

HF = HM, III
Holomorphic curves and the differential for
the ech/Heegaard Floer correspondence

ÇAĞATAY KUTLUHAN
 YI-JEN LEE
 CLIFFORD HENRY TAUBES

This is the third 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 the relationship between the differential on the embedded contact homology chain complex and the differential on the Heegaard Floer chain complex. The paper also describes the relationship between the various canonical endomorphisms that act on the homology groups of these two complexes.

53D42

1. Ech on Y and Heegaard Floer homology on M	3015
2. The approximations	3045
3. Cauchy–Riemann equations on $\mathbb{R} \times \mathcal{X}$	3059
4. Proof of Proposition 2.2 when $\Delta_p = 0$	3072
5. Analytic background for the $\Delta_p > 0$ cases	3086
6. Proof of Proposition 2.2: the $\Delta_p > 0$ case	3109
7. Cobordisms to the ech-HF submanifold moduli space	3127
8. Counting ech-HF submanifolds	3176
9. Proof of Theorem 1.1	3211
References	3217

The Heegaard Floer homology of a given compact and oriented three manifold as defined by Peter Ozsváth and Zoltán Szabó [13] is computed using a suitably chosen Morse function and associated pseudogradient vector field. This given manifold is denoted by M . The second paper in this series [9] used this Heegaard Floer data for M to construct geometric data on the connect sum of M with a certain number of copies of $S^1 \times S^2$. What follows uses Y to denote this connect sum but with orientation opposite from that on M . The geometric data on Y was used in [9] to define a stable Hamiltonian version of Michael Hutchings' embedded contact homology [5]. The generators of the latter chain complex were described in [9]. In particular, the \mathbb{Z} -module that serves as the chain complex for the embedded contact homology was written in [9] as the tensor product of the \mathbb{Z} -module that serves as the Heegaard Floer chain complex on M and a second, canonical factor. The pseudoholomorphic curves that are used to define the embedded contact homology differential were also described in [9]. This paper uses as input the latter's description of these curves to relate the differential on Y 's embedded contact homology chain complex to the differential on M 's Heegaard Floer chain complex. The relationship between the two differentials leads to an isomorphism between the embedded contact homology chain complex on Y and a tensor product of two factors, one being the Heegaard Floer homology on M and the other a certain canonical \mathbb{Z} -module.

The first paper in this series [8] uses the isomorphism provided by this paper as one of a triad of isomorphisms which compose to define an isomorphism between the Heegaard Floer homology on M and the Seiberg–Witten Floer homology on M . The input from this paper to [8] is summarized by [8, Theorem 2.3], which is restated here as part of Theorem 1.1. The proof of Theorem 1.1 and thus of [8, Theorem 2.3] constitutes almost all of this article.

Notation, definitions, constructions and results from [9] are used here freely, and so the reader should be familiar with [9]. Note in particular the following convention: Section numbers, equation numbers, and other references from [9] are distinguished from those in this paper by the use of the Roman numeral II as a prefix. For example, “Section II.1” refers to [9, Section 1]. Note also that the convention here as in [9] is to use c_0 to denote 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 used here and in [9] concerns a function that is denoted by χ . It is a fixed, nonincreasing function on \mathbb{R} that equals 1 on $(-\infty, 0]$ and equals 0 on $[1, \infty)$.

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. Yi-Jen Lee has also been supported by Hong Kong Research Grants Council grants GRF-401913, 14316516, 1405541.

1 Embedded contact homology on Y and Heegaard Floer homology on M

The first four subsections that follow provides a very brief summary of what is done in [9]. Section 1A describes Y and its relevant geometry and Section 1B describes the \mathbb{Z} -module that serves for the embedded contact homology chain complex on Y . Section 1C supplies a short primer on the almost complex geometry of $\mathbb{R} \times Y$. Section 1D briefly describes how the corresponding pseudoholomorphic curves are used to define the differential for the embedded contact geometry chain complex. This subsection also describes how these curves are used to define certain important endomorphisms of the of this chain complex.

Section 1E states the central result of this article, this being Theorem 1.1. It characterizes the embedded contact homology differential in terms of the differential that defines the Heegaard Floer homology of M . This theorem likewise characterizes the aforementioned endomorphisms of the embedded contact homology chain complex in terms of their analogs for the Heegaard Floer chain complex on M . As noted in the introductory remarks, Theorem 1.1 restates [8, Theorem 2.3]. Section 1F gives a brief look ahead at the proof of Theorem 1.1.

The last subsections, Sections 1G and 1H, supply some additional background from [9] for use in the proof of Theorem 1.1.

1A The manifold Y and its geometry

The manifold Y is diffeomorphic to the connect sum of M and $g + 1$ copies of $S^1 \times S^2$. A useful handle decomposition of Y is constructed from data on M that is

used to define M 's Heegaard Floer homology. The first item from this data set is a self-indexing Morse function, this denoted by f . The map f has image $[0, 3]$; it has one index 0 critical point, one index 3 critical point, some $G \geq 1$ index 1 critical points and the same number of index 2 critical points. The latter are on the respective $f = 1$ and $f = 3$ level sets. The level set $f = \frac{3}{2}$ is denoted by Σ , this being the Heegaard surface, a surface of genus G . The second item from the data set is the choice of a class in $H^2(M; \mathbb{Z})$ which defines a homomorphism from $H_2(M; \mathbb{Z})$ to $2\mathbb{Z}$. This class is denoted in what follows by c_{1M} . A $\text{Spin}^{\mathbb{C}}$ structure will be chosen in a moment, and its first Chern class will play the role of c_{1M} . The third item from this data set is a fiducial point in Σ , this denoted here by z_0 . The final item consists of an appropriate pseudogradient vector field for f . This vector field is denoted by \mathfrak{v} ; it is defined on the complement of f 's critical points and it is such that $\mathfrak{v}(f) = 1$. This vector field is chosen to obey various constraints; these are described in Sections II.1B, II.1C and II.1D. Note in particular that \mathfrak{v} is constrained so as to give what Ozsváth and Szabó in [13] deem to be a strongly admissible Heegaard diagram for the chosen class c_{1M} and the point z_0 . The data consisting of $(f, c_{1M}, z_0, \mathfrak{v})$ is said in what follows to be the Heegaard Floer data. Constants that depend on just this data are said to *depend solely on the Heegaard Floer data*.

The construction of Y requires the choice of a pairing between the set of index 1 critical points of f and the set of index 2 critical points of f . The resulting set of G pairs is denoted by Λ . An element $\mathfrak{p} \in \Lambda$ is written as an ordered pair of points with it understood that the first entry is the index 1 critical point of f and the second entry is the index 2 critical point of f .

The definition of Y also requires the choice of two additional positive numbers, these denoted by δ_* and R . The constant δ_* is from $(0, 1)$ and it is determined solely by the Heegaard Floer data. The constant R has the lower bound $-100 \ln \delta_*$. This constant R has no a priori upper bound, and the freedom to take R as large as needed is exploited in [9] and in the constructions to come in this article.

The construction of the geometry needed for the \mathbb{Z} -module that serves as the embedded contact geometry chain complex requires the choice of two additional positive numbers, these denoted by δ and χ_0 . The latter with R are not determined by the Heegaard Floer data. The trio (δ, χ_0, R) are constrained by the requirements that $\delta < \delta_*$, $\chi_0 < \delta^3$ and $R \geq -c_0 \ln \chi_0$. Note in particular that the choice of δ determines an upper bound for χ_0 , and that the choice of χ_0 subject to this upper bound then determines a lower bound for R . Constants δ , χ_0 and R that satisfy these bounds are said to be *appropriate*.

The remaining parts of this subsection describe first Y and then the geometry that is needed to define the embedded contact homology \mathbb{Z} -module.

Part 1 As noted above, the manifold Y is diffeomorphic to the connect sum of M with $G + 1$ copies of $S^1 \times S^2$. Section II.1A supplies a useful handle decomposition of Y as $M_\delta \cup \mathcal{H}_0 \cup_{\mathfrak{p} \in \Lambda} \mathcal{H}_\mathfrak{p}$ where M_δ is the complement in M of a certain set of $2(G + 1)$ disjoint balls whose centers are the critical points of f , and where \mathcal{H}_0 and each $\mathfrak{p} \in \Lambda$ version of $\mathcal{H}_\mathfrak{p}$ is a copy of $[-1, 1] \times S^2$. Figure 1 shows a sketch of M_δ (leaving out balls near the index 0 and index 3 critical points of f .) The sketch 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

What follows next summarizes from Section II.1A how the 1-handles are attached.

The handles $\{\mathcal{H}_\mathfrak{p}\}_{\mathfrak{p} \in \Lambda}$ Fix $\mathfrak{p} = (p_+, p_-) \in \Lambda$. The constant δ_* is chosen so that there are respective coordinate charts centered on the index 1 critical point p_+ and index 2 critical points p_- with coordinates (x, y, z) defined where $|x|^2 + |y|^2 + |z|^2 \leq (10\delta_*)^2$ and such that f appears as

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

Use $(r_+, (\theta_+, \varphi_+))$ to denote the standard spherical coordinates for the Euclidean coordinate chart centered on p_+ , and likewise use $(r_-, (\theta_-, \varphi_-))$ to denote the spherical coordinates for the coordinate chart centered on p_- . When $d \in (0, 10\delta_*)$, the ball in M given by $r_+ < d$ is said to be the *radius d coordinate ball* centered on p_+ , and the corresponding $r_- < d$ ball is said to be the *radius d coordinate ball* centered on p_- .

The handle \mathcal{H}_p is given coordinates $(u, (\theta, \phi))$, where (θ, ϕ) are the standard spherical coordinates on the S^2 factor, and where $u \in [-R - \ln(7\delta_*), R + \ln(7\delta_*)]$ is the coordinate for the interval factor. The handle \mathcal{H}_p is attached to the complement in M of the radius $e^{-2yR}(7\delta_*)^{-1}$ coordinate balls centered on p_+ and p_- via the identifications given by

$$(1-2) \quad (r_+ = e^{-(R-u)}, (\theta_+ = \theta, \varphi_+ = \phi)), \quad (r_- = e^{-(R+u)}, (\theta_- = \pi - \theta, \varphi_- = \phi)).$$

The part of \mathcal{H}_p where $1 - 3 \cos^2 \theta > 0$ is denoted by \mathcal{H}_p^+ . Any given constant u slice of \mathcal{H}_p^+ is an annular neighborhood of the equator in S^2 .

The handle \mathcal{H}_0 The constant δ_* is chosen so that respective coordinate charts centered on the index 0 and index 1 critical points of f have coordinates (x, y, z) that are defined where the coordinate functions obey $|x|^2 + |y|^2 + |z|^2 \leq (10\delta_*)^2$ and are such that f appears as

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

Use $(r_+, (\theta_+, \varphi_+))$ to denote the standard spherical coordinates for the Euclidean coordinate chart centered on the index 0 critical point of f , and use $(r_-, (\theta_-, \varphi_-))$ to denote the spherical coordinates for the coordinate chart centered on the index 0 critical point of f . When $d \in (0, 10\delta_*)$, the ball in M given by $r_+ < d$ is said to be the radius d coordinate ball about the index 0 critical point of f . The corresponding $r_- < d$ ball is said to be the radius d coordinate ball about the index 3 critical point of f .

The handle \mathcal{H}_0 is given coordinates $(u, (\theta, \phi))$ where (θ, ϕ) are the standard spherical coordinates on the S^2 factor, and where $u \in [-R - \ln(7\delta_*), R + \ln(7\delta_*)]$ is the coordinate for the interval factor. The handle \mathcal{H}_0 is attached to the complement in M of the radius $e^{-2R}(7\delta_*)^{-1}$ coordinate balls centered on the index 0 and index 3 critical points of f by the rule in (1-2).

The constant δ_* is chosen so that the respective radius $10\delta_*$ coordinate balls about any two distinct critical points of f are disjoint. Given $r \in (e^{-2R}(7\delta_*)^{-1}, 10\delta_*)$, the complement in M of the union of the radius r coordinate balls centered on M is denoted by M_r . The description of Y just given identifies M_r with a subset in Y . The latter is denoted also by M_r . The just-described identification is used implicitly in what follows to view these two incarnations of M_r as one and the same 3-manifold with boundary. In particular, this identification defines f as a function on the M_r part of Y , the latter also denoted by f .

Part 2 A stable Hamiltonian structure on Y consists of a pair (a, ω) where a is a 1-form, ω is a 2-form and these are such that $d\omega = 0$ and $da \in \text{Span}(\omega)$. Moreover, $a \wedge \omega$ is nowhere zero, and this 3-form defines the orientation for Y that is opposite to that defined for Y by M 's orientation of M_δ . Being nowhere zero, the 2-form ω defines a homomorphism from TY to T^*Y whose kernel is a real line bundle over Y . The 1-form a is nonzero on this line subbundle. This understood, let v henceforth denote the vector field that spans the kernel of ω and has pairing 1 with a .

Sections II.1B–II.1E describe a stable Hamiltonian structure for Y that is defined using the Heegaard Floer data $(f, c_{1M}, \mathfrak{v})$ and appropriate constants δ, χ_0 and R . The salient features of a, ω and v are summarized in a moment. This summary restates what is said in Section II.1E.

By way of notation, the upcoming formulas use functions χ, χ_+ and χ_- of $u \in \mathbb{R}$ given by $\chi = \chi_0 \chi(|u| - R - \ln \delta + 12)$ and $\chi_+ = \chi(-u - \frac{1}{4}R)$ and $\chi_- = \chi(u - \frac{1}{4}R)$. The formulas also employ functions f and g of the variable u given by

$$(1-4) \quad f = \chi + 2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)}) \quad \text{and} \quad g = \chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)}.$$

Their respective derivatives are denoted by f' and g' . (By way of a reminder, the function χ denotes here and subsequently a fixed, nonincreasing function on \mathbb{R} that equals 1 on $(-\infty, 0]$ and equals 0 on $[1, \infty)$.)

The bullets that follow supply the promised description of a, ω and v .

- **On M_δ** The 2-form ω on M_δ is nowhere zero on the kernel of the 1-form df and v here is the pseudogradient vector field \mathfrak{v} .
- **In the handle \mathcal{H}_0** The 2-form ω and the vector field v on \mathcal{H}_0 are

$$(1-5) \quad \omega = \sin \theta \, d\theta \wedge d\phi \quad \text{and} \quad v = \frac{1}{2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)})} \frac{\partial}{\partial u}.$$

- **In the handles $\{\mathcal{H}_p\}_{p \in \Lambda}$** Fix $p \in \Lambda$. The trio a, ω and v on \mathcal{H}_p are

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

Here, α is a certain positive function of the pair (u, θ) .

The next bullet concerns the cohomology class of the form w . This bullet refers to the direct-sum decomposition

$$(1-7) \quad H_2(Y; \mathbb{Z}) = H_2(M; \mathbb{Z}) \oplus (H_2(\mathcal{H}_0; \mathbb{Z}) \oplus \left(\bigoplus_{p \in \Lambda} H_2(\mathcal{H}_p; \mathbb{Z}) \right))$$

that comes via Mayer–Vietoris by writing $Y = M_\delta \cup \mathcal{H}_0 \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$. The summands in (1-7) that correspond to the various 1–handles are isomorphic to \mathbb{Z} , and any oriented, cross-sectional sphere is a generator. The convention in what follows is to orient these spheres with the 2–form $\sin \theta \, d\theta \, d\phi$.

- **The cohomology class of w** 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 $H_2(\mathcal{H}_0; \mathbb{Z})$, it has value zero on each $p \in \Lambda$ version of $H_2(\mathcal{H}_p; \mathbb{Z})$, and it acts on the $H_2(M; \mathbb{Z})$ –summand in (1-7) as the pairing with the chosen class c_{1M} .

A particular integral curve of the vector field v plays a distinguished role in the embedded contact homology story. This curve is described next.

- **The curve through z_0** There is a closed integral curve of v in $M_\delta \cup \mathcal{H}_0$ whose intersection with Σ is the chosen fiducial point z_0 . This curve is denoted by $\gamma^{(z_0)}$. It also has a single intersection with each cross-sectional sphere in \mathcal{H}_0 .

The final bullets introduce a pair of auxiliary 1–forms on Y that play central roles. The definition of the first of these 1–forms refers to the function f_* that is defined on any given $p \in \Lambda$ version of \mathcal{H}_p by the rule

$$(1-8) \quad f_* = (\chi_+ e^{2(u-R)} - \chi_- e^{-2(u+R)})(1 - 3 \cos^2 \theta).$$

The definition of the second of these 1–forms refers to the function χ_δ that is defined on any given $p \in \Lambda$ version of \mathcal{H}_p by the rule $\chi_\delta = \chi(|u| - R - \ln \delta + 10)$.

- **The 1–form ν_\diamond** The 1–form ν_\diamond is closed and is such that $\nu_\diamond \wedge w \geq 0$. Furthermore, $\nu_\diamond \wedge w = 0$ only where both $u = 0$ and $1 - 3 \cos^2 \theta = 0$ on each $p \in \Lambda$ version of \mathcal{H}_p . This 1–form equals df on M_δ , it is given by $\nu_\diamond = 2(\chi_+ e^{2(u-R)} + \chi_- e^{-2(u+R)}) \, du$ on \mathcal{H}_0 , and it is given by df_* on any given $p \in \Lambda$ version of \mathcal{H}_p .
- **The 1–form \hat{a}** The 1–form \hat{a} has pairing 1 with v and is such that $\hat{a} \wedge w > 0$. It is ν_\diamond on $M_\delta \cup \mathcal{H}_0$ and it is $\chi_\delta a + (1 - \chi_\delta)\nu_\diamond$ on any given $p \in \Lambda$ version of \mathcal{H}_p .

1B The embedded contact homology \mathbb{Z} -module

This subsection describes the \mathbb{Z} -module that serves as the chain complex for embedded contact homology. The subsection has four parts that briefly summarize material from Section II.1F and Section II.2.

Part 1 Fix a $\text{Spin}^{\mathbb{C}}$ structure on M and use c_{1M} now to denote the associated first Chern class in $H^2(M; \mathbb{Z})$. This class is used to construct the strongly admissible Heegaard diagram that is used to define the Heegaard Floer chain complex on M .

The \mathbb{Z} -module that serves for the Heegaard Floer chain complex on M for the chosen $\text{Spin}^{\mathbb{C}}$ structure can be defined with the help of a finite set that is denoted by \mathcal{Z}_{HF} . Any given element in \mathcal{Z}_{HF} is viewed here and in [9] as a suitably constrained, unordered G -tuple of integral curves of v . Let \hat{v} denote an element from \mathcal{Z}_{HF} . There are three constraints on \hat{v} : First, each constituent integral curve from \hat{v} runs from an index 1 critical point of f to an index 2 critical point of f . Second, no two distinct constituents share the same index 1 critical point or the same index 2 critical point. This being the case, \hat{v} defines a pairing between the set of index 1 critical points of f and the set of index 2 critical points of f . The third constraint demands that the G points that form $\bigcup_{v \in \hat{v}} (v \cap \Sigma)$ with the point z_0 define the chosen $\text{Spin}^{\mathbb{C}}$ structure in the manner that is described in [13].

The Heegaard Floer chain complex is the free \mathbb{Z} -module generated by the elements of the set $\mathcal{Z}_{\text{HF}} \times \mathbb{Z}$. This module is denoted by $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z})$. This interpretation of the Heegaard Floer chain complex is used by Robert Lipshitz in [10] to reformulate Heegaard Floer homology.

Part 2 The class c_{1M} is used to define Y and its stable Hamiltonian data (a, ω, v) . The \mathbb{Z} -module for the relevant version of embedded contact homology on Y is defined with the help of a set that is denoted by $\mathcal{Z}_{\text{ech}, M}$ and whose elements are suitably constrained, finite sets of closed integral curve of v that lie entirely in the union of the various $p \in \Lambda$ versions of \mathcal{H}_p and the $f \in (1, 2)$ part of M_δ . The set \mathcal{Z}_{ech} is described in the upcoming Part 3 of this subsection. What follows directly summarizes some of what is said in Section II.2 about the closed integral curves of v that lie entirely in the subset of Y just described.

Closed curves in $\bigoplus_{p \in \Lambda} \mathcal{H}_p$ Fix $p \in \Lambda$. There are precisely two integral curves of v that lie entirely in \mathcal{H}_p . These constitute the two components of the locus where both




Figure 2: The $u = 0$ sphere in \mathcal{H}_p .

$u = 0$ and $1 - 3 \cos^2 \theta = 0$. The curve with $\cos \theta = \frac{1}{\sqrt{3}}$ is denoted by $\hat{\gamma}_p^+$ and that where $\cos \theta = -\frac{1}{\sqrt{3}}$ is denoted by $\hat{\gamma}_p^-$. These curves are labeled on Figure 2.

Intersections with M_δ There exists a purely Heegaard Floer dependent constant $\kappa \geq 1$ whose significance is described in what follows. Construct Y with $\delta < \kappa^{-1} \delta_*^3$. Let γ denote a closed integral curve of v in $M_\delta \cup (\bigcup_p \mathcal{H}_p)$ that intersects M_δ . Then $\gamma \cap M_\delta$ consists of a finite set of segments of integral curves of v in the $f^{-1}((1, 2))$ part of M_δ . Each such segment lies in the radius $\kappa \delta$ tubular neighborhood of an integral curve of v that runs from an index 1 critical point of f to an index 2 critical point of f .

Intersections with $\bigcup_{p \in \Lambda} \mathcal{H}_p$ Let γ denote a closed integral curve of v in the space $M_\delta \cup (\bigcup_p \mathcal{H}_p)$ that intersects M_δ . Fix $p \in \Lambda$. The intersection of γ with \mathcal{H}_p consists of a finite set of segments. Let γ_* denote any one such segment. The following is true:

- (1-9) • γ_* sits where $1 - 3 \cos^2 \theta > 0$.
- γ_* runs from the $u = -R - \ln(7\delta_*)$ end of \mathcal{H}_p to the $u = R + \ln(7\delta_*)$ end.
- The function $\hat{h} = f(u) \cos \theta \sin^2 \theta$ is constant on γ_* .
- The coordinate u restricts as an affine coordinate to γ .
- The angle ϕ on γ changes according to the rule

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

The assertion about the closed integral curves in $\bigcup_{p \in \Lambda} \mathcal{H}_p$ summarizes Lemma II.2.1, the assertion about $\gamma \cap M_\delta$ summarizes Lemma II.2.4, and the assertion about $\gamma \cap (\bigcup_{p \in \Lambda} \mathcal{H}_p)$ summarizes Lemma II.2.2 and (II.2-5).

Part 3 Let \mathcal{O} denote the four-element set $\{0, 1, -1, \{1, -1\}\}$. The elements in $\mathcal{Z}_{\text{ech}, \mathcal{M}}$ enjoy a 1-to-1 correspondence

$$(1-10) \quad \mathcal{Z}_{\text{ech}, \mathcal{M}} = \mathcal{Z}_{\text{HF}} \times \left(\prod_{p \in \Lambda} (\mathbb{Z} \times \mathcal{O}) \right).$$

This correspondence is canonical when the various $p \in \Lambda$ factors of \mathbb{Z} are viewed as affine spaces modeled on \mathbb{Z} . A basepoint in a given $p \in \Lambda$ version is determined by the choice of a lift to \mathbb{R} of the $\mathbb{R}/2\pi\mathbb{Z}$ value ϕ coordinate in \mathcal{H}_p .

What follows describes the geometric meaning of the correspondence in (1-10). Write a given element in $\mathcal{Z}_{\text{HF}} \times (\prod_{p \in \Lambda} (\mathbb{Z} \times \mathcal{O}))$ as $(\hat{v}, (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$ with \hat{v} from \mathcal{Z}_{HF} and with any given $p \in \Lambda$ version of $(\mathfrak{k}_p, \mathcal{O}_p)$ in $\mathbb{Z} \times \mathcal{O}$. Let Θ denote the corresponding element in $\mathcal{Z}_{\text{ech}, \mathcal{M}}$. As noted in Part 2, each element in $\mathcal{Z}_{\text{ech}, \mathcal{M}}$ is a finite set of closed integral curves of v that lie in $M_\delta \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p)$. With this in mind, consider first the significance of the entry \hat{v} . The intersection of $\bigcup_{\gamma \in \Theta} \gamma$ with M_δ has G components, each being a segment of an integral curve of v that runs from the boundary of the radius δ coordinate ball about an index 1 critical point of f to the boundary of the radius δ coordinate ball centered on an index 2 critical point of f . The components $M_\delta \cap (\bigcup_{\gamma \in \Theta} \gamma)$ enjoy a 1-to-1 correspondence with the integral curves from \hat{v} with the correspondence such that a given segment from $M_\delta \cap (\bigcup_{\gamma \in \Theta} \gamma)$ lies in the radius $c_0\delta$ tubular neighborhood of its partner from \hat{v} . This version of c_0 depends only on the Heegaard Floer data.

To say more about the curves in Θ , fix $p \in \Lambda$. The intersection of $\bigcup_{\gamma \in \Theta} \gamma$ with \mathcal{H}_p has precisely one component that crosses \mathcal{H}_p from the $u = -R - \ln(7\delta_*)$ end to the end where $u = R + \ln(7\delta_*)$. More is said about this component in a moment. The remaining components (if any) are determined by \mathcal{O}_p using the following rule:

- (1-11) • If $\mathcal{O}_p = 0$, then Θ contains neither $\hat{\gamma}_p^+$ nor $\hat{\gamma}_p^-$.
- If $\mathcal{O}_p = 1$ or $\mathcal{O}_p = -1$, then Θ contains $\hat{\gamma}_p^+$ or $\hat{\gamma}_p^-$ respectively, but not both of them.
- If $\mathcal{O}_p = \{1, -1\}$, then Θ contains both $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$.

The endpoints of the segment of $(\bigcup_{\gamma \in \Theta} \gamma) \cap \mathcal{H}_p$ that crosses \mathcal{H}_p intersects the respective $u = R + \ln \delta$ and $u = -R - \ln \delta$ spheres in \mathcal{H}_p at a point whose spherical

coordinates (θ, ϕ) differ by at most $c_0\delta$ from the coordinates of the sphere’s intersection with $\bigcup_{\nu \in \hat{\nu}} \nu$. Note that this version of c_0 depends solely on the Heegaard Floer data. The entry ξ_p in Θ ’s label gives an indication of the total change in the coordinate ϕ along this segment. To elaborate, view p ’s entry of \mathbb{Z} in (1-10) as an affine space modeled on \mathbb{Z} . Suppose that Θ and Θ' are any two elements from $\mathcal{Z}_{\text{ech},M}$. Label these two elements as in (1-10) and assume that they have identical \mathcal{Z}_{HF} factor. Write their respective entries in p ’s factor of the affine copy of \mathbb{Z} as ξ_p and $\xi_p + k$ with $k \in \mathbb{Z}$. Let Δ denote the total change in ϕ along part of the curve from Θ that crosses \mathcal{H}_p , and let Δ' denote the analogous Θ' angle change. Then

$$(1-12) \quad \Delta' - \Delta = k + \epsilon,$$

where $|\epsilon| \leq c_0\delta$. As before, this version of c_0 depends only on the Heegaard Floer data.

See Section II.2 and in particular Corollary II.2.7 and Proposition II.2.8 for an expanded version of what was just said in this Part 3.

Part 4 The \mathbb{Z} -module for the relevant version of embedded contact homology is freely generated by a certain principal \mathbb{Z} -bundle over $\mathcal{Z}_{\text{ech},M}$. This bundle is denoted in what follows by $\hat{\mathcal{Z}}_{\text{ech},M}$ and the free \mathbb{Z} -module generated by $\hat{\mathcal{Z}}_{\text{ech},M}$ is denoted by $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$. Elements of this \mathbb{Z} -module are finite, integer-weighted formal sums of elements in $\hat{\mathcal{Z}}_{\text{ech},M}$.

The definition of $\hat{\mathcal{Z}}_{\text{ech},M}$ is given in a moment. By way of preliminaries, note that any given integral curve of ν is oriented by ν . This being the case, each closed integral curve of ν defines a closed 1-cycle in Y . When γ denotes the closed integral curve, then $[\gamma]$ is used to denote the corresponding cycle. The set of the closed integral curves from any given $\Theta \in \mathcal{Z}_{\text{ech},M}$ likewise defines a closed 1-cycle in Y , this being the cycle $\sum_{\gamma \in \Theta} [\gamma]$. The latter cycle is denoted by $[\Theta]$. The Poincaré dual of $2[\Theta]$ defines via the tautological pairing the homomorphism from $H_2(Y; \mathbb{Z})$ to $2\mathbb{Z}$ that acts as follows: it sends the $H_2(\mathcal{H}_0; \mathbb{Z})$ -summand in (1-7) to 0, it sends the generator of each $p \in \Lambda$ -labeled summand in (1-7) to 2, and it acts on the $H_2(M; \mathbb{Z})$ summand as the pairing with the given class c_{1M} .

A somewhat noncanonical description of $\hat{\mathcal{Z}}_{\text{ech},M}$ requires the choice of a fiducial element $\Theta_0 \in \mathcal{Z}_{\text{ech},M}$. This done, any given element $\hat{\Theta}$ can be viewed as an equivalence class of pairs (Θ, Z) with $\Theta \in \mathcal{Z}_{\text{ech},M}$ and with Z an element in the \mathbb{Z} -module $H_2(Y; [\Theta] - [\Theta_0])$. The equivalence relation is defined with the help of the closed

integral curve $\gamma^{(z_0)}$ that is described in the fifth bullet of Part 2 in Section 1B. Pairing with the Poincaré dual of $\gamma^{(z_0)}$ defines a homomorphism from the \mathbb{Z} -module of closed 2-cycles to \mathbb{Z} . This pairing is denoted by $[\gamma^{(z_0)}]^{Pd}(\cdot)$. The equivalence relation has $(\Theta, Z) \sim (\Theta', Z')$ if and only if $\Theta = \Theta'$ and also $[\gamma^{(z_0)}]^{Pd}(Z - Z') = 0$. The principal bundle projection map sends an equivalence class (Θ, Z) to Θ . The element $1 \in \mathbb{Z}$ acts to send (Θ, Z) to $(\Theta, Z + [S_0])$, where $[S_0]$ is the $u = 0$ sphere in \mathcal{H}_0 .

A different choice for Θ_0 produces a different principal \mathbb{Z} -bundle over $\mathcal{Z}_{ech,M}$, but the new one and the original are canonically isomorphic. To elaborate, suppose that Θ'_0 is a second choice. 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 the equivalence class of (Θ, Z) in the Θ_0 version of $\widehat{\mathcal{Z}}_{ech,M}$ to that of $(\Theta, Z + Z_0)$ in the Θ'_0 version. A different intersection pairing zero choice gives the same equivalence class and thus the same isomorphism.

As explained next, the existence of cycles with zero intersection pairing against $\gamma^{(z_0)}$ can be exploited to construct a canonical principal \mathbb{Z} -bundle isomorphism

$$(1-13) \quad \mathcal{Z}_{ech,M} \times \mathbb{Z} \rightarrow \widehat{\mathcal{Z}}_{ech,M}.$$

The isomorphism depicted here is defined by a certain section of $\widehat{\mathcal{Z}}_{ech,M}$ whose image corresponds via (1-13) to $\mathcal{Z}_{ech,M} \times \{0\}$. This section sends any given $\Theta \in \mathcal{Z}_{ech,M}$ to the equivalence class of a pair (Θ, Z) , where Z can be any 2-cycle in $H_2(Y; [\Theta] - [\Theta_0])$ that has pairing 0 against $\gamma^{(z_0)}$. A different choice of Z with intersecting pairing zero against $\gamma^{(z_0)}$ defines the same equivalence class and so the same section. This is why (1-13) is canonical. In fact, (1-13) is canonical in the following stronger sense: the canonical isomorphism between any two Θ_0 and Θ'_0 versions of $\widehat{\mathcal{Z}}_{ech,M}$ intertwines their respective versions of (1-13).

The image via (1-13) of the set $\mathcal{Z}_{ech,M} \times \{-\infty, \dots, -1\}$ defines a subfiber bundle in $\widehat{\mathcal{Z}}_{ech,M}$. The latter is denoted by $\widehat{\mathcal{Z}}^0_{ech,M}$. The free \mathbb{Z} -module generated by the elements in $\widehat{\mathcal{Z}}^0_{ech,M}$ plays a central role in Theorem 2.3 of [8]. This submodule is denoted here by $\mathbb{Z}(\widehat{\mathcal{Z}}^0_{ech,M})$.

Pairing with the class c_{1M} defines a linear functional from $H_2(M; \mathbb{Z})$ to $2\mathbb{Z}$. Let $p_M \in 2\mathbb{Z}$ denote the divisibility of the subgroup defined by the image. Rules laid out by Hutchings (see [4]) can be used here to give each generator of $\widehat{\mathcal{Z}}_{ech,M}$ a relative $\mathbb{Z}/(p_M\mathbb{Z})$ degree and so give $\mathbb{Z}(\widehat{\mathcal{Z}}_{ech,M})$ a relative $\mathbb{Z}/(p_M\mathbb{Z})$ grading.

1C The almost complex geometry of $\mathbb{R} \times Y$

An endomorphism of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ that serves as the embedded contact homology differential is defined using certain sorts of submanifolds in $\mathbb{R} \times Y$. The latter are pseudoholomorphic for a chosen almost complex structure. Part 1 of this subsection describes the allowed almost complex structures. Part 2 of the subsection summarizes some standard definitions.

Part 1 Section II.3A and (II.6-1) describe the constraints that delineate the set of allowed almost complex structures on $\mathbb{R} \times Y$. The first two constraints are the $\mathbb{R} \times Y$ versions of standard constraints that are used in all contact and symplectic versions of Floer homology. The remaining constraints are special to the situation at hand. By way of notation, the Euclidean coordinate on the \mathbb{R} factor of $\mathbb{R} \times Y$ is denoted by s .

Let J denote a given almost complex structure on $\mathbb{R} \times Y$. This almost complex structure is allowed if it has the properties listed in the seven bullets that follow:

- J maps the Euclidean tangent vector ∂_s to the \mathbb{R} factor of $\mathbb{R} \times Y$ to v .
- J is not changed by constant translations along the \mathbb{R} factor of $\mathbb{R} \times Y$.
- J preserves the kernel of the 1-form \widehat{a} , and its restriction to this 2-plane field defines the orientation given by w .

The next two bullets concern the restriction of J to any given $\mathfrak{p} \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_{\mathfrak{p}}$. The statement of the second refers to the vector fields

$$(1-14) \quad \begin{aligned} e_1 &= -6g \cos \theta \sin \theta \partial_u + (\chi + g')(1 - 3 \cos^2 \theta) \partial_{\theta}, \\ e_2 &= \partial_{\phi} + \sqrt{6} \chi_{\delta} f \cos \theta \sin^2 \theta v. \end{aligned}$$

These span the kernel of \widehat{a} where both u and $(1 - 3 \cos \theta) \sin \theta$ are not zero.

- J is unchanged by constant, $\mathbb{R}/2\pi\mathbb{Z}$ translations of the coordinate ϕ .
- $J e_1 = \sigma^{-1} e_2$, where σ is a positive function of u and θ .

The final two bullets concerns the restriction of J to $\mathbb{R} \times M_{\delta}$. The first of these refers to two sets of pairwise disjoint annuli in the Heegaard surface Σ . The annuli in the first set are labeled by the index 1 critical points of f , and those in the second are labeled by the index 2 critical points of f . Let p denote a given index 1 or index 2 critical point. The corresponding annulus is denoted respectively by T_{p+} or T_{p-} .

To say more about these annuli, let p denote an index 1 critical point of f . The annulus T_{p+} is the image via Lie transport along the integral curves of \mathfrak{v} of the annulus

in the radius δ_* coordinate ball centered on p where $1 - 3 \cos^2 \theta_+ > 0$. By way of a review from [9], the image of the central, $\theta_+ = \frac{\pi}{2}$, circle in this annulus is denoted by C_{p_+} ; this is the intersection between the Heegaard surface Σ and the ascending disk from the critical point p . Let h_+ denote the function $2e^{2t_+} \cos \theta_+ \sin^2 \theta_+$ on the radius $8\delta_*$ coordinate ball centered on p . Lie transport by \mathfrak{v} gives T_{p_+} coordinates (φ_+, h_+) with the former being $\mathbb{R}/2\pi\mathbb{Z}$ -valued. The restriction of ω to $T_{p_+} \subset \Sigma$ is given using these coordinates by $\sqrt{6} d\varphi_+ \wedge dh_+$.

Let p now denote an index 2 critical point of f . The annulus T_{p_-} is the image via Lie transport via \mathfrak{v} of the annular region where $1 - 3 \cos^2 \theta_- > 0$ in the boundary of the radius δ_* coordinate ball centered on p . The image in T_{p_-} of the central $\theta_- = \frac{\pi}{2}$ circle is denoted by C_{p_-} ; it is the intersection between Σ and the descending disk from p . Set $h_- = 2e^{2t_-} \cos \theta_- \sin^2 \theta_-$. Lie transport by \mathfrak{v} identifies gives T_{p_-} the coordinates (φ_-, h_-) . The 2-form ω on T_{p_-} is $-\sqrt{6} d\varphi_- \wedge dh_-$.

There is one more point to note regarding an intersection $T_{p_+} \cap T_{p'_-}$. The respective coordinates (φ_+, h_+) for T_{p_+} and (φ_-, h_-) for T_{p_-} are related on this intersection by the rule $(d\varphi_+, dh_+) = \pm(dh_-, d\varphi_-)$ with the $+$ sign taken when the pair of vectors $(\partial/\partial\varphi_+, \partial/\partial\varphi_-)$ define an oriented basis for $T\Sigma$ at the corresponding point in $C_{p_+} \cap C_{p'_-}$.

The union of the annuli in the set labeled by the index 1 critical points of f is denoted by T_+ , and of the union of the annuli from the set labeled by the index 2 critical points of f is denoted by T_- . The union of the index 1 critical point versions of C_{p_+} is denoted by C_+ and the corresponding union of the index 2 critical point versions of C_{p_-} is denoted by C_- .

The bullet that follows identifies the $f \in (1, 2)$ part of M_δ with $(1, 2) \times \Sigma$ in the manner just described:

- $J \partial/\partial\varphi_+ = \partial/\partial h_+$ on $\mathbb{R} \times (1, 2) \times T_+$ and $J \partial/\partial\varphi_- = -\partial/\partial h_-$ on $\mathbb{R} \times (1, 2) \times T_-$.

These two conditions are compatible on $T_+ \cap T_-$ because $(d\varphi_+, dh_+) = \pm(dh_-, d\varphi_-)$ on any given component of $T_+ \cap T_-$.

The final bullet refers to a certain residual subset in the C^∞ -Fréchet space of almost complex structures that obey the preceding bullets. This is the subset \mathcal{J}_{ech} from Theorem II.A.1:

- J comes from the residual set \mathcal{J}_{ech} .

The membership in the residual set \mathcal{J}_{ech} guarantees the vanishing of the cokernel of a Fredholm operator that is associated to certain sorts of pseudoholomorphic submanifolds.

An almost complex structure that obeys the first three bullets is compatible with the 2-form $\hat{\omega} = ds \wedge \hat{a} + \omega$. Said differently: The bilinear form $\hat{\omega}(\cdot, J(\cdot))$ defines a Riemannian metric on $\mathbb{R} \times Y$ when J obeys the first three bullets. Note that this metric gives both $\frac{\partial}{\partial s}$ and v norm 1, it makes them mutually orthogonal, and it makes both orthogonal to the kernel of \hat{a} . This metric also makes J into an orthogonal endomorphism of $T(\mathbb{R} \times Y)$. With J given, the metric $\hat{\omega}(\cdot, J(\cdot))$ is used implicitly in what follows to define norms and covariant derivatives on the various tensor bundles over $\mathbb{R} \times Y$.

An almost complex structure that obeys all seven of these bullets will be said to be a member of \mathcal{J}_{ech} .

Part 2 Let J denote an almost complex structure on $\mathbb{R} \times Y$. Assume for the moment that J obeys only the constraints from the first three bullets of Part 1. A proper subset $C \subset \mathbb{R} \times Y$ is said in what follows to be a J -holomorphic subvariety if it has the following properties:

- (1-15) • C has no isolated points and the complement of a finite set in C is a submanifold with J -invariant tangent space.
- The integral of ω over C is finite.

A J -holomorphic subvariety is said to be irreducible if the complement of any given finite set is connected.

A J -holomorphic subvariety may or may not be compact. If not, these conditions have the various standard implications [3; 18; 7] about the large $|s|$ part of the subvariety. To say more, let C denote a given, noncompact pseudoholomorphic subvariety. There exists $s_0 > 1$ such that the $|s| \geq s_0$ part of C is a disjoint union of embedded cylinders. The 1-form ds is nonzero on the tangent space of each such cylinder. A component cylinder of the $|s| \geq s_0$ part of C is said to be an *end* of C . An end of C where $s \geq s_0$ is said to be *positive* and an end where $s \leq -s_0$ is said to be *negative*. A constant $|s| \geq s_0$ slice of any given end is an embedded circle in Y . This circle appears as a braid in a small radius tubular neighborhood of a closed integral curve of v . As $|s|$ increases, the circle in question moves via an ambient isotopy so as to converge

pointwise as $|s| \rightarrow \infty$ as a multiple cover of the central integral curve of v . The closed integral curve in question is said to be *associated* to the given end.

The set of J -holomorphic subvarieties is given the topology that associates to any given J -holomorphic subvariety a basis of open neighborhoods of the following sort: Let C denote the given subvariety. The sets of C 's neighborhood basis are labeled by $(0, 1)$. A J -holomorphic subvariety C' is in a given $\varepsilon \in (0, 1)$ member of this basis when:

- (1-16) • $\sup_{z \in C} \text{dist}(z, C') + \sup_{z \in C'} \text{dist}(C, z) < \varepsilon$.
- Let μ denote a smooth 2-form on $\mathbb{R} \times Y$ with $|\mu| \leq 1$, with $|\nabla\mu| < \varepsilon^{-1}$ and with compact support where $|s| < \varepsilon^{-1}$. Then $|\int_C \mu - \int_{C'} \mu| < \varepsilon$.

The resulting topological space is called the *moduli space* of J -holomorphic subvarieties. The group \mathbb{R} has a continuous action on the moduli space, this given by the constant translations along the \mathbb{R} factor of $\mathbb{R} \times Y$. An irreducible, \mathbb{R} -invariant J -holomorphic subvariety is the product of \mathbb{R} with a closed integral curve of v .

Of particular interest in much of what follows are the moduli space components that contain elements that are characterized as follows. Let C denote a member.

- (1-17) • C is embedded.
- Distinct ends of C have distinct associated closed integral curves of v . This is also the case for distinct negative ends of C .
 - The constant $|s|$ slice of any given end are isotopic in the tubular neighborhood of the associated integral curve of v to this central integral curve.
 - The set of integral curves of v that are associated to the positive ends of C defines an element of $\mathcal{Z}_{\text{ech},M}$. This is also true for the negative ends.

A J -holomorphic subvariety that is described by the second, third and fourth bullets of (1-17) is said to be an *ech-subvariety*. An ech-subvariety is said to be an *ech-HF* subvariety if it lacks irreducible components that intersect $\mathbb{R} \times M_\delta$ in an $f = \text{constant}$ level set, or that intersect $\mathbb{R} \times \mathcal{H}_0$ in a $u = \text{constant}$ level set, or that intersect some $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$ in the $u = 0$ level set. Any such forbidden irreducible component is described completely by one of Propositions II.3.1–II.3.4. An ech-HF subvariety is said here to be an *ech-HF* submanifold if it obeys all four of the bullets in (1-17).

Let C denote an ech-subvariety. The element in $\mathcal{Z}_{\text{ech},M}$ that comes from the positive ends of C via the fourth bullet of (1-17) is denoted by Θ_{C+} ; the analogous negative

end element in $\mathcal{Z}_{\text{ech},M}$ is denoted by Θ_{C_-} . The smooth part of C is oriented by J , and so C 's image in Y via the projection defines a 2-cycle with boundary $[\Theta_{C_+}] - [\Theta_{C_-}]$. This 2-cycle is denoted by $[C]_Y$. Meanwhile, Hutchings (see [4]) gives rules for assigning an integer to C , its *ech index*. This ech index is denoted here by $I_{\text{ech}}(C)$.

Let $\hat{\Theta}'$ and $\hat{\Theta}$ denote a given pair from $\hat{\mathcal{Z}}_{\text{ech},M}$ and let k denote a given integer. Use $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$ to denote the set of ech subvarieties with membership characterized as follows: Write $\hat{\Theta}$ as a pair (Θ, Z) with $\Theta \in \mathcal{Z}_{\text{ech},M}$ and $Z \in H_2(Y; [\Theta] - [\Theta_0])$. Write in the analogous fashion $\hat{\Theta}' = (\Theta', Z')$. The subvariety C is a member of $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$ when

$$(1-18) \quad \Theta_{C_+} = \Theta, \quad \Theta_{C_-} = \Theta', \quad Z' = Z - [C]_Y \quad \text{and} \quad I_{\text{ech}}(C) = k.$$

The set $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$ is a union of components of the moduli space of J -holomorphic subvarieties. By way of a parenthetical remark, the set $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$ is empty unless the sum of $k \bmod (p_M)$ and the $\mathbb{Z}/p_M\mathbb{Z}$ grading of $\hat{\Theta}'$ equals the $\mathbb{Z}/p_M\mathbb{Z}$ grading of $\hat{\Theta}$.

1D The differential and the geometric endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$

Part 1 of this subsection supplies a brief description of the differential on $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ that defines the relevant version of embedded contact homology. As noted in Appendix II.A.1, rules laid out by Hutchings [6] can be used to define an action of the algebra $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ on the embedded contact homology \mathbb{Z} -module. Part 2 of this subsection describes the generators of this $\mathbb{Z}[\mathbb{U}] \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$ action. Part 3 talks about the grading of this \mathbb{Z} -module.

All that is said in what follows assumes that J comes from \mathcal{J}_{ech} .

Part 1 An endomorphism of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ is given by its action on the generators and the action on any given generator $\hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}$ results in a formal sum of the form

$$(1-19) \quad \hat{\Theta} \rightarrow \sum_{\hat{\Theta}' \in \hat{\mathcal{Z}}_{\text{ech},M}} N_{\hat{\Theta}', \hat{\Theta}} \hat{\Theta}',$$

with each coefficient an integer and where only finitely many coefficients are nonzero. The collection of integers $\{N_{\hat{\Theta}', \hat{\Theta}}\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}}$ defines the endomorphism.

Theorem II.A.1 asserts that the endomorphism of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$ that serves as the differential for embedded contact homology can be defined according to the rules laid out

by Hutchings (see [6; 5; 7]). These rules are summarized in Part 1 of Section 9B to come. Suffice it to say here that each $\{\hat{\Theta}', \hat{\Theta}\}$ version of the relevant version of $N_{\hat{\Theta}', \hat{\Theta}}$ is computed using the components of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$.

Theorem II.A.1 implies that $M_1(\hat{\Theta}', \hat{\Theta})$ is a smooth manifold with a finite set of components, each \mathbb{R} -equivariantly diffeomorphic to \mathbb{R} , and that each component of this space contributes either $+1$ or -1 to a sum that gives $N_{\hat{\Theta}', \hat{\Theta}}$.

The endomorphism that defines the differential for embedded contact homology is denoted in what follows by ∂_{ech} .

Part 2 As noted above, the homology of ∂_{ech} is a \mathbb{Z} -module with certain canonical endomorphisms that generate an action of $\mathbb{Z}(\mathbb{U}) \otimes (\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion}))$. The generators of this action are defined by endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$. As explained in Part 3 of Appendix II.A1, the endomorphism of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ that supplies the action of \mathbb{U} on the homology is defined with the help of a chosen point in either \mathcal{H}_0 or the part of M_δ where $f \in (0, 1) \cup (2, 3)$. Let y denote such a point. It follows from (II.A-6) and Theorem II.A.1 that a given $\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}$ version of the coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ is zero unless $\hat{\Theta}'$ and $\hat{\Theta}$ are related as follows: Write $\hat{\Theta}$ as (Θ, Z) with $\Theta \in \mathcal{Z}_{\text{ech}, M}$ and with $Z \in H_2(Y; [\Theta] - [\Theta_0])$. Then $\hat{\Theta}' = (\Theta, Z - [S])$, where $[S]$ here denotes the $u = 0$ sphere in \mathcal{H}_0 . The coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ in this case is 1.

What follows is also a consequence of Theorem II.A.1. The endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ that generate the action of $\wedge^*(H_1(Y; \mathbb{Z})/\text{torsion})$ on the homology are defined with the help of a chosen, suitably generic basis of cycles that generate $H_1(Y; \mathbb{Z})/\text{torsion}$. Fix such a basis and let $\hat{i} \subset Y$ denote a chosen basis element. Any given $\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}$ coefficient $N_{\hat{\Theta}', \hat{\Theta}}$ in the corresponding version of (1-19) is computed using the submanifolds from $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ that intersect $\{0\} \times \hat{i}$ and the corresponding intersection points. If γ is suitably generic, then the set of pairs consisting of a submanifold in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ and an intersection point with $\{0\} \times \hat{i}$ is a finite set. Moreover, each such intersection point contributes either $+1$ or -1 to a sum that gives $N_{\hat{\Theta}', \hat{\Theta}}$. The upcoming Section 9C explains how these ± 1 contributions are determined.

The upcoming Theorem 1.1 refers to an M -adapted 1-cycle basis for $H_1(Y; \mathbb{Z})/\text{torsion}$. The definition of M -adapted requires the introduction from Part 7 in Section II.1C of a certain finite set in the interior of $\Sigma \setminus (T_- \cup T_+)$. The set contains the fiducial point z_0 and $\dim(H_1(M; \mathbb{Z}))$ additional points. This set is denoted by \mathbb{Y} . Each $z \in \mathbb{Y}$

is the intersection point of Σ with a closed integral curve of ν . The latter curve is denoted by $\gamma^{(z)}$. Pairing with the Poincaré duals of the homology classes of the cycles in the set $\{[\gamma^{(z)}] - [\gamma^{(z_0)}]\}_{z \in \mathbb{Y} \setminus z_0}$ generates the dual in $\text{Hom}(H_2(Y; \mathbb{Z}); \mathbb{Z})$ of the $H_2(M; \mathbb{Z})$ -summand in (1-7).

An M -adapted basis is characterized as follows: The basis contains the cycle $[\gamma^{(z_0)}]$, it contains a set of cycles that can be labeled $\{\hat{i}^{(z)}\}_{z \in \mathbb{Y} \setminus z_0}$, and it is rounded out by a set of G cycles that can be labeled $\{\hat{i}_p\}_{p \in \Lambda}$. A given $z \in \mathbb{Y} \setminus z_0$ version of $\hat{i}^{(z)}$ lies entirely in the $M_{7\delta_*}$ part of Y . It is homologous to $[\gamma^{(z)}] - [\gamma^{(z_0)}]$ and it is obtained from the latter by first truncating the \mathcal{H}_0 portions of the curves $\gamma^{(z)}$ and $\gamma^{(z_0)}$ and then reconnecting the respective endpoints by arcs on the boundary of the radius $7\delta_*$ coordinate balls about the index 0 and index 3 critical points of f . A given $p \in \Lambda$ version of \hat{i}_p is disjoint from the $f \in [1, 2]$ part of $M_{7\delta_*}$, and it intersects the rest of $M_{7\delta_*}$ and \mathcal{H}_0 as a smooth curve that is transverse to the level sets of f in M_δ and the constant u spheres in \mathcal{H}_0 ; the orientation is such that it has intersection number 1 with the $u = 0$ sphere in \mathcal{H}_0 . Meanwhile, \hat{i}_p intersects $\bigcup_{p' \in \Lambda} \mathcal{H}_{p'}$ as the $\theta = 0$ arc in \mathcal{H}_p , and the orientation is such that it has intersection number -1 with each $u = 0$ sphere in \mathcal{H}_p .

1E The Heegaard Floer equivalence

A three-part digression follows directly to set the notation for Theorem 1.1.

Part 1 Use (1-10) and (1-13) to write $\hat{\mathcal{Z}}_{\text{ech}, M}$ as

$$(1-20) \quad \hat{\mathcal{Z}}_{\text{ech}, M} = \mathcal{Z}_{\text{ech}, M} \times \mathbb{Z} = \mathcal{Z}_{\text{HF}} \times \left(\prod_{p \in \Lambda} (\mathbb{Z} \times \mathbb{O}) \right) \times \mathbb{Z}.$$

The principal \mathbb{Z} -bundle action of \mathbb{Z} acts on the right-most \mathbb{Z} . This factor is now moved next to the \mathcal{Z}_{HF} factor to write (1-20) as

$$(1-21) \quad \hat{\mathcal{Z}}_{\text{ech}, M} = (\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \times \left(\prod_{p \in \Lambda} (\mathbb{Z} \times \mathbb{O}) \right).$$

The identification in (1-21) induces the tensor product decomposition

$$(1-22) \quad \mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M}) = \mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes \left(\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathbb{O}) \right).$$

This representation of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ is used implicitly by Theorem 1.1.

Part 2 The factor $\mathbb{Z}(\mathcal{Z}_{\text{ech},M} \times \mathbb{Z})$ is the \mathbb{Z} -module for the Heegaard Floer homology on M . The endomorphism that supplies the differential for this homology is denoted by ∂_{HF} . Theorem 1.1 describes the differential ∂_{ech} on $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ in terms of ∂_{HF} and certain endomorphisms that are induced on (1-22) by corresponding endomorphisms of the various $\mathfrak{p} \in \Lambda$ factors of $\mathbb{Z}(\mathbb{Z} \times \mathcal{O})$. A given $\mathfrak{p} \in \Lambda$ version is denoted by $\partial_{\mathfrak{p}}$. All are the same endomorphism of $\mathbb{Z}(\mathbb{Z} \times \mathcal{O})$, this being the endomorphism ∂_* that acts on the generating set as follows:

- (1-23) • $\partial_*(k, 0) = 0$ for each $k \in \mathbb{Z}$.
- $\partial_*(k, 1) = (k, 0) + (k + 1, 0)$ for each $k \in \mathbb{Z}$.
 - $\partial_*(k, -1) = (k, 0) + (k - 1, 0)$ for each $k \in \mathbb{Z}$.
 - $\partial_*(k, \{1, -1\}) = (k, -1) + (k + 1, -1) - (k, 1) - (k - 1, 1)$ for each $k \in \mathbb{Z}$.

As noted by Lemma 2.5 in [8], the homology of the chain complex $(\mathbb{Z}(\mathbb{Z} \times \mathcal{O}), \partial_*)$ is $\mathbb{Z} \oplus \mathbb{Z}$, and generators are the closed elements $(0, 0)$ and $(0, 1) - (1, -1)$.

The various versions of Heegaard Floer homology enjoy an action of the algebra $\mathbb{Z}[\mathbb{U}] \otimes \wedge^*(H_1(M; \mathbb{Z})/\text{torsion})$ whose generators can be defined by endomorphisms of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M} \times \mathbb{Z})$. Note in this regard that the generator of the action of $\mathbb{Z}[\mathbb{U}]$ sends any given pair $(\widehat{v}, k) \in \mathcal{Z}_{\text{ech},M} \times \mathbb{Z}$ to $(\widehat{v}, k - 1)$.

There is one other Heegaard Floer endomorphism that plays a role in what follows. The latter is defined by its action on the generators, and in doing so, it acts solely on the \mathcal{Z}_{HF} and ignores the \mathbb{Z} factor. This is the endomorphism that appears in Theorem 4.1 of [13] and Definition 8.1 of [14]. The latter is denoted here by $\partial_{\text{HF}0}$.

The upcoming formula for ∂_{ech} and the other endomorphisms of (1-22) use the following convention: Suppose that E and E' are graded chain complexes and that Δ and Δ' are respective graded endomorphisms of E and E' . The latter induce on $E \otimes E'$ an endomorphism, $\Delta + \Delta'$, that is defined by the following action on the reducible elements: Let ϵ and ϵ' denote respective elements of E and E' . Then $(\Delta + \Delta')(\epsilon \otimes \epsilon')$ is defined to be $\Delta\epsilon \otimes \epsilon' + (-1)^{\text{degree}(\Delta') \text{degree}(\epsilon)} \epsilon \otimes \Delta'\epsilon'$.

Part 3 Let $p_M \in 2\mathbb{Z}$ again denote the greatest divisor of the image of $H_2(M; \mathbb{Z})$ in \mathbb{Z} via the pairing homomorphism with c_{1M} . As noted previously, the \mathbb{Z} -module $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ has a relative $\mathbb{Z}/p_M\mathbb{Z}$ grading that is induced by a relative grading of its generators. The grading difference between given generators $\widehat{\Theta}'$ and $\widehat{\Theta}$ is denoted in what follows by $\text{gr}_{\text{ech}}(\widehat{\Theta}') - \text{gr}_{\text{ech}}(\widehat{\Theta})$.

As explained in [13; 14], the \mathbb{Z} -module $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z})$ has a relative $\mathbb{Z}/(p_M \mathbb{Z})$ grading that is induced by a relative $\mathbb{Z}/(p_M \mathbb{Z})$ grading of the set \mathcal{Z}_{HF} . The difference between the respective gradings of given elements $\hat{v}', \hat{v} \in \mathcal{Z}_{\text{HF}}$ is denoted by $\text{deg}_{\text{HF}}(\hat{v}') - \text{deg}_{\text{HF}}(\hat{v})$. Granted this notation, the difference between the gradings of corresponding elements (\hat{v}', k') and (\hat{v}, k) from $\mathcal{Z}_{\text{HF}} \times \mathbb{Z}$ is $\text{deg}_{\text{HF}}(\hat{v}') - \text{deg}_{\text{HF}}(\hat{v}) + 2(k' - k)$.

The module $\mathbb{Z}(\mathbb{Z} \times \mathcal{O})$ has an absolute \mathbb{Z} grading with values in the 3-element set $\{0, 1, 2\}$. The latter grading is induced by a grading of the generators that depends only on the factor $\mathcal{O} = \{0, 1, -1, \{1, -1\}\}$: The element 0 has grading zero, the elements -1 and 1 have grading 1, and the element $\{-1, 1\}$ has grading 2. The resulting grading map from $\mathbb{Z} \times \mathcal{O}$ to $\{0, 1, 2\}$ is denoted by $\text{gr}_{\mathcal{O}}(\cdot)$.

With the preceding as background, what follows is this paper’s central result:

Theorem 1.1 *Identify $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ with $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes \left(\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathcal{O})\right)$ as in (1-22).*

- *The differential ∂_{ech} appears as $\partial_{\text{ech}} = \partial_{\text{HF}} + \sum_{p \in \Lambda} \partial_p$.*
- *The \mathbb{U} -map acts as the map $((\hat{v}, k), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda}) \rightarrow ((\hat{v}, k - 1), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$.*
- *Use the M -adapted 1-cycle basis $\{[\gamma^{z_0}], \{\hat{i}^{(z)}\}_{z \in \mathbb{Y} \setminus z_0}, \{\hat{i}_p\}_{p \in \Lambda}\}$ to define endomorphisms of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$.*
 - (a) *The endomorphism defined by $[\gamma^{z_0}]$ acts as $\partial_{\text{HF}} - \partial_{\text{HF}0}$.*
 - (b) *The endomorphisms defined by cycles from $\{\hat{i}^{(z)}\}_{z \in \mathbb{Y} \setminus z_0}$ act only on the $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z})$ factor. In doing so, they induce a set of generators of the $\wedge^*(H_1(M; \mathbb{Z})/\text{torsion})$ action on the Heegaard Floer homology.*
 - (c) *The endomorphism that is defined by any given $p \in \Lambda$ version of \hat{i}_p acts as $\mathbb{I}_p \otimes \hat{i}_p$, where \mathbb{I}_p acts as the identity on the factors*

$$\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes \left(\bigotimes_{p' \in \Lambda - p} \mathbb{Z}(\mathbb{Z} \times \mathcal{O}) \right)$$

and \hat{i}_p is the degree -1 endomorphism that acts only on p ’s factor of $\mathbb{Z}(\mathbb{Z} \times \mathcal{O})$. It acts on this factor as the endomorphism that sends $(\mathfrak{k}_p, \mathcal{O}_p)$ to $(\mathfrak{k}_p, \mathcal{O}'_p)$ with coefficient either 1 or 0. The coefficient 1 appears if and only if both $\mathcal{O}_p = 1$ and $\mathcal{O}'_p = 0$, or both $\mathcal{O}_p = \{1, -1\}$ and $\mathcal{O}'_p = -1$.

- *Let $\hat{\Theta}' = ((\hat{v}', k'), (\mathfrak{k}'_p, \mathcal{O}'_p)_{p \in \Lambda})$ and $\hat{\Theta} = ((\hat{v}, k), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$ denote any two elements. Then*

$$\text{gr}_{\text{ech}}(\hat{\Theta}') - \text{gr}_{\text{ech}}(\hat{\Theta}) = \text{gr}_{\text{HF}}(\hat{v}') - \text{gr}_{\text{HF}}(\hat{v}) + 2(k' - k) + \sum_{p \in \Lambda} (\text{gr}(\mathcal{O}'_p) - \text{gr}(\mathcal{O}_p)).$$

The subsequent sections in this article contain the proof of Theorem 1.1. The next subsection gives an indication of what the proof involves.

1F A look ahead at the proof

Three fundamental observations serve as the foundation for the proof of Theorem 1.1. The first is provided by Lipshitz [10] and his theorem to the effect that the differential for Heegaard Floer homology can be defined using certain sorts of pseudoholomorphic subvarieties that reside in the $f^{-1}(1, 2)$ part of $\mathbb{R} \times M_\delta$. These are described in Section II.6 and their properties are summarized in the next subsection. The second observation is supplied by Propositions II.7.2 and II.7.3. The latter assert that the $\mathbb{R} \times M_\delta$ part of any of the relevant ech-HF subvariety looks very much like a subvariety of the sort considered by Lipshitz. The third observation is jointly supplied by Propositions II.4.5 and II.5.8. These two propositions jointly hint at a canonical form for the $\mathbb{R} \times (\bigcup_{p \in \Lambda} \mathcal{H}_p)$ part of any given ech-HF subvariety. The subsequent proof of Theorem 1.1 uses this view of an ech-HF subvariety as the union of a Heegaard Floer looking $\mathbb{R} \times M_\delta$ part and a roughly canonical $\mathbb{R} \times (\bigcup_{p \in \Lambda} \mathcal{H}_p)$ part to derive the decomposition given by the first bullet of Theorem 1.1, and likewise to prove the assertions of the remaining bullets.

A proof of Theorem 1.1 along the preceding lines must address the following fundamental question:

Fix a subvariety from [10]. As noted above, there is some set of ech-HF subvarieties that look much like it on $\mathbb{R} \times M_\delta$. What can be said about this set; in particular, can enough be said to justify the claims of Theorem 1.1?

As it turns out, only submanifolds need be considered, and the upcoming sections study the question just posed with regards to submanifolds. This is done by constructing the appropriate set of ech-HF submanifolds from a given submanifold from [10]. The construction has two parts. The following two parts of this subsection say a few things about the two parts of the construction and how they lead to Theorem 1.1.

Part 1 The first part of the construction starts with a submanifold from [10] and a suitably compatible pair of elements $\widehat{\Theta}', \widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}$; it then uses this data to build a canonical approximation to what would be an ech-HF submanifold from the moduli space $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$. This approximation exploits the dichotomy between what

is said in Propositions II.4.5 and II.5.8 and what is said in Proposition II.7.2 about the $\mathbb{R} \times (\bigcup_p \mathcal{H}_p)$ and $\mathbb{R} \times M_\delta$ parts of an ech-HF submanifold. In particular, the approximation consists of a set $\mathcal{C}_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$ where C_{S_0} denotes a J -holomorphic submanifold with boundary in $\mathbb{R} \times M_\delta$, and where any given $p \in \Lambda$ version of C_{p_0} denotes a J -holomorphic submanifold with boundary in $\mathbb{R} \times \mathcal{H}_p$. The submanifold C_{S_0} looks very much like one of the subvarieties from [10], and each $p \in \Lambda$ version of C_{p_0} is described by Propositions II.4.5 and II.5.8. The submanifold C_{S_0} has $2G$ boundary components, one on a certain $f \in (1 + \delta^2, 1 + \delta_*^2)$ level set in each $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$ and the other on a certain $f \in (2 - \delta_*^2, 2 - \delta^2)$ level set in each $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$. Meanwhile, any given $p \in \Lambda$ version of C_{p_0} has two boundary components, one on each of these same level sets of f in $\mathbb{R} \times \mathcal{H}_p$. However, the boundary components of C_{p_0} need not agree with the corresponding C_{S_0} boundary components on the relevant level sets of f .

The set of such approximations to would-be ech-HF submanifolds can be used to define an ersatz version of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$. This ersatz version can then be used to define coefficients of endomorphisms of $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, M})$ using Hutchings' rules. To say a bit more, note that the definition of the endomorphism coefficients using honest ech-submanifolds is along standard symplectic field theory lines in the sense that a family of Fredholm operators and a certain tautological \mathbb{R} action play the central roles. The space $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ parametrizes the right sort of family, and the tautological \mathbb{R} action is induced by the constant translations along the \mathbb{R} factor of $\mathbb{R} \times Y$. Granted this remark about the definitions, what follows is a key point: The corresponding ersatz moduli space that is constructed from the canonical approximations to ech-HF submanifolds has an analogous family of Fredholm operators and an analogous action of \mathbb{R} . This being the case, Hutchings' rules can also be used with the ersatz moduli spaces to define endomorphisms $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech}, M})$. Meanwhile, the canonical nature of the construction guarantees that the resulting versions of the endomorphisms relevant to Theorem 1.1 satisfy the conclusions of Theorem 1.1.

Section 2 describes in detail the canonical approximations to ech-HF submanifolds; Sections 3–6 construct them.

Part 2 The second part of the construction builds a cobordism between the ersatz version of a given $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ and the version with honest ech-HF submanifolds. The cobordism maps to the interval $[0, 1]$ with the inverse image of 0 giving the ersatz moduli space and that of 1 giving the version with honest ech-submanifolds. The

cobordism defines a smooth manifold with boundary such that the map to $[0, 1]$ is proper and smooth. A key point here is that the relevant family of Fredholm operators extends across the cobordism, as does the relevant \mathbb{R} action. Given that the approximation versions of Theorem 1.1's endomorphisms obey Theorem 1.1's conclusions, these last facts are seen to imply that Theorem 1.1's conclusions must hold with its endomorphisms defined using honest ech-submanifolds.

To say a bit more about this cobordism, consider for a moment one of the approximation sets, $\mathcal{C}_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$. As noted in Part 1, its elements are J -holomorphic manifolds with boundary with the boundaries lying on certain level sets of f . The boundary of C_{S_0} is determined solely by the given subvariety from [10]. The boundary of any given $p \in \Lambda$ version of C_{p_0} is constrained in part by that of C_{S_0} . Keeping this in mind, let τ denote the parameter in $[0, 1]$. The inverse image of τ in the cobordism consists of a set of the form $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ where C_S is a J -holomorphic submanifold with boundary in $\mathbb{R} \times M_\delta$, and where each $p \in \Lambda$ version of C_p is a J -holomorphic submanifold with boundary in $\mathbb{R} \times \mathcal{H}_p$. The submanifold C_S has $2G$ boundary components, these on the aforementioned level sets of f in $\bigcup_{p \in \Lambda} (\mathbb{R} \times \mathcal{H}_p)$. Meanwhile, each $p \in \Lambda$ version of C_p has two, one each on the $\mathbb{R} \times \mathcal{H}_p$ parts of these level sets. The parameter τ indicates the extent to which the two boundary components of any given $p \in \Lambda$ version of C_p agree with the relevant pair of boundary components of C_S . In the case when $\tau = 1$, they match up and so define an honest ech-subvariety. This is not necessarily true for $\tau < 1$.

By way of a hint as to the nature of the family of Fredholm operators, the operator for a given $\tau \in [0, 1]$ version of $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ is viewed as a set of $G + 1$ operators, with one defined by C_S and one by each $p \in \Lambda$ version of C_p . Thus, each is defined on a manifold with boundary and, as such, its definition requires the specification of some boundary conditions. These are local (as opposed to spectral) boundary conditions that couple the C_S operator to those defined by the various $p \in \Lambda$ versions of C_p so as to associate to \mathcal{C} a single Fredholm operator. The parameter τ determines the degree of coupling.

Section 7 constructs the cobordism space that interpolates between the ersatz moduli space and the space of ech-HF submanifolds.

Section 8 supplies the background needed to use the cobordism to compute the differential and other endomorphisms that appear in Theorem 1.1. Section 9 uses the properties of the cobordism space to complete the proof of Theorem 1.1.

1G The subvarieties used by Lipshitz

This subsection summarizes some of what is said in Section II.6 about the subvarieties that are used by Lipshitz in [10]. These are subvarieties in the $f \in (1, 2)$ part of $\mathbb{R} \times M$ that are best described by viewing this part of M as $(1, 2) \times \Sigma$ via the identification given by Lie transport with the pseudogradient vector field v .

The relevant subvarieties in $\mathbb{R} \times (1, 2) \times \Sigma$ are pseudoholomorphic for an almost complex structure with certain special properties. These are described in Section 1 of [10]. Let J denote an almost complex structure on $\mathbb{R} \times Y$ that obeys the constraints in Part 1 of Section 1C. The restriction of J to the $f \in (1, 2)$ part of $\mathbb{R} \times M_\delta$ can be extended to the whole of $\mathbb{R} \times (1, 2) \times \Sigma$ so as to give an almost complex structure of the sort considered by Lipshitz and, in particular, of the sort that is described by (II.6-1). Conversely, a suitably generic almost complex structure on $\mathbb{R} \times (1, 2) \times \Sigma$ that obeys (II.6-1) will serve for Lipshitz. Moreover, such an almost complex structure will restrict to the $\mathbb{R} \times M_\delta$ part of $\mathbb{R} \times (1, 2) \times \Sigma$ as the restriction of an almost complex structures on $\mathbb{R} \times Y$ that obeys the constraints in Section 1C. This understood, let J denote an almost complex structure on $\mathbb{R} \times Y$ that obeys the constraints in Part 1 of Section 1C and let J_{HF} denote an almost complex structure on $\mathbb{R} \times (1, 2) \times \Sigma$ that obeys (II.6-1). To say more about J_{HF} , note that the stable Hamiltonian 2-form ω appears on the M_δ part of $(1, 2) \times \Sigma$ as the pullback via the projection of an area form on Σ . Denote the latter by ω_Σ . This form extends in the obvious way to the whole of $(1, 2) \times \Sigma$. Let t denote the Euclidean coordinate on $(1, 2)$. The almost complex structure J_{HF} maps ∂_s to ∂_t , it preserves the level sets of t and it is compatible with the symplectic form $ds \wedge dt + \omega_\Sigma$. It also commutes with the ∂_s Lie derivative and it obeys the constraint given by the sixth bullet in Part 1 of Section 1C.

Lipshitz considers J_{HF} -holomorphic subvarieties in $\mathbb{R} \times (1, 2) \times \Sigma$ with eight special properties that are listed in a moment. For the purposes at hand, it is sufficient to consider the case where the subvariety in question is a smooth submanifold. The closure in $\mathbb{R} \times [1, 2] \times \Sigma$ of a submanifold with these properties is said here to be a *Lipshitz submanifold*. Let S_0 denote the interior of a Lipshitz submanifold.

Property 1 *The integral over S_0 of ω_Σ is finite. This is also the case for the integral of $ds \wedge dt$ over any subset of S_0 with bounded image in the \mathbb{R} factor of $\mathbb{R} \times (1, 2) \times \Sigma$.*

The second property refers to the G circles in Σ that form the latter's intersection with the ascending disks from the index 1 critical points of f , and the corresponding set of

G circles that form Σ 's intersection with the descending disks from the index 2 critical points of f . If p is a given index 1 or index 2 critical point of f , the corresponding circle is denoted by C_{p+} or C_{p-} as the case may be. The union of the index 1 critical point versions of C_{p+} is denoted by C_+ and the union of the index 2 critical point versions is denoted by C_- .

Property 2 *The J_{HF} -holomorphic submanifold S_0 is the interior of a properly embedded surface $\mathbb{R} \times [1, 2] \times \Sigma$ with $2G$ boundary components. Half of the boundary components are in $\mathbb{R} \times \{1\} \times C_+$ and no two of these lie in the same component. The other half are in $\mathbb{R} \times \{2\} \times C_-$ and, likewise, no two are in the same component.*

The surface with boundary in Property 2 is denoted by S . If p is an index 1 or index 2 critical point of f , then the corresponding boundary component of S is denoted by $\partial_p S$. It is a properly embedded copy of \mathbb{R} in $\mathbb{R} \times \{1\} \times C_{p+}$ or $\mathbb{R} \times \{2\} \times C_{p-}$ as the case may be.

With regard to notation, Lipshitz and also Section II.6 view what is denoted here by S as the image of a complex surface via a J_{HF} -holomorphic map, u . What is denoted by S here is denoted in [10] and in Section II.6 by $u(S)$.

The third property refers to elements from the set Z_{HF} . By way of a reminder, an element of Z_{HF} consists of a set of G integral curves of v with each starting at an index 1 critical point of f and ending at an index 2 critical point of f . Moreover, distinct curves from such a set have distinct start points and distinct endpoints.

Property 3 *The surface S is the complement of $2G$ points in a compact surface with boundary. The function s on S increases with no finite limit on sequences that limit to G of these points, and it decreases with no finite limit on sequences that limit to the remaining G points.*

This surface with boundary in question is denoted by \underline{S} . The G points of $\underline{S} \setminus S$ with neighborhoods where s is unbounded from above are said to be *positive points*, and the remaining points are said to be *negative points*.

The remaining properties of S are all consequences of the first three. The next two properties restate Lemmas II.6.2 and II.6.3. They refer to the two sets of G circles, C_- and C_+ , in Σ that are described in Section 1C. They also refer to their respective annular neighborhoods, T_- and T_+ , and the coordinate functions (φ_-, h_-) and (φ_+, h_+) on

these annular neighborhoods. These are also described in Section 1C. (The original definitions are in Parts 2 and 3 of Section II.2C). By way of a reminder, the $h_- = 0$ locus in T_- is C_- , and the $h_+ = 0$ locus in T_+ is C_+ .

Property 4 *Let p denote either an index 1 or index 2 critical point of f . The closure in \underline{S} of the corresponding boundary component $\partial_p S$ adds one positive point and one negative point from $\underline{S} \setminus S$. Meanwhile, $\partial_p S$ appears in $\mathbb{R} \times C_{p+}$ or $\mathbb{R} \times C_{p-}$ as a graph over \mathbb{R} of the form $x \rightarrow (s = x, \varphi_+ = \varphi^{S,p}(x))$ or $x \rightarrow (s = x, \varphi_- = \varphi^{S,p}(x))$ as the case may be. In either case, $\varphi^{S,p}: \mathbb{R} \rightarrow \mathbb{R}$ is a smooth map with bounded derivatives to any given order. Moreover, the $x \rightarrow \pm\infty$ limits of $\varphi^{S,p}$ exist and both are in $C_+ \cap C_-$.*

The next property describes the behavior of S near any given boundary component.

Property 5 *There exists $z_S > 0$ and $\kappa_S > 1$ with the following significance: Let p denote either an index 1 or index 2 critical point of f . Then a neighborhood of $\partial_p S$ in S appears as the image of a map from $\mathbb{R} \times (0, z_S)$ to $\mathbb{R} \times (1, 1 + z_S) \times T_{p+}$ or $\mathbb{R} \times (2 - z_S, 2) \times T_{p-}$ as the case may be. This map has the form*

- $(x, z) \rightarrow (s = x, t = 1 + z, \varphi_+ = \varphi(x, z), h_+ = \zeta(x, z))$ when p has index 1,
- $(x, z) \rightarrow (s = x, t = 2 - z, \varphi_- = \varphi(s, z), h_- = \zeta(s, z))$ 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

- $|\zeta(x, z)| + z^{-1} |\varphi(x, z) - \varphi^{S,p}(x)| < \kappa_S z,$
- $|\partial_x \zeta(s, z)| + z^{-1} |(\partial_x \varphi)(s, z) - (\partial_x \varphi^{S,p})(x)| < \kappa_S z.$

In addition, the pair φ and ζ have bounded derivatives to any given order on $\mathbb{R} \times [0, z_S)$.

The next property says more about the large $|s|$ part of S .

Property 6 *There exists $\kappa_S > 1$ such that the $s \leq -\kappa_S$ and $s \geq \kappa_S$ portions of S are disjoint unions of G half open rectangles. Those where $s \leq -\kappa_S$ are properly embedded submanifolds in $(-\infty, \kappa_S] \times [1, 2] \times (T_- \cap T_+)$ that appear as a graph over $(-\infty, \kappa_S] \times [1, 2]$ of a map to a component of $T_- \cap T_+$ with the following properties: Let q denote the point in $C_- \cap C_+$ that lies in the given component and let ψ denote the map. Then $\text{dist}(\psi, q) \leq e^{-|s|/\kappa_S}$. The derivatives of ψ to any given order are also bounded by a constant times this same exponential factor. Meanwhile, the components of the $s \geq \kappa_S$ have an analogous description as a graph over $[\kappa_S, \infty) \times [1, 2]$.*

This property and Property 2 lead directly to the next property.

Property 7 *The set of constant $s \in \mathbb{R}$ slices of S when viewed in M converge pointwise as $s \rightarrow \infty$ to define an element in \mathcal{Z}_{HF} . This is also true of the $s \rightarrow -\infty$ limit of the constant s slices.*

These two elements in \mathcal{Z}_{HF} are denoted respectively by \hat{v}_+ and \hat{v}_- .

The set of Lipshitz submanifolds is given the topology whereby an open neighborhood of a given subvariety S has a basis of open sets labeled by the positive numbers. The open set labeled by $\varepsilon \in (0, \infty)$ is characterized as follows: A submanifold S' is a member when:

$$(1-24) \quad \begin{aligned} & \bullet \sup_{z \in S} \text{dist}(z, S') + \sup_{z \in S'} \text{dist}(S, z) < \varepsilon. \\ & \bullet \text{ Let } \mu \text{ denote a smooth 2-form on } [-\frac{1}{\varepsilon}, \frac{1}{\varepsilon}] \times [1, 2] \times \Sigma \text{ with compact support,} \\ & \quad \text{with supremum norm 1 and with } |\nabla \mu| \leq \frac{1}{\varepsilon}. \text{ Then } |\int_{S'} \mu - \int_S \mu| \leq \varepsilon. \end{aligned}$$

The resulting topological space is denoted by \mathcal{A}_{HF} .

The group \mathbb{R} acts continuously on \mathcal{A}_{HF} via its action on $\mathbb{R} \times (1, 2) \times \Sigma$ as the group of constant translations along the \mathbb{R} factor. This \mathbb{R} action is free on the complement of the set $\{\mathbb{R} \times (\bigcup_{\hat{v} \in \mathcal{Z}_{\text{HF}}} \nu)\}_{\hat{v} \in \mathcal{Z}_{\text{HF}}}$ of 1-point components \mathcal{A}_{HF} . Lemmas II.6.6 and II.6.7 say more about the structure of \mathcal{A}_{HF} . These lemmas refer to a certain \mathbb{R} -linear, Fredholm incarnation of the $\bar{\partial}$ -operator that is canonically associated to any given Lipshitz subvariety. The relevant operator is described in Section II.6E for the case when the variety in question is a submanifold. The operator for a Lipshitz submanifold S is denoted by D_S . Let $N_S \rightarrow S$ denote the complex normal bundle of S , with the complex structure defined by J_{HF} and with the hermitian structure and thus holomorphic structure defined by the metric $w_\Sigma(\cdot, J_{\text{HF}}(\cdot))$. Let $T^{0,1}S$ denote the $(0, 1)$ part of $T^*S \otimes \mathbb{C}$. The operator D_S maps sections of N_S to sections of $N_S \otimes T^{0,1}S$ by the rule

$$(1-25) \quad \eta \rightarrow D_S \eta = \bar{\partial} \eta + v \eta + \mu \bar{\eta},$$

where v denotes a certain section of $T^{0,1}S$ and μ denotes a section of $N_S^2 \otimes T^{0,1}S$.

To describe the Fredholm domain of D_S , note that Property 4 can be used as in Section II.6E to identify the bundle N_S along the boundary of S with $T\Sigma$ along C_+ and C_- . Then, the Fredholm domain is the L_1^2 completion of the space of compactly supported sections of N_S that obey

$$(1-26) \quad \begin{aligned} & \bullet \eta \in TC_+ \text{ on the } t = 1 \text{ part of the boundary of } S, \\ & \bullet \eta \in TC_- \text{ on the } t = 2 \text{ part of the boundary of } S. \end{aligned}$$

Note in this regard that the definition of the L_1^2 -norm uses the metric $w_\Sigma(\cdot, J_{\text{HF}}(\cdot))$ to define the integration measure, inner products and covariant derivative on all tensor bundles constructed from N_S and TS . Meanwhile, the range space for this Fredholm incarnation of D_S is the space of square-integrable sections of $N_S \otimes T^{0,1}S$.

The final property speaks of this operator D_S .

Property 8 *The operator D_S has trivial cokernel.*

Given this last property, it follows from Lemma II.6.7 that the subspace of Lipschitz submanifolds in \mathcal{A}_{HF} has the structure of a smooth manifold whose dimension near any given submanifold S is the Fredholm index of D_S .

1H Coordinates for the $1 - 3 \cos^2 \theta > 0$ part of $\mathbb{R} \times \mathcal{H}_p$

The upcoming construction of ech-HF submanifolds exploits the parametrization of the $1 - 3 \cos^2 \theta > 0$ part each $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$ from Part 1 of Section II.4C. The parametrization is denoted by Ψ_p . The three parts of this subsection that follow define Ψ_p and list some of its important features.

Part 1 Fix $p \in \Lambda$. The upcoming description of Ψ_p requires introducing the coordinates (u, θ, ϕ) for \mathcal{H}_p and the function h of the variables u and θ given by

$$(1-27) \quad h = f(u) \cos \theta \sin^2 \theta,$$

with f as defined in (1-4). The 1-form dh is nowhere zero where $1 - 3 \cos^2 \theta > 0$. This function is also annihilated by the vector field v and so it has constant value along v 's integral curves in \mathcal{H}_p .

The definition of Ψ_p also involves the J -holomorphic submanifolds from Proposition II.3.2's space M_Σ and Proposition II.3.4's space \mathcal{M}_{p_0} . By way of a reminder, the space \mathcal{M}_Σ is \mathbb{R} -equivariantly diffeomorphic to $\mathbb{R} \times (1, 2)$. Each element is a compact submanifold that is diffeomorphic to Σ . A given $(s, t) \in \mathbb{R} \times [1 + \delta^2, 2 - \delta^2]$ element is the $(s = s, f = t)$ slice of $\mathbb{R} \times M_\delta$. An element parametrized by $\mathbb{R} \times (1, 1 + 7\delta_*^2)$ intersects the $u > 0$ portion of each $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$ where $1 - 3 \cos^2 \theta > 0$, and an element in the $\mathbb{R} \times (2 - 7\delta_*^2, 2)$ part of \mathcal{M}_Σ intersects the $u < 0$ and $1 - 3 \cos^2 \theta > 0$ part of each $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_p$. In each case, the intersection is a properly embedded annulus that can be parametrized by the functions ϕ and h . This parametrization is such that the range of h is symmetric with respect to multiplication by -1 , and such

that functions s and u restrict as ± 1 symmetric functions of \hat{h} . Thus, each constant \hat{h} slice of the annulus is a circle in some constant (s, u) sphere in $\mathbb{R} \times \mathcal{H}_p$. Taken together, these annuli from \mathcal{M}_Σ foliate the part of $\mathbb{R} \times \mathcal{H}_p$ where $u \neq 0$ and $1 - 3 \cos^2 \theta > 0$.

The space \mathcal{M}_{p_0} is \mathbb{R} -equivariantly diffeomorphic to \mathbb{R} . An element in the space \mathcal{M}_{p_0} is a properly embedded annulus in the part of $\mathbb{R} \times \mathcal{H}_p$ where $1 - 3 \cos^2 \theta > 0$ and $u = 0$. The pair (ϕ, \hat{h}) restrict as coordinates to this annulus such that \hat{h} defines a proper map to $(-(\chi_0 + 2e^{-2R})\frac{2}{3\sqrt{3}}, (\chi_0 + 2e^{-2R})\frac{2}{3\sqrt{3}})$. The coordinate s on the annulus is a symmetric function of \hat{h} that is unbounded from above on both ends of its domain. The member parametrized by $0 \in \mathbb{R}$ intersects the $(s = 0, u = 0)$ slice of $\mathbb{R} \times \mathcal{H}_p^+$ as the $\theta = \frac{\pi}{2}$ circle in S^2 . The annuli from \mathcal{M}_{p_0} foliate the $1 - 3 \cos^2 \theta > 0$ part of the $u = 0$ slice of $\mathbb{R} \times \mathcal{H}_p$.

Part 2 Introduce \mathcal{H}_p^+ to denote the $1 - 3 \cos^2 \theta > 0$ and $|u| < R + \ln \delta_*$ part of \mathcal{H}_p . The inverse of Ψ_p is an \mathbb{R} -equivariant embedding of $\mathbb{R} \times \mathcal{H}_p^+$ into

$$(1-28) \quad \mathbb{R} \times (-R - \ln \delta_*, R + \ln \delta_*) \times (\mathbb{R}/2\pi\mathbb{Z}) \times \left(-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2\right).$$

The image is denoted by $\mathbb{R} \times \mathcal{X}$. The coordinate functions on the space depicted in (1-28) and thus on $\mathbb{R} \times \mathcal{X}$ are denoted by $(x, \hat{u}, \hat{\phi}, h)$. The rules that follow define Ψ_p :

- (1-29) • Ψ_p sends a given $(x, \hat{u} \neq 0, \hat{\phi}, h)$ point in $\mathbb{R} \times \mathcal{X}$ to the $\phi = \hat{\phi}, \hat{h} = h$ point on the subvariety from \mathcal{M}_Σ that intersects the $\theta = \frac{\pi}{2}$ slice of $\mathbb{R} \times \mathcal{H}_p^+$ where $s = x$ and $u = \hat{u}$.
- Ψ_p sends a given $(x, \hat{u} = 0, \hat{\phi}, h)$ point in $\mathbb{R} \times \mathcal{X}$ to the $\phi = \hat{\phi}, \hat{h} = h$ point on the subvariety from \mathcal{M}_{p_0} that intersects the $\theta = \frac{\pi}{2}$ slice of $\mathbb{R} \times \mathcal{H}_p^+$ where $s = x$ and $u = 0$.

Formulas for the Ψ_p -pushforwards of the coordinate vector fields $\partial_x, \partial_{\hat{u}}, \partial_{\hat{\phi}}$ and ∂_h as given in (II.4-4) can be written as

$$(1-30) \quad \begin{aligned} &\bullet \Psi_{p*} \partial_x = \partial_s, \\ &\bullet \Psi_{p*} \partial_{\hat{u}} = v(v + \alpha^{-1} \sqrt{6} \chi \cos \theta \partial_\phi + \varpi \partial_s), \\ &\bullet \Psi_{p*} \partial_{\hat{\phi}} = \partial_\phi, \\ &\bullet \Psi_{p*} \partial_h = -\beta^{-1}(e_1 - \sigma^{-1} \sqrt{6} \chi_\delta f \cos \theta \sin^2 \theta \partial_s). \end{aligned}$$

Here, α is from (1-6), both v and β are certain positive functions of the pair (u, θ) , the function ϖ depends only on (u, θ) , the function σ is from the fifth bullet in Part 1 of Section 1C, and the vector field e_1 is defined by (1-14).

Part 3 The definition just given endows $\Psi_{\mathfrak{p}}$ with the properties listed in a moment. Let p denote the index 1 critical point of f from \mathfrak{p} and let p' the corresponding index 2 critical point. The list refers to the annuli T_{p+} and $T_{p'}$ in Σ and their respective coordinates (φ_+, h_+) and (φ_-, h_-) . These are introduced in Part 1 of Section 1C. The list also writes the $f \in (1, 2)$ part of $\mathbb{R} \times M_\delta$ as a subset of $\mathbb{R} \times (1, 2) \times \Sigma$, and it uses (s, t) to denote the Euclidean coordinates on $\mathbb{R} \times (1, 2)$. What follows next is the promised list:

- (1-31) • The constant (x, \hat{u}) surfaces in $\mathbb{R} \times \mathcal{X}$ are mapped by $\Psi_{\mathfrak{p}}$ to J -holomorphic submanifolds.
- The map $\Psi_{\mathfrak{p}}$ is equivariant with respect to the \mathbb{R} actions on $\mathbb{R} \times \mathcal{X}$ and $\mathbb{R} \times \mathcal{H}_{\mathfrak{p}}^+$ along their \mathbb{R} factors.
 - The map $\Psi_{\mathfrak{p}}$ 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 $\mathbb{R} \times \mathcal{H}_{\mathfrak{p}}^+$.
 - The $\hat{u} \geq \mathbb{R} + \ln \delta$ part of \mathcal{X} is

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

and $\Psi_{\mathfrak{p}}$ maps this part of $\mathbb{R} \times \mathcal{X}$ diffeomorphically onto

$$\mathbb{R} \times [1 + \delta^2, 1 + \delta_*^2] \times T_{p+}$$

by the rule

$$(x, \hat{u}, \hat{\phi}, h) \rightarrow (s = x, t = e^{2(\hat{u}-R)}, \varphi_+ = \hat{\phi}, h_+ = h).$$

- The $\hat{u} \leq -R - \ln \delta$ part of \mathcal{X} is

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

and $\Psi_{\mathfrak{p}}$ maps this part of $\mathbb{R} \times \mathcal{X}$ diffeomorphically onto

$$\mathbb{R} \times (2 - \delta_*^2, 2 - \delta^2) \times T_{p-}$$

by the rule

$$(x, \hat{u}, \hat{\phi}, h) \rightarrow (s = x, t = e^{-2(\hat{u}+R)}, \varphi_+ = \hat{\phi}, h_+ = -h).$$

The first three bullets of this list reproduce (II.4-3).

2 The approximations

As noted in Part 1 of Section 1F, each Lipschitz submanifold can be used to construct a corresponding set of approximations to ersatz ech-HF submanifolds. The set in question is parametrized by a subset $\widehat{\mathcal{Z}}^S \subset \widehat{\mathcal{Z}}_{\text{ech},M} \times \widehat{\mathcal{Z}}_{\text{ech},M}$ which is invariant with respect to the diagonal action of \mathbb{Z} . Each element in $\widehat{\mathcal{Z}}^S$ determines a corresponding ersatz ech-HF subvariety, this being a collection of $G + 1$ submanifolds with boundary that is denoted by $\mathcal{C}_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$. By way of a reminder, C_{S_0} is a submanifold with boundary in the $f \in (1, 2)$ part of $\mathbb{R} \times M_\delta$ and each $p \in \Lambda$ version of C_{p_0} is a submanifold with boundary in $\mathbb{R} \times \mathcal{H}_p^+$. The upcoming Section 2B defines $\widehat{\mathcal{Z}}^S$, and the remaining subsections describe the ersatz ech-HF submanifold that is associated to any given element in $\widehat{\mathcal{Z}}^S$. Section 2A describes the data needed to construct this association.

2A The parameters (δ, χ_0, R) and a new parameter, z_*

Section 1F does not mention one important point: The desired set of ersatz ech-HF submanifolds can be constructed from a given Lipschitz submanifold only if the parameter δ from the data set (δ, χ_0, R) that defines Y and its stable Hamiltonian geometry is sufficiently small. In particular, the chosen almost complex structure J_{HF} and the orbit in $\mathcal{A}_{\text{HF}}/\mathbb{R}$ of the chosen Lipschitz submanifold jointly determine an upper bound on δ . As noted in Section 1A, the latter determines an upper bound for χ_0 , and then χ_0 determines an upper bound for R .

The construction of the ersatz set of ech-HF submanifolds requires the specification of an additional parameter, this denoted by z_* . This z_* is a positive number whose maximum allowed value is determined by the orbit in $\mathcal{A}_{\text{HF}}/\mathbb{R}$ of the chosen Lipschitz submanifold. In any event, z_* is less than $e^{-32} \delta_*^2$. The choice of z_* must be made prior to choosing δ since the constructions that follow require the maximum allowed value for δ be less than $e^{-16} z_*^{1/2}$.

Additional refinements for allowed maximum of z_* and δ are stated as they are needed in the constructions to come. In any event, both are determined solely by the \mathbb{R} -orbit of the chosen Lipschitz surface. What follows are further comments on this issue that are of specific concern with regards to the proof of Theorem 1.1.

The first point is perhaps self-evident: a given finite set in $\mathcal{A}_{\text{HF}}/\mathbb{R}$ determines maximum values for z_* and δ such that the constructions to follow can be made using any Lipschitz surface from this chosen set of \mathbb{R} -orbits if z_* and δ are less than their allowed maxima.

As it turns out, the set of \mathbb{R} -orbits need not be finite to obtain this same conclusion. Saying more requires a digression to introduce the notion of a *weakly compact set*. Let $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$ denote an \mathbb{R} -invariant set of submanifolds. This set is said to be weakly compact when the following two requirements are met: First, integration of the 2-form w_Σ over the Lipschitz surfaces maps \mathcal{K} to a bounded subset in \mathbb{R} . Here is the second requirement: Fix any sequence from \mathcal{K} and there exists a surface $S \in \mathcal{K}$ and subsequence of the given sequence which, after renumbering as $\{S_n\}_{n=1,2,\dots}$, obeys

- $\sup_{z \in S \cap \left[-\frac{1}{n}, \frac{1}{n}\right] \times [1, 2] \times \Sigma} \text{dist}(z, S_n) + \sup_{z \in S_n \cap \left[-\frac{1}{n}, \frac{1}{n}\right] \times [1, 2] \times \Sigma} \text{dist}(S, z) < \frac{1}{n}$.
- Let μ denote a smooth 2-form on $\left[-\frac{1}{n}, \frac{1}{n}\right] \times [1, 2] \times \Sigma$ with compact support, with supremum norm 1 and with $|\nabla\mu| \leq \frac{1}{n}$. Then $\left| \int_S \mu - \int_{S_n} \mu \right| \leq \frac{1}{n}$.

Suppose \mathcal{K} is a given weakly compact subset of Lipschitz submanifolds. As it turns out, maximum values of z_* and δ can be chosen so that the constructions to come can be done using any surface from \mathcal{K} and values for z_* and δ that are less than these \mathcal{K} -dependent maxima. A data set $D = (z_*, \delta, \chi_0, R)$ that can be used for all Lipschitz submanifolds in \mathcal{K} is said in what follows to be \mathcal{K} -compatible. Note that the applications to the proof of Theorem 1.1 require only finite sets of \mathbb{R} -orbits of Lipschitz submanifolds.

The proof that \mathcal{K} -compatible data sets exist is a straightforward affair given how the maxima for z_* and δ are subsequently determined from any given Lipschitz submanifold. The proof is left to the reader save for what is said in the two parts that follow.

Part 1 Fix a Lipschitz submanifold S . The upper bounds for z_* and δ are determined by certain data that can be associated to S . The first two elements of this data set come from Property 5 of Section 1G. These are the constants z_S and κ_S . In particular, z_* is constrained to be less than $e^{-32} z_S$. The third element also comes via Property 5 of Section 1G. This is a bound for the C^6 -norm over any length 1 interval in \mathbb{R} of any index 1 and index 2 critical point version of the map $\varphi^{S,p}$. The data set also contains the C^6 -norms of the intrinsic and extrinsic curvatures of S , and a maximum for the allowed diameter of a tubular neighborhood of S in $\mathbb{R} \times [1, 2] \times \Sigma$.

The final element in the data set is a norm for a certain inverse of the operator D_S . To say more about what this means, recall from Section II.6E that D_S defines an \mathbb{R} -linear, Fredholm map from a certain Hilbert subspace of L^2_1 sections of the complex normal bundle of S to the L^2 Hilbert space of sections of the tensor product of this normal bundle with the $(0, 1)$ cotangent bundle of S . Property 5 of Section 1G implies that

this map is surjective. As a consequence, the operator D_S has an inverse that maps the range Hilbert space to the L^2 -orthogonal complement in the domain Hilbert space of the kernel of D_S . The latter map is continuous and so bounded; it is the desired inverse.

Part 2 The a priori bound on the integral of w_Σ over the submanifolds that form \mathcal{K} implies that all of the data listed in Part 1 lie in a compact set. The point being that the convergence criteria for membership in \mathcal{K} and the bound on the integral of w_Σ implies that the subsequence $\{S_n\}_{n=1,2,\dots}$ converges to the submanifold S in the C^∞ topology on compact subsets of $\mathbb{R} \times [1, 2] \times \Sigma$. Note in this regard that the constraint on J_{HF} given by the sixth bullet in Part 1 of Section 1C has the following consequence: when written as in the proof of Lemma II.6.3, the equations that define a Lipschitz subvariety in $\mathbb{R} \times (1, 2) \times \Sigma$ are \mathbb{C} -linear equations on neighborhoods of $\mathbb{R} \times \{1\} \times T_+$ and $\mathbb{R} \times \{2\} \times T_-$. This linearity is exploited in Part 3 of Section II.6B. In particular, only slight modifications to the arguments used in the proof of Lemma II.6.3 establish C^∞ convergence for the parts of $\{S_n\}$ near the boundary of $\mathbb{R} \times [1, 2] \times \Sigma$. Meanwhile, the C^∞ convergence in the interior of $\mathbb{R} \times [1, 2] \times \Sigma$ is proved using standard arguments about sequences of pseudoholomorphic curves. See for example [11].

The C^∞ convergence on compact subsets of $\mathbb{R} \times (1, 2) \times \Sigma$ to a surface in \mathcal{K} implies the desired a priori bound on all but one element of any $S \in \mathcal{K}$ version of the data set given in Part 1. The one element missing is the norm of the inverse of D_S . The needed bound on this norm can be derived using (II.6-15)–(II.6-17) to deal with the case when $\{S_n\}_{n=1,2,\dots}$ does not converge pointwise on the whole of $\mathbb{R} \times [1, 2] \times \Sigma$ to its limit. In the latter case the large n versions of S_n will have long, nearly \mathbb{R} -invariant cylinders (a consequence of Lemma II.5.6). Even so, (II.6-15)–(II.6-17) supply an $S \in \mathcal{K}$ -independent constant c_0 such that $\|D_S \eta\|_{L^2} \geq c_0^{-1} \|\eta\|_{L^2_1}$ if η is in the domain of D_S and has compact support on such a cylinder. This sort of bound plus the C^∞ convergence on compact subsets of $\mathbb{R} \times [1, 2] \times \Sigma$ implies the desired $S \in \mathcal{K}$ -independent bound on the norm of the inverse of D_S .

2B The set $\widehat{\mathcal{Z}}^S$

Fix a Lipschitz submanifold, S . As noted at the outset, the set of ersatz ech-HF submanifolds that are constructed from S is indexed by a subset $\widehat{\mathcal{Z}}^S \subset \widehat{\mathcal{Z}}_{ech, M} \times \widehat{\mathcal{Z}}_{ech, M}$. The two parts of this subsection describe the set $\widehat{\mathcal{Z}}^S$.

Part 1 The diagonal action of \mathbb{Z} on $\widehat{\mathcal{Z}}_{\text{ech},M} \times \widehat{\mathcal{Z}}_{\text{ech},M}$ preserves $\widehat{\mathcal{Z}}^S$ and gives it the structure of a principal \mathbb{Z} -bundle over its image in $\mathcal{Z}_{\text{HF}} \times \mathcal{Z}_{\text{HF}}$. The image is described in the upcoming Part 2 of this subsection. To say more about the fiber over this image, first write $\widehat{\mathcal{Z}}_{\text{ech},M}$ using (1-13) as $\mathcal{Z}_{\text{ech},M} \times \mathbb{Z}$. Introduce next n_S to denote the intersection number between S and the J_{HF} -holomorphic subvariety $\mathbb{R} \times (1, 2) \times z_0$. This is a nonnegative integer. A given element $((\Theta_-, k_-), (\Theta_+, k_+))$ from $\widehat{\mathcal{Z}}_{\text{ech},M} \times \widehat{\mathcal{Z}}_{\text{ech},M}$ sits in $\widehat{\mathcal{Z}}^S$ only if $k_+ = k_- + n_S$.

Part 2 Let \widehat{v}_- and $\widehat{v}_+ \in \mathcal{Z}_{\text{HF}}$ denote the elements that are defined by S as described in Property 6 of Section 1G. These sets define via (1-10) a corresponding set of elements in $\mathcal{Z}_{\text{ech},M}$, thus a subset in $\mathcal{Z}_{\text{HF}} \times \mathcal{Z}_{\text{HF}}$ of the form

$$(2-1) \quad \left(\widehat{v}_- \times \left(\prod_{p \in \Lambda} (\mathbb{Z} \times \mathcal{O}) \right) \right) \times \left(\widehat{v}_+ \times \left(\prod_{p \in \Lambda} (\mathbb{Z} \times \mathcal{O}) \right) \right).$$

The set $\widehat{\mathcal{Z}}^S$ will sit over a subset in (2-1). The latter is denoted in what follows by \mathcal{Z}^S .

The elements in (2-1) that lie in \mathcal{Z}^S are characterized by G conditions, one for each $p \in \Lambda$. As explained in a moment, a given pair $(\widehat{v}_-, (\mathfrak{k}_{p_-}, \mathcal{O}_{p_-})_{p \in \Lambda})$ and $(\widehat{v}_+, (\mathfrak{k}_{p_+}, \mathcal{O}_{p_+})_{p \in \Lambda})$ from (2-1) defines a corresponding set of integers, this denoted by $\{\mathfrak{m}_p\}_{p \in \Lambda}$. The given pair defines an element in \mathcal{Z}^S if and only if each $p \in \Lambda$ version of \mathfrak{m}_p , \mathcal{O}_{p_-} and \mathcal{O}_{p_+} obeys one of

- (2-2) • $\mathfrak{m}_p = 0$ and one of
- (a) $\mathcal{O}_{p_-} = \mathcal{O}_{p_+} = \{0\}$, or
 - (b) $\mathcal{O}_{p_-} = \{0\}$ and $\mathcal{O}_{p_+} = \{-1, 1\}$;
 - $\mathfrak{m}_p = -1$ and $\mathcal{O}_{p_-} = \{0\}$ and $\mathcal{O}_{p_+} = \{1\}$;
 - $\mathfrak{m}_p = 1$ and $\mathcal{O}_{p_-} = \{0\}$ and $\mathcal{O}_{p_+} = \{-1\}$.

Fix $p = (p, p') \in \Lambda$. To say more about \mathfrak{m}_p , let γ_{p_+} and γ_{p_-} denote the respective segments of the integral curves of v in \mathcal{H}_p^+ that are defined by the data $(\widehat{v}_-, (\mathfrak{k}_{p_-}, \mathcal{O}_{p_-})_{p \in \Lambda})$ and $(\widehat{v}_+, (\mathfrak{k}_{p_+}, \mathcal{O}_{p_+})_{p \in \Lambda})$. Fix $r \in [\frac{1}{2}\delta_*, \delta_*]$ so that projection map from $\mathbb{R} \times (1, 2) \times \Sigma$ to $(1, 2) \times \Sigma$ restricts to S so as to define a map that is transverse to surfaces in $(1, 2) \times \Sigma$ that correspond to the $f \in (1, 2)$ part of the boundary of the radius r coordinate balls centered at p and p' . These respective intersections define a pair of embedded arcs, one in the boundary of the radius r coordinate ball centered at p and the other in the boundary of the radius r coordinate ball centered at p' . The former starts at the point where the relevant integral curve of v from \widehat{v}_+ intersects the radius r

coordinate ball centered at p and ends at the point where the relevant integral curve of \mathfrak{v} from $\hat{\nu}_-$ intersects this radius r coordinate ball. Denote this arc by ν_{p_1} . The second arc starts from the point where the relevant integral curve of \mathfrak{v} from $\hat{\nu}_-$ intersects the radius r coordinate ball centered on p' and it ends where the relevant integral curve of \mathfrak{v} from $\hat{\nu}_+$ intersects the radius r coordinate ball centered on p_2 . Denote this second arc by ν_{p_2} .

As noted by Corollary II.2.6, the start point of ν_{p_1} has distance no greater than $c_0\delta$ from the point where γ_{p_+} intersects the boundary of the radius r coordinate ball centered at p and its endpoint has distance no greater than $c_0\delta$ from the point where γ_{p_-} intersects the boundary of the radius r coordinate ball centered at p . There is an analogous observation about the start and endpoints of ν_{p_2} .

Granted these observations, what follows defines a 1-cycle in \mathcal{H}_p^+ . Start where γ_{p_+} intersects the boundary of the radius r coordinate ball centered on p' and proceed along γ_{p_+} until it intersects the boundary of the radius r coordinate ball centered on p . Then proceed along a geodesic arc in this sphere of length $c_0\delta$ or less to the start point of ν_{p_1} . Proceed along ν_{p_1} to its endpoint and then along the geodesic arc in the sphere to the its intersection point with γ_{p_-} . Return to the boundary of the radius r coordinate ball centered on p' by traversing backwards along γ_{p_-} . Then proceed along the short geodesic in this sphere to the start point of ν_{p_2} , follow ν_{p_2} to its end, and then follow the short geodesic in this sphere to the nearby intersection point γ_{p_+} . Figure 3 shows the part of this 1-cycle in \mathcal{H}_p .




Figure 3

This closed 1–cycle defines a class in $H_1(\mathcal{H}_p^+; \mathbb{Z})$. The latter group is isomorphic to \mathbb{Z} with generator the equatorial circle in the $u = 0$ slice with the orientation given by $\frac{\partial}{\partial \phi}$. This understood, the closed 1–cycle defines an integer. This integer is m_p .

2C The submanifold C_{S_0}

Fix a Lipschitz submanifold S ; or if needed, take S from some chosen weakly compact subset $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$ of Lipschitz submanifolds. In any event, with S chosen, fix a pair $(\hat{\Theta}_-, \hat{\Theta}_+) \in \hat{\mathcal{Z}}^S$. This data labels an ersatz ech-HF submanifold, $\mathcal{C}_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$. This subsection describes C_{S_0} . The description is given in the first four parts of this subsection. An existence/uniqueness assertion is stated in Part 6 as Proposition 2.1. This proposition gives an indication of the role played by S . Part 5 of the subsection sets some background for Proposition 2.1.

Part 1 Use Lemmas II.6.3 and II.6.4 to find a constant $z_S < e^{-32} \delta_*^2$ such that the conclusions of Lemma II.6.3 holds and such that Lemma II.6.4 holds when $z \leq z_S$. With regards to Lemma II.6.4, choose z_S so as to guarantee the following: the composition of first projection from $\mathbb{R} \times (1, 2) \times \Sigma$ to $(1, 2) \times \Sigma$ and then the identification of the latter with the $f \in (1, 2)$ part of M sends the $f \leq 1 + z_S$ portion of S into the union of the radius $e^{-16} \delta_*$ coordinate balls centered on the index 1 critical points of f , and it sends the portion where $f \geq 2 - z_S$ into the union of the radius $e^{-16} \delta_*$ coordinate balls centered on the index 2 critical points of f . Note that z_S can be taken to be \mathcal{K} –compatible when \mathcal{K} is specified.

Fix $z_* \in (0, e^{-64} z_S)$ and then $\delta < e^{-16} z_*^{1/2}$. Some additional purely S –dependent (\mathcal{K} –compatible) constraints on the upper bounds for z_* and δ are given subsequently.

Part 2 The element C_{S_0} from \mathcal{C}_0 is a properly embedded, J –holomorphic submanifold with boundary in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. This surface has $2G$ boundary components, with G on $\mathbb{R} \times \{1\} \times \Sigma$ and G on $\mathbb{R} \times \{2\} \times \Sigma$. These former sets are mapped via the projection to Σ into pairwise distinct components of T_+ , and the latter are mapped via this projection into pairwise distinct components of T_- . In any event, each boundary component is an embedded copy of \mathbb{R} .

Part 3 This part describes the large $|s|$ part of C_{S_0} . To this end, let Θ_- and Θ_+ denote the respective basepoints in $\mathcal{Z}_{\text{ech}, M}$ for the chosen elements $\hat{\Theta}_-, \hat{\Theta}_+ \in \hat{\mathcal{Z}}^S$. What follows first describes the $s \ll -1$ behavior.

There exists $s_1 \geq 1$ which is such that the $s \leq -s_1$ part of C_{S_0} is a disjoint union of G graphs over $(-\infty, -s_1] \times [1 + z_*, 2 - z_*]$. Each such graph has the form

$$(2-3) \quad (s, t) \rightarrow (s, t, \psi(s, t)),$$

where ψ is a map from $(-\infty, -s_1] \times [1 + z_*, 2 - z_*]$ to $T_+ \cap T_-$. The images of these G maps are in distinct components, and each such component contains the intersection with Σ of an integral curve of v from Θ_- . Let q_* now denote such an intersection point and let ψ_* denote the map from (2-3) with image in the q_* component of $T_+ \cap T_-$. Then

$$(2-4) \quad \text{dist}(\psi_*(s, \cdot), q_*) \leq c e^{-|s|/c},$$

where $c \geq 1$ is a purely S -dependent (or \mathcal{K} -compatible) constant. Finally, if the given component of $T_+ \cap T_-$ is parametrized using the relevant index 1 critical point versions of the functions (φ_+, \hat{h}_+) , and if ψ is written with respect to these coordinates in terms of functions $(\varphi_+ = \varphi_*(s, t), \hat{h}_+ = \zeta_*(s, t))$, then the pair (φ_*, ζ_*) obey the Cauchy–Riemann equations $\partial_s \varphi_* - \partial_t \zeta_* = 0$ and $\partial_s \zeta_* + \partial_t \varphi_* = 0$.

The $s \geq s_1$ part of C_{S_0} has the analogous description with Θ_+ replacing Θ_- .

Part 4 This part describes the boundary behavior of C_{S_0} near any given boundary component. This involves a purely S -dependent (or \mathcal{K} -compatible) constant, κ_{S_*} , which is greater than 100. If z_S is chosen less than $\kappa_{S_*}^{-2}$ then what follows holds is true. Let p denote either an index 1 or index 2 critical point of f . When p has index 1, use (φ_+, \hat{h}_+) to parametrize T_{p_+} , and when p has index 2, use (φ_-, \hat{h}_-) to parametrize T_{p_-} . Then the part of C_{S_0} in $\mathbb{R} \times [1 + z_*, 1 + z_S] \times T_{p_+}$ or in $\mathbb{R} \times [2 - z_S, 2 - z_*] \times T_{p_-}$ is diffeomorphic to $\mathbb{R} \times [z_*, z_S]$ and parametrized via a map of the form

$$(2-5) \quad \begin{aligned} &\bullet (x, z) \rightarrow (s = x, t = 1 + z, \varphi_+ = \varphi^{S_0}(x, z), \hat{h}_+ = \zeta^{S_0}(x, z)) \text{ when } p \text{ has index 1,} \\ &\bullet (x, z) \rightarrow (s = x, t = 2 - z, \varphi_- = \varphi^{S_0}(s, z), \hat{h}_- = \zeta^{S_0}(s, z)) \text{ when } p \text{ has index 2.} \end{aligned}$$

The functions φ^{S_0} and ζ^{S_0} that appear here are \mathbb{R} -valued functions that obey the Cauchy–Riemann equations $\partial_x \varphi^{S_0} - \partial_z \zeta^{S_0} = 0$ and $\partial_x \zeta^{S_0} + \partial_z \varphi^{S_0} = 0$. In addition, the first and higher derivatives of these functions to any given order are bounded uniformly on $\mathbb{R} \times [z_*, z_S]$ by a constant that depends only on the given order and S (it is \mathcal{K} -compatible when \mathcal{K} is given). Finally $|\zeta^{S_0}(\cdot, z_*)| \leq \kappa_{S_*} z_*$.

Part 5 The upcoming Proposition 2.1 states an existence/uniqueness assertion for C_{S_0} . This part of the subsection supplies some background for this proposition.

Proposition 2.1 views the $f \in (1, 2)$ part of M as $(1, 2) \times \Sigma$ so as to describe C_{S_0} as a submanifold with boundary in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. The submanifold C_{S_0} in this guise is the $t \in [1 + z_*, 2 - z_*]$ portion of a properly embedded, J_{HF} -holomorphic submanifold in $\mathbb{R} \times [1, 2] \times \Sigma$, this denoted by S_* . The submanifold S_* has 2G boundary components, one in each index 1 critical point version of $\mathbb{R} \times \{1\} \times T_{p+}$, and, likewise, one in each index 2 critical point version of $\mathbb{R} \times \{2\} \times T_{p-}$.

The submanifold S_* is isotopic to S in a small radius tubular neighborhood of S . The description in Proposition 2.1 identifies this tubular neighborhood with a disk bundle in the normal bundle of S using an exponential map of the sort that is described in Section II.6E. What follows reviews some aspects of this sort of exponential map.

To start, recall from Section II.6E that S has a complex normal bundle, $N_S \rightarrow S$, and an exponential map that embeds a disk subbundle as a tubular neighborhood of S . The exponential map is denoted by ϵ_S and the disk subbundle by N_0 . The latter has radius ρ_S and its image in $\mathbb{R} \times [1, 2] \times \Sigma$ is a tubular neighborhood that contains the set of points with distance $c^{-1}\rho_S$ from S . The map ϵ_S embeds each fiber disk as a J_{HF} -holomorphic disk. If \mathcal{K} is a previously specified, weakly compact set of Lipschitz submanifolds, then ρ_S and c can be taken to be \mathcal{K} -compatible, as can the derivatives to any given order of the exponential map ϵ_S .

Although not stated as such in Section II.6E, the map ϵ_S can be chosen so as to respect the graph structure described in Section II.6C and Property 5 of Section 1G near the boundaries of S . In particular, ϵ_S can be chosen so that it maps any given fiber of N_0 over the $t \in (1, 1 + z_S)$ and $t \in (2 - z_S, z_S)$ portions of S as follows: The graph structure indicated by Property 5 of Section 1G identifies the bundle N_S over this part of S with the restriction to S of $T\Sigma$. In particular, the 1-forms $(d\varphi_+, dh_+)$ when $t \in (1, 1 + z_S)$ and $(d\varphi_-, -dh_-)$ when $t \in (2 - z_S, 2)$ with the underlying real bundle define an orientation-preserving isomorphism to the product \mathbb{R}^2 -bundle. Given this isomorphism, ϵ_S on these parts of S can and should be chosen so as to send any point in S parametrized by (x, z) and a pair (a, b) in the \mathbb{R}^2 factor of the product bundle to the point given by one of

$$(2-6) \quad \bullet \quad (s = x, t = 1 + z, \varphi_+ = \varphi(x, z) + a, h_+ = \zeta(x, z) + b),$$

$$\bullet \quad (s = x, t = 2 - z, \varphi_- = \varphi(x, z) + a, h_- = \zeta(x, z) - b).$$

The pair $(d\varphi_+, dh_{+})$ as defined by any given index 1 critical point of f are also used to write N_S as the product bundle over the corresponding $s \ll -1$ and $s \gg 1$ parts of S . By way of a reminder, the $s \ll -1$ part of S has G components, each projecting to a distinct component of $[1, 2] \times (T_+ \cap T_-)$, and these G components can be labeled by the index 1 critical points of f . The same can be said for the $s \gg 1$ part of S . In any event, the exponential map ϵ_S can and should be chosen so that it acts as in the top bullet of (2-6) on these large $|s|$ parts of S .

The construction of a map ϵ_S of this sort can be made using the techniques that are used to prove Lemma 5.4 in [19].

Proposition 2.1 refers to the Fredholm operator, D_S , that is described in Section II.6E; it is depicted in (1-25). By way of a reminder, this operator maps a certain Hilbert space of sections of N_S to the space of square-integrable sections of $N_S \otimes T^{0,1}S$. The Hilbert space for the domain is the Sobolev L^2_1 -norm completion of the subspace of sections that obeys the constraints in (II.6-12). The *kernel* of D_S refers to the sections of N_S in the domain Hilbert space that are annihilated by D_S . The L^2 inner product on sections of N_S is defined using the fiber metric on N_S and the integration measure on S that comes from the metric induced by its embedding in $\mathbb{R} \times [1, 2] \times \Sigma$.

Write the pair $\hat{\Theta}_-$ and $\hat{\Theta}_+$ from the chosen element in $\hat{\mathcal{Z}}^S$ as (Θ_-, k_-) and (Θ_+, k_+) . Proposition 2.1 refers to a number that is associated to each index 1 and each index 2 critical point of f by Θ_+ and another that is determined by Θ_- . When p is used to denote the critical point in question, the corresponding two numbers are denoted respectively by h_{p+} and h_{p-} . When p is an index 1 critical point of f , the numbers h_{p+} and h_{p-} denote the respective h_+ coordinates of the T_{p+} intersection point of an integral curve of v from Θ_+ and Θ_- with the $t = 1 + z_*$ slice of $(1, 2) \times \Sigma$. When p is an index 2 critical point of f , the numbers h_{p+} and h_{p-} denote the respective h_- coordinates of the T_{p-} intersection point of an integral curve of v from Θ_+ and Θ_- with the $t = 2 - z_*$ slice of $(1, 2) \times \Sigma$. By way of a parenthetical remark, it follows from what is said in Section II.2 that $|h_{p+}|$ and $|h_{p-}|$ are both bounded by $c_0\delta^2$.

With regards to Θ_- and Θ_+ , Proposition 2.1 uses \hat{v}_- and \hat{v}_+ to denote the respective HF-cycles that are used in (2-1) for their definition. Let $q \in \Sigma$ denote a given intersection point with an integral curve of v from either \hat{v}_- or \hat{v}_+ . Let q_* denote the corresponding, nearby intersection point of the corresponding segment of an integral curve of v from Θ_- or Θ_+ as the case may be. Note that q_* has distance at most $c_0\delta$ from q . The point q is an element in $C_+ \cap C_-$ and therefore in some index 1 critical point component

of C_+ . This critical point labels a corresponding $s \ll -1$ or $s \gg 1$ component of S . The latter component is denoted by \mathcal{E}_{S_q} . Proposition 2.1 writes the exponential map ϵ_S over \mathcal{E}_{S_q} as in the top line of (2-6) so as to view q_* as a section of N_S over \mathcal{E}_{S_q} .

Here is one final item of notation: the function $x \rightarrow \chi(x)$ maps \mathbb{R} to $[0, 1]$; it is nonincreasing, equal to 1 where $x \leq 0$ and equal to zero where $x \geq 1$.

Part 6 This final part of the subsection first states and then proves the existence/uniqueness proposition about C_{S_0} .

Proposition 2.1 *Fix a Lipschitz submanifold S or one from some specified weakly compact subset $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$ of Lipschitz submanifolds. There exist $\kappa_{S_*} \geq 100$ and $z_S \in (0, \kappa_{S_*}^{-2})$ that depend only on S (and are \mathcal{K} -compatible if relevant) such that what follows is true. Fix $z_* \in (0, e^{-32} z_S)$ and then $\delta \in (0, e^{-16} z_*^{1/2})$ and χ_0 and R . There exists a unique section η_* of N_0 that is characterized by:*

- The C^4 -norm of η_* is bounded by $\kappa_{S_*} \delta$.
- The restriction of η_* to the $t \in [1 + z_*, 2 - z_*]$ part of S is L^2 -orthogonal to the corresponding restriction of the elements in the kernel of the operator D_S .
- Let p denote either an index 1 or index 2 critical point of f . The pairing of the section η_* along the corresponding boundary component of S with the relevant 1-form $d\hat{h}_+$ or $d\hat{h}_-$ is the function on \mathbb{R} given by $\hat{h}_{p+}(1 - \chi) + \hat{h}_{p-}\chi$.
- Let $q \in \Sigma$ denote an intersection point with an integral curve of \mathfrak{v} from either $\hat{\nu}_-$ or $\hat{\nu}_+$ and let q_* denote the corresponding section of N_S over \mathcal{E}_{S_q} . The pointwise norm of $\eta_* - q_*$ converges to zero as $|s| \rightarrow \infty$ on \mathcal{E}_{S_q} . Moreover, given $k \geq 0$, there exists a purely S -dependent (or \mathcal{K} -compatible) constant c such that the derivatives to order k on \mathcal{E}_{S_q} are bounded by $c e^{-|s|/c}$.
- With $(1, 2) \times \Sigma$ viewed now as the $f \in (1, 2)$ part of M , use $C_{S_0} \subset \mathbb{R} \times M_\delta$ to denote the $t \in [1 + z_*, 2 - z_*]$ part of $S_* = \epsilon_S \circ \eta_*(S)$. This version of C_{S_0} obeys the properties listed in the preceding Parts 1–4 of this subsection.

Proof The proof that follows has three steps.

Step 1 Suppose that η_1 is a smooth section of N_0 that has the same large $|s|$ asymptotics and boundary behavior as the desired η_* . Assume that the pointwise norm of η_1 and those of its derivatives to sixth order are bounded by $c_0 \delta$. In addition, require that the surface $\epsilon_S \circ \eta_1$ be J_{HF} -holomorphic where $|s| \geq c$, where $c > 1$ is a purely

S -dependent (or \mathcal{K} -compatible) constant. A section η_1 satisfying these requirements is given in Step 2.

What follows constructs a section, η_2 , of N_0 in the Fredholm domain of the operator D_S such that $\eta_* = \eta_1 + \eta_2$ obeys all of the bullets of the proposition. To this end, keep in mind what is said by Part 2 of Section II.6E and by (II.6-10): If $\eta = \eta_1 + \eta_2$ is a section of N_S , then $\epsilon_S \circ \eta$ is J_{HF} -holomorphic if and only if η_2 obeys an equation that has the schematic form

$$(2-7) \quad \bar{\partial}\eta_2 + \tau_{1*}(\eta_2) \cdot \partial\eta_2 + \tau_{0*}(\eta_2) = \tau_*$$

where the notation is as follows: First, τ_* is a smooth section of $N_S \otimes T^{0,1}S$ with compact support where $|s| \leq c$ and C^5 -norm and L^2 -norm bounded by $c\delta$. Here, $c > 1$ again denotes a purely S dependent (or \mathcal{K} -compatible) constant. This term τ_* is determined by η_1 . Second, τ_{1*} and τ_{0*} are analogous to their counterparts (II.6-10). They differ from the latter by virtue of a dependence on η_1 , but even so, this difference has C^5 -norm bounded by $c\delta$ with c as just described. In particular, they obey $|\tau_{1*}(b)| \leq c|b|$ and $|\tau_{2*}(b) - \nu b - \mu \bar{b}| \leq c|b|^2$, where ν and μ are from (1-25) and c is as described above. In addition, their derivatives to any given order are bounded by purely S -dependent (or \mathcal{K} -compatible) constants.

Granted these last remarks, the equation for η_2 can be written as

$$(2-8) \quad D_S \eta_2 + \mathfrak{z}(\eta_2) = \tau_*$$

where $|\mathfrak{z}(\eta_2)| \leq c(|\eta_2|^2 + |\eta_2| |\nabla \eta_2|)$. Here again, $c \geq 1$ is a purely S -dependent (or \mathcal{K} -compatible) constant. With this last fact understood, and given the aforementioned bounds on the higher derivatives of τ_{0*} , τ_{1*} and τ_* , the existence and uniqueness of the desired solution to (2-8) follows via a standard application of the implicit function theorem.

Step 2 Consider now η_1 . What follows here describes η_1 on a component of the $s \ll -1$ part of S . To do this, return to the notation used in Part 3. Let $q \in C_+ \cap C_-$ denote the relevant point. As noted in Property 6 of Section 1G, the end \mathcal{E}_{S_q} can be viewed as a graph of a map from the $s \leq -c$ part of $\mathbb{R} \times [1, 2]$ into q 's component of $T_+ \cap T_-$. Use the coordinates (φ_+, \hat{h}_+) to write q as the origin in \mathbb{R}^2 and the corresponding map ψ to $T_+ \cap T_-$ as a map to \mathbb{R}^2 . With the normal bundle N_S identified with \mathbb{R}^2 as in the top line of (2-6), the section η_1 where $s \leq -c(s_1 + |\ln \delta|)$ is $q_* - \psi$. Note that this formula is such that $\epsilon_S \circ \eta_1$ is the J_{HF} -holomorphic surface $\mathbb{R} \times (1, 2) \times q_*$

on this part of S . There is an analogous formula for η_1 where $s \geq c(s_1 + |\ln \delta|)$. The desired behavior of η_1 near the boundary of S can be obtained in a straightforward fashion using the description of S given in Property 5 of Section 1G. A similarly straightforward use of “cutoff” functions will extend the section of N_S that is defined by the resulting formula for η_1 near the boundary of S and that given above on the $|s| \gg c$ on S so as to define the desired version of η_1 over the whole of S .

Step 3 The demands of the fifth bullet follow directly given the third and fourth bullets and given that η_2 obeys (2-8). As η_2 is in the Fredholm domain of D_S , the third bullet follows if η_2 is smooth up to the boundary. This can be proved using slightly modified versions of the arguments that are used to prove Lemma II.6.3. Meanwhile, the assertion made by the fourth bullet is proved using arguments that are little different from those used in Section II.6C. □

2D The submanifolds $\{C_{p_0}\}_{p \in \Lambda}$

The subsequent three parts of this subsection describe the salient features of the submanifolds that form the subset $\{C_{p_0}\}_{p \in \Lambda}$ from \mathcal{C}_0 . To this end, fix $p \in \Lambda$ so as to focus on the corresponding element C_{p_0} . The subsequent description uses $\mathcal{H}_{p_*}^+$ to denote the $e^{-2(R-|u|)}(1 - 3 \cos^2 \theta) \leq z_*$ part of \mathcal{H}_p^+ . Part 4 of the subsection states and then proves the existence/uniqueness assertion about C_{p_0} .

Part 1 What is denoted by C_{p_0} is a properly embedded submanifold with boundary in $\mathbb{R} \times \mathcal{H}_{p_*}^+$ with J -holomorphic interior. There are two boundary components, one on the $u > 0$ component of the boundary of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ and the other on the $u < 0$ component. Define $\Delta_p \in \{0, 1, 2\}$ as follows: If $m_p = 0$ and item (a) of the first bullet of (2-2) is relevant, then $\Delta_p = 0$. If item (b) is relevant, then $\Delta_p = 2$. If $m_p = 1$ or $m_p = -1$, then $\Delta_p = 1$. In the case $\Delta_p = 0$, the submanifold C_{p_0} is diffeomorphic to the product of \mathbb{R} with a closed interval. When $\Delta_p = 1$, the submanifold C_{p_0} is diffeomorphic to the complement of a single interior point in the product of \mathbb{R} with a closed interval. When $\Delta_p = 2$, it is diffeomorphic to the complement of two interior points in such a product.

To describe the large $|s|$ behavior of C_{p_0} , introduce γ_{p_-} and γ_{p_+} to denote the respective integral curve segments in \mathcal{H}_p^+ that come from Θ_- and Θ_+ , and introduce $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ to denote the respective closed integral curves of v in the $u = 0$ slice of \mathcal{H}_p that form the loci where $\cos \theta = \frac{1}{\sqrt{3}}$ and $\cos \theta = -\frac{1}{\sqrt{3}}$.

- (2-9) • Each constant $s \ll -1$ slice of C_{p_0} is a properly embedded arc in \mathcal{H}_p^+ , and these arcs converge pointwise as $s \rightarrow -\infty$ to γ_{p_-} .
- If $\Delta_p = 0$, then each constant $s \gg 1$ slice of C_{p_0} is a properly embedded arc in \mathcal{H}_p^+ , and these arcs converge in an isotopic fashion in \mathcal{H}_p as $s \rightarrow \infty$ to γ_{p_+} .
 - If $\Delta_p = 1$, then each constant $s \gg 1$ slice of C_{p_0} has two components.
 - (a) One component is a properly embedded arc in \mathcal{H}_p^+ , and these arcs converge in an isotopic fashion in \mathcal{H}_p as $s \rightarrow \infty$ to γ_{p_+} .
 - (b) The other component is an embedded circle, and these circles converge pointwise in \mathcal{H}_p as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$ when $m_p = -1$, and they converge in an isotopic fashion in \mathcal{H}_{p_*} as $s \rightarrow \infty$ to $\hat{\gamma}_p^-$ when $m_p = 1$.
 - If $\Delta_p = 2$, then each constant $s \gg 1$ slice of C_{p_0} has three components. One component is a properly embedded arc in \mathcal{H}_p^+ , and these arcs converge in an isotopic fashion in \mathcal{H}_p as $s \rightarrow \infty$ to γ_{p_+} . The other two components are embedded circles. One s -parametrized set of these circles converges pointwise in an isotopic fashion in \mathcal{H}_p as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$; the other set converges in an isotopic fashion in \mathcal{H}_{p_*} as $s \rightarrow \infty$ to $\hat{\gamma}_p^-$.

Part 2 This part says more about how C_{p_0} sits in $\mathbb{R} \times \mathcal{H}_{p_*}^+$. To this end, reintroduce from Section 1H the parametrization of \mathcal{H}_p^+ by the map Ψ_p . The domain of this map Ψ_p is an open subset in $\mathbb{R} \times (-R - \ln \delta_*, R + \ln \delta_*) \times (\mathbb{R}/2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$. As in Section 1H, the coordinates for the latter are written as $(x, \hat{u}, \hat{\phi}, h)$. By way of a reminder, the domain of Ψ_p is an open subset of the form $\mathbb{R} \times \mathcal{X}$. The Ψ_p -inverse image of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ is the subset of the domain $\mathbb{R} \times \mathcal{X}$ where $|\hat{u}| \leq R + \frac{1}{2} \ln z_*$. This being the case, it proves useful to introduce I_* to denote the interval $[-R - \frac{1}{2} \ln z_*, R + \frac{1}{2} \ln z_*]$ and restrict Ψ_p to the $\hat{u} \in I_*$ part of its domain.

The Ψ_p -inverse image of C_p is given as the image of a proper map from a certain domain in $\mathbb{R} \times I_*$ to $\mathbb{R} \times \mathcal{X}$. This map has the form

$$(2-10) \quad (x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi^{p_0}(x, \hat{u}), h = \zeta^{p_0}(x, \hat{u})),$$

where $(\varphi^{p_0}, \zeta^{p_0})$ is a map from a domain in $\mathbb{R} \times I_*$ to $\mathbb{R}/(2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$. The domain of this map is $\mathbb{R} \times I_*$ if $\Delta_p = 0$, it is the complement of a single point in the $\hat{u} = 0$ slice of $\mathbb{R} \times I_*$ if $\Delta_p = 1$, and it is the complement of two points in the $\hat{u} = 0$ slice if $\Delta_p = 2$.

Part 3 This part describes the behavior of C_{p_0} near its boundaries. Consider first the boundary on the $u > 0$ component of the $e^{-2(R-|u|)}(1 - 3 \cos^2 \theta) = z_*$ locus. To set the notation, note that the $f \geq 1 + \delta^2$ part of $M_\delta \cap \mathcal{H}_p^+$ can be parametrized using the coordinates $(t, \varphi_+, \hat{h}_+)$ from the relevant component of the $(1 + \delta^2, 1 + \delta_*^2) \times T_+$ portion of $(1, 2) \times \Sigma$.

The inverse image via Ψ_p of the $f \in (\delta^2, z_*]$ part of $\mathbb{R} \times \mathcal{H}_p^+$ corresponds to the part of $\mathbb{R} \times \mathcal{X}$ where $\hat{u} \in [R + \ln \delta, R + \frac{1}{2} \ln z_*]$. The coordinate functions $(x, \hat{u}, \hat{\phi}, h)$ on the $\hat{u} \in (R + \ln \delta, R + \frac{1}{2} \ln z_*)$ part of $\mathbb{R} \times \mathcal{X}$ are related to the coordinate functions $(s, t, \varphi_+, \hat{h}_+)$ using the rule $(s = x, t = 1 + e^{-2(R-|\hat{u}|)}, \varphi_+ = \hat{\phi}, \hat{h}_+ = h)$.

Granted the preceding, it follows that the functions $(\varphi^{p_0}, \zeta^{p_0})$ that appear in (2-10) can be viewed as functions of the coordinates (s, z) . Doing so writes the part of C_{p_0} in the $f \in (1 + \delta^2, 1 + z_*)$ part of $\mathbb{R} \times (M_\delta \cap \mathcal{H}_p^+)$ as the graph over $\mathbb{R} \times (\delta^2, z_*)$ given by the rule

$$(2-11) \quad (x, z) \rightarrow (s = x, t = 1 + z, \varphi_+ = \varphi^{p_0}, \hat{h}_+ = \zeta^{p_0}).$$

The fact that C_{p_0} is J -holomorphic implies that the pair $(\varphi^{p_0}, \zeta^{p_0})$ obey the Cauchy–Riemann equations: $\partial_x \varphi^{p_0} - \partial_z \zeta^{p_0} = 0$ and $\partial_x \zeta^{p_0} + \partial_z \varphi^{p_0} = 0$.

There is a corresponding picture of C_{p_0} on the $t \in [2 - z_*, 2 - \delta^2)$ portion of $\mathbb{R} \times (M_\delta \cap \mathcal{H}_{p_*}^+)$. This part of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ corresponds via Ψ_p to the

$$\hat{u} \in [-R - \frac{1}{2} \ln z_*, -R - \ln \delta)$$

part of $\mathbb{R} \times \mathcal{X}$. It is parametrized by coordinates $(x, z, \varphi_-, \hat{h}_-)$ with $z \in (\delta^2, z_*)$ related to the coordinate t by the rule $t = 2 - z$ and z related to \hat{u} by the rule $z = e^{-2(R+|\hat{u}|)}$. The part of C_{p_0} here is parametrized by viewing $(\varphi^{p_0}, \zeta^{p_0})$ as functions of (x, z) and writing $\varphi_- = \varphi^{p_0}$ and $\hat{h}_- = -\zeta^{p_0}$. The pair $(\varphi^{p_0}, \zeta^{p_0})$ obey here the Cauchy–Riemann equations when written as functions of $(x, -z)$.

What follows is now a crucial point: The functions φ^{p_0} is constrained on both boundary components of C_{p_0} as follows: Write φ^{p_0} on either boundary as a function of the coordinate $x \in \mathbb{R}$. Meanwhile, write φ^{S_0} from (2-5) on the relevant $\mathbb{R} \times \{1 + z_*\} \times T_+$ or $\mathbb{R} \times \{2 - z_*\} \times T_-$ part of the boundary of C_{S_0} as a function of x also. Then

$$(2-12) \quad \varphi^{p_0}(x, z_*) = \varphi^{S_0}(x, z_*).$$

There is no a priori constraint on the value of ζ^{p_0} on the boundaries of ∂C_{p_0} but for what is implied by (2-9).

Part 4 The proposition given below states the fundamental existence/uniqueness theorem for C_{p_0} . The proposition refers to the preceding Parts 1–3.

Proposition 2.2 *There exists a purely S –dependent (or \mathcal{K} –compatible) constant $\kappa > 1$ with the following significance: Define the geometry of Y with $z_* < \kappa^{-1}$ and $\delta < \kappa^{-1} z_*$. Fix an element $(\hat{\Theta}_-, \hat{\Theta}_+) \in \hat{\mathcal{Z}}^S$. In the case $\Delta_p = 1$, fix a $\hat{u} = 0$ point in $\mathbb{R} \times I_*$; and in the case $\Delta_p = 2$, fix a pair of $\hat{u} = 0$ points in $\mathbb{R} \times I_*$. There exists a unique pair $(\varphi^{p_0}, \zeta^{p_0})$ that are described by Parts 2 and 3 and are such that the Ψ_p –image of the surface given by (2-9) obeys the conditions stated for C_{p_0} in Part 1.*

The proof of Proposition 2.2 in the case when all $p \in \Lambda$ versions of Δ_p are zero is given separately in Section 4 because it has fewer components than the proof for the $\Delta_p > 1$ cases. The proof for general case is given in Section 6. The next section introduces certain analytic tools that are used in Section 4. Section 5 introduces some additional tools to handle the general case.

3 Cauchy–Riemann equations on $\mathbb{R} \times \mathcal{X}$

The Ψ_p –image of a graph of the form

$$(3-1) \quad (x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi(x, \hat{u}), h = \zeta(x, \hat{u}))$$

defines a J –holomorphic subvariety if and only if the pair of functions (φ, ζ) satisfy a certain nonlinear Cauchy–Riemann equation as functions of the coordinates (x, \hat{u}) . The purpose of this section is to describe this equation and to supply various tools that will be used subsequently to construct desired solutions.

3A Almost complex structures on $\mathbb{R} \times \mathcal{X}$

Suppose that J_p is a given almost complex structure on $\mathbb{R} \times \mathcal{H}_{p_*}^+$ with the property that the $\mathbb{R} \times \mathcal{H}_{p_*}^+$ part of any surface from Proposition II.3.2’s moduli space \mathcal{M}_Σ and Proposition II.3.4’s moduli space \mathcal{M}_{p_0} are J_p –holomorphic. Use Ψ_p to view J_p as an almost complex structure on $\mathbb{R} \times \mathcal{X}$.

The Ψ_p –inverse images of the surfaces from \mathcal{M}_Σ and those from \mathcal{M}_{p_0} are the constant (x, \hat{u}) slices of $\mathbb{R} \times \mathcal{X}$. This understood, the fact that they are J_p –holomorphic has the following implication: the J_p version of $T^{1,0}(\mathbb{R} \times \mathcal{X})$ must contain a form that can be written as

$$(3-2) \quad dx + i q_0 d\hat{u},$$

where q_0 is a \mathbb{C} -valued function with strictly positive real part. A second linearly independent 1-form for J_p 's version of $T^{1,0}(\mathbb{R} \times \mathcal{X})$ can always be written as

$$(3-3) \quad d\hat{\phi} + i q_1 dh + i \mu_0(dx - i \bar{q}_0 d\hat{u}),$$

where q_1 is a \mathbb{C} -valued function with strictly positive real part, and μ_0 is a \mathbb{C} -valued function.

Lemma 3.1 *Suppose that J_p obeys the $\mathbb{R} \times \mathcal{H}_{p*}^+$ versions of the first five bullets in Part 1 of Section 1C. Then q_0 , q_1 and μ_0 depend only on the coordinates \hat{u} and h , and both q_1 and μ_0 are real-valued.*

Proof Whether real or not, the functions q_0 , q_1 and μ_0 are invariant with respect to constant translations along the \mathbb{R} and $\mathbb{R}/2\pi\mathbb{Z}$ factors in $\mathbb{R} \times \mathcal{X}$ if J_0 is invariant with respect to the respective translations along the \mathbb{R} and $\mathbb{R}/2\pi\mathbb{Z}$ factors of $\mathbb{R} \times \mathcal{H}_{p*}^+$.

To see about the imaginary parts of q_1 and μ_0 define the adjoint action of J_p on the cotangent bundle by the following rule: Let $\langle \cdot, \cdot \rangle$ denote the pairing between covectors and vectors and let ϵ and w denote respective covectors and vectors over the same base-point. Then $\langle J_p^T \epsilon, w \rangle = \langle \epsilon, J_p w \rangle$. Note that J_p^T acts on (3-2) and on (3-3) as multiplication by i . With this in mind, use the identity $J_p \partial_s = v$ with the second bullet in (1-30) to see that J_p acts as multiplication by $-i$ on $\partial_x + i(v^{-1} \partial_{\hat{u}} - \alpha^{-1} \sqrt{6} \chi \cos \theta \partial_{\hat{\phi}} - \varpi \partial_x)$. This vector is therefore sent to zero when paired via $\langle \cdot, \cdot \rangle$ with (3-2). Such is the case if and only if $(1 - i\varpi) - q_0 v^{-1} = 0$ so $q_0 = v(1 - i\varpi)$. This same vector is also sent to zero when paired via $\langle \cdot, \cdot \rangle$ with (3-3), and this happens if and only if $2\mu_0 = \alpha^{-1} \sqrt{6} \chi \cos \theta$. To see about q_1 , use the fourth bullet in (1-30) and (II.3-9) to see that $\Psi_{p*} J_p \partial_h$ is proportional to ∂_ϕ , and so it follows from the third bullet of (1-30) that $J_p \partial_h$ is proportional to $\partial_{\hat{\phi}}$. Now use (3-3) to see that q_1 is real if and only if $\langle dh, J_p \partial_h \rangle = 0$. □

Assume in what follows that J_p obeys the assumptions of Lemma 3.1. With it understood that q_1 and μ_0 are real, (3-2) and (3-3) imply the following: A submanifold in $\mathbb{R} \times \mathcal{H}_{p*}^+$ given by the Ψ_p -image of the surface given by (3-1) is J_p -holomorphic if and only if the functions (φ, ζ) that appear in (3-1) obey a system of Cauchy-Riemann equations that can be written as

$$(3-4) \quad a_1 \partial_x \varphi - \partial_{\hat{u}} \zeta = 0 \quad \text{and} \quad a_2 \partial_x \zeta + \partial_{\hat{u}} \varphi + b = 0,$$

where a_1 , a_2 and b constitute a set of \mathbb{R} -valued functions with a_1 and a_2 strictly positive. Their respective values at any given point (x, \hat{u}) are obtained from an

eponymous set of the functions of the variables (\hat{u}, h) by setting $h = \zeta(x, \hat{u})$. This eponymous set is $a_1 = q_{0R}/q_1$, $a_2 = q_{0R}q_1$ and $b = -q_{0R}\mu_0$, where q_{0R} denotes the real part of q_0 .

Keep in mind that the expressions that appear on the left-hand sides of the two equations in (3-4) are only defined in the case that the map Ψ_p is defined on the graph in (3-1). This is to say that the graph must define a surface in $\mathbb{R} \times \mathcal{X}$. This requirement constitutes an implicit constraint on the absolute value of ζ at any given point (x, \hat{u}) . In particular, an assertion in the subsequent discussions that a given pair (φ, ζ) solves (3-4) in all cases implies that ζ obeys this implicit constraint.

3B Maps from $\mathbb{R} \times I_*$ to \mathbb{R}^2 and linear operators

The central concern for the rest of this section are first-order linear operators on $C^\infty(\mathbb{R} \times I_*; \mathbb{R}^2)$ that are described next. Let D denote the operator in question and let $(\varphi', \zeta') : \mathbb{R} \times I_* \rightarrow \mathbb{R}^2$ denote a given map. The respective first and second components of the map $D(\varphi', \zeta')$ are

$$(3-5) \quad \alpha_1 \partial_x \varphi' - \partial_{\hat{u}} \zeta' + b_1 \zeta' \quad \text{and} \quad \alpha_2 \partial_x \zeta' + \partial_{\hat{u}} \zeta' + b_2 \zeta',$$

where α_1 , α_2 , b_1 and b_2 are smooth functions of (x, \hat{u}) with the following four properties:

- (3-6) • α_1 and α_2 are everywhere positive and they both have uniform limits as $x \rightarrow \pm\infty$ to positive functions of \hat{u} .
- The function on \mathbb{R} given by the rule $x \rightarrow \sup_{\{x\} \times I_*} (|\partial_x \alpha_1| + |\partial_x \alpha_2| + |b_1|)$ limits uniformly as $|x| \rightarrow \infty$ with limit zero.
 - b_2 limits uniformly as $x \rightarrow -\infty$ to a function of \hat{u} . By the same token, b_2 limits uniformly as $x \rightarrow \infty$ to a function of \hat{u} .
 - The respective integrals over I_* of the $x \rightarrow \infty$ and $x \rightarrow -\infty$ limits of b_2 are nonzero and have the same sign.

Of interest in what follows is a Fredholm incarnation of the operator D given by (3-5) and (3-6) whose domain and range are certain Hilbert spaces of maps from $\mathbb{R} \times I_*$ to \mathbb{R}^2 that is characterized as follows: The domain Hilbert space for D is the L^2_1 completion of the subspace of smooth maps from $\mathbb{R} \times I_*$ to \mathbb{R}^2 whose elements are as follows: A given pair (φ', ζ') is in this subspace if and only if the following conditions are met:

- (3-7) • The pair has compact support on $\mathbb{R} \times I_*$.
- The function φ' vanishes on the $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundaries of $\mathbb{R} \times I_*$.

The square of the L^2_1 -norm in question assigns to (φ', ζ') the value

$$(3-8) \quad \int_{\mathbb{R} \times I_*} (|\partial_x \varphi', \partial_x \zeta'|^2 + |\partial_{\hat{u}} \varphi', \partial_{\hat{u}} \zeta'|^2 + |(\varphi', \zeta')|^2).$$

Here, $|\cdot|^2$ denotes the Euclidean inner product on \mathbb{R}^2 . The Hilbert space so defined is denoted in what follows by \mathbb{H} . The range space for this Fredholm version of D is the L^2 Hilbert space completion of the space of compactly supported elements in $C^\infty(\mathbb{R} \times I_*; \mathbb{R}^2)$; this is the completion that is defined using the inner product on $\mathbb{R} \times I_*$ whose square is the integral of $|(\varphi', \zeta')|^2 = |\varphi'|^2 + |\zeta'|^2$. This L^2 Hilbert space is denoted by \mathbb{L} .

The next proposition asserts the central fact about this Fredholm version of D .

Proposition 3.2 *The operator D as described by (3-5) and (3-6) with domain space \mathbb{H} and range space \mathbb{L} is Fredholm with index 0 and trivial kernel.*

Section 3C proves that D is Fredholm and Section 3D computes the index of D and proves that the kernel is trivial. The remainder of this subsection describes the relevant examples.

Let $\mathfrak{h} = (\varphi, \zeta) : \mathbb{R} \times I_* \rightarrow \mathbb{R}^2$ denote a pair of functions with a graph given by (3-1) that lies in $\mathbb{R} \times \mathcal{X}$. The pair \mathfrak{h} and the almost complex structure $J_{\mathfrak{p}}$ can be used to define a version of (3-5), this denoted by $D_{\mathfrak{h}}$. The definition is as follows: Let (φ', ζ') denote a given bounded map from $\mathbb{R} \times I_*$ to \mathbb{R}^2 . Take $t \in \mathbb{R}$ near zero and write the expressions on the right-hand sides of the two equations in (3-4) using pair $(\varphi + t\varphi', \zeta + t\zeta')$ in lieu of (φ, ζ) . View the result as a map from a neighborhood of 0 in \mathbb{R} to $C^\infty(\mathbb{R} \times I_*; \mathbb{R}^2)$. The derivative of this map at $t = 0$ is $D_{\mathfrak{h}}(\varphi', \zeta')$.

A more explicit description of $D_{\mathfrak{h}}$ is given in a moment. To this end, recall that the functions (a_1, a_2, b) that appear in (3-4) are obtained from an eponymous set of functions, (a_1, a_2, b) , of the coordinates (\hat{u}, h) for \mathcal{X} . Let (a_{1h}, a_{2h}, b_h) denote the functions on $\mathbb{R} \times I_*$ whose respective values at any given point $(x, \hat{u}) \in \mathbb{R} \times I_*$ are those of the partial derivative with respect to h at the point $(x, \hat{u}, h = \zeta(x, \hat{u}))$ of (a_1, a_2, b) . The respective first and second components of the \mathbb{R}^2 -valued function $D_{\mathfrak{h}}(\varphi', \zeta')$ can be written in terms of these partial derivatives as

$$(3-9) \quad \begin{aligned} & \bullet \quad a_1 \partial_x \varphi' - \partial_{\hat{u}} \zeta' + (a_{1h} \partial_x \varphi) \zeta', \\ & \bullet \quad a_2 \partial_x \zeta' + \partial_{\hat{u}} \varphi' + (a_{2h} \partial_x \zeta + b_{2h}) \zeta'. \end{aligned}$$

This observedly has the form depicted in (3-5).

What follows gives sufficient conditions on ζ for (3-6) to hold:

Lemma 3.3 *Suppose that $\zeta \in C^\infty(\mathbb{R} \times I_*)$ is a function with the following properties:*

- *The norm of ζ at any given $(x, \hat{u}) \in \mathbb{R} \times I_*$ is such that $(\hat{u}, \hat{\phi}, h = \zeta(x, \hat{u}))$ is in \mathcal{X} .*
- *$|\zeta|$ has respective limits as $x \rightarrow \pm\infty$, and both limits are less than $\chi_0 + 4e^{-2R}$.*
- *The function on \mathbb{R} given by the rule $x \rightarrow \sup_{\{x\} \times I_*} (|\partial_x \zeta| + |\partial_{\hat{u}} \zeta|)$ limits uniformly to zero as $|x| \rightarrow \infty$.*

Use ζ with a given bounded function φ to define the operator in (3-7). Then the corresponding version of $(\alpha_1, \alpha_2, \mathfrak{b}_1, \mathfrak{b}_2)$ obey the conditions in (3-6). In particular, this occurs if ζ comes from a pair whose corresponding graph in $\mathbb{R} \times \mathcal{X}$ is the Ψ_p -inverse image of a surface in $\mathbb{R} \times \mathcal{H}_{p*}^+$ that is J_p -holomorphic where $|s| \gg 1$ and also obeys the conditions in the first and second bullets of (2-9).

Proof The condition in the first bullet is required for defining D_h . Granted that D_h is well defined, the fact that α_1 and α_2 in (3-5) are positive follows from (3-9) since the functions a_1 and a_2 that appear in (3-4) are positive. The conditions in the second and third bullets imply that the $x \rightarrow \pm\infty$ limits of ζ exist and both are independent of \hat{u} . Moreover, the bound on these limits given by the second bullet imply that these respective values for h with a given value for $(\hat{u}, \hat{\phi})$ define a point in \mathcal{X} . It follows from the top bullet in (1-30) that the functions (a_1, a_2, b) are independent of x where $|x| \gg 1$ and so depend only on the coordinates \hat{u} and h where $|x| \gg 1$. This and the fact that ζ limits to a constant implies that α_1 and α_2 have uniform $x \rightarrow \pm\infty$ limits that are positive functions of \hat{u} . This also implies that \mathfrak{b}_2 has uniform $x \rightarrow \pm\infty$ limits that are functions of \hat{u} . The condition stated in the second bullet of (3-6) follows from the second bullet of Lemma 3.1 via the chain rule.

To prove the fourth bullet in (3-6), use the third bullet of Lemma 3.1 to see that the $x \rightarrow \pm\infty$ limits of \mathfrak{b}_2 are those of b_h , and so given by the function $\hat{u} \rightarrow b_{2h}(\hat{u}, h_\pm)$, where h_\pm are the corresponding $x \rightarrow \pm\infty$ limits of ζ . Meanwhile, the value of h in (1-27) is constant along any given integral curve of v in \mathcal{H}_p^+ and therefore $\{b_{2h}(\hat{u}, h_+)\}_{\hat{u} \in I_*}$ and $\{b_{2h}(\hat{u}, h_-)\}_{\hat{u} \in I_*}$ are the values of b_h along the Ψ_p -inverse image of integral curves of v . To say more, let γ denote the $\hat{u} \in I_*$ part of an integral curve of v in \mathcal{H}_{p*}^+ . The constant value of h on γ and the ϕ coordinate of the $\hat{u} = 0$ point on γ determines γ . If γ is parametrized by the coordinate u on \mathcal{H}_p^+ , then the

coordinate ϕ on γ changes via the rule in (II.2-5). With γ parametrized by \hat{u} , the change in ϕ along γ is given by

$$(3-10) \quad \frac{d\phi}{d\hat{u}} = -\frac{\sqrt{6} \chi \cos \theta}{f(1 - 3 \cos^2 \theta)} \frac{\partial u}{\partial \hat{u}},$$

where f is the function of u given in (1-4). Here, it is understood that θ is determined by u given the constant value of h on γ , and so it is determined by \hat{u} and the constant value of h . This understood, it follows that γ 's version of b_2 is given by

$$(3-11) \quad b_2 = \frac{\partial^2 \phi}{\partial h \partial \hat{u}} = -\frac{\partial}{\partial h} \left(\frac{\sqrt{6} \chi \cos \theta}{f(1 - 3 \cos^2 \theta)} \frac{\partial u}{\partial \hat{u}} \right).$$

Let $\Delta\phi = \phi(\gamma|_{\hat{u}=R+\frac{1}{2} \ln z_*}) - \phi(\gamma|_{\hat{u}=-R-\frac{1}{2} \ln z_*})$. This number depends on γ and so defines a function of the parameter h . The integral of b_2 over the domain I_* is $\frac{d}{dh}(\Delta\phi)$. To compute the latter, use (3-10) to write

$$(3-12) \quad \Delta\phi = -\sqrt{6} \int_{-R-\frac{1}{2} \ln z_*}^{R+\frac{1}{2} \ln z_*} \frac{\chi \cos \theta}{f(1 - 3 \cos^2 \theta)} du.$$

To compute the h -derivative of (3-12), introduce θ_0 to denote the value of θ at the $u = 0$ point along γ . This is determined by h by solving $h|_\gamma = (\chi_0 + 4e^{-2R}) \cos \theta_0 \sin^2 \theta_0$ with the constraint that the solution θ_0 is such that $1 - 3 \cos^2 \theta_0 > 0$. Lemma II.2.2 guarantees a unique solution. Meanwhile, θ along γ is determined at any given value of u by θ_0 via the rule $f(u) \cos \theta \sin^2 \theta = f(0) \cos \theta_0 \sin^2 \theta_0$. This understood, it follows from (3-12) using the chain rule that

$$(3-13) \quad \frac{d}{dh}(\Delta\phi) = -\sqrt{6} \int_{-R-\frac{1}{2} \ln z_*}^{R+\frac{1}{2} \ln z_*} \frac{\chi(1 + 3 \cos^2 \theta)}{f^2(1 - 3 \cos^2 \theta)^3} du.$$

As can be seen, the expression on the right-hand side is negative in all cases.

Consider now the final assertion of the lemma, which concerns the case where ζ comes from a pair (φ, ζ) whose large $|x|$ values define, via Ψ_p and the graph in (3-1), a surface in $\mathbb{R} \times \mathcal{H}_p^+$ that is J -holomorphic and has the asserted large $|s|$ behavior. The convergence condition implies that ζ limits uniformly as $x \rightarrow \infty$ to a constant, and likewise as $x \rightarrow -\infty$. By the same token, the function $\hat{u} \rightarrow \varphi(x, \hat{u})$ also limits uniformly as $x \rightarrow \infty$ to a function of \hat{u} , and likewise as $x \rightarrow -\infty$. Granted these uniform limits, and given that (φ, ζ) obey (3-4) at large $|x|$, the standard elliptic regularity theorems of the sort that can be found in Chapter 6 of Morrey's book [12] will prove that $\partial_x \zeta$ and $\partial_{\hat{u}} \zeta$ converge uniformly to zero as $|x| \rightarrow \infty$. □

3C Proof of Proposition 3.2: the Fredholm assertion

The proof that D is Fredholm has five steps.

Step 1 Let $\|\cdot\|$ denote the norm for \mathbb{L} . An operator such as D from \mathbb{H} to \mathbb{L} has finite-dimensional kernel and closed range if there exists $c \geq 1$ such that the following holds for all elements η in \mathbb{H} :

$$(3-14) \quad \bullet \quad \|D\eta\|^2 \geq c^{-1}\|d\eta\|^2 - c\|\eta\|^2.$$

- If η has support only where $|x| > c$, then $\|D\eta\|^2 \geq c^{-1}\|\eta\|^2$.

A standard argument using the Rellich lemma uses (3-14) to deduce that D has closed range and finite-dimensional kernel. The cokernel of D is isomorphic to the kernel of a certain formal L^2 adjoint which is also a bounded operator from \mathbb{H} to \mathbb{L} . The upcoming Step 5 explains why (3-14) holds for this formal adjoint, and so the cokernel of D is finite-dimensional.

Step 2 To prove what is asserted by the top bullet in (3-14), multiply the square of the left-most expression in (3-5) by α_1^{-1} and the square of the right-most by α_2^{-1} . Integrate the resulting expressions over $\mathbb{R} \times I_*$. Use $M(\eta)$ to denote the result of this integration. This number $M(\eta)$ is relevant because $M(\eta) > c_0^{-1}\|D\eta\|^2$. This understood, the bound that is asserted in the first bullet of the lemma is obtained with the help of an integration by parts to eliminate the term $(\partial_x\varphi'\partial_{\hat{u}}\zeta' - \partial_u\varphi'\partial_x\zeta')$ that appears in the integrand that defines $M(\eta)$. There are no boundary terms from the integration by parts because of the second bullet in (3-7). With this term absent, the desired bound follows directly using the triangle inequality.

Step 3 To see about the second bullet in (3-14), introduce α_{1-} , α_{2-} and \mathfrak{b}_{2-} to denote the respective $x \rightarrow -\infty$ limits of α_1 , α_2 and \mathfrak{b}_2 . Let Q^- denote the quadratic function on $C^\infty(I_*; \mathbb{R}^2)$ that is given by the rule

$$(3-15) \quad (\varphi', \zeta') \rightarrow \int_{I_*} (\alpha_{1-}^{-1}|\partial_{\hat{u}}\zeta'|^2 + \alpha_{2-}^{-1}|\partial_{\hat{u}}\varphi' + \mathfrak{b}_{2-}\zeta'|^2).$$

Restrict this form to the subspace of pairs (φ', ζ') with $\varphi' = 0$ at the boundary of the interval. On this restricted domain, the function Q^- is such that

$$(3-16) \quad Q^-(\varphi', \zeta') \geq c_0 R^{-2} \int_{I_*} (|\varphi'|^2 + |\zeta'|^2).$$

Indeed, this follows directly given that there is no pair (φ', ζ') with $\varphi = 0$ on the boundary of I_* and with $Q^-(\varphi', \zeta') = 0$. To see why no such pair exists, note that if

$Q^-(\varphi', \zeta') = 0$, then ζ' is constant and

$$(3-17) \quad \varphi'(\hat{u}) = \zeta' \int_{-R-\frac{1}{2} \ln z_*}^{\hat{u}} b_{2-}.$$

Since the integral over I_* of b_{2-} is nonzero, the right-hand side is not zero at $\hat{u} = R + \frac{1}{2} \ln z_*$ unless both ζ' and φ' are zero.

There is the analogous quadratic function, Q^+ , on $C^\infty(I_*; \mathbb{R}^2)$ that is defined by the $x \rightarrow \infty$ limits of a_1 , a_2 and b_2 . The latter also dominates what is written on the right-hand side of (3-16) when φ' is zero where $|\hat{u}| = R + \frac{1}{2} \ln z_*$.

Hold on to the Q^- and Q^+ versions of (3-16) for use in a moment.

Step 4 Suppose that $x_1 > 1$ and that $\eta = (\varphi', \zeta')$ has compact support that lies where $x < -x_1$. Integrate by parts as instructed in Step 3, but now write the resulting expression for $M(\eta)$ as

$$(3-18) \quad \int_{x < -x_1} Q^{-1}(\varphi', \zeta') + \int_{\mathbb{R} \times I_*} (a_1 |\partial_x \varphi'|^2 + a_2 |\partial_x \zeta'|^2) + 2 \int_{\mathbb{R} \times I_*} b_{2-} \zeta' \partial_x \zeta' + \epsilon,$$

where $|\epsilon| \leq \Delta(\|d\eta\|^2 + \|\eta\|^2)$ with Δ such that $\lim_{x_1 \rightarrow \infty} \Delta = 0$. Integrate by parts on the right-most integral in (3-18) to see that it is zero. Meanwhile, the left-most integral in (3-18) is no less than $c_0^{-1} R^{-2}(\|\varphi'\|^2 + \|\zeta'\|^2)$. Thus, what is written in (3-18) is greater than

$$(3-19) \quad \|\partial_x \eta\|^2 + c_0^{-1} R^{-2} \|\eta'\|^2 \quad \text{if } x_1 > c.$$

This last bound implies what is asserted by the second bullet in (3-14) for the case when η is supported where $x < -x_1$. But for notation, the same argument using Q^+ proves the second bullet of (3-14) for the case when η is supported where $x > x_1$.

Step 5 Up to a sign, the formal adjoint in question is defined by using integration by parts to rewrite inner products with $D\eta$ using the inner product on \mathbb{L} . To be explicit, the operator sends any given $\eta = (\varphi^\#, \zeta^\#)$ to the element in \mathbb{L} with respective components

$$(3-20) \quad \begin{aligned} &\bullet \quad -a_1 \partial_x \varphi^\# - \partial_{\hat{u}} \zeta^\# - (\partial_x a_1) \varphi^\#, \\ &\bullet \quad -a_2 \partial_x \zeta^\# + \partial_{\hat{u}} \varphi^\# + b_1 \varphi^\# + (b_2 - \partial_x a_2) \zeta^\#. \end{aligned}$$

Use $D^\#$ to denote the operator in (3-20). What follows explains why $D^\#$ obeys the assertions made by the two bullets in (3-14).

An argument much like that used in Step 2 proves that $\|D^\# \eta\|^2 \geq c^{-1} \|d\eta\|^2 - c \|\eta\|^2$. Meanwhile, the assumptions that $|\partial_x a_1|$ and $|\partial_x a_2|$ limit to zero as $|x| \rightarrow \infty$ imply

that $D^\#$ has the same form as that of D at large $|x|$ but for the sign changes in the derivative terms. As a consequence, the argument that proves the second bullet of (3-14) for D proves it for $D^\#$ as well.

3D Proof of Proposition 3.2: the kernel and cokernel

This subsection computes the Fredholm index of D and then its kernel dimension. Both are found equal to zero. The cokernel dimension is therefore zero as well. These computations are done in five steps.

Step 1 This step computes the index of D . This is done by deforming D to an operator whose index is readily computable. The discussion that follows concerns the case when the integral that is described in the fourth bullet of (3-6) is negative. A very much analogous discussion holds when the integral in question is positive.

The deformation is through a family of operators from \mathbb{H} to \mathbb{L} that all have the same schematic form as D . The family is parametrized by $[0, 1]$. Fix $r \in [0, 1]$, then the member parametrized by r sends any given $\eta = (\varphi', \zeta') \in \mathbb{H}$ to the element in \mathbb{L} whose respective components are

$$(3-21) \quad \bullet \quad ((1-r)a_1 + r)\partial_x\varphi' - \partial_{\hat{u}}\zeta' + (1-r)b_1\zeta',$$

$$\bullet \quad ((1-r)a_2 + r)\partial_x\zeta' + \partial_{\hat{u}}\varphi' + ((1-r)b_2 - r)\zeta'.$$

The $r = 0$ member is D and the $r = 1$ member sends η to the element in \mathbb{L} with respective components

$$(3-22) \quad \partial_x\varphi' - \partial_{\hat{u}}\zeta' \quad \text{and} \quad \partial_x\zeta' + \partial_{\hat{u}}\varphi' - \zeta'.$$

The index of D is the same as this $r = 1$ version. To see that the latter has index equal to 0, suppose first that η is such that what is written in (3-22) vanishes. Then φ' obeys the second-order equation $\partial_{\hat{u}}^2\varphi' + \partial_x^2\varphi' - \partial_{\hat{u}}\varphi' = 0$. Keeping in mind that $\varphi' = 0$ where $|\hat{u}| = R + \frac{1}{2} \ln z_*$, and that $|\varphi'|^2$ is integrable, the maximum principle demands that φ' vanish identically. This the case, then ζ' must be constant, and hence zero because $|\zeta'|^2$ is also integrable.

As noted in Step 5 of the preceding subsection, the cokernel of the operator defined by (3-22) is isomorphic to the kernel of latter's version of what is depicted in (3-20). This is the operator that sends any given η to the element in \mathbb{L} with respective components

$$(3-23) \quad -\partial_x\varphi' - \partial_{\hat{u}}\zeta' \quad \text{and} \quad -\partial_x\zeta' + \partial_{\hat{u}}\varphi' - \zeta'.$$

Granted this form for $D^\#$, the same maximum principle argument applies to prove that it lacks a nontrivial square-integrable kernel.

Step 2 This step proves that the kernel of D is trivial given a certain claim whose proof occupies the remaining steps. To this end, suppose that $\eta = (\varphi', \zeta') \in \mathbb{H}$ is such that $D\eta = 0$. Let $\Gamma \subset \mathbb{R} \times I_*$ denote the locus where $\zeta' = 0$. As explained in the subsequent steps, this set is nonempty, and it is either $\mathbb{R} \times I_*$ or it has the structure of a graph with the following properties:

- (3-24) • The interiors of the edges are the components of the locus in Γ where $d\zeta' \neq 0$, and each vertex is a critical point of the map to \mathbb{R}^2 defined by (φ', ζ') .
- Each edge is a C^1 -embedded, closed interval.
 - Each vertex has but a finite number of incident edges. No pair of distinct incident edges have the same outward pointing tangent vector at any given vertex.
 - Each interior vertex has an even number of incident edges; this number is at least 4.
 - Each edge is oriented by the restriction of $d\varphi$, and this is the orientation that is induced on the edge by viewing it as a boundary component of the $\zeta' < 0$ locus.

These last facts are not compatible with the fact that $\varphi' = 0$ where $|\hat{u}| = R + \frac{1}{2} \ln z_*$ and has limit zero as $|x| \rightarrow \infty$ unless $\Gamma = \mathbb{R} \times I_*$, in which case ζ' is everywhere zero and thus so is φ' . To see why Γ cannot be a graph, suppose to the contrary that Γ is described by (3-24). Let $U \subset \mathbb{R} \times I_*$ denote a component of the complement of Γ . Bullets 2–5 of (3-24) imply that ∂U is piecewise smooth, and so any given differential form can be integrated between points on ∂U . Meanwhile, either $\zeta' > 0$ in U or $\zeta' < 0$. In either case, the final bullet implies that $d\varphi'$ is positive on the smooth part of ∂U given a suitable orientation. As a consequence, φ' increases monotonically along ∂U . This is not possible for it precludes an endpoint of any component of ∂U where $|\hat{u}| = R + \frac{1}{2} \ln z_*$, and it precludes a noncompact component of ∂U , and it precludes a component with no boundary. The fact that $\Gamma \neq \emptyset$ precludes the case $U = \mathbb{R} \times I_*$.

Step 3 This step explains why (3-24) describes Γ given that $\Gamma \neq \emptyset$ and $\Gamma \neq \mathbb{R} \times I_*$. To this end, let $\Gamma' \subset \Gamma$ denote the subset where $d\zeta' \neq 0$. This is a smooth, 1-dimensional submanifold in $\mathbb{R} \times I_*$. It follows from (3-5) that $d\varphi' > 0$ on the tangent

line of Γ' if the latter is oriented so that $d\zeta'$ points towards the side where $\zeta' > 0$. It also follows that $d\varphi' = 0$ at the points in $\Gamma \setminus \Gamma'$. Thus $\Gamma - \Gamma'$ is a subset of the set of singular points of the map η to \mathbb{R}^2 . Let $p \in \Gamma \setminus \Gamma'$ and set $\lambda = (\varphi' - \varphi'(p)) + i\zeta'$. This \mathbb{C} -valued function vanishes at p . In addition, equation (3-5) when written for λ has the form $\bar{\partial}\lambda + v\lambda + \mu\bar{\lambda} = 0$ where v and μ are smooth \mathbb{C} -valued functions. Here, $\bar{\partial} = \frac{1}{2}(\partial_x + i\partial_{\hat{u}})$.

To exploit this equation for λ , introduce $w = x - x(p) + i(\hat{u} - \hat{u}(p))$, this being a \mathbb{C} -valued coordinate function for $\mathbb{R} \times I_*$. It follows from the equation for λ using Taylor's theorem with remainder that λ near p must have the form $\lambda = mw^q + \epsilon$ where m is a nonzero complex number, $q \geq 2$ is an integer and $|\epsilon| \leq c_0|w|^{q+1}$. Note that the unique continuation principle implies that q is finite. This depiction of λ implies what is asserted by (3-24) about the interior vertices of Γ . The argument for the boundary vertices is very much the same after using the Schwarz reflection trick from Theorem 24 in [2] to view any given boundary point as an interior point of a domain to which (φ', ζ') extend so as to solve a corresponding extension of (3-5).

Step 4 This step and Step 5 constitute a digression that is needed to explain why $\Gamma \neq \emptyset$. To start, let (a_{1*}, a_{2*}, b_*) denote either the $x \rightarrow \infty$ or $x \rightarrow -\infty$ limit of (a_1, a_2, b_2) . Introduce the operator $L: C^\infty(I_*; \mathbb{R}^2) \rightarrow C^\infty(I_*; \mathbb{R}^2)$ that is defined so as to send $\eta = (\iota, \lambda)$ to

$$(3-25) \quad L\eta = (-\partial_{\hat{u}}\lambda, \partial_{\hat{u}}\iota + b_*\lambda).$$

The relevant domain for L is the subspace in $C^\infty(I_*; \mathbb{R}^2)$ that consists of the pairs (σ, λ) with $\iota = 0$ at the boundary points of I_* . A pair (ι, λ) in this domain is said to be a *weighted eigenfunction* for L if

$$(3-26) \quad -\partial_{\hat{u}}\lambda = E a_{1*}\iota \quad \text{and} \quad \partial_{\hat{u}}\iota + b_*\lambda = E a_{2*}\lambda,$$

with $E \in \mathbb{R}$. The number E is said to be a *weighted eigenvalue*. Straightforward variations of standard arguments show the following: The set of weighted eigenvectors is discrete, has no accumulation points and is unbounded in both directions. What is said in Step 3 implies that 0 is not a weighted eigenvector. Moreover, at most a finite number of weighted eigenvectors share the same weighted eigenvalue. Third, if $\eta = (\iota, \lambda)$ and $\eta' = (\iota', \lambda')$ are weighted eigenvectors with different weighted eigenvalues, then

$$(3-27) \quad \int_{I_*} (a_{1*}\iota\iota' + a_{2*}\lambda\lambda') = 0.$$

Finally, the L^2 completion of the domain of L is spanned by the set of weighted eigenvectors.

Let ϑ denote a minimal spanning set of weighted eigenvectors, here chosen so that if $\eta = (\iota, \lambda) \in \vartheta$, then

$$(3-28) \quad \int_{I_*} (\mathfrak{a}_{1_*} \iota^2 + \mathfrak{a}_{2_*} \lambda^2) = 1.$$

Step 5 To see that $\Gamma \neq \emptyset$, first integrate the right-most equation in (3-5) on each slice of the form $\{x\} \times I_*$ to obtain

$$(3-29) \quad \partial_x \left(\int_{\{x\} \times I_*} \mathfrak{a}_2 \zeta' \right) - \int_{\{x\} \times I_*} (\partial_x \mathfrak{a}_2) \zeta' = - \int_{\{x\} \times I_*} \mathfrak{b}_2 \zeta'.$$

To exploit this identity, suppose now that ζ' is nowhere zero. No generality is lost by assuming that $\zeta' > 0$. If the integrals of the $x \rightarrow \pm\infty$ limits of \mathfrak{b}_2 are negative, then (3-29) is used at points where $x \gg 1$. If the integrals of the $x \rightarrow \pm\infty$ limits of \mathfrak{b}_2 are positive, then (3-29) is used at points where $x \ll -1$. Except for cosmetics, the argument for the latter case is identical to that for the former. Granted this, only the case where the integrals of these limits of \mathfrak{b}_2 are negative is considered in what follows.

To make something of (3-29), use arguments much like those in Section 2.3 of [7] to see that (φ', ζ') can be written for $x \gg 1$ as

$$(3-30) \quad (\varphi', \zeta') = c(e^{Ex}(\iota, \lambda) + \epsilon),$$

where the notation is as follows: First, $c \in (0, \infty)$. Second, $(\iota, \lambda) \in \vartheta$ is an element with negative, weighted eigenvalue, this being E . Third, ϵ is such that the function $x \rightarrow e^{Ex}|\epsilon|$ has limit zero as $x \rightarrow \infty$. Use (3-26) to see that if (σ, λ) is a weighted eigenvector, then λ has transversal zero locus. This understood, it follows from (3-30) that ζ' is positive where $x \gg 1$ if and only if (ι, λ) is such that $\lambda > 0$ at all interior points of I_* .

Granted the preceding, it follows from (3-26) and (3-30) that the weighted eigenvector that appears in (3-30) has $\lambda \geq 0$. It also follows from these equations that

$$(3-31) \quad \int_{\{x\} \times I_*} \mathfrak{b}_2 \zeta' < -r|E|(1 - c^{-1}) \int_{\{x\} \times I_*} \mathfrak{a}_2 \zeta' \quad \text{where } x \gg 1.$$

Use this last bound in (3-29) with the fact that $|\partial_x \mathfrak{a}_2| \rightarrow 0$ as $x \rightarrow \infty$ to deduce that

$$(3-32) \quad \partial_x \left(\int_{\{x\} \times I_*} \mathfrak{a}_2 \zeta' \right) > c^{-1} \int_{\{x\} \times I_*} \mathfrak{a}_2 \zeta' \quad \text{for } x \gg 1.$$

This inequality cannot hold if $|\zeta'|^2$ is integrable. As a consequence, the assumption that $\zeta' > 0$ is untenable.

3E The Banach spaces \mathbb{H}_* and \mathbb{L}_*

The norm that defines the Hilbert space \mathbb{H} does not control the supremum norm of its elements. This being the case, the inverse function theorem that is used in what follows employs a slightly stronger norm. The author learned the latter from Morrey’s book [12]. The definition requires the choice of a positive number that is less than $\frac{1}{100}$. Use ν in what follows to denote this number. This Banach space is denoted by \mathbb{H}_* . It is the closure of the space of pairs that obey (3-7) using a norm that is the sum of the L^2_1 -norm used for \mathbb{H} and a norm that is defined in a moment. The extra term in the norm for \mathbb{H}_* is the square root of the function that assigns to a given pair $\eta' = (\varphi', \zeta')$ the number

$$(3-33) \quad \sup_{(x, \hat{u}) \in \mathbb{R} \times I_*} \sup_{\rho \in (0, 1)} \rho^{-\nu} \int_{\text{dist}(\cdot, (x, \hat{u})) < \rho} |d\eta'|^2.$$

Here, $\text{dist}(\cdot, \cdot)$ denotes the Euclidean distance function. The norm on \mathbb{H}_* is denoted by $\|\cdot\|_{\mathbb{H}_*}$. The lemma below in part justifies the introduction of this space.

Lemma 3.4 *Elements in \mathbb{H}_* are Hölder continuous with exponent $\frac{1}{2}\nu$ and the inclusion map from \mathbb{H}_* into the corresponding Hölder Banach space is continuous. In particular, there exists a constant $\kappa > 1$ that depends only on ν and has the following significance: If $f \in \mathbb{H}_*$, then $|f| \leq \kappa \|f\|_{\mathbb{H}_*}$. In addition, $\lim_{|x| \rightarrow \infty} |f|$ exists and it is zero; thus, elements in \mathbb{H}_* have pointwise uniform limit zero as $|x| \rightarrow \infty$.*

Proof These assertions follow directly from Theorem 3.5.2 in Morrey’s book [12]. \square

A corresponding L^2 version of \mathbb{H}_* is defined to be the closure of the space of compactly supported elements in $C^\infty(\mathbb{R} \times I_*; \mathbb{R}^2)$ using the norm given by the sum of the L^2 -norm and that defined by replacing $d\eta'$ in (3-33) by η' . This last Banach space is denoted in what follows by \mathbb{L}_* . The norm on \mathbb{L}_* when needed is denoted by $\|\cdot\|_{\mathbb{L}_*}$.

Lemma 3.5 *An operator $D: \mathbb{H} \rightarrow \mathbb{L}$ of the sort described by (3-5) and (3-7) maps \mathbb{H}_* to \mathbb{L}_* , and its inverse restricts to \mathbb{L}_* so as to define a bounded linear operator from \mathbb{L}_* to \mathbb{H}_* .*

Proof Given Proposition 3.2, the assertion follows from Theorems 3.5.2 and 5.4.1 of Morrey’s book [12]. \square

By way of some parenthetical remarks about Lemmas 3.4 and 3.5: The spaces \mathbb{H}_* and \mathbb{L}_* are examples of what are now called *Morrey spaces*. Theorem 3.5.2 in Morrey's book [12] is an analog of a Sobolev embedding theorem concerning the regularity (with regards to a Hölder norm) of elements in a Morrey space. (It is now sometimes called "Morrey's lemma".) Theorem 5.4.1 in Morrey's book is an "elliptic regularity" theorem. In the present context, it says that if D of an L_1^2 element is in the Morrey space \mathbb{L}_* , then the element has more regularity than the typical L_1^2 element. In particular, it is \mathbb{H}_* and is, therefore, Hölder continuous. As Morrey's book is, by many accounts, not an easy read (due to the dense notation for the most part), some other references to Morrey spaces are [16; 1].

4 Proof of Proposition 2.2 when $\Delta_p = 0$

The first four subsections prove that there exists at least one pair $(\varphi^{p_0}, \zeta^{p_0})$ that satisfies the requirements of Proposition 2.2. By way of a look ahead, the existence proof uses an open/closed argument for a certain 1-parameter family of $|x| \rightarrow \infty$ asymptotic conditions and $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundary conditions for (3-4). The parameter space is the interval $[0, 1]$; the parameter $\{1\}$ boundary conditions are those required by Proposition 2.2. Meanwhile, the parameter $\{0\}$ case is designed so as to have an obvious solution. Use \mathcal{I} to denote the subset of parameter values in $[0, 1]$ for which (3-4) has a solution with the corresponding asymptotic conditions and boundary conditions. The set \mathcal{I} is proved to be both open and closed. This being the case, and as $\{0\} \in \mathcal{I}$, it follows that $\mathcal{I} = [0, 1]$ and there is at least one pair $(\varphi^{p_0}, \zeta^{p_0})$ that satisfies the requirements of Proposition 2.2.

The final subsection proves that this is the only pair of functions that satisfies all of Proposition 2.2's criteria. This uniqueness proof uses a nonlinear version of the argument that is used Section 3D to prove that the operator D in Proposition 3.2 has trivial kernel.

4A The 1-parameter family

The definition of the family of asymptotic/boundary conditions has three parts.

Part 1 Reintroduce γ_{p_-} and γ_{p_+} from Part 1 of Section 2D. By way of a reminder, these are the respective segments of integral curves of v in $\mathcal{H}_{p_*}^+$ that come from Θ_- and Θ_+ . Each parameter value $\tau \in [0, 1]$ also labels a segment of an integral curve of v that crosses $\mathcal{H}_{p_*}^+$. The corresponding segment is denoted γ_τ . The upcoming definition

uses $\varphi_+^{S_0}$ and $\varphi_-^{S_0}$ to denote the respective $\hat{u} = R + \frac{1}{2} \ln z_*$ and $\hat{u} = -R - \frac{1}{2} \ln z_*$ versions of the function $\varphi^{S_0}(\cdot, z_*)$ that appears in (2-11). The segment γ_τ is the unique integral curve of v in $\mathcal{H}_{p_*}^+$ that obeys the following three constraints:

- (4-1) • The segment γ_τ starts on the surface where $e^{-2(u+R)}(1 - 3 \cos^2 \theta) = z_*$ and it ends on the surface where $e^{2(u-R)}(1 - 3 \cos^2 \theta) = z_*$.
- If $\tau \in (0, 1)$, the ϕ coordinate of γ_τ at its start point is $\varphi_-^{S_0}(x = \frac{2\tau-1}{\tau(1-\tau)}, z_*)$ and the ϕ coordinate of γ_τ at its endpoint is $\varphi_+^{S_0}(x = \frac{2\tau-1}{\tau(1-\tau)}, z_*)$.
 - The segment $\gamma_{\tau=0}$ is γ_{p_-} and the segment $\gamma_{\tau=1}$ is γ_{p_+} .

Lemma II.2.2 supplies the desired segment γ_τ .

The next lemma addresses the continuity and differentiability of the family $\{\gamma_\tau\}_{\tau \in [0,1]}$. This lemma views each integral curve from the family $\{\gamma_\tau\}_{\tau \in [0,1]}$ as a map from $[0, 1]$ into $\mathcal{H}_{p_*}^+$ that pulls a back as a constant multiple of the Euclidean differential.

Lemma 4.1 *The assignment of the point $\gamma_\tau(\sigma)$ to any given pair $(\tau, \sigma) \in [0, 1] \times [0, 1]$ defines a smooth map from $[0, 1] \times [0, 1]$ into $\mathcal{H}_{p_*}^+$.*

Proof By construction, the map is continuous on $[0, 1] \times [0, 1]$ and smooth on $(0, 1) \times [0, 1]$. It follows from the fourth bullet of Proposition 2.1 using the chain rule that the map is smooth on $[0, 1] \times [0, 1]$. By the same token, if the map is continuous up to and along the $\{1\} \times [0, 1]$ boundary, then it is also smooth up to and along $\{1\} \times [0, 1]$. To see about continuity along this boundary, note that $\lim_{\tau=1} \gamma_\tau$ exists, and this limit is a segment of an integral curve of v that crosses $\mathcal{H}_{p_*}^+$. Let γ denote this limit. The issue is whether γ is γ_{p_+} . As explained next, such is the case because $m_p = 0$. To prove this, note that γ_{p_-} and γ_{p_+} concatenate with the $[0, 1]$ -parametrized paths $\tau \rightarrow \gamma_\tau(0)$ and $\tau \rightarrow \gamma_\tau(1)$ to define a piecewise smooth, closed 1-cycle in $\mathcal{H}_{p_*}^+$. The $m_p = 0$ condition implies that this 1-cycle is null-homotopic. Let ι denote this 1-cycle. Meanwhile, the paths γ_{p_-} , γ and the $[0, 1]$ -parametrized paths $\tau \rightarrow \gamma_\tau(0)$ and $\gamma \rightarrow \gamma_\tau(1)$ also concatenate to define a closed 1-cycle in $\mathcal{H}_{p_*}^+$. Use ι' to denote the latter. The 1-cycle ι' is also null-homotopic as it bounds the surface given by the closure of the image of the map from $[0, 1] \times [0, 1]$ that sends (τ, σ) to $\gamma_\tau(\sigma)$. Keeping this in mind, note that γ and γ_{p_+} have the same endpoints, and so the change, $\Delta\phi$, of the angle ϕ along γ must differ from that along γ_{p_+} by an integer multiple of 2π . This integer is zero if and only if ι and ι' are homotopic. This being the case, it follows by Lemma II.2.2 that $\gamma = \gamma_{p_+}$. \square

Part 2 Reintroduce the nonincreasing function, $\chi: \mathbb{R} \rightarrow [0, 1]$ which has value 1 on $(-\infty, 0]$ and value 0 on $[1, \infty)$. Given $\tau \in [0, 1]$, introduce x_τ to denote $\frac{2\tau-1}{\tau(1-\tau)}$ and

then define the function $\hat{x}_\tau: \mathbb{R} \rightarrow (-\infty, x_\tau)$ by the rule

$$(4-2) \quad x \rightarrow \hat{x}_\tau(x) = x\chi(x - x_\tau + 3) + x_\tau(1 - \chi(x - x_\tau + 3)).$$

This definition is such that $\hat{x}_\tau = x$ for $x < x_\tau - 3$ and $\hat{x}_\tau = x_\tau$ for $x > x_\tau + 1$. The derivative of this function \hat{x}_τ is nonnegative and bounded from above by 4. Its derivatives to any given order greater than 1 also enjoy τ -independent bounds.

Part 3 What follows are the parameter $\tau \in [0, 1]$ asymptotic/boundary conditions for (3-4):

$$(4-3) \quad \begin{aligned} &\bullet \lim_{x \rightarrow -\infty} (\varphi, \zeta)|_x = (\phi, \hat{h})|_{\gamma_{p-}} \text{ and } \lim_{x \rightarrow \infty} (\varphi, \zeta)|_x = (\phi, \hat{h})|_{\gamma_\tau}. \\ &\bullet \varphi(\cdot, \hat{u} = -R - \frac{1}{2} \ln z_*) = \varphi_-^{S_0}(\hat{x}_\tau(\cdot), z_*) \text{ and } \varphi(\cdot, \hat{u} = R + \frac{1}{2} \ln z_*) = \varphi_+^{S_0}(\hat{x}_\tau(\cdot), z_*). \end{aligned}$$

To say more about these conditions, note that the $\tau = 1$ version of the top bullet in (4-3) reproduces the first two bullets in (2-9), and the $\tau = 1$ version of the bottom bullet in (4-3) reproduces (2-12). Meanwhile, the $\tau = 0$ version of (4-3) demands that

$$(4-4) \quad \begin{aligned} &\bullet \lim_{x \rightarrow -\infty} (\varphi, \zeta)|_x = (\phi, \hat{h})|_{\gamma_{p-}} \text{ and } \lim_{x \rightarrow \infty} (\varphi, \zeta)|_x = (\phi, \hat{h})|_{\gamma_{p-}}, \\ &\bullet \varphi(x, \hat{u} = -R - \frac{1}{2} \ln z_*) = \phi(\gamma_{p-}|_{\hat{u} = -R - \frac{1}{2} \ln z_*}) \text{ and } \varphi(x, \hat{u} = R + \frac{1}{2} \ln z_*) = \phi(\gamma_{p-}|_{\hat{u} = R + \frac{1}{2} \ln z_*}). \end{aligned}$$

Note in particular that the equations in (3-4) with the boundary conditions in (4-4) are solved by the x -independent pair (φ, ζ) with $\varphi(\hat{u}) = \phi(\gamma_{p-}|_{\hat{u}})$ and ζ the constant function $\zeta = \hat{h}(\gamma_{p-})$. This is to say that the Ψ_p -image of the corresponding image of (3-1) is the J -holomorphic surface $\mathbb{R} \times \gamma_{p-}$.

4B Proof that \mathcal{I} is open

Let $\mathcal{I} \subset [0, 1]$ denote the set of parameters for which (3-4) has a solution that obeys the given parameter's version of (4-3). As noted at the end of the previous subsection, the set \mathcal{I} contains 0, and so it is not empty. This subsection proves that \mathcal{I} is open. The argument for this has four parts.

Part 1 For $\tau \in [0, 1]$, define the pair of functions $(\varphi_\tau, \zeta_\tau)$ on $\mathbb{R} \times \mathcal{I}_*$ using the rule

$$(4-5) \quad \begin{aligned} &\bullet \varphi_\tau(x, \hat{u}) = \chi(R + \frac{1}{2} \ln z_* + \hat{u})\varphi_-^{S_0}(\hat{x}_\tau(x), z_*) \\ &\quad + \chi(R + \frac{1}{2} \ln z_* - \hat{u})\varphi_+^{S_0}(\hat{u}_\tau(x), z_*) \\ &\quad + (1 - \chi(R + \frac{1}{2} \ln z_* + \hat{u}))\chi(x)\phi(\gamma_{p-}|_{\hat{u}}) \\ &\quad + (1 - \chi(R + \frac{1}{2} \ln z_* - \hat{u}))\chi(-x)\phi(\gamma_\tau|_{\hat{u}}). \\ &\bullet \zeta_\tau(x, \hat{u}) = \chi(x)\hat{h}(\gamma_{p-}) + \chi(-x)\hat{h}(\gamma_\tau). \end{aligned}$$

Here, the notation has $\phi(\gamma|\hat{u})$ with $\gamma = \gamma_{p-}$ or $\gamma = \gamma_\tau$ denoting the lift to \mathbb{R} of the coordinate ϕ on γ 's intersection with the Ψ_p -image of the $\hat{u} \in \mathcal{I}_*$ slice of $\mathbb{R} \times \mathcal{X}$. Note in this regard that γ has transversal intersection with this slice, this being a consequence of what is said by the second bullet in (1-30). The lift $\phi(\gamma|\hat{u})$ is chosen so that its value at the $\hat{u} = -R - \frac{1}{2} \ln z_*$ start point of γ is that of the function $\varphi_-^{S_0}$. With this choice, the value of this lift at the $\hat{u} = R + \frac{1}{2} \ln z_*$ endpoint of γ is that of $\varphi_+^{S_0}$. Note also that the function h is constant on the integral curves of v in \mathcal{H}_{p*}^+ ; what is written as $h(\gamma)$ in the lower bullet of (4-5) is the constant value of h on γ .

The pair $(\varphi_\tau, \zeta_\tau)$ obeys the parameter τ boundary condition given by (4-3).

Part 2 Reintroduce the Banach space \mathbb{H}_* from Section 3E. Let $\mathbb{B}_* \subset \mathbb{H}_*$ denote a small radius ball about the origin, chosen so that elements in \mathbb{B}_* have pointwise norm bounded by $c_0^{-1} \chi_0$. The norm is chosen so that any given (φ', ζ') in \mathbb{B}_* has $|\zeta'| \leq c_0^{-1} \chi_0$ at all points in $\mathbb{R} \times \mathcal{X}$. Lemma 3.4 supplies such a ball. Reintroduce the Banach space \mathbb{L}_* from Section 3E as well.

The rule that follows defines a map $[0, 1] \times \mathbb{B}_*$ to \mathbb{L}_* if \mathbb{B}_* has small radius. The desired map sends any given element $(\tau, (\varphi', \zeta'))$ to the pair of functions in \mathbb{L}_* with respective components

$$(4-6) \quad \bullet \quad a_1 \partial_x(\hat{\varphi} + \varphi_\tau) - \partial_{\hat{u}}(\hat{\zeta} + \zeta_\tau),$$

$$\bullet \quad a_2 \partial_x(\hat{\zeta} + \varphi_\tau) + \partial_{\hat{u}}(\hat{\varphi} + \varphi_\tau) + b.$$

To say more about the notation, the functions (a_1, a_2, b) are viewed as functions on $\mathbb{R} \times I_*$ that depend implicitly on $\zeta = \hat{\zeta} + \zeta_\tau$. As in (3-4), their values at any given $(x, \hat{u}) \in \mathbb{R} \times I_*$ are obtained from an eponymous set of $\hat{\varphi}$ -independent functions on \mathcal{X} by evaluating the latter at the point $(\hat{u}, h = \zeta(x, \hat{u}))$. The size constraint on the radius \mathbb{B}_* is needed to guarantee that the $(\hat{\varphi} + \varphi_\tau, \hat{\zeta} + \zeta_\tau)$ version of (3-1) defines a point in $\mathbb{R} \times \mathcal{X}$.

The map defined by (4-6) is denoted in what follows by \mathcal{F} . This map is designed so that any given $(\tau, (\hat{\varphi}, \hat{\zeta})) \in [0, 1] \times \mathbb{B}_*$ version of $(\varphi = \hat{\varphi} + \varphi_\tau, \zeta = \hat{\zeta} + \zeta_\tau)$ obeys (3-4) and the parameter τ version of (4-3) if and only $\mathcal{F}(\tau, (\hat{\varphi}, \hat{\zeta})) = 0$.

Part 3 The next lemma summarizes the contents of this step.

Lemma 4.2 Fix $\tau \in \mathcal{I}$ and let (φ, ζ) denote a corresponding solution to (3-4) with boundary values given by the parameter τ version of (4-3). There exists a neighborhood

$\mathcal{I}_\tau \subset [0, 1]$ of τ and a continuous map from \mathcal{I}_τ to \mathbb{B}_* of the following sort: Given $\tau' \in \mathcal{I}_\tau$, use $(\hat{\varphi}, \hat{\zeta}) \in \mathbb{B}_*$ to denote the corresponding element. Then:

- $\mathcal{F}(\tau', (\hat{\varphi} + \varphi_{\tau'}, \zeta')) = 0$ and so $(\hat{\varphi} + \varphi_{\tau'} \hat{\zeta} + \zeta_\tau)$ solves (3-4) with the parameter τ asymptotic/boundary conditions from (4-3).
- This solution for $\tau' = \tau$ is the given pair (φ, ζ) .

Proof It follows from Lemma 3.4 with the fourth bullet of Proposition 2.1 that \mathcal{F} defines a smooth map from $[0, 1] \times \mathbb{B}_*$ to \mathbb{L}_* and that any $\tau \in [0, 1]$ version of $\mathcal{F}(\tau, \cdot)$ defines a smooth map from \mathbb{B}_* to \mathbb{L}_* whose derivatives to any given order are bounded uniformly as τ varies in $[0, 1]$. The proof that \mathcal{F} is smooth on $[0, 1] \times \mathbb{B}_*$ invokes the fourth bullet of Proposition 2.1 to establish that the maps $\tau \rightarrow \varphi_{-}^{S_0}(\hat{x}_\tau(\cdot), z_*)$ and $\tau \rightarrow \varphi_{+}^{S_0}(\hat{x}_\tau(\cdot), z_*)$ from $[0, 1]$ to $C^\infty(\mathbb{R})$ are smooth on the interval $[0, 1]$. It follows from Lemma 3.5 that the differential along the \mathbb{B}_* component of $[0, 1] \times \mathbb{B}_*$ is an isomorphism from \mathbb{H}_* to \mathbb{L}_* . These last facts with the inverse function theorem prove the lemma. \square

Part 4 Granted that the solutions given by Lemma 3.5 are smooth, it then follows that the set \mathcal{I} is an open subset of $[0, 1]$. Meanwhile, the fact that these solutions are smooth follows using elliptic regularity arguments of the sort that can be found in Chapter 6 of [12]. Note in this regard that the equations in (3-4) are linear with constant coefficients on the part of $\mathbb{R} \times I_*$ where $|\hat{u}| > R + \frac{1}{2} \ln z_* + \ln \delta$. This being the case, standard boundary regularity arguments for the Laplace equation can be employed to prove that the solutions are smooth along the boundary of $\mathbb{R} \times I_*$.

4C Proof that \mathcal{I} is closed

The assertion that the set $\mathcal{I} \subset [0, 1]$ is a closed set is a consequence of the upcoming Lemma 4.3.

Lemma 4.3 uses the following notation: given a positive integer k and a function α on $\mathbb{R} \times I_*$, the lemma has $\nabla^{(k)}\alpha$ denoting the tensor of k^{th} -order partial derivatives of α .

Lemma 4.3 *There exists a purely S -dependent (or \mathcal{K} -compatible) constant $\kappa > 1$ with the following significance: Define the geometry of Y with $\delta \leq \kappa^{-1} z_*$. Then the space of solutions to (3-4) with asymptotic and boundary conditions given by versions of (4-3) is sequentially compact in the strong C^∞ sense. To elaborate, let $\{(\tau_n, (\varphi_n, \zeta_n))\}_{n=1,2,\dots}$ denote a sequence such that $\tau_n \in [0, 1]$ and such that*

(φ_n, ζ_n) is a solution to (3-4) with asymptotic/boundary conditions given by the τ_n -version of (4-3). There exists $\tau \in [0, 1]$ and a solution, (φ, ζ) , to (3-4) with asymptotic/boundary conditions given by the τ version of (4-3). Moreover, there exists a subsequence from the sequence (hence renumbered consecutively) such that $\{\tau_n\}_{n=1,2,\dots}$ converges to τ , and such that

$$\lim_{n \rightarrow \infty} \sup_{(x, \hat{u}) \in \mathbb{R} \times I_*} (|\nabla^{(k)}(\varphi - \varphi_n)| + |\nabla^{(k)}(\zeta - \zeta_n)|) = 0$$

for any given positive integer k .

Proof Suppose to start that the conditions stated in the subsequent equation are satisfied if $\delta < \kappa z_*$ for some purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$.

- (4-7) • There exists $\kappa_* > 1$ such that $1 - 3 \cos^2 \theta > \kappa_*^{-1}$ on the Ψ_p -image of any the graph that is defined via (3-1) by any given solution to (3-4) with asymptotic/boundary conditions given by any given $\tau \in [0, 1]$ version of (4-3).
- Given $\varepsilon > 0$, there exists $\kappa_\varepsilon > 1$ such that the following is true: Let (φ, ζ) denote a solution to (3-4) with asymptotic/boundary conditions given by some $\tau \in [0, 1]$ version of (4-3). Then $(\varphi, \zeta)|_x$ defines a pair of functions on I_* that differs by less than ε from its respective $x \rightarrow -\infty$ and $x \rightarrow \infty$ limit when $x < -\kappa_\varepsilon$ and $x > \kappa_\varepsilon$.

If (4-7) holds, then by-now standard elliptic regularity arguments as can be found in Chapter 6 of [12] prove Lemma 4.3. Note in this regard that the original set of three $\hat{\phi}$ -independent functions on \mathcal{X} that are used to define the (x, \hat{u}) -dependent coefficient functions (a_1, a_2, b_2) in (3-4) have uniformly bounded derivatives to any given order on the Ψ_p -inverse of any subset of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ where there is a positive lower bound for $1 - 3 \cos^2 \theta$. Note also that (3-4) is a linear equation with constant coefficients on the $|\hat{u}| > R + \frac{1}{2} \ln z_* + \ln \delta$ part of $\mathbb{R} \times I_*$ when written using the variables $(x, z = e^{-2(R-\hat{u})})$ on the positive \hat{u} part and $(-x, z = e^{-2(R+\hat{u})})$ on the negative \hat{u} part. This being the case, standard boundary regularity arguments for the Cauchy-Riemann equations can be employed to prove that the solutions are smooth along the boundary of $\mathbb{R} \times \mathcal{I}_*$.

Given what was just said, it remains to prove that (4-7) holds. This is done in five steps.

Step 1 A key input is a bound for the integral of the 2-form w over the Ψ_p -image in $\mathbb{R} \times \mathcal{H}_{p_*}^+$ of the graph of a solution to (3-4) with boundary values given by some parameter $\tau \in [0, 1]$ version of (4-3). The next lemma is used to derive such a bound. It has a second use in a subsequent step.

Lemma 4.4 *There exists a purely S -dependent (or \mathcal{K} -compatible) constant $\kappa \geq 1$ with the following significance: Fix $c \in (2\delta^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ and $z_{**} \in (\delta^2, z_*)$. Suppose that (φ, ζ) solves (3-4), obeys a given $\tau \in [0, 1]$ version of (4-3), and is such that $|\zeta| \leq c$ on the slice of $\mathbb{R} \times I_*$ where $|\hat{u}| = \mathbb{R} + \frac{1}{2} \ln z_{**}$. Then $|\zeta| \leq c + \kappa(z_* - z_{**})$ where $R + \frac{1}{2} \ln z_{**} \leq |\hat{u}| \leq R + \frac{1}{2} \ln z_*$.*

It is important to keep in mind for the subsequent applications of Lemma 4.4 that the constant κ from this lemma depends neither on δ_* nor on z_* when the latter are small.

Proof of Lemma 4.4 The argument that follows establishes the asserted upper bound for the $\hat{u} = R + \frac{1}{2} \ln z_*$ boundary given that $|\zeta| \leq c$ where $\hat{u} = R + \frac{1}{2} \ln z_{**}$. A completely analogous argument does the trick for the other boundary component of $\mathbb{R} \times I_*$. To start, write (φ, ζ) in terms of the coordinate x and a coordinate $z = e^{-2(R-\hat{u})}$. The \mathbb{C} -valued function $\varphi + i\zeta$ is a holomorphic function of $x + iz$ where $z \in [z_{**}, z_*]$ and, as a consequence, the function ζ is annihilated by the operator $\partial_x^2 + \partial_z^2$. Keep this fact in mind. Now define

$$(4-8) \quad R = 100 \sup_{\mathbb{R}} |\partial_x \varphi_+^{S_0}(\cdot, z_*)|.$$

This constant is purely S -dependent (or \mathcal{K} -compatible). With R in hand, use w to denote x -independent function on $\mathbb{R} \times [z_{**}, z_*]$ given by the rule $z \rightarrow w(z) = c + R(z - z_{**})$. This function is also harmonic. Its value where $z = z_{**}$ is greater than that of $|\zeta|$ and its $x \rightarrow \pm\infty$ limits are greater than those of $|\zeta|$. Meanwhile, its z -derivative where $z = z_*$ is greater than $|\partial_z \zeta|$ where $z = z_*$ because the Cauchy-Riemann equations identify $\partial_z \zeta$ with $\partial_x \varphi$, and φ where $z = z_*$ is given by $\varphi_+^{S_0}$ via (4-3). These various upper bounds with the maximum principle imply that $w \geq |\zeta|$ on the whole strip $\mathbb{R} \times [z_{**}, z_*]$. □

Step 2 This step states and then proves the desired bound on the integral of w .

Lemma 4.5 *There exists a purely S -dependent (or \mathcal{K} -compatible) $\kappa \geq 1$, and there exists $\kappa_\Theta \geq 1$ that depends only on S and Θ_\pm , and these have the following significance: Suppose that $\tau \in [0, 1]$ and that (φ, ζ) is a solution to the corresponding version of (3-4). Let $\mathbb{C} \subset \mathbb{R} \times \mathcal{H}_{p_*}^+$ denote the Ψ_p -image of the graph of (φ, ζ) . Then*

- $\int_C w \leq \kappa \delta_*^2,$
- $\int_{C \cap ([s_0, s_0+1] \times \mathcal{H}_{p_*}^+)} ds \wedge \hat{a} \leq \kappa_\Theta$ for all $s_0 \in \mathbb{R}.$

As in the case of Lemma 4.4, the constant κ supplied by Lemma 4.5 depends neither on δ nor on z_* when the latter are small.

Proof of Lemma 4.5 The fact that the integral of w over the $s \ll -1$ part of C is finite is proved in the next paragraph. But for obvious notational changes, the same argument proves the finiteness of the integral over the $s \gg 1$ part of C .

Let $\hat{u} \rightarrow (\varphi_{\gamma_{p-}}, \zeta_{\gamma_{p-}})|_{\hat{u}}$ denote the pair whose graph is the Ψ_p -inverse image of the J -holomorphic cylinder $\mathbb{R} \times \gamma_{p-}$. Note in this regard that $\zeta_{\gamma_{p-}}$ is the constant value of the function \hat{h} on γ_{p-} . Standard elliptic regularity theorems of the sort found in Chapter 6 of [12] in conjunction with (4-3) and (3-4) prove the following: fix $\varepsilon > 0$, then there exists $s_\varepsilon \geq 1$ such that any $s \geq s_\varepsilon$ version of $(\varphi, \zeta)|_s$ has C^1 distance at most ε from $(\varphi_{\gamma_{p-}}, \zeta_{\gamma_{p-}})$. This fact with (1-6), the fact that w is nonnegative on $\wedge^2 TC$ and Stokes' theorem imply the finiteness claim.

What follows next explains why the bound given by Lemma 4.5's first bullet holds. To start, use the just-described application of Stokes' theorem to identify the integral of w over C with the sum of the following two expressions:

$$(4-9) \bullet \int_{\gamma_\tau} \hat{h} d\phi - \int_{\gamma_{p-}} \hat{h} d\phi + \int_{\mathbb{R} \times \{\hat{u} = R + \frac{1}{2} \ln z_*\}} \zeta d\varphi - \int_{\mathbb{R} \times \{\hat{u} = -R - \frac{1}{2} \ln z_*\}} \zeta d\varphi.$$

$$\bullet \int_{\gamma_\tau} \chi(1 - 3 \cos^2 \theta) du - \int_{\gamma_{p-}} \chi(1 - 3 \cos^2 \theta) du.$$

To bound the left-most two terms in the top bullet, note first that the function ζ is constant on the integral curves of v in $\mathcal{H}_{p_*}^+$, its value being that of the function \hat{h} depicted in (1-27). As a consequence, the difference between these two terms can be written as

$$(4-10) \qquad \hat{h}(\gamma_\tau) \Delta\phi_{\gamma_\tau} - \hat{h}(\gamma_{p-}) \Delta\phi_{\gamma_{p-}},$$

where $\Delta\phi_{(\cdot)}$ is the change in the coordinate ϕ along the indicated integral curve. It follows from (1-4) that the two values of \hat{h} are bounded in absolute value by $c_0 \chi_0$. Meanwhile, the two values of $\Delta\phi_{(\cdot)}$ are determined by S and the $\Delta_p = 0$ constraint. It follows as a consequence that the two left-most terms in (4-9) are bounded by a purely S -dependent (or \mathcal{K} -compatible) multiple of χ_0 .

To bound the two right-most terms in the top bullet, keep in mind that $|\zeta| \leq \frac{4}{3\sqrt{3}} \delta_*^2$ in any event. This the case, the two right-most terms are bounded in absolute value by $c_0 \delta_*^2$ times the integral over \mathbb{R} of the function $x \rightarrow |\partial_x \varphi_-^{S_0}(x, z_*)| + |\partial_x \varphi_+^{S_0}(x, z_*)|$. The latter integrals are bounded by a purely S -dependent (or \mathcal{K} -compatible) constant, this by virtue of the fourth bullet in Proposition 2.1.

Turn now to the second bullet in (4-9). Since each of the two integrals that appear have size on the order of $\delta_*^2 R$, the task at hand is to explain why the magnitude of their

difference has an R -independent bound. To start this task, let γ denote for the moment a closed integral curve of v in the $|\hat{u}| \leq R + \frac{1}{2} \ln z_*$ part of $\mathcal{H}_{p_*}^+$. With the function u on \mathcal{H}_p^+ viewed as an affine parameter along γ , the restriction of the function θ to γ becomes a function of u , to be denoted by θ_γ . For the family $\{\gamma_\tau\}$, the corresponding family of functions of u depends smoothly on τ . This τ -dependence can be exploited (using the fundamental theorem of calculus) to bound the absolute value of what is written in the second bullet of (4-9) to be at most

$$(4-11) \quad 6 \int_0^1 \left(\int_{-R-\ln z_*}^{R+\ln z_*} \chi \left| \cos(\theta_{\gamma_\tau}(u)) \sin(\theta_{\gamma_\tau}(u)) \frac{d\theta_{\gamma_\tau}(u)}{d\tau} \right| du \right) d\tau.$$

The goal now is to bound the absolute value of the $d\theta_{\gamma_\tau}(u)/d\tau$ term in (4-11) by a number that is $\mathcal{O}(\frac{1}{R})$ so as to offset the $\mathcal{O}(R)$ length of the range of integration of the variable u in (4-11).

To do this, first use the third bullet of (1-9) to see that θ_{γ_τ} at any given value of u is determined by its value at $u = 0$. In particular, the chain rule can be invoked to write $d\theta_{\gamma_\tau}(u)/d\tau = (d\theta_\gamma(u)/d\theta_\gamma(0)|_{\gamma=\gamma_\tau}) d\theta_{\gamma_\tau}(0)/d\tau$. The factor of $d\theta_\gamma(u)/d\theta_\gamma(0)|_{\gamma=\gamma_\tau}$ is given below:

$$(4-12) \quad \sin(\theta_{\gamma_\tau}(u)) \frac{d\theta_\gamma(u)}{d\theta_\gamma(0)} \Big|_{\gamma=\gamma_\tau} = \sin(\theta_{\gamma_\tau}(0)) \frac{f(0)}{f(u)} \frac{1 - 3 \cos^2(\theta_{\gamma_\tau}(0))}{1 - 3 \cos^2(\theta_{\gamma_\tau}(u))}.$$

(This formula is a consequence of the third bullet in (1-9).) Noting that $|\cos \theta_{\gamma_\tau}|$ has its maximum at $u = 0$ (due to the third bullet of (1-9)), and that f has its minimum at $u = 0$ (see (1-4)), it follows as a consequence that $|d\theta_{\gamma_\tau}(u)/d\theta_{\gamma_\tau}(0)|$ has an R -independent upper bound.

Granted what is said in the preceding paragraph, a bound for (4-11) can be had given a bound for factor $|d\theta_{\gamma_\tau}(0)/d\tau|$. To obtain a suitable bound for this, let γ again denote a closed integral curve of v in the $|\hat{u}| \leq R + \frac{1}{2} \ln z_*$ part of $\mathcal{H}_{p_*}^+$. It follows from the last three bullets of (1-9) that γ is determined up to a $\phi \rightarrow \phi + \text{constant}$ rotation of the S^2 factor of \mathcal{H}_p^+ by the change in ϕ between its endpoints. This change is denoted by $\Delta\phi_\gamma$. Thus, $\theta_\gamma(0)$ is determined by $\Delta\phi_\gamma$, and therefore, by the chain rule,

$$\frac{d\theta_{\gamma_\tau}(0)}{d\tau} = \left(\frac{d\theta_\gamma(0)}{d(\Delta\phi_\gamma)} \Big|_{\gamma=\gamma_\tau} \right) \frac{d(\Delta\phi_{\gamma_\tau})}{d\tau}.$$

Since $\Delta\phi_{\gamma_\tau}$ is determined directly by S and the $\Delta_p = 0$ constraint, so is $d(\Delta\phi_{\gamma_\tau})/d\tau$. In particular, the norm of $d(\Delta\phi_{\gamma_\tau})/d\tau$ is bounded by a purely S -dependent (or \mathcal{K} -compatible) constant. The other factor in the chain rule formula for $d\theta_{\gamma_\tau}(0)/d\tau$ is

$d\theta_\gamma(0)/d(\Delta\phi_\gamma)|_{\gamma=\gamma_\tau}$. This is the inverse of $d(\Delta\phi_\gamma)/d\theta_\gamma(0)|_{\gamma=\gamma_\tau}$, which is

$$(4-13) \quad \frac{d(\Delta\phi_\gamma)}{d\theta_\gamma(0)} \Big|_{\gamma=\gamma_\tau} = -\sqrt{6} f(0) (1 - 3 \cos^2(\theta_{\gamma_\tau}(0)) \sin(\theta_{\gamma_\tau}(0))) \\ \times \int_{[-R-\ln\delta, R+\ln\delta]} \frac{\chi(u)}{f(u)^2} \frac{1 + 3 \cos(\theta_{\gamma_\tau}(u))}{(1 - 3 \cos^2(\theta_{\gamma_\tau}(u)))^3} du.$$

By way of an explanation, this formula is computed from the formula for $\Delta\phi$ in the fourth bullet of Lemma II.2.2 (which is obtained by integrating the formula in the fifth bullet of (1-9)) by differentiating with respect to $\theta_\gamma(0)$ and then using (4-12) after exchanging the order of integration with respect to u and differentiation.

The important takeaway from (4-13) is that the integral that appears on the right-hand side of (4-13) no smaller than $c_0^{-1} R$. Therefore, the norm of $d\theta_\gamma(0)/d(\Delta\phi_\gamma)|_{\gamma=\gamma_\tau}$ is *no greater* than $c_0 \frac{1}{R}$. Granted this $c_0 \frac{1}{R}$ upper bound, it follows from what was said about the other factors in the chain rule decomposition of $d\theta_{\gamma_\tau}(u)/d\tau$ that the norm of $d\theta_{\gamma_\tau}(u)/d\tau$ is bounded by $c_* \frac{1}{R}$ with c_* being purely S -dependent (or \mathcal{K} -compatible.) Because of this, the expression in (4-11) and thus the expression in (4-9)'s second bullet is at most $c_0 c_* \chi_0$.

Minor cosmetic changes to the arguments from Step 4 of the proof of Proposition II.5.1 in Section II.5B give the bound on the integral of $ds \wedge \hat{a}$. Note in this regard that the integration by parts used in these arguments has no boundary contributions from the $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundary of C because \hat{a} near the boundary is the 1-form df . \square

Step 3 This step states and then proves a refined bound on $|\zeta|$ when (φ, ζ) obey (3-4) with boundary values given by (4-3).

Lemma 4.6 *There exists a purely S -dependent (or \mathcal{K} -compatible) constant $\kappa \geq 1$ such that if $\delta^2 < \kappa^{-1} z_*$, then the following is true: Let (φ, ζ) denote a solution to (3-4) that obeys a given $\tau \in [0, 1]$ version of (4-3). Then $|\zeta|$ is bounded by κz_* where $|\hat{u}| > R + \frac{1}{2} \ln z_* - 8$.*

It is important to keep in mind that the version of κ from Lemma 4.6 does not depend on δ nor on z_* .

Proof of Lemma 4.6 Suppose that the lemma is false so as to derive some nonsense. Then there exists a sequence $\{D_n, (\hat{\Theta}_{n-}, \hat{\Theta}_{n+}), (\tau_n, (\varphi_n, \zeta_n))\}_{n=1,2,\dots}$ of the following sort: Each index n version of D_n is a data set of the form $(\delta_n, \chi_{0n}, R_n, J_n)$ suitable for defining the geometry of Y and (3-4), and such that $\delta_n < \frac{1}{n} z_*$. It is assumed here that all $n \in \{1, 2, \dots\}$ versions of D_n use the same data from M ; in particular, they define the

same almost complex structure and pseudogradient vector field \mathbf{v} on the complement of the their attaching handles. The pair $(\hat{\Theta}_{n-}, \hat{\Theta}_{n+})$ is an $\{m_p = 0\}_{p \in \Lambda}$ element from $\hat{\mathcal{Z}}^S$. Meanwhile, the pair (φ_n, ζ_n) obeys the D_n version of (3-4) with the asymptotic/boundary conditions determined via (4-3) by the data $(\Theta_- = \Theta_{n-}, \Theta_+ = \Theta_{n+})$ and $\tau = \tau_n$. In addition, there are points where $|\hat{u}| > R + \frac{1}{2} \ln z_* - 8$ at which $|\zeta_n| > 4\kappa_* z_*$ with κ_* here denoting the version of κ given by Lemma 4.4. Note that any such point lies in the radius $e^{-4} z_*^{1/2}$ coordinate ball centered at one or the other of the critical points from \mathfrak{p} . By passing to a subsequence and renumbering, arrange that this occurs for each n in the radius $e^{-4} z_*^{1/2}$ ball about a fixed critical point from \mathfrak{p} . The argument that follows discusses the case when the critical point in question has index 1. But for some sign changes, the same argument works for the index 2 critical point. Let p denote the index 1 critical point in question.

Let $\{C_{p0n}\}_{n=1,2,\dots}$ denote the corresponding sequence of submanifolds. Note that various index n version different data sets to define the geometry of Y and almost complex structure on $\mathbb{R} \times Y$. Even so, the following is true: Fix an integer N , and then all $n > N$ versions of the almost complex structure agree on the part of the radius δ_* coordinate ball centered at p where the radius is greater than $(N^{-1} z_*)^{1/2}$. This almost complex structure is denoted by J . Let z denote the function $e^{-2(R-|\hat{u}|)}$ and let U denote the part of the radius δ_* coordinate ball centered on p where the radius is greater than $e^{-100} z_*^{1/2}$ and less than z_* . Granted what was just said, each $n \geq c_0$ version of C_{p0n} intersects $\mathbb{R} \times U$ as a properly embedded, J -holomorphic submanifold. Let C_n denote this part of C_{p0n} .

For each n , let s_n denote a value for s of a point in C_n that corresponds to a point where $|\zeta_n| > 4\kappa_* z_*$. Translate C_n along the \mathbb{R} factor of $\mathbb{R} \times U$ by $-s_n$ so that such a point in C_n sits where $s = 0$ in the new submanifold. Let C'_n denote this new submanifold.

Let $s_0 \in \mathbb{R}$. Then the integral of $ds \wedge \hat{a}$ over the $s \in [s_0, s_0 + 1]$ part of C'_n is bounded by z_* since $\hat{a} = df$ here and C'_n has intersection number 1 or 0 with each constant (s, f) level set with $f \in (1, 1 + z_*)$. This understood, use Lemma 4.5 with Proposition II.5.5 to obtain a subsequence of $\{C'_n\}_{n=1,2,\dots}$ that converges on compact subsets in $\mathbb{R} \times U$ in the manner described by Proposition II.5.5. Let ϑ denote the resulting set of pairs consisting of an irreducible, J -holomorphic subvariety and positive integer weight.

As explained in a moment, the set ϑ must contain a pair whose subvariety component sits entirely in the $1 - 3 \cos^2 \theta = 0$ locus and is therefore the intersection of U with some element from Proposition II.3.3's moduli space \mathcal{M}_1 . Granted for the moment

that ϑ contains such a subvariety, it then follows from the manner of convergence described in Proposition II.5.5 that all large n versions of C_{p0n} must contain a loop that represents a nonzero multiple of the generator of $H_1(\mathbb{R} \times \mathcal{H}_{p_*}^+; \mathbb{Z})$. To elaborate, recall that such a generator can be taken to be any circle in U on which the coordinate s , the distance from p and the angle θ are constant. Such a circle is given by pushing a $1 - 3 \cos^2 \theta = 0$ circle where s and the distance from p are constant to the part of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ where $1 - 3 \cos^2 \theta$ is slightly positive. Meanwhile, a $1 - 3 \cos^2 \theta = 0$ circle of this sort is a constant radius slice of U 's intersection with any submanifold from \mathcal{M}_1 . As just noted, if ϑ contains U 's intersection with a submanifold from \mathcal{M}_1 , then there is a circle in each large n version of C_{p0n} that represents a nonzero multiple of the generator of the first homology of $\mathbb{R} \times \mathcal{H}_{p_*}^+$. But this conclusion is nonsense by virtue of the fact that C_{p0n} is diffeomorphic to $\mathbb{R} \times I_*$ and thus is contractible. This nonsense is what is required to prove the lemma.

What follows is the promised explanation for why ϑ contains U 's intersection with a submanifold from \mathcal{M}_1 . Lemma 4.4 is the key to the argument, for it implies that $|\zeta_n| > \kappa_* z_*$ where $z = 2\delta_n^2 < \frac{2}{n} z_*$. Keeping this in mind, write z and h in terms of the variables (r, θ, ϕ) where r is the distance to p . By way of reminder, the coordinate $z = r^2(1 - 3 \cos^2 \theta)$ and $h = r^2 \cos \theta \sin^2 \theta$. Thus,

$$(4-14) \quad h/z = \frac{\cos \theta \sin^2 \theta}{1 - 3 \cos^2 \theta} \quad \text{and} \quad z^2 + 6 \frac{h^2}{\sin^2 \theta} = r^4(1 + 3 \cos^4 \theta).$$

As a consequence, a $z < \frac{2}{n} z_*$ point where $|\zeta_n| > \kappa_* z_*$ is a point where

$$(4-15) \quad 1 - 3 \cos^2 \theta < \frac{1}{n} \frac{1}{3\sqrt{3}} \quad \text{and} \quad r > z_*^{1/2}.$$

This implies that ϑ contains a subvariety with a $1 - 3 \cos^2 \theta = 0$ point. Such a subvariety cannot have points where $1 - 3 \cos^2 \theta < 0$ as there are no such points in C_{p0n} . Thus, it must sit entirely in the $1 - 3 \cos^2 \theta = 0$ locus and so constitute the intersection between U and a submanifold from \mathcal{M}_1 . □

Step 4 Assume that the parameters (δ, χ_0, R) are such that Lemma 4.6 can be invoked. Let κ_{**} denote the version of κ from Lemma 4.6. Take z_* so that $100\kappa_{**}z_*$ is less than $10^{-6}\delta_*^2$.

Lemma 4.7 *There exists $\kappa \geq 1$ with the following significance: Let (φ, ζ) denote a solution to (3-4) that obeys a given version of (4-3). Then $1 - 3 \cos^2 \theta > \kappa^{-1}$ on the Ψ_p -image of the graph $(x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi(x, \hat{u}), h = \zeta(x, \hat{u}))$ in $\mathbb{R} \times \mathcal{X}$.*

The preceding lemma asserts the condition in the top bullet of (4-7).

Proof of Lemma 4.7 It follows from Lemma 4.6 and (4-14) that $1 - 3 \cos^2 \theta > c_0^{-1}$ on the Ψ_p -image of the $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundary of the graph where $c_0 \geq 1$ is a purely S -dependent (or \mathcal{K} -compatible) constant. Thus, if $1 - 3 \cos^2 \theta \leq c_0^{-1}$ on the Ψ_p -image of the graph, then this must occur in the interior. This requires that $\cos \theta$ take its maximum on the interior. Lemma II.4.8 asserts that this maximum can occur only where $\hat{u} = 0$.

With the preceding understood, suppose that the lemma is false so as to generate some nonsense. Granted this assumption, there exists a sequence $\{(\tau_n, (\varphi_n, \zeta_n))\}_{n=1,2,\dots}$ of the following sort: Each index n version of (φ_n, ζ_n) is a solution to (3-4) with asymptotic/boundary conditions given by the $\tau = \tau_n$ version of (4-3). Furthermore, there is some $u = 0$ point on the Ψ_p -image of the graph of (φ_n, ζ_n) where $1 - 3 \cos^2 \theta < \frac{1}{n}$. Given what is said in Lemma 4.5, an application of Proposition II.5.5 analogous to that used to prove Lemma 4.6 generates the same sort of conclusion: there is a loop in the Ψ_p -image of every large n graph that generates the first homology of $\mathbb{R} \times \mathcal{H}_{p*}^+$. As noted in the proof of Lemma 4.6, this is a nonsensical conclusion. \square

Step 5 The next lemma asserts the condition in the lower bullet of (4-7).

Lemma 4.8 *Given $\varepsilon > 0$, there exists $\kappa_\varepsilon > 1$ with the following significance: Suppose that $\tau \in [0, 1]$ and that (φ, ζ) is a solution to (3-4) with boundary values given by the parameter τ version of (4-3). There are \mathbb{R} -valued lifts of $\varphi(x, \hat{u})$, $\phi(\gamma_{p-})$ and $\phi(\gamma_\tau)$ such that*

- $|\varphi(x, \hat{u}) - \phi(\gamma_{p-}|\hat{u})| + |\zeta(x, \hat{u}) - h(\gamma_{p-})| < \varepsilon$ where $x < -\kappa_\varepsilon$,
- $|\varphi(x, \hat{u}) - \phi(\gamma_\tau|\hat{u})| + |\zeta(x, \hat{u}) - h(\gamma_\tau)| < \varepsilon$ where $x > \kappa_\varepsilon$.

Proof Suppose to the contrary that no such κ_ε exists so as to derive some nonsense. If this is the case, then there exists $\varepsilon_* > 0$ and a sequence $\{(\tau_n, (\varphi_n, \zeta_n))\}_{n=1,2,\dots}$ with $\tau_n \in [0, 1]$ and with (φ_n, ζ_n) a solution to (3-4) with asymptotic/boundary conditions given by the τ_n version of (4-3). Moreover, (φ_n, ζ_n) violates the $\varepsilon = \varepsilon_*$ conclusions of the lemma at all points where $x < -n$ or $x > n$. This said, no generality is lost by assuming that the sequence violates the conclusions where $x < -n$.

Construct a new solution to (3-4) by translating (φ_n, ζ_n) by the constant amount along the \mathbb{R} factor of $\mathbb{R} \times I_*$ so that the resulting pair violates the $\varepsilon = \varepsilon_*$ version of what

is asserted by the top bullet of the lemma at some point where $x = 0$. Let (φ'_n, ζ'_n) denote this new solution. The conditions given by the first bullet of (4-7) hold for the sequence $\{(\varphi'_n, \zeta'_n)\}_{n=1,2,\dots}$. This being the case, standard elliptic regularity arguments (see again Chapter 6 of [12]) prove that there is a subsequence that converges in the C^∞ -Fréchet topology on compact subsets of $\mathbb{R} \times I_*$. Let (φ, ζ) denote the limit. This pair obeys (3-4) and it obeys the $\tau = 0$ version of the condition given by the second bullet in (4-3). This is to say that the function φ on the boundary of $\mathbb{R} \times I_*$ is independent of the \mathbb{R} coordinate and its respective values on the two boundaries are those of $\phi(\gamma_{p-})$ on the relevant boundary of I_* .

The pair (φ, ζ) also satisfies the conditions given by the $\tau = 0$ version of the top bullet in (4-3). To see this, note that the integral of w over the Ψ_p -image of the graph given by this (φ, ζ) version of (3-1) is finite. As the image of the graph is J -holomorphic, it follows using Lemma II.5.6 that any given sufficiently large, constant $|s|$ slice of the Ψ_p -image of the graph must be everywhere very close to the $\hat{u} \in I_*$ segment of an integral curve of v in \mathcal{H}_{p*}^+ . Given the constant $|\hat{u}| = R + \frac{1}{2} \ln z_*$ value for φ , this segment must be from γ_{p-} .

Note next that the translation that defined $\{(\varphi'_n, \zeta'_n)\}_{n=1,2,\dots}$ guarantees that the solution (φ, ζ) is *not* the solution to (3-4) and with boundary values the $\tau = 0$ version of (4-3) that is given by the x -independent map $\hat{u} \rightarrow (\phi(\gamma_{p-}|\hat{u}), h(\gamma_{p-}))$.

The conclusion of the previous paragraph is nonsensical given the assertion that there is at most one solution to any given version of (3-4) with a given $\tau \in [0, 1]$ asymptotic/boundary conditions from (4-3). This assertion is proved in the next subsection; it is the latter's Lemma 4.9. □

This concludes the proof of Lemma 4.3. □

4D Uniqueness

The next lemma completes the proof of Lemma 4.8. It also proves the uniqueness assertion of Proposition 2.2.

Lemma 4.9 *Equation (3-4) has at most one solution whose boundary values are described by a given $\tau \in [0, 1]$ version of (4-3).*

Proof Suppose that $\tau \in [0, 1]$ and that $(\varphi^{(0)}, \zeta^{(0)})$ and $(\varphi^{(1)}, \zeta^{(1)})$ are two solutions to (3-4) with asymptotic/boundary conditions that are given by the parameter τ version

of (4-3). Introduce φ' to denote $\varphi^{(1)} - \varphi^{(0)}$. This function is zero on the boundary of $\mathbb{R} \times I_*$ and it has limit 0 as $|x| \rightarrow \infty$ on $R \times I_*$. Let $\zeta' = \zeta^{(1)} - \zeta^{(0)}$. The pair (φ', ζ') obeys an equation that can be written as $D_{\mathfrak{h}}(\varphi', \zeta') = 0$, where $D_{\mathfrak{h}}$ is described by a version of (3-5) and (3-6). Indeed, just such an equation arises by subtracting the $(\varphi^{(0)}, \zeta^{(0)})$ version of (3-4) from the corresponding $(\varphi^{(1)}, \zeta^{(1)})$, where it is understood that \mathbb{R} -valued lifts of $\varphi^{(0)}$ and $\varphi^{(1)}$ are chosen so their boundary values agree. The functions \mathfrak{a}_1 and \mathfrak{a}_2 that appear in this version of (3-5) are the functions $(x, \hat{u}) \rightarrow a_1(x, \hat{u}, \zeta^{(1)}|_{(x, \hat{u})})$ and $(x, \hat{u}) \rightarrow a_2(x, \hat{u}, \zeta^{(1)}|_{(x, \hat{u})})$ that appear in the $(\varphi^{(1)}, \zeta^{(1)})$ version of (3-4). Meanwhile, \mathfrak{b}_1 and \mathfrak{b}_2 are given by

$$(4-16) \bullet \quad \mathfrak{b}_1(x, \hat{u}) = \left[\int_0^1 a_{1h}(\cdot, \zeta^{(1)} + s(\zeta^{(0)} - \zeta^{(1)})) ds \right] \partial_x \varphi^{(0)} \\
+ \left[\int_0^1 a'_{1h}(\cdot, \zeta^{(1)} + s(\zeta^{(0)} - \zeta^{(1)})) ds \right] \partial_{\hat{u}} \varphi^{(0)}, \\
\bullet \quad \mathfrak{b}_2(x, \hat{u}) = b(x, \hat{u}) + \left[\int_0^1 a_{2h}(\cdot, \zeta^{(1)} + s(\zeta^{(0)} - \zeta^{(1)})) ds \right] \partial_x \varphi^{(0)} \\
+ \left[\int_0^1 a'_{2h}(\cdot, \zeta^{(1)} + s(\zeta^{(0)} - \zeta^{(1)})) ds \right] \partial_{\hat{u}} \varphi^{(0)}.$$

Given that $(\varphi^{(1)}, \zeta^{(1)})$ converges uniformly as $|x| \rightarrow \infty$, and given that this pair solves (3-4), standard elliptic regularity theorems as in Chapter 6 of [12] prove that the corresponding pair $(\partial_x \varphi^{(1)}, \partial_x \zeta^{(1)})$ converges uniformly to zero as $|x| \rightarrow \infty$. This implies that the version of $(\mathfrak{a}_1, \mathfrak{a}_2, \mathfrak{b}_1, \mathfrak{b}_2)$ just defined obeys the conditions in (3-6). Thus, Proposition 3.2 can be invoked to see that the just-defined version of $D_{\mathfrak{h}}$ has trivial kernel and so $(\varphi', \zeta') = 0$. □

5 Analytic background for the $\Delta_p > 0$ cases

This section prepares some analytic tools that are used in Section 6 to prove Proposition 2.2 when $\Delta_p > 0$. The analysis concerns two related issues that owe allegiance to item (b) in the third bullet of (2-9). This third bullet of (2-9) changes the domain of (φ, ζ) so as to be the complement of either one or two $\hat{u} = 0$ points in $\mathbb{R} \times I_*$. The first issue is of import with regards to the behavior of the pairs (φ, ζ) that arise in the $\Delta_p > 0$ versions of (2-10) near the missing $\hat{u} = 0$ points. Sections 5A and 5B are devoted to this topic. The second issue concerns the versions the operator D of of (3-5) that arise in the $\Delta_p > 0$ versions of (2-10). The domain and range spaces for the $\Delta_p > 0$ versions of D change to reflect the changed domain for the corresponding pair in (2-10) and the behavior of this pair near the missing $\hat{u} = 0$ point or points. The remaining subsections use what is said in Sections 5A and 5B to first define the new domain and range spaces for D , and then prove an analog of Proposition 3.2.

5A J -holomorphic ends and the $u = 0, 1 - 3 \cos^2 \theta = 0$ locus

This subsection describes the ends of J -holomorphic submanifolds whose constant s slices converge as $s \rightarrow \infty$ in an isotopic fashion to one or the other of the curves $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$. This is precisely the sort of end that appears $\Delta_p > 0$ cases of (2-9). What is said here concerns specifically the $\hat{\gamma}_p^+$ story as the story for the other curve can be obtained from the one told here by replacing θ with $\pi - \theta$ and changing some \pm signs at various points. The story here is told in four parts.

Part 1 Introduce by way of notation (s_+, ϕ_+) as coordinates for $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ so as to distinguish the latter from the eponymous factor in $\mathbb{R} \times \mathcal{H}_p$. A differential operator mapping $C^\infty(\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z}); \mathbb{R}^2)$ to itself is defined by the rule that sends a pair (a, b) to the pair with respective first and second components

$$(5-1) \bullet \partial_{s_+} a + \left(\sigma_0 \frac{3}{2\chi_0}\right) \partial_{\phi_+} b + \left(3\sigma_0 \frac{\chi_0 + 4e^{-2R}}{\chi_0^2}\right) a,$$

$$\bullet \partial_{s_+} b - \left(\frac{3}{4\chi_0\sigma_0}\right) \partial_{\phi_+} a - \left(\frac{4}{\chi_0\sigma_0} e^{-2R}\right) b.$$

Here, σ_0 is the value at $u = 0$ and $\theta = \theta_0$ of the function σ that is used for the fifth bullet of Part 1 in Section 1C. This operator is denoted by \mathfrak{D}_0 .

Pairs in the kernel of \mathfrak{D}_0 describe deformations of the J -holomorphic submanifold $\mathbb{R} \times \hat{\gamma}_p^+$ that are J -holomorphic to first order in the distance from $\mathbb{R} \times \hat{\gamma}_p^+$. In particular, pairs (a, b) with limit zero as $s_+ \rightarrow \infty$ describe the ends of J -holomorphic submanifolds whose constant s -slices converge as $s_+ \rightarrow \infty$ in an isotopic fashion to $\hat{\gamma}_p^+$. More is said about this in Parts 2-4. What follows directly talks about the kernel of \mathfrak{D}_0 .

The operator depicted in (5-1) has constant coefficients, and so the kernel has a basis whose elements are irreducible representations of the $\mathbb{R}/2\pi\mathbb{Z}$ action on the space of maps from $\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z})$ to \mathbb{R}^2 generated by ∂_{ϕ_+} . Using this Fourier mode decomposition makes an easy task of writing the kernel of (5-1). To say what this leads to, introduce

$$(5-2) \quad \lambda_1 = 3\sigma_0 \frac{\chi_0 + 4e^{-2R}}{\chi_0^2} \quad \text{and} \quad \lambda_2 = \frac{4}{\chi_0\sigma_0} e^{-2R}.$$

A basis for the kernel of \mathfrak{D}_0 is given by the ϕ_+ -independent elements

$$(5-3) \quad \eta_{0+} = (e^{-\lambda_1 s_+}, 0) \quad \text{and} \quad \eta_{0-} = (0, e^{\lambda_2 s_+}),$$

and then, for each $n \in \{1, 2, \dots\}$, elements that have the form

$$(5-4) \quad \begin{aligned} \eta_{n+} &= e^{-\lambda_{1n}s_+} (\cos n(\phi_+ - \phi_n), r_{1n} \sin n(\phi_+ - \phi_n)), \\ \eta_{n-} &= e^{\lambda_{2n}s_+} (\cos n(\phi_+ - \phi_n), -r_{2n} \sin n(\phi_+ - \phi_n)), \end{aligned}$$

where

$$\begin{aligned} \lambda_{1n} &= \frac{1}{2} \left(\left((\lambda_1 + \lambda_2)^2 + \frac{9n^2}{2\chi_0^2} \right)^{1/2} + \lambda_1 - \lambda_2 \right), \\ \lambda_{2n} &= \frac{1}{2} \left(\left((\lambda_1 + \lambda_2)^2 + \frac{9n^2}{2\chi_0^2} \right)^{1/2} + \lambda_2 - \lambda_1 \right), \end{aligned}$$

and r_{1n} and r_{2n} are certain specific positive constants. Meanwhile, $\phi_n \in \mathbb{R}/2\pi\mathbb{Z}$ can be any chosen angle.

Let \mathcal{S}_+ denote the linear span of $\{\eta_{n+}\}_{n=0,1,\dots}$ from (5-3) and (5-4). Let \mathcal{S}_- denote the linear span of $\{\eta_{n-}\}_{n=0,1,\dots}$ from (5-3) and (5-4). The elements from \mathcal{S}_+ limit to zero as $s_+ \rightarrow \infty$ and those from \mathcal{S}_- limit to zero as $s_+ \rightarrow -\infty$.

Part 2 The upcoming description of the ends of J -holomorphic submanifolds invokes some geometric constructions that are described next. To start, note that the restrictions of the coordinate functions s and ϕ parametrize $\mathbb{R} \times \widehat{\gamma}_p^+$. These functions on $\mathbb{R} \times \widehat{\gamma}_p^+$ are denoted by (s_+, ϕ_+) . These coordinates with an auxiliary set of Euclidean coordinates (θ_+, u_+) for a small radius disk in \mathbb{R}^2 can be used as coordinates for an $\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z})$ -invariant, tubular neighborhood in $\mathbb{R} \times \mathcal{H}_p$ of $\mathbb{R} \times \widehat{\gamma}_p^+$. This parametrization can be chosen so as to have the properties that are listed in the upcoming (5-5). The list uses $U_+ \subset \mathcal{H}_{p_*}^+$ to denote the constant s slices of this tubular neighborhood, this being an $\mathbb{R}/2\pi\mathbb{Z}$ -invariant tubular neighborhood of $\widehat{\gamma}_p^+$. The list also refers to respective $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ actions on the $(s_+, \phi_+, \theta_+, u_+)$ coordinate domain and on $\mathbb{R} \times U_+$. The action on the former are the constant translations of s_+ and ϕ_+ , and the action on the latter are the constant translations along the \mathbb{R} factor and the constant translations of the coordinate ϕ for the U_+ factor. The notation also uses θ_* to denote the angle with $\cos \theta_* = \frac{1}{\sqrt{3}}$, it being the value of θ on $\widehat{\gamma}_p^+$. Granted this notation, here are the parametrization's properties:

- (5-5) • The constant (s_+, ϕ_+) disks are J -holomorphic.
- The parametrization has $\theta = \theta_* + \theta_+$ and $u = u_+$
 - The parametrization is equivariant with respect to the $\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z})$ actions.
 - The coordinates (s_+, ϕ_+) equal (s, ϕ) on the $\theta_+ = 0, u_+ = 0$ cylinder.

A parametrization of this sort can be constructed using Lemma 5.4 from [19] with a little help from the inverse function theorem to arrange the condition in the second bullet.

Part 3 Granted these coordinates, a deformation of $\mathbb{R} \times \widehat{\gamma}_p^+$ can be parametrized as a graph via functions $(a, b): \mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z} \rightarrow \mathbb{R}^2$ as

$$(5-6) \quad (s_+, \phi_+) \rightarrow (s_+, \phi_+, \theta_+ = a(s_+, \phi_+), u_+ = b(s_+, \phi_+)).$$

If $\eta = (a, b)$ is defined on a given open set in $\mathbb{R} \times S^1$ and if $|\eta| \leq c_0^{-1}$, then the resulting graph over the given open set defines a J -holomorphic surface if and only if η obeys a nonlinear equation with the schematic form

$$(5-7) \quad \mathfrak{D}_0\eta + \tau_1 d\eta + \tau_0,$$

where τ_1 is a smooth map from a certain small radius disk about the origin in \mathbb{R}^2 to $\text{Hom}(T^*\mathbb{R}^2, \mathbb{R}^2)$, and where τ_0 is a smooth map from this same disk to \mathbb{R}^2 . These are such that $|\tau_1| \leq c_0 |\eta|$ and $|\tau_0| \leq c_0 |\eta|^2$.

By way of an example, the J -holomorphic cylinders that form Proposition II.3.4’s moduli space \mathcal{M}_{p_0} foliate the $u = 0$ slice of $\mathbb{R} \times \mathcal{H}_{p_0}^+$. These cylinders are ϕ -invariant. Each such cylinder has two ends; their constant s slices converge isotopically as $s \rightarrow \infty$ to the respective integral curves $\widehat{\gamma}_p^+$ and $\widehat{\gamma}_p^-$. The very large s parts of the end whose slices converge to $\widehat{\gamma}_p^+$ appears as a ϕ_+ -independent solutions to (5-7) that are defined for $s_+ \gg 1$ with pairs (a, b) such that $a > 0$ and $b = 0$. In particular, integrating the $u = 0$ version of the vector field in (II.3-10), or using arguments much like those in Section 2 of [7], finds that the relevant version of η can be written as

$$(5-8) \quad \eta = \alpha(e^{-\lambda_1 s_+} + e_1, 0), \quad \text{where } \alpha \in (0, \infty) \text{ and } |e_1| \leq c_0 |\alpha| e^{-(\lambda_1 + 1/c_0)s_+}.$$

Meanwhile, the large s part of the end of any given submanifold from Proposition II.3.4’s moduli space \mathcal{M}_{p_+} is described by (5-8) with $\alpha < 0$.

By way of a second example, the end in $\mathbb{R} \times \mathcal{H}_p^+$ from Proposition II.3.3’s moduli spaces \mathcal{M}_1 and \mathcal{M}_2 whose constant s slices converge as $s \rightarrow \infty$ to $\widehat{\gamma}_p^+$ are described where $s \ll -1$ by a ϕ_+ -invariant solution to (5-7) that is defined where $s_+ \ll -1$ and has the form

$$(5-9) \quad \eta = \alpha(0, e^{\lambda_2 s_+} + e_2), \quad \text{where } \alpha \in \mathbb{R} \setminus 0 \text{ and } |e_2| \leq c_0 |\alpha|^2 e^{-(\lambda_2 + 1/c_0)|s_+|}.$$

The $\alpha > 0$ cases describe the end of the submanifolds from \mathcal{M}_1 and the $\alpha < 0$ cases describe the end of the submanifolds from \mathcal{M}_2 .

A third example involves the submanifolds from Proposition II.3.2’s moduli space \mathcal{M}_Σ . Those parametrized as in the second bullet of Proposition II.3.2 by a pair (x, y) with y near 1 can be written using (5-5)–(5-7) as

$$(5-10) \quad \eta = (\alpha_+ e^{-\lambda_1 s_+} + e_+, \alpha_- e^{\lambda_2 s_+} + e_-),$$

where $\alpha_+ > 0$ and $\alpha_- > 0$. Here, e_+ and e_- are both ϕ_+ -invariant. In addition, their norms are such that $|e_+| \leq c_0 |\alpha_+| (|\alpha_+| + |\alpha_-|) e^{-(\lambda_1 + 1/c_0)s_+}$ and $|e_-| \leq c_0 |\alpha_-| (|\alpha_+| + |\alpha_-|) e^{-(\lambda_2 + 1/c_0)s_+}$. Note that this representation is valid only over a domain $I \times (\mathbb{R}/2\pi\mathbb{Z}) \subset \mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$, where I is a bounded interval whose endpoints are determined by α_+ and α_- . The left endpoint diverges as $\alpha_- \rightarrow 0$ and the right endpoint diverges as $\alpha_+ \rightarrow 0$. Meanwhile, a surface from \mathcal{M}_Σ parametrized by (x, y) with $y \sim 2$ appears as in (5-10) but with $\alpha_- < 0$.

Part 4 This part of the subsection directly addresses the issue of describing ends of J -holomorphic submanifolds using the kernel of \mathcal{D}_0 . As just noted, any such end whose large $s \gg 1$ slices sit in U_+ and converge to $\hat{\gamma}_p^+$ in an isotopic fashion as $s \rightarrow \infty$ is described by a solution to (5-7) that is defined where $s_+ \gg 1$ and has $s_+ \rightarrow \infty$ limit equal to zero. By the same token, any such end whose $s \ll -1$ slices sit in U_+ and converge to $\hat{\gamma}_p^+$ in an isotopic fashion is described by a solution to (5-7) defined where $s_+ \ll -1$ and converging to 0 as $s_+ \rightarrow -\infty$. The following proposition describes all such solutions to (5-7). By way of a reminder, what are denoted in the proposition by η_{n+} and η_{n-} are defined in (5-4).

Proposition 5.1 *There exists $\kappa \geq 1$ with the following significance: Fix $s_* \geq 1$.*

- *Suppose that η is a solution to (5-7) with domain $[s_*, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ that converges to 0 as $s_+ \rightarrow \infty$ and has pointwise norm bounded by κ^{-2} . There exists $n \geq 1$ such that η can be written as*

$$\eta = c_0 (e^{-\lambda_1 s_+} + e_1, 0) + c_n \eta_{n+} + \epsilon_n,$$

with $c_0, c_n \in (-\kappa^{-1}, \kappa^{-1})$, with e_1 as given by (5-8), and with

$$|\epsilon_n| \leq \kappa |c_n| (|c_n| + |c_0|) e^{-(\lambda_{1n} + 1/\kappa)s_+}.$$

Conversely, given $n \in \{1, 2, \dots\}$ and constants $c_0, c_n \in (-\kappa^{-1}, \kappa^{-1})$, there exists a solution to (5-7) that can be written in this way.

- Suppose that η is a solution to (5-7) with domain $(-\infty, -s_*) \times \mathbb{R}/2\pi\mathbb{Z}$ that converges to 0 as $s_+ \rightarrow -\infty$ and has pointwise norm bounded by κ^{-2} . There exists $n \geq 1$ such that η can be written as

$$\eta = c_0(0, e^{\lambda 2^{s_+}} + e_2) + c_n \eta_{n+} + \epsilon_n,$$

with $c_0, c_n \in (-\kappa^{-1}, \kappa^{-1})$, with e_2 given by (5-9), and with

$$|\epsilon_n| \leq \kappa |c_n| (|c_n| + |c_0|) e^{-(\lambda_{2n+1}/\kappa)|s_+|}.$$

Conversely, given $n \in \{1, 2, \dots\}$ and constants $c_0, c_n \in (-\kappa^{-1}, \kappa^{-1})$, there exists a solution to (5-7) that can be written in this way.

Moreover, in either case, the derivatives of η to any given order are square-integrable where $|s_+| > 2s_*$.

Proof The analysis from Section 2 and specifically Section 2.3 of [7] can be used but for one added comment to prove that any given solution to (5-7) with $s_+ \rightarrow \infty$ limit zero can be written as described. The extra comment concerns the derivation of the bounds on the norms for ϵ_0 and ϵ_1 . These bounds are obtained by projection η and the expression in (5-7) onto $\mathbb{R}/2\pi\mathbb{Z}$ -invariant subspace of maps from $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 . To elaborate, this projection is given by the map

$$(5-11) \quad \mathfrak{q} \rightarrow \Pi \mathfrak{q} = \frac{1}{2\pi} \int_{\mathbb{R}/2\pi\mathbb{Z}} \mathfrak{q}(\cdot, \phi_+) d\phi_+.$$

The use of such a projection is not discussed in Section 2 of [7]. Even so, the latter's arguments can be applied separately to the $\mathbb{R}/2\pi\mathbb{Z}$ -invariant part of (5-7) and the remainder with what are little more than notational changes to obtained the distinct bounds for the norms of ϵ_0 and on ϵ_1 .

The proof of the converse assertion in the first bullet is given below in two steps. The proof of the converse assertion in the second bullet is identical but for straightforward notational and cosmetic changes and so is not given. The proof that the derivatives to any given order are square-integrable invokes standard elliptic regularity theorems of the sort that can be found in Chapter 6 of [12].

Step 1 Use \mathbb{H} to denote now the L^2_1 completion of the space of smooth, \mathbb{R}^2 -valued functions on $[0, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ with compact support and which lie in Part 1's subspace \mathcal{S}_- on the boundary, $\{0\} \times \mathbb{R}/2\pi\mathbb{Z}$. Let \mathbb{L} denote the L^2 completion of the space of smooth, \mathbb{R}^2 -valued functions on $[0, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$. It is a straightforward

task using integration by parts to prove that \mathcal{D}_0 defines a Fredholm operator from \mathbb{H} to \mathbb{L} with trivial kernel and cokernel. As such, it has a bounded inverse. There is also a version here of the Hilbert spaces \mathbb{H}_* and \mathbb{L}_* that are defined in Section 3E. These are defined by completing the respective dense domains for \mathbb{H} and \mathbb{L} using for \mathbb{H}_* the $[0, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ analog of (3-33), and using for \mathbb{L}_* the analog that integrates the square of the norm of η' rather than that of its derivatives. The operator \mathcal{D}_0 also defines a bounded, linear map from \mathbb{H}_* to \mathbb{L}_* . The analog of Lemmas 3.4 and 3.5 holds in this case: The inverse of \mathcal{D}_0 maps $\mathbb{L}_* \subset \mathbb{L}$ to $\mathbb{H}_* \subset \mathbb{H}$ as a bounded operator. As \mathcal{D}_0 commutes with Π , the inverse also commutes with Π .

Step 2 Fix $T \geq 100$ and let β_T denote the function on $[0, \infty)$ given by $\beta(s - T)$. Fix $n \in \{1, 2, \dots\}$. Given the $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ version of Lemma 3.4, there exists $c_0 > 1$ and $T \geq 1$ such that following is true: Suppose that $|c_0| + |c_n| \leq c_0^{-2}$. Let $\mathbb{B} \subset \mathbb{H}_*$ denote the ball about the origin of radius c_0^{-1} . Reintroduce τ_0 and τ_1 from (5-7). Let $\eta_+ = (e^{-\lambda_1 s} + e_1, 0)$. A smooth map from $(\prod_2(-c_0^{-1}, c_0^{-1})) \times \mathbb{B}$ to \mathbb{L} is defined so as to send any given $((c_0, c_n), q)$ to

$$(5-12) \quad \mathcal{D}_0 q + \beta_T \tau_1 (c_0 \eta_+ + c_n \eta_{n+} + q) d(c_0 \eta_+ + c_n \eta_{n+} + q) + \beta_T \tau_0 (c_0 \eta_+ + c_n \eta_{n+} + q).$$

The differential of this map at $((c_0, 0), 0)$ along the \mathbb{H}_* factor is an isomorphism if $T \geq c_0^{-1}$. This being the case, the inverse function theorem finds c_0 and, for $T > c_0$, a smooth map $\epsilon_n: \prod_2(-c_0^{-1}, c_0^{-1}) \rightarrow \mathbb{B}$ such that the triple $((c_0, c_n), q = \epsilon_n(c_0, c_n))$ is mapped to 0 by (5-12). Moreover, this map is such that $|\epsilon_n| \leq c_0 |c_n| (|c_0| + |c_n|) e^{-\lambda_1 T}$. The techniques from Section 2.3 in [7] can be used to see that ϵ_n has the asserted norm bound. □

5B The kernel of \mathcal{D}_0 and graphs over $\mathbb{R} \times I_*$

Some of the J -holomorphic cylinders given by Proposition 5.1 via (5-5)–(5-7) will intersect $\mathbb{R} \times \mathcal{H}_{p*}^+$ and so intersect the image of Ψ_p . This subsection says something about the Ψ_p -inverse image of these intersections. Of particular interest are the cylinders where s is unbounded from above. There are five parts to what follows. Lemma 5.6 in Part 5 gives some indication as to why these cylinders are relevant.

Part 1 Fix $n \geq 1$ and a pair $(c_0, c_n) \in \mathbb{R}^2 \setminus 0$ whose absolute value is small enough to apply the first bullet of Proposition 5.1 to obtain a corresponding solution, η , to (5-7). Write η as (a, b) and use the latter in (5-5) and (5-6) to define a J -holomorphic cylinder in the $s \gg 1$ part of $\mathbb{R} \times U_+ \subset \mathbb{R} \times \mathcal{H}_p$. It follows from the second bullet of (5-5) that c_0 must be positive for this to occur.

Assume henceforth that $c_0 > 0$ and that $c_n \neq 0$. If this is so, then the large s part of the cylinder in question has algebraic intersection n with the large s parts of submanifolds from \mathcal{M}_{p_0} and from submanifolds from Proposition II.3.2's moduli space \mathcal{M}_Σ that come very near the $u = 0$ locus in $\mathbb{R} \times \mathcal{H}_p$ at large s . Indeed, this last point is a direct consequence of three facts: First, each such subvariety from \mathcal{M}_{p_0} appears as in (5-8), and those from \mathcal{M}_Σ appear as in (5-10). Second, the $c_0\eta_{0+} + \epsilon_0$ contribution to η has the form $(c_0e^{-\lambda_1 s_+} + \epsilon_0, 0)$ with ϵ_0 being ϕ -invariant and having the asserted norm bound. Third, the u_+ component of η_{1n} has $2n$ zeros on each large, constant s_+ circle in $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ and each such zero is transverse.

Part 2 The large s part of a cylinder in $\mathbb{R} \times \mathcal{H}_p^+$ parametrized via (5-5)–(5-7) and the first bullet of Proposition 5.1 with $c_0 > 0$ has $\eta = c_0(e^{-\lambda_1 s_+}, 0) + \mathcal{O}(e^{-(\lambda_1 + 1/c_0)s_+})$ and so looks to leading order like what is written in (5-8). By way of a reminder, the latter depicts a cylinder from Proposition II.3.2's moduli space \mathcal{M}_{p_0} . This being the case, what follows says more about the \mathcal{M}_{p_0} case of (5-8) in preparation for what is said in the next parts of the subsection about the $c_n \neq 0$ cases.

To set the stage, keep in mind that a cylinder from Proposition II.3.3's moduli space \mathcal{M}_{p_0} is ϕ -invariant and invariant with respect to the involution $\theta \rightarrow \pi - \theta$. Such a cylinder has two ends, and the constant s slices of these ends converge in an isotopic fashion as $s \rightarrow \infty$ to the respective integral curves $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$. The Ψ_p -inverse images of these cylinders from \mathcal{M}_{p_0} are the constant x slices of the $\hat{u} = 0$ locus in $\mathbb{R} \times \mathcal{X}$. The association of the value of x to the corresponding cylinder gives an \mathbb{R} -equivariant diffeomorphism between \mathbb{R} and \mathcal{M}_{p_0} . This diffeomorphism from \mathbb{R} to \mathcal{M}_{p_0} sends any given $y \in \mathbb{R}$ to the cylinder in \mathcal{M}_{p_0} whose $s = y$ slice is the $(u = 0, \theta = 0)$ circle in \mathcal{H}_p^+ .

Fix $y \in \mathbb{R}$ and let $\Sigma_y \in \mathcal{M}_{p_0}$ denote the corresponding cylinder. The function s on Σ_y has one critical value, this the $s = y$ locus. It restricts to both components of the complement of this locus as a proper map to (y, ∞) . The function $\cos \theta$ increases monotonically as a function of s with $s \rightarrow \infty$ limit $\frac{1}{\sqrt{3}}$ on one of these components. Meanwhile, $\cos \theta$ decreases monotonically on the other component with $s \rightarrow \infty$ limit equal to $-\frac{1}{\sqrt{3}}$. Let $\mathcal{E}_{0,y} \subset \Sigma_y$ denote the former component, this being the end whose constant s slices converge to $\hat{\gamma}_p^+$ and is given via (5-5)–(5-7) by using η as depicted in (5-8) for a suitable choice of α . Denote the $\mathcal{E}_{0,y}$ version of α by α_y .

Lemma 5.2 *Fix a pair $x, y \in \mathbb{R}$ and let Σ_x and Σ_y denote the corresponding surfaces from \mathcal{M}_{p_0} . Then $x - y = \frac{1}{\lambda_1} \ln(\alpha_x/\alpha_y)$.*

Proof This follows from (5-5)–(5-8) given the fact that Σ_x is obtained from Σ_y by translating the latter by $x - y$ along the \mathbb{R} factor of $\mathbb{R} \times \mathcal{H}_p^+$. \square

Part 3 Fix $n \in \{1, 2, \dots\}$ and $y \in \mathbb{R}$. Introduce by way of notation $\mathcal{E}_{n,y}$ to denote the large s part of a cylinder that is described via (5-5)–(5-7) and the first bullet of Proposition 5.1 with $c_0 = \alpha_y$ and $c_n \neq 0$. The next lemma describes the Ψ_p -inverse image of these sorts of cylinders.

Lemma 5.3 *There exists a constant $\kappa \geq 1$ with the following significance: Fix a positive integer n and $\phi_n \in \mathbb{R}/2\pi\mathbb{Z}$ so as to specify a particular version of η_{n+} from (5-4). Choose a real number y and set $c_0 = \alpha_y$; then choose $c_n \neq 0$ so as to define $\mathcal{E}_{n,y} \subset \mathbb{R} \times \mathcal{H}_p^+$ via (5-5)–(5-7) using η as in Proposition 5.1. Fix $(s_+, \phi_+) \in \mathbb{R} \times S^1$ with $s_+ \gg 1$.*

- The $\hat{\phi}$ and h coordinates of the Ψ_p^{-1} -image of the corresponding point in $\mathcal{E}_{n,y}$ are $\hat{\phi} = (\phi_+ - \phi_n)$ and

$$h = (\chi_0 + 4e^{-2R}) \times \frac{2}{3\sqrt{3}} (1 - 3\alpha_y^2 e^{-2\lambda_1 s_+} - 6\alpha_y^2 c_n e^{-(\lambda_1 + \lambda_{1n})s_+} \cos(n(\phi_+ - \phi_n)) + \dots),$$

where the unwritten term has two parts. The ϕ_+ -invariant part is bounded in absolute value by $e^{-(2\lambda_1 + 1/\kappa)s_+}$. The remainder is bounded in absolute value by $e^{-(\lambda_1 + \lambda_{1n} + 1/\kappa)s_+}$.

- The \hat{u} coordinate is

$$\hat{u}(s_+, \phi_+) = \frac{3}{2\chi_0\sigma_0} \left(\left((\lambda_1 + \lambda_2)^2 + \frac{9n^2}{2\chi_0^2} \right)^{1/2} + \lambda_1 + \lambda_2 \right) (1 + \dots) \times c_n e^{-\lambda'_{1n} s_+} \sin(n(\phi_+ - \phi_n)) + \dots,$$

where $\lambda'_{1n} = \lambda_{1n} + 12e^{-2R}(\sigma_0\chi_0)^{-1}$ and where the first unwritten factor is bounded in absolute value by κe^{-2R} and the second by $e^{-(\lambda'_{1n} + 1/\kappa)s_+}$. The absolute values of their derivatives are also bounded by these same respective factors.

- The x coordinate is

$$x(s_+, \phi_+) = y + \lambda_1^{-1} c_n e^{-(\lambda_{1n} - \lambda_1)s_+} \cos(n(\phi_+ - \phi_n)) + \dots,$$

where the unwritten term is bounded in absolute value by $e^{-(\lambda_{1n} - \lambda_1 + 1/\kappa)s_+}$. The absolute value of its derivatives is also bounded by this same factor.

A proof is given in a moment. What follows directly is a corollary of what is said by the second and third bullets of Lemma 5.3.

Corollary 5.4 Fix $s \gg 1$ and there is an open, contractible neighborhood $V_s \subset \mathbb{R} \times I_*$ of the point $(y, 0)$ with the following significance: The projection to $\mathbb{R} \times I_*$ of the Ψ_p -inverse image of where $s > s$ in $\mathcal{E}_{n,y}$ defines a proper, n -to-1 covering map onto $V_s \setminus (y, 0)$.

The rest of this part of the subsection is occupied with the:

Proof of Lemma 5.3 The claim in the first bullet follows from what is said in Proposition 5.1 and the fact that $\hat{\phi}$ and h are the respective pullbacks of ϕ and $f(u) \cos \theta \sin^2 \theta$. Use the identification $\theta = \theta_* + a(s_+, \phi_+)$ and $u = b(s_+, \phi_+)$ in the latter function with Taylor’s theorem to obtain the given expression for h .

The formulas for \hat{u} and x are derived in the four steps that follow. The arguments given take $\phi_n = 0$. The assertion in the general case follows directly from this case by applying a constant $\mathbb{R}/2\pi\mathbb{Z}$ translation.

Step 1 The $u = 0$ slice $\mathcal{E}_{n,y}$ is parametrized by s_+ via the rule

$$(5-13) \quad \theta(s_+) = \theta_* + \alpha_y(e^{-\lambda_1 s_+} + e_0 \pm c_n e^{-\lambda_{1n} s_+} + \epsilon_n,$$

with e_0 and ϵ_n as given by Proposition 5.1. The plus sign occurs at an angle $\phi_+ = 0 + \epsilon_+$ and the minus sign occurs at $\phi = \pi + \epsilon_-$, where $|\epsilon_+|$ and $|\epsilon_-|$ are both bounded by $c_0 e^{-(\lambda_{1n} + 1/c_0)s_+}$. From the vantage of $\mathbb{R} \times \mathcal{X}$, the $\phi \sim 0$ intersection locus correspond to points with

$$(5-14) \quad \hat{u} = 0, \quad \hat{\phi} = 0 + \epsilon_+ \quad \text{and} \quad h = (\chi_0 + 4e^{-2R}) \frac{2}{3\sqrt{3}} (1 - 3\alpha_y^2 e^{-2\lambda_1 s_+}) + \dots,$$

where the unwritten term in the expression for h is bounded in absolute value by $e^{-(2\lambda_1 + 1/c_0)s_+}$. The $u = 0$ locus in $\mathcal{E}_{n,y}$ with $\phi \sim \pi$ has $\hat{\phi}$ coordinate $\pi + \epsilon_-$ and h coordinate also given by (5-14).

Step 2 It follows from the definition given in (1-29) that the x and \hat{u} coordinates of the point in $\mathcal{E}_{n,y}$ can be determined from (5-5)–(5-7) by integrating the vector field given in (II.3-10) starting at the point $(s = s_+, \theta = \theta_* + a(s_+, \phi_+), b(s_+, \phi_+))$. The values of x and \hat{u} of this point on $\mathcal{E}_{n,y}$ are the respective s coordinates and u coordinates of the point on the relevant integral curve where $\theta = \frac{\pi}{2}$.

To see what results, let $\tau \rightarrow (s(\tau), \theta(\tau), u(\tau))$ denote for the moment a certain parametrization of this integral curve. Take $\tau = 0$ to be the start point. As θ increases along the curve, (II.3-9) implies that $|u|$ decreases from its initially small value as τ

increases. Now, define the parametrization of the curve by τ so that Taylor’s theorem applied to (II.3-9) writes the τ -derivative of u as

$$(5-15) \quad \frac{du}{d\tau} = -2\sqrt{3} e^{-2R} u(1 + \dots) \cos \theta \sin \theta,$$

where the unwritten term has absolute value bounded by $c_0 e^{-2\lambda_{1n}s}$. A second application of Taylor’s theorem writes

$$(5-16) \quad \frac{d\theta}{d\tau} = \frac{1}{2\sqrt{2}} (\chi_0 + 4e^{-2R} + \dots)(1 - 3 \cos^2 \theta),$$

where the unwritten terms are also bounded by $c_0 e^{-2\lambda_1 s}$. Given the very small $\tau = 0$ value for $\theta - \theta_*$ at the start, it follows from (5-16) that the value of τ where $\theta = \frac{\pi}{2}$ is given by

$$(5-17) \quad \tau = \frac{3}{\chi_0} \sigma_0 s_+ + \dots,$$

where the unwritten term has absolute value bounded by c_0 . Granted this, use Taylor’s theorem to approximate $\cos \theta \sin \theta$ in (5-15) by $\cos \theta_* \sin \theta_{0\setminus*} = \frac{2}{\sqrt{3}}$. Integration produces the formula given for \hat{u} in the second bullet of the lemma.

Step 3 To get an expression for the x coordinate of a given $u = 0$ point on $\mathcal{E}_{n,y}$, note that the value of θ on the end \mathcal{E}_y in \mathcal{M}_{p_0} ’s cylinder Σ_y is described at large s_+ by the $c_n = 0$ version of (5-13). For any given $x \in \mathbb{R}$, use Lemma 5.2 to see that the value of θ on the end \mathcal{E}_x in the corresponding Σ_x is described at large s_+ by the version of (5-13) that sets $c_n = 0$ and replaces s_+ by $s_+ - (x - y)$. Granted this last observation, use a first-order Taylor’s approximation to see that the value of x on $\mathcal{E}_{n,y}$ at a given very large s_+ and where $u = 0$ is obtained by solving

$$(5-18) \quad \alpha_y (e^{-\lambda_1 s_+} + e_1)(1 + \lambda_1(x - y) \dots) = \alpha_y ((e^{-\lambda_1 s_+} + e_1) \pm c_n e^{-\lambda_{1n} s_+}) + \dots,$$

where the unwritten term on the left-hand side involve higher powers of $(x - y)$ and a term with absolute value bounded by $e^{-s_+/c_0} |x - y|$. Meanwhile, the unwritten term on the right-hand side has absolute value bounded by $e^{-(\lambda_{1n} + 1/c_0)s_+}$. This last equation implies that the x coordinate of a given $s_+ \gg 1$ point on the $u = 0$ locus in $\mathcal{E}_{n,y}$ is given by

$$(5-19) \quad x(s_+) - y = \pm \lambda_1^{-1} c_n e^{-(\lambda_{1n} - \lambda_1)s_+} + \dots,$$

where the unwritten term has absolute value bounded by $e^{-(\lambda_{1n} - \lambda_1 + 1/c_0)s_+}$.

Step 4 Granted that $|u|$ decreases from its initially small value, it also follows from (II.3-9) that the value of x is very nearly the \mathbb{R} -parameter of the $\theta = \frac{\pi}{2}$ point on the

unique $\Sigma_{(\cdot)}$ surface that contains $(s_+, \theta = \theta_* + a(s_+, \phi_+))$. Given this observation, what is said the preceding steps implies directly the formula for x in the third bullet of the lemma. \square

Part 4 This part of the subsection concerns specifically the case where $n = 1$. The discussion here concerns the normal bundle to the large s part of the surface $\mathcal{E}_{1,y}$ when viewed using (5-5)–(5-7) and Proposition 5.1, and when viewed via Ψ_p as a submanifold in $\mathbb{R} \times \mathcal{X}$.

To start, use the almost complex structure J and the 2–form $\hat{w} = ds \wedge \hat{a} + w$ to define the Riemannian metric $\hat{w}(\cdot, J(\cdot))$. Let $N \rightarrow \mathcal{E}_{1,y}$ denote the normal bundle to the submanifold $\mathcal{E}_{1,y}$, this being the orthogonal complement in $T(\mathbb{R} \times \mathcal{H}_p)$ of $T\mathcal{E}_{1,y}$. View the large s part of $\mathcal{E}_{1,y}$ using (5-5)–(5-7) and Proposition 5.1 to see that pairing with the 1–forms $(d\theta_+, du_+)$ define an isomorphism between N and the product \mathbb{R}^2 –bundle. Meanwhile, this same part of $\mathcal{E}_{1,y}$ can be viewed as the Ψ_p –image of a surface in $\mathbb{R} \times \mathcal{X}$, and Corollary 5.4 implies that the 1–forms $(d\hat{\phi}, dh)$ also define an isomorphism between N and the product \mathbb{R}^2 –bundle.

These two product structures are related in the following way: Let η denote a map from the $s_+ \gg 1$ part of $\mathbb{R} \times S^1$ to \mathbb{R}^2 . Use the product structure defined by $(d\theta_+, du_+)$ to view η as a section of N over this part of $\mathcal{E}_{1,y}$. Meanwhile, use the $n = 1$ version of Corollary 5.4 to view this part of $\mathcal{E}_{1,y}$ as the Ψ_p –image of a graph of the sort depicted in (3-1) with the domain of the relevant version of the pair (φ, ζ) being the complement of $(y, 0)$ in an $\mathbb{R} \times I_*$ neighborhood of $(y, 0)$. With $\mathcal{E}_{1,y}$ viewed this way, the image via $(d\hat{\phi}, dh)$ of the section defined by η defines a map, η , from the domain of (φ, ζ) to \mathbb{R}^2 . The maps η and η are related via a rule given by

$$(5-20) \quad \eta|_{(s_+, \phi_+)} = U \cdot (\eta|_{(x(s_+, \phi_+), \hat{u}(s_+, \phi_+))}),$$

where U is a smooth map from the large s_+ part of $\mathbb{R} \times S^1$ to $GL(2; \mathbb{R})$ with positive determinant. Note that the latter component of $GL(2; \mathbb{R})$ deformation retracts on to the $SO(2)$ subgroup, and so the restriction of U to any given constant, large s_+ circle in $\mathbb{R} \times S^1$ has an integer degree that is independent of the chosen value for s_+ .

Lemma 5.5 *The map U just defined has degree 1.*

Proof The constant map from the domain of (φ, ζ) to \mathbb{R}^2 given by the element $(1, 0)$ corresponds via $(d\hat{\phi}, dh)$ to a section of N over $\mathcal{E}_{1,y}$, this being the orthogonal projection to N of the vector field $\partial_{\hat{\phi}}$. The latter generates the deformations of $\mathcal{E}_{1,y}$

that are given by the constant rotations of the $\hat{\phi}$ coordinate. Granted this, use Lemma 5.3 with (5-5) and (5-7) to see that these deformations are generated along the large s_+ part of the graph in (5-5) by the section of N that is defined by the orthogonal projection of the vector field

$$(5-21) \quad re^{-\lambda_{1n}s_+} \left(-\sin(\phi_+ - \phi_n) \partial_{\theta_+} + \frac{3}{2\chi_0\sigma_0} \left(\left((\lambda_1 + \lambda_2)^2 + \frac{9}{2\chi_0^2} \right)^{1/2} + \lambda_1 + \lambda_2 \right) \times \cos(\phi_+ - \phi_n) \partial_{u_+} \right) + \dots,$$

where unwritten terms are bounded in absolute value by $c_0 re^{-(\lambda_{1n} + 1/c_0)s_+}$. This last vector rotates once, counterclockwise in \mathbb{R}^2 as ϕ_+ changes from 0 to 2π . This implies that U has degree 1, as claimed. □

Part 5 The following lemma gives some hint as to the relevance of the cylinders that are described by the first bullet of Proposition 5.1.

Lemma 5.6 *Let $C \subset \mathbb{R} \times \mathcal{H}_{p*}^+$ denote a properly embedded, J -holomorphic submanifold, and let $\mathcal{E} \subset C$ denote an end where s is unbounded from above and whose constant s slices converge in an isotopic fashion to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$. Then the $s \gg 1$ part of \mathcal{E} can be parametrized via (5-5)–(5-7) by a map of the sort that is described by the first bullet in Proposition 5.1.*

Proof There exists $s_* > 1$ such that the $s \geq s_*$ part of \mathcal{E} is a proper submanifold with boundary in $[s_*, \infty) \times U_+$. Use (5-5)–(5-6) to view this submanifold using the coordinates $(s_+, \phi_+, \theta_+, u_+)$. The function s restricts to the $s \geq s_*$ part of \mathcal{E} as a proper function with no critical points. Granted that this is so, it follows that the projection to the $(\theta_+, u_+) = (0, 0)$ cylinder restricts to the large s_+ part of \mathcal{E} as a covering map. This covering map must have degree 1 because the constant s slices of \mathcal{E} are isotopic to $\hat{\gamma}_p^+$. This understood, the large s_+ part of \mathcal{E} has intersection number 1 with any given sufficiently large s_+ fiber of the projection to the $(\theta_+, u_+) = (0, 0)$ cylinder. This implies that the large s_+ part of \mathcal{E} can be written as the graph of a map from the large s_+ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 that is described by the first bullet of Proposition 5.1. □

5C Fredholm operators

This subsection introduces some new Fredholm domain and range spaces for certain operators of the sort that are described by (3-5) and (3-6). The upcoming Proposition 5.7 supplies the analog of Proposition 3.2 for the new Fredholm incarnations of these operators.

To set the stage, let $Q \in \mathbb{R} \times I_*$ denote either the complement of a single $\hat{u} = 0$ point or two $\hat{u} = 0$ points. Suppose that $\mathfrak{h} = (\varphi, \varsigma)$ maps the complement of Q in $\mathbb{R} \times I_*$ to \mathbb{R}^2 so as to define a graph in $\mathbb{R} \times \mathcal{X}$ via (3-1). Let $C_{\mathfrak{h}}$ denote the Ψ_p -image of this graph. Assume in what follows that the large $|s|$ part of $C_{\mathfrak{h}}$ is J -holomorphic, that it obeys the first bullet in (2-9), and that it obeys the $\Delta_p = 1$ or $\Delta_p = 2$ bullets of (2-9). The pair \mathfrak{h} has an associated version of the operator that is depicted in (3-9), this denoted by $D_{\mathfrak{h}}$. In what follows, D is used to denote an operator that is given by (3-5) and (3-6) with the extra condition

$$(5-22) \quad D = D_{\mathfrak{h}} \text{ on the complement of a compact set in } (\mathbb{R} \times I_*) \setminus Q.$$

Operators of this sort play a central role in the upcoming proof of the $\Delta_p > 0$ version of Proposition 2.2. Part 1 of what follows defines the new domain and range spaces. This first part of the subsection ends with Proposition 5.7. The subsequent parts of the subsection supply the proof of Proposition 5.7.

Part 1 Let $N \rightarrow C_{\mathfrak{h}}$ denote the normal bundle, this being the orthogonal complement to $TC_{\mathfrak{h}}$ in $T(\mathbb{R} \times \mathcal{H}_p)$ with orthogonality defined by the metric $\hat{\omega}(\cdot, J(\cdot))$. Here again, $\hat{\omega} = ds \wedge \hat{a} + w$. Identify $C_{\mathfrak{h}}$ with its inverse image via Ψ_p in $\mathbb{R} \times \mathcal{X}$. Having done so, use $(d\hat{\phi}, dh)$ to write a section of N as a map from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 . Granted this identification, a map from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 can be viewed as a section of N over $C_{\mathfrak{h}}$. A map with compact support defines a section of N with compact support, and vice versa. The aforementioned Riemannian metric defines a fiber metric for N and an associated metric compatible, covariant derivative for sections of N . It also defines a Riemannian metric on $TC_{\mathfrak{h}}$ and thus an area form. Use the fiber metric on N and $TC_{\mathfrak{h}}$, the covariant derivative on N , and integration with respect to this area form to define the L^2_1 inner product on the space of sections of N with compact support. This L^2_1 inner product gives an inner product on the space of compactly supported maps from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 . Use \mathbb{H} to denote the completion using this L^2_1 inner product of the subspace of compactly supported maps from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 whose first component is zero along the boundary of $\mathbb{R} \times I_*$. This Hilbert space \mathbb{H} will be the domain space for the desired Fredholm incarnation of D .

The range Hilbert space for the new incarnation of D is a certain L^2 inner product space. To set the stage for the definition of this inner product, introduce e^0 to denote the pullback of the \mathbb{C} -valued 1-form in (3-2) via Ψ_p^{-1} . The pullback of the latter to the graph of \mathfrak{h} defines a section, $e^{\mathfrak{h}}$, of $T^*_{\mathbb{C}}C_{\mathfrak{h}}$ and this section defines a polarization of $T^*_{\mathbb{C}}C_{\mathfrak{h}}$ since $e^{\mathfrak{h}} \wedge \bar{e}^{\mathfrak{h}} \neq 0$. Define $T^{0,1}C_{\mathfrak{h}}$ to be the span of $\bar{e}^{\mathfrak{h}}$. By way

of comparison, let $\mathcal{E} \subset C_{\mathfrak{h}}$ denote a J -holomorphic end whose constant $s \gg 1$ slices are circles, this an end whose large s slices converge in $\mathcal{H}_{\mathfrak{p}}$ to either $\widehat{\gamma}_{\mathfrak{p}}^-$ or $\widehat{\gamma}_{\mathfrak{p}}^+$. As J defines a complex structure on \mathcal{E} , so it defines a polarization of $T_{\mathbb{C}}^*\mathcal{E}$ as $T^{1,0}\mathcal{E} \oplus T^{0,1}\mathcal{E}$. This polarization is the same as that given by $\{e^{\mathfrak{h}}, \bar{e}^{\mathfrak{h}}\}$.

Use the metric defined by J and $\widehat{\omega}$ to define a hermitian metric on $T_{\mathbb{C}}^*(\mathbb{R} \times \mathcal{H}_{\mathfrak{p}}^+)$, and use the latter to define the norm of $e^{\mathfrak{h}}$. This norm is denoted by $|e^{\mathfrak{h}}|$. Let $|\cdot|_N$ denote the fiber norm described above on N . Reintroduce the functions a_1 and a_2 from (3-4). Use these functions and $|e^{\mathfrak{h}}|$ to define a norm on the space of maps from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 as follows: Let $\zeta = (\iota, o)$ denote a given map. View the map $(a_1^{-1}\iota, a_2^{-1}o)$ as a section of N over $C_{\mathfrak{h}}$. Set the norm of ζ to be $|(a_1^{-1}\iota, a_2^{-1}o)|_N |e^{\mathfrak{h}}|$. Use this pointwise norm and integration with respect to the area form on $C_{\mathfrak{h}}$ to define an L^2 inner product on the space of compactly supported maps from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 . The resulting Hilbert space is denoted by \mathbb{L} . This space \mathbb{L} is the new range Hilbert space.

By way of an explanation, the trivialization of the normal bundle of $C_{\mathfrak{h}}$ given by the 1-forms $(d\widehat{\phi}, dh)$ identifies the latter with the span of the vector fields $\{\partial_{\widehat{\phi}}, \partial_h\}$. The almost complex structure J preserves this span, and so endows N with the structure of a complex line bundle. The \mathbb{C} -valued 1-form $(a_1^{-1}\iota + ia_2^{-1}o)\bar{e}^{\mathfrak{h}}$ defines a section of $N \otimes T^{0,1}C_{\mathfrak{h}}$. The norm of this section as defined using the induced hermitian metric is the norm defined above for ζ .

Keep in mind for what follows that the norms that define \mathbb{H} and \mathbb{L} depend on the chosen pair \mathfrak{h} . Even so, the spaces \mathbb{H} and \mathbb{L} do not depend on \mathfrak{h} . This is so because the respective norms defined by pairs \mathfrak{h} and \mathfrak{h}' are commensurate.

Proposition 5.7 *Suppose that D is described by (3-5), (3-6) and (5-22). Then D extends as a Fredholm operator from \mathbb{H} to \mathbb{L} with index $\Delta_{\mathfrak{p}}$ and trivial cokernel.*

Proof The proof is contained in the subsequent parts of this subsection. Part 2 explains why D is Fredholm, Part 3 computes the index and Part 4 proves that the cokernel is trivial. □

Part 2 Use $\|\cdot\|_{\mathbb{L}}$ to denote the L^2 -norm that defined \mathbb{L} . Meanwhile, use $\|\cdot\|$ to define the L^2 -norm on sections of N and on sections of $N \otimes T^*C_{\mathfrak{h}}$. The covariant derivative on sections of N is denoted by ∇ . The operator D has closed range and finite-dimensional kernel if and only if there exists $c \geq 1$ such that if $\eta \in \mathbb{H}$, then:

$$(5-23) \quad \bullet \quad \|D\eta\|_{\mathbb{L}}^2 \geq c^{-1}\|\nabla\eta\|^2 - c\|\eta\|^2.$$

$$\bullet \quad \text{If } \eta \text{ has support only where } |\Psi_{\mathfrak{p}}^*s| > c, \text{ then } \|D\eta\|_{\mathbb{L}}^2 \geq c^{-1}\|\eta\|^2.$$

As in the case with Proposition 3.2, the finite-dimensionality of the cokernel follows if the formal, L^2 adjoint of D also obeys (5-23). Here, the L^2 -norm is that used to define \mathbb{L} . The proof that this is so differs only in notation for the proof that (5-23) holds for D and so will not be given.

To see about (5-23), it is sufficient to restrict attention to two sorts of compactly supported sections of N . With $s_1 > 1$ fixed, the first sort are those with no support where $s > 4s_1$ on an end $\mathcal{E} \subset C_{\mathfrak{h}}$ whose constant s slices are circles. The second sort are the sections with support only in the $s > 2s_1$ portion of such an end. The arguments in Section 3C establish the existence of an s_1 -dependent constant c that makes (5-23) true for all sections of the first sort. The proof that (5-23) for the sections with support where $s > 2s_1$ on an end \mathcal{E} as just described has three steps. These steps consider the case where the constant s slices of \mathcal{E} converge as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$. The argument for the other case is identical but for some sign changes.

Step 1 Let $y \in Q$ denote the point that corresponds to the end \mathcal{E} . Use what is said in Proposition 5.1, Lemma 5.3 and Corollary 5.4 to view the large s part of the end of \mathcal{E} via (5-5)–(5-7) with η in Proposition 5.1 defined using $n = 1$ and $c_0 = \alpha_y$, and with an appropriate choice for $c_1 \in \mathbb{R} \setminus 0$ and $\phi_1 \in \mathbb{R}/2\pi\mathbb{Z}$. Take s_1 so that the $s \geq s_1$ part of \mathcal{E} appears in this way. Nothing is lost by assuming that $D = D_{\mathfrak{h}}$ on the Ψ_p -inverse image of this part of \mathcal{E} .

Use $(d\theta_+, du_+)$ to identify N over the $s \geq s_1$ part of \mathcal{E} with the product bundle and so write a section of N with support on the $s > 2s_1$ part of \mathcal{E} as a map from the large s_+ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 . Let \mathfrak{x} denote such a map, but viewed as a section over \mathcal{E} of N . Multiply this section by ds_+ and use the parametrization of \mathcal{E} by (s_+, ϕ_+) and the complex line bundle structure on N defined by J to write the latter as a section of $N \otimes_{\mathbb{C}} (T_{\mathbb{C}}^*\mathcal{E})$. Use $(\mathfrak{x})_{0,1}$ to denote the $N \otimes T^{0,1}\mathcal{E}$ part of this section of $N \otimes T^{0,1}\mathcal{E}$.

Let η denote a map from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 with support only on the part of the domain that parametrizes the $s > 2s_1$ part of \mathcal{E} . Write the two components of $D_{\mathfrak{h}}\eta$ as (ι, o) and then view $(a_1^{-1}\iota, a_2^{-1}o)$ as a section of N over \mathcal{E} . With N viewed as a complex line bundle, multiply the latter by $\bar{e}^{\mathfrak{h}}$ to define a section of $N \otimes T_{\mathbb{C}}\mathcal{E}$. Denote this last section by $(D_{\mathfrak{h}}\eta)_{0,1}$.

Step 2 The lemma below is used in Step 3 to write $D_{\mathfrak{h}}$ near the point $(y, 0)$ in terms of the operator \mathfrak{D}_0 from (5-1).

Lemma 5.8 *There exists a first-order differential operator, \mathfrak{d} , on the space of maps from $[s_1, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 , a map $V: [s_1, \infty) \times \mathbb{R}/2\pi\mathbb{Z} \rightarrow \mathbb{C} \setminus 0$ and $\kappa > 1$ with the following properties:*

- *The coefficients of \mathfrak{d} bounded in absolute value by $\kappa e^{-s/\kappa}$*
- *The norms of both V and V^{-1} are bounded by κ .*
- *Let $\eta \in C^\infty((\mathbb{R} \times I_*) \setminus Q; \mathbb{R}^2)$ with support only on the part of the domain that parametrizes the $s > 2s_1$ part of \mathcal{E} . Let U denote \mathcal{E} 's version of the map to $GL(2; \mathbb{R})$ that appears in Lemma 5.5. Then $((\mathfrak{D}_0 + \mathfrak{d})(U\eta))_{0,1} = V(D_{\mathfrak{h}}\eta)_{0,1}$.*

Proof Introduce $U \subset (\mathbb{R} \times I_*) \setminus Q$ to denote the domain that parametrizes the $s > s_1$ part of \mathcal{E} . Let $\eta = (\varphi', \zeta')$ denote a map from U to \mathbb{R}^2 that is annihilated by $D_{\mathfrak{h}}$. Let $U' \subset U$ denote an open set with compact closure. For t near zero in \mathbb{R} , the Ψ_p -image of the graph of the map $(\varphi + t\varphi', \zeta + t\zeta')$ defines a deformation of $\Psi_p(U') \subset \mathcal{E}$ that is J -holomorphic to first order in t . Let $(a_{\mathfrak{h}}, b_{\mathfrak{h}})$ denote the map to \mathbb{R}^2 from the $s \gg 1$ portion of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ whose graph parametrizes \mathcal{E} via (5-6). Write $U\eta$ as (a', b') . The pair given by $(a_{\mathfrak{h}} + ta', b_{\mathfrak{h}} + tb')$ defines via (5-6) a deformation of $\Psi_p(U')$ that is J -holomorphic to first order in t if and only if (a', b') obeys an equation of the form $(\mathfrak{D}_0 + \mathfrak{d}_{\mathcal{E}})(a', b') = 0$ where $\mathfrak{d}_{\mathcal{E}}$ is a certain first-order differential operator whose coefficients are bounded by $c_0 e^{-s/c_0}$. It follows from this that there exists a map v from $[s_1, \infty) \times \mathbb{R}/2\pi\mathbb{Z} \rightarrow \mathbb{C} \setminus 0$ such that the assertion given by the third bullet of the lemma holds using $\mathfrak{d} = \mathfrak{d}_{\mathcal{E}}$ and for any smooth map η with support on U . The uniform bounds on v and v^{-1} can be derived using the chain rule from the formulas in Lemma 5.3. □

Step 3 Granted what is said in Lemma 5.8, it is sufficient to prove that there exists $c_0 \geq 1$ such that

$$(5-24) \quad \|\mathfrak{D}_0\eta\|_{L^2} \geq c_0^{-1} (\|d\eta\|_{L^2} + \|\eta\|_{L^2})$$

for all maps η with compact support on the $s_+ > 1$ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$. That this is so follows from the fact that the symmetric operator

$$(5-25) \quad (a, b) \rightarrow \left(\partial_\phi b + 2 \frac{\chi_0 + 4e^{-2R}}{\chi_0} a, -\partial_\phi a - \frac{1}{3} e^{-2R} b \right)$$

on $C^\infty(S^1; \mathbb{R}^2)$ has trivial kernel.

Part 3 This part of the subsection computes the Fredholm index of D . The computation has three steps.

Step 1 This step first defines from C_h a closed manifold with empty boundary that is diffeomorphic to the complement of a point in $\mathbb{R} \times S^1$. This manifold is denoted by Z in what follows. An oriented \mathbb{R}^2 -bundle is then defined over Z and D is shown to extend over Z as an operator that acts on sections of this bundle.

The manifold Z is defined by a suitable identification of the boundary components. To set the stage, note that any given version of D that obeys (3-5), (3-6) and (5-22) can be continuously deformed through a 1-parameter family of operators obeying (3-5), (3-6) and (5-22) to D_h . This deformation won't change the index. Granted that such is the case, assume that $D = D_h$.

Use (3-1) and the map Ψ_p to identify C_h with $(\mathbb{R} \times I_*) \setminus Q$. Fix $\varepsilon \in (0, \frac{1}{8})$ and introduce I' to denote the interval $[-R - \frac{1}{2} \ln((1 + \varepsilon)z_*), R + \frac{1}{2} \ln((1 + \varepsilon)z_*)]$. Granted the aforementioned identification, extend C_h as $(\mathbb{R} \times I') \setminus Q$. Having done so, introduce the function $t_+ = e^{-2(R-\hat{u})}$ where $\hat{u} > R + \ln \delta$ on I' and use the pair (x, t_+) to parametrize the part of $(\mathbb{R} \times I') \setminus Q$ where $\hat{u} \in (R + \frac{1}{2} \ln((1-\varepsilon)z_*), R + \frac{1}{2} \ln((1+\varepsilon)z_*)]$. Likewise introduce the function $t_- = -e^{-2(R+\hat{u})}$ and use (x, t_-) to parametrize the $\hat{u} < -R - \frac{1}{2} \ln((1-\varepsilon)z_*)$ portion of the domain $(\mathbb{R} \times I') \setminus Q$. Use these coordinates to identify the $t_+ \in [(1-\varepsilon)z_*, (1+\varepsilon)z_*]$ portion of $(\mathbb{R} \times I') \setminus Q$ with the portion of $(\mathbb{R} \times I_*) \setminus Q$ where $t_- \in [-(1+\varepsilon)z_*, -(1-\varepsilon)z_*]$ using the rule $t_- = -2z_* + t_+$. The slice of Z where $t_+ = z_*$ and so $t_- = -z_*$ in Z is said in what follows to be the z_* -locus. The complement of this z_* -locus in Z is the interior of $(\mathbb{R} \times I') \setminus Q$.

Define an oriented, \mathbb{R}^2 -bundle over Z as follows: the bundle is obtained from the product \mathbb{R}^2 -bundle over $(\mathbb{R} \times I') \setminus Q$ by identifying the point $((x, t_+), (\zeta_1, \zeta_2))$ with the point $((x, t_- = -2z_* + t_+), (\zeta_2, -\zeta_1))$. Use N_Z to denote this \mathbb{R}^2 -bundle.

As explained next, the operator D_h extends over the whole of Z as a differential operator on the space of sections of E . This is because D_h when written in terms of the coordinates (x, t_+) on the $\hat{u} \in (R + \ln \delta, R + \frac{1}{2} \ln z_*)$ part of $(\mathbb{R} \times I') \setminus Q$ is the standard Cauchy-Riemann operator, and this is also the case for D_h when written in terms of the coordinates (x, t_-) on the $\hat{u} = [-R - \frac{1}{2} \ln z_*, -R - \ln \delta)$ part of $(\mathbb{R} \times I') \setminus Q$. This extension of D_h to Z is denoted in what follows by \mathcal{D}_Z .

Step 2 This step defines a 1-parameter family of “matching conditions” for sections of N_Z with discontinuity on the z_* -locus in Z . The family is parametrized by the interval $[0, 1]$. A given parameter value is denoted by τ .

Fix $\tau \in [0, 1]$ and suppose that (φ', ζ') is a map from $(\mathbb{R} \times I_*) \setminus (y, 0)$ to \mathbb{R}^2 . This map is said to satisfy the τ -matching condition when the following is true: Let (φ'_-, ζ'_-) and

(φ_+, ζ_+) denote the respective \mathbb{R}^2 -valued functions that are defined by (φ', ζ') where the coordinate t_- is in $[-z_*, -\frac{1}{2}z_*]$ and where the coordinate t_+ is in $[\frac{1}{2}z_*, z_*]$. Then

$$(5-26) \quad \varphi'_+|_{t_+=z_*} = -\tau\zeta'_- \quad \text{and} \quad \varphi'_-|_{t_-=-z_*} = \tau\zeta'_+|_{t_+=z_*}.$$

For each $\tau \in [0, 1]$, define the Hilbert space \mathbb{H}_τ by copying the definition of the Hilbert space \mathbb{H} in Part 3 of this subsection but with the $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundary conditions used in Part 1 replaced by those in (5-26). Define the Hilbert space \mathbb{L} as in Part 1. The $\tau = 0$ version of \mathbb{H}_τ is the Hilbert space \mathbb{H} . The $\tau = 1$ version is a Hilbert space of sections of N_Z . Meanwhile, \mathbb{L} can be viewed as the closure of the space of sections of N_Z with respect to the L^2 -norm that is defined as in Part 1. Thus, \mathbb{H}_1 and \mathbb{L} can be viewed as respective L^2_1 and L^2 Hilbert spaces of sections of N_Z .

Lemma 5.9 *For each $\tau \in [0, 1]$, the operator $D_{\mathfrak{h}}$ defines a Fredholm map from \mathbb{H}_τ to \mathbb{L} .*

Proof The conditions in (5-23) must be established for $D_{\mathfrak{h}}$ on the dense domain of smooth, compactly supported maps from $(\mathbb{R} \times I') \setminus Q$ that obey (5-26). By the same token, these same conditions must be established for the formal L^2 adjoint. The argument that proves the analog of (5-23) for the formal L^2 adjoint of a given $\tau \in [0, 1]$ version of $D_{\mathfrak{h}}$ is identical to that just given but for cosmetics. Note in this regard that this adjoint has dense domain given by the compactly supported maps from $(\mathbb{R} \times I') \setminus Q$ that obeys (5-26).

The new issues with regards to (5-23) for $D_{\mathfrak{h}}$ do not concern $D_{\mathfrak{h}}$ near the points in Q ; they concern only the part of the argument that comes from Section 3C. This understood, consider the top line in (5-23). The top line is established in Section 3C using an integration by parts with the observation that the boundary terms are separately zero. The same integration by parts for $\tau \neq 0$ now yields respective $t_+ = z_*$ and $t_- = -z_*$ boundary terms that are not identically zero, but are opposite in sign. As a consequence, these terms add to zero and so make no contribution.

Consider next the lower line in (5-23). The key issue is whether (3-16) holds with $\tau \neq 0$. If this is so, then the argument used in Step 4 of Section 3C can be used here with only notational modifications to establish the desired result. To see about (3-16), use (3-17) to see that an element (φ', ζ') in the kernel of Q^- must be such that ζ' is constant and

$$(5-27) \quad \varphi'|_{t_+=z_*} = \varphi'|_{t_-=-z_*} + \zeta'c,$$

where $c \neq 0$ is the integral of b_{2-} over I_* . As can be seen using (3-13), this constant c is less than $-\chi_0 \delta^{-2} R$ and so significantly less than -2 . What with (5-26), this requires that

$$(5-28) \quad -\tau \zeta' = (\tau + c) \zeta',$$

and so $c = -2\tau$. Thus, (5-28) cannot hold. □

Step 3 The family of Hilbert spaces $\{\mathbb{H}_\tau\}_{\tau \in [0,1]}$ defines a smooth, Hilbert space bundle $\mathcal{H} \rightarrow [0, 1]$. Indeed, a nonisometric isomorphism from \mathbb{H}_τ to \mathbb{H}_0 can be defined as follows: Fix a compactly supported function $\mu: [\frac{1}{2}z_*, z_*] \rightarrow [0, 1]$ that is equal to 1 near z_* . Let (φ', ζ') denote a given element in \mathbb{H}_τ . Let (l', o') denote the image of this element in \mathbb{H}_0 . Then $(l', o') = (\varphi', \zeta')$ except where $t_+ > \frac{1}{2}z_*$ and where $t_- > -\frac{1}{2}z_*$. The pair (l', o') where $t_+ > \frac{1}{2}z_*$ is

$$(5-29) \quad (l', o')|_{(x, t_+)} = (\varphi'_+, \zeta'_+)|_{(x, t_+)} + (\tau \mu(t_+) (\zeta'_-|_{(x, t_- = -2z_* + t_+)}), 0).$$

A similar formula defines (l', o') where $t_- < -\frac{1}{2}z_*$.

The family of Fredholm operators $\{D_h: \mathbb{H}_\tau \rightarrow \mathbb{L}\}_{\tau \in [0,1]}$ defines a smooth section of the Fredholm homomorphisms from \mathcal{H} to the product Hilbert space bundle $[0, 1] \times \mathbb{L}$. This being the case, all members of this family have the same Fredholm index. In particular, the Fredholm index of D_h needed for Proposition 5.7 is that of \mathcal{D}_Z .

Given the latter observation, the arguments used in Sections 3d and 4b,d of [20] can be applied with only cosmetic changes to see that the Fredholm index of \mathcal{D}_Z on \mathbb{H}_1 is equal to Δ_p .

Part 4 This part of the subsection explains why D has trivial cokernel. This will follow with a proof that the kernel of D has dimension Δ_p . The proof that such is the case has five steps.

Step 1 With N viewed as the product bundle over the complement of Q in \mathbb{R}^2 , the operator D has the schematic form that is depicted in (3-5). Suppose that (φ', ζ') is in the kernel of D . Let Γ denote the locus in $(\mathbb{R} \times I_*) \setminus Q$ where $\zeta' = 0$. The argument used in Step 2 of Section 3D can be repeated here to see that Γ is described by (3-24) if it is not empty and not all of $(\mathbb{R} \times I_*) \setminus Q$.

The argument used in Step 5 of Section 3D can be repeated to see that Γ is not empty. This understood, assume that Γ is not all of $(\mathbb{R} \times I_*) \setminus Q$. The argument from this same step in Section 3D proves somewhat more about Γ . It proves, in particular, that Γ has

a nonzero, even number of edges with the following property: either x is unbounded on the edge, or the edge has an endpoint on the boundary of $\mathbb{R} \times I_*$, or the closure of the edge in $\mathbb{R} \times I_*$ is a point of Q .

Step 2 Fix $T \gg 1$ and define a closed, rectangular path in the interior $\mathbb{R} \times I_*$ with sides parallel to the axis such that the constant \hat{u} edges obey $|\hat{u}| = R + \frac{1}{2} \ln z_* - \frac{1}{T}$ and the constant x edges obey $|x| = T$. Orient this path so that a circumnavigation in the positive direction travels in the positive \hat{u} direction on the $x = T$ edge. Use R_T to denote this oriented rectangular path. If T is sufficiently large, then the restriction of (φ', ζ') to R_T defines a nowhere-zero map from R_T to \mathbb{R}^2 . Indeed, this can be seen for the constant \hat{u} edges by using the fact that (φ', ζ') obeys the Cauchy–Riemann equation where $\hat{u} > R + \ln \delta$ when written as function of $(x, t_+ = e^{-2(R-\hat{u})})$, and that it also obeys these equations where $\hat{u} < -R - \ln \delta$ when written as functions of $(x, t_- = -e^{-2(R+\hat{u})})$. Meanwhile, arguments much like those used to prove Proposition 2.4 in [7] prove that there are no zeros of (φ', ζ') where $|x| \gg 1$. Granted what was just said, the pair (φ', ζ') defines a map from R_T to $\mathbb{R}^2 \setminus \{0\}$ for all T sufficiently large. Each such large T map has a degree; they are all the same. As explained next, this degree is negative. To see this, note that the degree is equal to the intersection number between the image of R_T and any given outward directed ray in \mathbb{R}^2 , for example the positive x -axis. The path R_T intersects the positive x -axis where $\zeta' = 0$ and $\varphi' > 0$. These are all edges of Γ , and it follows from Step 1 that this set is nonempty when T is large. Meanwhile, (3-5) implies directly that each intersection point between the image of R_T and the positive x -axis has negative local intersection number.

Step 3 Suppose that q is a zero of (φ', ζ') in the interior of $(\mathbb{R} \times I_*) \setminus Q$. It follows from (3-5) that there are no zeros of (φ', ζ') save q in some small radius disk centered at q , and that (φ', ζ') has positive degree as a map from the boundary of this disk to $\mathbb{R}^2 \setminus \{0\}$.

Step 4 Let $U \subset \mathbb{R}^2 \times I_*$ denote a very small radius disk centered at $(y, 0) \in Q$ with the radius such that $D = D_{\mathfrak{h}}$ on U and such that the graph of (φ', ζ') over $U \setminus (y, 0)$ maps via Ψ_p to the very large s part of the corresponding end. Denote the latter by \mathcal{E} . It follows from the upcoming (5-29) that there are no zeros in (φ', ζ') in $U \setminus (y, 0)$ if U has small radius. Granted that U has such a small radius, the pair (φ', ζ') defines a map from the boundary of any concentric disk in U to $\mathbb{R}^2 \setminus \{0\}$. This map has a degree, this denoted by n_y . This degree is negative if Q has a single point. Indeed, this follows

from what is said in Steps 2 and 3. By the same token, if Q has two points, then the sum of the degrees of the maps at the two points is negative.

The pair (φ', ζ') on $U \setminus \{0\}$ defines a section of the bundle N over the large s part of \mathcal{E} . This section can be written with respect to the product structure for N on \mathcal{E} given by the basis $\{d\theta, du\}$. The resulting map to $\mathbb{R}^2 \setminus \{0\}$ from the large s part of \mathcal{E} is denoted in what follows by (a, b) . It follows from Lemma 5.5 that the pair (a, b) define a map with degree $n_y + 1$.

Step 5 As noted in Lemma 5.8, the operator $D_{\mathfrak{h}}$ on \mathcal{E}_0 can be decomposed as the sum $D_{\mathfrak{h}} = \mathfrak{D}_0 + \mathfrak{d}$ where \mathfrak{D}_0 is as given in (5-1) and where \mathfrak{d} is a first-order operator whose symbol and zeroth-order terms have norm bounded by $c_0 e^{-s/c_0}$. The arguments for Proposition 2.4 in [7] prove that any given element in the kernel of $D_{\mathfrak{h}}$ at large s on \mathcal{E} appear as follows for some $n \in \{0, 1, 2, \dots\}$:

$$(5-30) \quad e^{-\lambda_{1n}s}(\cos n(\phi - \phi_n), r_{1n} \sin n(\phi - \phi_n)) + \epsilon_n,$$

where each $n \geq 1$ version of λ_{1n} , r_{1n} and ϕ_n are as defined in (5-4) and where $\lambda_{10} = \lambda_1$. Meanwhile, ϵ_n is such that $|\epsilon_n| \leq e^{-(\lambda_{1n}+1/c_0)s}$. Note that each version of (5-30) defines a map from any constant and sufficiently large s circle in $\mathbb{R} \times S^1$ to $\mathbb{R}^2 \setminus \{0\}$. The $n = 0$ version has degree zero and all $n \geq 1$ versions have positive degree, this being n .

Suppose now that Q has a single point. Given that the $n_y + 1 \leq 0$ and $n_y < 0$, the pair (a, b) defined in Step 4 has nonpositive degree; it follows that it has degree zero. This must be true for any such pair arising from the kernel of $D_{\mathfrak{h}}$. If (a', b') is a second such pair, then a linear combination of the latter with (a, b) can be found such that the result defines an $n > 0$ version of (5-30). This is impossible if the linear combination is not identically zero. The preceding conclusion implies that the kernel of $D_{\mathfrak{h}}$ has dimension 1.

Suppose next that Q has two points. Denote these points as $(y, 0)$ and $(y', 0)$. Given that the degree in (5-30) is nonnegative, and given that $n_y + n_{y'} < 0$, it follows that only the cases $(n_y = 0, n_{y'} = -1)$, $(n_{y'} = 0, n_y = -1)$ and $(n_y = -1, n_{y'} = -1)$ can occur. The argument from the preceding paragraph can be repeated to see that the kernel of $D_{\mathfrak{h}}$ cannot contain two linearly independent elements which are such that both have $n_y = 0$ or both have $n_{y'} = 0$. This constraint is satisfied only if the kernel has dimension 2.

5D The Banach spaces \mathbb{H}_* and \mathbb{L}_*

There is an analog in the context of Proposition 5.7 of Section 3E’s Banach spaces \mathbb{H}_* and \mathbb{L}_* . To set the stage for the definitions, first reintroduce the notation used in (5-22). Use $\text{dist}(\cdot, \cdot)$ to denote the distance function on $C_{\mathfrak{h}}$ that is induced by the metric on $\mathbb{R} \times \mathcal{H}_{p_*}^+$. Let μ denote a smooth, nonincreasing function on $[0, \infty)$ with value 1 on $[0, \frac{1}{2}]$ and value 0 on $[1, \infty)$. Given $\rho > 0$ and $(x, \hat{u}) \in (\mathbb{R} \times I_*) \setminus (y, 0)$, use $\mu_{\rho, (x, \hat{u})}$ to denote the function $\mu(\rho^{-1} \text{dist}(\cdot, (x, \hat{u})))$. As in (3-33), fix $\nu \in (0, \frac{1}{100})$. The norm that defines \mathbb{H}_* is the sum of two terms. The first is the norm for \mathbb{H} , and the second is the square root of the function that assigns to a given smooth map in \mathbb{H} the value

$$(5-31) \quad \sup_{(x, \hat{u}) \in (\mathbb{R} \times I_*) \setminus (y, 0)} \sup_{\rho \in (0, 1)} \rho^{-\nu} \|\mu_{\rho, (x, \hat{u})} \nabla \eta'\|^2.$$

This norm is denoted by $\|\cdot\|_{\mathbb{H}_*}$. The Banach space \mathbb{H}_* is the completion of the set of smooth, compactly supported elements in \mathbb{H} using this norm.

The Banach space \mathbb{L}_* is the completion of the space of smooth, compactly supported sections of \mathbb{L} using the norm that is the sum of the norm $\|\cdot\|_{\mathbb{L}}$ with the norm whose square is the function that is given by replacing $\nabla \eta'$ in (5-31) by η' and by replacing the norm $\|\cdot\|$ by $\|\cdot\|_{\mathbb{L}}$.

The following lemma states the analogs of Lemmas 3.4 and 3.5 for these new versions of \mathbb{H}_* and \mathbb{L}_* . Given Proposition 5.7, the assertions also follow directly from Theorem 3.5.2 in [12].

Lemma 5.10 *Define \mathbb{H}_* and \mathbb{L}_* as above.*

- *Elements in \mathbb{H}_* are Hölder continuous with exponent $\frac{1}{2}\nu$ and the inclusion map from \mathbb{H}_* into the corresponding Hölder Banach space is continuous. In particular, there exists a constant $\kappa > 1$ that depends only on ν and has the following significance: if $f \in \mathbb{H}_*$, then $|f| \leq \kappa \|f\|_{\mathbb{H}_*}$.*
- *If (x, \hat{u}) is any given point in $(\mathbb{R} \times I_*) \setminus Q$, then $\lim_{\text{dist}(\cdot, (x, \hat{u})) \rightarrow \infty} |f|$ exists and it is zero; thus, elements in \mathbb{H}_* have pointwise uniform limit zero as $s \rightarrow \infty$ on $C_{\mathfrak{h}}$.*
- *Any operator D given by (3-5), (3-6) and (5-22) maps \mathbb{H}_* to \mathbb{L}_* , and its inverse restricts to \mathbb{L}_* so as to define a bounded linear operator from \mathbb{L}_* to \mathbb{H}_* .*

As was the case for \mathbb{H} and \mathbb{L} , the norms that define the Banach spaces \mathbb{H}_* and \mathbb{L}_* depend on the chosen pair \mathfrak{h} but the spaces do not.

6 Proof of Proposition 2.2: the $\Delta_p > 0$ case

The proof in the case when some $p \in \Lambda$ versions of Δ_p are 1 or 2 follows much the same path as that given in the preceding section for when Δ_p is zero. In particular, an open/closed argument is again used for a certain $[0, 1]$ -parametrized family of nonlinear, elliptic, first-order equations for a map from the complement of either one or two $\hat{u} = 0$ points in $\mathbb{R} \times I_*$ to $(\mathbb{R}/2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$. As in Section 4, the $\tau = 0$ member of this family is explicit, and the $\tau = 1$ member is the desired pair $(\varphi^{p_0}, \zeta^{p_0})$. The substantive differences are consequences of two related facts: The first is that the domain of $(\varphi^{p_0}, \zeta^{p_0})$ is now the complement in $\mathbb{R} \times I_*$ of the aforementioned $\hat{u} = 0$ point or points. The second stems from item (b) of the third bullet in (2-9); the latter prescribes the behavior of (φ, ζ) near the missing set in $\mathbb{R} \times I_*$. This prescribed behavior makes for a more complicated $\tau = 0$ member of the family. The new domain and the prescribed asymptotics requires versions of Proposition 5.7's operator D .

The arguments that follow discuss only the case when $\Delta_p = 1$ and $m_p = -1$ because the $\Delta_p = m_p = 1$ arguments and those when $\Delta_p = 2$ are identical but for cosmetic changes.

6A An approximation to $(\varphi^{p_0}, \zeta^{p_0})$

This subsection constructs an \mathbb{R} -parametrized family of maps from the complement of a $\hat{u} = 0$ point in $\mathbb{R} \times I_*$ to the space $(\mathbb{R}/2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ such that each member defines a graph in $\mathbb{R} \times \mathcal{X}$ whose Ψ_p -image has C_{p_0} 's large $|s|$ asymptotics and C_{p_0} 's behavior near the $|\hat{u}| = R + \frac{1}{2} \ln z_*$ boundaries of $\mathbb{R} \times I_*$. The family is parametrized by the \mathbb{R} coordinate of the missing $\hat{u} = 0$ point. The \mathbb{R} coordinate of this point is denoted by y . There are three parts to the construction.

Looking ahead, Section 6B explains how any one of these approximations can be used as the starting solution for a $[0, 1]$ -parametrized family of equations whose parameter 1 solution is the desired $(\varphi^{p_0}, \zeta^{p_0})$. The construction of this $[0, 1]$ -parametrized family requires the analytic tools that are supplied by Section 5.

Part 1 Construct the 1-parameter family of arcs $\{\gamma_\tau\}_{\tau \in [0,1]}$ as done in Part 1 of Section 4A. These are described in (4-1). The $m_p = 1$ condition implies that $\gamma_{\tau=1} \neq \gamma_{p_+}$. Even so, complete the constructions of Section 4 with the family $\{\gamma_\tau\}_{\tau \in [0,1]}$ to obtain a map from $\mathbb{R} \times I_*$ to $(\mathbb{R}/2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ whose graph in $\mathbb{R} \times \mathcal{X}$ has J -holomorphic image via Ψ_p . Use (φ_-, ζ_-) to denote \mathbb{R} -valued functions that define this map.

Let γ'_{p-} denote the integral curve of v in the $|\hat{u}| \leq R + \frac{1}{2} \ln z_*$ part of $\mathcal{H}_{p_*}^+$ with the following properties: It starts where $\hat{u} = -R - \frac{1}{2} \ln z_*$ at the same ϕ angle as γ_{p-} and it ends where $\hat{u} = R + \frac{1}{2} \ln z_*$ at the same ϕ angle as γ_{p-} . Let $\Delta\phi$ and $\Delta\phi'$ denote the respective angle changes along γ_{p-} and γ'_{p-} . These are given by the integral that appears in the fourth bullet of Lemma II.2.2. Require that $\Delta\phi' = \Delta\phi - 2\pi$. Redo the construction in Part 1 of Section 4A starting with the arc γ'_{p-} at $\tau = 0$. Use $\{\gamma'_\tau\}_{\tau \in [0,1]}$ to denote the resulting family. This family is such that $\gamma'_{\tau=1} = \gamma_{p+}$. Corresponding arcs γ_τ and γ'_τ have the same $\phi \in \mathbb{R}/2\pi\mathbb{Z}$ values where $\hat{u} = -R - \frac{1}{2} \ln z_*$ and also where $\hat{u} = R + \frac{1}{2} \ln z_*$. Even so, the corresponding versions of $\Delta\phi$ and $\Delta\phi'$ differ by -2π .

Redo the constructions of Section 4 with this second family $\{\gamma'_\tau\}_{\tau \in [0,1]}$ to obtain a second graph in $\mathbb{R} \times \mathcal{X}$ whose image via Ψ_p is J -holomorphic. Use (φ_+, ς_+) to denote the map from $\mathbb{R} \times I_*$ to $\mathbb{R} \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ that define this second graph. If necessary, add 2π times an integer to φ_+ so that φ_- and φ_+ agree where $\hat{u} = -R - \frac{1}{2} \ln z_*$. As a consequence, $\varphi_+ - \varphi_- = -2\pi$ at $\hat{u} = R + \frac{1}{2} \ln z_*$.

Part 2 Choose a smooth, nondecreasing map $w: [-1, 1] \rightarrow [0, 1]$ that is equal to 0 on the interval $[-1, -\frac{1}{8}]$, equal to 1 on $[\frac{1}{8}, 1]$, and is such that $w(-s) = 1 - w(s)$. Fix $\varepsilon > 0$ and introduce functions $w_{\varepsilon,y}^-$ and $w_{\varepsilon,y}^+$ mapping \mathbb{R} to $[0, 1]$ by the rules

$$(6-1) \quad w_{\varepsilon,y}^-(x) \rightarrow w\left(\frac{1}{\varepsilon}(x - y)\right) \quad \text{and} \quad w_{\varepsilon,y}^+(x) = w\left(-\frac{1}{\varepsilon}(x - y)\right).$$

These \hat{u} -independent functions are used to define a graph over $\mathbb{R} \times I_*$ that is smooth in the complement of the part of the $\hat{u} = 0$ locus where $|x - y| < \varepsilon$. The graph is defined by the pair of respective $\mathbb{R}/2\pi\mathbb{Z}$ - and \mathbb{R} -valued functions $(\varphi_{\varepsilon,y,0}, \varsigma_{\varepsilon,y,0})$ that are given by the two rules that follow:

$$(6-2) \quad \begin{aligned} &\bullet (\varphi_{\varepsilon,y,0} = \varphi_- + w_{\varepsilon,y}^-(\varphi_+ - \varphi_-), \varsigma_{\varepsilon,y,0} = \varsigma_- + w_{\varepsilon,y}^-(\varsigma_+ - \varsigma_-)) \text{ where } \hat{u} \leq 0. \\ &\bullet (\varphi_{\varepsilon,y,0} = \varphi_+ + w_{\varepsilon,y}^+(\varphi_- - \varphi_+ - 2\pi), \varsigma_{\varepsilon,y,0} = \varsigma_+ + w_{\varepsilon,y}^+(\varsigma_- - \varsigma_+)) \text{ where } \hat{u} \geq 0. \end{aligned}$$

What follows are two key properties of the pair $(\varphi_{\varepsilon,y,0}, \varsigma_{\varepsilon,y,0})$:

$$(6-3) \quad \begin{aligned} &\bullet \text{ The function } \varphi_{\varepsilon,y,0} = \text{ where } |\hat{u}| = R + \frac{1}{2} \ln z_* \text{ is the function } \phi^{S_0}(\cdot, z_*) \text{ in (2-12)}. \\ &\bullet \text{ The } |x - y| > \varepsilon \text{ parts of the graph of } (\varphi_{\varepsilon,y,0}, \varsigma_{\varepsilon,y,0}) \text{ defines via } \Psi_p \text{ a } J\text{-holomorphic submanifold with boundary in } \mathbb{R} \times \mathcal{H}_{p_*}^+ \text{ whose constant } |s| \text{ slices converge as } s \rightarrow -\infty \text{ to the arc } \gamma_{p-} \text{ and as } s \rightarrow \infty \text{ to the arc } \gamma_{p+}. \end{aligned}$$

Introduce C' to denote the image via Ψ_p of the graph of the pair $(\varphi_{\varepsilon,y,0}, \varsigma_{\varepsilon,y,0})$.

Part 3 Set $c_0 = \alpha_y$, fix $c_1 \in \mathbb{R} \setminus 0$ but small and fix $\phi_1 \in \mathbb{R}/2\pi\mathbb{Z}$ and use the resulting $n = 1$ version of η from Proposition 5.1 to define via (5-5)–(5-7) a J -holomorphic cylinder in the $s \gg 1$ part of $\mathbb{R} \times \mathcal{H}_{p_*}^+$. Use $\mathcal{E} \subset \mathbb{R} \times \mathcal{H}_{p_*}^+$ to denote this cylinder. The picture supplied by Lemma 5.3 has the following implication: There exists $s_1 > s_y$ such that the Ψ_p -inverse image of the complement of a compact set in the $s \geq s_1$ part of \mathcal{E} can be written as a graph over the complement of $(y, 0)$ in a small radius disk in $\mathbb{R} \times I_*$ about $(y, 0)$. This is to say that this part of $\Psi_p^{-1}(\mathcal{E})$ can be written as the graph

$$(6-4) \quad (x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi_y(x, \hat{u}), h = \zeta_y(x, \hat{u})),$$

where (φ_y, ζ_y) is a smooth map from the complement of $(y, 0)$ in such a small radius disk to $\mathbb{R} \times \left(-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2\right)$. Use ρ_y to denote the radius of this disk.

What follow are two additional observations that follow directly from Lemma 5.3: First, the function $(x, \hat{u}) \rightarrow \zeta_y(x, \hat{u})$ limits uniformly to $\frac{2}{3\sqrt{3}}(\chi_0 + 4e^{-2R})$ as $|x - y|^2 + |\hat{u}|^2$ limits to zero. The second concerns the map φ_y on circles where $|x - y|^2 + |\hat{u}|^2$ is constant. Fix any $r_* \in (0, \rho_y)$ and define the pair $(\varphi_{\varepsilon,y,0}, \zeta_{\varepsilon,y,0})$ using (6-8) with $\varepsilon \leq r_*$. Then the restriction of φ_y to the circle $|x - y|^2 + |\hat{u}|^2 = r_*^2$ defines a map from S^1 to S^1 that is homotopic to the restriction of $\varphi_{\varepsilon,y,0}$.

Let r now denote the radial coordinate on the disk of radius ρ_y in $\mathbb{R} \times I_*$ centered on $(y, 0)$. Reintroduce the function w and set w_y to be the function on this same disk given by $w(2\rho_y^{-1}r - 1)$. This function is 1 where $r \geq \frac{5}{8}\rho_y$ and it is 0 where $r \leq \frac{3}{8}\rho_y$.

Fix $\varepsilon < \frac{1}{8}\rho_y$ and use (6-8) to define the pair $(\varphi_{\varepsilon,y,0}, \zeta_{\varepsilon,y,0})$. The function $\varphi_{\varepsilon,y,0}$ can be written on the $r \in (\varepsilon, \rho_y)$ part of this disk as $\varphi_{\varepsilon,y,0} = \varphi_y + \varphi'_{\varepsilon,y,0}$, where $\varphi'_{\varepsilon,y,0}$ is an \mathbb{R} -valued function on this part of the disk.

The pairs $(\varphi_{\varepsilon,y,0}, \zeta_{\varepsilon,y,0})$ and (φ_y, ζ_y) with the function w_y are used next to define functions $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ on the complement of $(y, 0)$ in $\mathbb{R} \times I_*$. These are given by $(\varphi_{\varepsilon,y,0}, \zeta_{\varepsilon,y,0})$ on the complement of the radius ρ_y disk centered at $(y, 0)$, and given inside this disk by

$$(6-5) \quad (\varphi_{\varepsilon,y} = \varphi_y + w_y\varphi'_{\varepsilon,y,0}, \zeta_{\varepsilon,y} = \zeta_y + w_y\zeta_{\varepsilon,y,0}).$$

This pair is such that the $(\varphi, \zeta) = (\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ version of the following conditions are obeyed:

- (6-6) • The function φ where $|\hat{u}| = R + \frac{1}{2} \ln z_*$ is the function $\varphi^{S_0}(\cdot, z_*)$ that appears in (2-12).

- The graph $(x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi(x, \hat{u}), h = \zeta(x, \hat{u}))$ lies in $\mathbb{R} \times \mathcal{X}$, and as a consequence, this graph is in the domain of the map Ψ_p .
- There exists a purely S -dependent (or \mathcal{K} -compatible) constant $\kappa_c \geq 1$ with the following significance: Assume that $z_* < \kappa_c^{-1}$ and that $\delta^2 < \kappa_c^{-1} z_*$. The Ψ_p -image of the $|\hat{u}| \geq R + \frac{1}{2} \ln z_* - 6$ part of the graph is J -holomorphic where $1 - 3 \cos^2 \theta \leq \kappa_c^{-1}$.
- There is a constant $\kappa_{cc} > 1$ with the following property: the Ψ_p -image of the graph is J -holomorphic where either $|s| \geq \kappa_{cc}$ or $1 - 3 \cos^2 \theta \leq \kappa_{cc}^{-1}$.
- The Ψ_p -image of the $|x - y| \leq \varepsilon$ part of this graph is J -holomorphic.
- Each constant $s \ll -1$ slice of the Ψ_p -image of this graph consists of a single arc that is isotopic rel boundary in $\mathcal{H}_{p_*}^+$ to γ_{p_-} . The corresponding family of such arcs converges as an isotopy rel boundary to γ_{p_-} as $s \rightarrow -\infty$.
- Each constant $s \gg 1$ slice of the Ψ_p -image of this graph consists of two components:
 - (a) The first is an arc that is isotopic rel boundary in $\mathcal{H}_{p_*}^+$ to γ_{p_-} . The corresponding family of such arcs converges as an isotopy rel boundary to γ_{p_+} as $s \rightarrow -\infty$.
 - (b) The second is an embedded circle that is isotopic in $\mathcal{H}_{p_*}^+$ to $\hat{\gamma}_p^+$. The corresponding family of such circles converges as an isotopy rel boundary to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$.

By way of explanation for the third bullet, Lemma 4.6 supplies a purely S -dependent (or \mathcal{K} -compatible) version of κ_c such that the Ψ_p -image of the $|\hat{u}| \geq R + \frac{1}{2} \ln z_* - 8$ part of the graph of $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ is J -holomorphic where $1 - 3 \cos^2 \theta < \kappa_c^{-1}$. This fact is used in a moment. The constant κ_{cc} for the fourth bullet is supplied by Lemma 4.7. The remaining bullets follow directly from the definition of $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$.

Fix $\varepsilon \in (0, \frac{1}{8})$ so that (6-6) holds.

6B Deformations to $(\varphi^{p_0}, \zeta^{p_0})$

This subsection studies a $[0, 1]$ -parametrized family of equations for a map from $(\mathbb{R} \times I_*) \setminus (y, 0)$ that obeys (6-6). The initial equation is satisfied by $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ and a solution to the final equation can serve as $(\varphi^{p_0}, \zeta^{p_0})$ since the Ψ_p -image of its graph in $\mathbb{R} \times \mathcal{X}$ is J -holomorphic.

To define these equations, use $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ for (φ, ζ) in the left-hand side of (3-4) and write the resulting pair of functions on $(\mathbb{R} \times I_*) \setminus (y, 0)$ as (g_1, g_2) . These have compact support in $(\mathbb{R} \times I_*) \setminus (y, 0)$, a consequence of the third, fourth and fifth bullets of (6-6).

Use Lemma 4.6 to find a purely S -dependent (or \mathcal{K} -compatible) constant $r_c \geq 1$ such that the Ψ_p -image of the $|\hat{u}| \geq R + \frac{1}{2} \ln z_* - 8$ part of the graph of $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ is J -holomorphic where $1 - 3 \cos^2 \theta \leq r_c^{-1}$. Use Ψ_p to view the angle θ as a function on $\mathbb{R} \times \mathcal{X}$. Having done so, let χ_c denote function on $\mathbb{R} \times \mathcal{X}$ given by

$$(6-7) \quad \chi_c = 1 - \chi(2(R + \frac{1}{2} \ln z_* - 7 - |\hat{u}|))\chi(4r_c(1 - 3 \cos^2 \theta) - 3).$$

This function is equal to 1 where the Ψ_p -image of the graph of $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ is not J -holomorphic, and it is equal to zero on the part of this graph that is described in the third bullet of (6-6).

Use the fourth bullet of (6-6) to find $r > 1$ such that the Ψ_p -image of the graph of $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$ is J -holomorphic where $1 - 3 \cos^2 \theta < \frac{1}{r_1}$. Let $\chi_{cc}: \mathbb{R} \times \mathcal{X} \rightarrow [0, 1]$ denote the function $\chi(2 - 2r(1 - 3 \cos^2 \theta))$. This function equals 1 where $1 - 3 \cos^2 \theta > \frac{1}{r}$, and it vanishes where $1 - 3 \cos^2 \theta \leq \frac{1}{2r}$.

The $\tau \in [0, 1]$ member of the family of equations asks for a pair (φ, ζ) that obeys (6-12) and is such that

$$(6-8) \quad \begin{aligned} a_1 \partial_x \varphi - \partial_{\hat{u}} \zeta - (1 - \tau)(\chi_c \chi_{cc})|_{h=\zeta} g_1 &= 0, \\ a_2 \partial_x \zeta + \partial_{\hat{u}} \varphi + b - (1 - \tau)(\chi_c \chi_{cc})|_{h=\zeta} g_2 &= 0. \end{aligned}$$

Here, as in (3-4), what is written as a_1, a_2 and b are functions on $(\mathbb{R} \times I_*) \setminus (y, 0)$ that are obtained from the eponymous set of functions of the variables (\hat{u}, h) by setting $h = \zeta$.

An open/closed strategy is used in what follows to construct a smoothly parametrized family $\{(\varphi_\tau, \zeta_\tau)\}_{\tau \in [0,1]}$ such that each $\tau \in [0, 1]$ member obeys (6-6) and (6-8) and with the $\tau = 0$ member given by $(\varphi_{\varepsilon,y}, \zeta_{\varepsilon,y})$. The image via Ψ_p of the graph in $\mathbb{R} \times \mathcal{X}$ of $(\varphi_{\tau=1}, \zeta_{\tau=1})$ is J -holomorphic since the (g_1, g_2) terms in (6-8) are absent when $\tau = 1$. This being the case, this $\tau = 1$ member of the family serves for the desired $(\varphi^{p_0}, \zeta^{p_0})$. To set up the open/closed argument, use \mathcal{I} to denote the subset of points $\tau \in [0, 1]$ for which the corresponding version of (6-8) has a solution. Since $\tau = 0$ is in \mathcal{I} , this set is not empty. Part 1 of this subsection explains why \mathcal{I} is open. The remaining four parts explain why \mathcal{I} is closed. Given that \mathcal{I} is not empty, and both open and closed, this set can only be $[0, 1]$. Section 6C completes the proof of the $m_p = -1$

version of Proposition 2.2 by explaining why there is but a single $(\varphi^{p_0}, \zeta^{p_0})$ with the desired properties.

Part 1 This part of the subsection proves that \mathcal{I} is open. To this end, suppose that $\tau \in \mathcal{I}$ and let $\mathfrak{h} = (\varphi_\tau, \zeta_\tau)$ denote a corresponding pair that obeys (6-6) and (6-8). It follows from Lemma 5.9 that there exists a ball $\mathbb{B}_* \subset \mathbb{H}_*$ about the origin with two essential properties. To state them, fix for the moment $(\varphi', \zeta') \in \mathbb{B}_*$ and use (φ, ζ) to denote $(\varphi_\tau + \varphi', \zeta_\tau + \zeta')$. Here is the first property: the graph of (φ, ζ) is in $\mathbb{R} \times \mathcal{X}$. The second property is that the assignment to any given $(\varphi', \zeta') \in \mathbb{B}_*$ of the corresponding $(\varphi, \zeta) = (\varphi_\tau + \varphi', \zeta_\tau + \zeta')$ version of the expressions on the left-hand side of (6-8) defines a smooth map from \mathbb{B}_* to \mathbb{L}_* . Let $\mathcal{F}_{\mathbb{B}}$ denote this map.

Let $I \subset [0, 1]$ denote an open neighborhood of τ , and define a map $\mathcal{F}: I \times \mathbb{B}_* \rightarrow \mathbb{L}_*$ by the rule

$$(6-9) \quad (\tau', \eta) \rightarrow \mathcal{F}(\tau', \eta) = \mathcal{F}_{\mathbb{B}}(\eta) - (\tau - \tau')(\chi_c \chi_{cc})|_{h=\zeta}(\mathfrak{g}_1, \mathfrak{g}_2).$$

Fix $\tau' \in I$. The differential of \mathcal{F} at $(\tau, 0)$ along the \mathbb{B} factor of its domain is an operator D that obeys (3-5), (3-6) and (5-22). Lemma 5.10 asserts that D maps \mathbb{H}_* surjectively to \mathbb{L}_* and so the differential of \mathcal{F} at any such $(\tau', 0)$ point is an isomorphism. This with the inverse function theorem supplies a smooth map, q , from a neighborhood $I' \subset I$ of τ to \mathbb{B}_* such that $\mathcal{F}(\tau', q(\tau')) = 0$. This being the case, any given $\tau' \in I'$ version of the pair $(\varphi, \zeta) = (\varphi_\tau, \zeta_\tau) + q(\tau')$ obeys the τ' version of (6-8). Note in this regard that (φ, ζ) is smooth, a fact that can be proved in a straightforward fashion using standard elliptic regularity techniques, for example those in [12, Chapter 6].

Granted that (φ, ζ) is smooth, and granted that the pair $(\mathfrak{g}_1, \mathfrak{g}_2)$ has compact support $(\mathbb{R} \times I_*) \setminus (y, 0)$, it follows that (φ, ζ) is described by (6-6). Thus $I' \subset \mathcal{I}$ and so \mathcal{I} is open.

Part 2 This part of the subsection outlines the proof that \mathcal{I} is closed. The proof starts with a lemma which describes a compact set in $\mathbb{R} \times \mathcal{H}_{p_*}^+$ with the following significance: if $\tau \in [0, 1]$ and if (φ, ζ) obeys (6-6) and (6-8), then the Ψ_p -image of the graph of (φ, ζ) is J -holomorphic on the complement of this set. The proof then derives τ - and (φ, ζ) -independent bounds for the integral of w over such a graph, and for the integral of $ds \wedge \hat{a}$ over any subset of the graph where s is bounded. These integrals bounds are used with Proposition II.5.5 to control the part of the graph that lies in the complement of the Ψ_p -inverse image of the aforementioned compact set. The resulting

control over this part of the graph is used in conjunction with some standard elliptic regularity tools to obtain τ -independent pointwise bounds for the derivatives to any given order for (φ, ζ) .

Granted all of this preliminary work, the proof proceeds as follows: Fix a point, τ_0 , in the closure of \mathcal{I} . A sequence $\{\tau_n, (\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ is chosen with $\{\tau_n\}_{n=1,2,\dots} \subset \mathcal{I}$ converging to τ_0 and with any given $n \in \{1, 2, \dots\}$ version of (φ_n, ζ_n) obeying (6-6) and solving the $\tau = \tau_n$ version of (6-8). The control described in the preceding paragraph is used to obtain a subsequence of $\{(\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ that converges to a pair that obeys (6-6) and the $\tau = \tau_0$ version of (6-8).

The details of the arguments proving \mathcal{I} is closed occupy the remaining Parts 3–6 of this subsection.

Part 3 The second lemma gives the needed integral bounds for w and $ds \wedge \hat{a}$. To set the stage, reintroduce κ_c from the third bullet of (6-6), and let $\mathcal{W} \subset \mathbb{R} \times \mathcal{H}_{p_*}^+$ denote the set of points from the $|\hat{u}| > R + \frac{1}{2} \ln z_* - 6$ part of $\mathbb{R} \times \mathcal{X}$ where $1 - 3 \cos^2 \theta < \frac{1}{2} \kappa_c^{-1}$.

Reintroduce κ_{cc} from the fourth bullet of (6-6) and let $\mathcal{W}_* \subset \mathbb{R} \times \mathcal{H}_p^+$ denote the subset where both $1 - 3 \cos^2 \theta \leq \kappa_{cc}^{-1}$ and $|s| \geq \kappa_{cc}$.

Lemma 6.1 *There exists a purely S -dependent (or \mathcal{K} -compatible) constant $\kappa > 1$ with the following significance: Fix $\tau \in [0, 1]$ and suppose that $\mathfrak{h} = (\varphi, \zeta)$ is a pair that obeys (6-6) and (6-8). Use $C_{\mathfrak{h}}$ to denote the Ψ_p -image of the graph of \mathfrak{h} . Let $I \subset \mathbb{R}$ denote an interval of length 1. Then*

$$\int_{(C_{\mathfrak{h}} \cap \mathcal{W}) \cap (I \times \mathcal{H}_{p_*}^+)} (ds \wedge \hat{a} + w) \leq \kappa.$$

There is a (φ, ζ) and τ -independent constant $\kappa_* > 1$ such that

$$\int_{C_{\mathfrak{h}} \cap \mathcal{W}_*} w \leq \kappa_* \quad \text{and} \quad \int_{(C_{\mathfrak{h}} \cap \mathcal{W}_*) \cap (I \times \mathcal{H}_{p_*}^+)} ds \wedge \hat{a} \leq \kappa_*.$$

Proof The proof that follows is given in seven steps.

Step 1 This step, Step 2 and Step 3 consider the integral that involves $C_{\mathfrak{h}} \cap \mathcal{W}$. To this end, note that this integral has two regions of support: the part of the graph of (φ, ζ) where the function \hat{u} is greater than $R + \frac{1}{2} \ln z_* - 6$ and the part where $\hat{u} < -R - \frac{1}{2} \ln z_* + 6$. The argument that follows considers the former region as the argument for the other is identical but for some sign changes.

Reintroduce r_c from (6-7) and set $\kappa = \chi(4r_c(1 - 3 \cos^2 \theta) - 1)$, here viewed using Ψ_p as a function on $\mathbb{R} \times \mathcal{X}$. The latter is equal to 1 where $1 - 3 \cos^2 \theta < \frac{1}{4}r_c^{-1}$ and it is equal to 0 where $1 - 3 \cos^2 \theta > \frac{1}{2}r_c^{-1}$. Let s_0 denote the midpoint of the interval I and introduce L to denote the function $\chi(2|s - s_0| - 2)\chi(2(R + \frac{1}{2} \ln z_* - 6 - \hat{u}))$. This function is equal to 1 where both $s \in I$ and $\hat{u} \geq R + \frac{1}{2} \ln z_* - 6$ and it is equal to 0 where $|s - s_0| > \frac{3}{2}$ or where \hat{u} is less than $R + \frac{1}{2} \ln z_* - 6.5$.

Step 2 Consider the integral of $-\sqrt{6} d(L^2 \kappa^2 h d\hat{\phi})$ over the graph of (φ, ζ) . Note that the integrand is supported on the part where the Ψ_p -image is J -holomorphic. Moreover, the integrand is equal to $\Psi_p^* w$ on $(C_{\mathfrak{h}} \cap \mathcal{W}) \cap (I \times \mathcal{H}_{p_*}^+)$, this being a consequence of the formula for w in (1-6). To say more about the integrand, it proves useful to introduce the coordinate $v = e^{-2(R - \hat{u})}$ for the $\hat{u} > R + \ln \delta$ part of $(\mathbb{R} \times I_*) \setminus (y, 0)$. Using (x, v) now to parametrize the graph, the 2-form $-\sqrt{6} d(L^2 \kappa^2 h d\hat{\phi})$ appears as

$$(6-10) \quad \sqrt{6} \{L^2(\kappa^2 + 2\zeta \kappa \partial_{\mathfrak{h}} \kappa)(\partial_x \varphi \partial_v \zeta - \partial_x \zeta \partial_v \varphi) + 2\sqrt{6} \kappa L \zeta (\partial_v(\kappa L) \partial_x \varphi - \partial_x(\kappa L) \partial_v \varphi)\} dx \wedge dv.$$

Granted that the Ψ_p -image of the graph of (φ, ζ) is J -holomorphic on the support of the form $-\sqrt{6} d(L^2 \kappa^2 h d\hat{\phi})$, and given the properties of J in Section 1C, so (φ, ζ) obey the Cauchy-Riemann equations on the support of this form when viewed as functions of (x, v) . This is to say that $\partial_x \varphi - \partial_v \zeta = 0$ and $\partial_x \zeta + \partial_v \varphi$. Thus, the 2-form in (6-10) is

$$(6-11) \quad \sqrt{6} \{L^2(\kappa^2 + 2\zeta \kappa \partial_{\mathfrak{h}} \kappa)(|\partial_x \varphi|^2 + |\partial_v \varphi|^2) + 2\sqrt{6} \kappa L \varphi (\partial_v(\kappa L) \partial_x \varphi - \partial_x(\kappa L) \partial_v \varphi)\} dx \wedge dv.$$

What follows is a key observation: the function $h \partial_{\mathfrak{h}} \kappa$ is nonnegative, this being a consequence (1-27) and the definition of Ψ_p . Granted that this is so, the function that multiplies the form $dx \wedge dv$ in (6-11) is no less than

$$(6-12) \quad \sqrt{6} \left\{ \frac{1}{2} L^2 \kappa^2 (|\partial_x \varphi|^2 + |\partial_v \varphi|^2) - c \delta_*^4 (|\partial_x(\kappa L)|^2 + |\partial_v(\kappa L)|^2) \right\},$$

where $c \geq 1$ is a purely S -dependent (or \mathcal{K} -compatible) constant. Given the definitions of κ and L , it follows that this function is no less than $-c z_*^{-2}$ with $c \geq 1$ being another purely S -dependent (or \mathcal{K} -compatible) constant. This means that the integral of $-\sqrt{6} d(L^2 \kappa^2 h d\hat{\phi})$ over the complement of the part of the Ψ_p -inverse image of $(C_{\mathfrak{h}} \setminus \mathcal{W}) \cap (I \times \mathcal{H}_{p_*}^+)$ is bounded from below by $-c z_*^{-1}$ where c is again purely S -dependent (or \mathcal{K} -compatible).

Step 3 Stokes' theorem equates the integral of $-\sqrt{6} d(L^2 K^2 h d\hat{\phi})$ with the line integral

$$(6-13) \quad \int_{\mathbb{R}} (K^2 L^2 \zeta d\varphi)|_{\hat{u}=R+\frac{1}{2} \ln z_*}.$$

Given the boundary condition in the first bullet of (6-6) and given the fourth bullet of Proposition 2.1, this integral is no greater than a purely S -dependent (or \mathcal{K} -compatible) constant.

Lemma 6.1's bound for the integral of w over $(C_h \cap \mathcal{W}) \cap (I \times \mathcal{H}_p^+)$ follows directly from the conclusions of the preceding paragraph and from the conclusions of Step 3 because w is a nonnegative multiple of $dx \wedge d\hat{u}$ on the J -holomorphic part of C_h which contains $C_h \cap \mathcal{W}$.

The asserted bound for the integral of $ds \wedge \hat{a}$ over $(C_h \cap \mathcal{W}) \cap (I \times \mathcal{H}_p^+)$ follows from the fact that the latter form when pulled back by Ψ_p and written in terms of the coordinates (x, v) is $dx \wedge dv$. This the case, its integral is bounded by z_* .

Step 4 The set $\Psi_p^{-1}(\mathcal{W}_*)$ is a compact set in $\mathbb{R} \times \mathcal{X}$ and so the image in $R \times I_*$ of $\Psi_p^{-1}(\mathcal{W}_*)$ via the projection is compact. It follows from this that there exist $d \geq 1$ whose significance is explained next. To set the background, introduce W to denote the portion of $(\mathbb{R} \times I_*) \setminus (y, 0)$ where the following conditions are met:

$$(6-14) \quad |x| \leq d, \quad |x - y| + |\hat{u}| \geq d^{-1} \quad \text{and} \quad |\hat{u}| \leq R + \ln \delta - \frac{1}{8}.$$

Let $\tau \in [0, 1]$ and suppose that $h = (\varphi, \zeta)$ obeys (6-6) and (6-8). Then the Ψ_p -image of the graph of (φ, ζ) over the complement of W lies in the complement of \mathcal{W} . This last observation implies the existence of $c \geq 1$ such that:

- (6-15) • The functions a_1 and a_2 that appear in (6-6) are bounded from below on W by c^{-1} and bounded from above on W by c . Likewise, b is bounded on W by c .
- The metric on the $(x, \hat{u}) \in W$ part of the graph of (φ, ζ) coming from the Euclidean metric from $\mathbb{R} \times I_* \times \mathbb{R}/(2\pi\mathbb{Z}) \times (-\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ pushes forward via Ψ_p to a metric on $C_h \cap \mathcal{W}$ that is bounded respectively above and below c and c^{-1} times the metric that comes from $\mathbb{R} \times \mathcal{H}_{p_*}^+$.

The second bullet in particular implies that the Euclidean inner product on $\mathbb{R} \times \mathcal{X}$ can be used when deriving an upper bound for $-w$ and for $-ds \wedge \hat{a}$.

Step 5 Use (1-6) and (1-30) to write

$$(6-16) \quad \bullet \quad \Psi_p^* w = \sqrt{6} d\hat{\phi} \wedge dh + \alpha_w d\hat{u} \wedge dh,$$

$$\bullet \quad \Psi_p^*(ds \wedge \hat{a}) = \nu_0 dx \wedge d\hat{u} + \chi^2 \Gamma d\hat{\phi} \wedge dh + \alpha_x dx \wedge d\hat{\phi} + \alpha_{\hat{u}} d\hat{u} \wedge dh,$$

where ν_0 and Γ are positive and such that the following is true: there exists $c_0 \geq 1$ such that $\nu_0 > c_0^{-1}$ and $|\Gamma| + |\alpha_w| + |\alpha_x| + |\alpha_y| \leq c_0$ on $\Psi_p^{-1}(\mathcal{W})$.

Step 6 Suppose that $\tau \in [0, 1]$ and that $\mathfrak{h} = (\varphi, \zeta)$ obeys (6-12) and (6-13). Use the coordinates $(x, \hat{u}) \in (\mathbb{R} \times I_*) \setminus (y, 0)$ for the graph of (φ, ζ) to parametrize $\Psi_p^{-1}(C_{\mathfrak{h}})$. This coordinate map pulls back the form $d\hat{\phi} \wedge dh$ as

$$(6-17) \quad (\partial_x \varphi \partial_{\hat{u}} \zeta - \partial_x \zeta \partial_{\hat{u}} \varphi) dx \wedge d\hat{u}.$$

Use (6-8) to see that the latter expression can be written as

$$(6-18) \quad (a_1 |\partial_x \varphi|^2 + a_2 |\partial_x \zeta|^2 + (1 - \tau) \chi (\partial_x \varphi g_1 + \partial_x \zeta g_2)) dx \wedge d\hat{u}.$$

Since $dx \wedge d\hat{u}$ gives the proper orientation for $C_{\mathfrak{h}}$, what is written above with (6-6) and (6-16) imply that

$$(6-19) \quad w|_{\wedge^2 TC_{\mathfrak{h}}} \geq c_1^{-1} (|d\varphi|^2 + |d\zeta|^2) dx \wedge d\hat{u} - c_1 dx \wedge d\hat{u},$$

$$(ds \wedge \hat{a})|_{\wedge^2 TC_{\mathfrak{h}}} \geq -c_1 dx \wedge d\hat{u},$$

where $c_1 \geq 1$ enjoys a τ - and (φ, ζ) -independent upper bound on W . These formulas supply a (φ, ζ) - and τ -independent lower bound for the respective parts of $C_{\mathfrak{h}}$ where w and $ds \wedge \hat{a}$ are negative multiples of the area form.

Step 7 To see about the integral of w over the whole of $C_{\mathfrak{h}}$, note first that the argument used at the start of Lemma 4.5 has what are purely cosmetic modifications that prove that the integral of w over $C_{\mathfrak{h}}$ is finite. With (1-6) used to identify w on $\mathbb{R} \times \mathcal{H}_{p_*}^+$ as $w = d(\chi(1 - 3 \cos^2 \theta d\theta) - \sqrt{6} d(\hbar d\phi))$, this same argument justifies an application of Stokes' theorem to write the integral of w over $C_{\mathfrak{h}}$ as a sum of five terms. The first two are integrals over the arcs γ_{p_+} and γ_{p_-} . That over γ_{p_+} is given by (II.5-9) and that over γ_{p_-} is (-1) times the γ_{p_-} version of (II.5-9). Steps 2 and 3 in the proof of Proposition II.5.1 bound the total contribution from these two terms by a purely S -dependent (or \mathcal{K} -compatible) constant. The third term in the sum is the integral of $-\sqrt{6} \hbar d\phi$ over $\hat{\gamma}_p^+$. This is $\frac{4\pi\sqrt{2}}{3}(\chi_0 + 4e^{-2R})$. The last two terms are the integrals of $\hbar d\phi$ over the two boundary components of $C_{\mathfrak{h}}$. Up to an overall plus/minus sign, one is the integral of $\zeta \partial_x dx$ along the $\hat{u} = R + \frac{1}{2} \ln z_*$ boundary of $\mathbb{R} \times I_*$ and the other is the integral of $\varphi \partial_x \varphi dx$ along the $\hat{u} = -R - \frac{1}{2} \ln z_*$ boundary of $R \times I_*$. Both

integrals are bounded by the integral of $\frac{4}{\sqrt{3}}\delta_*^2|\partial_x\varphi^{S_0}|$, where φ^{S_0} is the value at z_* of the function that appears in (2-12). In particular, it follows from what is asserted by the fourth bullet of Proposition 2.1 that both versions of the latter integral are no greater than a purely S -dependent (or \mathcal{K} -compatible) constant.

A bound on the integral of $ds \wedge \hat{a}$ over $C_h \cap (I \times \mathcal{H}_{p_*}^+)$ is obtained by mimicking what is done in Step 4 of the proof of Proposition II.5.1. Note in this regard that the integration by parts done in the latter proof has no contributions from the boundary of C_h because the 1-form \hat{a} annihilates the tangent space of each level set of f in M_δ . □

Part 4 This first lemma below supplies a τ - and (φ, ζ) -independent $\mathcal{O}(z_*)$ upper bound for $|\zeta|$ near the boundary of $\mathbb{R} \times \mathcal{X}$. This lemma states what Lemma 4.6 states for the $\Delta_p = 0$ case. The second lemma uses what is said in Lemma 6.2 to obtain τ - and (φ, ζ) -independent, positive lower bounds for the function $1 - 3 \cos^2 \theta$ on various parts of the corresponding surface C_h . The latter are the analogs of those given in the $m_p = 0$ case by Lemma 4.7.

Lemma 6.2 *There exists a purely S -dependent (of \mathcal{K} -compatible) constant $\kappa \geq 1$ such that if $\delta^2 < \kappa^{-1}z_*$, then the following is true: Suppose that $\tau \in [0, 1]$ and that (φ, ζ) is described by (6-6) and obeys (6-8). Then $|\zeta|$ is bounded by κz_* where $|\hat{u}| > R + \frac{1}{2} \ln z_* - 6$.*

Proof Except for two modifications, the argument is identical to that given to prove Lemma 4.6. The first modification replaces the appeal to Lemma 4.5 with an appeal to the first inequality of Lemma 6.1. The second modification concerns the hypothetical nonsense loop in each large n version of C_{p_0n} , this being the loop that must define a nonzero homology class in $\mathbb{R} \times \mathcal{H}_{p_*}^+$. Although C_{p_0n} in this case does contain a loop that generates the homology of $\mathbb{R} \times \mathcal{H}_{p_*}^+$, the hypothetical nonsense loop would sit entirely in either the $u > 0$ or the $u < 0$ part of C_{p_0n} . Each of these parts is contractible in $\mathbb{R} \times \mathcal{H}_{p_*}^+$, so no such loop can exist. □

The next lemma gives the promised lower bounds for $1 - 3 \cos^2 \theta$.

Lemma 6.3 *There exists $\kappa > 1$ and, given $\varepsilon \in (0, 1]$, there exists $\kappa_\varepsilon > 1$, and these have the following significance: Suppose that $\tau \in [0, 1]$ and that (φ, ζ) is described by (6-6) and is a solution to (6-8).*

- $\cos \theta < \frac{1}{\sqrt{3}} - \kappa_\varepsilon^{-1}$ on the $|u| \geq \varepsilon$ part of C_h .

- $\cos \theta < \frac{1}{\sqrt{3}} - \kappa^{-1}$ on the $s < -\kappa$ part of C_h .
- $\cos \theta > -\frac{1}{\sqrt{3}} + \kappa^{-1}$ on the whole of C_h .

Proof The proof has seven steps.

Step 1 This step proves that there exists a (φ, ζ) - and τ -independent $\kappa_\varepsilon \geq 1$ such that $1 - 3 \cos^2 \theta > \kappa_\varepsilon^{-1}$ on the part of C_h where $|u| > \varepsilon$. This implies what is asserted by the first bullet. To start, use Lemma 6.2 to choose a (φ, ζ) - and τ -independent constant $r \geq 1$ such that the $1 - 3 \cos^2 \theta \leq \frac{1}{r}$ part of C_h is J -holomorphic and so that $1 - 3 \cos^2 \theta > \frac{1}{r}$ on the boundary of C_h and on the segments γ_{p+} and γ_{p-} that are used to describe the arc components of the large $|s|$ slices of C_h . Granted this, the proof of Lemma 4.7 can be copied to prove the existences of κ_ε but for two modifications. The first modification replaces the appeal to Lemma 4.5 with an appeal to Lemma 6.1. As in the proof of Lemma 6.2, the second modification concerns the hypothetical nonsense loop. This loop would sit entirely in either the $u > 0$ part of C_{p0n} or in the $u < 0$ part, and both parts are contractible. Thus, no such loop can exist.

Step 2 This step and Steps 3–6 explain why there exists a (φ, ζ) - and τ -independent constant $\kappa > 1$ such that $1 - 3 \cos^2 \theta > \kappa^{-1}$ on the $s < -\kappa$ part of C_h . The existence of such a constant implies what is asserted by the second bullet. Existence is proved by assuming to the contrary that no such constant exists so as to derive nonsense. Granted this assumption, there exists a sequence $\{\tau_n, (\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ with the following properties: First, any given $n \in \{1, 2, \dots\}$ version of $\tau_n \in [0, 1]$ and (φ_n, ζ_n) obeys (6-6) plus the $\tau = \tau_n$ version of (6-8). Furthermore, the corresponding $h = (\varphi_n, \zeta_n)$ version of C_h has a point where both $1 - 3 \cos^2 \theta < \frac{1}{n}$ and $s < -n$. Given what is said by Lemma II.4.8 nothing is lost by taking this point to be a local minimum of $1 - 3 \cos^2 \theta$ and thus a point where $u = 0$. No generality is lost by assuming that such a point occurs where $\cos \theta \sim \frac{1}{\sqrt{3}}$.

Introduce from Part 2 of Section 5A the tubular neighborhood $U_+ \subset \mathcal{H}_{p^*}^+$ of the $\cos \theta = \frac{1}{\sqrt{3}}$ and $u = 0$ integral curve of v , this being the loop $\hat{\gamma}_p^+$. Let r be as in Step 1, and use what is said in Step 1 to choose $\varepsilon < \frac{1}{r}$ so that points in C_h with $\cos \theta > \frac{1}{\sqrt{3}} - \varepsilon$ are mapped via the projection to $\mathcal{H}_{p^*}^+$ to a subset in U_+ with compact closure. Use $V \subset U_+$ to denote the subset of points where $1 - 3 \cos^2 \theta < \varepsilon$.

The ensuing discussion uses the coordinates $(s_+, \phi_+, \theta_+, u_+)$ for $\mathbb{R} \times U_+$ from Section 5B and (5-5). By way of reminder, the $(\theta_+ = 0, u_+ = 0)$ locus is $R \times \hat{\gamma}_p^+$, and any given constant s_+ and ϕ_+ disk is J -holomorphic.

For each $n \in \{1, 2, \dots\}$, choose a point in the $\mathfrak{h} = (\varphi_n, \zeta_n)$ version of $C_{\mathfrak{h}}$ with s_+ coordinate less than $-n$ and with $\cos \theta > \frac{1}{\sqrt{3}}(1 - \frac{1}{n})^{1/2}$. Use s_n to denote the value of s_+ at this point; and use $V_n \subset C_{\mathfrak{h}}$ to denote the component of this chosen point in the $\mathbb{R} \times V$ part of $C_{\mathfrak{h}}$. Fix $\varepsilon' \in (\frac{1}{4}\varepsilon, \frac{1}{2}\varepsilon)$ so that each index n version of V_n is transversal to the locus where $1 - 3 \cos^2 \theta = \varepsilon'$. Introduce $V' \subset V$ to denote the $1 - 3 \cos^2 \theta \leq \varepsilon'$ part and use V'_n to denote the connected component of the $\mathbb{R} \times V'$ part of V_n that contains the chosen point where $s_+ = s_n$ and $1 - 3 \cos^2 \theta < \frac{1}{n}$. Let $\partial V'_n \subset V_n$ denote the boundary of V'_n . The ensuing discussion here and in Steps 3–5 assumes that at least one of the following two conditions hold for an infinite subset of $n \in \{1, 2, \dots\}$:

- (6-20) • s_+ is bounded from above on V'_n .
 • There are points on $\partial V'_n$ where $s_+ \geq s_n$.

The case when neither condition holds when n is sufficiently large is treated in Step 6.

Assume now that one or the other of the conditions in (6-20) holds for all indices n . If the first condition holds, use s_{n+} to denote the maximum value of s_+ on V'_n . If the first condition fails but the second condition holds, use s_{n+} to denote the minimum value of s_+ on the $s \geq s_n$ part of $\partial V'_n$. Meanwhile, s_+ is bounded from below on V'_n in any event. If it is the case that $s \leq s_n$ on $\partial V'_n$ set s_{n-} to denote the maximum value of s_+ on the $s \leq s_n$ part of $\partial V'_n$. If $s > s_n$ on $\partial V'_n$, then set s_{n-} to be the minimum of s on V'_n . Extra arguments are needed when the following occurs:

- (6-21) Neither $\{s_n - s_{n+}\}_{n=1,2,\dots}$ nor $\{s_n - s_{n-}\}_{n=1,2,\dots}$ have convergent subsequences.

The next step assumes that one or both of these requirements is violated.

Step 3 Assume that (6-21) is violated. Pass to a subsequence (hence renumbered consecutively from 1) such that one of the sequences in question is convergent. For each large n in $\{1, 2, \dots\}$, translate V_n by $-s_n$ along the \mathbb{R} factor of $\mathbb{R} \times V_+$ and let V_{n*} denote the resulting J -holomorphic submanifold. This translate is a properly embedded submanifold in $\mathbb{R} \times V$. Moreover, s_+ on V_n takes value zero, and it is bounded either from above or below by some n -independent constant s_* .

Use Proposition II.5.5 with Lemma II.5.6 and Lemma 6.1 to obtain a subsequence of $\{V_{n*}\}_{n=1,2,\dots}$ (hence renumbered consecutively) that converges on compact subsets of $\mathbb{R} \times V_+$ in the manner dictated by Proposition II.5.5. The geometric limit is a closed, J -holomorphic subvariety of $\mathbb{R} \times V_+$ that sits where $1 - 3 \cos^2 \theta \leq 0$ and contains a point where $1 - 3 \cos^2 \theta = 0$. Moreover, s is bounded from either above or below on

this subvariety. But this is impossible because the bound $1 - 3 \cos^2 \theta \leq 0$ with a point of equality implies that the limit subvariety is $\mathbb{R} \times \widehat{\gamma}_p^+$.

Step 4 Now assume that (6-21) holds. Construct each index n version of V_{n*} as directed in Step 3. In this case, Lemma 6.1 with Proposition II.5.5 and Lemma II.5.6 provide a subsequence of $\{V_{n*}\}_{n=1,2,\dots}$ (hence relabeled consecutively from 1) that converges in the manner dictated by Proposition II.5.5 to $\mathbb{R} \times \widehat{\gamma}_p^+$.

Let π_+ denote the projection map $(s_+, \phi_+, \theta_+, u_+) \rightarrow (s_+, \phi_+)$ from $\mathbb{R} \times U_+$ to $\mathbb{R} \times \widehat{\gamma}_p^+$. The submanifold V'_n has positive local intersection numbers with the constant (s_+, ϕ_+) disks in $\pi(V'_n)$. This has the following consequence: Let $D \subset \pi(V'_n)$ denote a disk whose inverse image in V'_n is disjoint from the boundary. Then the restriction of π_+ to $\pi_+^{-1}(D) \cap V'_n$ is a finite-to-1, branched cover with purely positive ramification points. This observation has an important consequence that is described in a moment. To set the stage, suppose that $\nu \subset \mathbb{R} \times \widehat{\gamma}_p^+$ is an embedded, oriented loop with the following three properties: First, ν is -1 times the generator of $H_1(\mathbb{R} \times \widehat{\gamma}_p^+; \mathbb{Z})$. Second, $\nu \in \pi_+(V'_n)$ and $\pi_+^{-1}(\nu) \cap V'_n$ is disjoint from the boundary of V'_n . Third, ν does not contain any branch points of π_+ on V'_n . Granted these conditions, the projection map $\pi_+ : \pi_+^{-1}(\nu) \cap V'_n \rightarrow \nu$ must be 1-to-1 on each component of $\pi_+^{-1}(\nu) \cap V'_n$. This is proved in the next paragraph.

To prove this is, note that each component of $\pi_+^{-1}(\nu) \cap V'_n$ comes via the graph $(x, \widehat{u}) \rightarrow (x, \widehat{u}, \widehat{\phi} = \varphi_n(x, \widehat{u}), h = \zeta_n(x, \widehat{u}))$ of an embedded loop in $(\mathbb{R} \times I_*) \setminus (y, 0)$ that has positive linking number in $R \times I_*$ with the point $(y, 0)$. As each such loop is embedded, it must have linking number 1 with $(y, 0)$ and so its image via Ψ_{p*} must be -1 times the generator $H_1(\mathbb{R} \times \mathcal{H}_{p*}^+; \mathbb{Z})$. This would not be the case were $\pi_+^{-1}(\nu) \cap V_n$ a nontrivial covering map.

Step 5 Each index n version of V'_n has strictly positive and locally constant intersection number with the fibers of π_+ over $(s_{n-}, s_{n+}) \in \mathbb{R} \times \widehat{\gamma}_p^+$. Let m_n denote this intersection number. Granted what was said in the previous step, it must be the case that the $s_+ \in (s_{n-}, s_{n+})$ part of V'_n has m_n components and π_+ restricts to each component so as to map it diffeomorphically onto $(s_{n-}, s_{n+}) \times \widehat{\gamma}_p^+$. Fix a component of this part of V'_n whose closure has an $s_+ = s_{n+}$ point where $1 - 3 \cos^2 \theta = \varepsilon'$. Use A_n to denote the chosen component.

Let \mathcal{V}_n denote the translate of V_n by the constant factor $-s_{n+}$ of the s_+ coordinate. Let $\mathcal{V}'_n \subset \mathcal{V}_n$ denote the corresponding translate of V'_n and let $\mathcal{A}_n \subset \mathcal{V}'_n$ denote the

translate of A_n . The subvariety \mathcal{V}_n intersects $(s_{n-} - s_{n+}, 1) \times V_+$ as a J -holomorphic submanifold with an $s_+ = 0$ point in the closure of A_n where $1 - 3 \cos^2 \theta = \varepsilon'$.

Use Proposition II.5.5 and Lemma II.5.6 with Lemma 6.1 to find a subsequence of $\{\mathcal{V}_n\}_{n=1,2,\dots}$ (henceforth renumbered consecutively from 1) that converges on compact subsets of $(-\infty, 1) \times U_+$ to a nonempty, properly embedded J -holomorphic subvariety. Let \mathcal{V} denote this limit and let \mathcal{V}' and \mathcal{A} denote the respective subsets of \mathcal{V} that arise from the corresponding limits of $\{\mathcal{V}'_n\}_{n=1,2,\dots}$ and $\{\mathcal{A}_n\}_{n=1,2,\dots}$.

The function $1 - 3 \cos^2 \theta \geq 0$ on \mathcal{V} and so \mathcal{V} cannot contain a fiber of the projection π_+ for the latter has points where $1 - 3 \cos^2 \theta < 0$. This implies that $1 - 3 \cos^2 \theta = \varepsilon'$ at a point where $s_+ = 0$ in the closure of \mathcal{A} and so $1 - 3 \cos^2 \theta > 0$ on the $(-\infty, -1]$ part of \mathcal{A} . Meanwhile, the functions $1 - 3 \cos^2 \theta$ and u have limit 0 as s_+ and thus s limit to $-\infty$ on \mathcal{A} . In addition, \mathcal{A} has intersection number 1 with each $s_+ \leq -1$ fiber of π_+ .

Granted these properties, it follows that the $s \ll -1$ part of \mathcal{A} is given via (5-5)–(5-7) by a map η of the sort that is described by the second bullet of Proposition 5.1. This last conclusion is nonsense for the following reason: the function $1 - 3 \cos^2 \theta$ is positive on \mathcal{A} , but any given subvariety that comes via the second bullet of Proposition 5.1 has points where $1 - 3 \cos^2 \theta < 0$ and s is less than any specified value.

Step 6 This step considers the case where neither bullet in (6-20) is satisfied. If this is the case, then the $s_+ < s_{n-}$ part of V'_n will contain the end of $C_{\mathfrak{h}}$ whose constant s slices converge in an isotopic fashion as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$. This implies, in particular, that the map π_+ restricts to the $s < s_{n-}$ part of V'_n as a 1-to-1 diffeomorphism onto the $s \leq s_{n-}$ part of $\mathbb{R} \times \hat{\gamma}_p^+$. Let \mathcal{V}_n denote the translate of V_n that adds $-s_{n-}$ to the s_+ coordinate of each point. Use Lemma 6.1 with Proposition II.5.5 and Lemma II.5.6 to obtain a subsequence of $\{\mathcal{V}_n\}_{n=1,2,\dots}$ (hence renumbered consecutively from 1) that converges on compact subsets of $(1, \infty) \times V$ in the manner dictated by Proposition II.5.5. Let \mathcal{V} denote this limit and let $\mathcal{V}' \subset \mathcal{V}$ denote the part that comes as a limit from $\{V'_n\}$. Let $\mathcal{A} \subset \mathcal{V}'$ denote the $s_+ \geq 1$ part of \mathcal{V}' . A repeat of the arguments from the second to last paragraph in Step 5 prove that \mathcal{A} is an embedded cylinder with boundary with the following properties: First, $1 - 3 \cos^2 \theta > 0$ on \mathcal{A} , but also $1 - 3 \cos^2 \theta \leq \frac{1}{2}\varepsilon$ on \mathcal{A} . Second, both θ_+ and u_+ limit to 0 on \mathcal{A} as $s_+ \rightarrow \infty$. Third, the projection π_+ maps \mathcal{A} diffeomorphically to $[1, \infty) \times \hat{\gamma}_p^+$.

What is said in the first bullet of Proposition 5.1 and the first bullet in Lemma 6.4 implies that the constant ε can be chosen in advance so that \mathcal{A} is the graph of a smooth

map $(s_+, \phi_+) \rightarrow \eta = (\theta_+ = a_+(s_+, \phi_+), u_+ = b_+(s_+, \phi_+))$ that obeys (5-7) on $[1, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ and has the form

$$(6-22) \quad \eta = c_0(e^{-\lambda_1 s_+} + e_1, 0) + c_1 \eta_{1+} + \epsilon_1,$$

where $c_0 > 0$ and $c_1 > 0$, and e_1 is as given in (5-8) and η_{1+} is some ϕ_1 version of (5-4). Meanwhile, ϵ_1 is as described in the first bullet of Proposition 5.1.

The missing $\hat{u} = 0$ point in $\mathbb{R} \times I_*$ that defines the domain of (φ_n, ζ_n) has coordinates $(y, 0)$. Let α_y denote the strictly positive constant from Lemma 5.2 that is assigned to y . Let \mathcal{E}_n denote the index n version of the end $C_{\mathfrak{h}=(\varphi_n, \zeta_n)}$ whose constant s slices converge as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$. Lemmas 5.3 and 5.6 imply that the $s \gg 1$ part of \mathcal{E}_n is given via (5-5)–(5-6) by a map from Proposition 5.1’s first bullet that can be written as

$$(6-23) \quad \eta(n) = \alpha_y(e^{-\lambda_1 s_+} + e_1, 0) + c_{1n} \eta_{1+} + \epsilon_{1n},$$

where $c_{1n} \in \mathbb{R} \setminus 0$, where η_{1+} is given by some $\phi_1 = \phi_{1n}$ version of (5-4) and where ϵ_{1n} obeys the bounds given in Proposition 5.1 for the latter’s ϵ_1 . What follows is now a direct consequence of (6-22) and (6-23): given $\epsilon_* > 0$, there exists $n_* \geq 1$ such that

$$(6-24) \quad |\alpha_y - c_0 e^{-\lambda_1 |s_n - 1|} < \epsilon_*$$

when $n \geq n_*$. But this is nonsense given that $\alpha_y > 0$ and $\{s_n\}_{n=1,2,\dots}$ is unbounded from below.

Step 7 This step proves the third bullet of Lemma 6.3. The proof starts by assuming that the assertion is false so as to derive some nonsense. Granted there is no such κ , there is a sequence $\{\tau_n, (\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ with the following properties: First, any given $n \in \{1, 2, \dots\}$ version of $\tau_n \in [0, 1]$ and (φ_n, ζ_n) obeys (6-6) plus the $\tau = \tau_n$ version of (6-8). Furthermore, the corresponding $\mathfrak{h} = (\varphi_n, \zeta_n)$ version of $C_{\mathfrak{h}}$ has a point where both $\cos \theta < -\frac{1}{\sqrt{3}} + \frac{1}{n}$. Let $(s_n, p_n) \in \mathbb{R} \times \mathcal{H}_{p_*}^+$ denote such a point. Use Proposition II.5.5 and Lemma II.5.6 with Lemma 6.1 to see that the sequence $\{|s_n|\}_{n=1,2,\dots}$ cannot have convergent subsequences. What is said in the second bullet of Lemma 6.3 implies that $\lim_{n \rightarrow \infty} s_n = \infty$. This understood, the arguments used in Steps 2–6 can be used with only cosmetic changes to generate the desired nonsense. \square

Part 5 Let $\tau \in [0, 1]$ and let $\mathfrak{h} = (\varphi, \zeta)$ be as described in (6-6) and a solution to (6-8). The subvariety $C_{\mathfrak{h}}$ has an end whose constant s slices converge in an isotopic fashion as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$. The following is a consequence of Lemmas 5.3 and 5.10 and Proposition 5.1: There exist constants $s_{\mathfrak{h}} \geq 1$ and $c_{1\mathfrak{h}} \in \mathbb{R} \setminus 0$ and $\phi_{1\mathfrak{h}} \in \mathbb{R}/2\pi\mathbb{Z}$ such

that \mathcal{E} intersects the $s_+ \in [s_h, \infty)$ part of $\mathbb{R} \times U_+$ as a smooth, properly embedded submanifold with boundary on the $s_+ = s_h$ slice. Furthermore, this intersection is given by the graph of a smooth map as described in the first bullet of Proposition 5.1 with domain $[s_h, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ that has the form depicted in (6-23) with $c_1 = c_{1h}$ and with η_{1+} defined using $\phi_1 = \phi_{1h}$. The next lemma says something about the constants s_h and c_{1h} .

Lemma 6.4 *There exists $\kappa > 1$ with the following significance: Let $\tau \in [0, 1]$ and suppose that $h = (\varphi, \zeta)$ is described by (6-6) and that it obeys (6-8). Then the corresponding constant s_h can be chosen so that $s_h \leq \kappa$. Meanwhile, $|c_{1h}| \in [\kappa^{-1}, \kappa]$.*

Proof An upper bound for s_h is obtained using what are essentially the same arguments as those in Step 6 of the proof of Lemma 6.3. The salient difference in this case is that the assumption of no uniform upper bound gives a sequence $\{\tau_n, (\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ with the property that the corresponding sequence $\{s_{n-}\}_{n=1,2,\dots}$ is now unbounded from above instead of from below. This understood, the inequality in (6-24) is replaced by $|\alpha_y - c_0 e^{\lambda_1 s_{n-}}| < \varepsilon_*$, which cannot hold when n is large if $\{s_n\}$ diverges. □

The upper and lower bounds on c_h follow in a straightforward manner given the a priori bound on s_h . In fact, the upper bound follows from the constraint that $1 - 3 \cos^2 \theta$ is positive. The lower bound follows by assuming the contrary and deriving a contradiction from a limit submanifold that is described by (6-22) and obtained with the help of Proposition II.5.5 from a sequence $\{\tau_n, (\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ that has $|c_{1n}| < \frac{1}{n}$.

Let τ and $h = (\varphi, \zeta)$ be as described above. The $s \ll -1$ slices of C_h converge in an isotopic fashion to the arc γ_{p-} as $s \rightarrow -\infty$, and the $s \gg 1$ slices have a component that converges in an isotopic fashion as $s \rightarrow \infty$ to the arc γ_{p+} . The next lemma asserts that the convergence in both cases is suitably h - and τ -independent. This lemma is the analog of what Lemma 4.8 states for the $\Delta_p = 0$ case.

Lemma 6.5 *Given $\varepsilon > 0$, there exists $\kappa_\varepsilon > 1$ with the following significance: Suppose that $\tau \in [0, 1]$ and that (φ, ζ) is described by (6-6) and obeys (6-8). Then*

- $|\varphi(x, \hat{u}) - \phi(\gamma_{p-}|_{\hat{u}})| + |\zeta(x, \hat{u}) - h(\gamma_{p-})| < \varepsilon$ where $x < -\kappa_\varepsilon$,
- $|\varphi(x, \hat{u}) - \phi(\gamma_\tau|_{\hat{u}})| + |\zeta(x, \hat{u}) - h(\gamma_\tau)| < \varepsilon$ where $x > \kappa_\varepsilon$.

Proof What with Lemmas 6.3 and 6.4, the proof of Lemma 4.8 can be quoted in an essentially verbatim fashion to prove Lemma 6.5. □

Part 6 This last part of the subsection completes the proof that the set \mathcal{I} is closed. To start, suppose that $\tau \in [0, 1]$ and that $\mathfrak{h} = (\varphi, \zeta)$ is described by (6-6) and obeys (6-8). Lemmas 6.2–6.5 supply a τ - and \mathfrak{h} -independent disk $U \subset \mathbb{R} \times I_*$ centered on $(y, 0)$ and a compact set in $\mathbb{R} \times \mathcal{X}$ such that the graph of (φ, ζ) over $(\mathbb{R} \times I_*) \setminus U$ maps into this compact set and has uniform limits as $x \rightarrow \pm\infty$. This implies that the functions a_1, a_2 and b that appear in (6-8) have \mathfrak{h} - and τ -independent bounds for $(x, \hat{u}) \in (\mathbb{R} \times I_*) \setminus U$, and that a_1 and a_2 are bounded away from zero by \mathfrak{h} - and τ -independent, positive constants. As a consequence, standard elliptic regularity arguments of the sort that can be found in Chapter 6 of [12] can be employed to see that the absolute values of the derivatives of \mathfrak{h} on $(\mathbb{R} \times I_*) \setminus U$ to any given order have \mathfrak{h} - and τ -independent bounds. Lemma 6.2 ensures that the boundary values also enjoy \mathfrak{h} - and τ -independent bounds. Lemma 6.5 ensures that these τ - and \mathfrak{h} -independent derivative bounds hold uniformly as $x \rightarrow \pm\infty$.

Meanwhile, the part of $C_{\mathfrak{h}}$ that is parametrized via Ψ_{p_*} by U maps into the $\mathbb{R} \times U_+$ part of $\mathbb{R} \times \mathcal{H}_{p_*}$, and in particular the part where $s_+ \geq \kappa$ with κ as in Lemma 6.4. As a consequence, this part of $C_{\mathfrak{h}}$ can be described using (5-5)–(5-6) by a solution to (5-7) from the first bullet of Proposition 5.1 that has the form given in (6-23) with $c_{1\mathfrak{h}}$ as described in Lemma 6.5.

Granted all of this, suppose that $\{\tau_n\}_{n=1,2,\dots} \in \mathcal{I}$ converges to $\tau_0 \in [0, 1]$. For each index n , let (φ_n, ζ_n) denote a pair described by (6-6) and obeying the $\tau = \tau_n$ version of (6-8). Use the uniform bounds described in the preceding paragraphs for $\{(\varphi_n, \zeta_n)\}_{n=1,2,\dots}$ on $(\mathbb{R} \times I_*) \setminus U$ with the Arzelà–Ascoli theorem to obtain a subsequence that converges on $(\mathbb{R} \times I_*) \setminus U$ in the strong C^∞ topology to a pair (φ, ζ) whose graph over $(\mathbb{R} \times I_*) \setminus U$ maps into $\mathbb{R} \times \mathcal{X}$. Meanwhile, use the uniform bounds on the constants $\{c_{1\mathfrak{h}=(\varphi_n, \zeta_n)}\}_{n=1,2,\dots}$ to obtain a subsequence as above whose corresponding sequence of constants converges to a nonzero limit. Let $c_{1\mathfrak{h}}$ denote the latter. Granted this convergence, it follows that (φ, ζ) extends over $U \setminus (y, 0)$ to give a solution to the $\tau = \tau_0$ version of (6-8) that is described by (6-6). Thus $\tau_0 \in \mathcal{I}$.

6C Uniqueness

This subsection completes the proof of Proposition 2.2 by proving that there is only one pair $(\varphi^{p_0}, \zeta^{p_0})$ that obeys the conditions imposed by Proposition 2.2. This uniqueness assertion is one consequence of the lemma that follows.

Lemma 6.6 *Fix $y \in \mathbb{R}$ and $\tau \in [0, 1]$. There exists exactly one pair (φ, ζ) with the domain $(\mathbb{R} \times I_*) \setminus (y, 0)$ that obeys the conditions set forth in (6-6) and (6-8).*

Proof Suppose that $y \in \mathbb{R}$, that $\tau \in [0, 1]$ and that $(\varphi^{(0)}, \zeta^{(0)})$ and $(\varphi^{(1)}, \zeta^{(1)})$ are two pairs with domain $(\mathbb{R} \times I_*) \setminus (y, 0)$ that obey (6-6) and (6-8). Write

$$(6-25) \quad (\varphi^{(1)}, \varphi^{(1)}) = (\varphi^{(0)} + \varphi', \zeta^{(0)} + \zeta'),$$

where $\eta = (\varphi', \zeta')$ is a smooth map from $(\mathbb{R} \times I_*) \setminus (y, 0)$ to \mathbb{R} that obeys

$$(6-26) \quad \bullet \quad \varphi' = 0 \text{ where } |\hat{u}| = R + \frac{1}{2} \ln z_*,$$

$$\bullet \quad \lim_{|x| \rightarrow \infty} (\varphi', \zeta') = 0.$$

This pair obeys an equation of the form $D\eta = 0$ with D as described by (3-5) and (3-6) with coefficient functions a_1, a_2, b_1 and b_2 as described in the proof of Lemma 4.9.

Let C_0 and C_1 denote the respective $\mathfrak{h} = (\varphi^{(0)}, \zeta^{(0)})$ and $\mathfrak{h} = (\varphi^{(1)}, \zeta^{(1)})$ versions of $C_{\mathfrak{h}}$. These have corresponding ends where the constant s slices converge in an isotopic fashion to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$. These ends can be written via (5-5) and (5-6) in the $s_+ \geq c_0$ part of $\mathbb{R} \times U_+$ as graphs over $[s_+, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$. These respective maps, $\eta_{(0)}$ and $\eta_{(1)}$, are described by (6-23). Write $\eta_{(1)}$ as $\eta_{(0)} + \eta$ with $\eta = (a, b)$ here denoting a smooth map from $[c_0, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 with limit zero as $|s| \rightarrow \infty$. In particular, it follows from (6-23) that η can be written as $\eta = c\eta_{1+} + \epsilon$ where $c \in \mathbb{R}$ is nonzero if $\eta_{(0)} \neq \eta_{(1)}$, where η_{1+} is given by (5-4) with ϕ_1 determined by $\eta_{(0)}$ and $\eta_{(1)}$, and where $|\epsilon| \leq c_0 |c| e^{-(\lambda_{11} + 1/c_0)s_+}$. Note in particular that η , if not identically zero, it defines a degree 1 map from any constant $s_+ \gg 1$ circle in $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to $\mathbb{R}^2 \setminus 0$.

A change of variables relates the pair $\eta = (\varphi', \zeta')$ at points (x, \hat{u}) near $(y, 0)$ to the pair η . This formula takes the form $\eta = U(1 + \epsilon')\eta$, where U is the linear map from Lemma 5.5 and where $|\epsilon'| < c_0 e^{-s_+/c_0}$. With the preceding as background, use the arguments in Steps 1–4 of Part 4 in Section 5C to see that $\eta = (\varphi', \zeta')$, if not identically zero, defines a map from the boundary of any very small radius disk about the point $(y, 0)$ to $\mathbb{R}^2 \setminus \{0\}$ with negative degree. The relation $\eta = U(1 + \epsilon)\eta$ implies that η must define a nonpositive degree map to $\mathbb{R}^2 \setminus \{0\}$ from any sufficiently large s_+ circle in $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$, and in particular η cannot have degree 1.

This paradox is avoided only if $(\varphi^{(0)}, \zeta^{(0)}) = (\varphi^{(1)}, \zeta^{(1)})$. □

7 Cobordisms to the ech-HF submanifold moduli space

Section 2 describes sets of the form $C_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$ with C_{S_0} being a surface with boundary in $\mathbb{R} \times M_\delta$ and with each $p \in \Lambda$ version of C_{p_0} being a surface with

boundary in $\mathbb{R} \times \mathcal{H}_{p_*}^+$. The interiors of these surfaces are J -holomorphic. Each looks much like the portion of an ech-HF submanifold in the relevant part of $\mathbb{R} \times Y$. However, these surfaces with boundary do not necessarily fit together so as to define a closed surface in $\mathbb{R} \times Y$. As noted in Part 2 of Section 1F, the sets that are described in Section 2 form one boundary of a cobordism with two boundary components, the other being the moduli space of ech-HF submanifolds. This cobordism has an associated proper function mapping it to $[0, 1]$ with the inverse image over 1 being the boundary composed of ech-HF submanifolds. This section first describes and then constructs these cobordism spaces.

Section 2A describes a data set of the form (z_*, δ, χ_0, R) along with an almost complex structure J_{HF} for the construction of any given version of \mathcal{C}_0 . With (δ, χ_0, R) and J_{HF} specified, an almost complex structure for $\mathbb{R} \times Y$ is then chosen subject to the conditions given in Part 1 of Section 1C. This almost complex structure is again denoted by J .

The definitions in Section 2 and the constructions in the previous sections require the choice of an orbit in $\mathcal{A}_{\text{HF}}/\mathbb{R}$ of a given Lipschitz surface. As in Section 2, the latter determines an upper bound on z_* and δ . As explained in Section 1A, the choice of δ determines an upper bound for χ_0 and the choice for χ_0 determines one for R . The required upper bounds for z_* and δ may need some refinement in order to construct the cobordism space. The refined upper bounds are stated as needed for the various constructions that follow. In Section 2 and in what follows, the upper bounds in question for z_* and δ can be chosen so as to hold for all subvarieties chosen from a given finite or compact set in $\mathcal{A}_{\text{HF}}/\mathbb{R}$. Likewise, if $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$ is a given \mathbb{R} -invariant, weakly compact set, then the parameters z_* and δ can be chosen to be \mathcal{K} -compatible.

As in the previous sections, S is used to denote a chosen Lipschitz submanifold. Likewise, $(\hat{\Theta}_-, \hat{\Theta}_+) \in \hat{\mathcal{Z}}^S$ is chosen. Use (Θ_-, Θ_+) again to denote the corresponding pair in $\mathcal{Z}_{\text{ech}, M}$.

7A The cobordism space

A point in the cobordism space consists of a pair (τ, \mathcal{C}) where $\tau \in [0, 1]$ and where $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ is a set of $G + 1$ submanifolds with boundary in $\mathbb{R} \times Y$. The map to $[0, 1]$ sends (τ, \mathcal{C}) to τ . The three parts that follow describe \mathcal{C} . Part 1 describes C_S ; Part 2 describes the various $p \in \Lambda$ versions of C_p . Part 3 describes how τ enters the picture. The notation used in Section 2 is used here also.

Part 1 What is denoted by C_S is a properly embedded submanifold with boundary in the $f^{-1}([1 + z_*, 2 - z_*])$ part of $\mathbb{R} \times M_\delta$ whose interior is J -holomorphic. This submanifold is characterized in part by the four properties that are listed in what follows.

To set the stage for the statement of the first property, introduce the constant ρ_S , the disk bundle $N_0 \rightarrow S$, the map ϵ_S and the other notation from Part 5 of Section 2C. Let κ_0 denote the constant from Lemma II.6.5 and let κ denote the constant from Lemma II.6.6. Introduce κ_S to denote $10^6 \kappa \kappa_0$. Let $U \subset \mathbb{R} \times [1, 2] \times \Sigma$ denote the tubular neighborhood of S that is described in Lemma II.6.5.

Property 1 View C_S as a submanifold in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. As such, C_S lies in U and in the image via the exponential map ϵ_S of the radius $\kappa_S^{-1} \rho_S^2$ disk subbundle in N_0 . Moreover, C_S has intersection number 1 with the ϵ_S -image of each fiber of this disk bundle over the $t \in [1 + z_*, 2 - z_*]$ part of S .

The next property writes C_S as the image of a map from the $t \in [1 + z_*, 2 - z_*]$ part of S that has the form $\epsilon_S \circ \eta$ with η being a section of N_0 . This upcoming property also refers to the Fredholm operator D_S that is discussed in Parts 2–4 of Section II.6E and depicted in (1-25). The kernel of D_S is the vector space of sections in the domain Hilbert space that are annihilated by D_S .

Property 2 The section η is L^2 -orthogonal to the restriction of each element in the kernel of D_S to the part of S where $t \in [1 + z_*, 2 - z_*]$.

The third property speaks to the large $|s|$ behavior of C_S . The notation borrows from the fourth bullet of Proposition 2.1. By way of a reminder, let \hat{v}_- and \hat{v}_+ denote the respective HF-cycles that are used to define Θ_- and Θ_+ . Let q denote a given intersection point from either with Σ ; this a point in $C_- \cap C_+$. The corresponding integral curve of v from \hat{v}_- or \hat{v}_+ appears as $(1, 2) \times q$ when writing $f^{-1}(1, 2) \subset M$ as $(1, 2) \times \Sigma$. The point q labels a corresponding $s \ll -1$ or $s \gg 1$ end of S , this denoted by \mathcal{E}_{S_q} . This end of S is in q 's component of $\mathbb{R} \times [1, 2] \times (T_- \cap T_+)$. Note also that the functions (s, t) restrict as coordinate functions to \mathcal{E}_{S_q} . The normal bundle N_S over \mathcal{E}_{S_q} is identified with the product \mathbb{R}^2 -bundle in the manner that is described just prior to (2-6) and this identification is used to view a section as a map to \mathbb{R}^2 . Meanwhile, the exponential map ϵ_S over \mathcal{E}_{S_q} is written as in (2-6).

As in Proposition 2.1, q_* is used to denote the point in $T_- \cap T_+$ near q where the corresponding segment of an integral curve of v from Θ_- or Θ_+ intersects Σ . This

point q_* has distance $c_0\delta$ or less from q and so lies in q 's component of $T_- \cap T_+$. Writing ϵ_S over \mathcal{E}_{S_q} writes q_* as a section of the normal bundle N_S over \mathcal{E}_{S_q} .

Property 3 Let $q \in C_- \cap C_+$ denote an intersection point of an integral curve of \mathbf{v} from either \hat{v}_- or \hat{v}_+ . The section η over the $t \in [1 + z_*, 2 - z_*]$ part of \mathcal{E}_{S_q} converges pointwise as $s \rightarrow -\infty$ or $s \rightarrow \infty$ to q_* .

The final property views C_S as sitting in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. It talks about the behavior of C_S where t is near the endpoints of the interval $[1 + z_*, 2 - z_*]$. The statement of this property uses the notation from Part 1 of Section 1C. In particular, Part 1 of Section 1C uses the coordinates (φ_+, \hat{h}_+) for any given component of the region $T_+ \subset \Sigma$, and it uses the coordinate (φ_-, \hat{h}_-) for any given component of T_- .

Property 4 View C_S as a submanifold with boundary in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. As such, a neighborhood of its boundary has the following properties:

- The $[1 + z_*, 1 + z_S]$ portion of C_S has G components, with one mapping to each component of T_+ . A given component of this portion of C_S is the image of a map from $\mathbb{R} \times [z_*, z_S]$ to $\mathbb{R} \times [1 + z_*, 1 + z_S] \times \mathbb{R} / (2\pi\mathbb{Z}) \times (\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ that has the form

$$(x, z) \rightarrow (s = x, t = 1 + z, \varphi_+ = \varphi^S(x, z), \hat{h}_+ = \zeta^S(x, z)).$$

- The $[2 - z_S, 2 - z_*]$ portion of C_S has G components, with one mapping to each component of T_- . A given component of this portion of C_S is the image of a map from $\mathbb{R} \times [z_*, z_S]$ to $\mathbb{R} \times [2 - z_S, 2 - z_*] \times \mathbb{R} / (2\pi\mathbb{Z}) \times (\frac{4}{3\sqrt{3}}\delta_*^2, \frac{4}{3\sqrt{3}}\delta_*^2)$ of the form

$$(x, z) \rightarrow (s = x, t = 2 - z, \varphi_- = \varphi^S(x, z), \hat{h}_- = \zeta^S(x, z)).$$

In short, these four properties say that C_S looks much like the $\mathbb{R} \times M_\delta$ part of an ech-HF submanifold. Note in particular that these properties are satisfied if $C_S = C_{S_0}$ with the latter coming from a set of the sort that is described in Section 2.

Part 2 What is denoted by C_p is a properly embedded submanifold with boundary in $\mathbb{R} \times \mathcal{H}_{p_*}^+$ with J -holomorphic interior. There are two boundary components, one on the $u > 0$ component of the boundary of $\mathbb{R} \times \mathcal{H}_{p_*}^+$ and the other on the $u < 0$ component. The submanifold C_p is diffeomorphic to the complement of Δ_p interior points of the product of \mathbb{R} with a closed interval. What follows lists two additional properties.

Property 1 The large $|s|$ part of C_p is described by (2-9).

Property 2 The submanifold C_p is the Ψ_p -image of a graph in the $|\hat{u}| \leq R + \frac{1}{2} \ln z_*$ part of $\mathbb{R} \times \mathcal{H}_{p*}^+$ over a domain in $\mathbb{R} \times I_*$ having the form

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

The domain for the functions (φ^p, ζ^p) is $\mathbb{R} \times I_*$ when $\Delta_p = 0$, it is the complement of a single $\hat{u} = 0$ point when $\Delta_p = 1$, and it is the complement of two $\hat{u} = 0$ points when $\Delta_p = 2$.

In short, these properties say that C_p looks much like the $\mathbb{R} \times \mathcal{H}_{p*}^+$ part of an ech-HF submanifold.

Part 3 The parameter τ enters the story here. To set the stage, fix $p \in \Lambda$ and write the part of C_S in $\mathbb{R} \times \mathcal{H}_{p*}^+$ as in Property 4 in Part 1 using functions (φ^S, ζ^S) , and write C_p as in Property 2 of Part 2. Meanwhile, reintroduce from (2-5) the functions $(\varphi^{S_0}, \zeta^{S_0})$ that are defined by the surface C_{S_0} . The functions (φ^S, ζ^S) that define C_S and the pair (φ^p, ζ^p) that define C_p are constrained on the common $t = 1 + z_*$ and $t = 2 - z_*$ boundaries of their domains to obey

$$(7-1) \quad \begin{aligned} &\bullet \quad \zeta^S - \zeta^{S_0} = \tau(\zeta^p - \zeta^{S_0}), \\ &\bullet \quad \tau(\varphi^S - \varphi^{S_0}) = \varphi^p - \varphi^{S_0}. \end{aligned}$$

What follows are two remarks concerning these matching conditions. The first remark concerns the $\tau = 1$ version of (7-1): This version asserts that C_S and C_p fit seamlessly together across their common boundary in $\mathbb{R} \times \mathcal{H}_{p*}^+$. As a consequence, any given $\tau = 1$ version of $C = C_S \cup (\bigcup_{p \in \Lambda} C_p)$ is an ech-HF submanifold.

The second remark concerns the $\tau = 0$ case. A set $\mathcal{C}_0 = \{C_{S_0}, \{C_{p_0}\}_{p \in \Lambda}\}$ of the sort described in Section 2 obeys all of the $\tau = 0$ conditions. Moreover, it follows from Propositions 2.1 and 2.2 that these are the only sets that obey the $\tau = 0$ conditions.

A set $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ that is described by Parts 1 and 2 above and obeys a given $\tau \in [0, 1]$ version of (7-1) is said to be a (J, τ) -holomorphic submanifold.

7B The structure of the cobordism space

Introduce \mathcal{M}^* to denote the set of pairs of the form (τ, \mathcal{C}) with $\tau \in [0, 1]$ and with \mathcal{C} being a (J, τ) -holomorphic submanifold. This set is given the topology whereby open neighborhoods of a given element $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\})$ are generated by sets of the following sort: Fix $\varepsilon > 0$ and a compactly supported 2-form ν on $\mathbb{R} \times Y$.

The set in question contains a given $(\tau', \mathcal{C}' = \{C'_S, \{C'_p\}_{p \in \Lambda}\})$ if $|\tau - \tau'| < \varepsilon$ and if the conditions in (1-16) hold with the pair (C, \mathcal{C}') replaced by each pair from $\{(C_S, C'_S), \{(C_p, C'_p)\}_{p \in \Lambda}\}$.

The map from \mathcal{M}^* to $[0, 1]$ defined by the rule $(\tau, \mathcal{C}) \rightarrow \tau$ is denoted by π_I . A second map, this one from \mathcal{M}^* to a Euclidean space, also enters the story. The latter is denoted by p and its definition follows directly. To start, introduce $\Lambda_* \subset \Lambda$ to denote the subset of $\Delta_p \geq 1$ elements. The map p sends \mathcal{M}^* to $\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$. To give the rule that defines p , write a given element in \mathcal{M}^* as $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\})$. Each $p \in \Lambda$ version of C_p is defined by a pair of functions whose domain is the complement in $\mathbb{R} \times I_*$ of Δ_p with $\hat{u} = 0$. The \mathbb{R} coordinate of these missing $\hat{u} = 0$ points are the \mathbb{R} coordinates of $p(\tau, \mathcal{C})$ in p 's factor of $\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$ with it understood that when $\Delta_p = 2$, the first coordinate in the corresponding factor $\times_2 \mathbb{R}^2$ corresponds to the end of C_p where $\cos(\theta)$ limits to $\frac{1}{\sqrt{3}}$ as $s \rightarrow \infty$. The upcoming propositions set $n_* = \sum_{p \in \Lambda} \Delta_p$ and they set $\mathbb{R}^{n_*} = \times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$.

The following proposition describes the structure \mathcal{M}^* :

Proposition 7.1 *Fix a Lipschitz submanifold S such that D_S has trivial cokernel. There exists a purely S -dependent constant $\kappa \geq 1$ and, with $z_* < \kappa^{-1}$, there exists a constant κ_* that depends on z_* but is otherwise purely S -dependent with the following property: Use $\delta < \kappa_*^{-1} z_*$ with a pair $(\hat{\Theta}_-, \hat{\Theta}_+)$ from $\hat{\mathcal{Z}}^S$ to define \mathcal{M}^* . Fix $(\tau, \mathcal{C}) \in \mathcal{M}^*$.*

- There exist an integer $n \geq 0$, a neighborhood $U \subset \mathbb{R}^{n_*+n}$ of the origin, an open neighborhood $I \subset [0, 1]$ of τ , a smooth map $\mathfrak{f}: I \times U \rightarrow \mathbb{R}^n$ that sends $(\tau, 0)$ to the origin, and a topological embedding $\Phi: \mathfrak{f}^{-1}(0) \rightarrow \mathcal{M}^*$ onto an open set that sends the pair $(\tau, 0)$ to \mathcal{C} and is such that $\pi_I \circ \Phi$ gives the projection from $I \times U$ to I .
- The subspace of elements $\mathcal{M}_{\text{smooth}}^* \subset \mathcal{M}^*$ where \mathfrak{f} is a submersion is open and a smooth (n_*+1) -dimensional manifold with boundary. The maps π_I and p are smooth on this smooth subset.
- The integer n can be taken to be either 0 or 1 at points in $\mathcal{M}_{\text{smooth}}^*$, and it can be taken equal to zero at the points in $\mathcal{M}_{\text{smooth}}^*$ where $d\pi_I \neq 0$.
- An open neighborhood of the π_I -inverse image of 0 in $\mathcal{M}_{\text{smooth}}^*$ is mapped by $\pi_1 \times p$ diffeomorphically onto an open neighborhood of $\{0\} \times (\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R}))$ in $[0, 1] \times (\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R}))$.

Proposition 7.1 can be generalized in a straightforward manner to account for variations in the choice of the pair S . The formulation of this more general version is omitted.

The next proposition refers to the notion from Section 2A of a weakly compact subset of Lipshitz submanifolds. The proposition asserts that J , in particular, can be chosen so as to make Proposition 7.1's map f everywhere a submersion for a residual set of Lipshitz submanifolds from any given weakly compact subset of Lipshitz submanifolds.

To set the stage for the proposition, introduce the notion of a Lipshitz subvariety. The latter is a certain sort of 2-dimensional, J_{HF} -holomorphic subvariety in $\mathbb{R} \times [1, 2] \times \Sigma$. The definition is identical to that in Section 1G for a Lipshitz submanifold but for three items. First, the subvariety need not be a submanifold as it is allowed to have a finite number of interior singular points. Second, no irreducible component lies in a constant (s, t) slice of $\mathbb{R} \times [1, 2] \times \Sigma$. Third, Property 8 in Section 1G need not be obeyed. To say more about this last point, note that any given Lipshitz subvariety can be viewed as a pair, (S, u) , where S is a smooth complex curve with 2G boundary components and u a J_{HF} -holomorphic map from S into $\mathbb{R} \times [1, 2] \times \Sigma$ whose image is the subvariety in question. The pair (S, u) is described by the first six bullets in (II.6-2) and the modified version of the seventh bullet of (II.6-2) that requires u to embed the complement of a finite set of interior points. If u is an immersion, there is a holomorphic line bundle over S whose restriction to any given small radius disk is the normal bundle to its u -image. In this case an operator D_S that maps sections of the latter to sections of its tensor product with $T^{0,1}S$ which has the form depicted in (1-25). When u is not an embedding, there is an operator that plays the role of D_S and is denoted by D_S . This operator is obtained from what is denoted by $D\bar{d}$ in the proof of Proposition 3.4 in [10] by restricting the latter to elements of the form $(\xi, Y, 0)$. The latter is Fredholm when viewed as a bounded, linear map from the L^2_1 completion of its domain to the L^2 completion of its range. Because u is singular at only a finite number of points, all in the interior, this follows directly from what is said in Part 4 of Section II.6E.

Proposition 7.2 *Fix a countable set in $(\times_3(0, 1)) \times (1, \infty)$ of possible choices for the data (z_*, δ, χ_0, R) , then there is a C^∞ -residual set of allowed choices for J_{HF} for which the assertions that follow are true. Choose J_{HF} from this residual set.*

- *Let S denote a Lipshitz subvariety. Then D_S has trivial cokernel.*
- *Let \mathcal{K} denote a given \mathbb{R} -invariant, weakly compact subset of Lipshitz submanifolds, then there exists a constant $\kappa \geq 1$ that depends only on \mathcal{K} and, given*

$z_* < \kappa^{-1}$, there exists $\kappa_* > 1$ that depends on z_* and \mathcal{K} with the following significance: Choose a \mathcal{K} -compatible data set (z_*, δ, χ_0, R) from the given set with $z_* < \kappa^{-1}$ and $\delta < \kappa_*^{-1} z_*$. Use this data to define the geometry of Y .

- (a) There is a certain residual set of almost complex structures pursuant to the constraints given in Section 1C and there exists a residual subset in \mathcal{K} such that if J is chosen from the former and S from the latter then all $(\widehat{\Theta}_-, \widehat{\Theta}_+) \in \widehat{\mathcal{Z}}^S$ version of \mathcal{M}^* are such that
- (1) The corresponding $\mathcal{M}_{\text{smooth}}^*$ is the whole of \mathcal{M}^* and so \mathcal{M}^* is a smooth $(n_* + 1)$ -dimensional manifold with boundary, and $\pi_I \times p: \mathcal{M}^* \rightarrow [0, 1] \times (\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R}))$ is a smooth map.
 - (2) The critical values of π_1 are in $(0, 1)$, and only finitely many of them are critical values of π_I 's restriction to any given compact set in \mathcal{M}^* .
- (b) If \mathcal{K} is an open set, then the various versions of $\mathcal{M}^*|_{\tau=1}$ as defined by the elements in \mathcal{K} and a given choice for J define a smooth manifold such that the tautological map to \mathcal{K} is smooth.

Propositions 7.1 and 7.2 are proved in Section 7E.

By way of a parenthetical remark, the tools that are developed in this section can be used to strengthen both the fourth bullet of Proposition 7.1 and item (a)(2) of the second bullet of Proposition 7.2. With regards to the fourth bullet of Proposition 7.1, the open neighborhood of $\pi^{-1}(0)$ can be taken to be the π_1 -inverse image of an open neighborhood of 0 in $[0, 1]$. The strengthened version of item (a)(2) of the second bullet of Proposition 7.2 asserts that the set of critical values of π_1 on the whole of \mathcal{M}^* is finite. The proofs of these strengthened versions do not involve any new technology. Even so, a full presentation is lengthy and so these stronger versions are not proved here.

The final proposition implies that π_I is in all cases a proper map. This proposition refers to the strong C^∞ topology on spaces of sections and maps. This topology is defined as follows: An open neighborhood of a given section or map is indexed by a positive integer and a positive number. Let q denote the given section or map and let (k, ε) denote a given positive integer and positive number. Elements in the corresponding neighborhood of q have C^k distance less than ε from q over the whole of q 's domain.

The upcoming proposition also refers to the normal bundle and various associated notions for a submanifold that is described by Part 2 of Section 7A, an example being

some $p \in \Lambda$ version of C_p . Let C denote the relevant submanifold. The fiber metric on C 's normal bundle is defined by the ambient metric on $\mathbb{R} \times \mathcal{H}_{p_*}^+$ that comes from the chosen almost complex structure J and the compatible 2-form $ds \wedge \hat{a} + w$. The latter metric defines a metric on C and the covariant derivative on sections of the normal bundle and tensor bundles over C . The normal bundle can also be endowed with an exponential map that embeds a constant radius disk subbundle into $\mathbb{R} \times \mathcal{H}_{p_*}^+$. This exponential map gives the canonical identification between the zero section and C , its differential on the zero section is the identity map, and it maps the disk bundle over the boundary of C to the boundary of $\mathbb{R} \times \mathcal{H}_{p_*}^+$.

Proposition 7.3 *Fix a Lipschitz submanifold S . There exists a purely S -dependent constant $\kappa \geq 1$, and with $z_* < \kappa^{-1}$, there exists a constant κ_* that depends on z_* but is otherwise purely S -dependent with the following property: Fix $z_* < \kappa^{-1}$ and $\delta < \kappa_*^{-1} z_*$, and then a pair $(\hat{\Theta}_-, \hat{\Theta}_+) \in \hat{\mathcal{Z}}^S$ to define \mathcal{M}_* . The map $p: \mathcal{M}^* \rightarrow \times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$ is proper in the following strong sense: Let*

$$\{(\tau_n, C_n = \{C_{Sn}, \{C_{pn}\}_{p \in \Lambda})\}_{n=1,2,\dots} \subset \mathcal{M}^*$$

be any given sequence with fixed p -image. There is an element $(\tau, C = \{C_S, \{C_p\}_{p \in \Lambda})$ in \mathcal{M}^ with the given p -image and a subsequence (hence renumbered consecutively from 1) with the properties listed next:*

- *For each $n \in \{1, 2, \dots\}$, let η_n denote the section of the disk bundle N_S that defines C_{Sn} . The resulting sequence of sections $\{\eta_n\}_{n=1,2,\dots}$ converges over the $t \in [1 + z_*2 - z_*]$ part of S in the strong C^∞ topology to the section that defines C_S .*
- *Fix $p \in \Lambda$. There exists a sequence of sections $\{\eta_{pn}\}_{n=1,2,\dots}$ of the disk subbundle of C_p 's normal bundle that converges to zero in the strong C^∞ -topology on $C^\infty(C_p; N)$ and is such that each index n version of C_{pn} is the image of the composition of the exponential map with the corresponding section η_{pn} .*

This proposition is proved in Section 7D.

7C Boundary conditions for the Cauchy–Riemann equations

This subsection constitutes a digression to introduce the analytic tools that are needed to handle the boundary-matching conditions given by (7-1). By way of background, the pair (φ^S, ζ^S) that appears in the index 1 critical point version of (7-1) obeys

the standard Cauchy–Riemann equations in coordinates (x, z) for $\mathbb{R} \times [z_*, z_S]$ that are defined by the rule $s = x, t = 1 + z$. Meanwhile, the pair (φ^p, ζ^p) obey these equations on the domain $\mathbb{R} \times [\delta^2, z_*]$ if z is identified with $e^{-2(R-\hat{u})}$. There are analogous Cauchy–Riemann equations near an index 2 critical point version of (7-1). The Cauchy–Riemann equations here are obeyed by $(\varphi^S, -\zeta^S)$ using the coordinates $(x, -z)$ for $\mathbb{R} \times [z_*, z_S]$ that are defined by writing (s, t) as $s = x$ and $t = 2 - z$. Meanwhile, the pair $(\varphi^p, -\zeta^p)$ obey the same Cauchy–Riemann equations in terms of coordinates (x, z) on the domain $\mathbb{R} \times [\delta^2, z_*]$ when z is defined by $z = e^{-2(R+\hat{u})}$. These coordinate identifications are used implicitly in the rest of this section and in the subsequent sections. The pair $(\varphi^{S_0}, \zeta^{S_0})$ or $(\varphi^{S_0}, -\zeta^{S_0})$, as the case may be, obeys the Cauchy–Riemann equations on the domain $\mathbb{R} \times [e^{-8} z_*, z_S]$ and thus on both sides of the $z = z_*$ locus where (7-1) holds. The fact that all of these pairs obey a linear equation near the $z = z_*$ locus explains the focus in this subsection on the coupled, linear boundary value problem that is described next.

The boundary value problem is that for pairs (φ_+, ζ_+) and (φ_-, ζ_-) which obey the Cauchy–Riemann equations on the respective domains $\mathbb{R} \times [z_*, z_S]$ and $\mathbb{R} \times [e^{-8} z_*, z_*]$. This is to say that

$$(7-2) \quad \partial_x \varphi_{\pm} - \partial_z \zeta_{\pm} = 0 \quad \text{and} \quad \partial_x \zeta_{\pm} + \partial_z \varphi_{\pm} = 0$$

on the relevant domain, and so $\eta_+ = \varphi_+ + i\zeta_+$ and $\eta_- = \varphi_- + i\zeta_-$ are holomorphic functions of the complex coordinate $x + iz$. Their boundary values are constrained on the common boundary of their respective domains by a given $\tau \in [0, 1]$ version of

$$(7-3) \quad \zeta_+ = \tau \zeta_- \quad \text{and} \quad \tau \varphi_+ = \varphi_- \quad \text{where } z = z_*.$$

They are also constrained so that

$$(7-4) \quad \lim_{|x| \rightarrow \infty} (|\zeta_{\pm}| + |\varphi_{\pm}|) = 0.$$

The five parts that follow in this subsection discuss various aspects of this coupled, linear boundary value problem.

Part 1 This part of the subsection describes energy bounds that hold for the pairs just described. These are summarized by the next lemma. The lemma uses the notation $\|\cdot\|$ to denote the L^2 -norm of a given function with it understood that the integration domain is the domain where the function is defined.

Lemma 7.4 Suppose that (φ_+, ζ_+) and (φ_-, ζ_-) are pairs of compactly supported, smooth functions that are defined on the domains $\mathbb{R} \times [z_*, z_S]$ and $\mathbb{R} \times (e^{-8} z_*, z_*)$, respectively, and that obey (7-3) on the common boundary of their domains of definition. Let η_{\pm} denote the \mathbb{C} -valued functions $\varphi_{\pm} + i\zeta_{\pm}$. Then $\|\bar{\partial}\eta_+\|^2 + \|\bar{\partial}\eta_-\|^2 = \frac{1}{4}\|d\eta_+\|^2 + \frac{1}{4}\|d\eta_-\|^2$.

Proof Integration by parts finds that

$$(7-5) \quad \|\bar{\partial}\eta_+\|^2 + \|\bar{\partial}\eta_-\|^2 = \frac{1}{4}\|d\eta_+\|^2 + \frac{1}{4}\|d\eta_-\|^2 + \int_{\mathbb{R} \times z_*} (\varphi_- \partial_x \zeta_- - \varphi_+ \partial_x \zeta_+).$$

Use (7-3) to see that the boundary integral is zero. □

Part 2 This part describes a version of the maximum principle for pairs $(\varphi_{\pm}, \zeta_{\pm})$ that obey (7-2)–(7-4). Such is the content of the next lemma.

Lemma 7.5 Suppose that (φ_+, ζ_+) and (φ_-, ζ_-) are pairs of smooth functions that are defined on the respective domains $\mathbb{R} \times [z_*, z_S]$ and $\mathbb{R} \times [e^{-8} z_*, z_*]$ and obey (7-2)–(7-4). Define functions φ and ζ on the domain $\mathbb{R} \times [e^{-8} z_*, z_S]$ by the rule

- $\varphi = \varphi_-$ where $z \in [e^{-8} z_*, z_S]$ and $\varphi = \tau\varphi_+$ where $z \in [z_*, z_S]$,
- $\zeta = \tau\zeta_-$ where $z \in [e^{-8} z_*, z_S]$ and $\zeta = \zeta_+$ where $z \in [z_*, z_S]$.

If either function is not identically zero, then neither function can have a local maximum or minimum on $\mathbb{R} \times (e^{-8} z_*, z_S)$.

Proof Consider, for example, ζ . The function ζ_- is harmonic, as is ζ_+ . Thus, neither can have local extremal values in the interior of their domain of definition. If $\tau = 0$, then ζ is zero where $z \leq z_*$. Suppose for the sake of argument that 0 is a local minimum of ζ . As $\zeta_+ = 0$ at $z = z_*$, its z -derivative must be nonnegative and positive at some $z = z_*$ points if ζ_+ is not identically zero. If this is so then (7-2) implies that φ_+ at $z = z_*$ is a nondecreasing function of x with positive derivative at some points, so (7-4) cannot be true.

Suppose next that $\tau > 0$ and suppose that ζ takes a local maximum or minimum at a given point (x_*, z_*) . The simplest case to consider is that where $\partial_z \zeta_+ \neq 0$ at this point. Suppose for the sake of argument that $\partial_z \zeta_+ < 0$. Then the function $z \rightarrow \tau\zeta_-(x_*, z)$ is a decreasing function of z for z near to but slightly less than z_* . Thus, $\tau\zeta_-$ and hence ζ will not have a local maximum at (x_*, z_*) .

To see about the general case, suppose for the sake of argument that (x_*, z_*) is a local maximum for ζ . What follows generates some nonsense from this assumption. To set up the notation, introduce q to denote the value of ζ at this point. No generality is lost by assuming that $q \neq 0$. Introduce r to denote the Euclidean distance on $\mathbb{R} \times [e^{-8}z_*, z_S]$ from (x_*, z_*) . Fix $\varepsilon \in (0, q)$ but very small, chosen in particular so that $q - \varepsilon$ is a regular value of both ζ_+ and $\tau\zeta_-$. Let U denote the component of the set where $\zeta \geq q - \varepsilon$ that contains the point (x_*, z_*) . This set is compact. Introduce ∂U_+ and ∂U_- to denote the respective $z \geq z_*$ and $z \leq z_*$ components of the boundary of U . These are smooth arcs. It follows from (7-2) that these arcs are oriented by the respective 1-forms $d\varphi_+$ and $d\varphi_-$. In particular, the respective integrals

$$(7-6) \quad \int_{\partial U_+} d\varphi_+ \quad \text{and} \quad \tau \int_{\partial U_-} d\varphi_-$$

are positive. But this last conclusion is nonsense because they sum to zero as a consequence of (7-2) with the fundamental theorem of calculus (use Stokes' theorem along an arc). □

Part 3 This part explains how a bound $|\zeta_+|$ for $z > \frac{1}{2}z_S$ and on $|\zeta_-|$ for $z < \frac{1}{2}z_*$ can be used to obtain τ -independent bounds at $z = z_*$. The following lemma makes a quantitative statement:

Lemma 7.6 *There exists a z_* - and z_S -independent constant $\kappa \geq 1$ with the following significance: Suppose that $(\varphi_{\pm}, \zeta_{\pm})$ are pairs of functions as in Lemma 7.5 that obey (7-2)–(7-4). Fix constants $r_* > 0$ and $r_S > 0$ such that $|\zeta_+| \leq r_S$ for $z > \frac{1}{2}z_S$ and $|\zeta_-| \leq r_*$ for $z < \frac{1}{2}z_*$. Then $|\zeta_-| \leq \kappa(r_* + z_*r_S/z_S)$ where $z = z_*$.*

Proof Let $\Delta \subset \mathbb{R} \times [e^{-8}z_*, \frac{1}{2}z_*]$ denote a disk of radius $\frac{1}{4}z_*$, and introduce Δ' to denote the concentric disk with radius $\frac{1}{8}z_*$. Use of the standard Green's function for the Laplacian with a cutoff function that is 1 on Δ' and zero on the complement of Δ will find that

$$(7-7) \quad |\nabla\zeta_-| \leq c_0 r_* z_*^{-1}$$

at the origin of Δ . It follows as a consequence of (7-2) that $|\nabla\varphi_-| \leq c_0 r_* z_*^{-1}$ where $z \leq \frac{1}{4}z_*$. Much the same argument finds that $|\nabla\varphi_+| \leq c_0 r_S z_S^{-1}$ where $z \geq \frac{1}{4}z_S$. Hold on to these bounds for the moment.

The two pairs of functions $(\varphi'_{\pm}, \zeta'_{\pm}) = (\partial_x\varphi_{\pm}, \partial_x\zeta_{\pm})$ obey (7-2) and they also obey (7-3). As explained in a moment, Lemma 7.5 can be invoked using the function φ' which is defined to be $\partial_x\varphi_-$ where $z \leq z_*$ and $\tau\partial_x\varphi_+$ where $z \geq z_*$. Granted

Lemma 7.5, it follows from the conclusions of the preceding paragraph that $\partial_x \varphi$ at $z = z_*$ can have absolute value no greater than $c_0(r_* z_*^{-1} + r_S z_S^{-1})$. This with the Cauchy–Riemann equations imply that $|\partial_z \zeta_-| \leq c_0(r_* z_*^{-1} + r_S z_S^{-1})$ at $z = z_*$.

Let w denote the function on $[\frac{1}{4} z_*, z_*]$ given by the rule

$$(7-8) \quad z \rightarrow w(z) = r + c_0(r_* z_*^{-1} + r_S z_S^{-1})z.$$

View w as an x -independent function on $\mathbb{R} \times [\frac{1}{4} z_*, z_*]$. As such, it is harmonic, and the maximum principle implies that $w(z) \geq \zeta_-(z)$ on its domain of definition. Since this is the case, it follows that $\zeta_- \leq w$ at all $z \in [\frac{1}{4} z_*, z_*]$ and so $c_0 r$ at $z = z_*$. This gives the asserted upper bound for ζ_- where $z = z_*$. The exact same argument with $-\zeta_-$ replacing ζ_- gives the asserted upper bound at $z = z_*$ for $-\zeta_-$.

Return now to the assertion that Lemma 7.5 can be invoked using φ' . There is no issue if it is known a priori that $|\partial_x \zeta_{\pm}|$ and $|\partial_x \varphi_{\pm}|$ limit uniformly to zero as $|x| \rightarrow \infty$. If this has not been established, the argument proceeds as follows: Let $\beta: \mathbb{R} \rightarrow [0, 1]$ denote a smooth function with compact support with integral equal to 1. Given $\varepsilon > 0$, introduce $\beta_\varepsilon: \mathbb{R} \rightarrow [0, 1]$ to denote the function given by the rule $x \rightarrow \varepsilon^{-1} \beta(\varepsilon^{-1} x)$. Consider now the pair $(\varphi_{\varepsilon_{\pm}}, \zeta_{\varepsilon_{\pm}})$ given by the mollifying formula

$$(7-9) \quad (\varphi_{\varepsilon_{\pm}}, \zeta_{\varepsilon_{\pm}})|_{(x,z)} = \int_{\mathbb{R}} \beta_\varepsilon(x') (\varphi_{\pm}, \zeta_{\pm})|_{(x'+x,z)} dx'.$$

Every $\varepsilon > 0$ version of $(\varphi_{\varepsilon_{\pm}}, \zeta_{\varepsilon_{\pm}})$ obeys (7-2)–(7-4). This is also the case for their partial derivatives to any given order with respect to x . This understood, define for $\varepsilon > 0$ the mollified function φ'_ε given by $\partial_x \varphi_{\varepsilon_-}$ where $z \leq z_*$ and $\tau \partial_x \varphi_{\varepsilon_+}$ where $z \geq z_*$. Lemma 7.5 holds for this function. Meanwhile, the family $\{\varphi'_\varepsilon\}_{\varepsilon > 0}$ converges uniformly as $\varepsilon \rightarrow 0$ on compact subsets of the domain $\mathbb{R} \times [e^{-8} z_*, z_S]$ to the function $\partial_x \varphi$. This understood, take the $\varepsilon \rightarrow 0$ limit of $|\partial_x \varphi_\varepsilon|$ to see that the assertion of Lemma 7.5 holds for $\partial_x \varphi$. □

Part 4 This part says something about a priori estimates near the $z = z_*$ locus for solutions to (7-2)–(7-4).

Lemma 7.7 Fix $k \in \{1, 2, \dots\}$, then there exists a z_* - and z_S -independent constant $\kappa \geq 1$ with the following significance: Suppose that $(\varphi_{\pm}, \zeta_{\pm})$ are smooth functions that obey (7-2)–(7-4). Fix $r_* > 0$ and $r_S > 0$ such that $|\zeta_+| \leq r_S$ for $z > \frac{1}{2} z_S$ and $|\zeta_-| \leq r_*$ for $z < \frac{1}{2} z_*$. Then the norms of the derivatives of (φ_-, ζ_-) to order k where $z > \frac{1}{4} z_*$ and those of (φ_+, ζ_+) to order k where $z < \frac{1}{4} z_S$ are bounded by $\kappa(r_* z_*^{-k} + r_S z_S^{-k})$.

Proof The argument given in the proof of Lemma 7.6 bounds $|\partial_x \varphi_-|$ where $z \in [\frac{1}{4}z_*, z_*]$ by $c_0(r_* z_*^{-1} + r_S z_S^{-1})$, and an analogous argument bounds $|\partial_x \zeta_+|$ where $z \in [z_*, \frac{1}{4}z_S]$. Much the same argument bounds $|\partial_x^k \varphi_-|$ by $c_k(r_* z_*^{-k} + r_S z_S^{-k})$ for $z \in [\frac{1}{4}z_*, z_*]$ and $|\partial_x^k \zeta_+|$ by this same constant where $z \in [z_*, \frac{1}{4}z_S]$. An iterative bootstrapping argument uses the bound for $|\partial_x^k \varphi_-|$ to obtain the desired bound on $|\partial_x^{k-1} \zeta_-|$, and it uses the bound for $|\partial_x^k \zeta_+|$ to obtain the desired bound on $|\partial_x^{k-1} \varphi_+|$. What follows describes how this works for $|\partial_x^{k-1} \zeta_-|$. The argument for $|\partial_x^{k-1} \varphi_+|$ is identical but for the notation.

Note first that $|\partial_x^{k-1} \zeta_-|$ is a priori bounded by $c_{k-1}(r_* z_*^{-k+1} + r_S z_S^{-k+1})$ where $z = \frac{1}{4}z_*$. This the case, note next that $\partial_z(\partial_x^{k-1} \zeta_-) = \partial_x^k \varphi_-$ because of (7-2). The assumed bound on $\partial_x^k \varphi_-$ implies that $|\partial_z \partial_x^{k-1} \zeta_-| \leq c_k(r_* z_*^{-k} + r_S z_S^{-k})$ where $z = z_*$. Granted that such is the case, use the fact that $\partial_x^{k-1} \zeta_-$ is harmonic to invoke the maximum principle for the harmonic function

$$(7-10) \quad (x, z) \rightarrow (\partial_x^{k-1} \zeta_-)|_{(x,z)} - c_{k-1}(r_* z_*^{-k+1} + r_S z_S^{-k+1}) - c_k(r_* z_*^{-k} + r_S z_S^{-k})z$$

on $\mathbb{R} \times [\frac{1}{4}z_*, z_*]$ to obtain the desired upper bound on $\partial_x^{k-1} \zeta_-$. The desired lower bound is obtained by this argument by replacing ζ_- with $-\zeta_-$.

Bounds on the partial derivatives to order k in both the variables x and z are obtained via the Cauchy–Riemann equations from those for just the partial derivatives with respect to x . □

7D Proof of Proposition 7.3

The argument for the proposition when all $p \in \Lambda$ versions of Δ_p are zero is very much like that given in Section 4C for the proof of Lemma 4.3. This version of the argument is given in the first six parts of this subsection. Part 7 adds what is needed to prove the proposition when some $p \in \Lambda$ versions of Δ_p are 1 or 2.

Part 1 This part derives an upper bound for the integral of the 2–form w over the J –holomorphic submanifolds with boundary from any given element in \mathcal{M}^* . The lemma below states such a bound:

Lemma 7.8 *There exists a purely S –dependent $\kappa \geq 1$ with the following significance: Define the space \mathcal{M}^* using $z_* \leq \kappa^{-1}$, $\delta < \kappa^{-2} z_*$ and a $\{\Delta_p = 0\}_{p \in \Lambda}$ pair from $\widehat{\mathcal{Z}}^S$. Suppose that $\tau \in [0, 1]$ and $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ is a (J, τ) –holomorphic submanifold.*

Let $I \subset \mathbb{R}$ denote an interval of length 1. Then

$$\int_{C_S} w + \sum_{p \in \Lambda} \int_{C_p} w \leq \kappa \quad \text{and} \quad \int_{C_S \cap (I \times M_\delta)} ds \wedge d\hat{a} + \sum_{p \in \Lambda} \int_{C_p \cap (I \times \mathcal{H}_{p_*}^+ \cap M_\delta)} ds \wedge d\hat{a} \leq \kappa.$$

Meanwhile, $\sum_{p \in \Lambda} \int_{C_p \cap (I \times \mathcal{H}_{p_*}^+)} ds \wedge h\hat{a} \leq \kappa_*$, where κ_* is a \mathcal{C} - and τ -independent constant.

Proof The argument is much like that used for Lemma 4.5. To start, note that what are essentially cosmetic modifications to the arguments used at the beginning of the Lemma 4.5’s proof can be used to prove that the integrals in question are finite, and that Stokes’ theorem in various guises can be used to compute them.

With the preceding in mind, use (1-6) to write the 2-form w as $w = w' + \sum_{p \in \Lambda} db_p$ where any given $p \in \Lambda$ version of the 1-form b_p has compact support in \mathcal{H}_p^+ and has the form $b_p = \chi(1 - 3 \cos^2 \theta) du - N\sqrt{6} f \cos \theta \sin^2 \theta d\phi$ with N being the function of u given by the rule $u \rightarrow N(u) = \chi(|u| - R + \ln \delta_*)$. The form w' is zero on each $p \in \Lambda$ version of C_p , and its support on C_S is disjoint from the boundary of C_S . Given that w' is closed, Stokes’ theorem with what is said in Corollary II.2.6 can be used to see that its integral over C_S differs by no more than $c_0\delta$ from its integral over S .

Fix $p \in \Lambda$ and use Stokes’ theorem to write the integral of db_p as the sum of the integrals given in (4-9) as defined using $(\varphi, \zeta) = (\varphi^p, \zeta^p)$. As explained in the proof of Lemma 4.5, the sum of the left-most two terms in the top bullet of (4-9) and the two terms in the lower bullet of (4-9) are bounded by a purely S -dependent (or \mathcal{K} -compatible) constant. The as yet unspoken for terms in the top bullet of (4-9) are

$$(7-11) \quad \int_{\mathbb{R} \times \{\hat{u} = R + \frac{1}{2} \ln z_*\}} \zeta^p d\varphi^p - \int_{\mathbb{R} \times \{\hat{u} = -R - \frac{1}{2} \ln z_*\}} \zeta^p d\varphi^p.$$

Meanwhile, integration by parts identifies the integral of b_p over C_S with the sum

$$(7-12) \quad \int_{\gamma'_{p+}} Nh d\phi - \int_{\gamma'_{p-}} Nh d\phi - \int_{\mathbb{R} \times \{\hat{u} = R + \frac{1}{2} \ln z_*\}} \zeta^S d\varphi^S + \int_{\mathbb{R} \times \{\hat{u} = -R - \frac{1}{2} \ln z_*\}} \zeta^S d\varphi^S,$$

where γ'_{p+} and γ'_{p-} are the parts of the integral curves of v that extend γ_{p+} and γ_{p-} , and lie in the $f \geq 1 + z_*$ and $f < 2 - z_*$ parts of the radius $4\delta_*$ coordinate balls centered on the index 1 and 2 critical points from p . The left-most two integrals in (7-12) are zero as ϕ is constant in the integral curve segments in question.

To say more about the integrals in (7-11) and the two right-most integrals in (7-12) note first that the corresponding $\hat{u} = R + \frac{1}{2} \ln z_*$ integrals in (7-11) and in (7-12) come with opposite signs. This is also the case for the corresponding $\hat{u} = -R - \frac{1}{2} \ln z_*$ integrals. With this understood, one can use (7-1) to identify the sum of the various $\mathfrak{p} \in \Lambda$ versions of (7-11) and (7-12) with the sum

$$(7-13) \quad -(1 - \tau) \int_{\mathbb{R} \times \{\hat{u} = R + \frac{1}{2} \ln z_*\}} (\zeta^{S_0} d\varphi^S - \zeta^{\mathfrak{p}} d\zeta^{S_0}) \\ + (1 - \tau) \int_{\mathbb{R} \times \{\hat{u} = -R - \frac{1}{2} \ln z_*\}} (\zeta^{S_0} d\varphi^S - \zeta^{\mathfrak{p}} d\varphi^{S_0}).$$

The integrals in (7-13) of the 1-form $\zeta^{\mathfrak{p}} d\varphi^{S_0}$ are bounded by $c\delta_*^2$ with c here denoting a purely S -dependent (or \mathcal{K} -compatible) constant. This is because $|\zeta^{\mathfrak{p}}| \leq \frac{4}{\sqrt{3}}\delta_*^2$ and because of what is said by the fourth bullet in Proposition 2.1, which implies that the integral of $|d\varphi^{S_0}|$ is finite on both $\mathbb{R} \times \{\hat{u} = R + \frac{1}{2} \ln z_*\}$ and $\mathbb{R} \times \{\hat{u} = -R - \frac{1}{2} \ln z_*\}$.

Bounds are obtained below for the norms of the integrals that involve $\zeta^{S_0} d\varphi^S$ by writing them using Stokes' theorem as the sum of two integrals whose integrands are supported in the $z \geq z_*$ part of the radius $4\delta_*$ coordinate balls about \mathfrak{p} 's critical points. The term with support in the index 1 critical point coordinate ball is

$$(7-14) \quad -(1 - \tau) \int_{z \geq z_*} (\zeta^{S_0} dN \wedge d\varphi^S + N d\zeta^{S_0} \wedge d\varphi^S),$$

and the term with support in the index 2 critical point coordinate ball has the same form but no minus sign in front. (Remember that N is given in terms of the variable u by the rule $N(u) = \chi(|u| - R + \ln \delta_*)$, so it is zero where $|u| \geq R - \ln \delta_* + 1$.)

To deal with integral in (7-14), fix for the moment a number $m > 1$ and use the triangle inequality to bound the integrand by the sum of $m(|\zeta^{S_0}|^2 + |d\zeta^{S_0}|^2)(|dN|^2 + |N|^2)$ and $\frac{1}{m}|d\varphi^S|^2$. Let O denote any given open subset of the domain of the integral in (7-14). The absolute value of the integral over O of the integrand in (7-14) is no greater than the sum of the integrals over O of these two functions $m(|\zeta^{S_0}|^2 + |d\zeta^{S_0}|^2)(|dN|^2 + |N|^2)$ and $\frac{1}{m}|d\varphi^S|^2$. The integral of the former over O has an upper bound that is independent of O . This is because of what is said by the fourth bullet of Proposition 2.1 concerning the exponential rate of convergence. As a consequence, the integral over O and thus over the domain in (7-14) of the function $m(|\zeta^{S_0}|^2 + |d\zeta^{S_0}|^2)(|dN|^2 + |N|^2)$ necessarily has an upper bound of the form $m c_*$ with c_* being a purely S -dependent (or \mathcal{K} -compatible) number.

Two facts are needed to deal with the integral of $\frac{1}{m}|d\varphi^S|^2$: first, the 2-form w on C_S , when written in the radius $4\delta_*$ coordinate ball in terms of the pair (φ^S, ζ^S) , is given by $\sqrt{6}(\partial_x\varphi^S\partial_v\zeta^S - \partial_v\varphi^S\partial_x\zeta^S)dx \wedge dv$. Second, the pair (φ^S, ζ^S) obeys the Cauchy–Riemann equations as functions of (x, v) . These two facts imply that w can be written on this part of C_S as $\sqrt{6}(|\partial_x\varphi^S|^2 + |\partial_{\bar{u}}\varphi^S|^2)dx \wedge dv$. It follows as a consequence that the integral of $\frac{1}{m}|d\varphi^S|^2$ over the domain O is no greater than the integral of $\frac{1}{m}w$ over O .

Taking $m = 1000$, it now follows from what was said in the preceding two paragraphs that the integral in (7-14) is no greater than $\frac{1}{1000} \int_{C_S} w + c$ where c is a purely S -dependent (or \mathcal{K} -compatible) constant. This bound with those derived previously imply that

$$(7-15) \quad \int_{C_S} w + \sum_{p \in \Lambda} \int_{C_p} w \leq \frac{1}{1000} \int_{C_S} w + c,$$

with c being purely S -dependent (or \mathcal{K} -compatible). The inequality in (7-15) implies what is asserted by the first bullet of Lemma 7.8.

Turn now to the integrals of $ds \wedge \hat{a}$. With the $f \in (1, 2)$ part of M written as $(1, 2) \times \Sigma$, the form \hat{a} is dt with t the Euclidean coordinate on the $(1, 2)$ factor. This being the case, the integral of $ds \wedge \hat{a}$ over $C_S \cap (I \times M_\delta)$ is no greater than G and the integral of $ds \wedge \hat{a}$ over the $I \times M_\delta$ part of any given $p \in \Lambda$ version of C_p is no greater than z_* . Minor cosmetic changes to the arguments from Step 4 of the proof of Proposition II.5.1 in Section II.5B give the bound on the integral of $ds \wedge \hat{a}$ over the whole of $C_p \cap (I \times \mathcal{H}_{p_*}^+)$. Note that the integration by parts that is used in this Step 4 does not lead to boundary terms because \hat{a} annihilates the tangent space to the boundary of $\mathcal{H}_{p_*}^+$. □

Part 2 This part supplies an upper bound for the distance in $\mathbb{R} \times [1, 2] \times \Sigma$ from any given \mathcal{M}^* version of C_S and S . Here again, the $f \in (1, 2)$ part of M_δ is identified with $(1, 2) \times \Sigma$. This part of the section also supplies a positive bound for the function $1 - 3 \cos^2 \theta$ on any given \mathcal{M}^* and $p \in \Lambda$ version of C_p . These bounds are summarized in the respective lemmas that follow.

Lemma 7.9 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$ and, given $z_* < \kappa^{-1}$, there exists $\kappa_* > 1$ that depends only on z_* but is otherwise purely S -dependent (or \mathcal{K} -compatible) with the following significance: Define the space \mathcal{M}^* using $z_* < \kappa^{-1}$, $\delta < \kappa^{-2} z_*$ and a $\{\Delta_p = 0\}_{p \in \Lambda}$ pair from $\hat{\mathcal{Z}}^S$. Suppose that*

$(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}) \in \mathcal{M}^*$. Then C_S is the image via the exponential map ϵ_S of a section of the radius $\frac{1}{100}\kappa_S^{-1}\rho_S^2$ subbundle of N_0 .

The next lemma concerns the function $1 - 3 \cos^2 \theta$ on the various $p \in \Lambda$ versions of C_p .

Lemma 7.10 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$ and, given $z_* < \kappa^{-1}$, there exists $\kappa_* > 1$ that depends only on z_* but is otherwise purely S -dependent (or \mathcal{K} -compatible) with the following significance: Define \mathcal{M}^* using $z_* < \kappa^{-1}$ and $\delta < \kappa_*^{-1}z_*$ and a $\{\Delta_p = 0\}_{p \in \Lambda}$ pair from $\widehat{\mathcal{Z}}^S$. There exists $\kappa_{**} > 1$ such that if $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}) \in \mathcal{M}^*$ is a given element, then $1 - 3 \cos^2 \theta$ is greater than κ_{**}^{-1} on all $p \in \Lambda$ versions of C_p .*

The proofs of Lemmas 7.9 and 7.10 require a preliminary lemma to supply an a priori bound on the norm of any $p \in \Lambda$ version of $|\zeta^S|$ on the common boundary of C_S and C_p . The proof of Lemma 7.10 requires in addition a bound for any $p \in \Lambda$ version of $|\zeta^p|$ where $z = e^{64\delta^2}$. This is the content of the upcoming Lemma 7.11. Lemma 7.11 refers to the coordinate z that is defined on either component of the $|\widehat{u}| \geq R + \ln \delta$ portion of $\mathcal{H}_{p_*}^+$ by the rule $z = e^{-2(R-|\widehat{u}|)}$.

Lemma 7.11 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$ and, given $z_* < \kappa^{-1}$ and $\varepsilon > 0$, there exists $\kappa_\varepsilon > 1$ that depends on z_* and ε but is otherwise purely S -dependent (or \mathcal{K} -compatible) such that the following is true: Define the space \mathcal{M}^* using $z_* < \kappa^{-1}$, $\delta < \kappa_\varepsilon^{-1}z_*$ and a $\{\Delta_p = 0\}_{p \in \Lambda}$ pair from $\widehat{\mathcal{Z}}^S$. Let $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}) \in \mathcal{M}^*$, and let $p \in \Lambda$.*

- *The pair (φ^S, ζ^S) is such that $|\zeta^S| \leq \varepsilon\delta_*^2$ where $z = z_*$, this the common boundary of C_S and C_p .*
- *The pair (φ^p, ζ^p) is such that $|\zeta^p| \leq \varepsilon\delta_*^2$ where $z = e^{64\delta^2}$.*

This lemma is proved in Part 4. Given the lemma, what follows in this Part 2 is the proof of Lemma 7.9. Part 3 contains the proof of Lemma 7.10.

Proof of Lemma 7.9 The proof has five steps.

Step 1 Write t on $[1, 1 + z_S]$ as $1 + z$ with $z \in [0, z_S]$. Likewise write t as $t = 2 - z$ on $[2 - z_S, 2]$ with z again in $[0, z_S]$. Let $c_S > 1$ denote the version of the constant κ that is supplied by Lemma 7.11. Given $\varepsilon > 0$ and $z_* < \min(c_S^{-1}z_S, \varepsilon\delta_*^2)$, choose δ

so that the conclusions of Lemma 7.11 can be invoked. Having made such choices, fix $p \in \Lambda$. The function ζ^{S_0} that appears in p 's version of (7-1) has norm bounded by $c z_*$, where $c \geq 1$ is purely S -dependent (or \mathcal{K} -compatible); this follows from Proposition 2.1. It follows from what is said in Corollary II.2.6 that the $|x| \rightarrow \infty$ limits of $|\zeta^S|$ are bounded by $c\delta$. As a consequence, these limits are bounded by $c\varepsilon\delta_*^2$. Here again, c is purely S -dependent (or \mathcal{K} -compatible).

Step 2 The pair (φ^S, ζ^S) is a solution to the Cauchy–Riemann equations on the domain $\mathbb{R} \times [z_*, z_S]$. As a consequence, ζ^S is a harmonic functions on this domain. This is to say that it is annihilated by the Laplacian $\partial_x^2 + \partial_z^2$. With the preceding understood, fix $c > 1$ and introduce the function

$$(7-16) \quad z \rightarrow \frac{z - z_*}{z_S - z_*} (\rho_S^2 - c\varepsilon\delta_*^2) + c\varepsilon\delta_*^2.$$

This is a harmonic function on $\mathbb{R} \times [z_*, z_S]$, and, if $c \geq c$ with c purely S -dependent (or \mathcal{K} -compatible), then this function is greater than $|\zeta^S|$ on the boundaries and at large $|x|$. Granted that such is the case, the maximum principle demands that this function be greater than $|\zeta^S|$ on the whole of $\mathbb{R} \times [z_*, z_S]$. Choose c in (7-16) so as to be purely S -dependent (or \mathcal{K} -compatible) and so that this last conclusion holds.

Step 3 Suppose now that there is no κ as claimed by Lemma 7.9 so as to derive some nonsense. Granted this assumption, there is a sequence $\{(D_n, \eta_n)\}_{n=1,2,\dots}$ of the following sort: First, D_n is a data set with elements $((\hat{\Theta}_{n-}, \hat{\Theta}_{n+}), z_{*n}, \delta_n, \chi_{0n}, R_n)$ that is suitable for defining the geometry of Y and \mathcal{M}^* , and is such that the conclusions of the top bullet in Lemma 7.11 can be invoked with $\varepsilon = \frac{1}{n}$ and $z_{*n} \in (0, \frac{1}{n}z_S)$ and $\delta_n \leq \frac{1}{n^2}z_{*n}$. Second, η_n is a section of the bundle N_S over the $t \in [1 + z_{*n}, 2 - z_{*n}]$ part of S with the properties given in the list below:

- (7-17) • η_n is L^2 -orthogonal to the $t \in [1 + z_{*n}, 2 - z_{*n}]$ restriction of the elements in the kernel of the operator D_S .
- The norm of $|\eta_n|$ at some point in its domain is greater than $\frac{1}{100}\kappa_S^{-1}\rho_S^2$.
 - The $|s| \rightarrow \infty$ limit of $|\eta_n|$ is bounded by $\frac{1}{n}\delta_*^2$.
 - The composition $\epsilon_S \circ \eta_n$ has J -holomorphic image.
 - The absolute value of the η_n analog of the function ζ^S is bounded where $z \in [z_{*n}, z_S]$ by the $z = z_{*n}$ and $\varepsilon = \frac{1}{n}$ version of the function in (7-16).

As noted in Part 2 of Section II.6E, the third bullet in (7-17) implies that $\eta = \eta_n$ obeys an equation with the schematic form

$$(7-18) \quad \bar{\partial}\eta + \tau_1(\eta) \cdot \partial\eta + \tau_0(\eta) = 0,$$

where the notation is that in (II.6-10). By way of a reminder, $\bar{\partial}$ 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, the map $\tau_1: N_0 \rightarrow N_S \otimes \text{Hom}(T^{1,0}S; T^{0,1}S)$ and the map $\tau_0: N_0 \rightarrow N_S \otimes T^{0,1}S$ are smooth, fiber-preserving maps that vanish along the zero section.

Step 4 The properties listed in Step 3 together with (7-18) imply via standard elliptic regularity arguments that there is a subsequence of $\{\eta_n\}_{n=1,2,\dots}$, hence renumbered consecutively from 1, that is described either by Case 1, Case 2 or Case 3 given below.

Case 1 The subsequence converges uniformly on compact domains in S to a nontrivial section of N_0 over the whole of S with the following properties:

- (a) It obeys (7-18) and the conditions in the second and third bullets of (II.6-12).
- (b) Its norm is no greater than $\kappa_S^{-1} \rho_S^2$
- (c) It is L^2 -orthogonal to $\text{kernel}(D_S)$

The first two conditions above imply that the section is described by Lemma II.6-6. This understood, the third condition implies that the section is identically zero. Given that the limit section is asserted to be nontrivial, Case 1 cannot describe $\{\eta_n\}_{n=1,2,\dots}$.

Case 2 below uses terminology from Section 1G and Sections II.6C and II.6E.

Case 2 There is a negative point $q \in \underline{S} \setminus S$ and a sequence $\{q_n\}_{n=1,2,\dots} \subset S_0$ of points that converges in \underline{S} to q such that $|\eta_n|(q_n) > \frac{1}{100} \kappa_S^{-1} \rho_S^2$. Write a neighborhood of q in S as in Part 4 of Section II.6E. This done, view η_n on this part of S as a \mathbb{C} -valued function as done in Section II.6E. There is an unbounded, decreasing sequence $\{x_n\} \in (-\infty, x_0]$ such that the translated sequence with n^{th} member η'_n given by $\eta'_n|_x = \eta_n|_{x-x_n}$ converges on compact domains of $\mathbb{R} \times [1, 2]$ to a nontrivial function with the following properties:

- (a) It is holomorphic.
- (b) Its imaginary part vanishes on $\mathbb{R} \times \{1\}$.
- (c) Its real part vanishes on $\mathbb{R} \times \{2\}$.
- (d) It is bounded.

Properties (a)–(c) plus (II.6-18) are incompatible with property (d). This being the case, the subsequence that is given at the end of Step 4 is not described by Case 2.

Case 3 This is the analog of Case 2 where the point q is a positive point of $\underline{S} \setminus S$. The analogous conclusion applies: the subsequence that is given at the end of Step 4 is not described by Case 3 either.

Step 5 Neither Case 1 nor Case 2 nor Case 3 describe the subsequence given by Step 4. This contradicts what is said at the end of Step 4. The contradiction is avoided if and only if Lemma 7.9 is true. □

Part 3 This part contains the:

Proof of Lemma 7.10 The proof has five steps.

Step 1 Let $c_S > 1$ denote the larger of the versions of constant κ that are supplied by Lemmas 7.11 and 7.9. Given $\varepsilon > 0$ and $z_* < \max(c_S^{-1} z_S, \varepsilon^2 \delta_*^2)$, choose $\delta < (c_S + \kappa_{\varepsilon^2 \delta_*^2})^{-1} z_*$ so as to invoke the conclusions of Lemmas 7.11 and 7.9 using $\varepsilon^2 \delta_*^2$ in lieu of ε . Having made such a choice, suppose that there exists $\varepsilon_* \in (0, 1)$ with the following property:

(7-19) *Let $(\tau, C = \{C_S, \{C_p\}_{p \in \Lambda}\})$ denote an element from the resulting version of \mathcal{M}^* . Fix $p \in \Lambda$. Then $|\zeta^p| \leq (1 - \varepsilon_*) \frac{4}{3\sqrt{3}} \delta_*^2$ where $z = z_*$.*

Granted (7-19), it then follows from (4-14) that $1 - 3 \cos^2 \theta > c_0^{-1} \varepsilon_*$ on the boundary of any $(\{C_S, \{C_p\}_{p \in \Lambda}\}, \tau) \in \mathcal{M}^*$ and $p \in \Lambda$ version of C_p . With a bound of this sort in hand, a repeat of the arguments for Lemma 4.7 using Lemma 7.8 in lieu of Lemma 4.5 proves Lemma 7.10. The steps that follow give an existence proof for a suitable ε_* .

What with the top line in (7-1) and the top bullet of Lemma 7.11, the assertion made by (7-19) holds automatically for $\tau > c_0^{-1} \varepsilon^2$. This the case, only small values of τ are of any concern. Even so, no upper bound for τ is assumed in the remaining steps.

Step 2 Fix $r \in (0, 1)$ and let $D \subset \mathbb{R} \times ((1+r)z_*, z_S)$ denote a disk of radius rz_* . The function $|\zeta^S|$ is bounded by what is written in (7-16), and because ζ^S is harmonic on D , this implies that the norm of $d\zeta^S$ at the center of D is no greater than $c_0(r^- z_*)^{-1} \varepsilon^2 \delta_*^4$. This with the Cauchy–Riemann equations implies that $|d\varphi^S|$ is bounded where $z = (1+r)z_*$ by $c_0(r^- z_*)^{-1} \varepsilon^2 \delta_*^4$. Meanwhile, $|\zeta^p| < \frac{4}{3\sqrt{3}} \delta_*^2$. As ζ^p is harmonic on the domain $\mathbb{R} \times [\delta^2, z_*]$, it follows that $|d\zeta^p|$ and thus $|d\varphi^p|$ are bounded by $c_0 z_*^{-1} \delta_*^2$ where $z = \frac{1}{2} z_*$.

Step 3 Define a function φ' on $\mathbb{R} \times [e^{-8} z_*, z_S]$ as follows: Set $\varphi' = \partial_x(\varphi^p - \varphi^{S_0})$ on the $z \leq z_*$ part of this domain, and set $\varphi' = \tau \partial_x(\varphi^S - \varphi^{S_0})$ on the $z \geq z_*$ part.

Define next a function w of the coordinate z on $\mathbb{R} \times [\frac{1}{2}z_*, (1+r)z_*]$ by the rule

$$(7-20) \quad z \rightarrow w(z) = z_*^{-1} (r^{-1} \varepsilon^2 \delta_*^4 + z_*^{-1} \delta_*^2 (z - (1+r)z_*)).$$

The function w is harmonic and there exists a constant $c_w \leq c_0$ such that $\varphi' - c_w w$ is less than zero where $z = \frac{1}{2}z_*$ and $z = (1+r)z_*$. As explained in the next step, the maximum principle can be used to infer that $\varphi' < c_w w$ at all points in $\mathbb{R} \times [\frac{1}{2}z_*, (1+r)z_*]$. Granted this, it then follows using Lemma 7.11 that

$$(7-21) \quad |\partial_x \varphi^p| \leq c_0 z_*^{-1} (r^{-1} \varepsilon^2 \delta_*^4 + r) \quad \text{where } z = z_*.$$

Thus, $|\partial_z \zeta^p| \leq c_0 z_*^{-1} (r^{-1} \varepsilon^2 \delta_*^4 + r)$ where $z = z_*$. Taking $r = \varepsilon \delta_*^2$ finds $|\partial_x \zeta^p| \leq c_0 z_*^{-1} \varepsilon \delta^2$.

Step 4 What follows considers the case when both $\lim_{|x| \rightarrow \infty} |\partial_x(\varphi^p - \varphi^{S_0})|$ and $\lim_{|x| \rightarrow \infty} |\partial_x(\varphi^S - \varphi^{S_0})|$ are zero. The general case is handled using mollifiers as done in the proof of Lemma 7.6. Granted this assumption about the $|x| \rightarrow \infty$ limits, it follows that the function $\varphi' - c_w w$ is negative at large $|x|$. This function is harmonic where $z \neq z_*$ and so it lacks local maxima and local minima in $\mathbb{R} \times (\frac{1}{2}z_*, z_*)$ and in $\mathbb{R} \times (z_*, (1+r)z_*)$.

To see about local maxima or minima where $z = z_*$, note that the function on \mathbb{R} given by

$$(7-22) \quad x \rightarrow q(x) = z_*^{-2} \delta_*^2 x$$

is a conjugate harmonic function for w when viewed as a z -independent function on the domain $\mathbb{R} \times [\frac{1}{2}z_*, (1+r)z_*]$. This is to say that the pair (q, w) obey the Cauchy-Riemann equations. This understood, define the function ζ' on $\mathbb{R} \times [\frac{1}{2}z_*, (1+r)z_*]$ by setting it to equal $\tau(\zeta^p - c_w \tau^{-1} q)$ where $z \leq z_*$, and to equal $\zeta^S - c_w q$ where $z \geq z_*$. With the pair (φ', ζ') in hand, repeat the argument in the paragraph preceding (7-6) with φ' playing the role of ζ and ζ' the role of $-\varphi'$ to rule out local extreme points for $\varphi' - c_w w$.

An analogous argument with the signs of w and q reversed rules out local extreme points for $\varphi' + c_w w$.

Step 5 The lower bullet of Lemma 7.11 asserts that $|\zeta^p| \leq \varepsilon \delta_*^2$ where $z = e^{64} \delta^2$. Meanwhile, Step 4 finds that $|\partial_z \zeta^p| \leq c_0 z_*^{-1} \varepsilon \delta_*^2$ where $z = z_*$. With this in mind, let w now denote the x -independent function on $\mathbb{R} \times [2\delta^2, z_*]$ given by

$$(7-23) \quad z \rightarrow w(z) = z_* + z_*^{-1} z.$$

There exists $c_0 \geq 1$ such that the function $\zeta^p - c_0 \varepsilon \delta_*^2 w$ on $\mathbb{R} \times [2\delta^2, z_*]$ has the following properties: It is negative where $|x| \gg 1$ and where $z = e^{64}\delta^2$. Meanwhile, its z -derivative is negative where $z = z_*$. Granted these facts, a version of the maximum principle implies that $\zeta^p < c_0 \varepsilon \delta_*^2$ where $z = z_*$. The analogous argument using $-\zeta^p$ in lieu of ζ^p proves that $\zeta^p > -c_0 \varepsilon \delta_*^2$ where $z = z_*$. A suitable choice of ε establishes what is asserted by (7-19). □

Part 4 This part contains the:

Proof of Lemma 7.11 The proof has four steps. The first three steps prove the assertion made by the lower bullet in the lemma.

Step 1 To prove the assertion made by the lower bullet, assume the contrary to generate some nonsense. If lower bullet is false, there is a sequence $\{(D_n, (C_n, \tau_n))\}_{n=1,2,\dots}$ of the following sort: First, D_n is a set with elements $((\hat{\Theta}_{n-}, \hat{\Theta}_{n+}), z_{*n}, \delta_n, \chi_{0n}, R_n, J_n)$ that are suitable for defining the geometry of Y and the corresponding version of \mathcal{M}^* . This data is such that the pair $(\hat{\Theta}_{n-}, \hat{\Theta}_{n+})$ is a $\{\Delta_p = 0\}_{p \in \Lambda}$ element in $\hat{\mathcal{Z}}^S$. The constants z_{*n} and δ_n satisfy $z_{*n} < \frac{1}{n} z_S$ and $\delta_n < \frac{1}{n} z_{*n}$. Meanwhile, $(C_n = \{C_{Sn}, \{C_{pn}\}), \tau_n)$ is an element in the corresponding version of \mathcal{M}^* . In addition, there exists $p \in \Lambda$ such that the C_{pn} version of the function ζ^p has absolute value greater than $\varepsilon \delta_*^2$ at some point where $z = e^{64}\delta_n^2$.

For each $n \in \{1, 2, \dots\}$, fix a point in C_{pn} where $z = e^{64}\delta_n^2$ and where $|\zeta^p| \geq \varepsilon \delta_*^2$. Let q_n denote a given such point. The point q_n projects to the $\delta = \delta_n$ version of M_δ and the image lies in the radius $2\delta_*$ coordinate ball about one or the other of the critical points from p . It follows from (4-14) that all sufficiently large n versions of q_n project in M_δ so as to have coordinate radius greater than $\frac{1}{25}\varepsilon^{1/2}\delta_*$. The image of q_n also lies where the function f differs from either 1 or 2 by less than $\frac{1}{n} z_{*n}$.

For each n , let U_n denote the following part of the union of the respective radius $2\delta_*$ coordinate balls centered on the index 1 and index 2 critical points from p : It is the part where the radius is greater than $\frac{1}{100}\varepsilon^{1/2}\delta_*$ and where the function either $|f - 1|$ or $|f - 2|$ is less than $\frac{1}{2} z_{*n}$. If n is large, $C_{pn} \cap (\mathbb{R} \times U_n)$ is a nonempty, J -holomorphic submanifold; it is nonempty as it contains the point q_n .

Pass to a subsequence so that after renumbering consecutively from 1, the projection to M_δ of the resulting set of points $\{q_n\}_{n=1,2,\dots}$ lies very near one critical point of p . For each n , let $C_n \subset (C_{pn} \cap (\mathbb{R} \times U_n))$ denote the component that contains q_n . Translate C_n by a constant amount along the \mathbb{R} factor in $\mathbb{R} \times U_n$ so that the point q_n is

moved to where the \mathbb{R} coordinate is 0. Let $\{C'_n\}_{n=1,2,\dots}$ denote the resulting, translated sequence of submanifolds.

Step 2 This step reviews some relevant geometry. To start, let $p \in \mathfrak{p}$ denote the critical point that is described at the end of the last step. The critical point p labels an irreducible component of the locus $C_{1+} \cup C_{2-}$ in Σ . Use C to denote this irreducible component. As noted in Section 1C, the circle C has an annular neighborhood, \mathbb{T} , with the following property: The identification between $f^{-1}(1, 2) \subset M$ and $(1, 2) \times \Sigma$ identifies the \mathbb{C} -valued 1-form $d\phi + i d\hbar$ on the part of $f^{-1}(1, 2) \cap M_\delta$ in the radius $2\delta_*$ coordinate ball centered on p with a 1-form on an annular neighborhood of C in \mathbb{T} . The latter 1-form is holomorphic of type $(1, 0)$.

Let \widehat{U} denote the part of the radius $\frac{3}{2}\delta_*$ coordinate ball centered on p where the radius is greater than $\frac{1}{100}\varepsilon^{1/2}\delta_*$ and where either $|f - 1| < e^{-100}\delta_*^2$ or $|f - 2| < e^{100}\delta_*^2$ as the case may be. The map from \widehat{U} to Σ that is given by the flow along the integral curves of v identifies \widehat{U} with $(-e^{-100}\delta_*^2, e^{-100}\delta_*^2) \times U$, where $U \subset \mathbb{T}$ can be written as $U_1 \setminus U_2$ where U_1 is an annular neighborhood of C and where $U_2 \subset U_1$ is a smaller width annular neighborhood of C .

Let z denote the Euclidean coordinate on $(-e^{-100}\delta_*^2, e^{-100}\delta_*^2)$. The identification just described extends so as to identify $\mathbb{R} \times \widehat{U}$ with $\mathbb{R} \times (-e^{-100}\delta_*^2, e^{-100}\delta_*^2) \times U$. Let x denote the Euclidean coordinate on the \mathbb{R} factor. The identification just described identifies $T^{1,0}(\mathbb{R} \times \widehat{U})$ with the span of the pair $(dx + i dz, d\phi + i d\hbar)$. This integrable complex structure is observedly compatible with the symplectic form $dx \wedge dz + d\phi \wedge d\hbar$.

Step 3 Each large n version of U_n lies in \widehat{U} ; it appears in the coordinates just defined as the subset where $z \in (-\frac{1}{2}z_{*n}, \frac{1}{2}z_{*n})$. Given such large n , set $\alpha_n = 2z_{*n}^{-1}e^{-100}\delta_*^2$ and define a diffeomorphism $\psi: \mathbb{R} \times U_n \rightarrow \mathbb{R} \times \widehat{U}$ by the rule

$$(7-24) \quad \psi_n(x, z, \phi, \hbar) = (\alpha_n x, \alpha_n z, \phi, \hbar).$$

This diffeomorphism is J -holomorphic.

It follows as a consequence that $\psi_n(C'_n)$ is a properly embedded, J -holomorphic submanifold in $\mathbb{R} \times \widehat{U}$. This submanifold sits entirely where $1 - 3 \cos^2 \theta > 0$, this being the locus where $z = 0$. Even so the point $\psi_n(q_n)$ has distance bounded by $c_0 \frac{1}{n}$ from this locus. Moreover, $\psi_n(q_n)$ has distance at least c_0^{-1} from the boundary of the closure of $\mathbb{R} \times \widehat{U}$ inside $\mathbb{R} \times Y$. Note as well that $\psi_n(q_n)$ sits on the $x = 0$ locus. Given that the 2-form w on \widehat{U} appears as $-2\sqrt{6} d\hbar \wedge d\phi$, it follows from Lemma 7.8 that $\psi_n^* w = w$. As a consequence, if $I \subset \mathbb{R}$ is any unit-length interval, there is an I -

and n -independent bound on the integral of w over $\psi_n(C'_n) \cap (I \times \widehat{U})$. Meanwhile, the integral of $dx \wedge dz$ over $\psi_n(C'_n) \cap (I \times \widehat{U})$ is no greater than $2e^{-100}\delta_*^2$. Indeed, this follows from the fact that the corresponding C_{pn} is given as a graph over the $|\widehat{u}| \leq R + \frac{1}{2} \ln z_*$ part of $\mathbb{R} \times \mathcal{H}_{p+}$. To elaborate, $dx \wedge dz$ on $I \times \widehat{U}$ is the pullback via the projection map of its namesake on $I \times (-e^{-100}\delta_*^2, e^{-100}\delta_*^2)$. Meanwhile, the projection to this product restricts to $\psi_n(C'_n) \cap (I \times \widehat{U})$ to define a 1-to-1 map into $I \times (-e^{-100}\delta_*^2, e^{-100}\delta_*^2)$ because C_{pn} is a graph. This implies that the integral of the 2-form $dx \wedge dt$ over $\psi_n(C'_n) \cap (I \times \widehat{U})$ cannot be greater than its integral over the whole of $I \times (-e^{-100}\delta_*^2, e^{-100}\delta_*^2)$, which is $2e^{-100}\delta_*^2$.

Granted these observations, invoke Proposition II.5.5 using the sequence $\{\psi_n(C'_n)\}_{n \gg 1}$ to obtain a subsequence that converges on compact subsets of $\mathbb{R} \times \widehat{U}$ in the manner dictated by Proposition II.5.5 to a weighted J -holomorphic subvariety in $\mathbb{R} \times \widehat{U}$. Let ϑ denote this subvariety. Given what is said about $\psi_n(q_n)$, the set ϑ must contain a pair whose subvariety component is the $x = 0, z = 0$ locus in $\mathbb{R} \times \widehat{U}$.

The latter conclusion constitutes the required nonsense because the existence of such a pair in ϑ has the same implications as its existence in the analogous version of ϑ given by Step 1: there is a circle in each large n version of C_{pn} whose image via the projection to \mathcal{H}_{p+}^+ defines a nonzero generator of the latter's first homology. This nonsense proves what is asserted by the lower bullet of Lemma 7.11.

Step 4 With ε chosen, fix z_* and δ to invoke the lower bullet of Lemma 7.12 as a guarantee that $|\zeta^p| < \frac{1}{4}\varepsilon\delta_*^2$ where $z = 2\delta^2$. Require in addition that $z_* < \frac{1}{4}\varepsilon_S\delta_*^2$ and that δ is chosen less than $c_0^{-1}e^{-100}\varepsilon\delta_*^2$. Define the x -independent function w on $\mathbb{R} \times [2\delta^2, z_S]$ by the rule

$$(7-25) \quad z \rightarrow \frac{1}{2}\varepsilon\delta_*^2 + z z_S^{-1} \rho_S^2.$$

Meanwhile, let ζ denote the function on $\mathbb{R} \times [2\delta^2, z_S]$ that is given by $\tau\zeta^p - w$ where $z \leq z_*$ and given by $\zeta^S + (1 - \tau)\zeta^{S_0} - w$ where $z > z_*$. It follows from (7-1) that this function is continuous across the $z = z_*$ locus. Meanwhile, the lower bullet of Lemma 7.11 and Property 1 in Part 1 of Section 7A with Corollary II.2.6 imply that $\zeta \leq 0$ where $z = e^{64}\delta^2$, where $z = z_S$ and where $|x| \gg 1$. As ζ is a harmonic function where $z \neq z_*$, it has no local maxima where $z \neq z_*$. As explained in a moment, it has no local maxima where $z = z_*$. Granted that such is the case, it follows that $\zeta^S \leq \varepsilon\delta_*^2$ where $z = z_*$.

To see why ζ has no local maxima where $z = z_*$, define the function φ on the domain $\mathbb{R} \times [2\delta^2, z_S]$ as follows: Set φ equal to $\varphi^p - x z_S^{-1} \rho_S^2$ where $z \leq z_*$ and set φ to equal

$\tau\varphi^S + (1 - \tau)\varphi^{S_0} - xz_S^{-1}\rho_S^2$ where $z \geq z_*$. Note that the pair $(xz_S^{-1}\rho_S^2, w)$ obey the Cauchy–Riemann equations. This understood, repeat the argument in the paragraph preceding (7-6) using the just-defined version of (φ, ζ) to rule out local maxima for ζ on the $z = z_*$ locus.

To prove that $\zeta^S \geq -\varepsilon\delta_*^2$ where $z = z_*$, repeat the preceding argument with the sign of w reversed and with $xz_S\rho_S^2$ added to φ^p and to $\tau\varphi^S + (1 - \tau)\varphi^{S_0}$ rather than subtracted when defining φ . □

Part 5 This part proves that \mathcal{M}^* is compact in a topology that is slightly weaker than the one defined at the outset of Section 7B. This topology is defined as follows: The open neighborhoods of a given element $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\})$ are generated by sets that are labeled by data of the form (ε, ν, V, I) where ε is a positive number, ν is a smooth, compactly supported 2–form on $\mathbb{R} \times Y$ and $V \subset \mathbb{R} \times Y$ is an open set with compact closure. Meanwhile, $I \subset [0, 1]$ is an open neighborhood of τ . The corresponding open set in \mathcal{M}^* consists of pairs of the form $(\tau', \mathcal{C}' = \{C'_S, \{C'_p\}_{p \in \Lambda}\})$ with $\tau' \in I$ and with \mathcal{C}' obeying

$$(7-26) \quad \bullet \quad \sup_{z \in C \cap V} \text{dist}(z, C' \cap V) + \sup_{z \in C' \cap V} \text{dist}(z, C \cap V) < \varepsilon,$$

$$\bullet \quad \left| \int_C \nu - \int_{C'} \nu \right| < \varepsilon.$$

Part 6 to come proves convergence in Section 7B’s topology on \mathcal{M}^* and convergence in the strong C^∞ topology as asserted by the two bullets of Proposition 7.3.

The arguments that follow assume that the data that defines the geometry of Y and then \mathcal{M}^* is such that the conclusion of Lemma 7.11 holds with a given, small choice for ε . The arguments require ε to be less than an S –independent, positive number. The subsequent arguments also assume that the defining data is such that the conclusions of Lemmas 7.9 and 7.10 hold.

To set up the arguments, suppose that $\{(\tau_n, \mathcal{C}_n = \{C_{Sn}, \{C_{pn}\}_{p \in \Lambda}\})\}_{n=1,2,\dots} \in \mathcal{M}^*$ is a given sequence. No generality is lost by taking a sequence for which the corresponding sequence $\{\tau_n\}_{n=1,2,\dots}$ converges. Use $\tau \in [0, 1]$ to denote the limit. The six steps that follow prove that this sequence has a subsequence that converges to some (J, τ) –holomorphic submanifold.

Step 1 Fix $p \in \Lambda$ and fix attention on one or the other of the $|\hat{u}| \geq R + \frac{1}{2} \ln z_* - 8$ parts of $\mathcal{H}_{p_*}^+$. Use the coordinates (x, z, ϕ, \hat{h}) here. For each $n \in \{1, 2, \dots\}$, write the relevant part of C_{Sn} as a graph using functions (φ_n^S, ζ_n^S) . Meanwhile, write the relevant part of C_{pn} as a graph using functions (φ_n^p, ζ_n^p) .

Invoke Lemma 7.7 to conclude the following: The sequence $\{(\tau_n, C_n)\}_{n=1,2,\dots}$ has a subsequence (hence renumbered consecutively from 1) with the following property: The sequence of pairs $\{(\varphi_n^S, \varsigma_n^S)\}_{n=1,2,\dots}$ and the sequence of pairs $\{(\varphi_n^p, \varsigma_n^p)\}_{n=1,2,\dots}$ converge in the C^∞ topology on compact subsets of $\mathbb{R} \times [\frac{1}{4}z_*, \frac{1}{4}z_S]$ to respective pairs denoted by (φ^S, ς^S) and (φ^p, ς^p) . The latter obey (7-1).

In addition, the surface defined where $z \in [z_*, \frac{1}{4}z_S]$ in $\mathbb{R} \times \mathcal{H}_p$ by the graph of the pair (φ^S, ς^S) when viewed as a surface in $\mathbb{R} \times (1, 2) \times \Sigma$ lies in the radius $\kappa_S^{-1} \rho_S^2$ tubular neighborhood of S . Meanwhile, the $|\hat{u}| \in [R + \frac{1}{2} \ln z_* - \ln 2, R + \frac{1}{2} z_*]$ part of the graph of the pair (φ^p, ς^p) defines via Ψ_p a surface in portion of $\mathbb{R} \times \mathcal{H}_{p_*}^+$.

Step 2 Fix $p \in \Lambda$. For each $n \in \{1, 2, \dots\}$, use $(\varphi_n^p, \varsigma_n^p)$ to denote the pair that defined C_{pn} as a graph over $\mathbb{R} \times I_*$. It follows from Lemma 7.10 that the corresponding positive integer sequence of coefficient functions that appear in the various $(\varphi, \varsigma) \in \{(\varphi_n^p, \varsigma_n^p)\}_{n=1,2,\dots}$ versions of (3-4) are obtained from the restriction of their eponymous brethren on \mathcal{X} to a compact set. The partial derivatives of the latter set of functions to any given order have uniformly bounded absolute values on such a compact set.

Granted these last observations, and granted the conclusions from Step 1, standard elliptic regularity arguments in Chapter 6 of [12] can be applied to find a subsequence of $\{(\varphi_n^p, \varsigma_n^p)\}_{n=1,2,\dots}$, hence renumbered consecutively from 1, that converges in the C^∞ topology on compact subsets of $\mathbb{R} \times I_*$ to a pair, (φ^p, ς^p) , that obeys (3-4). Given Lemma 7.10, it follows that the graph of this pair is in the domain of the map Ψ_p and so defines a J -holomorphic surface $\mathbb{R} \times \mathcal{H}_{p_*}^+$. Let C_p denote the latter surface.

Step 3 Each $n \in \{1, 2, \dots\}$ version of C_{S_n} is defined by a section, η_n , of the bundle N_S over the $t \in [1 + z_*, 2 - z_*]$ part of S . As noted by Lemma 7.9, each C_{S_n} is the image via ϵ_S of a section, η_n , of the radius $\frac{1}{100} \kappa_S^{-1} \rho_S^2$ subbundle in N_0 .

The maps τ_0 and τ_1 that appear in (7-18) enjoy uniform bounds on their derivatives to any given order on the radius $\frac{1}{50} \kappa_S^{-1} \rho_S^2$ subbundle in N_0 . As each $\eta \in \{\eta_n\}_{n=1,2,\dots}$ obeys (7-18), it follows using what is said in Step 1 with the aforementioned elliptic regularity arguments from Chapter 6 in [12] that the sequence $\{\eta_n\}_{n=1,2,\dots}$ has a subsequence (now renumbered consecutively from 1) that converges in $C^\infty(S; N_0)$ on compact subsets of the $t \in [1 + \frac{1}{2} z_*, 2 - \frac{1}{2} z_*]$ part of S to a smooth section of N_0 with norm no greater than $\frac{1}{100} \kappa_S^{-1} \rho_S^2$. Let η denote this section.

It follows from Lemma 7.9 that the composition $\epsilon_S \circ \eta$ defines a J -holomorphic surface in $\mathbb{R} \times [1 + \frac{1}{2} \ln z_*, 2 - \frac{1}{2} \ln z_*]$ that lies in the radius $\kappa_S^{-1} \rho_S$ tubular neighborhood

of S . Let C_S denote this surface. Meanwhile, Step 1 has the following additional implication: the set $\mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}$ obeys the parameter τ version of the matching conditions given by (7-1).

Step 4 It follows from Steps 1–3 that (τ, \mathcal{C}) satisfies all of the requirements for membership in \mathcal{M}^* except perhaps for the conditions on the $|s| \rightarrow \infty$ limits of the various surfaces that form \mathcal{C} . This step with Steps 5 and 6 prove that this last requirement is met. The assertion that any given sequence in \mathcal{M}^* has a subsequence that limits in the manner described by Steps 1–3 to an element in \mathcal{M}^* verifies the claim that \mathcal{M}^* is compact in the topology that is defined by (7-26).

To start the story on the large $|s|$ behavior, let η again denote the section of N_0 over the $t \in [1 + z_*, 2 - z_*]$ part of S that defines C_S . The manner of convergence of the sequence $\{\eta_n\}_{n=1,2,\dots}$ to η , and the fact that C_S and each submanifold from $\{C_{S_n}\}_{n=1,2,\dots}$ are J -holomorphic, implies that the integral of w over C_S is finite, and this limit is no larger than the lim-sup of the sequence whose n^{th} component is the integral of w over C_{S_n} . As noted previously, Lemma 7.8 provides an upper bound for this lim-sup.

What with Lemma II.5.6, these last observations imply that each very large $|s|$ slice of C_S is very close to a union of curves in $[1 + z_*, 2 - z_*] \times \Sigma$ whose members project to points in Σ . Moreover, given that η lies in the radius $\frac{1}{100} \kappa_S^{-1} \rho_S^2$ subbundle of N_0 , the following must be true: If $s \gg 1$, then the corresponding set of components of the large s slice of Σ has $\text{genus}(\Sigma)$ elements, and this set enjoys a 1-to-1 correspondence with the points in $C_{1+} \cap C_{2-}$ that define the $s \rightarrow \infty$ limit of S . This is such that each component of a given constant s slice of C_S lies in the radius $c_0 \sqrt{\varepsilon} \delta_*$ disk neighborhood of the corresponding point in $C_{1+} \cap C_{2-}$. Indeed, this last bound on the radius follows from Lemma 7.11 because the latter implies that any integral curve segment in question has endpoints in the radius $\sqrt{\varepsilon} \delta_*$ coordinate ball centered on some index 1 and index 2 critical point of f . Use the correspondence just described to label the components of the large, constant s slices of $C_S|_s$. Introduce now $d(s)$ to denote the diameter of the image in Σ via projection from $[1 + z_*, 2 - z_*] \times \Sigma$ of the component of $C_S|_s$ with a given label. Then $\lim_{s \rightarrow \infty} d(s) = 0$.

There is the analogous story for the constant $s \ll -1$ slices. Each has $\text{genus}(\Sigma)$ components, and the components enjoy an analogous 1-to-1 correspondence with the points in Σ that define the $s \rightarrow -\infty$ limit of S .

Step 5 Fix $p \in \Lambda$. This step considers the large s slices of C_p . Lemma 7.8 supplies an upper bound for the sequence whose n^{th} member is the integral of w over C_{pn} . The

fact that the sequence $\{(\varphi_n^p, \zeta_n^p)\}_{n=1,2,\dots}$ converges smoothly on compact domains in $\mathbb{R} \times I_*$ to (φ^p, ζ^p) implies that the integral of w over C_p is finite. As a consequence, each very large $|s|$ slice of C_p must be everywhere close to the segment of some integral curve of v in the part of $\mathcal{H}_{p_*}^+$ where $\hat{u} \in I_*$. When viewed via Ψ_p in terms of the functions (φ^p, ζ^p) , this means the following: If $|x|$ is very large, then the function $\hat{u} \rightarrow \zeta^p(x, \hat{u})$ on I_* is nearly constant. In particular, this implies the following: given $r > 0$, and then given s sufficiently large, there exists a segment γ of an integral curve of v in the $\hat{u} \in I_*$ part of $\mathcal{H}_{p_*}^+$ such that

$$(7-27) \quad \sup_{q \in C^p|_s} \text{dist}(q, \gamma) + \sup_{q \in \gamma} \text{dist}(C^p|_s, q) < r.$$

Fix a critical point in p so as to consider the corresponding version of (7-1). Fix some $x \in \mathbb{R}$ with $x \gg 1$ so that the corresponding constant $s = s(x)$ slices of C_S and C_p are very near respective segments of integral curves of v . Given that $|\varphi^S - \varphi^{S_0}| < c_0 \sqrt{\varepsilon} \delta_*$ at the given value of x on the common boundary of C_S and C_p , the lower bullet in (7-1) asserts that $|\varphi^p - \zeta^{S_0}| \leq c_0 \sqrt{\varepsilon} \delta_*$ at the given value of x on the boundary of I_* . Let $\gamma^x \subset \mathcal{H}_{p_*}^+$ denote the $\hat{u} \in I_*$ part of an integral curve of v that lies very close to the given constant $s(x)$ slice of C_p . The values of the angle ϕ on the respective boundary points of γ^x and γ_{p+} differ by at most $c_0 \sqrt{\varepsilon} \delta_*$. This has the following consequence: Let $\Delta\phi^x$ denote the change in the angle ϕ along γ^x and let $\Delta\phi^+$ denote the corresponding angle change along γ_{p+} . Then there exists $m \in \mathbb{Z}$ such that

$$(7-28) \quad |\Delta\phi^x - \Delta\phi^+ + 2\pi m| \leq c_0 \sqrt{\varepsilon} \delta_*.$$

To see that $m = 0$ in (7-28) if $\varepsilon < c_0^{-1}$, note that the same constant s slice of any given large n version of C_{pn} must be everywhere close to γ^x also. This understood, define a closed curve in C_{pn} as follows: this constant x slice of C_{pn} with the $x' > x$ parts of the boundary of C_{pn} and the arc γ_{p+} concatenate to define a closed loop in $\mathcal{H}_{p_*}^+$. This loop is null-homologous: it is homologous to the image via the projection from $\mathbb{R} \times \mathcal{H}_{p_*}^+$ of a loop in C_{pn} and the latter space is contractible. Such a loop is null-homologous if and only if the integer m that appears in (7-28) is zero when $\varepsilon < c_0^{-1}$. Note in particular that this bound on ε is purely S -dependent.

As usual, there is an analogous description of the $s \ll -1$ part of C_p .

Step 6 Let $\mathcal{L} \subset [0, 1)$ denote the following set: A number $D \in [0, 1)$ lies in \mathcal{L} if the following is true: Fix any $p \in \Lambda$. Then for any given but sufficiently large $s \in \mathbb{R}$, the slice $C_p|_s$ has distance less than D from γ_{p+} .

What follows in a moment proves that $0 \in \mathcal{L}$. There is an analogous, $s \rightarrow -\infty$ version of \mathcal{L} , and the analogous argument proves that the latter also contains 0. These last facts with (7-1) imply that the constant s slices of C_S and each $\mathfrak{p} \in \Lambda$ version of $C_{\mathfrak{p}}$ converge as $s \rightarrow \pm\infty$ to the segments of integral curves of v that are defined by $(\widehat{\Theta}_-, \widehat{\Theta}_+)$. Given what is said by Steps 1–3, this means that $(\tau, \mathcal{C}) \in \mathcal{M}^*$.

To prove that $0 \in \mathcal{L}$, it is enough to prove the following: if $D \in \mathcal{L}$, then so is $\frac{1}{2}D$. To see that such is the case, let $T_D \subset \mathcal{H}_{\mathfrak{p}_*}^+$ denote the $\widehat{u} \in I_*$ part of the radius D disk neighborhood of $\gamma_{\mathfrak{p}_+}$. Let $\gamma \subset T_{2D}$ denote an integral curve of v . It follows as a consequence of (II.2-4) and (4-12) that

$$(7-29) \quad |\widehat{h}(\gamma_{\mathfrak{p}_+}) - \widehat{h}(\gamma)| < c_0 \chi_0 R^{-1}D.$$

This last bound implies that all sufficiently large x values of $\zeta^{\mathfrak{p}}$ are such that

$$(7-30) \quad |\zeta^{\mathfrak{p}}|_x - \widehat{h}(\gamma_{\mathfrak{p}_+})| \leq c_0 \chi_0 R^{-1}D.$$

Given (7-1), this last point implies that $|\zeta^S|_x - \widehat{h}(\gamma_{\mathfrak{p}_+})| \leq c_0 \chi_0 R^{-1}D$ for all sufficiently large x . Use this with (II.2-6) to deduce that any given $\mathfrak{p} \in \Lambda$ version of $|(\varphi^S - \varphi^{S_0})|_x$ is bounded by $c_0 \chi_0 R^{-1}D$ for all sufficiently large x on the common boundary of C_S and $C_{\mathfrak{p}}$. What with (7-2), this implies that

$$(7-31) \quad |(\varphi^{\mathfrak{p}} - \varphi^{S_0})|_x| \leq c_0 \chi_0 R^{-1}D \quad \text{for all large } x.$$

The preceding inequality has the following implication: If s is sufficiently large, then there is a segment of an integral curve of v in the $\widehat{u} \in I_*$ part of $\mathcal{H}_+^{\mathfrak{p}}$ that is described by some $r \leq c_0 \chi_0 R^{-1}D$ version of (7-27). This implies in particular that $\frac{1}{2}D \in \mathcal{L}$.

As usual, the $s \rightarrow -\infty$ limit of the constant s slices of $C|_s$ behaves in an analogous fashion.

Part 6 This step completes the proof of the $\{\Delta_{\mathfrak{p}} = 0\}_{\mathfrak{p} \in \Lambda}$ case of Proposition 7.3 by verifying that the sequence $\{(\tau_n, \mathcal{C}_n = \{C_{S_n}, \{C_{\mathfrak{p}n}\}_{\mathfrak{p} \in \Lambda})\}_{n=1,2,\dots}$ has a subsequence that converges in the appropriate manner.

The upcoming Lemma 7.12 plays a central role in this argument. Lemma 7.12 assumes that the data chosen to define the geometry of Y and \mathcal{M}^* is such that the conclusions of Lemma 7.11 holds with a constant ε chosen less than a certain purely S -dependent constant. The data is also such that the conclusions of Lemmas 7.9 and 7.10 hold.

Let $(\Theta_-, \Theta_+) \in \mathcal{Z}_{\text{ech}, \mathcal{M}} \times \mathcal{Z}_{\text{ech}, \mathcal{M}}$ denote the pair that lies under $\{\widehat{\Theta}_-, \widehat{\Theta}_+\}$. Since $\Delta_{\mathfrak{p}}$ is zero for all $\mathfrak{p} \in \Lambda$, both lack integral curves of v from $\{\widehat{\gamma}_{\mathfrak{p}}^+, \widehat{\gamma}_{\mathfrak{p}}^-\}_{\mathfrak{p} \in \Lambda}$. Lemma 7.12

and the subsequent arguments view the element Θ_- as a set of segments of integral curves of v written as $\{\Gamma_{S_-}, \{\gamma_{p_-}\}_{p \in \Lambda}\}$, where Γ_{S_-} denotes the union of the parts of the curves in Θ_- where $f \in [1 + z_*, 2 - z_*]$. The set Θ_+ is likewise written as $\{\Gamma_{S_+}, \{\gamma_{p_+}\}_{p \in \Lambda}\}$.

Lemma 7.12 *Given $r > 0$, there exists $s_r > 1$ with the following significance: Suppose that $(\tau, C = \{C_S, \{C_p\}_{p \in \Lambda}\})$ is a given element from \mathcal{M}^* . If $s < -s_r$, then*

$$\sup_{q \in (C_S \cup (\bigcup_{p \in \Lambda} C_p))|_s} \text{dist}\left(q, \left(\Gamma_{S_-} \cup \left(\bigcup_{p \in \Lambda} \gamma_{p_-}\right)\right)\right) + \sup_{q \in (\Gamma_{S_-} \cup (\bigcup_{p \in \Lambda} \gamma_{p_-}))} \text{dist}\left(\left(C_S \cup \left(\bigcup_{p \in \Lambda} C_p\right)\right)|_s, q\right) < r.$$

If $s > s_r$, then

$$\sup_{q \in (C_S \cup (\bigcup_{p \in \Lambda} C_p))|_s} \text{dist}\left(q, \left(\Gamma_{S_+} \cup \left(\bigcup_{p \in \Lambda} \gamma_{p_+}\right)\right)\right) + \sup_{q \in (\Gamma_{S_+} \cup (\bigcup_{p \in \Lambda} \gamma_{p_+}))} \text{dist}\left(\left(C_S \cup \left(\bigcup_{p \in \Lambda} C_p\right)\right)|_s, q\right) < r.$$

Lemma 7.12 is proved in a moment.

Lemma 7.12 with what is said in Parts 1–5 imply that \mathcal{M}^* is compact in the topology that is defined in Section 7B. Lemma 7.12 has the following additional implication: Given the lemma, standard elliptic regularity arguments using (7-8) and the various $p \in \Lambda$ versions of (3-4) can be applied to prove that the topology as defined in Section 7A on \mathcal{M}^* is the same as the strong C^∞ topology. To elaborate, these tools can be used to bootstrap from the uniform L^∞ convergence that is asserted by Lemma 7.12 at large s to prove strong convergence of the sort asserted by Proposition 7.3 but with respect to some Hölder topology with exponent $\nu > 0$. The tools are used again with (7-8) and (3-4) to prove the strong $C^{1+\nu}$ convergence, then again to prove strong $C^{2+\nu}$ convergence, and so on. The sorts of tools needed can be found in Chapter 6 of [12].

This equivalence between the topology from Section 7A and the strong C^∞ topology implies what is asserted by the two bullets of Proposition 7.3.

Proof of Lemma 7.12 Suppose that the lemma is false so as to generate nonsense. Granted that such is the case, there exists $r > 0$ and a sequence $\{(\tau_n, C_n)\}_{n=1,2,\dots}$ of the following sort: Each $n \in \{1, 2, \dots\}$ version of (τ_n, C_n) is an element in \mathcal{M}^* and

such that the condition stated by Lemma 7.12 fails to hold for $s = s_n \geq n$ with some fixed choice of r . Given the conclusions of the preceding Parts 1–5 of this subsection, no generality is lost by requiring that the sequence $\{(\tau_n, C_n)\}_{n=1,2,\dots}$ converge in the topology on \mathcal{M}^* from Part 5 to the element $(\tau, C) \in \mathcal{M}^*$. The derivation of nonsense from the existence of such a sequence has two steps.

Step 1 Let $\hat{\nu}_+$ denote the HF-cycle that is used to define Θ_+ . Write $f^{-1}(1, 2) \subset M$ as $(1, 2) \times \Sigma$ so as to identify this set of integral curves with a set given by $(1, 2) \times \hat{q}$ with \hat{q} denoting a certain set of G distinct elements from $C_{1+} \cap C_{2-}$. Let S_- denote the submanifold $\mathbb{R} \times [1, 2] \times \hat{q}_-$ in $\mathbb{R} \times [1, 2] \times \Sigma$. This S_- is a Lipschitz submanifold. Introduce \mathcal{M}_-^* to denote the S_- version of the space \mathcal{M}^* as defined using $\hat{\Theta}_-$ for both the $s \rightarrow -\infty$ and $s \rightarrow \infty$ limit conditions on its constituent elements. For any given $\tau \in [0, 1]$, the set $\{\mathbb{R} \times \Gamma_{S_-}, \{\mathbb{R} \times \gamma_{p_-}\}_{p \in \Lambda}\}$ is a (J, τ) -holomorphic submanifold. Thus \mathcal{M}_-^* is nonempty. These are the only elements in \mathcal{M}_-^* . To prove this, let $(\tau, C = \{C_S, \{C_p\}_{p \in \Lambda}\})$ denote a given element \mathcal{M}_-^* . Compute the sum whose constituent terms are the integral of w over C_S and the integral of w over the respective $p \in \Lambda$ versions of C_p . If this sum is zero, then $C = \{\mathbb{R} \times \Gamma_{S_-}, \{\mathbb{R} \times \gamma_{p_-}\}_{p \in \Lambda}\}$. The fact that this sum of integrals is indeed zero is proved in the next paragraph.

To see that the sum of integrals is zero, note that the form w is exact on a neighborhood in Y of the union of the sets $\{\mathcal{H}_{p_*}^+\}_{p \in \Lambda}$ with a uniform radius tubular neighborhood of the arcs that form Γ_{S_-} . This being the case, integration by parts as in the proof of Lemma 7.8 writes the sum of the integrals of w over the constituents of C as a sum of three terms: The first is the integral of the antiderivative 1-form over the union of the curves that form Θ_- . The second is minus the integral of this same 1-form over the same union of curves. The third is itself a sum, this sum indexed by Λ with any given term given by (7-13). As both φ^{S_0} and ζ^{S_0} are constant in each $p \in \Lambda$ version of (7-13), each $p \in \Lambda$ contribution to the third term is zero.

Step 2 For each $n \in \{1, 2, \dots\}$, translate each element in C_n by $-s_n$ along the \mathbb{R} factor in either $\mathbb{R} \times [1, 2] \times \Sigma$ or the appropriate $p \in \Lambda$ version of $\mathbb{R} \times \mathcal{H}_{p_*}^+$. Let C'_n denote the corresponding set. Minor modifications of the arguments given in Parts 1–5 prove that the sequence $\{(\tau'_n, C'_n)\}_{n=1,2,\dots}$ converges in the topology given by (7-26) to some element in the space \mathcal{M}_-^* . Given the assumptions about the initial sequence, this element cannot be $(\tau, \{\mathbb{R} \times \Gamma_{S_-}, \{\mathbb{R} \times \gamma_{p_-}\}_{p \in \Lambda}\})$ as at least one of the submanifolds of each $n \in \{1, 2, \dots\}$ version of C'_n contains a point that has distance at least r from a point in the corresponding submanifold from the set $\{\mathbb{R} \times \Gamma_{S_-}, \{\mathbb{R} \times \gamma_{p_-}\}_{p \in \Lambda}\}$. But this is nonsense given what is said in Step 1. \square

Part 7 This part adds what is needed to the arguments in Parts 1–6 so as to prove Proposition 7.3 when some $p \in \Lambda$ versions of Δ_p are 1 or 2. The arguments given in Parts 1–6 can be seen as having three components. The first component is summarized by Lemma 7.8; this component gives bounds on the integrals of w and $ds \wedge \hat{a}$. The second component is summarized by Lemma 7.11; this component controls the behavior the relevant submanifolds where they intersect $\mathbb{R} \times M_\delta$. The final component controls the behavior of the remaining portions of the constituent subvarieties. This component comprises Lemmas 7.9 and 7.10 and Parts 5 and 6. The three steps that follow speak to these three components in the case when some $p \in \Lambda$ versions of Δ_p are 1 or 2.

Step 1 The analog of Lemma 7.8 in the general case makes the same assertion as the original with the $\{\Delta_p = 0\}_{p \in \Lambda}$ condition omitted. The proof copies the arguments that prove Lemma 6.1 to control the integral of w on the relevant parts of any given (τ, C) and $p \in \Lambda$ version of C_p and it copies what is said in Part 1's proof of Lemma 7.8 to control the integral of w over the (τ, C) version of C_S . What is said in Part 1 with regards to (7-13) holds whatever the value of Δ_p .

Step 2 The analog of Lemma 7.11 in the general case is identical to its namesake but for the absence of the $\{\Delta_p = 0\}_{p \in \Lambda}$ assumption. There is but one change in the proof. The arguments for lemma when $\Delta_p = 0$ derive nonsense from a certain a priori assumption about the sequence any given $p \in \Lambda \setminus \Lambda_*$ version of the sequence $\{C_{pn}\}_{n=1,2,\dots}$ as they find a loop in all large n versions of C_{pn} that generates the first homology of $\mathbb{R} \times \mathcal{H}_{p_*}^+$. This nonsense comes from the fact that C_{pn} is contractible if $p \in \Lambda \setminus \Lambda_*$. The space C_{pn} is not contractible if $\Delta_p \neq 0$, but even so the existence of the corresponding loop in C_{pn} is nonsense: the loop in question sits in a component of the $|u| > 0$ part of C_{pn} , and each such component is contractible.

Step 3 The analog of Lemma 7.9 in the case when some $p \in \Lambda$ versions of Δ_p are nonzero is identical to its namesake but for the absence of the $\{\Delta_p = 0\}_{p \in \Lambda}$ assumption. The proof of the analog is identical to that given for Lemma 7.9 with it understood that Lemma 7.11 holds when various $p \in \Lambda$ versions of $\Delta_p > 0$.

The analog of Lemma 7.10 replaces the latter with two lemmas. The first one is much like Lemma 6.3.

Lemma 7.13 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$ and given $z_* < \kappa^{-1}$, there exists $\kappa_* > 1$ that depends only on z_* but is otherwise purely S -dependent (or \mathcal{K} -compatible), and these have the following significance: Define \mathcal{M}^**

using $z_* < \kappa^{-1}$, using $\delta < \kappa_*^{-1} z_*$ and using a pair from $\widehat{\mathcal{Z}}^S$. There exists $\kappa_{**} > 1$ and, given $\varepsilon \in (0, 1]$, there exists $\kappa_\varepsilon > 1$ such that any given $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}) \in \mathcal{M}^*$ has the properties listed next:

- If $\Delta_p = 0$, then $1 - 3 \cos^2 \theta > \kappa_{**}^{-1}$.
- If $\Delta_p > 0$, then
 - (a) $1 - 3 \cos^2 \theta > \kappa_{**}^{-1}$ on the $s < -\kappa$ part of C_p .
 - (b) $1 - 3 \cos^2 \theta > \kappa_\varepsilon^{-1}$ on the $|u| \geq \varepsilon$ part of C_p .
- If $\Delta_p = 1$ and $m_p = -1$, then $\cos \theta > -\frac{1}{\sqrt{3}} + \kappa_{**}^{-1}$ on the whole of C_p
- If $\Delta_p = 1$ and $m_p = 1$, then $\cos \theta < \frac{1}{\sqrt{3}} - \kappa_{**}^{-1}$ on the whole of C_p

Proof Given that the assertions of Lemma 7.11 are true when some $p \in \Lambda$ versions of $\Delta_p > 0$, the proof of Lemma 7.13 is obtained by using the arguments for the proof of Lemma 7.10 to prove what is asserted in the first bullet, and by using the arguments for the proof of Lemma 6.3 or its $(\Delta_p = 1, m_p = 1)$ or $\Delta_p = 2$ incarnations to prove the assertions of the second and third bullets. □

To set the stage for the second lemma, suppose for the sake of argument that $p \in \Lambda$ has $\Delta_p = 1$ and $m_p = -1$. There are, in this case, constants $s_p \geq 1$ and $c_{1_p} \in \mathbb{R} \setminus 0$ and $\phi_{1_p} \in \mathbb{R}/2\pi\mathbb{Z}$ with the following significance: The complement of a certain compact set of the $s \geq 1$ part of C_p has two components. One is a strip diffeomorphic to $[s_p, \infty) \times I_*$ whose image via the projection to $\mathcal{H}_{p_*}^+$ is very close to γ_{p_+} . The other is a cylinder whose image via this projection lies in the tubular neighborhood U_+ of $\widehat{\gamma}_p^+$ that is described in Section 5A. Let \mathcal{E} denote this end of C_p . Reintroduce the coordinates $(s_+, \phi_+, \theta_+, u_+)$ for $\mathbb{R} \times U_+$ as defined in (5-5). Then \mathcal{E} sits in the $s_+ \in [s_p, \infty)$ part of $\mathbb{R} \times U_+$ as a smooth, properly embedded submanifold with boundary on the $s_+ = s_p$ slice. Furthermore, this intersection is given by the graph of a smooth map as described in the first bullet of Proposition 5.1 with domain $[s_p, \infty) \times \mathbb{R}/2\pi\mathbb{Z}$ that has the form depicted in (6-23) with $c_1 = c_\varepsilon$ and with η_{1_+} defined using $\phi_1 = \phi_\varepsilon$. There is a completely analogous picture when $(\Delta_p = 1, m_p = 1)$ and when $\Delta_p = 2$.

Lemma 7.14 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa > 1$ and, given $z_* < \kappa^{-1}$, there exists $\kappa_* > 1$ that depends only on z_* but is otherwise purely S -dependent (or \mathcal{K} -compatible), and these have the following significance: Define \mathcal{M}^* using $z_* < \kappa^{-1}$, using $\delta < \kappa_*^{-1} z_*$ and using a pair from $\widehat{\mathcal{Z}}^S$. There exists $\kappa_{**} \geq 1$ such that if $(\tau, \{C_S, \{C_p\}_{p \in \Lambda}\})$ is from \mathcal{M}^* and if $p \in \Lambda$ is such that $\Delta_p > 0$, then*

$s_p < \kappa_{**}$ and $\kappa_{**}^{-1} \leq |c_\mathcal{E}| \leq \kappa_{**}$ for each end $\mathcal{E} \subset C_p$ whose constant $s \gg 1$ slices converge as $s \rightarrow \infty$ to either $\widehat{\gamma}_p^+$ or $\widehat{\gamma}_p^-$.

Proof But for notation, the argument is the same as that used to prove Lemma 6.4. \square

With Lemmas 7.13 and 7.14 in hand, the remaining arguments for the third component differ only cosmetically from the arguments given in Parts 5 and 6 of this section and so no more will be said.

7E The structure of \mathcal{M}^*

There are four parts to this subsection; Part 4 contains the proofs of Propositions 7.1 and 7.2. Parts 1–3 set up the necessary machinery.

Part 1 Let $(\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\})$ denote a given element of \mathcal{M}^* . This step associates a certain Fredholm operator to (τ, \mathcal{C}) whose cokernel provides a specific version of \mathbb{R}^n for use in Proposition 7.1. This operator is also used in the proof of Proposition 7.2. The operator is denoted in what follows by $D_{\mathcal{C}}$.

A dense domain for $D_{\mathcal{C}}$ consists of a direct sum of function spaces with the first summand labeled by S and the others labeled by the set Λ . An element in the summand labeled by S is, among other things, a smooth and compactly supported section of the normal bundle N_S on the $t \in [1 + z_*, 2 - z_*]$ part of S . These sections are further constrained to be L^2 -orthogonal to the restriction of the L^2 kernel of (1-25)'s operator D_S to the $t \in [1 + z_*, 2 - z_*]$ part of S . Meanwhile, any given $p \in \Lambda$ labeled summand consists of a certain sort of compactly supported, smooth sections of the normal bundle to C_p . The respective elements in the S -summand and any given $p \in \Lambda$ summand are further constrained on the common boundary of C_p and C_S . The next paragraph describes this constraint.

Write the intersection of C_S with \mathcal{H}_p^+ as a graph in the manner of Property 4 in Part 1 of Section 7A so as to identify a section of the normal bundle of C_S on either boundary component as a map from \mathbb{R} to \mathbb{R}^2 . Use x for the \mathbb{R} coordinate. The coordinate z for the $[z_*, z_S]$ factor is $z = t - 1$ for the index 1 critical point side of \mathcal{H}_p^+ and $z = 2 - t$ for the index 2 critical point side. With a section of S given, write the components of the corresponding map as $(x, z) \rightarrow (\varphi^{S'}, \zeta^{S'})|_{(x,z)}$. View C_p via Ψ_p as a graph in $\mathbb{R} \times \mathcal{X}$ in the manner of Property 2 in Part 2 of Section 7A, and use the 1-forms $(d\widehat{\phi}, dh)$ to identify the normal bundle with the product \mathbb{R}^2 -bundle. Having done so,

a section of this normal bundle becomes a map to \mathbb{R}^2 from the domain in $\mathbb{R} \times I_*$ of the pair (φ^p, ζ^p) . Introduce the coordinate $z = e^{-2(R-\hat{u})}$ for the $\hat{u} \geq R + \ln \delta$ part of the domain and $z = e^{-2(R+\hat{u})}$ for the $\hat{u} \leq -R - \ln \delta$ part. Granted this notation, the following constraint holds where $z = z_*$ in either case:

$$(7-32) \quad \begin{aligned} &\bullet \quad \zeta^{S'} = \tau \zeta^{p'}, \\ &\bullet \quad \tau \varphi^{S'} = \varphi^{p'}. \end{aligned}$$

What follows describes the operator \mathcal{D}_C . To start, \mathcal{D}_C acts diagonally with respect to the labeling of the summands of the domain and range spaces. To define its action on the S -labeled summand, write C_S as in Section 7A in terms of a section of the normal bundle over the $t \in [1 + z_*, 2 - z_*]$ part of S . Let η denote this section. The action of \mathcal{D}_C on the S -labeled summand is that of an operator that is denoted in what follows by D_η . Let η' denote a section of N_S over the part of S where $t \in [1 + z_*, 2 - z_*]$. The operator D_η sends η' to

$$(7-33) \quad \bar{\partial}\eta' + \tau_1(\eta) \cdot \partial\eta' + (\nabla_{\eta'} \tau_1)|_\eta \partial\eta + (\nabla_{\eta'} \tau_0\eta)|_\eta,$$

where the notation uses $\nabla_{\eta'}$ to denote the directional derivative along the fiber of N_S in the direction given by η' . The terms that involve $\nabla_{\eta'}$ are zeroth-order and \mathbb{R} -linear. By way of an example, the operator in (7-18) is the $\eta = 0$ version of (7-33).

To define the action of \mathcal{D}_C on the remaining summands, fix $p \in \Lambda$ and let (φ^p, ζ^p) denote the pair that defines C_p . Use D_p to denote the $(\varphi = \varphi^p, \zeta = \zeta^p)$ version of the operator that is depicted in (3-6). With the elements in the p -labeled summand viewed as maps from the relevant domain in $\mathbb{R} \times I_*$ to \mathbb{R}^2 , the action of \mathcal{D}_C on the p -labeled summand is given by D_p .

To say slightly more about the operators from the set $\{D_\eta, \{D_p\}_{p \in \Lambda}\}$, keep in mind that C_S and each $p \in \Lambda$ version of C_p is J -holomorphic. This being the case, the normal bundle of each can be viewed as a complex line bundle. Meanwhile, the Riemannian metric defined by J and the compatible 2-form $ds \wedge \hat{u} + \omega$ endow these submanifolds and their normal bundles with holomorphic structures. This understood, let N denote the normal bundle to a given $C \in \{C_S, \{C_p\}_{p \in \Lambda}\}$ but viewed now as a complex line bundle. Writing C_S in terms of the section η identifies the $t \in [1 + z_*, 2 - z_*]$ part of the normal bundle N_S with the C_S version of N . The S -labeled part of the domain of \mathcal{D}_C can be viewed as a section of this version of N . Do so and D_η appears as a first-order differential operator that maps $C^\infty(C_S; N)$ to $C^\infty(C_S; N \otimes T^{0,1}C_S)$ which has the schematic form of the operator in (1-25). Meanwhile, each $p \in \Lambda$ version

of D_p maps the space $C^\infty(C_p; N)$ to $C^\infty(C_p; N \otimes T^{0,1}C_p)$ and it also can be written in the form depicted by (1-25).

The Banach space domain for \mathcal{D}_C is obtained by completing the dense domain described above using the norm whose square is the sum of the squares of certain L^2_1 -norms on the constituent summands. The L^2_1 -norm on the S -labeled summand is that defined by the covariant derivative and metric on S . With C_S viewed as the image of ϵ_S of the section η , this norm is equivalent to the one defined by the induced Riemannian metric on TC_S and the induced metric and covariant derivative for the normal bundle of C_S . The L^2_1 -norm on any given p -labeled summand is defined using the induced Riemannian metric for TC_p and the induced metric and covariant derivative for the normal bundle to C_p .

The operator \mathcal{D}_C maps the Banach space just defined to the Banach space that is defined by completing the respective space of compactly supported sections of each $C \in \{C_S, \{C_p\}_{p \in \Lambda}\}$ version of $C^\infty(C; N \otimes T^{0,1}C)$ using the induced L^2 inner products. In the case $C = C_S$, this is the same as completing a space of compactly supported sections of the appropriate bundle over the $t \in [1 + z_*, 2 - z_*]$ part of S using the L^2 -norm on S .

The respective domain and range Banach spaces for \mathcal{D}_C are denoted in what follows by \mathbb{H}_τ and \mathbb{L} . The norm on \mathbb{H}_τ is denoted by $\|\cdot\|_1$, and the L^2 -norm on either \mathbb{H}_τ or \mathbb{L} is denoted by $\|\cdot\|$. The operator \mathcal{D}_C defines a bounded operator from \mathbb{H}_τ to \mathbb{L} .

Respective domain and range Banach spaces with slightly stronger norms are also needed in what follows. The domain version is denoted in what follows by \mathbb{H}_{τ_*} . The norm that defines \mathbb{H}_{τ_*} is the sum of the L^2_1 -norm defined above and the norm whose square is given on each $p \in \Lambda$ -labeled summand by the C_p analog of (5-31), and is given on the S -labeled summand by the analog of (5-31) for sections of the normal bundle of S . Lemma 5.10 has an analog for S ; they assert that elements in \mathbb{H}_{τ_*} are Hölder continuous and the associated map from \mathbb{H}_{τ_*} to the relevant Hölder space is continuous. The strengthened range Hilbert space is denoted by \mathbb{L}_* ; the square of the norm that defines this space is given on the various summands by replacing in (5-31) the length of the covariant derivative of a given section with that of the section itself. The operator \mathcal{D}_C defines also a bounded map from \mathbb{H}_{τ_*} to \mathbb{L}_* .

Part 2 The lemma that follows states what is needed concerning the operator \mathcal{D}_C .

Lemma 7.15 *There is a purely S -dependent (or \mathcal{K} -compatible) $\kappa \geq 1$ such that if \mathcal{M}^* is defined using $\rho_S \leq \kappa^{-1}$, $z_* \leq \kappa^{-1}$, $\delta < \kappa^{-2}z_*$ and any pair from $\widehat{\mathcal{Z}}^S$, then the*

following is true: Use any given $(\tau, C) \in \mathcal{M}^*$ to define the operator \mathcal{D}_C and the spaces \mathbb{H}_τ and \mathbb{L} . The operator \mathcal{D}_C is a Fredholm operator from \mathbb{H}_τ to \mathbb{L} with index equal to $\sum_{p \in \Lambda} \Delta_p$. Moreover, elements in the kernel of \mathcal{D}_C are C^∞ elements in \mathbb{H}_{τ_*} , and elements in the cokernel of \mathcal{D}_C are represented by smooth elements in \mathbb{L}_* .

The remainder of this Part 2 contains the:

Proof of Lemma 7.15 The proof of this lemma has five steps.

Step 1 Just as in Part 4 of Section II.6E and in the proofs of Propositions 3.2 and 5.7, the assertion that \mathcal{D}_C has closed range and finite-dimensional kernel follows if there exists a constant $c > 1$ such that the following two conditions hold:

- (7-34) • $\|\mathcal{D}_C \mathfrak{h}\|^2 \geq c^{-1} \|\mathfrak{h}\|_1^2 - c \|\mathfrak{h}\|^2$ for any $\mathfrak{h} \in \mathbb{H}_\tau$.
- There exists $s_1 > 1$ such that if $\mathfrak{h} \in \mathbb{H}$ has S -summand with support where $|s| > s_1$ and each Λ -labeled summand has support where $|s| \geq s_1$, then $\|\mathcal{D}_C \mathfrak{h}\|^2 \geq c^{-1} \|\mathfrak{h}\|^2$.

There is a purely S -dependent (or \mathcal{K} -compatible) constant $c > 1$ such that if both $\rho_S < c^{-1}$ and $\delta < c^{-1}$, then the conditions in (7-34) hold for those $\mathfrak{h} \in \mathbb{H}_\tau$ with the following property: The closure of the support of \mathfrak{h} is disjoint from the boundary of the submanifolds from \mathcal{C} . Indeed, if a section of N_S has compact support on the interior of C_S then this follows from what is said in Parts 3 and 4 of Section II.6E given that what is denoted in (7-33) by τ_1 obeys $|\tau_1(\eta)| \leq c_0^{-1} |\eta|$. If a section has support in some $p \in \Lambda$ version of C_p , then this follows from Propositions 3.2 and 5.7.

To see about the top bullet in (7-34) when the support of the summands of \mathfrak{h} intersect the boundaries of the defining domains, write \mathcal{C} as $\{C_S, \{C_p\}_{p \in \Lambda}\}$ and fix attention on a given $p \in \Lambda$. Suppose that $\mathfrak{h} \in \mathbb{H}_\tau$ is in the dense domain, and write its S -summand near the intersection of $\mathbb{R} \times M_\delta$ with $\mathbb{R} \times \mathcal{H}_{p_*}^+$ as a pair of functions on the domain $\mathbb{R} \times [z_*, z_S]$ as done in (7-32). These are denoted by $(\varphi^{S'}, \zeta^{S'})$. Likewise, write an element in the p -summand as a pair $(\varphi^{p'}, \zeta^{p'})$, these being functions on $\mathbb{R} \times [\delta^2, z_*]$. Both D_η and D_p appear here as the Cauchy–Riemann operator that acts on a given (φ', ζ') to give $(\partial_x \varphi' - \partial_z \zeta', \partial_x \zeta' + \partial_z \varphi')$. Given that (7-32) holds, so does Lemma 7.4. The conclusions of this lemma as applied for all $p \in \Lambda$ prove the top bullet in (7-34).

The argument for the second bullet in (7-34) when the closure of the support of a given element in \mathbb{H}_τ intersects the boundaries of the submanifolds from \mathcal{C} occupies Steps 2 and 3 of what follows. These steps focus on the case where the respective summands in a given element from the dense domain of \mathbb{H}_τ have support where $s \ll -1$ on C_S

and each $p \in \Lambda$ version of C_p . With regards to C_p , no generality is lost by assuming that there exists a purely S -dependent (or \mathcal{K} -compatible) constant $c \geq 1$ such that the support of the C_p -summand lies where $1 - 3 \cos^2 \theta > c^{-1}$. This is because of what was said at the outset: (7-34) holds for sections which are supported on closed sets in C_p 's interior. But for notation, the argument works when the summands have support where $s \gg 1$ on the various submanifolds from \mathcal{C} .

Step 2 As in Part 4 of Section II.6E and the proofs of Propositions 3.2 and 5.7, the second bullet in (7-34) is obeyed if a certain nonnegative-definite, quadratic form has trivial kernel. The quadratic form is defined on a certain space of sections of the normal bundle in Y to the union of the curves in Θ_- . This step defines the relevant space of sections and the quadratic form. The quadratic form in question is denoted by \mathcal{Q} in what follows and the domain by \mathbb{V}_τ .

To give the definitions, first decompose the curves from Θ_- that are not contained in a single $p \in \Lambda$ version of C_p so as to define a set of segments of the form $\{\Gamma_S, \{\gamma_{p_-}\}_{p \in \Lambda}\}$ where $\Gamma_S \subset M_\delta$ is the union of the G segments that form the intersection of the union of the curves from Θ_- with the $f \in [1 + z_*, 2 - z_*]$ part of M_δ .

View the $f \in [1, 2]$ part of M_δ as a subset of $[1, 2] \times \Sigma$. Doing so identifies Γ_S as $[1, 2] \times \Lambda^S$ where $\Lambda^S \subset \Sigma$ is a set of N distinct points. Fix a holomorphic coordinate centered on each point in Λ^S . Doing so identifies a section of the normal bundle of Γ^S with a pair of real functions on the interval $[1 + z_*, 2 - z_*]$, these corresponding to the real and imaginary parts of the complex coordinate. A section of the normal bundle in Y to a segment from Γ_S that concatenates with a given $p \in \Lambda$ version of γ_{p_-} appears at a shared endpoint as a pair of functions on $[z_*, z_S]$, these denoted by $(\varphi^{S'}, \zeta^{S'})$.

Fix $p \in \Lambda$ and use the map Ψ_p to view γ_{p_-} as the locus in $\mathbb{R} \times \mathcal{X}$ where the coordinates are such that $x = 0$ and $h = h(\gamma_{p_-})$. Doing so identifies any given section of the normal bundle to γ_{p_-} with a pair of functions on I_* . Such a pair is written as $(\varphi^{p'}, \zeta^{p'})$. Near the boundary of I_* these can be written as a pair of functions on the interval $[\delta^2, z_*]$ by writing z on this interval as $z = e^{-2(R-|\hat{u}|)}$.

Given this notation, the domain, \mathbb{V}_τ , for \mathcal{L} consists of the direct sum of $G + 1$ function spaces, the first labeled by S and the others labeled by Λ . The S -labeled summand consists of a suitably constrained set of sections of the normal bundle to Γ_S . Meanwhile, any given $p \in \Lambda$ -labeled summand consists of a suitably constrained set of pairs of functions on I_* . The constraints on the components of the various summands are the boundary constraints given by the various $p \in \Lambda$ versions of (7-32).

The quadratic form \mathcal{Q} is the sum of quadratic forms that are defined on the various summands of \mathbb{V}_τ . To define the contribution to this sum from the S -labeled summand, view an element of the latter as a map from $[1 + z_*, 2 - z_*]$ to \mathbb{R}^2 suitably constrained on its boundary. Let $t \rightarrow \eta'(t)$ denote such a map. The value of the S -labeled quadratic form on η is

$$(7-35) \quad \mathcal{Q}_S(\eta') = \int_{1+z_*}^{2-z_*} \left| \frac{d}{dt} \eta' \right|^2 dt.$$

Let p denote a given element in Λ . The contribution to \mathcal{Q} from the summand labeled by p is given by the expression in (3-15) with it understood that the coefficients a_{1-} , a_{2-} and b_{2-} are defined using γ_{p-} . In particular, b_{2-} is the $\gamma = \gamma_{p-}$ version of what appears on the right-hand side of (3-11).

Step 3 Part 4 of Section II.6E and Section 3C prove that \mathcal{Q} is positive definite in the case $\tau = 0$. This understood, assume $\tau > 0$ in what follows. To see that \mathcal{Q} is positive definite when $\tau > 0$, suppose that $h \in \mathbb{V}_\tau$ and $\mathcal{Q}(h) = 0$. Fix $p \in \Lambda$ and write the corresponding component of h as $(\varphi^{p'}, \zeta^{p'})$. It follows from (3-15) that $\zeta^{p'}$ is constant along I_* , and it follows from (3-17) that $\varphi^{p'}$ increases along I_* . To elaborate, let $\varphi_-^{p'}$ denote the $\hat{u} = -R - \frac{1}{2} \ln z_*$ value of $\varphi^{p'}$ and let $\varphi_+^{p'}$ denote the value at $\hat{u} = R + \frac{1}{2} \ln z_*$. Then (3-15) and (3-13) imply that

$$(7-36) \quad \varphi_+^{p'} - \varphi_-^{p'} = c_p \chi_0^{-1} R \zeta^{p'},$$

where $c_p > c_0^{-1}$. Let η' denote the component of h in the S -summand of \mathbb{V}_τ . It follows from (7-35) that the latter is constant.

Let $\gamma \in \Gamma_S$ denote the segment whose start point is the $\hat{u} = R + \frac{1}{2} \ln z_*$ endpoint of γ_{p-} . Write η' near the start point of γ as $(\varphi_-^\gamma, \zeta_-^\gamma)$, so as to see the implications of what was just said. This pair of functions is independent of the parameter z on the interval $[z_*, z_S]$. The constraint in (7-32) demands that

$$(7-37) \quad \zeta_-^\gamma = \tau \zeta^{p'} \quad \text{and} \quad \varphi_-^\gamma = \tau^{-1} \varphi_+^{p'}.$$

The endpoint of γ is the $\hat{u} = -R - \frac{1}{2} \ln z_*$ boundary of some $q \in \Lambda$ version of γ_q . Introduce the analogous pair $(\varphi_+^\gamma, \zeta_+^\gamma)$. It follows from what is said in the proof of Proposition II.2.7 that these are determined by $(\varphi_-^\gamma, \zeta_-^\gamma)$ via a formula of the form

$$(7-38) \quad \zeta_+^\gamma = -a_\gamma \zeta_-^\gamma - b_\gamma \varphi_-^\gamma \quad \text{and} \quad \varphi_+^\gamma = c_\gamma \zeta_-^\gamma + d_\gamma \varphi_-^\gamma,$$

where the coefficients here are such that $b_\gamma \neq 0$ and such that

$$(7-39) \quad \begin{pmatrix} a_\gamma & b_\gamma \\ c_\gamma & d_\gamma \end{pmatrix}$$

has determinant -1 . Meanwhile, it follows from (7-32) that the pair $(\varphi_-^{q'}, \zeta_-^{q'})$ is given by

$$(7-40) \quad \zeta_-^{q'} = -\tau^{-1} \zeta_+^\gamma \quad \text{and} \quad \varphi_-^{q'} = \tau \varphi_+^\gamma.$$

The simplest case to analyze is that where $q = p$ and so the endpoint of γ is the start point of γ_{p-} . If this is the case, then (7-36), (7-37), (7-38) and (7-40) require that

$$(7-41) \quad \begin{aligned} \zeta^{p'} &= a_\gamma \zeta^{p'} + b_\gamma \tau^{-2} (\varphi_-^p + c_p \chi_0^{-1} R \zeta^{p'}), \\ \varphi_-^{p'} &= c_\gamma \tau^2 \zeta^{p'} + d_\gamma (\varphi_-^p + c_p \chi_0^{-1} R \zeta^{p'}). \end{aligned}$$

It follows from (7-41) that $\mathfrak{h} \neq 0$ if and only if the matrix

$$(7-42) \quad \begin{pmatrix} a_\gamma + \tau^{-2} b_\gamma c_p \chi_0^{-1} R & \tau^{-2} b_\gamma \\ \tau^2 c_\gamma + d_\gamma c_p \chi_0^{-1} R & d_\gamma \end{pmatrix}$$

has $+1$ as an eigenvalue. But its determinant is -1 and its trace is larger than $c_0^{-1} \chi_0^{-1} R$, so one eigenvalue has norm greater than $c_0^{-1} \chi_0 R$ and the other has norm less than $c_0 \chi_0 R^{-1}$.

The argument in the general case is much the same. In this case, $\mathfrak{h} \neq 0$ if and only if the product of some $k \in \{1, \dots, G\}$ versions of (7-42) have $+1$ as an eigenvalue. The details of the linear algebra are straightforward and so left to the reader.

Step 4 Granted that \mathcal{D}_C has closed range, it follows that its cokernel is isomorphic to the kernel of its adjoint, this a bounded operator from \mathbb{L} to \mathbb{H}_τ . The kernel of the latter is isomorphic to the kernel of the formal, L^2 adjoint of \mathcal{D}_C . The proof that such is the case amounts to a standard application of linear, elliptic regularity arguments as applied on the interiors of the domains of the summands of \mathbb{L} and an appeal to Lemma 7.7 to deal with the boundary conditions for these domains. This formal L^2 adjoint is denoted by $\mathcal{D}_C^\#$. The latter is elliptic, first-order and \mathbb{R} -linear, with leading-order symbol given by the adjoint of the $\bar{\partial}$ symbol. Its kernel obeys (7-32). The kernel of $\mathcal{D}_C^\#$ is finite-dimensional if the $\mathcal{D}_C^\#$ version of (7-34) holds. The proof that such is the case is, but for notation, identical to the proof just given for the \mathcal{D}_C version.

To see about the index of \mathcal{D}_C , fix $\tau' \in [0, 1]$. The operator \mathcal{D}_C also defines a bounded map from $\mathbb{H}_{\tau'}$ to \mathbb{L} , and the argument given above that the map $\mathcal{D}_C: \mathbb{H}_\tau \rightarrow \mathbb{L}$ is

Fredholm can be repeated using now the τ' version of (7-32) to see that \mathcal{D}_C is a Fredholm map from $\mathbb{H}_{\tau'}$ to \mathbb{L} also. Meanwhile, the family $\{\mathbb{H}_{\tau'}\}_{\tau' \in [0,1]}$ defines a smooth, Banach space bundle over $[0, 1]$. The proof that this is so mimics what is said in Step 3 of Part 3 in Section 5C. Granted this smoothness, and granted that the set of Fredholm maps between Banach spaces is an open subset in the Banach space of bounded maps, it follows that the index of \mathcal{D}_C as a map from any given τ' version of $\mathbb{H}_{\tau'}$ is independent of τ' . In particular, the index of the \mathbb{H}_{τ} version is the same as its index on the $\mathbb{H}_{\tau'=0}$ version.

The $\tau' = 0$ matching conditions do not couple the various summands that form $\mathbb{H}_{\tau'}$; they give separate conditions on each summand. As a consequence, the index of \mathcal{D}_C on $\mathbb{H}_{\tau'}$ is the sum of the indices of the operators on the various summands with certain boundary conditions. The operator on the S -summand is D_η with the boundary conditions $\zeta^{S'} = 0$. The latter has index zero because it is given up to a term with small norm by the operator discussed in Section II.6E on a restricted domain, the orthogonal complement of its kernel. Meanwhile, any given $\mathfrak{p} \in \Lambda$ version of $D_{\mathfrak{p}}$ with the boundary condition $\varphi^{\mathfrak{p}'} = 0$ is of the form that is considered by Propositions 3.2 and 5.7, and the latter have index $\Delta_{\mathfrak{p}}$. Thus, the index of the $\mathbb{H}_{\tau'=0}$ version of \mathcal{D}_C is $\sum_{\mathfrak{p} \in \Lambda} \Delta_{\mathfrak{p}}$.

Step 5 The assertion that elements in the kernel of \mathcal{D}_C are smooth and in \mathbb{H}_{τ_*} follows using standard linear elliptic estimates for smoothness in the interiors of the submanifolds from \mathcal{C} . Lemma 7.7 gives smoothness and the \mathbb{H}_{τ_*} -norm bound near the boundaries. The same argument as applied to the formal L^2 adjoint $\mathcal{D}_C^\#$ proves that the elements in the kernel of the latter operator are smooth and are in \mathbb{L}_* . These last remarks imply what is asserted by the last sentence in the statement of Lemma 7.15. \square

Part 3 This part of the subsection starts with a lemma that is used subsequently to say more about the behavior at large $|s|$ of the various submanifolds from any given (τ, \mathcal{C}) from \mathcal{M}^* . The lemma is also invoked to prove both Propositions 7.1 and 7.2.

To set the notation for this lemma, a domain in $C_S \cup \left(\bigcup_{\mathfrak{p} \in \Lambda} C_{\mathfrak{p}}\right)$ is said to be *semi-bounded* when the coordinate s is unbounded from above and bounded from below, or else unbounded from below and bounded from above. When U is a semibounded domain, use $\mathbb{H}_{\tau;U,loc}$ to denote the vector space whose elements are as follows: An element in \mathfrak{k} consists of a set $(\mathfrak{k}_S, \{\mathfrak{k}_{\mathfrak{p}}\}_{\mathfrak{p} \in \Lambda})$ where \mathfrak{k}_S is a locally L^2_1 section of the normal bundle of $C_S \cap U$, and where any given $\mathfrak{p} \in \Lambda$ version of $\mathfrak{k}_{\mathfrak{p}}$ is a section of the normal bundle of $C_{\mathfrak{p}} \cap U$. Moreover, with the normal bundle of C_S identified with

the bundle N_S over the $t \in [1 + z_*, 2 - z_*]$ part of S as in Part 1, the ξ_S and various $p \in \Lambda$ versions of ξ_p obey (7-32) on the boundaries of their domain of definition.

Fix $(C, \tau) \in \mathcal{M}^*$. Suppose U is a semibounded domain and that $h = (h_S, \{h_p\}_{p \in \Lambda}) \in \mathbb{H}_{\tau;U,loc}$ and that it obeys an equation of the form

$$(7-43) \quad \mathcal{D}_C h + t[h] + z = 0,$$

where z and $t[\cdot]$ are as follows: First, z is a fixed, smooth element in \mathbb{L} . Meanwhile, $t[\cdot] = (t_S, \{t_p\}_{p \in \Lambda})$ defines a map from a subspace of uniformly pointwise bounded elements in $\mathbb{H}_{\tau;U,loc}$ to locally square-integrable sections of the appropriate bundles over C_S and the various $p \in \Lambda$ versions of C_p . The element $t[\xi]$ depends explicitly on the domain coordinates and implicitly on the chosen $\xi \in \mathbb{H}_{\tau;U,loc}$. It has smooth dependence on the domain coordinates. Note that $t[\xi]$ need not be a local function of the entries of ξ . The map $\xi \rightarrow t[\xi]$ must have a certain additional requirement whose statement uses the following notation: Let W denote an open set with closure in U . Use $\|\cdot\|_W$ to denote the L^2 -norm on W , and $\|\cdot\|_{W;1}$ to denote the L^2_1 -norm on W . Finally, use $\|\cdot\|_{W;\infty}$ to denote the L^∞ -norm on W . The additional requirement involves parameters $n \geq 1$ and $r > 0$. What follows is the additional requirement:

$$(7-44) \quad \text{Suppose that } W \subset U \text{ is a given semibounded subdomain, and suppose that } h \text{ and its first derivatives are square-integrable on } W. \text{ Suppose in addition that } \|h\|_{W;\infty} \leq n^{-1}. \text{ Then } \|t[h]\|_W^2 \leq (r + n\|h\|_{W;\infty})\|h\|_{W;1}^2.$$

Assume that z and t are as just described.

Lemma 7.16 *There exists $\kappa > n$ that such that if $r < \kappa^{-1}$, then the following is true: Suppose that h is defined on a semibounded domain where it obeys (7-43). Suppose in addition that the entries of h have absolute value bounded by κ^{-1} , and that their first derivatives are bounded. Then h and its first derivatives are square-integrable on some semibounded subdomain in its original domain of definition.*

Proof Write the S -labeled component of h as η' and each $p \in \Lambda$ -labeled component as $(\varphi^{p'}, \zeta^{p'})$. This done, fix $\varepsilon > 0$ and define h_ε as follows: its S -labeled component is $e^{-\varepsilon|s|}\eta'$ and any given $p \in \Lambda$ -labeled component is $e^{-\varepsilon|s|}(\varphi^{p'}, \zeta^{p'})$. Note that h' obeys (7-32). The assertion that h is pointwise bounded implies that h_ε is square-integrable on a semibounded domain. The derivatives of h_ε are likewise square-integrable on some semibounded domain.

Let m denote the supremum norm of h . Given (7-44), the fact that h obeys (7-43) has the following implication: the restriction of h_ε to a suitable semibounded subdomain

$W \subset U$ obeys

$$(7-45) \quad \|\mathcal{D}_C h_\varepsilon\|_{\mathcal{W}}^2 \leq c_0(r + nm)\|h_\varepsilon\|_{\mathcal{W};1}^2 + c_0\|\mathfrak{h}\|^2.$$

Note that all of the integrals in (7-45) are finite. Reintroduce c and s_1 from (7-34), and suppose that $W' \subset W$ is a semibounded subdomain such that $|s| > 100s_1$ at each point, and such that each point in W' has distance at least 1 from some point in W . Then (7-34), (7-44) and (7-45) imply that

$$(7-46) \quad c^{-1}\|h_\varepsilon\|_{\mathcal{W};1}^2 \leq (c_0(r + nm) - \frac{1}{100}c^{-4})\|h_\varepsilon\|_{\mathcal{W};1}^2 + \mathfrak{w}(h) + c_0\|\mathfrak{h}\|^2,$$

where $\mathfrak{w}(h)$ is finite and ε -independent. If $r \leq c_0^{-1}c^{-4}$ and $m < c_0^{-1}n^{-1}c^{-4}$, then (7-46) supplies an ε -independent bound for $\|h_\varepsilon\|_{\mathcal{W};1}^2$. The existence of such a bound implies that h is square-integrable on an unbounded domain. \square

The next lemma states one consequence of Lemma 7.16. To set the stage for this lemma, decompose Θ_- into segments so as to define a set $\{\Gamma_S, \{\gamma_{p-}\}_{p \in \Lambda}, \{\mathfrak{o}_p\}_{p \in \Lambda}\}$ where Γ_S denotes a disjoint union of segments of integral curves of v in the $f \in [1 + z_*, 2 - z_*]$ part of M_δ , and where each $p \in \Lambda$ version of \mathfrak{o}_p is the subset of integral curves from $\{\widehat{\gamma}_p^+, \widehat{\gamma}_p^-\}_{p \in \Lambda}$. If $T \gg 1$, then $(-\infty, -T] \times \Gamma_S$ can be written as a section of the normal bundle of S over the portion where both $t \in [1 + z_*, 2 - z_*]$ and $s \leq -T$. Use η_Θ to denote this section. Meanwhile, each $p \in \Lambda$ version of $(-\infty, -T] \times \gamma_{p-}$ can be written via Ψ_p as the graph of a pair of functions, these denoted by $(\varphi_\Theta^p, \zeta_\Theta^p)$. The function ζ_Θ^p is the constant value of \widehat{h} on γ_{p-} and φ_Θ^p is the function of \widehat{u} that gives the ϕ -angle along γ_{p-} . An analogous η_Θ and set $\{(\varphi_\Theta^p, \zeta_\Theta^p)\}_{p \in \Lambda}$ are defined using Θ_+ with it understood that these are defined respectively where both $t \in [1 + z_*, 2 - z_*]$ and $s \gg 1$ on S and where $x \gg 1$ on each $p \in \Lambda$ version of $\mathbb{R} \times I_*$.

Fix next a smooth function on $(\frac{1}{2}z_*, \frac{3}{2}z_*)$ with compact support and which is equal to 1 where $z = z_*$. Use β here to denote the chosen function. Fix $p \in \Lambda$ and introduce the functions $(\varphi^{S_0}, \zeta^{S_0})$ on \mathbb{R} that appear in (7-1). This data is used to modify η_Θ and each $p \in \Lambda$ version of $(\varphi_\Theta^p, \zeta_\Theta^p)$ near the common boundary of their domains of definition. The modification requires writing η_Θ where t differs from 1 or 2 by less than z_S as a pair of constant functions, $(\varphi_\Theta^S, \zeta_\Theta^S)$, on $\mathbb{R} \times [z_*, z_S]$. Likewise, view $(\varphi_\Theta^p, \zeta_\Theta^p)$ near the boundary of $\mathbb{R} \times I_*$ as constant functions on $\mathbb{R} \times [\delta^2, z_*]$. The modification is given by the replacements

$$(7-47) \quad \begin{aligned} &\bullet (\varphi_\Theta^S, \zeta_\Theta^S) \rightarrow (1 - \beta)(\varphi_\Theta^S, \zeta_\Theta^S) + \beta(\varphi^{S_0}, \zeta^{S_0}), \\ &\bullet (\varphi_\Theta^p, \zeta_\Theta^p) \rightarrow (1 - \beta)(\varphi_\Theta^p, \zeta_\Theta^p) + \beta(\varphi^{S_0}, \zeta^{S_0}). \end{aligned}$$

Let $\hat{\eta}_\Theta$ and $\{(\hat{\varphi}_\Theta^p, \hat{\zeta}_\Theta^p)\}_{p \in \Lambda}$ denote the modified versions. These obey (7-1).

To make contact with Lemma 7.16, fix $(\tau, \mathcal{C}) \in \mathcal{M}^*$ and let η and $\{(\varphi^p, \zeta^p)\}_{p \in \Lambda}$ denote the defining data for \mathcal{C} , the former as described in Part 1 of Section 1A and each $p \in \Lambda$ version of (φ^p, ζ^p) as described in Part 2 of Section 1A. Fix $x_* > 1$ so as to be larger than the absolute value of the \mathbb{R} coordinates of the deleted $\hat{u} = 0$ points from $\mathbb{R} \times I_*$ that define the domain of the $\Delta_p > 0$ versions of (φ^p, ζ^p) .

Define $h_{\mathcal{C}} = (h_{\mathcal{C},S}, \{h_{\mathcal{C},p}\}_{p \in \Lambda})$ with $h_{\mathcal{C},S}$ the section of N_S over the $t \in [1 + z_*, 2 - z_*]$ part of S given by $\eta - \hat{\eta}_\Theta$, and with any given $p \in \Lambda$ version of h_p being the map from the domain of (φ^p, ζ^p) given by $(x, \hat{u}) \rightarrow \chi(x_* - |x| + 2)(\varphi^p - \hat{\varphi}_\Theta^p, \zeta^p - \hat{\zeta}_\Theta^p)|_{(x, \hat{u})}$. View $h_{\mathcal{C},S}$ as a section of the normal bundle of C_S and view each $p \in \Lambda$ version of $h_{\mathcal{C},p}$ as a section of the normal bundle of C_p . This done, $h_{\mathcal{C}}$ has support on two disjoint, semibounded domains, one where $s \gg 1$ and the other where $s \ll -1$. It obeys (7-32) on each. Moreover, given $r > 0$, there exists $s_r > 1$ such that $|h_{\mathcal{C}}| \leq r$ where $|s| \geq s_r$.

This $h_{\mathcal{C}}$ also obeys a version of (7-43) on an $s \gg 1$ semibounded domain, and also on an $s \ll -1$ semibounded domain. This is a consequence of (II.6-10) and (3-4).

The function $h \rightarrow t[h]$ in this case is a local function of h and its derivatives; this is to say that its value at any given point is determined by the point in question and the value of the relevant component of h and its derivatives at this same point. In particular the relevant version of $t[\cdot]$ has the schematic form

$$(7-48) \quad t[h] = t_1[h] \cdot \nabla h + t_0[h],$$

where any given component of t_1 and t_0 is a nonlinear, local, smooth function of the corresponding component of h and the relevant domain coordinates. These functions are such that:

- (7-49) • $|t_1(h)| \leq c_0 |h|$ and $|t_0(h)| \leq c_0 |h|^2$.
- The first derivatives of $t_0(\cdot)$ at $h = 0$ with respect to variations in h and the domain variables are zero, and the analogous first derivatives of t_1 and the second derivatives of t_0 are bounded by c_0 where $|h| \leq c_0^{-1}$.
 - In general, the derivatives of t_0 and t_1 with respect to variations in h and the domain variables to any given order are bounded if $|h| \leq c_0^{-1}$.

In addition, the first derivatives of $t_0(\cdot)$ at $h = 0$ are zero, and the first derivatives of t_1 and the second derivatives of t_0 are bounded by c_0 where $|h| \leq c_0^{-1}$.

Granted this background, what follows gives the first of the promised applications of Lemma 7.16.

Lemma 7.17 *Let $(\tau, \mathcal{C}) \in \mathcal{M}^*$ and define $\mathfrak{h}_{\mathcal{C}}$ as above using either Θ_- or Θ_+ . Then $\mathfrak{h}_{\mathcal{C}}$ and its derivatives to any given order are square-integrable on a semibounded domain.*

Proof The fact that $\mathfrak{h}_{\mathcal{C}}$ is square-integrable follows from Lemma 7.16. The assertion concerning its derivatives follows using standard elliptic regularity techniques with Lemma 7.7 as applied to the version of (7-43) that uses (7-48) for \mathfrak{t} . (As before, the elliptic regularity techniques from Chapter 6 in [12] will do the trick.) \square

Part 4 This part of the subsection contains the proofs of Propositions 7.1 and 7.2. These are taken in order.

Proof of Proposition 7.1 The proof has four steps.

Step 1 This step sets the groundwork for an application of the implicit function theorem. To this end, fix $(\tau, \mathcal{C}) \in \mathcal{M}^*$ and let $(\tau', \mathcal{C}') \in \mathcal{M}_{\tau}$ denote an element in a neighborhood of (τ, \mathcal{C}) with the neighborhood chosen so that τ' is very close to τ and so that each point in any submanifold from \mathcal{C}' is very close to the corresponding submanifold from \mathcal{C} and vice versa. Let $\{\eta, \{(\varphi^{\mathfrak{p}}, \zeta^{\mathfrak{p}})\}_{\mathfrak{p} \in \Lambda}\}$ denote the data that defines \mathcal{C} with η as described in Part 1 of Section 7A and with each $\mathfrak{p} \in \Lambda$ version of $(\varphi^{\mathfrak{p}}, \zeta^{\mathfrak{p}})$ as described in Part 2 of Section 7A. A corresponding data set is used to define (τ', \mathcal{C}') , but the latter is written now as $\{\eta + \eta', \{(\varphi^{\mathfrak{p}} + \varphi^{\mathfrak{p}'}, \zeta^{\mathfrak{p}} + \zeta^{\mathfrak{p}'})\}_{\mathfrak{p} \in \Lambda}$. It is a consequence of the final assertion of Lemma 7.17 that η' when viewed as a section over C_S of the latter's normal bundle is a smooth, L_1^2 section. Meanwhile, Lemma 7.17 with the final assertion of Proposition 5.1 imply that each $\mathfrak{p} \in \Lambda$ version of $(\varphi^{\mathfrak{p}'}, \zeta^{\mathfrak{p}'})$ defines a smooth, L_1^2 section of the normal bundle of $C_{\mathfrak{p}}$ in $\mathbb{R} \times \mathcal{H}_{\mathfrak{p}^*}^+$.

Use \mathfrak{h} to denote the set $\{\eta', (\varphi^{\mathfrak{p}'}, \zeta^{\mathfrak{p}'})\}_{\mathfrak{p} \in \Lambda}$. This version of \mathfrak{h} does not obey (7-32) unless $\tau = \tau'$; but if it did, then it would define an element in \mathbb{H}_{τ^*} , this a consequence of Lemma 7.17 and Proposition 5.1. In any event, \mathfrak{h} obeys a version of (7-43) with $\mathfrak{t}[\cdot]$ given by (7-48) with \mathfrak{t}_0 and \mathfrak{t}_1 described by (7-49).

The failure of (7-32) is rectified in the next step by incorporating $\tau' - \tau$ in a new definition of \mathfrak{h} and compensating with a corresponding $(\tau' - \tau)$ -dependent term added to $\mathfrak{t}[\cdot]$. The resulting version of \mathfrak{t} is not a local function of \mathfrak{h} .

Step 2 This step supplies the new definitions of \mathfrak{h} and \mathfrak{t} . To start, fix $\mathfrak{p} \in \Lambda$ and write η near the common boundary of $C_{\mathfrak{p}}$ and C_S as functions (φ^S, ζ^S) on $\mathbb{R} \times [z_*, z_S]$ in the manner of (7-1). Define η' as above and write the latter near this same boundary as a pair of functions $(\varphi^{S'}, \zeta^{S'})$ on $\mathbb{R} \times [z_*, z_S]$. View $(\varphi^{\mathfrak{p}}, \zeta^{\mathfrak{p}})$ and also the pair $(\varphi^{\mathfrak{p}'}, \zeta^{\mathfrak{p}'})$ near this boundary as functions on $\mathbb{R} \times [\delta^2, z_*]$. In this guise, the primed pairs obey the following where $z = z_*$:

$$(7-50) \quad \zeta^{S'} = \tau \zeta^{\mathfrak{p}'} + (\tau' - \tau)(\zeta^{\mathfrak{p}} + \zeta^{\mathfrak{p}'} - \zeta^{S_0}),$$

$$(7-51) \quad \varphi^{\mathfrak{p}'} = \tau \varphi^{S'} + (\tau' - \tau)(\varphi^S + \varphi^{S'} - \varphi^{S_0}).$$

To obtain something that obeys (7-32), reintroduce the function β from the proof of Lemma 7.16 and (7-47). Define pairs $(\hat{\varphi}^{S'}, \hat{\zeta}^{S'})$ and $(\hat{\varphi}^{\mathfrak{p}'}, \hat{\zeta}^{\mathfrak{p}'})$ as follows:

$$(7-52) \quad \bullet \quad \hat{\varphi}^{S'} = \varphi^{S'} \text{ and } \hat{\zeta}^{S'} = \zeta^{S'} - (\tau' - \tau)\beta(\zeta^{\mathfrak{p}} + \zeta^{\mathfrak{p}'} - \zeta^{S_0})|_{z_* - (z - z_*)} \text{ where } z > z_*.$$

$$\bullet \quad \hat{\varphi}^{\mathfrak{p}'} = \varphi^{\mathfrak{p}'} - (\tau' - \tau)\beta(\varphi^S + \varphi^{S'} - \varphi^{S_0})|_{z_* - (z - z_*)} \text{ where } z \leq z_*.$$

Extend the latter as η' and $(\varphi^{\mathfrak{p}'}, \zeta^{\mathfrak{p}'})$ over the rest of their respective domains. Use these to define the promised new version of \mathfrak{h} . The entries of this new version obey (7-32) and it follows from Lemma 7.17 that this new version is in \mathbb{H}_{τ_*} .

This new version also obeys a version of (7-43), but with a nonlocal version of $\mathfrak{t}[\cdot]$ and also with a $\mathfrak{z} \neq 0$ term proportional to $(\tau' - \tau)$. To elaborate, the nonlocal version of \mathfrak{t} is obtained from the version from Step 1 by adding a term that is supported near the boundaries of C_S and $\bigcup_{\mathfrak{p} \in \Lambda} C_{\mathfrak{p}}$. Fix a given $\mathfrak{p} \in \Lambda$. The portion of this term that lies where $z > z_*$ is the pair of functions with respective left- and right-hand components

$$(7-53) \quad (\tau' - \tau)\partial_z(\beta \zeta^{\mathfrak{p}'}|_{z_* - (z - z_*)}) \quad \text{and} \quad -(\tau' - \tau)\beta(\partial_x \zeta^{\mathfrak{p}'})|_{z_* - (z - z_*)}.$$

The part that lies where $z < z'$ has components

$$(7-54) \quad -(\tau' - \tau)\beta(\partial_x \varphi^{S'})|_{z_* - (z - z_*)} \quad \text{and} \quad (\tau' - \tau)\partial_z(\beta \varphi^{S'}|_{z_* - (z - z_*)}).$$

Meanwhile, the \mathfrak{z} term has support near these same boundaries, where it is given by replacing $\zeta^{\mathfrak{p}'}$ with $(\zeta^{\mathfrak{p}} - \zeta^{S_0})$ in (7-53) and $\varphi^{S'}$ with $(\varphi^S - \varphi^{S_0})$ in (7-54).

Step 3 Given the preceding definition of \mathfrak{t} , it follows that the left-hand side of the corresponding version of (7-43) defines a smooth map to \mathbb{L}_* from the product of a ball about the origin in \mathbb{H}_{τ_*} with an interval centered on $\tau \in [0, 1]$. Use \mathfrak{F} to denote this map. The inverse function theorem finds a ball, \mathcal{B} , about the origin in the kernel

of \mathcal{D}_C , an interval, I , centered on τ in $[0, 1]$, and a smooth map, $\mathfrak{b}: \mathcal{B} \times I \rightarrow \mathbb{H}_{\tau_*}$, and these are such that:

- (7-55) • $\mathfrak{b}(0, \tau) = 0$ and $\nabla^{\mathbb{H}} \mathfrak{b}|_{(0, \tau)} = 0$.
- Suppose that $(\mathfrak{h}_0, \tau') \in \mathcal{B} \times I$. Then $\mathfrak{h}_0 + \mathfrak{b}(\mathfrak{h}_0, \tau') \in \mathbb{B}$ and we have $(1 - \Pi)\mathfrak{F}(\mathfrak{h}_0 + \mathfrak{b}(\mathfrak{h}_0, \tau'), \tau') = 0$.
 - Let $\mathbb{B}' \subset \mathbb{B}$ denote the concentric, half-radius ball. Suppose that $(\mathfrak{h}, \tau') \in \mathbb{B}' \times I$ and suppose that $(1 - \Pi)\mathfrak{F}(\mathfrak{h}', \tau') = 0$, then $\mathfrak{h} = \mathfrak{h}_0 + \mathfrak{b}(\mathfrak{h}_0, \tau)$ with $\mathfrak{h}_0 \in \mathfrak{B}$.

A map, \mathfrak{f} , from $\mathcal{B} \times I$ to the cokernel of \mathcal{D}_C is defined now by the rule

(7-56)
$$(\mathfrak{h}_0, \tau) \rightarrow \mathfrak{f}((\mathfrak{h}_0, \tau')) = \Pi\mathfrak{F}(\mathfrak{h}_0 + \mathfrak{b}(\mathfrak{h}_0, \tau'), \tau').$$

The map from $\mathcal{B} \times I$ to $\mathbb{B} \times I$ given by $(\mathfrak{h}_0, \tau') \rightarrow (\mathfrak{h}_0 + \mathfrak{b}(\mathfrak{h}_0, \tau'), \tau')$ embeds $\mathfrak{f}^{-1}(0) \subset \mathcal{B}$ homeomorphically onto an open set in $\mathbb{B} \times I$ of solutions to (7-43) that contains $\mathbb{B}' \times I$.

This map \mathfrak{f} is the map required by Proposition 7.1, and the embedding just described gives the homeomorphism Φ .

Step 4 The claim made by the second bullet of Proposition 7.1 that \mathcal{M}^* is smooth where \mathfrak{f} is a submersion and the claim that π_I is a smooth on this same set are standard consequences of the inverse function theorem as used in Step 3. The proof that p is continuous follows from Proposition 5.1 and Lemma 5.2.

To see about the derivatives of p , let $\mathfrak{p} \in \Lambda$ be such that $\Delta_{\mathfrak{p}} > 0$. For the sake of argument, suppose that $\mathcal{E} \subset C_{\mathfrak{p}}$ is an end where the $s \gg 1$ portion maps to the tubular neighborhood $U_+ \subset \mathcal{H}_{\mathfrak{p}}$ of the integral curve $\widehat{\gamma}_{\mathfrak{p}}^+$ via the projection from $\mathbb{R} \times \mathcal{H}_{\mathfrak{p}^*}^+$. Use the coordinates $(s_+, \phi_+, \theta_+, u_+)$ for $\mathbb{R} \times U_+$ and parametrize the $s \gg 1$ part of \mathcal{E} as in Proposition 5.1. It is a consequence of Proposition 5.1 that the operator $D_{\mathfrak{p}}$ on \mathcal{E} appears as an operator on the space of \mathbb{R}^2 -valued functions of the coordinates (s_+, ϕ_+) . An essentially verbatim repeat of the proof of the first bullet of Lemma 5.8 proves that the latter has the form $\mathfrak{D}_0 + \mathfrak{d}$, where \mathfrak{d} is a first-order differential operator on the space of maps from to \mathbb{R}^2 whose coefficients are bounded in absolute value by $c_0 e^{-s/c_0}$.

Now let $\mathfrak{h}_0 \in \mathcal{B} \subset \ker(\mathcal{D}_C)$ denote a given element, and let r denote its norm. The $C_{\mathfrak{p}}$ component of \mathfrak{h}_0 appears using this parametrization as a square-integrable map, $\eta_{\mathfrak{p}}$, from the very large s_+ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 that obeys an equation that has the

schematic form $\mathcal{D}_0\eta_p + \partial\eta_p = 0$. Granted that this is so, then the techniques used in Section 2.3 in [7] can be employed to prove that η_p can be written as in

$$(7-57) \quad \eta_p = r((c e^{-\lambda_1 s_+}, 0) + \epsilon),$$

where $c \in \mathbb{R}$ and where $|\epsilon| \leq c_0 e^{-(\lambda_1 + 1/c_0)s_+}$. Meanwhile, (7-43), (7-48) and (7-49) and the techniques from Section 2.3 in [7] can be employed in a straightforward manner to see that $\mathfrak{b} = \mathfrak{b}(\mathfrak{h}_0, \tau')$ from (7-55) can be written on \mathcal{E} as a square-integrable map from the very large s_+ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 that is bounded by $c_0(r^2 + \tau'^2)e^{-\lambda_1 s_+}$. Granted Lemma 5.2, this last observation implies that p is a C^1 map on $\mathcal{M}_{\text{smooth}}^*$. The proof that p has derivatives to any given order has a similar flavor and is omitted.

The assertions of the third bullet follows from (7-56) because the latter equation depicts a version of f with the property that $|f(\mathfrak{h}_0)| \leq c_0 |\mathfrak{h}_0|^2$. It follows as a consequence that f is a submersion at $(0, \tau)$ only if $\dim(\text{cokernel}(\mathcal{D}_C) \leq 1$, and that if f is a submersion at $(0, \tau) \in \mathcal{B} \times I$ and $d\pi_I \neq 0$ at $(0, \tau)$, then $\text{cokernel}(\mathcal{D}_C) = 0$.

The proof of the fourth bullet starts with Propositions 2.1 and 2.2 for they jointly asserted that $\pi_1^{-1}(0)$ is mapped diffeomorphically by the map p to $\times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$ given bounds for z_* and δ that are purely S -dependent (or \mathcal{K} -compatible). Meanwhile, Propositions 2.1, 3.2 and 5.7 imply that the \mathcal{D}_C along $\pi_1^{-1}(0)$ has trivial cokernel. This fact with (7-56) implies what is asserted by the fourth bullet. □

Proof of Proposition 7.2 The assertion made by the first bullet is proved in Section 3 of [10]. To prove item (a) of the second bullet, note first that given Propositions 3.2 and 5.7, it is enough to consider the case of the part of \mathcal{M}^* where $\tau > 0$. With this restriction on τ understood, the allowed variations of the almost complex structure on the portion of $\mathbb{R} \times Y$ in the $f^{-1}([1 + \delta_*^2, 2 - \delta_*^2])$ part of $\mathbb{R} \times M_\delta$ constitute a sufficiently large set for applying standard Smale–Sard arguments as done in Section 3 of [10]. In particular, straightforward modifications to the arguments from this same section of [10] prove what is asserted by this item.

By way of a parenthetical remark, note that the variations in J that are allowed on any $\cup_{p \in \Lambda} \mathbb{R} \times \mathcal{H}_p$ may not form a set that is large enough to invoke the Smale–Sard theorem. This is because the almost complex structures here are constrained to be invariant with respect to both the group of constant translations along the \mathbb{R} factor and the group of constant rotations of the angle ϕ . The set of allowed variations of J on the $f \in (1 + \delta_*^2, 2 - \delta_*^2)$ part of $\mathbb{R} \times M_\delta$ is sufficiently large precisely because Lipschitz allows almost complex structures on $\mathbb{R} \times [1, 2] \times \Sigma$ that depend on the coordinate

$t \in [1, 2]$. In fact, the set of t -independent almost complex structures is likely too small for the applications in Section 3 of [10]. The assertion made by item (b) of the second bullet is proved using the Smale–Sard theorem using the aforementioned arguments from [10]. The details are also straightforward and also omitted. \square

8 Counting ech-HF submanifolds

The section starts with an existence assertion for ech-HF submanifolds, and then a sort of uniqueness assertion. The existence result is stated as Proposition 8.1 and the uniqueness result is stated by Proposition 8.2. These two propositions are in Section 8A. Sections 8B–8F explain, among other things, how to count the ech-HF submanifolds that are provided by Proposition 8.1.

8A Existence and uniqueness of ech-HF subvarieties

Fix a countable set in $(\times_3(0, 1)) \times (1, \infty)$ of possible choices for the data (z_*, δ, χ_0, R) and then choose J_{HF} from Proposition 7.2’s residual set.

To set the stage for the upcoming Propositions 8.1 and 8.2, fix a finite or weakly compact subset $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$. With \mathcal{K} in hand, select the data set (z_*, δ, χ_0, R) and the almost complex structure J as described by Propositions 7.1–7.3 so that their conclusions can be assumed.

Proposition 8.1 assumes implicitly that a submanifold S has been chosen from \mathcal{K} and choice of $(\widehat{\Theta}_-, \widehat{\Theta}_+) \in \widehat{\mathcal{Z}}^S$ has been made so as to define the corresponding version of \mathcal{M}^* . By way of a reminder, \mathcal{M}^* is a smooth manifold with boundary and $\pi_1 \times p$ is a smooth, proper map. Supposing that $y \in \times_{p \in \Lambda} (\times_{\Delta_p} \mathbb{R})$ is a regular value of the map p , use \mathcal{M}_y^* to denote $p^{-1}(y)$. The latter is a smooth, 1-dimensional manifold with boundary. The fourth bullet of Proposition 7.1 asserts that there is one and only one component with a boundary point on $\pi_1^{-1}(0)$. Meanwhile, it follows from what is said in Proposition 7.2 that the differential of π_I at each point in $\pi_I^{-1}(1)$ is surjective, and it follows from Propositions 7.2 and 7.3 that there is at most a finite set of points in $\pi_I^{-1}(1)$.

The manifold \mathcal{M}_y^* is orientable because this is the case for any 1-manifold. Orientations for the components of \mathcal{M}_y^* are defined by requiring that π_I be orientation-preserving where it is increasing and orientation-reversing where it is decreasing.

Assign to any $\pi_1^{-1}(1)$ point in \mathcal{M}_y^* the weight $+1$ if the differential of π_I at the point is orientation-preserving, and assign -1 if not.

Proposition 8.1 *The set $\pi_1^{-1}(1)$ in \mathcal{M}_y^* is nonempty. Moreover, the sum of the ± 1 weights of these elements is equal to 1.*

Given that the $\pi_1^{-1}(1)$ points in \mathcal{M}_y^* are ech-HF submanifolds, this proposition supplies an existence theorem for ech-HF submanifolds.

Proof of Proposition 8.1 If z_* and δ are chosen small, then Proposition 7.3 asserts that \mathcal{M}_y^* is a compact, oriented 1-manifold with boundary. Each component has either 0 or 2 boundary points. There is one component with a boundary point where $\pi_I = 0$ and the latter must have a second boundary point, thus where $\pi_I = 1$. As π_I is in no case greater than 1 and as its differential is nonzero at this point, this point has weight $+1$. There is a finite set of other components. Those with boundary points must have both boundary points where $\pi_I = 1$. As the differential of π_I is nonzero at both points, one must have weight $+1$ and the other weight -1 . Granted this accounting, the sum of the weights of the elements in \mathcal{M}_y^* is equal to 1. □

Proposition 8.2 uses \mathcal{K} to denote the following subset of \mathcal{A}_{HF} : a Lipschitz submanifold S is in \mathcal{K} if and only if the operator D_S has Fredholm index no greater than 1. The quotient space \mathcal{K}/\mathbb{R} is finite, this a consequence of Lemma 5.4 and Corollary 7.2 in [10]. To say more about notation, suppose that $(\hat{\Theta}', \hat{\Theta})$ is a chosen pair from $\times_2 \mathcal{Z}_{\text{ech}, M}$. Given a data set (δ, χ_0, R) to define the geometry of Y , and given an almost complex structure, J , subject to the constraints in Part 1 of Section 1C, the proposition refers to the space $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ defined in Part 2 of Section 1C. Any given element in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ will have a union of components that form an ech-HF submanifold. This ech-HF submanifold part is either \mathbb{R} -invariant or not. If the ech-HF submanifold part is \mathbb{R} -invariant, then this element must have a single component from some $p \in \Lambda$ version of Proposition II.3.4's moduli spaces \mathcal{M}_{p-} and \mathcal{M}_{p+} . The element can also contain \mathbb{R} -invariant cylinder components from the set $\bigcup_{p \in \Lambda} \{\mathbb{R} \times \hat{\gamma}_p^+, \mathbb{R} \times \hat{\gamma}_p^-\}$. If the ech-HF submanifold is not \mathbb{R} -invariant, then it sits in some $(\hat{\Theta}_-, \hat{\Theta}_+) \in \times_2 \mathcal{Z}_{\text{ech}, M}$ version of $\mathcal{M}_1(\hat{\Theta}_-, \hat{\Theta}_+)$ with $(\hat{\Theta}_-, \hat{\Theta}_+)$ as described in (2-2). The preceding observations about $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ follow from what is said in Propositions II.3.1–II.3.4 about ech indices.

Proposition 8.2 *Fix the data (z_*, δ, χ_0, R) and J as described by Propositions 7.1–7.3 with a suitably large choice for their respective versions of κ and κ_* , and with any*

choice of Lipschitz submanifold from \mathcal{K} . Fix $(\hat{\Theta}_-, \hat{\Theta}_+) \in \times_2 \mathcal{Z}_{\text{ech}, M}$ that obey (2-2) and are such that $\mathcal{M}_1(\hat{\Theta}_-, \hat{\Theta}_+) \neq \emptyset$. Then $\sum_{p \in \Lambda} \Delta_p \leq 1$ and what follows is true: Let C denote a given ech-HF submanifold from $\mathcal{M}_1(\hat{\Theta}_-, \hat{\Theta}_+)$. There exists a unique $S \in \mathcal{K}$ such that $(\hat{\Theta}_-, \hat{\Theta}_+) \in \hat{\mathcal{Z}}^S$, and there exists a unique $(1, \mathcal{C})$ in the corresponding version of \mathcal{M}^* whose respective components are C 's intersections with the $f^{-1}(1 + z_*, 2 - z_*)$ portion of $\mathbb{R} \times Y$ and the various $p \in \Lambda$ versions of $\mathbb{R} \times \mathcal{H}_{p_*}^+$.

Proof Delete the assertion $\sum_{p \in \Lambda} \Delta_p \leq 1$ from Proposition 8.2, and suppose that the resulting weaker proposition is false. If this is the case, then there is a sequence $\{D_n, C_n\}_{n=1,2,\dots}$ whose constituents will now be described. First, what is denoted by D_n is a data set that can be written as $((\hat{\Theta}_{n-}, \hat{\Theta}_{n+}), (z_{*n}, \delta_n, \chi_{0n}, R_n, J_n))$ where $z_{*n} < \frac{1}{n}$ and $\delta_n < n^{-2} z_{*n}$. The latter with χ_{0n} and R_n are suitable for defining the geometry of Y . Meanwhile, J_n is an almost complex structure on the $(\delta_n, \chi_{0n}, R_n)$ version of $\mathbb{R} \times Y$ as described in Section 1C. In addition J_n with $(z_{*n}, \delta_n, \chi_{0n}, R_n)$ are such that Propositions 7.1–7.3 can be invoked using any Lipschitz submanifold from \mathcal{K} . Meanwhile, $\hat{\Theta}_{n-}$ and $\hat{\Theta}_{n+}$ are elements in the index n version of $\hat{\mathcal{Z}}_{\text{ech}, M}$ that are defined in part by respective HF-cycles that can be assumed to be independent of the index n . The $\hat{\Theta}_{n-}$ and also $\hat{\Theta}_{n+}$ elements from $\{\hat{\gamma}_p^+ \hat{\gamma}_p^-\}_{p \in \Lambda}$ are also independent of n . What is denoted by C_n signifies an ech-HF submanifold defined by the data D_n from a submanifold in $\mathcal{M}_1(\hat{\Theta}_{n-}, \hat{\Theta}_{n+})$ but C_n is not from some $\pi_1^{-1}(1)$ element in the index n version of \mathcal{M}^* . No generality is lost by assuming that C_n 's version of the set $\{\Delta_p\}_{p \in \Lambda}$ is independent of n . The twelve steps that follow derive nonsense with such a sequence.

The assertion that $\sum_{p \in \Lambda} \Delta_p \leq 1$ follows from Proposition 7.1 if $\mathcal{M}_1(\hat{\Theta}_-, \hat{\Theta}_+)$ has an ech-HF submanifold from the corresponding version of \mathcal{M}^* .

Step 1 The submanifold C_n has a normal bundle which also inherits a holomorphic line bundle structure. Use N_{C_n} to denote this bundle. There is an associated first-order operator that maps sections of N_{C_n} to sections of $N_{C_n} \otimes T^{0,1}C_n$. This operator also has the form depicted on the right-hand side of (1-25). Use \hat{D}_{C_n} to denote this operator. This is Fredholm when mapping the L_1^2 space of sections N_{C_n} to the L^2 space of sections of $N_C \otimes T^{0,1}C$. These respective domain and range spaces are denoted by \mathbb{H}_1 and \mathbb{L} in what follows. Note that this Fredholm incarnation of \hat{D}_{C_n} has index 1 and trivial cokernel.

Step 2 What follows is a consequence of the assumption that $\{z_{*n}\}_{n=1,2,\dots}$ has limit zero: Given ε , the conclusions of Proposition II.7.2 can be invoked for all sufficiently

large n versions of C_n . This being the case, fix some small $z_* > 0$. For n large, z_{*n} will be less than z_* . For such n , use $\mathcal{H}_{p_*n}^+$ to denote the version of $\mathcal{H}_{p_*}^+$ that is defined using z_* and the data set $(\delta_n, \chi_{0n}, R_n)$. Use Ψ_{pn} to denote the version of the map Ψ_p that is defined using z_* and the data set $(\delta_n, \chi_{0n}, R_n, J_n)$.

Proposition II.5.8 and Lemma II.4.7 can be invoked when n is large to conclude the following: Fix $p \in \Lambda$ and let C_{pn} denote the intersection between C_n and $\mathbb{R} \times \mathcal{H}_{p_*n}^+$. This is a smooth, properly embedded submanifold with boundary. Moreover, it is the image via the map Ψ_{pn} of a graph in the $(z_*, \delta_n, \chi_{0n}, R_n, J_n)$ version of $\mathbb{R} \times \mathcal{X}$ which is defined by a map to \mathbb{R}^2 from the complement of Δ_p points in $\mathbb{R} \times I_*$ where $\hat{u} = 0$. This map has the form

$$(8-1) \quad (x, \hat{u}) \rightarrow (x, \hat{u}, \hat{\phi} = \varphi^{pn}(x, \hat{u}), h = \zeta^{pn}(x, \hat{u})).$$

Note that C_{pn} obeys all of the requirements listed in Part 2 of Section 7A. Therefore, the fact that C_n is not from the index n version of \mathcal{M}^* is not due to properties of its intersection with $\mathbb{R} \times \mathcal{H}_{p_*}^+$.

Step 3 With n assumed large, introduce C_{Sn} to denote C_n 's intersection with the $t \in [1 + z_*, 2 - z_*]$ part of $\mathbb{R} \times M_{\delta_n}$. When viewed in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$, this C_{Sn} is a smooth, properly embedded submanifold with boundary. A neighborhood of each component of the boundary of C_{Sn} can be depicted as a graph of the sort that is described by Property 4 in Section 7A with the constant z_S being n -dependent now. Use $(\varphi^{Sn}, \zeta^{Sn})$ to denote C_{Sn} 's version of the functions (φ^S, ζ^S) that appear in this Property 4. Note in particular that (7-1) holds by virtue of the fact that C_{Sn} attaches seamlessly along its boundaries with the boundaries of $\bigcup_{p \in \Lambda} C_{pn}$ to give the surface C_n . It follows from this last remark that C_n obeys the requirements from Part 3 in Section 7A for membership in the index n version of \mathcal{M}_* .

Granted this, and granted what is said in the final paragraph of Step 2, C_n 's lack of membership in the index n version of \mathcal{M}^* must be due to some property of C_{Sn} . To see what this might be, view C_{Sn} as a submanifold in $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$, and suppose that there exists a Lipschitz submanifold, $S \in \mathcal{K}$ such that Property 1 in Part 1 of Section 7A is obeyed. If this is the case, then Proposition II.7.3 supplies a \mathcal{K} -compatible, and, in particular, n -independent constant $c > 1$ with the following significance: If $z_* < c^{-1}$, then both Properties 1 and 2 in Part 1 of Section 7A are obeyed with some perhaps different choice for S from \mathcal{K} . In any event, Property 1

implies Property 3 in Part 1 of Section 7A, and Property 4 in Part 1 of Section 7A follows if $z_* < c^{-1}$ as well.

These last observations lead to the following conclusion:

(8-2) *There exists $\varepsilon_* > 0$ with the following significance: if $z_* < \varepsilon_*$, then no sufficiently large n version C_{S_n} lies entirely in the radius ε_* tubular neighborhood of any submanifold from \mathcal{K} .*

This assertion leads to the desired nonsense as it is proved false in the upcoming Step 6.

Step 4 The restriction of N_{C_n} to C_{S_n} is denoted by N_{S_n} and the restriction to C_{pn} is denoted by N_{pn} . Let η denote a given, smooth section of the bundle N_{C_n} over C_n . The restriction of the section η to C_{S_n} and to each $p \in \Lambda$ version of C_{pn} defines the $G+1$ -tuple $(\eta_S, \{\eta_p\}_{p \in \Lambda})$ with η_S denoting a section of N_{S_n} and with each $p \in \Lambda$ version of C_{pn} denoting a section of N_{pn} . Taking this view of $C^\infty(C_n; N_{C_n})$ leads to the equivalent definition of the space \mathbb{H} as the L^2_1 space of sections of N_{C_n} given in the next paragraph.

The space \mathbb{H} is the completion of a subspace of

$$C^\infty(C_{S_n}, N_{S_n}) \oplus \left(\bigoplus_{p \in \Lambda} C^\infty(C_{pn}, N_{pn}) \right).$$

The subspace consists of elements with compact support and with boundary values as follows: Let $(\eta_S, \{\eta_p\}_{p \in \Lambda})$ denote an element in the subspace. Given $p \in \Lambda$, view a neighborhood of C_{S_n} near a given critical point from p as in Step 3. With this view understood, write η_S on the corresponding $z = z_*$ boundary as a pair of functions of x , these denoted by $(\varphi^{S'}, \zeta^{S'})$. Meanwhile, write η_p on the contiguous boundary of C_{pn} as a pair of functions of x , these denoted by $(\varphi^{p'}, \zeta^{p'})$. Then $\zeta^{S'} = \zeta^{p'}$ and $\varphi^{p'} = \zeta^{S'}$. The relevant completion of this subspace is defined by the respective L^2_1 -norms on the spaces of compactly supported sections of N_{S_n} and each $p \in \Lambda$ version of C_{pn} .

The range space \mathbb{L} for \widehat{D}_{C_n} can be viewed as the completion of the space of compactly supported sections of $C^\infty(C_{S_n}, N_{S_n} \otimes T^{0,1}S_n) \oplus (\bigoplus_{p \in \Lambda} C^\infty(C_{pn}, N_{pn} \otimes T^{0,1}C_{pn}))$ using the norm that is defined by the respective L^2 -norms for each summand.

Use D_{S_n} to denote the restriction of \widehat{D}_{C_n} to C_{S_n} . Fix $p \in \Lambda$ and view C_{pn} as the graph of $(\varphi^{pn}, \zeta^{pn})$. The restriction of the operator \widehat{D}_{C_n} to C_{pn} is given by the $\mathfrak{h} = (\varphi^{pn}, \zeta^{pn})$ version of (3-9) with the functions a_1, a_2 and b defined by the index n data set. The latter incarnation of \widehat{D}_{C_n} is denoted by D_{pn} in what follows.

Step 5 For each $\tau \in [0, 1]$, use \mathbb{H}_τ now to denote the Banach space that is obtained by the L^2_1 completion of the subspace of compactly supported sections in

$$C^\infty(C_{S_n}, N_{S_n} \otimes T^{0,1} S_n) \oplus \left(\bigoplus_{p \in \Lambda} C^\infty(C_{pn}, N_{pn} \otimes T^{0,1} C_{pn}) \right)$$

whose boundary values obey

$$(8-3) \quad \zeta^{S'} = \tau \zeta^{p'} \quad \text{and} \quad \varphi^{p'} = \tau \varphi^{S'}.$$

The operator $\mathcal{D}_{C_n} = (D_{S_n}, \{D_{pn}\}_{p \in \Lambda})$ acts as a bounded operator from each $\tau \in [0, 1]$ version of \mathbb{H}_τ to \mathbb{L} . The arguments for Lemma 7.15 can be used with almost no changes to prove that \mathcal{D}_{C_n} defines a Fredholm operator from each $\tau \in [0, 1]$ version of \mathbb{H}_τ to \mathbb{L} and that the index of each such Fredholm incarnation of \mathcal{D}_{C_n} is equal to the Fredholm index of the $\tau = 1$ version, this being 1 since the $\tau = 1$ version is $\widehat{\mathcal{D}}_{C_n}$.

Consider now the $\tau = 0$ version. The latter is a direct sum of $G + 1$ Fredholm operators. The first of these is D_{S_n} acting on the L^2_1 completion of the subspace of compactly supported sections of N_{S_n} that obey the following boundary condition: Write a given section on a given boundary component as a pair of functions, $(\varphi^{S'}, \zeta^{S'})$. Then the boundary condition asserts only that $\zeta^{S'} = 0$. Note in particular that this boundary condition makes no reference to any $p \in \Lambda$. The range space for this Fredholm operator is the L^2 completion of the space of compactly supported sections of $N_{S_n} \otimes T^{0,1} C_{S_n}$. Use $\text{index}(D_{S_n})$ in what follows to denote the Fredholm index of this Fredholm incarnation of D_{S_n} .

Meanwhile, each $p \in \Lambda$ labels an operator in the aforementioned direct sum. The latter is D_{pn} acting on the L^2_1 completion of the space of compactly supported sections of N_{pn} whose boundary values are as follows: Write a section on the a boundary component as $(\varphi^{p'}, \zeta^{p'})$. Then $\varphi^{p'} = 0$. Note that this condition makes no reference to S or to the other elements in Λ . The range space for this Fredholm incarnation of D_{pn} is the L^2 completion of the space of compactly supported sections of $N_{pn} \otimes T^{0,1} C_{pn}$. This incarnation of D_{pn} is described by Proposition 5.7; it has trivial cokernel and kernel dimension equal to Δ_p .

What was said in the three previous paragraphs gives

$$(8-4) \quad \text{index}(D_{S_n}) + \sum_{p \in \Lambda} \Delta_p = 1.$$

This understood, it follows that $\text{index}(D_{S_n}) \leq 1$ and that $\text{index}(D_{S_n}) \leq 0$ if any $p \in \Lambda$ version of Δ_p is positive.

Step 6 Invoke Proposition II.7.2 and use Corollary 7.2 in [10] to find a subsequence of $\{(D_n, C_n)\}_{n=1,2,\dots}$ (hence renumbered consecutively from 1) and a sequence $\{\varepsilon_n\}_{n=1,2,\dots} \in (0, 1)$ with the following three properties: The latter subsequence is decreasing and converges in $[0, 1]$ to 0. Second, Proposition II.7.2 applies to each C_n with $\varepsilon = \varepsilon_n$. Third, the various index n versions of the relevant broken, singular admissible sets are identical. Let Ξ denote this set. The lemma that follows uses Ξ to compute the number $\text{index}(D_{S_n})$.

The notation used by the lemma writes a given element in Ξ as $((S, u), \vartheta_\Sigma)$, where (S, u) denotes a Lipschitz subvariety and ϑ_Σ denotes a finite set of constant (s, t) slices of $\mathbb{R} \times (1, 2) \times \Sigma$. A given slice can appear more than once in ϑ_Σ . The lemma uses n_Σ to denote the number of elements in ϑ_Σ . As noted in the paragraph prior to Proposition 7.2, the subvariety pair (S, u) has an associated Fredholm operator, this denoted by D_S .

Lemma 8.3 *There exists an n -independent constant $\kappa > 1$ such that if $z_* < \kappa^{-1}$ and if n is sufficiently large, then $\text{index}(D_{S_n}) = \sum_{((S,u),\vartheta_\Sigma) \in \Xi} (\text{index}(D_S) + 2n_\Sigma)$.*

This lemma is proved in a moment. Accept it for now.

What with (8-4), this lemma implies that either Ξ has just one component, and the latter has $\vartheta_\Sigma = \emptyset$, or else there exists a non- \mathbb{R} -invariant Lipschitz submanifold, S , with $\text{index}(D_S) \leq 0$. As this is precluded by Proposition 7.2, it follows as a consequence that Ξ has but a single element with $\vartheta_\Sigma = \emptyset$.

Granted this last conclusion, invoke the second bullet of Proposition 7.3 to see that all sufficiently large n versions of C_n violate what is asserted in (8-2). This observation constitutes the desired nonsense.

Step 7 This step and Steps 8–10 contain the:

Proof of Lemma 8.3 Fix $Z = ((S, u), \vartheta_\Sigma) \in \Xi$ and let $|Z|$ denote the union of $u(S)$ with the curves from ϑ_Σ . Use the data given in Proposition II.7.2 to obtain a subsequence of $\{C_n\}_{n=1,2,\dots}$, hence renumbered from 1, and a sequence $\{s_n\}_{n=1,2,\dots}$ with the following property: Fix $n \in \{1, 2, \dots\}$ and use X_n to denote $[-4n, 4n] \times [1 + \frac{1}{n}, 2 - \frac{1}{n}] \times \Sigma$. View X_n for the moment as a subset of $\mathbb{R} \times M_{\delta_n}$. Translate the surface C_n by s_n along the \mathbb{R} factor of $\mathbb{R} \times [1, 2] \times \Sigma$ so that each point in $C_n \cap X_n$ has distance at most $\frac{1}{n}$ from some point in $|Z| \cap X_n$, and vice versa. Furthermore, if μ is any 2-form on X_n with $|\mu| \leq 1$ and $|\nabla \mu| < n$, then the integral of μ over $C_n \cap X_n$ differs from $\int_{u(S) \cap X_n} \mu + \sum_{\Sigma' \in \vartheta_\Sigma} \int_{\Sigma'} \mu$ by less than $\frac{1}{n}$.

There exists, by assumption, a constant $z_S > 0$ such that there are G components of the $t \in [1, 1 + z_S]$ part of S and each is described by Property 5 in Part 1 of Section 1G. The $t \in [2 - z_S, 2]$ part of S consists of an analogous set of G components. Meanwhile, there exists $s_S \geq 1$ such that the $|s| > s_S$ part of S is described by Property 6 in Part 1 of Section 1G. Assume that $z_* < 10^{-4} z_S$. Then the various components of the s_n -translate of C_{S_n} where $t \in [1 + z_*, 1 + \frac{1}{2} z_S]$, where $t \in [2 - \frac{1}{2} z_S, 2 - z_*]$ and where $|s| \in [2s_S, 2n]$ obey the conclusions of Property 1 of Part 1 in Section 7A.

Granted what was just said, the arguments in Section 4 of [10] will write $\text{index}(D_{S_n})$ as a sum of various contributions that can be readily identified with the terms in Lemma 8.3's sum. To do this, focus again on a given $((S, u), \vartheta_\Sigma) \in \Xi$. Truncate the s_n -translate of C_{S_n} on the slices where $|s| = 3s_n$. The result, when n is large, has $2G$ constant s boundary arcs that run from the $t = 1 + z_*$ boundary to the $t = 2 - z_*$ boundary of $\mathbb{R} \times [1 + z_*, 2 - z_*] \times (T_- \cap T_+)$. Attach to each such arc a properly embedded, infinite strip that is a graph over either the $s \leq -3s_n$ or $s \geq 3s_n$ part of $\mathbb{R} \times [1 + z_*, 2 - z_*]$ of a smooth map to $T_- \cap T_+$ that converges as $|s| \rightarrow \infty$ at an exponential rate in $|s|$ to the nearby $C_- \cap C_+$ point. Let Z_n denote the resulting properly embedded submanifold with boundary in the manifold with boundary $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$. This submanifold has a corresponding version of the operator in (1-25) which is Fredholm when viewed as a linear map between the Z_n analogs of the Banach space domain and range spaces that were defined for D_{S_n} . The associated Fredholm index is denoted in what follows by $\text{index}(D_{Z_n})$.

Standard gluing theorems can be used to prove that $\text{index}(D_{S_n})$ is the sum of the various $Z \in \Xi$ versions of $\text{index}(D_{Z_n})$ when n is large. See for example, Lemma 9.6 in [7] for a statement in an analogous context but where the operator is defined on a manifold without boundary. The corresponding lemma for the case at hand is proved using arguments that differ only cosmetically. (These gluing theorems are geometric expressions of the excision property that is obeyed by the index of Fredholm elliptic operators on manifolds.)

The next steps explains why $\text{index}(D_{Z_n}) = \text{index}(D_S) + 2n_\Sigma$ when n is large.

Step 8 Because the $t \in [1 + z_*, 1 + z_S] \cup [2 - z_S, 2 - z_*]$ part of Z_n is a graph over the analogous part of S which is very close to S when $z_* \ll 1$, the arguments from Section 4 of [10] and the erratum of [10] can be applied directly to Z_n to prove the equality $\text{index}(D_{Z_n}) = \text{index}(D_S) + 2n_\Sigma$. To elaborate on this, first identify $\mathbb{R} \times [1 + z_*, 2 - z_*] \times \Sigma$ with $\mathbb{R} \times [1, 2] \times \Sigma$ by the diffeomorphism from $[1 + z_*, 2 - z_*]$

to $[1, 2]$ that maps t to $\frac{t-3z_*}{1-2z_*}$. Use this diffeomorphism to identify Z_n with a properly embedded submanifold with boundary in $\mathbb{R} \times [1, 2] \times \Sigma$. This submanifold can be written as a graph over S very near $\mathbb{R} \times \{1\} \times \Sigma$ and very near $\mathbb{R} \times \{2\} \times \Sigma$. As a consequence, it can be isotoped very near these boundaries so that resulting submanifold (to be denoted by Z'_n) agrees with S near the boundary components $\mathbb{R} \times \{1\} \times \Sigma$ and $\mathbb{R} \times \{2\} \times \Sigma$. This submanifold Z'_n has a corresponding $D_{Z'_n}$ when n is large that is Fredholm. Its domain is defined by the same boundary conditions that define D_S , and its index is the same as that of D_{Z_n} . (This is because it can be readily homotoped to D_{Z_n} through a 1-parameter family of first-order Fredholm operators.)

The surface Z'_n defines a class (to be denoted by A') in what Lipshitz denotes by $\pi(\vec{x}, \vec{y})$. The notation from [10] has \vec{x} signifying a vector with G components, the components being the points in $C_+ \cap C_-$ that label the negative points of $\bar{S} \setminus S$. Meanwhile, \vec{y} is likewise a vector with G components, these being the points in $C_+ \cap C_-$ that label the positive points of $\bar{S} \setminus S$. (The closure \bar{S} of S is described in Section 1G.) The class A' is what Lipshitz calls a positive class. In the present context, this means that Z'_n has only positive intersection numbers with a certain (finite) set of holomorphic submanifolds in $\mathbb{R} \times (1, 2) \times \Sigma$. These submanifolds all have the form $\mathbb{R} \times [1, 2] \times \{z\}$ with z coming from a certain finite set in $\Sigma \setminus (T_- \cup T_+)$, and the intersections are positive because Z_n is holomorphic near the intersection points.

Lipshitz gives a formula for $\text{index}(D_{Z'_n})$ at the beginning of the erratum of [10]:

$$(8-5) \quad \text{index}(D_{Z'_n}) = G - \chi(Z'_n) + 2e(A'),$$

with $\chi(Z'_n)$ being the Euler characteristic of Z'_n and with $e(A')$ being what Lipshitz calls the *Euler measure* of the class A' . (See also equation (6) in [10].) The definition of $e(\cdot)$ is such that $e(A') = e(S) + n_\Sigma \chi(\Sigma)$, which is $e(S) + n_\Sigma(2 - 2G)$ (this formula is invoked in the argument for Corollary 4.5 in [10].) Meanwhile, because Z_n is holomorphic, the Gromov convergence as n gets large that results in the pointwise convergence of $\{Z_n\}_{n=1,2,\dots}$ to the union of S with n_Σ constant (s, t) copies of Σ can be used to relate $\chi(Z'_n)$ to $\chi(S)$. In particular, the formula that relates them has the same form as that for a similar degeneration were each Z_n and S a compact holomorphic submanifold in some complex surface (for example, the product of two Riemann surfaces). The relationship has the same form as in the compact case because each of the n_Σ copies of Σ intersects S in a set with compact closure in the interior of S . The relation is

$$(8-6) \quad \chi(Z'_n) = \chi(S) + n_\Sigma 2G - 2n_\Sigma(G - 1).$$

More is said by way of an explanation for this formula in a moment. What was said about $e(A')$ being $e(S) + n_\Sigma(2 - 2G)$ in conjunction with the formula in (8-6) and writes $\text{index}(D_{Z'_n})$ as $G - \chi(S) + 2e(S) + 2n_\Sigma$. This identifies $\text{index}(D_{Z'_n})$ with $\text{index}(D_S) + 2n_\Sigma$ because S has its corresponding version of (8-5) that asserts $\text{index}(D_S) = G - \chi(S) + 2e(S)$.

Step 9 The argument for (8-6) is simplest when each of the constant (s, t) versions of Σ that appear in ϑ_Σ have multiplicity 1. In this event, the first homology of Z'_n is related to that of S (because of the manner of the Gromov convergence) as follows: The surface S contributes $H_1(S)$ homology classes to $H_1(Z'_n)$. This accounts for the $\chi(S)$ factor in (8-6). Meanwhile, each copy of Σ adds $2G$ additional homology cycles (which accounts for the $n_\Sigma 2G$ factor in (8-6)). In addition, all but one of the intersection points between S and Σ add 2 new homology cycles, and the surface S has G intersections (counted with multiplicity) with each constant (s, t) version of Σ . (This accounts for the $2n_\Sigma(G - 1)$ factor in (8-6).)

Step 10 The argument for (8-6) in the general case is (co)homological. By way of some background for what is to come, suppose for the moment that Z is a compact, pseudoholomorphic surface in an almost complex 4-manifold. Let $[Z]$ denote the pushforward of Z 's fundamental class to the homology of the 4-manifold and let \hat{e} denote the Poincaré dual of $[Z]$. This is a 2-dimensional cohomology class in the ambient 4-manifold. The adjunction formula says that $-\chi(Z) = \hat{e}([Z]) + c_1([Z])$ with c_1 denoting the first Chern class of $(2, 0)$ part of the complexified cotangent bundle of the ambient manifold. An analog of the adjunction formula in what follows obtains (8-6).

Let \mathcal{X} denote the manifold with corners $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times \Sigma$ and let Z_{n_0} denote the part of Z_n in \mathcal{X} . This is the part of Z_n that is a properly embedded, pseudoholomorphic submanifold. Also, if n is large, then Z_{n_0} has transversal intersection with the boundary and corners of \mathcal{X} . Introduce by way of notation $S_\mathcal{X}$ to denote $S \cap \mathcal{X}$. This is a properly embedded submanifold with boundary also. The submanifold Z_{n_0} near $\partial\mathcal{X}$ is very close to $S_\mathcal{X}$ when n is large. In particular, it can be written near $\partial\mathcal{X}$ as the image of $S_\mathcal{X}$ near $\partial\mathcal{X}$ of an isotopy that preserves $\partial\mathcal{X}$ and moves points only a very small amount.

The Poincaré dual of the fundamental class of Z_{n_0} is a class in $H^2(\mathcal{X}; \mathbb{Z})$, which will be denoted by \hat{e}_n in what follows. The intersection number between Z_{n_0} and any given compact, oriented surface in the interior of \mathcal{X} gives the pairing between \hat{e}_n and the

pushforward to $H_2(\mathcal{X}; \mathbb{Z})$ of the fundamental class of the surface. The pairings of \hat{e}_n with the fundamental classes of Z_{n_0} and S are not a priori defined because the latter have boundaries. Even so, pairings with these fundamental classes can be defined as the intersection numbers between an allowed isotopic copy of Z_{n_0} with the original and, likewise, with S . The only a priori constraint on the isotopic copy of Z_{n_0} is that it should be disjoint on $\partial\mathcal{X}$ from Z_{n_0} and S . This is the case, for example, if the boundary of the isotopic copy is disjoint from $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$. However, additional constraints are described in the next paragraph.

The isotopic copy of Z_{n_0} can and should be constructed from the time 1 flow of a vector field on a neighborhood of $\partial\mathcal{X}$ in \mathcal{X} that is tangent to $\partial\mathcal{X}$ along $\partial\mathcal{X}$. It is important for what follows to choose this vector field so that it is normal to Z_{n_0} along ∂Z_{n_0} and normal to $S_{\mathcal{X}}$ along $\partial S_{\mathcal{X}}$. The time 1 flow of the vector field should also have two additional properties: First, it should define an isotopy that pushes the boundary of $S_{\mathcal{X}}$ out of $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$. Second, the boundaries of the isotopic versions of Z_{n_0} and $S_{\mathcal{X}}$ must remain isotopic in the complement in $\partial\mathcal{X}$ of the boundary of the domain $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$. Let ∂_* denote the vector field in question. Such a vector field and isotopy can be constructed to have the desired properties for $S_{\mathcal{X}}$ and then, for large n , it will also have the desired properties for Z_{n_0} because Z_{n_0} is very close to $S_{\mathcal{X}}$ near $\partial\mathcal{X}$ when n is large.

By virtue of what is said in Sections II.6B and II.6C, the vector field can be constructed using cutoff functions from local lifts of a vector field on a neighborhood of $T_- \cup T_+$ in Σ . See Figure 4 for a schematic picture.




Figure 4

In this diagram, the top vertical line segment indicates the image in Σ of the projection of part of ∂S that is mapped to C_- , and the left horizontal line segment indicates the corresponding image in Σ of the part of ∂S that is mapped to C_+ . Their intersection is

a point in $C_- \cap C_+$. The arrows indicate the direction of a vector field that can be lifted to give ∂_* near the relevant boundary components of $\partial S_{\mathcal{X}}$ (which is very near ∂S if z_* is small.) (If another part of ∂S projects to this same neighborhood of $C_- \cap C_+$, then its image will appear as a 180° rotation of the preceding picture, in which case the 180° rotated arrows indicate the direction of a vector field that can be lifted to give ∂_* near this other part of ∂S .)

Let Z_{n_1} denote the isotopic version of Z_{n_0} that is described in the preceding paragraphs, and let $S_{\mathcal{X}_1}$ denote the corresponding isotopic version of $S_{\mathcal{X}}$. What follows is a consequence of ∂Z_{n_1} being isotopic to $\partial S_{\mathcal{X}_1}$ in the complement in $\partial \mathcal{X}$ of the part of $\partial \mathcal{X}$ from $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$: the intersection numbers of Z_{n_1} and $S_{\mathcal{X}_1}$ with Z_{n_0} obey the formula $\langle Z_{n_1}, Z_{n_0} \rangle = \langle S_{\mathcal{X}_1}, Z_{n_0} \rangle + n_\Sigma \langle \Sigma, Z_{n_0} \rangle$. (Here and below $\langle \cdot, \cdot \rangle$ denotes the intersection number between, oriented, 2-dimensional submanifolds of \mathcal{X} that intersect only in the interior of \mathcal{X} .) Meanwhile, the fact that ∂Z_{n_0} is isotopic to $\partial S_{\mathcal{X}}$ in the part of $\partial \mathcal{X}$ from $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$ has the following implication: The number $\langle S_{\mathcal{X}_1}, Z_{n_0} \rangle + n_\Sigma \langle \Sigma, Z_{n_0} \rangle$, which is $\langle Z_{n_1}, Z_{n_0} \rangle$, is the same as $\langle S_{\mathcal{X}_1}, S_{\mathcal{X}} \rangle + 2n_\Sigma \langle \Sigma, S_{\mathcal{X}} \rangle$. (Keep in mind that $\langle \Sigma, \Sigma \rangle = 0$.) \square

Step 11 Fix a properly embedded, oriented surface with boundary (and corners) in \mathcal{X} whose fundamental class is Poincaré dual to the first Chern class of $T^{2,0} \mathcal{X}$ (which is isomorphic to the pullback via the projection of $T^{1,0} \Sigma$). This surface can and should be chosen so that it is disjoint near $\partial \mathcal{X}$ from $[-3s_n, 3s_n] \times [1 + z_*, 2 - z_*] \times (T_- \cup T_+)$. Denote such a surface by \mathfrak{S} . By definition, there is a section of $T^{2,0} \mathcal{X}$ that is nowhere zero on $\mathcal{X} \setminus \mathfrak{S}$ and vanishes transversely along \mathfrak{S} . Denote this section by \mathfrak{s} . Because of the $\partial \mathcal{X}$ constraint, there are well-defined intersection numbers between the surface \mathcal{Y} and both Z_{n_0} and S_{n_0} . These intersection numbers are used to define the pairing between $c_1(T^{2,0} \mathcal{X})$ and Z_{n_0} and $S_{\mathcal{X}}$. These are related by the rule $\langle \mathfrak{S}, Z_{n_0} \rangle = \langle \mathfrak{S}, S_{\mathcal{X}} \rangle + n_\Sigma \langle \mathfrak{S}, \Sigma \rangle$. (Note in this regard that $\langle \mathfrak{S}, \Sigma \rangle = 2G - 2$.)

Step 12 Now the interior of S is pseudoholomorphic and this implies the following version of the adjunction formula: $-\chi(S) = \langle S_{\mathcal{X}_1}, S \rangle + \langle \mathfrak{S}, S \rangle + \mathfrak{R}_S$ with \mathfrak{R}_S denoting an integer correction term that can be computed directly from the homotopy class along $\partial S_{\mathcal{X}}$ of the nonzero section \mathfrak{s} of $T^{2,0} \mathcal{X} \setminus 0$ and the homotopy class of the nonzero vector normal vector field ∂_* along ∂Z_n . By the same token, the number $-\chi(Z_n) = \langle Z_{n_1}, Z_n \rangle + \langle \mathfrak{S}, Z_n \rangle + \mathfrak{R}_n$ with \mathfrak{R}_n denoting an integer correction term that can be computed directly from the homotopy class along ∂Z_n of the nonzero section \mathfrak{s} of $T^{2,0} \mathcal{X} \setminus 0$ and the homotopy class of the nonzero vector field ∂_* along ∂Z_n . Now,

the key point is that $\partial S_{\mathcal{X}}$ is isotopic to ∂Z_n along $\partial \mathcal{X}$ and that this isotopy identifies the respective homotopy classes of \mathfrak{s} and the respective homotopy classes of ∂_* when n is sufficiently large. This implies that $\mathfrak{R}_S = \mathfrak{R}_n$ (keep in mind that both of these correction factors are integers). Therefore, the writing of $\langle Z_{n_1}, Z_n \rangle$ as $\langle S_{\mathcal{X}_1}, S_{\mathcal{X}} \rangle + 2n_{\Sigma} \langle \Sigma, S_{\mathcal{X}} \rangle$ and the writing of $\langle \mathfrak{S}, Z_n \rangle$ as $\langle \mathfrak{S}, S_{\mathcal{X}} \rangle + n_{\Sigma} \langle \mathfrak{S}, \Sigma \rangle$ gives the formula in (8-6) because $\langle \Sigma, S_{\mathcal{X}} \rangle = G$ and $\langle \mathfrak{S}, \Sigma \rangle = 2G - 2$. \square

8B Quillen's construction and orientations

As explained in [6; 5; 7, Section 9], the differential for embedded contact homology is defined using certain dimension one moduli spaces of J -holomorphic subvarieties in $\mathbb{R} \times Y$. The definition involves a ± 1 weight that is defined by comparing two orientations that can be defined for these moduli spaces. The first is defined by the \mathbb{R} action that is induced by the constant translations along the \mathbb{R} factor of $\mathbb{R} \times Y$. The second is defined using notions that were introduced by Quillen [15] about determinant line bundles for parametrized families of Fredholm operators. Section 9 uses what is said here and in Sections 8C–8E to show that the weight used for the embedded contact homology differential differs by a purely S -dependent sign from the weight used in Proposition 8.1. This subsection describes the relevant version of Quillen's construction of orientations. The story is told in five parts.

Part 1 To say more about the Fredholm operator that is used to define the differential for embedded contact homology, fix $(\tau = 1, \mathcal{C}) \in \mathcal{M}^*$ and let C denote the corresponding ech-HF submanifold. The operator in question is the operator $\mathcal{D}_{\mathcal{C}}$ from Lemma 7.15 acting here on a slightly larger domain. The range space is the same as for the original. The domain is denoted here by \mathbb{H}_S . The space \mathbb{H}_S is defined just as $\mathbb{H}_{\tau=1}$ in Part 1 of Section 7E but for the following: Elements in the S -labeled summand of $\mathbb{H}_{\tau=1}$ are required to be L^2 -orthogonal on the $t \in [1 + z_*, 2 - z_*]$ part of S to the restriction of the kernel of D_S . This last condition is not imposed on the elements in the S -labeled summand of \mathbb{H}_S . In any case the $\tau = 1$ version of (7-32) is imposed. If $z_* \leq c^{-1}$ with $c \geq 1$ purely S -dependent (of \mathcal{K} -compatible), then there is a canonical isomorphism between \mathbb{H}_S and $\text{kernel}(D_S) \oplus \mathbb{H}_{\tau=1}$.

This version of $\mathcal{D}_{\mathcal{C}}$ with domain \mathbb{H}_S is denoted in what follows by $\widehat{\mathcal{D}}_{\mathcal{C}}$. The notation here is meant to indicate that the latter operator can be defined intrinsically as an operator on C . As noted in the proof of Proposition 8.2, it has the form of what is depicted on the right-hand side of (1-25). This intrinsic definition identifies the domain

Hilbert space with the space of L^2_1 sections of C 's normal bundle and the range Hilbert space to the space of L^2 sections of the tensor product of this normal bundle with $T^{0,1}C$. This intrinsic definition does not reference (7-32).

Part 2 This part of the subsection summarizes Quillen's construction. To start, suppose for the moment that \mathcal{H}_1 and \mathcal{H}_2 are Hilbert spaces and $\mathcal{D}: \mathcal{H}_1 \rightarrow \mathcal{H}_2$ is a Fredholm operator with positive index and trivial cokernel. Define $\mathbb{D}\text{et}(\mathcal{D})$ to be the real line given by the top exterior power of the kernel of \mathcal{D} .

Suppose next that \mathcal{Y} denotes a smooth, finite-dimensional manifold and that \mathcal{H}_1 and \mathcal{H}_2 are Hilbert space bundles over \mathcal{Y} . (Most of what is said here generalizes readily to the case when \mathcal{Y} is a Hilbert manifold.) Let \mathcal{D} now denote a continuous section of $\text{Hom}(\mathcal{H}_1, \mathcal{H}_2)$ whose restriction to each fiber is Fredholm. What follows describes a real line bundle that is defined over \mathcal{Y} in a canonical fashion by \mathcal{D} . To this end, note first that the manifold \mathcal{Y} has a locally finite cover with the following property: Let $U \subset \mathcal{Y}$ denote a set from this cover. Then there exist a nonnegative integer, n , and a bundle homomorphism $L: U \times \mathbb{R}^{2n} \rightarrow \mathcal{H}_2|_U$ such that

$$(8-7) \quad \mathcal{D} + L: \mathcal{H}_1|_U \oplus (U \times \mathbb{R}^{2n}) \rightarrow \mathcal{H}_2|_U$$

restricts to each fiber as a linear map with positive index and trivial cokernel. This understood, define the real line bundle $\mathbb{D}\text{et}|_U \rightarrow U$ to be the bundle whose fiber at any given $y \in U$ is the top exterior power of the kernel of $(L + \mathcal{D})|_y$. As explained by Quillen, different choices for the integer n and, given n , for the homomorphism L , subject to the condition that $\mathcal{D} + L$ have positive index and trivial cokernel, give canonically isomorphic versions of $\mathbb{D}\text{et}|_U$. It follows as a consequence that these line bundles over the open sets of the given cover patch together over the pairwise intersections to define a real line bundle over \mathcal{Y} . It also follows (by taking subdivisions) that two covers of \mathcal{Y} with the requisite properties supply isomorphic bundles. This being the case, the construction just described defines from \mathcal{D} a canonical real line bundle over \mathcal{Y} . This is the bundle $\mathbb{D}\text{et}$.

The particular version of $\mathbb{D}\text{et}$ that is used in the definition of the embedded contact homology differential is defined over the moduli spaces of ech-HF submanifolds. Let C denote the ech-HF submanifold that is associated to a given pair $(1, \mathcal{C}) \in \mathcal{M}^*$. The fiber of the relevant version of $\mathbb{D}\text{et}$ over C is the determinant line of the operator $\widehat{\mathcal{D}}_C$.

Part 3 The versions of $\mathbb{D}\text{et}$ that arise in what follows are orientable. This part of the subsection sets up the conventions that are used in the subsequent parts that concern

choices of orientations. If V and V' denote an ordered pair of oriented vector spaces, then their direct sum has a canonical orientation that is obtained as follows: Let n and n' denote the respective dimensions of V and V' . Choose respective bases $\{v_1, \dots, v_n\}$ for V and $\{v'_1, \dots, v'_{n'}\}$ for V' such that $v_1 \wedge \dots \wedge v_n$ defines the orientation for $\det(V) = \bigwedge^n V$ and $v'_1 \wedge \dots \wedge v'_{n'}$ defines the orientation for $\det(V') = \bigwedge^{n'} V'$. The orientation for $\det(V \oplus V') = \bigwedge^{n+n'} (V \oplus V')$ is defined by

$$(8-8) \quad v_1 \wedge \dots \wedge v_n \wedge v'_1 \wedge \dots \wedge v'_{n'}.$$

The oriented tensor product of the lines $\det(V)$ and $\det(V')$ is defined to be the line $\det(V \oplus V')$ with the orientation given by (8-8). This oriented line is denoted by $\det(V) \otimes \det(V')$. The oriented lines $\det(V') \otimes \det(V)$ and $\det(V) \otimes \det(V')$ are isomorphic as oriented lines if and only if nn' is an even number.

It follows as a consequence of what was just said that ordering issues are minimized in any given situation when one or both of V and V' have even dimensions; this is why (8-7) uses only even-dimensional Euclidean spaces. In particular, the restriction to even dimensions in (8-7) makes it easier to compare orientations for $\mathbb{D}\det$.

What follows is meant to provide an abstract but relevant illustration. The vector space \mathbb{R}^{2n} in here and in subsequent parts of this subsection always denotes the eponymous vector space with a standard orientation, chosen once and for all time. Suppose that \mathcal{D} is a Fredholm operator with trivial cokernel and positive index. Choose an orientation for $\ker(\mathcal{D})$ so as to orient the line $\mathbb{D}\det(\mathcal{D})$. Let \mathcal{H}_2 denote the range space for \mathcal{D} , let n denote any given positive integer and let $L: \mathbb{R}^{2n} \rightarrow \mathcal{H}_2$ denote any given map. The kernel of $\mathcal{D} + L$ is canonically isomorphic to $\ker(\mathcal{D}) \oplus \mathbb{R}^{2n}$, and so the oriented line $\det(\ker(\mathcal{D}))$ is canonically isomorphic as an oriented line to $\det(\ker(\mathcal{D} \oplus L))$. As a consequence, the orientation on $\mathbb{D}\det(\mathcal{D})$ defined via its identification with $\det(\ker(\mathcal{D}))$ is the same as that defined by its identification with $\det(\ker(\mathcal{D} \oplus L))$.

For example, suppose now that \mathcal{D} has trivial cokernel and zero index. Fix nonnegative integers n and n' whose sum is at least 1, and fix linear maps $L: \mathbb{R}^{2n} \rightarrow \mathcal{H}_2$ and $L': \mathbb{R}^{2n'} \rightarrow \mathcal{H}_2$. The kernel of $(\mathcal{D} + L) + L': (\mathcal{H}_1 \oplus \mathbb{R}^{2n}) \oplus \mathbb{R}^{2n'} \rightarrow \mathcal{H}_2$ is canonically isomorphic to $\mathbb{R}^{2n} \oplus \mathbb{R}^{2n'}$ and the kernel of $(\mathcal{D} + L') + L: (\mathcal{H}_1 \oplus \mathbb{R}^{2n'}) \oplus \mathbb{R}^{2n} \rightarrow \mathcal{H}_2$ is canonically isomorphic to $\mathbb{R}^{2n'} \oplus \mathbb{R}^{2n}$. The orientation on $\mathbb{D}\det(\mathcal{D})$ that comes by identifying the latter with the oriented, top exterior power of $\mathbb{R}^{2n} \oplus \mathbb{R}^{2n'}$ is the same as that defined by identifying $\mathbb{D}\det(\mathcal{D})$ with the top exterior power of $\mathbb{R}^{2n'} \oplus \mathbb{R}^{2n}$.

To end this illustration, suppose that \mathcal{D}' is a second Fredholm operator with trivial cokernel and positive index. Orient $\text{kernel}(\mathcal{D}')$ so as to orient the line $\mathbb{D}\text{et}(\text{kernel}(\mathcal{D}'))$ and so orient $\mathbb{D}\text{et}(\mathcal{D}')$. Because the integer $2n$ in (8-7) is even, the orientation of the oriented line $\mathbb{D}\text{et}(\mathcal{D}) \otimes \mathbb{D}\text{et}(\mathcal{D}')$ is insensitive to the choice for n and L in (8-7) when defining $\mathbb{D}\text{et}(\mathcal{D})$ subject to the constraint that $\mathcal{D} + L$ has positive index and trivial cokernel.

Part 4 Suppose $(\tau = 1, \mathcal{C}) \in \mathcal{M}^*$ and let C denote the corresponding ech-HF submanifold. As noted in Part 1, the operator that is used to define the embedded contact homology differential is the intrinsically defined operator $\widehat{\mathcal{D}}_C$. Of particular concern with regards to the differential is a certain orientation for the real line $\mathbb{D}\text{et}(\widehat{\mathcal{D}}_C)$. The operator $\widehat{\mathcal{D}}_C$ can be viewed as the operator \mathcal{D}_C with domain space $\mathbb{H}_S = \text{kernel}(D_S) \oplus \mathbb{H}_{\tau=1}$. It follows as a consequence of what is said in Parts 2 and 3 that

$$(8-9) \quad \mathbb{D}\text{et}(\widehat{\mathcal{D}}_C) = \mathbb{D}\text{et}(D_S) \otimes \mathbb{D}\text{et}(\mathcal{D}_C).$$

This understood, orientations for the line $\mathbb{D}\text{et}(D_S)$ and for the line $\mathbb{D}\text{et}(\mathcal{D}_C)$ orient the line of interest, $\mathbb{D}\text{et}(\widehat{\mathcal{D}}_C)$. The rest of this Part 4 explains how to relate the line $\mathbb{D}\text{et}(\mathcal{D}_C)$, and thus the line $\mathbb{D}\text{et}(\widehat{\mathcal{D}}_C)$, to Proposition 8.1.

Suppose that (τ, \mathcal{C}) is any given pair in \mathcal{M}^* . Use \mathcal{D}_C in this case to denote the operator from Lemma 7.15 acting on the Banach space \mathbb{H}_τ . The range Banach space is \mathcal{C} 's version of \mathbb{L} . These respective Banach spaces are the fibers over \mathcal{M}^* of a pair of smooth, Banach space bundles, \mathcal{H}_1 and \mathcal{H}_2 . Meanwhile, the various (τ, \mathcal{C}) versions of \mathcal{D}_C define the fibers of a section, \mathcal{D} , of $\text{Hom}(\mathcal{H}_1, \mathcal{H}_2)$ which is Fredholm on each fiber. Note in this regard that the smooth variation can be proved using what is said in Step 3 from Part 4 of Section 7D (with arguments that mimic those from Step 3 from Part 3 of Section 5C). The section \mathcal{D} has its associated determinant line bundle, $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$. The $(\tau = 1, \mathcal{C})$ version of $\mathbb{D}\text{et}(\mathcal{D}_C)$ that appears in (8-9) is the fiber over $(1, \mathcal{C})$ of $\mathbb{D}\text{et}(\mathcal{D})$.

Hold on to $\mathbb{D}\text{et}(\mathcal{D})$ for a moment and reintroduce the map \mathfrak{F} as in (7-55). The differential of \mathfrak{F} at the point $(\tau, 0)$ defines a Fredholm map from $\mathbb{R} \times \mathbb{H}_\tau \rightarrow \mathbb{L}$ of the form $\mathcal{D}_C + L_C$ where $L_C: \mathbb{R} \rightarrow \mathbb{L}$ is a linear map. Given the choice for J , the operator $\mathcal{D}_C + L_C$ has trivial cokernel and kernel dimension equal to $\sum_{p \in \Lambda} \Delta_p + 1$. This kernel is canonically isomorphic to $T\mathcal{M}^*$. To summarize,

$$(8-10) \quad \text{kernel}(\mathcal{D}_C + L_C) = \text{kernel}(\mathfrak{F}_*) = T\mathcal{M}^*|_{(\tau, \mathcal{C})}.$$

The differential of π_1 at (τ, C) appears here as the restriction to $\ker(\mathcal{D}_C + L_C)$ of the projection map from $\mathbb{R} \times \mathbb{H}_\tau$ to the \mathbb{R} factor. Meanwhile, $\bigwedge^{\text{top}} \ker(\mathcal{D}_C + L_C)$ is the line $\mathbb{D}\text{et}(\mathcal{D}_C + L_C)$. Thus, an orientation for $T\mathcal{M}^*|_{(\tau, C)}$ is induced from an orientation for $\mathbb{D}\text{et}(\mathcal{D}_C + L_C)$ and vice versa.

At a point (τ, C) where τ is a regular value of π_1 , the determinant line $\mathbb{D}\text{et}(\mathcal{D}_C + L_C)$ is isomorphic to $\mathbb{R} \otimes \mathbb{D}\text{et}(\mathcal{D}_C)$. As the differential of π_1 at (τ, C) is an isomorphism, the differential of π_1 at (τ, C) identifies the \mathbb{R} factor in this tensor product as the oriented line $(-\infty, \infty)$. Reintroduce now the line $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$. It follows as a consequence of what was just said that an orientation for the line $\mathbb{D}\text{et}(\mathcal{D})$ defines an orientation for the $\pi_1 = \tau$ level set in \mathcal{M}^* , and in particular for the $\tau = 1$ boundary of \mathcal{M}^* . Note in this regard that this level set is a union of points when each $p \in \Lambda$ version of Δ_p are zero. Meanwhile, an orientation of a point is, by definition, a choice of $+1$ or -1 . The $+1$ appears if the orientation of $T\mathcal{M}^*$ given by $\mathbb{D}\text{et}(\mathcal{D})$ at the point in question is that given by the differential of π_1 . By the same token, the orientation given by $\mathbb{D}\text{et}(\mathcal{D})$ to the $\tau = 0$ point in \mathcal{M}^* is $+1$ if the orientation of $T\mathcal{M}^*$ at this point agrees with that given by the differential of π_1 .

Part 5 Suppose that $\sum_{p \in \Lambda} \Delta_p = 0$. An orientation for $\mathbb{D}\text{et}(\mathcal{D})$ induces an orientation on $T\mathcal{M}^*|_{\tau=1}$ which agrees or not with that used in Proposition 8.1; but agreement or not is the same at all points in $\mathcal{M}^*|_{\tau=1}$. It follows from (8-10) that agreement occurs if and only if there is agreement for the sole point of $\mathcal{M}^*|_{\tau=0}$. Meanwhile, Section 8C describes a completely canonical orientation for $\mathbb{D}\text{et}(\mathcal{D})$ when $\sum_{p \in \Lambda} \Delta_p = 0$. As can be seen readily from the definition in Section 8C, the resulting orientation for $\mathcal{M}^*|_{\tau=0}$ agrees with that used in Proposition 8.1. This being the case, the $\mathbb{D}\text{et}(\mathcal{D})$ orientation for $\mathcal{M}^*|_{\tau=0}$ also agrees with Proposition 8.1's orientation when $\sum_{p \in \Lambda} \Delta_p = 0$.

In the case $\sum_{p \in \Lambda} \Delta_p > 0$, choose once and for all an ordering of Λ_* up to even permutations. Such a choice orients $(\times_{p \in \Lambda_*} \mathbb{R})$. Suppose that $y \in (\times_{p \in \Lambda_*} \mathbb{R})$ is a regular value of the map p . Then the tangent space to \mathcal{M}_y^* is isomorphic at any given point (τ, C) to the kernel of the p 's differential. Meanwhile, the normal bundle of \mathcal{M}_y^* in \mathcal{M}^* is mapped isomorphically by p 's differential to $(\times_{p \in \Lambda_*} \mathbb{R})$. This understood, Proposition 8.1's orientation for \mathcal{M}_y^* and p 's orientation of the normal bundle to \mathcal{M}_y^* orients the tangent space to any given smooth level set of π_1 and in particular the tangent space to $\mathcal{M}^*|_{\tau=1}$. Granted this last observation, (8-10) implies that the orientation of $\mathcal{M}^*|_{\tau=1}$ given by the differential of $\pi_1 \times p$ either agrees with or disagrees with the orientation induced by an orientation of $\mathbb{D}\text{et}(\mathcal{D})$; but agreement or not is the same at all

points. Moreover, agreement occurs if and only if the corresponding two orientations of $\mathcal{M}^*|_{\tau=0}$ agree.

Meanwhile, Section 8D describes a completely canonical orientation for the line $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$ which consistently orients the level sets of π_1 in \mathcal{M}^* . A look ahead at what is said in Section 8D shows that this $\mathbb{D}\text{et}(\mathcal{D})$ orientation for $\mathcal{M}^*|_{\tau=0}$ agrees with the one defined by Proposition 8.1 and so the $\mathbb{D}\text{et}(\mathcal{D})$ orientation for $\mathcal{M}^*|_{\tau=1}$ also agrees with Proposition 8.1's orientation of $\mathcal{M}^*|_{\tau=1}$.

8C The canonical orientation when $\{\Delta_p = 0\}_{p \in \Lambda}$

This subsection describes the promised canonical orientation for $\mathbb{D}\text{et}(\mathcal{D})$ in the case when all $p \in \Lambda$ versions of Δ_p are zero. The description has four parts.

Part 1 The Banach space \mathbb{H}_τ in this case can be viewed as a completion of a dense domain whose typical element is written as $(\eta'_S, \{(\varphi^{p'}, \zeta^{p'})\}_{p \in \Lambda})$ where η'_S is a compactly supported section of N_S over the $t \in [1 + z_*, 2 - z_*]$ part of S that is orthogonal to the restriction of the kernel of D_S . Meanwhile, each $p \in \Lambda$ version of $(\varphi^{p'}, \zeta^{p'})$ is a compactly supported map from $\mathbb{R} \times I_*$ to \mathbb{R}^2 . The parameter τ enters through the boundary constraint in (7-32). The Banach space norm is that induced by the L^2_1 -norm on sections of the $t \in [1 + z_*, 2 - z_*]$ part of S and the L^2_1 -norm on maps from $\mathbb{R} \times I_*$. Meanwhile, the range space \mathbb{L} is the corresponding L^2 completion of a dense domain with typical element $(\eta^\#_S, \{(\varphi^{p\#}, \zeta^{p\#})\}_{p \in \Lambda})$ where $\eta^\#_S$ is a section of an appropriate 2-plane bundle of the $t \in [1 + z_*, 2 - z_*]$ part of S , and where each $p \in \Lambda$ version of $(\varphi^{p\#}, \zeta^{p\#})$ is a map from $\mathbb{R} \times I_*$ to \mathbb{R}^2 . The Banach space norm is the norm induced by the L^2 -norm on the $t \in [1 + z_*, 2 - z_*]$ part of S and on $\mathbb{R} \times I_*$.

As can be seen from preceding descriptions, the Banach space \mathbb{H}_τ depends only on τ (with S fixed) and the Banach space \mathbb{L} depends only on S . What is done in Step 3 from Part 3 of Section 5C can be mimicked to see that the assignment to $\tau \in [0, 1]$ of the Banach space \mathbb{H}_τ defines a smooth Banach space bundle over the interval $[0, 1]$. The latter's π_1 -pullback over \mathcal{M}^* is the Banach space bundle \mathcal{H}_1 that was defined in Part 4 of the previous subsection. Meanwhile, the bundle \mathcal{H}_2 from this same Part 4 is the product bundle $\mathcal{M}^* \times \mathbb{L}$.

Part 2 The Hilbert space bundles \mathcal{H}_1 and \mathcal{H}_2 extend over the product $\mathcal{M}^* \times [0, 1]$ as follows: The bundle \mathcal{H}_2 extends as the product bundle. The bundle \mathcal{H}_1 extends as the pullback via the projection to the square $[0, 1] \times [0, 1]$ of the bundle whose fiber

at any given point (τ, r) is the Hilbert space $\mathbb{H}_{r\tau}$. The homomorphism \mathcal{D} from \mathcal{H}_1 to \mathcal{H}_2 over \mathcal{M}^* given by $(\tau, \mathcal{C}) \rightarrow \mathcal{D}_{\mathcal{C}}$ extends over $\mathcal{M}^* \times [0, 1]$ so that the restriction to each fiber is Fredholm. The section at a given $((\tau, \mathcal{C}), r)$ is the operator $\mathcal{D}_{\mathcal{C}}$ with the parameter τ in (7-32) replaced by $r\tau$. This extended bundle homomorphism is also denoted by \mathcal{D} . An orientation over $\mathcal{M}^* \times \{0\}$ for the associated real line bundle $\mathbb{D}\text{et}(\mathcal{D})$ defines an orientation for $\mathbb{D}\text{et}(\mathcal{D})$ over the whole of $\mathcal{M}^* \times [0, 1]$ and thus over $\mathcal{M}^* \times \{1\}$. Of course, the converse of this last assertion is also true.

To define an orientation for $\mathbb{D}\text{et}(\mathcal{D})$ over $\mathcal{M}^* \times \{0\}$, fix $(\tau, \mathcal{C}) \in \mathcal{M}^*$. The version of $\mathcal{D}_{\mathcal{C}}$ on \mathbb{H}_0 is a direct sum of operators. In particular, there is one for each $\mathfrak{p} \in \Lambda$. In each case, it is a version of (3-5) with coefficients that obey (3-6) and are given by (3-9). The elements in the dense domain of this \mathfrak{p} -summand operator are constrained to obey the boundary conditions in (3-7). There is the remaining S -labeled summand. The relevant operator here is that depicted in (7-33) with domain given by the orthogonal complement of the restriction kernel of D_S to the part of S where $t \in [1 + z_*, 2 - z_*]$. Elements in the domain also obey the $\tau = 0$ version of the top line in (7-32).

As explained in Step 1 from Section 3D, each $\mathfrak{p} \in \Lambda$ version of (3-5) in the aforementioned direct sum is homotopic through a family of Fredholm operators to a canonical operator, this given by (3-22) with the homotopy given by (3-21). The operators in this family all have trivial kernel and cokernel. This understood, then, a once and forever choice for the orientation of the determinant line of the operator in (3-22) with boundary conditions given by (3-7) orients the determinant line for each of $\mathfrak{p} \in \Lambda$ version of (3-5).

Part 3 This part elaborates on what was said at the end of Part 2. To start, introduce \mathcal{O} to denote the set whose elements are 6-tuples of functions on $\mathbb{R} \times I_*$ that are of the form $(\alpha_1, \alpha_2, \beta_1, \beta_2)$ which are suitable for use in (3-5). In particular, they must obey the constraints in (3-6). No generality is lost for what follows by restricting to the case where the integrals in the third bullet are negative. The set \mathcal{O} is given the topology that is induced by its inclusions into two topological function spaces of maps from $\mathbb{R} \times I_*$ to \mathbb{R}^6 . The first topology is the C^∞ -Fréchet space topology with it understood that convergence means convergence in the various C^k topologies on compact subsets. The second topology is the strong C^1 topology. The space \mathcal{O} with this topology is contractible since the constraints in (3-6) form a convex set.

Each point in \mathcal{O} defines an operator to which Proposition 3.2 applies. In particular, each such operator has trivial kernel and cokernel. Let $\mathbb{D}\text{et}$ denote the corresponding

determinant line bundle. It follows from what is said in Part 2 of the previous subsection that this version of $\mathbb{D}et$ has a canonical identification with the oriented line $\wedge^2 \mathbb{R}^2$. Use this orientation. Such a choice gives the determinant line of the operators parametrized by the elements in \mathcal{O} a canonical orientation. Note that the space \mathcal{O} and the line $\mathbb{D}et$ see nothing of the spaces M and Y or their geometry. They do depend on the parameters z_* and R only to the extent that these define I_* . Even so, the allowed choices form a contractible set, so the orientation of $\mathbb{D}et$ just described is truly canonical and universal.

Part 4 As noted, the other operator that enters the direct sum giving \mathcal{D}_C on \mathbb{H}_0 is depicted in (7-33) but acting on the orthogonal complement of restriction of the kernel of D_S to the $t \in [1 + z_*, 2 - z_*]$ part of S . The boundary conditions are given by the $\tau = 0$ version of the top line in (7-32). This operator with the domain as indicated is denoted here by D_{η_0} . Meanwhile, D_S^\perp is used here to denote the restriction of the operator D_S (given in (II.6-11)) to the orthogonal complement of its kernel. The boundary conditions are given by (6-12). The Fredholm version needed here is the one described in Part 3 of Section II.6E.

The operator D_S^\perp has trivial kernel and cokernel. As noted previously, this is the case for D_{η_0} when $z_* < c^{-1}$, where c is purely S -dependent (or \mathcal{K} -compatible). It is a straightforward task to prove that D_{η_0} is homotopic via an essentially canonical family of Fredholm operators to D_S^\perp with each operator in the family having trivial kernel and cokernel. The details are omitted but for the following description of the family: The family is the concatenation of two 1-parameter families. The first modifies the operator on the fixed domain via the family parametrized by $[0, 1]$ with the $\mu \in [0, 1]$ version of the operator defined by (7-33) with $\mu\eta$ replacing η . The version with $\mu = 0$ is the restriction of the operator D_S to the $t \in [1 + z_*, 2 - z_*]$ part of S . The second part of the homotopy keeps the operator fixed as D_S but changes the domain Hilbert space by introducing a parameter $\mu \in [0, 1]$ and restricting D_S to the $[1 + \mu z_*, 2 - \mu z_*]$ portion of S . The boundary conditions that define any given $\mu < 1$ Hilbert space are the $t = 1 + \mu z_*$ and $t = 2 - \mu z_*$ analogs of those for the $\mu = 1$ member. Likewise, the orthogonality condition with regards to the kernel of D_S is changed only to the extent that the orthogonality is defined by integration over the $t \in [1 + \mu z_*, 2 - \mu z_*]$ part of S .

Granted all of this, it then follows that the line bundle $\mathbb{D}et(\mathcal{D})$ along $\{0\} \times \mathcal{M}^*$ has a completely canonical orientation given a choice of orientation for the determinant

line of the operator D_S^\perp . It follows from what is said in Part 2 of the previous subsection that such an orientation is canonically defined by identifying $\mathbb{D}\text{et}(D_S^\perp)$ with $\wedge^2 \mathbb{R}^2$. Moreover, it follows as a consequence of what is said in Parts 2 and 3 of the previous subsection that this orientation for $\mathbb{D}\text{et}(D_S^\perp)$ is induced from an orientation for $\mathbb{D}\text{et}(D_S)$ that comes by writing D_S as $D_S = D_S^\perp + L_S$, where L_S is the map from $\text{kernel}(D_S)$ to the range of D_S that sends all elements to zero.

Part 5 Given what is said in Part 3 of Section 5B, the respective orientations for $\mathbb{D}\text{et}(D_S^\perp)$ and for each $p \in \Lambda$ version of the determinant line of the operator in (3-5) defines a completely canonical orientation for $\mathbb{D}\text{et}(\mathcal{D})$ along $\mathcal{M}^* \times \{0\}$. As noted above, the latter defines a completely canonical orientation for $\mathbb{D}\text{et}(\mathcal{D})$ on $\mathcal{M}^* \times [0, 1]$, and thus to $\mathbb{D}\text{et}(\mathcal{D})$ along $\mathcal{M}^* \times \{1\}$. The latter is the canonical orientation promised in Part 5 of the previous subsection.

8D Canonical orientations when $\sum_{p \in \Lambda} \Delta_p > 0$

This subsection describes the promised canonical orientation for $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$ in the general case. The description that follows has three parts.

Part 1 An orientation for $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$ is defined by first extending the family of operators to a family with parameter space $\mathcal{M}^* \times [0, 1]$. This extension is defined by making the domain for any given $((\tau, \mathcal{C}), r)$ version of the operator depend on the parameter for the extra $[0, 1]$ factor. The range space is kept constant, and the operator itself stays as $\mathcal{D}_\mathcal{C}$. To say more about the r -dependence of the domain, fix a pair $((\tau, \mathcal{C}), r) \in \mathcal{M}^* \times [0, 1]$. The domain for the corresponding Fredholm operator is identical to that when $r = 1$ but for one item: the boundary conditions in (7-32) replace τ with $r\tau$. This extended family defines a homomorphism over $\mathcal{M}^* \times [0, 1]$ between the corresponding extensions of the Banach space bundles \mathcal{H}_1 and \mathcal{H}_2 . The respective extensions of these bundles over $\mathcal{M}^* \times [0, 1]$ are also denoted by \mathcal{H}_1 and \mathcal{H}_2 , and the homomorphism between them by \mathcal{D} . The latter version of \mathcal{D} has its associated determinant line, $\mathbb{D}\text{et}(\mathcal{D})$. Because the factor $[0, 1]$ is contractible, an orientation for $\mathbb{D}\text{et}(\mathcal{D})|_{r=0}$ canonically induces one for $\mathbb{D}\text{et}(\mathcal{D})|_{r=1}$ and thus for the line of interest, $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{M}^*$.

Part 2 To define a canonical orientation for $\mathbb{D}\text{et}(\mathcal{D})$ over the $r = 0$ boundary of $\mathcal{M}^* \times [0, 1]$, focus for the moment on a given pair $((\tau, \mathcal{C} = \{C_S, \{C_p\}_{p \in \Lambda}\}), r = 0)$ on this boundary. The relevant version of (7-32) is such that there is no coupling between

the summands that define the domain space for the operator \mathcal{D}_C . As such, the kernel and cokernel of \mathcal{D}_C is the direct sum of the respective kernel and cokernel for the operator D_{η_0} from Part 4 of the previous subsection, and the respective kernels and cokernels for each $C \in \{C_p\}_{p \in \Lambda}$ version of an operator that is described either by Proposition 3.2 or by Proposition 5.7.

Part 4 of the previous subsection asserts that D_{η_0} is homotopic through a family of Fredholm operators with trivial kernel and trivial cokernel to the operator D_S^\perp . In particular, this same Part 4 finds that $\mathbb{D}\text{et}(D_{\eta_0})$ is canonically isomorphic to $\mathbb{D}\text{et}(D_S^\perp)$ and thus canonically oriented by the identification $\mathbb{D}\text{et}(D_S^\perp) = \wedge^2 \mathbb{R}^2$ from Part 2 of Section 8B.

Meanwhile, if $p \in \Lambda$ and $\Delta_p = 0$, then the corresponding $C = C_p$ contribution to \mathcal{D}_C is an operator that is described in Parts 2 and 3 of the previous subsection. In particular, these parts of Section 8C endow the corresponding determinant line with a completely canonical orientation by identifying it with $\wedge^2 \mathbb{R}^2$.

In the case when $\Delta_p > 0$, the contribution of C_p to \mathcal{D}_C is an operator of the sort that is described by Proposition 5.7. Use D in what follows to denote this operator. What follows describes an absolutely canonical orientation for $\mathbb{D}\text{et}(D)$.

Proposition 5.7 asserts that D has Δ_p -dimensional kernel and trivial cokernel. This being the case, its determinant line is oriented by an orientation of its kernel. To orient the kernel of D , suppose that $\mathcal{E} \subset \mathbb{R} \times \mathcal{H}_{p*}^+$ is an end of C_p whose constant s slices converge as $s \rightarrow \infty$ to $\hat{\gamma}_p^+$, this the $(\hat{u} = 0, \cos \theta = \frac{1}{\sqrt{3}})$ integral curve of v . Reintroduce the notation from Sections 5A and 5C so as to talk about the kernel of D on \mathcal{E} . As explained in Step 5 of Part 4 from Section 5C, if $\Delta_p = 1$ and $m_p = -1$, or if $\Delta_p = 2$, there is an element in the kernel of D that can be written at large values of s_+ on \mathcal{E} as a map from the large s_+ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 that has the form

$$(8-11) \quad (s_+, \phi_+) \rightarrow c_{\mathcal{E}}(e^{-\lambda_1 s_+}, 0) + c_{\mathcal{E}}\mathbf{e},$$

where $c_{\mathcal{E}} \in \mathbb{R}$ and where $|\epsilon| \leq c_0 e^{-(\lambda_1 + 1/c_0)s_+}$.

In the case $\Delta_p = 1$ and $m_p = -1$, the kernel of D is oriented by the unique element with $c_{\mathcal{E}} = 1$. In the case $\Delta_p = 1$ and $m_p = 1$, there is an analogous orientation for the kernel of D , this defined by the $c_{\mathcal{E}} = 1$ element with $c_{\mathcal{E}}$ now defined by the analog of (8-11) for the end of C_p whose constant s slices converge to $\hat{\gamma}_p^-$ as $s \rightarrow \infty$.

If $\Delta_p = 2$, what is said in Step 5 of Part 4 of Section 5C implies that the kernel of D has a unique basis of the form (η_+, η_-) with the following properties: The element η_+

is given by the $c_\varepsilon = 1$ version of (8-11) on the end of C_p whose constant s slices converge to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$. Meanwhile, $|\eta_+| \leq c_0 e^{-(\lambda_1+1/c_0)s}$ on the end of C_p whose large s slices converge to $\hat{\gamma}_p^-$ as $s \rightarrow \infty$. The situation is reversed for η_- ; it is given by the $c_\varepsilon = 1$ version of (8-11) on the end of C_p whose constant s slices converge to $\hat{\gamma}_p^-$ as $s \rightarrow \infty$, and it is bounded in absolute value by $c_0 e^{-(\lambda_1+1/c_0)s}$ on the end whose constant s slices converge to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$.

The canonical nature of these orientations is explained in Part 3.

Part 3 This part of the subsection explains the sense in which Part 2’s orientation for the kernel of D is completely canonical.

To put things in a sufficiently general framework, use \mathcal{O} now to denote the set whose elements have the form $(D, Q, \gamma_{p+}, \gamma_{p-}, \mathfrak{h} = (\varphi, \zeta), D)$ with entries as follows: What is denoted by D is a data set (z_*, δ, χ_0, R) that can be used to define \mathcal{H}_p and the subspace $\mathcal{H}_{p_*}^+$. Meanwhile, $Q \subset \mathbb{R}$ consists of a single point if $\Delta_p = 1$; it consists of two labeled points if $\Delta_p = 2$, one point labeled with a plus sign and the other with a minus sign. The pair γ_{p+} and γ_{p-} are integral curves of v in $\mathcal{H}_{p_*}^+$ with one boundary point on the $u < 0$ boundary of $\mathcal{H}_{p_*}^+$ and the other boundary point is on the $u > 0$ boundary. What is denoted by \mathfrak{h} signifies a smooth map from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 of the sort that is described at the start of Section 5C; and D denotes an operator of the sort described by Proposition 5.7. The next two paragraphs define a topology for \mathcal{O} .

Let $\hat{d} = (D, Q, \gamma_{p+}, \gamma_{p-}, \mathfrak{h} = (\varphi, \zeta), D)$ denote a given element in \mathcal{O} . A basis for the neighborhoods of \hat{d} is indexed by $(\varepsilon, \mathcal{V}, \mathcal{W})$ where the notation is as follows: First, ε is a positive number, and less than 10^{-4} times the distance between the points from Q if $\Delta_p = 2$. The definition of \mathcal{V} and \mathcal{W} requires the introduction of the set $Q_\varepsilon \subset \mathbb{R} \times I_*$, this the set of points with distance less than ε from Q . What is denoted by \mathcal{V} is an open neighborhood of (φ, ζ) in the C^∞ -Fréchet topology on the space of smooth maps from $(\mathbb{R} \times I_*) \setminus Q_\varepsilon$ to \mathbb{R}^2 . To describe \mathcal{W} , let (a_1, a_2, b_1, b_2) denote the coefficient functions that define D via (3-5).

What is denoted by \mathcal{W} is an open neighborhood of (a_1, a_2, b_1, b_2) in the C^∞ -Fréchet topology on the space of maps from the domain $(\mathbb{R} \times I_*) \setminus Q_\varepsilon$ to \mathbb{R}^4 .

Let $\mathcal{U} \subset \mathcal{O}$ denote the neighborhood of \hat{d} with the given indexing set. A point $\hat{d}' = (D', Q', (\gamma'_{p+}, \gamma'_{p-})\mathfrak{h}', D')$ from \mathcal{O} lies in \mathcal{U} when the conditions listed next are met. Each entry of D' has distance less than ε from the corresponding entry of D . Corresponding points from Q and Q' have distance less than $10^{-1/4}\varepsilon$ from each

other, and each of the endpoints of γ'_{p+} and γ'_{p-} has distance less than ε from the corresponding endpoint of the respective segments γ_{p+} and γ_p . The map h' from $(\mathbb{R} \times I_*) \setminus Q'$ to \mathbb{R}^2 lies in \mathcal{V} , and the map from this same domain to \mathbb{R}^4 given by sending any given (x, \hat{u}) to the D' version of (a_1, a_2, b_1, b_2) lies in \mathcal{W} .

It follows from what is said in Proposition 5.1, Lemmas 5.9 and 3.3 and Steps 3–4 in Section 3C that the assignment to any given data set $\hat{d} \in \mathcal{O}$ of its operator D defines a continuous section of a vector bundle over \mathcal{O} of the form $\text{Hom}(\mathcal{H}_1, \mathcal{H}_2)$ where \mathcal{H}_1 and \mathcal{H}_2 are Banach space bundles over \mathcal{O} . Let \mathcal{D} denote here this section of $\text{Hom}(\mathcal{H}_1, \mathcal{H}_2)$. There is the corresponding determinant line bundle, $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{O}$.

The lemma that follows makes the salient observations about \mathcal{O} and $\mathbb{D}\text{et}(\mathcal{D})$.

Lemma 8.4 *Define the space \mathcal{O} as above.*

- *The line bundle $\mathbb{D}\text{et}(\mathcal{D}) \rightarrow \mathcal{O}$ is orientable.*
- *\mathcal{O} is path connected if $\Delta_p = 1$ and it has two path connected components if $\Delta_p = 2$. In the latter case the components are distinguished by whether the $+$ labeled point from Q is greater than or less than the $-$ labeled point.*

This lemma is proved in a moment. The next paragraph explains how this lemma leads to a canonical orientation for the versions of $\mathbb{D}\text{et}(\mathcal{D})$ that arise in Part 2 of this subsection.

Let $\hat{d} \in \mathcal{O}$ and let D denote \hat{d} 's operator. Proposition 5.7 guarantees that D has trivial cokernel and Δ_p -dimensional kernel. The paragraph subsequent to (5-11) at the end of Part 2 can be repeated to define a canonical basis for the kernel of D . Lemma 8.4 guarantees that the various $\hat{d} \in \mathcal{O}$ versions of this basis define a canonical isomorphism between $\mathbb{D}\text{et}(\mathcal{D})$ and the oriented product bundle $\mathcal{O} \times \mathbb{R}$. It is in this sense that the versions of $\mathbb{D}\text{et}(\mathcal{D})$ from Part 2 have a completely canonical orientation.

Proof of Lemma 8.4 To prove the first bullet, look first at Proposition 5.7 to see that each $\hat{d} \in \mathcal{O}$ version of D has Δ_p -dimensional kernel. This being the case, these kernels fit together as \hat{d} varies in \mathcal{O} to define a Δ_p -dimensional vector bundle $\ker(\mathcal{D}) \rightarrow \mathcal{O}$. The basis given in the paragraphs subsequent to (5-11) define an isomorphism from this bundle to the product bundle.

The proof of the second bullet has eight steps. The first step reviews some background material. In the $\Delta_p = 1$ case, the remaining seven steps construct a continuous path parametrized by $[0, 7]$ that starts at a given element $\hat{d} \in \mathcal{O}$ and ends at any given second

element $\hat{\delta}' \in \mathcal{O}$. If $\Delta_p = 2$, these same steps construct a continuous path parametrized by $[0, 6]$ that starts at a given element $\hat{\delta} \in \mathcal{O}$ and ends at any given second element $\hat{\delta}' \in \mathcal{O}$ with the following property: Let Q and Q' denote the respective $\hat{\delta}$ and $\hat{\delta}'$ versions of the \pm labeled points in \mathbb{R} . The labels of the largest points in Q and Q' agree. These last seven steps construct the desired path as an end-to-end concatenation of seven $[0, 1]$ -parametrized paths. By way of notation, each step writes the start point of its segment as $(D, Q, \gamma_{p+}, \gamma_{p-}, \mathfrak{h} = (\varphi, \zeta), D)$. They all write the point $\hat{\delta}'$ in similar fashion using primes to distinguish respective components that differ.

Step 1 Suppose that \mathcal{E} is an end in $\hat{\delta}$'s submanifold $C_{\mathfrak{h}}$ whose constant s slices converge to $\hat{\gamma}_p^+$ as $s \rightarrow \infty$. Use \mathcal{E}' to denote the corresponding end of $C_{\mathfrak{h}'}$. Parametrize both as in Section 5A and Proposition 5.1; use η and η' denote the respective parametrizing maps. The map η is defined in part by data $(c_{n=1}, \phi_{n=1})$, where $c_{n=1} \in \mathbb{R} \setminus 0$ and $\phi_1 \in \mathbb{R}/2\pi\mathbb{Z}$. The constant c_1 can be assumed positive, for if not, there is an equivalent parametrization with ϕ_1 replaced by $\phi_1 + \pi$ and c_1 replaced by $-c_1$. Let $c'_1 > 0$ and ϕ'_1 denote the corresponding η' versions of these parameters with $c'_1 > 0$ also. There are corresponding versions of c_1 and c'_1 when the respective constant s slices of \mathcal{E} and \mathcal{E}' converge to $\hat{\gamma}_p^-$ as $s \rightarrow \infty$.

Step 2 The first path moves $\hat{\delta}$'s subset Q to the corresponding $\hat{\delta}'$ subset Q' . Such a path is readily constructed from a compactly supported isotopy of $\mathbb{R} \times I_*$ that moves only the \mathbb{R} coordinate of any given point, moves a neighborhood of each point from Q by a rigid translation, and takes Q to Q' . The $\tau = 1$ member of this path is now denoted by $\hat{\delta}$.

Step 3 This step constructs a $[0, 1]$ -parametrized path $\tau \rightarrow \hat{\delta}_\tau$ in \mathcal{O} whose $\tau = 0$ member is $\hat{\delta}$ and whose $\tau = 1$ member is given by $\hat{\delta}_1 = (D, Q, \gamma_{p+}, \gamma_{p-}, \mathfrak{h}_1 = (\varphi_1, \zeta_1), D_1)$ where $(\varphi_1, \zeta_1) = (\varphi, \zeta)$ except in very small radius disks about the point or points in Q . In a somewhat small radius disk $(\varphi_1, \zeta_1) = (\varphi', \zeta')$. Likewise, $D_1 = D$ on the first disk and $D_1 = D'$ on the smaller radius disk. To do this write $C_{\mathfrak{h}}$ and $C_{\mathfrak{h}'}$ on corresponding ends \mathcal{E} and \mathcal{E}' using the maps η and η' as in Step 1. Let (c_1, ϕ_1) and \mathfrak{e}_1 denote the data that appears in Proposition 5.1's depiction of η , and let (c'_1, ϕ'_1) and \mathfrak{e}'_1 denote the corresponding η' data set. Write $\phi'_1 = \phi_1 + \iota$ with $\iota \in [0, 2\pi)$. Set $\phi_{1\tau} = \phi_1 + \tau\iota$ for $\tau \in [0, 1]$. Meanwhile, set $c_{1\tau} = c_1 + \tau(c'_1 - c_1)$ and set $\mathfrak{e}_{1\tau} = \mathfrak{e}_1 + \tau(\mathfrak{e}'_1 - \mathfrak{e}_1)$ so as to define

$$(8-12) \quad \eta_{*\tau} = \alpha_y(e^{-\lambda_1 s} + e_1, 0) + c_{1\tau} e^{-\lambda_{11} s} (\cos(\phi_+ - \phi_{1\tau}), r_{11} \sin n(\phi_+ - \phi_{1\tau})) + \mathfrak{e}_{1\tau}.$$

The arguments from the Steps 1–3 of the proof of Proposition 5.1 can be used in an almost verbatim fashion to find $s_* > 1$ and a smooth, $[0, 1]$ -parametrized family of maps, $\tau \rightarrow q_{*\tau}: [s_*, \infty) \times \mathbb{R}/2\pi\mathbb{Z} \rightarrow \mathbb{R}^2$, with the following three properties: First, $\eta_\tau = \eta_{*\tau} + q_{*\tau}$ obeys (5-7). Second, $|\eta_\tau - \eta_{*\tau}| \leq \frac{1}{1000} \min\{c_1, c_2\}e^{-\lambda_{11}s_*}$. Third, $q_{*0} = q_{*1} = 0$.

Given the family $\{\eta_\tau\}_{\tau \in [0,1]}$, use what is said in Lemma 5.3 to write each $\tau \in [0, 1]$ version as a map $(x, \hat{u}) \rightarrow (\varphi_\tau, \zeta_\tau)|_{(x,\hat{u})}$ with image \mathbb{R}^2 and domain the complement of the given point from Q in a small radius disk in $\mathbb{R} \times I_*$ about this point. Use r to denote the radius of this disk and use χ_r to denote the function on $\mathbb{R} \times I_*$ given by the rule $(x, \hat{u}) \rightarrow \chi(\frac{4}{r}((x - y)^2 + \hat{u}^2)^{1/2} - 1)$. If Q has two points, do this for both.

Extend the corresponding family of maps defined in the radius r disk or disks about the points in Q over the whole of $(\mathbb{R} \times I_*) \setminus Q$ as h on the complement of the disk or disks and as the relevant $(\varphi_\tau, \zeta_\tau)$ version of $h_\tau = (1 - \chi_\tau)(\varphi, \zeta) + \chi_\tau(\varphi_\tau, \zeta_\tau)$ in each disk. Use $\{h_\tau\}_{\tau \in [0,1]}$ to denote this family of maps from $(\mathbb{R} \times I_*) \setminus Q$ to \mathbb{R}^2 and for each $\tau \in [0, 1]$. Use D_τ to denote the h_τ version of (3-9) and use \hat{d}_τ to denote $(D, Q, \gamma_{p+}, \gamma_{p-}, h_\tau, D_\tau)$. The assignment $\tau \rightarrow \hat{d}_\tau$ defines a continuous path in \mathcal{O} with the desired $\tau = 0$ and $\tau = 1$ members. Use $\hat{d} \in \mathcal{O}$ henceforth to denote the $\tau = 1$ member of this family.

Step 4 Use cutoff functions in the manner of Parts 1–3 of Section 6A to construct a continuous, $[0, 1]$ -parametrized path $\tau \rightarrow \hat{d}_\tau = (D, Q, \gamma_\tau, \gamma_{p-}, h_\tau, D_\tau)$ in \mathcal{O} whose $\tau = 0$ member is \hat{d} and whose $\tau = 1$ member is such that γ_{p-} and γ_1 have the same respective ϕ angles at their endpoints, but are such that the respective change in ϕ differs by -2π . Meanwhile, each $\tau \in [0, 1]$ version of h_τ agrees with h in the radius $\frac{1}{2}$ disk or disks centered on the points in Q . In addition, the $\hat{u} = -R - \frac{1}{2} \ln z_*$ boundary values of h_1 have constant $\hat{\phi}$ component, and the corresponding $\hat{u} = R + \frac{1}{2} \ln z_*$ component winds once around $\mathbb{R}/2\pi\mathbb{Z}$ in the anticlockwise direction as x varies from $-\infty$ to ∞ . By way of a parenthetical remark, the Ψ_p -image of the graph in $\mathbb{R} \times \mathcal{X}$ of h_τ can be guaranteed J -holomorphic only where $|s| \gg 1$. The needed modifications to what is done in Parts 1–3 of Section 6A are straightforward and are omitted. What is denoted by D_τ is the h_τ version of the operator in (3-9). Use $\hat{d} \in \mathcal{O}$ now to denote the $\tau = 1$ member of this path.

Step 5 Cutoff functions in the manner of Parts 1–3 of Section 6A are now used to construct a continuous, $[0, 1]$ -parametrized path $\tau \rightarrow \hat{d}_\tau = (D_\tau, Q, \gamma_\tau, \gamma_{p-, \tau}, h_\tau, D_\tau)$ in \mathcal{O} whose $\tau = 0$ member is \hat{d} and whose $\tau = 1$ member is such that $D_{\tau=1} = D'$.

Meanwhile, the ϕ angles of the boundary points of both $\gamma_{p-, \tau}$ and γ_τ are independent of τ ; Lemma II.2.2 is invoked to arrange this. The restriction of h_τ to a very small disk or disks centered on the point or points in Q is also independent of τ . As in the previous steps, D_τ denotes the h_τ version of (3-9). Use $\hat{\delta}$ now to denote the $\tau = 1$ member of this path.

Step 6 The constructions in Parts 1–3 of Section 6A are used yet again, this time to construct a continuous, $[0, 1]$ -parametrized path $\tau \rightarrow \hat{\delta}_\tau = (D', Q, \gamma_\tau, \gamma_{p-, \tau}, h_\tau, D_\tau)$ which is such that $\gamma_{p-, \tau=1} = \gamma'_{p-}$ and $\gamma_{\tau=1}$ is the γ'_{p-} analog of what is denoted by γ_1 in Step 4. This is done by moving the endpoints while invoking Lemma II.2.2. The map h_τ restricts to a very small radius disk or disks centered on the point or points in Q to be independent of τ . As before, D_τ denotes the h_τ version of (3-9). Use $\hat{\delta} \in \mathcal{O}$ henceforth to denote the $\tau = 1$ member of this fifth segment.

Step 7 The construction in Step 4 is run in reverse to construct a $[0, 1]$ -parametrized path $\tau \rightarrow \hat{\delta}_\tau$ in \mathcal{O} that moves γ_{p+} so that $\hat{\delta}_1 = (D', Q, \gamma'_{p+}, \gamma'_{p-}, h_1, D_1)$ with the path such that each $\tau \in [0, 1]$ member of h_τ is again independent of τ on some small radius disk or disks centered on the point or points in Q . The operator D_τ is the h_τ version of (3-9). As in the previous steps, use $\hat{\delta} \in \mathcal{O}$ to denote the $\tau = 1$ version of this path.

Step 8 This final leg of the path, denoted by $[0, 1] \rightarrow \hat{\delta}_\tau$, is a family of data sets that have the form $(D', Q, \gamma'_{p+}, \gamma'_{p-}, h_\tau, D_\tau)$. The $\tau = 1$ member is $\hat{\delta}'$. The family is defined using a suitable 1-parameter family of cutoff functions to homotope h to h' . The details are omitted as they contain no novelties. □

8E Canonical orientations for $\mathcal{M}_{p\pm}$

Proposition II.3.4 introduces for each $p \in \Lambda$ a pair of moduli spaces, \mathcal{M}_{p+} and \mathcal{M}_{p-} , whose constituents are embedded disks in the part of the $\hat{u} = 0$ locus of $\mathbb{R} \times \mathcal{H}_p$ where $1 - 3 \cos^2 \theta < 0$. Those in \mathcal{M}_{p+} sit where $\cos \theta > \frac{1}{\sqrt{3}}$ and their constant s slices converge as $s \rightarrow \infty$ in an isotopic fashion to $\hat{\gamma}_p^+$. Those in \mathcal{M}_{p-} sit where $\cos \theta < -\frac{1}{\sqrt{3}}$ and their constant s slices converge as $s \rightarrow -\infty$ to $\hat{\gamma}_p^-$. What follows talks solely about \mathcal{M}_{p+} as the \mathcal{M}_{p-} story is identical save for changing θ to $\pi - \theta$.

Let $C \in \mathcal{M}_{p+}$ denote a given curve. By way of a reminder from Section II.3C, the curve C is invariant under the action of S^1 that rotates the angle ϕ on \mathcal{H}_p and so the vector field ∂_ϕ is tangent to C . The vector field depicted in (II.3-10) is proportional to $J\partial_\phi$ and so is also tangent to C also. This being the case, C is foliated by the integral curves of the latter vector field except for the single point on C where it and ∂_ϕ

are zero. This point is the $\theta = 0$ point on C and the infimum of s on C . The latter point is the only critical point of s on ϕ . Use $y \in \mathbb{R}$ to denote this point. The association to each curve in \mathcal{M}_{p_+} of the minimum of s on the curve defines an \mathbb{R} -equivariant diffeomorphism from \mathcal{M}_{p_+} to \mathbb{R} .

Let $U_+ \subset \mathcal{H}_p$ denote the tubular neighborhood of $\hat{\gamma}_p^+$ that is described in Part 2 of Section 5A. The $s \gg y$ part of the given curve C lies in $\mathbb{R} \times U_+$ and so can be described using the coordinates $(s_+, \phi_+, \theta_+, u_+)$ as defined in (5-5). As such, it appears as the graph of a function $(s_+, \phi_+) \rightarrow (s_+, \phi_+, a, b)$, where $\eta = (a, b)$ is given by (5-8) with $\alpha = \alpha_y < 0$.

Since C is not \mathbb{R} -invariant, the normal projection of the vector field ∂_s along C supplies a canonical element to the kernel of C 's version of the operator \mathcal{D}_C . Denote the latter by η_C . Since the vector field in (3-10) is not proportional to ∂_s along C , this canonical element η_C is nowhere zero.

To say more about η_C , use (5-5) with (5-8) to identify the normal bundle to C along $C \cap (\mathbb{R} \times U_+)$ with the product bundle using the 1-forms $(d\theta_+, du_+)$. Granted this identification, and with $C \cap (\mathbb{R} \times U_+)$ parametrized as a graph in the manner just described, it follows from (5-8) that η_C appears as a map from the $s_+ \gg 1$ part of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 that can be written as

$$(8-13) \quad (s_+, \phi_+) \rightarrow -\lambda_+ \alpha_y (e^{-\lambda_1 s_+} + \epsilon, 0),$$

where $|\epsilon| \leq |\alpha_y| e^{-(\lambda_1 + 1/c_0)s_+}$.

A parenthetical remark subsequent to Proposition II.3.4 asserts that the cokernel of \mathcal{D}_C is trivial. This assertion is proved in a moment. It implies that the kernel of \mathcal{D}_C is 1-dimensional, this being the span of η_C . The identification between the kernel of \mathcal{D}_C and the tangent space at C to \mathcal{M}_{p_+} gives the canonical section η_C , and this section defines the desired canonical orientation.

To see why $\text{cokernel}(\mathcal{D}_C) = 0$, first identify the cokernel with the kernel of C 's formal L^2 adjoint. This done, use the previously described parametrizations of C and C 's normal bundle on $C \cap (\mathbb{R} \times U_+)$ to view an element in the cokernel of this adjoint operator as a map from the $s_+ \gg 1$ portion of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 . To see what such an element looks like, view \mathcal{D}_C on this part of C as an operator of the form $\mathfrak{D}_0 + \mathfrak{d}$ with \mathfrak{D}_0 as in (5-1) and with \mathfrak{d} a first-order operator on the space of maps from this large s_+ portion of $\mathbb{R} \times \mathbb{R}/2\pi\mathbb{Z}$ to \mathbb{R}^2 with ϕ -invariant coefficients that are bounded in absolute value by $c_0 |\alpha_y| e^{-(\lambda_1 + 1/c_0)s_+}$. With \mathcal{D}_C written as $\mathfrak{D}_0 + \mathfrak{d}$, an element in

the kernel of \mathcal{D}_C 's formal L^2 adjoint can be written on C 's intersection with $\mathbb{R} \times U_+$ as

$$(8-14) \quad (s_+, \phi_+) \rightarrow c_0(0, e^{-\lambda_2 s_+} + \epsilon_0) + c_n(\eta_{n-}|_{(-s_+, \phi_+)} + \epsilon_n),$$

where the notation is as follows: First, $c_0 \in \mathbb{R}$, $c_n \in \mathbb{R}$ and one of these is not zero. Meanwhile, η_{-n} is as defined in (5-4). What is written as ϵ_0 in (8-14) is ϕ -independent and obeys $|\epsilon_0| \leq c_0 e^{-(\lambda_2+1/c_0)s_+}$, and what is written as ϵ_n obeys $|\epsilon_n| \leq c_0 e^{-(\lambda_{2n}+1/c_0)s_+}$.

Granted that a cokernel element is a section of $N \otimes T^{0,1}C$, and granted that C is a disk, the claim that the cokernel of \mathcal{D}_C is trivial follows from (8-14) because the latter forces any nonzero element in the kernel of C 's adjoint to vanish at some point on C with positive local degree. This sort of vanishing is not possible by virtue of the fact that the formal L^2 adjoint of \mathcal{D}_C differs from that of $\bar{\partial}$ by a zeroth-order endomorphism.

8F Canonical orientations for the $I_{\text{ech}} = 1$ moduli spaces

Assume that the defining data for the geometry of Y and $\mathbb{R} \times Y$ are such that what is said in Sections 1–7 and 8A–8E hold. This subsection defines a canonical orientation for the $I_{\text{ech}} = 1$ moduli spaces, $\{\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}}$. The desired orientations are defined with the help of a given orientation for the real line bundle over \mathcal{A}_{HF1} whose fiber over any given surface S is the determinant line for the operator D_S . This line bundle is denoted by $\mathbb{D}\text{et}(\hat{D})$. A given orientation for $\mathbb{D}\text{et}(\hat{D})$ is assumed in the four parts that follow. Also needed is a chosen ordering for the set Λ . The resulting ordered set is written as $\{p_1, \dots, p_G\}$ when the ordering is relevant.

Part 1 Fix a pair $(\hat{\Theta}', \hat{\Theta}) \in \hat{\mathcal{Z}}_{\text{ech}, M} \times \hat{\mathcal{Z}}_{\text{ech}, M}$ and write these two elements respectively as $((\hat{v}', k'), (\xi'_p, O'_p)_{p \in \Lambda})$ and $((\hat{v}, k), (\xi_p, O_p)_{p \in \Lambda})$. It follows from Propositions 8.1 and 8.2 that the moduli space $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is nonempty if and only if one of the two conditions listed in the upcoming (8-15) hold. By way of notation, (8-15) introduces $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$ to denote the space of Lipschitz submanifolds with the following three properties: If $S \in \mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$, then the constant s slices of S converge in an isotopic fashion as $s \rightarrow \infty$ to the arcs that form \hat{v} , and they converge in an isotopic fashion as $s \rightarrow -\infty$ to the arcs that form \hat{v}' . In addition, the surface S has intersection number $k - k'$ with the arc $\gamma^{(z_0)}$. Finally, the operator D_S has Fredholm

index 1. The following also uses Δ_p to denote the number of elements in a given $p \in \Lambda$ version of O_p :

- (8-15) • $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k)) \neq \emptyset$ and each $p \in \Lambda$ version of (\mathfrak{k}_p, O_p) equals (\mathfrak{k}'_p, O'_p) .
- $(\hat{v}', k') = (\hat{v}, k)$ and there exists precisely one $p \in \Lambda$ such that $(\mathfrak{k}_p, O_p) \neq (\mathfrak{k}'_p, O'_p)$. In this case, $\Delta'_p = \Delta_p - 1$ and one of the following holds:
 - (a) $\mathfrak{k}'_p = \mathfrak{k}_p$.
 - (b) $\mathfrak{k}'_p = \mathfrak{k}_p \pm 1$.

The upcoming Part 2 specifies the desired orientations for the case of the first bullet in (8-15), Part 3 considers the case of item (a) of the second bullet, and Part 4 speaks to the case of item (b) of the second bullet.

Part 2 Suppose that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is described by the first bullet in (8-15). Any given element $\vartheta \in \mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is a disjoint union of components with some union from a version of \mathcal{M}^* that is defined by a surface $S \in \mathcal{A}_{\text{HF1}}$. The other components are \mathbb{R} -invariant cylinders of the form $\mathbb{R} \times \hat{\gamma}_p^+$ or $\mathbb{R} \times \hat{\gamma}_p^-$ for various $p \in \Lambda$. The contribution to ϑ from any given p depends on O_p ; either none, one or both of these p -labeled cylinders can be present.

Let $\mathcal{C} = (1, C)$ denote \mathcal{M}^* part of ϑ . This submanifold may also be a union of components, but in any event, precisely one such component is not \mathbb{R} -invariant. In any event, the tangent space to the curve C is canonically identified with the kernel of the corresponding version of the operator \hat{D}_C and it is therefore oriented by a choice of orientation for the line $\mathbb{D}\text{et}(\hat{D}_C)$. The desired orientation for the latter is supplied by (8-9) using the given orientation for $\mathbb{D}\text{et}(D_S)$ and the canonical orientations for the line $\mathbb{D}\text{et}(D_C)$ given in Section 8C.

Part 3 Suppose that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is described by item (a) of the second bullet in (8-15). This version of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ contains but a single \mathbb{R} -orbit. Let ϑ denote a given point on this orbit. The element ϑ has some union of \mathbb{R} -invariant components that define an element in a version of \mathcal{M}^* . The latter is defined by an HF-cycle. Part 5 in Section II.2B associates an *orientation sign*, either $+1$ or -1 , to each integral v in this cycle. Let N_+ denote the number of positive orientation signs.

The given point ϑ also contains a union of \mathbb{R} -invariant cylinders, each of the form $\mathbb{R} \times \hat{\gamma}_{p'}^+$ or $\mathbb{R} \times \hat{\gamma}_{p'}^-$ for various $p' \in \Lambda \setminus p$. The contribution to ϑ from any given such p' depends on $O_{p'}$; either none, one or both of these p -labeled cylinders can be

present. The nature of the remaining components depends on which of the cases listed below occurs:

- (8-16) • $O'_p = \{0\}$ and $O_p = \{1\}$.
 • $O'_p = \{-1\}$ and $O_p = \{1, -1\}$.
 • $O'_p = \{0\}$ and $O_p = \{-1\}$.
 • $O'_p = \{1\}$ and $O_p = \{1, -1\}$.

In the case of the first bullet, the pair p contributes to ϑ an element from Proposition II.3.4's moduli space \mathcal{M}_{p+} ; and in the case of the second bullet, it contributes such an element and also the cylinder $\mathbb{R} \times \widehat{\gamma}_p^-$. In either case, let C_+ denote the element from \mathcal{M}_{p+} . The tangent space at ϑ of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is canonically isomorphic to that of \mathcal{M}_{p+} at C_+ . Meanwhile, $T\mathcal{M}_{p+}$ has its canonical orientation from Section 8E. Use σ to denote the orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ that comes from this canonical orientation for $T\mathcal{M}_{p+}$. This may or may not be the desired orientation. To say if it is or not, introduce $k \in \{1, \dots, G\}$ to denote the label for p when Λ is written as $\{p_1, \dots, p_G\}$ and introduce N to denote $N = \sum_{1 \leq k' < k} \Delta_{k'}$. The desired orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is $(-1)^{N_+ + N} \sigma$.

In the case of the third bullet, the pair p contributes to ϑ an element from Proposition II.3.4's moduli space \mathcal{M}_{p-} ; and in the case of the fourth bullet, it contributes such an element and also the cylinder $\mathbb{R} \times \widehat{\gamma}_p^+$. In either case, let C denote the element from \mathcal{M}_{p-} . The tangent space at ϑ of $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is canonically isomorphic to that of \mathcal{M}_{p-} at C . The tangent bundle to \mathcal{M}_{p-} also has a canonical orientation from Section 8E. Use σ now to denote the orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ that comes from this canonical orientation for $T\mathcal{M}_{p-}$. Reintroduce $k \in \{1, \dots, G\}$ to denote the label for p when Λ is written as $\{p_1, \dots, p_G\}$ and $N = \sum_{1 \leq k' < k} \Delta_{k'}$. The desired orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is $(-1)^{N_+ + N} \sigma$ in the case of the third bullet in (8-16) and it is $(-1)^{N_+ + N + 1} \sigma$ in the case of the fourth bullet.

Part 4 This part deals with item (b) of the second bullet in (8-15). Consider first the case where $\xi'_p = \xi_p + 1$. In this case, either the first or the second bullet in (8-16) holds. In any event, a given ϑ from $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ consists of a union of components. Some subset of these define an element $\mathcal{C} = (1, C)$ from a version of \mathcal{M}^* . The pair p contributes the \mathbb{R} -invariant cylinder $\mathbb{R} \times \widehat{\gamma}_p^-$ if and only if the second bullet in (8-16) is relevant. The various $p' \in \Lambda \setminus p$ contribute either none, one or both \mathbb{R} -invariant cylinders from the set $\{\mathbb{R} \times \widehat{\gamma}_{p'}^-, \mathbb{R} \times \widehat{\gamma}_{p'}^+\}$, this depending as usual on $O_{p'}$. This understood, the tangent space to $\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ at ϑ is canonically isomorphic to

the kernel of the operator \widehat{D}_C . The latter is oriented via (8-9) using the given orientation for the relevant version of S and the canonical orientations supplied by Sections 8C and 8D for $\mathbb{D}et(\mathcal{D}_C)$. Note in this regard that S is \mathbb{R} -invariant and so D_S has trivial kernel and cokernel. Use σ to denote the orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ at ϑ that comes via the aforementioned identification with the kernel of \widehat{D}_C . Use k again to denote the label for \mathfrak{p} when Λ is written as $\{1, \dots, G\}$ and use N again to denote $\sum_{1 \leq k' < k} \Delta_{k'}$. The desired orientation for ϑ 's component of $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is $(-1)^{N_+ + N} \sigma$.

Suppose next that $\mathfrak{k}'_{\mathfrak{p}} = \mathfrak{k}_{\mathfrak{p}} - 1$. The story here is almost identical to that just told but for two salient and very much related changes. First, either the third or the fourth bullets in (8-16) hold. In the case of the fourth bullet the element ϑ contains the \mathbb{R} -invariant cylinder $\mathbb{R} \times \widehat{\gamma}_{\mathfrak{p}}^+$. In either case, let σ again denote the orientation for $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ at ϑ that comes via the canonical identification with the kernel of the relevant version of \widehat{D}_C with it understood that the kernel of the latter is oriented using (8-9) as before. Reintroduce the integer N . The desired orientation for ϑ 's component of $T\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})$ is $(-1)^{N_+ + N} \sigma$ if the third bullet of (8-16) is relevant, and it is $(-1)^{N_+ + N + 1} \sigma$ if the fourth bullet in (8-16) is relevant.

8G Coherent orientations

The definition of the Heegaard Floer differential requires the specification of an orientation for the low-dimensional components of \mathcal{A}_{HF} . There are constraints on the choice that are described in Section 6 of [10]. A choice that obeys the constraints is said to be a *coherent system of orientations*. The definition of the embedded contact homology differential likewise requires the specification of suitably constrained orientations for the low-dimensional components of \mathcal{M}_{ech} . Orientations that obey the latter constraints are also said to constitute a coherent system. The constraints are given in Section 9.5 of [7]. Proposition 8.5 in the upcoming Part 2 of this subsection makes a precise the assertion that a coherent system of orientations for \mathcal{A}_{HF} leads to a suitably compatible coherent system of orientations for \mathcal{M}_{ech} . Part 1 of the subsection sets up the needed background information.

Part 1 A closed integral curve of v is said to be hyperbolic if the associated linearized return map in $SL(2; \mathbb{R})$ has two real eigenvalues with neither equal to 1 or -1 . The integral curve is said to be positive when these eigenvalues are positive.

As noted previously, Section 9.5 in [7] defines what is meant by a coherent orientation for the components of \mathcal{M}_{ech} . There are four constraints which are labeled (OR1)–(OR4)

in this section of [7]. Those labeled (OR1) and (OR4) are normalization constraints that set the orientation in specific instances. The salient constraints are those expressed by (OR2) and (OR3). The former asserts that the orientation should be compatible with end-to-end concatenation of the subvarieties in \mathcal{M}_{ech} . The condition expressed by (OR3) constrains the orientation over the components of \mathcal{M}_{ech} whose elements consist of disjoint unions of two or more submanifolds. This (OR3) constraint requires an a priori choice of ordering for the ends of any given element whose constant s slices converge as $s \rightarrow \infty$ to positive hyperbolic integral curves of v . Also needed is a choice of ordering for those ends that converge as $s \rightarrow -\infty$ to positive hyperbolic integral curves of v .

Section 8F describes orientations for $\{\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}}$, these being the 1-dimensional components of \mathcal{M}_{ech} . The definition in Section 8F makes no reference to an ordering of the relevant ends of the constituent submanifolds. Even so, Proposition 8.5 refers to the Section 8F orientations when describing a coherent system of orientations for \mathcal{M}_{ech} . This referral implicitly invokes the ordering given in a moment for the ends of the elements in $\{\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}}$.

It proves useful to first define an ordering for the positive hyperbolic integral curves of v in any given element from $\mathcal{Z}_{\text{ech}, M}$. To this end, let Θ denote a given element from $\mathcal{Z}_{\text{ech}, M}$ and let $\Theta_+ \subset \Theta$ denote the subset of positive, hyperbolic integral curves. Introduce $\Theta_{\text{ech}, M} \subset \Theta$ to denote the subset of closed integral curves of v that cross one or more of the handles $\{\mathcal{H}_p\}_{p \in \Lambda}$. Given $\gamma \in \Theta_{\text{ech}, M}$, let n_γ denote the smallest of the labels of those $p \in \Lambda$ with $\mathcal{H}_p \cap \gamma \neq \emptyset$. Order $\Theta_{\text{ech}, M}$ so that the corresponding ordered set of integers $\{n_\gamma\}_{\gamma \in \Theta_{\text{ech}, M}}$ is increasing. Use $\Theta_{\text{ech}, M}^+ \subset \Theta_{\text{ech}, M}$ to denote the corresponding ordered subset of positive hyperbolic elements. Let $\sigma_{\Theta, p}$ denote the subset from the set $\{\hat{\gamma}_p^+, \hat{\gamma}_p^-\}$ that come from Θ . If this set has two elements, order it as just written. Use this convention to order Θ^+ as

$$(8-17) \quad \Theta^+ = \{\Theta_{\text{ech}, M}^+, \sigma_{p_1}, \dots, \sigma_{p_G}\}.$$

Let $\hat{\Theta}'$ and $\hat{\Theta}$ denote elements in $\hat{\mathcal{Z}}_{\text{ech}, M}$ and let C denote a given element in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$. What follows directly describes the convention for the ordering of C 's ends for the case described by the first bullet in (8-15). The submanifold C may have more than one component, but only one is not \mathbb{R} -invariant. The remaining components are \mathbb{R} -invariant cylinders. Some union of the latter with the component that is not \mathbb{R} -invariant defines an ech-HF submanifold. Write this ech-HF submanifold as $C_1 \cup C_2$ with C_1 denoting the non- \mathbb{R} -invariant component. The desired ordering places the

ends of C_1 before those of C_2 . Meanwhile, the ends of C_1 are ordered amongst themselves so as to be consistent with the ordering of the relevant Θ or Θ' version of (8-17)'s set $\Theta_{\text{ech},M}^+$, and likewise those of C_2 . All of the remaining \mathbb{R} -invariant cylinders in C are from the set $\{\mathbb{R} \times \hat{\gamma}_p^+, \mathbb{R} \times \hat{\gamma}_p^-\}_{p \in \Lambda}$. Their ends are ordered after the ends from $C_1 \cup C_2$. The union of the respective $s \gg 1$ and $s \ll -1$ ends from these sorts of \mathbb{R} -invariant cylinders is then ordered so as to be consistent with the ordering given by the relevant version of (8-17)'s ordered set $\{\sigma_{p_1}, \dots, \sigma_{p_G}\}$.

The convention in the case described by the second bullet in (8-16) orders the relevant subset of $s \gg 1$ ends of C so as to give the ordering in (8-17), and likewise for the $s \ll -1$ ends of C .

Part 2 This part explains how a coherent system of orientations for \mathcal{A}_{HF} leads to one for \mathcal{M}_{ech} and in particular for the components of $\{\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}}$. The assertion that such is the case is given by Proposition 8.5. The proposition refers to the notion from Section 2A of a weakly compact $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$. As noted in [10] and implied by Lemma II.6.6, the tangent space to \mathcal{A}_{HF} at any given Lipschitz submanifold S has a canonical identification with the kernel of the corresponding operator D_S . This being the case, an orientation for the tangent space of \mathcal{A}_{HF} is neither more nor less than an orientation for the real line bundle $\mathbb{D}\text{et}(\hat{D}) \rightarrow \mathcal{A}_{\text{HF}}$.

Proposition 8.5 *Suppose that a coherent system of orientations has been chosen for \mathcal{A}_{HF} and thus for the line bundle $\mathbb{D}\text{et}(\hat{D})$. There exists a weakly compact set $\mathcal{K} \subset \mathcal{A}_{\text{HF}}$ that contains all elements in $\mathcal{A}_{\text{HF}1}$ and has the following significance: Choose a \mathcal{K} -compatible data set $\mathbb{D} = (z_*, \delta, \chi_0, R)$ from the collection described in Proposition 7.2 for a suitable choice of κ , and choose the almost complex structure on $\mathbb{R} \times Y$ pursuant to the constraints in Proposition 7.1–7.3, 8.1 and 8.2. Use this data to define $\mathcal{Z}_{\text{ech},M}$ and \mathcal{M}_{ech} . There exists a coherent system of orientations for \mathcal{M}_{ech} whose restriction to $\{\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}}$ agrees with that defined in Section 8F using the orientations given by $\mathbb{D}\text{et}(\hat{D})$.*

Proof With the goal a proof of Theorem 1.1, coherent orientation systems are needed only for the one-dimensional components and certain two-dimensional components of \mathcal{A}_{HF} and \mathcal{M}_{ech} . To keep this long paper from being even longer, the coherence for the orientations of \mathcal{M}_{ech} will be verified only for these relevant components.

The various $p \in \Lambda$ versions of $\hat{\gamma}_p^+$ and $\hat{\gamma}_p^-$ are all positive hyperbolic closed integral curves of v . Proposition II.2.7 characterizes the other positive hyperbolic integral curves

of v that can appear in any given element from $\mathcal{Z}_{\text{ech}, \mathcal{M}}$. With the preceding understood, it is a straightforward task to verify that the orientations for the various elements in $\{\mathcal{M}_1(\widehat{\Theta}', \widehat{\Theta})\}_{\widehat{\Theta}', \widehat{\Theta} \in \widehat{\mathcal{Z}}_{\text{ech}, \mathcal{M}}}$ given in Section 8F using Part 1's ordering of the relevant ends of the constituent submanifolds obey the (OR3) constraint from Section 9.5 in [7]. The (OR3) constraint is the only relevant constraint on the 1-dimensional components of \mathcal{M}_{ech} .

Consider next the case where $C \subset \mathcal{M}_{\text{ech}}$ is a 2-dimensional component. There are two sorts of components to consider. The first is the union of cross-sectional 2-sphere in the handle \mathcal{H}_0 with \mathbb{R} -invariant cylinders. The sphere component is oriented using the convention (OR1) from Section 9.5 of [7]. The constraint in (OR3) of [7] is obeyed as long as the ordering for the set of $s \gg 1$ ends is the same as that for the set of $s \ll -1$ ends. The other constraints in Section 9.5 of [7] are not relevant.

The other sort of 2-dimensional component contains submanifolds that lie entirely in $\mathbb{R} \times (M_\delta \cup (\bigcup_{p \in \Lambda} \mathcal{H}_p))$. The only salient constraint to consider for this case is that given by (OR2). The concern with (OR2) arises when a 2-dimensional component of \mathcal{M}_{ech} has the following property: For fixed $\varepsilon > 0$, the component has two or more open subsets that are described by Proposition II.7.2 using distinct versions of Ξ of the form $\{Z_1, Z_2\}$ where $Z_1 = \{(S_1, u_1), \emptyset\}$ and $Z_2 = \{(S_2, u_2), \emptyset\}$ are such that S_1 and S_2 come from $\mathcal{A}_{\text{HF}1}$. What with Lemma 9.6 in [7], the end-to-end concatenation using any such version of Ξ orients the relevant component. The constraint (OR2) requires that all such orientations agree.

To see about (OR2), note that end-to-end concatenations of the pair S_1 and S_2 from any given such Ξ supply Lipschitz submanifolds in a 2-dimensional component of \mathcal{A}_{HF} . The construction is described in Appendix B of [10]. Moreover, what is said in this appendix implies a Heegaard Floer version of Lemma 9.6 in [7]. This analog orients the relevant 2-dimensional component of \mathcal{A}_{HF} given orientations for $\mathbb{D}\text{et}(S_1)$ and $\mathbb{D}\text{et}(S_2)$. With the preceding understood, suppose for a moment that all relevant versions of Ξ define in this way the same 2-dimensional component of \mathcal{A}_{HF} . The corresponding set of orientations for this component will agree if the orientations for the components of \mathcal{A}_{HF} constitute a coherent system. Meanwhile, an appropriate choice for the set \mathcal{K} has the following property: having chosen a 2-dimensional component of \mathcal{M}_{ech} , then there is but one 2-dimensional components of \mathcal{A}_{HF} .

Granted the preceding, a straightforward modification to what is said in Step 3 in the proof of Proposition 7.1 for a suitable version of (7-43) proves the following: if the

set \mathcal{K} and Proposition 7.2's constant κ are chosen appropriately, then the isomorphism given by (8-9) is compatible with respect to end-to-end concatenations of ech-HF subvarieties on the one hand and Heegaard Floer subvarieties from \mathcal{A}_{HF1} on the other. This fact implies that the (OR2) constraint in Section 9.5 of [7] is obeyed by the 2-dimensional components of \mathcal{M}_{ech} . \square

9 Proof of Theorem 1.1

This section uses the results from the previous two sections to prove the assertions made by the various bullets in Theorem 1.1.

9A The grading of the ech chain complex

This section addresses the assertion made by the fourth bullet of Theorem 1.1. The proof of this fourth bullet has five parts.

Part 1 Suppose that $(\hat{\Theta}', \hat{\Theta})$ is an ordered pair of elements from $\hat{\mathcal{Z}}_{\text{ech},M}$. The grading difference $\text{gr}_{\text{ech}}(\hat{\Theta}') - \text{gr}_{\text{ech}}(\hat{\Theta})$ is equal modulo the integer p_M to -1 times the ech index $I(\cdot)$ of a suitable relative 2-cycle. In particular, suppose that $k \in \{0, 1, 2, \dots\}$ and that C is an element in $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$. Then the aforementioned grading difference is equal to $-k$ modulo p_M . The formula in Definition 4.3 of [4] defines $I(C)$.

In some of the cases considered in the subsequent parts of the proof, the integer $I(C)$ is equal to the Fredholm index that is defined in equation (4.3) of [4], this being a consequence of the fact that all integral curves of v from elements in $\mathcal{Z}_{\text{ech},M}$ are hyperbolic. The equivalence between the ech index and the Fredholm index is used at times in the arguments that follow.

Part 2 Suppose that $(\hat{\Theta}', \hat{\Theta})$ lie over the same element, $\Theta \in \mathcal{Z}_{\text{ech},M}$. As such, they can be written respectively as $((\hat{v}', k'), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$ and $((\hat{v}', k), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$. If $k > k'$, then there is an ech-subvariety in $\mathcal{M}_{2(k-k')}(\hat{\Theta}', \hat{\Theta})$ which is a union of an \mathbb{R} -invariant part and $k - k'$ distinct spheres from Proposition II.3.1's moduli space \mathcal{M}_0 . It follows as a consequence that $\text{gr}_{\text{ech}}(\hat{\Theta}') - \text{gr}_{\text{ech}}(\hat{\Theta}) = 2(k' - k)$. Granted this fact, it is enough to consider the assertion of the fourth bullet only for those cases where $k = k'$.

Part 3 Suppose that $(\hat{\Theta}', \hat{\Theta})$ are given respectively as $((\hat{v}, k), (\mathfrak{k}_p, \mathcal{O}'_p)_{p \in \Lambda})$ and $((\hat{v}, k), (\mathfrak{k}_p, \mathcal{O}_p)_{p \in \Lambda})$. Suppose in addition that there is but one $p \in \Lambda$ where $\mathcal{O}'_p \neq \mathcal{O}_p$

and that one of the following conditions hold:

$$(9-1) \quad \begin{cases} O_p = \{1, -1\}, \\ O_p = \{-1\} \text{ and } O_{p'} = \{0\}, \\ O_p = \{1\} \text{ and } O_{p'} = \{0\}. \end{cases}$$

Let k denote the integer $\text{gr}(O_p) - \text{gr}(O'_p) \in \{1, 2\}$. There is in this case an element in $\mathcal{M}_k(\hat{\Theta}', \hat{\Theta})$, this being a union of \mathbb{R} -invariant cylinders with either one element from Proposition II.3.4's moduli space \mathcal{M}_{p+} , or one element from the latter's \mathcal{M}_{p-} in the case $k = 1$, or one from each in the case $k = 2$. It follows as a consequence that the ech grading difference $\text{gr}_{\text{ech}}(\hat{\Theta}') - \text{gr}_{\text{ech}}(\hat{\Theta})$ is equal to $\text{gr}(O'_p) - \text{gr}(O_p)$.

Given what was said in Part 3, repeated applications of this last observation justify the assertion that it is sufficient to consider the assertion of fourth bullet of Theorem 1.1 only for cases with $k = k'$ and with all $p \in \Lambda$ versions of O'_p and O_p equal.

Part 4 Suppose that $(\hat{\Theta}', \hat{\Theta})$ are given respectively as $((\hat{v}, k), (\xi'_p, O_p)_{p \in \Lambda})$ and $((\hat{v}, k), (\xi_p, O_p)_{p \in \Lambda})$. Suppose in addition that each $p \in \Lambda$ version of $O_p = \{0\}$, and that there is but one $p \in \Lambda$ with $\xi'_p \neq \xi_p$. Let $q \in L$ denote the exception and take $\xi'_q = \xi_q + 1$. Introduce $\hat{\Theta}''$ to denote $((\hat{v}, k), (\xi_q, O_q = \{1\}), (\xi_p, O_p)_{p \in \Lambda \setminus q})$. It follows from Proposition 8.1 that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta}'') \neq \emptyset$ and so $\text{gr}_{\text{ech}}(\hat{\Theta}') = \text{gr}_{\text{ech}}(\hat{\Theta}'') - 1$. Meanwhile, it follows from what is said in Part 3 that $\mathcal{M}_1(\hat{\Theta}, \hat{\Theta}'') \neq \emptyset$ also. Thus, $\text{gr}_{\text{ech}}(\hat{\Theta})$ also equals $\text{gr}_{\text{ech}}(\hat{\Theta}'') - 1$.

Repeated applications of this last observation justify the assertion that it is enough to consider the fourth bullet only for cases with $k = k'$ and with all $p \in \Lambda$ versions of (ξ_p, O_p) equal to (ξ'_p, O'_p) .

Part 5 Fix $n \in \{0, 1, 2, \dots\}$ and suppose that $(\hat{\Theta}', \hat{\Theta})$ are given respectively by $((\hat{v}, k), (\xi_p, O_p)_{p \in \Lambda})$ and by $((\hat{v}', k - n), (\xi_p, O_p)_{p \in \Lambda})$. It follows from Lemma 4.1 in [10] that there exists n and an almost complex structure J'_{HF} for $\mathbb{R} \times [1, 2] \times \Sigma$ with the following properties: First, both depend only on the Heegaard Floer data. Second, J'_{HF} obeys Lipshitz's requirements and those in Section II.6.1. Third, there is a J'_{HF} version of a Lipshitz submanifold, S' , whose constant s slices converge in an isotopic fashion to the arcs in \hat{v} as $s \rightarrow \infty$, and to the arcs in \hat{v}' as $s \rightarrow -\infty$. Fourth, this Lipshitz submanifold has intersection number n with the $f \in [1, 2]$ part of the curve $\gamma^{(z_0)}$ that is described in the fifth bullet of Part 2 in Section 1B. Fifth, the corresponding operator D'_S has trivial cokernel. Let ι denote the Fredholm index of D'_S .

Given that n , J'_{HF} and S depend only on the Heegaard Floer data, it follows that the data set (z_*, δ, χ_0, R) can be taken without loss of generality so that the following is true: there is an almost complex structure J' for $\mathbb{R} \times Y$ that is of the sort defined in Section 1C whose restriction to $\mathbb{R} \times M_\delta$ is J'_{HF} , and is such that Propositions 7.1–7.3, 8.1 and 8.2 can be invoked using the Lipschitz submanifold S' . In particular, Proposition 8.1 implies that $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta}) \neq \emptyset$.

This last observation implies that $\text{deg}_{\text{ech}}(\hat{\Theta}') - \text{deg}_{\text{ech}}(\hat{\Theta}) = \text{gr}_{\text{HF}}(\hat{v}') - \text{gr}_{\text{HF}}(\hat{v}) - 2n$. The assertion from of the fourth bullet follows from this and what is said in Parts 2–4.

9B The ech differential

This subsection proves the assertion made by the first bullet in Theorem 1.1. By way of a reminder the endomorphism on the chain complex $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech}, M})$ that defines the differential has the form depicted in (1-19); thus saying something about the differential requires saying something about the various integers from the relevant version of the set $\{N_{\hat{\Theta}', \hat{\Theta}}\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech}, M}}$. The definition of any given $N_{\hat{\Theta}', \hat{\Theta}}$ is reviewed in Part 1 of what follows. Parts 2–4 say what is needed about these integers to deduce the first bullet of Theorem 1.1.

Part 1 Fix $(\hat{\Theta}', \hat{\Theta}) \in \hat{\mathcal{Z}}_{\text{ech}, M} \times \hat{\mathcal{Z}}_{\text{ech}, M}$. The corresponding integer $N_{\hat{\Theta}', \hat{\Theta}}$ that is used in (1-19) to define the embedded contact homology differential is given by a sum that is indexed by the components of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})/\mathbb{R}$ whereby each component contributes either $+1$ or -1 . Whether $+1$ or -1 is determined by comparing two orientations of the given component. To say more, keep in mind that each component of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is a 1–dimensional manifold with a free action of \mathbb{R} , with the action given by the constant translations along the \mathbb{R} factor of $\mathbb{R} \times Y$. The generator of this action is a nowhere-zero vector field on each component. This vector field defines an orientation for each component, this denoted by $\hat{d}_{\text{ech}, \mathbb{R}}$. The second orientation is given by a coherent orientation for \mathcal{M}_{ech} . The latter orientation is denoted by $\hat{d}_{\text{ech}, Q}$. Write $\hat{d}_{\text{ech}, \mathbb{R}}$ as $N^C \hat{d}_{\text{ech}, Q}$ with $N^C \in \{1, -1\}$. The assignment $C \rightarrow N^C$ is a locally constant, $\{1, -1\}$ –valued function on $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$. The value of this function N^C on C 's component of $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is the component's contribution to the sum that defines $N_{\hat{\Theta}', \hat{\Theta}}$.

Part 2 There is an analogous definition of the endomorphism of $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z})$ that defines the differential for Heegaard Floer homology. To elaborate, any given endomorphism of $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z})$ is defined by its action on the generating set, and so by a

rule of the form

$$(9-2) \quad (\hat{v}', k') \rightarrow \sum_{(\hat{v}, k) \in \mathcal{Z}_{\text{HF}} \times \mathbb{Z}} M_{(\hat{v}', k'), (\hat{v}, k)}(\hat{v}, k),$$

where $M_{\cdot, \cdot}$ is in all cases an integer. The paragraphs that follow define these integers when the endomorphism in question is the differential ∂_{HF} .

Fix ordered pairs (\hat{v}', k') and (\hat{v}, k) of elements from $\mathcal{Z}_{\text{HF}} \times \mathbb{Z}$. Having done so, reintroduce the subspace $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k)) \subset \mathcal{A}_{\text{HF}}$ from Part 1 of Section 8F. The integer $M_{(\hat{v}', k'), (\hat{v}, k)}$ for the version of (9-2) that defines the differential for Heegaard Floer homology is a sum that is indexed by the components of $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$ with each component contributing either $+1$ or -1 .

To define these ± 1 contributions, note that $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$ has a finite set of components, with each a copy of \mathbb{R} . Each such copy of \mathbb{R} is a free orbit of the \mathbb{R} action that comes from the constant translations along the \mathbb{R} factor of $\mathbb{R} \times [1, 2] \times \Sigma$. The generator of this \mathbb{R} action orients each component. This orientation is denoted by $\hat{\partial}_{\text{HF}, \mathbb{R}}$. A second orientation is that supplied by a given coherent orientation for \mathcal{A}_{HF} . This orientation is denoted by $\hat{\partial}_{\text{HF}, \mathcal{Q}}$. Let $S \in \mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$ denote a given surface and write $\hat{\partial}_{\text{HF}, \mathbb{R}}$ at S as $Z^S \hat{\partial}_{\text{HF}, \mathcal{Q}}$ where $Z^S \in \{-1, 1\}$ is constant on the component of S in $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))/\mathbb{R}$. The value of Z^S is this component's contribution to $M_{(\hat{v}', k'), (\hat{v}, k)}$.

Part 3 Fix a coherent system of orientations for \mathcal{A}_{HF} to define the coefficients in the version of (9-2) that defines the Heegaard Floer differential. Use this same coherent system in Proposition 8.5 to define the coherent system of orientations that is used to define the embedded contact homology differential.

Fix a pair, $(\hat{\Theta}', \hat{\Theta})$, from $\hat{\mathcal{Z}}_{\text{ech}, M}$. The corresponding $N_{\hat{\Theta}', \hat{\Theta}}$ is zero unless $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is nonempty and therefore $(\hat{\Theta}', \hat{\Theta})$ is described by (8-15). Consider first the case given by the first bullet in (8-15). It follows from Proposition 7.2 that the elements in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ are labeled in part by the surfaces in $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$. Let S denote a given such surface and let $\mathcal{M}_1^S(\hat{\Theta}', \hat{\Theta}) \subset \mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ denote the corresponding subset. It follows directly from Propositions 8.1, 8.2 and 8.5 plus what is said in Part 5 of Section 8B that

$$(9-3) \quad \sum_{C \in \mathcal{M}_1^S(\hat{\Theta}', \hat{\Theta})/\mathbb{R}} N^C = Z^S.$$

This implies directly that $N_{\hat{\Theta}', \hat{\Theta}} = M_{(\hat{v}', k'), (\hat{v}, k)}$.

Write $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ as $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M}) = \mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes (\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathcal{O}))$ as done in Theorem 1.1. The conclusion of the preceding paragraph implies that writing $\mathbb{Z}(\widehat{\mathcal{Z}}_{\text{ech},M})$ in this way makes ∂_{ech} appear as $\partial_{\text{ech}} = \partial_{\text{HF}} + L$ with L acting solely on the $(\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathcal{O}))$ factor.

Part 4 The endomorphism L is defined by those $N_{\widehat{\Theta}', \widehat{\Theta}}$ with $(\widehat{\Theta}', \widehat{\Theta})$ as described in the second bullet of (8-15). It follows that L can be written as $\sum_{p \in \Lambda} L_p$ where L_p acts only on the $p \in \Lambda$ factor of $\mathbb{Z}(\mathbb{Z} \times \mathcal{O})$ in $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes (\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathcal{O}))$. This is because there is but one pair $p \in \Lambda$ with $(\xi_p, \mathcal{O}_p) \neq (\xi'_p, \mathcal{O}'_p)$. Moreover, the relevant version of L_p acts on the given generator (ξ_p, \mathcal{O}_p) to give an integer-weighted sum of generators with the weight being zero unless $\Delta'_p = \Delta_p - 1$ in which case either item (a) or item (b) in (8-15) must occur. This being the case, it follows from Propositions 8.1 and 8.2 with what is said in Part 5 of Section 8B that the corresponding integer weight is either 1 or -1 .

Consider first the case when item (a) of the second bullet is obeyed. Proposition 8.5 and what is said in Part 3 of Section 8B determine these signs:

- (9-4) • In the case of the first three bullets in (8-16), the sign is $(-1)^{N_+ + N}$.
 • In the case of the fourth bullet in (8-16), the sign is $(-1)^{N_+ + N + 1}$.

Suppose next that item (b) of the second bullet holds. Proposition 8.5 and what is said in Part 4 of Section 8F determine that the sign is again given by (9-4). Note in this regard that the fourth bullet in (8-16) can occur only if $\xi'_p = \xi_p - 1$.

What is written in (9-4) is consistent with what is claimed by Theorem 1.1 if and only if the integers N_+ are such that $(-1)^{N_+} = \varepsilon(-1)^{\text{deg}_{\text{HF}}(\widehat{v}, k)}$ with $\varepsilon \in \{-1, 1\}$ being independent of both \widehat{v} and k . Since the Heegaard Floer degree changes by an even integer as k varies, it is enough to verify that this is so for any given value of k . That such is the case follows from Proposition 4.8 in [10] and equation (9) in [10]. The latter expresses an equality that was derived by J Rasmussen in [17].

9C The endomorphisms from the second and third bullets of Theorem 1.1

The assertion made by the second bullet about the action of the \mathbb{U} -map on the chain complex $\mathbb{Z}(\mathcal{Z}_{\text{HF}} \times \mathbb{Z}) \otimes (\bigotimes_{p \in \Lambda} \mathbb{Z}(\mathbb{Z} \times \mathcal{O}))$ follows directly from what is said in the first paragraph of Part 2 in Section 9B. The assertions about the endomorphisms in the third bullet of Theorem 1.1 are discussed in the subsequent two parts of this subsection. The first part briefly reviews the definitions of the coefficients that appear in the corresponding versions of (1-19).

Part 1 This part of the subsection explains how to the endomorphisms in the third bullet of Theorem 1.1 are defined. To this end, let \hat{i} denote one of the cycles from the set $\{\hat{i}_p\}_{p \in \Lambda}$ and let $Q_{\text{ech}}^{\hat{i}}$ denote the corresponding endomorphism of $\mathbb{Z}(\hat{\mathcal{Z}}_{\text{ech},M})$. This endomorphism is described by a version of (1-19), and thus defined by the integers $\{N_{\hat{\Theta}', \hat{\Theta}}\}_{\hat{\Theta}', \hat{\Theta} \in \hat{\mathcal{Z}}_{\text{ech},M}}$. Let $(\hat{\Theta}', \hat{\Theta})$ denote a given pair from $\hat{\mathcal{Z}}_{\text{ech},M}$. The corresponding integer $N_{\hat{\Theta}', \hat{\Theta}}$ is nonzero only if $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ is nonempty. Each component in the latter set contributes an integer to a sum whose value is $N_{\hat{\Theta}', \hat{\Theta}}$. This understood, let $C \subset \mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ denote a given element, and let $[C]$ denote the corresponding relative 2-cycle in $H_2(Y; [\Theta] - [\Theta'])$ defined by the image of C in Y via the projection from $\mathbb{R} \times Y$. The cycle \hat{i} has been chosen so as to be disjoint from the integral curves of v that appear in elements from $\mathcal{Z}_{\text{ech},M}$, and so there is a well-defined pairing between \hat{i} and $[C]$ with values in \mathbb{Z} . Use $\langle \hat{i}, [C] \rangle$ to denote this pairing. Reintroduce the sign $N^C \in \{1, -1\}$ from Part 1 of Section 9C. The contribution of C 's component in $\mathcal{M}_1(\hat{\Theta}', \hat{\Theta})$ to the sum for $N_{\hat{\Theta}', \hat{\Theta}}$ is $\langle \hat{i}, [C] \rangle N^C$.

By way of comparison, what follows summarizes from Section 8 of [10] the Heegaard Floer version of the endomorphisms that are defined by the cycles from the set $\{\gamma^{(z_0)}, \{\hat{i}^{(z)}\}_{z \in \mathbb{Y} \setminus z_0}\}$. Let \hat{i} now denote one of the cycles from the latter set, and let $Q_{\text{HF}}^{\hat{i}}$ denote the corresponding endomorphism. This endomorphism is defined by a version of (9-2). A given coefficient $M_{(\hat{v}', k'), (\hat{v}, k)}$ is nonzero only if $\mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$ is nonempty. If so, then each component of this space contributes an integer to a sum whose value is $M_{(\hat{v}', k'), (\hat{v}, k)}$. Let S denote a given Lipschitz surface from this space. The image of S in $\Sigma \times [0, 1]$ via the projection from $\mathbb{R} \times \Sigma \times [0, 1]$ has a well-defined intersection pairing with \hat{i} , this denoted by $\langle \hat{i}, S \rangle$. Reintroduce $z^S \in \{1, -1\}$ from Part 1 of Section 9B. The component of S contributes $\langle \hat{i}, S \rangle z^S$ to the sum that computes $M_{(\hat{v}', k'), (\hat{v}, k)}$.

Part 2 Consider first the statements made by items (a) and (b) of the third bullet in Theorem 1.1. To this end, fix $\hat{\Theta}'$ and $\hat{\Theta}$ from $\hat{\mathcal{Z}}_{\text{ech},M}$. Given $S \in \mathcal{A}_{\text{HF1}}((\hat{v}', k'), (\hat{v}, k))$, reintroduce from Part 3 in Section 9B the subspace $\mathcal{M}_1^S(\hat{\Theta}', \hat{\Theta})$. Given the definitions in Part 1, the assertion made by item (b) of the third bullet in Theorem 1.1 follows from (9-3) if $\langle \hat{i}, [C] \rangle = \langle \hat{i}, S \rangle$ when $C \in \mathcal{M}_1^S(\hat{\Theta}', \hat{\Theta})$. The latter equality is a consequence of Proposition 8.2.

Consider next the statement made by item (c) of the third bullet in Theorem 1.1. To this end, fix $p \in \Lambda$ so as to see about the action of the $\hat{i} = \hat{i}_p$ version of $Q_{\hat{i}}$. The cycle \hat{i}_p is disjoint from $M_\delta \cup (\bigcup_{p' \in \Lambda \setminus p} \mathcal{H}_{p'})$ and as a consequence, it must act as $\mathbb{I}_p + \hat{i}_p$. This

being the case, at issue is the precise form for \hat{l}_p . The cycle \hat{l}_p lies in the $\cos \theta > \frac{1}{\sqrt{3}}$ of \mathcal{H}_p , and as a consequence any given integer $N_{\hat{\Theta}', \hat{\Theta}}$ from the Q_i version of (1-19) is zero unless $\hat{\Theta}'$ and $\hat{\Theta}$ are described by item (a) of the second bullet in (8-15) and the first two bullets in (8-16). With this point understood, item (c) of the third bullet in Theorem 1.1 follows directly from Propositions 8.1 and 8.5 plus what is said in Part 5 of Section 8B.

References

- [1] **D R Adams**, *Morrey spaces*, Springer (2015) MR
- [2] **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
- [3] **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
- [4] **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
- [5] **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 Gastesi, editors), volume II, Hindustan Book Agency, New Delhi (2010) 1022–1041 MR
- [6] **M Hutchings, M Sullivan**, *Rounding corners of polygons and the embedded contact homology of T^3* , Geom. Topol. 10 (2006) 169–266 MR
- [7] **M Hutchings, C H Taubes**, *Gluing pseudoholomorphic curves along branched covered cylinders, II*, J. Symplectic Geom. 7 (2009) 29–133 MR
- [8] **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
- [9] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, II: Reeb orbits and holomorphic curves for the ech/Heegaard Floer correspondence*, Geom. Topol. 24 (2020) 2855–3012
- [10] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. 10 (2006) 955–1096 MR Correction in 18 (2014) 17–30
- [11] **D McDuff, D Salamon**, *J-holomorphic curves and quantum cohomology*, University Lecture Series 6, Amer. Math. Soc., Providence, RI (1994) MR
- [12] **C B Morrey, Jr.**, *Multiple integrals in the calculus of variations*, Grundle Math. Wissen. 130, Springer (1966) MR

- [13] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. 159 (2004) 1027–1158 MR
- [14] **P Ozsváth, Z Szabó**, *An introduction to Heegaard Floer homology*, from “Floer homology, gauge theory, and low-dimensional topology” (D A Ellwood, P S Ozsváth, A I Stipsicz, Z Szabó, editors), Clay Math. Proc. 5, Amer. Math. Soc., Providence, RI (2006) 3–27 MR
- [15] **D Quillen**, *Determinants of Cauchy–Riemann operators on Riemann surfaces*, Funktsional. Anal. i Prilozhen. 19 (1985) 37–41 MR In Russian; translated in Funct. Anal. Appl. 19 (1985) 31–34
- [16] **H Rafeiro, N Samko, S Samko**, *Morrey–Campanato spaces: an overview*, from “Operator theory, pseudo-differential equations, and mathematical physics” (Y I Karlovich, L Rodino, B Silbermann, IM Spitkovsky, editors), Oper. Theory Adv. Appl. 228, Springer (2013) 293–323 MR
- [17] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) MR arXiv
- [18] **R Siefring**, *Relative asymptotic behavior of pseudoholomorphic half-cylinders*, Comm. Pure Appl. Math. 61 (2008) 1631–1684 MR
- [19] **CH 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
- [20] **CH Taubes**, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$* , Geom. Topol. 6 (2002) 657–814 MR

*Department of Mathematics, University at Buffalo
Buffalo, NY, United States*

*Institute of Mathematical Sciences, 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: 14 September 2016