

Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: Properties and existence

CLIFFORD HENRY TAUBES

This is the first of at least two articles that describe the moduli spaces of pseudoholomorphic, multiply punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$ as defined by a certain natural pair of almost complex structure and symplectic form. This article proves that all moduli space components are smooth manifolds. Necessary and sufficient conditions are also given for a collection of closed curves in $S^1 \times S^2$ to appear as the set of $|s| \rightarrow \infty$ limits of the constant $s \in \mathbb{R}$ slices of a pseudoholomorphic, multiply punctured sphere.

53D30; 53C15, 53D05, 57R17

1 Introduction

This is the first of at least two articles that describe the moduli spaces of multiply punctured, pseudoholomorphic spheres for a very natural symplectic form and compatible almost complex structure on $\mathbb{R} \times (S^1 \times S^2)$. In this regard, the symplectic form and attending almost complex structure arise when considering 4 dimensional, compact Riemannian manifolds with an associated self-dual harmonic 2–form. To elaborate, if the metric is suitably generic, then the zero locus of the harmonic form is an embedded union of circles and the harmonic 2–form defines a symplectic structure on the complement of this locus (see, for example, Honda [11] or Gay and Kirby [3]). In addition, the complement of any given component of the zero locus is in an open set that is diffeomorphic to $(0, \infty) \times (S^1 \times S^2)$. As explained in [17], the given self-dual 2–form can be modified on the complement of its zero locus so as to give a symplectic form on this complement that restricts to any of these $(0, \infty) \times (S^1 \times S^2)$ subsets as either the symplectic form from $\mathbb{R} \times (S^1 \times S^2)$ considered here, or that of its push-forward by a free, symplectic $\mathbb{Z}/2\mathbb{Z}$ action.

With the preceding understood, remark next that there is some evidence (see [15]) that pseudoholomorphic curves for certain almost complex structures, compatible with this new symplectic form, code information about the smooth structure on the underlying 4 dimensional manifold. And, if such is the case, then a program to decipher this code

will almost surely need knowledge of the multi-punctured sphere pseudoholomorphic curve moduli spaces on the whole of $\mathbb{R} \times (S^1 \times S^2)$. For example, these moduli spaces will arise in a definition of a smooth 4-manifold invariant that uses any sort of refined version of the Eliashberg–Givental–Hofer symplectic field theory [2]. (Some refinement would have to be made since the symplectic form that arises on $\mathbb{R} \times (S^1 \times S^2)$ comes from an overtwisted contact structure on $S^1 \times S^2$.)

This article provides an introduction to the multi-punctured sphere moduli spaces, a description of some of their local properties, an introduction to techniques used in the sequel article, and an existence proof for the various components. The afore-mentioned sequel describes the components of the multi-punctured sphere moduli spaces in great detail with the help of an explicit parametrization. The reader is also referred to a sort of prequel to this series, this the article [18] that describes the pseudoholomorphic disks, cylinders and certain of the 3-holed spheres in $\mathbb{R} \times (S^1 \times S^2)$.

Acknowledgements Before turning to the details, there is a debt to acknowledge: In hindsight, the approach in these articles most probably owes a great deal to the author's subconscious remembering of old conversations with both Helmut Hofer and Michael Hutchings.

The author is supported in part by the National Science Foundation.

1.A The symplectic and contact geometry of $\mathbb{R} \times (S^1 \times S^2)$

An introduction to the relevant geometry is in order. To start, introduce standard coordinates (s, t, θ, φ) for $\mathbb{R} \times (S^1 \times S^2)$ where s is the Euclidean coordinate for the \mathbb{R} factor, $t \in \mathbb{R}/(2\pi\mathbb{Z})$ is the coordinate for the S^1 factor and $(\theta, \varphi) \in [0, \pi] \times \mathbb{R}/(2\pi\mathbb{Z})$ are standard spherical coordinates for the 2-sphere factor. The symplectic form that is used here on $\mathbb{R} \times (S^1 \times S^2)$ comes as the 'symplectification' of a contact 1-form on $S^1 \times S^2$, this the 1-form

$$(1-1) \quad \alpha \equiv -(1 - 3 \cos^2 \theta) dt - \sqrt{6} \cos \theta \sin^2 \theta d\varphi.$$

To be explicit, here is the symplectic form:

$$(1-2) \quad \omega = d(e^{-\sqrt{6}s} \alpha).$$

Note that the convention is that the $s \rightarrow \infty$ end of $\mathbb{R} \times (S^1 \times S^2)$ is the concave side end and the $s \rightarrow -\infty$ is the convex side end. (The concave side end is the half that appears in the 4-manifold context.) It proves convenient at times to write the form ω as

$$(1-3) \quad \omega = dt \wedge df + d\varphi \wedge dh,$$

where

$$(1-4) \quad f \equiv e^{-\sqrt{6}s} (1 - 3 \cos^2 \theta) \quad \text{and} \quad h \equiv \sqrt{6} e^{-\sqrt{6}s} \cos \theta \sin^2 \theta.$$

The almost complex structure that defines here the notion of a pseudoholomorphic subvariety is specified by the relations

$$(1-5) \quad J \cdot \partial_t = g \partial_f \text{ and } J \cdot \partial_\varphi = \sin^2 \theta g \partial_h,$$

where $g = \sqrt{6} e^{-\sqrt{6}s} (1 + 3 \cos^4 \theta)$. This almost complex structure is ω -compatible. In fact, the form $g^{-1} \omega(\cdot, J(\cdot))$ on the tangent bundle of $\mathbb{R} \times (S^1 \times S^2)$ is the standard product metric, $ds^2 + dt^2 + d\theta^2 + \sin^2 \theta d\varphi^2$. As remarked earlier, this almost complex structure is not integrable.

Note that J is invariant under an $\mathbb{R} \times (S^1 \times S^1)$ subgroup of the product metric's group of isometries, $\mathbb{R} \times S^1 \times SO(3)$. Here, the \mathbb{R} factor in this subgroup acts as the constant translations along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times S^2)$, the first S^1 factor in the subgroup acts to rotate the S^1 factor of $\mathbb{R} \times (S^1 \times S^2)$, while the second S^1 factor rotates the 2-sphere about the axis where $\theta \in \{0, \pi\}$. Thus, the \mathbb{R} action is generated by the vector field ∂_s and the two S^1 actions are respectfully generated by the vector fields ∂_t and ∂_φ . This particular $S^1 \times S^1$ subgroup of the metric isometry group is denoted below as T .

1.B The pseudoholomorphic subvarieties

Following the lead of Hofer [4; 5; 6] and Hofer–Wysocki–Zehnder [8; 7; 10], a pseudoholomorphic subvariety in $\mathbb{R} \times (S^1 \times S^2)$ is defined here as follows:

Definition 1.1 A pseudoholomorphic subvariety $C \subset \mathbb{R} \times (S^1 \times S^2)$ is a non-empty, closed subset with the following properties:

- The complement in C of a countable, nowhere accumulating subset is a 2-dimensional submanifold whose tangent space is J -invariant.
- $\int_{C \cap K} \omega < \infty$ when $K \subset \mathbb{R} \times (S^1 \times S^2)$ is an open set with compact closure.
- $\int_C d\alpha < \infty$.

A pseudoholomorphic subvariety is said to be 'reducible' if the removal of a finite set of points makes a set with more than one connected component.

Note that [18] uses the term 'HWZ variety' for what is defined here to be a pseudoholomorphic subvariety.

If $C \subset \mathbb{R} \times (S^1 \times S^2)$ is an irreducible pseudoholomorphic subvariety, then C defines a canonical ‘model curve’, this a complex curve C_0 that comes with a proper, pseudoholomorphic map to $\mathbb{R} \times (S^1 \times S^2)$ that is almost everywhere 1–1 and has image C . A multi-punctured, pseudoholomorphic sphere is, by definition, an irreducible, pseudoholomorphic subvariety whose model curve is a multiply punctured sphere.

1.C The ends of pseudoholomorphic subvarieties

The set of pseudoholomorphic subvarieties comes with a natural topology whose description occupies the next subsection. This subsection constitutes a digression of sorts to introduce various facts about the large $|s|$ portions of pseudoholomorphic subvarieties that are used both to define this topology and to characterize the resulting space of subvarieties. This digression has three parts.

Part 1 Pseudoholomorphic subvarieties are quite well behaved at large $|s|$. In particular, as demonstrated in [18, Section 2], any given irreducible pseudoholomorphic subvariety C has the following property:

- (1–6) *There exists $R > 1$ such that the $|s| \geq R$ portion of C is a finite disjoint union of embedded cylinders to which the function s restricts as an unbounded function without critical points. Moreover, the constant $|s|$ slices of any such cylinder converge in $S^1 \times S^2$ as $|s| \rightarrow \infty$ to a closed orbit of the Reeb vector field*

$$\hat{\alpha} \equiv (1 - 3 \cos^2 \theta) \partial_t + \sqrt{6} \cos \theta \partial_\varphi.$$

In addition, this convergence is such that any constant $|s|$ slice defines a closed braid in a tubular neighborhood of the limit closed orbit.

The notion of convergence used here can be characterized as follows: The diameter of a tubular neighborhood of the limit closed orbit that contains any given $|s| \geq \mathbb{R}$ slice can be taken to be a function of $|s|$ that decreases to zero at an exponential rate as $|s|$ diverges.

The closed orbits of $\hat{\alpha}$ are called Reeb orbits. As noted in [18], they can all be listed; and here is the full list:

- *The $\theta = 0$ and $\theta = \pi$ loci.*
- *The others are labeled by data $((p, p'), \iota)$ where $\iota \in \mathbb{R}/(2\pi\mathbb{Z})$ and where p and p' are relatively prime integers that are subject to the following constraints:*
 - (a) *At least one is non-zero.*
 - (b) *$|\frac{p'}{p}| > \frac{\sqrt{3}}{\sqrt{2}}$ when $p < 0$.*

The Reeb orbit that is labeled by this data is the locus where $p't - p\varphi = \iota$ and where θ is the unique point in $(0, \pi)$ for which

$$(1-7) \quad p'(1 - 3 \cos^2 \theta) - p\sqrt{6} \cos \theta = 0 \text{ and } p' \cos \theta \geq 0.$$

The convention in (1-7) takes a pair of integers to be relatively prime when they have no common, positive integer divisor save for the number 1. For example, $(0, -1)$ and $(0, 1)$ are deemed to be relatively prime, as is $(-1, -1)$. In this regard, any pair $P = (p, p')$ that obeys the constraints in the second point above defines, via (1-7), a unique angle between 0 and π .

The pair p and p' can be alternatively defined as the respective integrals over the Reeb orbit of the 1-forms $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ using the orientation from $\hat{\alpha}$.

Note that all $\theta \in (0, \pi)$ Reeb orbits come in smooth, 1-parameter families. In this regard, a (p, p') Reeb orbit is fixed by the subgroup of T generated by $p\partial_t + p'\partial_\varphi$ while its corresponding family is obtained from its translates by the action of T . Meanwhile, the Reeb orbits where $\theta = 0$ and $\theta = \pi$ are T -invariant.

Part 2 Granted the preceding, it then follows that any end of C whose associated limit Reeb orbit lies where $\theta \in (0, \pi)$ determines a triple $(\varepsilon, (p, p'))$, where $\varepsilon \in \{+, -\}$ and where (p, p') are a pair of integers. To elaborate, ε is $+$ for a concave side end and $-$ for a convex side one. Meanwhile, the pair (p, p') is a positive, integer multiple of the relatively prime pair of integers that classifies the end's limiting Reeb orbit. In particular, they are the respective integrals of the 1-forms $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ over any constant $|s|$ slice of the end with the latter oriented so its homology class in a tubular neighborhood of the limit Reeb orbit is a positive multiple of the class of the Reeb orbit. For example, if $\gamma \subset S^1 \times S^2$ is a (p, p') Reeb orbit, then $\mathbb{R} \times \gamma$ is a pseudoholomorphic cylinder and in this case the integer pair that is associated to either end is (p, p') .

Of course, an end $E \subset C$ of the sort just described also determines an element, $\iota_E \in \mathbb{R}/(2\pi\mathbb{Z})$, this the angular parameter in (1-7) that helps to specify its associated limit Reeb orbit. A convex side end also determines a real number, c_E . This comes about as follows: Let θ_E denote the $|s| \rightarrow \infty$ limit of θ on E . The arguments from [18, Section 2] can be used to prove that the function θ on any end of C can be written as

$$(1-8) \quad \theta = \theta_E + c_E e^{-\zeta|s|} + o(e^{-(\zeta+\varepsilon)|s|}),$$

where c_E is constant while $\zeta = \sqrt{6} \sin^2 \theta_E (1 + 3 \cos^2 \theta_E) / (1 + 3 \cos^4 \theta_E)$, and ε is positive and also determined a priori by θ_E . In this regard, if $c_E = 0$, then the leading

order term in (1–8) is both above and below θ_E over the large and constant $|s|$ slices of E unless θ is constant on E . In the latter case, E is part of some $\mathbb{R} \times \gamma$ with γ a Reeb orbit. Note that c_E is always zero for a concave side end (see the proof of the second point in [18, (4.21)]).

Part 3 As will now be explained, an end of C whose limit Reeb orbit has $\theta = 0$ or π but does not coincide with the corresponding $\theta = 0$ or $\theta = \pi$ cylinder can also be assigned a discrete triple $(\varepsilon, (p, p'))$ as well as an angular parameter and a real number. To begin the explanation, note first that $\varepsilon \in \{+, -\}$ has the same meaning as before, $+$ when the end is on the convex side and $-$ otherwise. Meanwhile p and p' are the respective integrals over any sufficiently large and constant $|s|$ slice of the 1-forms $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ using the pull-back of $-dt$ to define the orientation. In this regard, the following fact from [18, Section 2] is used:

(1–9) *Let C denote an irreducible pseudoholomorphic subvariety. If $\gamma \subset S^1 \times S^2$ is a Reeb orbit, and if $C \not\subset \mathbb{R} \times \gamma$, then C has at most a finite number of intersections with $\mathbb{R} \times \gamma$.*

As can be proved using results from [18, Sections 2 and 3], any pair (p, p') that arises in the manner just described from an end of C where the $|s| \rightarrow \infty$ limit of θ is 0 or π is constrained by the following rules:

- (1–10) • $p < 0$ in all cases.
- $\frac{p'}{p} < -\frac{\sqrt{3}}{\sqrt{2}}$ for $\theta = 0$ concave side ends and $\frac{p'}{p} > -\frac{\sqrt{3}}{\sqrt{2}}$ for $\theta = 0$ convex side ends.
 - $\frac{p'}{p} > \frac{\sqrt{3}}{\sqrt{2}}$ for $\theta = \pi$ concave side ends and $\frac{p'}{p} < \frac{\sqrt{3}}{\sqrt{2}}$ for $\theta = \pi$ convex side ends.

To explain the angle and real number to assign an end $E \subset C$ where the $|s| \rightarrow \infty$ limit of θ is 0, introduce the functions

$$(1-11) \quad a = |h|^{1/2} \cos \varphi \quad \text{and} \quad b \equiv |h|^{1/2} \sin \varphi.$$

The analysis from [18, Section 3] can be used to prove that an end of the sort under consideration can be parametrized in an orientation preserving fashion at large $|s|$ by a complex parameter $z \in \mathbb{C} - 0$ via a map that sets

$$(1-12) \quad s - it = p \ln(z) \quad \text{and} \quad a - ib = \hat{c}_E z^{\pm p} (1 + o(1)),$$

where the $+$ sign is used when the $|s| \rightarrow \infty$ limit of θ on the end is 0 and the $-$ sign when this limit is π . Note that (1–12) is valid for a concave side end only where $|z| \ll 1$

and only where $|z| \gg 1$ for a convex side one. In any case, the constant $\hat{c}_E \in \mathbb{C} - 0$ while the term indicated by $o(1)$ and its derivatives limits to zero as $|\ln |z|| \rightarrow \infty$. Granted the preceding, the real number, c_E , and the angle, ι_E , assigned to E are the respective real and imaginary parts of $\ln(\hat{c}_E)$.

An end $E \subset C$ where the $|s| \rightarrow \infty$ limit of θ is π has an analogous orientation preserving parametrization by $z \in \mathbb{C} - 0$ that is obtained from the preceding by using the fact that the almost complex geometry is invariant under the involution of $S^1 \times S^2$ that acts to send $(t, (\theta, \varphi)) \rightarrow (t + \pi, (\pi - \theta, -\varphi))$.

1.D The moduli spaces

Fix a finite set whose elements are of the following sort. Each element is a 4-tuple of the form $(\delta, \varepsilon, (p, p'))$ with $\delta \in \{-1, 0, 1\}$, $\varepsilon \in \{+, -\}$ and $(p, p') \in \mathbb{Z} \times \mathbb{Z}$. A given 4-tuple is allowed to appear more than once in this set. Let \hat{A} denote the set that is obtained by augmenting this chosen set of 4-tuples with a single pair, (c_+, c_-) , of non-negative integers. Such data sets are used in what follows to label subsets of the set of pseudoholomorphic subvarieties.

Given a non-negative integer, ζ , and a set \hat{A} as just described, let $\mathcal{M}_{\hat{A}, \zeta}$, denote the set of irreducible pseudoholomorphic subvarieties in $\mathbb{R}(S^1 \times S^1)$ with the following three properties: First, if $C \in \mathcal{M}_{\hat{A}, \zeta}$, then C 's model curve has genus ζ . Second, there is a 1 – 1 correspondence between the 4-tuples in \hat{A} and the set of ends of C so that when E is any given end of C , then its corresponding 4-tuple in \hat{A} is as follows: The component δ is 1, 0 or -1 in the respective cases that the $|s| \rightarrow \infty$ limit of θ on E is 0, neither 0 nor π , or π . Meanwhile, the component ε is $+$ when E is a concave side end and $-$ otherwise. Finally, the pair (p, p') from the 4-tuple is the integral over any sufficiently large $|s|$ slice of E of the pair of 1-forms $(\frac{1}{2\pi} dt, \frac{1}{2\pi} d\varphi)$. Said succinctly, the 4-tuples from the set \hat{A} describe the discrete asymptotic data of the ends of any subvariety in $\mathcal{M}_{\hat{A}, \zeta}$. Finally, C has intersection number c_+ with the $\theta = 0$ cylinder and c_- with the $\theta = \pi$ cylinder.

The genus 0 subvarieties are the multiply punctured spheres. The $\zeta = 0$ version of $\mathcal{M}_{\hat{A}, \zeta}$, is denoted below as $\mathcal{M}_{\hat{A}}$.

Give $\mathcal{M}_{\hat{A}, \zeta}$, the topology where a basis for the neighborhoods of a given $C \in \mathcal{M}_{\hat{A}, \zeta}$ are the subsets that consist of those $C' \in \mathcal{M}_{\hat{A}, \zeta}$ with

$$(1-13) \quad \sup_{z \in C} \text{dist}(z, C') + \sup_{z \in C'} \text{dist}(C, z) < k.$$

Here, k is some fixed positive real number. The topological space that results is a ‘moduli space’. In this regard, note that the definition of a moduli space given here

differs from that in [18] in the case that \hat{A} has 4-tuples with first component equal to ± 1 since the definition in [18] does not constrain the p' component of the 4-tuple.

In any event, with the data set \hat{A} fixed, the structure of the corresponding moduli space $\mathcal{M}_{\hat{A}}$ of multi-punctured spheres is of prime interest in this article. In this regard, the first significant result in this article is summarized by the following theorem:

Theorem 1.2 *The multi-punctured sphere moduli space $\mathcal{M}_{\hat{A}}$ has the structure of a finite dimensional, smooth manifold.*

This theorem is proved in Section 2. This same section also describes various local coordinate charts for $\mathcal{M}_{\hat{A}}$. In particular some are obtained using the $\mathbb{R}/(2\pi\mathbb{Z})$ and real valued parameters that are defined from the $|s| \rightarrow \infty$ limits on various ends.

As explained in Section 2, the dimension of $\mathcal{M}_{\hat{A}}$ is determined by the set \hat{A} . A formula is given in Proposition 2.5. This proposition provides a formula for the ‘formal’ dimension for any given $\zeta > 0$ version of $\mathcal{M}_{\hat{A},\zeta}$, in terms of \hat{A} and ζ . To elaborate, note first that Section 2 proves that any $\zeta > 0$ version of $\mathcal{M}_{\hat{A},\zeta}$, is a finite dimensional variety in the sense that any given point has a neighborhood that is homeomorphic to the zero locus near the origin of a smooth map between two Euclidean spaces. The difference between the dimensions of the domain and range Euclidean spaces is independent of the chosen point in $\mathcal{M}_{\hat{A},\zeta}$. This difference is taken to be the formal dimension of $\mathcal{M}_{\hat{A},\zeta}$.

1.E When $\mathcal{M}_{\hat{A}}$ is non-empty

This subsection provides necessary and sufficient conditions on \hat{A} so as to guarantee a non-empty version of $\mathcal{M}_{\hat{A}}$. In this regard, it follows from what has been said already that $\mathcal{M}_{\hat{A}} = \emptyset$ unless the constraints listed next are obeyed. A set, \hat{A} , of the sort under consideration that obeys these constraints is said here to be an *asymptotic data set*.

Here is the first asymptotic data set constraint: Each $(\delta, \varepsilon, (p, p')) \in \hat{A}$ must obey:

- (1–14) • If $\delta = 0$ and $p < 0$, then $|\frac{p'}{p}| > \frac{\sqrt{3}}{\sqrt{2}}$.
- If $\delta = 1$, then $p < 0$. In addition, $\frac{p'}{p} < -\frac{\sqrt{3}}{\sqrt{2}}$ when $\varepsilon = +$, and $\frac{p'}{p} > -\frac{\sqrt{3}}{\sqrt{2}}$ when $\varepsilon > -$.
 - If $\delta = -1$, then $p < 0$. In addition, $\frac{p'}{p} > \frac{\sqrt{3}}{\sqrt{2}}$ when $\varepsilon = +$, and $\frac{p'}{p} < \frac{\sqrt{3}}{\sqrt{2}}$ when $\varepsilon > -$.

Two more constraints come via Stokes' theorem as applied to line integrals of dt and $d\varphi$:

$$(1-15) \quad \sum_{(\delta, \varepsilon, (p, p')) \in \hat{A}} \varepsilon p = 0 \quad \text{and} \quad \sum_{(\delta, \varepsilon, (p, p')) \in \hat{A}} \varepsilon p' = -(\mathfrak{c}_+ - \mathfrak{c}_-).$$

Here is the next asymptotic data set constraint:

(1-16) *If \hat{A} has two 4-tuples and $\mathfrak{c}_+ = \mathfrak{c}_- = 0$, then the 4-tuples have relatively prime integer pairs.*

To explain, note that when \hat{A} has $\mathfrak{c}_+ = \mathfrak{c}_- = 0$ and two 4-tuples, then $\mathcal{M}_{\hat{A}}$ has only cylinders. All such spaces are described in [18, Section 4], and all obey (1-16).

The final two asymptotic data set constraints involve the set $\Lambda_{\hat{A}} \subset [0, \pi]$ that consists of the angle 0 when $\mathfrak{c}_+ > 0$ or \hat{A} has a $(1, \dots)$ element, the angle π when $\mathfrak{c}_- > 0$ or \hat{A} has a $(-1, \dots)$ element, and the angles that are defined via (1-7) from the $(0, \dots)$ elements in \hat{A} . Granted this definition, here are the last two constraints:

- (1-17) • *If $\Lambda_{\hat{A}}$ has one angle, then \hat{A} has $\mathfrak{c}_+ = \mathfrak{c}_- = 0$ and two 4-tuples, $(0, +, P)$ and $(0, -, P)$ with P relatively prime.*
- *If $\Lambda_{\hat{A}}$ has more than one angle, then neither extremal angle arises via (1-7) from an integer pair of any $(0, +, \dots)$ element in \hat{A} .*

With regards to the first point here, note that any moduli space where the corresponding $\Lambda_{\hat{A}}$ has one angle contains only \mathbb{R} invariant cylinders. The final point arises by virtue of the fact noted previously that the constant c_E in (1-8) is zero when E is a concave side end of a subvariety where the $s \rightarrow \infty$ limit of θ is in $(0, \pi)$.

In all that follows, \hat{A} refers to an asymptotic data set where $\Lambda_{\hat{A}}$ has more than one angle. As it turns out, $\mathcal{M}_{\hat{A}}$ is nonempty if and only if the data from \hat{A} can be used to construct a certain linear graph with labeled edges. The latter graph is denoted in what follows by $L_{\hat{A}}$. The following three part digression describes what is involved.

Part 1 To set the stage, introduce $\Lambda_{\hat{A}}$ to denote the set in $[0, \pi]$ that consists of the distinct angles that come via (1-7) from the integer pairs of the $(0, \dots)$ elements in \hat{A} together with the angle 0 when $\mathfrak{c}_+ > 0$ or when \hat{A} has a $(1, \dots)$ element, and the angle π when $\mathfrak{c}_- > 0$ or when \hat{A} has a $(-1, \dots)$ element.

Part 2 In the present context, a linear graph is viewed as a finite set of distinct points in $[0, \pi]$ with two or more elements. In particular, each graph has at least one edge.

The points in the set are the vertices of the graph, and the intervals that connect adjacent points are the edges. Thus, such a graph has two monovalent vertices and some number of bivalent ones. As a point in $[0, \pi]$, each vertex has a canonical angle assignment. These angles should coincide with the angles in $\Lambda_{\hat{A}}$.

Part 3 As noted at the outset, the edges of the graph $L_{\hat{A}}$ are labeled. In particular, each edge is labeled by an ordered pair of integers subject to the set of six constraints that appear in the upcoming list (1–18).

The notation used is as follows: When e designates an edge, then Q_e or (q_e, q_e') is used to denote its corresponding ordered pair of integers. An edge is said to ‘start a graph’ when its smallest angle is the smallest angle on its graph. By the same token, an edge is said to ‘end a graph’ when its largest angle is the largest angle on its graph.

Here are the constraints:

- (1–18) • *If e ends the graph at an angle in $(0, \pi)$, then $-Q_e$ is the sum of the pairs from each of the $(0, -, \dots)$ elements in \hat{A} that define this maximal angle via (1–7).*
- *If π is the largest angle on e , then Q_e is obtained using the following rule: First, subtract the sum of the integer pairs from the $(-1, -, \dots)$ elements in \hat{A} from the sum of those from the $(-1, +, \dots)$ elements, and then subtract $(0, c_-)$ from the result.*
 - *If e starts the graph at an angle in $(0, \pi)$, then Q_e is the sum of the pairs from each of the $(0, -, \dots)$ elements in \hat{A} that define this minimal angle via (1–7).*
 - *If 0 is the smallest angle on e , then Q_e is obtained using the following rule: First, subtract the sum of the integer pairs from the $(1, +, \dots)$ elements in \hat{A} from the sum of those from the $(1, -, \dots)$ elements and then subtract $(0, c_+)$ from the result.*
 - *Let o denote a bivalent vertex, let θ_0 denote its angle, and let e and e' denote its incident edges with the convention that θ_0 is the largest angle on e . Then $Q_e - Q_{e'}$ is obtained by subtracting the sum of the integer pairs from the $(0, -, \dots)$ elements in \hat{A} that define θ_0 via (1–7) from the sum of the integer pairs from the $(0, +, \dots)$ elements in \hat{A} that defined θ_0 via (1–7).*
 - *Let \hat{e} denote an edge. Then $pq_{\hat{e}'} - p'q_{\hat{e}} > 0$ in the case that (p, p') is an integer pair that defines the angle of a bivalent vertex on \hat{e} . Moreover, if $q_{\hat{e}'} < 0$ and if neither vertex on \hat{e} has angle 0 or π , and if the version*

of (1–7)’s integer p' for one of the vertex angles is positive, then both versions of p' are positive.

A graph of the sort just described is deemed to be an ‘positive line graph’ for \hat{A} . The next theorem explains its significance.

Theorem 1.3 *Suppose that \hat{A} is an asymptotic data set. Then $\mathcal{M}_{\hat{A}}$ is non-empty if and only if \hat{A} has a positive line graph.*

Note that Theorem 3.1 provides equivalent necessary and sufficient criteria for a non-empty $\mathcal{M}_{\hat{A}}$.

To explain something of the nature of Theorem 1.3, note that a graph much like $L_{\hat{A}}$ can be constructed from any subvariety in $\mathcal{M}_{\hat{A}}$. As explained in the next section, the edges of the latter graph are in 1–1 correspondence with the components of the complement in the subvariety of the singular and/or non-compact constant θ loci. Meanwhile, the integer pair that is assigned to any given edge is obtained by integrating the pair $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about any constant θ slice of the corresponding component. However, a graph from a subvariety can differ from T in one fundamental aspect. Although the graph as defined by the subvariety is contractible and connected, it might not be linear.

However, as is proved in the subsequent sections, if $L_{\hat{A}}$ obeys the constraints in (1–18), then there exists either a subvariety in $\mathcal{M}_{\hat{A}}$ that supplies precisely this graph $L_{\hat{A}}$, or there is a sequence of subvarieties in $\mathcal{M}_{\hat{A}}$ whose graphs converge to $L_{\hat{A}}$ in a suitable sense.

The conditions in (1–18) are more or less direct consequences of the manner in which the graph $L_{\hat{A}}$ is designed to encode aspect of the topology of the constant θ slices of a pseudoholomorphic subvariety. Indeed, the only subtle constraint is the final one. The latter can be seen as necessary by considering intersections between the given pseudoholomorphic subvariety and versions of $\mathbb{R} \times \gamma$ where $\gamma \subset S^1 \times S^2$ is an open subset with compact closure in the integral curve of the Reeb vector field. Such submanifolds are pseudoholomorphic, and the final constraint in (1–18) follows from the fact that the local intersection numbers with such submanifolds are necessarily positive. In this regard, keep in mind that local intersection numbers between any two pseudoholomorphic subvarieties are positive (see, for example, McDuff [13].)

As indicated, subvarieties in $\mathcal{M}_{\hat{A}}$ can be constructed granted only the constraint in (1–18). Observations by Michael Hutchings (see also Hutchings–Sullivan [12]) led the author to think that the constraints in (1–18) are sufficient as well as necessary.

1.F Outline of the remaining sections

The remainder of this article is organized as follows: Section 2 contains a proof of Theorem 1.2. This section also describes some useful coordinate systems on $\mathcal{M}_{\hat{A}}$, in particular, some that are obtained using the topology of the constant θ loci on the pseudoholomorphic subvarieties. The latter play a prominent role in the later sections and in the sequel to this article. The final subsection of this article explains how the θ -level sets are used to assign a contractible, connected graph to each subvariety in $\mathcal{M}_{\hat{A}}$.

Section 3 starts the proof of Theorem 1.3 with a result of the following sort: A reasonably general class of immersed subvarieties with the large $|s|$ asymptotics of a subvariety from $\mathcal{M}_{\hat{A}}$ can be deformed to give a subvariety from $\mathcal{M}_{\hat{A}}$. The latter result with the constructions in Section 4 provide necessary and sufficient conditions for the existence of subvarieties in $\mathcal{M}_{\hat{A}}$ that differ from those in Theorem 1.3. These alternative conditions are stated as Theorem 3.1. Arguments, strategies and constructions from Section 3 also play prominent roles in the sequel to this article, as does Theorem 3.1.

Section 4 completes the proof of Theorem 3.1 by exhibiting subvarieties to start the deformations that are described in Section 3. The final section explains how Theorem 1.3 follows from Theorem 3.1.

2 Dimensions and regularity

As remarked in the opening section, the interest here is in the moduli space of pseudoholomorphic, multiply punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$. In this regard, a given component is labeled by an asymptotic data set, \hat{A} , subject to the constraints in (1–14)–(1–17). The added constraints on \hat{A} from Theorem 1.3 are not assumed in this section.

As in the introduction, $\mathcal{M}_{\hat{A}}$ denotes the part of the moduli space that is labeled by \hat{A} . Among other things, this subsection establishes that $\mathcal{M}_{\hat{A}}$ is a smooth manifold and derives a formula in terms of the data from \hat{A} for its dimension. The final three parts of the subsection describe various useful constructions that are subsequently used to study $\mathcal{M}_{\hat{A}}$. In particular, these include various local coordinate charts for $\mathcal{M}_{\hat{A}}$ that can be constructed from the $\mathbb{R}/(2\pi\mathbb{Z})$ and real valued parameters that are associated to the ends of the subvarieties in $\mathcal{M}_{\hat{A}}$.

The subsection starts by introducing a somewhat more general context for the subsequent discussions.

2.A Admissible almost complex structures

Arguments in the next section require facts about the moduli spaces of pseudoholomorphic subvarieties in $\mathbb{R} \times (S^1 \times S^2)$ as defined by almost complex structures that differ from J . The almost complex structures that arise are deemed here to be ‘admissible’, and they are distinguished by three salient features: First, the almost complex structure is ‘tamed’ by all sufficiently small, constant and positive r versions of the symplectic form $d(e^{-rs}\alpha)$. In this regard, the form $d(e^{-rs}\alpha) \equiv \mu$ tames an almost complex structure, J' , when the quadratic function on $T(\mathbb{R} \times (S^1 \times S^2))$ that sends any given vector v to $\mu(v, J'(v))$ is positive on the complement of the zero section. Second, the almost complex structure sends ∂_s to $(1 + 3 \cos^4 \theta)^{-1/2} \hat{\alpha}$, where $\hat{\alpha}$ is the Reeb vector field in (1–6). Third, the given almost complex structure agrees with J on the complement of some compact subset of $\mathbb{R} \times (S^1 \times S^2)$. The almost complex structure J is, of course, admissible as it is compatible with all $r > 0$ versions of $d(e^{-rs}\alpha)$ and thus tamed by all of them.

As defined, the set of admissible almost complex structures should be viewed as a Frechét space with the topology defined so that a given sequence, $\{J_\alpha\}$, of such structures converges to a given admissible J' when there is C^∞ convergence on compact sets and when there exists some fixed compact subset such that each $J_\alpha = J$ on its complement.

As it turns out, this space of admissible almost complex structures is contractible. To see why, note first that if J' is admissible, and $w \in \text{kernel}(\alpha)$ is tangent to a constant s slice of $\mathbb{R} \times (S^1 \times S^2)$, then $J'w$ must be of the form $w' + a\partial_s + b\hat{\alpha}$ where a and b can be any pair of real numbers and where w' is also annihilated by α and tangent to $S^1 \times S^2$. In this regard, $d\alpha(w, w')$ must be positive when $w \neq 0$. Written in this way, the space of admissible almost complex structures manifestly deformation retracts onto the subspace of those that map the tangents to $S^1 \times S^2$ in the kernel of α to themselves. Meanwhile, the latter subspace is contractible since $SL(2; \mathbb{R})/SO(2)$ is contractible.

Consider now the following generalization of Definition 1.1:

Definition 2.1 Let J' denote an admissible almost complex structure. A non-empty subset, C , in $\mathbb{R} \times (S^1 \times S^2)$ is a J' -pseudoholomorphic subvariety if it is closed and has the following properties:

- The complement in C of a countable, nowhere accumulating subset is a 2-dimensional submanifold whose tangent space is J' -invariant.
- $\int_{C \cap K} \omega < \infty$ when $K \subset \mathbb{R} \times (S^1 \times S^2)$ is an open set with compact closure.

- $\int_C d\alpha < \infty$.

A subvariety is called ‘pseudoholomorphic’ below without reference to a particular almost complex structure only in the case that J is the unnamed almost complex structure. Unless stated to the contrary, a subvariety should be assumed irreducible.

Here is the simplest, yet very important example: Let $\gamma \subset S^1 \times S^2$ denote a closed orbit of the vector field $\hat{\alpha}$ from (1–6). Let J' denote any admissible almost complex structure. Then $\mathbb{R} \times \gamma \subset \mathbb{R} \times (S^1 \times S^2)$ is a J' -pseudoholomorphic subvariety.

With Definition 2.1 understood, the next proposition summarizes results from [18, Propositions 2.2 and 2.3] that are germane to the situation at hand.

Proposition 2.2 *Suppose that J' is an admissible almost complex structure and that C is a J' -pseudoholomorphic subvariety. Then, there exists $R > 1$ such that the $|s| \geq R$ portion of C is a finite disjoint union of embedded cylinders to which the function s restricts as an unbounded function without critical points. Moreover,*

- *The constant $|s|$ slices of any such cylinder converge in $S^1 \times S^2$ as $|s| \rightarrow \infty$ to some closed orbit in $S^1 \times S^2$ of the Reeb vector field $\hat{\alpha}$ from (1–6). In this regard, there exists $\kappa > 0$ such that the function of $|s|$ that assigns the maximum distance from the large $|s|$ slices of E to the limit closed orbit of $\hat{\alpha}$ is bounded by a constant multiple of $e^{-\kappa|s|}$.*
- *This convergence is such that any sufficiently large and constant $|s|$ slice defines a closed braid in any given tubular neighborhood of the limit closed orbit. In this regard, all sufficiently large $|s|$ slices are disjoint from the limit Reeb orbit unless the subvariety is the product of \mathbb{R} with the Reeb orbit.*
- *The subvariety C is the image of a complex curve via a proper, J' -pseudoholomorphic map into $\mathbb{R} \times (S^1 \times S^2)$ that is 1–1 on the complement of a finite set of points.*

With Proposition 2.2 understood, define an ‘end’ of a J' -pseudoholomorphic subvariety to be any of the cylinders that appear in Proposition 2.2. The ends of any given J' -pseudoholomorphic subvariety comprise a set with a natural 1–1 correspondence to some ‘asymptotic data set’ as defined in the previous section. In this regard, the correspondence is defined for the J' -pseudoholomorphic subvariety in the same manner as with a J -pseudoholomorphic one. Note as well that a J' -pseudoholomorphic subvariety, if irreducible, can be assigned a ‘genus’, the genus of its model curve.

Granted the preceding, suppose now that \hat{A} is an asymptotic data set and that ζ a non-negative integer. Define the J' version of the moduli spaces $\mathcal{M}_{\hat{A}, \zeta}$ as done for

the J version in Subsection 1.D; this is the set of irreducible, J' -pseudoholomorphic subvarieties with genus ς whose set of ends are in 1–1 correspondence with the data set \hat{A} . Use $\mathcal{M}_{\hat{A}, \varsigma, J'}$ to denote this set. Give this set the topology where a basis for the open neighborhoods of a given subvariety have the form in (1–13).

The following subsumes Theorem 1.2:

Theorem 2.3 *Let J' denote an admissible almost complex structure and let \hat{A} denote an asymptotic data set. Then the multipunctured sphere moduli space $\mathcal{M}_{\hat{A}, 0, J'}$ is a smooth manifold.*

The proof of this theorem occupies the next three parts of the subsection.

2.B The local structure of the moduli spaces

This section contains what is essentially a review of material from [18, Sections 2, 3 and 4]. Before starting the review, agree to fix an admissible almost complex structure J' , an asymptotic data set \hat{A} subject to the constraints in (1–14)–(1–17), and fix a non-negative integer ς . Set $\mathcal{M} \equiv \mathcal{M}_{\hat{A}, \varsigma, J'}$. This subsection describes the local structure around points in \mathcal{M} .

The following three propositions summarize the story on the local structure of \mathcal{M} about any given subvariety. The proofs are straightforward and mostly cosmetic modifications of arguments from [18, Sections 3 and 4]. An outline is given at the end of this subsection but the details are left to the reader.

Proposition 2.4 *There exists a positive integer \hat{I} that depends only on \hat{A} and ς ; and given $C \in \mathcal{M}$, there is a positive integer, n , a smooth map, f , from an origin centered ball in $\mathbb{R}^{\hat{I}+n}$ to \mathbb{R}^n that maps 0 to 0, and a homeomorphism from $f^{-1}(0)$ onto a neighborhood of C in \mathcal{M} that maps the origin to C .*

The integer \hat{I} that appears here is the ‘formal dimension’ of \mathcal{M} . The formula for \hat{I} given in the next proposition uses ς and some of the data from \hat{A} . The data of particular use in this regard are listed next. First on the list is the integer

$$(2-1) \quad c_{\hat{A}} \equiv c_+ + c_-.$$

In this regard, note that $c_{\hat{A}} \geq 0$ as it counts the number of intersections, each weighted with its multiplicity, between any given $C \in \mathcal{M}$ with the $\theta = 0$ and $\theta = \pi$ cylinders. To elaborate, let C_0 denote the model curve for C and let $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ denote its attending J' -pseudoholomorphic map onto C . As the points in C_0 where

the pull-back of θ is either 0 or π are isolated, the closure of a small disk about each such point will intersect the $\theta = 0$ or $\theta = \pi$ cylinder only at its origin. Thus, the ϕ -image of each such disk has a well defined intersection number with the $\theta \in \{0, \pi\}$ locus. This number is positive because any two J' -pseudoholomorphic subvarieties have only positive local intersection numbers at their intersection points. The sum of these local intersection numbers is the integer $c_{\hat{A}}$ in (2-1).

Second on the list are the integers N_- and N_+ , these the respective number of elements of the form $(0, -, \dots)$ and $(0, +, \dots)$ in \hat{A} . Finally, use \hat{N} to denote the number of elements in \hat{A} whose first entry is 1 or -1 . With this notation set, consider:

Proposition 2.5 *The integer \hat{I} from Proposition 2.4 is given by the formula*

$$(2-2) \quad \hat{I} = N_+ + 2(N_- + \hat{N} + C_{\hat{A}} + \zeta - 1).$$

This proposition is also proved below

A point $C \in \mathcal{M}$ is called a regular point when the $n = 0$ case of Proposition 2.4 applies. The next proposition concerns the subset of regular points in \mathcal{M} .

Proposition 2.6 *The set of regular points in \mathcal{M} has the structure of a smooth manifold of dimension \hat{I} . Moreover, if $C \in \mathcal{M}$ is a regular point, then Proposition 2.4's homeomorphism between the \hat{I} -dimensional ball and a neighborhood of C in \mathcal{M} defines a smooth coordinate chart.*

The remainder of this subsection sketches the argument for the preceding propositions. To begin, fix a smooth Riemannian metric, g' , on $\mathbb{R} \times (S^1 \times S^2)$ for which J' is orthogonal. In this regard, such a metric can and should be chosen so that ∂_s and $\hat{\alpha}$ have norm 1, and such that g' agrees with $ds^2 + dt^2 + d\theta^2 + \sin^2 \theta d\varphi^2$ on the complement of a compact subset of $\mathbb{R} \times (S^1 \times S^2)$.

Now fix $C \in \mathcal{M}$ and let C_0 again denote the model curve for C . For simplicity, assume that the J' -pseudoholomorphic map ϕ from C_0 is an immersion. The story when C is not immersed is similar in most respects to that given below and is summarized briefly in Subsection 2.D. The reader is referred to [18, Section 3] for a more detailed account.

Granted that C is immersed, there exists a pull-back normal bundle, $N \rightarrow C_0$; its fiber at any given point is the g' -orthogonal complement in $T(\mathbb{R} \times (S^1 \times S^2))|_C$ to TC_0 at the image point in C . This bundle inherits a complex, hermitian line bundle structure from J' and the metric g' . The latter structure endows N with the structure of a holomorphic bundle over C_0 . In addition, there is a disk bundle $N_1 \subset N$ of

some fixed radius, r_1 , together with an immersion, $e: N_1 \rightarrow \mathbb{R} \times (S^1 \times S^2)$, onto a regular neighborhood of C that restricts to the zero section as the map from C_0 . In this regard, e is chosen so as to embed any given fiber of N_1 as a pseudoholomorphic disk. Also, the differential of e is uniformly bounded, and it defines along the zero section a g' -isometric map from $TC_0 \otimes N$ to the pull-back of TX .

Let η denote a smooth section of N_1 . Then $e(\eta)$ is a pseudoholomorphic subvariety if and only if η satisfies a certain differential equation, one with the schematic form

$$(2-3) \quad \bar{\partial}\eta + \nu\eta + \mu\bar{\eta} + \mathcal{R}_0(\eta) + \mathcal{R}_1(\eta) \cdot \partial\eta = 0,$$

where the notation is as follows: First, ν and μ are bounded sections of $T^{0,1}C_0$ and $N^2 \otimes T^{0,1}C_0$ respectively. Second, \mathcal{R}_0 is a smooth (but not complex analytic), fiber preserving map from $N_1 \otimes N_1$ to $N \otimes T^{0,1}C_0$. Meanwhile, \mathcal{R}_1 is a smooth (but not complex analytic), fiber preserving map from N_1 to $\text{Hom}(T^{1,0}C_0, T^{0,1}C_0)$. These two maps obey

$$(2-4) \quad |\mathcal{R}_0(\eta)| \leq c|\eta|^2 \quad \text{and} \quad |\mathcal{R}_1(\eta)| \leq c|\eta|.$$

where c is a constant.

The linear part of (2-3) defines the first order, operator D_C , an \mathbb{R} -linear map from the space of sections of N to those of $N \otimes T^{0,1}C_0$. Thus,

$$(2-5) \quad D_C\eta \equiv \bar{\partial}\eta + \nu\eta + \mu\bar{\eta}.$$

The operator D_C induces a bounded, Fredholm operator between various weighted, Sobolev space completions of certain subspaces of sections of N and $N \otimes T^{0,1}C_0$. In particular, the completions of interest are defined as follows: Fix a positive, but very small real number, κ ; an upper bound can be deduced from the data in \hat{A} . Now, fix a smooth non-negative function, r , on C_0 with the following properties:

- $r = -\kappa|s|$ on any end in C_0 that provides an element in \hat{A} with first component 0.
- If $E \subset C_0$ is an end that contributes an element of the form $(\pm 1, \pm, (p, p'))$, then

$$(2-6) \quad r = \left(-\kappa + \left| \frac{p'}{p} \right| - \sqrt{\frac{3}{2}} \right) |s| \text{ on } E.$$

With such a function chosen, define respective domain and range Hilbert spaces for D_C to be the completions of the spaces of smooth sections of N and $N \otimes T^{0,1}C_0$ for which the quadratic functionals

$$(2-7) \quad \eta \rightarrow \int_{C_0} e^r (|\nabla\eta|^2 + |\eta|^2) \quad \text{and} \quad \eta \rightarrow \int_{C_0} e^r |\eta|^2,$$

are finite. The respective functionals in (2–7) define the Hilbert space norms on the domain and range. These are both denoted here as $\|\cdot\|$.

[18, Lemma 3.3 in Section 3] asserts that the operator D_C defines a Fredholm operator from the domain Hilbert space to the range. The vector spaces $\text{kernel}(D_C)$ and $\text{cokernel}(D_C)$ refer to the respective kernel and cokernel of this Fredholm operator. Arguments from [18, Sections 3d and 4b,d] prove that the index of this Fredholm operator is the integer \hat{I} .

Meanwhile, almost verbatim copies of the constructions from [18, Section 3c] allow [18, Proposition 3.2] to construct a ball, $B \subset \text{kernel}(D_C)$ about the origin, a smooth map, $f: B \rightarrow \text{cokernel}(D_C)$, and a smooth map $F: B \rightarrow C^\infty(N_1)$ with the following properties: First, $\|f(\eta)\| \leq c\|\eta\|^2$ and $|F(\eta) - \eta| + |\nabla(F(\eta) - \eta)| \leq c\|\eta\|^2$. Second, F composes with the exponential map $e: N_1 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ so as to map $f^{-1}(0)$ homeomorphically onto a neighborhood of C in \mathcal{M} .

The conclusions of these last two paragraphs restate those of Propositions 2.4 and 2.5. The conclusions of Proposition 2.6 follow as a formal consequence of the role played by the implicit function theorem in the construction of F . In this regard, keep in mind that when $C \in \mathcal{M}$ is a regular point, then the use of the map F provides the identification

$$(2-8) \quad T\mathcal{M}|_C = \text{kernel}(D_C).$$

It proves useful in subsequent arguments to write D_C explicitly on the ends of C . For this purpose, let E denote a given end. Constructions from [18, Section 2] parametrize E by coordinates $(\rho, \tau) \in [0, \infty) \times \mathbb{R}/(2\pi\mathbb{Z})$ and trivialize N over E as $E \times \mathbb{R}^2$ so that D_C becomes an operator of the form

$$(2-9) \quad \partial_\rho + \begin{pmatrix} -\zeta' & -\partial_\tau \\ \partial_\tau & -\zeta \end{pmatrix} + \hat{\sigma}_1 + \hat{\sigma}_2 \cdot \partial_\rho + \hat{\sigma}_3 \cdot \partial_\tau,$$

acting on 2–component column vectors. To elaborate, ζ is a positive constant for concave side ends and negative for convex side ones. In all cases, the value of ζ is determined by the element from \hat{A} that labels E . Meanwhile, $\zeta' = \zeta$ when the $|s| \rightarrow \infty$ limit of θ on E is 0 or π . Otherwise, $\zeta' = 0$. Finally, $\hat{\sigma}_{1-3}$ are smooth, 2×2 matrix valued functions whose components with their derivatives decay to zero as $\rho \rightarrow \infty$ faster than $e^{-\kappa'\rho}$ with κ' a positive constant. Note that the coordinates (ρ, τ) are such that the function $|s|$ restricts to E as a multiple of ρ . Also, $d\rho \wedge d\tau$ orients E . Meanwhile, the trivialization chosen for N is such that when $\gamma \subset S^1 \times S^2$ is a Reeb orbit where $\theta \notin \{0, \pi\}$ and $E \subset \mathbb{R} \times \gamma$, then the column vector with 0 in its top entry and 1 in its lower entry is a positive multiple of the projection to γ 's normal bundle of the vector field $-\partial_\theta$. The column vector with 1 in its top entry and 0 in its

lower entry defines a deformation of γ along its orbit under the action of the group, T , of isometries of $\mathbb{R} \times (S^1 \times S^2)$ generated by the vector fields ∂_t and ∂_φ .

2.C The moduli space for multi-punctured spheres near immersed varieties

Restrict attention now to the genus zero case. Thus, $\mathcal{M} = \mathcal{M}_{\hat{A},0}$ contains only multi-punctured spheres. Using the notation in Proposition 2.5, the model curve of each $C \in \mathcal{M}$ has $N_+ + N_- + \hat{N}$ punctures. Here is the first key observation about \mathcal{M} :

Proposition 2.7 *The space \mathcal{M} is a smooth manifold of dimension $N_+ + 2(N_- + \hat{N} + c_{\hat{A}} - 1)$ on some neighborhood of any given subvariety with only immersion singularities. In this regard, the operator D_C has trivial cokernel for each immersed subvariety in \mathcal{M} and thus each such subvariety is a regular point of \mathcal{M} .*

To explain the terminology, a subvariety is said to have only immersion singularities when the tautological map from its model curve is an immersion.

An analogous assertion for the non-immersed submanifolds is given in Subsection 2.D.

The proof of Proposition 2.7 requires a preliminary digression to introduce new pairings between the fundamental class of a pseudoholomorphic subvariety and certain classes from $H^2(\mathbb{R} \times (S^1 \times S^2); \mathbb{Z})$. In this regard, keep in mind that the fundamental class of a non-compact subvariety does not canonically define a linear functional on this second cohomology.

In this digression, C denotes any given J' -pseudoholomorphic subvariety without restriction on its genus or its singularities. To start, note that Subsection 3.a in [18] defines an integer valued pairing between the fundamental class of an irreducible, pseudoholomorphic subvariety and its Poincaré dual. It also defines an integer valued pairing between the fundamental class of such a subvariety and the first Chern class for the given almost complex structure on $\mathbb{R} \times (S^1 \times S^2)$. With C denoting the subvariety in question, these integer pairings are respectively denoted by $\langle e, [C] \rangle$ and $\langle c_1, [C] \rangle$; they enter both in [18, Proposition 3.1] to compute the Euler characteristic of C , and in [18, Proposition 3.6].

To elaborate, $\langle e, [C] \rangle$ is a ‘self-intersection’ number that is defined by pushing C off of itself using a fiducial push-off on its ends. In this regard, the fiducial push-off used to define $\langle e, [C] \rangle$ pushes along any section of N over the large $|s|$ part of C that is homotopic there through nowhere zero sections to a particular standard section. To obtain this standard section, note first that ∂_θ is not tangent to any Reeb orbit that

is obtained as the $|s| \rightarrow \infty$ limit of the constant s slices of any end of C_0 where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$. Thus, the projection of ∂_θ along the large $|s|$ part of C to its local normal bundle defines a non-vanishing section, η_0 , of N over any such end. This η_0 is used for the standard section over any end of C where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$. A different section is used on those ends where $\lim_{|s| \rightarrow \infty} \theta \in \{0, \pi\}$.

Meanwhile, $\langle c_1, [C] \rangle$ is defined using a certain standard section for the restriction to the large $|s|$ part of C of the J -version of the canonical bundle for $\mathbb{R} \times (S^1 \times S^2)$. The standard section over an end $E \subset C_0$ is $(dt + \frac{i}{g}df) \wedge (\sin^2 \theta d\varphi + \frac{i}{g}dh)$ in the case that the $|s| \rightarrow \infty$ limit of θ on E is neither 0 nor π . A different section is used on the other ends.

If C is not a $\theta = 0$ or $\theta = \pi$ cylinder, then η_0 is nowhere zero at large $|s|$ on C and so can be used on all ends of C to define a pairing between C 's fundamental class and its Poincaré dual. This new pairing is denoted here by $\langle e, [C] \rangle_*$. Of course, the new pairing is identical to the old when \hat{A} has no elements with first component ± 1 .

If C is not a $\theta = 0$ or a $\theta = \pi$ cylinder, then $(dt + \frac{i}{g}df) \wedge (\sin^2 \theta d\varphi + \frac{i}{g}dh)$ is also non-zero at large $|s|$ on C , and so this section can be used on all of C 's ends to define a new pairing of C with c_1 . This new pairing is denoted by $\langle c_1, [C] \rangle_*$. This new pairing has the virtue that it is equal to minus the number of intersections (counted with multiplicity) between C and the locus where θ is 0 or π .

To end the digression, let m_C denote the number of double points of any compactly supported perturbation of C that gives an immersed, symplectic curve with purely double point immersion singularities with positive local intersection numbers. (Such a perturbation always exists.) It then follows from [18, Proposition 3.1 and Proposition 4.1] using observations from [18, Section 4c] that

$$(2-10) \quad \begin{aligned} -\chi(C) &= \langle e, [C] \rangle_* + \langle c_1, [C] \rangle_* - 2m_C. \\ \hat{I} &= \langle e, [C] \rangle_* - \langle c_1, [C] \rangle_* - 2m_C + N_- + \hat{N}. \end{aligned}$$

These last formulas end the digression.

Proof of Proposition 2.7 The proof has five steps.

Step 1 By definition, no C in any $\mathcal{M}_{\hat{A}, \zeta}$ is a $\theta = 0$ or $\theta = \pi$ cylinder. Thus, the section η_0 as defined above is nowhere zero at large $|s|$ on C . Define $\text{Deg}(N)$ to be the usual algebraic count of the zero's of any section of N that has no zeros where $|s|$ is large and is homotopic on the large $|s|$ slices of C_0 through non-vanishing sections to η_0 . The following lemma identifies $\text{Deg}(N)$:

Lemma 2.8 *Let \hat{A} be an asymptotic data set and ζ a non-negative real number. If $C \in \mathcal{M}_{\hat{A}, \zeta}$ is an immersed subvariety, then*

$$(2-11) \quad \text{Deg}(N) = N_+ + N_- + \hat{N} + c_{\hat{A}} + 2\zeta - 2.$$

Proof of Lemma 2.8 By definition, the integer $\text{Deg}(N)$ and the pairing $\langle e, [C] \rangle_*$ are related by the formula $\text{Deg}(N) = \langle e, [C] \rangle_* - 2m_C$. This being the case, add the two lines in (2-10) to obtain

$$(2-12) \quad -\chi(C) + \hat{I} = 2\text{Deg}(N) + N_- + \hat{N}.$$

Upon rearrangement and division by 2, this last equation gives (2-11). □

Step 2 Now suppose that $C \in \mathcal{M}_{\hat{A}, 0} \equiv \mathcal{M}_{\hat{A}}$. Let η be an element in $\text{kernel}(D_C)$. It follows from the large $|s|$ picture of D_C in (2-9) that each such η is bounded on C and has a well defined $|s| \rightarrow \infty$ limit on each end of C . Moreover, the form of (2-9) also implies that η has finitely many zeros on C_0 and so has a well defined degree at large $|s|$ on each end of C_0 as measured with respect to the degree zero standard of η_0 and with the orientation of the constant ρ circles reversed. (Imagine gluing a disk to such a circle and then the degree is the degree as viewed from the origin of the glued disk.) Moreover, the form for D_C given in (2-9) implies that this large $|s|$ degree on each end of C_0 is non-negative. Indeed, a column vector with negative degree at all large ρ that is annihilated by the operator in (2-9) will grow in size as $\rho \rightarrow \infty$ faster than allowed as a member of $\text{kernel}(D_C)$. (The growth is faster than the exponential of a positive, constant multiple of ρ whose value is determined by the data from \hat{A} . Meanwhile, the function r that appears in (2-7) is defined so as to rule out elements with such growth.)

When $E \in \text{End}(C)$, use $\delta_E(\eta)$ to denote the degree of η at large $|s|$ on E .

Meanwhile, if $z \in C_0$ is a point where η vanishes, use $\text{deg}_z(\eta)$ to denote the local degree of η . Note that $\text{deg}_z(\eta) > 0$ at each zero of η as can be proved using the fact that $D_C \eta = 0$.

The following identity is now a consequence of the various definitions:

$$(2-13) \quad \sum_E \delta_E(\eta) + \sum_{\{z: \eta(z)=0\}} \text{deg}_z(\eta) = \text{Deg}(N)$$

For reference later, note that right hand equality in (2-13) is valid even if C is not everywhere J' -pseudoholomorphic and η is not a solution to a particular differential equation. Indeed, (2-13) holds provided only that C is an immersed subvariety with

the large $|s|$ asymptotics of a J' -pseudoholomorphic subvariety, and that the section η of C 's pull-back normal bundle has no large $|s|$ zeros. Of course, in this general context, there need not exist sign constraints on $\delta_{(\cdot)}(\eta)$ and $\deg_{(\cdot)}(\eta)$.

Step 3 The next point to make is that there exists a subspace, K_0 , of codimension no greater than N_+ in $\text{kernel}(D_C)$ with the property that $|\eta| \rightarrow 0$ as $|s| \rightarrow \infty$ on each concave side end of C_0 where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$. Indeed, the upper bound on the codimension of K_0 follows from the assertion that the requirement of a non-zero limit on a concave side end of C_0 defines a codimension 1 condition on $\text{kernel}(D_C)$. And, the latter claim follows from the form of D_C given by (2-9) since any 2-component column vector that is bounded, has non-zero limit as $\rho \rightarrow \infty$ and is annihilated by (2-9) must limit to zero or to the column vector with zero in the lower entry and 1 in the upper. In this regard, remember that the constant ζ in (2-9) is positive for concave side ends.

To summarize: The subspace K_0 has $\dim(K_0) \geq 2(N_- + \hat{N} + c_A - 1)$ with a strict inequality when $\text{cokernel}(D_C)$ is non-trivial.

Step 4 The subspace K_0 also has the following key property: If $\eta \in K_0$ is non-trivial, then the degree $\delta_{(\cdot)}(\eta) \geq 1$ on each concave side end. Indeed, this is another consequence of the form for D_C given in (2-9) because a 2-component column vector with zero degree at all sufficiently large ρ that is annihilated by (2-9) will converge to a non-zero multiple of the column vector with zero in the lower entry and 1 in the upper.

With this positivity noted, then (2-11) and (2-13) imply that

$$(2-14) \quad \sum_{\{z:\eta(z)=0\}} \deg_z(\eta) \leq N_- + \hat{N} + c_A - 2$$

if η is a non-trivial element of K_0 . In this regard, note that (2-14) is an equality if and only if η has degree $\delta_{(\cdot)}(\eta) = 1$ on each concave side end of C_0 where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$ and $\delta_{(\cdot)}(\eta) = 0$ on all other ends. (Remember that $\delta_{(\cdot)}(\eta) \geq 0$ on all ends.)

Step 5 Now take a set $\Omega \subset C_0$ of $N_- + \hat{N} + c_A - 1$ distinct points. These points define a subset $K_1 \subset K_0$ of codimension no greater than $2(N_- + \hat{N} + c_A - 1)$ whose elements vanish at each point in Ω . By virtue of the dimension count in Step 3, above, this subspace K_1 is non-trivial if $\text{cokernel}(D_C)$ is non-trivial. However, by virtue of (2-14), and by virtue of the fact that $\deg_{(\cdot)}(\eta) > 0$ at each of its zeros, the subspace K_0 has no elements that vanish at more than $N_- + \hat{N} + c_A - 1$ points of C_0 . Thus

K_1 must be trivial. This being the case, $\text{cokernel}(D_C) = \{0\}$, and it follows that \mathcal{M} is smooth near C_0 of the asserted dimension. \square

2.D The story on \mathcal{M} near non-immersed subvarieties

This subsection explains why \mathcal{M} is a smooth manifold on neighborhoods of its non-immersed subvarieties. There are five parts to the story.

Part 1 To start, suppose that an admissible J' has been fixed and that C is a subvariety in the J' version of \mathcal{M} . Let $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ again denote the tautological map from C 's model curve onto C . Since ϕ is J' -pseudoholomorphic, it fails as an immersion only where its differential is zero. Use $\Xi \subset C_0$ to denote the set of points where this occurs. By virtue of Proposition 2.2, the set Ξ is compact, and standard arguments about the local structure of pseudoholomorphic subvarieties (as can be found in McDuff [13], and McDuff and Salamon [14]) prove that Ξ is a finite set. To be more explicit, results from [13] (see also Ye [19]) can be used to prove that there is a holomorphic coordinate u on disk in C_0 centered at a given point in Ξ and complex coordinates (x, y) on a ball in $\mathbb{R} \times (S^1 \times S^2)$ centered at the image of the given point that give ϕ the form

$$(2-15) \quad \phi(u) = (au^{q+1}, 0) + o(|u|^{q+2})$$

where a is a non-zero complex number and $q \geq 1$ is an integer. As a consequence, the pull-back by ϕ of the $(1, 0)$ part of the complexified tangent space to $\mathbb{R} \times (S^1 \times S^2)$ canonically splits as $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) = W \oplus N$ where W and N are complex line bundles that are characterized as follows: The differential of ϕ provides a complex linear map from $T_{1,0}C$ into W and N restricts to $C_0 - \Xi$ as its pull-back normal bundle.

Part 2 This step constitutes a digression to elaborate on the assertion in Proposition 2.4 in the present case. To begin the digression, note that a deformation of the map ϕ that moves image points only slightly can always be written as the image via an exponential map of a section of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. In this regard, an exponential map restricts to the zero section as ϕ and it embeds a ball in each of the fibers that is centered at the origin and has a base-point independent radius. Note that an exponential map can be chosen to embed a disk about the origin in each fiber of each of the W and N subbundles as J' -pseudoholomorphic submanifolds in $\mathbb{R} \times (S^1 \times S^2)$. These disks can be assumed to have base-point independent radii.

A section of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ defines J' -pseudoholomorphic map from C_0 into $\mathbb{R} \times (S^1 \times S^2)$ if and only if it obeys a certain non-linear differential equation whose

linearization along the zero section has the form $\widehat{D}\eta = 0$ where \widehat{D} is a first order, \mathbb{R} -linear operator with the symbol of the Cauchy–Riemann operator $\bar{\partial}$. Here, \widehat{D} maps the space of sections of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ to those of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) \otimes T^{0,1}C_0$. The operator \widehat{D} is described in [18, Part 2 of Section 3b]. Note in particular that \widehat{D} maps the sections of the W summand of $\phi_*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ to sections of $W \otimes T^{0,1}C_0$. In particular, if v is a section of $T_{0,1}C_0$, then $\widehat{D}(\phi_*v) = \phi_*(\bar{\partial}v)$. Meanwhile, \widehat{D} followed by the orthogonal projection onto $N \otimes T^{0,1}C_0$ acts as the operator in (2–5) on sections over $C_0 - \Xi$ of the N summand of $\phi_*T_{0,1}(\mathbb{R} \times (S^1 \times S^2))$.

This operator \widehat{D} is Fredholm when mapping a certain L^2_1 Hilbert space completion of a particular subspace of sections of $\phi^*T_{0,1}(\mathbb{R} \times (S^1 \times S^2))$ to a corresponding L^2 Hilbert space completion of one of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) \otimes T^{0,1}C_0$. These completions are defined by norms on the W and N summands of these bundles that are straightforward analogs of those that are depicted in (2–7). In this regard, the norms on the W summand force the sections to have limit zero as $|s| \rightarrow \infty$, while those on the N summands are weighted exactly as depicted in (2–7).

Now deformations of ϕ that preserve the J' -pseudoholomorphic condition are not of primary interest. Rather, the interest is in deformations of ϕ that are J' -pseudoholomorphic for an appropriately deformed complex structure on C_0 . The description of the latter requires the introduction of the vector space of first order deformations of the complex structure on C_0 . In particular, when there are three or more ends to C_0 , this last vector space has dimension $N_+ + N_- + \widehat{N} - 3$, and it is the quotient of a suitably constrained (at large s) space of sections of $T_{1,0}C_0 \otimes T^{0,1}C_0$ by the image of $\bar{\partial}$. This the case, fix a vector space V of smooth sections of $T_{1,0}C_0 \otimes T^{0,1}C_0$ that projects isomorphically to said quotient.

All this understood, let D_C now denote $(1 - \prod) \cdot \underline{D}$ where \prod denotes the orthogonal projection in D 's range space onto ϕ_*V . Here, ϕ_*V denotes the image of V under the tautological map that is defined by the differential of ϕ , thus a vector subspace of the summand $W \otimes T^{0,1}C_0$ of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) \otimes T^{0,1}C_0$. The operator D_C is viewed here as mapping the domain of \widehat{D} to the image of $(1 - \prod)$ in the range Hilbert space for \widehat{D} .

The operator D_C as just described plays the role for the non-immersed subvarieties that is played by its namesake in the story in the immersed case following Proposition 2.6 and in the previous subsection. In particular, arguments from [18, Section 3] prove the following: A neighborhood of C in \mathcal{M} is homeomorphic to $f^{-1}(0)$ where f is a certain smooth map from a ball in the kernel of D_C to the cokernel of D_C that vanishes with its first derivatives at the origin. Moreover, C is a regular point of \mathcal{M} when $\text{cokernel}(D_C) = \{0\}$ in which case \mathcal{M} is a smooth manifold near C and (2–8)

holds. Arguments from [18, Section 3] also prove that the index of D_C is equal to \hat{I} from (2–2).

The identification between $f^{-1}(0)$ and \mathcal{M} uses the exponential map from a uniform radius ball subbundle of $\phi^*T(\mathbb{R} \times (S^1 \times S^2))$ into $\mathbb{R} \times (S^1 \times S^2)$. The identification also involves a certain smooth map from the domain of f to the domain of D_C . The latter map, F , is smooth in the C^∞ topology, it maps 0 to 0 and it is the identity to first order. The embedding of $f^{-1}(0)$ into \mathcal{M} identifies any $\lambda \in f^{-1}(0)$ with the image in $\mathbb{R} \times (S^1 \times S^2)$ of the composition of exponential map with $F(\lambda)$.

Granted all of this, here is the promised analog of Proposition 2.7:

Proposition 2.9 *The space \mathcal{M} is a smooth manifold of dimension $N_+ + (N_- + \hat{N} + c_{\hat{A}} - 2)$ on some neighborhood of any given non-immersed subvariety. In particular the operator D_C as just described has trivial cokernel at each such subvariety and each is a regular point of \mathcal{M} .*

Part 3 This part of the subsection contains the following proof.

Proof of Proposition 2.9 It follows from the discussion in the preceding part of the subsection that it is sufficient to prove that the dimension of the kernel of D_C is equal to its index.

To start this task, recall that the restriction of the operator D_C to the elements in its domain that are sections of W maps this subspace to the subspace of its domain whose elements are sections of $W \otimes T^{0,1}C_0$. Denote this restricted operator as D^W . Meanwhile, the composition of D_C and then pointwise orthogonal projection to $N \otimes T^{0,1}C_0$ restricts to the subspace of sections of N in D_C 's domain to give a differential operator that is denoted in what follows as D^N . Since D_C is Fredholm, so are D^W and D^N . Furthermore, the sum of their indices is the index of D_C . In this regard, note that D_C followed by pointwise orthogonal projection to the summand $W \otimes T^{0,1}C_0$ defines a zeroth order operator from D^N 's domain to the range space of D^W .

To give formulas for the indices of D^W and D^N , associate to each $z \in \Xi$ the integer $q \equiv q_z$ that appears in the relevant version of (2–15), then set $\wp \equiv \sum_{z \in \Xi} q_z$. An analysis much like that used in [18, Section 3] for D_C 's namesake in (2–7) finds that D^W has index $2\wp$ and D^N has index $\hat{I} - 2\wp$. This is all relevant to D_C by virtue of the following observation: If the cokernel of D^W is trivial, then $\text{kernel}(D_C)$ is isomorphic to the direct sum of the kernels of D^W and D^N . Thus, it is enough to prove that both D^W and D^N have trivial cokernel.

Consider first the case of D^W . Were the kernel to have $2\wp + 1$ linearly independent vectors, then by virtue of (2–15), this kernel would have a non-trivial vector from the image via ϕ_* of TC_0 . This would provide a vector field on C_0 . Let v denote the latter. Then $\bar{\partial}v \in V$ and this implies that $\bar{\partial}v$ must be zero since V is defined so as to project isomorphically onto the cokernel of $\bar{\partial}$. But then v must be zero because its $|s| \rightarrow \infty$ limit is zero on all ends of C_0 . Thus, D^W has index $2\wp$ and trivial cokernel.

Consider next the case of D^N . To this end, note that the bundle N has an associated degree that is defined with respect to the large $|s|$ section that was used to give (2–8) in the immersed case. As a consequence of (2–15), the degree of N so defined is $N_+ + (N_- + \hat{N} + c_{\hat{A}} - 2) - \wp$. Now, Corollary 2.11 to come asserts that $\wp \leq N_- + \hat{N} + c_{\hat{A}} - 2$ in all cases. Grant this bound. As D^N has index $N_+ + 2(N_- + \hat{N} + c_{\hat{A}} - 1) - 2\wp$, its index is at least $N_+ + 2$. Since this index is positive, the argument given in the previous subsection can be taken in an essentially verbatim fashion to prove that D^N has trivial cokernel and thus its kernel dimension is $N_+ + 2(N_- + \hat{N} + c_{\hat{A}} - 1) - 2\wp$. \square

2.E The critical points of θ

This subsection serves as a digression of sorts to describe various key properties of the pull-back of θ to any given subvariety in \mathcal{M} . The discussion here has four parts. In this regard, note that the last part has the promised Corollary 2.11.

Part 1 Let J' denote an admissible almost complex structure and let C be a J' -pseudoholomorphic subvariety. Assume that C is not an \mathbb{R} -invariant cylinder, but there is no need to assume that C is a multiply punctured sphere. This part explains why the only local maxima and minima of θ 's pull-back to C 's model curve occur where $\theta = 0$ or $\theta = \pi$.

To see why such is the case, consider a point $z \in C_0$ where $\theta \in (0, \pi)$ and $d\theta = 0$. The point $\phi(z)$ sits in some pseudoholomorphic disk D' whose tangent space is everywhere spanned by ∂_s and the vector field $\hat{\alpha}$. Note that θ is constant on D' . Since C is not \mathbb{R} -invariant, the closure of D' intersects C only at $\phi(z)$ if its radius is small, so assume that such is the case. Then D' has a well defined, intersection number with the ϕ -image of any sufficiently small radius disk in C_0 centered on z , and this is positive because both D' and the image of the disk in C_0 are pseudoholomorphic. Were $\theta(z)$ a local maximum or minimum of θ on C , then a sufficiently small radius version of D' could be pushed in the respective ∂_θ or $-\partial_\theta$ directions so that the resulting isotopy has the following two properties: First, it avoids the ϕ -image of the boundary of any sufficiently small radius disk in C_0 centered on z . Second, it results in a disk that is entirely disjoint from the ϕ -image of the whole any sufficiently small radius

disk centered at z . Such an isotopy is precluded by the positive intersection number between D' and the ϕ -image of the disks in C_0 centered on z .

Part 2 Continuing the story from Part 1, define the degree of vanishing of $d\theta$ at z to be one less than the intersection number between D' and the ϕ -image of any sufficiently small radius disk in C_0 centered at z . Denote this number by $\deg(d\theta|_z)$. What follows is an equivalent definition of this number.

To start, note that by virtue of the fact that θ is constant on the integral curves of the vector field $\hat{\alpha}$ but not so on C , a neighborhood in $\mathbb{R} \times (S^1 \times S^2)$ of $\phi(z)$ has complex coordinates, (x, y) , with the following four properties: First, $(0, 0)$ is the point z . Second, dx and dy span $T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$ at $\phi(z)$. Third, the $y = 0$ disk is J' -pseudoholomorphic. Finally, the constant x disks centered where $y = 0$ are J' -pseudoholomorphic disks whose tangent planes are everywhere spanned by the vector fields ∂_s and $\hat{\alpha}$. Thus, θ is constant on the $x = \text{constant}$ disks and so a function of x only. This understood, then θ can be written as $\theta = \text{Re}(\sigma x) + o(|x|^2)$ with σ a non-zero constant

Meanwhile, by virtue of the fact that ϕ is J' -pseudoholomorphic, there is a complex coordinate, u , for a small radius disk in C_0 centered at z such that ϕ pulls the coordinate x back as $\phi^*x = au^{p+1} + o(|u|^{p+2})$ with $a \in \mathbb{C}$ and with p an integer of size at least 1. The integer p is the degree of vanishing of $d\theta$ at z because $d\theta$ pulls back near z as

$$(2-16) \quad d\theta = p \text{Re}(\sigma au^p du) + o(|u|^{p+1}).$$

Part 3 Note that C_0 has at most a finite number of critical points in any given compact subset. This is a consequence of (2-16) and the fact that θ 's extremal critical points occur where C intersects the $\theta = 0$ or $\theta = \pi$ cylinders.

The next point is that $d\theta$ is non-vanishing at all sufficiently large values of $|s|$ on C_0 . To see that such is the case, note first that if E is an end of C where the $|s| \rightarrow \infty$ limit of θ is either 0 or π , then (1-12) implies that there are no large $|s|$ critical points of θ on E . The analysis used in [18, Sections 2 and 3] also serves to prove this in the case that the $|s| \rightarrow \infty$ limit of θ on E is in $(0, \pi)$. To be more precise with regard to the latter case, these techniques in [18] find coordinates (ρ, τ) for E such that ρ is a positive multiple of s , $\tau \in \mathbb{R}/(2\pi\mathbb{Z})$, and $d\rho \wedge d\tau$ is positive. Moreover, when written as a function of ρ and τ , the function θ has the form

$$(2-17) \quad \theta(\rho, \tau) = \theta_E + e^{-r\rho}(b \cos(n(\tau + \sigma)) + \hat{\delta}),$$

where the notation is as follows: First, θ_E is the $|s| \rightarrow \infty$ limit of θ on E . Second, $r > 0$ when E is on the concave side and $r < 0$ when E is on the convex side of C . Third, b is a non-zero real number, n is a non-negative integer, but strictly positive if E is on the concave side of C_0 , and $\sigma \in \mathbb{R}/(2\pi\mathbb{Z})$. Finally, $\hat{\rho}$ and its first derivatives limit to zero as $|\rho| \rightarrow \infty$.

In what follows, the integer n that appears in (2–17) is denoted as $\deg_E(d\theta)$. In case that E is an end of C where $\lim_{|s| \rightarrow \infty} \theta$ is 0 or π , define $\deg_E(d\theta)$ to be zero.

Part 4 This part starts with the following key proposition:

Proposition 2.10 *Let C be a J' -pseudoholomorphic subvariety that is not \mathbb{R} -invariant and introduce k_C to denote the number of points in C_0 where θ is either zero or π . Then*

$$(2-18) \quad \sum_E \deg_E(d\theta) + \sum_z \deg(d\theta|_z) = N_+ + N_- + \hat{N} + k_C + 2\zeta - 2$$

where the first sum on the left hand side is indexed by the ends of C , and the second sum on the left hand side is indexed by the set of non-extremal critical points of θ on C_0 .

Proof of Proposition 2.10

This is a standard Euler class calculation given that all of the extremal points of θ 's pull-back to C_0 occur where $\theta = 0$ or $\theta = \pi$. The k_C summand on the right hand side of (2–18) accounts for the singular behavior of $d\theta$ at the points in C_0 where θ is either 0 or π .

This proposition has three immediate consequences. These are stated in the upcoming corollary. This corollary refers to the integer \wp that is defined from the singular points of ϕ_* as in the proof of Proposition 2.9. In particular, \wp can be defined for any J' -pseudoholomorphic subvariety that is not \mathbb{R} -invariant. To be precise, each zero of ϕ_* on such a subvariety has a version of (2–15) with an attending integer q . Then \wp is obtained by adding the resulting set of integers. The corollary that follows also refers to the integer \wp_* that is obtained by restricting the sum for \wp to the integers that are associated to the zeros of ϕ_* that lie where $\theta = 0$ or $\theta = \pi$. \square

Corollary 2.11 *Suppose that C is a J' -pseudoholomorphic subvariety that is not \mathbb{R} -invariant. Then*

- *The number of non-extremal critical points of θ 's pull-back to the model curve C_0 is greater than $N_- + \hat{N} + k_C + 2\zeta - 2$.*

- $\wp_* \leq c_{\hat{A}} - k_C$.
- $\wp \leq N_- + \hat{N} + c_{\hat{A}} + 2\zeta - 2$.

Proof of Corollary 2.11 The asserted bound on the critical points follows directly from (2–18) by virtue of the fact that $\deg_E(d\theta)$ is in all cases non-negative and at least 1 on the $(0, +, \dots)$ elements in \hat{A} .

As for the bounds on \wp_* and \wp , remark first that $d(\cos \theta)$ pulls back as zero at the singular points of ϕ_* . Now, to argue for the bound on \wp_* , focus attention on a point $z \in C_0$ where ϕ_* is zero and θ is either zero or π . Let q_z denote the integer that appears in the corresponding version of (2–15). Then it follows from (2–15) that z contributes a factor of at least $q_z + 1$ to the count for c_A . This observation implies the bound for \wp_* .

Granted the \wp_* bound, then the asserted bound on \wp then follows from the fact that $\wp - \wp_*$ is no greater than the second sum on the left hand side of (2–18). \square

2.F Local parametrizations for points in \mathcal{M}

A close reading of the proofs of Propositions 2.7 and 2.9 indicate that certain natural functions on the subvarieties in $\mathcal{M}_{\hat{A}}$ can serve as local coordinates. For example, the proof suggest that such is the case for the $\mathbb{R}/(2\pi\mathbb{Z})$ parameters that characterize concave side ends where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$. The purpose of this subsection is to prove that such is the case, and also to provide an expanded list of local coordinates.

To begin, partition \hat{A} into the disjoint subsets whereby any two elements in the same subset are identical and any two elements from distinct subsets are distinct. Let Λ denote this list of subsets. Label the elements in each subset in the partition Λ by consecutive integers starting at 1.

Now, let \mathcal{M}^Λ denote the set of elements of the form $(C, L \equiv \{L_\lambda\}_{\lambda \in \Lambda})$ where $C \in \mathcal{M}_{\hat{A}}$ and where any given L_λ is a 1–1 correspondence between the subset λ and the set of ends of C that contribute elements to λ . This space \mathcal{M}^Λ has a natural topology whereby the evident projection to $\mathcal{M}_{\hat{A}}$ is a covering map. It is fair to view \mathcal{M}^Λ as a moduli space of subvarieties with labeled ends.

The point of this is that a pair consisting of a subset λ from the partition Λ and an element $r \in \lambda$, define a function,

$$(2-19) \quad \varpi_{\lambda,r}: \mathcal{M}^\Lambda \rightarrow \mathbb{C}^*$$

that maps a given (C, L) to $\exp(c_{E(r)} + i\iota_{E(r)})$. Here, $E(r) \subset C$ is the end that L_λ assigns to r , while $c_{(\cdot)}$ and $\iota_{(\cdot)}$ are the continuous parameters that are associated to the

given end. In this regard, note that when λ consists of elements of the form $(0, +, \dots)$, then any $r \in \lambda$ version of $\varpi_{\lambda,r}$ maps to the unit circle in \mathbb{C} . Let

$$(2-20) \quad \varpi_+ : \mathcal{M}^\Lambda \rightarrow \times_{N_+} S^1$$

denote the product of the N_+ versions of $\varpi_{\lambda,r}$ where $\lambda \subset \hat{A}$ is a such a $(0, +, \dots)$ element.

Now consider the following:

Proposition 2.12 *Fix integers $b \in \{0, \dots, N_-\}$ and $c \leq N_- + \hat{N} + c_{\hat{A}} - b - 2$, then fix a size b subset of $(0, -, \dots)$ elements from \hat{A} and a c -element subset in $(0, \pi)$. Use B for the former and θ for the latter, and use $\mathcal{M}^B[\Theta] \subset \mathcal{M}^\Lambda$ to denote the subset of subvarieties with the following three properties: First, θ 's pull-back to the associated model curve has precisely c non-extremal critical points. Second, θ is the corresponding c -element set of critical values. Finally, the versions of $c_{(\cdot)}$ from (1-8) vanish for the ends from B but for no other convex side ends. If non-empty, this set $\mathcal{M}^B[\Theta]$ is a smooth submanifold of \mathcal{M}^Λ of dimension $N_+ + b + c + 2$. Furthermore, choose any map $\varpi_{\lambda,r}$ for which r is a $(\pm 1, \dots)$ from \hat{A} or a $(0, -, \dots)$ element from $\hat{A} - B$. Then, the map $\varpi_+ \times (\times_{(\lambda',r'):r' \in B} \varpi_{\lambda',r'}) \times \varpi_{\lambda,r}$ restricts over $\mathcal{M}[\Theta]$ as a smooth submersion to $(\times_{N_+} S^1) \times (\times_b S^1) \times \mathbb{C}^*$.*

Note that in the case $B = \emptyset$, then $\mathcal{M}^B[\Theta]$ is denoted by $\mathcal{M}[\Theta]$ and viewed as a submanifold of dimension $N_+ + c + 2$ in \mathcal{M} .

Local coordinates near any given subvariety in $\mathcal{M}^B[\Theta]$ can be obtained in the following manner: Let $(C, L) \in \mathcal{M}^B[\Theta]$ and let $\{z_1, \dots, z_c\} \subset C_0$ denote a labeling of the non-extremal critical points of θ 's pull-back to C_0 . This set denoted by $\text{Crit}(C)$ in what follows. Let $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ again denote the tautological map. There is an open neighborhood of C in \mathcal{M} whose subvarieties enjoy the following property: The model curve of each such subvariety can be viewed as the image in the normal bundle $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ of a section with everywhere very small norm. This is in accord with the story from Proposition 2.4 and the discussion prior to Proposition 2.9. If C' is in such a neighborhood and comes from a point near C in $\mathcal{M}^B[\Theta]$, then the non-extremal critical points of θ 's pull-back to the model curve of C' can be put in 1-1 correspondence with the points in the set $\{z_1, \dots, z_c\}$ by associating any given critical point of θ on C'_0 with the closest critical point in C_0 as measured by distance in $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. This understood, use $\{z'_1, \dots, z'_c\}$ to denote the corresponding labeled set of non-extremal critical points of θ 's pull-back to the model curve of C' . Note that the degree of vanishing of $d\theta$ at any given z'_k on the model curve for C' is identical to the degree that $d\theta$ vanishes on C' 's model curve at z_k .

To continue, suppose $z \in \{z_1, \dots, z_n\}$. Fix respective \mathbb{R} -valued lifts, \hat{t} and $\hat{\varphi}$, of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued functions t and φ that are defined on a neighborhood in $\mathbb{R} \times (S^1 \times S^2)$ of the image of z . Then set

$$(2-21) \quad v \equiv (1 - 3 \cos^2 \theta) \hat{\varphi} - \sqrt{6} \cos \theta \hat{t}.$$

The c versions of (2-21) define c functions, $\{v_1(\cdot), \dots, v_c(\cdot)\}$, on a neighborhood of C in $\mathcal{M}^B[\Theta]$ as follows: The value of v_k on a given subvariety C' is the value of the z_k version of (2-21) at the image in $\mathbb{R} \times (S^1 \times S^2)$ of the C' version of the θ -critical point z'_k .

To obtain $N_+ + b + 2$ additional functions, the map ϖ_+ from (2-20) pulls back an affine coordinate from each of the S^1 factors from its range space. Let $\{\varpi_{+1}, \dots\}$ denote an ordering of the resulting N_+ element set of such functions. Use $\{\varpi_{-1}, \dots, \varpi_{-b}\}$ to denote an ordering of the b affine function that are pulled back by viewing $\times_{(\lambda, e'): r' \in B} \varpi_{\lambda', r'}$ as a map from $\mathcal{M}^B[\Theta]$ to $\times_b S^1$. Finally, either choose one of the following: A pair (λ, r) with $\lambda \in \Lambda$ and r some $(\pm 1, \dots)$ element from \hat{A} . Or, a pair (λ, r) with $\lambda \in \Lambda$ and $r \in \hat{A} - B$ some $(0, -, \dots)$ element. Or, a point $z \in C_0$ where θ is 0 or π . In the first two cases, use $\varpi_{\lambda, r}$ from (2-19) to pull-back the standard complex coordinate on \mathbb{C} and call this coordinate ϖ' . For the third case, note that by virtue of (2-18), each C' from $\mathcal{M}^B[\Theta]$ near C has an unambiguous point in its model curve that corresponds to z and also maps very near to z in the $\theta \in \{0, \pi\}$ locus. Let z' denote the latter. Because C' comes from $\mathcal{M}^B[\Theta]$, the contribution from z' to the intersection number between C' and the $\{0, \pi\}$ locus is the same as that from z to C' 's intersection number with the $\{0, \pi\}$ locus. Use ϖ' in this case to denote the complex valued function on C' 's neighborhood that assigns to any given C' the value of the complex coordinate on the $\theta \in \{0, \pi\}$ locus at the image of z' .

Proposition 2.13 *The functions $\{v_j : 1 \leq j \leq c\}$, $\{\varpi_{+\alpha}\}$, $\{\varpi_{-\alpha}\}$ and ϖ' together define local coordinates on a neighborhood of C in $\mathcal{M}^B[\Theta]$.*

Proof of Propositions 2.12 and 2.13

The proof is given in five parts.

Part 1 This part provides some comments on the dimension count for $\mathcal{M}^B[\Theta]$. For this purpose, let $C \in \mathcal{M}^B[\Theta]$ and let $\text{Crit}(C)$ denote the set of non-extremal critical points of θ 's pull-back to C 's model curve. As noted earlier, if $C' \in \mathcal{M}^B[\Theta]$ is near C , then corresponding critical points of θ 's pull-back to the model curves for C and C' must have the same values for θ and also for $\deg(d\theta|_{(\cdot)})$. This places $2 \sum_{z \in \text{Crit}(C)} \deg(d\theta|_z) - c$ constraints on the subvarieties that are near C in $\mathcal{M}^B[\Theta]$.

Requiring that they also have the same values for the functions v_j adds c more constraints.

There are more constraints that come from C' 's ends. To see these, let $E \subset C$ denote an end that corresponds to some $(0, \dots)$ element in \hat{A} . A given C' near to C in \mathcal{M} has a corresponding end E' . It follows from (2–17) and (2–18) that θ 's pull-back to the model curve of C' has critical points at large values of $|s|$ on E' in the case that $\deg_{E'}(d\theta) < \deg_E(d\theta)$. Taking this into account finds an additional set of $2(\sum_E \deg_E(d\theta) - N_+) - b$ constraints on the elements in $\mathcal{M}^B[\Theta]$.

Even more constraints arise from the local intersection numbers of the subvarieties with the $\theta \in \{0, \pi\}$ locus. To elaborate, let z denote a point in C' 's model curve that maps to a point in this locus, and let p denote z 's contribution to the intersection number $c_{\hat{A}}$. If $C' \in \mathcal{M}^B[\Theta]$ is near C , then (2–18) requires C' to have a corresponding point very near z in its model curve which contributes p to the intersection number between C' and the $\theta \in \{0, \pi\}$ locus. All such intersection points thus account for an additional $2(c_A - k_C)$ constraints on the subvarieties in $\mathcal{M}^B[\Theta]$.

When totalled, the number of constraints that must be satisfied for placement in $\mathcal{M}^B[\Theta]$ is

$$(2-22) \quad 2\left(\sum_E \deg_E(d\theta) + \sum_{z \in \text{Crit}(C)} \deg(d\theta|_z) + c_{\hat{A}} - k_C\right) - 2N_+ - b - c.$$

What with (2–18), this number is $N_+ + b + c$ less than the dimension of \mathcal{M} .

Part 2 Granted this count, the assertion that $\mathcal{M}^B[\Theta]$ is a manifold of the asserted dimension with the given local coordinates is proved using an application of the implicit function theorem. The application requires the introduction of the linearized constraints on the tangent space to C in \mathcal{M} . For this purpose, identify $T\mathcal{M}|_C = \text{kernel}(D_C)$ and let $K^* \subset \text{kernel}(D_C)$ denote the subspace of vectors that satisfy all of the linear constraints and also annihilate all of the functions that are listed in Proposition 2.13. Both Propositions 2.12 and 2.13 follow from the implicit function theorem if K^* is trivial.

The identification given below of K^* requires a preliminary digression to introduce some notation that concerns a given $\lambda \in \text{kernel}(D_C)$. First, λ^N is used below to denote λ 's image in N via the projection from $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. Second, $\langle d\theta, \lambda \rangle$ is used to denote the pairing on the complement of the $\theta = 0$ and π points between λ and $d\theta$ when the latter form is viewed in $\phi^*T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$. Finally, when z is a point in C' 's model curve where θ is 0 or π , then $r(\lambda)$ denotes the projection of λ to the holomorphic tangent bundle of the $t = \text{constant}$ pseudoholomorphic subvariety

through the given point. These subvarieties are described in [18, Section 4a], and for the present purposes, it is enough to know that these subvarieties are embedded, they foliate $\mathbb{R} \times (S^1 \times S^2)$, and the tangent bundle to any such subvariety on the $\theta = 0$ or π locus is normal to the locus.

What follows are the conditions for membership in K^* :

- (2–23) • If ϖ' is defined by an element $r \in \widehat{A}$, then $\deg_E(\lambda^N) > 0$ on the end $E \subset C$ that corresponds to r .
- If $E \subset C$ is an end where $\deg_E(d\theta) > 0$, then $\delta_E(\lambda^N) \geq \deg_E(d\theta)$.
 - If $z \in C_0$ is a non-extremal critical point of the pull-back of θ and u is a local holomorphic coordinate for a disk in C_0 centered at z , then $\langle d\theta, \lambda \rangle = o(|u|^k)$ near z with $k \geq \deg(d\theta|_z)$.
 - If $z \in C_0$ is a point where $\theta = 0$ or π , let p denote z 's contribution to $c_{\widehat{A}}$. Let u denote a complex coordinate for a disk in C_0 centered at z . If $p \geq 2$, then $r(\lambda) = o(|u|^k)$ near z with $k \geq p - 1$.
 - Suppose that $z \in C_0$ is a point where $\theta = 0$ and that z is used to define ϖ' . If $\phi_*|_z$ is zero, then $\lambda|_z$ must also vanish. If $\phi_*|_z$ is non-zero, then $\eta \sim |u|^k$ near z where $k \geq p$.

Note that the final three constraints only involve η where ϕ_* is non-zero.

Part 3 As there are $\dim(\text{kernel}(D_C))$ conditions in (2–23), a proof that they are linearly independent proves that $K^* = \{0\}$. The argument for linear independence invokes two observations that concern a section, ξ , of the W summand in $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. To set the stage, let z denote a given point in C_0 and let q_z denote the integer that appears in z 's version of (2–15). Thus, $q_z = 0$ if $\phi_*|_z \neq 0$ and $q_z > 0$ otherwise.

Here is the first observation: If $z \in \text{Crit}(C)$ and u is a complex coordinate for a disk centered in C_0 with center z , then $\langle d\theta, \xi \rangle = o(|u|^k)$ near z with $k \geq \deg(d\theta|_z) - q_z$. Moreover, if $\langle d\theta, \xi \rangle = o(|u|^k)$ with $k \geq \deg(d\theta|_z)$, then ξ near z is the image via ϕ_* of a section of $T_{1,0}C_0$. This observation follows from (2–15) and (2–16).

The second observation concerns a point $z \in C_0$ where θ is 0 or π , and it involves the integer, p_z , that z contributes to $c_{\widehat{A}}$. Here is the observation: Let u denote a complex coordinate for a disk in C_0 centered at z . Then $r(\xi) = o(|u|^k)$ where $k \geq p_z - q_z - 1$ near z , and if $\xi|_z = 0$, then $r(\xi) = o(|u|^k)$ with $k \geq p_z - q_z$. Moreover, if $r(\xi) = o(|u|^k)$ with $k \geq p_z - 1$, then ξ near z is the image via ϕ_* of a section of $T_{1,0}C_0$. This observation follows from (2–15).

Part 4 The analysis of the conditions in (2–23) begins by considering those that involve only the projection λ^N of λ . In particular, such is the case for the first two. As is explained next, some of the others also involve only λ^N . In particular, suppose first that $z \in \text{Crit}(C)$. As noted previously, $\deg(d\theta|_z) \geq q_z$ and so, by virtue of Part 3’s first observation, the vanishing to order $\deg(d\theta|_z)$ of $\langle d\theta, \lambda \rangle$ at z puts $2(\deg(d\theta|_z) - q_z)$ constraints on λ^N as it requires that $\lambda^N = o(|u|^k)$ near z with $k \geq (\deg(d\theta|_z) - q_z)$.

Consider next a point $z \in C_0$ where θ is 0 or π . As noted previously, $p_z \geq q_z + 1$, and so by virtue of Part 3’s second observation, the vanishing to order $p_z - 1$ of $r(\lambda)$ at z forces λ^N to be $o(|u|^k)$ near z where $k \geq p_z - q_z - 1$. Of course, this is a constraint only in the case that $p_z > q_z + 1$. In the case that z is used to define ϖ' , the final point in (2–23) forces λ^N to be $o(|u|^k)$ near z with $k \geq p_z - q_z$.

If λ^N is to satisfy all of these constraints, then

$$(2-24) \quad \sum_E \delta_E(\lambda^N) + \sum_z \deg_z(\lambda^N) \\ \geq \sum_E \deg_E(d\theta) + \sum_{z \in \text{Crit}(C)} (\deg(d\theta|_z) - q_z) + \sum_{z: \theta=0 \text{ or } \pi} (p_z - q_z - 1) + 1,$$

and this, according to (2–18), is greater than $\deg(N) = N_+ + N_- + \widehat{N} + c_{\widehat{A}} - 2 - \wp$. As a consequence, $\lambda^N = 0$ if $\lambda \in K^*$.

Part 5 As just noted, any $\lambda \in K^*$ is a section of W . Granted the first observation from Part 3, the third condition in (2–23) implies that such λ is the image via ϕ_* near each $z \in \text{Crit}(C)$ of a section of $T_{1,0}C_0$. What with second observation from Part 3, the fourth condition in (2–23) implies that λ is also in the image of ϕ_* near each point in C_0 where θ is 0 or π . Thus, λ is in the image of ϕ_* on the whole of C_0 since any $\theta \in (0, \pi)$ zero of ϕ_* is a zero of $\phi^*d\theta$. However, as noted in the proof of Proposition 2.9, the kernel of D^W has only 0 from the image of ϕ_* . Thus, $K^* = \{0\}$ as required. \square

2.G Slicing curves by θ level sets

This subsection constitutes a digression of sorts to discuss some algebraic and geometric issues that arise in conjunction with the use of the critical points of θ ’s pull-back to construct coordinates on $\mathcal{M}_{\widehat{A}}$. Some of these issues appear both implicitly and explicitly in the subsequent sections of this article, and they play a central role in the sequel to this article. In any event, the subsection starts by examining the nature of the θ -level sets in any given subvariety from \mathcal{M} . These level sets are then used to associate to

each such variety a certain connected, contractible graph with labeled vertices and labeled edges. The discussion of the constant θ loci is contained in Part 1–Part 3 of this subsection, while Part 4 contains the definition of the associated graph. In all of what follows, it is assumed that the subvariety C in question is not an \mathbb{R} –invariant cylinder, thus not of the form $\mathbb{R} \times \gamma$ where $\gamma \subset S^1 \times S^2$ is a Reeb orbit.

Part 1 To begin the story here, suppose that \hat{A} is an asymptotic data set, J' is an admissible almost complex structure, and C is a subvariety from the J' –version of $\mathcal{M}_{\hat{A}}$. Let C_0 again denote the model curve for C . Introduce now the locus, $\Gamma \subset C_0$, which is defined as follows: The components of this set consist of the level sets of θ on C_0 that are either zero dimensional, singular or non-compact. In particular, Γ contains all of the critical points of θ on C_0 .

To continue, note that any given component of Γ can be viewed as the embedded image in C_0 of an oriented graph with labeled edges and vertices. To elaborate, the zero dimensional components of Γ are the points in C_0 where θ is 0 or π . In particular each zero dimensional component is a graph with a single vertex, the latter labeled by a non-zero integer whose absolute value is the contribution of the given $\theta = 0$ or π point to $c_{\hat{A}}$. The sign of the integer is positive when the θ value is 0 and the integer is negative when the θ value is π .

Each singular point in a non-point like component of Γ is a critical point of θ . These points constitute the vertices of the corresponding graph. The components of the complement of these singular points constitute the edges in the graph. In this regard, these edges are henceforth referred to as ‘arcs’ so as not to confuse them with the edges in the graph that is defined subsequently in Part 4 from C . These arcs are oriented by the pull-back of the 1–form $x \equiv (1 - 3 \cos^2 \theta)d\varphi - \sqrt{6} \cos \theta dt$. Note that this 1–form is nowhere zero on the smooth portion of any given θ level set in C_0 for the following reason: The differential of the contact form in (1–1) is $\sqrt{6} \sin \theta d\theta \wedge x$ and because J' is admissible and C is J' –pseudoholomorphic, this form is positive on TC_0 save at the critical points of θ where it vanishes.

Because a given vertex in a non-point like component of Γ is a critical point of θ , it has an even number, at least 4, of incident arcs. This follows from the form of $d\theta$ in (2–16). Moreover, half of the incident arcs are oriented to point towards the vertex and half are oriented to point away. Indeed, a circumnavigation of a small radius circle about the critical point will alternately meet inward pointing and outward pointing arcs. For example, if $\Gamma_* \subset \Gamma$ is a compact, singular component with a single, non-degenerate, non-extremal critical point, then the associated graph has a single vertex and looks like the figure ‘8’.

Meanwhile, the complement of the θ critical points in a given non-compact component of Γ has an even number of unbounded arcs. Indeed, this follows from (2–17). In particular, any given end of C where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π has the following property: Let n denote the integer that appears in (2–17) for the given end. Then any sufficiently large and constant $|s|$ slice of the end intersects precisely $2n$ components of Γ and this intersection is transverse. Moreover, a circumnavigation of the constant $|s|$ slice meets components whose orientations alternate towards increasing $|s|$ and towards decreasing $|s|$.

By way of an example, suppose that E is a concave side end of C_0 where $\lim_{|s| \rightarrow \infty} \theta \in (0, \pi)$, and suppose that this limit is distinct from all other $|s| \rightarrow \infty$ limits of θ on C . Suppose as well that this limit is distinct from all of the critical values of θ on C_0 . Then the large $|s|$ portion of E will intersect precisely one component of Γ , the latter a smooth, properly embedded copy of \mathbb{R} whose large $|s|$ portions are properly embedded in the large $|s|$ part of E .

Part 2 By virtue of the definition of Γ , any given component $K \subset C_0 - \Gamma$ is a cylinder to which $d\theta$ and x pullback without zeros. In fact, θ and the restriction of x to the constant θ level sets of K can be used to give coordinates to such a cylinder. To elaborate, let (θ_0, θ_1) denote the range of θ on K . Next, let q and q' denote the respective integrals around the constant θ slices of K (as oriented by x) of the closed forms $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$. Then K can be parametrized by the open cylinder $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ so that the restriction to K of the tautological immersion of C_0 into $\mathbb{R} \times (S^1 \times S^2)$ has a rather prescribed form. To be more specific, let $\sigma \in (\theta_0, \theta_1)$ and $v \in \mathbb{R}/(2\pi\mathbb{Z})$ denote the coordinates for the cylinder. Written using these coordinates, the tautological immersion involves two smooth functions on $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$, these denoted by a and w ; and it sends any given point (σ, v) to the point where

$$(2-25) \quad \begin{aligned} (s = a, t = qv + (1 - 3 \cos^2 \theta)w \pmod{2\pi\mathbb{Z}}, \\ \theta = \sigma, \varphi = q'v + \sqrt{6} \cos \theta w \pmod{2\pi\mathbb{Z}}). \end{aligned}$$

Note that the J' -pseudoholomorphic nature of the immersion of K requires that the pair (a, w) obey a certain non-linear differential equation. For example, in the case where $J' = J$, this equation reads

$$(2-26) \quad \begin{aligned} \alpha_Q a_\sigma - \sqrt{6} \sin \sigma (1 + 3 \cos^2 \sigma) w a_v &= -\frac{1 + 3 \cos^4 \sigma}{\sin \sigma} \left(w_v - \frac{1}{1 + 3 \cos^4 \sigma} \beta \right) \\ (\alpha_Q w)_\sigma - \sqrt{6} \sin \sigma (1 + 3 \cos^2 \sigma) w w_v &= \frac{1}{\sin \sigma} a_v, \end{aligned}$$

Here, $\alpha_Q = \alpha_Q(\sigma)$ is the function

$$(2-27) \quad \alpha_Q = (1 - 3 \cos^2 \sigma)q' - \sqrt{6} \cos \sigma q,$$

and $\beta = q(1 - 3 \cos^2 \sigma) + q' \sqrt{6} \cos \sigma \sin^2 \sigma$. In these equations and below, Q denotes the pair (q, q') .

By the way, α_Q is necessarily positive on (θ_0, θ_1) by virtue of the fact that the parametrization in (2-25) of K pulls back the exterior derivative of the contact form α as

$$(2-28) \quad \sqrt{6} \sin \sigma \alpha_Q(\sigma) d\sigma \wedge dv.$$

In this regard, keep in mind that the form $d\alpha$ is non-negative on J' -pseudoholomorphic 2-planes in $\mathbb{R} \times (S^1 \times S^2)$. Moreover, $d\alpha$ is zero on such a plane only if the latter is spanned by ∂_s and the Reeb vector field $\hat{\alpha}$.

This last conclusion has the following converse: Suppose that (a, w) are any given pair of functions on $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$. Then, the resulting version of (2-25) immerses the points in its domain where α_Q is positive.

Here is one final remark about any map having the form given in (2-25): Suppose that α_Q is positive on (θ_0, θ_1) . Now, let (a, w) denote any given pair of functions on the cylinder $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$. By virtue of the fact that the coordinates t and φ are defined only modulo $2\pi\mathbb{Z}$, the image cylinder in $\mathbb{R} \times (S^1 \times S^2)$ via the map in (2-25) is unchanged under the action of $\mathbb{Z} \times \mathbb{Z}$ on the space of function pairs (a, w) whereby a given integer pair $N = (n, n')$ acts to send (a, w) to (a^N, w^N) with the latter given by

$$(2-29) \quad \begin{aligned} a^N(\sigma, v) &= a\left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_Q(\sigma)}\right) \\ \text{and} \quad w^N(\sigma, v) &= w\left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_Q(\sigma)}\right) + 2\pi \frac{qn' - q'n}{\alpha_Q(\sigma)}. \end{aligned}$$

Part 3 This part of the subsection discusses the behavior of the parametrization in (2-25) at points near the boundary of the closure of the parametrizing cylinder.

To start, remark that if a given $\theta_* \in \{\theta_0, \theta_1\}$ is not achieved by θ on the closure of K , then there exists $\varepsilon > 0$ such that the portion of K where $|\theta - \theta_*| \leq \varepsilon$ is properly embedded in an end of C . In particular, the constant θ slices of this portion of K are isotopic to the constant $|s|$ slices when θ is very close to θ_* . Moreover, if $\theta_* \notin \{0, \pi\}$, then such an end is on the convex side of C and the associated integer n that appears in (2-17) is zero.

On the other hand, if $\theta_* \in \{\theta_o, \theta_1\}$ is neither 0 nor π and if θ_* is achieved on the closure of K , then the complement of the θ critical points in the $\theta = \theta_*$ boundary of this closure is the union of a set of disjoint, embedded, open arcs. The closures of each such arc is also embedded. However, the closures of more than two arcs can meet at any given θ -critical point. Conversely, every arc in any given component of Γ is entirely contained in the boundary of the closures of precisely two components of $C_0 - \Gamma$.

This decomposition of the $\theta = \theta_*$ boundary of K into arcs is reflected in the behavior of the parametrizations in (2–25) as σ approaches θ_* . To elaborate, each critical point of θ on the $\theta = \theta_*$ boundary of the closure of K labels one or more distinct points on the $\sigma = \theta_*$ circle in the cylinder $[\theta_o, \theta_1] \times \mathbb{R}/(2\pi\mathbb{Z})$. These points are called ‘singular points’. Meanwhile, each end of C that intersects the $\theta = \theta_*$ boundary of the closure of K in a set where $|s|$ is unbounded also labels one or more distinct points on this same circle. The latter set of points are disjoint from the set of singular points. A point from this last set is called a ‘missing point’.

The complement of the set of missing and singular points is a disjoint set of open arcs. Each point on such an arc has a disk neighborhood in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$ on which the parametrization in (2–25) has a smooth extension as an embedding into $\mathbb{R} \times (S^1 \times S^2)$ onto a disk in C . This last observation is frequently used in subsequent arguments from this article and from the sequel.

As might be expected, the set of arcs that comprise the complement of the singular and missing points are in 1–1 correspondence with the set of arcs that comprise the $\theta = \theta_*$ boundary of the closure K . In particular, the extension to (2–25) along any given arc in the $\sigma = \theta_*$ boundary of $[\theta_o, \theta_1] \times \mathbb{R}/(2\pi\mathbb{Z})$ provides a smooth parametrization of the interior of its partner in the $\theta = \theta_*$ boundary of the closure of K .

By way of an example, consider the case that θ_* is a critical value of θ on C_0 that is realized by a single critical point with the latter non-degenerate. Assume further that θ_* is not an $|s| \rightarrow \infty$ limit of θ on C . Thus, the critical locus is a ‘figure 8’. In this case, there are three components of $C_0 - \Gamma$ with boundary on this locus, one whose boundary maps to the top circle in the figure 8, another whose boundary maps to the lower circle, and a third whose boundary traverses the whole figure 8. The first two have but one singular point on the $\sigma = \theta_*$ boundary of any parametrizing domain, while the third has two singular points.

For a second example, suppose that θ_* is neither 0 nor π and is the $|s| \rightarrow \infty$ limit of θ on a convex side end, E , for which $\deg_E(d\theta) = 1$. Suppose, in addition, that θ_* is not the $|s| \rightarrow \infty$ limit of θ on any other end of C . In this case, the corresponding $\theta = \theta_*$ component of Γ is a properly embedded copy of \mathbb{R} . Furthermore, there are two

components of $C_0 - \Gamma$ whose closures lie in this $\theta = \theta_*$ component of Γ , and both have just a single missing point on the $\sigma = \theta_*$ boundaries of any of their parametrizing domains.

In the case that $\theta_* \in \{\theta_0, \theta_1\}$ is either 0 or π and θ takes value θ_* on the closure of K , then the map in (2-25) extends to the $\sigma = \theta_*$ boundary of the cylinder as a smooth map that sends this boundary to a single point. This extended map factors through a pseudoholomorphic map of a disk into $\mathbb{R} \times (S^1 \times S^2)$ with the $\sigma = \theta_*$ circle being sent to the disk's origin.

Part 4 The graph assigned to a given C from $\mathcal{M}_{\hat{A}}$ is denoted here by T_C . As remarked at the outset, this is a connected, contractible graph with labeled edges and vertices. In this regard, the edges of T_C are in 1-1 correspondence with the components of $C_0 - \Gamma$. If e denotes an edge, then e is labeled by an integer pair, $Q_e \equiv (q_e, q_e')$, these being the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about the constant θ slices of the corresponding component of $C_0 - \Gamma$. Here, as in Part 1, these slices are oriented using the pull-back of the 1-form $(1 - 3 \cos^2 \theta) d\varphi - \sqrt{6} \cos \theta dt$.

The multivalent vertices of T_C are in 1-1 correspondence with the subsets of a certain partition of the components of Γ . To define this partition, first introduce a new graph, G , as follows: The vertices of G are in 1-1 correspondence with the components of Γ . Meanwhile, an edge connects two vertices of G when there is an end of C with the following property: Every sufficiently large and constant $|s|$ slice of the end intersects at least one arc from each of the corresponding components of Γ . With G understood, then the components of G naturally partition the set of components of Γ . For example, every compact component of Γ defines its own, single point set in this partition.

The set of multivalent vertices in T_C are in 1-1 correspondence with this partition of Γ . If o is a multivalent vertex of T_C , then the incident edges to o label the components of $C_0 - \Gamma$ whose closure intersects that part of Γ that is assigned to o .

The monovalent vertices in T_C are in 1-1 correspondence with the elements in the union of three distinct sets. The first set consists of points where C intersects the $\theta = 0$ and $\theta = \pi$ loci. In this regard, a given $\theta = 0$ or $\theta = \pi$ point can label more than one monovalent vertex of T_C . To elaborate, suppose that a small ball about such a point intersects C in some $k \geq 0$ irreducible components that all meet at the given point. Then this point labels k monovalent vertices of T_C . The second set consists of the ends of C where the $|s| \rightarrow \infty$ limit of θ is either 0 or π . The third set consists of the convex side ends of C where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π and where the integer n in (2-17) is zero. Said differently, the second and third sets consist of those ends of C where the $|s| \rightarrow \infty$ limit of θ is not achieved at any sufficiently large value of $|s|$. A given monovalent vertex lies on the most obvious edge.

Note that various versions of T_C will be defined here and in the sequel to this article that differ in the complexity of the labels that are assigned to the vertices. In all versions, the vertex label contains an angle in $[0, \pi]$, this the obvious one available. Elements from \hat{A} are part of labels that are used in the next subsequent sections of this article. The most sophisticated labeling occurs in the sequel to this article where any given multivalent vertex label is a certain sort of graph, this defined from the components of Γ that are contained in the corresponding partition subset.

3 Existence

This and the remaining sections derive necessary and sufficient conditions that insure that any given $\mathcal{M}_{\hat{A}}$ is non-empty. The result is a proof of Theorem 1.3. The strategy used here is to construct proper immersions of multi-punctured spheres with the correct $|s| \rightarrow \infty$ asymptotics and then deform them so that the result is pseudoholomorphic.

It is assumed here that an asymptotic data set \hat{A} has been specified that obeys a certain set of auxiliary constraints that differ from those stated in Theorem 1.3. Section 5 explains why the constraints listed here are satisfied if and only if \hat{A} satisfies the conditions in Theorem 1.3.

3.A An associated graph

Granted that \hat{A} has been specified, fix a partition of the set of $(0, +, \dots)$ elements in \hat{A} subject to the following constraint: The integer pairs from any two elements in the same partition subset define the same angle via (1-7). Let \wp denote the given partition. What follows is a description of a contractible graph, T , with labeled vertices and labeled edges that is defined using \hat{A} and \wp . This graph is used in the subsequent construction as a blueprint of sorts for constructing the initial subvariety in $\mathbb{R} \times (S^1 \times S^2)$.

The story on T begins with the remark that T has $N_- + \hat{N} + c_{\hat{A}}$ monovalent vertices, N_{\wp} bivalent vertices and $N_- + \hat{N} + c_{\hat{A}} - 2$ trivalent vertices. Here N_{\wp} denotes the number of sets in the partition \wp . A subset of monovalent vertices with $N_- + \hat{N}$ elements are labeled by assigning a 1-1 correspondence between the vertices in the subset and the subset of \hat{A} whose elements are either of the form $(\pm 1, \dots)$ or $(0, -, \dots)$. Of those that remain, c_+ are labeled (1) and c_- by (-1) .

The bivalent vertices are labeled by assigning a 1-1 correspondence between the set of such vertices and elements in \wp . Thus, each bivalent vertex is labeled by a partition subset.

The labeling just described associates an angle in $[0, \pi]$ to each monovalent and each bivalent vertex; this is the angle 0 when the vertex has label either $(1, \dots) \in \hat{A}$ or (1) , the angle π when the label is either $(-1, \dots) \in \hat{A}$ or (-1) . Meanwhile, a monovalent vertex with label $(0, -, \dots)$ from \hat{A} is assigned the angle that is defined via (1–7) by the integer pair from this element. Finally, a bivalent vertex is assigned the angle that is defined via (1–7) by the integer pair from any element in its corresponding partition subset.

As for the trivalent vertices, each is labeled by an angle in $(0, \pi)$ so that no two vertices are assigned the same angle, and none are assigned an angle that is associated to any monovalent or bivalent vertex.

There are three further constraints on the angle assignments to the vertices of T . Here is the first:

Constraint 1

- (a) *The vertices that share an edge have distinct angle assignments.*
- (b) *The angle assigned to any given multivalent vertex is neither a largest nor a smallest angle in the set of angles that are assigned those vertices on its incident edges.*

Each edge of T is labeled by a non-trivial, ordered pair of integers. If e denotes an edge, then its assigned pair is denoted here by Q_e or $(q_e, q_{e'})$. What follows is the second constraint.

Constraint 2 *These integer pair assignments to the edges are constrained as follows:*

- (3–1) • *If o is a monovalent vertex on e with a 4–tuple label from \hat{A} , then $Q_e = \pm P_o$ where P_o is the integer pair from the label. Here, the $+$ sign appears if and only if one of the following hold:*
 - (a) *o 's angle is in $(0, \pi)$ and it is the smaller of e' 's vertex angles.*
 - (b) *o is labeled by either a $(1, -, \dots)$ or a $(-1, +, \dots)$ element in \hat{A} .*
- *If o is a monovalent vertex with label $\delta \in (\pm 1)$, then its incident edge, e , has $q_e = 0$ and $q_{e'} = -1$.*
- *If o is a bivalent vertex with incident edges e and e' with the convention that e connects o to a vertex with smaller angle label, then $Q_e - Q_{e'} = P_o$ where the integer pair P_o is the sum of the pairs from the elements that comprise o 's partition subset.*

- If o is a trivalent vertex with incident edges e, e' and e'' , then $Q_e - Q_{e'} - Q_{e''} = 0$ given that the angle labels of the vertices opposite o on e' and e'' lie on the same side of the angle that labels o in $(0, \pi)$.

With the collection $\{Q_e\}$ now defined, here is the third constraint on the angle assignments to the vertices in T :

Constraint 3 Let e denote any given edge of T and let $\theta_o < \theta_1$ denote the angles that are assigned the vertices on e . Then

$$q'_e(1 - 3 \cos^2 \theta) - q_e \sqrt{6} \cos(\theta) \geq 0$$

at all $\theta \in [\theta_o, \theta_1]$ with equality if and only if θ is either θ_o or θ_1 and the corresponding vertex is monovalent with label $(0, -, \dots)$ from \hat{A} .

A graph T that obeys all of the preceding constraints is said here to be a ‘moduli space graph’ for \hat{A} .

Two moduli space graphs are said here to be isomorphic when there is an isomorphism of the unlabeled graphs that preserves the labelings of the vertices and edges.

To explain the relevance of such a graph to $\mathcal{M}_{\hat{A}}$, remember that any $C \in \mathcal{M}_{\hat{A}}$ defines a graph T_C as described in Part 4 of Subsection 2.G. In particular, if $\mathcal{M}_{\hat{A}}$ is non-empty, then according to Propositions 2.12 and 2.13, the version of T_C that is defined by any sufficiently generic choice of $C \in \mathcal{M}_{\hat{A}}$ defines just such a moduli space graph for a particular choice of \wp . To elaborate, the aforementioned propositions guarantee an open and dense subset in $\mathcal{M}_{\hat{A}}$ whose subvarieties have the following property: All of the critical points of the function $\cos(\theta)$ on the corresponding model curve are non-degenerate, and there are $N_- + \hat{N} + c_{\hat{A}} - 2$ non-extremal critical values with no two identical and none equal to an $|s| \rightarrow \infty$ limit of θ . If C is such a generic subvariety, then the corresponding graph T_C has only monovalent, bivalent and trivalent vertices. The angles of the vertices satisfy the requirements in Constraint 1. Meanwhile, the requirements of Constraint 2 are also met when the partition assigned to any given bivalent vertex consists of the 4-tuples in \hat{A} that label those ends of the subvariety that contain the very large $|s|$ parts of the corresponding components of the locus Γ . As explained in Subsection 2.G, the third constraint is met automatically.

Henceforth T_C denotes the just described moduli space version of the graph from Part 4 of Subsection 2.G in the case that $C \in \mathcal{M}_{\hat{A}}$ satisfies the stated genericity requirement.

As just indicated, if $\mathcal{M}_{\hat{A}}$ is non-empty, then it has a moduli space graph. The theorem that follows states this fact and its converse:

Theorem 3.1 *The space $\mathcal{M}_{\hat{A}}$ is non-empty if and only if \hat{A} has a moduli space graph. Moreover, if T is a moduli space graph, then there is a subvariety in $\mathcal{M}_{\hat{A}}$ whose version of $T_{(\cdot)}$ is isomorphic to T .*

The criteria here for $\mathcal{M}_{\hat{A}}$ were suggested by observations of Michael Hutchings.

The proof starts by assuming the existence of a moduli space graph for \hat{A} and ends with the conclusion that $\mathcal{M}_{\hat{A}}$ is non-empty. This proof occupies the remainder of Section 3 and all of Section 4.

3.B Parametrizations of cylinders in $\mathbb{R} \times (S^1 \times S^2)$

As remarked above, the graph T is used as a blue-print for the construction of a properly immersed, multi-punctured sphere in $\mathbb{R} \times (S^1 \times S^2)$. All of the monovalent vertices with 4-tuple labels from \hat{A} correspond to ends of the surface; here, the label from \hat{A} on a given vertex is used to specify the asymptotic behavior of the corresponding end. The monovalent vertices with either (1) or (-1) labels correspond to intersection points between the surface and the respective $\theta = 0$ or $\theta = \pi$ cylinders. Meanwhile, the trivalent vertices label the index one critical points of the restriction to the surface of the function $\cos(\theta)$. The angle label of a vertex gives the value of θ at the corresponding critical point. The only local maxima or minima of $\cos(\theta)$ on the surface are relegated to the intersection with the respective $\theta = 0$ and $\theta = \pi$ cylinders.

Meanwhile, any given edge of T labels an open, cylindrical component of the surface where θ ranges between the values given by the edge's end vertices. The ordered integer pair that is associated to the edge specifies the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ around any constant θ slice as oriented by $(1 - 3 \cos^2 \theta) d\varphi - \sqrt{6} \cos \theta dt$. The incidence relations at the vertices direct the manner in which the edge labeled cylinders attach to form a closed surface. In this regard, the closure of any component cylinder whose edge label ends in a (1) or (-1) labeled monovalent vertex is a disk that intersects the respective $\theta = 0$ or $\theta = \pi$ locus.

The basic building blocks for the surface are thus the edge labeled cylinders. When $e \subset T$ denotes an edge, its corresponding cylinder is denoted by K_e . Let o and o' denote the angles that label the end vertices of e with the convention that the θ -label of o is less than that of o' . These labels are respectfully denoted here by θ_o and $\theta_{o'}$.

The parametrization of K_e is via a map from its 'parametrizing cylinder', this being the interior of $[\theta_o, \theta_{o'}] \times \mathbb{R}/(2\pi\mathbb{Z})$. With σ denoting the coordinate on $[\theta_o, \theta_{o'}]$ and v an affine coordinate on $\mathbb{R}/(2\pi\mathbb{Z})$, the parametrizing map can be written using two

functions on the open cylinder, (a_e, w_e) . To be precise, the parametrizing map sends any given point (σ, v) to the $Q \equiv Q_e$ and $(a, w) \equiv (a_e, w_e)$ version of (2–25). Thus,

$$(3-2) \quad \begin{aligned} (s = a_e, t = q_e v + (1 - 3 \cos^2 \sigma) w_e \pmod{2\pi}), \\ \theta = \sigma, \varphi = q'_e v + \sqrt{6} \cos \sigma w_e \pmod{2\pi}) \end{aligned}$$

Unless specified to the contrary, a ‘parametrization’ is one given as in (3–2).

For future reference, note that the 2–form $d\sigma \wedge dv$ orients K_e . Also, note that the map into $\mathbb{R} \times (S^1 \times S^2)$ as defined by (3–2) defines an immersion. As is explained next, this is a consequence of the positivity of the $Q = Q_e \equiv (q_e, q'_e)$ version in (2–27) of the function α_Q . To see why, first let ϕ denote the parametrizing map to $\mathbb{R} \times (S^1 \times S^2)$ and reintroduce the contact 1–form α from (1–1) and (1–2). Then $\phi^* d\alpha$ is the 2–form that appears in (2–28). Granted that such is the case, the rank of ϕ^* is two on the parametrizing cylinder.

The next section provides a specific version of each (a_e, w_e) . These versions are chosen to meet the following five criteria:

- (3–3) • *The collection $\{(a_e, w_e)\}_{e \in T}$ are constrained near the boundaries of their corresponding parametrizing cylinders so as to insure that the closure of $\cup_{e \in T} K_e$ is the image in $\mathbb{R} \times (S^1 \times S^2)$ via a proper immersion of an oriented, multiply punctured sphere.*
- *The singularities of this immersion are transversal double points with positive local intersection number.*
- *The critical points on the multi-punctured sphere of the pull-back of θ are non-degenerate, the 1–form $J \cdot d\theta$ pulls back as zero at these critical points, and the symplectic form ω pulls back as a positive form at these critical points. Moreover, no critical point maps to an immersion point of C .*
- *The subvariety has the asymptotics of a subvariety from $\mathcal{M}_{\hat{A}}$.*
- *The level sets of the pull-back of θ to the multiply punctured sphere defines a moduli-space graph that is isomorphic to the given graph T .*

The precise meaning of the fourth point is given in the definition that follows.

Definition 3.2 A subvariety $C \subset \mathbb{R} \times (S^1 \times S^2)$ is said to have the asymptotics of a subvariety from $\mathcal{M}_{\hat{A}}$ when the requirements listed below are met.

Requirement 1: There is a compact subset in C whose complement is a disjoint union of embedded cylindrical submanifolds in $\mathbb{R} \times (S^1 \times S^2)$ that are in 1–1 correspondence with the elements in \hat{A} . Such a cylinder is called an ‘end’ of C .

Requirement 2: Let $E \subset C$ denote any given end and let $(\delta, \varepsilon, (p, p'))$ denote its label from \widehat{A} . Then the following conditions are satisfied:

- (a) The function s restricts to E as a smooth function without critical points with εs bounded from below on the closure of E . Moreover, the 1-form α in (1-1) has nowhere zero pull-back on each constant $|s|$ slice of E .
- (b) The restriction to E of the function θ has a unique $|s| \rightarrow \infty$ limit; and the latter equals 0 when $\delta = 1$, it equals π when δ is -1 , and it is given by P via (1-7) when $\delta = 0$. Moreover, when $\delta = \pm 1$, this convergence makes the $\kappa \equiv |\delta \sqrt{\frac{3}{2}} + \frac{p'}{p}|$ version of $e^{\kappa|s|} \sin \theta$ converge to a unique limit as $|s| \rightarrow \infty$.
- (c) The integers p and p' are the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about any given constant $|s|$ slice of E when the latter is oriented by the 1-form α .
- (d) Any given anti-derivative on E for the restriction of the 1-form $p' dt - p d\varphi$ has a unique $|s| \rightarrow \infty$ limit.
- (e) Let $N_E \rightarrow E$ denote the normal bundle to E . Define $\prod_J: TE \rightarrow N_E$ to be the composition of the map $J: TE \rightarrow T(\mathbb{R} \times (S^1 \times S^2))|_E$ with the projection to N_E . Define the norm of \prod_J , the covariant derivative of \prod_J , and the latter's norm using the metrics and connections on TE and N_E that are induced by the metric on $\mathbb{R} \times (S^1 \times S^2)$. Then the norm of \prod_J and that of its covariant derivative limit to zero as $|s| \rightarrow \infty$ on E .

As is explained subsequently, a collection of pairs $\{(a_e, w_e)\}_{e \in T}$ that meet the criteria in (3-3) will serve as a starting point for the deformation to a pseudoholomorphic subvariety. The remainder of this section assumes that such a collection has been specified.

3.C A preliminary deformation

Let C denote the closure of $\cup_e K_e$ as defined using the given collection of pairs, $\{(a_e, w_e)\}$. This subsection begins the construction a family of deformations of C whose end member is a subvariety in $\mathcal{M}_{\widehat{A}}$. In particular, a preliminary deformation of C is constructed here so that the result is pseudoholomorphic with respect to an admissible almost complex structure.

To start the task, return to the observation that $d\alpha$ pulls back as a non-zero 2-form to any given version of K_e . This, being the case, it follows that $d\alpha$ is non-negative on TC . With the third point in (3-3), the last observation has the following consequence: There exists $r > 0$ such that the symplectic form $d(e^{-rs}\alpha)$ is uniformly positive on TC . This is to say that its pull-back to the multi-punctured sphere is a multiple of the

induced area form that is positive and uniformly bounded away from zero. Indeed, by virtue of the third point, any positive r version of this form is positive on TC near the images of the critical points of the pull-back of θ . Meanwhile, the form is positive for small r on any given compact subset in the complement of these same critical points. At large $|s|$ on any given end of C , both $-ds \wedge \alpha$ and $d\alpha$ are positive.

Now, specify a positive real number, ε . Granted the second and fourth points of (3–3), there are standard constructions that provide $\mathbb{R} \times (S^1 \times S^2)$ with an almost complex structure, J_0 , with the following properties:

- (3–4)
- The subvariety C is J_0 -pseudoholomorphic.
 - $J_0 = J$ where the distance to C is greater than ε .
 - $J_0 \partial_s = \frac{1}{(1+3 \cos^4 \theta)^{1/2}} \hat{\alpha}$ where the distance to any singular point of C or critical point of the restriction of θ is greater than ε .
 - Both $J_0 - J$ and its covariant derivative converge uniformly with limit zero as $|s| \rightarrow \infty$ on $\mathbb{R} \times (S^1 \times S^2)$.
 - All sufficiently small but positive r versions of $d(e^{-rs} \alpha)$ tame J_0 .

Having specified J_0 , fix a Riemannian metric on $\mathbb{R} \times (S^1 \times S^2)$ to be called g_0 , one with the following three properties: First J_0 acts as a g_0 -isometry. Second, g_0 converges uniformly as $|s| \rightarrow \infty$ to the metric $ds^2 + dt^2 + d\theta^2 + \sin^2 \theta d\varphi^2$, and its covariant derivative (as defined by the latter metric) converges uniformly to zero as $|s| \rightarrow \infty$. Finally, g_0 agrees with the latter metric where $J_0 = J$.

The almost complex structure J_0 would be admissible in the sense given prior to Definition 2.1 were the third point to hold on the whole of $\mathbb{R} \times (S^1 \times S^2)$, and were J_0 and J to agree on the nose outside of some compact subset of $\mathbb{R} \times (S^1 \times S^2)$. This part of the subsection describes how to move C slightly so that the result is an immersed subvariety that is pseudoholomorphic for an admissible almost complex structure. The construction of such a deformation is done in three steps. A fourth step explains how this can be done so that the version of Subsection 2.G's graph $T_{(c)}$ for the resulting subvariety gives the starting moduli space graph T .

Step 1 This first step specifies J_0 in a more precise manner near the images of the critical points of the pullback of θ . To start, note that if $z \in C_0$ is a critical point of the pullback of θ , then there is a small radius, embedded disk, $D \subset \mathbb{R} \times (S^1 \times S^2)$, that is contained in C and centered at the image of z . By virtue of the third point in (3–3), the vectors ∂_s and $\hat{\alpha}$ span TD at the image of z . Thus, J_0 must map one to a multiple of the other at this particular point. This being the case, there is a constant,

c , with the following significance: For all sufficiently small yet positive ε , an almost complex structure J_0 can be found that obeys (3–4) and is such that

$$(3-5) \quad |J - J_0| < C\varepsilon \quad \text{and} \quad |\nabla(J - J_0)| < C$$

at points with distance ε or less from the image in $\mathbb{R} \times (S^1 \times S^2)$ of any critical point of the pullback of θ .

Step 2 This step modifies both C and J_0 near the singular points of C so that (3–5) is also obeyed at each of the latter points. To explain how this is done, let $D \subset C$ for the moment denote an embedded disk whose closure is disjoint from the image of any critical point of the pullback of θ . Let z denote the center point of D . Fix some J -pseudoholomorphic disk, $D' \subset \mathbb{R} \times (S^1 \times S^2)$, with center z whose tangent space at z is spanned by ∂_θ and $J \cdot \partial_\theta$. Having chosen such a disk, there exists $\rho > 0$ and complex coordinates (x, y) centered at z , defined for $|x| < \rho$ and $|y| < \rho$, such that D' is the $y = 0$ disk, and such that ∂_x and \hat{a} are tangent to each $x = \text{constant}$ disk. Thus θ is constant on each of the latter. In these coordinates, the disk D can be viewed as the image of a neighborhood of the origin in \mathbb{C} to \mathbb{C}^2 that maps the complex coordinate u on \mathbb{C} as

$$(3-6) \quad u \rightarrow (x = u, y = au + b\bar{u} + o(|u|^2)),$$

where a and b are complex numbers. Note that the x -coordinate of the map can be defined in this way by virtue of the fact that θ is a function only of x .

Consider now deforming D in a manner that will now be described. To start, pick some small, positive δ with the property that the disk of radius 4δ in \mathbb{C} is mapped via (3–6) some distance from the boundary of D . Let β denote a favorite, smooth function on $[0, \infty)$ that is identically 1 on $[0, 1]$, vanishes on $[2, \infty)$ and is non-increasing. With β chosen, consider the deformed disk, $D(\delta)$, that is defined by the image of the map

$$(3-7) \quad u \rightarrow (x = u, y = au + \beta\left(\frac{1}{\delta}|u|\right)b\bar{u} + o(|u|^2)).$$

The image of this new map agrees with the old where $|u| > 2\delta$. The new subvariety will be immersed and pseudoholomorphic for an almost complex structure that also obeys the constraints on J_0 in (3–4). However, such an almost complex structure exists that agrees with J near the point z .

Now suppose that z is a singular point of C . Thus, there are two versions of D with center at z , these now denoted by D_1 and D_2 . No generality is lost by assuming here the respective closures of D_1 and D_2 are disjoint save for the shared point z . Each such disk is described by a map as in (3–6) using respective (a_1, b_1) and (a_2, b_2) versions of the pair (a, b) of complex numbers. For sufficiently small δ , each of D_1

and D_2 has their corresponding deformation as given in (3–7). The claim here is that the resulting disks, D_{1*} and D_{2*} , intersect only at z , transversely, and with positive intersection number. To explain, remark that any point in $D_{1*} \cap D_{2*}$ is the image of a point $u \in \mathbb{C}$ where

$$(3-8) \quad (a_1 - a_2)u = (b_1 - b_2)\beta\bar{u} + o(|u|^2).$$

Since $0 \leq \beta \leq 1$, this can happen at non-zero u when δ is small only if $|b_1 - b_2| \geq |a_1 - a_2|$. However, the latter inequality is forbidden by the fact that the D_1 and D_2 have transversal intersection at z with positive intersection number.

Thus, the new disks, D_{1*} and D_{2*} , intersect transversely only at z with positive intersection number. Moreover, each is J -pseudoholomorphic at z . This understood, if such a deformation is made for each singular point of C , then the result is the image (henceforth named C) of C_0 via an immersion that is pseudoholomorphic for a new version of the almost complex structure J_0 , one that obeys (3–4) and also obeys (3–5) at each of the singular points of the immersion and at the image of each of the critical points of the pull-back of θ .

Step 3 At this point, the stage is set to deform the newest version of C so that the result is pseudoholomorphic for an admissible almost complex structure.

To begin describing the latter deformation, let $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ now denote the tautological immersion with image C . Introduce the bundle $N \rightarrow C_0$ to denote the pull-back normal bundle; this defined so that the fiber over any given $z \in C_0$ is the g_0 -normal bundle at $\phi(z)$ to the ϕ image of any given sufficiently small radius disk in C_0 with z as center. As in the case for J -pseudoholomorphic subvarieties, the almost complex structure J_0 and the metric g_0 together endow N with the structure of a complex line bundle with a Hermitian and thus holomorphic structure. In addition, there exists $\delta > 0$, a disk subbundle, $N_1 \subset N$ of g_0 -radius δ , and an ‘exponential’ map $e: N_1 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ with the following properties: First, e is an immersion that restricts to the zero section as the map ϕ . Second, e is a g_0 -isometry along the zero section. Third, the differential of e is uniformly bounded. Finally, e embeds each fiber disk in N_1 as a J_0 -pseudoholomorphic disk.

With e chosen, there exists some $\delta_0 \in (0, \delta)$ with the following significance: If η is a section of N_1 with suitable decay at large $|s|$, and if both $|\eta|$ and $|\nabla\eta|$ are both everywhere less than δ_0 , then $\phi' \equiv e \circ \eta$ will immerse C_0 as a pseudoholomorphic subvariety for a complex structure, J'_0 , that also obeys the constraints in (3–5). Moreover, if the norms of $|\eta|$ and $|\nabla\eta|$ are small, then the critical values and critical points of the pullback of θ via the new immersion will hardly differ from those of the original. This is an important point in subsequent arguments, so keep it in mind.

In any event, the plan is to find such a section η with a corresponding J'_0 that is admissible. For this purpose, remark that there exists a constant, c , and, given some very small, but positive constant ε , there exists an admissible complex structure, J' , that has the following properties:

- (3–9) • $J' = J_0$ except where $|s| > \frac{1}{\varepsilon}$ and where the distance to any singular point of C or image of a critical point of the pullback of θ is less than ε .
- $|J_0 - J'| < c|J - J_0|$ and $|\nabla(J - J')| \leq c(|\nabla(J - J_0)| + |J - J_0|)$ where $|s| > \frac{1}{\varepsilon}$.
- $|J_0 - J'| \leq c \cdot \varepsilon$ and $|\nabla(J_0 - J')| \leq c$ where the distance is less than ε to any singular point of c or to the image of any critical point θ 's pullback.

With ε now chosen very small (an upper bound appears below), fix an admissible J' that obeys the constraints in (3–9). This done, the plan here is to search for an immersion, $\phi': C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$, whose image is a J' -pseudoholomorphic subvariety. Thus, $\prod J' \cdot d\phi' = 0$, where \prod is the projection to the normal bundle of the immersion. The sought for deformation of C is obtained by composing e with a suitable section of the bundle N_1 . In particular, if η is a section of N_1 , then the condition on η can be written schematically as

$$(3-10) \quad D_C \eta + \mathcal{R}_0(\eta) + \mathcal{R}_1(\eta) \cdot \partial \eta + \gamma(\eta) + \hat{\tau} = 0,$$

where the notation is as follows: First, D_C is the (C, J_0) version of the operator that is depicted in (2–5) while \mathcal{R}_0 and \mathcal{R}_1 are the (C, J_0) versions of their namesakes from (2–3) and (2–5). Meanwhile, γ is a smooth, fiber preserving map from N_1 to $N_1 \otimes T^{0,1}C_0$ that obeys

$$(3-11) \quad |\gamma(\eta)| \leq c'|J' - J_0|(1 + |\nabla \eta|) + |\nabla(J' - J_0)||\eta|,$$

where c' is a constant that can be taken to be independent of the choice of ε and J' . Note that \mathcal{R}_0 , \mathcal{R}_1 and γ are defined on some small, positive and constant radius disk subbundle of N , and the latter can be taken equal to N_1 with no loss of generality. Finally, $\hat{\tau}$ is a linear map from a certain finite dimensional vector subspace of $C^\infty(N \otimes T^{0,1}C)$ back into the latter space whose image has compact support where the distance to any singular point of C or image of a critical point of θ 's pull-back is large. The form of $\hat{\tau}$ is described momentarily.

The operator D_C in (3–10) has the same sort of Fredholm extension as a bounded linear operator between the C versions of the range and domain spaces that appear in (2–7). Note that the index of this Fredholm version of D_C is the integer \hat{I} in (2–2). In the present context, it may well be the case that D_C has a non-trivial cokernel.

If $\text{cokernel}(D_C) = 0$, then $\hat{\tau}$ in (3–10) can be discarded. If $\text{cokernel}(D_C)$ has positive dimension, then $\hat{\tau}$ is necessary. To elaborate, $\hat{\tau}$ can be any linear map from $\text{cokernel}(D_C)$ into $C^\infty(N \otimes T^{0,1}C_0)$ with the following properties: First, the support of the image of $\hat{\tau}$ is compact and with all points in the support mapped to points with distance at least one from singular point of C or image of a critical point of θ 's pullback. Second, the orthogonal projection in D_C 's range Hilbert space composes with $\hat{\tau}$ to give the identity map on $\text{cokernel}(D_C)$.

With the preceding understood, let H_R denote the orthogonal complement in the range space of D_C to the D_C 's cokernel and let \prod denote the orthogonal projection in this range Hilbert space onto H_R . Meanwhile, use H_D to denote the orthogonal complement in the domain space of D_C to its kernel. Note that D_C restricts to H_D to define a bounded, invertible map onto H_R . The inverse of the latter map is denoted below as $(D_C)^{-1}$.

To continue, let $H' \subset H_D$ denote the subset of smooth elements with pointwise norm no greater than half the radius of the disk bundle N_1 . Finally, define the smooth map $Y: H' \times \text{cokernel}(D_C) \rightarrow H_D$ by the rule

$$(3-12) \quad Y(\eta, \lambda) = -(D_C)^{-1} \prod [\mathcal{R}_0(\eta) + \mathcal{R}_1(\eta) \cdot \partial\eta + \gamma(\eta) + \hat{\tau}(\lambda)].$$

By design, if $\eta = Y(\eta, \lambda)$, then η solves (3–10) provided that

$$(3-13) \quad \lambda = -\left(1 - \prod\right) [\mathcal{R}_0(\eta) + \mathcal{R}_1(\eta) \cdot \partial\eta + \gamma(\eta)].$$

The existence of such a pair (η, λ) is guaranteed when ε is very small. Indeed, the analysis used in [18, Section 3c and the proof of its Proposition 3.2] can be used here to construct a version of the contraction mapping theorem to prove the following:

Lemma 3.3 *Given small $\varepsilon' > 0$, then all sufficiently small ε versions of the fixed point equation $\eta = Y(\eta, \lambda)$ have a unique solution with λ given by (3–13) and with ε' bounding both the Hilbert space norm and pointwise C^1 -norm of η .*

Step 4 This last step explains how to make the preceding construction result in a subvariety whose version of Subsection 2.G's graph $T_{(\cdot)}$ is the given moduli space graph T . To start, take note that given T , there exists a positive constant, δ_T , with the following significance: Suppose that T' is a labeled graph, isomorphic to T save for the fact that its trivalent vertex angle assignments differ. Even so, suppose that there is a 'quasi' isomorphism that identifies the underlying graphs so as to pair like labeled monovalent and bivalent vertices, match edge labels and pair trivalent vertices only if their respective angle assignments differ by less than δ_T . This graph T' is also a

moduli space graph for \hat{A} . If $\delta \in (0, \delta_T)$, say that a graph T' is “ δ close to T ” when such a quasi-isomorphism pairs trivalent angles so that all such pairs differ by less than δ .

If δ is small, then graphs that are δ close to T can be parametrized by a cube of side length 2δ in the product of $(N_- + \hat{N} + c_{\hat{A}} - 2)$ copies of $(0, \pi)$; this the cube centered on the angle assignments for the trivalent vertices of T . Let B_δ denote this cube.

There are now three remarks to make: First, with both ε from (3–9) and ε' from Lemma 3.3 taken to be very small, the constructions just given in Step 1–Step 3 can be made for any graph that comes from a point in B_δ . In this way, each point in B_δ produces a subvariety, and thus a version of Subsection 2.G’s graph $T_{(\cdot)}$. Second, with $\varepsilon_1 > 0$ fixed and then both ε from (3–9) and ε' from Lemma 3.3 even smaller, any chosen point from $B_{\delta/2}$ provides a version of $T_{(\cdot)}$ that is isomorphic to T via an isomorphism with the following additional property: It pairs trivalent vertices so that the resulting angle assignments define a point in B_δ with distance ε_1 or less from the initially chosen point. Finally, the constructions in Step 1–Step 3 can be made so that the result of all this is a continuous map from $B_{\delta/2}$ to B_δ .

Now take ε_1 very small. Granted the preceding three observations, some starting point in $B_{\delta/2}$ gives $T_{(\cdot)} = T$ for the simple reason that $B_{\delta/2}$ doesn’t retract onto its boundary.

3.D The deformation to a J –pseudoholomorphic subvariety

Let J' be an admissible almost complex structure, let ϑ denote an unordered set of N_+ points in S^1 , and let $\mathcal{M}_{\hat{A}}[\Theta, \vartheta]$ denote the subset in the J' –version of $\mathcal{M}_{\hat{A}}[\Theta]$ that consists of subvarieties whose inverse images in \mathcal{M}^Λ are sent to the points in ϑ by (2–21)’s map ϖ_+ . The respective sets Θ and ϑ are deemed ‘generic’ when the following conditions apply:

- (3–14) • The set θ has $N_- + \hat{N} + c_{\hat{A}} - 2$ elements, and these elements are pairwise disjoint and none arises via (1–7) from an integer pair of any $(0, \dots)$ element in \hat{A} .
- The set ϑ contains N_+ distinct elements.

According to Proposition 2.12, any generic θ and ϑ version of $\mathcal{M}[\Theta, \vartheta]$ is a submanifold of $\mathcal{M}_{\hat{A}}$.

Now, if T is given by the graph from Subsection 2.G of a subvariety from the J' version of $\mathcal{M}_{\hat{A}}$, then the subvariety is in the version of $\mathcal{M}[\Theta]$ where θ ’s angles are

those assigned to the trivalent vertices in T . Thus, θ is generic if T is generic. More to the point, Proposition 2.12 finds a subvariety in a generic ϑ version of $\mathcal{M}[\Theta, \vartheta]$ whose graph is also isomorphic to T .

With the preceding understood, and granted what has been said in the previous subsections, there exists an admissible almost complex structure J' , a generic pair (Θ, ϑ) and a subvariety C in the J' version of $\mathcal{M}[\Theta, \vartheta]$ whose graph from Subsection 2.G is isomorphic to the graph T .

The remainder of this subsection explains how C is used to construct a J -pseudoholomorphic subvariety in the J -version of $\mathcal{M}[\Theta, \vartheta]$ whose version of $T_{(\cdot)}$ is isomorphic to the given moduli space graph T . The description of such a deformation is broken into six steps.

Step 1 This step explains the strategy for obtaining the desired subvariety. To begin, choose a continuously parametrized family, $\{J^a\}_{a \in [0,1]}$, in the space of admissible almost complex structures whose initial element, J^0 , is J' , and whose final element, J^1 , is J . Having made such a choice, an attempt is made to construct a corresponding family, $\{C^a\}_{a \in [0,1]}$, of subvarieties in $\mathbb{R} \times (S^1 \times S^2)$ that has $C^0 = C$ and is such that any given C^a is a J^a -pseudoholomorphic subvariety in the J^a version of $\mathcal{M}_{\hat{A}}$. In particular, the goal is to construct such a family where each $a \geq 0$ version of C^a is in the submanifold $\mathcal{M}[\Theta, \vartheta]$ from the J^a version of $\mathcal{M}_{\hat{A}}$ and whose corresponding version of $T_{(\cdot)}$ is the given moduli space graph T .

To proceed, introduce the set, \mathfrak{f} , of points $r \in [0, 1]$ for which C^a exists for every value of $a \leq r$. This \mathfrak{f} is non-empty since it contains 0. The next step explains why \mathfrak{f} is open. Modulo a technical proposition, an argument is given in the third step that proves the following: Either \mathfrak{f} is closed, or else the submanifold $\mathcal{M}[\Theta, \vartheta]$ in the J -version of $\mathcal{M}_{\hat{A}}$ contains the desired subvariety. If \mathfrak{f} is closed, then $1 \in \mathfrak{f}$ and the submanifold $\mathcal{M}[\Theta, \vartheta]$ in the J -version of $\mathcal{M}_{\hat{A}}$ contains the desired subvariety. Thus, the desired conclusion follows in either case.

A bit more work will establish that \mathfrak{f} is, in fact, closed. Moreover, the resulting parametrized family $\{C^a\}_{a \in [0,1]}$, can be constructed so that the parametrization varies continuously with the parametrization, or smoothly in the case that the parametrization $a \rightarrow J^a$ is smooth. However, this extra work is left to the reader.

(By the way, the terms ‘continuous’ and ‘smooth’ for the parametrization that sends $a \rightarrow C^a$ are defined as follows: The parametrization is continuous if there exists a multi-punctured sphere, C_0 , with a continuous map $\Phi: C_0 \times [0, 1] \rightarrow \mathbb{R} \times (S^1 \times S^2)$ such that each $\Phi(\cdot, a)$ is a smooth, proper immersion with image C^a that is 1–1 to its

image on the complement of a finite set. The parametrization is smooth when there is such a map ϕ that is smooth.)

Step 2 The proof that f is open makes fundamental use of the generalization of Proposition 2.6 that follows. The proof of this proposition, like that of Proposition 2.6, is much like that of [18, Proposition 3.2] and thus is omitted.

Proposition 3.4 *Let J' be an admissible almost complex structure, let \hat{A} be any given asymptotic data set, and let C' be a subvariety in the J' version of $\mathcal{M}_{\hat{A}}$. Let C_0 denote the model curve for C' and let $\phi: \mathbb{R} \times (S^1 \times S^2)$ denote its attending J' -pseudoholomorphic map. Then, there exists a constant $\kappa \geq 1$, a ball $B \subset \text{kernel}(D_{C'})$, an open neighborhood, \mathcal{U} , of J' in the space of admissible almost complex structures, and a smooth map F , from $\mathcal{U} \times B$ to $C^\infty(\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)))$ with the following properties:*

- $|F(J', \eta) - \eta| + |\nabla(F(J', \eta) - \eta)| \leq \kappa \|\eta\|^2$.
- The exponential map on the C' version of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ composes with F to give a smooth map, $\Phi: \mathcal{U} \times B \times C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$.
- With $(J'', \eta) \in \mathcal{U} \times \text{kernel}(D_{C'})$ fixed, then $\Phi(J'', \eta, \cdot)$ maps C_0 onto a J'' -pseudoholomorphic subvariety.
- As η varies in B with J'' fixed, the resulting family of subvarieties defines an embedding, $\psi_{J''}$, from B onto an open set in the J'' version of $\mathcal{M}_{\hat{A}}$. In particular, if C'' is in the J'' version of $\mathcal{M}_{\hat{A}}$ and if

$$(3-15) \quad \sup_{z \in C'} \text{dist}(z, C'') + \sup_{z \in C''} \text{dist}(C', z) < \frac{1}{\kappa},$$

then C'' is in the image of $\psi_{J''}$.

With Proposition 3.4 in hand, what follows explains why f is open. To begin, suppose that $\tau \in \mathfrak{f}$ and that $\tau < 1$. Let $e: \phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) \rightarrow \mathbb{R} \times (S^1 \times S^2)$ denote an exponential map of the sort that is described in Subsection 2.D. For values of a that are somewhat greater than τ , the subvariety C^a is the image of the composition of e with a suitably chosen, smooth section, η^a , of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. The section η^a has very small norm when $a \sim \tau$ and, in any event, is an element in the domain Hilbert space for the C^τ version of the operator D_C as described in Subsection 2.D. The properties of η^a are summarized by the next lemma. An immediate corollary is that f is open.

To prepare for the lemma, first note that it refers to the functions

$$\{v_j : 1 \leq j \leq c\}, \{\varpi_{+\alpha} : 1 \leq \alpha \leq N_+\} \text{ and } \varpi_{\lambda,r}$$

that appear in the C^τ version of Proposition 2.13. In this regard, be aware that the domain of definition of these functions extends in a straightforward manner to include any subvariety in $\mathbb{R} \times (S^1 \times S^2)$ that has the asymptotics of a subvariety from $\mathcal{M}_{\hat{A}}$ and is the image via the exponential map of a pointwise small section of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ with pointwise small covariant derivative. For example, these functions are defined for the $a < \tau$ versions of C^a when a is sufficiently close to τ .

Lemma 3.5 *Given $c \in \{0, \dots, N_- + \hat{N} + c_{\hat{A}} - 2\}$ and a c element subset $\Theta \subset (0, \pi)$, suppose that $\tau \in [0, 1)$ and that C^τ is any element of the J^τ version of $\mathcal{M}[\Theta]$. Then, there exists some $\delta > 0$ and a continuous map from $(-\delta, \delta)$ into the intersection of the domain of D_C with the C^τ version of $C^\infty(\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)))$ such that the image of any given $\tau' \in (-\delta, \delta)$ is a section, $\eta^{\tau+\tau'}$, whose composition with the exponential map sends C_0 onto a $J^{\tau+\tau'}$ -pseudoholomorphic subvariety from the $J^{\tau+\tau'}$ version of the space $\mathcal{M}[\Theta]$. Moreover, the following is also true: Suppose that the 1-parameter family $\{C^a\}$ is continuously defined along the interval $[0, \tau]$.*

- The map $\tau' \rightarrow \eta^{\tau+\tau'}$ can then be constructed as a continuous map with domain $(-\delta, \delta)$ such that each $a \in (\tau-\delta, \tau]$ version of C^a is the image of the composition of the exponential map with the corresponding η^a .
- If the family $\{J^a\}_{a \in [0, \tau+\delta]}$ is smoothly parametrized, and if the original family $\{C^a\}_{a \in [0, \tau]}$ is smoothly parametrized on $[0, \tau]$, then the map $\tau' \rightarrow \eta^{\tau+\tau'}$ can be constructed to be smooth on the whole of $(-\delta, \delta)$.
- If a given subset of the functions $\{v_j : 1 \leq j \leq c\}, \{\varpi_{+\alpha} : 1 \leq \alpha \leq N_+\}$ and $\varpi_{\lambda,r}$ are constant on C^a for values of a near to but less than τ . Then the map $\tau' \rightarrow \eta^{\tau+\tau'}$ can be constructed so that the same subset of these functions are constant on the $a > \tau$ subvarieties as well.

In any event, the given 1-parameter family, $\{C^a\}_{a \in [0, \tau]}$, has an extension that is parametrized by the points in the interval $[0, \tau + \delta]$.

Proof of Lemma 3.5 Granted Propositions 2.7, 2.9, 2.13 and 3.4, all of the assertions are proved by using various straightforward applications of the implicit function theorem. \square

Step 3 This step explains why the submanifold $\mathcal{M}[\Theta, \vartheta]$ in the J -version of $\mathcal{M}_{\hat{A}}$ contains the desired subvariety when \mathfrak{f} is not closed. To begin, suppose that $\tau \leq 1$ and that the family $a \rightarrow C^a$ has been defined for $a < \tau$ with each C^a in the appropriate version of the submanifold $\mathcal{M}[\Theta, \vartheta]$. The issue here is whether the domain of definition

for the 1-parameter family extends to the parameter value $a = \tau$ as well. The focus here is thus on the convergence or lack thereof for the sequence $\{C^a\}_{a < \tau}$ as $a \rightarrow \tau$.

The next proposition asserts some facts about sequences of subvarieties of the sort that is under consideration. The following is a direct corollary: Either \mathfrak{f} is closed or else the subset $\mathcal{M}[\Theta, \vartheta]$ in the J -version of $\mathcal{M}_{\hat{A}}$ is non-empty and contains a subvariety whose corresponding graph from Subsection 2.G gives the graph T .

Proposition 3.6 *Assume here that θ and ϑ are generic in the sense of (3–14). Let J' be an admissible almost complex structure, and suppose that $\{(J_j, C_j)\}_{j=1, \dots}$ is a sequence of pairs of the following sort: First, $\{J_j\}$ is a sequence of admissible almost complex structures that converges to J' . Meanwhile, each C_j is in the subset $\mathcal{M}[\Theta, \vartheta]$ from the J' version of $\mathcal{M}_{\hat{A}}$ and each has graph $T_{(\cdot)}$ giving T . Then one of the following two assertions hold:*

- A** *There exists a subvariety C' in the subset $\mathcal{M}[\Theta, \vartheta]$ of the J' version of $\mathcal{M}_{\hat{A}}$ with graph $T_{C'} = T$, a subsequence of $\{C_j\}$ and a corresponding sequence, $\{\eta_j\}$, of sections of a fixed radius ball subbundle in the C' version of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ such that composition of the exponential map with any given η_j sends the model curve of C' onto C_j . Moreover, the sequence of supremum norms over C' of the elements in $\{\eta_j\}$ limits to zero as $j \rightarrow \infty$, as do the analogous sequences of norms of the higher derivatives.*
- B** *There exists a subvariety C in the subset $\mathcal{M}[\Theta, \vartheta]$ of the J version of $\mathcal{M}_{\hat{A}}$ with graph $T_C = T$ a subsequence of $\{C_j\}$ and a corresponding sequence, $\{\eta_j\}$, of sections of a fixed radius ball subbundle in the C version of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ such that composition of the exponential map with any given η_j sends the model curve of C onto a translate of C_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. As before, the sequence of supremum norms over C of the elements in $\{\eta_j\}$ limits to zero as $j \rightarrow \infty$, as do the analogous sequences of norms of the higher derivatives.*

The proof of Proposition 3.6 exploits convergence theorems that are modified versions of assertions from Hofer, Wysocki and Zehnder [8] about the behavior of limits of pseudoholomorphic curves. (See also Bourgeois, Eliashberg, Hofer, Wysocki and Zehnder [1], which appeared during the preparation of this article.) The next proposition summarizes the needed results. Note that it makes no assumptions about θ and ϑ or any given moduli space graph such as T .

Proposition 3.7 *Let $\{J_j\}$ denote a sequence of admissible almost complex structures with the following two properties: First, the derivatives of each such endomorphism*

to any fixed, non-negative order are bounded over the whole of $\mathbb{R} \times (S^1 \times S^2)$ by a j -independent constant. Second, there is an admissible almost complex structure, J' , such that the restriction of $\{J_j\}$ to any given compact set in $\mathbb{R} \times (S^1 \times S^2)$ converges in the C^∞ topology to the corresponding restriction of J' . Next, let $\{C_j\}$ denote a sequence where each C_j is in the corresponding J_j version of $\mathcal{M}_{\hat{\lambda}}$. Then, there exists a subsequence of $\{C_j\}$ (hence renumbered by consecutive integers starting from 1) and a finite set, Ξ , of pairs of the form (S, n) where n is a positive integer and S is an irreducible, J' -pseudoholomorphic multi-punctured sphere; and these have the following properties:

- $\lim_{j \rightarrow \infty} \int_{C_j} \varpi = \sum_{(S,n) \in \Xi} n \int_S \varpi$ for each compactly supported 2-form ϖ .
- If $K \subset \mathbb{R} \times (S^1 \times S^2)$ is compact, then the following limit exists and is zero:

$$(3-16) \quad \lim_{j \rightarrow \infty} \left(\sup_{z \in C_j \cap K} \text{dist}(z, \cup_{\Xi} S) + \sup_{z \in (\cup_{\Xi} S) \cap K} \text{dist}(z, C_j) \right)$$

The proof is given momentarily.

Step 4 A portion of the proof of the previous proposition, as well as subsequent arguments in this article and in the sequel, require the lifting of certain submanifolds of $\Sigma \equiv \cup_{(S,n) \in \Xi} S$ to large j versions of C_j . This step explains how these liftings are done.

To start, it is necessary to first pass to a subsequence of $\{C_j\}$ where the corresponding sequence of sets of critical points of θ and sequence of sets of singular points converge in $\mathbb{R} \times (S^1 \times S^2)$. Let $Y_j \subset C_j$ denote the set of critical points of θ and singular points of C_j , and let $Y_o \subset \sigma$ denote the limit of $\{Y_j\}$. Next, let $\Sigma_* \subset \sigma$ denote the union of the irreducible components that are not of the form $\mathbb{R} \times \gamma$ where $\gamma \subset S^1 \times S^2$ is a Reeb orbit. Now, let $Y_* \subset \Sigma_*$ denote the union of the set Y_o , the critical points of θ on the subvarieties that comprise Σ_* and the singular points of σ that lie in Σ_* . Note that the latter set may contain points that are not points of Y_o . In any event, Y_* is a finite set.

Suppose next that $K \subset \Sigma_* - Y_*$ is a given compact set. Such a set K has a tubular neighborhood, $U_K \subset \mathbb{R} \times (S^1 \times S^2)$ with projection $\pi: U_K \rightarrow K$ whose fibers are disks on which θ is constant and that are pseudoholomorphic for any admissible almost complex structure. Indeed, the fiber of the projection to any given $p \in K$ is a disk centered on p inside $\mathbb{R} \times \gamma^p$, where $\gamma^p \subset S^1 \times S^2$ is a small segment of the integral curve of the Reeb vector field through the image of p . In this regard, note that $\mathbb{R} \times \gamma^p$ intersects K transversely at p by virtue of the fact that p is not a critical point of θ on σ . With K fixed, then each sufficiently large j version of C_j will have proper

intersection with U_K and intersect each fiber precisely n times, with each a transversal intersection and local intersection number $+1$. Such is the case precisely because the large j version of C_j has no θ critical points in U_K . In any event, here is a fundamental consequence: The projection, $\pi: C_j \cap U_K \rightarrow K$ defines a smooth, proper, degree n covering map. In particular, any compact, embedded arc in $\Sigma_* - Y_*$ has lifts to C_j under the projection π .

The lifts just described can be extended as lifts of arcs in a somewhat larger set in Σ_* . To define this set, let $Y \subset Y_*$ denote the subset of points that are either singular points of σ , critical points of θ on σ , or limits of convergent sequences of the form $\{p_j\}$ where p_j is either a critical point of θ on C_j or a non-immersion singular point of C_j . Then a smooth arc in a compact subset of $\Sigma - Y$ has a well defined lift to all large j versions of C_j that extends the lift defined in the preceding paragraph. In the discussion that follows, such a lift is deemed a ‘ θ -preserving preimage’ in C_j .

Step 5 Here is the proof of Proposition 3.7:

Proof of Proposition 3.7

This result is essentially from Hofer [4] or Hofer–Wysocki–Zehnder [8] (see also the article by Bourgeois, Eliashberg, Hofer, Wysocki and Zehnder [1] that appeared during the writing of this article). Here is the basic idea: There is a bound on the area of the intersection of C_j with any given $[s - 1, s + 1] \times (S^1 \times S^2)$ subcylinder that is independent of both s and j . The existence of such a bound can be deduced from the fact that the integral of $d\alpha$ on C_j is finite. The existence of this local area bound is the key observation. The existence of the asserted limit data set Ξ is deduced from the latter using arguments that are very similar to those used for the convergence theorems about sequences of pseudoholomorphic curves on compact symplectic manifold. Given the local area bound, a somewhat different proof of the convergence assertion can be obtained using [16, Proposition 3.3], but without the observation that each subvariety from Ξ is a multipunctured sphere.

What follows explains why each subvariety from Ξ is a sphere with punctures. In this regard, it is enough to consider only the non- \mathbb{R} invariant subvarieties. To start this chore, take any given irreducible component S from Σ_* , and note that it is enough to prove that some non-zero multiple of any class in the first homology of the model curve for S is generated by loops on the ends of S , thus ‘end-homologous’. For this purpose, fix a generator of the first homology of the model curve for S , and let $\tau \subset S$ be the image of an embedded representative that is disjoint from $Y \cap S$. In addition, fix $R \gg 1$ so that the $|s| > R$ part of S lies out on the ends of S and is disjoint from

τ . In this regard, remark that given R , there exists $\varepsilon > 0$ such that the $|s| \leq R$ portion of the ends of S have pairwise disjoint, radius ε tubular neighborhoods.

Take the index i to be very large, and let $\tau_i \subset C_i$ denote a connected, θ -preserving preimage of τ . Then τ_i is homologous in a regular neighborhood of S to some non-zero multiple of τ . However, τ_i is also homologous in C_i to a union of curves on the ends of C_i . Intersect this homology with the $|s| \leq R$ part of C_i and then deform the latter back to S in a small radius, regular neighborhood. The result is a homology between a non-zero multiple of τ and a union of curves on the ends of S as well as a union of circles that are very close to points in $S \cap Y$. As the inverse image of the latter circles are null-homologous in the model curve, so the chosen generator of the model curve's first homology is end-homologous. \square

Step 6 With the proof of Proposition 3.7 now complete, remark that it may well be the case that each subvariety from Ξ is an \mathbb{R} -invariant cylinder, thus of the form $\mathbb{R} \times \gamma$ where $\gamma \subset S^1 \times S^2$ is an orbit of the Reeb vector field. Item B of Proposition 3.6 holds when all subvarieties from Ξ are \mathbb{R} -invariant cylinders, and Item A of Proposition 3.6 holds when such is not the case. To explain how this dichotomy comes about, suppose first that there exists some subvariety from Ξ that is not \mathbb{R} -invariant. In this case, the proof of Proposition 3.6 proceeds to establish that Ξ consists of a single pair, and that the latter has the form $(S, 1)$ with S in the submanifold $\mathcal{M}[\Theta, \vartheta]$ from the J' version of $\mathcal{M}_{\hat{\lambda}}$. This implies that the graph T_S from Subsection 2.G can be labeled as a moduli space graph. The fact that the T_S is isomorphic to T is seen as an automatic consequence of the strengthened versions of (3-16) that appear in the next three subsections.

In the case that all subvarieties from Ξ are \mathbb{R} -invariant, the argument for Item B of Proposition 3.6 proceeds as follows: The original sequence $\{C_i\}$ is now replaced by a new sequence, $\{C'_i\}$, where each C'_i is obtained by translating the corresponding C_i along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times S^2)$. It should be evident from the description given below that the resulting sequence of translations (as elements in \mathbb{R}) does not have convergent subsequence. In any event, each C'_i is pseudoholomorphic for the translated almost complex structure, this denoted by J'_i . The latter sequence has a subsequence that converges on compact sets in $\mathbb{R} \times (S^1 \times S^2)$ to the almost complex structure J . This understood, the sequence of translations is chosen so as to insure that $\{C'_i\}$ converges as described in Proposition 3.7, but with a limit data set of pairs that contains one whose subvariety component is not \mathbb{R} -invariant. The argument proceeds from here as in the previous case: It demonstrates that Proposition 3.7's limit data set of pairs is a single pair, this of the form $(S, 1)$, where S is in the submanifold $\mathcal{M}[\Theta, \vartheta]$ of the J -version of $\mathcal{M}_{\hat{\lambda}}$ whose corresponding graph is isomorphic to T .

The translation, $s_j \in \mathbb{R}$, that takes C_j to C'_j is defined as follows: Chose a fixed angle, $\theta_* \in (0, \pi)$, with the following properties: First, θ_* is strictly between the minimal and maximal angles that are defined by the elements of \hat{A} . Thus the $\theta = \theta_*$ locus in C_j is non-empty. Second, no pair (p, p') makes the $\theta = \theta_*$ version of (1–7) hold. Now, define s_j so that the translate $s \rightarrow s + s_j$ moves C_j so that the result, C'_j , has a point on its $\theta = \theta_*$ locus where $s = 0$.

Granted this definition, it then follows from the $\{C'_j\}$ version of (3–16) that the resulting limit data set, $\{(S, n)\}$, contains some S on which θ takes value θ_* . Such a subvariety can not be an \mathbb{R} -invariant cylinder.

3.E Convergence

With Proposition 3.7 and with what has been said so far, Proposition 3.6 follows directly as a corollary to

Proposition 3.8 *Assume that θ and ϑ are generic. Let $\{J_i\}$ denote a sequence of admissible, almost complex structures with the following three properties: First, the derivatives of each J_j to any fixed, non-negative order are bounded over $\mathbb{R} \times (S^1 \times S^2)$ by a j -independent constant. Second, there is a constant, L , such that $J_j = J$ on the complement of some length L subcylinder in $\mathbb{R} \times (S^1 \times S^2)$. Finally, there is an admissible, almost complex structure, J' , such that the restriction of $\{J_j\}$ to any compact set in $\mathbb{R} \times (S^1 \times S^2)$ converges in the C^∞ topology to the corresponding restriction of J' . Let $\{C_j\}$ denote a sequence where each C_j is in the submanifold $\mathcal{M}[\Theta, \vartheta]$ of the J_j -version of $\mathcal{M}_{\hat{A}}$, and where each has its graph from Subsection 2.G giving T . Now suppose that $\{C_j\}$ converges as described in Proposition 3.7 with limit data set Ξ . In this regard, assume that Ξ contains at least one subvariety that is not of the form $\mathbb{R} \times \gamma$ where γ is a Reeb orbit in $S^1 \times S^2$. Then Ξ consists of a single pair, this pair has the form $(S, 1)$, S is in the submanifold $\mathcal{M}[\Theta, \vartheta]$ from the J' -version of $\mathcal{M}_{\hat{A}}$, and the graph of S is isomorphic to T .*

The remainder of this subsection and the next two subsections are occupied with the proof of this proposition. In this regard, note that there are various ways to prove this proposition, in particular, some using mostly differential equation techniques such as can be found in Hofer [4; 5; 6] and Hofer–Wysocki–Zehnder [8; 7; 10] and the very recent Bourgeois–Eliashberg–Hofer–Wysocki–Zehnder [1]. The proof offered below relies almost entirely on arguments that are of a topological nature. In any event, the arguments used below are exploited in various modified forms in the sequel to this article.

The proof starts with a proof that the convergence assertion in (3–16) holds even in the case that $K = \mathbb{R} \times (S^1 \times S^2)$. This first part of the proof occupies the remainder of this subsection.

Part 1 of the Proof of Proposition 3.8

The proof that the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3–16) holds starts here by making the assumption that the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3–16) is false. The proof proceeds to derive a patently nonsensical conclusion.

To start this derivation, let $\Sigma \equiv \cup_{(S,n) \in \Xi} S$, let E denote any end from σ , and let $\mathbb{R} \times \gamma_*$ denote the translationally invariant cylinder with the same large $|s|$ asymptotics of as E . This is to say that the $|s| \rightarrow \infty$ limits of E converge in $S^1 \times S^2$ to the Reeb orbit γ_* . Now, fix some very small but positive number ε and there exists a value, s_0 , of s on E and a pair of sequences, $\{s_{j+}\} \subset [s_0, \infty)$ and $\{s_{j-}\} \in (-\infty, s_0]$ with the following properties:

- (3–17) • *If E is on the concave side of σ , then $\{s_{j+}\}$ has no convergent subsequences; and if E is on the convex side of σ , then $\{s_{j-}\}$ has no convergent subsequences.*
- *For each index j , the intersection of C_j with the $s \in [s_{j-}, s_{j+}]$ portion of the radius ε tubular neighborhood of $\mathbb{R} \times \gamma_*$ has an irreducible component, C_{j*} , where $|s|$ takes both the values s_{j-} and s_{j+} and whose points have distance $\frac{1}{4}\varepsilon$ or less from $\mathbb{R} \times \gamma_*$.*
 - *For each index j , there exists a subinterval, $I_j \subset [s_{j-}, s_{j+}]$ such that*
 - (a) *The sequence whose j' th element is the length of I_j diverges as $j \rightarrow \infty$.*
 - (b) *The sequence whose j' th element is the maximum distance from the $s \in I_j$ portion of C_{j*} to $\mathbb{R} \times \gamma_*$ limits to zero as $j \rightarrow \infty$.*

The fact that all of this can be arranged is a straightforward consequence of the manner of convergence that is dictated by Proposition 3.7.

Now comes a key point: Because each C_j is irreducible, when ε is small, there must exist an infinite sequence of positive integers j (hence renumbered consecutively from 1) and at least one end $E \subset \sigma$ with the following properties:

- (3–18) *Values for s_{j+} and s_{j-} can be chosen for use in E' 's version of (3–17) so that both the $s = s_{j+}$ and $s = s_{j-}$ loci in C_{j*} contain some point with distance $\frac{1}{4}\varepsilon$ from $\mathbb{R} \times \gamma_*$.*

With (3–18) understood, there are two cases to consider. □

Case 1 There is an end $E \subset \sigma$ where (3–18) holds and where the value, θ_* , of θ on γ_* is neither 0 nor π . The derivation of nonsense in this case is a four step affair.

Step 1 This step starts with a crucial lemma.

Lemma 3.9 *There exists $\delta > 0$ such that for all large j , the angle θ takes both the values $\theta_* + \delta$ and $\theta_* - \delta$ on the $s = s_{j-}$ locus in C_{j*} . Meanwhile, θ takes at least one of these values on the $s = s_{j+}$ locus in C_{j*} .*

The proof of this lemma is given below.

Granted Lemma 3.9, the ‘mountain pass’ lemma with the third point in (3–17) implies that there is a critical point of θ on each large j version of C_{j*} with critical value equal to θ_* .

As is explained next, the relatively prime integer pair (p, p') that is defined by θ_* via (1–7) must be proportional to either $Q_{e'}$ or $Q_{e''}$ where e', e'' with e label the three edges in T that are incident to the vertex that labels the critical point with critical value θ_* . Here, the convention for distinguishing e from e' and e'' is as follows: The respective vertices on e' and e'' that lie opposite that with angle label θ_* have angle labels on the same side of θ_* in $(0, \pi)$. This last conclusion exhibits the required nonsense since it requires the vanishing at $\theta = \theta_*$ of either the $Q = Q_{e'}$ or $Q = Q_{e''}$ version of α_Q .

To see why (p, p') are proportional to one of Q_e or $Q_{e'}$, let j be very large and let $K_e, K_{e'}$ and $K_{e''}$ denote the components in the complement of the C_j version of Γ in C_j ’s model curve. Slice C_{j*} into two pieces near the $s = \frac{1}{2}s_j$ locus. This then slices C_j into two parts, where one part, C_{j+} , contains the larger s portion of the sliced component of C_{j*} . Let C_{j-} denote the other part. By virtue of the fact that θ spreads uniformly to both sides of θ_* on C_{j-} , the latter must contain most of both K_e and one of $K_{e'}$ or $K_{e''}$. Agree to distinguish the latter as $K_{e'}$. Meanwhile, C_{j+} contains most of $K_{e''}$. In this regard, the portions that are missing in either of the three cases are portions where θ is everywhere very close θ_* .

Now, recall from Subsection 2.G that the $\theta = \theta_*$ part of the Γ -locus in C_j ’s model curve has the form of a ‘figure 8’, where one of the circles is the $\theta = \theta_*$ boundary of the closure of $K_{e'}$ and the other that of the closure of $K_{e''}$. This implies that any given constant θ circle in $K_{e''}$ is homologous to the union of a constant θ circle in $K_{e'}$ and a constant θ circle in K_e . Take these circles to have θ value that differ by order one from θ_* . This the case, the obvious ‘pair of pants’ in C_j ’s model curve with these three constant θ circles as boundary provides a homology. Now, this pair of pants is

sliced in two pieces by the $s \sim \frac{1}{2}s_j$ cut. In particular, the C_{j+} piece contains one of the boundary circles in $K_{e''}$ and the C_{j-} piece contains the other two boundary circles. This understood, it then follows from the definitions of $Q_{e''}$ and (p, p') as integrals of $\frac{1}{2\pi}dt$ and $\frac{1}{2\pi}d\varphi$ that these integer pairs are proportional.

Step 2 This step, Step 3 and Step 4 contain the following proof.

Proof of Lemma 3.9

Translate each C_j in $\mathbb{R} \times (S^1 \times S^2)$ by sending s to $s + s_{j-}$ in the \mathbb{R} factor, and let C_j' denote the corresponding subvariety. The sequence $\{C_j'\}$ converges in the manner dictated by Proposition 3.7 with some limit data set Ξ' . It follows from (3–18) that Ξ' contains an irreducible subvariety with a concave side end, E' , with the following property: Given some very large R , a value of s on E' , there is an infinite subsequence from $\{C_j\}$ (hence renumbered from 1) such that the $s \rightarrow s + s_{j-}$ translates of the $s_{j-} + R$ slices of C_{j*} converge pointwise to the $s = R$ slice of E' . Let γ' denote the Reeb orbit that is the limit of the the $|s| \rightarrow \infty$ slices of E' . Step 3 proves that $\gamma' = \gamma_*$.

Granted that $\gamma' = \gamma_*$, there are two possibilities: Either E' sits as a sub-cylinder in $\mathbb{R} \times \gamma_*$ or not. If not, then E' is a concave side end of some subvariety from Ξ' that is not \mathbb{R} -invariant. This understood, it follows from (2–17) that θ takes values both above and below θ_* on any constant $|s|$ slice of E' . (Remember that the $n = 0$ case of (2–17) is reserved solely for convex side ends.) Thus θ must take values on the $s = s_{j-}$ slice of each large j version of C_{j*} that differ from θ_* in both directions by some j -independent, non-zero amount.

Suppose, on the other hand that E' is contained in $\mathbb{R} \times \gamma_*$. As there are points on C_{j*} where $s \sim s_{j-}$ with distances at least $\frac{1}{4}\varepsilon$ from $\mathbb{R} \times \gamma_*$, the convergence described by Proposition 3.7 requires a subvariety from Ξ' that is not \mathbb{R} -invariant and contains a disk with the following property: The center point is on $\mathbb{R} \times \gamma_*$ and all other points are limit points of sequences whose j' th element is in the $s \rightarrow s + s_{j-}$ translate of C_{j*} . Since θ has no local maximum or minimum on such a disk, it thus follows that θ must take values on the $s = s_{j-}$ slice of each large j version of C_{j*} that differ from θ_* in both directions by some j -independent, non-zero amount.

To establish the asserted behavior of θ where s is near s_{j+} on C_j , translate each C_j in $\mathbb{R} \times (S^1 \times S^2)$ by sending $s \rightarrow s + s_{j+}$ in the \mathbb{R} -factor. Let C_j' now denote the result of this new translation. Invoke Proposition 3.7 once again to describe the convergence of this new version of $\{C_j'\}$, using Ξ' to denote the new limit data set. In this case, there is a convex side end, E' , with the following property: Given some very large R , a value of s on E' , there is an infinite subsequence from $\{C_j\}$ (hence renumbered

from 1) such that the $s \rightarrow s + s_{j+}$ translates of the $s_{j+} - R$ slices of C_{j*} converge pointwise to the $s = -R$ slice of E' . Let γ' denote the Reeb orbit that is the limit of the the $|s| \rightarrow \infty$ slices of E' . Step 3 and Step 4 prove that $\gamma' = \gamma_*$. Granted this, then the argument for the desired conclusion that θ takes some value on the $s = s_{j+}$ slice of C_{j*} that differs from θ_* by a j -independent amount is much the same as that given in the preceding paragraph. In fact, here is the only substantive difference: The argument now only finds values of θ on the $s \sim s_{j+}$ locus in C_{j*} that differ in at least one direction from θ_* by a non-zero, j -independent amount. This is because the integer n that appears in a convex side end version of (2-17) can vanish.

Step 3 But for one claim, this step proves that with ε small, the Reeb orbits γ' and γ_* agree in all of their Step 2 incarnations. To start, note first that with ε small, the Reeb orbit γ' sits in a tubular neighborhood of γ_* , and so γ' must be a translate of γ_* by some element in the group $T = S^1 \times S^1$ whose distance is $o(\varepsilon)$ from the identity.

To make further progress, let (p, p') again denote the relatively prime pair of integers that θ_* defines via (1-7). As can be readily verified, the 1-form $p d\varphi - p' dt$ is exact on a tubular neighborhood of γ_* , and this form pulls back as zero on any translate of γ_* by the group T . Moreover, on such a tubular neighborhood, the values of any chosen anti-derivative of this 1-form distinguish the T -translates of γ_* . In this regard, let f be an anti-derivative with value zero on γ_* . Then

$$(3-19) \quad \int_{\gamma} f(p dt + p' \sin^2 \theta_* d\varphi) = f|_{\gamma} 2\pi(p^2 + p'^2 \sin^2 \theta_*),$$

when γ is any translate of γ_* by an element from a small radius ball about the identity in T . Thus, the value of the integral on the left side here will distinguish γ' from γ_* if these two orbits are not one and the same.

With the preceding understood, take j large, and let s_{j0} and s_{j1} denote any two regular values of s on C_{j*} , chosen so that $J_j = J$ on the cylinder where $s \in [s_{j0}, s_{j1}]$. As is explained in Step 4, the respective integrals of $f(p dt + p' \sin^2 \theta_* d\varphi)$ over the $s = s_{j0}$ and $s = s_{j1}$ slices of C_{j*} agree by virtue of the fact that C_{j*} is pseudoholomorphic. The use of this last fact is simplest in the case that there exists some j -independent, positive number, R , such that $J_j = J$ where $s \in [s_{j-} + R, s_{j+} - R]$. If such is the case, then take any fixed $r > R$ such that each large j version of $s_{j-} + r$ is a regular value of s on C_{j*} . Granted this, take s_{j0} to be $s_{j-} + r$. For r large and then j very large, the convergence as described in Proposition 3.7 guarantees that the integral of $f(p dt + p' \sin^2 \theta_* d\varphi)$ over the $s = s_{j0}$ slice of C_{j*} is very close to the $\gamma = \gamma'$ version of the right hand side in (3-19). Meanwhile, take s_{j1} to be some generic value of s from the interval I_j from the third point of (3-17). As the latter slice of C_{j*} is

very close to $\mathbb{R} \times \gamma_*$, the integral of the form $f(pdt + p' \sin^2 \theta_* d\varphi)$ about such a slice will be very close to the $\gamma = \gamma_*$ version of the right hand side of (3–19). As r can be as large as desired, and likewise j , it thus follows that $\gamma' = \gamma_*$.

Of course, it may well be the case that there is no j -independent choice of R that excludes points with $J_j \neq J$ from where $s \in [s_{j-} + R, s_{j+} - R]$. There exists in this case, some fixed $R > 0$ such that $J_j = J$ save where $s \in [-R, R]$. If the lower bound for s on all large j versions of the interval I_j contain points where $s < -R$, then the argument given in the previous paragraph works just fine. The situation is different if the lower bound of s on all large j versions of I_j is greater than $-R$. Assuming that such is the case, first fix some large value of r so that for every sufficiently large j , each of $s_{-j} + r$, $-(R + r)$ and $R + r$ are regular values of s on C_{j*} . This done, first take $s_{0j} = s_{-j} + r$ and $s_{1j} = -(R + r)$ to establish that the integral of $f(pdt + p' \sin^2 \theta_* d\varphi)$ about the $s = -(R + r)$ slice of C_{j*} is very close to the $\gamma = \gamma'$ version of the right hand side of (3–19). Next, take s_{0j} to equal $R + r$ and take s_{1j} to be a generic value in I_j so as to establish that the integral of $f(pdt + p' \sin^2 \theta_* d\varphi)$ about the $s = (R + r)$ slice of C_{j*} is nearly the $\gamma = \gamma_*$ version of the right hand side of (3–19).

It remains now to establish that the integrals of $f(pdt + p' \sin^2 \theta_* d\varphi)$ along the respective $s = -(R + r)$ and $s = (R + r)$ slices of C_{j*} are, for very large j , very close to each other. To argue for this, note that the sequence $\{C_j\}$ converges according to Proposition 3.7 to a limiting J' -pseudoholomorphic subvariety. The sequence $\{C_{j*}\}$ thus has a subsequence that converges on every subcylinder of the form $[-R + r, R + r] \times (S^1 \times S^2)$ to a component, C'_* , of this subvariety. In this regard, all points of C'_* , must have distance less than $\frac{1}{4}\varepsilon$ from $\mathbb{R} \times \gamma_*$. Furthermore, the constant $|s|$ slices of the convex side ends of C'_* must converge as $|s| \rightarrow \infty$ to γ_* . Meanwhile, those of its concave side ends must converge as $|s| \rightarrow \infty$ to γ' . But, this then implies that $\gamma' = \gamma_*$ since the angle θ is the same on γ' as on γ_* and so has to be constant on C'_* .

Step 4 To tie up the final loose end, suppose that γ_* is a Reeb orbit where the value, θ_* , of θ is neither 0 nor π . Let (p, p') denote the relatively prime pair of integers that θ_* defines via (1–7). Let $U \subset S^1 \times S^2$ denote a tubular neighborhood of γ_* and let f denote an anti-derivative of $pd\varphi - p'dt$ on U . Now, suppose that $s_+ > s_-$, and that C_* is closed, J -pseudoholomorphic subvariety inside $[s_- - 1, s_+ + 1] \times U$. Thus, C_* has compact intersection with the subcylinder $[s_-, s_+] \times U$. Suppose that both s_- and s_+ are regular values of the restriction of s to C_* . Proved here is the assertion that the respective integrals of the 1-form $f(pdt + p' \sin^2 \theta_* d\varphi)$ over the $s = s_-$ and $s = s_+$ slices of C_* agree.

For this purpose, use Stokes' theorem to write the difference of the two integrals as the integral over C_* 's intersection with $[s_-, s_+] \times U$ of the form $df \wedge (pdt + p' \sin^2 \theta_* d\varphi)$. Written out, the latter is $-(p^2 + p'^2 \sin^2 \theta_*) dt \wedge d\varphi$. Now, as C_* is J -pseudoholomorphic, the latter integral is identical to that of the 2-form $-(p^2 + p'^2 \sin^2 \theta_*) ds \wedge \frac{1}{\sin \theta} d\theta$. A second application of Stokes' theorem establishes the assertion. \square

Case 2 This case assumes that the $|s| \rightarrow \infty$ limits of θ is 0 or π on every end in σ where (3–18) holds. Note that the arguments below consider only the case where the aforementioned limit of θ is 0. The argument for the case where the limit is π is identical but for some minor notational modifications.

To start the derivation of nonsense in this case, remark that the Reeb orbit γ_* in (3–17) is the $\theta = 0$ locus. Let θ_{j-} denote the maximum value of θ on the $s = s_{j-}$ slice of C_{j*} , and let θ_{j+} denote the maximum of θ on the $s = s_{j+}$ slice. Since both of these slices have points with distance $\frac{1}{4}\varepsilon$ from the $\theta = 0$ locus, it follows that for fixed, small ε and large j , both θ_{j-} and θ_{j+} are greater than $\frac{1}{100}\varepsilon$. This understood, let θ_{j*} denote the minimum value of the function on the interval $[s_{j-}, s_{j+}]$ that assigns to any given $a \in [s_{j-}, s_{j+}]$ the maximum value of θ on the $s = a$ slice of C_{j*} . In this regard, $\theta_{j*} > 0$, but by virtue of the third point in (3–17), $\lim_{j \rightarrow \infty} \theta_{j*}$ is zero.

Granted the preceding, the mountain pass lemma dictates that all large j versions of C_{j*} have a critical point of θ where $\theta = \theta_{j*}$. Of course, the latter conclusion is nonsense as all critical values of θ that are neither 0 nor π lie in the fixed, j -independent set θ .

3.F Part 2 of the proof of Proposition 3.8

This second part of the proof establishes that the set Ξ contains just one element. The argument here is presented in four steps.

Step 1 The first point to make is that the set of $|s| \rightarrow \infty$ limits of θ on σ are identical to the set of such limits on any given C_j . Of course, this follows from the fact that the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3–16) is valid here. In the case that no angle from any $(0, -, \dots)$ element in \hat{A} is the same as that from a $(0, +, \dots)$ element, this last conclusion rules out an \mathbb{R} -invariant component of σ where θ differs from either 0 or π .

To rule out such components in any case, suppose for the sake of argument that one were present. Denote the latter as $\mathbb{R} \times \gamma_*$ where γ_* is a Reeb orbit, and use θ_* to denote the value of θ on γ_* . Keep in mind that θ_* comes both from a $(0, +, \dots)$ and

a $(0, -, \dots)$ element in \hat{A} . Let $\varepsilon > 0$ be very small and let $U \subset S^1 \times S^2$ denote the radius ε tubular neighborhood of γ_* . The argument that follows proves that θ on all large j versions of C_j has a critical value that differs by at most a uniform multiple of ε from θ_* . Of course, this conclusion is nonsense for small ε since Proposition 3.8 assumes that no angle from θ coincides with an angle from a $(0, \dots)$ element in \hat{A} .

To start the argument, note that Proposition 3.7 dictates that there is an irreducible component of each sufficiently large j version of $C_j \cap (\mathbb{R} \times U)$ with one rather specific property. To describe this property, let C_{j*} denote the component in question. Here is the property: This C_{j*} contains two ends of C_j where the $|s| \rightarrow \infty$ limit of θ is θ_* , one on the convex side end and the other on the concave side.

To see where these observations lead, remark that by virtue of the assumptions in Proposition 3.8, the integer $\deg_E(d\theta)$ from (2–16) is zero on each convex side end of each C_j where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . This is to say that the invariant c_E that appears in (1–8) is non-zero on any such end. This understood, it follows that there exists some $s_{j0} > 0$ such that either $\theta > \theta_*$ or else $\theta < \theta_*$ where $|s| > s_{j0}$ on any convex side end in any large j -version of C_{j*} . Meanwhile, by virtue of the fact that the integer n that appears in (2–17) is non-zero on any concave side end of C_j where $\lim_{|s| \rightarrow \infty} \theta = \theta_*$, so θ takes on the value θ_* at points in C_{j*} .

Now, suppose that $\nu: [0, \infty) \rightarrow C_{j*}$ is a smooth, proper map with $\nu(0)$ on the $\theta = \theta_*$ locus and with the image of ν on a convex side end in C_{j*} at all sufficiently large values of its domain. Associate to ν the value of θ where $|\theta - \theta_*|$ is maximal. Next, minimize the latter over all such ν . By virtue of what was said in the previous paragraph, the resulting min-max angle is not equal to θ_* . Moreover, the mountain pass lemma guarantees that the latter is a critical value of θ on C_j and so an angle in θ .

However, since C_{j*} is connected, and θ takes value θ_* on C_{j*} , the large j versions of this critical value can differ by at most some j -independent multiple of ε from θ_* since the paths in C_{j*} stay in the tubular neighborhood U .

Step 2 This step rules out the existence of either a $\theta = 0$ or a $\theta = \pi$ cylinder as an irreducible component of σ by again producing nonsense, a positive critical value of θ that is either too small or too large to be in θ . Suppose, for the sake of argument, that the $\theta = 0$ cylinder is an irreducible component of σ . Only this case is discussed, as the $\theta = \pi$ argument is identical save for some notation.

To see why no $\theta = 0$ cylinder can appear, fix some small, positive ε such that 100ε is less than the smallest angle in θ , and let $U \subset S^1 \times S^2$ denote the radius ε tubular neighborhood of the $\theta = 0$ Reeb orbit.

In this case, Proposition 3.7 provides an irreducible component of each large j version of $C_j \cap (\mathbb{R} \times U)$ with one very particular property. To state the latter, let C_{j*} again denote the component. Here is the salient property: This C_{j*} contains two ends of C_j where the $|s| \rightarrow \infty$ limit of θ is 0, one on the convex side of C_j and the other on the concave side.

Granted the preceding, take j very large. If C_{j*} intersects the $\theta = 0$ cylinder, let s_j denote the largest value that is taken on by s at any $\theta = 0$ points in C_{j*} . If there are no such points, set s_j to equal $-\infty$. Set $C_{j*}' \subset C_{j*}$ to denote that portion where $s > s_j$.

Now, let $\tau: \mathbb{R} \rightarrow C_{j*}'$ be any smooth map with the following two properties: First, $\lim_{r \rightarrow \infty} \tau^*s|_r = \infty$ and $\lim_{r \rightarrow -\infty} \tau^*s|_r = s_j$. Second, $\lim_{r \rightarrow -\infty} \tau^*\theta = 0$. Associate to τ the maximum value of $\tau^*\theta$, and then let θ_{*j} denote the infimum of these maxima over the set of all such maps from \mathbb{R} to C_{j*}' . The mountain pass lemma now guarantees that θ_{*j} is a non-extremal critical value of θ . As $\theta_{*j} > 0$ it is a point in θ . However, it is also the case that θ_{*j} is bounded by 100ε if ε is small, and this is nonsense as 100ε is smaller than the minimal angle in θ .

Step 3 This step and the next complete the proof that Ξ contains but a single element. The argument begins by assuming, to the contrary, that Ξ has more than one component so as to derive some nonsense. In this case, the nonsensical conclusion finds distinct critical values of θ on each large j version of C_j that are closer than the minimal separation between the points in θ .

To start this derivation, reintroduce the set $Y \subset \sigma$ as defined in Step 4 of Subsection 3.D. As remarked earlier, Y is a finite set. Fix a pair of points in $\Sigma - Y$ that lie on distinct irreducible components and choose a θ -preserving preimage of each point and so obtain, for each large j , a pair, z_j and z_j' , of points in C_j . Associate to each path in C_j between z_j and z_j' the supremum of $|s|$ along the path, and let r_j denote the infimum of the resulting subset of $[0, \infty)$. As is explained next, the sequence $\{r_j\}$ is bounded.

To see that there must be such a bound, suppose to the contrary that this sequence diverges. Since neither $\{z_j\}$ nor $\{z_j'\}$ diverges, there is a path in each large j version of C_j between z_j and z_j' that avoids all convex side ends of C_j where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . Indeed, if each such end is defined so that the values of $|s|$ are everywhere greater than its value on either z_j or z_j' , then any path between these points must exit any such end that it enters.

Meanwhile, let $E \subset C_j$ denote a concave side end where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π , and let θ_* denote said limit. A path between z_j and z_j' can also be chosen to avoid E even if it must cross the $\theta = \theta_*$ locus. This is a consequence of the

following observations: First, as is described in Subsection 2.G, the intersection of the $\theta = \theta_*$ locus with E contains just two connected components. These are either the two ends of a single component of the $\theta = \theta_*$ locus, a properly embedded copy of \mathbb{R} , or ends of distinct components. In the latter case, the other ends of the corresponding copies of \mathbb{R} are in other ends of C_j . In any event, if there is no path from z_j to z_j' that crosses the $\theta = \theta_*$ locus at reasonable values of s , it must be the case that the infimum of s on some component of the large j version of this locus is very large. Were such to occur, then the whole of this component would lie very close to some end of σ , thus in $\mathbb{R} \times U$ where U is a small radius tubular neighborhood in $S^1 \times S^2$ of a $\theta = \theta_*$ Reeb orbit. Since ϑ is generic, this means that the two components of the intersection of the $\theta = \theta_*$ locus with E lie on the same component of this locus. This understood, let (p, p') denote the relatively prime pair of integers that θ_* defines via (1–7). As remarked previously, the 1–form $pd\varphi - p'dt$ is exact on U . In particular, the integral of $pd\varphi - p'dt$ from one end to the other of all large j versions of the $\theta = \theta_*$ component in question must then be zero. However, the latter integral can not be zero because the pointwise restriction of $pd\varphi - p'dt$ to such a component is nowhere zero.

With the preceding understood, the only way that $\{r_j\}$ can diverge is if all paths between the large j versions of z_j and z_j' have their large values of $|s|$ where θ is nearly 0 or nearly π . However, such an event is ruled out using the mountain pass lemma. Indeed, under the circumstances just described, this lemma would provide a non-extremal critical point of θ on every large j version of C_j , one whose critical value was either too close to 0 or too close to π to come from θ .

Step 4 This part of the proof argues that Ξ has but one element, and makes use of the following auxiliary lemma:

Lemma 3.10 *Let $Q \equiv (q, q')$ denote a pair of integers and let $\theta_0 < \theta_1$ denote a pair of angles such that the function $\alpha_Q(\cdot)$ from (2–27) is positive on $[\theta_0, \theta_1]$. Given $\varepsilon > 0$, there exists $\varepsilon' \in (0, \varepsilon)$ such that the following is true: Let J_* denote an admissible almost complex structure, and let ϕ denote an immersion of $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ into $\mathbb{R} \times (S^1 \times S^2)$ that is J_* -pseudoholomorphic and defined using a pair of functions, (a_*, w_*) , by the rule*

$$\begin{aligned} (s = a_*, t = qv + (1 - 3 \cos^2 \sigma)w_* \pmod{2\pi}), \\ \theta = \sigma, \varphi = q'v + \sqrt{6} \cos \sigma w_* \pmod{2\pi}) \end{aligned}$$

Then, any two points in $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ with ϕ -image in the complement of any given, radius ε ball in $\mathbb{R} \times (S^1 \times S^2)$ are the endpoints of a continuous path whose ϕ

image lies in the complement of the concentric, radius ε' ball. Moreover, if the two points have the same σ -coordinate, then such a path exists on which σ is constant, and if they have the same v -coordinate, such a path exists on which v is constant.

Proof of Lemma 3.10

This is the case simply because the variation of θ in a radius ε ball and the integral of $(1 - 3 \cos^2 \theta)d\varphi - \sqrt{6} \cos \theta dt$ on any constant θ slice of such a ball are $O(\varepsilon)$ for small ε .

To resume the proof that Ξ has a single element, the next point to make is that if such isn't the case, then σ contains a pair of points that have the same θ value but lie in distinct, irreducible components. Indeed, were there no such pair, then the subvarieties that comprise σ would be pairwise disjoint, and this last conclusion is incompatible with the results of the previous steps. This understood, fix some very small, but positive ε and then fix points p_0 and p_1 in σ that lie on distinct irreducible components, have the same θ value and have distance at least 2ε from any point in the set Y . In particular, choose the θ value to be different than any $|s| \rightarrow \infty$ limit of θ on σ . When the index j is very large, these two points are very close to respective θ -preserving preimages, p_{0j} and p_{1j} in C_j . As the model curve for C_j is connected, there exists a smooth path, $\gamma_j \subset C_j$, that connects p_{0j} to p_{1j} . \square

Next, introduce, as in Subsection 2.G, the C_j version of the model curve, C_0 , and the corresponding locus $\Gamma \subset C_0$. Suppose now that p_{0j} and p_{1j} lie in the same component of the C_j version of $C_0 - \Gamma$. According to Lemma 3.10, the path γ_j can be chosen to be an arc on the constant θ locus that avoids all points of Y by some fixed, j -independent amount. Moreover, as this θ value is not one of the $|s| \rightarrow \infty$ limits of θ on σ , the large j versions of such a path must lie everywhere very close to σ . Thus, any large j version of γ_j has a well defined projection to give a path in σ that avoids all points in Y and runs from p_0 to p_1 . As this is impossible, it must therefore be the case that p_{0j} and p_{1j} lie on distinct components of the C_j version of $C_0 - \Gamma$.

To see that the latter case is also impossible, fix some very small but positive ε and use Lemma 3.10 to construct, for each large index j , a path $\gamma_j \subset C_j$ that runs from p_{0j} to p_{1j} and is an end to end concatenation of two kinds of paths. Paths of the first kind avoid the radius ε balls about the points of Y . Meanwhile, a path of the second kind is an arc on some constant θ locus that is contained in a radius ε ball about a point in Y and passes through a critical point of θ on C_j .

Remember now that each critical point on C_i is non-degenerate and the critical values of distinct critical points are distinct. Thus, if an arc portion of γ_j contains a critical

point of θ , there is a unique component of the C_j version of $C_0 - \Gamma$ whose image in C_j contains this arc in its closure. This understood, it follows from Lemma 3.10 that there exists an j -independent choice for its constant ε' such that the endpoints of this same arc are connected by a path that also avoids the radius ε' balls about the points of Y . Thus, each large j version of C_j contains a second path, ν_j , that connects p_{0j} to p_{1j} and avoids the radius ε' balls about the points of Y . Of course, no such path exists if p_0 and p_1 are in distinct, irreducible components of σ .

3.G Part 3 of the proof of Proposition 3.8

This part of the proof establishes that the one element in Ξ has the form $(S, 1)$ with $S \in \mathcal{M}[\Theta, \vartheta]$. As such, the graph T_S from Subsection 2.G can be labeled as a moduli space graph and it is explained here why the latter is isomorphic to T . The argument for all of this is given below in nine steps.

Step 1 Let S denote the subvariety from Ξ 's one pair. This step establishes that every non-extremal critical value of θ on the model curve of S is in the set θ . In fact, the argument below proves that every non-extremal critical value of θ on the model curve of S maps to a point in S that is a limit point of a sequence whose j 'th element is the image of a critical point on the model curve of C_j .

To start the argument, let S_0 denote the model curve for S and let $z \in S_0$ be a non-extremal critical point of θ and let θ_* denote the associated critical value. Also, set $k \equiv \deg(d\theta|_z) + 1$. Now, if $D \subset S_0$ is a small radius disk that is centered at z , then the $\theta = \theta_*$ locus in S_0 will intersect the boundary of D transversely in $2k$ points. This can be seen, for example, using the local coordinate on D that appears in (2-16). If the radius of D is sufficiently small, then the tautological map from S_0 to $\mathbb{R} \times (S^1 \times S^2)$ will embed the closure of D . Furthermore, the image of z will be the only point from Y in the image of this closure. Take any such small radius for D .

Let $B \subset \mathbb{R} \times (S^1 \times S^2)$ be a ball that contains the closure of the image of D with center at the image of z . Introduce \hat{t} and $\hat{\varphi}$ to denote the respective anti-derivatives on B of dt and $d\varphi$ that vanish at z . Then, let $\hat{v} \equiv (1 - 3 \cos^2 \theta)\hat{\varphi} - \sqrt{6} \cos \theta \hat{t}$. As can be seen using the parametrizations provided by (2-25), the pair (θ, \hat{v}) pulls-back to D as bona fide coordinates on the complement of z . Likewise, they pull back to the model curve of any large j version of C_j as coordinates on the complement of the θ -critical points in the inverse image of B .

Note for use below that \hat{v} is annihilated by ∂_s and the Reeb vector field \hat{a} . As a consequence, the values of both θ and \hat{v} are constant on the set of θ -preserving preimages of any given point in $S \cap B$.

As can be seen using the parametrization from (2–16), there exists an embedded circle, $\gamma \subset D - z$, with the following properties: First, γ intersects the $\theta = \theta_*$ locus transversely in $2k$ points. Second, θ has k local maxima and k local minima on γ , and the values of θ at the local maxima are identical, as are the values at the local minima.

When the index j is very large, then γ has a θ -preserving preimages in C_j . Let γ_j denote the embedded loop in C_j 's model curve that maps to one of these preimages. Suppose, for the sake of argument that there is no θ -critical point in the component in C_j 's model curve of the preimage of B that contains γ_j . Thus, (θ, \hat{v}) provide local coordinates on an open set in C_j 's model curve that contains γ_j .

To continue, start at the local maximum of θ on γ_j with smallest \hat{v} value and traverse γ_j in the direction of increasing \hat{v} . Since γ_j is embedded, successive local maxima must occur at successively larger values of \hat{v} . Such must also be the case for the successive minima. This is impossible if γ_j is embedded, for when the largest value of \hat{v} is attained, the traverse must then return to its start without crossing itself even as it crosses values of \hat{v} that are achieved at the local maxima and local minima.

Thus, C_j 's model curve has a critical point that maps to B . Since the radius of D , thus that of B , can be as small as desired, so θ_* must come from θ .

Step 2 As will now be explained, an argument much like the one just given proves that each critical point of θ on S_0 is non-degenerate. A modification is also used here to prove that Ξ is $(S, 1)$ as opposed to (S, n) with $n > 1$ in the case that θ has a non-extremal critical point on S_0 . Finally, a slightly different modification proves that every element of θ is a critical value of θ on S_0 in the case that $n = 1$.

Here is the proof that $n = 1$: Suppose first that $z \in S_0$ is a critical point of θ , and let $D \subset S_0$ be as before, a very small radius disk that contains z . Let γ be as before. Suppose that j is large and that γ has more than one θ -preserving preimage in the model curve of C_j . As is explained next, this assumption leads to a contradiction. To start, remark that the θ -preserving preimage of γ in C_j is contractible. Indeed, this can be seen using the parametrizations provided in Subsection 2.G with the fact that each such preimage maps to a small radius ball in $\mathbb{R} \times (S^1 \times S^2)$. Being contractible, each preimage of γ is the boundary of an embedded disk in the model curve of C_j . In this regard, the argument given in Step 1 implies that each such disk contains the $\theta = \theta_*$ critical point. This implies that the preimages are nested. In particular, one such preimage, call it γ_0 , bounds a disk that contains all of the others. Let D_0 denote the latter disk. Note that the function θ must take its maxima and minima on the boundary of D_0 since its only critical point is in the interior. Now, let $\gamma_1 \neq \gamma_0$ denote

a hypothetical second θ -preserving preimage of γ . Since $\gamma_1 \subset D_0$ and since the maximum value of θ on γ_1 is also the maximum value of θ on γ_0 , it follows that γ_1 must intersect γ_0 .

Up now is the proof that the non-extremal critical points of θ on S_0 are non-degenerate. To start, let z denote the critical point in question and let γ again be as before. Reintroduce the integer k from Step 1. Thus, $k = 1$ if and only if z is a non-degenerate critical point. In any event, a circumnavigation of γ meets k local maxima of θ and k local minima with all local maxima having the same θ -value, and likewise all local minima. For large values of j , let γ_j denote the θ -preserving preimage of γ in C_j . As argued in the preceding paragraph, γ_j bounds an embedded disk in the model curve of C_j that contains the $\theta = \theta_*$ critical point. Let D_0 denote the latter

To continue, use the descriptions from Subsection 2.G to find an embedded disk, U , in the model curve of C_j with the following properties: The disk contains the $\theta = \theta_*$ critical point, it contains γ_j , and its intersection with the $\theta = \theta_*$ locus consists of four properly embedded, half open arcs that meet only at their common endpoint, this being the $\theta = \theta_*$ critical point. These arcs are called ‘legs’ of the $\theta = \theta_*$ locus. In any event, here is one more requirement for U : The complement of the $\theta = \theta_*$ locus in this disk consists of four open sets, each an embedded disk.

Now, let $U' \subset U$ denote any one of the four components of the complement of the $\theta = \theta_*$ locus. The closure of U' in U intersects γ_j in some number of properly embedded, disjoint, closed arcs. In this regard, there are $2k + 2$ such arcs amongst the four components, with at least one in each. Now, any given such arc in U' together with the stretch of the boundary of U' between its two endpoints defines a piecewise smooth circle in U which is the boundary of the closure of an embedded disk. If the interior of a second arc lies in this disk, then the second arc is said to be nested with respect to the first. Of course, no arc in U' can be nested with respect to another by virtue of the fact that all local maxima of θ on γ_j have the same θ -value, as do all local minima.

There is one more point here to keep in mind: One and only one component arc in $\gamma_j \cap U'$ ‘encircles’ the $\theta = \theta_*$ critical point in the following sense: This critical point is contained in the segment of the boundary of the closure of U' that lies between the arc’s two endpoints. Indeed, this is because the disk D_0 contains the $\theta = \theta_*$ critical point. Note that the endpoints of the latter arc lie on distinct legs of the $\theta = \theta_*$ locus. Such an endpoint on a given leg of the $\theta = \theta_*$ locus is nearer to the $\theta = \theta_*$ critical point than any other arc endpoint on the given leg. This is a consequence of the fact that the arcs that comprise $\gamma_j \cap U'$ are not nested.

Granted all of the above, each component of the complement of $\theta = \theta_*$ locus in U has its one arc that encircles the $\theta = \theta_*$ critical point, so there are four such ‘encircling’ arcs in all. Moreover, by virtue of what is said about endpoints in the preceding paragraph, these four arcs concatenate to define a closed loop in U . This loop must thus be γ , and so $k = 1$ as claimed.

What follows is the proof that every angle from θ is a critical value of θ on S_0 in the case that the integer n that is paired with S in θ is equal to 1. To start, let $\theta_* \in \theta$ and assume that θ_* is not a critical value of θ on S_0 so as to derive some nonsense.

To start the derivation, choose a very small, but positive constant δ , subject to the following constraints: First, neither the $\theta = \theta_* + \delta$ nor $\theta_* - \delta$ loci in S_0 should contain points of Y . Second, no critical values of θ on S_0 lie in the interval $[\theta_* - \delta, \theta_* + \delta]$.

The parametrizations described in Subsection 2.G can now be used to first find a positive constant, ε , and then construct for any sufficiently large j , an embedded circle in C_j 's model curve with four special properties: First, the circle bounds a disk in the model curve that contains the $\theta = \theta_*$ critical point. Second, θ has two local maxima on the circle, both where $\theta = \theta_* + \delta$; and θ has two local minima on the circle both where $\theta = \theta_* - \delta$. Finally, the tautological map to $\mathbb{R} \times (S^1 \times S^2)$ embeds the circle. Finally, all points in the circle's image have $|s| < \frac{1}{\varepsilon}$ and distance ε or more from all points of the set Y .

Any given large j version of such a circle has its θ -preserving projection as an embedded circle in S . (This is where the $n = 1$ assumption is used. If $n > 1$, then this circle will not be embedded.) The latter has its inverse image circle, γ , in S_0 . Now, γ is null-homotopic since the integrals of dt and $d\varphi$ over the original circle in the model curve of C_j are zero. Meanwhile, γ is embedded, it lies where $|\theta - \theta_*| < 2\delta$, and θ 's restriction to γ has two local maxima, both with the same θ -value, and two local minima, also with the same θ -value. A repetition of one of Step 1's arguments now proves that there is no such loop. This nonsense thus proves that θ_* must be a critical value of θ on S_0 .

Step 3 This step investigates the concave side ends of S where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . In particular, this step establishes $\deg_{(c)}(d\theta) = 1$ for all such ends. A variation of the latter argument is then used to prove that the integer n that Ξ associates to S is equal to 1 if S has a concave side end where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$. A second variation of the argument proves that there is precisely one such end of S for every $(0, +, \dots)$ element in \hat{A} . With regards to this last point, keep in mind that the ends of S define an unordered set of points in S^1 via (2-19)'s map ϖ_+ , and the latter is the same as ϑ up to multiplicity. This follows directly from the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3-16).

To start the analysis, let $E \subset S$ denote a concave side end where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . As will now be explained, $\deg_E(d\theta) = 1$. To prove that such is the case, suppose $\deg_E(d\theta) > 1$ so as to derive nonsense. Thus, let $k > 1$ denote $\deg_E(d\theta)$. Let θ_* now denote the $|s| \rightarrow \infty$ limit of θ on E . As can be seen using (2–17), there exists $s_0 \geq 0$ such that the $\theta = \theta_*$ locus intersects the $s \geq s_0$ portion of E as a disjoint union of $2k$ properly embedded copies of $[s_0, \infty)$, with the diffeomorphism given by the function s itself. Moreover, it follows from (2–17) that if $\delta > 0$ and is sufficiently small, there exists an embedded circle, $\nu \subset E$, with the following properties: First, this loop ν has transversal intersections with the $\theta = \theta_*$ locus. Second, $|\theta - \theta_*| < \delta$ on ν . Third, θ 's restriction to ν has precisely k local maxima and k local minima. Finally, all local maxima have the same θ value, this greater than θ_* , and all local minima have the same θ value, this less than θ_* .

The circle ν has θ -preserving preimages in every large j version of C_j . Let ν_j denote one of the latter. Because the variation of θ on ν_j is small, the $\theta < \theta_*$ portion of ν_j is contained in a single component of the C_j version of $C_0 - \Gamma$. Call this component K . For the same reason, the $\theta > \theta_*$ part of ν_j is entirely in a single component also. Use K' for the latter. Note that the closure of the $|\theta - \theta_*| < 2\delta$ portion of $K \cup K'$ is diffeomorphic to a closed cylinder with some number of punctures, all on the $\theta = \theta_*$ circle.

To proceed, now view the loop ν_j sitting in this abstract cylinder. Here, it sits as an embedded, null-homotopic loop. To explain, introduce the relatively prime pair of integers, (p, p') , that θ_* defines via (1–7). The 1-form $p d\varphi - p' dt$ is exact near E and so integrates to zero around ν_j . Meanwhile, this form has non-zero integral around any essential loop in the unpunctured cylinder since this form restricts as a positive form on the $\theta = \theta_*$ locus in C_j . To summarize: As a loop in the abstract cylinder, ν_j is embedded, it is null homotopic, it intersects the $\theta = \theta_*$ locus in $2k$ points, it has k local maxima all with the same value of θ , this greater than θ_* , and k local minima, all with the same value of θ , this less than θ_* . Granted all of this, the argument from Step 1's second to last paragraph can be borrowed with only minor cosmetic changes to obtain a contradiction unless $k = 1$.

Given that there exists a concave side end $E \subset S$ where $\lim_{|s| \rightarrow \infty} \theta \notin \{0, \pi\}$, what follows proves that the integer n that Ξ associates to S is equal to 1. For this purpose, construct a loop, ν , as just described. The loop ν again has θ -preserving preimages in every large j version of C_j . If $n > 1$ and if there are less than n such preimages, then one of them has the following properties: The restriction of θ to the loop has more than one local maximum, and more than one local minimum. Moreover, all local maxima have the same θ value, and all local minima have the same θ value. Finally, the form $p d\varphi - p' dt$ integrates to zero over this loop. The argument given

in the preceding paragraph shows that this is impossible. Thus, there are n disjoint, θ -preserving preimages of ν in every large j version of C_j , each mapping via the θ -preserving projection to ν as a diffeomorphism.

To see that $n = 1$, first note that all θ -preserving preimages of ν must lie in the closure of the union of the same two components of the C_j version of $C_0 - \Gamma$. Indeed, such is the case because the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3-16) holds and because the set ϑ is both j -independent and has N_+ elements.

To continue, let K and K' again denote the two relevant components of the C_j version of $C_0 - \Gamma$ and once again view the $\theta \in [\theta_* - 2\delta, \theta_* + 2\delta]$ part of the union of the closures of K and K' as a closed, multi-punctured cylinder. Viewed in this cylinder, any θ -preserving preimage of ν must encircle one or more of the punctures. Were this otherwise, then the preimage would be null-homotopic in the $\theta \in (\theta_* - 2\delta, \theta_* + 2\delta)$ part of C_j and thus the same part of $\mathbb{R} \times (S^1 \times S^2)$. But such a loop represents the same homotopy class as ν , a non-zero class in the $\theta \in (0, \pi)$ portion of $\mathbb{R} \times (S^1 \times S^2)$.

With the preceding understood, pick a point on the $\theta = \theta_* - 2\delta$ boundary circle of the punctured cylinder, and then draw a smooth path from this chosen point to any given puncture. A loop in the interior of the punctured cylinder that encircles the given puncture has non-zero intersection number with this path. Draw a specific path as follows: Let s_0 denote the maximum of the function s on ν . Start the path at the puncture and draw it to decrease s until the latter is equal to $2s_0$. Call this path γ_o . Note that when j is very large, the whole of this path is in a very small radius tubular neighborhood of a single end of S , and thus it is far from any other end of C_j . This follows because the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3-16) holds and because the set ϑ is both j -independent and has N_+ elements. When j is very large, the $s = 2s_0$ endpoint of γ_o will be very close to one particular end, $E' \subset S$, and so its θ -preserving projection to E' is well defined. Draw an s -decreasing path in E' from the latter point to a point where both s and θ are less than their minimal values on ν . If $E = E'$, then have this path intersect ν transversely at a single point. In any event, call this path $\gamma_{E'}$. When j is large, one of the θ -preserving preimages of $\gamma_{E'}$ attaches to the $s = 2s_0$ endpoint of γ_o . This understood, continue the concatenation of the latter preimage with γ_o as a path from where θ is less than its minimum on ν to the chosen point on the $\theta = \theta_* - 2\delta$ boundary circle of the cylinder. Make θ decrease monotonically on this continuation.

Only one of the paths just described will intersect any θ -preserving preimage of ν . This is the path that at large s is very close to E . Moreover, only one θ -preserving preimage of γ_E can intersect any given θ preserving preimage of ν . Thus, only one θ -preserving preimage of ν encircles a puncture, and so $n = 1$.

Here is why S does not have a pair of ends whose constant s slices limit as $s \rightarrow \infty$ to the same $\theta = \theta_*$ closed Reeb orbit in $S^1 \times S^2$: Were this otherwise, let E and E' denote the two ends involved, and define loops $\nu \subset E$ and $\nu' \subset E'$ as just described. Each has a θ -preserving preimage in the same multi-punctured cylinder in C_j . However, only one will intersect a path as described above from a puncture to the $\theta = \theta_* - 2\delta$ boundary circle of the multi-punctured cylinder.

Step 4 This step proves that $\deg_{(C)}(d\theta) = 0$ for each convex side end of S where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . A variation of the latter argument also proves that there is only one such end for each $(0, -, \dots)$ element in \hat{A} .

To start, suppose that $E \subset S$ is a convex side end of the sort in question, and suppose that $\deg_E(d\theta)$ is non-zero. Let θ_* denote the $|s| \rightarrow \infty$ limit of θ on E . Now (2–17) guarantees that given any $\delta > 0$, there exists R with the following significance: First, the function $|\theta - \theta_*|$ is less than δ on the $|s| > R$ part of E . Second, θ takes values both greater than θ_* and less than θ_* on any constant $|s| \geq R$ slice of E .

To see that such an event is nonsensical, take δ very small. Granted this, take j very large, and there is but one component of the C_j version of $C_0 - \Gamma$ that maps very close to the $s \geq R$ portion of E . Indeed, such is the case because small δ guarantees that there are no components of the C_j version of Γ where $\theta \in (\theta_*, \theta_* + 2\delta)$. Use K to denote the component of the C_j version of $C_0 - \Gamma$ in question. This component must contain a convex side end of C_j where the $|s| \rightarrow \infty$ limit of θ is θ_* . Since $\deg_{(C)}(d\theta) = 0$ on such an end, the description of K offered in Subsection 2.G finds that the function θ must be either strictly less than θ_* or strictly greater than θ_* on K . Of course, this is impossible when j is large. Indeed, when j is large, then K maps very close to the $s \geq R$ portion of E and so to points where θ is greater than and to points where θ is less than θ_* .

To see that there is but one end of S for each $(0, -, \dots)$ element in \hat{A} , suppose for the sake of argument that there were two, E and E' . Let θ_* denote the common $|s| \rightarrow \infty$ limit of θ on E and E' . Fix some $\delta > 0$ and very small, and then let $\nu \subset E$ and $\nu' \subset E'$ denote the respective loci where $|\theta - \theta_*| = \delta$. By virtue of (2–17), these loci are embedded circles. Each has its θ -preserving preimage in any sufficiently large j version of C_j . An argument from the preceding paragraph can be readily modified to prove that these preimages must lie in the same component of the C_j version of $C_0 - \Gamma$. As such, they must coincide.

Note that the same argument proves the following: Suppose that the integer n that Ξ pairs with S is greater than 1, and suppose that $E \subset S$ is a convex side end on which the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . Let $\nu \subset E$ denote an embedded circle in

E that is homologically non-trivial. Then there is but one θ -preserving preimage of ν in every sufficiently large j version of C_j ; and the latter maps back to ν as an n to 1 covering map.

Step 5 This step investigates the nature of the ends of S where the $|s| \rightarrow \infty$ limit of θ is either 0 or π . In particular, it is proved here that such ends are naturally in 1–1 correspondence with the set of $\{\pm 1, \dots\}$ elements in \hat{A} unless S is either a disk or a cylinder and the integer that Ξ pairs with S is greater than 1.

The discussion starts with the following claim:

(3–20) *There exists $R \geq 0$ such that when j is large, the intersections of C_j with the $\theta \in \{0, \pi\}$ cylinders occur where $|s| \leq R$.*

This claim is proved momentarily. Note first that it has the following corollary: In the case that the integer that Ξ pairs with S is 1, the respective intersection numbers between S and the $\theta = 0$ and $\theta = \pi$ cylinders are those prescribed by c_+ and c_- . Indeed, this corollary follows using the parametrizations from (2–25) given Proposition 3.7 and given that θ contains all non-extremal critical values of θ on the model curve of S .

To see why (3–20) holds, fix a very small but positive number, δ , chosen so that there are no elements of θ that lie where $\theta < 2\delta$ and none where $\pi - \theta < 2\delta$. Also, choose δ so that no angle as defined via (1–7) from the integer pair of any $(0, \dots)$ element in \hat{A} lies either between 2δ and 0 or between $\pi - 2\delta$ and π .

Having chosen δ , then choose R so that the $|s| \geq \frac{1}{2}R$ part of S is contained in the ends of S . Moreover, choose R so that the variation of θ on the $|s| \geq \frac{1}{2}R$ part of any end of S is very much smaller than δ . Granted this, suppose that j_0 is such that the $|s| = \frac{1}{2}R$ locus in any end of S has its full set of θ -preserving preimages in all $j \geq j_0$ versions of C_j .

Now suppose, for the sake of argument, that some $j \geq j_0$ version of C_j intersects the $\theta = 0$ locus at a point where $|s| > R$. This point is the image of a point in the closure of a particular component of the C_j version of $C_0 - \Gamma$. Let K denote the latter. The 1-form dt must pull back to K as an exact form. However, as indicated in the preceding paragraph, there is some end of S where $\lim_{|s| \rightarrow \infty} \theta = 0$ whose $|s| = \frac{1}{2}R$ slice has a θ -preserving preimage in K . Since the 1-form dt is not exact on such a slice, so its pull-back to K can not be exact. Thus, there is no such K .

A very minor modification of the argument just given also proves the following: Any given integer pair that appears in some $(1, +, \dots)$ element in \hat{A} is n times that of an

integer pair that is defined by a concave side end of S where the $|s| \rightarrow \infty$ limit of θ is 0, and vice-versa. Here, n is the integer that Ξ pairs with S . Of course, the analogous assertion holds for $(1, -\dots)$ elements and convex side ends where the $\lim_{|s| \rightarrow \infty} \theta = 0$. Likewise, a similar assertion holds for $(-1, \dots)$ elements in \hat{A} and ends of S where the $|s| \rightarrow \infty$ limit of θ is π .

Note that this correspondence assigns precisely one end of S to each end of every large j version of C_j . Indeed, if not then there exists some very small $\varepsilon > 0$ and two disjoint $\theta = \varepsilon$ circles in S , or two disjoint $\theta = \pi - \varepsilon$ circles in S whose θ -preserving preimages in all sufficiently large j versions of C_j lie in the same component of the C_j version of $C_0 - \Gamma$. This means that the preimages coincide, an impossibility when j is large.

The next point to make is that this correspondence is a 1–1 correspondence unless S is either a disk or a cylinder of a certain sort. The assertion that the correspondence is 1–1 follows from the following claim: Two ends of any large j -version of C_j can not both lie very close to the same end of S . To see why the latter claim holds, remark first that the occurrence of two ends very close to the same end of S can occur only in the case that the integer n is greater than 1. This is because distinct ends of C_j that are convex or have $|s| \rightarrow \infty$ limit of θ either 0 or π lie in distinct components of the C_j version of $C_0 - \Gamma$. Now, if $n \neq 1$, then it follows from (2–18) and from what has been said in previous steps that S has at most two ends, and neither is a concave side end unless the corresponding $|s| \rightarrow \infty$ limit of θ is 0 or π . In particular, S is either a cylinder or a disk.

By the way, if S is a disk, then (2–18) requires that S have a single, transversal intersection with one but not both of the $\theta = 0$ or $\theta = \pi$ cylinders. In this case, the integral of dt over any constant θ circle in S must be zero, and so the large $|s|$ slices of S converge in $S^1 \times S^2$ to one of the two $\cos^2 \theta = \frac{1}{3}$ Reeb orbits. In particular, the sign of $\cos \theta$ on this orbit is the same as its sign at the zero of $\sin \theta$. In any event, any large and constant $|s|$ slice of S is isotopic to the $|s| \rightarrow \infty$ limit Reeb orbit.

On the other hand, if S is a cylinder, then (2–18) requires that it be disjoint from both the $\theta = 0$ and $\theta = \pi$ cylinders. In this case, it must have at least one convex side end where the $|s| \rightarrow \infty$ limit of θ avoids 0 and π . Indeed, if not, then the fact that the restriction of θ to S has no extremal critical values in $(0, \pi)$ would require the $|s| \rightarrow \infty$ limit of θ to be 0 on one end and π on the other. Were this the case, the whole of S could be parametrized as in (2–25) by $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$. However, this is impossible because the corresponding function α_Q as defined in (2–27) would then vanish at some value of σ that is realized on the parametrizing cylinder.

Step 6 This step proves that the integer n that Ξ pairs with S is equal to 1 in the cases that S is a disk as described in the preceding step. The cases where S is a cylinder are discussed in Step 7 and Step 8.

The argument has five parts, with the first four constituting a digression to set the stage. In what follows, keep in mind that S is a J' -pseudoholomorphic disc, an element in the J' -version of the moduli space $\mathcal{M}_{\hat{A}'}$ where \hat{A}' has only the 4-tuple $(0, -, (0, 1))$. In particular, S intersects the $\theta = 0$ cylinder transversely in a single point, there are no non-extremal critical points of θ on S , and the $|s| \rightarrow \infty$ limit of the constant $|s|$ slices on S converge to a Reeb orbit where $\cos \theta = \sqrt{1/3}$.

Part 1 For j large, Proposition 3.4 provides C^∞ -small deformations of S that results in a J_j -pseudoholomorphic subvariety in the J_j -version of $\mathcal{M}_{\hat{A}'}$. Any such subvariety intersects the $\theta = 0$ cylinder transversely, also at a single point and there are no non-extremal critical points of θ on any such S_j .

Proposition 2.12 and Proposition 2.13 in conjunction with Proposition 3.4 provide two different parametrizations of the subsets of the respect J' and, for large j , J_j versions of $\mathcal{M}_{\hat{A}'}$ whose subvarieties are everywhere close to S in $\mathbb{R} \times (S^1 \times S^2)$. The first parametrizes the constituent subvarieties by the point where they intersect the $\theta = 0$ cylinder, and the other by their large $|s|$ asymptotics on their one end.

To be more explicit about the parametrization by points in the $\theta = 0$ cylinder, let z_0 denote the point where S intersects this cylinder. Then respective neighborhoods that contain the subvarieties that are pointwise near S in the J' version and, for large j , in the J_j version of $\mathcal{M}_{\hat{A}'}$ can be parametrized as follows: The coordinates for the parametrization are the points that lie in a j -independent disk centered at z_0 in the $\theta = 0$ cylinder. The parametrization provides a 1-1 correspondence that assigns a subvariety in the relevant moduli space to the point where it intersects the $\theta = 0$ cylinder.

Part 2 The second parametrization uses the large $|s|$ asymptotics on the subvariety. To be more precise, first note that the $\cos \theta = \sqrt{1/3}$ Reeb orbits are parametrized by the constant value on the Reeb orbit of the coordinate, t , on the S^1 factor in $S^1 \times S^2$. Let τ_0 denote the value for the orbit that is obtained as the $|s| \rightarrow \infty$ limit of the constant $|s|$ slices of S . Meanwhile, let c_0 denote the constant that appears in the version of (1-8) that is relevant for the one end of S . In this regard, note that $c_0 < 0$.

According to the aforementioned propositions, there exists some $\delta > 0$ such that respective neighborhoods of subvarieties that are pointwise near S in the J' version and, for large j , in the J_j version of $\mathcal{M}_{\hat{A}'}$ can be parametrized as follows: The

parametrization uses those $(c, \tau) \in (-\infty, 0) \times \mathbb{R}/(2\pi\mathbb{Z})$ where $|c - c_0|^2 + |\tau - \tau_0|^2 < \delta^2$. In particular, the parametrization provides a 1–1 correspondence that assigns a subvariety to a pair (c, τ) when the subvariety's one end provides $c_E = c$ in (1–8), while the large $|s|$ slices of the subvariety converge as $|s| \rightarrow \infty$ to the Reeb orbit where $\cos \theta = \sqrt{1/3}$ and $t = \tau$. In this regard, note that the assigned value for τ is simply the value on the end in question of Section 1's parameter $\iota_{(\cdot)}$.

Part 3 Both of the preceding parametrizations are compatible with the notion of convergence as given in Proposition 3.7. Indeed, suppose that either a point in the $\theta = 0$ cylinder is fixed in the parametrizing disk about z_0 , or else a point (c, τ) is fixed with distance less than δ from (c_0, τ_0) . Now, in either case, let S_j denote the subvariety in the J_j version of $\mathcal{M}_{\hat{A}}$ that is parametrized by the given point. Meanwhile, let S' denote the corresponding J' -version. Then the sequence $\{S_j\}$ converges in the sense of Proposition 3.7 with $(S', 1)$ in the role of Ξ , and with (3–16) valid for $K = \mathbb{R} \times (S^1 \times S^2)$.

Part 4 The part of the argument is summarized by the following lemma:

Lemma 3.11 *There exists $\varepsilon > 0$ with the following significance: If j is large and if (c, τ) has distance less than δ from (c_0, τ_0) but $|\tau - \tau_0| > \frac{1}{4}\delta$, then all points in the (c, τ) subvariety from the J_j -version of $\mathcal{M}_{\hat{A}}$ have distance at least ε from S .*

The digression ends with the following proof.

Proof of Lemma 3.11 Granted the contents in Part 2 of this digression, it is sufficient to prove the following: If (c, τ) is δ close to (c_0, τ_0) and if $\tau \neq \tau_0$, then the (c, τ) subvariety in the J' -version of $\mathcal{M}_{\hat{A}}$ is disjoint from S . To establish this claim, let S' denote the (c, τ) subvariety. If S' intersects S , then it has positive intersection number with S . Thus, any deformation of S' has positive intersection number with S if the intersections with S along the way remain in a fixed, compact subset of $\mathbb{R} \times (S^1 \times S^2)$. This understood, push S' along the vector field ∂_s . In this regard, there is a compact subset of $\mathbb{R} \times (S^1 \times S^2)$ that contains all the putative intersections between this deformation and S . Moreover, the same subset suffices no matter how far S' is pushed along ∂_s . Indeed, such is the case because $\tau \neq \tau_0$ and s is bounded from above on S .

If pushed far enough, the resulting subvariety at values of s that are achieved on S is very close to the pseudoholomorphic cylinder that is defined by the $((0, 1), \tau)$ Reeb orbit. As θ has no local maximum on S , it follows that S must be disjoint from this Reeb orbit, and so disjoint from a large push of S' along ∂_s .

Part 5 With the preliminaries over, what follows are the final arguments to prove that the integer paired to S by Ξ is 1. To start, let ε denote the constant from Lemma 3.11. When j is large, then every point in C_j has distance less than $\frac{1}{100}\varepsilon$ from S . Also, when j is large, the parametrizations from Part 1 and Part 2 above provide some pair (c, τ) that is very close to (c_0, τ_0) and have the following significance: This pair parametrizes a disk in the J_j version of $\mathcal{M}_{\hat{A}}$ that intersects C_j at a point on the $\theta = 0$ cylinder. Let S_j denote the latter disk. Now, the $|s| \rightarrow \infty$ limit of the constant $|s|$ slices of C_j converge as a multiple cover to some $\cos \theta = \frac{1}{\sqrt{3}}$ Reeb orbit with parameter τ_j very close to τ_0 . If $\tau_j = \tau$ and if the integer paired to S in Ξ is greater than 1, then use Propositions 2.12 and 2.13 to find a new subvariety, C_j' , from the J_j -version of $\mathcal{M}_{\hat{A}}$ that intersects S_j on the $\theta = 0$ cylinder, has asymptotic parameter $\tau_j' \neq \tau$, and lies entirely in the radius $\frac{1}{100}\varepsilon$ tubular neighborhood of S . In particular, S_j and C_j' have positive intersection number. Now, move S_j in its moduli space to some S_j' whose corresponding parameters (c', τ') obey $|\tau' - \tau_0| > \frac{1}{2}\delta$. Do so by a path $r \rightarrow (c, \tau(r))$ where $\tau(0) = \tau$ and where $|\tau(r) - \tau_j'|$ is strictly increasing. The intersection number between S_j' and C_j' is thus the same as that between S_j and C_j' . However, according to Lemma 3.11, the subvarieties S_j' and C_j' are disjoint. This is a contradiction and so Ξ must pair S with 1. \square

Step 7 This step and Step 8 prove that the integer paired with S by Ξ is 1 in the case that S is a cylinder with no concave side ends where $\lim_{|s| \rightarrow \infty} \theta$ is not in $(0, \pi)$ and at least one convex side end where the analogous limit is in $(0, \pi)$. In this regard, the arguments are, with minor modifications, a reprise of those given previously in Step 6.

This step considers the case that neither 0 nor π is the $|s| \rightarrow \infty$ limits of θ on S . In this case, S is a subvariety in the J' version of the moduli space $\mathcal{M}_{\hat{A}'}$ where \hat{A}' consists of two elements, $(0, -, (p, p'))$ and $(0, -, (-p, -p'))$. Here, p and p' are relatively prime integers and such that $\frac{p'}{p} \geq \sqrt{3}/2$. No generality is lost by taking p and p' to be positive. The argument here has three parts.

Part 1 As in the case considered by Step 6, there are two parametrizations for neighborhoods that contain the subvarieties that are pointwise near S from the J' version and, for large j , from the J_j version of $\mathcal{M}_{\hat{A}'}$. In this case, a description of these parametrizations requires the use of the respective pairs $(c_{+0}, \tau_{+0}) \in (0, \infty) \times \mathbb{R}/(2\pi\mathbb{Z})$ and $(c_{-0}, \tau_{-0}) \in (-\infty, 0) \times \mathbb{R}/(2\pi\mathbb{Z})$ to parametrize the asymptotics of the (p, p') and $(-p, -p')$ ends of S . Here, the parameters $\tau_{\pm 0}$ parametrize the respective (p, p') and $(-p, -p')$ Reeb orbits that are obtained as $|s| \rightarrow \infty$ limits of the constant $|s|$ slices of S . Meanwhile, $c_{\pm 0}$ are the respective versions of the parameter c_E from (1–8).

It then follows from Propositions 2.13 and 3.4 that there exists $\delta > 0$ such that one of the parametrizations in question is by the subset in $(0, \infty) \times \mathbb{R}/(2\pi\mathbb{Z})$ of points (c, τ) where $|c - c_{+0}|^2 + |\tau - \tau_{+0}|^2 < \delta$, and the other is by the subset in $(-\infty, 0) \times \mathbb{R}/(2\pi\mathbb{Z})$ that consists of the points (c, τ) where $|c - c_{-0}|^2 + |\tau - \tau_{-0}|^2 < \delta$. The first parametrization assigns a subvariety to the pair (c, τ) that describe the asymptotics of its (p, p') end, and the second those of its $(-p, -p')$ end. Thus, τ is again the value of Section 1's parameter $\iota_{(\cdot)}$ on the end in question. Call the first parametrization the 'plus' parametrization and call the second one the 'minus' parametrization.

Part 2 The following observation is the analog of that made in Lemma 3.11: There exists $\varepsilon > 0$ with the following significance: If (c, τ) has distance less than δ from (c_{+0}, τ_{+0}) and if $|\tau - \tau_{+0}|$ is greater than $\frac{1}{4}\delta$, then the large j version of the subvariety that is parametrized by (c, τ) via the plus parametrization lies outside the radius ε tubular neighborhood of S . Of course, the analogous statement holds for the minus parametrization when (c, τ) is δ close to (c_{-0}, τ_{-0}) and $|\tau - \tau_{-0}| > \frac{1}{4}\delta$.

To prove the assertion in this case, it is enough to consider, as in the proof of Lemma 3.11, the intersections between S and the J' subvariety, S' , that is parametrized by (c, τ) via the appropriate parametrization. For this purpose, note that if (c, τ) gives S' by the plus parametrization and if (c', τ') gives S' via the minus one, then $\tau' \neq \tau_{-0}$ if and only if $\tau \neq \tau_{+0}$.

Having said this, consider deforming S' by pushing it along the vector field ∂_s . Such a deformation keeps the intersections with S in a compact set of $\mathbb{R} \times (S^1 \times S^2)$ because $\tau \neq \tau_{+0}$ and $\tau' \neq \tau_{-0}$. Thus, the intersection number between the resulting subvariety and S is that between S' and S . Of course, the latter is zero if and only if S is disjoint from S' . Now, by virtue of the fact that the function s is bounded from above on S , if S' is pushed far enough along the vector field ∂_s , then the portion of the resulting subvariety where s has values that are also achieved on S has two components, each very close to an \mathbb{R} -invariant cylinder. Of course, one of these cylinders is the product of \mathbb{R} with the Reeb orbit parametrized by $((p, p'), \tau)$, and the other the product of \mathbb{R} with the Reeb orbit parametrized by $((-p, -p'), \tau')$. Now, as θ has neither maxima nor minima on S , it follows that S stays a uniform distance from both of these constant θ cylinders. Thus, the deformation of S' is disjoint from S and so S' is also disjoint from S .

Part 3 Granted all of the proceeding, take j very large, and in particular, large enough so that C_j is contained in the radius $\frac{1}{100}\varepsilon$ tubular neighborhood of S . The point now is that there exists some S_j in the J_j version of $\mathcal{M}_{\hat{\lambda}}$ that intersects C_j . This follows

using Proposition 3.4, the lead observation in Part 2, and the observation made towards the end of Part 2 that distinct subvarieties in any given version of $\mathcal{M}_{\hat{A}}$ are disjoint.

If the integer paired with S by Ξ were greater than 1, then any subvariety S_j from the J_j version of $\mathcal{M}_{\hat{A}}$ that intersects C_j must do so in a finite set of points. This understood, pick such an S_j that is parametrized by some (c, τ) via the plus parametrization, and some (c', τ') via the minus parametrization. Now use Proposition 2.13 to find a subvariety $C_j' \in \mathcal{M}_{\hat{A}}$ with the following properties: First, it lies in the radius $\frac{1}{100}\varepsilon$ tubular neighborhood of S and it intersects S_j . Second, the $|s| \rightarrow \infty$ limit of the constant $|s|$ slices converge to Reeb orbits that are distinct from both the $((p, p'), \tau)$ and $((-p, -p'), \tau)$ Reeb orbits. Note that by virtue of S_j and C_j' intersecting in a finite set of points, the latter have positive intersection number between them.

Having chosen C_j' , now deform S_j by pushing it along the vector field ∂_s . The argument given at the end of the previous part works as well here to establish that result of a large push has the same intersection number with C_j' as does S_j , but is also disjoint from C_j' . As these two constraints are mutually exclusive, it follows that Ξ assigns 1 to S .

Step 8 This step proves that the integer paired with S by Ξ is 1 in the case where S is a cylinder with one end where the $|s| \rightarrow \infty$ limit of θ is in $\{0, \pi\}$ and where the other end is a convex side end where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . Granted the discussion in the previous step, the simplest case to consider is that where both ends of S are convex side. In this case, the argument from the previous step translates with almost no essential changes to handle this case. In fact, the only slight substantive difference arises in from, the different meanings of Section 1's parameters (c_E, ι_E) in the cases that E is an end where $\lim_{|s| \rightarrow \infty} \theta$ is or is not one of 0 or π . In any event, the details for this case are left to the reader.

Turn instead to the case where S has a concave side end where the $|s| \rightarrow \infty$ limit of θ is either 0 or π . Again, save for notation, no generality is lost by taking this limit to be 0. The argument for this case differs somewhat from that in the preceding case and in Step 7 because the function s on S ranges over the whole of \mathbb{R} . In particular, a somewhat different argument must be used to establish that distinct subvarieties in any given version of $\mathcal{M}_{\hat{A}}$ are disjoint. In particular, where in Step 7 (and in the proof of Lemma 3.11), the subvariety S' was pushed along the vector field ∂_s , the argument now pushes S' along the vector field $-\partial_\theta$ to values of θ very near zero. Make this change and then the rest of the argument amounts to little more than a notationally changed version of that given previously.

Step 9 Here is a summary of what has been established by the preceding steps: First, the one element in Ξ has been shown, in all cases, to have the form $(S, 1)$. As for S , its ends are known to be canonically in 1–1 correspondence with the elements in \hat{A} . In addition, the value of $\deg_{(\cdot)}(d\theta)$ on all $(0, +, \dots)$ ends of S has been shown to be 1, and its value on all $(0, -, \dots)$ ends has been shown to be zero. The set ϑ arises from the asymptotic data for the ends of S that correspond to $(0, +, \dots)$ elements of \hat{A} . Meanwhile, the angles in θ are now known to be in 1–1 correspondence with the critical values of θ on S , and each non-extremal critical point of θ has been proved non-degenerate. The arguments given above also prove that there are $N_- + \hat{N} + c_{\hat{A}} - 2$ such critical points in all. Finally, the respective numbers, counting multiplicity, of the intersections between S and the $\theta = 0$ and $\theta = \pi$ cylinders are c_+ and c_- .

Granted all of this, it follows directly that $S \in \mathcal{M}_{\hat{A}}[\Theta, \vartheta]$. Moreover, since the $K = \mathbb{R} \times (S^1 \times S^2)$ version of (3–16) holds, it follows directly that the graph T_S from Subsection 2.G when labeled as a moduli space graph is isomorphic to the graph T .

4 Constrained punctured spheres

This section completes some unfinished business from Section 3 by finishing the proof of Theorem 3.1. This is done with the specification of a collection $\{(a_e, w_e)\}_{e \in T}$ that meets the criteria that are laid out in Subsection 3.B and (3–3).

What follows is a brief outline of the manner in which $\{(a_e, w_e)\}$ are specified. Subsection 4.A starts the story with a description of any given pair (a_e, w_e) at points in the parametrizing cylinder that are comparatively far from the boundary circles. This description involves a set of two positive but very small numbers, $\{\rho_{e0}, \rho_{e1}\}$, that are constrained in the subsequent subsections plus a function, ε_e , of the coordinate σ on the closed parametrizing cylinder. In this regard, ε_e is strictly positive. Keep in mind throughout that all of the subsequent constraints involve only upper bounds on ε_e , ρ_{e0} and ρ_{e1} . No positive lower bounds arise.

The mid-cylinder definition of (a_e, w_e) also involves three additional functions of the coordinate σ on the closed parametrizing cylinder, these denoted by a_e^0 , w_e^0 and v_e . These three have no essential role until Subsection 4.E, and until then, they are unconstrained save for their boundary values. However, substantive constraints do arise on a_e^0 and w_e^0 in the final section so as to insure that distinct versions of $K_{(\cdot)}$ intersect transversely with $+1$ local intersection numbers.

Subsections 4.B, 4.C and 4.D specify $\{(a_e, w_e)\}$ near the boundaries of the parametrizing cylinders. In this regard, Subsection 4.B specifies these pairs near boundaries that

correspond to the monovalent vertices in T . Subsection 4.C does this same task near the boundaries that correspond to the bivalent vertices in T ; and Subsection 4.D gives the specifications for boundaries that correspond to the trivalent vertices in T . The criteria in Definition 3.2 are addressed in Subsections 4.B and 4.C. In this regard, note that the definitions in Subsection 4.C are relevant only to the case where partition for the graph T has only single element subsets.

These subsections also provide constraints on the relevant versions of $(\varepsilon_e, \rho_{e0}, \rho_{e1})$, but all are of the following sort: An upper bound appears for the values of ε_e near each boundary circle of the parametrizing domain. A particular choice for ε_e then determines upper bounds for ρ_{e0} and ρ_{e1} . As remarked above, no positive lower bounds arise. Mild constraints on $\{(a_e^0, w_e^0, v_e^0)\}$ occur in these subsections.

Sections 4.B, 4.C and 4.D also address the nature of the singular points in the resulting versions of $K_{(\cdot)}$. In particular, they prove that any singular point in the closure of any given version of $K_{(\cdot)}$ arises as the transversal intersections of two disks with $+1$ local intersection number. Note that these subsections do not address the nature of the intersections between versions of $K_{(\cdot)}$ with distinct edge labels.

Subsection 4.E, addresses this last issue by explaining how to modify the original choices for $\{(\varepsilon_e, \rho_{e0}, \rho_{e1}, a_e^0, w_e^0, v_e^0)\}$ subject to all previously noted constraints to guarantee that distinctly labeled versions of $K_{(\cdot)}$ have transversal intersections with $+1$ local intersection number.

Granted the results from Section 3, the discussion in Subsection 4.E completes the proof of Theorem 3.1 in the case that the partition for T has only single element subsets. Subsection 4.F completes the proof of Theorem 3.1 in the cases where the latter assumption does not hold.

4.A Parametrizations in the mid-cylinder

To start, fix a number, δ that is positive but less than $\frac{1}{1000}$ times the difference between the maximal and minimal angle labels of the vertices on every edge of T .

Now, let e denote a given edge in T , and let θ_o and $\theta_1 > \theta_o$ denote the angles that are assigned to the vertices of T that lie on e . Fix a positive numbers $\rho_0 \equiv \rho_{e0}$, $\rho_1 \equiv \rho_{e1}$ but constrained so that both are much smaller than δ . In addition, choose a similarly small, strictly positive function $\varepsilon \equiv \varepsilon_e$ on $[\theta_o, \theta_1]$. The constructions that follow assume that ρ_0 , ρ_1 and ε are all very small.

Let σ denote the coordinate on $[\theta_o, \theta_1]$ and let v denote the usual affine coordinate on $\mathbb{R}/(2\pi\mathbb{Z})$. At values of $\sigma \in [\theta_o + 2\rho_0, \theta_1 - 2\rho_1]$, the pair (a_e, w_e) are given by

$$(4-1) \quad \begin{aligned} a_e(\sigma, v) &= a_e^0(\sigma) + \varepsilon(\sigma) \cos(v + v_e^0(\sigma)) \\ \text{and } w_e(\sigma, v) &= w_e^0(\sigma) - \varepsilon(\sigma) \sin(v + v_e^0(\sigma)) \end{aligned}$$

where a_e^0 , w_e^0 and v_e^0 functions on $[\theta_o, \theta_1]$. These functions are as yet unconstrained.

One point to verify at the outset is whether the use of (4-1) leads via (3-2) to an embedding in $\mathbb{R} \times (S^1 \times S^2)$ of the $\sigma \in [\theta_o + 2\rho_0, \theta_1 - 2\rho_1]$ portion of the parametrizing cylinder. That such is the case when ε is small is one consequence of the following lemma.

Lemma 4.1 *Suppose that $\theta_o < \theta_1$ are angles in $[0, \pi]$ and that $Q = (q, q')$ is an integer pair such that $\alpha_Q(\sigma) > 0$ when $\sigma \in (\theta_o, \theta_1)$. Now, suppose that a^0 , w^0 and v^0 are smooth functions on $[\theta_o, \theta_1]$, and suppose that ε is a strictly positive function of σ and constrained so that $\varepsilon\alpha_Q < \frac{1}{2}$ at all points. Use this data to define the functions*

$$a \equiv a^0 + \varepsilon \cos(v + v^0) \quad \text{and} \quad w \equiv w^0 - \varepsilon \sin(v + v^0)$$

on the cylinder $(\theta_o, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$. The pair (a, w) then define an embedding of the cylinder $(\theta_o, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ into $\mathbb{R} \times (S^1 \times S^2)$ via the map in (3-2).

Proof of Lemma 4.1 Suppose for the moment that the functions a and w that are used in (2-25) are any given pair of functions on (θ_o, θ_1) . As remarked in Subsection 2.G, the resulting map then defines an immersion of the parametrizing cylinder when $\alpha_Q > 0$. This understood, the issue is whether two distinct points in the domain are mapped to the same point in the range. To analyze this last issue, note that points (σ, v) and (σ', v') from the parametrizing cylinder $(\theta_o, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ are mapped to the same point in $\mathbb{R} \times (S^1 \times S^2)$ if and only if both $\sigma = \sigma'$ and there exists an integer pair $N = (n, n')$ such that

$$(4-2) \quad \begin{aligned} v' &= v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)} \pmod{2\pi\mathbb{Z}}, \\ a_e\left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)}\right) &= a_e(\sigma, v), \\ w_e\left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)}\right) &= w_e(\sigma, v) - 2\pi \frac{nq'_e - n'q_e}{\alpha_{Q_e}(\sigma)}. \end{aligned}$$

Here, and below, α_N denotes the function $\sigma \rightarrow (1 - 3 \cos^2 \sigma)n' - \sqrt{6} \cos \sigma n$.

To analyze the condition in (4–2), fix a pair $Z \equiv (z, z') \in \mathbb{Z} \times \mathbb{Z}$ of integers such that $z q_{e'} - z' q_e = m$ with m used here to denote the greatest common divisor of the ordered pair of integers that comprise $Q_e \equiv (q_e, q_{e'})$. With Z fixed in this way, then N can be written as $N = \hat{x}(z, z') + \frac{\hat{y}}{m}(q_e, q_{e'})$ with \hat{x} and \hat{y} integers. This notation allows the third point in (4–2) to be written as

$$(4-3) \quad \alpha_{Q_e}(\sigma)w_e \left(\sigma, v - 2\pi \hat{x} \frac{\alpha_Z(\sigma)}{\alpha_{Q_e}(\sigma)} - 2\pi \frac{\hat{y}}{m} \right) = \alpha_{Q_e}(\sigma)w_e(\sigma, v) - 2\pi m\hat{x}.$$

In particular, this last condition implies that there are at most a finite number of possible values for \hat{x} that can appear at any given value of $\sigma \in (\theta_o, \theta_1)$. Meanwhile, the second condition in (4–2) requires that

$$(4-4) \quad a_e \left(\sigma, v - 2\pi \hat{x} \frac{\alpha_Z(\sigma)}{\alpha_{Q_e}(\sigma)} - 2\pi \frac{\hat{y}}{m} \right) = a_e(\sigma, v).$$

Together, (4–3) and (4–4) imply that only \hat{y} 's reduction modulo m is relevant. Thus, there is a finite set of possible values for (\hat{x}, \hat{y}) that need be considered for immersion points in any given compact subset of $(\theta_o, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$.

Now consider the additional ramifications of (4–3) and (4–4) in the case that a and w are as described in the lemma. The first point to make is that with ε constrained as indicated, only $\hat{x} = 0$ can appear in (4–3). This understood, it then follows that both (4–3) and (4–4) can hold simultaneously only if $\hat{y} = 0 \pmod{m\mathbb{Z}}$. Thus v and v' in (4–2) agree mod $(2\pi\mathbb{Z})$ and (4–1) is an embedding. \square

4.B Parametrizations near boundary circles with a monovalent vertex label

Suppose here that $e \subset T$ is an edge and that θ_o and $\theta_o < \theta_1$ are the angles that are assigned to the vertices on e . Suppose, in addition that $o \in e$ is a monovalent vertex from T . For the sake of argument, suppose that the latter is assigned the angle θ_o . The story when the assigned angle is θ_1 is identical save for notation and some sign changes and so the latter case is not presented.

There are three separate cases to consider, these depending on the label given the vertex o . These cases are considered in turn below. In what follows, β denotes a favorite smooth function on $[0, \infty)$ that takes value 1 on $[0, 1]$, value 0 on $[2, \infty)$, and has negative derivative on $(1, 2)$. Having chosen β , and granted that $\rho > 0$ and $\theta_* \in [0, \pi]$, introduce the function

$$(4-5) \quad \beta' \equiv \beta \left(\frac{1}{\rho^4} |\sigma - \theta_*| \right)$$

Case 1 In this case, o is assigned a $(1, \pm, \dots)$ label in \hat{A} . This is to say that in the respective $+$ and $-$ cases, the image in $\mathbb{R} \times (S^1 \times S^2)$ of the $\sigma < 2\rho_{e0}$ part of the parametrizing cylinder should have the asymptotics of a concave side or convex side end of a J -pseudoholomorphic subvariety where the $|s| \rightarrow \infty$ limit of θ is 0.

In the remainder of this Case 1 discussion, ρ denotes ρ_{e0} . In this regard, $\rho \equiv \rho_{e0}$ along with $\theta_* = 0$ are to be used for defining the function β' via (4-5).

To extend the definition in (4-1) of (a_e, w_e) to the $\sigma < 2\rho$ portion of the parametrizing cylinder, first constrain the functions ε , a_e^0 , w_e^0 and v_e^0 that appear in (4-1) to be constant where $\sigma < 2\rho$. This understood, extend the definition in (4-1) to the points where $\sigma < 2\rho$ by setting

$$(4-6) \quad \begin{aligned} a_e &= \frac{1}{\kappa} \beta' \ln \sigma + a_e^0 + (\varepsilon(1 - \beta') + \sigma\beta') \cos(v + v_e^0). \\ w_e &= (1 - \beta')w_e^0 - (\varepsilon(1 - \beta') + \sigma\beta') \sin(v + v_e^0). \end{aligned}$$

Here,

$$(4-7) \quad \kappa \equiv \frac{qe'}{q_e} + \sqrt{\frac{3}{2}},$$

According to Lemma 4.1, any small ε version of (4-1) embeds the $\sigma \in (0, \theta_1 - 2\rho)$ portion of the parametrizing cylinder in $\mathbb{R} \times (S^1 \times S^2)$. Moreover, as $|a_e| \rightarrow \infty$ uniformly as $\sigma \rightarrow 0$, any such version of (4-1) defines a proper embedding of this same portion of the parametrizing cylinder. Thus, the only issue to consider is whether the $\sigma \rightarrow 0$ asymptotics are correct. In particular, the key point here is to verify (1-12), and the latter task is straightforward so left to the reader.

Case 2 In this case, the vertex 0 is assigned the element (1) from \hat{A} . To start, once again set $\rho \equiv \rho_{e0}$ when referring to the function β' in (4-5), also set $\theta_* = 0$. It is also to be understood here that a_e^0 , w_e^0 and v_e^0 from (4-1) are again constrained to be constant where $\sigma \leq 2\rho$. Granted these conventions, the extension of the pair (a_e, w_e) to the points where $\sigma < 2\rho$ is given by

$$(4-8) \quad \begin{aligned} a_e &= a_e^0 + \varepsilon(1 - \beta') \cos(v + v_e^0) \\ \text{and } w_e &= w_e^0 - \varepsilon(1 - \beta') \sin(v + v_e^0). \end{aligned}$$

The reader is left to verify that the resulting extension of (4-1) to the closed cylinder $[0, \theta_1 - 2\rho_1] \times \mathbb{R}/(2\pi\mathbb{Z})$ maps it onto an embedded, closed disk in $\mathbb{R} \times (S^1 \times S^2)$ that intersects the $\theta = 0$ circle transversely and with intersection number $+1$ with respect to the latter's symplectic orientation.

Granted that the extension given in (4–8) maps onto an embedded disk, note that at the latter’s intersection point with the $\theta = 0$ cylinder, the restriction of the symplectic form on its tangent space is positive. Indeed, this can be seen from the fact that the symplectic form pulls back along the $\sigma = 0$ circle in the parametrizing cylinder to the form $\sigma d\sigma \wedge dv$.

Case 3 *In this case, the monovalent vertex is assigned some $(0, -, \dots)$ element from \hat{A} . As in the previous cases, set ρ to equal ρ_{e0} . Use this value for ρ and use $\theta_* = \theta_0$ for defining the function β' .*

The parametrization given below requires that a_e^0 , w_e^0 and v_e^0 from (4–1) are constant where $\sigma \in [\theta_0, \theta_0 + 2\rho]$. Granted that such is the case, extend the definition of (a_e, w_e) to the $\sigma < \theta_0 + 2\rho$ portion of the parametrizing cylinder using the rule

$$(4-9) \quad \begin{aligned} a_e &\equiv \frac{1}{\zeta} \beta' \ln(\sigma - \theta_0) + a_e^0 + (\varepsilon(1 - \beta') + (\sigma - \theta_0)\beta') \cos(v + v_e^0), \\ w_e &= (1 - \beta')w_e^0 - (\varepsilon(1 - \beta') + (\sigma - \theta_0)\beta') \sin(v + v_e^0), \end{aligned}$$

where $\zeta = \sqrt{6 \sin^2 \theta_0 (1 + 3 \cos^2 \theta_0) / (1 + 3 \cos^4 \theta_0)}$. It is left as another exercise for the reader to verify that (4–9) and (3–2) together define a proper embedding of the $\sigma \in (\theta_0, \theta_1 - 2\rho_{e1}]$ portion of $(\theta_0, \theta_1) \times \mathbb{R} / (2\pi\mathbb{Z})$ into $\mathbb{R} \times (S^1 \times S^2)$ as submanifold whose large $|s|$ asymptotics meet the requirements of Definition 3.2 to be those of a convex side end in some J -pseudoholomorphic subvariety.

4.C Parametrizations near boundary circles with a bivalent vertex label

In this subsection, o denotes a bivalent vertex in T whose associated partition subset has but a single element. In what follows, e and e' are the two incident edges to o with the convention that o 's angle label, θ_o , is the greater of the two angles that label the vertices on e , and so the lesser of the two that label the vertices on e' .

The story starts with a preliminary digression to set the stage. To begin the digression, note that $\rho_{e1} = \rho_{e'0}$ in what follows, and ρ denotes either. Take $\rho \ll \delta$. The function β' now refers to the version in (4–5) with this same value for ρ and with θ_* set to equal θ_o .

Require that both the e and e' versions of ε in (4–1) are constant where $|\sigma - \theta_o| < \delta$ and that these constants agree. Require that a_e^0 is constant where $\sigma > \theta_o - 2\rho$, that $a_{e'}^0$ is constant where $\sigma < \theta_o + 2\rho$, and that these two constants also have the same value. Use a_0 for the latter. In addition, require that both w_e^0 and $w_{e'}^0$ are zero where σ is within 2ρ of θ_o . Finally, require that both v_e^0 and $v_{e'}^0$ are constant where σ

is respectively greater than $\theta_o - 2\rho$ and less than $\theta_o + 2\rho$ with equal value, and for notational convenience, take the constant to equal 0. To obtain the case where the constant value for v_e^0 and $v_{e'}^0$ is non-zero, replace v in what follows by $v - v_e^0$.

To proceed with the digression, introduce $P_0 \equiv (p_0, p_0')$ to denote the integer pair from o 's element in \hat{A} , and set

$$(4-10) \quad x_0 \equiv q_{e'} p_0 - q_e p_0'.$$

Note that x_0 is a positive integer. To explain, remember that (p_0, p_0') is a positive multiple of the pair (p, p') that is defined by θ_o via (1-7). This understood, positivity of x_0 is a consequence of the positivity of the $Q = Q_e$ version of $\alpha_Q(\theta_o)$. Note that the formula for x_0 in (4-10) can be written with the pair $Q_{e'}$ replacing Q_e ; this a consequence of (3-1).

The next task for this digression is to define certain 'polar' coordinates for respective neighborhoods of $(\theta_o, 0)$ in both the e and e' versions of the parametrizing cylinder. In both cases, the 'radial' coordinate is denoted as r ; it takes values in $[0, 3\rho)$. Meanwhile, the angular coordinate is denoted as τ ; it takes values in $[-\pi, 0]$ on e' 's version of the parametrizing cylinder, and it takes values in $[0, \pi]$ on the e ' version. To describe the coordinate transformation from (r, τ) coordinates to the standard coordinates, it is necessary to fix an \mathbb{R} -valued anti-derivative, \hat{v} , for dv that is defined near 0 in $\mathbb{R}/(2\pi\mathbb{Z})$ and vanishes at 0. Thus, v is the $\text{mod}(2\pi)$ reduction of \hat{v} .

With the preceding understood, here is the coordinate transformation between the (σ, \hat{v}) and the variables (r, τ) for e' 's parametrizing cylinder:

$$(4-11) \quad \begin{aligned} \sigma &= \theta_* + \varepsilon r \sin(\tau). \\ \hat{v} &= \left(1 - \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)}\right) \tau + \frac{1}{\alpha_{Q_e}(\sigma)} r \cos(\tau). \end{aligned}$$

In this regard, keep in mind that $\tau \in [-\pi, 0]$. To verify that (r, τ) are bona fide coordinates, use Taylor's theorem with remainder while referring to (3-1) and the first point in (4-11) to write $\hat{v} = r(c_0 \varepsilon \tau \sin(\tau) + \cos(\tau)) + 0(\rho r)$ with c_0 a positive constant that is determined by θ_o . In particular, the Jacobian of the map $(r, \tau) \rightarrow (\sigma, v)$ therefore has the form $-r(1 - c_0 \varepsilon \sin^2(\tau)) + 0(\rho r)$, and this is negative if ε is small and ρ is very small.

Meanwhile, the coordinate transformation between the (σ, \hat{v}) and (r, τ) coordinates for the e' version of the parametrizing cylinder is given as follows:

$$(4-12) \quad \begin{aligned} \sigma &= \theta_* + \varepsilon r \sin(\tau). \\ \hat{v} &= \left(\frac{\alpha_{Q_e}(\sigma)}{\alpha_{Q_{e'}}(\sigma)} - 1\right) \tau + \frac{1}{\alpha_{Q_{e'}}(\sigma)} r \cos(\tau). \end{aligned}$$

In this case, the coordinate τ ranges in $[0, \pi]$.

The digression continues with the introduction of a certain function, v_* , a function that is defined where $r > \frac{1}{8}\rho$ in the $|\sigma - \theta_o| < 3\rho$ portion of $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$. To define v_* , it is necessary to view v as taking values in $[0, 2\pi]$. This understood, set

$$(4-13) \quad v_* = \left((1 - \beta') + \frac{\alpha_{Q_e}(\sigma)}{\alpha_{Q_{e'}}(\sigma)} \beta' \right) v.$$

As its final task, this digression introduces two versions of a function, β_* , one on e' 's version of the parametrizing cylinder and the other on the e' version. In both cases, β_* is defined to be zero on the complement of the set where r is defined and less than 3ρ . Meanwhile, where $r \leq 3\rho$, this function is set equal to $\beta(\frac{1}{\rho}r)$.

With the digression now over, what follows are the rules for extending the definition of (a_e, w_e) to the $\sigma > \theta_o - 2\rho$ portion of e' 's version of the parametrizing cylinder. With τ viewed as taking values in $[-\pi, 0]$, set

$$(4-14) \quad \begin{aligned} a_e &= -\beta_* \ln(r) + a_0 + \varepsilon \left(\beta_* + (1 - \beta_*) \cos(v_*) \right). \\ w_e &= -\varepsilon(1 - \beta_*) \sin(v_*) \\ &\quad + x_0 \beta' \left(\frac{1}{\alpha_{Q_e}} \beta_* \left(\tau - \frac{1}{2\alpha_{Q_{e'}}} r \cos(\tau) \right) - \frac{1}{2\alpha_{Q_e}} (1 - \beta_*) v_* \right). \end{aligned}$$

Here, one must view v and σ as functions of r and τ where $r < 3\rho$. As for $(a_{e'}, w_{e'})$, view τ as taking values in $[0, \pi]$ and set

$$(4-15) \quad \begin{aligned} a_{e'} &= -\beta_* \ln(r) + a_0 + \varepsilon \left(\beta_* + (1 - \beta_*) \cos(v) \right). \\ w_{e'} &= -\varepsilon(1 - \beta_*) \sin(v) \\ &\quad + x_0 \beta' \left(\frac{1}{\alpha_{Q_{e'}}} \beta_* \left(\tau + \frac{1}{2\alpha_{Q_e}} r \cos(\tau) \right) + \frac{1}{2\alpha_{Q_{e'}}} (1 - \beta_*) v \right). \end{aligned}$$

In this last equation, v and σ must again be viewed as functions of r and τ where $r < 3\rho$ with v taking values in $[0, 2\pi]$.

As is proved below, these extensions have the following three special properties: First, the union of the images in the $|\theta - \theta_o| < 3\rho$ portion of $\mathbb{R} \times (S^1 \times S^2)$ of the e and e' parametrizing cylinders fit along the $\theta = \theta_o$ locus so as to define a smooth, properly immersed, thrice punctured sphere. Second, the closure of this thrice punctured sphere in $\mathbb{R} \times (S^1 \times S^2)$ has the asymptotics as dictated by Definition 3.2 of a $\lim_{|s| \rightarrow \infty} \theta = \theta_o$, concave side end of a J -pseudoholomorphic subvariety. Third, this thrice punctured sphere has a finite number of singular points, all are transversal double points, and all have $+1$ local intersection number.

The remainder of this subsection is divided into four parts, with the first two containing the proofs of the first two of the preceding assertions. The final two parts contain the proof of the third assertion.

Part 1 This part addresses the manner in which the images of the two parametrizing cylinders match along the $\theta = \theta_o$ locus in $\mathbb{R} \times (S^1 \times S^2)$. In this regard, it follows from (2–25) that these images define a smoothly immersed surface near the $\theta = \theta_o$ locus provided that the following is true: Let $v_1 \in (0, 2\pi)$. Then, there exists an integer pair, $N = (n, n')$, and extensions of the definitions of (a_e, w_e) and $(a_{e'}, w_{e'})$ to some neighborhood in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$ of (θ_*, v_1) so that

$$(4-16) \quad \begin{aligned} a_e \left(\sigma, \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} v + 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)} \right) &= a_{e'}(\sigma, v), \\ w_e \left(\sigma, \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} v + 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)} \right) \\ &= w_{e'}(\sigma, v) + \frac{1}{\alpha_{Q_e}(\sigma)} (q_{e'}' q_{e'} - q_e q_{e'}') v + \frac{2\pi}{\alpha_{Q_e}(\sigma)} (q_{e'}' n - q_e n'). \end{aligned}$$

To verify that (4–16) holds, let $U \subset (0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$ denote the complement of the point $(\theta_o, 0)$. Now observe that the formulae in (4–14) and (4–15) make perfectly good sense on some neighborhood in U of U 's intersection with the $\sigma = \theta_o$ circle. In particular, where $r \leq 3\rho$, the formula in (4–14) makes good sense where $\tau \in (-\frac{3\pi}{2}, \frac{\pi}{2})$ and that in (4–15) makes good sense where $\tau \in (-\frac{\pi}{2}, \frac{3\pi}{2})$. Meanwhile, where r is either undefined or greater than 2ρ , both formula make good sense where $|\sigma - \theta_o| < \rho^4$. Use these extensions of (4–14) and (4–15) to provide extensions for use in (4–16) of the domains of (a_e, w_e) and $(a_{e'}, w_{e'})$.

Consider now (4–16) at a point (θ_o, v_1) where both $|\sigma - \theta_o| < \rho^4$ and r is either undefined or greater than 2ρ . In this case, the $N = (0, 0)$ version of (4–16) holds by virtue of two facts: First, $\beta' = 1$ where $|\sigma - \theta_o| < \rho^4$ and thus

$$(4-17) \quad v_* \left(\frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} v \right) = v.$$

Second, (3–1) equates $-x_0$ with $q_{e'}' q_{e'} - q_e q_{e'}'$.

Consider next the story for a point (θ_o, v_1) where $r < 3\rho$ that sits on the $\cos(\tau) = 1$ ray. To begin, note that (4–15) and (4–16) require the evaluation of $(a_{e'}, w_{e'})$ at parameters (r, τ) that give σ in (4–12), and that gives the \mathbb{R} -valued parameter \hat{v} as it ranges over some interval of small length in $(0, 2\pi)$ that contains v_1 . In particular, let v' denote a point in such an interval and take σ to lie within ρ^4 of θ_o . Let (r', τ') denote the values for (r, τ) that give (σ, v') .

Now suppose that $N = (0, 0)$ again so that the pair (a_e, w_e) in (4-16) are to be evaluated at parameters $r = r_*$ and $\tau = \tau_*$ that give the value of σ used for $(a_{e'}, w_{e'})$ and give v' for the \mathbb{R} -valued parameter \hat{v} . As such, (4-11) and (4-12) require that (r_*, τ_*) is determined by (r', τ') via the identities

$$\begin{aligned}
 (4-18) \quad & r_* \sin(\tau_*) = r' \sin(\tau'), \\
 & \left(1 - \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)}\right) \tau_* + \frac{1}{\alpha_{Q_e}(\sigma)} r_* \cos(\tau_*) \\
 & = \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} \left[\left(\frac{\alpha_{Q_e}(\sigma)}{\alpha_{Q_{e'}}(\sigma)} - 1\right) \tau' + \frac{1}{\alpha_{Q_{e'}}(\sigma)} r' \cos(\tau') \right]
 \end{aligned}$$

In particular, if ε is small and ρ very small, then (4-19) requires $r_* = r'$ and $\tau_* = \tau'$.

Granted this last conclusion, the top equality in (4-16) now follows from (4-14) and (4-15). Meanwhile, since $\beta' = 1$ at the given value of σ , the lower equality holds provided that

$$\begin{aligned}
 (4-19) \quad & x_0 \beta_* \left[\left(\frac{1}{\alpha_{Q_e}} - \frac{1}{\alpha_{Q_{e'}}}\right) \tau' - \frac{1}{\alpha_{Q_e} \alpha_{Q_{e'}}} r \cos(\tau') \right] - x_0 (1 - \beta_*) \frac{1}{2\alpha_{Q_e}} (2v') \\
 & = \frac{1}{\alpha_{Q_e}} (-x_0) v';
 \end{aligned}$$

and such is the case by virtue of (4-12).

The last case to consider is that where the point (θ_o, v_1) sits where $r < 3\rho$ on the $\cos(\tau) = -1$ ray. Supposing that σ is within ρ^4 of θ_o and v' is very close to v_1 , then the pair $(a_{e'}, w_{e'})$ in (4-16) must be evaluated at parameters values r' for r and τ' for τ that give the point (σ, v') via (4-12). In this regard, note that $\tau' \sim \pi$.

Now suppose that $N = Q_{e'}$ in (4-16). This being the case, then the pair (a_e, w_e) are to be evaluated at parameters $r = r_*$ and $\tau = \tau_*$ with $\tau_* \sim -\pi$ that give the value of σ used for $(a_{e'}, w_{e'})$, but now give the \mathbb{R} -valued parameter \hat{v} that obeys

$$(4-20) \quad \hat{v} = \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} (v' + 2\pi) \pmod{2\pi\mathbb{Z}}.$$

As such, (4–11) and (4–12) require that (r_*, τ_*) is determined by (r', τ') via the identities

$$\begin{aligned}
 (4-21) \quad & r_* \sin(\tau_*) = r' \sin(\tau'), \\
 & \left(1 - \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)}\right) \tau_* + \frac{1}{\alpha_{Q_e}(\sigma)} r_* \cos(\tau_*) = \\
 & \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} \left[\left(\frac{\alpha_{Q_e}(\sigma)}{\alpha_{Q_{e'}}(\sigma)} - 1 \right) \tau' + \frac{1}{\alpha_{Q_{e'}}(\sigma)} r' \cos(\tau') \right] \\
 & \qquad \qquad \qquad + \left(\frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} - 1 \right) 2\pi.
 \end{aligned}$$

When ε is small and ρ is very small, then (4–21) requires $r_* = r'$ and $\tau_* = \tau' - 2\pi$.

Granted the preceding, the top equality in (4–16) again follows from (4–14) and (4–15) straight away. Meanwhile, the lower equality holds provided that

$$\begin{aligned}
 (4-22) \quad & x_0 \beta_* \left[\left(\frac{1}{\alpha_{Q_e}} - \frac{1}{\alpha_{Q_{e'}}} \right) \tau' - \frac{1}{\alpha_{Q_e} \alpha_{Q_{e'}}} r \cos(\tau') \right] \\
 & - 2\pi x_0 \beta_* \frac{1}{\alpha_{Q_e}} - x_0 (1 - \beta_*) \frac{1}{2\alpha_{Q_e}} (2v' + 4\pi) \\
 & \qquad \qquad \qquad = \frac{1}{\alpha_{Q_e}} (-x_0)(v' + 2\pi).
 \end{aligned}$$

And this last equation does indeed hold; this is another consequence of (4–12).

Part 2 This part discusses the image of the points in the respective e and e' parametrizing cylinders that are close to the $(\sigma = \theta_o, v = 0)$ point. The result is a verification that the images of this portion of the two parametrizing cylinders fit together so as to define a submanifold that has the required large $|s|$ asymptotics as dictated by Definition 3.2.

To start the discussion, note that the respective images of the $\theta \geq \theta_o$ and $\theta \leq \theta_o$ parts of the complement of $(\theta_o, 0)$ in a closed, small radius disk about this point define a properly immersed surface with boundary in $\mathbb{R} \times (S^1 \times S^2)$. In this regard, let e denote the union of the respective images in $\mathbb{R} \times (S^1 \times S^2)$ of the portions of the e and e' versions of the parametrizing cylinder where r is defined and where $r \leq \rho^4$.

Here is the first observation: With both ρ and ε chosen to be very small, then the map from $(0, 3\rho] \times \mathbb{R} / (2\pi\mathbb{Z})$ to $\mathbb{R} \times (S^1 \times S^2)$ that uses the appropriate pair of (4–11) and (4–14) or (4–12) and (4–15) defines a smooth, proper map. As is explained some paragraphs hence, this map embeds the subset of the disk where $\beta_* = \beta' = 1$. In particular, E is embedded.

To verify the second requirement, note that the 1-form ds pulls back to the (r, τ) disk via the parametrizing map as $-\frac{1}{r}dr$ where $\beta_* = 1$. Thus s has no critical points on E . Moreover, as $r = e^{-s}$, so θ has a unique $s \rightarrow \infty$ limit on E . The condition on the pull-back of the contact 1-form in (1-1) is considered momentarily. The fact that p_0 and p_0' give the respective integrals of the pull-backs of $\frac{1}{2\pi}dt$ and $\frac{1}{2\pi}d\varphi$ follow from the relationship between Q_e and $Q_{e'}$ given in (3-1).

To see about the requirement for the 1-form $p_0d\varphi - p_0'dt$, first use (4-11) and (4-14) where $\beta_* = \beta' = 1$ and $\tau \in [-\pi, 0]$ to write

$$(4-23) \quad p_0\varphi - p_0't = x_0 \left[\left(1 - \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} \right) \tau + \frac{r}{\alpha_{Q_e}(\sigma)} \cos(\tau) \right] - \alpha_{P_o} \frac{1}{\alpha_{Q_e}} x_0 \left[\tau - \frac{1}{2\alpha_{Q_{e'}}} r \cos(\tau) \right] \pmod{2\pi\mathbb{Z}}.$$

Now use (3-1) to conclude that

$$(4-24) \quad p_0\varphi - p_0't = x_0 \frac{1}{2} \left(\frac{1}{\alpha_{Q_e}(\sigma)} + \frac{1}{\alpha_{Q_{e'}}(\sigma)} \right) r \cos(\tau) \pmod{2\pi\mathbb{Z}}$$

where $\beta_* = \beta' = 1$ and $\tau \in [-\pi, 0]$. Meanwhile, the same expression for $p_0\varphi - p_0't$ appears when (4-12) and (4-15) are used where $\beta_* = \beta' = 1$ and $\tau \in [0, \pi]$. Thus, the right hand side of (4-24) without the ‘ $\pmod{2\pi\mathbb{Z}}$ ’ proviso provides an anti-derivative for $p_0d\varphi - p_0'dt$ with a unique $s \rightarrow \infty$ limit on E .

Return next to the question of the contact form in (1-1). The fact that it has nowhere zero pull-back at large $|s|$ on e can be readily deduced from the following three facts: First, the vectors $(p_0', -p_0)$ and $((1 - 3 \cos^2 \theta_o), \sqrt{6} \cos \theta_o \sin^2 \theta_o)$ are linearly independent in \mathbb{R}^2 ; this a consequence of the θ_o version of (1-7). Second, (4-24) implies that the 1-form $p_0d\varphi - p_0'dt$ is $o(e^{-s})$ on the large s slices of E . Finally, the $(q, q') = Q_e$ and $(q, q') = Q_{e'}$ versions of $qd\varphi - q'dt$ differ from $x_0d\tau$ by $o(e^{-s})$ on the respective portions of the large s slices in E where $\tau \in [-\pi, 0]$ and where $\tau \in [0, \pi]$.

Here is the promised explanation as to why the punctured (r, τ) disk is embedded where $\beta_* = \beta' = 1$. To see that such is the case, note first that two points are mapped to the same point only if the corresponding pull-backs of θ agree. Moreover, the corresponding pull-backs of a chosen \mathbb{R} lift of the right hand side of (4-24) must also agree. These two requirements can be met only if the two points are one and the same.

The final issue concerns the size of the projection \prod_J as defined in the last line of Definition 3.2's second requirement. The fact is that this projection and its covariant

derivative are both $o(e^{-s})$ at large s on E , this by virtue of (4–24) and the formulae for θ in the top lines in (4–11) and (4–12).

Part 3 This and the remaining part of the discussion in this subsection analyze the singular points of the $|\theta - \theta_0| < 3\rho$ portion of the closure of $K_e \cup K_{e'}$. In this regard, it follows from what has been said so far that this closure is an immersed surface with compact singular set. Indeed, this happens because both the e and e' parametrizations extend across the $\sigma = \theta_0$ circle save at the missing point where $\sigma = \theta_0$ and $v = 0$. With the preceding understood, the task at hand is to verify that there are but a finite number of singular points, all regular double points and all with $+1$ local intersection number.

To start the story, remark that points in K_e and $K_{e'}$ are disjoint as they have distinct θ values. Thus, a pair of points (σ, v) and (σ', v') in the extended domain of either the e or e' parametrizing maps are sent to the same point in $\mathbb{R} \times (S^1 \times S^2)$ if and only if $\sigma = \sigma'$ and the respective e or e' versions of the conditions in (4–2) are satisfied. This noted, the discussions that follows in this Part 3 and in Part 4 focus exclusively on points in the closure of K_e ; thus points in the image of the complement of $(\theta_0, 0)$ in e' 's version of the closed parametrizing cylinder. The analysis for $K_{e'}$ is very much the same and so omitted.

The rest of the story from this Part 3 is summarized by the following lemma:

Lemma 4.2 *There exists ε_0 and given $\varepsilon \in (0, \varepsilon_0)$, there exist positive constants δ and $\rho_* \ll \varepsilon$ such that when $\rho < \rho_*$, then the closure of K_e is smooth near any point with an inverse image in the portion of the closed parametrizing cylinder that lie where $\beta_* > 0$, or where $\beta' < \delta$, or where $\beta' > 1 - \delta$.*

The proof of this lemma occupies the remainder of Part 3.

Proof of Lemma 4.2 The proof is facilitated by introducing a relatively prime pair of integers, $Z \equiv (z, z')$, such that $zq_{e'} - z'q_e = m$ where m denotes the greatest common divisor of $(q_e, q_{e'})$. With Z so specified, any given integer pair $N = (n, n')$ can be written as $N = \hat{x}(z, z') + \frac{\hat{y}}{m}(q_e, q_{e'})$ with \hat{x} and \hat{y} integers. Writing N in this way makes the third point in N' 's version of (4–2) into the condition in (4–3) and the second point into the condition in (4–4). Note that \hat{x} does not depend on the choice for Z but \hat{y} does. In any event, only values for \hat{y} that lie in $\{0, \dots, m-1\}$ need be considered.

The rest of the proof is broken into four steps.

Step 1 This step derives the lower bound for β' on the inverse image of a singularity. For this purpose, note that with ε small in (4–14) and $\beta' \leq \varepsilon^2$, then $|w_e| < 2\varepsilon$ and

so only the case that $\hat{x} = 0$ case can possibly arise in (4-3). However, this case then requires $\hat{y} = 0$ as well since the form of a_e and w_e in (4-14) precludes any other $\hat{x} = 0$ solutions to both (4-3) and (4-4). The argument here is essentially the same as one given in Part 1, above.

Step 2 This step proves the following assertion: If ε is small and then ρ is very small, the closure of K_e is smooth at the image of any point in the closed parametrizing cylinder that lies where r is defined and less than $\frac{3}{4}\rho$.

To start the proof, suppose that ζ and ζ' are two points in the closed parametrizing cylinder that map to the same point in $\mathbb{R} \times (S^1 \times S^2)$ and are such that one lies in the indicated region. As the respective values for a_e must agree at the two points, (4-14) demands that both points lie where $r < \rho$ when ε is small. In particular, $\beta_* = 1$ at both points, so they both lie where r is defined, and their respective r coordinates must agree. As the respective values of σ also agree at the two points, (4-11) demands that their respective τ coordinates either agree or are interchanged by the involution of $[-\pi, 0]$ that sends τ to $-\pi - \tau$. The latter must be the case if the two points are distinct.

Meanwhile, (4-3) and (4-4) requires that the respective $\mathbb{R}/(2\pi\mathbb{Z})$ coordinates of the two points are related by

$$(4-25) \quad v(\zeta') = v(\zeta) - 2\pi\hat{x} \frac{\alpha_Z(\sigma)}{\alpha_{Q_e}(\sigma)} - 2\pi \frac{\hat{y}}{m} \pmod{2\pi}.$$

with \hat{x} and \hat{y} integers and with $Z = (z, z')$ as in (4-3). To see what this implies, note that as $\beta' \neq 0$, it follows that $|\sigma - \theta_o| < 2\rho^4$ and so

$$(4-26) \quad \frac{\alpha_Z(\sigma)}{\alpha_{Q_e}(\sigma)} = \frac{z'p_o - zp'_o}{x_o} + o(\rho^4),$$

where the term that is indicated by $o(\rho^4)$ is bounded by $\kappa\rho^4$ where κ depends only on Q_e and P_0 . As the ratio on the right hand side is a rational number with denominator no larger than x_0 , the equality in (4-25) demands that

$$(4-27) \quad v(\zeta') = v(\zeta) + o(\rho^4) \pmod{2\pi}$$

where the term $o(\rho^4)$ has the same significance as in (4-26).

To continue, note next that (4-27) is consistent with the lower line in (4-11) only in the case that the value of $|r \cos(\tau)|$ at the two points is bounded by $\kappa\rho^4$ where κ again depends only on P_0 and Q_e . On the other hand, as β' is neither 0 nor 1, the top line in (4-11) requires that $|r \sin(\tau)| > \varepsilon^{-1}\rho^4$. In particular, $r \geq \varepsilon^{-1}\rho^4$. This then means that the value of $|\cos(\tau)|$ at the two points is bounded by $\kappa\varepsilon$. As such, the values of τ

at these points must have the form $\tau = \frac{\pi}{2} \pm \varepsilon \cdot \zeta$, where $\zeta > 0$ is bounded solely in terms of P_0 and Q_e . This understood, use of the lower line in (4–14) will establish that the respective values of w_e at the two points differ by an amount that is bounded by $\kappa\varepsilon$, where κ is as before, a constant that depends only on P_0 and Q_e . However, if such is the case, then small ε is consistent with (4–3) only if $\hat{x} = 0$. This and (4–27) require that $\hat{y} = 0$ also, and so $\zeta = \zeta'$.

Step 3 As is argued subsequently in this step, if ε is small and ρ is very small, then the closure of K_e is smooth near any point with an inverse image that lies where r is defined and has value in $(\frac{1}{2}\rho, 3\rho)$. To see why this is true, take ζ and ζ' to be points in the closed parametrizing cylinder that are mapped to the same point in $\mathbb{R} \times (S^1 \times S^2)$, are such that $\beta' > 0$, and that one lies where r is defined and is in the indicated range. In this case, the top line of (4–14) requires that the respective values of $\beta_* \ln(r)$ differ by less than ε at the two points. Thus, both points lie where r is greater than $\frac{1}{4}\rho$. Furthermore, as one of them lies where $r \leq 3\rho$, the top line in (4–14) requires that $1 - \kappa^{-1}\rho^2 > \cos(v_*) > 1 - \kappa\rho^2$; at both; here $\kappa \geq 1$ is a constant that depends only on P_0 and Q_e .

Meanwhile, (4–3) and (4–4) require (4–25), and as $\beta' \neq 0$, so (4–26) and (4–27) hold here too. However, if such is the case, then both lie where $\sin(v_*) \geq \kappa^{-1}\rho$ or both lie where $\sin(v_*) < -\kappa^{-1}\rho$. Here again, $\kappa \geq 1$ depends only on P_0 and Q_e when ρ is small. As a result, (4–11) and (4–14) insure that the respective values of w_e at the two points differ by no more than $\kappa'\rho$ where κ' is again a constant that depends only on P_0 and Q_e . Granted this, then only the $\hat{x} = 0$ case of (4–3) can arise when ρ is small, and this then requires that $\hat{y} = 0 \pmod{m}$ also. Thus, $v(\zeta') = v(\zeta)$ and the points are one and the same.

Step 4 This step proves the following: If ε is small, there then exists $\delta' > 0$ with the following significance: When ρ is very small, the parametrizing map embeds the portion of e' 's parametrizing cylinder where $\beta' > 1 - \delta'$ and $\beta_* = 0$.

To start the proof, suppose that (σ, v) and (σ, v') are two points that lie in the indicated portion of the closed parametrizing cylinder. Suppose, in addition, that $v \neq v'$ and that these two points are sent to the same point in $\mathbb{R} \times (S^1 \times S^2)$. This being the case, the equality between the respective values of a_e at the two points requires that $v' = 2\pi - v \pmod{2\pi\mathbb{Z}}$. As the two points (σ, v) and (σ, v') are distinct, so $v \neq \pi \pmod{2\pi\mathbb{Z}}$. Thus, at the expense of choosing which to call v and which to call v' , as well as a \mathbb{Z} lift of \hat{y} , one can assume that v and v' have respective \mathbb{R} -lifts \hat{v} and

\hat{v}' with $\hat{v} \in (\pi, 2\pi)$ and with $\hat{v}' = 2\pi - \hat{v}$. Meanwhile, (4-3) requires that

$$(4-28) \quad \varepsilon \sin(\hat{v}) - x_0 \frac{1}{2\alpha_{Q_{e'}}} \beta' \hat{v}' = -\varepsilon \sin(\hat{v}) - x_0 \frac{1}{2\alpha_{Q_{e'}}} \beta' \hat{v} - 2\pi m \frac{1}{\alpha_{Q_e}} \hat{x}.$$

To see the implications of these constraints, it proves convenient to introduce

$$(4-29) \quad \wp \equiv -\left(\hat{x} \frac{z' p_o - z p'_o}{x_o} - \frac{\hat{y}}{m}\right).$$

As will now be explained, $\wp \in (0, 1)$. Indeed, as $\beta' \neq 0$, so (4-26) holds and can be used to write

$$(4-30) \quad \hat{v} = \pi(1 + \wp + \hat{x}\hat{u}) \quad \text{and} \quad \hat{v}' = \pi(1 - \wp - \hat{x}\hat{u})$$

where \hat{u} is a function of σ that obeys $|\hat{u}| \leq \kappa\rho^4$ and $|\partial_\sigma \hat{u}| \leq \kappa$ with κ a constant that is determined by P_0 and Q_e .

With \wp as just defined, the constraint in (4-28) can be rearranged to read

$$(4-31) \quad \beta' \wp - 2m \frac{\hat{x}}{x_o} - \frac{1}{\pi} \varepsilon \sin(2\pi\wp) + \hat{x}u = 0,$$

where u is a function of σ that vanishes at θ_o and is determined a priori by P_0 and Q_e . In particular $|u| \leq \kappa\rho^4$ where $\beta' \neq 0$ and $|\partial_\sigma u| \leq \kappa$ where κ is a constant that is determined a priori by P_0 and Q_e . As will now be explained, the desired bound for $(1 - \beta')$ follows from this last equation. To see why, note that by virtue of (4-30), small ρ insures that the distance between \hat{v} from either π or 2π is bounded from zero by a constant that depends only on P_0 and Q_e . In particular, $|\sin(\hat{v})|$ is bounded away from zero by such a positive constant. As $|\sin(\hat{v})|$ is uniformly bounded away from zero, and as \wp is a fraction between 0 and 1 whose denominator is $m x_0$, it follows that there are no (\hat{x}, \hat{y}) versions of (4-31) when ε is small, $(1 - \beta') < \varepsilon^2$, and ρ is very small. \square

Part 4 This part of the discussion identifies the singular points and verifies that they have positive local intersection number. In this regard, remember Lemma 4.2 which states that the inverse images of the singular points lie where $\beta_* = 0$ and $\beta' \in (\delta, (1 - \delta))$.

To begin, fix \hat{x} and \hat{y} so that $\wp \in (0, 1)$ as defined in (4-29). Then view the small ε and very small ρ version of (4-31) as an equation for $\sigma \in (\theta_0 - 3\rho, \theta_o)$. Because the values of β' range over $[0, 1]$, this equation has a solution only if $0 \leq m\hat{x} \leq \frac{1}{2}x_0$. Moreover, if a solution exists, then it is unique. To explain, remember that with ε fixed and small, Lemma 4.2 provides a constant $\delta \equiv \delta(\varepsilon) \in (0, 1)$ such that a solution

to (4–31) must occur at a value for σ where $\delta < \beta' < (1 - \delta)$. Thus, there exists a constant, $b \equiv b(\varepsilon) \in (0, 1)$ such that

$$(4-32) \quad \partial_\sigma \beta' > b \frac{1}{\rho^4}$$

at any value of σ where (4–31) holds. This last fact implies the uniqueness of any solution to (4–31) in the case that ρ is very small.

The next point to make is that the solutions to the various (\hat{x}, \wp) versions of (4–31) result in singularities of K_e of the simplest sort: Each singular point is the center of a small radius ball in $\mathbb{R} \times (S^1 \times S^2)$ that intersects K_e as the union of two embedded disks meeting only at the origin. Here is why: According to (4–30), $\hat{v} = \pi(1 + \wp) + o(\rho^4)$ and as a consequence, when ρ very small, (4–4) guarantees that any pair of versions of (4–31) with different values for \wp will yield disjoint singular points in K_e . Meanwhile, two versions of (4–31) that are defined using the same choice for \wp but with different choices for \hat{x} give corresponding singular points in K_e at distinct θ values.

With the singular points identified, the next task is to verify that each self intersection is transversal and has positive intersection number. This task requires a suitable expression for the push-forward by the parametrizing map of the vector fields ∂_σ and ∂_v . Such a formula is given in (4–33) below. In this regard, the following notation is used: The pull-back via the parametrizing map of $(1 - 3 \cos^2 \theta)$ is denoted as c , and that of $\sqrt{6} \cos \theta$ as c' . Note also that the push-forwards of ∂_σ and ∂_v are not notationally distinguished from the originals. In addition, the label e on (a_e, w_e) is suppressed so that a subscript on the resulting a or w can be used to indicate the partial derivatives in the direction labeled by the subscript. The label e is also suppressed so that the integer pair Q_e appears as (q, q') . With this notation set, here are the promised formulae for the push-forwards of ∂_σ and ∂_v :

$$(4-33) \quad \begin{aligned} \partial_\sigma &= (cw)_\sigma \partial_t + (c'w)_\sigma \partial_\varphi + a_\sigma \partial_s + \partial_\theta. \\ \partial_v &= (q + cw_v) \partial_t + (q' + c'w_v) \partial_\varphi + a_v \partial_s. \end{aligned}$$

Now let (a, w) and (a', w') denote the respective versions of the parametrizing functions that come from the two sheets that are involved at the given intersection point. Likewise, use $(\partial_\sigma, \partial_v)$ and $(\partial_{\sigma'}, \partial_{v'})$ to denote the corresponding versions of (4–33). The convention used below takes the unprimed pair as the image via the parametrizing map of the point (σ, \hat{v}) with \hat{v} as in (4–30). Meanwhile, the primed pair is the image of (σ, \hat{v}') with \hat{v}' also from (4–30).

To establish transversality for the self-intersection point and to obtain the local intersection number, first write $\partial_\sigma \wedge \partial_v \wedge \partial_{\sigma'} \wedge \partial_{v'}$ as $\tau(\partial_s \wedge \partial_t \wedge \partial_\theta \wedge \partial_\varphi)$ with $\tau \in \mathbb{R}$.

As demonstrated below, τ is non-zero; thus, the intersection is transversal. Granted this, the sign of τ is the sign to take for the local sign of the self intersection point. With regard to the upcoming expression for τ , note that (4–14) finds $a_\sigma = a'_\sigma = 0$ and $w_\nu = w'_\nu$ at the intersection point. Here is τ :

$$(4-34) \quad \tau = -(a_\nu - a'_\nu) \left[(c(w - w'))_\sigma (q' + c'w_\nu) - (c'(w - w'))_\sigma (q + cw_\nu) \right].$$

To evaluate (4–34), use (4–14) to deduce that when ρ is very small, then

$$(4-35) \quad \begin{aligned} a_\nu - a'_\nu &= 2\varepsilon \sin(\pi(\wp + \hat{x}\hat{u})) > 0, \\ w_\sigma - w'_\sigma &= -\pi\wp x_0 \frac{1}{\alpha_{Q_e}(\sigma)} \partial_\sigma \beta' + o(1), \end{aligned}$$

where the term denoted by $o(1)$ is bounded by a constant that depends only on P_0 and Q_e . Granted (4–35) and granted that $w - w'$ is of the order of unity, any very small ρ version of (4–34) has the form $\tau = 2\pi\varepsilon\wp \sin(\pi\wp)x_0 \partial_\sigma \beta' + o(1)$, where the term designated as $o(1)$ is again bounded by a constant that depends only on P_0 and Q_e . As (4–32) guarantees that $\partial_\sigma \beta'$ is very large when ρ is very small, so $\tau > 0$ as required for a transversal intersection with positive local intersection number.

4.D Parametrizations near boundary circles with a trivalent vertex label

In this subsection, o denotes a trivalent vertex in T while e , e' and e'' denote the three incident edges to o . In this regard, it is assumed here that only one of these edges labels a cylinder in $\mathbb{R} \times (S^1 \times S^2)$ where $\theta < \theta_o$. The discussion for the case when $\theta > \theta_o$ on only one of K_e , $K_{e'}$ and $K_{e''}$ is not presented since it is identical to the discussion that follows but for some obvious cosmetic alterations. This understood, the edges e , e' and e'' are distinguished as follows: First, $\theta < \theta_o$ only on K_e . Second, in the case that $Q_{e'}$ and $Q_{e''}$ are not proportional (and thus not proportional to Q_e), take e' so as to make

$$(4-36) \quad q_{e'}q_{e''} - q_{e'}'q_{e''} < 0.$$

The story starts with a preliminary digression to set the stage. To begin the digression, assume that $\rho_{e1} = \rho_{e'0} = \rho_{e''0}$ in what follows, and use ρ denote any of the three. Here again, take $\rho \ll \delta$. Now the function β' refers to the version in (4–5) with this same value for ρ and with θ_* set to equal θ_o .

Require the e , e' and e'' versions of ε in (4–1) to be constant and much less than δ near the respective circles where $|\sigma - \theta_o| = 2\rho$, and require that the three constants agree. Likewise, require that a_e^0 , $a_{e'}^0$ and $a_{e''}^0$ are constant near these circles; the values for these constants are specified below. In addition, require that w_e^0 , $w_{e'}^0$ and $w_{e''}^0$

vanish near the respective circles where $|\sigma - \theta_o| = 2\rho$. Finally, require that each of v_e^0 , $v_{e'}^0$, and $v_{e''}^0$ is constant where $|\theta - \theta_o| \sim 2\rho$; the values of the latter are also specified below.

As σ approaches θ_o , the functional form of a and w must be modified for each of the three edges to accommodate the first three constraints in (3–3). To ease the proliferation of subscripts in the subsequent discussion, agree now to use α , α' and α'' to denote the respective $Q = Q_e$, $Q_{e'}$ and $Q_{e''}$ versions of $\alpha_Q(\sigma)$. Also, to avoid possible confusion, the coordinate v is used only to denote the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate on the parametrizing cylinder for K_e . The corresponding coordinates for the $K_{e'}$ and $K_{e''}$ cylinders are denoted below by v' and v'' . Finally, to avoid a proliferation of primes, when $N = (n, n')$ and $K = (k, k')$ are pairs of integers (or real numbers), then $[N, K]$ is used below to denote $nk' - n'k$. For example, the distinction between the edges e' and e'' that appears in (4–36) can now be written as $[Q_{e'}, Q_{e''] < 0$.

The first step to defining the three versions of (a, w) where σ is near θ_o is to specify each near the point or points on the $\sigma = \theta_o$ boundary of its parametrizing cylinder that will map to the point in the mutual intersection of the closures of K_e , $K_{e'}$ and $K_{e''}$. In this regard, there are two such ‘singular points’ on the $\sigma = \theta_o$ circle in the e version of the parametrizing cylinder, and one each on the $\sigma = \theta_o$ circle in the e' and e'' versions of the parametrizing cylinder. Here, the singular points on the e version of the $\sigma = \theta_o$ circle have $\mathbb{R}/(2\pi\mathbb{Z})$ coordinates 0 and $2\pi\frac{\alpha'}{\alpha}$, the singular point on the e' version of this circle has $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate 0, while that on the e'' circle has $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate $2\pi\frac{\alpha'}{\alpha''}$.

To define the versions of (a, w) near these singular points, let (x, y) denote Cartesian coordinates for \mathbb{R}^2 . Divide a small radius, open disk centered at the origin in the (x, y) plane into four open sets, the four components of the complement of the locus where $x^2 = y^2$. The closures of the two components where $|y| > |x|$ will be identified with respective open neighborhoods of the points $v = 0$ and $v = 2\pi\frac{\alpha'}{\alpha}$ on the $\sigma = \theta_o$ circle in the parametrizing cylinder for K_e . Meanwhile, the closure of the $x > |y|$ component will be identified with an open neighborhood of the $v' = 0$ point on the the $\sigma = \theta_o$ circle in the parametrizing cylinder for $K_{e'}$. Finally, the closure of the $x < -|y|$ portion will be identified with an open neighborhood of the $v'' = 2\pi\frac{\alpha'}{\alpha''}$ point on the $\sigma = \theta_o$ circle in the parametrizing cylinder for $K_{e''}$. These identification are made as follows: In all cases,

$$(4-37) \quad \sigma = \theta_o + x^2 - y^2.$$

Thus, $\sigma = \theta_o$ where $x^2 = y^2$. Meanwhile,

$$(4-38) \quad \begin{aligned} v &= \frac{1}{\alpha}xy \quad \text{where } y > |x|, & v &= \frac{1}{\alpha}xy + 2\pi \frac{\alpha'}{\alpha} \quad \text{where } y < -|x|, \\ v' &= \frac{1}{\alpha'}xy \quad \text{where } x > |y|, & v'' &= \frac{1}{\alpha''}xy + 2\pi \frac{\alpha'}{\alpha''} \quad \text{where } x < -|y|. \end{aligned}$$

Here, all assignments are defined modulo $2\pi\mathbb{Z}$.

The next step parametrizes the three versions of (a, w) using the coordinates x and y . This is done as follows: In all cases,

$$(4-39) \quad a \equiv x.$$

Meanwhile,

$$(4-40) \quad \begin{aligned} w_e &\equiv y && \text{where } y > |x|. \\ w_e &\equiv y - 2\pi \frac{[Q_{e'}, Q_{e''}]}{\alpha} && \text{where } y < -|x|. \\ w_{e'} &\equiv y + \frac{[Q_e, Q_{e'}]}{\alpha\alpha'}xy. \\ w_{e''} &\equiv y + \frac{[Q_e, Q_{e''}]}{\alpha\alpha''}xy + 2\pi \frac{[Q_e, Q_{e''}]}{\alpha''}. \end{aligned}$$

Granted (4-37)–(4-40), the e, e' and e'' versions of (3-2) now define a smooth map to $\mathbb{R} \times (S^1 \times S^2)$ from each of the four components of the complement in a small disk about the origin in the x - y plane of the $x^2 - y^2 = 0$ locus. As explained below, the resulting maps fit together across this locus so as to define a smooth, symplectic embedding of the whole of some smaller radius disk into $\mathbb{R} \times (S^1 \times S^2)$.

The preceding formulae write (a_e, w_e) near the points where $v = 0$ and $v = 2\pi \frac{\alpha'}{\alpha}$ on the $\sigma = \theta_o$ boundary circle of K_e 's parametrizing cylinder. They also give $(a_{e'}, w_{e'})$ near the $v' = 0$ point on the $\sigma = \theta_o$ in the parametrizing cylinder for $K_{e'}$, and they give $(a_{e''}, w_{e''})$ near the point where $v'' = 2\pi \frac{\alpha'}{\alpha''}$ on the $\sigma = \theta_o$ boundary of the parametrizing domain for $K_{e''}$. The third step parametrizes the three versions of (a, w) on a neighborhood of the rest of the relevant $\sigma = \theta_o$ circle. Note that in the equations that follow,

$$(4-41) \quad \beta_* \equiv \beta \left(\frac{1}{\rho} (x^2 + y^2) \right).$$

To start, consider (a_e, w_e) . In this case, take

(4-42)

$$a_e = \beta_* x + (1 - \beta_*) \varepsilon \left(\cos \left(v - \pi \frac{\alpha'}{\alpha} \right) - \cos \left(\pi \frac{\alpha'}{\alpha} \right) \right)$$

$$\text{and } w_e = \beta_* y + \beta' \left(\beta \frac{[Q_e, Q_{e'}]}{\alpha \alpha'} x y + \frac{[Q_e, Q_{e'}]}{\alpha'} v \right) - (1 - \beta_*) \varepsilon \sin \left(v - \pi \frac{\alpha'}{\alpha} \right)$$

where $0 \leq v \leq 2\pi \frac{\alpha'}{\alpha}$ and $|\sigma - \theta_o| < \rho$. Where $2\pi \frac{\alpha'}{\alpha} \leq v \leq 2\pi$ and $|\sigma - \theta_o| < \rho$, take

$$(4-43) \quad a_e = \beta_* x + (1 - \beta_*) \varepsilon \left(\cos \left(v - \pi \frac{\alpha'}{\alpha} \right) - \cos \left(\pi \frac{\alpha'}{\alpha} \right) \right)$$

$$\text{and } w_e = \beta_* y + \beta' \left(\beta_* \frac{[Q_e, Q_{e''}]}{\alpha \alpha''} x y - 2\pi \frac{[Q_e, Q_{e''}]}{\alpha''} + \frac{[Q_e, Q_{e''}]}{\alpha''} v \right) - (1 - \beta_*) \varepsilon \sin \left(v - \pi \frac{\alpha'}{\alpha} \right).$$

The e' version of (a, w) requires the introduction of a σ -dependent family of diffeomorphisms of the constant σ circles in the e' version of the parametrizing cylinder. The latter is denoted by $\phi_{e'}$. View it as a map from the parametrizing cylinder to $\mathbb{R}/(2\pi\mathbb{Z})$ with the property that its restriction to any fixed value of σ define a diffeomorphism of $\mathbb{R}/(2\pi\mathbb{Z})$. As such, it need only have the following two properties:

$$(4-44) \quad \begin{aligned} \phi_{e'}(\sigma, v') &= \frac{\alpha'}{\alpha} (v' - \pi) & \text{where both } |\sigma - \theta_o| < 2\rho^4 \text{ and } x^2 + y^2 > \frac{1}{2}\rho. \\ \phi_{e'}(\sigma, v') &= v' & \text{where } |\sigma - \theta_o| \geq 3\rho^4. \end{aligned}$$

An analogous e'' version is denoted by $\phi_{e''}$. The latter should obey

$$(4-45) \quad \begin{aligned} \phi_{e''}(\sigma, v'') &= \frac{\alpha''}{\alpha} v'' - \pi \frac{\alpha'}{\alpha} & \text{where both } |\sigma - \theta_o| < 2\rho^4 \text{ and } x^2 + y^2 > \frac{1}{2}\rho. \\ \phi_{e''}(\sigma, v'') &= v'' & \text{where } |\sigma - \theta_o| \geq 3\rho^4. \end{aligned}$$

With $\phi_{e'}$ and $\phi_{e''}$ as above, what follows are the e' and e'' versions of (a, w) :

$$(4-46) \quad \begin{aligned} a_{e'} &= \beta_* x + (1 - \beta') \rho + (1 - \beta_*) \varepsilon (\cos(\phi_{e'}) - \cos(\phi_{e'}|_{v'=0})). \\ w_{e'} &= \beta_* \left(y + \frac{[Q_e, Q_{e'}]}{\alpha \alpha'} x y \right) - (1 - \beta_*) \varepsilon \sin(\phi_{e'}). \end{aligned}$$

Meanwhile,

$$(4-47) \quad \begin{aligned} a_{e''} &= \beta_* x - (1 - \beta') \rho + (1 - \beta_*) \varepsilon (\cos(\phi_{e''}) - \cos(\phi_{e''}|_{v''=0})). \\ w_{e''} &= \beta_* \left(y + \frac{[Q_e, Q_{e''}]}{\alpha \alpha''} x y \right) - (1 - \beta_*) \varepsilon \sin(\phi_{e''}) + 2\pi \beta' \frac{[Q_e, Q_{e''}]}{\alpha''}. \end{aligned}$$

Three tasks lie ahead. Here is the first: Verify that the closures of K_e , $K_{e'}$ and $K_{e''}$ fit together where they meet along the $\theta = \theta_o$ locus to define a smoothly immersed surface. The second task is to verify that the singular points are of the simplest sort: Each centers a small radius ball that intersects the surface as a pair of embedded disks that meet transversely at a single point. Here is the final task: Verify the positivity of the local intersection number at each singular point.

The three tasks are addressed next.

Task 1 The first step here is to verify that (4–37)–(4–40) in conjunction with (3–2) define maps to $\mathbb{R} \times (S^1 \times S^2)$ that fit together across the $x^2 = y^2$ locus to provide a smooth, symplectic embedding of some small radius disk about the origin in the x – y plane. To be precise here, the resulting map from a small radius disk can be written so that it sends a pair (x, y) with norm $(x^2 + y^2)^{1/2} \ll \rho^4$ to the point with coordinates

$$(4-48) \quad \left(s = x, t = q_e \frac{1}{\alpha(\theta)} xy + (1 - 3 \cos^2 \theta) y, \right. \\ \left. \theta = \theta_o + x^2 - y^2, \varphi = q_{e'} \frac{1}{\alpha(\theta)} xy + \sqrt{6} \cos(\theta) y \right)$$

Here, both t and φ are defined modulo $2\pi\mathbb{Z}$. Note that the differential at the origin of the map in (4–48) sends ∂_x to ∂_s and ∂_y to the Reeb vector field in (1–6). Thus, it symplectically embeds a small radius disk about the origin.

By definition, the map in (4–48) agrees with that given where $y > |x|$ using (3–2) and (4–37)–(4–40). The verification that it agrees with the maps from (3–2) and (4–37)–(4–40) on the other components of the complement of the $x^2 = y^2$ locus is left to the reader except for the following comment: Algebraic manipulations can rewrite the various maps to $\mathbb{R} \times (S^1 \times S^2)$ from the other components so that they appear exactly as depicted in (4–48) but for the addition to t and φ of some integer multiple of 2π .

The next step is to verify that the maps that are defined using (3–2) in conjunction with (4–42)–(4–47) extend to the $\sigma = \theta_o$ circle of each parametrizing cylinder so that the union of the $|\theta - \theta_o| < 3\rho$ portion of the resulting images defines the image via a proper immersion of the complement in S^2 of three pairwise disjoint, closed disks. This is done as follows: Suppose that $v \in (0, 2\pi \frac{\alpha'}{\alpha})$. As written, the formulae in (4–42) extend the definition of (a_e, w_e) to a small radius disk centered on (θ_o, v) in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$. Likewise, when $v \in (2\pi \frac{\alpha'}{\alpha}, 2\pi)$, then (4–43) extend the definition of (a_e, w_e) to a small radius disk centered on (θ_o, v) in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$. Meanwhile, if $v' \in (0, 2\pi)$, then (4–44) together with (4–46) extend the definition of $(a_{e'}, w_{e'})$ to a small radius disk about (θ_o, v') . Finally, if $v'' \neq 2\pi \frac{\alpha'}{\alpha''} \pmod{2\pi}$, then (4–45)

and (4-47) extend the definition of $(a_{e''}, w_{e''})$ to a small radius disk about (θ_o, v'') in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$.

With these extensions understood, suppose again that $v \in (0, 2\pi\frac{\alpha'}{\alpha})$. Then the extended (a_e, w_e) define via (3-2) an immersion into $\mathbb{R} \times (S^1 \times S^2)$ of a small radius disk centered at (θ_o, v) . Set $v' \equiv \frac{\alpha}{\alpha'}v$. Then (4-44) and (4-46) define via (3-2) an immersion into $\mathbb{R} \times (S^1 \times S^2)$ of a small radius disk centered at (θ_o, v') . Now, let $\sigma \in (0, \pi)$ and $v \in \mathbb{R}/2\pi$ be such that $|\theta_o - \sigma| \ll \rho^4$ and such that the extension of (a_e, w_e) is defined at the point (σ, v) and that of $(a_{e'}, w_{e'})$ is defined at $(\sigma, \frac{\alpha}{\alpha'}v)$. It then follows from (4-42), (4-44) and (4-46) that

(4-49)

$$a_e(\sigma, v) = a_{e'}\left(\sigma, \frac{\alpha}{\alpha'}v\right) \quad \text{and} \quad w_e(\sigma, v) = w_{e'}\left(\sigma, \frac{\alpha}{\alpha'}v\right) + \frac{1}{\alpha'}(q_{e'}'q_{e'} - q_e q_{e'})v$$

These last equalities imply that the e and e' versions of the extended maps parametrize open subsets of a single immersed surface, this the union of the $\theta < \theta_o + 3\rho$ portion of the closure of $K_{e'}$, the closure of the portion of K_e in the image of points (σ, v) with $\sigma > \theta_o - 3\rho$ and $v \in (0, 2\pi\frac{\alpha}{\alpha'})$, and the image via (4-48) of a small radius disk centered at the origin in the (x, y) plane.

A similar argument using (4-43), (4-45) and (4-47) proves the analogous statement for the union of the $\theta < \theta_o + 3\rho$ portion of the closure of $K_{e''}$, the closure of the portion of K_e that is in the image of points (σ, v) with $\sigma > \theta_o - 3\rho$ and $v \in (2\pi\frac{\alpha}{\alpha'}, 2\pi)$, and the image via (4-48) of a small radius disk centered at the origin in the (x, y) plane. The details of the latter argument are left to the reader.

Task 2 The task is to describe all of the immersion points. This task is accomplished in five steps.

Step 1 Note that $a_{e'} > 0$ except at $v' = 0 \in \mathbb{R}/(2\pi\mathbb{Z})$ and $x = 0$, while $a_{e''} < 0$ save at $v'' = 2\pi\frac{\alpha'}{\alpha}$ and $x = 0$. However, both of these points correspond to the origin in the x - y coordinate disk. Thus, the closures of the portions of $K_{e'}$ and $K_{e''}$ where $|\sigma - \theta_o| < 3\rho$ are disjoint save for the image of the origin in the x - y coordinate disk. As a consequence, it is sufficient to focus separately on the singularities in the respective closures of K_e , $K_{e'}$ and $K_{e''}$.

As will now be explained, if ε and ρ are small, then the closures of the $\theta < \theta_o + 3\rho$ portions of $K_{e'}$ and $K_{e''}$ lack singular points. To argue in the case of $K_{e'}$, note that the variation in $w_{e'}$ is not greater than a multiple of $\rho + \varepsilon$, so with the latter very small, only the $\hat{x} = 0$ case of (4-3) and (4-4) can appear. Furthermore, no $\hat{x} = 0$ and $\hat{y} \neq 0 \pmod{m}$ versions of (4-3) and (4-4) can occur in this case with one of

v' and $v' - 2\pi \frac{\hat{y}}{m}$ very close to 0 in $\mathbb{R}/(2\pi\mathbb{Z})$. Indeed, such is the case because $a_{e'}$ achieves its unique maximum on any given constant σ circle at the origin in $\mathbb{R}/(2\pi\mathbb{Z})$. This understood, the existence of any $\hat{x} = 0$ and $\hat{y} \neq 0 \pmod{m}$ solutions to (4-3) and (4-4) is precluded by virtue of two facts: First, $\phi_{e'}$ is a diffeomorphism. Second, if the respective values of the cosine function agree at two distinct points in $\mathbb{R}/(2\pi\mathbb{Z})$, then the corresponding values of the sine function do not.

Except for notational changes, the argument just given also proves the assertion that the $\theta < \theta_o - 3\rho$ part of the closure of $K_{e''}$ is embedded.

Step 2 Turn next to the case of K_e . In this regard, keep in mind that points (σ, v) and (σ', v') in the parametrizing domain are mapped to the same point if and only if the conditions in (4-2) are obeyed for some integer pair N . Equivalently, the conditions in (4-3) and (4-4) are obeyed for some pair $(\hat{x}, \hat{y}) \in \mathbb{Z} \times \mathbb{Z}/(m\mathbb{Z})$, and

$$(4-50) \quad v' = v - 2\pi \hat{x} \frac{\alpha_Z(\sigma)}{\alpha_{Q_e}(\sigma)} - 2\pi \frac{\hat{y}}{m} \pmod{2\pi\mathbb{Z}}$$

To start the story for K_e , note that $a_e > 0$ when $v \in (0, 2\pi \frac{\alpha'}{\alpha})$, while $a_e < 0$ in the case that $v \in (2\pi \frac{\alpha'}{\alpha}, 2\pi)$. Thus, the respective images of the maps that are defined via (3-2) by (4-42) and (4-43) are disjoint except where their domains overlap, where $v = 0$ and $v = 2\pi \frac{\alpha'}{\alpha}$. On both of these loci, $a_e = 0$. In any event, it is sufficient to consider separately the cases where $v \in [0, 2\pi \frac{\alpha'}{\alpha}]$ and where $v \in [2\pi \frac{\alpha'}{\alpha}, 2\pi]$, but taking care not to double count any immersion points that occur where $v = 0$ or $v = 2\pi \frac{\alpha'}{\alpha}$.

The next point to make is that the values of either β_* or β' at any point mapping to a singular point must be non-zero when ε and ρ are small. Indeed, with a reference to (4-42), an argument given previously establishes the existence here of a positive lower bound for $\beta_* + \beta'$ that depends only on θ_o , Q_e and $Q_{e'}$. Thus, solutions to (4-3) and (4-4) where $\sigma \in (\theta_o - 3\rho, \theta_o]$ can occur only where $|\theta_o - \sigma| < 2\rho^4$, or in the image of a point where $x^2 + y^2 \leq 4\rho^2$ via the map in (4-48).

The subsequent discussion involves the number $\wp_0 \in [0, 1)$ that is defined for each pair $(\hat{x}, \hat{y}) \in \mathbb{Z} \times \mathbb{Z}/(m\mathbb{Z})$ by the condition

$$(4-51) \quad \wp_0 = \hat{x} \frac{\alpha_Z(\theta_o)}{\alpha(\theta_o)} + \frac{\hat{y}}{m} \pmod{\mathbb{Z}}.$$

Because any small ε and ρ version of $|w_e|$ in (4-42) is a priori bounded by a constant that depends only on θ_o , Q_e and $Q_{e'}$, so the set of pairs (\hat{x}, \hat{y}) that allow a solution to (4-3) and (4-4) where $\beta_* > 0$ or where $\beta' > 0$ has size bounded by θ_o , Q_e and $Q_{e'}$. Thus, the set of values of \wp_0 that can arise from such pairs has a corresponding

upper bound to its size. Moreover, the possible values for \wp_0 in this set are determined a priori by θ_o , Q_e and $Q_{e'}$.

With \wp_0 so defined, introduce the function of σ given by

$$(4-52) \quad \wp \equiv \wp_0 + \hat{x} \left(\frac{\alpha_Z(\sigma)}{\alpha(\sigma)} - \frac{\alpha_Z(\theta_o)}{\alpha(\theta_o)} \right).$$

Note that if σ is such that $\beta' > 0$, then $|\wp - \wp_0| \leq \kappa |\hat{x}| \rho^4$ where κ is determined solely by θ_o , Q_e and $Q_{e'}$. Note as well that the relation in (4-50) between v' and v can be summarized succinctly by the formula $v' = v - 2\pi\wp \pmod{2\pi\mathbb{Z}}$.

Step 3 Suppose now that (σ, v) is a solution to some given (\hat{x}, \hat{y}) version of (4-3) and (4-4) with $v \in [\frac{1}{\alpha}\rho^2, 2\pi\frac{\alpha'}{\alpha} - \frac{1}{\alpha}\rho^2]$ and $\beta'(\sigma) > 0$. Granted (4-42), the condition in (4-4) is equivalent to

$$(4-53) \quad \cos \left(v - \pi \frac{\alpha'}{\alpha} - 2\pi\wp \right) = \cos \left(v - \pi \frac{\alpha'}{\alpha} \right)$$

If ρ is small and σ fixed, there are at most two values for v that lie in the indicated range and satisfy (4-53). To elaborate, without the constraint on the domain of v , there are precisely two solutions to equation in (4-53) for the given value of σ ; one is the point $\pi(\frac{\alpha'}{\alpha} + \wp)$, and the other is $\pi(\frac{\alpha'}{\alpha} + \wp - 1)$. Moreover, if ρ is sufficiently small and $\frac{\alpha'}{\alpha} < \frac{1}{2}$, then at most one of these lies in the required interval. If $\frac{\alpha'}{\alpha} > \frac{1}{2}$, then at least one of the two is in this interval.

In any event, if $\pi(\frac{\alpha'}{\alpha} + \wp)$ lies in $[\frac{1}{\alpha}\rho^2, 2\pi\frac{\alpha'}{\alpha} - \frac{1}{\alpha}\rho^2]$, then so does $\pi(\frac{\alpha'}{\alpha} - \wp)$; and if $\pi(\frac{\alpha'}{\alpha} + \wp - 1)$ lies in $[\frac{1}{\alpha}\rho^2, 2\pi\frac{\alpha'}{\alpha} - \frac{1}{\alpha}\rho^2]$, then so does $\pi(\frac{\alpha'}{\alpha} + 1 - \wp)$. This said, note that if any given (\hat{x}, \hat{y}) 's version of \wp puts $\pi(\frac{\alpha'}{\alpha} + \wp)$ in $[\frac{1}{\alpha}\rho^2, 2\pi\frac{\alpha'}{\alpha} - \frac{1}{\alpha}\rho^2]$, then the corresponding version of \wp for the pair $(-\hat{x}, m - \hat{y})$ has $\pi(\frac{\alpha'}{\alpha} + \wp - 1)$ in this same interval. The converse is also true. Moreover, this correspondence does not alter the corresponding intersecting disks in K_e since it amounts to switching v with v' . Thus, it is enough to consider the case that $\pi(\frac{\alpha'}{\alpha} + \wp)$ lies in the desired interval $[\frac{1}{\alpha}\rho^2, 2\pi\frac{\alpha'}{\alpha} - \frac{1}{\alpha}\rho^2]$. Note that when ρ is small, such is the case if and only if

$$(4-54) \quad \wp_0 < \frac{\alpha'}{\alpha} |_{\sigma=\theta_o}$$

To start the analysis, use (4-42) to write (4-3) as

$$(4-55) \quad \varepsilon \sin(\pi\wp) = \pi m \frac{1}{\alpha'} \left(\beta' \wp \frac{[Q_e, Q_{e'}]}{m} - \frac{\alpha'}{\alpha} \hat{x} \right).$$

This last equation should be viewed as a condition on σ . In particular, because β' takes values in $[0, 1]$, and because $[Q_e, Q_{e'}] = -[Q_{e'}, Q_{e''}] \geq 0$, any small ε and $\rho < \varepsilon^2$ version of (4-55) can be satisfied by some $\sigma \in [\theta_o - 3\rho, \theta_o]$ if and only if

$$(4-56) \quad 0 < \hat{x} \frac{\alpha'}{\alpha} |_{\sigma=\theta_o} < \wp_0 \frac{[Q_e, Q_{e'}]}{m}.$$

Moreover, this solution occurs at a value of σ where

$$(4-57) \quad \beta' \in (\kappa, (1 - \kappa)).$$

$$\beta' \wp \frac{[Q_e, Q_{e'}]}{m} - \frac{\alpha'}{\alpha} \hat{x} > \frac{1}{\kappa} \varepsilon.$$

Here, $\kappa \in (0, 1)$ is a constant that depends only on θ_o , Q_e and $Q_{e'}$. Finally, if ε is small, if $\rho < \varepsilon^2$, and if (4-56) holds, then there is a unique choice of σ that solves (4-55).

There is one last point to make for the cases when \wp_0 obeys (4-54): Distinct values for the pair (\hat{x}, \hat{y}) produce disjoint singular points in K_e . To explain, note first that by virtue of the fact that the respective values of a_e that arise must be equal, two choices for (\hat{x}, \hat{y}) can produce the same singular point in K_e only if the corresponding values for \wp_0 agree. Granted this, if the resulting singular points have the same θ coordinate, then (4-55) demands that the respective values for \hat{x} agree. Thus, so do the values for \hat{y} .

Step 4 The story in this step concerns the cases where (4-3) and (4-4) hold with a value of v either in $[2\pi \frac{\alpha'}{\alpha} - \frac{1}{\alpha} \rho^2, 2\pi \frac{\alpha'}{\alpha}]$ or in $[0, \frac{1}{\alpha} \rho^2]$. As is explained next, no solutions to (4-3) and (4-4) with such values for v result in K_e singularities if ε and ρ are small, and if θ_o is suitably generic.

To begin the explanation, assume for the moment only that v is within $\frac{1}{\alpha} \rho^2$ of either 0 or $2\pi \frac{\alpha'}{\alpha}$. If $v \in [2\pi \frac{\alpha'}{\alpha} - \frac{1}{\alpha} \rho^2, 2\pi \frac{\alpha'}{\alpha}]$, then (4-4) requires that \wp_0 is either 0 or equal to the value of $\frac{\alpha'}{\alpha}$ at θ_o . If $v \in [0, \frac{1}{\alpha} \rho^2]$, then \wp_0 is either zero or equal to the value of $1 - \frac{\alpha'}{\alpha}$ at θ_o . In this regard, note that when ε and ρ are small, the case where $\wp_0 = 0$ requires $\hat{x} = 0$ and thus $\hat{y} = 0$ as well. Thus, the $\wp_0 = 0$ case does not lead to a singularity in K_e . Meanwhile, at the risk of replacing the pair (\hat{x}, \hat{y}) with $(-\hat{x}, m - \hat{y})$, it is sufficient to study the case where $v \in [2\pi \frac{\alpha'}{\alpha} - \frac{1}{\alpha} \rho^2, 2\pi \frac{\alpha'}{\alpha}]$ and where \wp_0 is equal to the θ_o value of $\frac{\alpha'}{\alpha}$.

To see what this last condition implies, write $Q_{e'} = [Q_{e'}, Q_e]Z + \frac{w}{m} Q_e$ where $w \in \mathbb{Z}$. Doing so identifies

$$(4-58) \quad \frac{\alpha'}{\alpha} = [Q_{e'}, Q_e] \frac{\alpha Z}{\alpha} + \frac{w}{m}$$

and so the condition on \wp_0 requires that

$$(4-59) \quad ([Q_{e'}, Q_e] - \hat{x}) \frac{\alpha_Z}{\alpha} = 0 \pmod{\left(\frac{1}{m}\mathbb{Z}\right)}.$$

Now, if θ_o is suitably generic, then the ratio α_Z/α will be irrational, and so the only solution to (4-59) is that where $\hat{x} = [Q_{e'}, Q_e]$. However, with ε and ρ small, a glance at (4-42) shows that such a value for \hat{x} is incompatible with the condition in (4-52) unless $[Q_{e'}, Q_e]$ and \hat{x} both vanish. Indeed, such is the case because

$$(4-60) \quad w_e(\sigma, v - 2\pi\wp) - w_e(\sigma, v) = 2\pi\beta' \frac{[Q_e, Q_{e'}]}{\alpha} + o(\varepsilon + \rho),$$

and this has the same sign as $-\hat{x}$ in the case that $\hat{x} = [Q_{e'}, Q_e]$, neither are zero and both ε and ρ are small.

To rule out the case that both $[Q_{e'}, Q_e] = 0$ and $\hat{x} = 0$, note that (4-42) demands that the resulting singularity in K_e is the image of two points in the (x, y) plane via the map in (4-48) where one has the form $(0, y)$ and the other $(0, -y)$. Moreover, (4-4) demands that y obey

$$(4-61) \quad \beta_* |y| = -(1 - \beta_*) \varepsilon \sin\left(\pi \frac{w}{m}\right)$$

where $w \in \{1, \dots, m-1\}$ is the integer that appears in (4-58). Since $\pi \frac{w}{m} = \frac{\alpha'}{\alpha}$ in this case and since $\frac{\alpha'}{\alpha} \in (0, 1)$, the right hand side of (4-61) is non-positive and the left hand side is non-negative. Since the two sides can not vanish simultaneously, there are no values of y that make (4-61) hold.

In the case that θ_o is special and so there is an $\hat{x} \neq [Q_{e'}, Q_e]$ solution to (4-59), there may well be solutions to (4-3) and (4-4) with $v \in [2\pi \frac{\alpha'}{\alpha} - \frac{1}{\alpha} \rho^2, 2\pi \frac{\alpha'}{\alpha}]$. These can be analyzed with much the same machinery as used for when v is further from either 0 or $2\pi \frac{\alpha'}{\alpha}$. To keep an already long story from getting longer, this task is left to the reader, as is the task of verifying that the resulting singularities of K_e are transversal with local self-intersection number 1.

Step 5 This final step characterized the singularities in the closure of K_e that lie where $|\theta_o - \sigma| < 3\rho$ and $v \in [2\pi \frac{\alpha'}{\alpha}, 2\pi]$. The story here is much as in Step 3 and Step 4. First, when θ_o is sufficiently generic, there are no solutions that lie in the image via the map in (4-48) of points where $\beta_* \neq 0$. Second, if ε and ρ are sufficiently small, then all singularities in K_e must lie where $\beta'(\theta)$ is bounded away from zero by a constant that depends only on θ_o , Q_e and $Q_{e'}$. Third, if (σ, v) and (σ', v') map to the same point in K_e , then $\sigma = \sigma'$ and $v' = v - 2\pi\wp \pmod{2\pi\mathbb{Z}}$ where \wp is defined as in (4-52).

Here, $\wp_0 \in [0, 1)$ is defined from some pair (\hat{x}, \hat{y}) as in (4-51). In addition, (4-53) must hold. As a consequence, if v is taken to be a real number in $[2\pi \frac{\alpha'}{\alpha}, 2\pi]$, then v must have one of two forms: Either $v = \pi(1 + \frac{\alpha'}{\alpha} + \wp)$ or else $v = \pi(\frac{\alpha'}{\alpha} + \wp)$. In the former case, $1 - \wp > \frac{\alpha'}{\alpha}$ and in the latter, $\wp > \frac{\alpha'}{\alpha}$. For essentially the same reasons as before, it is only necessary to consider one of these two sorts of cases, and so the discussion below makes the assumption that $v = \pi(1 + \frac{\alpha'}{\alpha} + \wp)$ and that $1 - \wp > \frac{\alpha'}{\alpha}$. In this regard, note that when ρ is small, then the latter condition holds if and only if $1 - \wp_0 > \frac{\alpha'}{\alpha}$.

With this last assumption understood, then $v' = \pi(1 + \frac{\alpha'}{\alpha} - \wp)$ is also in $[2\pi \frac{\alpha'}{\alpha}, 2\pi]$ and $v > v'$. This being the case, a referral to (4-43) finds (4-3) equivalent to the condition

$$(4-62) \quad -\varepsilon \sin(\wp) = \pi m \frac{1}{\alpha''} \left(\beta' \frac{[Q_e, Q_{e''}]}{m} \wp - \frac{\alpha''}{\alpha} \hat{x} \right).$$

As with its analog in (4-55), this should be viewed as an equation for σ . As such, it has a solution if and only if

$$(4-63) \quad \frac{[Q_e, Q_{e''}]}{m} \wp_0 < \hat{x} \frac{\alpha''}{\alpha} \Big|_{\sigma=\theta_o} < 0.$$

Moreover, when ε is small and ρ is very small, then the solution is unique and it occurs where

$$(4-64) \quad \begin{aligned} &\beta' \in (\kappa, 1 - \kappa). \\ &\beta' \frac{[Q_e, Q_{e''}]}{m} \wp - \frac{\alpha''}{\alpha} \hat{x} < \frac{1}{\kappa} \varepsilon, \end{aligned}$$

where $\kappa \in (0, 1)$ is a constant that depends only on θ_o , Q_e and $Q_{e'}$.

An argument from Step 3 also applies here to prove that each singular point in the $\beta' > 0$ and $v \in [2\pi \frac{\alpha'}{\alpha}, 2\pi]$ portion of K_e lies in a ball whose intersection with K_e is the union of two embedded disks that meet only at their centers.

Task 3 The task here is to verify that small ε and very small ρ guarantees that the singularities in the $\theta > \theta_o - 3\rho$ portion of K_e are those of transversally intersecting disks with positive local intersection number. For this purpose, use (a, w) and (a', w') to denote the respective versions of the parametrizing functions that come from the two disk that are involved at the given intersection point. Use $(\partial_\sigma, \partial_v)$ and $(\partial_{\sigma'}, \partial_{v'})$ to denote the corresponding versions of the push-forward that are depicted in (4-33). Consider first the case where the inverse image in the parametrizing cylinder of the singular point in the unprimed disk is a point (σ, v) with $v = \pi(\frac{\alpha'}{\alpha} + \wp)$ and with

$\wp < \frac{\alpha'}{\alpha}$. In this case, the primed pair is the image of (σ, v') with $v' = \pi(\frac{\alpha'}{\alpha} - \wp)$. Thus, both v and v' lie in the interval $(0, 2\pi\frac{\alpha'}{\alpha})$.

In the present situation, the intersection is transversal if τ as defined in (4–34) is non-zero; and then the sign of τ is the local sign at the intersection point. In the case at hand, a referral to (4–42) finds that

$$(4-65) \quad \begin{aligned} a_v - a'_v &= -2\varepsilon \sin(\pi\wp) < 0. \\ w_\sigma - w'_\sigma &= 2\pi\wp \frac{[Q_e, Q_{e'}]}{\alpha'} \beta'_\sigma + o(1). \end{aligned}$$

Here, the term that is designated as $o(1)$ is uniformly bounded no matter how small ε and ρ . By virtue of (4–57), β'_ρ is bounded from below by $\kappa'\rho^{-4}$, with $\kappa' > 0$ depending solely on θ_o , Q_e and $Q_{e'}$. Meanwhile, $[Q_e, Q_{e'}] = -[Q_{e'}, Q_{e''}]$ and thus is positive. Therefore, τ is positive and so the local sign at the intersection point is $+1$.

To continue with the assigned task, the second case to consider are those intersections where $v = \pi(\frac{\alpha'}{\alpha} + 1 + \wp)$ and $1 - \wp > \frac{\alpha'}{\alpha}$. Here, $v' = \pi(\frac{\alpha'}{\alpha} + 1 - \wp)$. Thus, the self-intersections that occur in the case at hand occur at values of v and v' that lie in the interval $(2\pi\frac{\alpha'}{\alpha}, 2\pi)$. This noted, then referral to (4–43) finds

$$(4-66) \quad \begin{aligned} a_v - a'_v &= 2\varepsilon \sin(\pi\wp) > 0. \\ w_\sigma - w'_\sigma &= 2\pi\wp \frac{[Q_e, Q_{e''}]}{\alpha''} \beta'_\sigma + o(1), \end{aligned}$$

where the term designated as $o(1)$ has the same properties as its analog in (4–65). In this case, the second line in (4–66) is very negative when ρ is small, so τ is again positive and the local intersection number is equal to 1.

4.E Intersections between distinct cylinders

The purpose of this next to last subsection is to complete the proof of Theorem 3.1 in the case that all partition sets that define T have a single element. In this regard, the task here is to verify that the intersections between cylinders K_e and $K_{e'}$ when $e \neq e'$ are distinct edges in the graph T are transversal with positive local intersection number.

To see how such a guarantee can be made, suppose that e is any given edge. The immersion constructed in the preceding subsections that defines K_e involved the specification of data $\{\rho_{e0}, \rho_{e1}, \varepsilon_e, a_e^0, w_e^0, v_e^0\}$ where $\varepsilon_e, a_e^0, w_e^0$ and v_e^0 are functions of the coordinate σ on the parametrizing cylinder, and where the other two are constant and positive. Although subsequent subsections gave upper bounds for ε_e near the boundaries of the parametrizing cylinder there is no positive lower bound near these

boundaries. With ε_e chosen near boundaries of the parametrizing cylinder, upper bounds were then specified for the constants ρ_{e0} and ρ_{e1} , but there were no positive lower bounds. As is explained below, the transversality of all intersections between all pairs of distinct K_e and $K_{e'}$ is guaranteed with careful choices for the various versions of $\{\varepsilon_e, \rho_{e0}, \rho_{e1}, a_e^0, w_e^0, v_e^0\}$. As should be evident from details below, all new constraints on $\{\varepsilon_e, \rho_{e0}, \rho_{e1}, a_e^0, w_e^0, v_e^0\}$ are compatible with those given in the previous subsections.

The discussion in the remainder of this subsection is divided into eight parts.

Part 1 There is an immediate issue that arises when edges e' and e'' have monovalent vertices that share an angle assignment. Of concern is to choose the $e = e'$ and $e = e''$ versions of $\{\varepsilon_e, \rho_{e0}, \rho_{e1}, a_e^0, w_e^0, v_e^0\}$ so as to keep the resulting versions of $K_{e'}$ and $K_{e''}$ disjoint at values of θ that approach the common vertex angle assignment. The discussion of this issue addresses the respective cases where the angle label in question is 0, π , and then neither 0 nor π .

The case of angle 0

Let o' and o'' denote the respective vertices on e' and e'' with angle label 0. Now, o' has a label from \hat{A} , either of the form $(1, +, \dots)$, $(1, -, \dots)$ or simply $\{1\}$. In the first case, $s \rightarrow \infty$ on $K_{e'}$ as $\theta \rightarrow 0$, in the second case, the function s has a finite limit as $\theta \rightarrow 0$, and in the third case, $s \rightarrow -\infty$ on $K_{e'}$ as $\theta \rightarrow 0$. Of course, o'' has one of these three sorts of labels also.

Now, suppose that o' is labeled by an element of the form $(1, +, \dots)$ from \hat{A} and o'' by either an element $\{1\}$ from \hat{A} or one of the form $(1, -, \dots)$. In this case, make both $a_{e'}^0$ and $a_{e''}^0$ constant at values of σ where either the e' or e'' version of the relevant (4–6) or (4–8) holds. Choose these constants so that $a_{e'}^0 \gg a_{e''}^0$; this then makes $K_{e'}$ disjoint from $K_{e''}$ where either version of which ever of (4–6) or (4–8) holds. A similar choice for $a_{e'}^0$ and $a_{e''}^0$ guarantees this same conclusion when o' is labeled by $\{1\}$ and o'' by $(1, -, \dots)$.

Suppose next that $[Q_{e'}, Q_{e''] < 0$ and that both o' and o'' are labeled by $(1, +, \dots)$ elements from \hat{A} . Granted this, take $\rho_{e'0} \gg \rho_{e''0}$, and take $a_{e'}^0$ and $a_{e''}^0$ to be constant where either version of (4–6) holds with $a_{e'}^0 \gg a_{e''}^0$. This makes $K_{e'}$ and $K_{e''}$ disjoint where either version of (4–6) holds. In the case that $[Q_{e'}, Q_{e''] < 0$ and both o' and o'' are labeled by an $(-1, -, \dots)$ element from \hat{A} , now take $\rho_{e'0} \ll \rho_{e''0}$, and take $a_{e'}^0$ and $a_{e''}^0$ to be constant where either version of (4–6) holds, but keep $a_{e'}^0 \gg a_{e''}^0$.

To continue, suppose that $[Q_{e'}, Q_{e''] = 0$. The simplest case has both o' and o'' labeled by (1). Here, it is sufficient to take $a_{e'}^0$ and $a_{e''}^0$ to be constant where either

version of (4–8) holds, but with $|a_{e'}^0 - a_{e''}^0| > 1$. Such a choice makes $K_{e'}$ and $K_{e''}$ disjoint where either version of (4–8) holds.

Here is the story when $[Q_{e'}, Q_{e''] = 0$ and both o' and o'' are labeled by $(1, +, \dots)$ elements from \hat{A} , or both by $(1, -, \dots)$ elements from \hat{A} . In either case, make $\rho_{e'0} = \rho_{e''0}$, both $\varepsilon_{e'}$ and $\varepsilon_{e''}$ much less than 1. Then, make both $a_{e'}^0$ and $a_{e''}^0$ constant where either version of (4–6) holds, but with $|a_{e''}^0| \gg |a_{e'}^0| > 1$. This then guarantees that $K_{e'}$ is disjoint from $K_{e''}$ where (4–6) holds, and both $\varepsilon_{e'}$ and $\varepsilon_{e''}$ much less than 1 to insure that $K_{e'}$ and $K_{e''}$ are disjoint where either version of (4–6) holds.

The case of angle π

Each of the angle 0 subcases just described has a very evident angle π analog and vice versa. The stories for the corresponding angle 0 and angle π subcases are identical save for some notation and sign changes. This understood, the angle π cases are left to the reader save for the following equations that give the angle π versions of (4–6) and (4–7):

$$(4-67) \quad \begin{aligned} a_e &= \frac{1}{\kappa} \beta' \ln(\pi - \sigma) + a_e^0 + (\varepsilon(1 - \beta') + (\pi - \sigma)\beta') \cos(v + v_e^0). \\ w_e &= (1 - \beta')w_e^0 - (\varepsilon(1 - \beta') + (\pi - \sigma)\beta') \sin(v + v_e^0). \end{aligned}$$

Here, $\varepsilon \equiv \varepsilon_e$ and

$$(4-68) \quad \kappa \equiv -\frac{q'_e}{q_e} + \sqrt{\frac{3}{2}},$$

Note that the angle π version of (4–8) has the same form as the original.

The case of neither 0 nor π

Let e' and e'' again denote the two edges that are involved. The first point to make is that $K_{e'} \cap K_{e''} = \emptyset$ if the vertex o' has the smaller angle label of the two vertices on e' while o'' has the larger of the two angle labels of the vertices on e'' . This understood, consider the case where the angle labels of o' and o'' are either both the smaller of the two angle labels on their incident edges. In this case, take $\rho_{0e'} = \rho_{0e''}$, with both much less than 1. Likewise, choose ε_e and $\varepsilon_{e'}$ to be very small. Finally, take both $a_{e'}^0$ and $a_{e''}^0$ to be constant with $|a_{e'}^0 - a_{e''}^0| > 2\pi$ where (4–9) is valid. This makes K_e and $K_{e'}$ disjoint where (4–9) holds.

Part 2 There is also an issue to address in the case that two bivalent vertices in T have the same angle assignment. Let o denote the first and let e and e' denote its incident edges using the usual convention where θ_o is the larger of the angles that are

assigned to the vertices on e . Let \hat{o} denote a vertex with $\theta_{\hat{o}} = \theta_o$, and let \hat{e} and \hat{e}' denote the corresponding incident edges. Let Y_o denote the closure of $K_e \cup K_{e'}$ and let $Y_{\hat{o}}$ denote that of $K_{\hat{e}} \cup K_{\hat{e}'}$. Since $s \rightarrow \infty$ on both Y_o and $Y_{\hat{o}}$ along certain paths where θ limits to θ_o , these two subvarieties may well intersect where θ is near θ_o . The goal is to insure that the intersection points are transversal with positive intersection number.

For this purpose, keep in mind that both Y_o and $Y_{\hat{o}}$ converge in $S^1 \times S^2$ as multiple covers of $\theta = \theta_*$ Reeb orbits. Let γ_o and $\gamma_{\hat{o}}$ denote the latter. Note that γ_o is determined by the $\theta = \theta_o$ value for the parameter v_e^0 , and $\gamma_{\hat{o}}$ is determined in an analogous fashion by $v_{\hat{e}}^0$. In particular, if the respective constant values of $v_e^0 = v_{e'}^0$ and $v_{\hat{e}}^0 = v_{\hat{e}'}^0$ near the $\sigma = \theta_o$ circles in the relevant parametrizing domains are chosen to be unequal and sufficiently generic, then γ_o and $\gamma_{\hat{o}}$ will be distinct Reeb orbits. Choose these two angles to insure that such is the case.

To continue, take $\rho_{e1} \ll \rho_{\hat{e}1}$ and take a_e^0 and $a_{\hat{e}}^0$ both constant with $a_{\hat{e}}^0 \gg a_e^0$ at points in the respective parametrizing cylinders where σ is within $3\rho_{\hat{e}}$ of θ_o . In particular, choose $a_e^0 \gg a_e^0 - 2 \ln(\rho_{e1})$. Likewise, make $a_{e'}^0$ and $a_{\hat{e}'}^0$ both constant with $a_{\hat{e}'}^0 \gg a_{e'}^0$ at points where $\sigma < \theta_o + 3\rho_{\hat{e}}$ in their respective parametrizing cylinders. If $\rho_{\hat{e}1}$ is sufficiently small, then these choices have the following consequences: First, all intersections between Y_o and $Y_{\hat{o}}$ occur at points in Y_o at very large s , in particular where the o version of the coordinates (r, τ) are defined and where the corresponding $\beta_* = 1$. More to the point, these intersection points occur where Y_o looks very much like a multiple cover of the \mathbb{R} -invariant cylinder $\mathbb{R} \times \gamma_o$. Meanwhile, these intersection points occur in $Y_{\hat{o}}$ where the \hat{o} version of β_* is zero.

Granted the preceding, keep in mind the following: Let I denote an arc with compact closure in an orbit of the Reeb vector field. Then $\mathbb{R} \times I$ has transversal intersections with the closure of any given version of $K_{(\cdot)}$, and that these intersection points have positive local intersection number. As can be verified using (4–33), this is a consequence of the positivity of the relevant version of the function α_Q .

Now, as remarked, if $\rho_{e1} \ll \rho_{\hat{e}1}$ and if $a_e^0 \ll a_{\hat{e}}^0$, then Y_o looks very much like the cylinder $\mathbb{R} \times \gamma_o$ where it intersects $Y_{\hat{o}}$. Meanwhile, neighborhoods in $Y_{\hat{o}}$ of the intersection points are constant translates along \mathbb{R} in $\mathbb{R} \times (S^1 \times S^2)$ of a standard embedding. This understood, it should not come as a surprise that these intersections are also transversal and have positive intersection number. It is left to the reader as an exercise with (4–14), (4–15) and (4–33) to verify that such is the case.

Part 3 This part of the discussion provides an overview of the strategy that is used below to control the remaining intersections between distinct versions of $K_{(\cdot)}$.

To start, suppose that e' and e'' are edges of T , and that $\theta_{*0} < \theta_{*1}$ are values of θ on both $K_{e'}$ and $K_{e''}$. In addition, suppose that both the $e = e'$ and $e = e''$ versions of the pair (a_e, w_e) are given by (4–1) when $\sigma \in [\theta_{*0}, \theta_{*1}]$. It then follows from (3–2) that $K_{e'}$ and $K_{e''}$ are disjoint provided that

$$(4-69) \quad |a_{e'}^0 - a_{e''}^0| > \varepsilon_{e'} + \varepsilon_{e''} \text{ when } \sigma \in [\theta_{*0}, \theta_{*1}].$$

The pair $K_{e'}$ and $K_{e''}$ are said below to be ‘well separated’ at a given value, θ_* , of θ if (4–1) describes the $e = e'$ and $e = e''$ versions of (a_e, w_e) and if the inequality in (4–69) holds at $\sigma = \theta_*$. So as to avoid repetitive qualifiers, the respective portions of two versions of $K_{(c)}$ where θ has a given range are also deemed ‘well separated’ in the event that one or both such portions is empty.

The strategy used below keeps the various versions of $K_{(c)}$ pairwise well separated as much as possible. To implement the strategy, first fix some positive constant ρ , smaller than the constant δ that was introduced in Subsection 4.A. Thus, ρ is much smaller than $\frac{1}{1000}$ times the difference between the larger and smaller of the angles that label the vertices on any given edge of T . Agree to make sure that all versions of ρ_{e0} and ρ_{e1} are much less than ρ . The plan is to keep the versions of $K_{(c)}$ pairwise well separated at angles with distance 2ρ or more from the angles that label T 's vertices. To be precise, the various versions of a_e^0 are taken to be locally constant on the complement in $[0, \pi]$ of the points with distance 2ρ or less from the finite set of angles that label the vertices of T . Of course, these constant values are chosen to insure that (4–69) is pairwise obeyed.

Granted the preceding, it is worth noting in advance those values of θ where well separation must be abandoned. It proves useful for this purpose to have on hand a particular proper immersion of T into the rectangle $[-1, 1] \times [0, \pi]$. To define this immersion, first map the monovalent and bivalent vertices to the boundary of the rectangle $[-1, 1] \times [0, \pi]$ in the following manner: Each vertex whose label from \hat{A} has the form $(\cdot, -, \dots)$ is placed on $\{-1\} \times [0, \pi]$ by using its angle label for the $[0, \pi]$ factor. The analogous rule places each vertex from \hat{A} of the form $(\cdot, +, \dots)$ on $\{1\} \times [0, \pi]$. Put each monovalent vertex with label (1) from \hat{A} on $(-1, 1) \times \{0\}$, and put each with a (-1) label on $(-1, 1) \times \pi$. In this regard, if e is the incident edge to such a vertex, set the horizontal coordinate of the vertex equal to the value of $\tanh(a_e^0)$ at either $\sigma = 0$ or $\sigma = \pi$ as the case may be. If o is any given trivalent vertex, use θ_o to denote its angle label, and place o on $(-1, 1) \times \theta_o$.

To finish the construction, it is necessary to identify each edge of T with an arc in the rectangle that runs between the relevant vertices. This is to be done so that the interior of each arc avoids the boundary of the rectangle and also avoids all vertices. In

addition, the horizontal coordinate on the rectangle must restrict without critical points to each arc.

The vertical coordinate on the arc labeled by a given edge e at any given interior point is written below as $\tanh(s_e)$ with $s_e \in \mathbb{R}$.

If e' and e'' are distinct edges and if their representative arcs can be drawn without interior intersections, then $K_{e'}$ and $K_{e''}$ can be kept well separated. Indeed, this can be done along the following lines: Use σ to denote the vertical coordinate in the rectangle. As σ has nowhere zero derivative on each arc, so the corresponding version of $s_{(\cdot)}$ can be viewed as a function of σ . This understood, identify $a_{e'}^0$ and $a_{e''}^0$ with $s_{e'}$ and $s_{e''}$. If $\varepsilon_{e'}$ and $\varepsilon_{e''}$ are made very small, the resulting $K_{e'}$ will then be well separated from $K_{e''}$.

As might be expected, crossing of these edge labeled arcs may be unavoidable. To identify the necessary arc crossings, draw the edge labeled arcs by starting at the top edge of the rectangle, $[-1, 1] \times \{\pi\}$, and proceeding downwards. To conform to what has been said already, all edge labeled arcs will be drawn as vertical arcs except perhaps where the horizontal coordinate has distance 2ρ or less to an angle that labels a vertex in T . Of course, distinctly labeled vertical arcs will have distinct horizontal coordinates.

With the arcs drawn in this manner, the following are the only circumstances that may require one arc to cross another:

- (4–70) • *If edges e and \hat{e} have monovalent vertices that share an angle assignment in $(0, \pi)$, then a crossing of their arcs may be necessary to keep the a_e^0 and $a_{\hat{e}}^0$ assignments compatible with those given already in Part 1.*
- *Let o denote a monovalent vertex with a label $(0, -, \dots)$ and suppose that e is the incident edge. Let \hat{e} denote a second edge whose vertices are assigned angles that are distinct from θ_o . The arcs labeled by e and \hat{e} must cross in the case that s_e and $s_{\hat{e}}$ are both defined with $s_e > s_{\hat{e}}$ at the relevant $\sigma \in \{\theta_o \pm 2\rho\}$. It follows from (4–9) that no crossing is necessary if $s_e < s_{\hat{e}}$ at this value of σ .*
 - *Let o denote a bivalent vertex and suppose that e is an edge that is incident to o . The arc labeled by e must cross that labeled by some other edge \hat{e} if s_e and $s_{\hat{e}}$ are both defined with $s_e < s_{\hat{e}}$ at the relevant $\sigma \in \{\theta_o \pm 2\rho\}$. It follows from (4–14) and (4–15) that no crossing is necessary if $s_e > s_{\hat{e}}$ at this value of σ .*
 - *Let o denote a trivalent vertex that connects by two incident edges to vertices with larger angle label. Denote these two edges by e' and e'' . Then the respective arcs labeled by e' and e'' cross in the case that $[Q_{e'}, Q_{e''] < 0$*

and $s_{e'} < s_{e''}$ at $\sigma = \theta_o + 2\rho$. It follows from (4–46) and (4–47) that no crossing is necessary if both $[Q_{e'}, Q_{e''] < 0$ and $s_{e'} > s_{e''}$ at $\sigma = \theta_o + 2\rho$, or if $[Q_{e'}, Q_{e''] = 0$.

- Let o, e' and e'' be as in the previous point. Let \hat{e} denote a third edge and suppose that $s_{\hat{e}}$ is defined at $\sigma = \theta_o + 2\rho$ and suppose that it lies between $s_{e'}$ and $s_{e''}$. Then the arc labeled by \hat{e} must cross either that labeled by e' or that labeled by e'' .
- Let o and o' both denote vertices with angle label 0. Then the respective arcs that are labeled by the incident edges to o and o' may have to cross to keep the $a_{e'}^0$ and $a_{e''}^0$ assignments compatible with those given already in Part 1 of this discussion.

Part 5–Part 8 below address these various cases.

Part 4 This part of the story relates two observations that are subsequently exploited in the case that $K_{e'}$ and $K_{e''}$ can not be kept well separated.

Observation 1 This observation concerns an example where (4–69) holds at $\sigma = \theta_{*0}$ and $\sigma = \theta_{*1}$, fails in between, yet $K_{e'}$ and $K_{e''}$ remain disjoint. In particular, if $Q_{e'}$ is proportional to $Q_{e''}$, then the respective signs of $a_{e'}^0 - a_{e''}^0$ can differ at $\sigma = \theta_{*1}$ and at $\sigma = \theta_{*0}$ with $K_{e'}$ still disjoint from $K_{e''}$.

To explain, note that whether or not $Q_{e'}$ and $Q_{e''}$ are proportional, a point (σ, v') in the parametrizing cylinder for $K_{e'}$ and a point (σ, v'') in the parametrizing cylinder for $K_{e''}$ are sent via the relevant versions of (3–2) to the same point in $\mathbb{R} \times (S^1 \times S^2)$ if and only if the following holds: There is an \mathbb{R} -valued lift, \hat{v}' , of v' , a corresponding lift, \hat{v}'' , of v'' , and an integer pair $N = (n, n')$ such that

$$\begin{aligned}
 \alpha_{Q_{e'}} \hat{v}' &= \alpha_{Q_{e''}} \hat{v}'' - 2\pi \alpha_N \\
 a_{e'}(\sigma, v') &= a_{e''}(\sigma, v'') \\
 w_{e''}(\sigma, v') &= w_{e''}(\sigma, v'') - \frac{1}{\alpha_{Q_{e'}}} [Q_{e'}, Q_{e''}] \hat{v}'' + 2\pi \frac{1}{\alpha_{Q_{e'}}} [Q_{e'}, N]
 \end{aligned}
 \tag{4-71}$$

Now, if $Q_{e'}$ is proportional to $Q_{e''}$ then the middle term in the lowest line above is zero. Such being the case, let κ denote the maximum of the $Q = Q_{e'}$ version of α_Q over the interval $[\theta_{*0}, \theta_{*1}]$. Then the third point in (4–71) can not be met if $0 < |w_{e'} - w_{e''}| < 2\pi \frac{1}{\kappa}$ in this interval. This last condition can be achieved by suitable choices of $w_{e'}^0$ and $w_{e''}^0$, if ε_e and $\varepsilon_{e''}$ are small. Of course, if the third condition in (4–71) can not be met, then no amount of variation in $a_{e'}^0$ and $a_{e''}^0$ on $[\theta_{*0}, \theta_{*1}]$ will make $K_{e'}$ intersect $K_{e''}$.

Observation 2 Intersections between $K_{e'}$ and $K_{e''}$ are allowed when transversal with +1 local intersection number. This can always be arranged at values of θ in $(\theta_{*0}, \theta_{*1})$ if the following three conditions hold: First, $[Q_{e'}, Q_{e''}] < 0$. Second, (4–69) holds at both $\sigma = \theta_{*0}$ and at $\sigma = \theta_{*1}$. Finally, $a_{e'}^0 - a_{e''}^0$ is positive at $\sigma = \theta_{*0}$ but negative at $\sigma = \theta_{*1}$.

To explain this last claim, note that when both ε_e and $\varepsilon_{e'}$ are sufficiently small, then (4–33) dictates that the intersections between $K_{e'}$ and $K_{e''}$ where $\theta \in (\theta_{*0}, \theta_{*1})$ are transversal with sign that of $(a_{e'}^0 - a_{e''}^0)_\sigma [Q_{e'}, Q_{e''}]$. Thus, if the given conditions are satisfied, then the variation of $a_{e'}^0$ and $a_{e''}^0$ over $[\theta_{*0}, \theta_{*1}]$ can be arranged to guarantee the given conclusions.

Part 5 This part of the discussion considers the cases in the first and second points of (4–70). In the situation outlined by the first point, any two edges that are involved will have respective versions of $Q_{(\cdot)}$ that are proportional. With this understood, then the first observation in the preceding Part 4 can be used to keep the corresponding versions of $K_{(\cdot)}$ disjoint in spite of the crossing of the corresponding arcs in $[-1, 1] \times [0, \pi]$.

Turn now to the second point in (4–70). Suppose that o is a monovalent vertex in T with label $(0, -, \dots)$ from \hat{A} . Suppose first that θ takes values that are greater than θ_o on K_e . If, as assumed, $a_e^0 > a_{\hat{e}}^0$ at $\sigma = \theta_o + 2\rho$, then a suitable modification of a_e^0 can guarantee that the $\theta \in (\theta_o, \theta_o + 2\rho)$ intersection points between K_e and $K_{\hat{e}}$ are transversal and have +1 local intersection number.

To explain, take ρ_{e0} to be very small, and let $\theta_{*0} \equiv \theta_o + 2\rho_{e0}$ and $\theta_{*1} \equiv \theta_o + 2\rho$. The assumption here is that $a_e^0 - a_{\hat{e}}^0$ is negative at θ_{*1} . The goal then is to modify a_e^0 inside the interval $(\theta_{*0}, \theta_{*1})$ so that $a_e^0 - a_{\hat{e}}^0$ is positive at θ_{*0} . For this purpose, keep in mind that the pair Q_e is equal to $m(p, p')$ with m a positive integer and with (p, p') the relatively prime pair of integers that θ_o defines via (1–7). As such, there is a positive number, κ , such that

$$(4-72) \quad \alpha_{Q_{\hat{e}}}(\theta_o) = \kappa [Q_e, Q_{\hat{e}}]$$

Since the right hand side of (4–72) is positive, so $[Q_e, Q_{\hat{e}}] > 0$. This understood, if \hat{e} and e are respectively renamed as e' and e'' , then the final observations in Part 4 can be applied here to find the desired modification of a_e^0 .

Consider next the case that θ takes values on K_e that are less than θ_o . In this case, set $\theta_{*0} = \theta_o - 4\rho_{e1}$ and set $\theta_{*1} = \theta_o - 2\rho$. By assumption, $a_e^0 - a_{\hat{e}}^0 > 0$ at $\sigma = \theta_{*0}$, and the goal is to modify a_e^0 inside the interval $(\theta_{*0}, \theta_{*1})$ so that $a_e^0 - a_{\hat{e}}^0$ is negative at $\sigma = \theta_{*1}$. For this purpose, note that Q_e in this case is equal to $-m(p, p')$ with m a positive integer and with (p, p') as before. Thus, (4–72) holds with $\kappa < 0$ and so

$[Q_{\hat{e}}, Q_e] > 0$. In this case, agree to relabel \hat{e} as e'' as e as e' . This done, then the observations in Part 4 again apply to give the desired modification of a_e^0 .

Part 6 This part considers the third point in (4–70). Consider here the case that o is a bivalent vertex in T . Let e and e' denote the edges of T that contain o with the convention taken here that θ takes values that are less than θ_o on K_e . Let ρ_* denote the equal values of ρ_{e1} and $\rho_{e'0}$. Now, suppose that \hat{e} is a third edge and that θ_o is a value of θ on $K_{\hat{e}}$. As noted in (4–70), in the case that $a_{\hat{e}}^0 < a_e^0$ where $\sigma = \theta_o - 2\rho$ then it is a straightforward consequence of (4–14) and (4–15) that the values of $a_{\hat{e}}^0$ can be modified if necessary where $|\sigma - \theta_o| < 2\rho$ so as to guarantee that $K_{\hat{e}}$ is disjoint from the $|\theta - \theta_o| \leq 2\rho$ part of the closure of $K_e \cup K_{e'}$.

On the other hand, if $a_{\hat{e}}^0 - a_e^0 > 0$ where $\sigma = \theta_o - 2\rho$, then it may not be possible to modify $a_{\hat{e}}^0$ where $|\sigma - \theta_o| \leq 2\rho$ so that $K_{\hat{e}}$ avoids the $\theta \in [\theta_o - 2\rho, \theta_o + 2\rho]$ portion of the closure of $K_e \cup K_{e'}$. However, as is explained next, there are modifications that guarantee that all intersection points here are transversal with +1 local intersection number.

To see how this such modifications come about, suppose that $a_{\hat{e}}^0 > a_e^0$ where $\sigma = \theta_o - 2\rho$. In this case, keep a_e^0 constant where $\sigma \in [\theta_o - 2\rho, \theta_o - \rho]$ but make $a_{\hat{e}}^0$ a non-decreasing function of σ in this interval so that the result is constant near $\theta_o - \rho$ and is such that $a_{\hat{e}}^0 \gg -4 \ln(\rho_*)$ at $\sigma = \theta_o - \rho$. In particular, make this constant value greater than $-4 \ln(\rho_*)$ plus the supremum of the values of a_e^0 and $a_{e'}^0$ on the $|\sigma - \theta_*| < 2\rho$ portions of their parametrizing cylinders. This done, keep $a_{\hat{e}}^0$ constant on $[\theta_o - \rho, \theta_o + \rho]$, thus huge. Now, the larger this constant value for $a_{\hat{e}}^0$, the closer $K_e \cup K_{e'}$ is to an \mathbb{R} -invariant cylinder where it comes near $K_{\hat{e}}$. With this in mind, the argument used in Part 2 can be repeated in the case at hand to guarantee that the intersection points between $K_{\hat{e}}$ and the closure of the $\theta \in [\theta_o - 2\rho, \theta_o + 2\rho]$ part of $K_e \cup K_{e'}$ are transversal with +1 local intersection numbers if the constant value of $a_{\hat{e}}^0$ on the interval $[\theta_o - \rho, \theta_o + \rho]$ is sufficiently large.

Part 7 This part considers the fourth and fifth points in (4–70). In the case of the fourth point, the second observation of Part 4 can be employed to defined $a_{e'}^0$ and $a_{e''}^0$ where $\sigma \in [\theta_o + \rho, \theta_o + 2\rho]$ to the following effect: First, all $\theta \in [\theta_o, \theta_o + 2\rho]$ intersections between $K_{e'}$ and $K_{e''}$ occur where $\theta \in [\theta_o + \rho, \theta_o + 2\rho]$, are transversal, and have +1 local intersection number. Second, $a_{e'}^0 - a_{e''}^0 > \varepsilon_{e'} + \varepsilon_{e''}$ at $\sigma = \theta_o + \rho$.

Consider next the situation that is described in the fifth point of (4–70). To start, set the convention so that $[Q_{e'}, Q_{e''}] \leq 0$. Use e to denote the third of the incident edges to o , and use ρ_* to denote the common values of ρ_{e1} , $\rho_{e'0}$ and $\rho_{e''0}$.

There are two cases to consider. The first case has $a_{e''}^0 > a_{\hat{e}}^0 > a_{e'}^0$, at $\theta = \theta_o + 2\rho$. With the first two observations of Part 4 in mind, there is no cause for concern in this case if either $[Q_{\hat{e}}, Q_{e''}] \leq 0$ or $[Q_{\hat{e}}, Q_{e'}] \geq 0$. An argument that rules out the possibility that both inequalities fail simultaneously invokes an identity that concerns a set of four ordered pairs of real number: Denote the four ordered pairs as $\{A_k\}_{k=0,1,2,3}$, and here is the identity:

$$(4-73) \quad [A_1, A_2][A_3, A_0] + [A_2, A_3][A_1, A_0] + [A_3, A_1][A_2, A_0] = 0.$$

In this last equation, the bracket between a pair $A = (a, a')$ and another, $B = (b, b')$ is again defined by the rule $[A, B] = ab' - a'b$. This last identity is now applied with A_0 equal to the value of $(1 - 3 \cos^2 \theta, \sqrt{6} \cos \theta)$ at $\theta = \theta_o + 2\rho$, $A_1 = Q_{e'}$, $A_2 = Q_{e''}$ and $A_3 = Q_{\hat{e}}$. With these assignments, (4-73) is equivalent to the assertion that

$$(4-74) \quad [Q_{e'}, Q_{e''}] \alpha_{Q_{\hat{e}}} + [Q_{e''}, Q_{\hat{e}}] \alpha_{Q_{e'}} + [Q_{\hat{e}}, Q_{e'}] \alpha_{Q_{e''}} = 0.$$

Since the various versions of α_Q are positive and since $[Q_{e'}, Q_{e''}] \leq 0$, this equality rules out the possibility that both $[Q_{\hat{e}}, Q_{e''}] > 0$ and $[Q_{\hat{e}}, Q_{e'}] < 0$.

In the second case, $a_{e'}^0 > a_{\hat{e}}^0 > a_{e''}^0$, at $\sigma = \theta_o + 2\rho$. If $[Q_{e''}, Q_{\hat{e}}] \leq 0$ or if $[Q_{\hat{e}}, Q_{e'}] \leq 0$, then the second observation in Part 4 above can be applied to suitably modify $a_{\hat{e}}^0$ where $\sigma \in [\theta_o - 2\rho, \theta_o + 2\rho]$ so that the resulting version of $K_{\hat{e}}$ intersects the portion of the closure of $K_e \cup K_{e'} \cup K_{e''}$ where $\theta \in [\theta_o - 2\rho, \theta_o + 2\rho]$ transversally with +1 local intersection numbers. Of course, it may well be that both of these inequalities go the wrong way.

What follows explains the story when both $[Q_{e''}, Q_{\hat{e}}] > 0$ and $[Q_{\hat{e}}, Q_{e'}] > 0$. The first step here is to make $\varepsilon_e, \varepsilon_{e'}$ and $\varepsilon_{e''}$ constant where $|\sigma - \theta_o| < 2\rho$. Meanwhile, decrease $\varepsilon_{\hat{e}}$ between $\theta_o + 2\rho$ and $\theta_o + \rho$ so that the result is constant near $\theta_o + \rho$ and very much smaller than $(\varepsilon_e \rho_*)^6$. Extend $\varepsilon_{\hat{e}}$ to $[\theta_o - 2\rho_*, \theta_o + \rho]$ as this constant. Next, vary $a_{\hat{e}}^0$ as σ decreases from $\theta_o + \rho$ to $\theta_o + \frac{1}{2}\rho$ so that the result is constant and zero near $\theta_o + \frac{1}{2}\rho$. Then, extend $a_{\hat{e}}^0$ to the interval $[\theta_o - 2\rho_*, \theta_o + \frac{1}{2}\rho]$ as zero. It follows from the vertex o versions of the formulae in (4-1), (4-46) and (4-47) that $a_{e'}^0$ and $a_{e''}^0$ can be chosen so that $K_{e'}$, $K_{e''}$ and $K_{\hat{e}}$ are pairwise disjoint where $\theta \in [\theta_o + 2\rho_*, \theta_o + 2\rho]$.

In the case that $[Q_{\hat{e}}, Q_e] \neq 0$, vary $w_{\hat{e}}^0$ as σ decreases from $\theta_o + 2\rho$ to $\theta_o + \rho$ so that the result is constant and zero near $\theta_o + \rho$. In the case that $[Q_{\hat{e}}, Q_e] = 0$, introduce κ to denote the maximum value of the $Q = Q_e$ version of α_Q on the $\sigma \in [\theta_o - 2\rho, \theta_o]$ portion of K_e 's parametrizing cylinder. In this case, vary $w_{\hat{e}}^0$ as σ decreases from $\theta_o + 2\rho$ to $\theta_o + \rho$ so that the result is constant near $\theta_o + \rho$ and equal to $\frac{\pi}{2\kappa}$. Extend $w_{\hat{e}}^0$ as a constant to the interval $[\theta_o - 2\rho_*, \theta_o + \rho]$.

Now, if $\varepsilon_{\hat{e}}$ is very much smaller than $(\varepsilon_e \rho_*)^6$ on $[\theta_o - 2\rho_*, \theta_o + \rho]$ it then follows from the formulae in (4-42), (4-43), (4-46) and (4-47) that any intersection between $K_{\hat{e}}$ and the portion of the closure of $K_e \cup K_{e'} \cup K_{e''}$ where $\theta \in [\theta_o - 2\rho_*, \theta_o + 2\rho]$ occurs in the region of the latter surface that is parametrized via (3-2) and (4-48) by the radius ρ_* disk centered in the x - y plane. In fact, these intersections must occur at points whose x - y coordinates are within $4\varepsilon_{\hat{e}}$ of zero when $\varepsilon_{\hat{e}}$ is very small.

To continue, note that the origin in the x - y plane maps to a point in the closure of the union $K_e \cup K_{e'} \cup K_{e''}$ where the tangent plane is parallel to the tangent plane of $\mathbb{R} \times \gamma$, where $\gamma \subset S^1 \times S^2$ is a small portion of an integral curve of the Reeb vector field. Now, as an observation of Part 2 recalls, $K_{\hat{e}}$ intersects such a surface transversely with $+1$ local intersection numbers. This suggests that the intersections of the small $\varepsilon_{\hat{e}}$ version of $K_{\hat{e}}$ with the $\theta \in [\theta_o - 2\rho_*, \theta_o + 2\rho]$ portion of the closure of $K_e \cup K_{e'} \cup K_{e''}$ will be transversal with $+1$ local intersection numbers. It is left as an exercise with (4-33) and (4-48) to verify that such is indeed the case.

There is still more to do because as things stand now, both a_e^0 and $a_{\hat{e}}^0$ are zero where $\sigma = \theta_o - 2\rho_*$. Note that the $\theta \in [\theta_o - 2\rho, \theta_o - 2\rho_*]$ portions of the cylinders K_e and $K_{\hat{e}}$ will be disjoint as long as $\varepsilon_{\hat{e}}$ is very much smaller than ε and both the pairs $(a_{\hat{e}}^0, w_{\hat{e}}^0)$ and (a_e^0, w_e^0) are kept at their $\sigma = \theta_o - 2\rho_*$ values as σ decreases further to $\theta_o - 2\rho$. Even so, such an extension is not consistent at $\sigma = \theta_o - 2\rho$ with (4-69).

The desired extension of (a_e^0, w_e^0) keeps the latter constant on $[\theta_o - 2\rho, \theta_o - 2\rho_*]$. Meanwhile, the extension of $(a_{\hat{e}}^0, w_{\hat{e}}^0)$ employs the first and second observations in Part 4. To be more explicit, $a_{\hat{e}}^0$ is either increased or decreased from zero as σ decreases so that it is constant near $\theta_o - 2\rho$, but with a value that obeys $|a_{\hat{e}}^0| > \varepsilon_e + \varepsilon_{\hat{e}}$. In this regard, $a_{\hat{e}}^0$ is increased in the case that $[Q_{\hat{e}}, Q_e] < 0$ and it is decreased when $[Q_{\hat{e}}, Q_e] > 0$. In the case that $[Q_{\hat{e}}, Q_e] = 0$, either a decrease or increase is permissible. It then follows using (4-1) and (4-33) that such a version of $a_{\hat{e}}^0$ can be constructed to insure that $K_{\hat{e}}$ and K_e are disjoint at values of θ near $\theta_o + 4\rho$ and that they intersect transversally where $\theta \in [\theta_o - 2\rho, \theta_o - 2\rho_*]$ with $+1$ local intersection numbers. In this regard, note that this can be done in the case that $[Q_{\hat{e}}, Q_e] = 0$ without introducing any intersections between $K_{\hat{e}}$ and K_e .

Part 8 This last part of the subsection discusses the final point in (4-70). To explain the situation here, let o' and o'' denote distinct vertices of T with angle label 0, and let e' and e'' denote the corresponding incident edges. Suppose first that the \hat{A} label of o' has the form $(1, +, \dots)$ while o'' has either (1) or a label of the form $(1, -, \dots)$ from \hat{A} . In this case, there is no need for an arc crossing if $s_{e'} > s_{e''}$ at $\sigma = 2\rho$. Such is also true when $s_{e'} > s_{e''}$ at $\sigma = 2\rho$ and the \hat{A} label of o' is labeled by (1) while

that of o'' has the form $(1, -, \dots)$. However, in either case, the corresponding arcs must cross where $\sigma < 2\rho$ if $s_{e'} < s_{e''}$. Make such a crossing where $\sigma \in [\rho, 2\rho]$ and the second observation in Part 4 can be applied to choose $a_{e'}^0$ and $a_{e''}^0$ on $[\rho, 2\rho]$ so as to keep all $\theta < 2\rho$ intersections between $K_{e'}$ and $K_{e''}$ where $\theta \in [\rho, 2\rho]$, all transversal, and all with $+1$ local intersection number.

The case that both o' and o'' have a $(+1)$ label is the simplest of those where o' and o'' have the same sort of label from \hat{A} . In this case, $[Q_{e'}, Q_{e''}] = 0$, and so the first observation in Part 4 can be used to keep $K_{e'}$ disjoint from $K_{e''}$ if the corresponding e' and e'' arcs must cross at some point where $\sigma < 2\rho$.

In the case that o' and o'' both have either a $(1, +, \dots)$ label or a $(1, -, \dots)$ label, there are two subcases to consider. In the case that the e' and e'' arcs must cross where $\sigma \leq 2\rho$, then make such a crossing where $\sigma \in [\rho, 2\rho]$. In the case that $Q_{e'}$ is not proportional to $Q_{e''}$, the second observation in Part 4 is used to choose $a_{e'}^0$ and $a_{e''}^0$ on $[\rho, 2\rho]$ so as to keep all $\theta < 2\rho$ intersections between $K_{e'}$ and $K_{e''}$ where $\theta \in [\rho, 2\rho]$, all transversal, and all with local $+1$ intersection number. In the case that $[Q_{e'}, Q_{e''}] = 0$, then the first observation in Part 4 can be used to choose $(a_{e'}^0, w_{e'}^0)$ and $(a_{e''}^0, w_{e''}^0)$ on $[\rho, 2\rho]$ so as to keep $K_{e'}$ disjoint from $K_{e''}$ where $\theta < 2\rho$.

4.F The case when \wp has sets with two or more elements

This last subsection considers now the general case where the graph T is defined by a partition with sets that have more than one element. In what follows, \wp will denote such a partition with chosen cyclic orderings of its subsets. The discussion here is broken into four parts. The first three parts serve to specify the collection $\{(a_e, w_e)\}$ and the remaining part verifies that the collection meets all requirements.

Part 1 The purpose of this first part of the discussion is to construct from T a canonical moduli space graph to which the constructions in Subsection 4.C apply. This new graph is denoted by \hat{T} . The latter is isomorphic to T as an abstract graph via an isomorphism that preserves the labels of all edges and all but the bivalent vertices. The isomorphism also preserves the angles of the corresponding pairs of bivalent vertices.

Here is how the graphs T and \hat{T} differ: Suppose that $o \in T$ is a bivalent vertex, and let $\hat{o} \in \hat{T}$ denote its partner. The vertex o is labeled by a cyclic ordering of a partition subset, say $\wp_o \in \wp$. Meanwhile, \hat{o} is labeled by the data $(0, +, P_{\hat{o}})$ where $P_{\hat{o}}$ is the sum of the integer pairs from the elements in \wp_o .

A referral to Subsection 3.A shows that \hat{T} is a bona fide moduli space graph. Moreover, the discussion in Subsection 4.C applies to \hat{T} . Let $\{(a_{\hat{o}}, w_{\hat{o}})\}$ denote the resulting

data set for \widehat{T} as constructed in the preceding Sections 4.B–4.E. The required set $\{(a_e, w_e)\}$ for T is constructed now either starting directly from $\{(a_{\widehat{e}}, w_{\widehat{e}})\}$, or from the J -pseudoholomorphic curve that Theorem 3.1 provides from $\{(a_{\widehat{e}}, w_{\widehat{e}})\}$. The former approach is taken below and the latter is left as an exercise for the reader.

Part 2 Suppose that e is a given edge of T , and let o' and o denote the vertices from T that lie on e with the convention taken that $\theta_{o'} < \theta_o$. The corresponding edge, \widehat{e} , in \widehat{T} has the corresponding bounding vertices \widehat{o} and \widehat{o}' with $\theta_{\widehat{o}} = \theta_o$ and $\theta_{\widehat{o}'} = \theta_{o'}$. This understood, the functions (a_e, w_e) on $[\theta_{o'}, \theta_o] \times \mathbb{R}/(2\pi\mathbb{Z})$ are set equal to $(a_{\widehat{e}}, w_{\widehat{e}})$ at all points except in the case that one of o and o' is a bivalent vertex. In the latter case, the equality still holds except at values of (σ, v) that are very close to those of the missing point for the \widehat{T} -parametrization on the relevant boundary circle. In any event, the required data $\{\varepsilon_e, \rho_{e0}, \rho_{e1}, a_e^0, w_e^0, v_e^0\}$ for (a_e, w_e) are declared equal to their \widehat{T} counterparts.

To be more explicit about the differences between (a_e, w_e) and $(a_{\widehat{e}}, w_{\widehat{e}})$, suppose for the sake of argument that the vertex $o \in T$ is bivalent. Let \widehat{e} denote the respective partner to e in \widehat{T} . Let ρ denote the constant value of $\rho_{\widehat{e}1}$ where σ is near θ_o . For convenience of notation, assume that $v_e^0 = 0$ where σ is near θ_o . This understood, then the equality between (a_e, w_e) and $(a_{\widehat{e}}, w_{\widehat{e}})$ holds where $\sigma \sim \theta_o$ except possibly at values of (σ, v) with distance ρ^8 or less from the missing point on the $\sigma = \theta_o$ circle.

Part 3 To describe (a_e, w_e) near the point where $\sigma = \theta_0$ and $v = 0$, it is necessary to parametrize a neighborhood of the point $(\sigma = \theta_0, v = 0)$ in the parametrizing cylinder for e by the coordinates (r, τ) with $r \leq 3\rho$ and with $\tau \in [-\pi, 0]$. For this purpose, it proves necessary to introduce the complex coordinate $z \equiv r e^{i\tau}$. Also required is the choice of a parameter $\delta \in (0, \rho^7)$.

To obtain the desired parametrization, write $P_{\widehat{o}} = m_{\widehat{o}}(p, p')$ with p and p' the relatively prime integers defined via (1–7) by θ_o and with $m_{\widehat{o}} \geq 1$. Next, let n denote the number of elements in \wp_o and suppose that \wp_o has been given a linear ordering. Use the latter to label the integer pairs from its elements as $\{m_1(p, p'), \dots, m_n(p, p')\}$ with each m_j a positive integer. Thus, $\sum_j m_j = m_{\widehat{o}}$.

Let $0 = b_1 < b_2 < \dots < b_n < \delta\rho^8$ now denote a chosen set of very small real numbers, and introduce the complex function

$$(4-75) \quad \eta(z) = \beta \left(\frac{1}{\delta\rho} r \right) \prod_{1 \leq j \leq n} (z - b_j)^{m_j} + \left(1 - \beta \left(\frac{1}{\delta\rho} r \right) \right) z^{m_{\widehat{o}}}.$$

Note that with ρ small, and any choice for $\delta \in (0, \rho^7)$, the zeros of η consist of the points in the set $\{b_j\}$. An argument for the function η is needed on the lower half

plane and also at points on the real axis where $z \notin \{b_j\}$. To be precise, take the branch that gives

$$(4-76) \quad \arg(\eta) = m_{\hat{\delta}}\tau \text{ at points where } |z| \geq 2\rho^8.$$

With the preceding in hand, here is how to write σ and v in terms of r and τ :

$$(4-77) \quad \begin{aligned} \sigma &= \theta_{\hat{\delta}} + \varepsilon r \sin(\tau). \\ \hat{v} &= \left(1 - \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)}\right) \frac{1}{m_{\hat{\delta}}} \arg(\eta) + \frac{1}{\alpha_{Q_e}(\sigma)} r \cos(\tau). \end{aligned}$$

As in the analogous (4-11), the coordinate \hat{v} is \mathbb{R} -valued and reduces modulo 2π to v . Meanwhile, e' is the second of o 's incident edges.

With v_* as in (4-13) set

$$(4-78) \quad \begin{aligned} a_e &= -\beta_* \frac{1}{m_{\hat{\delta}}} \ln(|\eta|) + a_{\hat{\delta}} + \varepsilon(\beta_* + (1 - \beta_*) \cos(v_*)). \\ w_e &= -\varepsilon(1 - \beta_*) \sin(v_*) \\ &\quad + x_{\hat{\delta}}\beta' \left(\frac{1}{\alpha_{Q_e}} \beta_* \left(\frac{1}{m_{\hat{\delta}}} \arg(\eta) - \frac{1}{2\alpha_{Q_{e'}}} r \cos(\tau) \right) - \frac{1}{2\alpha_{Q_e}} (1 - \beta_*) v_* \right). \end{aligned}$$

Here, v is viewed as taking values in $[0, 2\pi]$. In addition, both σ and v are to be viewed where $\beta_* > 0$ as functions of r and τ .

To define $(a_{e'}, w_{e'})$ near the point where $\sigma = \theta_{\hat{\delta}}$ and $v = 0$ on the e' version of the parametrizing cylinder, first write the cylinder's coordinates σ and v near this point in terms of $r \in (0, 3\rho)$ and $\tau \in [0, \pi]$ using the rule

$$(4-79) \quad \begin{aligned} \sigma &= \theta_* + \varepsilon r \sin(\tau). \\ \hat{v} &= \left(\frac{\alpha_{Q_e}(\sigma)}{\alpha_{Q_{e'}}(\sigma)} - 1 \right) \frac{1}{m_{\hat{\delta}}} \arg(\eta) + \frac{1}{\alpha_{Q_{e'}}(\sigma)} r \cos(\tau). \end{aligned}$$

Here, $\arg(\eta)$ is again defined by (4-76). With (4-79) set, define

$$(4-80) \quad \begin{aligned} a_{e'} &= -\beta_* \frac{1}{m_{\hat{\delta}}} \ln(|\eta|) + a_{\hat{\delta}} + \varepsilon(\beta_* + (1 - \beta_*) \cos(v)). \\ w_{e'} &= -\varepsilon(1 - \beta_*) \sin(v) \\ &\quad + x_0\beta' \left(\frac{1}{\alpha_{Q_{e'}}} \beta_* \left(\frac{1}{m_{\hat{\delta}}} \arg(\eta) + \frac{1}{2\alpha_{Q_e}} r \cos(\tau) \right) + \frac{1}{2\alpha_{Q_e}} (1 - \beta_*) v \right). \end{aligned}$$

Here v is again viewed as taking values in $[0, 2\pi]$, and it with σ are viewed as functions of r and τ where $\beta_* > 0$.

Part 4 With what has been said in Section 3, Theorem 3.1 now follows from the following claim:

Very small δ versions of the definition just given of $\{(a_e, w_e)\}$ satisfy the criteria in (3–3).

The justification for this claim is given next in three steps.

Step 1 To start, let o denote a bivalent vertex in T and let e and e' denote the incident edges. The first point to note is that if δ is very small, then minor modifications of the arguments from Subsection 4.C prove that the change of variables from (σ, v) to (r, τ) is invertible on both the e and e' versions of the parametrizing cylinders.

Step 2 The closures of K_e and $K_{e'}$ fit together to define a smoothly immersed surface near points with $\theta \sim \theta_{\hat{\delta}}$ provided that the following is true: Let $v_1 \in (0, 2\pi)$ obey

$$(4-81) \quad v_1 \notin \left\{ \frac{1}{\alpha_{Q_e}(\theta_o)} b_1, \dots, \frac{1}{\alpha_{Q_e}(\theta_o)} b_n \right\}.$$

Then, there exists an integer pair, $N = (n, n')$, and extensions of the definitions of (a_e, w_e) and $(a_{e'}, w_{e'})$ to some neighborhood in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$ of (θ_*, v_1) so that (4–16) holds.

The verification of this condition proceeds just as in Subsection 4.C at points v_1 that differ by more than $2\delta\rho$ from either 0 or 2π . In the case that v_1 does not obey this condition, the equations in (4–77)–(4–80) directly give the required extensions of (a_e, w_e) and $(a_{e'}, w_{e'})$. This understood, there are various cases to consider depending on whether

$$(4-82) \quad \begin{aligned} &v_1 < 2\pi, \\ \text{or } &\frac{1}{\alpha_{Q_e}(\theta_{\hat{\delta}})} b_k < v_1 < \frac{1}{\alpha_{Q_e}(\theta_{\hat{\delta}})} b_{k+1} \quad \text{for some } k \in \{1, \dots, n-1\}, \\ \text{or } &v_1 > \frac{1}{\alpha_{Q_e}(\theta_o)} b_N. \end{aligned}$$

In the left most case, take the integer pair $N = Q_{e'}$. In the k 'th version of the middle case in (4–82), take $N = Q_{e'} + \sum_{1 \leq j \leq k} m_j(p, p')$, and in the right most case, take $N = Q_e$. Note that when comparing this last case with the case in Subsection 4.C, the $N = Q_e$ version of (4–16) is indistinguishable from the $N = 0$ version. It is left to the reader to confirm that (4–16) holds with the values of N as above. In this regard, note that the functions β_* and β' that appear in (4–13) and (4–77)–(4–80) are equal to 1 at the relevant points.

Step 3 There are three more issues to examine vis à vis (3–3). The first is that of the asymptotics as laid out in Definition 3.2. The verification that these are as required is a straightforward task using (3–2) with (4–77)–(4–80). The second is to verify that the data $\{(a_e, w_e)\}$ define a moduli space graph that is isomorphic to the given graph T . This is also straightforward and so the details are omitted.

The final issue concerns the singularities that lie in the closure of $\cup_e K_e$. A very small choice for δ also simplifies the analysis. To explain, let o denote any given bivalent vertex. Then (3–2) and (4–77)–(4–80) define a smooth map, ϕ_0 , from some multiply punctured version of the $|z| < 2\delta\rho$ disk in \mathbb{C} into $\mathbb{R} \times (S^1 \times S^2)$. If δ is very small, then s is huge on the image of each such ϕ_0 . Thus, any singularity in $\cup_e K_e$ that is not already present in its \hat{T} analog is a singularity of the image of some ϕ_0 . However, as explained next, each ϕ_0 is an embedding when δ is small. Hence, the closure of $\cup_e K_e$ meets all of (3–3)’s requirements.

To prove that ϕ_0 is an embedding, note first that points z and z' in the domain of ϕ_0 are mapped to the same point only if they have the same imaginary part. Indeed, otherwise, the images will have distinct θ values. Meanwhile, use of (3–2) with (4–77)–(4–80) finds that (4–24) still holds. Granted that this is the case, it then follows that the real parts of z and z' must agree as well if $\phi_0(z) = \phi_0(z')$. Thus, $z = z'$.

5 Proof of Theorem 1.3

The purpose of this last section is to prove Theorem 1.3. In this regard, the proof is obtained from Theorem 3.1 by demonstrating that \hat{A} has a positive line graph if and only if it has a moduli space graph. The implication from positive line graph to moduli space graph is proved in the first subsection. The reversed implication is proved in the second.

5.A From a positive line graph to a moduli space graph

Suppose that \hat{A} has a positive line graph, $L_{\hat{A}}$. The goal is to use the data from $L_{\hat{A}}$ to construct a labeled, contractible graph, T , as described in Subsection 3.A. The construction starts with the graph $L_{\hat{A}}$ and successively modifies it to obtain T . This construction of T occupies the seven parts of this subsection that follow.

Part 1 The purpose of this part of the subsection is to explain why the edges of a positive line graph obeys Constraint 2 in Subsection 3.A. Here is a formal statement to this effect:

Lemma 5.1 Let L denote a positive line graph for \widehat{A} , let $e \in L$ denote an edge, and let $\theta_o < \theta_1$ denote the angles that are assigned the vertices on e . Then

$$q_e'(1 - 3 \cos^2 \theta) - q_e \sqrt{6} \cos(\theta) \geq 0$$

at all $\theta \in [\theta_o, \theta_1]$ with equality if and only if θ is either θ_o or θ_1 and the corresponding vertex is monovalent with angle in $(0, \pi)$.

Proof of Lemma 5.1

The verification of the claim is given in five steps.

Step 1 This first step considers the claim at the vertex angles on e . To start, remark that the stated inequality holds at any bivalent vertex angle on e because the final point in (1–18) requires

$$(5-1) \quad pq_e' - p'q_e > 0$$

when (p, p') defines the angle of the vertex via (1–7). In this regard, keep in mind that (1–7) writes p as a positive multiple of $(1 - 3 \cos^2 \theta)$ and p' as the same multiple of $\sqrt{6} \cos \theta$.

When the vertex is monovalent with angle in $(0, \pi)$, then the condition in the lemma holds at the vertex angle in as much as the first and third points in (1–18) assert that (q_e, q_e') is proportional to the pair that defines the angle via (1–7).

Meanwhile, the required inequality can be seen to hold at $\theta = \theta_o$ when the latter is 0 by using (1–14) with the fourth point in (1–18). Likewise, (1–14) and the second point in (1–18) imply the lemma's assertion when $\theta = \theta_1$ and $\theta_1 = \pi$.

Step 2 This and the remaining steps verify the Lemma's inequality at the angles that lie strictly between θ_o and θ_1 . To start this process, let $Q = (q, q') \neq (0, 0)$ denote a given ordered pair of integers. Then the function $\alpha_Q(\theta)$ on $[0, \pi]$ vanishes only at that angles θ_Q and θ_{-Q} that are respectively defined via (1–7) by Q and by $-Q$. In this regard, note that θ_Q can be defined when $q < 0$ provided that $|\frac{q'}{q}| > \sqrt{3/2}$. Meanwhile, θ_{-Q} can be defined when $q > 0$ provided $|\frac{q'}{q}| > \sqrt{3/2}$. Thus, at least of one of θ_Q and θ_{-Q} exists in all cases and both exist only in the case that $|\frac{q'}{q}| < \sqrt{3/2}$.

To continue, note that the derivative of any given $Q = (q, q')$ version of α_Q is

$$(5-2) \quad \sqrt{6} \sin \theta (q + \sqrt{6} \cos \theta q').$$

In particular, this derivative is positive at $\theta = \theta_Q$ and negative at $\theta = \theta_{-Q}$. As a consequence, the desired inequality is satisfied for the given edge e if and only if one of the conditions listed next hold:

- (5-3)
- Both θ_{Q_e} and θ_{-Q_e} are defined, $\theta_{Q_e} < \theta_{-Q_e}$, and both $\theta_{Q_e} \leq \theta_o$ and $\theta_{o'} \leq \theta_{-Q_e}$.
 - Both θ_{Q_e} and θ_{-Q_e} are defined, $\theta_{-Q_e} < \theta_{Q_e}$ and either $\theta_{Q_e} \leq \theta_o$ or $\theta_{o'} \leq \theta_{-Q_e}$.
 - θ_{-Q_e} is not defined and $\theta_o \geq \theta_{Q_e}$.
 - θ_{Q_e} is not defined and $\theta_{o'} \leq \theta_{-Q_e}$.

These last constraints are analyzed with the help of the following observation: Suppose that P and Q are non-trivial integer pairs and suppose that both θ_P and θ_Q exist. Then $\theta_Q < \theta_P$ if and only if one of the following holds:

- (5-4)
- $q' \geq 0$, $p' \leq 0$ and at least one is non-zero.
 - If p' and q' have the same sign, then $q'p - qp' > 0$.

Step 3 This step and the next assume that $\theta_o > 0$ and $\theta_{o'} < \pi$. For this purpose, let (p_o, p'_o) and $(p_{o'}, p'_{o'})$ denote the respective integer pairs that define these angles via (1-7).

This step considers the case when the Q_e version of $|\frac{q'}{q}|$ is greater than $\sqrt{3/2}$. Thus, the third and fourth options in (5-3) are moot. The first option in (5-3) holds when $q_e' > 0$ and the second when $q_e' < 0$. Suppose first that $q_e' > 0$. If $p'_o \leq 0$ then the θ_o requirement is met by virtue of the first point in (5-4). Meanwhile, (5-1) together with the second point in (5-4) guarantee the θ_o requirement in the case that $p'_o > 0$. The first point in (5-4) guarantees the $\theta_{o'}$ requirement if $p'_{o'} \geq 0$. If $p'_{o'} < 0$, then the $\theta_{o'}$ requirement is guaranteed by the o' version of (5-1) using the second point in (5-4).

Now suppose that $q_e' < 0$. If $p'_o < 0$, then (5-1) and the second point in (5-4) guarantee the θ_o requirement. If $p'_o \geq 0$, then the θ_o requirement fails. In this case, the $\theta_{o'}$ requirement holds due to the second part of the final point in (1-18). Indeed, it would fail automatically were $p'_{o'} \leq 0$, but this is not allowed. On the other hand, if $p'_{o'} > 0$, then the $\theta_{o'}$ requirement follows from the o' version of (5-1) using the second point in (5-4).

Step 4 Granted that $\theta_o > 0$ and $\theta_{o'} < \pi$, this step assumes that the Q_e version of $|\frac{q'}{q}|$ is less than $\sqrt{3/2}$. In this case, neither of the first two points in (5-3) hold. In the case that $q_e > 0$, only the third point is possible to satisfy. If q_e' and p'_o have the same sign, then the requirement is met by virtue of (5-1) and the second point in (5-4). The requirement is also met if $p'_o \leq 0$ and $q_e' \geq 0$. Of course, the requirement can not be met if $p'_o > 0$ and $q_e' < 0$. However, as will now be explained, such signs for p'_o

and q_e' never appear. Indeed, were these the correct signs, then (5-1) would demand p_o to be negative and

$$(5-5) \quad \sqrt{\frac{3}{2}} > \left| \frac{q_e'}{q_e} \right| \geq \left| \frac{p_o'}{p_o} \right|.$$

However, this violates the condition in (1-7).

Suppose next that $q_e < 0$ and so only the fourth point in (5-3) is relevant. In the case that $p_{o'} \geq 0$ and $q_e' \geq 0$, then the desired inequality is insured by the first point in (5-4). If $p_{o'}$ and q_e' have different signs, the desired inequality is insured by (5-1) and the second point in (5-4). Meanwhile, the case where both $p_{o'} < 0$ and $q_e' < 0$ can not occur because (5-1) again requires that p_o is negative while satisfying (5-3).

Step 5 This step considers the case that $\theta_o = 0$. The argument for the case when $\theta_{o'} = \pi$ is omitted because it is identical to that given here save for some cosmetic changes. In the $\theta_o = 0$ case, it is necessary to verify either the second or fourth of the options in (5-3). In this regard, the first point is that $-Q_e$ in all cases defines an angle via (1-7). Indeed, this is a consequence of the fact noted in Step 1 that Lemma 5.1's inequality holds at $\theta = 0$.

In the case that Q_e defines an angle via (1-7), then Lemma 5.1's inequality at $\theta = 0$ requires that $q_e' < -(\sqrt{3/2})|q_e|$ and so the angle defined by Q_e via (1-7) is greater than that defined by $-Q_e$. Moreover, neither is less than $\theta_{o'}$. Indeed, the angle defined by Q_e via (1-7) must be greater than $\theta_{o'}$ since the condition in Lemma 5.1 holds near $\theta = 0$. Since $q_e' < 0$, this last point, the o' version of (5-1) and the second point in (5-4) are consistent only if $p_{o'} > 0$. Given that $p_{o'} > 0$, then the o' version of (5-1) and the second point in (5-4) establish the claim.

In the case that Q_e does not define an angle via (1-7), then $q_e < 0$ and the absolute value of the ratio of q_e' to q_e is less than $\sqrt{3/2}$. If q_e' and $p_{o'}$ have opposite signs, or if both are positive, then (5-4) guarantees the conditions for the fourth option in (5-3). On the other hand, in no case can both q_e' and $p_{o'}$ be negative when q_e is negative. Here is why: Were all three negative, then the o' version of (5-1) would require $p_{o'} < 0$ also. As such, this same version of (5-1) would declare the ratio of $p_{o'}$ to $p_{o'}$ to be less than that of q_e' to q_e . By assumption the latter is less than $\sqrt{3/2}$, and thus the former would be less than $\sqrt{3/2}$. But this conclusion with $p_{o'} < 0$ violates the given fact that $(p_{o'}, p_{o'})$ defines an angle via (1-7). \square

Part 2 Suppose that the maximal angle on $L_{\hat{A}}$ is less than π . Let θ_o denote this angle. This step describes a modified version of $L_{\hat{A}}$, a graph that is isomorphic to $L_{\hat{A}}$

except perhaps at angles that are very close to θ_o . This new graph has some number of added monovalent vertices, all with angle θ_o , these labeled by the $(0, -, \dots)$ elements from \hat{A} whose integer pairs define θ_o via (1-7). The modification of $L_{\hat{A}}$ is denoted below as T_1 .

To start the description, let $o \in L_{\hat{A}}$ denote the monovalent vertex with the largest angle on $L_{\hat{A}}$. Let \hat{e} denote the incident edge to o . In the case that the element $(0, -, -Q_{\hat{e}})$ is in \hat{A} , no modification occurs and $T_1 = L_{\hat{A}}$. If this 4-tuple is not in \hat{A} , then $-Q_{\hat{e}}$ is equal to a sum of some $n > 1$ pairs, $P_1 + \dots + P_n$, where each such pair is a positive multiple of the relatively prime pair that defines θ_o via (1-7) and where each $(0, -, P_k)$ is in \hat{A} .

To proceed in this case, choose $n - 1$ angles $\theta_1 < \theta_2 \dots \theta_{n-1} < \theta_o$ that are all greater than the smallest angle on e .

Modify $L_{\hat{A}}$ so that the resulting graph has $n - 1$ trivalent vertices at these chosen angles. Label the incident edges to the k' th such vertex as $e, e',$ and e'' using the convention that e connects the vertex o to a vertex with smaller angle, while e' and e'' connect to vertices with larger angle. Also take the convention that any given $k \leq n - 2$ version of the edge e'' is the same as the $(k + 1)'$ st version of the edge e . In particular, e' is capped with a $\theta = \theta_o$ monovalent vertex. This is also the case for e'' when $k = n - 1$.

Here are the integer pair assignments: In the case $k = 1$, the edge integer pair assignments are $Q_e = Q_{\hat{e}}, Q_{e'} = -P_1$ and $Q_{e''} = Q_{\hat{e}} + P_1$. In the case where $k > 1$, these pair assignments are $Q_e = Q_{\hat{e}} + \sum_{j < k} P_j$, while $Q_{e'} = -P_k$ and $Q_{e''} = Q_{\hat{e}} + \sum_{j \leq k} P_j$.

By virtue of Lemma 5.1, this graph obeys the moduli space graph constraints where it differs from L , thus at angles that are less than the minimal angle on \hat{e} . In this regard, Constraint 2 in Subsection 3.A is obeyed for T_1 because any given T_1 version of α_Q is a positive multiple of a corresponding $L_{\hat{A}}$ version that obeys the constraint in Lemma 5.1 for the relevant interval. Moreover, the $\theta = \theta_o$ monovalent vertices on T_1 are in 1-1 correspondence with the subset of $(0, -, \dots)$ elements in \hat{A} whose integer pair component defines θ_o via (1-7).

Part 3 This part of the construction describes the analogous operation on $L_{\hat{A}}$ when its largest angle is π . The resulting version of T_1 is isomorphic to $L_{\hat{A}}$ except at angles near π where it may have some trivalent vertices and more than one $\theta = \pi$ monovalent vertex. In particular, the labels of its $\theta = \pi$ monovalent vertices account for the $(-1, \dots)$ elements in \hat{A} .

To start, let n and n' denote the respective numbers of $(-1, -, \dots)$ and $(-1, +, \dots)$ elements in \hat{A} . If $n > 0$, label the $(-1, -, \dots)$ elements from \hat{A} from 1 to n , and if $n' > 0$, label the $(-1, +, \dots)$ elements from 1 to n' . Let $\{P_k^-\}_{1 \leq k \leq n}$ and $\{P_k^+\}_{1 \leq k \leq n'}$ denote the corresponding set of integer pair components.

Two trivial cases can be dispensed with straight away; that where \hat{A} has but a single $(-1, \dots)$ element and $c_- = 0$, and that where \hat{A} has no $(-1, \dots)$ elements and $c_- = 1$. No modification of $L_{\hat{A}}$ is necessary in either of these cases. Thus, T_1 is equal to $L_{\hat{A}}$ in both of these cases.

In the general case, the modification adds $n + n' + c_- - 1$ trivalent vertices with successively larger angles, all near π . To be precise here, the incident edges to any given vertex can be designated by e , e' and e'' so that e connects the vertex in question to one with a smaller angle while e' and e'' connect to vertices with larger angles. In all cases, the edge designated as e' is capped by a monovalent vertex with angle $\theta = \pi$. Such is also the case for the version of e'' that attaches to the trivalent vertex with the largest angle.

The edge labels for the incident edges to the trivalent vertices are obtained via an inductive process using the following rules: For the trivalent vertex with the smallest angle label, set $Q_e = \sum_k P_k^+ - \sum_k P_k^- - (0, c_-)$. Now, granted this, label this vertex as number 1 and label the remaining trivalent vertices by consecutive integers starting from 2 in order of increasing angle. Granted this numbering system, the first c_- of the trivalent vertices have $Q_{e'} = (0, -1)$. If $n' = 0$, then the remaining $n - 1$ have $Q_{e'} = P_j$ for the vertex numbered $c_- + j$ when $1 \leq j \leq n - 1$. If $n' > 0$, then $Q_{e'} = P_j$ for the vertex numbered $c_- + j$ when $1 \leq j \leq n$. Use $Q_{e'} = -P_j$ for the vertex numbered $c_- + n + j$ with $1 \leq j \leq n' - 1$. With regard to these assignments, note that the convention that $[Q_{e'}, Q_{e''}] \leq 0$ is not necessarily observed.

By virtue of Lemma 5.1, the graph T_1 obeys the moduli space graph conditions where it differs from $L_{\hat{A}}$, thus at the angles that are achieved on $L_{\hat{A}}$'s largest angled edge. Moreover, this new graph has the desired property: Its $\theta = \pi$ monovalent vertices account for all of the $(-1, \dots)$ elements in \hat{A} plus c_- elements with label (-1) .

By the way, the positivity requirement in Constraint 2 in Subsection 3.A can be deduced from the following observations: If θ is less than its value at the first trivalent vertex, then the relevant Q is that of an edge in $L_{\hat{A}}$ whose version of Constraint 2 in Subsection 3.A holds for the same value of θ . To argue for this constraint in the case that θ is near π , note that if the relevant version of $Q = (q, q')$ obeys $q' < (\sqrt{3/2})q$, then $\alpha_Q(\theta) > 0$ if θ is nearly π since (1-14) guarantees its positivity at $\theta = \pi$. This is to say that if the trivalent vertex angles are very near π , it is enough that $q_{e'} < (\sqrt{3/2})q_e$

for each incident edge to each trivalent vertex. This last requirement is met by virtue of the conditions in (1–14).

Part 4 Let θ_o now denote the minimal angle on $L_{\hat{A}}$. Of course, this is also the minimal angle on T_1 . If $\theta_o > 0$, then an upside down version of the discussion in Part 2 (a verbatim repetition save for some evident cosmetic changes) modifies T_1 by adding trivalent vertices with angles just slightly greater than θ_o and monovalent vertices at θ_o to construct a new graph, T_2 , with the following property: First there exists some $\delta > 0$ such that T_2 obeys the moduli space graph conditions at angles $\theta \in [\theta_o, \theta_o + \delta]$. Moreover, it has only trivalent vertices at angles in $(\theta_o, \theta_o + \delta)$ and it has as many $\theta = \theta_o$ monovalent vertices as there are $(0, -, \dots)$ elements in \hat{A} whose integer pairs define θ_o via (1–7). Moreover, these elements label the $\theta = \theta_o$ monovalent vertices in T_2 . Meanwhile, T_2 is isomorphic to T_1 at angles $\theta > \theta_o + \frac{1}{2}\delta$.

In the case that $\theta_o = 0$, the upside down version of the discussion of Part 3 modifies T_1 by adding only trivalent vertices at angles near 0 and monovalent vertices with angle equal to 0. This version of T_2 obeys the moduli space graph conditions where it differs from T_1 and thus where it differs from $L_{\hat{A}}$. Moreover, the set of $\theta = 0$ vertices in T_2 has a two subset partition: The first subset accounts for the $(1, \dots)$ elements in \hat{A} , and the second contains c_+ vertices with the label (1).

Part 5 Suppose now that o is a bivalent vertex in $L_{\hat{A}}$ whose angle is defined via (1–7) by one or more integer pairs from the collection of $(0, -, \dots)$ elements in \hat{A} . In this regard, consider only the case when there are no pairs from $(0, +, \dots)$ elements in \hat{A} that define θ_o . Described here is a modification to T_2 at angles very close to θ_o that replaces the bivalent vertex o with one or more trivalent and monovalent vertices that account for those $(0, -, \dots)$ elements in \hat{A} with integer pair giving θ_o . To begin, let e denote the incident edge to o on which θ_o is maximal, and let e' denote the incident edge on which θ_o is minimal. Fix some very small and positive number, δ . The modification proceeds in two steps.

The first step constructs a graph, \hat{T} , that lacks a bivalent vertex at θ_o , having one trivalent vertex at $\theta_o - \delta$ and one monovalent vertex at θ_o . To be more explicit, let $\hat{o} \in \hat{T}$ denote its trivalent vertex at $\theta_o - \delta$ and let \hat{e} , \hat{e}' and \hat{e}'' denote its three incident edges. The labeling convention here is such that \hat{o} has the largest angle of the two vertices on \hat{e} , and the smallest angle of the two vertices on \hat{e}' and on \hat{e}'' . In this regard, \hat{e}'' contains the added monovalent vertex with angle θ_o . In addition, as the integer pair assigned to \hat{e}'' is $Q_e - Q_{e'}$, so $-Q_{\hat{e}''}$ is the sum of the integer pairs that define θ_o via (1–7).

To describe the rest of \widehat{T} , agree to designate the three components of $\widehat{T} - \widehat{o}$ as $\widehat{T}_{\widehat{e}}$, $\widehat{T}_{\widehat{e}'}$ and $\widehat{T}_{\widehat{e}''}$, so that the subscript indicates that the component contains the interior of its labeling edge. Let T_{2e} and $T_{2e'}$ denote the analogously labeled components of $T_2 - o$. Then $\widehat{T}_{\widehat{e}}$ is isomorphic to T_{2e} and $\widehat{T}_{\widehat{e}'}$ to T_{2e} . With regards to such isomorphisms, the convention taken here and subsequently is that an isomorphism between labeled graphs with some open edges must preserve all labeling of vertices and edges, but it need not match the angles of any 'absent' vertices.

If \widehat{A} has a single $(0, -, \dots)$ element whose integer pair defines θ_o , then the graph T_3 is set equal to \widehat{T} . If there is more than one such element, the graph \widehat{T} is further modified by employing the construction in Part 2 with the edge \widehat{e}'' playing the role of $L_{\widehat{A}}$. Thus, the modification replaces the edge \widehat{e}'' with a subgraph whose monovalent vertices account for all of the $(0, -, \dots)$ elements in \widehat{A} with integer pairs that define θ_o via (1-7). This subgraph has one less trivalent vertex than it has monovalent vertices. These can be assigned distinct angles, all between θ_o and $\theta_o - \delta$.

As will now be explained, any sufficiently small δ version of the graph T_3 obeys the moduli space graph conditions where it differs from T_2 and thus where it differs from $L_{\widehat{A}}$. To begin, remark that the positivity of the $Q = Q_e$ and $Q = Q_{e'}$ versions of the function α_Q imply that these functions are positive for small δ on the edges \widehat{e} and \widehat{e}' of \widehat{T} . If the $Q = Q_{\widehat{e}}$ version of α_Q is positive on $[\theta_o - \delta, \theta_o)$ and vanishes at θ_o , then the arguments from Part 2 settle the claim that T_3 obeys the moduli space graph conditions where it differs from T_2 . Granted this, remark that the $Q = Q_{\widehat{e}}$ version of α_Q is zero at θ_o because $Q_{\widehat{e}} = Q_e - Q_{e'}$, and according to the fifth point in (1-18), this pair is -1 times a pair that defines θ_o via (1-7). Moreover, as explained subsequent to (5-2), the derivative of α_Q at its zero is negative when the angle of the zero is θ_{-Q} . Since this is the case at hand, the $Q = Q_{\widehat{e}}$ version of α_Q is positive on the half open interval $[\theta_o - \delta, \theta_o)$ when δ is small.

Part 6 Suppose next that o is a bivalent vertex whose angle is defined via (1-7) by an integer pair from some $(0, +, \dots)$ element in \widehat{A} . Consider first the case when no integer pairs from $(0, -, \dots)$ elements in \widehat{A} define this angle. In this case, the modification to T_2 amounts to adding some data to the label of the bivalent vertex o so as to make the label that of a bivalent vertex in a moduli space graph. In this regard, o 's label must be a partition subset for some partition of the set of $(0, +, \dots)$ elements whose integer pairs define θ_o via (1-7). Take the one set partition and assign o this set.

Consider now the case where θ_o is also defined via (1-7) by integer pairs from both $(0, +, \dots)$ elements in \widehat{A} and $(0, -, \dots)$ elements in \widehat{A} . Let P_+ denote the sum of the integer pairs from the former set and let P_- denote the sum of those from the latter.

Note that both P_+ and P_- define θ_o via (1-7). What follows describes a modification of T_2 so as to obtain a graph, T_3 , with one bivalent vertex with angle θ_o , one or more monovalent vertices with angle θ_o , and some trivalent vertices with angles near θ_o . This graph T_3 will satisfy the moduli space graph conditions where it differs from T_2 and its bivalent and monovalent vertices will account for all of the $(0, \dots)$ elements in \hat{A} whose integer pair component defines θ_o via (1-7). The modification here is very similar to that described in Part 5. In particular, there are two steps, with the first modifying T_2 by adding a single trivalent vertex at an angle just less than θ_o and adding a monovalent vertex with angle θ_o . This preliminary modification also has a bivalent vertex with angle θ_o . Let \hat{T} denote this new graph. If δ is positive but very small, then \hat{T} can be constructed so that it has a trivalent vertex, \hat{o} , with angle $\theta_o - \delta$. The three incident edges, \hat{e} , \hat{e}' , and \hat{e}'' are such that \hat{o} has the larger of the angles of the vertices on \hat{e} . As before, the component $\hat{T}_{\hat{e}} \subset \hat{T} - \hat{o}$ is isomorphic to the component $T_{2e} \subset T_2 - o$. Meanwhile, \hat{e}' connects \hat{o} to the bivalent vertex at angle θ_o in \hat{T} while \hat{e}'' connects \hat{o} to the monovalent vertex with angle θ_o . The label for \hat{e}'' is $-P_-$, while that for \hat{e}' is $Q_e + P_-$. Note that the open graph $\hat{T}_{\hat{e}'} - \hat{e}'$ is isomorphic to $T_{2e'}$.

The graph \hat{T} must now be modified so that the result, T_3 , obeys the moduli space graph conditions where it differs from T_2 . First of all, this involves replacing \hat{e}'' by a subgraph with some number of monovalent vertices and one less number of trivalent vertices, with the subgraph chosen so that its monovalent vertices have angle θ_o and account for those $(0, -, \dots)$ elements in \hat{A} whose integer pair defines θ_o via (1-7). This procedure is exactly that used in the previous step to go from the latter's \hat{T} to T_3 . Note that Constraint 2 in Subsection 3.A is obeyed on all of the edges in this subgraph. Indeed, the argument for this is a verbatim repetition of the one that proves the analogous claim in Part 5. The final task in the construction of T_3 is to grant a label to the bivalent vertex at angle θ_o . In this case, the label must be a partition of the set of those $(0, +, \dots)$ elements in \hat{A} whose integer pair defines θ_o via (1-7). As before, take the 1-set partition. Note that this is forced by the fact that $Q_{\hat{e}'} - Q_{e'} = P_+$ which is the sum of the integer pairs from this same set of elements. By the way, note that the \hat{e}' version of Constraint 2 in Subsection 3.A is obeyed when δ is small by virtue of two facts: First, the $Q = Q_e$ version of α_Q is bounded away from zero on $[\theta_o - \delta, \theta_o]$. Second, the $Q = P_+$ version of α_Q is zero at θ_o and so is very small on this interval when δ is small.

Part 7 Apply the constructions in Part 5 and Part 6 simultaneously to all of the bivalent vertices. The result is a moduli space graph for \hat{A} .

5.B From moduli space graph to positive line graph

Now suppose that \hat{A} has a moduli space graph, $T_{\hat{A}}$. The goal is to obtain from $T_{\hat{A}}$ a positive line graph for \hat{A} . This is accomplished in a sequential fashion using the various ‘moves’ that are described in Part 1, below. These moves are used to eliminate trivalent vertices. To picture this process, imagine a trivalent vertex as the point in a partially unzipped zipper where two edges are joined as one. The modifications amount to closing in a sequential fashion all of these zippers. Part 2 of the subsection provides the details for how these moves are used.

The modifications to $T_{\hat{A}}$ will result in graphs that are not moduli space graphs. Even so, these graphs have labeled edges and vertices that obey certain constraints. These constraints are listed below, and a graph that obeys them is deemed a ‘positive graph’.

A positive graph, T , is a connected, contractible graph with at least one edge and with labeled vertices and edges. The vertices of T are either monovalent, bivalent or trivalent. Each is labeled with an angle in $[0, \pi]$. These angles are constrained as follows:

- (5–6)
- The two vertices on any given edge have distinct angles.
 - The angle of any given multivalent vertex is neither the largest nor the smallest of the angles of the vertices on its incident edges.
 - Any vertex angle in $(0, \pi)$ is defined via (1–7) by an integer pair.

Each edge of T is labeled by an integer pair. If e is an edge, then $Q_e = (q_e, q_{e'})$ denotes its integer pair. These are constrained as follows:

- (5–7)
- Let $o \in T$ denote a monovalent vertex with angle in $(0, \pi)$ and let e denote its incident edge. Then $\pm Q_e$ defines θ_o via (1–7) with the $+$ sign taken if and only if θ_o is the smaller of the two angles of the vertices.
 - Let $o \in T$ denote a bivalent vertex and let e and e' denote its incident edges. If $Q_e \neq Q_{e'}$, then either $Q_e - Q_{e'}$ or $Q_{e'} - Q_e$ defines θ_o via (1–7).
 - Let $o \in T$ denote a trivalent vertex, and let e , e' and e'' denote its incident edges. Then $Q_e = Q_{e'} + Q_{e''}$ with the convention that θ_o lies between the angle of the vertex opposite o on e and the angles of the vertices opposite o on both e' and e'' .
 - Let e denote any given edge of T and let $\theta_o < \theta_1$ denote the angles that are assigned the vertices on e . Then

$$q_{e'}(1 - 3 \cos^2 \theta) - q_e \sqrt{6} \cos(\theta) \geq 0$$

at all $\theta \in [\theta_o, \theta_1]$ with equality if and only if θ is either θ_o or θ_1 , the angle in question is in $(0, \pi)$, and the corresponding vertex is monovalent.

Each positive graph that appears below is related to the asymptotic data set \hat{A} in a manner that is described momentarily. For this purpose, it is necessary to assign an integer pair to each vertex with angle in $(0, \pi)$. If o is such a vertex, then P_o is used to denote its integer pair. Here are the assignments: If o is monovalent, then $P_o = \pm Q_e$, where e denotes o 's incident edge and where the $+$ sign is taken if and only if θ_o is the smaller of the two angles of e 's vertices. If o is bivalent, then $P_o = Q_e - Q_{e'}$ where e and e' are o 's incident edges with the convention here that θ_o is the larger of the two angles of the vertices on e . If o is a trivalent vertex, set $P_o = 0$.

What follows describes how \hat{A} enters the picture:

- (5–8) • *The sum of the integer pairs that are associated to the edges with a $\theta = \pi$ vertex is obtained from \hat{A} by the following rule: First, subtract the sum of the integer pairs from the $(-1, -, \dots)$ elements in \hat{A} from the sum of those from the $(-1, +, \dots)$ elements, and then subtract $(0, c_-)$ from the result.*
- *The sum of the integer pairs that are associated to the edges with a $\theta = 0$ vertex is obtained from \hat{A} by the following rule: First, subtract the sum of the integer pairs from the $(1, +, \dots)$ elements in \hat{A} from the sum of those from the $(1, -, \dots)$ elements and then subtract $(0, c_+)$ from the result.*
- *Let $\theta \in (0, \pi)$. Then, the sum of the integer pairs that are associated to the bivalent vertices at angle θ minus the sum of those pairs that are associated to the monovalent vertices at angle θ is obtained from \hat{A} by the following rule: Subtract the sum of the integer pairs from the $(0, -, \dots)$ elements in \hat{A} that define θ via (1–7) from the sum of the integer pairs from the $(0, +, \dots)$ elements in \hat{A} that defined θ via (1–7).*

A positive graph that obeys (5–8) is called a ‘positive graph for \hat{A} ’. According to Lemma 5.1, a positive line graph for \hat{A} is a linear positive graph for \hat{A} , that is, one with no trivalent vertices. Lemma 5.2 below proves the converse. Note that $T_{\hat{A}}$ itself is a positive graph for \hat{A} .

Part 1 To set the stage here and in Part 2, let T now denote a given positive graph. Let o denote a trivalent vertex in T and let e , e' and e'' denote the three incident edges to o with the usual convention taken to distinguish e . This is to say that the angle θ_o lies between the angle of the vertex opposite o on e and both the angle of the vertex opposite o on e' and that of the angle opposite o on e'' . The edges e' and e'' are distinguished when $Q_{e'}$ is not proportional to $Q_{e''}$ by making $[Q_{e'}, Q_{e''}] \equiv q_{e'}q_{e''}' - q_{e'}'q_{e''}$ negative. Also, keep in mind the following two conventions from the previous subsection that

concern the three components of $T - o$: The first is with regards to their labeling, this as T_e , $T_{e'}$ and $T_{e''}$ with the labeling such that the closure of any one of the three contains its labeling edge. The other convention concerns the notion of an isomorphism between one of these components and some other non-compact graph with labeled vertices and edges. In particular, the isomorphism must send vertices to vertices and edges to edges so as to respect the labeling. However, such an isomorphism has no need to respect the angle of the absent vertex on the open edge.

With these conventions set, what follows in this Part 1 are the moves that are used to modify a given positive graph for \hat{A} so as to eliminate the trivalent vertices. In all cases, the modified graph is a positive graph for \hat{A} . There are two versions to each move listed below, one for the case that e connects the given vertex to a vertex with a larger angle, and one for the case that the connection is to a vertex with a smaller angle. Only the first version is presented since the two versions differ only cosmetically.

Note that the first three moves modify the original graph so as to decrease the angle that is assigned to the given trivalent vertex. (In the omitted version where e connects to a vertex with smaller angle, the corresponding moves will increase the angle of the given trivalent vertex.) The remaining four moves modify the graph so as to eliminate the given trivalent vertex.

To set the stage for the moves, agree to let o denote the trivalent vertex in question and let θ_o denote its original angle assignment. In what follows, $\theta_{\hat{o}}$ denotes the larger of the two angles that label the vertices that lie opposite o on e' and e'' . Keep in mind that $\theta_{\hat{o}}$ is less than θ_o . A distinguishing feature of the geometry here is that neither the $Q = Q_{e'}$ nor $Q_{e''}$ versions of α_Q can vanish on $(\theta_{\hat{o}}, \theta_o)$. Since $Q_e = Q_{e'} + Q_{e''}$, the $Q = Q_e$ version of α_Q is also positive on $(\theta_{\hat{o}}, \theta_o)$. As a result, T can be modified without either compromising the positive graph conditions or changing its topology by giving o any angle in $(\theta_{\hat{o}}, \theta_o)$.

With the preceding understood, the first three moves describe cases where T is modified so that the result has a trivalent vertex with angle just less than $\theta_{\hat{o}}$.

Move 1 Assume here that $\theta_{\hat{o}}$ labels just one vertex on $e' \cup e''$ and that this vertex is bivalent. In this case T is modified to produce a new positive graph for \hat{A} , this denoted by T_* . The graph T_* has a trivalent vertex, o_* , with angle θ_* just less than $\theta_{\hat{o}}$ and a bivalent vertex with angle $\theta_{\hat{o}}$. The integer pair component of the latter's label is the same as that of the $\theta_{\hat{o}}$ labeled vertex on $e' \cup e''$.

To continue the description, note that o_* has incident edges e_* , e'_* , e''_* where e_* is the only one of the three that connects o_* to a vertex with a larger angle. The latter is the aforementioned bivalent vertex with angle $\theta_{\hat{o}}$. Write the components of $T_* - o_*$

as T_{*e} , $T_{*e'}$ and $T_{*e''}$. These graphs are related to T_e , $T_{e'}$ and $T_{e''}$ as follows: In the case that e'' has the $\theta_{\hat{o}}$ labeled vertex, then, $T_{e''} - e''$ and $T_{*e''}$ are isomorphic as non-compact graph with labeled vertices and edges. Meanwhile $T_{e'}$ and $T_{*e'}$ are likewise isomorphic, as are the pair T_e and $T_{*e} - e_*$. The analogous isomorphisms hold when e' has the $\theta_{\hat{o}}$ bivalent vertex.

Move 2 This move is relevant to the case that both e' and e'' have bivalent vertices with angle $\theta_{\hat{o}}$. In this case, T is again modified to produce a new positive graph for \hat{A} . This graph has a trivalent vertex at angle just less than $\theta_{\hat{o}}$ and a single bivalent vertex at angle $\theta_{\hat{o}}$ that sits on the incident edge e_* to o_* . Here, the notational convention for the incident edges to o_* are as in Move 1. The integer pair component of the label for this bivalent vertex is the sum of the integer pairs that label the bivalent vertices on e' and e'' . In this regard, the components $T_{*e'}$ and $T_{*e''}$ of $T_* - o_*$ are respectively isomorphic to $T_{e'} - e'$ and $T_{e''} - e''$ from $T - o$. Meanwhile, $T_{*e} - e_*$ is isomorphic to T_e . The verification that the version of T_* with θ_* nearly $\theta_{\hat{o}}$ is a positive graph for \hat{A} requires only the verification of the fourth point in (5–7) for the edges that touch e_* . In this regard, the positivity of the relevant versions of α_Q follow from the positivity at $\theta_{\hat{o}}$ of the $Q_{e'}$ and $Q_{e''}$ versions.

Move 3 This move is relevant to when there is a single vertex $\hat{o} \in e' \cup e''$ with angle $\theta_{\hat{o}}$, that this vertex is trivalent, and that it has a single incident edge that connects it to a vertex with angle less than $\theta_{\hat{o}}$. Agree to relabel the edge between o and \hat{o} as e_0 . Now, label the other two incident edges to o as e_1 and e_2 with the convention that $e = e_1$, while labeling the other two incident edges to \hat{o} as e_3 and e_4 with the convention that e_4 connects \hat{o} to a vertex with angle less than $\theta_{\hat{o}}$.

Let T_* denote the modified graph. It has a trivalent vertex, o_* , at angle just less than $\theta_{\hat{o}}$, and another, \hat{o}_* at angle just greater than θ_o . These two are connected by an edge, e_{*0} . The remaining two incident edges to o_* connect the latter to vertices with smaller angles, while the remaining two incident edges to \hat{o}_* connect it to vertices with larger angles. The integer pair assigned to e_{*0} is the sum of those assigned to e_2 and e_4 , this being also the sum of those assigned to e_1 and e_3 . Meanwhile, $T_* - e_{*0}$ is isomorphic to $T - e_0$.

The fact that T_* is a positive graph for \hat{A} follows directly with the verification of the fourth condition in (5–7). And, the latter follows when o_* has angle nearly $\theta_{\hat{o}}$ and \hat{o}_* has angle nearly θ_o from the fact that the inequality is strictly obeyed by e_0 on $[\theta_{\hat{o}}, \theta_o]$, and by the other incident edges to o and \hat{o} on the relevant intervals in $[0, \pi]$.

The remaining moves describe modifications to T that either remove a given trivalent vertex, or replace it with either one monovalent vertex or one bivalent vertex.

Move 4 Suppose here that only one vertex on $e' \cup e''$ has angle $\theta_{\hat{o}}$, and that the latter is monovalent. This move explains how T is modified so as to replace the trivalent and monovalent vertices with a single bivalent vertex.

To begin the story, remark that $\theta_{\hat{o}}$ must be greater than 0 as $e' \cup e''$ has a vertex with a smaller angle. Moreover, because the e' version of α_Q is positive at $\theta_{\hat{o}}$ and because $[Q_{e'}, Q_{e''}] < 0$, the vertex on $e' \cup e''$ with angle $\theta_{\hat{o}}$ must sit on e'' . The graph T is modified at angles near $\theta_{\hat{o}}$ by removing $e'' - o$ so as to replace o with a bivalent vertex in the modified graph. To elaborate, let T_* denote the new graph. It has a bivalent vertex, o_* , at angle $\theta_{\hat{o}}$. Use e_* and e'_* to denote its incident edges with the convention that e_* connects o_* to a vertex with angle less than $\theta_{\hat{o}}$. Label the components of $T_* - o_*$ as T_{*e} and $T_{*e'}$ with the convention that T_{*e} contains the interior of e_* . Then T_{*e} is isomorphic to the component $T_{e'}$ of $T - o$, and $T_{*e'}$ is isomorphic to the component T_e . Because the $Q = Q_{e'}$ version of α_Q is positive at $\theta_{\hat{o}}$, this is also the case for the $Q = Q_e$ version. This then implies that T_* obeys the fourth constraint in (5–7). Thus, T_* is a positive graph since it also obeys the first three conditions in (5–7). Meanwhile, the T_* version of (5–8) is obeyed by virtue of the fact that the integer pair for the vertex o_* is $-Q_{e''}$.

Move 5 This move is relevant when both e' and e'' have vertices with angle $\theta_{\hat{o}}$ with one bivalent and the other monovalent. In this regard, note that $Q_{e'}$ and $Q_{e''}$ can not lie on the same line in \mathbb{R}^2 in this case. Thus, with the $[Q_{e'}, Q_{e''}] < 0$ convention, the bivalent vertex must lie on e' . In this case, the graph T is modified by eliminating both the trivalent vertex and the monovalent vertex on e'' . To elaborate here, let T_* again denote the new graph. It has a bivalent vertex with angle $\theta_{\hat{o}}$. Let o_* denote the latter, and let e_* and e'_* denote its incident edges with the convention that e_* connects o_* to a vertex with a smaller angle label. Then the component T_{*e} of $T_* - o_*$ is isomorphic to the component $T - e'$ that contains vertices with angles less than $\theta_{\hat{o}}$. Meanwhile, the component $T_{*e'}$ of $T_* - o_*$ is isomorphic to the component T_e of $T - o$.

With the labeling as describe, T_* is a positive graph for \hat{A} . Indeed, the only substantive issue here is that raised by the fourth point in (5–7). In this regard, the e'_* version of this inequality holds because it is strictly obeyed by the $Q = Q_e$ version of α_Q at θ_* . Meanwhile, the e_* version of the inequality holds because it holds for the version that is labeled by the edge that connects the bivalent vertex on e' to a smaller angled vertex.

Move 6 This and the remaining moves are relevant to the cases where $\theta_{\hat{o}}$ is the angle of a monovalent vertex on e' and a monovalent vertex on e'' . This move considers the case where the angle is in $(0, \pi)$.

In this case, a new graph, T_* , is obtained from T by removing $(e' \cup e'') - o$, thus replacing o by a monovalent vertex with angle $\theta_{\hat{o}}$. To elaborate, the graph T_* has a monovalent vertex, o_* , with angle $\theta_{\hat{o}}$. In addition, $T_* - o_*$ is isomorphic to T_e . In this regard, keep in mind that both $Q_{e'}$ and $Q_{e''}$ define $\theta_{\hat{o}}$ via (1–7). Thus, they are positive multiples of each other. This understood, then Q_e must also define $\theta_{\hat{o}}$ via (1–7). The fourth point in (5–7) holds on T_* because it holds on T and because the $Q = Q_e$ version of α_Q vanishes at $\theta_{\hat{o}}$ and is positive on $(\theta_{\hat{o}}, \theta_o]$.

Move 7 This considers the case that $\theta_{\hat{o}} = 0$. In all of these cases, T_* has a $\theta = 0$ monovalent vertex whose complement is isomorphic to T_e . The verification that T_* is a positive graph for \hat{A} is straightforward and so left to the reader.

Part 2 This last part of the subsection explains how the preceding moves can be used to construct a positive line graph from \hat{A} given its original moduli space graph $T_{\hat{A}}$. In this regard, keep in mind that $T_{\hat{A}}$ is a positive graph for \hat{A} . The forthcoming Lemma 5.2 asserts that a linear positive graph for \hat{A} is a positive line graph. This understood, the task at hand is to modify $T_{\hat{A}}$ using Move 1–Move 7 so as to obtain a positive graph for \hat{A} that lacks trivalent vertices.

Lemma 5.2 *A linear, positive graph for \hat{A} is a positive line graph for \hat{A} .*

Proof of Lemma 5.2 The only substantive issue here is that raised by the fourth point in (1–18). In this regard, suppose that L is a linear, positive graph for \hat{A} and that e is an edge in L . The condition on the positivity of $pq_{e'} - p'q_e > 0$ in the case that (p, p') is an integer pair that defines the angle of a bivalent vertex on \hat{e} follows directly from fourth constraint in (5–7). Now, suppose that $q_{e'} < 0$. Let o and o' denote the two vertices on e with the convention that o 's angle is less than that of o' . Assume first that both these vertices have angles in $(0, \pi)$. If such is the case, then $p'_{o'} > 0$ requires $p'_o > 0$ because both angles must be smaller than $\frac{\pi}{2}$ if the larger is.

For the sake of argument, suppose that $p'_o > 0$ but that $p'_{o'} < 0$. The $Q = Q_e$ version of θ_Q must be greater than θ_o since the former is greater than $\frac{\pi}{2}$ and the latter less than $\frac{\pi}{2}$. However, according to the second point in (5–4), this same θ_Q is less than $\theta_{o'}$. Thus, α_Q vanishes between θ_o and $\theta_{o'}$ and this violates the last item in (5–7). The other possibility is that where either θ_o is 0 or $\theta_{o'} = \pi$. In the case that $\theta_o = 0$, then the argument just given has the $Q = Q_e$ version of θ_Q less than $\theta_{o'}$ when it exists. If available, the $Q = -Q_e$ version of this angle is also less than $\theta_{o'}$ by virtue of the first point in (5–4). In either case, this means that the $Q = Q_e$ version of α_Q has a zero in $(0, \theta_{o'})$. In the case that $\theta_{o'} = \pi$, the $Q = Q_e$ version of θ_Q is greater than θ_o

when it exists. This is another consequence of the first point of (5–4). When available, the $Q = -Q_e$ version of θ_Q is also greater than θ_o ; this a consequence of the second point in (5–4). Thus the $Q = Q_e$ version of α_Q has a zero in (θ_o, π) .

The remainder of this section describes an algorithm that uses Move 1–Move 7 from Part 1 to change $T_{\hat{A}}$ into a linear, positive graph for \hat{A} and thus produce the desired positive line graph for \hat{A} . The algorithm has four steps. \square

Step 1 Suppose that T is any given positive graph for \hat{A} . Let V_+ denote the set of trivalent vertices with only one incident edge that contains a larger angle vertex. Likewise, define V_- to be the set of trivalent vertices with only one incident edge that contains a smaller angle vertex. Let n_{T_+} denote the number of elements in V_+ and let n_{T_-} denote the corresponding number in V_- . If V_+ is empty, go to Step 3. If not, let $o \in V_+$ be a vertex whose angle is the smallest of those from the vertices in V_+ . Move 1–Move 7 can now be used to successively modify T so that the result, T' , is a positive graph for \hat{A} with $n_{T'_+} = n_{T_+} - 1$ and with $n_{T'_-} = n_{T_-}$. Indeed, Move 1–Move 3 successively decrease the angle of the relevant trivalent vertex, this by an amount that is bounded uniformly away from zero. Thus, only finitely many applications of Move 1–Move 3 are possible. The subsequent moves all eliminate a trivalent vertex. In any event, when T' is produced, go to Step 2.

Step 2 Repeat Step 1 using T' now instead of T .

Note that Step 1 and Step 2 ultimately result in a positive graph for \hat{A} whose version of V_+ is empty and whose version of V_- has the same number of elements as does that of $T_{\hat{A}}$.

Step 3 The input to this step is a positive line graph for \hat{A} whose version of V_+ is empty. Let T now denote the latter. If V_- is also empty, then stop because T is the desired linear graph. If $V_- \neq \emptyset$, let o denote the trivalent graph with the largest angle. Successively apply the up side down versions of Move 1–Move 7 to o . The result is a new, positive graph for \hat{A} with no elements in its version of V_+ and one less trivalent vertex than T . Denote this graph by T' . Go to Step 4.

Step 4 Repeat Step 3 using T' now instead of T .

References

- [1] **F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, E Zehnder**, *Compactness results in symplectic field theory*, *Geom. Topol.* 7 (2003) 799–888 MR2026549
- [2] **Y Eliashberg, A Givental, H Hofer**, *Introduction to symplectic field theory*, *Geom. Funct. Anal.* (2000) 560–673 MR1826267
- [3] **D T Gay, R Kirby**, *Constructing symplectic forms on 4-manifolds which vanish on circles*, *Geom. Topol.* 8 (2004) 743–777 MR2057780
- [4] **H Hofer**, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, *Invent. Math.* 114 (1993) 515–563 MR1244912
- [5] **H H W Hofer**, *Dynamics, topology, and holomorphic curves*, *Doc. Math.* (1998) 255–280 MR1648034
- [6] **H Hofer**, *Holomorphic curves and dynamics in dimension three*, from: “Symplectic geometry and topology (Park City, UT, 1997)”, *IAS/Park City Math. Ser.* 7, Amer. Math. Soc., Providence, RI (1999) 35–101 MR1702942
- [7] **H Hofer, K Wysocki, E Zehnder**, *Properties of pseudo-holomorphic curves in symplectisations. II. Embedding controls and algebraic invariants*, *Geom. Funct. Anal.* 5 (1995) 270–328 MR1334869
- [8] **H Hofer, K Wysocki, E Zehnder**, *Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics*, *Ann. Inst. H. Poincaré Anal. Non Linéaire* 13 (1996) 337–379 MR1395676 See [9] for correction
- [9] **H Hofer, K Wysocki, E Zehnder**, *Correction to: “Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics”*, *Ann. Inst. H. Poincaré Anal. Non Linéaire* 15 (1998) 535–538 MR1632925
- [10] **H Hofer, K Wysocki, E Zehnder**, *Properties of pseudoholomorphic curves in symplectizations. III. Fredholm theory*, from: “Topics in nonlinear analysis”, *Progr. Nonlinear Differential Equations Appl.* 35, Birkhäuser, Basel (1999) 381–475 MR1725579
- [11] **K Honda**, *An openness theorem for harmonic 2-forms on 4-manifolds*, *Illinois J. Math.* 44 (2000) 479–495 MR1772422
- [12] **M Hutchings, M Sullivan**, *The periodic Floer homology of a Dehn twist*, *Algebr. Geom. Topol.* 5 (2005) 301–354 (electronic) MR2135555
- [13] **D McDuff**, *The local behaviour of holomorphic curves in almost complex 4-manifolds*, *J. Differential Geom.* 34 (1991) 143–164 MR1114456
- [14] **D McDuff, D Salamon**, *J-holomorphic curves and quantum cohomology*, *University Lecture Series* 6, American Mathematical Society, Providence, RI (1994) MR1286255
- [15] **C H Taubes**, *The geometry of the Seiberg–Witten invariants*, *Doc. Math.* (1998) 493–504 MR1648099

- [16] **C H Taubes**, *The structure of pseudo-holomorphic subvarieties for a degenerate almost complex structure and symplectic form on $S^1 \times B^3$* , *Geom. Topol.* 2 (1998) 221–332 MR1658028
- [17] **C H Taubes**, *Seiberg–Witten invariants, self-dual harmonic 2-forms and the Hofer–Wysocki–Zehnder formalism*, from: “Surveys in differential geometry”, *Surv. Differ. Geom.*, VII, Int. Press, Somerville, MA (2000) 625–672 MR1919438
- [18] **C H Taubes**, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$* , *Geom. Topol.* 6 (2002) 657–814 (electronic) MR1943381
- [19] **R Ye**, *Gromov’s compactness theorem for pseudo holomorphic curves*, *Trans. Amer. Math. Soc.* 342 (1994) 671–694 MR1176088

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

`chtaubes@math.harvard.edu`

Proposed: Rob Kirby

Seconded: Peter Ozsváth, Yasha Eliashberg

Received: 6 April 2004

Accepted: 9 May 2006