \times \mathcal{X}$ are mapped by Ψ_p to J -holomorphic submanifolds.
- The map Ψ_p is equivariant with respect to the \mathbb{R} action on $\mathbb{R} \times \mathcal{X}$ that translates along the first factor, and the corresponding action on $\mathbb{R} \times \mathcal{H}_p^+$.
 - The map Ψ_p is equivariant with respect to the $\mathbb{R}/2\pi\mathbb{Z}$ action that translates the coordinate $\hat{\phi}$ on $\mathbb{R} \times \mathcal{X}$ and translates the coordinate ϕ on $\mathbb{R} \times \mathcal{H}_p^+$.

These properties, and especially the first, motivate the introduction of Ψ_p . Note in particular that the constant $(x, \hat{u} \neq 0)$ surfaces in $\mathbb{R} \times \mathcal{X}$ are mapped to the intersections with $\mathbb{R} \times \mathcal{H}_p^+$ of the surfaces from \mathcal{M}_Σ . Meanwhile, the constant $(x, \hat{u} = 0)$ surfaces are mapped by Ψ_p to the surfaces in \mathcal{M}_{p_0} .

The formulas for the Ψ_p -pushforwards of the coordinate vector fields ∂_x , $\partial_{\hat{u}}$, $\partial_{\hat{\phi}}$ and ∂_h are also worth keeping in mind. These are written using the notation from Property 3 in Section 1E and (3-3) as

- (4-4) • $\Psi_{p*}\partial_x = \partial_s$,
- $\Psi_{p*}\partial_{\hat{u}} = \nu(v + \alpha^{-1}\sqrt{6}\chi \cos \theta \partial_\phi + \varpi \partial_s)$,
 - $\Psi_{p*}\partial_{\hat{\phi}} = \partial_\phi$,
 - $\Psi_{p*}\partial_h = -\beta^{-1}(e_1 - \sigma^{-1}\sqrt{6}\chi_\delta f \cos \theta \sin^2 \theta \partial_s)$,

where ν , ϖ and β are certain functions of the pair (u, θ) with ν and β positive.

Part 2 This part of the subsection states the promised proposition that describes the intersection of an ech-HF subvariety with a given $p \in \Lambda$ version of \mathcal{H}_p^+ . To set the stage, fix $z_* \in (0, \frac{1}{4}\delta_*^2)$. The proposition refers to the subspace $\mathcal{M}_{\Sigma,*} \subset \mathcal{M}_\Sigma$ that is parametrized by $(x, y) \in \mathbb{R} \times (1, 2)$ with either $y \leq 1 + z_*$ or with $y \geq 2 - z_*$. It also refers to ν 's integral curves $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ in \mathcal{H}_p where $(u = 0, \cos \theta = \frac{1}{\sqrt{3}})$ and $(u = 0, \cos \theta = -\frac{1}{\sqrt{3}})$. By way of notation, the proposition introduces I_* to denote $[-R - \frac{1}{2} \ln z_*, R + \frac{1}{2} \ln z_*]$.

The proposition also introduces the following terminology with regards to any given $p \in \Lambda$ version of the circles $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$. Let $\hat{\gamma}$ denote one or the other of these

integral curves of v . The curve $\hat{\gamma}$ is said to be associated to an end of C in \mathcal{H}_p^+ when the following occurs: the large $|s|$ part of an end of C lies entirely in \mathcal{H}_p^+ , and its constant $|s|$ slices converge as $|s| \rightarrow \infty$ to $\hat{\gamma}$. Note in this regard that $\hat{\gamma}$ is not associated to an end of C in \mathcal{H}_p^+ if $\mathbb{R} \times \gamma$ is a component of C .

Proposition 4.5 Fix $p \in \Lambda$ and $z_* \in (0, \frac{1}{4}\delta_*^2)$. Suppose that C is an ech-HF subvariety which has intersection number 1 in \mathcal{H}_p^+ with each submanifold from $\mathcal{M}_{\Sigma, *}$. Then the Ψ_p -inverse image of C in $\mathbb{R} \times \mathcal{X}$ intersects the $\hat{u} \in I_*$ part of $\mathbb{R} \times \mathcal{X}$ as a graph of the form

$$(x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi^p(x, \hat{u}), h = \zeta^p(x, \hat{u})),$$

where the pair (φ^p, ζ^p) denotes a smooth map from certain domain in $\mathbb{R} \times I_*$ to $(\mathbb{R}/2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$. This domain is as follows:

- If neither $\hat{\gamma}_p^+$ nor $\hat{\gamma}_p^-$ is associated to an end of C in $\mathbb{R} \times \mathcal{H}_p^+$, then the domain of (φ^p, ζ^p) is the whole strip $\mathbb{R} \times I_*$.
- If just one of $\hat{\gamma}_p^+$ or $\hat{\gamma}_p^-$ is associated to an end of C in $\mathbb{R} \times \mathcal{H}_p^+$, then the domain of (φ^p, ζ^p) is the complement in $\mathbb{R} \times I_*$ of a single $\hat{u} = 0$ point.
- If both $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ are associated to ends of C in $\mathbb{R} \times \mathcal{H}_p^+$, then the domain of (φ^p, ζ^p) is the complement in $\mathbb{R} \times I_*$ of two $\hat{u} = 0$ points.

This proposition implies that C 's intersection with the image of Ψ_p is a smooth, properly embedded submanifold with boundary in this image.

The proof of this proposition is given in the next part of this subsection.

Part 3 The proof of Proposition 4.5 invokes two auxiliary lemmas. Both are proved in Part 4. The first of the lemmas is very specific to the situation at hand.

Lemma 4.6 Let C denote an ech-HF subvariety.

- If neither $\hat{\gamma}_p^+$ nor $\hat{\gamma}_p^-$ is associated to an end of C in $\mathbb{R} \times \mathcal{H}_p^+$, then C has intersection number 1 with each subvariety from \mathcal{M}_{p_0} .
- If just one of $\hat{\gamma}_p^+$ or $\hat{\gamma}_p^-$ is associated to an end of C in $\mathbb{R} \times \mathcal{H}_p^+$, then C has intersection number zero with one subvariety from \mathcal{M}_{p_0} and intersection number 1 with each of the others.
- If both $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ are associated to ends of C in $\mathbb{R} \times \mathcal{H}_p^+$, then C has intersection number zero with two subvarieties from \mathcal{M}_{p_0} and intersection number 1 with each of the others.

The next lemma contains a fundamental observation about pseudoholomorphic curves and fibrations with pseudoholomorphic fiber. To set the stage for this lemma, suppose that U is an open set with compact closure in a smooth 4–manifold with an almost complex structure. A subset $C \subset U$ is said to be a *pseudoholomorphic subvariety in U* when the following conditions are met: First, the closure of U is contained in an open set and this open set contains a closed subset that lacks point components and intersects U as C . Second, this larger closed subset contains a finite set whose complement is a smooth 2–dimensional submanifold. Third, the tangent space of this submanifold is mapped to itself by the almost complex structure.

Lemma 4.7 *Let U denote an open set in a 4–manifold with almost complex structure. Let S denote a 2–dimensional surface and suppose that the closure of U is contained in an open set with a proper fibration to S with pseudoholomorphic fibers. Let C denote a pseudoholomorphic subvariety in U with no open set in a fiber. Assume in addition that C has intersection number at most 1 with each fiber. Then C is a smooth submanifold and the fibering map to S restricts to C as a diffeomorphism onto an open set in S .*

Granted for the moment these two lemmas, here is the promised:

Proof of Proposition 4.5 The composition of first the inverse of Ψ_p and then the projection from $\mathbb{R} \times \mathcal{X}$ to $\mathbb{R} \times (-R - \ln \delta_*, R + \ln \delta_*)$ is a submersion with J –holomorphic fibers. Granted that this is the case, then the depiction by the proposition of the Ψ_p –inverse image of C follows from Lemmas 4.6 and 4.7 with the fact that the pair (ϕ, \hbar) restrict as coordinates to the $\mathbb{R} \times \mathcal{H}_p^+$ part of surfaces from M_Σ and they also restrict as coordinates to each subvariety from \mathcal{M}_{p_0} . □

Part 4 This part of the subsection contains the proofs of Lemmas 4.6 and 4.7.

Proof of Lemma 4.6 The proof has seven steps.

Step 1 The set Θ_{C_+} contains one integral curve of v whose intersection with \mathcal{H}_p runs from the $u = -R - \ln(7\delta_*)$ end to the $u = R + \ln(7\delta_*)$ end. This curve intersects the $u = 0$ slice of \mathcal{H}_p , and this implies that each $x \gg 1$ labeled submanifolds in \mathcal{M}_{p_0} intersect C precisely once. The analogous argument using Θ_{C_-} proves that the $x \ll -1$ labeled submanifolds in \mathcal{M}_{p_0} also intersect C precisely once. If C is as described by the first bullet of Lemma 4.6, then the continuity of the intersections as x varies in \mathbb{R} implies that each submanifold in \mathcal{M}_{p_0} must intersect C exactly once.

Step 2 Before considering the second and third bullets, use (3-11) with Lemma 2.3 to see that the restriction of Ψ_p^*s to the $\hat{u} = 0$ slice of the domain of Ψ_p has the form

$$(4-5) \quad \Psi_p^*s = x - \mu \ln(1 - 3 \cos^2 \hat{\theta}) + w(\hat{\theta}),$$

where $\mu > 0$ is a constant, $w(\cdot)$ is a smooth function on \mathbb{R} and $\hat{\theta}$ is the implicit function of the coordinate h on $\mathbb{R} \times \mathcal{X}$ that is given by the unique $\theta \in (0, \pi)$ with $1 - 3 \cos^2 \theta > 0$ and is such that $h = (\chi_0 + 4e^{-2R}) \cos \theta \sin^2 \theta$. The pullback of s to where \hat{u} is nonzero has the same schematic form of what is written on the right-hand side of (4-5), but with μ a function of \hat{u} , with $\hat{\theta}$ a function of h and \hat{u} , and with w now a function of \hat{u} and $\hat{\theta}$. When using (4-5) in what follows, keep in mind that the very small $|u|$ part of \mathcal{H}_p has inverse image via Ψ_p in the very small $|\hat{u}|$ part of Ψ_p 's domain (but not vice versa).

Step 3 This step explains why neither $\hat{\gamma}_p^+$ nor $\hat{\gamma}_p^-$ are associated to negative ends of C . Suppose to the contrary that $1 - 3c^2 \rightarrow 0$ as $s \rightarrow -\infty$ on an end $\mathcal{E} \subset C$ in $\mathbb{R} \times \mathcal{H}_p$. This can occur only if $u \rightarrow 0$ on the constant s slices of \mathcal{E} . This requires that \hat{u} limit to zero on $\Psi_p^{-1}(\mathcal{E})$. This observation implies that there exists a path $\nu \subset \mathbb{R} \times I_*$ of the following sort: First, ν is in the domain of the functions (φ^p, ζ^p) so ν defines a corresponding path in C . Denote this path by $\nu_{\mathcal{E}}$. This path lies in \mathcal{E} , and s is unbounded from below on $\nu_{\mathcal{E}}$. Granted these properties, use (4-5) to see that $s \rightarrow -\infty$ on $\nu_{\mathcal{E}}$ only if $x \rightarrow -\infty$ on ν . Since $u \rightarrow 0$ on $\nu_{\mathcal{E}}$ as $s \rightarrow -\infty$, so $\hat{u} \rightarrow 0$ on ν as $x \rightarrow -\infty$ on ν . The existence of such a path is not possible given that the $s \ll -1$ part of C in $\mathbb{R} \times \mathcal{H}_p$ has a second $s \rightarrow -\infty$ limit of its constant s slices. The latter is a segment of an integral curve of v from Θ_{C_-} . This segment crosses the $u = 0$ slice transversely, and so its inverse image via Ψ_p has transverse intersection with the $\hat{u} = 0$ slice of Ψ_p 's domain. It follows as a consequence that the $x \ll -1$ part of a path such as ν must have more than one inverse image in $\Psi_p^{-1}(C)$. This nonsense proves that neither $\hat{\gamma}_p^+$ nor $\hat{\gamma}_p^-$ are associated a negative end of C .

Step 4 Suppose that $\hat{\gamma}_p^+$ or $\hat{\gamma}_p^-$ are associated to a positive end of C . Let \mathcal{E} denote the latter. The $s \gg 1$ part of \mathcal{E} can be parametrized by the correspondingly large s part of $\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z})$ as

$$(4-6) \quad (s, \phi) \rightarrow (s, u = b(s, \phi), \theta = \theta_* + a(s, \phi), \phi),$$

where (a, b) are smooth functions of their arguments that limit uniformly to 0 as $s \rightarrow \infty$.

Because the integral curves involved are nondegenerate, the functions (a, b) must behave at large s as

$$(4-7) \quad a = e^{-\lambda s}(a_*(\phi) + \epsilon_a) \quad \text{and} \quad b = e^{-\lambda s}(b_*(\phi) + \epsilon_b),$$

where λ is a positive constant and (a_*, b_*) are smooth functions on $\mathbb{R}/2\pi\mathbb{Z}$ that are not simultaneously zero. Meanwhile, ϵ_a and ϵ_b are smooth functions of (s, ϕ) with limit zero as $s \rightarrow \infty$. This can be proved using the arguments used to prove Proposition 2.4 in [10]. The arguments in the latter reference identify (a_*, b_*) as eigenvectors of a symmetric, first-order differential operator on $C^\infty(S^1; \mathbb{R}^2)$ that is defined by the Lie transport by v on the normal bundle to the relevant integral curve and the metric $\omega(\cdot, J\cdot)$ on this bundle. The constant λ is the corresponding eigenvalue.

In the cases at hand, the integral curves are orbits of the $\mathbb{R}/2\pi\mathbb{Z}$ action that translates ϕ and both v and J are invariant with respect to this action. This implies that any given eigenfunction of the relevant operator must have the form

$$(4-8) \quad a_*(\phi) = a_n \cos(n\phi - \phi_n) \quad \text{and} \quad b_*(\phi) = b_n \sin(n\phi - \phi_n),$$

where a_n and b_n are \mathbb{R} -valued constants and $\phi_n \in \mathbb{R}/2\pi\mathbb{Z}$.

Step 5 In the case at hand, only the $n = 0$ version of (4-8) can occur because $\theta - \theta_*$ does not change sign on \mathcal{E} . In this case, (a_*, b_*) are constants that are not both zero. Note that $a_* \geq 0$ if γ_p^+ is associated to \mathcal{E} and $a_* \leq 0$ if γ_p^- is associated to \mathcal{E} . As explained next, the only case that can occur here is that where $a_* \neq 0$ and $b_* = 0$.

To see why this is, note first that the argument given in Step 3 can be repeated no matter the value of b_* to see that x stays bounded on $\Psi_p^{-1}(\mathcal{E})$. What follows is a consequence. There exists a bounded set $I \subset \mathbb{R}$ with the following property: given $T \geq 1$ and $\epsilon > 0$, there exists (x, \hat{u}) with $x \in I$ and $|\hat{u}| < \epsilon$ that are in the domain of (φ^p, ζ^p) and parametrize a point on \mathcal{E} where $s > T$.

Suppose for the sake of argument that $b_* > 0$. This implies via (4-6)–(4-8) that u is positive on the $s \gg 1$ part of \mathcal{E} . Therefore, the $s \gg 1$ part of \mathcal{E} is parametrized via Ψ_p only by points where \hat{u} is positive. Let $\{x_n\}_{n=1,2,\dots}$ denote a sequence of points in \mathbb{R} such that there exists $\hat{u} \in I_*$ with $\hat{u} \in (0, \frac{1}{n})$ and such that (x_n, \hat{u}) is a point in the domain of (φ^p, ζ^p) that parametrizes and $s > n$ point in \mathcal{E} . Let $x_* \in I$ denote a limit point of $\{x_n\}_{n=1,2,\dots}$. The interval $\{x_*\} \times [-1, 0)$ in $\mathbb{R} \times I_*$ lies in the domain of (φ^p, ζ^p) and so corresponds via Ψ_p to an embedded arc in C . The function s is bounded on this arc if b_* is positive, and so the arc has a limit on the $u = 0$ slice of C .

A disk about the limit point is parametrized via (φ^p, ζ^p) by an open neighborhood of $(x_*, 0)$ in $\mathbb{R} \times I_*$. But if this is the case, then C will have two or more intersections in $\mathbb{R} \times \mathcal{H}_p^+$ with some elements in $\mathcal{M}_{\Sigma, *}$.

Step 6 Assume here that only one of $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ is associated to an end of C . Let $\theta_* \in (0, \pi)$ denote the angle of the component of the $1 - 3c^2 = 0$ locus in question.

Use $x_+ \in \mathbb{R}$ to denote the maximum value of x on the complement of the domain of (φ^p, ζ^p) , and let $x_- \in \mathbb{R}$ denote the minimum value. Use the fact that $a_* \neq 0$ and $b_* = 0$ in (4-6)–(4-8) to write (4-5) at points $(x, 0) \in \mathbb{R} \times I_*$ with $x > x_+$ as

$$(4-9) \quad (1 - \lambda\mu)\Psi_{p*}s = x - \mu \ln(2\sqrt{2}|a_*|) + w(\theta_*) + \dots,$$

where the unwritten terms have limit zero as $s \rightarrow \infty$. Since $x > x_+$, this equation requires that $\lambda\mu = 1$ and it identifies $x_+ = \mu \ln(2\sqrt{2}|a_*|) + w(\theta_*)$. The analogous argument finds this same value for x_- . It follows that the domain of (φ^p, ζ^p) is the complement in $\mathbb{R} \times I_*$ of a single $\hat{u} = 0$ point. This proves what is asserted by the second bullet of Lemma 4.6.

Note that the argument in Step 5 proves the following: given $T > 1$ there exists positive ε_T such that the complement of the origin of the radius ε_T disk in $\mathbb{R} \times I_*$ about this missing point parametrizes only points in C from the $s > T$ part of \mathcal{E} .

Step 7 Assume in this last step that both $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ are associated to ends of C . Let \mathcal{E}_+ and \mathcal{E}_- denote the associated domains. The domain of (φ^p, ζ^p) in this case is the complement in $\mathbb{R} \times I_*$ of some bounded set $I \subset \mathbb{R}$ on the $\hat{u} = 0$ slice. The third bullet asserts that I consist of two points. To see why this is, note first that the function s on the image via Ψ_p of the domain of (φ^p, ζ^p) must have an index 1 critical point that corresponds to the saddle between \mathcal{E}_+ and \mathcal{E}_- . To elaborate, let \mathcal{Y} denote the space of compact, embedded paths in C of the following sort: Both ends are parametrized by $\hat{u} = 0$ points in the domain of (φ^p, ζ^p) . In addition, one end of the path lies in \mathcal{E}_+ and the other lies in \mathcal{E}_- . The function s achieves its minimum on each path in \mathcal{Y} , and the set of these minima is bounded from above. Let s_* denote the supremum.

Fix a path, $\nu_C \subset \mathcal{Y}$. This path is parametrized by a corresponding path, $\nu \subset \mathbb{R} \times I_*$ in the domain of (φ^p, ζ^p) . This path must cross the $\hat{u} = 0$ locus at a point with components from I on both sides. For, if not, the path can be homotoped rel endpoints in the domain of (φ^p, ζ^p) to a path whose image is in \mathcal{Y} has s everywhere greater than s_* .

Granted that I has at least two components, the argument used in Step 6 can be repeated to prove that there are precisely two components and each is a single point. \square

Proof of Lemma 4.7 Use $\pi: U \rightarrow S$ to denote the fibering map. The map π is 1-to-1 and continuous on C because pseudoholomorphic subvarieties have locally positive intersection number. The differential of π at each smooth point of C must be invertible. To see why, note that this differential is zero where C is tangent to a fiber. But there can be no tangencies. This can be proved by using the local structure of intersecting pseudoholomorphic curves to see that C must have intersection at least two with the fibers of π if it has a tangency. See eg Proposition 2.6 in [14] or Lemma A.7 in [24]. The local picture given in these references also imply that C must have intersection number greater than 1 with a fiber through any singular point. \square

Part 5 The last part of the subsection first states and then proves a lemma about the values of the function θ on pseudoholomorphic subvarieties in $\mathbb{R} \times \mathcal{H}_p$.

Lemma 4.8 *Let $U \subset \mathbb{R} \times \mathcal{H}_p$ denote an open set with compact closure and let $C \subset U$ denote an embedded, pseudoholomorphic subvariety in U . If the function θ is not constant on C , then any local minima or maxima of θ on C has $\cos \theta = \pm 1$ or $u = 0$.*

Proof The proof has four steps. It is assumed in these steps that θ is not constant on C .

Step 1 This step proves that the function $\cos \theta$ on C cannot take a local maximum or local minimum where $1 - 3 \cos^2 \theta = 0$ and $u \neq 0$. To this end, assume $p \in C$ is a point where $\cos \theta = \frac{1}{\sqrt{3}}$ and $u \neq 0$. The locus where $\cos \theta = +\frac{1}{\sqrt{3}}$ in $\mathbb{R} \times \mathcal{H}_p$ is foliated by the $\mathbb{R} \times \mathcal{H}_p$ parts of the J -holomorphic submanifolds that are described in Proposition 3.3. Suppose that S is one of these manifolds and S intersects C at a local maximum or minimum of θ on C . Denote this intersection point by p . Since intersection points between J -holomorphic subvarieties do not accumulate, the point p is an isolated point of $S \cap C$. As a consequence, there exists $\delta > 0$ and a small disk in S centered at p with all of its boundary points having distance at least δ from C . This disk has J -holomorphic deformations where $\cos \theta > \frac{1}{\sqrt{3}}$, and it has J -holomorphic deformations where $\cos \theta < \frac{1}{\sqrt{3}}$. For example, these deformations can be taken to be disks in subvarieties from Proposition 3.1 and thus where $\cos \theta > \frac{1}{\sqrt{3}}$, or disks in subvarieties from Proposition 3.2 and thus where $\cos \theta < \frac{1}{\sqrt{3}}$. If such a deformation moves the original disk by a small amount, the new disk must still intersect C because local intersection numbers between J -holomorphic subvarieties are positive. It follows

as a consequence that $\cos \theta$ takes values on either side of $\frac{1}{\sqrt{3}}$ on any neighborhood of p in C . Much the same argument shows that $\cos \theta$ cannot take a local maximum or minimum on C with value $-\frac{1}{\sqrt{3}}$ where u is nonzero.

Step 2 Suppose again that θ is not constant on C and that $p \in C$ is a point where $\cos \theta = 0$ and $u \neq 0$. The locus in \mathcal{H}_p where $\cos \theta = 0$ is foliated by integral curves of v . As a consequence, there is a J -holomorphic subvariety of the form $\mathbb{R} \times \gamma$ with γ being a $\cos \theta = 0$ integral curve of v that intersects C at the point p . Meanwhile, there are integral curves of v as near as desired to γ in \mathcal{H}_p where $\cos \theta > 0$ and likewise where $\cos \theta < 0$. An argument much like that given in Step 1 explains why C must intersect the corresponding product with \mathbb{R} at points very close to p . It follows as a consequence that C does not take on a local maximum or minimum at p .

Step 3 It remains to prove that $\cos \theta$ has no local extremum on C in the complement of C 's intersection with the three loci where $u = 0$, where $\cos \theta \sin^2 \theta = 0$ and where $1 - 3 \cos^2 \theta = 0$. This is done by applying the maximum principle to a second-order equation for θ 's restriction to this part of C .

To start the argument, introduce $\{e^1, e^2\}$ to denote the dual 1-forms to the vector fields e_1 and e_2 that are depicted in (3-3). These are given by

$$(4-10) \quad \begin{aligned} e^1 &= \beta^{-1}(f(1 - 3 \cos^2 \theta) d\theta - f' \cos \theta \sin \theta du), \\ e^2 &= \beta_*^{-1}(d\phi + \sqrt{6} \chi \alpha^{-1} \cos \theta \hat{a}), \end{aligned}$$

where β is the positive function of (u, θ) that appears in (4-4) and β_* is a second, strictly positive function of (u, θ) . It follows from (4-3) and from the formulas in Property 3 of Section 1E that the 1-form $d\theta$ can be written as

$$(4-11) \quad \begin{aligned} d\theta &= (\chi + g')(1 - 3 \cos^2 \theta) e^1 + \sqrt{6} \alpha^{-1} \chi_\delta f' f \cos^2 \theta \sin^3 \theta e^2 \\ &\quad + \alpha^{-1} f' \cos \theta \sin \theta \hat{a}. \end{aligned}$$

Let $J^T: T^*(\mathbb{R} \times Y) \rightarrow T^*(\mathbb{R} \times Y)$ denote the adjoint of J . Then

$$(4-12) \quad \begin{aligned} J^T \cdot d\theta &= -\sigma(\chi + g')(1 - 3 \cos^2 \theta) e^2 + \sqrt{6} \sigma^{-1} f' f \cos^2 \theta \sin^3 \theta e^1 \\ &\quad + \alpha^{-1} f' \cos \theta \sin \theta ds. \end{aligned}$$

This can be written as

$$(4-13) \quad J^T \cdot d\theta = \sigma(\chi + g')(1 - 3 \cos^2 \theta) d\phi + b_1 d\theta + b_2 du + \alpha^{-1} f' \cos \theta \sin \theta ds,$$

where b_1 and b_2 are functions of u and θ . It follows from this that $d(J^T \cdot d\theta)$ can be written schematically as

$$(4-14) \quad d(J^T \cdot d\theta) = \mathfrak{z}_1 \wedge d\theta + (1 - 3 \cos^2 \theta) b_3 du \wedge d\phi + \cos \theta \sin \theta b_4 du \wedge ds,$$

where \mathfrak{z}_1 is a 1-form and b_3 and b_4 are functions of θ and u .

Step 4 Let λ denote a given function of u and θ that is defined and nowhere zero on the complement of the three loci where $u = 0$, $1 - 3 \cos^2 \theta = 0$ and $\cos \theta \sin^2 \theta = 0$. Use (4-13) and (4-14) to see that

$$(4-15) \quad \lambda^{-1} d(\lambda J^T \cdot d\theta) = \mathfrak{z}_2 \wedge d\theta + (b_3 + \lambda^{-1} \lambda_u \sigma(\chi + g'))(1 - 3 \cos^2 \theta) du \wedge d\phi + (b_4 + \lambda^{-1} \lambda_u \alpha^{-1} f' \cos \theta \sin \theta) du \wedge ds.$$

Here, λ_u is shorthand for $\partial_u \lambda$ and \mathfrak{z}_2 is another 1-form.

At the points where $f' \cos \theta \sin^2 \theta \neq 0$, use (4-10) to write

$$(4-16) \quad du = -(f' \cos \theta \sin \theta)^{-1} \beta e^1 + b_5 d\theta,$$

where b_5 is a function of u and θ . This done, (4-15) can be written as

$$(4-17) \quad \lambda d(\lambda J^T \cdot d\theta) = \mathfrak{z}_3 \wedge d\theta + (b_3 + \lambda^{-1} \lambda_u \sigma(\chi + g'))(1 - 3 \cos^2 \theta) du \wedge d\phi - (b_6 + \lambda^{-1} \lambda_u \beta \alpha^{-1}) e^1 \wedge ds.$$

Here \mathfrak{z}_3 is a 1-form and b_6 is a function of u and θ . The restriction to TC of $e^1 \wedge ds$ is that of $\sigma e^2 \wedge \hat{a}$. Modulo forms with a factor of $d\theta$, the latter is

$$(4-18) \quad \beta_*^{-1} \sigma(\chi + g')(1 - 3 \cos^2 \theta) d\phi \wedge du.$$

Noting the function multiplying $d\phi \wedge du$, it follows from (4-17) that up to terms with $d\theta$, the right-hand side of (4-17) restricts to TC as

$$(4-19) \quad (1 + \beta_*^{-1} \beta \alpha^{-1})(b_7 + \lambda^{-1} \lambda_u) \sigma(\chi + g')(1 - 3 \cos^2 \theta) du \wedge d\phi,$$

where b_7 is yet another function of u and θ . Granted (4-19), note that λ can be chosen on the complement of the three loci where $u = 0$, $\cos \theta \sin^2 \theta = 0$ and $1 - 3 \cos^2 \theta = 0$ so as to make (4-19) equal to zero. This done, the restriction of (4-17) to TC on the complement of these loci has the form

$$(4-20) \quad d(J^T d\theta) + \mathfrak{R} \cdot d\theta,$$

where \mathfrak{R} is a homomorphism from T^*C to $\wedge^2 T^*C$. What with (4-10), the claim made in the lemma follows using the strong maximum principle. \square

5 Quantitative concerns

This section is concerned with various cohomological and numerical invariants that can be associate to any given ech-HF subvariety. These invariants are used in the last subsection to give refined versions of Lemmas 4.3 and 4.4 and Proposition 4.5.

5A Homological considerations

This subsection introduces some homological considerations that play a central role in subsequent discussions. The subsection has six parts.

Part 1 This part sets the stage. Introduce by way of notation Y_\diamond to denote the subset of Y given by the union of the following:

- (5-1) • The $f^{-1}(1, 2)$ part of M_δ .
- $\bigcup_{p \in \Lambda} \mathcal{H}_p^+$.
 - Each $p \in \Lambda$ version of $\hat{\gamma}_p^+ \cup \hat{\gamma}_p^-$, this the $\{u = 0, 1 - 3 \cos^2 \theta = 0\}$ locus in \mathcal{H}_p .

(By way of a reminder, what is denoted by \mathcal{H}_p^+ here and subsequently is the part of \mathcal{H}_p where $|u| < R + \ln \delta$ and where $1 - 3 \cos^2 \theta > 0$.) Thus, Y_\diamond is obtained from the $f^{-1}(1, 2)$ part of M by removing the balls of radius δ balls centered at the index 1 and 2 critical points of M and then gluing in the portion of each $p \in \Lambda$ part of \mathcal{H}_p where $|\cos \theta|$ is less than $\frac{1}{\sqrt{3}}$, and then adding to this the two integral curves of v in each $p \in \Lambda$ version of \mathcal{H}_p where $u = 0$ and where $\cos \theta$ is either $\frac{1}{\sqrt{3}}$ or $-\frac{1}{\sqrt{3}}$.

The subspace Y_\diamond is introduced because $\mathbb{R} \times Y_\diamond$ contains all ech-HF subvarieties. In fact, Proposition 4.1 says that a given ech-HF subvariety is in the interior of $\mathbb{R} \times Y_\diamond$ unless it intersects a set described by the third bullet in (5-1), and it intersects such a set if and only if it contains some $p \in \Lambda$ version of the \mathbb{R} -invariant cylinder $\mathbb{R} \times \hat{\gamma}_p^+$ or the corresponding cylinder $\mathbb{R} \times \hat{\gamma}_p^-$. In any event, of interest in what follows is the homology of Y_\diamond .

Let Θ denote an element in $\mathcal{Z}_{\text{ech}, M}$. The set Θ defines a class in $H_1(Y_\diamond; \mathbb{Z})$ as follows: Orient each of its constituent closed integral curves using v so as to view each as a closed 1-cycle. If γ denotes the given curve, then $[\gamma]$ is used to denote this 1-cycle. The class defined by Θ is that of $\sum_{(\gamma, 1) \in \Theta} [\gamma]$. This class is denoted by $[\Theta]$ in what follows.

Let C denote an ech-HF subvariety. The corresponding classes $[\Theta_{C-}]$ and $[\Theta_{C+}]$ must define the same class in $H_1(Y_\diamond; \mathbb{Z})$ because the image of C in Y_\diamond via the projection defines a 2-chain whose boundary is the difference between these two classes. This observation places a constraint on the possible ech-HF subvarieties. Meanwhile, any two ech-HF with the same limit sets Θ_{C-} and Θ_{C+} define by their difference a closed 2-cycle in Y_\diamond and thus a class in $H_2(Y_\diamond; \mathbb{Z})$.

Part 2 The homology of Y_\diamond can be determined using the Mayer-Vietoris sequence for the decomposition $Y_\diamond = U \cup V$ with U being the $f^{-1}(1, 2)$ part of M_δ and with V being $\bigcup_{p \in \Lambda} (\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-)$. Note that $U \cap V$ is the union of the $1 - 3 \cos^2 \theta > 0$ part of disjoint spherical annuli about each of the index 1 and index 2 critical points of f . These annuli have inner radius δ and outer radius $7\delta_*$.

The relevant portions of this sequence are

$$(5-2) \quad 0 \rightarrow \mathbb{Z} \rightarrow H_2(Y_\diamond) \rightarrow \bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z}) \xrightarrow{\hat{O}} H_1(\Sigma) \oplus \left(\bigoplus_{p \in \Lambda} \mathbb{Z} \right) \xrightarrow{P} H_1(Y_\diamond) \xrightarrow{Q} \bigoplus_{p \in \Lambda} \mathbb{Z} \rightarrow 0.$$

Here and in what follows, \mathbb{Z} coefficients are implicit. By way of explanation, note that $H_2(\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-)$ and $H_2(U \cap (\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-))$ are zero, the former because the space $\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-$ deformation retracts onto a cylinder, and the latter because the intersection of U with $\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-$ deformation retracts onto two circles. Meanwhile, $U = f^{-1}(1, 2) \cap M_\delta$ deformation retracts onto $\Sigma = f^{-1}(\frac{3}{2})$ and so $H_2(U) = \mathbb{Z}$. This \mathbb{Z} is the left-most term in (5-2). Meanwhile, each $p \in \Lambda$ version of $H_1(U \cap (\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-))$ is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$. The direct sum of the latter is the $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ term preceding the map \hat{O} . The aforementioned deformation retract identifies $H_1(\Sigma)$ with $H_1(U)$, which is why the latter appears in the summand preceding the map P . The $\bigoplus_{p \in \Lambda} \mathbb{Z}$ summand in this same term is $H_1(\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-)$. Note for reference in what is to come that the generator of $H_1(\mathcal{H}_p^+ \cup \hat{\gamma}_p^+ \cup \hat{\gamma}_p^-)$ and of any given component of $H_1(U \cap \mathcal{H}_p^+)$ can be taken to be the $\theta = \frac{\pi}{2}$ circle in a constant u slice of \mathcal{H}_p^+ with the orientation defined by ∂_ϕ .

To say something about the homomorphisms in (5-2), consider first that on the right from $H_1(Y_\diamond)$ that is labeled by Q . Any oriented loop in $H_1(Y_\diamond)$ that crosses a given $p \in \Lambda$ version of \mathcal{H}_p^+ so as to have intersection number 1 with the $u = 0$ slice defines an element in $H_1(Y_\diamond)$ that maps via Q to the generator of p 's summand of the right-most

term in (5-2). By way of an example, suppose that $\Theta \in \mathcal{Z}_{\text{ech},M}$. Then the image of $[\Theta]$ via Q is the class $(1, 1, \dots, 1)$ in this $\bigoplus_{p \in \Lambda} \mathbb{Z}$ term on the far right in (5-2).

Granted what was just said, the difference between the classes in $H_1(Y_\diamond)$ that come from any two elements in $\mathcal{Z}_{\text{ech},M}$ define a class in $H_1(\Sigma) \oplus (\bigoplus_{p \in \Lambda} \mathbb{Z})$ via the homomorphism P . These classes are determined by the homomorphism

$$(5-3) \quad \hat{O}: \bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z}) \rightarrow H_1(\Sigma) \oplus \left(\bigoplus_{p \in \Lambda} \mathbb{Z} \right).$$

To say something about \hat{O} , introduce for each $p \in \Lambda$ the classes Z_{p_1} and Z_{p_2} in $H_1(\Sigma)$ that are generated by the intersection of Σ with the respective ascending disk from the index 1 component of p and descending disk from the index 2 component. Then \hat{O} sends an element $(n_{p_1}, n_{p_2})_{p \in \Lambda} \in \bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ to the element whose respective $H_1(\Sigma)$ and $\bigoplus_{p \in \Lambda} \mathbb{Z}$ summands are

$$(5-4) \quad \sum_{p \in \Lambda} (n_{p_1} Z_{p_1} + n_{p_2} Z_{p_2}) \quad \text{and} \quad \{(n_{p_1} - n_{p_2})\}_{p \in \Lambda}.$$

What follows are some additional remarks. To start, let \hat{O}_Σ denote the homomorphism from $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ to $H_1(\Sigma)$ that is given by the left-most term in (5-4). The cokernel of \hat{O}_Σ is isomorphic to $H_1(M)$ and its kernel is isomorphic to $H_2(M)$. This can be seen by using the homology Mayer–Vietoris sequence for M that comes from splitting M along Σ 's incarnation as the $f = \frac{3}{2}$ level set. It follows from this that the image of \hat{O} is isomorphic to $\ker(H_1(\Sigma) \rightarrow H_1(M)) \oplus H_2(M)$. Note however that this isomorphism is not canonical. It follows from this last observation that the image of P is isomorphic to $H_1(M) \oplus ((\bigoplus_{p \in \Lambda} \mathbb{Z})/H_2(M))$.

Part 3 Suppose that C is an ech-HF subvariety. Because the class $[\Theta_{C_+}] - [\Theta_{C_-}]$ in $H_1(Y_\diamond)$ is zero, there are classes in the $H_1(\Sigma) \oplus (\bigoplus_{p \in \Lambda} \mathbb{Z})$ term in (5-2) that map to the class $[\Theta_{C_+}] - [\Theta_{C_-}]$ via the homomorphism P in (5-2). As explained here, C determines a canonical choice for such a class with the $H_1(\Sigma)$ summand equal to 0. The entry in any given $p \in \Lambda$ summand is denoted here by m_p^C ; the latter is either 0, 1 or -1 . The explanation follows.

Fix $p \in \Lambda$ and introduce γ_{p_+} to denote the segment in \mathcal{H}_p^+ of the integral curve from Θ_{C_+} that crosses \mathcal{H}_p . Likewise define γ_{p_-} using Θ_{C_-} . Let p_1 and p_2 denote here the respective index 1 and index 2 critical points from p . Given $r \in (\delta_*, 7\delta_*)$, let $S_{1,r}$ and $S_{2,r}$ denote the respective boundaries of the radius r coordinate balls centered on

p_1 and p_2 . Choose r so that C has transversal intersection with both $\mathbb{R} \times S_{1,r}$ and $\mathbb{R} \times S_{2,r}$. If r is sufficiently generic, then both $C \cap (\mathbb{R} \times S_{1,r})$ and $C \cap (\mathbb{R} \times S_{2,r})$ will be disjoint from both $\mathbb{R} \times \gamma_{p+}$ and $\mathbb{R} \times \gamma_{p-}$. Assume that this is the case also.

Note that the intersection of C with $\mathbb{R} \times S_{1,r}$ has precisely one unbounded component; this is a smooth arc with s unbounded from above and below. The function s restricts to this arc as an affine parameter where $|s| \gg 1$. The constant s points on the arc converge in $S_{1,r}$ as $s \rightarrow \infty$ to the point where the integral curve segment γ_{p+} intersects $S_{1,r}$, and they converge as $s \rightarrow -\infty$ to the point where γ_{p-} intersects $S_{1,r}$. The intersection of C with $\mathbb{R} \times S_{2,r}$ has a similar arc component.

With the preceding as background, introduce c_1 and c_2 to denote the respective images in \mathcal{H}_p^+ of $C \cap (\mathbb{R} \times S_{1,r})$ and $C \cap (\mathbb{R} \times S_{2,r})$. Neither c_1 nor c_2 intersects γ_{p+} , and neither intersects γ_{p-} . As a consequence, the closure of c_1 adds two points to c_1 ; these are the points where γ_{p+} and γ_{p-} intersect the radius r coordinate ball centered on p_1 . Likewise, the closure of c_2 adds the points where γ_{p+} and γ_{p-} intersect the radius r coordinate ball centered on p_2 .

Define a cycle in \mathcal{H}_p^+ as follows: Start where γ_{p+} intersects $S_{2,r}$ and proceed along γ_{p+} in \mathcal{H}_p^+ until it intersects the boundary of $S_{1,r}$. Then proceed along c_1 until it ends where γ_{p-} intersects $S_{1,r}$. Travel in the reverse direction along γ_{p-} until it intersects $S_{2,r}$. Close up the cycle by proceeding along c_2 to its endpoint where γ_{p+} intersects $S_{2,r}$. This closed cycle is denoted by v_p^C . The cycle v_p^C defines a class in $H_1(\mathcal{H}_p^+ \cup \widehat{\gamma}_p^+ \cup \widehat{\gamma}_p^-)$, and thus an integer. This is the integer m_p^C . The latter does not depend on the precise choice for r .

Given the definition above of m_p^C , what is said in Proposition 4.5 has the following consequence:

- (5-5) • $m_p^C = 0$ if either none or both of $\{\widehat{\gamma}_p^+, \widehat{\gamma}_p^-\}$ are associated to ends of C in \mathcal{H}_p^+ .
- $m_p^C = -1$ if $\widehat{\gamma}_p^+$ is associated to an end of C in \mathcal{H}_p^+ .
- $m_p^C = 1$ if $\widehat{\gamma}_p^-$ is associated to an end of C in \mathcal{H}_p^+ .

To see that the class $(0, (m_p^C)_{p \in \Lambda}) \in H_1(\Sigma) \oplus (\bigoplus_{p \in \Lambda} \mathbb{Z})$ maps via P to $[\Theta_{C+}] - [\Theta_{C-}]$, define a class $M_\Sigma^C \subset H_1(\Sigma)$ as follows: Use $[\gamma_p^C]$ to denote the $p \in \Lambda$ version of the cycle that was described in the preceding paragraph. Then $[\Theta_{C+}] - [\Theta_{C-}] - \sum_{p \in \Lambda} [\gamma_p^C]$ is a closed cycle in the $f^{-1}(1, 2)$ portion of M_δ and so defines a class in $H_1(\Sigma)$. The latter is M_Σ^C , and it follows from the Mayer–Vietoris definition of P that $(M_\Sigma^C, (m_p^C)_{p \in \Lambda})$ is

sent by P to the class $[\Theta_{C_+}] - [\Theta_{C_-}]$. However, the chain $[\Theta_{C_+}] - [\Theta_{C_-}] - \sum_{p \in \Lambda} [\gamma^C_p]$ in $f^{-1}(1, 2) \cap M_r$ is the boundary of $C \cap (f^{-1}(1, 2) \cap M_r)$ and so represents the zero class in $H_1(\Sigma)$.

It is a straightforward task to verify that the assignment $C \rightarrow (m_p^C)_{p \in \Lambda}$ is constant as C varies in its moduli space component.

Part 4 Let C denote an ech-HF subvariety. This part of the subsection associates a set of $1 + b_1$ nonnegative integers to the ech-HF moduli space component that contains C . To define these integers, reintroduce the set $\mathbb{Y} \subset \Sigma \setminus (\tau_+ \cup \tau_-)$ from Parts 6 and 7 of Section 1C. By way of a reminder, \mathbb{Y} contains the fiducial point z_0 and b_1 additional points. As noted in Section 2A, each point is the intersection point between Σ and a closed integral curve of ν . The curve through any given $z \in \mathbb{Y}$ is denoted by $\gamma^{(z)}$. Use $n_C^{(z)}$ to denote the intersection number between the subvariety C and the J -holomorphic cylinder $\mathbb{R} \times \gamma^{(z)}$. This is a nonnegative integer, and the collection $\{n_C^{(z)}\}_{z \in \mathbb{Y}}$ is the desired set of integers. The z_0 version plays a distinguished role at times, and is denoted by n_C .

Part 5 Suppose as before that C is an ech-HF subvariety. Let \hat{v}_+ and \hat{v}_- denote the respective HF cycles that are used to define Θ_{C_+} and Θ_{C_-} . View these as collections of G oriented arcs in the $f \in [1, 2]$ part of M that run from the index 1 critical points of f to the index 2 critical points of f . Let $[\hat{v}_+]$ and $[\hat{v}_-]$ denote the corresponding 1-chains. The next paragraph defines from the subvariety C a 2-chain in M with boundary the closed 1-cycle $[\hat{v}_+] - [\hat{v}_-]$. Here and in what follows, a “chain” in M or Y of a given dimension indicates that the given subset can be viewed as a linear combination of simplices of the indicated dimension, and as such is viewed as a chain in the appropriate singular homology chain complex. Let $H_2(M; [\hat{v}_+] - [\hat{v}_-])$ denote the affine space (over \mathbb{Z}) of equivalence classes of 2-chains in M with boundary $[\hat{v}_+] - [\hat{v}_-]$; two chains with boundary $[\hat{v}_+] - [\hat{v}_-]$ are equivalent when their difference is a boundary. The chain constructed in the next paragraph from C will be an element in $H_2(M; [\hat{v}_+] - [\hat{v}_-])$.

To define the desired 2-chain from C , fix $p \in \Lambda$ and reintroduce the spheres $S_{1,r}$ and $S_{2,r}$ from Part 3. Remove the $|u| \leq R + \ln r$ portion of C 's intersection with $\mathbb{R} \times \mathcal{H}_p^+$ so as to obtain a subvariety with boundary on $\mathbb{R} \times S_{1,r}$ and $\mathbb{R} \times S_{2,r}$. Now view $S_{1,r}$ as a sphere in M centered on the index 1 critical point of f . Cone the boundary curve $c_1 \subset S_{1,r}$ to the index 1 critical point in question by a cone that lies

where $f \geq 1$ in the radius r coordinate ball centered on this index 1 critical point. Do the same for the boundary curve $c_2 \subset S_{2,r}$. Do this for all $p \in \Lambda$. The resulting chain with C 's image in the complement of the radius r coordinate balls about the index 1 and index 2 critical points of f defines a 2-chain whose boundary is a closed 1-chain which is homologous in the radius $c_0\delta$ tubular neighborhood of $\bigcup_{v \in \Theta_{C_+} \cup \Theta_{C_-}} v$ to the cycle $[\hat{v}_+] - [\hat{v}_-]$. Choose a 2-chain that gives such a homology and add it to the chain that was just defined by the coning of C 's image in M_δ . Introduce by way of notation $[C]_M$ to denote the image of the latter 2-chain in $H_2(M; [\hat{v}_+] - [\hat{v}_-])$.

The intersection number of this chain with the closed integral curve of v through the point z_0 in Σ is the integer n_C that is defined in Part 4.

Suppose that C and C' are two ech-HF subvarieties such that Θ_{C_+} and $\Theta_{C'_+}$ are defined using the same HF cycle, and such that Θ_{C_-} and $\Theta_{C'_-}$ are also defined using the same HF cycles. The difference $[C]_M - [C']_{M'}$ is a closed 2-cycle in M .

5B Integral bounds

Let C denote an ech-HF subvariety. This subsection supplies bounds on the integral over C of w , and the integral of $ds \wedge \hat{a}$ over compact subsets of C . These are stated in the next proposition. The proposition refers to an integer that can be associated to any given ech-HF moduli space component. This integer is called the *ech index*. Definition 2.14 in [7] supplies a formula for the index. This index is denoted here by $I_{\text{ech}}(C)$. Note in this regard that the image of C in Y via the projection from $\mathbb{R} \times Y$ defines a class in the relative homology group $H_2(Y; [\Theta_{C_+}] - [\Theta_{C_-}])$.

Proposition 5.1 *There exists $\kappa > 1$ that is independent of the choice for J and the data R, χ_0 , and δ that is used to construct Y , and which has the following significance: Let C denote an ech-HF subvariety and let $I_{\text{ech}}(C)$ denote the ech index of its moduli space component. Then:*

- $0 \leq \int_C w \leq \kappa(1 + |I_{\text{ech}}(C)|)$.
- $(1 - 2\delta^2)G < \int_{C \cap ([s, s+1] \times M_\delta)} ds \wedge \hat{a} < G$.

There exists $\kappa_J \geq 1$ that depends on J and the data R, χ_0 and δ and which has the following significance: Let C be as described above. Then

- $0 \leq \int_{C \cap ([s, s+1] \times Y)} ds \wedge \hat{a} \leq \kappa_J(1 + |I_{\text{ech}}(C)|) + \sum_{(y,1) \in \Theta_{C_-}} \int_Y a$.

With regards to the proof, note that no generality is lost by considering only that case where C has no components from the set $\{\mathbb{R} \times \widehat{\gamma}_p^+, \mathbb{R} \times \widehat{\gamma}_p^-\}_{p \in \Lambda}$. This assumption is made implicitly in all that follows, lemmas included.

Part 1 of what follows supplies a key input to the proof, Lemma 5.2. Part 2 of what follows contains the proof Proposition 5.1. Parts 3 and 4 of the subsection contain the proof of Lemma 5.2. Note that the lower bounds by zero in the first and third bullets follow from the fact that J is compatible with $ds \wedge \widehat{a} + w$ and so no more will be said about them.

Part 1 The upcoming Lemma 5.2 supplies a key input to the proof of Proposition 5.1. This lemma refers to the integer n_C that is defined in Part 4 of Section 5A. The lemma also refers to an element in the $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ summand of (5-2) that is associated to C 's moduli space component. What follows directly gives the definition.

To start things off, reintroduce the respective 1-chains $[\widehat{v}_+]$ and $[\widehat{v}_-]$ that are defined by the HF cycles which are used in Proposition 2.8 to label Θ_{C_+} and Θ_{C_-} . As noted in Part 5 of Section 5A, the existence of C requires that $[\widehat{v}_+] - [\widehat{v}_-]$ is zero in $H_1(M; \mathbb{Z})$. The vanishing of this has the following implication: There exists a 2-chain in the $f^{-1}([1, 2])$ part of M whose boundary is $[\widehat{v}_+] - [\widehat{v}_-]$ that depends only on the Heegaard Floer data. Moreover, 2-chains with certain desirable properties can be constructed using the fact that the components of the 1-cycle $[\widehat{v}_+] - [\widehat{v}_-]$ are smoothly embedded arcs that intersect in pairs at their endpoints, these being index 1 and index 2 critical points of f . In particular, there exists a smooth map from an oriented surface with boundary into M whose image represents the 2-cycle with the following properties, where we use S to denote the surface and $\mathfrak{f}: S \rightarrow M$ to denote the map:

- (5-6) • The boundary of S has $2G$ marked points. These are mapped in a 1-to-1 fashion to the index 1 and index 2 critical points of f .
- The complement in S of the $2G$ marked boundary points is mapped to the $f^{-1}(1, 2)$ part of M .
 - The complement in the boundary of these $2G$ marked point is identified by \mathfrak{f} with the interiors of the segments that form $\widehat{v}_+ \cup \widehat{v}_-$.
 - The map \mathfrak{f} is transversal to the boundary of each radius $r \in [\delta_*, 8\delta_*]$ coordinate ball about the index 1 and index 2 critical points of f , and the image of \mathfrak{f} intersects each such boundary sphere as an immersed arc.

- Let $\iota \subset \partial S$ denote a component of the complement of the 2G marked boundary points. The map f restricts to a neighborhood in S of ι as an embedding which has the vector field v tangent its image.

Let $[S]$ denote the 2-chain in M given by $f(S)$.

Fix $p \in \Lambda$ and write p as (p_1, p_2) as in Part 3 of Section 5A. The Mayer-Vietoris construction of (5-2) identifies the left-hand \mathbb{Z} in the p -labeled summand of $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ with $H_1(T_{p_1^+})$. Likewise, it identifies the right-most \mathbb{Z} with $H_1(T_{p_2^-})$. The specification of an integer in each index 1 critical point version of $H_1(T_{p_1^+})$ and in each index 2 critical point version of $H_1(T_{p_2^-})$ gives an element in the $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ summand. The 2-chain $[S]$ is used to define these integers.

Reintroduce the notation from Section 5A, and in particular the chain $[C]_M$, the constant $r \in (\delta_*, 7\delta_*)$ and the spheres $S_{1,r}$ and $S_{2,r}$. The 2-chain $[C]_M - [S]$ is a closed 2-cycle whose intersection with the sphere $S_{1,r}$ lies in its $1 - 3 \cos^2 \theta$ part. As such, it defines a closed 1-cycle in $T_{p_1^+}$ and thus a class in $H_1(T_{p_1^+})$. Use $n_{p_1}^C$ to denote the corresponding integer. The intersection of $[C]_M - [S]$ with $S_{2,r}$ defines an analogous integer, which is denoted by $n_{p_2}^C$. The desired element $(n_{p_1}^C, n_{p_2}^C)_{p \in \Lambda} \in \bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ is the desired element.

Lemma 5.2 *Fix $I \geq 1$; there exists $\kappa > 1$ that depends only on the Heegaard Floer data and which has the following significance: Suppose that C is an ech-HF subvariety with $I_{\text{ech}}(C) \leq I$. Then $\sum_{z \in \mathbb{Y}} n_C^{(z)} + \sum_{p \in \Lambda} (|n_{p_1}^C| + |n_{p_2}^C|) \leq \kappa$.*

Lemma 5.2 is proved in Parts 3 and 4 of this subsection. Assume this lemma for now.

Part 2 This part of the subsection contains the:

Proof of Proposition 5.1 The proof has four steps. Steps 1–3 prove the assertion made by the first bullet of Proposition 5.1. Step 4 proves the assertions that are made by the second and third bullets. Appeals to Lemma 5.2 enter only in Steps 1 and 3.

Step 1 As noted in (1-19), the 2-form w can be written as da on the complement of the intersection of C with the part of M_δ that corresponds to the $(1, 2) \times (\bigcup_{z \in \mathbb{Y}} D_z)$ part of $(1, 2) \times \Sigma$. Granted this, and given that C has strictly positive intersection number with each $z \in \mathbb{Y}$ version of $\mathbb{R} \times \gamma^{(z)}$, an integration by parts with (1-16) writes

$$(5-7) \quad \int_C w = \sum_{\gamma \in \Theta_{C^+}} \int_\gamma a - \sum_{\gamma \in \Theta_{C^-}} \int_\gamma a + 2 \sum_{z \in \Lambda} n_C^{(z)} R_z.$$

Each $z \in \mathbb{Y}$ version of $n_C^{(z)}$ is nonnegative, and Lemma 5.2 says that each is bounded by a constant that depends only on $I_{\text{ech}}(C)$ and the Heegaard Floer data. Thus, the right-most term in (5-7) is also bounded by a constant that depends only on the Heegaard Floer data.

The integrals in (5-7) over the various closed orbits of v can be broken into their contributions from the parts in the 1–handles and the parts in M_δ . It is a consequence of what is said in Corollary 2.6 that the M_δ contribution to the integrals in (5-7) is no greater than a constant the depends solely on the Heegaard Floer data.

Fix $p \in \Lambda$. The contribution to the integrals on the right side of (5-7) from \mathcal{H}_p can be written as a sum of various terms. There is, first of all, a contribution from each integral curve in Θ_{C_+} or Θ_{C_-} from the set $\{\hat{\gamma}_p^+, \hat{\gamma}_p^-\}$. A look at Property 3 in Section 1E finds that each such curve in Θ_{C_+} contributes $\frac{4\sqrt{2}}{3}\pi(\chi_0 + 4e^{-2R})$ to the right-hand side of (5-7), and each in Θ_{C_-} contributes minus this amount to the right-hand side of (5-7).

Let γ_{p+} and γ_{p-} denote the segments of the integral curves in Θ_{C_-} and Θ_{C_+} that cross \mathcal{H}_p^+ . The remaining contribution to the right-hand side of (5-7) from \mathcal{H}_p is given by

$$(5-8) \quad \int_{\gamma_{p+}} a - \int_{\gamma_{p-}} a.$$

It follows from the depiction of a in (1-8)–(1-10) and the definition of the function \hat{h} in (4-1) that a can be written as $a = \chi(u)(1 - 3 \cos^2 \theta) du - \sqrt{6} \hat{h} d\phi + df_*$ near γ_{p+} and γ_{p-} . This depiction of a can be used to write integral a over γ_{p+} as

$$(5-9) \quad \int_{\gamma_{p+}} \chi(1 - 3 \cos^2 \theta) du - \sqrt{6} \hat{h}_+ \Delta_+,$$

where \hat{h}_+ here denotes the constant value of the function \hat{h} on γ_{p+} , and where Δ_+ denotes the change in the angle ϕ along γ_{p+} . The latter is given by the integral in the fourth bullet of Lemma 2.2. (It follows from (c) of the second bullet of Lemma 2.2 that the integral of df_* over γ_{p+} is zero.) The integral of a over γ_{p-} is given by the analogous expression in terms of an integral and the corresponding \hat{h}_- and Δ_- .

Step 2 This step rewrites (5-8). This is done by first writing (5-8) as the difference between the γ_{p+} and γ_{p-} versions of (5-9). To start the task of rewriting the latter expression, let γ denote for the moment a segment of an integral curve of v in \mathcal{H}_p^+ with $|\hat{h}| < (\chi_0 + 2e^{-2R})\frac{2}{3\sqrt{3}}$. Recall from Lemma 2.2 that this last bound is the

necessary and sufficient condition for γ to cross \mathcal{H}_p^+ from its $u = -R - \ln(7\delta_*)$ end to its $u = R + \ln(7\delta_*)$ end. View θ along the segment γ as a function of u that depends implicitly on the constant value of \hbar along the segment. The assignment to γ of the integral $\int_\gamma \chi(1 - 3 \cos^2 \theta) du$ defines a function of \hbar . Let $\hbar \rightarrow X(\hbar)$ denote this function. Use the chain rule with (4-1) to see that $\frac{d}{d\hbar} X = \sqrt{6} \Delta$, where $\Delta = \Delta(\hbar)$ is given by the integral in the fourth bullet of Lemma 2.2. By way of an explanation for the derivative identity, note that the dependence of $X(\hbar) = \int_\gamma \chi(1 - 3 \cos^2 \theta) du$ on \hbar can be seen by using coordinates (u, \hbar) on \mathcal{H}_p^+ in lieu of (u, θ) . This is to say that θ should be viewed as being a function u and \hbar . This writes the integral $\int_\gamma \chi(1 - 3 \cos^2 \theta) du$ that defines $X(\hbar)$ as an integral between the fixed u limits. Thus, $X(\hbar) = \int_{-R-\ln \delta}^{R+\ln \delta} \chi(u)(1 - 3 \cos^2(\theta(u, \hbar))) du$. Granted this depiction of X , differentiating with respect to \hbar can be done under the integral. This gives the integral with respect to u of $6 \cos \theta \sin \theta \frac{\partial \theta}{\partial \hbar}$. Writing $\frac{\partial \theta}{\partial \hbar} = \left(\frac{\partial \hbar}{\partial \theta}\right)^{-1}$ leads directly to the identity $\frac{d}{d\hbar} X = \sqrt{6} \Delta$ because (4-1) implies that $\frac{\partial \hbar}{\partial \theta} = -f(u) \sin \theta (1 - 3 \cos^2 \theta)$.

Use $\frac{d}{d\hbar} X = \sqrt{6} \Delta$ to write the difference between the γ_{p+} and γ_{p-} versions of (5-9) as

$$(5-10) \quad \sqrt{6} \left(\int_{\hbar_-}^{\hbar_+} \Delta d\hbar - \hbar_+ \Delta_+ + \hbar_- \Delta_- \right).$$

With (5-10) in hand, use the mean value theorem to write the integral that appears in (5-10) as $\Delta(\hbar_0)(\hbar_+ - \hbar_-)$, where \hbar_0 is a number between \hbar_- and \hbar_+ . This done, (5-10) can be written as

$$(5-11) \quad \sqrt{6} (\Delta(\hbar_0) - \Delta_+) \hbar_+ - \sqrt{6} (\Delta(\hbar_0) - \Delta_-) \hbar_-.$$

This last expression is the desired rewrite of (5-8).

Step 3 Use the expression in the fourth bullet of Lemma 2.2 to see that the derivative of the function $\hbar \rightarrow \Delta(\hbar)$ is given by

$$(5-12) \quad -\sqrt{6} \int_{[-R-\ln \delta, R+\ln \delta]} \frac{\chi(u)}{f(u)^2} \frac{1 + 3 \cos^2(\theta(u))}{(1 - 3 \cos^2(\theta(u)))^3} du.$$

This implies in particular that the function $\hbar \rightarrow \Delta(\hbar)$ is a monotonic function of \hbar and so the number $\Delta(\hbar_0)$ that appears in (5-11) is between Δ_- and Δ_+ . Thus, both $|\Delta(\hbar_0) - \Delta_+|$ and $|\Delta(\hbar_0) - \Delta_-|$ are bounded by $|\Delta_+ - \Delta_-|$. It follows as a consequence that the absolute value of the expression in (5-11) is no greater than $c_0 \chi_0 |\Delta_+ - \Delta_-|$.

Use Proposition 2.8 to label Θ_{C_+} and Θ_{C_-} . This done, introduce \mathfrak{k}_{p_+} and \mathfrak{k}_{p_-} to denote the respective integers that are used in these labels for p 's factor of $\mathbb{Z} \times O$. Given that $|m_p| \leq 1$, the bound given by Lemma 5.2 implies that $|\mathfrak{k}_{p_+} - \mathfrak{k}_{p_-}|$ is bounded by a constant that depends only on $I_{ech}(C)$ and the Heegaard Floer data. This bound implies that $|\Delta_+ - \Delta_-|$ is also bounded by a constant that depends only on $I_{ech}(C)$ and the Heegaard Floer data. This fact, with what is said in the preceding paragraph, implies that the absolute value of what is written in (5-8) is bounded by a constant that depends solely on $I_{ech}(C)$ and the Heegaard Floer data. This fact, with what is said in the first step, proves the assertion made by the first bullet of Proposition 5.1.

Step 4 To prove the second bullet, note first that C intersects $\mathbb{R} \times M_\delta$ only where $f \in (1, 2)$. This understood, write this part of M_δ as $(1, 2) \times \Sigma$. The 1-form \hat{a} appears in this guise as dt where t is the Euclidean coordinate on the $(1, 2)$ factor. The integral depicted in the second bullet is that of the pullback of $ds \wedge dt$ via the projection map to $[s, s + 1] \times (1, 2)$ of the intersection between C and the part of $[s, s + 1] \times (1, 2) \times \Sigma$ that lies in $[s, s + 1] \times M_\delta$. Each constant (s, t) slice of $\mathbb{R} \times (1, 2) \times \Sigma$ is pseudoholomorphic, and so it follows from Lemma 4.2 that C has at most G intersection points (counting multiplicities) with each such slice. This implies that the integral in question is at most G .

By way of comparison, the projection induced map to $[s, s + 1] \times (1 + \delta^2, 2 - \delta^2)$ from C 's intersection with the $[s, s + 1] \times M_\delta$ part of $[s, s + 1] \times (1 + \delta^2, 2 - \delta^2) \times \Sigma$ is a proper map of degree G and so the integral of $ds \wedge \hat{a}$ over $C \cap ([s, s + 1] \times (1 + \delta^2, 2 - \delta^2) \times \Sigma)$ is exactly $(1 - 2\delta^2)G$. The latter integral is strictly less than the integral of $ds \wedge \hat{a}$ over $C \cap ([s, s + 1] \times M_\delta)$ because $ds \wedge \hat{a}$ is nonnegative on TC .

To prove the final assertion of Proposition 5.1, use the second bullet of (3-4) to bound

$$(5-13) \quad \int_{C \cap ([s, s+1] \times Y)} ds \wedge \hat{a} = r \int_{C \cap ([s, s+1] \times Y)} (ds \wedge a + w).$$

To exploit (5-13), write

$$(5-14) \quad \int_{C \cap ([s, s+1] \times Y)} ds \wedge a = \int_{[s, s+1]} \left(\int_{C \cap (\{s'\} \times Y)} a \right) ds',$$

and integrate by parts to write

$$(5-15) \quad \int_{C \cap (\{s'\} \times Y)} a = \int_{C \cap ((-\infty, s'] \times Y)} da + \sum_{(\gamma, 1) \in \Theta_{C_-}} \int_\gamma a.$$

By construction, $da = Hw$, where H is a smooth function on Y . This understood, what is written on the right-hand side of (5-13) is no greater than

$$(5-16) \quad r_* \int_C w + \sum_{(\gamma, 1) \in \Theta_{C_-}} \int_\gamma a,$$

where $r_* \geq 1$. An appeal to the first bullet in Proposition 5.1 finishes the proof. \square

Part 3 This part of the subsection proves Lemma 5.2 modulo an auxiliary result, this being Lemma 5.3. Lemma 5.3 is proved in Part 4 of this subsection. Lemma 5.3 is one of two key inputs to the proof. The argument for Lemma 4.14 of [17] is the second key input.

Proof of Lemma 5.2 The proof has five steps.

Step 1 Let C denote a given ech-HF subvariety and let Z denote a 2-chain in $M_\delta \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$ that defines an element in $H_2(Y; [\Theta_{C_+}] - [\Theta_{C_-}])$. A chain of this sort has an associated ech index. Let $[C]_Y$ denote the 2-chain that is defined by C 's image in Y via the projection from $\mathbb{R} \times Y$. The chain defined by $[C]_Y - Z$ is a closed 2-cycle in Y . Hutchings explains in Section 2 of [7] why

$$(5-17) \quad I_{\text{ech}}(C) - I_{\text{ech}}(Z) = \langle c_1, [C]_Y - Z \rangle,$$

where the notation is such that $\langle \cdot, \cdot \rangle: H^2(Y; \mathbb{Z}) \otimes H_2(Y; \mathbb{Z}) \rightarrow \mathbb{Z}$ denotes the canonical pairing and c_1 denotes the first Chern class of the chosen $\text{Spin}^{\mathbb{C}}$ structure on Y , this being the class depicted in (1-29).

Lemma 5.3 *There is a constant κ that depends only on the Heegaard Floer data, and a 2-chain $Z \in H_2(Y; [\Theta_{C_+}] - [\Theta_{C_-}])$ with the following properties:*

- Z intersects only the union of $f^{-1}(1, 2)$ part of M_δ and the $p \in \Lambda$ versions of \mathcal{H}_p^+ .
- The part of Z that lies in M_δ is obtained from the chain $[S]$ by adding a chain with support in the radius $\kappa\delta$ tubular neighborhood of the integral curves of \mathfrak{v} from the HF cycles \hat{v}_+ and \hat{v}_- .
- $|I_{\text{ech}}(Z)| \leq \kappa$.

This lemma is proved in Part 4 of the subsection. The subsequent steps refer to a chosen 2-chain Z that is described by Lemma 5.3.

Step 2 It follows from (1-29) that the first Chern class c_1 can be represented by a closed form with support disjoint from $\bigcup_{p \in \Lambda} \mathcal{H}_p$. Observations from Parts 4 and 6 of Section 1C are used to construct such a form whose intersection with the $f^{-1}((1, 2))$ portion of M_δ is a 2–form that is obtained from a form on Σ via the Lie transport defined by \mathfrak{v} . This form is denoted in what follows by w_{Σ_0} . It is nonnegative on Σ and it has the following properties, where we use terminology from Parts 4 and 6 of Section 1C:

- (5-18) • The form w_{Σ_0} has support in the complement of $T_+ \cup T_-$.
- The integral of w_{Σ_0} over any fundamental domain is strictly positive.
 - The integral of w_{Σ_0} over Σ is 2.
 - The integral of w_{Σ_0} any given periodic domain in Σ is equal to the value of c_{1M} on the corresponding homology class in $H_2(M; \mathbb{Z})$.

What is said in Lemma 5.3 about Z and what was just said about the representative form for the class c_1 have the following consequence: The pairing on the right-hand side of (5-17) can be computed as described next. Use Lie transport by the pseudogradient vector field \mathfrak{v} to map $[C]_M$ and $[S]$ into Σ . Their respective images define 2–chains in Σ ; these are denoted in what follows by C_Σ and S_Σ . The right-hand side of (5-17) is the integral of the 2–form w_{Σ_0} over $C_\Sigma - S_\Sigma$. Note in this regard that neither C_Σ nor is a closed 2–chain. Even so, their respective boundaries lie in $T_+ \cup T_-$, which is disjoint from the support of w_{Σ_0} . The fact that their respective boundaries lie in $T_+ \cup T_-$ implies S_Σ that both have a well-defined intersection number with each point in \mathbb{Y} . The intersection number between C_Σ and any given $z \in \mathbb{Y}$ is the integer $n_C^{(z)}$, and the absolute value of that between S_Σ and z is bounded solely by the Heegaard Floer data.

The difference $C_\Sigma - S_\Sigma$ is not closed either, because $[C]_M - [S]$ intersects the critical index 1 and index 2 critical points of f . More to the point, C_Σ and S_Σ can be modified in $T_+ \cup T_-$ so that the resulting chains, C_{Σ_0} and S_{Σ_0} , are such that:

- (5-19) • The intersection number between $C_{\Sigma_0} - S_{\Sigma_0}$ and any given $z \in \mathbb{Y}$ differs from $n_C^{(z)}$ by a constant that depends only on the Heegaard Floer data.
- $\partial(C_{\Sigma_0} - S_{\Sigma_0}) = \sum_{p \in \Lambda} (n_{p_1}^C Z_{p_1} + n_{p_2}^C Z_{p_2})$.

The integral of w_{Σ_0} over $C_{\Sigma_0} - S_{\Sigma_0}$ also computes the right-hand side of (5-17) because the support of w_{Σ_0} is disjoint from $T_+ \cup T_-$.

Lemma 5.2 asserts a particular relationship between $I_{\text{ech}}(C)$ on the one hand and the integers that form the sets $\{n_C^{(z)}\}_{z \in \mathbb{Y}}$ and $\{(n_{p_1}^C, n_{p_2}^C)\}_{p \in \Lambda}$ on the other. The chain $C_{\Sigma_0} - S_{\Sigma_0}$ is the bridge that brings these two seemingly disparate notions together by virtue of two facts: the boundary of $C_{\Sigma_0} - S_{\Sigma_0}$ defines the two sets of integers via (5-19), and the integral of w_{Σ_0} over this 2-cycle computes (5-17). This bridge leads to a proof of Lemma 5.2 using what are essentially the same arguments that are used to prove Lemma 4.14 in [17]. The remaining steps supply these arguments.

Step 3 If z is a point of \mathbb{Y} , let $\gamma^{(z)}$ again denote the corresponding closed integral curve of v through z . Use $[\gamma^{(z)}]$ to again denote the corresponding closed 1-cycle and $[\gamma^{(z)}]^{\text{Pd}}$ to denote the Poincaré dual in $H^2(Y; \mathbb{Z})$ of the resulting homology class. The span of the set $\{[\gamma^{(z)}]^{\text{Pd}}\}_{z \in \mathbb{Y}} \in H^2(Y; \mathbb{Z})$ is a $1+b_1$ -dimensional subspace that is dual via the pairing $\langle \cdot, \cdot \rangle$ to the $H_2(\mathcal{H}_0; \mathbb{Z}) \oplus H_2(M; \mathbb{Z})$ summand in (1-27). The set \mathbb{Y} also has its fiducial point, z_0 , and the span of the set $\{[\gamma^{(z)}]^{\text{Pd}} - [\gamma^{(z_0)}]^{\text{Pd}}\}_{z \in \mathbb{Y} - z_0}$ is dual to the $H_2(M; \mathbb{Z})$ summand in (1-27).

It proves useful at this point to introduce some specific cycles to represent the homology in the summand $H_2(M; \mathbb{Z}) \oplus H_2(\mathcal{H}_0; \mathbb{Z})$. In particular, the classes in the $H_2(M; \mathbb{Z})$ summand are represented by periodic domains in Σ as done in the proof of Lemma 1.1. Let \mathcal{P} denote such a domain, and let $h(\mathcal{P}) \in H_2(M; \mathbb{Z})$ denote the corresponding homology class. The cycle $[\Sigma]$ of the Heegaard surface completes the desired set of cycles.

Each cycle from the set $\{[\gamma^{(z)}]^{\text{Pd}}\}_{z \in \mathbb{Y}}$ has pairing 1 with $[\Sigma]$. The elements from $\{[\gamma^{(z)}]^{\text{Pd}} - [\gamma^{(z_0)}]^{\text{Pd}}\}_{z \in \mathbb{Y} - z_0}$ have pairing 0 with $[\Sigma]$, and so the span of this set is dual to the span of the classes defined by the periodic domains. This understood, fix a basis, $\{\mathcal{P}_z\}_{z \in \mathbb{Y} - z_0}$, of periodic domains such that the corresponding set $\{h(\mathcal{P}_z)\}_{z \in \mathbb{Y} - z_0}$ is dual (over \mathbb{Q}) to the elements in the set $\{[\gamma^{(z)}]^{\text{Pd}} - [\gamma^{(z_0)}]^{\text{Pd}}\}_{z \in \mathbb{Y} - z_0}$.

Let \mathcal{P} again denote a periodic domain. The boundary of \mathcal{P} is a closed 1-cycle in Σ that can be written as

$$(5-20) \quad \partial\mathcal{P} = \sum_{p \in \Lambda} (n_{p_1}^{\mathcal{P}} Z_{p_1} + n_{p_2}^{\mathcal{P}} Z_{p_2}).$$

The resulting element $(n_{p_1}^{\mathcal{P}}, n_{p_2}^{\mathcal{P}})_{p \in \Lambda} \in \bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ is mapped by the homomorphism \hat{O} in (5-3) to an element with $H_1(\Sigma)$ entry zero. In fact, the elements in $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ that come from the boundaries of periodic domains span the linear space of such elements.

Step 4 As (5-20) indicates, the class $(n_{p_1}^{\mathcal{P}}, n_{p_2}^{\mathcal{P}})_{p \in \Lambda}$ in $\bigoplus_{p \in \Lambda} (\mathbb{Z} \oplus \mathbb{Z})$ is mapped by the homomorphism \widehat{O} of (5-3) to an element in $H_1(\Sigma) \oplus (\bigoplus_{p \in \Lambda} \mathbb{Z})$ with $H_1(\Sigma)$ entry zero. What follows is a consequence of this and what said in Step 3: the 2-chain $C_{\Sigma_0} - S_{\Sigma_0}$ is homologous rel $C_+ \cup C_-$ to

$$(5-21) \quad (n_C + \kappa_0^S)[\Sigma] + \sum_{z \in \mathbb{Y} \setminus z_0} q_z^C \mathcal{P}_z,$$

where κ_0^S is the intersection number between S_{Σ_0} and $\gamma^{(z_0)}$ and where each $z \in \mathbb{Y} \setminus z_0$ version of q_z^C is an integer. Granted this depiction, use (5-18) to see that the right-hand side of (5-17) is equal to

$$(5-22) \quad 2(n_C + \kappa_0^S) + \sum_{z \in \mathbb{Y} \setminus z_0} q_z^C \left(\sum_{1 \leq i \leq N} \kappa_i^z A_i \right),$$

where the notation is as follows: First, \mathcal{P}_z is written in terms of fundamental domains as $\mathcal{P}_z = \sum_{1 \leq i \leq N} \kappa_i^z \mathcal{D}_i$. Meanwhile, any given $i \in \{1, \dots, N\}$ version of A_i is the integral of w_{Σ_0} over the domain \mathcal{D}_i .

To exploit this formula, note that both C_{Σ_0} and S_{Σ_0} are homologous rel $C_+ \cup C_-$ to integer-weighted sums of periodic domains; for the purposes at hand, C_{Σ_0} can be replaced by a 2-chain of the form $n_C \mathcal{D}_0 + \sum_{1 \leq i \leq N} \kappa_i^C \mathcal{D}_i$ and S_{Σ_0} by $\sum_{0 \leq i \leq N} \kappa_i^S \mathcal{D}_i$. This done, (5-17) and (5-22) imply that

$$(5-23) \quad I_{\text{ech}}(C) = 2n_C A_0 + \sum_{1 \leq i \leq n} \kappa_i^C A_i + \tau,$$

where $|\tau| \leq c_0$ with c_0 determined solely by the Heegaard Floer data. This is so because S , and so S_{Σ_0} , is determined by the Heegaard Floer data, and because of what is said about $I_{\text{ech}}(Z)$ by Lemma 5.3.

Step 5 The 2-chain C_{Σ_0} is an example of what Ozsváth and Szabó call a *positive* 2-chain, which is to say that the integer coefficients $\{\kappa_i^C\}_{1 \leq i \leq N}$ and n_C are non-negative. Here is why: Any given κ_i^C is the intersection number between C and a J -holomorphic subvariety defined in $\mathbb{R} \times M_\delta$. To elaborate, let $z_i \in \mathcal{D}_i$ denote the point described in Part 6 of Section 1C. Let $\gamma^{(i)} \in Y$ denote the segment of the integral curve of v that contains z_i and sits in M_δ . Then κ_i^C is the intersection number between C and $\mathbb{R} \times \gamma^{(i)}$. As noted previously, the integer n_C is nonnegative for the same reason.

The preceding observation, the fact that each $i \in \{0, \dots, N\}$ version of A_i is positive and (5-23) have the following consequence: there exists $c_0 \geq 1$ that is determined

solely by the Heegaard Floer data such that

$$(5-24) \quad |n_C| + \sum_{1 \leq i \leq N} K_i^C \leq c_0(1 + I_{\text{ech}}(C)).$$

What with (5-19), this bound on the elements in the set $\{K_i^C\}_{1 \leq i \leq N}$ and on n_C implies what is asserted by Lemma 5.2 because it gives an a priori bound on the coefficients that appear when writing the 2-chain c_{Σ_0} as $n_C \mathcal{D}_0 + \sum_{1 \leq i \leq N} K_i^C \mathcal{D}_i$. □

Part 4 This part of the subsection contains the proof of Lemma 5.3. The proof starts in a moment. What follows directly is meant to give an idea of the strategy.

The argument for Lemma 5.3 starts by constructing certain elements in $\mathcal{Z}_{\text{ech}, M}$ from the HF cycles \widehat{v}_+ and \widehat{v}_- and the pair (S, f) . These are denoted by Θ_+ and Θ_- . To say more, first write Θ_{C_+} using the notation from Proposition 2.8 as $(\widehat{v}_+, (\xi_{p_+}, O_{p_+})_{p \in \Lambda})$, and likewise write Θ_{C_-} as $(\widehat{v}_-, (\xi_{p_-}, O_{p_-})_{p \in \Lambda})$. The respective elements Θ_+ and Θ_- are defined as

$$(5-25) \quad \Theta_+ = (\widehat{v}_+, (\xi_{p_0}, O_{p_+})_{p \in \Lambda}) \quad \text{and} \quad \Theta_- = (\widehat{v}_-, (0, O_{p_-})_{p \in \Lambda}),$$

where each $p \in \Lambda$ version of ξ_{p_0} is chosen with the help of S and f . The construction is such that $|\xi_{p_0}| \leq c_0$ where c_0 depends only on the Heegaard Floer data.

The preceding bound on the integers $\{\xi_{p_0}\}$ implies that the integral curve of v from either Θ_+ or Θ_- that crosses any given $p \in \Lambda$ version of \mathcal{H}_p does so with a (δ, χ_0, R) -independent bound on the change in the ϕ coordinate. This fact is used to construct a 2-chain $Z_0 \subset Y$ with the following properties:

- $$(5-26) \quad \bullet \quad Z_0 \in H_2(Y; [\Theta_+] - [\Theta_-]).$$
- $|I_{\text{ech}}(Z_0)| \leq c_0$.
 - Z_0 lies in the union of the $f^{-1}(1, 2)$ part of M_δ and the $1 - 3 \cos^2 \theta \geq 0$ part of each $p \in \Lambda$ version of \mathcal{H}_p .
 - $Z_0 \cap M_\delta$ differs from $[S] \cap M_\delta$ only in the radius $c_0 \delta$ tubular neighborhood of the elements from \widehat{v}_+ and \widehat{v}_- .

Moreover, the constant c_0 that appears here depends only on the Heegaard Floer data.

With Z_0 in hand, a pair of 2-chains are constructed, one in $H_2(Y; [\Theta_{C_+}] - [\Theta_+])$ and the other in $H_2(Y; [\Theta_-] - [\Theta_{C_-}])$. These are denoted by Z_+ and Z_- . Their properties

are described in a moment. In the meantime, note that the chain $Z = Z_+ + Z_0 + Z_-$ defines an element in $H_2(Y; [\Theta_{C_+}] - [\Theta_{C_-}])$. As noted in Section 2 of [7],

$$(5-27) \quad I_{\text{ech}}(Z) = I_{\text{ech}}(Z_+) + I_{\text{ech}}(Z_0) + I_{\text{ech}}(Z_-).$$

What follows says something about Z_+ and Z_- :

- (5-28) • Both lie in the union of the $f^{-1}(1, 2)$ part of M_δ with $\bigcup_{p \in \Lambda} \mathcal{H}_p$.
- The parts of Z_+ and Z_- in M_δ lie in the respective radius c_0 tubular neighborhoods of the curves from \widehat{v}_+ and \widehat{v}_- .
 - $|I_{\text{ech}}(Z_+)| + |I_{\text{ech}}(Z_-)| \leq c_0$.

As in (5-25), the constant c_0 depends only on the Heegaard Floer data. Whereas the calculation of $I_{\text{ech}}(Z_0)$ is straightforward, those of $I_{\text{ech}}(Z_+)$ and $I_{\text{ech}}(Z_-)$ are quite delicate when $\sup_{p \in \Lambda} |\mathfrak{k}_{p+}|$ and $\sup_{p \in \Lambda} |\mathfrak{k}_{p-}|$ are large because they involve a cancellation of two terms that are on the order of these numbers.

The chain Z as just described satisfies the requirements of Lemma 5.3.

Proof of Lemma 5.3 The proof has eight steps.

Step 1 The element Θ_- is given in (5-14). The element $\Theta_+ \in \mathcal{Z}_{\text{ech}, \mathcal{M}}$ is given by the formula in (5-25) modulo the definition of the integers $\{\mathfrak{k}_{p_0}\}_{p \in \Lambda}$. This step defines these integers. To start, suppose that $\Theta = (\widehat{v}_+, (\mathfrak{k}_{p_0}, O_{p_+})_{p \in \Lambda})$ is a given element in $\mathcal{Z}_{\text{ech}, \mathcal{M}}$. Fix $r \in (\delta_*, 7\delta_*)$ and fix $p \in \Lambda$. Write p as (p_1, p_2) and reintroduce $S_{1,r}$ and $S_{2,r}$ to denote the respective boundary spheres of the radius r coordinate balls centered around p_1 and p_2 . Let c_1^S and c_2^S denote the respective intersections between $f(S)$ and the spheres $S_{1,r}$ and $S_{2,r}$.

Let γ_{p+}^Θ denote the segment of the integral curve of v from Θ that crosses \mathcal{H}_p^+ and let γ_{p-}^Θ denote the corresponding integral curve of v from Θ_- . The intersection point of γ_{p+}^Θ with $S_{1,r}$ has distance at most $c_0\delta$ from one endpoint of c_1^S and that of γ_{p-}^Θ has distance at most $c_0\delta$ from the other endpoint. The intersection points of γ_{p+}^Θ and γ_{p-}^Θ with $S_{2,r}$ are similarly close to the respective endpoints of c_2^S .

Granted the preceding, define a 1-cycle in \mathcal{H}_p^+ as follows: Start where γ_{p+}^Θ intersects $S_{2,r}$ and proceed along γ_{p+}^Θ in \mathcal{H}_p^+ until it intersects the boundary of $S_{1,r}$. Then proceed along an arc of length at most $c_0\delta$ in this sphere to the nearby endpoint of c_1^S . Continue along c_1^S until it ends, and then proceed along an arc of length at most $c_0\delta$ to the point where γ_{p-}^Θ intersects $S_{1,r}$. Travel in the reverse direction along γ_{p-}^Θ until it

intersects $S_{2,r}$. Proceed next along an arc of length at most $c_0\delta$ to the nearby endpoint of c_2^S , and then along c_2^S to its endpoint near $\gamma_{p+}^\ominus \cap S_{2,r}$. Close up the cycle by taking an arc of length at most $c_0\delta$ to $\gamma_{p+}^\ominus \cap S_{2,r}$. This closed 1-cycle is denoted by ν_p^S .

The 1-cycle ν_p^S determines a class in $H_1(\mathcal{H}_p^+ \cup \widehat{\gamma}_p^+ \cup \widehat{\gamma}_p^-)$, and thus an integer. In particular, there is a unique choice for \mathfrak{k}_{p_0} that makes this integer equal to the integer $m_p^C \in \{-1, 0, 1\}$ that is defined in (5-5). The corresponding set of such choices is used in (5-25) to define the element Θ_+ . Note in this regard that the set $\{\mathfrak{k}_{p_0}\}_{p \in \Lambda}$ is determined up to an error of size at most ± 1 solely by the pair S and f , and thus solely by the Heegaard Floer data.

Step 2 This step describes the chain Z_0 . Consider first its intersection with M_r , this being the complement in M of the radius r coordinate spheres centered at the index 1 and index 2 critical points of f . The chain here is the image of S by a map that is constructed by modifying f near the boundary of S .

To say more, let ν denote for the moment a segment $M_r \cap \widehat{\nu}_+$. What follows is a consequence of what is said in Corollary 2.6. There exists an isotopy of M_r with the following four properties: The isotopy moves only points in the radius $c_0\delta$ tubular neighborhood of ν , and these points are moved at most distance $c_0\delta$. The isotopy preserves the relevant versions of $S_{1,r}$ and $S_{2,r}$ that contain the endpoints of ν . The end member of the isotopy maps ν to the nearby segment of M_r 's intersection with the integral curves from Θ_+ . The derivatives of the isotopy to order 5 are bounded by a constant that depends only on the Heegaard Floer data. Choose such an isotopy, and likewise choose an isotopy of this sort for the other components $M_r \cap (\widehat{\nu}_+ \cup \widehat{\nu}_-)$. Compose f with the resulting set of 2G isotopies and use f_r to denote the resulting map from $f^{-1}(M_r)$ to M_r . The 2-chain Z_0 intersects M_r as the f_r -image of $f^{-1}(M_r)$.

Fix $p \in \Lambda$ so as to consider the part of Z_0 in \mathcal{H}_p . The simplest case to consider is that where neither $\widehat{\gamma}_p^+$ nor $\widehat{\gamma}_p^-$ is associated to an end of C in \mathcal{H}_p^+ . This implies that the integer m_p^C is equal to zero. Here is a consequence: Use the arcs γ_{p+}^\ominus and γ_{p-}^\ominus and the arcs that form the intersections of $f_r(f^{-1}(M_r))$ with $S_{1,r}$ and $S_{2,r}$ to define a version of the 1-cycle ν_p^S from Step 1. There is a smooth map from $[-1, 1] \times [-R - \ln r, R + \ln r]$ into \mathcal{H}_p^+ with the properties that follow (the map is denoted by f_p):

- (5-29) • f_p sends $[-1, 1] \times (-R - \ln r, R + \ln r)$ to the $|u| < R + \ln r$ part of \mathcal{H}_p^+ .
- f_p restricts to the boundary components $\{-1, 1\} \times [-R - \ln r, R + \ln r]$ as diffeomorphisms onto γ_{p-}^\ominus and γ_{p+}^\ominus , respectively.

- f_p restricts to the boundary components $[-1, 1] \times \{-R - \ln r, R + \ln r\}$ as immersions onto $f_r(f^{-1}(S_{2,r}))$ and $f_r(f^{-1}(S_{1,r}))$, respectively.
- f_p restricts to a neighborhood of $\{-1, 1\} \times [-R - \ln r, R + \ln r]$ as an embedding which has the vector field v tangent to its image.

We let the intersection of Z_0 with \mathcal{H}_p^+ be the f_p -image of $[-1, 1] \times [-R - \ln r, R + \ln r]$.

Consider next the case when the closed integral curve $\widehat{\gamma}_p^+$ is associated to an end of C in \mathcal{H}_p^+ . In this case, $m_p^C = -1$ and the cycle v_p^S is homotopic in \mathcal{H}_p^+ to -1 times the generator of its first homology. This implies that it is homotopic in $\mathcal{H}_p^+ \cup \widehat{\gamma}_p^+ \cup \widehat{\gamma}_p^-$ to $\widehat{\gamma}_p^+$. With this understood, what follows describes the intersection of Z_0 with \mathcal{H}_p^+ .

The 2-chain Z_0 is the image of a map into $\mathcal{H}_p^+ \cup \widehat{\gamma}_p^+$ from the complement in the strip $[-1, 1] \times [-R - \ln r, R + \ln r]$ of a disk centered at $(\frac{1}{2}, 0)$ with radius bounded by $\frac{1}{1000}$. Let $D \subset [-1, 1] \times [-R - \ln r, R + \ln r]$ denote such a disk. The map in this case is also denoted by f_p . This map obeys the conditions in (5-29) with it understood that the domain in the first bullet is the complement of D in $[-1, 1] \times [-R - \ln r, R + \ln r]$. The boundary of the closure of D is mapped by f_p to the cycle $\widehat{\gamma}_p^+$ and this boundary is the only part of the domain that lies where $1 - 3 \cos^2 \theta = 0$. The behavior of f_p near the boundary of D is described by (5-30) below, which parametrizes a neighborhood of the boundary of D by coordinates (ρ, φ) with ρ the Euclidean coordinate on a neighborhood of 0 in $[0, 1)$ and $\varphi \in \mathbb{R}/2\pi\mathbb{Z}$. The boundary is the locus where $\rho = 0$. The angle in $[0, \pi]$ where $\cos \theta = \frac{1}{\sqrt{3}}$ is denoted by θ_* .

(5-30) f_p embeds a neighborhood of the boundary of D in \mathcal{H}_p as the map that sends (ρ, φ) to $(u = 0, \theta = \theta_* - \frac{1}{1000}\rho, \phi = \varphi)$.

The intersection of Z_0 with \mathcal{H}_p in the case when only $\widehat{\gamma}_p^-$ is associated to an end of C in \mathcal{H}_p^+ is the image of a smooth map from the complement in $[-1, 1] \times [-R - \ln r, R + \ln r]$ of a disk centered at $(-\frac{1}{2}, 0)$. The intersection of Z_0 with \mathcal{H}_p in the case when both $\widehat{\gamma}_p^+$ and $\widehat{\gamma}_p^-$ are associated to ends of C in \mathcal{H}_p^+ is the image of a smooth map from the complement in $[-1, 1] \times [-R - \ln r, R + \ln r]$ of a disk centered at $(\frac{1}{2}, 0)$ and a disk centered at $(-\frac{1}{2}, 0)$.

The map in either case is denoted by f_p ; it obeys (5-29) with the domain in the first bullet suitably modified, and it obeys a suitable version of (5-30) near the boundary of a deleted disk.

The map f_p in all cases depends on the parameters (δ, χ_0, R) because $\mathcal{H}_p^+, v, \gamma_{p-}^\ominus$ and γ_{p+}^\ominus depend on them. Even so, f_p can be chosen so that this dependence is irrelevant

to the purposes at hand. This is said formally by the next lemma. To set the stage, suppose that (δ', χ'_0, R') is a data set that can be used as in Section 1 to define Y and its stable Hamiltonian structure. Assume that $\delta' \leq \delta$, $\chi'_0 \leq \chi_0$ and $R' \leq R$. The lemma uses \mathcal{H}'_p^+ to denote the (δ', χ'_0, R') version of the handle \mathcal{H}_p^+ .

Lemma 5.4 *Suppose that f_p is as described by the (δ, χ_0, R) version of (5-30). There exists a diffeomorphism from \mathcal{H}'_p^+ to \mathcal{H}_p^+ whose composition with f_p is described by the (δ', χ'_0, R') version of (5-29).*

Proof This follows in a straightforward manner from two facts: First, the arcs $f_r(f^{-1}(S_{1,r}))$ and $f_r^{-1}(f^{-1}(S_{2,r}))$ are essentially independent of (δ, χ_0, R) . Of course they depend on these parameters where the distance is bounded by $c_0\delta$ from their endpoints, this where the isotopies that give f_r from f are involved. However, given the bounds on the derivatives of these isotopies, this dependence is of no consequence. The second relevant fact is that the integer k_p is independent of (δ, χ_0, R) . With the preceding understood, fix (δ', χ'_0, R') and identify the $|u| \leq R' + \ln r$ part of \mathcal{H}_p with $[-1, 1] \times S^2$ via the map that sends the coordinates (u, θ, ϕ) to $(\frac{1}{R' + \ln r}u, \theta, \phi)$. The (δ', χ'_0, R') versions of the arcs $f_r(f^{-1}(S_{1,r}))$ and $f_r^{-1}(f^{-1}(S_{2,r}))$ appear now as arcs on the respective $u = 1$ and $u = -1$ boundary spheres. Meanwhile, the (δ', χ'_0, R') version of the vector field v appears as a smooth vector field in $[-1, 1] \times S^2$ and the arcs γ_{p+}^\ominus and γ_{p-}^\ominus appear as arcs that cross from the $u = -1$ boundary sphere to the $u = 1$ boundary sphere.

The formulas in Property 3 of Section 1E have the following implication: the vector field v varies smoothly in $[-1, 1] \times S^2$ as the parameters (δ', χ'_0, R') vary. Meanwhile, Lemma 2.2 implies that the arcs γ_{p+}^\ominus and γ_{p-}^\ominus vary via ambient, compactly supported isotopies in the $1 - 3 \cos^2 \theta > 0$ part of $[-1, 1] \times S^2$ as the parameters (δ', χ'_0, R') vary. This fact and what is said about the isotopies that define f_r from f imply what is asserted by the lemma. □

Step 3 This step describes Z_+ . The chain Z_- has, but for notation, the same sort of description. This being the case, only Z_+ is discussed.

Consider first the part of Z_+ in M_r . This part lies in the radius $c_0\delta$ tubular neighborhood of the M_r parts of the integral curves that form $\widehat{\nu}_+$. To say more, let ν denote the M_r part of one of these curves. The segment ν has distance at most $c_0\delta$ from a component of the M_r part of an integral curve of v from Θ_{C_+} , and it has distance at most $c_0\delta$ from an integral curve of v from Θ_+ . Let $\gamma_{M_+}^C$ and $\gamma_{M_+}^\ominus$ denote these

respective segments. These are also integral curves of \mathfrak{v} . Both start on the boundary of the radius r coordinate ball centered at an index 1 critical point of f and both end on the boundary of a similar ball centered at one of f 's index 2 critical points. With $f^{-1}(1, 2)$ viewed in the usual way as $(1, 2) \times \Sigma$, both $\gamma_{M_+}^C$ and $\gamma_{M_+}^\ominus$ project to Σ as points with distance $c_0\delta$ or less from the point in $C_+ \cap C_-$ that is defined by ν . Fix an arc of length $c_0\delta$ or less in Σ from the $\gamma_{M_+}^\ominus$ point to the $\gamma_{M_+}^C$ point. Denote the latter by ι . The rectangle $[1, 2] \times \iota$ when viewed in M intersects M_r as an embedded rectangle that is foliated by integral curves of \mathfrak{v} with one edge on $\gamma_{M_+}^C$, the opposing edge on $\gamma_{M_+}^\ominus$, a third edge on the boundary of the radius r coordinate ball centered on the relevant index 1 critical point of f , and the latter's opposing edge on the boundary of the radius r coordinate ball centered on the relevant index 2 critical point of f . There is an analogous rectangle for each component of $\widehat{\nu}_+$. The union of these rectangles forms the intersection of Z_+ with the closure of M_r .

Fix $\mathfrak{p} \in \Lambda$ so as to consider the intersection between Z_+ and $\mathcal{H}_\mathfrak{p}$. If $\widehat{\gamma}_\mathfrak{p}^+$ is in Θ_{C_+} , then it is also in Θ_+ . In this case, Z_+ contains as a component the degenerate 2-chain that is given by the projection to $\mathcal{H}_\mathfrak{p}$ from $[-1, 1] \times \mathcal{H}_\mathfrak{p}$ of the cylinder $[-1, 1] \times \widehat{\gamma}_\mathfrak{p}^+$. There is an analogous component if $\widehat{\gamma}_\mathfrak{p}^-$ is in Θ_{C_+} and Θ_+ . There is, in all cases, another component of Z_+ in the $|u| \leq R + \ln r$ part of $\mathcal{H}_\mathfrak{p}$, this denoted by $Z_{\mathfrak{p}_+}$. What follows describes the latter.

Let $S_{1,r}$ and $S_{2,r}$ again denote the respective boundary spheres of the radius r coordinate balls that are centered on the index 1 and index 2 critical points from \mathfrak{p} . Let $c_{+1} \subset S_{1,r}$ and $c_{+2} \subset S_{2,r}$ denote the short arcs that form the intersection of these spheres with the M_r part of Z_+ . The concatenation of these two arcs with γ_{C_+} and γ_+ defines a closed cycle which is denoted by $\nu_{\mathfrak{p}_+}$ in $\mathcal{H}_\mathfrak{p}^+$. This cycle will not be homotopically trivial in $\mathcal{H}_\mathfrak{p}^+$ unless $\gamma_{\mathfrak{p}_+}$ and $\gamma_{\mathfrak{p}_+}^\ominus$ have the same Proposition 2.8 integer label, this being $\mathfrak{k}_{\mathfrak{p}_0}$.

The simplest case by far is that here $\nu_{\mathfrak{p}_+}$ is null-homotopic. In this event, there is an embedded rectangle in the $|u| \leq R + \ln r$ part of $\mathcal{H}_\mathfrak{p}^+$ that is foliated by integral curves of \mathfrak{v} with one edge $\gamma_{\mathfrak{p}_+}$, the opposite edge $\gamma_{\mathfrak{p}_+}^\ominus$, a third edge c_{+1} and the opposite edge c_{+2} . This embedded rectangle is $Z_{\mathfrak{p}_+}$.

Suppose next that $\nu_{\mathfrak{p}_+}$ is not null-homotopic in $\mathcal{H}_\mathfrak{p}^+$. It proves convenient in this case to represent $Z_{\mathfrak{p}_+}$ as the image of a submanifold in the $|u| \leq R + \ln r$ part of $[-1, 1] \times \mathcal{H}_\mathfrak{p}^+$ via the projection to the $\mathcal{H}_\mathfrak{p}^+$ factor. The submanifold in question is the graph of a map from $[-1, 1] \times [-R - \ln r, R + \ln r]$ to S^2 . To describe the latter map,

first note that the cycle ν_{p+} can be viewed as a graph of a map from the boundary of $[-1, 1] \times [-R - \ln r, R + \ln r]$ to S^2 . The image of the latter map on the $u = -R - \ln r$ part of the boundary is the arc c_{+2} and its image on the $u = R + \ln r$ part of the boundary is the arc c_{+1} . The map on the boundary component $\{1\} \times [-R - \ln r, R + \ln r]$ is the map $u \rightarrow (\theta(u), \phi(u))$ that parametrizes the v -integral curve γ_{p+} . The map on the component $\{-1\} \times [-R - \ln r, R + \ln r]$ is the corresponding parametrizing map for the v -integral curve γ_{p+} . Use q_p to denote the map. The extension to the interior of the rectangle $[-1, 1] \times [-R - \ln r, R + \ln r]$ can be made so as to have the properties listed below in (5-31). The list uses ξ_{p+} to denote the Proposition 2.8 label of γ_{p+} , and also refers to the intersection number between the image of q_p and the $\theta = 0$ point in S^2 . This intersection number is well defined by virtue of the fact that the boundary of $[-1, 1] \times [-R - \ln r, R + \ln r]$ is mapped by q_p to the $1 - 3 \cos^2 \theta > 0$ portion of S^2 .

- (5-31) • q_p maps to the $\cos \theta > -\frac{1}{\sqrt{3}}$ portion of S^2 .
- q_p restricts to the boundary components of $[-1, 1] \times [-R - \ln r, R + \ln r]$ as described above.
 - q_p restricts to a neighborhood of the boundary of $[-1, 1] \times [-R - \ln r, R + \ln r]$ so that variation of u with constant $[-1, 1]$ factor parametrizes part of an integral curve of v .
 - The image via q_p of $[-1, 1] \times \{0\}$ lies where $1 - 3 \cos^2 \theta > 0$.
 - The image of q_p has $|\xi_{p+} - \xi_{p0}|$ intersections with the $\theta = 0$ point in S^2 . Each such intersection is transverse, and all have the same intersection number.

The construction of a map of this sort is straightforward and so the task is left to the reader.

Step 4 This step summarizes Hutchings' definition of I_{ech} as a prelude to the calculations that are done in the upcoming Step 5. To set the background, suppose that Θ is a finite set whose typical element is a pair (γ, m) with γ a closed integral curve of v and m a positive integer. Assume that no two pairs from Θ share the same closed integral curve. Let Θ' denote another such set, and suppose that $[\Theta] - [\Theta']$ bounds in Y . Let Z denote a 2-chain with boundary $[\Theta] - [\Theta']$. Definition 2.14 in [7] supplies a formula for computing $I_{ech}(Z)$ that is written as

$$(5-32) \quad I_{ech}(Z) = c_\tau(Z) + Q_\tau(Z) + \mu_\tau(\Theta) - \mu_\tau(\Theta').$$

What follows describes the meaning of the various terms in this equation.

To start, introduce the 2–plane bundle $K^{-1} \subset TY$, this being the kernel of a with orientation given by w . The subscript τ in (5-32) refers to a chosen homotopy class of the trivialization of K^{-1} along the integral curves that form Θ and Θ' . The individual terms in (5-32) depend on τ but their sum does not. As explained next, there is an almost canonical choice for τ in the present circumstances.

Consider first τ along a given $\mathfrak{p} \in \Lambda$ version of either $\widehat{\gamma}_{\mathfrak{p}}^{+}$ or $\widehat{\gamma}_{\mathfrak{p}}^{-}$. The 1–forms $\{d\theta, du\}$ map K^{-1} isomorphically to the trivial bundle on such a curve, and so defines the desired version of τ .

Consider next the case when $\gamma \subset M_{\delta} \cup (\bigcup_{\mathfrak{p} \in \Lambda} \mathcal{H}_{\mathfrak{p}}^{+})$ is a closed integral curve of v from an element in $\mathcal{Z}_{\text{ech}, M}$. Fix $\mathfrak{p} \in \Lambda$ so as to describe a trivialization of K^{-1} along γ 's intersection with $\mathcal{H}_{\mathfrak{p}}$. To this end, reintroduce the function h that is depicted in (4-1). The 1–forms $\{d\phi, dh\}$ map $K^{-1}|_{\gamma}$ isomorphically to the trivial bundle along $\gamma \cap \mathcal{H}_{\mathfrak{p}}^{+}$. Let $\tau_{\mathfrak{p}}$ denote this trivialization. It follows from what is said by Constraint 3 in Part 3 of Section 1C that the trivializations that form the collection $\{\tau_{\mathfrak{p}}\}_{\mathfrak{p} \in \Lambda}$ extend as trivializations of K^{-1} along $\gamma \cap ((M_{\delta} - \Sigma) \cup (\bigcup_{\mathfrak{p} \in \Lambda} \mathcal{H}_{\mathfrak{p}}^{+}))$. Note that each component of this intersection is a segment that passes through precisely one $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}^{+}$. It also follows from this same Constraint 3 that these trivializations do not agree when compared along Σ . Agreement can be obtained by suitably rotating the fiber of the trivial bundle $\gamma \times \mathbb{R}^2$ along γ 's intersection with the $f \in (\frac{5}{4}, \frac{3}{2})$ part of $\gamma \cap M_{\delta}$. In particular, the rotation on any given segment of this intersection can be chosen to have angle either $\frac{\pi}{2}$ or $-\frac{\pi}{2}$ depending on whether the $+$ or $-$ sign appears in Constraint 3 for the relevant component of $T_{+} \cap T_{-}$. This convention defines a homotopy class of trivialization of K^{-1} along the whole of γ , and this gives the desired version of τ .

With τ defined, consider next the term in (5-31) that is denoted by $\mu_{\tau}(\Theta)$. The definition in the case when $\Theta \in \mathcal{Z}_{\text{ech}, M}$ is straightforward by virtue of the fact that all of the closed integral curves from Θ are hyperbolic. In this case, these terms are defined as follows: Let γ denote one of the integral curves of v from Θ . Lie transport of $K^{-1}|_{\gamma}$ by v along a circumnavigation of γ rotates a basis vector in the chosen trivialization by an angle $n\pi$ with $n \in \mathbb{Z}$. This integer n is $\mu_{\tau}(\Theta)$. It follows from the formula in Property 3 in Section 1E that $n = 0$ when γ is some $\mathfrak{p} \in \Lambda$ version of either $\widehat{\gamma}_{\mathfrak{p}}^{+}$ or $\widehat{\gamma}_{\mathfrak{p}}^{-}$. Suppose next that γ is a closed orbit from Θ that corresponds to an irreducible, HF cycle. It follows from Property 3 in Section 1E, from what is said in Lemma 2.2, and from what is said about τ in the preceding paragraph that $|\mu_{\tau}(\Theta)|$ is no greater than G .

The terms $c_\tau(Z)$ and $Q_\tau(Z)$ depend on the chain Z . What follows describes these terms for the case when Θ and Θ' are from $\mathcal{Z}_{\text{ech},M}$. To start, it is necessary to first choose a properly embedded, oriented surface with boundary in $[-1, 1] \times Y$ with certain properties. Use \mathcal{S} to denote this surface.

- (5-33) • The components of $\mathcal{S} \cap (\{-1\} \times Y)$ are the integral curves of ν from Θ' , and those from $\mathcal{S} \cap (\{1\} \times Y)$ are the integral curves of ν from Θ .
- The surface \mathcal{S} is transversal to $\{-1, 1\} \times Y$.
 - The image of \mathcal{S} in Y via the projection defines a 2-chain that gives the same class as Z in $H_2(Y; [\Theta] - [\Theta'])$.

To define $c_\tau(Z)$, choose a section along \mathcal{S} of K^{-1} that restricts to $\mathcal{S} \cap (\{-1, 1\} \times Y)$ as a basis vector for a trivialization defined by τ . A sufficiently generic choice will vanish in a transversal fashion, and so will have a well-defined Euler number. This number is independent of the chosen section, and it is independent of \mathcal{S} except to the extent that the constraints in (5-33) are imposed. This Euler number is $c_\tau(Z)$.

To define $Q_\tau(Z)$, let $N \rightarrow \mathcal{S}$ denote the latter's normal bundle. The constraints given by the first and second bullets in (5-33) imply that N along $\mathcal{S} \cap (\{-1, 1\} \times Y)$ is canonically isomorphic to K^{-1} . Choose a section of N that restricts to $\mathcal{S} \cap (\{-1, 1\} \times Y)$ as a basis vector for a trivialization defined by τ . A sufficiently generic choice will vanish transversally and so have a well-defined Euler number. The latter is independent of the chosen section, and it is also independent of \mathcal{S} except to the extent that (5-33) is imposed. This Euler number is $Q_\tau(Z)$.

Step 5 This step, Step 6 and Step 7 derive bounds for the absolute values of $I_{\text{ech}}(Z_0)$, $I_{\text{ech}}(Z_+)$ and $I_{\text{ech}}(Z_-)$. To this end, the simplest case to consider is $I_{\text{ech}}(Z_0)$. It follows directly from what is said in Step 2 about $Z_0 \cap M_r$ and from Lemma 5.4 that the absolute value of each term that appears in (5-32) is bounded by c_0 , where c_0 depends only on the Heegaard Floer data. Thus $|I_{\text{ech}}(Z_0)|$ is also bounded by c_0 with c_0 dependent only on the Heegaard Floer data.

Consider next the case of $I_{\text{ech}}(Z_+)$. As noted in Step 4, the respective absolute value of the relevant version of $\mu_\tau(\Theta)$ and $\mu_\tau(\Theta')$ is bounded by c_0 , where c_0 depends only on the Heegaard Floer data. This is not true for either $c_\tau(Z)$ or $Q_\tau(Z)$.

To compute $c_\tau(Z_+)$ and $Q_\tau(Z_+)$, it is necessary to first choose a surface \mathcal{S} that is described by (5-32). If $\gamma \in \Theta_{C_+}$ is from the set $\{\hat{\gamma}_p^+ \cup \hat{\gamma}_p^-\}_{p \in \Lambda}$, then $[-1, 1] \times \gamma$

is a component of \mathcal{S} . The boundaries of the remaining components account for the remaining integral curves in Θ_{C_+} and Θ_{C_-} .

The remaining components of \mathcal{S} lie in $[-1, 1] \times (M_r \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p))$. Each such component is a closed annulus; thus each is diffeomorphic to $[-1, 1] \times S^1$. These annuli are in 1-to-1 correspondence with the integral curves from Θ_{C_+} that come from the HF cycle \hat{v}_+ . Let $\gamma^C \in \Theta_C$ denote one of these integral curves, and let $\gamma^\Theta \subset \Theta_+$ denote the corresponding curve. Use $\mathcal{S}_\gamma \subset \mathcal{S}$ to denote the associated annular component of \mathcal{S} . The respective boundary components $\{-1\} \times S^1$ and $\{1\} \times S^1$ of \mathcal{S}_γ appear in $\{-1\} \times Y$ and in $\{1\} \times Y$ as $\{-1\} \times \gamma^\Theta$ and $\{1\} \times \gamma^C$.

An annulus can be viewed as a union of rectangles with disjoint interior such that each rectangle has one edge on one of the boundary circles and the other edge on the other boundary circle. The two other edges intersect the boundary circle only at the corners. A description of \mathcal{S}_γ of this sort is presented in what follows.

The part of \mathcal{S}_γ in $[-1, 1] \times M_r$ is a disjoint union of rectangles. To describe a given such rectangle, recall first from Step 4 that Z_+ intersects M_r as a disjoint union of embedded rectangles. A subset of these rectangles that form $Z_+ \cap M_r$ have one edge on a component of $\gamma^C \cap M_r$ and the opposing edge on the nearby component of $\gamma^\Theta \cap M_r$. A given rectangle is foliated by integral curves of \mathfrak{v} ; these integral curves intersect Σ as an arc, ι , of length bounded by $c_0\delta$ in a component of $T_+ \cup T_-$. This rectangle is in the $f \in (1, 2)$ part of M_r , and so it can be viewed as the graph from a rectangle in $[-1, 1] \times (1, 2)$ to $T_+ \cup T_-$. The latter rectangle has one edge on $\{-1\} \times (1, 2)$ and the opposite edge on $\{1\} \times (1, 2)$. These edges map to M_r as γ^Θ and γ^C . The other two edges map to the boundary of M_r . Any given slice of the $[-1, 1]$ factor maps to an integral curve of \mathfrak{v} . Identify $(1, 2) \times \Sigma$ with the $f^{-1}(1, 2)$ part of M and this rectangle appears as a submanifold in $[-1, 1] \times M_r$. The latter is a component of the intersection of \mathcal{S}_γ with $[-1, 1] \times M_r$.

Each of the remaining rectangles that form \mathcal{S}_γ is the $|u| \leq R + \ln r$ part of the intersection between \mathcal{S}_γ and a $p \in \Lambda$ version of $[-1, 1] \times \mathcal{H}_p$. Each rectangle of this sort is a suitable parametrization of the relevant $p \in \Lambda$ version of the graph in (5-31).

Step 6 With \mathcal{S}_γ understood, consider now the latter's contribution to the Z_+ version of what is denoted by $c_\tau(\cdot)$ in (5-32). It proves useful for this purpose to first choose a particular section of the dual to K^{-1} over \mathcal{S}_γ . The pair of sections of T^*Y that identify K^{-1} with the trivial bundle over γ^C and γ^Θ extend τ without a zero over the rectangles that form $Z_+ \cap M_r$. One of these sections can be taken so that it restricts

as $d\hbar$ to the $|u| \geq R + \ln r$ part of any given $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_{\mathfrak{p}}$. Let η_r denote the latter. The 1-form $d\hbar$ is defined over the whole of $\mathcal{H}_{\mathfrak{p}}$ and so defines by restriction a section of $(K^{-1})^*$ over $\mathcal{H}_{\mathfrak{p}}$. This understood, the pullback to $\mathcal{S}_{\mathcal{Y}}$ via the projection to Y of the 1-form η_r over the $[-1, 1] \times M_r$ part of $\mathcal{S}_{\mathcal{Y}}$, and the pullback to $\mathcal{S}_{\mathcal{Y}}$ of $d\hbar$ over the part of $\mathcal{S}_{\mathcal{Y}}$ in any given $\mathfrak{p} \in \Lambda$ version of $[-1, 1] \times \mathcal{H}_{\mathfrak{p}}$ defines a section over $\mathcal{S}_{\mathcal{Y}}$ of $(K^{-1})^*$. This section is denoted in what follows by η .

The section η of $(K^{-1})^*$ has no zeros on the $[-1, 1] \times M_r$ part of $\mathcal{S}_{\mathcal{Y}}$. Fix $\mathfrak{p} \in \Lambda$ so as to consider the zero locus of η on the $[-1, 1] \times \mathcal{H}_{\mathfrak{p}}$ part of $\mathcal{S}_{\mathcal{Y}}$. It is assumed in what follows that this part of $\mathcal{S}_{\mathcal{Y}}$ is not empty. The 1-form $d\hbar$ annihilates $K^{-1}|_{\mathcal{H}_{\mathfrak{p}}}$ at the points where it is proportional to a . This locus consists of the $\theta \in \{0, \pi\}$ arcs and the circles where both $u = 0$ and $1 - 3 \cos^2 \theta = 0$. All of these zeros of $d\hbar|_{K^{-1}}$ are transversal. Note that the projection to $\mathcal{H}_{\mathfrak{p}}$ of the graph defined by (5-31) intersects only the $\theta = 0$ arc.

The fifth bullet in (5-30) and what is said in the preceding paragraph have the following implication: The bundle K^{-1} over $\mathcal{S}_{\mathcal{Y}}$ has a smooth section with the large $|s|$ asymptotics required to compute $c_{\tau}(\cdot)$ and which vanishes only at the points in the $\mathfrak{p} \in \Lambda$ versions of $\mathcal{S}_{\mathcal{Y}} \cap ([-1, 1] \times \mathcal{H}_{\mathfrak{p}})$ where $\theta = 0$. Moreover, the section vanishes transversally at each such point, and all of these points make the same contribution to c_{τ} . This section also generates the kernel of $d\hbar$ on $\mathcal{S}_{\mathcal{Y}} \cap ([-1, 1] \times \mathcal{H}_{\mathfrak{p}})$. The section in question is denoted by \mathfrak{z} .

To compute the contribution of the zeros of \mathfrak{z} to c_{τ} , introduce Euclidean coordinates $z_1 = \sin \theta \cos \phi$ and $z_2 = \sin \theta \sin \phi$ for a neighborhood in S^2 of the $\theta = 0$ point. The bundle K^{-1} on $\mathcal{H}_{\mathfrak{p}}$ is trivialized near the $\theta = 0$ arc by the pair of 1-forms $\{dz_1, dz_2\}$. It follows from the formula for w in Property 3 of Section 1E that the pair $\{-dz_1, dz_2\}$ define an orientation-preserving identification of K^{-1} near the $\theta = 0$ point with the product \mathbb{R}^2 -bundle. Here, \mathbb{R}^2 is oriented in the standard manner. This identification of K^{-1} with the product bundle near the $\theta = 0$ arc is assumed implicitly in what follows.

Let $(s_0, u_0) \in [-1, 1] \times [-R - \ln r, R + \ln r]$ denote a point that is sent to the $\theta = 0$ point by the map $q_{\mathfrak{p}}$ from (5-31). For the purposes at hand, no generality is lost by assuming that the $q_{\mathfrak{p}}$ near (s_0, u_0) is the map

$$(5-34) \quad (s, u) \rightarrow (z_1 = s - s_0, z_2 = u - u_0),$$

and that \mathfrak{z} appears on a neighborhood of (x_0, u_0) as the map to \mathbb{R}^2 given by

$$(5-35) \quad (s, u) \rightarrow (u - u_0, s - s_0).$$

These formulas imply that each zero of \mathfrak{z} on $S^{-1} \cap ([-1, 1] \times \mathcal{H}_p)$ contributes -1 to $c_\tau(Z_+)$.

Step 7 Consider next the $Z = Z_+$ version of the term Q_τ in (5-32). The section \mathfrak{z} that was just used to compute $c_\tau(Z_+)$ will be used to compute $Q_\tau(Z_+)$ as well. Its use in this regard exploits the following observation: Suppose that X_1 and X_2 are smooth manifolds, $g: X_1 \rightarrow X_2$ a smooth map and $\mathcal{V} \subset X_1 \times X_2$ the graph of g . Then the differential of the projection to X_2 identifies the normal bundle to \mathcal{V} with $TX_2|_{\mathcal{V}}$.

As explained above, a given rectangle in $S_\gamma \cap ([-1, 1] \times M_r)$ can be viewed as a graph of a map from a rectangle in $[-1, 1] \times (1, 2)$ to $T_+ \cup T_- \subset \Sigma$. Doing so identifies the normal bundle to this part of S_γ with $T\Sigma$. Meanwhile, K^{-1} is also identified as $T\Sigma$ here. This understood, the section \mathfrak{z} defines a section of the normal bundle to S_γ on this rectangle.

Fix $p \in \Lambda$. If nonempty, then the $|u| \leq R + \ln r$ part of $S_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ is the graph of the map q_p that is depicted in (5-31). This being the case, the differential of the projection map to S^2 from $[-1, 1] \times [-R - \ln r, R + \ln r] \times S^2$ identifies the normal bundle to this part of S_γ with TS^2 . The differential of this projection also defines a homomorphism from $K^{-1}|_{\mathcal{H}_p}$ to TS^2 . This homomorphism is not an isomorphism only where v is tangent to the S^2 factor in \mathcal{H}_p . A look at Property 3 in Section 1E finds that such is the case where $1 - 3 \cos^2 \theta = 0$. The bundle K^{-1} on this locus is spanned by ∂_u and $3f \partial_\theta - g \partial_\phi$, and so the kernel of the homomorphism to TS^2 is spanned by ∂_u .

The image of \mathfrak{z} by this homomorphism to TS^2 defines a section of the normal bundle of S_γ over its intersection with $[-1, 1] \times \mathcal{H}_p$. This normal bundle section can be modified slightly with no added zeros near the $|u| = R + \ln r$ boundary of its domain so that the result extends the normal bundle section that is defined by \mathfrak{z} over the $|u| \leq R + \ln r$ part of $S_\gamma \cap ([-1, 1] \times M_r)$. Here is why: the intersection of S_γ with the $|u| = R + \ln r$ spheres in $[-1, 1] \times \mathcal{H}_p$ lies where the angle θ differs from $\frac{\pi}{2}$ by at most $c_0 \delta$, and the 1-form a on this part of \mathcal{H}_p differs from $2e^{2(|u|-R)} du$ by a term with absolute value bounded by $c_0 \delta$.

It follows as a consequence of what was just said that there exists a section of the normal bundle of S_γ that has the requisite large $|s|$ asymptotics for computing Q_τ and whose zero locus consists solely of the various $p \in \Lambda$ versions of the $|u| < R + \ln r$ part of $S_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ where $\theta = 0$. Note that this can be done despite the fact that the homomorphism from K^{-1} to TS^2 has rank 1 on the $1 - 3 \cos^2 \theta = 0$ locus.

The assertion about the $1 - 3 \cos^2 \theta = 0$ locus can be justified by straightforward arguments using five facts noted previously. The first is that S_γ intersects each sphere

where $|u| = R + \ln r$ only where θ has distance at most $c_0\delta$ from $\frac{\pi}{2}$. The second fact is that each such intersection is an arc with length bounded by $c_0\delta$. The third fact is that the bundle K^{-1} where $|u| = R + \ln r$ and θ is close to $\frac{\pi}{2}$ is very nearly identical to TS^2 . In particular, it is very nearly spanned by the pair $\{\partial_\theta, \partial_\phi\}$. The fourth fact is that the vector field $3f\partial_\theta - g\partial_\phi$ is in K^{-1} along the locus where $1 - 3\cos^2\theta = 0$. The final fact is that this vector field has nonzero image in TS^2 via the homomorphism from K^{-1} .

To see about the contribution to $Q_\tau(Z_+)$ from a $\theta = 0$ point in the $|u| < R + \ln r$ part of $\mathcal{S}_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$, parametrize S^2 near the $\theta = 0$ point as before using the coordinates (z_1, z_2) . The normal bundle to TS^2 near the $\theta = 0$ point has the oriented basis $\{\frac{\partial}{\partial z_1}, \frac{\partial}{\partial z_2}\}$, and this basis is used implicitly in what follows to identify TS^2 near the $\theta = 0$ point with the product \mathbb{R}^2 -bundle. With the preceding kept in mind, parametrize the $|u| < R + \ln r$ part of $\mathcal{S}_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ as the graph of the map q_p in (5-33) and let (s_0, u_0) again denote a $\theta = 0$ point. It follows from (5-35) that the normal bundle section defined by \mathfrak{z} here appears as the map to \mathbb{R}^2 given by

$$(5-36) \quad (s, u) \rightarrow (-u + u_0, s - s_0).$$

The determinant of the differential of this map is $+1$. As a consequence, each $\theta = 0$ point in the $|u| < R + \ln r$ part of $\mathcal{S}_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ contributes $+1$ to $Q_\tau(Z_+)$.

Step 8 As noted in Step 5, the absolute value of the contribution of $\mu_\tau(Z_+)$ to the $Z = Z_+$ version is bounded by a constant that depends only on the Heegaard Floer data. As seen in Step 6, each $\theta = 0$ point in the $|u| < R + \ln r$ part of $\mathcal{S}_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ contributes -1 to $c_\tau(Z_+)$, and as seen in Step 7, each such point contributes $+1$ to $Q_\tau(Z_+)$. This implies that the $|u| < R + \ln r$ part of $\mathcal{S}_\gamma \cap ([-1, 1] \times \mathcal{H}_p)$ contributes *zero* to the sum $c_\tau(Z_+) + Q_\tau(Z_+)$ that appears in the $Z = Z_+$ version of (5-32). Granted this, it then follows that $|I_{\text{ech}}(Z_+)|$ is bounded by a constant that depends solely on the Heegaard Floer data. □

5C Applications of Gromov compactness

This subsection has two parts. Part 1 states the version of Gromov’s compactness theorem for J -holomorphic submanifolds that is used in Part 2. Part 2 gives the applications. Of particular interest for [12] is the upcoming Proposition 5.8 from Part 2. This lemma states conclusions that are stronger versions of those stated by Lemmas 4.3 and 4.4.

Part 1 To set the notation for the statement of the compactness theorem, suppose that $U \subset Y$ is a given open set with compact closure. Given $\varepsilon > 0$, introduce U_ε to denote the set of points in Y with distance ε or less from U . The statement of the compactness theorem also introduces the following terminology: Let $I \subset \mathbb{R}$ denote a given bounded, open interval, and let U be as just described. A subset $C \subset I \times U$ is said to be a J -holomorphic subvariety if there is an open neighborhood in $\mathbb{R} \times U_\varepsilon$ of the closure of $I \times U$ and a set in this neighborhood with the following properties: The intersection of this set with $I \times U$ is C , the set has no isolated components, and the complement of a finite set of points from this set is a finite-area submanifold with J -invariant tangent space. The term *weighted J -holomorphic subvariety* in $I \times U$ denotes a finite set of pairs consisting of an irreducible, J -holomorphic subvariety in $I \times U$ and positive integer with it understood that no two pairs share the same subvariety.

Proposition 5.5 Fix an open set $U \subset Y$, an open, bounded interval $I \subset \mathbb{R}$ and $\mathcal{L} \geq 1$. Given this data, there exists $\kappa > 1$ with the following significance: Let $\{C_n\}_{n=1,2,\dots}$ denote a sequence of J -holomorphic subvarieties in $\mathbb{R} \times U$. Suppose this sequence has the following property: there exists $\varepsilon > 0$ such that

$$\int_{C_n \cap ([x-2, x+2] \times U_\varepsilon)} (ds \wedge \hat{a} + w) \leq \mathcal{L}$$

for each $x \in I$ and for each $n \in \{1, 2, \dots\}$. Then there exists a weighted J -holomorphic subvariety ϑ and a subsequence of the original sequence (henceforth renumbered consecutively from 1) with the following properties:

- $\lim_{n \rightarrow \infty} (\sup_{z \in C_n} \text{dist}(z, \bigcup_{(C,m) \in \vartheta} (C \cap (I \times U))) + \sup_{z \in C \cap (I \times U)} \text{dist}(z, C_n)) = 0$.
- Let μ denote a smooth, bounded 2-form defined on a neighborhood of $I \times U$ in $\mathbb{R} \times Y$. Then $\lim_{n \rightarrow \infty} (\int_{C_n} \mu - \sum_{(C,m) \in \vartheta} m \int_C \mu) = 0$.

Proof The existence of a weighted J -holomorphic subvariety and subsequence with the required properties follows from the standard sort of compactness theorems for sequences of J -holomorphic subvarieties in 4-dimensional manifolds. See for example Section 6 of [20]. (Convergence assertions are proved in [26; 27; 5] given a priori bounds on the topology of the curves in the set $\{C_n\}_{n=1,2,\dots}$) □

The next lemma states an important corollary to Proposition 5.5.

Lemma 5.6 Fix an open set $U \subset Y$, $\mathcal{L} \geq 1$ and $\varepsilon > 0$. Given this data, there exists $\kappa > 1$ with the following significance: Let $C \subset \mathbb{R} \times Y$ denote a J -holomorphic subvariety. Suppose that $x \in \mathbb{R}$ is such that:

- $\int_{C \cap ([x-2, x+2] \times U_1)} \omega < \kappa^{-1}$.
- $\int_{C \cap ([x-2, x+2] \times U_1)} ds < a \leq \mathcal{L}$.

Let $C' \subset C \cap ([x - 1, x + 1] \times U)$ denote a connected component. Then there is an integral curve of v such that each point of $(C'|_s \cap U)$ for each $s \in [x - 1, x + 1]$ has distance no greater than ε from the intersection of this curve with U , and vice versa. Let γ denote this integral curve. There exists a positive integer m which is less than \mathcal{L} and such that if μ is a smooth 2-form on $[x - 1, x + 1] \times U$ with $\|\mu\|_\infty = 1$ and $\|\nabla\mu\|_\infty \leq \varepsilon^{-1}$, then

$$\left| \int_{C'} \mu - m \int_{[x-1, x+1] \times \gamma} \mu \right| \leq \varepsilon.$$

Proof The conditions given by the two bullets bound the area of C 's intersection with $[x - 2, x + 2] \times U_1$. This understood, the assertion can be proved by assuming the contrary and invoking Proposition 5.5 so as to obtain a nonsense conclusion: there exists a sequence of J -holomorphic subvarieties whose n^{th} member obeys the assumptions with $\kappa = n$ and $\varepsilon = n^{-1}$ that converges as described in Proposition 5.1 to a weighted J -holomorphic subvariety with the following property: the subvariety component from each constituent pair is an \mathbb{R} -invariant cylinder. For more details, see for example the proofs of Lemma 3.6 and Corollary 4.7 in [23], as these state an analogous assertion in the context where a is a contact 1-form and $\omega = da$. □

Part 2 The next lemma states a crucial result for the proof of the upcoming Proposition 5.8 and for use in [12].

Lemma 5.7 *Given $\mathcal{I} > 0$ and $\rho \in (0, \delta_*)$, there exists $\kappa > 1$ with the following significance: Fix data δ, χ_0, R and J suitable for defining the geometry of Y and $\mathbb{R} \times Y$ with $\delta < \frac{1}{2}\rho$. Let $C \subset \mathbb{R} \times Y$ denote an ech-HF subvariety with $I_{\text{ech}}(C) \leq \mathcal{I}$ or $\int_C \omega \leq \mathcal{I}$.*

Let c_C denote the maximum value of $\cos^2(\theta)$ on C 's intersection with the portion of the radius δ_* coordinate ball centered on each of the index 1 and index 2 critical points of f where the radius is greater than ρ . Then $1 - 3c_C^2 > \kappa^{-1}$.

Proof Since Proposition 4.1 says that $1 - 3\cos^2\theta > 0$ on C , the issue is whether there is a uniform lower bound given \mathcal{I} and ρ . Were there no such κ , one could find a countably infinite sequence of ech-HF subvarieties that obeyed the assumptions of the lemma, but whose n^{th} member violated the conclusions with $\kappa = n$. Use

the bounds from the two bullets in Proposition 5.1 to invoke Proposition 5.5 with $U = M_{\rho/2}$. Proposition 5.5 finds a subsequence that converges on compact subsets of part of $\mathbb{R} \times M_{\rho/2}$ in a pointwise manner to a J -holomorphic subvariety in $\mathbb{R} \times M_{\rho}$. The limit subvariety must have an irreducible component that intersects the $f = 1$ or $f = 2$ locus because this is the $1 - 3 \cos^2 \theta = 0$ locus in the radius $7\delta_*$ ball about any index 1 or index 2 critical point of f . Such a component would be the intersection of $\mathbb{R} \times M_{\rho/2}$ with a submanifold from either $\mathcal{M}_{f=1}$ or $\mathcal{M}_{f=2}$ because it follows from Proposition 3.3 that the submanifolds in these spaces foliate the $f = 1$ and $f = 2$ loci in $\mathbb{R} \times M_{\rho/2}$. Let C_f (which is in $\mathbb{R} \times Y$) denote this component of the limit. Either C_f is disjoint from the other irreducible components of the limit or it is not. If it is not disjoint, then another irreducible component of the limit in $\mathbb{R} \times M_{\rho/2}$ must intersect the $f < 1$ or $f > 2$ part of $\mathbb{R} \times M_{\delta}$ because otherwise it would be described by Proposition 3.3 and thus disjoint from C_f . Since $f < 1$ or $f > 2$ on this hypothetical component of the limit, it would also be the case that $f < 1$ or $f > 2$ on all sufficiently large n submanifolds from the subsequence. This is ruled out by the first bullet of Proposition 4.1. On the other hand, if C_f is disjoint from the rest of the limit, then a winnowing of the chosen subsequence and renumbering will result in a new subsequence with the following property: There exists some $s \in \mathbb{R}$ such that any given large n member of the subsequence has a component that intersects $\mathbb{R} \times M_{\delta}$ with all points having distance at most n^{-1} from $\{s\} \times C_f$. Moreover, all points in $\{s\} \times C_f$ will have distance at most n^{-1} to this component of the n^{th} element from the subsequence. However, such a component of an admissible J -holomorphic subvariety will intersect at least one submanifold from \mathcal{M}_{Σ} , and so it will intersect all submanifolds from \mathcal{M}_{Σ} unless it is an actual element in \mathcal{M}_{Σ} . This last option is also ruled out for the following reasons: Let C_n denote the n^{th} element in the subsequence and let $C'_n \subset C_n$ denote the component in question. Since C'_n lies very close to C_f , it can't intersect all elements in \mathcal{M}_{Σ} , and it can't be an element in \mathcal{M}_{Σ} because, being that C'_n and C_f are very close when n is large, this would run afoul of Lemma 4.4 when n is large. \square

The next proposition subsumes some of what is asserted by Lemmas 4.3 and 4.4. This proposition identifies the moduli space \mathcal{M}_{Σ} of Proposition 3.2 with $\mathbb{R} \times (1, 2)$ using the diffeomorphism described by Proposition 3.2.

Proposition 5.8 *Given $\mathcal{I} > 0$ and $\rho \in (0, \delta_*]$, there exists $\kappa > 1$ and $z_{\diamond} \in (0, \rho^2)$ with the following significance: Fix data δ , χ_0 , R and J that are suitable for defining*

the geometry of Y and $\mathbb{R} \times Y$ with $\delta < \frac{1}{100}\rho$. Let $C \subset \mathbb{R} \times Y$ denote an ech-HF subvariety with either $I_{\text{ech}}(C) \leq \mathcal{I}$ or $\int_C \omega \leq \mathcal{I}$.

- C intersects the $f \leq 1 + z_\diamond$ part of $\mathbb{R} \times M_\delta$ in the union of the radius ρ_\diamond coordinate balls centered at the index 1 critical points of f . In addition, C intersects each subvariety from the $\mathbb{R} \times (1, 1 + z_\diamond)$ part of \mathcal{M}_Σ precisely once in each $\mathfrak{p} \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$ and this intersection point lies where $u \in (0, R + \ln \rho)$
- C intersects the $f \geq 2 - z_\diamond$ part of $\mathbb{R} \times M_\delta$ in the union of the radius ρ coordinate balls centered at the index 2 critical points of f . In addition, C intersects each subvariety from the $\mathbb{R} \times (1, 1 + z_\diamond)$ part of \mathcal{M}_Σ precisely once in each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ and this intersection point lies where $u \in (-R - \ln \rho, 0)$.

Proof Given $\rho \in (0, \delta_*)$, fix δ, χ_0, R and J so as to define the geometry of Y and $\mathbb{R} \times Y$, and so that $\delta < \frac{1}{100}\rho$. Suppose C is an ech-HF subvariety. Suppose that $\varepsilon > 0$ is such that the following is true: Let p denote either an index 1 or index 2 critical point of f . Then $1 - 3 \cos^2 \theta$ is greater than ε on the part of C in the radius δ_* coordinate ball centered on p where the radius is greater than ρ . If p is an index 1 critical point then $f > 1 + \rho^2\varepsilon$ on C 's intersection with this part of the radius δ_* coordinate ball centered at p , and if p is an index 2 critical point of f , then $f < 2 - \rho^2\varepsilon$ on C 's intersection with this part of the radius δ_* coordinate ball centered at p . The next paragraph describes some consequences of these observations.

Suppose that $(x, y) \in (1, 1 + \rho^2\varepsilon)$ and let $S_{(x,y)} \in \mathcal{M}_\Sigma$ denote the corresponding element from the moduli space of Proposition 3.2. By way of a reminder, this submanifold intersects $\mathbb{R} \times M_\delta$ as $\{x\} \times f^{-1}(y)$. If y is very close to 1, then C has precisely one intersection point with $S_{(x,y)}$ in each $\mathfrak{p} \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$ and this intersection point lies where $u \in (0, R + \ln \rho)$. As y is increased, the corresponding intersection point will move in a continuous fashion if it doesn't leave $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$. However, it can't leave $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$ because of what is said in the preceding paragraph: it can't leave until $y > 1 + \rho^2\varepsilon$. Since the intersection point stays in $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$, each subvariety from the $\mathbb{R} \times (1, 1 + \rho^2\varepsilon)$ part of \mathcal{M}_Σ intersects C precisely once in each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ and this intersection point lies where $u \leq R + \ln \rho_\diamond$. The analogous argument shows that C intersects each subvariety from the $\mathbb{R} \times (2 - \rho^2\varepsilon, 2)$ part of $\mathbb{R} \times \mathcal{H}_\mathfrak{p}$ precisely once in each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ and that the latter intersection point occurs where $u > -R - \ln \rho$.

Granted the preceding, let κ denote the constant that is supplied by Lemma 5.7 for the given values of \mathcal{I} and ρ . It follows from what is said in Lemma 5.7 and in the previous two paragraphs that the conclusions of Proposition 5.8 hold with $z_\diamond < \rho^2 \kappa^{-1}$. \square

6 Heegaard Floer curves

This section constitutes a digression to describe in more detail the sorts of pseudoholomorphic subvarieties that arise in Lipshitz's formulation [13] of Heegaard Floer geometry. What is said here summarizes and in some cases elaborates on observations made by Lipshitz. The discussion in this section may be of independent interest to those using the constructions in [13]. In any event, the results here are used in Section 7 and in [12].

6A The Heegaard Floer geometry

This first subsection sets up the geometry that is used by Lipshitz to define his version of pseudoholomorphic curves. This subsection has three parts.

Part 1 Introduce $M_{(1,2)}$ as shorthand to denote the $f^{-1}(1, 2)$ part of M . Define a closed 2-form w_f on $M_{(1,2)}$ as follows: Set, $w_f = w$ on $M_\delta \cap M_{(1,2)}$. Use (1-13) to define w_f on the radius 2δ coordinate ball centered on any given index 1 critical point of f . Use the corresponding $(t_-, \theta_-, \varphi_-)$ version to define w_f on the radius 2δ coordinate ball centered on any given index 2 critical point of f . The resulting 2-form is annihilated by the Lie derivative of the pseudogradient vector field \mathfrak{v} .

Lie transport by \mathfrak{v} identifies $M_{(1,2)}$ with $(1, 2) \times \Sigma$ as in Part 5 of Section 1C. By way of review, this identification writes f as the Euclidean coordinate t on the $(1, 2)$ factor, and \mathfrak{v} as the dual vector field $\frac{\partial}{\partial t}$. Meanwhile, the 2-form w_Σ appears as the area form on Σ that is defined in Part 5 of Section 1C. The identification between $M_{(1,2)}$ and $(1, 2) \times \Sigma$ is used implicitly in much of what follows.

An almost complex structure on $\mathbb{R} \times M_{(1,2)}$ is defined by the rules listed in a moment. This almost complex structure is denoted in what follows by J_{HF} . To set the notation, reintroduce from Part 3 of Section 1C the function $h_+ = 2e^{2t_+} \cos \theta_+ \sin^2 \theta_+$ on the radius $8\delta_*$ coordinate ball centered on a given index 1 critical point of f . The functions (f, φ_+, h_+) define coordinates on the $1 - 3 \cos^2 \theta_+ > 0$ part of this coordinate ball. The list refers to the dual vector fields for these coordinates. There is the analogous function

$h_- = 2e^{2t-} \cos \theta_- \sin^2 \theta_-$, and coordinates (f, φ_-, h_-) and with the corresponding dual vector fields on the $1 - 3 \cos^2 \theta_- > 0$ part of the radius $8\delta_*$ coordinate ball centered on a given index 2 critical point of f . The list also refers to the subsets T_+ and T_- of Σ that are introduced in Part 4 of Section 1C.

- (6-1) • $J_{HF} \cdot \partial_s = v$.
- J_{HF} is invariant with respect to constant translations along the \mathbb{R} factor of $\mathbb{R} \times M$.
 - J_{HF} preserves the kernel of df and doing so, it defines the orientation given by w_Σ .
 - Write $M_{(1,2)}$ as $(1, 2) \times \Sigma$. The almost complex structure J_{HF} is invariant with respect to translations along the $(1, 2)$ factor of $\mathbb{R} \times (1, 2) \times (T_+ \cap T_-)$.
 - $J_{HF} \frac{\partial}{\partial \varphi_+} = \frac{\partial}{\partial h_+}$ on \mathbb{R} 's product with the radius $8\delta_*$ coordinate ball centered on any index 1 critical point of f , and $J_{HF} \frac{\partial}{\partial \varphi_-} = -\frac{\partial}{\partial h_-}$ on \mathbb{R} 's product with the radius $8\delta_*$ coordinate ball centered on any index 2 critical point of f .

What follows are comments about these rules. With $M_{(1,2)}$ written as $(1, 2) \times \Sigma$, the third bullet requires that J_{HF} preserve the $T\Sigma$ factor of $T(\mathbb{R} \times (1, 2) \times \Sigma)$ and that it define the orientation on $T\Sigma$ given by the area form w_Σ . Compatibility between the rule in the fifth bullet and the rules in the third and fourth bullets is ensured by Constraint 3 in Section 1C. The final comment concerns the 2-form $ds \wedge df + w_\Sigma$ on $\mathbb{R} \times M_{1,2}$. This 2-form is symplectic, and J_{HF} is a compatible, almost complex structure. The identification of $M_{(1,2)}$ with $(1, 2) \times \Sigma$ writes this symplectic form as $ds \wedge dt + w_\Sigma$.

Part 2 The identification $M_{(1,2)} = (1, 2) \times \Sigma$ identifies the geometry just defined with the geometry that Lipshitz introduces in the Section 1 of [13]. Lipshitz introduces in this section of his paper certain sorts of almost complex structure on $\mathbb{R} \times [1, 2] \times \Sigma$, and J_{HF} is an example. The form $ds \wedge df + w_f$ on $\mathbb{R} \times M_{(1,2)}$ when viewed as a 2-form on the manifold $\mathbb{R} \times (1, 2) \times \Sigma$ is what Lipshitz calls a *split* symplectic form.

What follows is a converse of sorts to the preceding observations. Let J_* denote an almost complex structure on $\mathbb{R} \times [1, 2] \times \Sigma$ of the sort that Lipshitz considers in Section 1 of [13]. This almost complex structure obeys all but possibly the fifth bullet in (6-1). The rule given by the fifth bullet in (6-1) can be imposed without compromising either what is said by Lipshitz [13] or by Ozsváth and Szabó in any of their many papers

(see [17; 18]). The following says more on this point with regards to the Heegaard Floer framework of Ozsváth and Szabó [17]. The assumption made in the fifth bullet is constraining the $[0, 1]$ -parametrized families of almost complex structures on $\text{Sym}^G(\Sigma)$ that are considered in Section 3.2 of [17] very near the subspaces in $\text{Sym}^G \Sigma$ from the sets $\{\alpha_i \times \Sigma^{G-1}\}_{i=1,2,\dots,G}$ and $\{\beta_i \times \Sigma^{G-1}\}_{i=1,2,\dots,G}$. (The notation here is from [17].) These constraints are consistent with what is said in Definition 3.1 of [17] in the sense that they lead to open sets V in $\text{Sym}^G(\Sigma)$ that can be used in Definition 3.1 to define corresponding almost complex structure spaces $\mathcal{J}(j, \eta, V)$. Note in this regard that the $[0, 1]$ dependence of the almost complex structures in [17] is used only to ensure that certain genericity conditions hold, and the genericity assumptions in [17] (such as those that are used to prove Theorem 3.4 in [17]) are consistent with the constraint imposed by the fifth bullet in (6-1). (This is so because the nonconstant holomorphic strips that are considered in [17] do not lie entirely in the union of the sets $\{\alpha_i \times \Sigma^{G-1}\}_{i=1,2,\dots,G}$ and $\{\beta_i \times \Sigma^{G-1}\}_{i=1,2,\dots,G}$.) In any event, the rule given by the fifth bullet can be weakened considerably without changing the conclusions of the propositions and lemmas that follow. For example, it is most likely sufficient to assume that J_{HF} when viewed as an almost complex structure on $\mathbb{R} \times (1, 2) \times \Sigma$ extends smoothly to a neighborhood of $\mathbb{R} \times \{1\} \times C_+$ and to $\mathbb{R} \times \{2\} \times C_-$. This said, bullet five is kept as is.

The rule (J2) in Section 1 of [13] requires somewhat more than what is required by that in the fourth bullet of (6-1). However, the extra conditions in Lipshitz's rule (J2) are not required for what is done in his paper, nor are they truly necessary for the work of Ozsváth and Szabó. In any event, the rule (J2) from [13] can be imposed here without compromising anything done in what follows or in [12].

Part 3 Let J_{HF} denote an almost complex structure that obeys the rules in (6-1). Lipshitz introduces in [13, Section 1] certain J_{HF} -holomorphic maps into $\mathbb{R} \times [1, 2] \times \Sigma$. The domain of such a map is a Riemann surface, S , with punctured boundary. In particular, there is compact, Riemann surface, \underline{S} , with smooth boundary which is such that S is the complement in \underline{S} of a set of $2G$ labeled points in the boundary of \underline{S} . Half of these points are labeled as *negative* and half as *positive*. The map, $u: S \rightarrow \mathbb{R} \times [1, 2] \times \Sigma$, is constrained so as to satisfy seven conditions, these labeled as (M0)–(M6) in Section 1 of [13]. These conditions are stated in the upcoming (6-2). The statement reintroduces from Part 2 of Section 1C the sets C_+ and C_- in Σ , these the respective union of Σ 's intersection with the ascending disks from the index 1 critical points of f and the descending disks from the index 2 critical points of f . What follows are, in order, Lipshitz's conditions:

- (6-2) • The map u is smooth on S and it is J_{HF} -holomorphic on the interior of S .
- The map sends the boundary of S to the union of $\mathbb{R} \times \{2\} \times C_-$ and $\mathbb{R} \times \{1\} \times C_+$.
 - There are no components of S in the fiber of the projection to Σ .
 - Each component of $(\mathbb{R} \times \{2\} \times C_-) \cup (\mathbb{R} \times \{1\} \times C_+)$ contains the image via u of one and only one boundary component of S .
 - The pullback via u of the coordinate s limits to $-\infty$ on sequences in S that converge in \underline{S} to a negative $\underline{S} \setminus S$ point, and s limits to ∞ on sequences in S that converge in \underline{S} to a positive $\underline{S} \setminus S$ point.
 - The integral of $u^* w_\Sigma$ over S is finite, and there exists $\kappa_u \geq 1$ with the following property: Let $I \subset \mathbb{R}$ denote an interval of length 1. Then the integral of $u^*(ds \wedge dt)$ over $u^{-1}(I \times [1, 2] \times \Sigma)$ is bounded by κ_u .
 - The map u is an embedding.

A pair (S, u) as just described is said in [13] and in what follows to be an *admissible pair*. A *singular admissible set* is a set of the form $((S, u), \vartheta_\Sigma)$ where (S, u) is as described above except for the fact that the last bullet in (6-2) need not be obeyed. Even so, u must be almost everywhere 1-to-1, and its image has but a finite number of singular points with none on the boundary. Meanwhile, ϑ_Σ consists of a finite set of pairs whose first component is a fiber of the projection from $\mathbb{R} \times (1, 2) \times \Sigma$ to $\mathbb{R} \times (1, 2)$ and whose second component is a positive integer. However, distinct pairs from ϑ_Σ are distinct fibers. The set ϑ_Σ can be empty.

6B The points in S near its boundary

This subsection says more about the behavior of an admissible pair (S, u) near the boundary of S . The results below elaborate on various remarks in [13]. Note that what is said here is also valid if $u(S)$ has a finite number of interior singular points. The discussion that follows has three parts.

Part 1 Let p denote an index 1 or index 2 critical point of f . The lemma that follows describes the function u^*s on the boundary of S . The lemma uses ∂S to denote the boundary of S .

Lemma 6.1 *Let $\partial S \subset S$ denote the boundary. The coordinate function s from the \mathbb{R} factor of $\mathbb{R} \times [1, 2] \times \Sigma$ pulls back via u to ∂S as a proper function with no critical points.*

Proof The fact that $u^* ds$ cannot vanish on $T(\partial S)$ owes allegiance to the fact that u embeds ∂S into a level sets of t . The argument that follows explains why this so is in the case when the boundary component in question is labeled by an index 1 critical point of f . Except for notation, the same argument works for the components with index 2 critical point labels.

To start, let p denote a given index 1 critical point and let $\partial_p S$ denote the corresponding component of ∂S . The map u sends a neighborhood of $\partial_p S$ in S into a neighborhood of $\mathbb{R} \times \{1\} \times C_{p_+}$ in $\mathbb{R} \times [1, 2] \times \Sigma$. The latter neighborhood is chosen to have the form $\mathbb{R} \times [1, 1 + z_p) \times T_+$ with $z_p > 0$. Given that z_p is small, the functions (φ_+, \tilde{h}_+) from Part 3 of Section 1C define coordinates for T_+ and so the coordinate s for \mathbb{R} , t for $[1, 1 + z_p)$ and the pair (φ_+, \tilde{h}_+) supply coordinates for this neighborhood. The almost complex structure endomorphism J_{HF} maps $\frac{\partial}{\partial s}$ to $\frac{\partial}{\partial t}$, and by dint of the fifth bullet of (6-1), it maps $\frac{\partial}{\partial \varphi_+}$ to $\frac{\partial}{\partial \tilde{h}_+}$.

Suppose that $u^* ds$ vanishes at some point on $\partial_p S$ so as to derive some nonsense. Given that u is J_{HF} -holomorphic, the latter condition requires that $u^* dt$ is also zero at this point. As a consequence, the function $s + i t$ pulls back to S with a critical point at this boundary point. The ramifications of this fact are explained in a moment. To set the stage, fix a holomorphic identification between a neighborhood of this boundary point and the half disk $D_+ \subset \mathbb{C}$ whose points have norm less than 1 and nonnegative imaginary part. Use ζ to denote the complex coordinate on \mathbb{C} and thus on D_+ .

The pullback via u of $s + i t$ is holomorphic on the interior of S . As explained to the author by Curt McMullen, a version of the Schwarz reflection trick from Theorem 24 in [1] can be used to write the pullback of $s + i t$ on D_+ as

$$(6-3) \quad u^*(s + i t) = a_0 \zeta^n + \mathcal{O}(|\zeta|^{n+1}),$$

with $a_0 \in \mathbb{C} \setminus 0$ and $n > 1$. However, (6-3) implies the $t = 1$ level set in S intersects the interior of S , and this cannot happen because it would violate the maximum principle.

Since $u^* ds \neq 0$ and since $|s|$ must limit to ∞ as it converges to any point of $\underline{S} \setminus S$, it follows that $u^* s$ must be proper on each component of ∂S . □

Part 2 It follows from what is said in Part 1 that the pullback via u of the coordinates functions s and t on $\mathbb{R} \times [1, 2] \times \Sigma$ give coordinates on a neighborhood of ∂S . This part says more about the map u on such a neighborhood.

To this end, let p again denote an index one critical point of f . The boundary component $\partial_p S$ can be parametrized by the pullback of the function s . In particular, the image via u of $\partial_p S$ can be parametrized by \mathbb{R} as a graph

$$(6-4) \quad s \rightarrow (s, t) = (1, \varphi^S(s)),$$

where $\varphi^S: \mathbb{R} \rightarrow C_{p+}$ is a smooth map.

Lemma 6.2 *The closure of the corresponding boundary component $\partial_p S$ in \underline{S} adds precisely one negative and one positive point of $\underline{S} \setminus S$. Moreover, the map φ^S extends continuously to this closure so as to map these points of $\underline{S} \setminus S$ to $C_+ \cap C_-$.*

Proof Let ζ again denote the standard complex coordinate for \mathbb{C} and $D_+ \subset \mathbb{C}$ denote the half disk where $|\zeta| < 1$ and $\text{im}(\zeta) \geq 0$. Let $q_S \in \underline{S}$ denote a negative boundary point. There is a neighborhood of q_S in \underline{S} whose closure intersects $\underline{S} \setminus S$ only in q_S , and is diffeomorphic to the half disk D_+ with q corresponding to the origin. This diffeomorphism can be chosen so that the complex structure on \underline{S} near q_S maps to the standard complex structure on \mathbb{C} . It proves useful to employ the map $\zeta \rightarrow \frac{1}{\pi} \ln \zeta + i$ so as to holomorphically identify the $|\zeta| \leq e^{-1}$ part of $D_+ \setminus 0$ with the strip $(-\infty, -1] \times [1, 2]$. The boundary component $(-\infty, -1] \times \{1\}$ is mapped by u to a component of $\mathbb{R} \times \{1\} \times C_+$. Let p denote the index 1 critical point of u that labels this component. Meanwhile, the boundary component, $(-\infty, -1] \times \{2\}$, is mapped by u to a component of $\mathbb{R} \times \{1\} \times C_+$ or of $\mathbb{R} \times \{2\} \times C_-$.

Given that the integral of w_Σ over S is finite, it follows using Lemma 5.7 that the constant s slices of the images via u of $(-\infty, 0) \times [1, 2]$ converge in $[1, 2] \times \Sigma$ to a curve of the form $[1, 2] \times q$ with q some point in Σ . The $t = 1$ end of the u -image of each such slice must lie in C_{p+} and the $t = 2$ end must therefore lie in C_- . Let p' denote the index 2 critical point of f that labels this component. The $s \rightarrow -\infty$ convergence of the image of these slices requires that $q \in C_{p+} \cap C_{p'}$. This convergence of the slices with what was just said about their limit implies what is asserted for the lemma about the $s \rightarrow -\infty$ behavior of ∂S and the map φ^S . Very much the same argument implies what is said about the $s \rightarrow \infty$ behavior of ∂S and φ^S . □

Part 3 Let p denote an index 1 or index 2 critical point of f . The next lemma says something about the manner in which the restriction of u to the constant $u^*(t)$ slices of S near $\partial_p S$ approximate the map φ^S . When p is an index 1 critical point, the lemma refers to the annular neighborhood T_{p+} of C_{p+} with its coordinates (φ_+, h_+) .

When p is an index 2 critical point, it refers to T_{p-} and the corresponding coordinates (φ_-, \hbar_-) . The map φ^S has domain \mathbb{R} and so can be viewed as an \mathbb{R} -valued map. The lemma views it as such.

Lemma 6.3 *Let p denote an index 1 or index 2 critical point of the function f . There exists $z_p \in (0, 1)$ and a neighborhood of $\partial_p S$ in S with the following properties: The composition of first u and then the projection from $\mathbb{R} \times [1, 2] \times \Sigma$ to Σ maps this neighborhood into T_{p+} or T_{p-} as the case may be. Moreover, the neighborhood has a parametrization by $\mathbb{R} \times [0, z_p)$ whereby u appears as*

- $(s, z) \rightarrow (s, t = 1 + z, \varphi_+ = \varphi(s, z), \hbar_+ = \zeta(s, z))$ when p has index 1,
- $(s, z) \rightarrow (s, t = 2 - z, \varphi_- = \varphi(s, z), \hbar_- = \zeta(s, z))$ when p has index 2,

where $\varphi(\cdot)$ and $\zeta(\cdot)$ are maps from $\mathbb{R} \times [0, z_p)$ to \mathbb{R} that obey

$$|\zeta(s, z)| + z^{-1}|\varphi(s, z) - \varphi^S(s)| < \kappa z,$$

with κ a positive constant. In addition,

$$|\partial_s \zeta(s, z)| + z^{-1}|(\partial_s \varphi)(s, z) - (\partial_s \varphi^S)(s)| < \kappa z.$$

In general, the absolute values of the derivatives of $\varphi(\cdot)$, $\zeta(\cdot)$ and $\varphi^S(\cdot)$ to any given specified order enjoy an s and z independent upper bound on $\mathbb{R} \times [0, z_p)$.

Proof Fix $T \geq 1$; then Lemma 6.2 with fact that u is J -holomorphic implies that the lemma’s assertions hold for a neighborhood of the part of $\partial_p S$ where $|u^* s| \leq T$. This understood, it is sufficient to prove that the constants κ and z_p can be chosen to be independent of T . The argument that follows takes p to be an index 1 critical point of f . But for cosmetics, the same argument works for the index 2 critical points.

To start, let $q^S \in C_{p+} \cap C_-$ denote the $s \rightarrow -\infty$ limit on $\partial_p S$. View q^S as a point in $\mathbb{R}/2\pi\mathbb{Z}$ using the coordinate φ_+ . As noted in Part 1 above, the constant s slices of S that intersect $u(\partial_p S)$ converge as $s \rightarrow \infty$ to the arc $[1, 2] \times q^S$ in $[1, 2] \times \Sigma$. As a consequence, given $\varepsilon > 0$, there exists T_ε such that the part of $u(S)$ with $s < -T_\varepsilon$ that intersects $u(\partial_p S)$ has distance ε or less from this arc. In particular, it follows from (6-2) that if $\varepsilon < c_0^{-1}$, then this part of $u(S)$ lies in $[1, 2] \times T_{p+}$, and, in particular, it lies in the part of $\mathbb{R} \times [1, 2] \times T_{p+}$ where $|\varphi_+ - q^S| + |\hbar_+| < c_0 \varepsilon$.

It follows from these last observations and their $s \rightarrow \infty$ analogs that a version of z_p can be chosen so as to satisfy the demands of the lemma except possibly for the uniform

bounds on φ and ζ . The following says precisely what this means: Given $\varepsilon > 0$, there exists $z_{p,\varepsilon}$ and a neighborhood of $\partial_p S$ with two properties: First, the composition of u and then projection to Σ maps this neighborhood into T_{p+} . Second, the neighborhood has a parametrization by $\mathbb{R} \times [0, z_p)$ whereby u appears as the indicated graph with $\varphi(\cdot)$ and $\zeta(\cdot)$ smooth and such that $|\varphi(s, z) - \varphi^S(s)| + |\zeta(s, z)| < \varepsilon$.

The map $(s, z) \rightarrow \zeta(s, z) = \varphi(s, z) + i\zeta(s, z)$ is holomorphic on $\mathbb{R} \times (0, z_{p,\varepsilon})$ because u is J_{HF} -holomorphic. Fix a Euclidean half disk $D_+ \subset \mathbb{R} \times [0, z_{p,\varepsilon})$ of radius $\frac{1}{8}z_{p,\varepsilon}$ with center on any given $(s_0, 0)$ point. The Schwarz reflection argument from Theorem 24 in [1] can be repeated here to extend ζ as a holomorphic function on the disk in \mathbb{C} where the complex coordinate $s + iz$ is such that $|s - s_0|^2 + |z|^2 \leq \frac{1}{64}z_{p,\varepsilon}^2$. The boundary values on the full disk of this extended function $s + iz \rightarrow \zeta(s + iz)$ are uniformly bounded. It follows as a consequence of the Cauchy integral formula that its derivatives to any given order are uniformly bounded on the concentric disk with radius $\frac{1}{16}z_{p,\varepsilon}$ with this bound independent of s_0 . This implies in particular that

$$(6-5) \quad |\zeta(s, z)| + z^{-1}|\varphi(s, z) - \varphi^S(s)| \leq c_0|z|$$

in this smaller disk. This is the inequality asserted by the lemma for the norms of $\varphi - \varphi^S$ and ζ . A similar argument gives the asserted bound for the norms of the respective s -derivatives, and likewise the asserted bound on the norms of the derivatives to any given order. □

6C Neighborhoods of $\underline{S} \setminus S$

This subsection says more about the map u near a point in $\underline{S} \setminus S$. The discussion that follows talks about a neighborhood of a given negative point. There is an analogous story for neighborhoods of the positive points. The discussion here has two parts.

Part 1 Let q_S denote a given negative point of $\underline{S} \setminus S$, and let p and p' denote the respective index 1 and index 2 critical points of f that label the nearby components of ∂S . Parametrize the $\underline{S} \setminus q_S$ part of a neighborhood of q_S by $(-\infty, 0) \times [1, 2]$ as in the proof of Lemma 6.2. As noted in this same proof, the constant s slices in the u -image of the domain converge in $[1, 2] \times \Sigma$ to a curve of the form $[1, 2] \times q$ with $q \in C_{p+} \cap C_{p'}$. This fact has the following consequence: there exists $s_0 > 1$ such that the image via u of the complement of some compact set in $(-\infty, -1] \times [1, 2]$ has intersection number 1 with each constant (s, t) slice of $(-\infty, s_0] \times [1, 2] \times \Sigma$. Since u is pseudoholomorphic, the image via u of the complement of a compact set in

$(-\infty, -1] \times [1, 2]$ is a pseudoholomorphic subvariety. Taking this subvariety to be the pseudoholomorphic subvariety C in Lemma 4.7 leads to the following observation: The composition of first the map u and then projection to $(-\infty, 2s_0] \times [1, 2]$ restricts to its inverse image in $(-\infty, 0) \times [1, 2]$ as a diffeomorphism. Moreover, this diffeomorphism is complex analytic by virtue of the fact that J_{HF} maps $\frac{\partial}{\partial s}$ to $\frac{\partial}{\partial t}$. This understood, there exists $s_1 \geq 2s_0$ such that the map u when written using these coordinates can be written as a graph over $(-\infty, s_1] \times [1, 2]$ using a map, ψ , from $(-\infty, s_1] \times [1, 2]$ into $T_{p_+} \cap T_{p'_-}$. This is to say that q_S has a neighborhood in \underline{S} whose intersection with S can be parametrized by $(-\infty, s_1] \times [1, 2]$ such that the map u appears as the map

$$(6-6) \quad (s, t) \rightarrow (s, t, \psi(s, t)).$$

The map ψ is smooth; it maps $(-\infty, s_1] \times \{1\}$ to C_{p_+} and it maps $(-\infty, s_1] \times \{2\}$ to $C_{p'_-}$; and it is such that $\lim_{s \rightarrow \infty} \sup_{t \in [1, 2]} \text{dist}(\psi(s, t), q) = 0$.

Part 2 It proves useful to holomorphically identify the domain of ψ to the complement of 0 in the $|\zeta| \leq e^{-\frac{\pi}{2}s_1}$ part of the first quadrant of \mathbb{C} , the quadrant where both $\text{re}(\zeta) \geq 0$ and $\text{im}(\zeta) \geq 0$. This part of \mathbb{C} is denoted by D_{++} . This identification of the domain of ψ is done using the map

$$(6-7) \quad (s, t) \rightarrow \zeta = e^{\frac{\pi}{2}(s+i(t-1))}.$$

Use U to denote q 's component of $T_{p_+} \cap T_{p'_-}$. The coordinates (φ_+, \hbar_+) can be used as Euclidean coordinates on U , and no generality is lost by assuming that q has coordinate $\varphi_+ = 0$. This parametrization is defined for $|\varphi_+| + |\hbar_+| \leq c_0^{-1}$. It follows from what is said in Part 3 of Section 1C that the respective intersections of C_{p_+} and $C_{p'_-}$ with U appear in terms of the coordinates (φ_+, \hbar_+) as the $\hbar_+ = 0$ axis and the $\varphi_+ = 0$ axis. Meanwhile, the first and fourth bullets in (6-1) say that J_{HF} on $\mathbb{R} \times [1, 2] \times U$ acts to send $\frac{\partial}{\partial s}$ to $\frac{\partial}{\partial t}$ and $\frac{\partial}{\partial \varphi_+}$ to $\frac{\partial}{\partial \hbar_+}$.

Using this notation, the map ψ takes the form

$$(6-8) \quad \zeta \rightarrow (\varphi_+ = \varphi(\zeta), \hbar_+ = \varsigma(\zeta)),$$

where φ and ς are \mathbb{R} -valued functions. What was said above about J_{HF} implies that the \mathbb{C} -valued function $\lambda = \varphi + i\varsigma$ is holomorphic on $D_{++} \setminus 0$. What was said subsequent to (6-6) about ψ says the following about λ : First, $\lim_{\zeta \rightarrow 0} \lambda = 0$. Second, the imaginary part of λ is zero on the $\text{im}(\zeta) = 0$ boundary of D_{++} and the real part of λ is zero on the $\text{re}(\zeta) = 0$ boundary of D_{++} .

With the preceding understood, a two-fold application of the Schwarz reflection trick from Theorem 24 in [1] extends the domain of the map λ from D_{++} as a holomorphic map whose domain is the complement of the origin in the full $|\zeta| \leq e^{-\frac{\pi}{2}s_1}$ disk in \mathbb{C} . This extended version of λ also limits to zero as ζ limits to 0.

It follows from what was just said that λ and thus ψ extend across the origin of \mathbb{C} so as to define a holomorphic map from a disk about the origin in \mathbb{C} to D with the origin going to q . Such a map has the form

$$(6-9) \quad \zeta \rightarrow w(\zeta) = r\zeta^n + \mathcal{O}(\zeta^{n+1}),$$

where $n \in \{1, 2, \dots\}$ and where $r \in \mathbb{C} \setminus 0$. This last observation with (6-7) and (6-8) supply the desired picture of u near the given point in $\underline{S} \setminus S$.

6D The $\mathbb{R} \times M_{(1,2)}$ -image of u

Of interest in Section 7 and in [12] is the u -image of the interior of S with the latter viewed as a subset of $\mathbb{R} \times M_{(1,2)}$. The u -image of the part near ∂S is of particular concern in this subsection, in Lemma 6.4.

To set notation, introduce S_0 to denote $S \setminus \partial S$. The composition of first u and then projection from $\mathbb{R} \times M_{(1,2)}$ to $M_{(1,2)}$ defines a map from S_0 to $M_{(1,2)}$ that is denoted by u_M . Lemma 6.4 refers to the spherical coordinate $\theta = \theta_+$ or $\theta = \theta_-$ that is introduced in Section 1A for the radius δ_* coordinate ball centered at a given critical point of f with index respectively 1 or 2.

Lemma 6.4 *Let (S, u) denote an admissible pair with it understood that $u(S)$ can have interior singular points. View $u(S_0)$ as a J_{HF} -holomorphic subvariety in $\mathbb{R} \times M_{(1,2)}$. There exists $\kappa \geq \delta_*^{-1}$ with the following significance: Fix $z \in (0, \kappa^{-4})$.*

- *The function $1 - 3 \cos^2 \theta$ is bounded below by κ^{-1} on the part of $u_M(S_0)$ that lies in the radius δ_* coordinate ball centered on any given index 1 or index 2 critical point of f .*
- *The part of $u_M(S_0)$ where $f < 1 + z$ lies in the union of the radius $\kappa z^{1/2}$ coordinate balls centered on the index 1 critical points of f . The part where $f > 2 - z$ lies in the union of the radius $\kappa z^{1/2}$ coordinate balls centered on the index 2 critical points of f .*

Proof Fix an index 1 critical point of f and introduce the function h from (4-1) for the corresponding version of \mathcal{H}_p . Lemma 6.3 bounds

$$\frac{|h|}{|f-1|} = \frac{\cos \theta \sin^2 \theta}{1-3 \cos^2 \theta}$$

on the part of $u_M(S_0)$ that lies in the radius δ_* coordinate ball centered at p . In the notation of Lemma 6.2, this is $|\zeta(s, z)|/z$. Lemma 6.3 also bounds the function $|h|/|2-f| = \cos \theta \sin^2 \theta / (1-3 \cos^2 \theta)$ on the part of $u_M(S_0)$ that lies in the radius δ_* coordinate ball centered at any index 2 critical point of f . These bounds imply what is asserted by the first bullet of the lemma. The bound in the first bullet implies the second bullet. For example, if p is an index 1 critical point of f , and if q is near p and $f(q) - 1$ is small with $1 - 3 \cos^2 \theta \geq \kappa^{-1}$, then the distance from p to q is bounded by $\kappa^{1/2}(f(q) - 1)^{1/2}$. \square

6E Linear operators and admissible pairs

The purpose of this subsection is to summarize some of what is said in [13] about the operator that defines the first-order deformations of an admissible pair. The discussion here has six parts. In what follows, (S, u) is a given admissible pair.

Part 1 The image via u of S_0 is a J -holomorphic submanifold of $\mathbb{R} \times (1, 2) \times \Sigma$ and so its normal bundle inherits from J a complex line bundle structure. Use $N_S \rightarrow S_0$ to denote the pullback of this normal bundle. The map u extends as an embedding of the closed surface S into $\mathbb{R} \times [1, 2] \times \Sigma$. What follows describes the extension of N_S to ∂S .

To start, let p denote an index 1 critical point of f , and let $\partial_p S$ denote the corresponding boundary component. Lemma 6.3 describes the map u on a certain parametrized neighborhood in S of $\partial_p S$. With respect to this parametrization, the normal bundle appears as the pullback via u of the restriction to the neighborhood τ_{p+} of $T\Sigma$. This bundle has its basis of sections $\left\{ \frac{\partial}{\partial \varphi_+}, \frac{\partial}{\partial h_+} \right\}$ and the bundle complex structure maps the section $\frac{\partial}{\partial \varphi_+}$ to $\frac{\partial}{\partial h_+}$. There is the analogous basis and description of the complex line bundle structure for N_S on a neighborhood of an index 2 critical point version of $\partial_p S$.

Note that the trivializations given above for respective index 1 and index 2 critical point versions of $\partial_p S$ are compatible on a small neighborhood of a given point in $\underline{S} \setminus S$, this being a consequence of what is said in the fourth bullet of (6-1).

Given that u is an embedding, it maps suitable neighborhoods of distinct negative points in $\underline{S} \setminus S$ to distinct points in $T_- \cap T_+$, and it likewise maps neighborhoods of distinct positive points in $\underline{S} \setminus S$ to distinct points in $T_- \cap T_+$. These facts about u have an important consequence that is described in a moment. The statement of this consequence refers to a certain hermitian fiber metric on N_S . This fiber metric is induced by the Riemannian metric on $\mathbb{R} \times [1, 2] \times \Sigma$ that is defined by the symplectic form $ds \wedge dt + w_\Sigma$ and its compatible almost complex structure J_{HF} .

Here is the promised consequence: There exists $\rho_S > 0$ and an *exponential* map ϵ_S from N_S to $\mathbb{R} \times [1, 2] \times \Sigma$ that embeds the radius ρ_S disk subbundle as a neighborhood of $u(S)$ that contains the set of points in $\mathbb{R} \times [1, 2] \times \Sigma$ with distance $c_0^{-1} \rho_S$ or less from S . In addition, this number ρ_S and the map ϵ_S can be chosen so that the image of each fiber disk of radius ρ_S is a J_{HF} -holomorphic disk in $\mathbb{R} \times [1, 2] \times \Sigma$. Such a choice proves useful in subsequent sections. A version of ϵ_S with these properties can be constructed by mimicking what is done in Lemma 5.4 of [20]. Use $N_0 \subset N_S$ to denote this radius ρ_S disk subbundle.

Part 2 Suppose that $u': S \rightarrow \mathbb{R} \times [1, 2] \times \Sigma$ is a smooth map that obeys Lipshitz's conditions (M1)–(M6) in [13], which is to say it is described by the bullets in (6-2) except that it need not be J_{HF} -holomorphic. If $\sup_{y \in S} |u(y) - u'(y)| < c_0^{-1} \rho_S$, then the image of u' will lie in $\epsilon_S(N_0)$. If u' also has algebraic and geometric intersection 1 with each fiber of N_S , then u' factors as $\epsilon_S \circ \eta$ where η is a smooth section of N_S . Assuming that this is the case, the surface $u'(S_0)$ will be J_{HF} -holomorphic if and only if η obeys an equation that has the schematic form

$$(6-10) \quad \bar{\partial}\eta + \tau_1(\eta) \cdot \partial\eta + \tau_0(\eta) = 0,$$

where the notation is as follows: First, $\bar{\partial}$ signifies the d -bar operator on sections of N_S as defined using the hermitian metric to give the bundle a holomorphic structure. What is written as ∂ is the adjoint operator. Meanwhile, $\tau_1: N_0 \rightarrow N_S \otimes \text{Hom}(T^{1,0}S; T^{0,1}S)$ and $\tau_0: N_0 \rightarrow N_S \otimes T^{0,1}S$ are smooth, fiber-preserving maps that vanish along the zero section. If η obeys (6-10), then the pair (S, u') is admissible if it is understood that S has the complex structure that is induced via J_{HF} by the embedding u' .

As explained next, the nonlinear differential operator that is depicted on the left-hand side of (6-10) supplies an \mathbb{R} -linear operator from $C^\infty(S; N_S)$ to $C^\infty(S; N_S \otimes T^{0,1}S)$. Write the left-hand side of (6-10) as a map $\mathcal{F}: C^\infty(S; N_0) \rightarrow C^\infty(S; N_S \otimes T^{0,1}S)$. Fix $\eta \in C^\infty(S; N_S)$. The operator in question sends η to $D_S \eta = \lim_{\epsilon \rightarrow 0} \epsilon^{-1} \mathcal{F}(\epsilon \eta)$.

This operator has the form

$$(6-11) \quad D_S \eta = \bar{\partial} \eta + v \eta + \mu \bar{\eta},$$

where v is a smooth, bounded section of $C^\infty(S; T^{0,1}S)$ and μ is a smooth, bounded section of $C^\infty(S; N_S^2 \otimes T^{0,1}S)$.

Part 3 The operator D_S has a Fredholm version that comes by suitably choosing its domain. This domain is a certain completion of a linear subspace of $C^\infty(S; N_S)$. The latter subspace is denoted by $C_\partial^\infty(S; N_S)$. Sections in $C_\partial^\infty(S; N_S)$ are constrained on ∂S and in their behavior as points in $\underline{S} \setminus S$ are approached. The constraints that specify membership in $C_\partial^\infty(S; N_S)$ are given in a moment. The statement of the constraints refers implicitly to the identification from Part 1 of N_S along an index 1 critical point version of $\partial_p S$ with the pullback via u of $T\Sigma|_{C_{p+}}$. Note in particular that TC_{p+} sits in this bundle as a real 1–dimensional subbundle. Likewise, N_S along an index 2 critical point version of $\partial_p S$ is identified with the pullback via u of $T\Sigma|_{C_{p-}}$, and TC_{p-} is likewise viewed as a real line in the latter.

Here are the constraints: a section $\eta \in C^\infty(S; N_S)$ sits in the subspace $C_\partial^\infty(S; N_S)$ if and only if

- (6-12) • η has compact support,
- η along an index 1 critical point version of $\partial_p S$ sits in u^*TC_{p+} ,
 - η along an index 2 critical point version of $\partial_p S$ sits in u^*TC_{p-} .

The completion of $C_\partial^\infty(S; N_S)$ that defines the domain of S uses the Sobolev L^2_1 inner product. Likewise, the image Hilbert space uses the L^2 inner product on the space of compactly supported sections of $N_S \otimes T^{0,1}S$. To elaborate, use ∇ to denote the covariant derivative on sections of N_S , its tensor powers and tensor products with powers of T^*S that is induced by the Riemannian metric on $\mathbb{R} \times [1, 2] \times \Sigma$. The respective norms that define the domain Hilbert space and image Hilbert space for D_S are the square roots of the functions on $C_\partial^\infty(N; S)$ and $C^\infty(N \otimes T^{0,1}S)$ that assign to a given section η the value

$$(6-13) \quad \int_S (|\nabla \eta|^2 + |\eta|^2) \quad \text{and} \quad \int_S |\eta|^2.$$

The index of this Fredholm version of D_S is computed by Lipshitz [13]. Note in this regard that Lipshitz uses weighted versions of these norms. The upcoming Part 4 explains why the weights are not necessary.

An admissible pair (S, u) is said in what follows to be *unobstructed* when the corresponding Fredholm version of D_S has trivial cokernel. Lipshitz proves in Section 3 of [13] that there is a C^∞ -residual set of allowed complex structures on $\mathbb{R} \times [1, 2] \times \Sigma$ with only unobstructed admissible pairs.

Part 4 A bounded linear operator between Banach spaces is Fredholm if it has closed range, finite-dimensional kernel and finite-dimensional cokernel. Let \mathbb{H}_S and \mathbb{L}_S denote the respective domain and image Hilbert spaces for D_S . The fact that D_S has closed range in \mathbb{L}_S and also finite-dimensional kernel follows from two facts that are asserted and then proved in a moment. The notation used has $\|\cdot\|$ denoting the L^2 -norm whose square is depicted in (6-13). What is denoted by η refers to an element in $C^\infty_\partial(N; S)$. What is denoted by c is a constant that depends on S and is greater than 1. What follows are the two promised facts:

- (6-14) • $\|D_S \eta\|^2 \geq c^{-1} \|d\eta\|^2 - c \|\eta\|^2$.
- There exists $s_* > 1$ such that if η has support where $|s| > s_*$, then $\|D_S \eta\|^2 \geq c^{-1} \|\eta\|^2$.

A standard argument using the Rellich lemma uses (6-14) to deduce that D_S has closed range and finite-dimensional kernel. Given that the range is closed, the cokernel is isomorphic to the kernel of the adjoint of D_S ; this a bounded operator that maps \mathbb{L}_S to \mathbb{H}_S . Integration by parts and some standard arguments about elliptic regularity identify the kernel of the latter operator with that of a certain bounded, differential operator that also maps \mathbb{H}_S to \mathbb{L}_S . The operator in question is denoted by $D_S^\#$; as explained below, what is asserted by (6-14) is true for $D_S^\#$, and thus its kernel is also finite-dimensional.

Turn now to (6-14). The first bullet is derived using a standard integration by parts. The conditions from the second and third bullets of (6-12) ensure the vanishing of boundary terms. To see about the second bullet in (6-14), parametrize the map u on a neighborhood of a given negative point in $\underline{S} \setminus S$ as described in Part 1 of Section 1C. By way of a reminder, the parametrization is by $(-\infty, s_1] \times [1, 2]$ and is such that the map u appears as a graph $(s, t) \rightarrow (s, t, \varphi_+ = \varphi(s, t), \hat{h}_+ = \zeta(s, t))$ where (φ, ζ) are \mathbb{R} -valued functions that map the domain into the relevant component of $T_+ \cap T_-$.

The coordinate φ_+ can be assumed to be \mathbb{R} -valued on the image of the pair (φ, ζ) , and such that the point $(\varphi_+ = 0, \hat{h}_+ = 0)$ is the nearby point in $C_+ \cap C_-$. The locus C_+ appears as the $\hat{h}_+ = 0$ axis and C_- as the $\varphi_+ = 0$ axis. This understood, the pair

(φ, ζ) is such that $\varphi + i\zeta$ is a holomorphic function of $s + it$. In addition, this pair is constrained so that $\lim_{s \rightarrow \infty} (|\varphi| + |\zeta|) = 0$, so that $\zeta = 0$ where $t = 1$, and so that $\varphi = 0$ where $t = 2$.

A section of the normal bundle N_S over this part of S pulls back via this parametrization as a map from $(-\infty, s_1] \times [1, 2]$ to \mathbb{C} with it understood that the \mathbb{C} -valued 1-form $d\varphi + i d\zeta$ is used to identify the normal bundle with the product complex line bundle. Write a such a map η as $\eta_1 + i\eta_2$ with η_1 and η_2 both real-valued functions. Then the map η comes from $C^\infty(S; N)$ if $\eta_2 = 0$ on $(-\infty, s_1] \times \{1\}$ and $\eta_1 = 0$ on $(-\infty, s_1] \times [1, 2]$. Meanwhile, the operator D_S here is $\frac{1}{2}(\partial_s + i\partial_t)$ and so

$$(6-15) \quad D_S \eta = \frac{1}{2}(\partial_s \eta_1 - \partial_t \eta_2) + \frac{i}{2}(\partial_s \eta_2 + \partial_t \eta_1).$$

As a consequence, if η has support where $s > s_1$, then

$$(6-16) \quad \|D_S \eta\|^2 = \frac{1}{4}(\|\partial_s \eta_1\|^2 + \|\partial_s \eta_2\|^2 + \|\partial_t \eta_1\|^2 + \|\partial_t \eta_2\|^2).$$

With (6-16) understood, introduce C_*^∞ to denote the subspace of smooth, \mathbb{R} -valued functions on $[1, 2]$ that vanish at one of the endpoints of $[1, 2]$. The function

$$(6-17) \quad \zeta \rightarrow Q_*(\zeta) = \int_1^2 |\partial_t \zeta|^2 dt$$

on C_* is bounded away from zero in the sense that

$$(6-18) \quad Q_*(\zeta) > \frac{\pi^2}{4} \int_1^2 |\zeta|^2 dt.$$

This fact with (6-16) and its analog for the positive points in $\underline{S} \setminus S$ imply what is asserted by the second bullet in (6-14).

The operator $D^\#$ is the formal, L^2 adjoint of D_S , and so it is a first-order operator with leading derivative term given by the formal L^2 adjoint of $\bar{\partial}$. This being the case, an integration by parts can be employed to prove the $D^\#$ version of the first bullet in (6-14). The proof of the second bullet is just like that for D_S .

Part 5 Let \mathcal{A}_{HF} denote the space of admissible pairs (S, u) where the topology is defined as follows: The open neighborhoods of a pair (S, u) are generated by a basis whose sets are labeled a positive real number, ε . An admissible pair (S', u') is in the corresponding set when two conditions are met:

$$(6-19) \quad \bullet \quad \sup_{z \in S} \text{dist}(u(z), u'(S')) + \sup_{z \in S'} \text{dist}(u(S), u'(z)) < \varepsilon.$$

- $|\int_S u^* \mu - \int_{S'} u'^* \mu| < \varepsilon$ if μ is any 2-form on $\mathbb{R} \times [1, 2] \times \Sigma$ with $|\mu| \leq 1$, $|\nabla \mu| \leq \frac{1}{\varepsilon}$ and compact support on $|s| < \frac{1}{\varepsilon}$.

Suppose that (S, u) is an admissible pair. What follows is a direct consequence of the definition in (6-19). The pair (S, u) has a neighborhood in \mathcal{A}_{HF} with the following property: if (S', u') is in this neighborhood, then $u'(S')$ is in the radius ρ_S tubular neighborhood of $u(S)$ and it has geometric intersection number 1 with each fiber. In particular, it can be written as the image of S via the map $\varepsilon_S \circ \eta$ where η is a section of N_0 that obeys (6-10) and the second two bullets of (6-12). Moreover, $|\eta|$ has limit zero as $|u^* s| \rightarrow \infty$. The next lemma is a consequence of these last observations.

Lemma 6.5 *Let (S, u) denote an admissible pair. There is a neighborhood of (S, u) in the space \mathcal{A}_{HF} and, for each $k \in \{0, 1, 2, \dots\}$, a constant κ_k with the following significance: Suppose that (S', u') is in this neighborhood. Then $u'(S')$ can be written as the image of $\varepsilon_S \circ \eta$ where η is a section of N_0 that obeys (6-10), the second two bullets in (6-12), and is such that $\lim_{|s| \rightarrow \infty} |\eta| = 0$. In addition,*

$$|\nabla^k \eta| \leq e^{-|u^* s|/\kappa_0} \sup_z |\eta|.$$

Proof Without the exponential factor, a bound given by the lemma for the derivatives of η follows directly using the sorts of elliptic regularity theorems that can be found in Chapter 6 of Morrey’s book [16]. Note in this regard that it follows from what is said in Section 6B and in Part 2 above that (6-10) near the boundary of S can be written so as to appear as the standard, linear Cauchy–Riemann equations. This being the case, regularity near the boundary can be proved using classical linear techniques.

The bound with the exponential factor can be proved using what is asserted in the second bullet of (6-14). Here is a brief sketch of how this is done: Fix $\varepsilon > 0$ and $T \gg 1$. Let $\chi: \mathbb{R} \rightarrow [0, 1]$ denote a smooth function which is equal to 1 where $s < 0$ and equal to 0 where $s > 1$. Take the square of the norm of the expression on the left-hand side of (6-10), multiply by $e^{-\varepsilon|u^* s|} \chi(T - |u^* s|)$ and integrate over S . Integrate by parts using the boundary conditions given by the second bullet in (6-12). Use the fact that $|\eta|$ limits to zero as $|u^* s| \rightarrow \infty$ with the second bullet in (6-12) to see that the $\varepsilon \rightarrow 0$ limit of the resulting integral can be taken, and that this limit implies the bound

$$(6-20) \quad \int_{|u^* s| > T} |\eta|^2 \leq c_0 \int_{|u^* s| \in [T-1, T]} |\eta|^2.$$

The sequence $\{T = n\}_{n=1,2,\dots}$ of such inequalities implies directly that

$$(6-21) \quad \int_{|u^*s|>T} |\eta|^2 \leq c_0 e^{-|u^*s|/c_0} \sup_s |\eta|.$$

This last bound with the aforementioned elliptic regularity techniques will give the bounds that are asserted by the lemma. \square

Part 6 The next lemma speaks to the relationship between the kernel of the operator D_S in (6-11) and solutions to (6-10). The lemma refers to the kernel and cokernel of the Fredholm version of D_S that maps \mathbb{H}_S to \mathbb{L}_S . The lemma also uses \mathbb{I}_S to denote the tautological inclusion map from $\text{kernel}(D_S)$ into \mathbb{H}_S .

Lemma 6.6 *Let (S, u) denote a given admissible pair. There exists a constant $\kappa \geq 1$, a neighborhood $U \subset \text{kernel}(D_S)$ of the origin, and smooth maps $f: U \rightarrow \text{cokernel}(D_S)$ and $q: U \rightarrow \mathbb{H}_S \cap C^\infty(S; N_0)$ with the following properties:*

- $f(0) = 0$, $q(0) = 0$ and $dq|_0 = \mathbb{I}_S$.
- Let \mathcal{N} denote the space of smooth sections of N_0 that obey (6-10), the second and third bullets of (6-12) and have pointwise norm bounded by $\kappa^{-1}\rho_S$. Then q maps $f^{-1}(0)$ homeomorphically onto \mathcal{N} .
- The space \mathcal{N} parametrizes a neighborhood of (S, u) in \mathcal{A}_{HF} in the following way: Let $\eta \in \mathcal{N}$. The corresponding admissible pair (S', u') is defined by the condition that $u'(S') = (\epsilon_S \circ \eta)(S)$. This is to say that $S' = S$ but with the complex structure that is induced by the J_{HF} -holomorphic embedding $\epsilon_S \circ \eta$.

Proof The assertions of the first two bullets are straightforward applications of the inverse function theorem following standard arguments pioneered by Kuranishi. This is done by viewing the expression on the left-hand side of (6-10) as defining a map from a ball in a certain Banach space to a second Banach space. The only subtlety is that the spaces \mathbb{H}_S and \mathbb{L}_S cannot be used because the L^2_1 -norm does not control the supremum norm. Even so, this is a standard issue in the theory of J -holomorphic curves and is dealt with by using a domain and range that give slightly stronger control over the derivatives of the section; see for example [15]. The third bullet follows directly from what is said in Part 5. \square

The next lemma is a corollary to the preceding lemma. It refers to the \mathbb{R} action on \mathcal{A}_{HF} where the generator $1 \in \mathbb{R}$ acts on any given pair (S, u) to give the pair (S, u^1)

where u_1 is the composition of first u and then translation by 1 along the \mathbb{R} factor of $\mathbb{R} \times [1, 2] \times \Sigma$.

Lemma 6.7 *Let $\mathcal{A}_{\text{smooth}}$ denote the subspace of the moduli space of admissible pairs that contains the unobstructed pairs. The subspace $\mathcal{A}_{\text{smooth}}$ is an open, \mathbb{R} -invariant set with the structure of a smooth manifold. Moreover, if $(S, u) \in \mathcal{A}_{\text{smooth}}$, then the corresponding version of the map \mathfrak{f} of Lemma 6.6 defines a smooth embedding from a small radius ball about the origin in $\text{kernel}(D_S)$ onto an open neighborhood of (S, u) in $\mathcal{A}_{\text{smooth}}$.*

Proof The assertion that a neighborhood in \mathcal{A}_{HF} of a given pair (S, u) is homeomorphic to a neighborhood of 0 in the kernel of D_S follows directly from Lemma 6.6. The smooth structure on this neighborhood is defined by requiring that this same homeomorphism define a diffeomorphism. The inverse function theorem can be used to verify that the associated transition functions are C^∞ . □

7 Ech-HF subvarieties and Heegaard Floer curves

This section explains how the $\mathbb{R} \times M_\delta$ part of any given ech-HF subvariety determines pseudoholomorphic subvarieties of the sort used in Lipshitz’s reformulation [13] of Heegaard Floer homology. The observations of this section are summarized by Proposition 7.1 in Section 7A, Proposition 7.2 in Section 7B and Proposition 7.3 in Section 7C.

7A Ech-HF subvarieties and singular admissible sets

Singular admissible sets enter the embedded contact homology story via the upcoming Proposition 7.1. To set things up for this proposition, it is necessary to add one additional constraint to those in Section 3A for the sorts of almost complex structures on $\mathbb{R} \times Y$ that are used to define the notion of an ech-HF subvariety. To this end, fix an almost complex J_{HF} that obeys (6-1). What follows is the extra constraint on the almost complex structure J of Section 3A:

$$(7-1) \quad J = J_{\text{HF}} \quad \text{on the } f^{-1}(1, 2) \text{ part of } \mathbb{R} \times M_\delta.$$

Note in this regard that Constraints 1 and 2 in Section 3A are equivalent on $\mathbb{R} \times M_\delta$ to what is required by the first and second bullets in (6-1). Constraint 3 in Section 3A is

equivalent to the rule in the third bullet of (6-1). Meanwhile, Constraints 4 and 5 in Section 3A are compatible with the constraints in the fourth bullet of (6-1).

A given singular admissible set $Z = ((S, u), \vartheta_\Sigma)$ defines an integral, 2-dimensional current on any given compact subset of $\mathbb{R} \times [1, 2] \times \Sigma$, this given by integration over the 2-chain $[u(S)] + \sum_{(Z,m) \in \vartheta_\Sigma} m[Z]$. This current is denoted by $[Z]$. The pairing of any given 2-form μ with $[Z]$ is denoted by $[Z](\mu)$. The support of the current $[Z]$ is the subvariety $u(S) \cup (\bigcup_{(Z,m) \in \vartheta_\Sigma} Z)$. The latter is denoted by $|Z|$. Note that integration over $|Z|$ does not give $|Z|$ unless all integers from ϑ_Σ are equal to 1.

Proposition 7.1 *Fix an almost complex structure J_{HF} that obeys the conditions in (6-1). Given $\mathcal{I} \geq 1$ and $\varepsilon > 0$, there exists $\kappa > \varepsilon^{-1}$ with the following significance: Define the geometry of Y using data δ, χ_0 and R with $\partial < \kappa^{-2}$, and fix an almost complex structure, J , on $\mathbb{R} \times Y$ that obeys the constraints in Section 3A plus (7-1). Fix an ech-HF subvariety C with $I_{ech}(C) < \mathcal{I}$ or with $\int_C \omega < \mathcal{I}$. There is a singular admissible set, $Z = ((S, u), \vartheta_\Sigma)$, with:*

- $\sup_{z \in (C \cap ([-1/\varepsilon, 1/\varepsilon] \times M_\varepsilon))} \text{dist}(z, |Z|) + \sup_{z \in (|Z| \cap ([-1/\varepsilon, 1/\varepsilon] \times M_\varepsilon))} \text{dist}(z, C) < \varepsilon.$
- Let μ denote a smooth 2-form on $\mathbb{R} \times M_\varepsilon$ with sup norm bounded by 1, covariant derivative norm bounded by ε^{-1} and compact support where $|s| < \varepsilon^{-1}$.

Then

$$\left| \int_C \mu - [Z](\mu) \right| \leq \varepsilon.$$

Proof The proof that follows has eight steps.

Step 1 Suppose that no such κ exists. There would necessarily exist a sequence $\{(D_n, C_n)\}_{n=1,2}$, of pairs of the following sort: First, any given $n \in \{1, 2, \dots\}$ version of D_n is a data set consisting of 4-tuples $(\delta_n, \chi_{0,n}, R_n, J_n)$ such that $\delta_n < 1/n$, and such that the geometry of Y and $\mathbb{R} \times Y$ can be defined as in Section 1A and so that the results in the previous sections can be invoked using $\delta = \delta_n, \chi_0 = \chi_{0,n}, R = R_n$ and $J = J_n$. Second, for each such integer n , what is denoted by C_n is an ech-HF subvariety as defined by J_n . This subvariety is such that $I_{ech}(C_n) \leq \mathcal{I}$ or that $\omega(C_n) < \mathcal{I}$. However, there is no pair (u, S) that satisfies the $C = C_n$ version of what is required by the two bullets of the lemma.

Step 2 With Proposition 5.1 in mind, use Proposition 5.5 to find a finite set, ϑ , and a subsequence of $\{C_n\}_{n=1,2,\dots}$, henceforth renumbered consecutively from 1, with properties that are described in this step and in Steps 3–5.

Consider first the properties of $\{C_n\}_{n=1,2,\dots}$. The $s \rightarrow \infty$ limit of the constant s slices of a given $n \in \{1, 2, \dots\}$ version of C_n defines an element in $\mathcal{Z}_{\text{ech},M}$ and thus an HF cycle of length G . By the same token, the $s \rightarrow -\infty$ limit of the constant s slices of C_n also defines an HF cycle of length G . The sequence $\{C_n\}_{n=1,2,\dots}$ is chosen so that the same HF cycle is defined by the $s \rightarrow \infty$ limit of the constant s slices of each $n \in \{1, 2, \dots\}$ version of C_n , and so that the same HF cycle is defined by the $s \rightarrow -\infty$ limit of the constant s slices of each such C_n . These respective cycles are denoted by \widehat{v}_+ and \widehat{v}_- .

Step 3 Additional properties of $\{C_n\}_{n=1,2,\dots}$ are stated in a moment. To set the stage, use Proposition 5.8 to find $z_\diamond < \delta_*^2$ so that the conclusions of Proposition 5.8 hold for each $n \in \{1, 2, \dots\}$ version of C_n . If n is sufficiently large, then $\delta_n < \frac{1}{100} z_\diamond^{1/2}$ and so C_n intersects each $\mathbb{R} \times \{1 + z_\diamond\}$ member of \mathcal{M}_Σ precisely once in the radius δ_* coordinate ball centered on each index 1 critical point of f . The subvariety C_n likewise intersects each $\mathbb{R} \times \{2 - z_\diamond\}$ member of \mathcal{M}_Σ precisely once in the radius δ_* coordinate ball centered on each index 2 critical point of f .

What follows is a consequence of the preceding observations: Fix $z \in (0, z_\diamond]$. Then all sufficiently large n versions of C_n intersect each $\mathbb{R} \times \{1 + z\}$ member of \mathcal{M}_Σ precisely once in the radius δ_* coordinate ball centered on each index 1 critical point of f . Meanwhile, C_n intersects each $\mathbb{R} \times \{2 - z\}$ member of \mathcal{M}_Σ precisely once in the radius δ_* coordinate ball centered on each index 2 critical point of f .

To say more about these intersections, fix z as above. The level set in question is an annulus, as can be seen using the coordinates $(t, \varphi_+, \widehat{h}_+)$ if p has index 1 or the coordinates $(t, \varphi_-, \widehat{h}_-)$ if p has index 2. In the former case, this is the $t = 1 + z$ level set, and the annulus is that where $|\widehat{h}_+| < \frac{4}{3\sqrt{3}} \delta_*^2$. In the latter case, the level set has $t = 2 - z$ and the annulus is that where $|\widehat{h}_-| < \frac{4}{3\sqrt{3}} \delta_*^2$. In either case, let Σ_p^z denote the part of this level set in the radius δ_* coordinate ball centered at p . It follows from Lemma 4.7 that the intersection of C_n with $\mathbb{R} \times \Sigma_p^z$ is a smoothly embedded arc that appears as the graph of a map from \mathbb{R} to Σ_p^z . The $s \rightarrow -\infty$ limit of the latter map has distance $c_0 \delta_n$ or less from the intersection between Σ_p^z and the relevant integral curve of \mathfrak{v} from \widehat{v}_- , and the $s \rightarrow \infty$ limit has distance $c_0 \delta_n$ or less from the intersection between Σ_p^z and the relevant integral curve of \mathfrak{v} from \widehat{v}_+ .

Proposition 5.8 has the following additional consequence: fix $\rho \in (0, \delta_*]$; there exists $z_\rho \in (0, z_\diamond]$ such that all sufficiently large n versions of C_n intersect Σ_p^z in the radius ρ coordinate ball centered on p .

Step 4 Given the assumed bound on $\{I_{\text{ech}}(C_n)\}_{n=1,2,\dots}$, the two bullets of Proposition 5.1 can be used to apply Proposition 5.5 with U first being M_{δ_*} , then $M_{\delta_*/2}$, and successively for $U = M_{\delta_*/k}$ with $k \in \{1, 2, \dots\}$. These successive applications of Proposition 5.5 produce the following: First, a finite set ϑ whose typical element is a pair (Z, m) with $Z \subset \mathbb{R} \times M_{(1,2)}$ an irreducible, J_{HF} -holomorphic subvariety. These are such that no two pairs from ϑ share the same subvariety. Moreover:

$$(7-2) \quad \bullet \quad \sum_{(Z,m) \in \vartheta} \int_{Z \cap ([s, s+1] \times M_{(1,2)})} ds \wedge df = G \text{ for each } s \in \mathbb{R}.$$

$$\bullet \quad \sum_{(Z,m) \in \vartheta} m \int_Z \omega_\Sigma \leq \mathcal{I}.$$

Also produced is a subsequence (henceforth renumbered consecutively from 1) of $\{C_n\}_{n=1,2,\dots}$ such that if $\mathcal{K} \subset \mathbb{R} \times M_{(1,2)}$ is any given compact set, then:

(7-3) • We have

$$\lim_{n \rightarrow \infty} \left(\sup_{z \in (C_n \cap \mathcal{K})} \text{dist} \left(z, \bigcup_{(Z,m) \in \vartheta} Z \right) + \sup_{z \in \bigcup_{(Z,m) \in \vartheta} Z \cap \mathcal{K}} \text{dist}(z, C_n) \right) = 0.$$

• Let μ denote a smooth 2-form with compact support on \mathcal{K} . Then

$$\lim_{n \rightarrow \infty} \left(\int_{C_n} \mu - \sum_{(Z,m) \in \vartheta} m \int_Z \mu \right) = 0.$$

Let $\vartheta_\Sigma \subset \vartheta$ denote the subset of pairs whose subvariety component is mapped to a constant f slice of $M_{(1,2)}$ via the projection from $\mathbb{R} \times M_{(1,2)}$. Use ϑ_* to denote $\vartheta \setminus \vartheta_\Sigma$.

Step 5 Let p denote an index 1 or index 2 critical point of f . Fix $z \in (0, z_\diamond]$. Each sufficiently large n version of C_n has precisely one intersection point with $\mathbb{R} \times \Sigma_p^z$. This with (7-3) has two immediate implications: First, there is precisely one pair from ϑ_* whose image via the projection from $\mathbb{R} \times M_{(1,2)}$ intersects Σ_p^z . Second, each pair from ϑ_* has integer component 1. This being the case, view ϑ_* now as a set of distinct J_{HF} holomorphic subvarieties. What follows is a consequence of Lemma 4.2 and (7-3): View $M_{(1,2)}$ as $(1, 2) \times Z$ and $\mathbb{R} \times M_{(1,2)}$ as $\mathbb{R} \times (1, 2) \times \Sigma$. The subvariety $Z_0 = \bigcup_{Z \in \vartheta_*} Z$ has intersection number G with each constant $\mathbb{R} \times (1, 2)$ slice of $\mathbb{R} \times (1, 2) \times \Sigma$.

The preceding observations with Lemma 4.7 imply the following: there exists $\rho_* \in (0, \delta_*]$ such that if p is any given index 1 or index 2 critical point of f , then exactly one irreducible component of the subvariety Z_0 is mapped via the projection to $M_{(1,2)}$

into the radius ρ_* coordinate ball centered at p , and its intersection with the product of \mathbb{R} with this ball is a submanifold of $\mathbb{R} \times M_{(1,2)}$. Take ρ_* such that no submanifold from ϑ_Σ projects to the union of these radius ρ_* coordinate balls.

The next observation is a consequence both of Lemma 4.7 and what is said at the end of Step 3. Given $\rho \in (0, \rho_*]$, there exists $z_\rho \in (0, z_\diamond]$ such that if $z \in (0, z_\rho]$, then the subvariety from ϑ_* whose projection to $M_{(1,2)}$ intersects a given version of Σ_p^z does so in the radius ρ coordinate ball centered on the critical point p .

View $\mathbb{R} \times M_{(1,2)}$ as $\mathbb{R} \times (1, 2) \times \Sigma$ and view Z_0 as a subvariety in $\mathbb{R} \times (1, 2) \times \Sigma$. Given Lemma 4.7, what was said in the preceding two paragraphs has the following reinterpretation: There exists $z_1 \in (0, 1)$ such that the $\mathbb{R} \times (1, 1 + z_1)$ part of Z_0 is a smoothly embedded submanifold of $\mathbb{R} \times (1, 2) \times \Sigma$ that consists of G components, and each component is a graph over its image in $\mathbb{R} \times (1, 2)$. Moreover, this graph is defined by a map $\lambda: \mathbb{R} \times (1, 2) \rightarrow \Sigma$ whose image is in T_+ and is such that

$$(7-4) \quad \lim_{t \rightarrow 1} \sup_{s \in \mathbb{R}} \text{dist}(\lambda(s, t), C_+) = 0.$$

Each circle from C_+ has the limit points of precisely one $\mathbb{R} \times (1, 1 + z_1)$ component of Z_0 .

There is a completely analogous story for the $\mathbb{R} \times (2 - z_1, 2)$ part of Z_0 with it understood that T_- and C_- appears in lieu of T_+ and C_+ .

Items (7-2) and what is said above can be used with Lemma 5.6 to deduce the following: There exists $s_1 \geq 1$ such that the $s > s_1$ part of Z_0 is a disjoint union of G smooth, open strips. Each such strip is a graph over its image in $(s_1, \infty) \times (1, 2)$. Such a graph is defined by a map, $\psi: (s_1, \infty) \times (1, 2) \rightarrow \Sigma$ with image in $T_+ \cap T_-$. In addition, there exists a point $q \in C_+ \cap C_-$ such that

$$(7-5) \quad \lim_{s \rightarrow \infty} \sup_{t \in (1,2)} \text{dist}(\psi(s, t), q) = 0.$$

Finally, the G points so defined in $C_+ \cap C_-$ pair the set of index 1 critical points of f with the set of index 2 critical points of f .

The analogous conclusion holds for the $s < -s_1$ portion of Z_0 .

Step 6 The union of $Z_0 \subset \mathbb{R} \times (1, 2) \times \Sigma$ with the constant $\mathbb{R} \times (1, 2)$ slices from ϑ_Σ define a J_{HF} -holomorphic subvariety in $\mathbb{R} \times (1, 2) \times \Sigma$ with at worst a finite set of singular points. The fact that holomorphic structures on punctured surfaces extend over the punctures implies that there is a *model curve* for Z_0 . This model curve is a

smooth complex curve with a J_{HF} -holomorphic, almost-everywhere 1-to-1 map to $\mathbb{R} \times (1, 2) \times \Sigma$ whose image is Z_0 . Let S_0 denote this model curve and let u denote the map. The complement of a finite set in S_0 is identified by u with the complement of the singular points in Z_0 and this identification defines the complex structure on this part, and hence all of S_0 . As explained in the next steps, S_0 is the interior of a surface with boundary S , to which u extends so that (S, u) with ϑ_Σ defines a singular admissible set. Given what is said by (7-3), the existence of (S, u) puts the lie to the assumption that the conclusions of Proposition 7.1 do not hold for the sequence $\{C_n\}_{n=1,2,\dots}$. This contradiction is avoided only if Proposition 7.1 is true.

Step 7 This step describes S and the compact surface \underline{S} . The complement of S_0 in S is the disjoint union of $2G$ copies of \mathbb{R} , each these in 1-to-1 correspondence with the index 1 critical points of f . Let p denote a given index 1 critical point of p . And let $\partial_p S$ denote the corresponding boundary component of S . A neighborhood of $\partial_p S$ in S is described in the next paragraph.

To start the description, let Z_p denote the component of the $t < 1 + z_1$ part of Z_0 that lies in $\mathbb{R} \times (1, 1 + z_1) \times T_{p+}$ and let $S_{0p} \subset S_0$ denote $u^{-1}(Z_p)$. The map u identifies S_{0p} with Z_p . Since Z_p is a graph over $\mathbb{R} \times (1, 1 + z_1)$, this identification with Z_p followed by the projection identifies S_{0p} with $\mathbb{R} \times (1, 1 + z_1)$. This identification extends to a neighborhood of S_{0p} in S so as to identify the latter with $\mathbb{R} \times [1, 1 + z_1)$. This identification writes the component ∂S_p as $\mathbb{R} \times \{1\}$.

Let p now denote an index 2 critical point and let Z_p in this case denote the component of the $t > 2 - z_1$ part of Z_0 in $\mathbb{R} \times (2 - z_1, 2) \times T_{p-}$. The u -inverse image of Z_p is denoted by S_{0p} and it is identified with $\mathbb{R} \times (2 - z_1, 2)$ via the composition of u and the projection from Z_p to the $\mathbb{R} \times (2 - z_1, 2)$ factor. A neighborhood in S of this part of S_0 has the corresponding identification with $\mathbb{R} \times (2 - z_1, 2]$ with $\partial_p S$ appearing as $\mathbb{R} \times \{2\}$.

Consider next \underline{S} . The complement of S in \underline{S} consists of $2G$ points, half labeled positive and half labeled negative. What follows describes the positive points. There is a completely analogous description of the negative points, which is left to the reader. The set of positive points in $\underline{S} \setminus S$ is in 1-to-1 correspondence with the G points in $C_+ \cap C_-$ that appear in the G versions of (7-5). Let q denote such a point and let q_S denote its partner in $\underline{S} \setminus S$. Let p and p' denote the respective index 1 and index 2 critical point versions of C_{p+} and $C_{p'-}$ whose intersection contains q . The point q_S is contained in the closure of $\partial_p S$ and $\partial_{p'} S$. To elaborate, let Z_{0q+} denote the component

of the $s \geq s_1$ part of Z_0 that appears with q in (7-5). Let $S_{0q_+} \subset S_0$ denote the u -inverse image of Z_{0q_+} . The map u identifies S_{0q_+} with Z_{0q_+} and, as a consequence, the projection from Z_{0q_+} to $(s_1, \infty) \times (1, 2)$ identifies S_{0q_+} with $(s_1, \infty) \times (1, 2)$. With this identification understood, it follows that S_{0q_+} has a neighborhood in S that is diffeomorphic as a manifold with boundary to $(s_1, \infty) \times [1, 2]$ with the boundary component $(s_1, \infty) \times \{1\}$ in $\partial_p S$ and the component $(s_1, \infty) \times \{2\}$ in $\partial_{p'} S$.

Let ζ denote the standard complex coordinate on \mathbb{C} and introduce D_+ to denote the $|z| < e^{-\pi s_1}$ part of the closed upper half plane. Identify the $(s_1, \infty) \times [1, 2]$ with $D_+ \setminus 0$ via the map

$$(7-6) \quad (s, t) \rightarrow e^{-\pi(s+it)}.$$

This identification extends to a neighborhood of q_S so as to identify the latter with D_+ .

Step 8 Look at (6-2) to see that the only remaining task is that of extending the map u to the boundary of S . Note in this regard that (7-4) and its $t \rightarrow 2$ analog imply the following: Let p denote an index 1 critical point of f . If u extends to $\partial_p S$, then it maps the latter to $\mathbb{R} \times C_{p_+}$. By the same token, let p now denote an index 2 critical point of f . If u extends to $\partial_p S$, then it maps $\partial_p S$ to $\mathbb{R} \times C_{p_-}$. In either case, if the extension is smooth, then it is de facto an embedding of $\partial_p S$ by virtue of what is said in Step 5 about the map u on S_{0p} . By way of a reminder, let p denote an index 1 critical point of f . Step 5 writes a neighborhood of ∂S_p in S as $\mathbb{R} \times [1, 1 + z_1)$ and writes u on the $\mathbb{R} \times (1, 1 + z_1)$ part as the graph of the map λ from $\mathbb{R} \times (1, 1 + z_1)$ to T_{p_+} . The extension of u to $\mathbb{R} \times \{1\}$ is given by an extension of λ which maps $\mathbb{R} \times \{1\}$ to C_{p_+} .

The assumption in the fifth bullet of (6-1) implies that J_{HF} on $\mathbb{R} \times (1, 1 + z_1) \times T_+$ defines an s and t independent complex structure on T_+ such that $\varphi_+ + i\hbar_+$ defines a holomorphic parameter on a neighborhood of any given point. Parametrize a neighborhood of $\partial_p S$ in S as $\mathbb{R} \times [1, 1 + z_1)$ in the manner described above. Use s as before to denote the Euclidean coordinate on the \mathbb{R} factor of $\mathbb{R} \times [1, 1 + z_1)$, but write the Euclidean coordinate t on the $[1, 1 + z_1)$ factor as $1 + z$, where z is the Euclidean coordinate on $(0, z_1)$. Use the coordinates (φ_+, \hbar_+) to write λ as the map

$$(7-7) \quad (x, z) \rightarrow (\varphi_+ = \varphi^S(s, z), \hbar_+ = \zeta^S(s, z)),$$

where φ^S and ζ^S are smooth functions. The fact that u is J -holomorphic implies in particular that $\varphi^S + i\zeta^S$ must obey the Cauchy-Riemann equations, which is to say

that

$$(7-8) \quad \partial_s \varphi^S - \partial_z \zeta^S = 0 \quad \text{and} \quad \partial_z \zeta^S + \partial_s \varphi^S = 0.$$

Equation (7-4) asserts that $\lim_{z \rightarrow 0} \sup_{s \in \mathbb{R}} |\zeta^S| = 0$.

The Cauchy–Riemann equations in (7-8) and the fact that $|\zeta^S| \rightarrow 0$ as $z \rightarrow 0$ implies via a version of the Schwarz reflection trick from Theorem 24 in [1] that the pair (φ^S, ζ^S) extend smoothly to the $z = 0$ boundary of the domain. □

7B Ech-HF subvarieties and broken, singular admissible sets

An ech-HF subvariety that obeys the conditions of Proposition 7.1 can define in principle a finite collection of singular admissible sets. The members of such a collection obey certain constraints that are described Part 1 of the subsection. A collection that obeys these constraints is said in what follows to be a *broken, singular admissible set*. The upcoming Proposition 7.2 in Part 2 of the subsection explains how an ech-HF subvariety determines such a collection.

Part 1 Let $N \in \{1, 2, \dots\}$ and let $\Xi = \{((S_1, u_1), \vartheta_{\Sigma_1}), \dots, ((S_N, u_N), \vartheta_{\Sigma_N})\}$ denote an ordered collection of N singular admissible sets. This set Ξ is said to be a broken, singular admissible set when the constraints given in the upcoming (7-9) are met. These constraints refer solely to the case when $N > 1$. A broken admissible set with one element is neither more nor less than an admissible set. The notation in (7-9) uses \widehat{v}_{k+} and \widehat{v}_{k-} when $k \in \{1, \dots, N\}$ to denote the respective collections of G points in $C_+ \cap C_-$ that are defined by the positive and negative points of $\underline{S}_k \setminus S_k$. List (7-9) also refers to two currents in $\mathbb{R} \times [1, 2] \times \Sigma$ that are defined on any given compact set by a singular admissible set. The first of these currents is the current $[Z]$ from Proposition 7.1. The second is given by integration over the $s \leq 0$ part of the 2-chain that defines $[Z]$, it denoted by $[Z_{<0}]$. Note that both $[Z]$ and $[Z_{<0}]$ have finite pairing with the 2-form ω_Σ . These pairings are denoted by $[Z](w)$ and $[Z_{<0}](w)$.

Here are the promised constraints:

- (7-9) • No element $((S, u), \vartheta_\Sigma) \in \Xi$ is such that both $\vartheta_\Sigma = \phi$ and $u(S)$ is \mathbb{R} -invariant.
- Each $Z = ((S, u), \vartheta_\Sigma) \in \Xi$ is such that $[Z_{<0}](w) = \frac{1}{2}[Z](w)$.
- $\widehat{v}_{(k-1)+} = \widehat{v}_{k-}$ for each $k \in \{2, \dots, N\}$.

To give an idea as to how to view these constraints, fix $k \in \{2, \dots, N\}$. What is said in (7-9) implies that the constant, $s \gg 1$ slices of the u_{k-1} -image of S_{k-1} are very close to the constant, $s \ll -1$ slices of the u_k -image of S_k . This fact suggests that the union of a suitable set of $k \in \{1, \dots, N\}$ dependent translates in $\mathbb{R} \times [1, 2] \times \Sigma$ of the subvarieties that are given by Ξ may well be everywhere very close to a subvariety that is defined by a single admissible set of the form $\{(S, u), \vartheta_\Sigma\}$ where S is obtained from the disjoint union of the surfaces from the set $\{S_k\}_{1 \leq k \leq N}$ by identifying a suitable neighborhood of the positive points in each $k \in \{2, \dots, N\}$ version of \underline{S}_{k-1} with a suitable neighborhood of the negative points in \underline{S}_k .

Note that any singular admissible set $Z = ((S, u), \vartheta_\Sigma)$ such that $u(S)$ is not \mathbb{R} -invariant is a member of a 1-parameter family of such sets (parametrized by \mathbb{R}) that differ only in the sense that one u map is obtained from another by composing with a constant translation along the \mathbb{R} factor of $\mathbb{R} \times [1, 2] \times \Sigma$. The constraint in the second bullet of (7-9) selects a unique member of this family.

Part 2 The focus here is the upcoming Proposition 7.2. The statement of this proposition uses the following notation: Let $Z = \{(S, u), \vartheta_\Sigma\}$ denote a singular admissible set and let $s \in \mathbb{R}$ denote a given point. The J_{HF} -holomorphic map from S to $\mathbb{R} \times [1, 2] \times \Sigma$ that is obtained by composing first u and then translation by s along the \mathbb{R} factor is denoted by u^s . The translation of the elements in ϑ_Σ by s along the \mathbb{R} factor defines a new set of constant (s, t) slices of $\mathbb{R} \times [1, 2] \times \Sigma$ that is denoted in by ϑ_Σ^s . The new singular admissible set $((S, u^s), \vartheta_\Sigma^s)$ is denoted below by Z^s .

The proposition also refers to a set of intervals $\{I_1, \dots, I_N\} \subset \mathbb{R}$ that is defined using a constant ε and a corresponding set of real numbers $\{s_1, \dots, s_N\}$ with $s_{k-1} < s_k - \frac{3}{2\varepsilon}$. If $N = 1$, then $I_1 = \mathbb{R}$. If N is bigger than 1, then the interval I_1 is $(-\infty, s_2 - \frac{1}{2\varepsilon})$, the interval I_N is $(s_{N-1} + \frac{1}{2\varepsilon}, \infty)$, and each $k \in \{2, \dots, N - 1\}$ version of I_k is $(s_{k-1} + \frac{1}{2\varepsilon}, s_{k+1} - \frac{1}{2\varepsilon})$. Note that the union of these intervals is \mathbb{R} .

Proposition 7.2 *Fix an almost complex structure J_{HF} that obeys the conditions in (6-1). Given $\mathcal{I} \geq 1$ there exists a constant $\kappa_{\mathcal{I}} \geq 1$, and given also $\varepsilon > 0$, there exists $\kappa > \varepsilon^{-1}$ with the following significance: Define the geometry of Y using data δ, χ_0 and \mathbb{R} with $\delta < \kappa^{-2}$, and fix an almost complex structure, J , on $\mathbb{R} \times Y$ that obeys the constraints in Section 3A plus (7-1). Let C denote an ech-HF subvariety with either $I_{ech}(C) \leq \mathcal{I}$ or with $\int_C \omega \leq \mathcal{I}$. There exists a broken, singular admissible set $\Xi = \{Z_1, \dots, Z_N\}$ with at most $\kappa_{\mathcal{I}}$ elements and a corresponding set $\{s_1, s_2, \dots, s_N\} \in \mathbb{R}$ with the properties listed below:*

- $s_{k-1} < s_k - \frac{3}{2\varepsilon}$ for each $k \in \{2, \dots, N\}$.
- $\sup_{z \in (C \cap (I_k \times M_\varepsilon))} \text{dist}(z, |Z_k^{s_k}|) + \sup_{z \in (|Z_k^{s_k}| \cap (I_k \times M_\varepsilon))} \text{dist}(z, C) < \varepsilon$ for each $k \in \{1, \dots, N\}$.
- Let μ denote a smooth 2-form on $\mathbb{R} \times M_\varepsilon$ with norm bounded by 1, covariant derivative norm bounded by ε^{-1} and compact support on an interval of length $\frac{2}{\varepsilon}$ in some $k \in \{1, \dots, N\}$ version of I_k . Then $|\int_C \mu - [Z_k^{s_k}](\mu)| \leq \varepsilon$.

Note that this proposition in the case $N = 1$ says more than what is said by Proposition 7.1 by virtue of the fact that $I_1 = \mathbb{R}$.

Proof The proof is obtained by applying Proposition 7.1 to suitable translates of C along the \mathbb{R} factor of $\mathbb{R} \times M$. The bound on N is obtained using Proposition 5.1 to bound the sum of the various $Z \in \Xi$ versions of $[Z](w)$. Note in addition that the set of singular admissible sets with a given bound for this pairing is compact in the sense that any given sequence with this bound converges in the sense given by the bullets in Proposition 7.2 to a broken, singular admissible set. Proposition 5.1 and Lemma 5.7 are used to deal with the noncompact ends of I_1 and I_N . As there are no hidden subtleties in the arguments, the details are omitted. \square

7C Canonical admissible pairs for ech-HF submanifolds

This section supplies a refinement to the versions of Propositions 7.1 and 7.2 for the case when Ξ has but a single element which is a true admissible pair. This refinement is used in [12] to choose such a pair in a “canonical” fashion. The upcoming Proposition 7.3 states the refinement. What follows directly supplies some necessary background.

Let $\mathcal{A}_{\text{smooth}}$ denote the subspace in the moduli space of unobstructed admissible pairs. Let $S \rightarrow \mathcal{A}_{\text{smooth}}$ denote the fiber bundle whose fiber over a given element (S, u) is the surface S . Sitting over S is the universal normal bundle, this the complex line bundle whose fiber over $S|_{(S, u)}$ is the normal bundle $N \rightarrow S$. Let $\mathcal{N} \rightarrow \mathcal{A}_{\text{smooth}}$ denote this bundle.

To continue, suppose that $\mathcal{K} \subset \mathcal{A}_{\text{smooth}}$ has compact image in $\mathcal{A}_{\text{smooth}}/\mathbb{R}$. The techniques used for Lemma 5.4 of [20] work for parametrized families of surfaces and, in particular, they can be used to obtain the following data: a constant radius disk subbundle $\mathcal{N}_0 \subset \mathcal{N}|_{\mathcal{K}}$, a smooth map, \mathbb{R} -equivariant map $e^{\mathcal{K}}: \mathcal{N}_0 \rightarrow \mathbb{R} \times [1, 2] \times \Sigma$ and a constant $\rho_{\mathcal{K}} > 0$ with the following property: the map $e^{\mathcal{K}}$ restricts to a given $(S, u) \in \mathcal{K}$ fiber $S|_{(S, u)}$ as an exponential map of the sort described in Part 1. In

particular, the choice $\epsilon_S = \epsilon^{\mathcal{K}}|_{(S,u)}$ embeds each fiber as a J_{HF} -holomorphic disk such that the corresponding disk bundle $N_0 = \mathcal{N}_0|_{(S,u)}$ is embedded onto a neighborhood of S that contains all points in $\mathbb{R} \times [1, 2] \times \Sigma$ with distance less than $\rho_{\mathcal{K}}$ from S . The parts that follow use implicitly $N_0 = \mathcal{N}_0|_{(S,u)}$ and $\epsilon_S = \epsilon^{\mathcal{K}}|_{(S,u)}$ when referring to an exponential map for a given pair $(S, u) \in \mathcal{K}$. In particular, this version is used in Proposition 7.3 when defining the $(S, u) \in \mathcal{K}$ version of the operator D_S , and it is used to view submanifolds near $u(S)$ as submanifolds of the corresponding version of N_0 .

Proposition 7.3 *Fix an \mathbb{R} -invariant compact set $\mathcal{K} \subset \mathcal{A}_{\text{smooth}}$ and there exists $\kappa \geq 1$; and given $\mathcal{I} \geq 1$ and $\epsilon \in (0, \kappa^3)$, there exists $\kappa_* \geq \kappa^4$ with the following significance: Define the geometry of Y and $\mathbb{R} \times Y$ using $\delta < \kappa_*^{-1}$ and an almost complex structure that obeys the constraints in Section 3A and also (7-1). Let C denote an embedded, ech-HF subvariety with either $I_{\text{ech}}(C) < \mathcal{I}$ or with $\int_C w < \mathcal{I}$.*

- Suppose that the conclusions of Proposition 7.1 hold for C using $\Xi = ((S', u'), \phi)$ with $(S', u') \in \mathcal{K}$. There exists a unique $(S, u) \in \mathcal{A}_{\text{smooth}}$ such that the intersection of C with $[-\frac{1}{\epsilon}, \frac{1}{\epsilon}] \times M_\epsilon$ can be written as $\epsilon_S \circ \eta$ where η is a section of $N_0 \rightarrow S$ with the following properties:
 - (a) The conclusions of Proposition 7.1 hold using the given value of ϵ and (S, u) .
 - (b) The norm of η is bounded by $\kappa\epsilon$.
 - (c) The section η obeys (6-17) on the $(u^*s, u^*t) \in [-\frac{1}{\epsilon}, \frac{1}{\epsilon}] \times [1 + \epsilon^2, 2 - \epsilon^2]$ part of S , and the respective restrictions of η and the kernel of D_S are L^2 -orthogonal on this same part of S .

- Suppose that the conclusions of Proposition 7.2 hold for C using $\Xi = ((S', u'), \phi)$ with $(S', u') \in \mathcal{K}$. There exists a unique $(S, u) \in \mathcal{A}_{\text{smooth}}$ such that the intersection of C with $\mathbb{R} \times M_\epsilon$ can be written as $\epsilon_S \circ \eta$ where η is a section of $N_0 \rightarrow S$ with the following properties:
 - (a) The conclusions of Proposition 7.2 hold using the given value of ϵ and (S, u) .
 - (b) The norm of η is bounded by $\kappa\epsilon$.
 - (c) The section η obeys (6-17) on the $u^*t \in [1 + \epsilon^2, 2 - \epsilon^2]$ part of S and the respective restrictions of η and the kernel of D_S are L^2 -orthogonal on this same part of S .

The proof of this proposition is given in a moment. By way of a parenthetical remark, what is said in this proposition is not needed if (S', u') is in a 0- or 1-dimensional

component of $\mathcal{A}_{\text{smooth}}$ because the latter is a single orbit of the \mathbb{R} action. As it turns out, only those cases where the relevant part of $\mathcal{A}_{\text{smooth}}$ is a single \mathbb{R} -orbit are needed for the constructions in [12].

The three steps that follow directly provide background for the proof Proposition 7.3. These steps and the subsequent proof use c to denote a constant with value greater than 1 that depends only on \mathcal{K} . Its precise value increases between subsequent appearances.

Step 1 Fix an open, \mathbb{R} -invariant set in $\mathcal{A}_{\text{smooth}}$ with compact closure in $\mathcal{A}_{\text{smooth}}/\mathbb{R}$ whose interior contains \mathcal{K} . Let \mathcal{K}_* denote the closure of this \mathbb{R} -invariant set. Use in what follows the \mathcal{K}_* version of the subbundle \mathcal{N}_0 with the corresponding exponential map $e^{\mathcal{K}_*}$ and constant $\rho_{\mathcal{K}_*}$ to define any given $(S, u) \in \mathcal{K}_*$ versions of the bundle N_0 and the corresponding exponential map e_S . Since \mathcal{K}_* is chosen given \mathcal{K} , the constant $\rho_{\mathcal{K}_*}$ can be assumed to depend only on \mathcal{K} . For the same reason, the derivatives of e_S to any given order can be assumed to depend only on \mathcal{K} .

The use of a smoothly varying exponential map implies that the family of vector spaces $\{\text{kernel}(D_S) : (S, u) \in \mathcal{K}_*\}$ fit together as a smooth vector bundle of \mathcal{K}_* .

Step 2 The remarks in this step are for the most part direct consequences of Lemmas 6.6 and 6.7 using the fact that the various (S, u) versions of N_0 and e_S come from \mathcal{N}_0 and $e^{\mathcal{K}_*}$. In particular, the specific remarks that concern the $(S, u) \in \mathcal{K}_*$ versions of D_S and $\text{kernel}(D_S)$ follow from Lemmas 6.6 and 6.7 and what is said in Step 1 using standard techniques for studying small perturbations of Fredholm operators.

The first remark concerns the size of the ball B of Lemma 6.6: There exists $\rho_* > c^{-1}$ such that if $(S, u) \in \mathcal{K}$, then (S, u) has a neighborhood in \mathcal{K}_* that is parametrized by a ball $B \subset \text{kernel}(D_S)$ of radius ρ_* . Let q denote the map given by Lemma 6.6. Fix $\eta_0 \in B$ and let (S', u') denote the admissible pair that corresponds to η_0 . By way of a reminder, the 2-dimensional surface S' is the same as S , but its complex structure is now induced by u' . Meanwhile, $u' = e_S \circ q(\eta_0)$. The surface S' has its complex normal bundle, N' , and the identification between $u'(S')$ and the image of S via the map $e_S \circ q(\eta_0)$ supplies an \mathbb{R} -linear identification of the bundle N' with the normal bundle $N \rightarrow S$.

The bundle $N' \rightarrow S'$ has its disk subbundle N'_0 and the corresponding exponential map $e_{S'}: N'_0 \rightarrow \mathbb{R} \times [1, 2] \times \Sigma$. The aforementioned identification of N' with N identifies N'_0 with a subbundle in N . The latter is not N_0 in general. With N'_0 viewed in N , the exponential map $e_{S'}$ can be written on a somewhat smaller radius subbundle

$N'_1 \subset N'_0$ as a composition of first an embedding $\mathfrak{h}_{S'}: N'_1 \rightarrow N_0$ and then ϵ_S . Note in this regard that $\mathfrak{h}_{S'}$ is neither linear nor basepoint-preserving in general. In any event, the following can be assumed:

- (7-10) • Let ρ denote the constant fiber radius of the disk bundle \mathcal{N}_0 . The subbundle N'_1 has radius $\rho - c|\eta_0|$.
- The $\mathfrak{h}_{S'}$ -image of N'_1 contains the subbundle $N_1 \subset N_0$ of radius $\rho - c^2|\eta_0|$.
- The map from $B \times N_1$ to N_0 that restricts to any given $\eta_0 \in B$ slice of the domain as $(\mathfrak{h}_{S'} \circ \epsilon_{S'}^{-1}) \circ \epsilon_S$ is such that
 - (a) its differential along B at the origin is -1 times the identity,
 - (b) the norms of its derivatives to any given order are bounded by c .

Step 3 Suppose again that $(S, u) \in \mathcal{K}$ and that $(S', u') \in \mathcal{K}_*$ is parametrized as in Lemmas 6.6 and 6.7 by a given $\eta_0 \in B$. The identification between N' and N identifies the vector space $\text{kernel}(D_{S'})$ with a subspace, $H_{S'}$, in $C^\infty(S; N) \cap L^2(S; N)$.

Let $\Pi': L^2(S'; N') \rightarrow \text{kernel}(D_{S'})$ denote the $L^2(S'; N')$ orthogonal projection. With N' viewed as N and with $\text{kernel}(D_{S'})$ identified with $H_{S'}$, this projection will not in general agree with the $L^2(S; N)$ orthogonal projection. Even so, the following can be said: Let $\Pi: L^2(S; N) \rightarrow \text{kernel}(D_S)$ denote the $L^2(S; N)$ orthogonal projection. There is a linear isomorphism $H_{S'}: \text{kernel}(D_{S'}) \rightarrow \text{kernel}(D_S)$ such that

$$(7-11) \quad \|(H_{S'}\Pi' - \Pi)(\cdot)\|_{L^2(S; N)} \leq c\|\eta_0\|_{L^2(S; N)} \cdot \|\cdot\|_{L^2(S'; N')}.$$

Moreover, the linear map $(H_{S'}\Pi' - \Pi)$ varies smoothly as η_0 varies in B , and, in particular, its differential and that of its derivatives to any given order have norm bounded by c .

Proof of Proposition 7.3 The proof of the first bullet is identical but for the introduction of an extra \mathbb{R} -dependent cutoff function. This being the case, no more is said here about the first bullet. Note that the roles of (S', u') and (S, u) have been inverted in the proof. This is to avoid a proliferation of “primed” symbols. This understood, it is assumed in what follows that $(S, u) \in \mathcal{K}$ is such that the assumptions of Proposition 7.2 hold for C with $\Xi = ((S, u), \phi)$. Meanwhile, it is assumed that the conclusions of the second bullet of Proposition 7.3 hold for a unique pair $(S', u') \in \mathcal{A}_{\text{smooth}}$. The proof has two steps. The first finds the desired (S', u') and the second proves that there is only one such (S', u') .

Step 1 A purely \mathcal{K} -dependent upper bound for ε can be chosen so that C 's intersection with the $u^*t \in [1 + \varepsilon^2, 2 - \varepsilon^2]$ part of $\mathbb{R} \times [1, 2] \times M_\varepsilon$ is in the image via ϵ_S of the

bundle N_1 . This being the case, there is a smooth section, η , of N_1 over this same part of S with norm bounded by $c\varepsilon$ and such that this part of C can be written as $\varepsilon_S \circ \eta$. View η as a piecewise smooth section of N over the whole of S by declaring it to be zero on the part of S where u^*t is not in $[1 + \varepsilon^2, 2 - \varepsilon^2]$. This done, use $\eta_C \in \text{kernel}(D_S)$ to denote $\Pi\eta$.

Let $\eta_0 \in B$ and let (S', u') denote the corresponding admissible pair. The subvariety C can also be written $\varepsilon_{S'} \circ \eta'$ where η' is a section of $\mathfrak{h}_{S'}^{-1}(N_1) \subset N'_1$ over the $u^*t \in [1 + \varepsilon^2, 2 - \varepsilon^2]$ part of S' . Extend η' over the remainder of S' as zero. It follows from (7-10) and (7-11) that

$$(7-12) \quad H_{S'}\Pi'\eta' = \eta_C - \eta_0 + \tau(\eta),$$

where $\tau(\eta)$ has L^2 -norm bounded by $c\|\eta_0\|_{L^2(S;N)}\|\eta\|_{L^2(S';N')}$.

What with (7-10) and (7-11), the implicit function theorem finds $c \geq 1$ that makes the following true: If $\varepsilon < c^{-1}$, then there exists a $\eta_0 \in B$ such that the corresponding η' obeys $\Pi'\eta' = 0$. Moreover, η_0 differs from η_C by at most $c\varepsilon^2$.

Step 2 Suppose next that (S'_1, u'_1) and (S'_2, u'_2) satisfy the conclusions of the second bullet of the proposition. Given that \mathcal{K}_* is compact, there exists $c \geq 1$ such that if $\varepsilon < c^{-1}$, then the conclusions of Proposition 7.2 require that $u'_1(S'_1)$ and $u'_2(S'_2)$ be written respectively as $\varepsilon_S(q(\eta_{01}))$ and $\varepsilon_S(q(\eta_{02}))$ where η_{01} and η_{02} are points in B with norm bounded by $c\varepsilon$.

Let η'_1 and η'_2 denote the respective (S'_1, u'_1) and (S'_2, u'_2) versions of η' . The left-hand sides of the respective (S'_1, u'_1) and (S'_2, u'_2) versions of (7-12) are zero by definition, and, as a consequence, subtracting these two versions of (7-12) yields an equation of the form

$$(7-13) \quad \eta_{02} - \eta_{01} + \tau_{1,2} = 0,$$

where $|\tau_{1,2}| \leq c\varepsilon|\eta_{02} - \eta_{01}|$. An equation of this sort implies that $\eta_{01} = \eta_{02}$ if $\varepsilon \leq c^{-1}$. □

Appendix

The purpose of this appendix is to supply a proof of the assertion that a certain twisted version of embedded contact homology on Y can be defined by the rules laid out by Hutchings in [8]. This assertion is the content of Theorem 2.2 in [11]. The first part of

this appendix briefly describes the relevant chain complex, its differential and various other important endomorphisms. The upcoming Theorem A.1 restates the assertion that these endomorphisms are well defined. The remaining parts of this appendix explain why Theorem A.1 is true.

A1 The ech chain complex

This subsection describes the embedded contact homology chain complex of interest, the differential, and various other relevant endomorphisms. The discussion here has four parts.

Part 1 The relevant embedded contact homology chain complex is the free \mathbb{Z} -module that is generated by a certain principal \mathbb{Z} -bundle over the set $\mathcal{Z}_{\text{ech},M}$ of Proposition 2.8. This principal \mathbb{Z} -bundle will be denoted by $\widehat{\mathcal{Z}}_{\text{ech},M}$; it is defined in a moment.

To set the stage for the definition, let Θ denote an element in $\mathcal{Z}_{\text{ech},M}$. This element defines the 1-cycle $[\Theta]$ in Y by the formula $[\Theta] = \sum_{(\gamma,m) \in \Theta} m[\gamma]$ with $[\gamma]$ denoting the cycle that is defined by orienting the loop γ using the vector field v . Supposing that Θ and Θ' are two elements in $\mathcal{Z}_{\text{ech},M}$, let $H_2(Y; [\Theta] - [\Theta'])$ denote the equivalence classes of singular 2-chains with boundary equal to the difference cycle $[\Theta] - [\Theta']$. The equivalence relation posits that 2-chains A and B are equivalent when their difference is a boundary. Note that $H_2(Y; [\Theta] - [\Theta'])$ is an affine space (over \mathbb{Z}) modeled by $H_2(Y; \mathbb{Z})$.

The depiction of $\widehat{\mathcal{Z}}_{\text{ech},M}$ that follows writes it in a noncanonical way because it requires first choosing a fiducial element in $\mathcal{Z}_{\text{ech},M}$ to be denoted by Θ_0 . The choice of Θ_0 is the noncanonical part. Having chosen Θ_0 , define $\widehat{\mathcal{Z}}_{\text{ech},M}$ to be the set of equivalence classes of pairs of the form (Θ, Z) where $\Theta \in \mathcal{Z}_{\text{ech},M}$ and where Z is an element in $H_2(Y; [\Theta] - [\Theta_0])$. To state the equivalence relation, introduce, by way of notation, $\gamma^{(z_0)}$ to denote the closed integral curve of v in Y that contains the fiducial point $z_0 \in \Sigma$. Pairing with the Poincaré dual of $\gamma^{(z_0)}$ defines a homomorphism from the \mathbb{Z} -module of closed 2-cycles in Y to \mathbb{Z} . This pairing is denoted by $\langle \gamma^{(z_0)}, \cdot \rangle$. The equivalence relation posits

$$(A-1) \quad (\Theta, Z) \sim (\Theta', Z') \iff \Theta = \Theta' \quad \text{and} \quad \langle \gamma^{(z_0)}, Z - Z' \rangle = 0.$$

The principal bundle \mathbb{Z} action is defined by the rule whereby the element $1 \in \mathbb{Z}$ acts to send a given element (Θ, Z) to $(\Theta, Z + [S_0])$, where $[S_0]$ is the cycle given by the $u = 0$ sphere in \mathcal{H}_0 .

The depiction of $\widehat{\mathcal{Z}}_{\text{ech},M}$ using Θ_0 can be used to write $\widehat{\mathcal{Z}}_{\text{ech},M}$ as $\mathcal{Z}_{\text{ech},M} \times \mathbb{Z}$ using the principal \mathbb{Z} -bundle isomorphism from $\widehat{\mathcal{Z}}_{\text{ech},M}$ to $\mathcal{Z}_{\text{ech},M} \times \mathbb{Z}$ that sends the equivalence class of a given element (Θ, Z) in $\widehat{\mathcal{Z}}_{\text{ech},M}$ to the pair $(\Theta, \langle \gamma^{(z_0)}, Z \rangle)$ in $\mathcal{Z}_{\text{ech},M} \times \mathbb{Z}$. (The intersection number between $\gamma^{(z_0)}$ and Z is well defined because the boundary of Z is disjoint from $\gamma^{(z_0)}$.)

A different choice for Θ_0 produces a canonically isomorphic principal \mathbb{Z} -bundle over $\mathcal{Z}_{\text{ech},M}$. To elaborate, suppose that Θ'_0 is a second choice from $\mathcal{Z}_{\text{ech},M}$. Any cycle in $H_2(Y; [\Theta_0] - [\Theta'_0])$ has a well-defined intersection pairing with the curve $\gamma^{(z_0)}$. Adding a suitable multiple of the $f = \frac{3}{2}$ level set in M_δ will give a cycle in $H_2(Y; [\Theta_0] - [\Theta'_0])$ with zero intersection pairing against $\gamma^{(z_0)}$. Let Z_0 denote such a cycle. The isomorphism in question sends an the equivalence class of (Θ, Z) in the Θ_0 version of $\widehat{\mathcal{Z}}_{\text{ech},M}$ to that of $(\Theta, Z + Z_0)$ in the Θ'_0 version.

Since different Θ_0 versions of $\widehat{\mathcal{Z}}_{\text{ech},M}$ are canonically isomorphic, they define an abstract principal \mathbb{Z} -bundle over $\mathcal{Z}_{\text{ech},M}$, which is henceforth what $\widehat{\mathcal{Z}}_{\text{ech},M}$ denotes. The choice of Θ_0 specifies only a particular isomorphism with the product principal \mathbb{Z} -bundle.

The embedded contact homology chain complex of interest is the free \mathbb{Z} -module generated by $\widehat{\mathcal{Z}}_{\text{ech},M}$. This module is denoted in what follows by $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$. At the risk of pedantry, any given element in $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ is a finite, integer-weighted linear combination of generators.

The linear functional on $H_2(M; \mathbb{Z})$ given by the pairing with the class c_{1M} maps to a subgroup of $2\mathbb{Z}$. Let p_M denote the divisibility of this subgroup. With (A-1) understood, rules laid out by Hutchings [7] can be employed in the context at hand to give each generator of $\widehat{\mathcal{Z}}_{\text{ech},M}$ a relative $\mathbb{Z}/(p_M\mathbb{Z})$ degree and so give the module $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ a relative $\mathbb{Z}/(p_M\mathbb{Z})$ grading.

Part 2 The differential that defines the homology of this chain complex is given by a certain endomorphism of the module $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$. This endomorphism is defined by its action on each generators. The corresponding homology enjoys an action of the algebra $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ whose generators are also defined by endomorphisms of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$.

In general, a given endomorphism of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ is defined by its action on the generating set. This is to say that it is determined by a rule that assigns to any given element

$\hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}$ a formal sum of the form

$$(A-2) \quad \hat{\Theta} \rightarrow \sum_{\hat{\Theta}' \in \hat{\mathcal{Z}}_{\text{ech},M}} N_{\hat{\Theta}', \hat{\Theta}} \hat{\Theta}'$$

where any given coefficient is an integer, and where only finitely many are nonzero. The endomorphism is thus defined by the corresponding set of coefficients. The definition of those that give the differential is described below. The definition of those that generate the action of $\mathbb{Z}(\mathbb{U}) \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ is described in the upcoming Part 3.

According to the rules laid out by Hutchings [8], the coefficients that appear in (A-2) for the case of the differential are defined using certain sets of J -holomorphic subvarieties in $\mathbb{R} \times Y$. In particular, each ordered pair $\{\hat{\Theta}', \hat{\Theta}\} \subset \hat{\mathcal{Z}}_{\text{ech},M}$ labels such a set, this denoted in what follows by $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$. To give the criteria for membership, write $\hat{\Theta}$ as (Θ, Z) and $\hat{\Theta}'$ as (Θ', Z') . A given subvariety C is a member of this set when four conditions are satisfied, the first three being

$$(A-3) \quad \Theta_{C+} = \Theta, \quad \Theta_{C-} = \Theta' \quad \text{and} \quad I_{\text{ech}}(C) = 1.$$

To state the fourth, introduce $[C]_Y$ to denote the element in $H_2(Y; [\Theta] - [\Theta'])$ that is defined by the image of C in Y via the projection from $\mathbb{R} \times Y$. The fourth condition asserts that

$$(A-4) \quad Z = [C]_Y + Z'.$$

The group \mathbb{R} acts on $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ via constant translations of its elements along the \mathbb{R} factor of $\mathbb{R} \times Y$. With this in mind, the coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ in the version of (A-2) that gives the differential can be calculated if $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})/\mathbb{R}$ is a finite set, and if a certain Fredholm operator that is associated to each element in this set has trivial cokernel.

Part 3 The other relevant endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ generate an action of the algebra $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ on the resulting homology. The coefficients that appear in the relevant versions of (A-2) are also computed using certain sets of J -holomorphic subvarieties on $\mathbb{R} \times Y$.

The definition of an endomorphism that defines the action of \mathbb{U} on the homology of requires first the choice of a point in Y that does not lie on any closed integral curve of v from elements in $\mathcal{Z}_{\text{ech},M}$. Choose such a point in \mathcal{H}_0 or in M_δ where $f \notin [1, 2]$. Let y denote this point. According to Hutchings (see Section 11 of [9]), a given $\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}$ labeled coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ in the corresponding version of (A-2) is

computed from data supplied by the set, $\mathcal{M}_{2,y}(\hat{\Theta}', \hat{\Theta})$, of J -holomorphic subvarieties with membership defined as follows: A subvariety C is a member when (A-4) holds and when the $I_{\text{ech}}(C) = 2$ version of (A-3) holds. In addition, C must contain the point $(0, y) \in \mathbb{R} \times Y$. Let $\mathcal{M}_0(\hat{\Theta}', \hat{\Theta})$ denote the set of J -holomorphic subvarieties that obey (A-4) and the $I_{\text{ech}} = 0$ version of (A-3). It follows from Proposition 4.1 and what is said in Section 3 that the set of subvarieties that form $\mathcal{M}_{2,y}(\hat{\Theta}', \hat{\Theta})$ enjoy a 1-to-1 correspondence with the set that forms $\mathcal{M}_0(\hat{\Theta}', \hat{\Theta})$ because a subvariety in the former is the union of some subvariety in the latter with the sphere from the moduli space of Proposition 3.1 that contains the point $(0, y)$. As noted parenthetically after Proposition 3.1, each sphere in the latter's moduli space has $I_{\text{ech}} = 2$. This identification of the respective coefficient in question can be calculated if the following conditions hold: The space $\mathcal{M}_0(\hat{\Theta}', \hat{\Theta}) = \emptyset$ when $\hat{\Theta}' \neq \hat{\Theta}$ and $\mathcal{M}_0(\hat{\Theta}, \hat{\Theta})$ has but the one element $C = \bigcup_{\gamma \in \Theta} (\mathbb{R} \times \gamma)$. In addition, the cokernel of a certain Fredholm operator that is associated to each sphere in Proposition 3.1's moduli space must have trivial cokernel.

A set of endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ that generate the desired action of the algebra $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$ on the homology of the chain complex is defined using a chosen, suitably generic set of smooth 1-cycles in Y that represent a basis for $H_1(Y; \mathbb{Z})/\text{torsion}$. Each cycle from this set must be disjoint from the integral curves of v that come from elements in $\mathcal{Z}_{\text{ech},M}$. Let γ denote such a cycle. With γ chosen, any given $\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}$ version of the coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ for the corresponding endomorphism of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ is computed using the subset of subvarieties in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ that intersect $\{0\} \times \gamma \subset \mathbb{R} \times Y$ and their corresponding intersection points. See also what is said in Section 11 of [9].

Part 4 The theorem that follows asserts that the desired version of embedded contact homology can be defined in the manner outlined above. It also asserts that this version of contact homology admits the desired action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(M; \mathbb{Z})/\text{torsion}))$. The theorem uses \mathcal{J} to denote the C^∞ -Fréchet space of almost complex structures on $\mathbb{R} \times Y$ with the following two properties: each $J \in \mathcal{J}$ obeys the constraints given in Section 3A, and each obeys (7-1) for some almost complex structure J_{HF} that obeys (6-1).

Theorem A.1 *Fix the data δ , χ_0 and R to define the geometry of Y . There exists a residual set $\mathcal{J}_{\text{ech}} \subset \mathcal{J}$ with the following significance: Fix an almost complex structure $J \in \mathcal{J}_{\text{ech}}$ to define the almost complex geometry of $\mathbb{R} \times Y$. Then the*

endomorphisms that define the desired embedded contact homology differential on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ and the endomorphism that define the generators of the desired action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ on the resulting homology can be defined by Hutchings' rules. The differential decreases the relative $\mathbb{Z}/(p_M \mathbb{Z})$ of each generator by 1 and so the resulting homology has a relative $\mathbb{Z}/(p_M \mathbb{Z})$ grading. The action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ is such that \mathbb{U} decreases the relative grading degree by 2 and the generators of the action of $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$ decrease it by 1. The action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ on the homology is independent of the chosen point in $\mathcal{H}_0 \cup (M_\delta \setminus f^{-1}([1, 2]))$ that is used to define the \mathbb{U} -endomorphism, and it is independent of the cycles that are used to define the endomorphisms that give the action of $H_1(Y; \mathbb{Z})/\text{torsion}$.

The proof of this theorem is given in a moment. By way of a look ahead, the proof has four aspects. The first verifies that (A-5) and (A-6) hold when J comes from a certain residual subset of \mathcal{J} . The second verifies that the endomorphism that describes the differential has square zero. The third verifies that the endomorphisms of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ that are meant to define the action of \mathbb{U} and of $H_1(Y; \mathbb{Z})$ induce an action of the algebra $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$. The fourth verifies that this action is independent of the chosen point in Y and the chosen 1-cycles that are used to construct the defining endomorphisms of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$. The first aspect of the proof constitutes Section A2, and the remaining aspects constitute Section A3. What is said about the relative grading follows from the definitions of the endomorphisms.

A2 Proof of Theorem A.1: the endomorphisms

Fix an almost complex structure $J \in \mathcal{J}$. What is said in Part 2 of Section A1 about the definition of the differential on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ can be summarized as follows: The endomorphism that defines the differential can be computed according to the rules laid out by Hutchings when the following conditions are met:

- (A-5) • Fix $\widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech},M}$. All but a finite set of $\widehat{\Theta}' \in \widehat{\mathcal{Z}}_{\text{ech},M}$ versions of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ are empty.
 - Each $\widehat{\Theta}', \widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech},M}$ version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is a finite set of \mathbb{R} -orbits.
 - The Fredholm operator associated to each such \mathbb{R} -orbit has trivial cokernel.

What is said in Part 3 of Section A1 concerning the \mathbb{U} endomorphism can be summarized as follows: The desired endomorphism can be defined using the rules laid out by Hutchings if there exists $y \in Y$ such that the following conditions hold:

- (A-6) • If $\hat{\Theta}, \hat{\Theta}' \in \hat{\mathcal{Z}}_{\text{ech}, M}$ are distinct, then $\mathcal{M}_0(\hat{\Theta}', \hat{\Theta}) = \phi$.
- If $\hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}$ then $\mathcal{M}_0(\hat{\Theta}, \hat{\Theta})$ has but the one element $C = \bigcup_{\gamma \in \Theta} (\mathbb{R} \times \gamma)$.
 - The relevant Fredholm operator for any given sphere from the moduli space of Proposition 3.1 has trivial cokernel.

What is said in this same part of Section A1 also implies the following: The endomorphisms on $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ that generate the action of $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$ on the homology can be defined by Hutchings' rules from a chosen, suitably generic basis set of 1-cycles when (A-5) holds.

The three parts that follow explain why (A-5) and (A-6) hold for the almost complex structures in a certain residual subset of \mathcal{J} .

Part 1 This part and Part 2 address the cokernel issue. To this end, let C denote a given J -holomorphic subvariety in $\mathbb{R} \times Y$ with $\Theta_{C+} = \Theta$ and $\Theta_{C-} = \Theta'$. The Fredholm operator in question is the standard deformation operator for J -holomorphic subvarieties. If C is immersed, then it has the same form as the operator that is depicted in (6-11) on each component of C . The operator in this case defines an \mathbb{R} -linear, Fredholm operator with domain the L^2_1 Hilbert space of sections of C 's complex normal bundle and range the L^2 Hilbert space of sections of the tensor product of the latter bundle with the complex line bundle $T^{0,1}C$. The operator in the general case is described in Section 4 of [10]. In any event, the operator is denoted by \mathcal{D}_C in what follows.

The set of almost complex structures on $\mathbb{R} \times Y$ is given a C^∞ -Fréchet manifold structure by viewing it as a submanifold in $C^\infty(\mathbb{R} \times Y; \text{End}(T(\mathbb{R} \times Y)))$. The set \mathcal{J} likewise inherits the structure of a C^∞ -Fréchet manifold. A standard argument (see eg the proof of Theorem 1.8 in [6] or what is said in [10]) can be used to find a C^∞ -residual subset in the space of almost complex structures satisfying (3-4) and Constraints 1 and 2 in Section 3A whose members are characterized as follows: Take any two elements from $\mathcal{Z}_{\text{ech}, M}$, and any pseudoholomorphic subvariety whose positive and negative ends are associated to the chosen elements. Then the corresponding Fredholm operator has trivial cokernel.

A close look at the proof (see eg the appendix in [24]) gives a slightly weaker statement if the residual set is to sit in \mathcal{J} . In particular, arguments that mimic those in Section 3 of [13] can be used to find a C^∞ -residual set in \mathcal{J} whose members are characterized

as follows:

(A-7) Fix an almost complex structure from this set. Take any two elements from $\mathcal{Z}_{\text{ech},M}$, and any pseudoholomorphic subvariety whose positive and negative ends are associated to the chosen elements. Then the corresponding Fredholm operator has trivial cokernel if the subvariety has no irreducible components from the space \mathcal{M}_Σ of Proposition 3.2.

As explained next, the third bullet in (A-5) and all of (A-6) hold if J is from this residual set. Part 2 of this subsection describes what is involved in the proof that (A-7) characterizes a residual subset of \mathcal{J} .

Take $J \in \mathcal{J}$ so that (A-7) holds. Consider first (A-5). It follows from (1.1) in [7] that each subvariety in any given $\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}$ version of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is an embedded submanifold. This being the case, it lacks irreducible components from \mathcal{M}_Σ and so the conclusions of (A-7) apply. The third bullet in (A-5) restates these conclusions.

To see about (A-6), choose the point y in either \mathcal{H}_0 or in M_δ where f is in either $(0, 1)$ or $(2, 3)$. Hutchings' index inequality in (1.1) of [7] gives the first two bullets in (A-6). The assertion given by the third bullet follows using degree arguments along the lines of those used in Section 2.3 of [22]. These degree arguments do not require (A-7).

Part 2 This part of the subsection explains why \mathcal{J} has a C^∞ -residual subset whose members are characterized by (A-7). It proves useful in this regard to introduce a foliation of \mathcal{J} with any given leaf consisting of almost complex structures that differ only with regards to the choice of J_{HF} in (7-1). Thus, all of the almost complex structures in the given leaf define the same almost complex structure on each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$. The plan is to find a C^∞ -residual subset obeying (A-7) in each leaf of this foliation. Said roughly, (A-7) holds on each leaf of this foliation because the only subvariety components of any Θ from $\mathcal{Z}_{\text{ech},M}$ that cannot be modified by varying J_{HF} are those that sit entirely in some $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ or those that form a level set of f in the $f^{-1}(1, 2)$ part of M_δ . The former are described by the first and second bullets of Proposition 3.4, and the fourth bullet of Lemma 3.6 says that they are unobstructed. The latter are from the space \mathcal{M}_Σ of Proposition 3.2. What is said in Section 4 implies that all other subvariety components of Θ are paired with integer 1, they all cross each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ and they cross the $f^{-1}(1, 2)$ part of M_δ . They are affected by variations of J_{HF} because almost none of their tangent planes at their points in $f^{-1}(1, 2) \cap M_\delta$ are parallel to f -level sets.

To say more about how to find the desired C^∞ -residual sets, fix $J' \in \mathcal{J}$ for the moment, fix a pair of elements from $\mathcal{Z}_{\text{ech},M}$ and then fix a J' -holomorphic subvariety whose respective positive and negative ends are associated to the chosen elements. Let C' denote this subvariety and let $\mathcal{D}_{C'}$ denote the associated Fredholm operator as defined using the chosen almost complex structure J' . The complement in C' of a finite set of points will be embedded and so have a complex normal bundle. An endomorphism chosen from $T\mathcal{J}|_{J'}$ on this part of C' defines a homomorphism from $T^{0,1}C'$ to N with N denoting C' 's normal bundle viewed as a complex line bundle using J' (and $T^{0,1}C'$ is also defined using J'). The endomorphism from $T\mathcal{J}|_{J'}$ therefore defines a section of $T^{0,1}C' \otimes N$. Thus, if its support is on the smooth part of C' , it defines a section of the domain space of \mathcal{D}_C . With this point in mind, let $\mathcal{T} \subset T\mathcal{J}$ denote a given subbundle. The elements of $\mathcal{T}|_{J'}$ with support on the smooth part of C' define in the manner outlined above, a subspace of elements in the range space of $\mathcal{D}_{C'}$. This subspace is denoted in what follows by $\mathcal{R}_{\mathcal{T}}|_{C'}$. Let $\Pi_{C'}$ denote the tautological projection from the range Hilbert space of $\mathcal{D}_{C'}$ to the cokernel of $\mathcal{D}_{C'}$. A subbundle $\mathcal{T} \subset T\mathcal{J}$ is said to be *large* if the following is true:

(A-8) Suppose that $J' \in \mathcal{J}$ and that C' is an ech-HF subvariety as defined by J' . Then $\Pi_{C'}$ is surjective on the subspace $\mathcal{R}_{\mathcal{T}}|_{C'}$.

Suppose that $\mathcal{T} \subset T\mathcal{J}$ is large, and also tangent to a foliation of \mathcal{J} . Fix $J' \in \mathcal{J}$ and let \mathcal{F} denote the leaf of this foliation through J' . Let \mathcal{F}^* denote the set of pairs of the form (J, C) such that $J \in \mathcal{F}$ and C is a J' -holomorphic, ech-HF subvariety. For example, if C' is as described above, then $(J', C') \in \mathcal{F}^*$. It then follows from (A-8) via the implicit function theorem that some neighborhood of (J', C') in \mathcal{F}^* has the structure of a smooth (Hilbert) manifold. The argument for this is similar to arguments in Chapter 3 of [15]. The Smale–Sard theorem can then be invoked (much like what is done in Chapter 3 of [15]) when \mathcal{T} is large to prove that there is a C^∞ residual set of J in \mathcal{F} near J' that are regular values for the projection map from \mathcal{F}^* to \mathcal{F} . Since the condition that J be a regular value for this projection map means, in particular, that the cokernel of \mathcal{D}_C is trivial at all pairs $(J, C) \in \mathcal{F}^*$, the following is a consequence of the ubiquity of regular values: Fix $J' \in \mathcal{F}$, an integer $k \geq 0$ and $\varepsilon > 0$. Let $\mathcal{F}_{J'}$ denote the leaf of the foliation through J' . There exist endomorphisms $j_0 \in \mathcal{T}|_{J'}$ and j_1 of $T(\mathbb{R} \times Y)$ such that:

- (A-9) • $J = J' + j_0 + j_1 \in \mathcal{F}$ is such that (A-7) holds.
- The C^k -norm of j_0 is less than ε and that of j_1 is at most ε times the C^k -norm of j_0 .

This says that any given $J' \in \mathcal{J}$ can be modified by changing its corresponding J_{HF} subject to (6-1) so that the result obeys (A-7) if \mathcal{T} is large.

The fact that \mathcal{T} is large in this case is proved using Aronszajn’s unique continuation principle [2]; the argument for this differs only notationally from an argument in Chapter 3.4 of [15] to prove Proposition 3.4.1 in [15].

Part 3 This part of the subsection explains why the first two bullets of (A-5) hold when J is characterized by (A-7). The first bullet follows directly from (5-5) and Lemma 5.2. In particular, Lemma 5.2 supplies an a priori upper bound for the integer n_C which is the intersection number between C and the J -holomorphic cylinder $\mathbb{R} \times \gamma^{(z_0)}$, and given n_C , it then follows from (5-5) that there at most finitely many $\hat{\Theta}' \in \hat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}$ such that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta}) \neq \emptyset$.

To see about the second bullet, it proves useful to reintroduce the foliation \mathcal{F} from Part 2 and the corresponding space \mathcal{F}^* . Suppose that $J \in \mathcal{F}$. Viewing $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ as a union of components of $\mathcal{F}^*|_J$ endows $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ with a topology. This topology is to be understood in what follows. Now, if J is also a regular value of the projection map from \mathcal{F}^* to \mathcal{F} , then it follows from the implicit function theorem that $\mathcal{F}^*|_J$ has the structure of a smooth manifold with the tangent space to $\mathcal{F}^*|_C$ at any given $(J, C) \in \mathcal{F}^*$ being canonically identified with the kernel of \mathcal{D}_C . In particular, this implies that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is a discrete, disjoint union of 1-dimensional manifolds because the Fredholm index of \mathcal{D}_C is 1 when $C \in \mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ (see eg Theorem 1.7 in [7] or Proposition 1.2 in [8] or Proposition 11.2 in [9]) and the cokernel of \mathcal{D}_C is trivial when J is a regular value of the projection map from \mathcal{F}^* to \mathcal{F} . Since \mathbb{R} acts on this space without fixed points, each manifold component must be an orbit of the \mathbb{R} action. Granted the preceding, the second bullet in (A-5) follows if $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})/\mathbb{R}$ is a finite set. But for one additional comment, the proof that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})/\mathbb{R}$ is a finite set uses the bounds in Proposition 5.1 as input for arguments that differ only in notation from those that prove assertion (a)(i) of Theorem 1.8 in [6]. (See also Lemma 7.19 in [10].)

The extra comment is needed to address the fact that the situation here corresponds to the case of Theorem 1.8 in [6] when what Hutchings denotes by d is the same as the genus of Σ . The analog here of Hutchings’ surface Σ is the Heegaard surface, and the analog here of d is the number of segments that form the M_δ part of the union of the integral curves of v from any given $\Theta \in \hat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}$. The versions here of d and genus Σ are equal, but assertion (a)(i) of Theorem 1.8 of [6] as stated requires $d > \text{genus}(\Sigma)$. Even so, the $I_{\text{ech}} = 1$ and $d = \text{genus}(\Sigma)$ version of assertion (a)(i)

of this Theorem 1.8 is true (Hutchings, personal communication, 2010). In fact, the argument for this case is identical to that given in [6] for the $d > \text{genus}(\Sigma)$ cases. To say slightly more about why this is, note that the proof of assertion (a)(i) invokes an inequality for the embedded contact homology index; this is the inequality written in Case 2 of the proof of Lemma 9.5 in [6]. The proof that assertion (a)(i) of Theorem 1.8 in [6] holds requires that this inequality have negative right-hand side, and it is this requirement that determines the relative sizes of d and the genus of Σ in the statement of assertion (a)(i). Meanwhile, the right-hand side of this inequality is observedly negative when $I_{\text{ech}} = 1$ and $d = \text{genus}(\Sigma)$.

By way of comparison, there is a Heegaard Floer analog of the $I_{\text{ech}} = 1$ and $d = \text{genus}(\Sigma)$ version of assertion (i)(a) of Theorem 1.8 in [6], this being Lemma 8.2 in [13]. What is more, the proof of the latter invokes the Heegaard Floer analog of the inequality that appears in Case 2 of the proof of Lemma 9.4 of [6].

A3 The square of the differential and the action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$

Fix an almost complex structure for which (A-7) holds so as to use Hutchings' rules to define the endomorphism of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, M})$ that is meant to be the desired differential. Part of what follows explains why this endomorphism has square zero. Fix a point, y , in either \mathcal{H}_0 or in the part of M_δ where $f \notin [1, 2]$ to define the \mathbb{U} endomorphism. Likewise, fix a basis of smooth cycles for $H_1(Y; \mathbb{Z})/\text{torsion}$ to define the endomorphisms that are meant to generate the $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$ action on the embedded contact homology. What follows also explains why these endomorphisms generate the desired action of $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$, and why this action is independent of the chosen point y and of the chosen basis of 1-cycles.

These assigned tasks are dealt with in a moment. This paragraph and the next explain why the standard arguments (in particular, those in Section 7.3 of [10] that prove $\delta^2 = 0$ for embedded contact homology) cannot be copied verbatim without some initial constructions to set the appropriate stage. The explanation starts by pointing out that the $\delta^2 = 0$ proof in Section 7.3 of [10] (and similar arguments for other versions of Floer homology, for example those in [15]) use the moduli spaces with $I_{\text{ech}} = 2$ to prove the vanishing of δ^2 . In particular, these moduli spaces can be used to prove that $\delta^2 = 0$ when the complement of a compact set in any such moduli space is a union of components whose J -holomorphic subvarieties can be constructed by the "end-to-end" gluing of an ordered pair of $I_{\text{ech}} = 1$ pseudoholomorphic subvarieties. These pairs

are constrained so that the $s \rightarrow \infty$ limit in $\mathbb{R} \times Y$ of the first element in the pair determines the same element in $\mathcal{Z}_{\text{ech},M}$ as the $s \rightarrow -\infty$ limit in $\mathbb{R} \times Y$ of the second element in the pair. See for example Lemma 7.23 in [10]. (An ordered pair of $I_{\text{ech}} = 1$ subvarieties that obeys the preceding constraint is an example of a *broken trajectory*.) If the required $I_{\text{ech}} = 2$ moduli spaces have additional ends, then they can't be used to prove $\delta^2 = 0$ without arguments that go beyond those in Section 7.3 of [10]. In the case at hand, when J comes from the residual set described by (A-7), the $I_{\text{ech}} = 2$ moduli spaces may well have additional ends. In particular, the proof of Lemma 7.23 in [10] cannot be used directly here because J -holomorphic subvarieties from the space \mathcal{M}_Σ of Proposition 3.2 have negative Fredholm (and ech) index. (See Lemma 3.6.) As a consequence, a sequence in an $I_{\text{ech}} = 2$ moduli space might converge pointwise and as a current to a reducible J -holomorphic subvariety with one or more irreducible components coming from \mathcal{M}_Σ .

By way of a look ahead, this problem with \mathcal{M}_Σ is circumvented in what follows by computing δ^2 using a perturbation of J on the $\mathbb{R} \times M_\delta$ part of Y that slightly violates the constraint in the third bullet of (6-1). There is no analog of \mathcal{M}_Σ for a suitably generic perturbation of this sort because a subvariety with a negative index version of the operator \mathcal{D}_C cannot appear when J is a regular value of the relevant analog of the projection map from the space \mathcal{F}^* to the space \mathcal{F} . (These spaces are defined in Part 2 of Section A2.)

The assigned tasks from the first paragraph of this subsection begin with the introduction of a nested exhaustion of $\mathcal{Z}_{\text{ech},M}$ by subsets $\{\mathcal{Z}_{\text{ech},M}^L\}_{L \geq 1}$ that are defined as follows: Define the length of any given $\Theta \in \mathcal{Z}_{\text{ech},M}$ to be the sum of the integrals of \hat{a} over the closed integral curves of v from Θ . The subset $\mathcal{Z}_{\text{ech},M}^L$ contains the elements with length less than L . Any given $\mathcal{Z}_{\text{ech},M}^L$ is a finite set and the union of all of them is $\mathcal{Z}_{\text{ech},M}$. Moreover, if $\Theta \in \mathcal{Z}_{\text{ech},M}^L$ and if $C \subset \mathbb{R} \times Y$ is a J -holomorphic subvariety with $\Theta_{C_+} \in \mathcal{Z}_{\text{ech},M}^L$ and with $\Theta_{C_-} \in \mathcal{Z}_{\text{ech},M}$, then Θ_{C_-} is also in $\mathcal{Z}_{\text{ech},M}^L$.

For any given $L \geq 1$, use $\hat{\mathcal{Z}}_{\text{ech},M}^L$ to denote the restriction of $\hat{\mathcal{Z}}_{\text{ech},M}$ to $\mathcal{Z}_{\text{ech},M}$. The set $\hat{\mathcal{Z}}_{\text{ech},M}^L$ generates the corresponding submodule $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M}^L) \subset \mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$. It follows from what is at the end of the preceding paragraph that each such submodule is mapped to itself by the differential, the \mathbb{U} -endomorphism, and by the endomorphisms that are meant to generate the action of $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$. This being the case, it is sufficient to prove that $\delta^2 = 0$ on each $L \geq 1$ version of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M}^L)$. Likewise, it is sufficient to prove the assertions in Theorem A.1 about the other endomorphisms for their restrictions to each $L \geq 1$ version of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M}^L)$.

With the preceding understood, fix $L \geq 1$. What follows explains why the differential has square zero on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L)$. The arguments for the assertions that concern the other endomorphisms amount to straightforward variations of that given below and so the latter are omitted.

To start, by (A-5), the set of \mathbb{R} -orbits that form $\bigcup_{\widehat{\Theta}', \widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L} \mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is finite. What follows is a consequence of this fact. Fix a C^∞ neighborhood of J in the space of almost complex structures that obey (3-4) and also Constraints 1 and 2 in Section 3A. If J' is in such a neighborhood, and if $J - J'$ has sufficiently small C^6 -norm, then there is a J' analog of any given $\widehat{\Theta}', \widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L$ version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$, and the elements in the latter enjoy a 1-to-1 correspondence with those in J 's version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$. The correspondence is such that if C is a submanifold from the J version, and C' is from the J' version, then C' sits in a small radius tubular neighborhood of C and is isotopic to C in this neighborhood by an isotopy whose C^k -norm for any given $k \in \{1, 2, \dots\}$ is bounded by a k -dependent multiple of the C^{k+3} -norm of $J' - J$. This partner C' can be constructed using perturbation-theoretic methods because the cokernel of \mathcal{D}_C is trivial. Standard elliptic estimates (much like those in Section 3.3 of [15]) can be used to get the desired a priori bounds on the norms of the isotopy.)

These perturbation methods construct an injective map from the J version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ to the J' version. Meanwhile, a limiting argument that differs only in notation from what is said in the proof of assertion (a)(i) from Theorem 1.8 in [6] proves this injective map is also surjective onto the J' version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$.

Let C' denote a submanifold from the J' version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$. The fact that this submanifold is C^6 -close to its partner in the J version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ implies that the operator $\mathcal{D}_{C'}$ also has trivial cokernel. More to the point, it implies that the coefficient $N_{\widehat{\Theta}', \widehat{\Theta}}$ that appears in (A-2) in the case of the differential can be computed using the submanifolds from the J' version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$. As a consequence, the endomorphism that defines the action of the differential on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L)$ can be computed using J' in lieu of J . This is important by virtue of what is said in the second paragraph of Part 1 of Section A2: The almost complex structure J' can be chosen so that every J' -holomorphic subvariety whose respective positive and negative ends are associated to elements in $\mathcal{Z}_{\text{ech}, \mathcal{M}}$ is such that its corresponding Fredholm operator has trivial cokernel. This being the case, the arguments in [10] can be repeated with only notational changes to see that the J' version of the differential has square zero on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L)$. As noted a moment ago, latter endomorphism is identical to the J version, and so J 's differential also has square zero on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}^L)$.

This use of a modified almost complex structure to show that $\delta^2 = 0$ is analogous to what is done in Section 8.2 of [13]. Meanwhile, [13] pays homage to a construction in [17].

List of symbols

Section 1A

f	This is the Morse function on M . See (1-1), (1-2) and (1-4).
G	The genus of the Heegaard surface. It is the number of index 1 (or index 2) critical points of f .
Λ	The set of pairs of index 1 and 2 critical points.
$(r_+, \theta_+, \varphi_+)$	Coordinates near index 1 and 2 critical points.
$(r_-, \theta_-, \varphi_-)$	
δ_*	A positive number, being the radius of an outer ball about the critical points.
\mathfrak{p}	This denotes a pair from Λ .
$\mathcal{H}_{\mathfrak{p}}$	The 1–handle at the critical points labeled by the pair \mathfrak{p} .
\mathcal{H}_0	The 1–handle attached at the index 0 and 3 critical points.
R	A large positive number, specifying the length of the handles.
(u, θ, φ)	Coordinates on the handles. The function u increases along the handle and (θ, φ) are spherical coordinates on the constant u slices.
M_r	For $r > 0$, this is the complement in M of the radius r balls about the critical points of f .

Section 1B

δ	A positive number less than δ_* , the radius of an inner ball about the critical points.
χ_0	A positive number less than δ^3 . Used to specify the Hamiltonian structure.
χ	A nonincreasing function on \mathbb{R} equaling 1 on $(-\infty, 0]$ and 0 on $[1, \infty)$.
χ	A nonnegative function of the coordinate u on the handles. See (1-5).

χ_+ and χ_-	Nonnegative functions of the coordinate u on the handles. See (1-6).
f_*	A function on the index 1-2 handles that extends f .
M_δ	The complement in M of the radius δ balls about the critical points of f .

Section 1C

\mathfrak{v}	The pseudogradient vector field for f on M .
Σ	The Heegaard surface, $f^{-1}(\frac{3}{2})$.
C_{p_+}	The intersection of Σ with the ascending disk from the index 1 critical point p .
$C_{p'_-}$	The intersection of Σ with the descending disk from the index 2 critical point p' .
z_0	A basepoint in Σ . See Part 4 of Section 1C.
T_{p_+}	A small diameter, annular neighborhood in Σ of C_{p_+} .
$T_{p'_-}$	A small diameter, annular neighborhood in Σ of $C_{p'_-}$.
\hat{h}_+ and \hat{h}_-	Functions defined in Part 3 of Section 1C.
T_+	The union of the T_{p_+} annuli.
T_-	The union of the $T_{p'_-}$ annuli.
c_{1M}	A class in $H^2(M; \mathbb{Z})/\text{torsion}$ used to define a strongly admissible Heegaard diagram.
\mathcal{D}_0 and $\{\mathcal{D}_i\}$	Fundamental domains for the Heegaard diagram.
\mathcal{P}	A periodic domain for the Heegaard diagram.
w_Σ	An area form for the Heegaard surface. See Lemma 1.1.
b_1	The first Betti number of M .
a_Σ	An antiderivative of w_Σ specified in Lemma 1.2.
\mathbb{Y}	A certain set of $b_1 + 1$ points in Σ . See Lemma 1.2.

Section 1D

\mathbb{Y}_0	A set that corresponds to \mathbb{Y} near the index 0 critical point.
----------------	---

Section 1E

a	The 1-form that defines in part the stable Hamiltonian structure on Y .
-----	---

w	The 2–form that defines in part the stable Hamiltonian structure on Y .
f and g	Functions of the coordinate u on the 1–handles. See (1-24).
v	The kernel of the 2–form w .
v_\diamond	A closed 1–form obeying $v_\diamond \wedge w \geq 0$ with equality only on certain closed orbits of v on the 1–handles connecting index 1 and 2 critical point pairs.

Section 1F

\mathcal{Z}_{HF}	A set of generators for the Heegaard Floer chain complex.
\hat{v}	A collection of integral curves of v that define an element in \mathcal{Z}_{HF} .
Spin ^ℂ structure constraints	See (1-28).
Θ	A generator of the embedded contact homology chain complex on Y .
$[\gamma]$	For a closed, oriented loop γ , its homology class (usually in $H_1(Y; \mathbb{Z})$).
$[\gamma]^{\text{Pd}}$	The Poincaré dual of $[\gamma]$.
K^{-1}	The kernel of the 1–form a : A 2–plane bundle on Y .
$e_{K^{-1}}$	The Euler class of K^{-1} . See (1-29) and (1-30).
\mathcal{Z}_{ech}	A collection of pairs consisting of a closed integral curve of v and an integer that obey (1-29).
$\mathcal{Z}_{\text{ech}, M}$	The subset of \mathcal{Z}_{ech} that is used to generate the embedded contact homology chain complex. (Hyperbolic integral curves are only paired with integer 1.) See Proposition 2.8.
s	The Euclidean coordinate on \mathbb{R} in the context of $\mathbb{R} \times Y$ or $\mathbb{R} \times M$.
J –holomorphic subvariety	See (1-31).
ω	A 2–form that tames J . See (1-32).

Section 2B

$\frac{d\theta}{du}$	See (2-3).
$\frac{d\phi}{du}$	See (2-5).

index 1-2 cycle Certain collections of integral curves of \mathfrak{v} in M connecting index 1 and index 2 critical points of f .

orientation sign See Part 5 of Section 2B.

Section 2C

$\widehat{\gamma}_p^+$ and $\widehat{\gamma}_p^-$ The $u = 0$ and $\cos^2 \theta = \frac{1}{3}$ integral curves of v in the handle \mathcal{H}_p indexed by the pair \mathfrak{p} of index 1 and index 2 critical points of f . See also Lemma 2.1.

Section 3A

J An almost complex structure. See Section 3A.

$\{e_1, e_2\}$ An orthonormal basis for the kernel of a on the Λ indexed 1–handles.

moduli space of J –holomorphic subvarieties See (3-7).

\mathbb{R} –equivariant See Part 3 of Section 3A.

Section 3B

\mathcal{M}_0 A moduli space of J –holomorphic subvarieties. See Proposition 3.1.

\mathcal{M}_Σ A moduli space of J –holomorphic subvarieties. See Proposition 3.2.

\mathcal{M}_1 and \mathcal{M}_2 Moduli spaces of J –holomorphic subvarieties. See Proposition 3.3.

Section 3E

$\mathcal{M}_{p^+}, \mathcal{M}_{p^-}, \mathcal{M}_{p_0}$ Moduli spaces of J –holomorphic subvarieties in the $\mathfrak{p} \in \Lambda$ version of \mathcal{H}_p . See Proposition 3.4.

ech index A number assigned by Hutchings to certain J –holomorphic subvarieties. It is denoted by I_{ech} .

Fredholm index The index of the operator whose kernel is the formal first-order deformations of a J –holomorphic subvariety.

obstructed A J –holomorphic subvariety is obstructed when the operator that defines the Fredholm index has nonzero cokernel in a Banach space.

Section 4A

ech-HF subvariety The subvarieties that are used to define the differential in embedded contact homology. They are described in Section 4.

Section 4C

h A certain function on the handles indexed by Λ . See (4-1).
 Ψ_p A certain coordinate like chart for part of the handle \mathcal{H}_p labeled by $p \in \Lambda$. See (4-3) and (4-4).
 \mathcal{X} The domain of Ψ_p .

Section 5A

Y_\diamond A part of Y . See (5-1).
 $[\Theta]$ A class in $H_1(Y; \mathbb{Z})$ defined by the element Θ from $\mathcal{Z}_{\text{ech}, M}$.
 $[\Theta_{C_-}]$ and $[\Theta_{C_+}]$ The classes in $H_1(Y; \mathbb{Z})$ that are defined from the respective $s \rightarrow \pm\infty$ limits of the constant s slices of the ech-HF subvariety C .
 \hat{O} A certain homomorphism in a Mayer–Vietoris sequence. See (5-2) and (5-3).
 Z_{p_1} and Z_{p_2} Classes in $H_1(\Sigma)$ determined by a given $p \in \Lambda$.
 (n_{p_1}, n_{p_2}) Integers specifying certain homology classes. See (5-4).
 m_p^C An integer defined by the ech-HF subvariety C and the given $p \in \Lambda$. See (5-5).
 $n_C^{(z)}$ An integer defined by a point $z \in \mathbb{Y}$ and the ech-HF subvariety C . This is an intersection number between C and a certain curve through z .
 n_C The z_0 version of $n_C^{(z)}$.
 $H_2(M; [\hat{v}_+] - [\hat{v}_-])$ A set of equivalence class of certain 2–chains in M .
 $[C]_M$ An element in $H_2(M; [\hat{v}_+] - [\hat{v}_-])$ that is defined by the ech-HF subvariety C .

Section 6A

$M_{(1,2)}$ The part of M where $f(\cdot) \in (1, 2)$.
 w_f An extension of w to $M_{(1,2)}$ from $M_{(1,2)} \cap M_\delta$.

J_{HF}	The almost complex structure on $\mathbb{R} \times M_{(1,2)}$ used to define Lipshitz's version of Heegaard Floer homology. See (6-1).
t	The coordinate on the interval $(1, 2)$. It is used when identifying $M_{(1,2)}$ as $(1, 2) \times \Sigma$ via the diffeomorphism given by flowing from Σ along the integral curves of \mathfrak{v} .
S	A Riemann surface with punctured boundary that is mapped into $\mathbb{R} \times [1, 2] \times \Sigma$ to define the differential in Lipshitz's version of Heegaard Floer homology.
\underline{S}	The Riemann surface S is the complement of $2G$ labeled points in the compact surface with boundary \underline{S} .
u	A map from S into $\mathbb{R} \times [1, 2] \times \Sigma$. See (6-2).
admissible pair	A pair (S, u) that is described in Part 3 of Section 6A and, in particular, obeys conditions in (6-2).
ϑ_{Σ}	A finite set of pairs whose first component is a level set of f in $M_{(1,2)}$ and whose second component is a positive integer. Distinct pairs from Σ have distinct level sets.
singular admissible set	A data set $((S, u), \vartheta_{\Sigma})$.

Section 6B

$\partial_p S$	Boundary component of S corresponding to the critical point p of f .
----------------	--

Section 6D

S_0	This is the interior of S , thus $S \setminus \partial S$.
-------	---

Section 6E

N_S	The normal bundle to the image in $\mathbb{R} \times [1, 2] \times \Sigma$ via the holomorphic map u of S .
N_0	A small radius disk bundle in N_S .
ϵ_S	A certain exponential map from N_0 to $\mathbb{R} \times [1, 2] \times \Sigma$.
D_S	A certain perturbation of the $\bar{\partial}$ operator on S by an \mathbb{R} -linear term defined by a holomorphic map u from S to $\mathbb{R} \times [0, 1] \times \Sigma$. See (6-11). Its kernel is the space of formal, first-order deformations of the holomorphic map u .

$C_0^\infty(S; N_S)$	A Banach space of sections of the normal bundle to the image via u of S . See (6-12).
unobstructed	The admissible pair (S, u) is unobstructed when D_S has trivial cokernel in a certain Banach space.
\mathcal{A}_{HF}	The space of admissible pairs. See (6-19).
$\mathcal{A}_{\text{smooth}}$	The subspace of admissible pairs that are unobstructed.

Section 7A

$[Z]$	When $Z = ((S, u), \vartheta_\Sigma)$, this is a 2-dimensional current defined by Z .
$ Z $	The support of the current Z , $u(S) \cup (\bigcup_{(Z,m) \in \vartheta} Z)$.

Section 7B

\mathbb{E}	A broken singular admissible set as defined in (7-9).
--------------	---

Section A1

$\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}$	A certain principal \mathbb{Z} -bundle over $\mathcal{Z}_{\text{ech}, \mathcal{M}}$.
$H_2(Y; [\Theta] - [\Theta'])$	When $\Theta, \Theta' \in \mathcal{Z}_{\text{ech}, \mathcal{M}}$, the equivalence classes of singular 2-chains with boundary $[\Theta] - [\Theta']$.
$\gamma^{(z_0)}$	The closed integral curve of ν in Y containing z_0 .
$\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}})$	The embedded contact homology chain complex.
p_M	Divisibility of the image in \mathbb{Z} of $c_{1M}: H_2(M; \mathbb{Z}) \rightarrow \mathbb{Z}$.
$\widehat{\Theta}$	An element in $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}})$.
$\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$	A certain space of J -holomorphic subvarieties in $\mathbb{R} \times Y$ that is used to compute the differential in embedded contact homology.
\mathbb{U}	A certain endomorphism of embedded contact homology.
$\mathcal{M}_{2,y}(\widehat{\Theta}', \widehat{\Theta})$	A certain space of J -holomorphic subvarieties in $\mathbb{R} \times Y$ that is used to compute the endomorphism \mathbb{U} .
$\mathcal{M}_0(\widehat{\Theta}', \widehat{\Theta})$	A certain space of J -holomorphic subvarieties in $\mathbb{R} \times Y$ that is bijective with $\mathcal{M}_{2,y}(\widehat{\Theta}', \widehat{\Theta})$.

Section A2

\mathcal{D}_C	A Fredholm operator associated to the ech-HF subvariety C . It is depicted in (6-11) for the case when C is immersed.
-----------------	---

\mathcal{J}	A certain space of almost complex structures that obeys (3-1) and (7-1).
\mathcal{F}	A certain foliation of \mathcal{J} . (The all of the complex structures in any given leaf of \mathcal{F} have the same almost complex structure on the 1–handle parts of Y .)

Section A3

broken trajectory	A finite set of J –holomorphic subvarieties whose $s \rightarrow \pm\infty$ limits are suitably constrained. See Section A3 for a precise definition.
$\mathcal{Z}_{\text{ech},M}^L$	The subset of $\mathcal{Z}_{\text{ech},M}$ with length (suitably defined) less than L .
$\widehat{\mathcal{Z}}_{\text{ech},M}^L$	The restriction of $\widehat{\mathcal{Z}}_{\text{ech},M}$ to $\mathcal{Z}_{\text{ech},M}^L$.
$\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M}^L)$	The free \mathbb{Z} –module generated by $\widehat{\mathcal{Z}}_{\text{ech},M}^L$.

References

- [1] **L V Ahlfors**, *Complex analysis: an introduction to the theory of analytic functions of one complex variable*, 3rd edition, McGraw-Hill, New York (1978) MR
- [2] **N Aronszajn**, *A unique continuation theorem for solutions of elliptic partial differential equations or inequalities of second order*, J. Math. Pures Appl. 36 (1957) 235–249 MR
- [3] **M Gromov**, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. 82 (1985) 307–347 MR
- [4] **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 MR
- [5] **C Hummel**, *Gromov’s compactness theorem for pseudo-holomorphic curves*, Progress in Mathematics 151, Birkhäuser, Basel (1997) MR
- [6] **M Hutchings**, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. 4 (2002) 313–361 MR
- [7] **M Hutchings**, *The embedded contact homology index revisited*, from “New perspectives and challenges in symplectic field theory” (M Abreu, F Lalonde, L Polterovich, editors), CRM Proc. Lecture Notes 49, Amer. Math. Soc., Providence, RI (2009) 263–297 MR
- [8] **M Hutchings**, *Embedded contact homology and its applications*, from “Proceedings of the International Congress of Mathematicians” (R Bhatia, A Pal, G Rangarajan, V Srinivas, M Vanninathan, P Ganesi, editors), volume II, Hindustan Book Agency, New Delhi (2010) 1022–1041 MR

- [9] **M Hutchings, M Sullivan**, *Rounding corners of polygons and the embedded contact homology of T^3* , *Geom. Topol.* 10 (2006) 169–266 MR
- [10] **M Hutchings, C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders, II*, *J. Symplectic Geom.* 7 (2009) 29–133 MR
- [11] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, I: Heegaard Floer homology and Seiberg–Witten Floer homology*, *Geom. Topol.* 24 (2020) 2829–2854
- [12] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, III: Holomorphic curves and the differential for the ech/Heegaard Floer correspondence*, *Geom. Topol.* 24 (2020) 3013–3218
- [13] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, *Geom. Topol.* 10 (2006) 955–1096 MR Correction in 18 (2014) 17–30
- [14] **D McDuff**, *The local behaviour of holomorphic curves in almost complex 4–manifolds*, *J. Differential Geom.* 34 (1991) 143–164 MR
- [15] **D McDuff, D Salamon**, *J–holomorphic curves and quantum cohomology*, University Lecture Series 6, Amer. Math. Soc., Providence, RI (1994) MR
- [16] **C B Morrey, Jr.**, *Multiple integrals in the calculus of variations*, Grundle. Math. Wissen. 130, Springer (1966) MR
- [17] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math.* 159 (2004) 1027–1158 MR
- [18] **P Ozsváth, Z Szabó**, *Heegaard diagrams and Floer homology*, from “International Congress of Mathematicians” (M Sanz-Solé, J Soria, J L Varona, J Verdera, editors), volume II, Eur. Math. Soc., Zürich (2006) 1083–1099 MR
- [19] **R Siefring**, *Relative asymptotic behavior of pseudoholomorphic half-cylinders*, *Comm. Pure Appl. Math.* 61 (2008) 1631–1684 MR
- [20] **C H Taubes**, *SW \Rightarrow Gr: from the Seiberg–Witten equations to pseudo-holomorphic curves*, *J. Amer. Math. Soc.* 9 (1996) 845–918 MR Reprinted in “Seiberg–Witten and Gromov invariants for symplectic 4–manifolds” (R Wentworth, editor), First International Press Lecture Series 2, International, Somerville, MA (2000) 1–97
- [21] **C H Taubes**, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$* , *Geom. Topol.* 6 (2002) 657–814 MR
- [22] **C H Taubes**, *Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: properties and existence*, *Geom. Topol.* 10 (2006) 785–928 MR
- [23] **C H Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology, IV*, *Geom. Topol.* 14 (2010) 2819–2960 MR
- [24] **C H Taubes**, *Tamed to compatible: symplectic forms via moduli space integration*, *J. Symplectic Geom.* 9 (2011) 161–250 MR

- [25] **V Turaev**, *Torsion invariants of Spin^c -structures on 3-manifolds*, Math. Res. Lett. 4 (1997) 679–695 MR
- [26] **J G Wolfson**, *Gromov’s compactness of pseudo-holomorphic curves and symplectic geometry*, J. Differential Geom. 28 (1988) 383–405 MR
- [27] **R Ye**, *Gromov’s compactness theorem for pseudo holomorphic curves*, Trans. Amer. Math. Soc. 342 (1994) 671–694 MR

*Department of Mathematics, University at Buffalo
Buffalo, NY, United States*

*Department of Mathematics, The Chinese University of Hong Kong
Shatin, NT, Hong Kong*

*Department of Mathematics, Harvard University
Cambridge, MA, United States*

kutluhan@buffalo.edu, yjlee@math.cuhk.edu.hk,
chtaubes@math.harvard.edu

Proposed: Tomasz Mrowka
Seconded: András I Stipsicz, Bruce Kleiner

Received: 17 February 2012
Revised: 3 November 2015