

Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: Moduli space parametrizations

CLIFFORD HENRY TAUBES

This is the second of 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. The first article in this series described the local structure of the moduli spaces and gave existence theorems. This article describes a stratification of the moduli spaces and gives explicit parametrizations for the various strata.

53D30; 53C15, 53D05, 57R17

1 Introduction

This is the second of two articles that describe the moduli spaces of multiply punctured, pseudoholomorphic spheres in $\mathbb{R} \times (S^1 \times S^2)$ as defined using the almost complex structure, J , for which

$$\begin{aligned} J \cdot \partial_s &= \frac{1}{\sqrt{6(1 + 3 \cos^4 \theta)^{1/2}}} ((1 - 3 \cos^2 \theta) \partial_t + \sqrt{6} \cos \theta \partial_\varphi) \\ (1-1) \quad J \cdot \partial_\theta &= \frac{1}{\sqrt{6(1 + 3 \cos^4 \theta)^{1/2}}} \left(-\sqrt{6} \cos \theta \sin \theta \partial_t + (1 - 3 \cos^2 \theta) \frac{1}{\sin \theta} \partial_\varphi \right). \end{aligned}$$

Here, s is the Euclidean coordinate on the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$, while t is the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate on the S^1 factor and $(\theta, \varphi) \in [0, \pi] \times \mathbb{R}/(2\pi\mathbb{Z})$ are the usual spherical coordinates on the S^2 factor. A change of coordinates shows that this almost complex structure is well defined near the $\theta = 0$ and $\theta = \pi$ cylinders, and that the latter are pseudo-holomorphic with a suitable orientation.

This almost complex structure arises naturally in the following context: A smooth, compact and oriented 4 dimensional manifold with non-zero second Betti number has a 2-form that is symplectic on the complement of its zero set, this a disjoint set of embedded circles (see, eg Taubes [13], Honda [10], Gay and Kirby [3]). There are indications that certain sorts of closed, symplectic surfaces in the complement of the zero set of such a 2-form code information about the differential structure of the 4-manifold

(Taubes [12]). Meanwhile, [13] describes the complement of any such vanishing circle in one of its tubular neighborhoods as diffeomorphic to $(0, \infty) \times (S^1 \times S^2)$ via a diffeomorphism that makes all of the relevant symplectic surfaces pseudoholomorphic with respect to either the almost complex structure in (1–1) or its push-forward by the two-fold covering map that sends (t, θ, φ) to the same point as $(t + \pi, \pi - \theta, -\varphi)$.

The investigation of these symplectic surfaces in the differential topology context lead the author to study the pseudoholomorphic subvarieties in $\mathbb{R} \times (S^1 \times S^2)$ in general with a specific focus on the multiply punctured spheres. This series of articles reports on this study. In particular, the first article in this series [15] defined a topology on the set of pseudoholomorphic subvarieties and defined them as elements of moduli spaces that are much like those introduced by Hofer [4; 5; 6], Hofer–Wysocki–Zehnder [7; 8; 9] and Eliashberg–Hofer–Givental [2] in a closely related context. The multiply punctured sphere moduli spaces were then proved to be smooth manifolds and formulae were given for their dimensions. Finally, [15] describes necessary and sufficient conditions to guarantee the existence of a moduli space component with prescribed $|s| \rightarrow \infty$ asymptotics in $\mathbb{R} \times (S^1 \times S^2)$. The details of much of this are summarized below for the benefit of those who have yet to see [15].

This second article in the series describes the multiply punctured sphere moduli spaces in much more detail as it describes the set of components and provides something of a parametrization for each component.

Note that the article [14], a prequel to this series, provided this information for certain disk, cylinder and thrice-punctured sphere moduli space components.

1.A The background

The almost complex structure in (1–1) is compatible with the symplectic form

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

where α is the following contact 1-form on $S^1 \times S^2$:

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

In this regard, the standard product metric on $\mathbb{R} \times (S^1 \times S^2)$ is related to the bilinear form $\omega(\cdot, J\cdot)$ using the rule

$$(1-4) \quad \frac{1}{\sqrt{6(1 + 3 \cos^4 \theta)^{1/2}}} e^{\sqrt{6}s} \omega(\cdot, J\cdot) = ds^2 + dt^2 + d\theta^2 + \sin^2 \theta d\varphi^2.$$

On a related note, the form ω is self-dual and harmonic with respect to the product metric in (1–4), this a consequence of the various strategically placed factors of $\sqrt{6}$.

Following Hofer, a pseudoholomorphic subvariety in $\mathbb{R} \times (S^1 \times S^2)$ is defined to be a closed subset, C , that lacks isolated points and has the following properties:

- (1–5)
- 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’ when the removal of a finite set of points disconnects it.

As explained in [14, Section 2], any pseudoholomorphic subvariety intersects some sufficiently large R version of the $|s| \geq R$ portion of $\mathbb{R} \times (S^1 \times S^2)$ as an embedded union of disjoint, half–open cylinders. In particular, if E denotes such a cylinder, then the restriction of s to E defines a smooth, proper function with no critical points. Moreover, the limit as $|s| \rightarrow \infty$ of the constant s slices of E converge in $S^1 \times S^2$ pointwise as a multiple cover of some ‘Reeb orbit’, this an embedded, closed orbit of the vector field

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

A subset $E \subset C$ of the sort just described is said to be an ‘end’ of C .

In the case that an irreducible subvariety is not a $\theta = 0$ or $\theta = \pi$ cylinder, considerations of the convergence of the constant s slices of any given end to the limiting Reeb orbit led in [15, Section 1] to the association of a 4–tuple of asymptotic data to the end in question. To elaborate, such a 4–tuple has the form $(\delta, \varepsilon, (p, p'))$ with δ either -1 , 0 or 1 ; with ε one of the symbols $\{-, +\}$; and with (p, p') being an ordered pair of integers that are not both zero and are constrained to obey:

- (1–7)
- $|\frac{p'}{p}| > \sqrt{\frac{3}{2}}$ in the case that $p < 0$ and $\delta = 0$.
 - If $\delta = \pm 1$, then $p < 0$.
 - If $\delta = 1$, then $\frac{p'}{p} < -\sqrt{\frac{3}{2}}$ if $\varepsilon = +$, and $\frac{p'}{p} > -\sqrt{\frac{3}{2}}$ if $\varepsilon = -$.
 - If $\delta = -1$, then $\frac{p'}{p} > \sqrt{\frac{3}{2}}$ if $\varepsilon = +$, and $\frac{p'}{p} < \sqrt{\frac{3}{2}}$ if $\varepsilon = -$.

To explain the meaning of the 4–tuple assignment, first note that the angle θ is constant on any integral curve of the vector field in (1–6). This understood, the case $\delta = 0$ signifies that this constant value of θ on the limit Reeb orbit for the given end is an angle in $(0, \pi)$. Meanwhile, the case $\delta = 1$ signifies that this constant value is 0 , and

$\delta = -1$ signifies that the $|s| \rightarrow \infty$ limit of θ on the end is π . In all of these cases, the appearance of $+$ for the parameter ε signifies that $s \rightarrow \infty$ on the given end, while the appearance of $-$ for ε signifies that $s \rightarrow -\infty$ on the end. An end with $\varepsilon = +$ is said to be a ‘concave side’ end, and an $\varepsilon = -$ end is said to be a ‘convex side’ end. Finally, the integers p and p' are the respective integrals of the 1-forms $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ around any given constant $|s|$ slice of the end granted that the latter are oriented by the pull-back of the 1-form $-\alpha$.

In the case that $\delta = 0$, the $|s| \rightarrow \infty$ limit of θ on the end in question is determined by the integer pair (p, p') as this limiting angle obeys

$$(1-8) \quad p'(1 - 3 \cos^2 \theta) - p\sqrt{6} \cos \theta = 0 \text{ and } p' \cos \theta \geq 0.$$

In this regard, keep in mind that any ordered pair, (p, p') , of integers with at least one non-zero defines a unique angle in $(0, \pi)$ via (1-8) in the case that $p \geq 0$. In the case that $p < 0$, such a pair defines an angle in $(0, \pi)$ if and only if $|\frac{p'}{p}| > \sqrt{\frac{3}{2}}$. The angle so defined is also unique.

In the case that $\delta = \pm 1$, the pair (p, p') determine the rate of convergence of the angle θ to its limiting value of 0 or π . To be explicit, results from [14, Sections 2 and 3] can be used to verify that

$$(1-9) \quad \sin \theta = \sqrt{6\hat{c}} e^{(\sqrt{\frac{3}{2}} + \delta \frac{p'}{p})s} (1 + o(1))$$

as $|s| \rightarrow \infty$ on the given end with \hat{c} some positive constant.

Aside from the collection of 4-tuples from its ends, the subvariety also defines a pair, (ζ_-, ζ_+) , of non-negative integers, these being the respective numbers (counting multiplicity) of intersections between the $\theta = 0$ and $\theta = \pi$ cylinders and the subvariety. In this regard, keep in mind that these cylinders are pseudoholomorphic. Also keep in mind a consequence of the analysis in [14, Section 2]: There are at most a finite set of intersection points between any two distinct, irreducible pseudoholomorphic subvarieties in $\mathbb{R} \times (S^1 \times S^2)$. Finally, keep in mind that any intersection point between distinct pseudoholomorphic subvarieties has positive local intersection number (McDuff [11]).

Granted what has just been said, an irreducible pseudoholomorphic subvariety that is not a $\theta = 0$ or $\theta = \pi$ cylinder defines an example of what is called here an asymptotic data set, this the set whose elements consist of the ordered pair (ζ_-, ζ_+) and the collection of 4-tuples from its ends. In general, the term ‘asymptotic data set’ refers to a certain sort of set that consists of one ordered pair of non-negative integers, here (ζ_-, ζ_+) ; and some number of 4-tuples that have the form $(\delta, \varepsilon, (p, p'))$ where

δ is -1 , 0 or 1 , ε is either $-$ or $+$, and (p, p') is an ordered pair of integers that obey the rules in (1-7). A set, \hat{A} , as just described is deemed an asymptotic data set in the event that it obeys five additional constraints. Here are the first two:

$$(1-10) \quad \sum_{(\delta, \varepsilon, (p, p')) \in \hat{A}} \varepsilon p = 0 \quad \text{and} \quad \sum_{(\delta, \varepsilon, (p, p')) \in \hat{A}} \varepsilon p' + \zeta_+ + \zeta_- = 0.$$

The constraint in (1-10) follows when \hat{A} comes from an irreducible pseudoholomorphic subvariety by using Stokes' theorem when considering the integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ on any sufficiently large, constant $|s|$ slice of the subvariety.

Here is the third constraint:

$$(1-11) \quad \text{If } \hat{A} \text{ has two 4-tuples and } \zeta_+ = \zeta_- = 0, \text{ then the 4-tuple integer pairs are relatively prime.}$$

Indeed, any \hat{A} with two 4-tuples and $\zeta_+ = \zeta_- = 0$ labels a moduli space of pseudoholomorphic cylinders. All of the latter are described in [14, Section 4] and none violate (1-11).

The two remaining constraints refer to a set, $\Lambda_{\hat{A}}$, of distinct angles in $[0, \pi]$ that is defined from \hat{A} . This set contains the angle 0 if $\zeta_+ > 0$ or if \hat{A} contains a $(1, \dots)$ element, it contains the angle π if $\zeta_- > 0$ or if \hat{A} contains a $(-1, \dots)$ element, and its remaining angles are those that are defined through (1-8) by the integer pairs from the $(0, \dots)$ elements in \hat{A} . This understood, here are the remaining constraints:

$$(1-12) \quad \begin{aligned} &\bullet \Lambda_{\hat{A}} \text{ is a single angle if and only if the latter is in } (0, \pi), \hat{A} \text{ has only two 4-tuples, and these are } (0, +, P) \text{ and } (0, -, P) \text{ with } P = (p, p') \text{ being a relatively prime integer pair.} \\ &\bullet \text{ If } \Lambda_{\hat{A}} \text{ has more than one angle, then neither of its maximal or minimal elements is defined via (1-8) by the integer pairs of any } (0, +, \dots) \text{ element from } \hat{A}. \end{aligned}$$

The constraints in (1-12) are consequences of two facts noted in [14, (4.21)] and in [15, Section 2.E] about the pull-back of the angle θ when the latter is non-constant on an irreducible pseudoholomorphic subvariety: First, this pull-back has neither local maxima nor local minima in $(0, \pi)$. Second, if its $s \rightarrow \infty$ limit on a concave side end is neither 0 nor π , then its restriction to any constant s slice of the end takes values both greater and less than this limit.

Given an asymptotic data set \hat{A} , use $\mathcal{M}_{\hat{A}}$ to denote the set of irreducible, pseudoholomorphic, multiply punctured spheres that give rise to \hat{A} . Grace $\mathcal{M}_{\hat{A}}$ with the topology

where a basis for the neighborhoods of any given $C \in \mathcal{M}_{\hat{A}}$ is given by sets that are indexed by positive real numbers where the version that is defined by $\kappa > 0$ consists of those $C' \in \mathcal{M}_{\hat{A}}$ with

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

The following theorem restates [15, Theorem 1.2 and Proposition 2.5]:

Theorem 1.1 *If \hat{A} is an asymptotic data set and $\mathcal{M}_{\hat{A}}$ is non-empty, the latter is a smooth manifold of dimension*

$$N_+ + 2(N_- + \hat{N} + \zeta_- + \zeta_+ - 1),$$

where N_+ , N_- and \hat{N} are the respective numbers of $(0, + \dots)$, $(0, - \dots)$ and $(\pm 1, \dots)$ elements in \hat{A} .

Note that by virtue of (1-12), the sum $N_- + \hat{N} + \zeta_- + \zeta_+$ is at least 1 for any asymptotic data set. As explained in [14], the story when $N_- + \hat{N} + \zeta_- + \zeta_+ = 1$ is as follows:

$$(1-14) \quad \mathcal{M}_{\hat{A}} \text{ consists of } \mathbb{R}\text{-invariant cylinders when } N_- + \hat{N} + \zeta_- + \zeta_+ = 1.$$

The components of the moduli spaces of pseudoholomorphic, \mathbb{R} -invariant cylinders are of two sorts. First, there are two single element components, the $\theta = 0$ cylinder and the $\theta = \pi$ cylinder. The second sort of component is a circle. In this regard, each circle component is labeled by a relatively prime pair of integers, (p, p') , with at least one non-zero and with $|\frac{p'}{p}| > \sqrt{\frac{3}{2}}$ in the case that $p < 0$. The circle consists of the subvarieties of the form $\mathbb{R} \times \gamma$ where $\gamma \subset S^1 \times S^2$ is an orbit of the vector field $\hat{\alpha}$ where θ is given by the pair (p, p') via (1-8). The circle parameter on the moduli space can be taken to be the constant value in $\mathbb{R}/(2\pi\mathbb{Z})$ along γ of $p't - p\varphi$.

A description of $\mathcal{M}_{\hat{A}}$ in the case that $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ is given in the next subsection. Subsection 1.C starts the story that is told in subsequent sections about the general case.

1.B The space $\mathcal{M}_{\hat{A}}$ when its dimension is $N_+ + 2$

This subsection is divided into two parts. The first provides an explicit parameterization for $\mathcal{M}_{\hat{A}}$ in the case that $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ and so $\dim(\mathcal{M}_{\hat{A}}) = N_+ + 2$. The second part describes some aspects of the subvarieties that map near the frontier in the parametrizing space.

Before starting, note that the story in this case is simpler than for cases when the dimension is greater than $N_+ + 2$ by virtue of the fact that function θ lacks critical points

with θ values in $(0, \pi)$ on the model curves of subvarieties in the $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ versions of $\mathcal{M}_{\hat{A}}$. This is a consequence of [15, Proposition 2.11]. The same proposition implies that the tautological map from any such model curve to $\mathbb{R} \times (S^1 \times S^2)$ is an immersion that is transversal to the $\theta = 0$ and $\theta = \pi$ loci.

Part 1 The story here and also in the $N_- + \hat{N} + \zeta_- + \zeta_+ > 2$ case presented later requires the introduction of a graph with labeled vertices and labeled edges. The simplicity in the case at hand stems from the fact that the graph in question is linear.

To describe the graph in the $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ case, first agree to view a linear graph as a closed interval in $[0, \pi]$ whose vertices define a finite subset that includes the endpoints. Let $T^{\hat{A}} \subset [0, \pi]$ denote the graph in question. Each vertex of $T^{\hat{A}}$ is labeled with its corresponding angle, and the set of these vertex angles is the set $\Lambda_{\hat{A}}$ as defined in the previous subsection. Meanwhile, each edge of $T^{\hat{A}}$ is labeled by an integer pair using the rules that follow. In the statement of these rules and subsequently, the letter ‘ e ’ is used to signify an edge, and Q_e is used to signify an integer pair that is associated to the edge e . Here are the rules:

- (1–15) • *If e starts the graph at an angle in $(0, \pi)$, then Q_e is the integer pair from the $(0, -, \dots)$ element in \hat{A} that define this minimal angle via (1–8).*
- *If 0 is the smallest angle on e , then one and only one of the following situations arise: There is a $(1, -, \dots)$ element in \hat{A} , or there is a $(1, +, \dots)$ element in \hat{A} , or $\zeta_+ = 1$. In these respective cases, Q_e is the integer pair from the $(1, -, \dots)$ element in \hat{A} , or minus the integer pair from the $(1, +, \dots)$ element in \hat{A} , or $(0, -1)$ when $\zeta_+ = 1$.*
 - *Let o denote a bivalent vertex, let θ_o denote its angle, and let e and e' denote its incident edges with the convention that θ_o is the largest angle on e . Then $Q_e - Q_{e'}$ is the sum of the integer pairs from the $(0, +, \dots)$ elements in \hat{A} that define θ_o via (1–8).*

According to [15, Theorem 1.3], the moduli space $\mathcal{M}_{\hat{A}}$ is non-empty if and only if the following condition holds:

- (1–16) *Let \hat{e} denote an edge in $T^{\hat{A}}$. 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 all of*

- (a) $q_{\hat{e}'} < 0$,
- (b) *neither vertex on \hat{e} has angle 0 or π and*

- (c) the version of (1–8)’s integer p' for one of the vertex angles on \hat{e} is positive,

hold, then both versions of p' are positive.

This condition is assumed in what follows.

The graph $T^{\hat{A}}$ is now used as a blueprint of sorts to define a space, $O^{\hat{A}}$, that plays a fundamental role in what follows. The definition of the latter follows in three steps.

Step 1 Let $\hat{A}_+ \subset \hat{A}$ denote the set of $(0, +, \dots)$ elements and let

$$(1-17) \quad \mathbb{R}^{\hat{A}} \subset \text{Maps}(\hat{A}_+; \mathbb{R})$$

denote the subspace where distinct elements in \hat{A}_+ have distinct images in $\mathbb{R}/(2\pi\mathbb{Z})$ when their respective integer pair components define the same angle in (1–8).

Let \mathbb{R}_- denote an extra copy of the affine line \mathbb{R} .

Step 2 View the space $\text{Maps}(\hat{A}_+; \mathbb{Z})$ as a group using addition in \mathbb{Z} to give the composition law. Of interest is an action of this group on

$$(1-18) \quad \mathbb{R}_- \times \text{Maps}(\hat{A}_+; \mathbb{R}).$$

This action is trivial on \mathbb{R}_- . To describe the action on the second factor in (1–18), note first that $\text{Maps}(\hat{A}_+; \mathbb{Z})$ is generated by a set $\{z_u : u \in \hat{A}_+\}$ where $z_u(\hat{u}) = 0$ unless $\hat{u} = u$ in which case $z_u(u) = 1$. The action of z_u on a given $x \in \text{Maps}(\hat{A}_+; \mathbb{R})$ is as follows: First, $(z_u \cdot x)(\hat{u}) = x(\hat{u})$ in the case that the integer pair from \hat{u} defines an angle via (1–8) that is less than that defined by u ’s integer pair. Such is also the case when the two angles are equal and $\hat{u} \neq u$. Meanwhile, $(z_u \cdot x)(u) = x - 2\pi$. Finally, if the integer pair from \hat{u} defines an angle that is greater than that defined by u ’s integer pair, then $(z_u \cdot x)(\hat{u})$ is obtained from $x(\hat{u})$ by adding

$$(1-19) \quad -2\pi \frac{p_u' p_{\hat{u}} - p_u p_{\hat{u}}'}{q_{\hat{e}}' p_{\hat{u}} - q_{\hat{e}} p_{\hat{u}}'}$$

where (p_u, p_u') is the integer pair entry of the element u , $(p_{\hat{u}}, p_{\hat{u}}')$ is that of \hat{u} , and \hat{e} is the edge in $T^{\hat{A}}$ whose largest angle vertex has the angle that is defined via (1–8) by this same $(p_{\hat{u}}, p_{\hat{u}}')$.

The action just described of $\text{Maps}(\hat{A}_+; \mathbb{Z})$ commutes with the action of $\mathbb{Z} \times \mathbb{Z}$ that is defined as follows: An integer pair $N = (n, n')$ acts as translation by $-2\pi(n'q_e - nq_e')$ on \mathbb{R}_- ; here (q_e, q_e') is the integer pair that is associated to the edge in $T^{\hat{A}}$ with

the smallest angle vertex. Meanwhile, N acts on any $x \in \text{Maps}(\hat{A}_+; \mathbb{R})$ so that the resulting map, $N \cdot x$, sends any given $u \in \hat{A}_+$ to the point that is obtained from $x(u)$ by adding

$$(1-20) \quad -2\pi \frac{n' p_{\hat{u}} - n p'_{\hat{u}}}{q e' p_{\hat{u}} - q e p'_{\hat{u}}}.$$

Step 3 Granted the definitions in the preceding steps, set

$$(1-21) \quad O^{\hat{A}} \equiv [\mathbb{R}_- \times \mathbb{R}^{\hat{A}}] / [(\mathbb{Z} \times \mathbb{Z}) \times \text{Maps}(\hat{A}_+; \mathbb{Z})].$$

As is explained in Subsection 3.A, $O^{\hat{A}}$ is a smooth manifold. Of particular interest here is its quotient by the evident action of the group, $\text{Aut}^{\hat{A}}$, whose elements are the 1–1 self maps of \hat{A}_+ that permute only elements with identical 4–tuples. To elaborate, this group action is induced from the action on $\mathbb{R}^{\hat{A}}$ whereby a given $\iota \in \text{Aut}^{\hat{A}}$ acts by composition. Thus,

$$(1-22) \quad (\iota \cdot x)(u) = x(\iota(u)) \text{ for each } u \in \hat{A}_+.$$

In what follows, $\hat{O}^{\hat{A}} \subset O^{\hat{A}}$ denotes the subset of points with trivial $\text{Aut}^{\hat{A}}$ stabilizer.

The following theorem explains the relevance of these constructs:

Theorem 1.2 *There is a diffeomorphism from $\mathcal{M}_{\hat{A}}$ to $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$.*

There are diffeomorphisms between $\mathcal{M}_{\hat{A}}$ and $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ that grant direct geometric interpretations to the \mathbb{R} parameter and to all of the parameters in the $\hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ factor. Indeed, Theorem 3.1 asserts that the diffeomorphism in question can be chosen so as to intertwine the \mathbb{R} action on $\mathcal{M}_{\hat{A}}$ that translates the subvarieties by a constant amount along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times S^2)$ with the \mathbb{R} action on its factor in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. Furthermore, Propositions 3.4 and 3.5 describe diffeomorphisms of the latter sort that interpret the image in $\hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ for any given subvariety in terms of the $|s| \rightarrow \infty$ limits on the subvariety and its behavior near the 0 and π loci.

For example, what follows describes how this $\hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ data determines Reeb orbit limits for a particular choice of diffeomorphism in Theorem 1.2. To set the stage, keep in mind that θ is constant on any Reeb orbit and that the Reeb orbits at a given $\theta \in (0, \pi)$ are distinguished as follows: Let (p, p') denote the relatively prime integer pair that defines θ via (1–8). The $\mathbb{R}/(2\pi\mathbb{Z})$ valued function $p\varphi - p't$ is constant on any Reeb orbit at angle θ , and its values distinguish these orbits.

The set of $s \rightarrow \infty$ limits in $(0, \pi)$ of θ 's restriction to any subvariety in $\mathcal{M}_{\hat{A}}$ is the set of angles for the bivalent vertices in $T^{\hat{A}}$. Here is how to obtain information about the corresponding Reeb orbit limits of the $s \rightarrow \infty$ slices: Each element in \hat{A}_+ whose integer pair defines a given angle in $\Lambda_{\hat{A}}$ labels a map from $O^{\hat{A}}$ to $\mathbb{R}/(2\pi\mathbb{Z})$ that is obtained as follows: First take the $\mathbb{R}/(2\pi\mathbb{Z})$ reduction of the image via a given map in $\mathbb{R}^{\hat{A}}$ of the element in \hat{A}_+ and then multiply the result by $(q_e' p - q_e p')$ where e denotes the edge that ends at the given vertex and (p, p') denotes the relatively prime pair of integers that defines the angle in question. Now, let n denote the number of \hat{A}_+ elements whose integer pair defines the given angle. The unordered set of n points in $\mathbb{R}/(2\pi\mathbb{Z})$ so defined by the image in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ of any chosen subvariety from $\mathcal{M}_{\hat{A}}$ is precisely the set of values for $p\varphi - p't$ on those Reeb orbits at the given angle that arise as $s \rightarrow \infty$ limits of the constant s slices of the chosen subvariety.

Part 2 The space $\hat{O}^{\hat{A}}$ is compact in the case that the pairs from distinct $(0, +, \dots)$ elements in \hat{A} define distinct angles via (1–8). Thus, $\mathcal{M}_{\hat{A}}/\mathbb{R}$ is compact in this case. However, when two or more $(0, +, \dots)$ elements of \hat{A} have integer pairs that give the same angle, then the quotient $\mathcal{M}_{\hat{A}}/\mathbb{R}$ is no longer compact. This part of the subsection concerns the latter situation. In particular, what follows is a brief introduction to the sorts of subvarieties that inhabit the frontier of $\mathcal{M}_{\hat{A}}/\mathbb{R}$. A more detailed description of the frontier is contained in Section 9.

To set the stage, remark that $\hat{O}^{\hat{A}}$ sits in $O^{\hat{A}}$ and the latter sits in the compact space, $\underline{O}^{\hat{A}}$, that is obtained by replacing $\mathbb{R}^{\hat{A}}$ in (1–21) with the whole of $\text{Maps}(\hat{A}_+; \mathbb{R})$. As the $\text{Aut}^{\hat{A}}$ action on $O^{\hat{A}}$ extends to one on $\underline{O}^{\hat{A}}$, the space $\underline{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$ provides a compactification of $\mathcal{M}_{\hat{A}}/\mathbb{R}$. This compactification is geometrically natural since each point in $\underline{O}^{\hat{A}}$ can be used to parametrize some pseudoholomorphic, multiply punctured sphere. However, a point in $\underline{O}^{\hat{A}} - \hat{O}^{\hat{A}}$ together with a point in \mathbb{R} always parametrizes a sphere with less than $N_+ + N_- + \hat{N}$ punctures. Even so, the subvarieties that are parametrized by the points in $\underline{O}^{\hat{A}} - \hat{O}^{\hat{A}}$ are geometric limits of sequences in $\mathcal{M}_{\hat{A}}$.

To elaborate on this last remark, suppose that $y \in \mathbb{R} \times \underline{O}^{\hat{A}}$ and that C is the pseudoholomorphic, punctured sphere that is parametrized by y . Meanwhile, let $\{y_j\}$ denote any sequence in $\mathbb{R} \times \underline{O}^{\hat{A}}$ that converges to y and let $\{C_j\}$ denote the corresponding sequence of pseudoholomorphic subvarieties. As explained in Subsection 9.C, the latter converges pointwise to C in $\mathbb{R} \times (S^1 \times S^2)$ in the sense that

$$(1-23) \quad \lim_{j \rightarrow \infty} \left(\sup_{z \in C} \text{dist}(z, C_j) + \sup_{z \in C_j} \text{dist}(C, z) \right) = 0.$$

In fact, more is true: Let C_0 denote the model curve for C , this a punctured sphere together with a proper, almost everywhere 1–1 pseudoholomorphic map to $\mathbb{R} \times (S^1 \times S^2)$ whose image is C . As noted above, C has at worst immersion singularities. Thus, the punctured sphere C_0 has a canonical ‘pull-back’ normal bundle, $N \rightarrow C_0$, with an exponential map from a fixed radius disk subbundle in N to $\mathbb{R} \times (S^1 \times S^2)$ that immerses this disk bundle as a regular neighborhood of C . Let $N_1 \subset N$ denote this disk bundle. Now, suppose that R is very large number, chosen to insure, among other things, that the $|s| \geq R$ part of C lies far out on the ends of C . Let C_{0j} denote the corresponding model curve for C_j . Then the $|s| \leq R$ portion of each sufficiently large j version of C_{0j} maps to the corresponding portion of N_1 so that the composition with the exponential map gives the tautological map to the $|s| \leq R$ part of C_j . Meanwhile, the composition of this map to N_1 with the projection from N_1 to C_0 defines a proper covering map of the $|s| \leq R$ part of C_{0j} over this same part of C_0 . In this regard, the degree of this covering can be 1 or greater; in all cases, the covering is unramified.

To describe the $|s| \geq R$ part of C_j , suppose that $E \subset C$ is an end and let $\gamma \subset S^1 \times S^2$ denote the corresponding closed Reeb orbit. This is to say that the constant $|s|$ slices of E converge pointwise to γ as $|s| \rightarrow \infty$. Thus, E lies in a small, constant radius tubular neighborhood of one of the very large $|s|$ sides of $\mathbb{R} \times \gamma$ as an embedded cylinder. A component of the $|s| \geq R$ part of C_j lies in this same tubular neighborhood. When y is in $O^{\hat{A}}$, then the corresponding part of C_{0j} is a cylinder. However, if y is in the boundary of $O^{\hat{A}}$, there is at least one end of C where the corresponding part of C_{0j} is a sphere with at least three punctures. In any case, there is a tubular neighborhood projection onto $\mathbb{R} \times \gamma$ that restricts to the nearby part of C_j so as to define a proper, possibly ramified covering map from the corresponding $|s| \geq R$ part of C_{0j} onto the relevant $|s| \geq R$ part of $\mathbb{R} \times \gamma$.

More is said in Subsection 9.C about this compactification.

1.C The case that $\mathcal{M}_{\hat{A}}$ has dimension greater than $N_+ + 2$

The discussion here is meant to provide a brief overview of the story in the case that $N_- + \hat{N} + \zeta_- + \zeta_+ = k + 2 > 2$.

Necessary and sufficient conditions for a non-empty $\mathcal{M}_{\hat{A}}$ are given in [15, Theorem 1.3]. As remarked previously, when non-empty, then $\mathcal{M}_{\hat{A}}$ is a smooth manifold whose dimension is $N_+ + 2k + 2$. The structure of $\mathcal{M}_{\hat{A}}$ in the $k > 0$ case is rather more complicated than in the $k = 0$ case. To simplify matters to some extent, the space $\mathcal{M}_{\hat{A}}$ is viewed here as an open subset in a somewhat larger space whose extra elements consist of what are called ‘multiply covered’ subvarieties. In this regard, a multiply

covered subvariety consists of an equivalence class of elements of the form (C_0, ϕ) where C_0 is a connected, complex curve with finite Euler characteristic and ϕ is a proper, pseudoholomorphic map from C_0 to $\mathbb{R} \times (S^1 \times S^2)$ whose image is a pseudoholomorphic subvariety in the sense of (1–5). Here, the equivalence relation puts $(C_0, \phi) \sim (C_0, \phi')$ in the case that ϕ' is obtained from ϕ by composing with a holomorphic diffeomorphism of C_0 . If ϕ is almost everywhere 1–1 onto its image, then (C_0, ϕ) defines a pseudoholomorphic subvariety as described in (1–5). The added elements consist of the equivalence classes of pairs (C_0, ϕ) where ϕ maps C_0 to $\phi(C_0)$ with degree greater than 1. Of interest here is the case where C_0 and $\phi(C_0)$ are punctured spheres. Moreover, with \hat{A} given, then C_0 must have $N_+ + N_- + \hat{N}$ punctures. Furthermore, the ends of C_0 and the points where θ 's pull-back is 0 and π must define the data set \hat{A} . In what follows $\mathcal{M}_{\hat{A}}^*$ is used to denote this larger space.

The topology on $\mathcal{M}_{\hat{A}}^*$ is defined as follows: A basis for the neighborhoods of any given (C_0, ϕ) are sets $\{\mathcal{U}_\kappa\}_{\kappa>0}$, where \mathcal{U}_κ consists of the equivalence classes of elements (C_0', ϕ') that obey the following: There exists a diffeomorphism $\psi: C_0 \rightarrow C_0'$ such that

$$(1-24) \quad \sup_{z \in C} (\text{dist}(\phi(z), (\phi' \circ \psi)(z)) + r_z(\psi)) < \kappa.$$

Here, $r_z(\psi)$ is the ratio of the norm at z of the $\text{Hom}(T_{0,1}C_0; T_{1,0}C_0')$ part of ψ 's differential to that of the $\text{Hom}(T_{1,0}C_0; T_{1,0}C_0')$ part. The theorem that follows describes the local topology of $\mathcal{M}_{\hat{A}}^*$. This theorem introduces the subspace $\mathcal{R} \subset \mathcal{M}_{\hat{A}}^*$ whose elements are the equivalence classes (C_0, ϕ) where ϕ agrees with its pull-back under some non-trivial holomorphic diffeomorphism. Thus $\phi = \phi \circ \psi$ where $\psi: C_0 \rightarrow C_0$ is a non-trivial, holomorphic diffeomorphism.

The description that follows of $\mathcal{M}_{\hat{A}}^*$ speaks of a ‘smooth orbifold’. This term is used here to denote a Hausdorff space with a locally finite, open cover by sets of the form B/G , where B is a ball in a Euclidean space and where G is finite group acting on B . Moreover, these local charts are compatible in the following sense: Let U and U' denote two sets from the cover that overlap. Let $\lambda: B \rightarrow U$ and $\lambda': B' \rightarrow U'$ denote the quotient maps and let $\Omega \subset B$ and $\Omega' \subset B'$ denote respective components of the λ and λ' inverse images of $U \cap U'$. Then, there is a diffeomorphism, $h: \Omega \rightarrow \Omega'$ such that $\lambda = \lambda' \circ h$. A smooth map between smooth orbifolds is a map that lifts near any given point in the domain to give a smooth, equivariant map between Euclidean spaces. A diffeomorphism is a smooth homeomorphism with smooth inverse.

Theorem 1.3 *The space $\mathcal{M}_{\hat{A}}^*$ has the structure of a smooth manifold of dimension $N_+ + 2(N_- + \hat{N} + \zeta_{\hat{A}} - 1)$ on the complement of \mathcal{R} and, overall, it has the structure of a smooth orbifold. Moreover, $\mathcal{M}_{\hat{A}}$ embeds in $\mathcal{M}_{\hat{A}}^*$ as an open set.*

This theorem is proved in Section 5.

A more detailed view of $\mathcal{M}_{\hat{A}}^*$ is provided via a decomposition as a stratified space where each stratum intersects $\mathcal{M}_{\hat{A}}^* - \mathcal{R}$ as a smooth submanifold that extends across \mathcal{R} as an orbifold. The details of this are in Sections 5–9. What follows is a brief outline of the story.

The strata of $\mathcal{M}_{\hat{A}}^*$ are indexed in part by a subset, B , of $(0, -, \dots)$ elements from \hat{A} along with a non-negative integer, c , that is no greater than $N_- + \hat{N} + \zeta_- + \zeta_+ - 2 - |B|$. Here, $|B|$ denotes the number of elements in the set B . With B and c fixed, introduce $\mathcal{S}_{B,c} \subset \mathcal{M}_{\hat{A}}^*$ to denote those equivalence classes (C_0, ϕ) for which the ϕ pull-back of θ has precisely c critical points where θ is neither 0 nor π , and where the following condition holds:

- (1–25) *The 4–tuples in B correspond to the ends in C_0 where s is unbounded from below, where the $s \rightarrow -\infty$ limit of θ is neither 0 nor π , and where this limit is achieved at all sufficiently large values of $|s|$.*

In general, $\mathcal{S}_{B,c}$ is a union of strata. To describe the latter, introduce $d \equiv N_+ + |B| + c$ and use I_d to denote the space of d unordered points in $(0, \pi)$, thus the d –fold symmetric product of the interval $(0, \pi)$. The space I_d has a stratification whose elements are labeled by the partitions of d as a sum of positive integers. A partition of d as $d_1 + \dots + d_m$ corresponds to the subset in I_d where there are precisely m distinct θ values, with d_1 points supplying one θ value, d_2 supplying another, and so on. If $\mathfrak{d} = (d_1, \dots, d_m)$ denotes such a partition, use $I_{d,\mathfrak{d}}$ to denote the corresponding stratum.

Now define a map, $f: \mathcal{S}_{B,c} \rightarrow I_d$ by assigning to any given pair (C_0, ϕ) the angles of the critical values of $\phi^*\theta$ in $(0, \pi)$ along with the angles that are defined via (1–8) using the integer pairs from the $(0, +, \dots)$ elements in \hat{A} and from the 4–tuples from B . Each stratum in \mathcal{S}_c has the form $f^{-1}(I_{d,\mathfrak{d}})$ where \mathfrak{d} is a partition of the integer d . The stratum $f^{-1}(I_{d,\mathfrak{d}}) \subset \mathcal{S}_{B,c}$ is denoted by $\mathcal{S}_{B,c,\mathfrak{d}}$ in what follows.

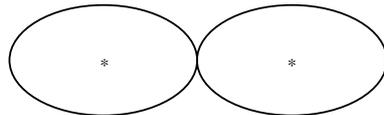
As explained in Sections 6 and 8, each component of each stratum is homeomorphic to a space that has the schematic form $\mathbb{R} \times \mathbb{O} / \mathbb{A}ut$ where $\mathbb{A}ut$ is a certain finite group and \mathbb{O} is a certain tower of circle bundles over a product of simplices. The locus \mathcal{R} appears here as the image of the points where the $\mathbb{A}ut$ action is not free.

By way of an example, suppose that the integer pairs from the $(0, +, \dots)$ elements in \hat{A} define distinct angles via (1–8). In this case, the version of \mathbb{O} for any component of the $N_+ + 2(k + 1)$ dimensional stratum is a tower of circle bundles over a k –dimensional product of simplices. Moreover, $\mathbb{A}ut$ is always trivial when $k > 0$. To say more,

reintroduce the set $\Lambda_{\hat{A}}$ and let $\{\theta_-, \theta_+\}$ denote the respective minimal and maximal angles in $\Lambda_{\hat{A}}$. Then any k -dimensional simplex product that appears in the definition of an $N_+ + 2(k + 1)$ dimensional version of \mathbb{O} is a component of the space of k distinct angles in $(\theta_-, \theta_+) - \Lambda_{\hat{A}}$.

In this example, the codimension 1 strata consist of the subvarieties from $\mathcal{M}_{\hat{A}}^*$ where the pull-back of θ to the model curve has k non-degenerate critical points with either $k - 1$ distinct critical values, none in $\Lambda_{\hat{A}}$, or k distinct critical values with one in $\Lambda_{\hat{A}} - (\theta_-, \theta_+)$. What follows is meant to give a rough picture of the manner in which the coincident top dimensional strata join along the codimension 1 stratum.

In the case where there is a critical value in $\Lambda_{\hat{A}}$, the picture depends on where the angle comes from. For example, in the case that the angle comes from some $(0, +, \dots)$ element in \hat{A} , three coincident top dimensional strata glue across the codimension 1 strata via a fiber bundle version of the ‘pair of pants’ that joins two circles to one. This happens fiberwise as respective fiber circles for two of the associated top strata are joined across a codimension 1 stratum to a fiber circle in the third. The following is a schematic drawing:



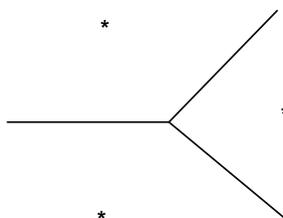
(1-26)

In this picture, the interior of each oval, minus its center point, and the exterior of the union of the ovals correspond to the three codimension 0 strata. The two ovals minus their intersection point correspond to the codimension 1 strata; the central intersection point is a codimension two stratum.

If the angle comes from some $(0, -, \dots)$ element, then there are two incident top dimensional strata; and the gluing comes from an identification between a fiber circle in one top strata and a corresponding circle in the other. There is no fancy stuff here.

The story for the case where two critical values coincide is more involved by virtue of the fact that there are various ways for this to happen. In the first, the critical values coincide but the two critical points are in mutually disjoint components of the critical locus. In this case again, two codimension 0 strata glue across the codimension 1 stratum via an identification of a fiber circle in one stratum with that in the other.

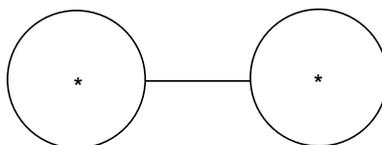
In the remaining cases, the two critical points share the same component of the critical locus. What follows is a schematic picture for the first of these cases:



(1-27)

Here, the complement of the three points and the three rays corresponds to three codimension 0 strata. The rays minus the origin correspond to codimension 1 strata. The origin and the point at ∞ correspond to codimension 2 strata.

The following drawing depicts the second of the cases under consideration.



(1-28)

There are three top dimensional strata depicted here; these correspond to the interiors of the two circles minus their centers, and the exterior region. Meanwhile, the codimension 1 strata correspond to the interior of the horizontal arc, and the complement in the two circles of the endpoints of this same arc. The end points of the horizontal arc depict codimension 2 strata.

Section 9 contains additional details about all of this. Section 9 also describes a compatible, stratified compactification of $\mathcal{M}^*_{\hat{A}}$.

1.D A table of contents for the remaining sections

Section 2 constitutes a digression of sorts to accomplish two tasks. The first is to explain how to use the level sets of the function θ on a given subvariety to construct a certain sort of graph with labeled edges and vertices. The second task explains how this graph can be used to define a canonical set of parametrizations for the subvariety. As explained in a later section, the associated graph and canonical parametrizations provide the map that identifies a given component of a given stratum in $\mathcal{M}^*_{\hat{A}}$ with a particular version of $\mathbb{R} \times \mathbb{O} / \text{Aut}$.

Section 3 elaborates on the story told above for the case where $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$. In particular, this section describes a map that provides the diffeomorphism in Theorem 1.2. The proof that the map is a diffeomorphism is started in this section.

Section 4 starts with a digression to describe how the parametrizations from Section 2 can be used to distinguish distinct elements in $\mathcal{M}^*_{\hat{A}}$. These results are then used to prove that the map from Section 3 is 1–1. The final part of this section proves that the map is proper. This completes the proof of Theorem 1.2.

Section 5 proves Theorem 1.3; and it proves that each component of each stratum of $\mathcal{M}^*_{\hat{A}}$ is a suborbifold.

Section 6 uses the graphs from Section 2 to parametrize certain sorts of slices of the strata of $\mathcal{M}^*_{\hat{A}}$. A graph T determines a certain submanifold, $\mathcal{M}^*_{\hat{A},T}$; these are the fibers for a map that fibers a given component of a given strata over a product of simplices. Theorems 6.2 and 6.3 identify each $\mathcal{M}^*_{\hat{A},T}$ as the product of \mathbb{R} with the quotient by a certain finite group of an iterated tower of circle bundles over simplices. This section ends with description of a particular map that realizes the identification in Theorem 6.2.

Section 7 completes the proof of Theorems 6.2 and 6.3. The arguments here first use the results from the beginning of Section 4 to prove that the map from Section 6 is one-to-one. The map is then proved to be proper. The proof of the latter assertion requires the analysis of limits of sequences in $\mathcal{M}^*_{\hat{A}}$, and the analysis of such limits is facilitated by the use of Section 2’s canonical parametrizations. In particular, the latter are used to replace some of the hard analysis in similar compactness theorems from Bourgeois–Eliashberg–Hofer–Wysocki–Zehnder [1] with topology.

Section 8 describes the stratification of $\mathcal{M}^*_{\hat{A}}$ in greater detail. In particular, this section describes each component of each stratum as a fiber bundle over a product of simplices with the typical fiber being the space $\mathcal{M}^*_{\hat{A},T}$ from Section 6. The section then explains how the graphs that arise can be used to classify the strata.

Section 9 is the final section. This section describes how the codimension 0 strata fit together across the codimension 1 strata and associated codimension 2 strata. This section also describes a certain stratified compactification of $\mathcal{M}^*_{\hat{A}}$. In particular, Subsection 9.C says more about Subsection 1.B’s compactification of the $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ moduli spaces. The section ends with a description of the neighborhoods of the codimension 1 strata that are added to make the compactification of $\mathcal{M}^*_{\hat{A}}$ in the cases where $N_- + \hat{N} + \zeta_- + \zeta_+ > 2$.

Acknowledgements David Gay pointed out to the author that the use of the critical values of θ to decompose $\mathcal{M}^*_{\hat{A}}$ has very much a geometric invariant theory flavor. He also pointed out that the techniques that are described in Subsection 4.A are similar to those used by Grigory Mikhalkin in other contexts. The author also gratefully acknowledges the insight gained from conversations with Michael Hutchings.

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

2 Background material

The purpose of this section is to elaborate on various notions from [15] that are used extensively in the subsequent subsections to construct the parametrizations of any given component of the multi-punctured sphere moduli space. The first subsection explains how to use a subvariety in such a moduli space to define a certain graph with labeled edges and vertices. The second describes some useful parametrizations of various cylinders in any given pseudoholomorphic subvariety. The third explains how a graph from the first subsection is used to designate as ‘canonical’ some of the parametrizations from the second subsection.

2.A Graphs and subvarieties

The purpose of this subsection is to elaborate on the part of the discussion in [15, Section 2.G] that describes a method of associating a graph to a pair (C_0, ϕ) where C_0 is a complex curve and ϕ is a pseudoholomorphic map from C_0 into $\mathbb{R} \times (S^1 \times S^2)$ whose image is a pseudoholomorphic subvariety as defined in (1–5). As discussed in [15, Section 2.E], the pull-back of θ to C_0 has no local extreme points where its value is in $(0, \pi)$. This understood, let $\Gamma \subset C_0$ denote the union of the non-compact or singular components of the level sets of the pull-back of θ .

The graph assigned to (C_0, ϕ) is denoted here as T . Its edges are in a 1–1 correspondence with the components of $C_0 - \Gamma$. Given the correspondence and an edge, e , then K_e is used in what follows to denote e 's component of $C_0 - \Gamma$. Each edge is labeled by an ordered pair of integers, and when e denotes a given edge, Q_e or (q_e, q_e') is used to denote the labeling integer pair. These integers are the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\phi$ about any constant θ slice of K_e when these components are oriented by the pull-back of the form

$$(2-1) \quad x = (1 - 3 \cos^2 \theta) d\phi - \sqrt{6} \cos \theta dt.$$

In this regard, note that this pull-back is non-zero along any such slice because the form in (2–1) is a non-zero multiple of $J \cdot d\theta$. (Here, the action of J on the cotangent bundle is dual to its action on the tangent bundle.) The integral of x around any constant θ slice of K_e gives the $Q \equiv (q, q') = Q_e$ version of the function

$$(2-2) \quad \alpha_Q(\theta) = q'(1 - 3 \cos^2 \theta) - q \sqrt{6} \cos \theta.$$

The $Q = Q_e$ version of α_Q is strictly positive on the closure of K_e .

The monovalent vertices in T are in a 1–1 correspondence with the following:

- (2–3)
- The points in C_0 where θ is either 0 or π .
 - The ends of C_0 where the $|s| \rightarrow \infty$ limit of θ is 0 or π .
 - The convex side ends of C_0 where the $|s| \rightarrow \infty$ limit of θ is not achieved at finite $|s|$.

With regards to the last two points, the discussion in [15, Section 1.E] noted the existence of some $R \gg 1$ such that the $|s| \geq R$ portion of C_0 is a disjoint union of embedded cylinders to which $|s|$ restricts as an unbounded, proper map to $[R, \infty)$ without critical points. Each such cylinder is called an ‘end’. The end is said to be on the convex side of C_0 when s is negative on the end. Otherwise, the end is said to be on the concave side of C_0 . The angle θ on any such end has a unique limit as $|s| \rightarrow \infty$.

To say more about the manner in which this limit is approached, note first that the analysis used in [14, Sections 2 and 3] proves a version of (1–9) for a given end where the $|s| \rightarrow \infty$ limit of θ is in $(0, \pi)$. In fact, the techniques from [14] can be used to find coordinates (ρ, τ) for the end such that ρ is equal to 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

$$(2-4) \quad \theta(\rho, \tau) = \theta_E + e^{-r\rho}(b \cos(n_E(\tau + \sigma)) + \hat{\theta}),$$

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

The convex side ends that arise in the third point of (2–3) are those using the $n_E = 0$ version of (2–4).

The multivalent vertices of T are in a 1–1 correspondence with the sets that comprise a certain partition of the collection of non-point like components in Γ . To describe these partition subsets, define a graph, G , whose vertices are the non-point like components of Γ , and where two vertices share an edge if the restrictions of θ to the corresponding components of Γ agree and if both of these components lie in the closure of some component of $C_0 - \Gamma$. The set of components of G defines the desired partition of the set of non-point like components of Γ . In this regard, note that each compact, non-point like component of Γ defines its own partition subset.

As remarked at the outset, each vertex in T has a label. These labels are explained next. To start, a monovalent vertex that corresponds to a point in C_0 where $\theta = 0$ is labeled by (p') , where p' is the positive integer that gives the degree of tangency at the intersection point between the $\theta = 0$ cylinder and the image of any sufficiently small radius disk about the given point in C_0 . Meanwhile, a vertex that corresponds to a point in C_0 where $\theta = \pi$ is labeled by $(-p')$ where $p' \geq 1$ is the analogous degree of tangency.

A monovalent vertex that corresponds to an end where the $|s| \rightarrow \infty$ limit of θ is 0 or π is labeled by a 4-tuple of the form $(\delta, \varepsilon, (p, p'))$, where $\delta = 1$ if the θ limit is 0 and $\delta = -1$ if the θ limit is π . Meanwhile, $\varepsilon \in \{+, -\}$ with $+$ appearing when the end is on the concave side and $-$ appearing when the end is on the concave side. Finally, the ordered pair (p, p') are the integers that appear in (1–9).

A monovalent vertex that corresponds to a convex side end where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π is labeled by a 4-tuple of the form $(0, -, (p, p'))$ where the ordered pair of integers is either $+$ or $-$ the pair that labels its incident edge. The sign here is that of the constant b in (2–4).

The labeling of the multivalent vertices is more involved. To elaborate, each such vertex is first assigned the angle in $(0, \pi)$ of the components of its corresponding partition subset. In addition, each is assigned a certain graph of its own, this also a graph with labeled vertices and edges. When o denotes a vertex, its graph is denoted by $\underline{\Gamma}_o$. The graph $\underline{\Gamma}_o$ is a certain closure of the union, Γ_o , of the components of the partition subset that is assigned to o . The vertices of $\underline{\Gamma}_o$ that lie in Γ_o consist of the critical points of θ on Γ_o . The vertices in $\underline{\Gamma}_o - \Gamma_o$ are in 1–1 correspondence with the set of ends of C_0 where all sufficiently large $|s|$ slices intersect Γ_o . Each vertex in $\underline{\Gamma}_o$ is labeled with an integer. This integer is zero when the vertex lies in Γ_o . The integer is positive when the vertex corresponds to a concave side end of C_0 and it is negative when the vertex corresponds to a convex side end. In either case, the absolute value of this integer is the greatest common divisor of the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ around any given constant $|s|$ slice in the corresponding end.

So as not to confuse the edges in $\underline{\Gamma}_o$ with o 's incident edges in T , those in $\underline{\Gamma}_o$ are called 'arcs'. The arcs are in 1–1 correspondence with the components of the complement in Γ_o of the Γ_o 's θ critical points. Note that each arc is oriented by the 1–form that appears in (2–1). As can be seen in [15, (2.16)] and (2–4) here, each vertex in $\underline{\Gamma}_o$ has an even number of incident half-arcs and with half oriented so as to point towards the vertex and half oriented so as to point away. Only vertices with non-zero integer label can have two incident half arcs. Each arc is also labeled by two incident edges to the vertex o ; these correspond to the two components of $C_0 - \Gamma$ whose closure contains

the arc's image in the locus Γ_o . Thus, one labeling edge connects o to a vertex with smaller angle and the other to a vertex with larger angle.

An isomorphism between graphs T and T' of the sort just described consists of, among other things, a homeomorphism between the underlying 1-complexes. However, such an isomorphism must map vertices to vertices and edges to edges so as to respect all labels. In particular, if o is a bivalent vertex in T and o' its image in T' , then the isomorphism induces an isomorphism between $\underline{\Gamma}_o$ and $\underline{\Gamma}_{o'}$ that preserves their vertex labels and arc orientations and respects the labeling of their arcs by pairs of incident edges to the vertices. An automorphism of a given graph T is an isomorphism from T to itself.

As just described, the edges of T correspond to the components of $C_0 - \Gamma$, the vertices to certain ends of C_0 and the singular and non-compact level sets of $\phi^*\theta$. Meanwhile, the labeling of the edges and vertices correspond to other aspects of the ends of C_0 and the $\phi^*\theta$ level sets. This understood, a 'correspondence' in (C_0, ϕ) of a graph T of the sort just described signifies in what follows a particular choice for the identification that were just described between the geometry of $\phi^*\theta$ on C_0 and the various edges, vertices and so on of T . Note that any one correspondence of T in (C_0, ϕ) is obtained from any other by the use of an automorphism on T . The use of a subscript, 'C', on T , thus T_C , signifies in what follows a graph T together with a chosen correspondence in a specified pair (C_0, ϕ) .

If T has a correspondence in (C_0, ϕ) , and if ψ is a holomorphic diffeomorphism of C_0 , then T also has a correspondence in $(C_0, \phi' \equiv \phi \circ \psi)$. More to the point, there is an automorphism $\delta: T \rightarrow T$ and respective correspondences for T in (C_0, ϕ) and in (C_0, ϕ') such that the inverse image via ψ and δ intertwine the two correspondences. For example, if $e \subset T$ is an edge and K_e is its corresponding component of the $\phi^*\theta$ version of $C_0 - \Gamma$, then $\psi^{-1}(K_e)$ is the component of the $\phi' * \theta$ version of $C_0 - \Gamma$ that corresponds to $K_{\delta(e)}$.

Granted what has just been said, the isomorphism type of graph with a correspondence in a given pair (C_0, ϕ) depends only on the image of (C_0, ϕ) in $\mathcal{M}^*_{\hat{A}}$.

2.B Preferred parametrizations

Let $K \subset C_0 - \Gamma$ denote a component. Since the θ level sets in K are circles, the angle θ and an affine parameter on the θ level sets can be used to parametrize K by an open cylinder. In this regard, there is a set of preferred parametrizations. To set the stage for their description, introduce the ordered pair $Q \equiv (q, q')$ to denote the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ around any given constant θ slice of K using

the orientation that is defined by the pull-back of the 1-form in (2-1). Next, let θ_0 denote the infimum of θ on K and let θ_1 denote the supremum. In what follows, σ is used to denote the linear coordinate on (θ_0, θ_1) and v is used to denote an affine coordinate for the circle $\mathbb{R}/(2\pi\mathbb{Z})$.

Definition 2.1 A preferred parametrization for K is a diffeomorphism from the cylinder $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ to K whose composition with the tautological immersion of K into $\mathbb{R} \times (S^1 \times S^2)$ can be written in terms of smooth functions a and w on the cylinder as the map that pulls back the coordinates (s, t, θ, φ) as

$$(2-5) \quad \begin{aligned} & \bullet \quad s = a, \\ & \bullet \quad t = qv + (1 - 3 \cos^2 \sigma)w \pmod{2\pi\mathbb{Z}}, \\ & \bullet \quad \theta = \sigma, \\ & \bullet \quad \varphi = q'v + \sqrt{6} \cos \sigma w \pmod{2\pi\mathbb{Z}}. \end{aligned}$$

The set of preferred parametrizations for K is in all cases non-empty. Indeed, to construct such a parametrization, start by fixing a transversal to the constant θ circle in K , this a properly embedded, open arc in K that is parametrized by the restriction of θ . Next, use θ as one coordinate and, for the other, use the line integral of x/α_Q along the constant θ circles from the chosen arc. A preferred coordinate system can be obtained from the latter by adding an appropriate function of θ to the second coordinate.

Listed next are six important properties of the preferred parametrizations.

Property 1 All preferred parametrizations pull the exterior derivative of the contact 1-form α back as $\sin \sigma \alpha_Q(\sigma) d\sigma \wedge dv$ where α_Q is the function on $[0, \pi]$ that appears in (2-2). This implies that α_Q is positive on the interval (θ_0, θ_1) , and that it is positive at an endpoint of this interval if the value there of θ is attained on the closure of K .

With regards, to α_Q , note as well that the 1-form x in (2-1) pulls back to any constant σ circle in the parametrizing cylinder as $\alpha_Q dv$.

Property 2 The fact that K is pseudoholomorphic puts certain demands on the pair (a, w) . In particular, K is pseudoholomorphic in $R \times (S^1 \times S^2)$ if and only if

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

Here, β is defined to be the function $p(1 - 3 \cos^2 \sigma) + p' \sqrt{6} \cos \sigma \sin^2 \sigma$. In this equation and subsequently, the subscripts σ and v denote the partial derivatives in the

indicated direction. Because α_Q is nowhere zero on (θ_0, θ_1) the system in (2–6) is a non-linear version of the Cauchy–Riemann equations.

Here is an important consequence of the second equation in (2–6):

(2–7) *The function $\sigma \rightarrow \alpha_Q(\sigma) \int_{\mathbb{R}/(2\pi\mathbb{Z})} w(\sigma, v) dv$ is constant on the interval (θ_0, θ_1) .*

Property 3 Consider the behavior of K where θ is near a given $\theta_* \in \{\theta_0, \theta_1\}$. If θ_* 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_0 . 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_0 and the associated integer n_E that appears in (2–4) is zero.

Supposing still that $\theta_* \in (0, \pi)$, let (p, p') denote the relatively prime integer pair that defines θ_* via (1–8). Then $(q, q') = m(p, p')$ with m a non-zero integer, and it follows from (2–5) and (2–7) that the mod $(2\pi\mathbb{Z})$ reduction of any $\sigma \in (\theta_0, \theta_1)$ version of

$$(2-8) \quad \frac{1}{2\pi m} \alpha_Q(\sigma) \int_{\mathbb{R}/(2\pi\mathbb{Z})} w(\sigma, v) dv$$

is the $\mathbb{R}/(2\pi\mathbb{Z})$ parameter that distinguishes the Reeb orbit limit of the $\theta \rightarrow \theta_*$ circles in K .

Property 4 If $\theta_* \in \{\theta_0, \theta_1\}$ 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.

This decomposition of the $\theta = \theta_*$ boundary of K into arcs is reflected in the behavior of the parametrizations in (2–5) 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 $[\theta_0, \theta_1] \times \mathbb{R}/(2\pi\mathbb{Z})$. These points are called ‘singular points’. Meanwhile, each end of C_0 that intersects the $\theta = \theta_*$ boundary of the closure of K in a set where $|s|$ is unbounded also labels one or more 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 (2–5) has a smooth extension as an embedding into $\mathbb{R} \times (S^1 \times S^2)$ onto a disk in C_0 .

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 (2–5) along any given arc in the $\sigma = \theta_*$ boundary of $[\theta_0, \theta_1] \times \mathbb{R}/(2\pi\mathbb{Z})$ provides a smooth parametrization of the interior of its corresponding arc in the $\theta = \theta_*$ boundary of the closure of K .

Property 5 It is pertinent to what follows to say more about the behavior of the parametrization near the singular points on the $\sigma = \theta_*$ circle in the case that $\theta_* \in \{\theta_0, \theta_1\}$ is a value of θ on the closure of K .

To set the stage, let z_* denote any given point in C_0 . Let \hat{t} and $\hat{\varphi}$ denote the functions, defined on a ball centered at the image of z_* in $S^1 \times S^2$, that vanish at the latter point and whose respective differentials are dt and $d\varphi$. Now introduce

$$(2-9) \quad r \equiv \frac{\sin \theta}{(1 + 3 \cos^4 \theta)^{1/2}} \left((1 - 3 \cos^2 \theta) \hat{\varphi} - \sqrt{6} \cos \theta \hat{t} \right),$$

a function that is defined on the given ball about the image of z_* in $S^1 \times S^2$. Next, let $D_* \subset C_0$ be a small radius disk with center z_* on which the pull-back of r is well defined. Note that $d\theta \wedge dr$ is zero on D_* only at the critical points of θ . Thus, r and θ define local coordinates on the complement in D_* of its θ critical points.

When a component $K \subset C_0 - \Gamma$ whose closure contains z_* is given a preferred parametrization, there is a function, \hat{v} , that is defined on any given contractible component of $K \cap D_*$ whose differential pulls back to $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ as dv . This understood, then (2–5) identifies r on such a component of $K \cap D_*$ as

$$(2-10) \quad r = \frac{\sin \theta}{(1 + 3 \cos^4 \theta)^{1/2}} \alpha_{Q_K}(\theta) (\hat{v} - v^*) - \frac{\sin \theta}{(1 + 3 \cos^4 \theta)^{1/2}} \sqrt{6} (\cos \theta - \cos \theta_*) (1 + 3 \cos \theta \cos \theta_*) w^*.$$

Here, v^* and w^* are constants. In particular, if z_* is in K , then $v^* = \hat{v}(z_*)$ and w^* is the value of w at the point in $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ that parametrizes z_* via (2–5). Such is also the case when z_* is on the boundary of K and is not a critical point of θ provided that the radius of D_* is sufficiently small. With regard to this last case, note that a small radius guarantees the following: The boundary of K 's closure intersects D_* as an embedded arc and any chosen \hat{v} extends to this arc as a smooth function.

According to (2–9), the 1–form dr can be written as $dr = J \cdot d\theta + rd\theta$ where $|r|$ vanishes at the image of z_* . This and the fact that ϕ is pseudoholomorphic have the

following consequence: There exists a holomorphic coordinate, u , on D_* such that the pull-back of the complex function $\theta + ir$ has the form

$$(2-11) \quad \theta + ir = \theta_* + u^{m+1} + \mathcal{O}(|u|^{m+2}).$$

Here, $m \geq 0$ is the degree at z_* of the zero of $d\theta$. The term in (2-11) designated as $\mathcal{O}(|u|^{m+2})$ is such that its quotient by $|u|^{m+2}$ is bounded as u limits to zero. Now, given that D_* has small radius, (2-11) indicates that K intersects D_* in a finite number of components, each contractible. This noted, then r is given on each such component by a version of (2-10) where v^* is determined by the choice for \hat{v} and the $R/(2\pi\mathbb{Z})$ coordinate of the given given $\sigma = \theta_*$ singular point. Meanwhile, w^* is the limiting value of the function w at the singular point that maps to z_* . Together, (2-10) and (2-11) describe the behavior of a preferred parametrization near this singular point.

Property 6 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 (2-5) 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.

To say a bit more about this case, note that the the pair (q, q') has $q = 0$ and $q' < 0$. In this regard, $-q'$ is the intersection number between the image of the aforementioned disk and the $\theta \in \{0, \pi\}$ locus. As $q = 0$, the assertion in (2-7) implies that the mod $(2\pi\mathbb{Z})$ reduction of any $\sigma \in (\theta_0, \theta_1)$ version of the expression in (2-8) gives the t -coordinate of the intersection point between the disk and the $\theta = \{0, \pi\}$ locus.

To end this subsection, note that the set of preferred parametrizations of a given component $K \subset C_0 - \Gamma$ constitutes a single orbit for an action of the group $\mathbb{Z} \times \mathbb{Z}$. To elaborate, let ψ denote a parametrization for K , and let $N \equiv (n, n')$ denote an ordered pair of integers. Let ϕ_N denote the diffeomorphism of $(\theta_0, \theta_1) \times \mathbb{R}/(2\pi\mathbb{Z})$ that pulls (σ, v) back as

$$(2-12) \quad \phi_N^*(\sigma, v) = \left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_Q(\sigma)} \right).$$

Then $\psi^N \equiv \psi \circ \phi^N$ is also a preferred parametrization. In this regard, if (a, w) are the pair that appears in ψ 's version of (2-5), then the ψ^N version is the pair (a^N, w^N) given by

$$(2-13) \quad \begin{aligned} a^N(\sigma, v) &= a \left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_Q(\sigma)} \right) \quad \text{and} \\ w^N(\sigma, v) &= w \left(\sigma, v - 2\pi \frac{\alpha_N(\sigma)}{\alpha_Q(\sigma)} \right) + 2\pi \frac{qn' - q'n}{\alpha_{Q_e}(\sigma)}. \end{aligned}$$

The assignment of the pair (N, ψ) to ψ^N defines a transitive action of $\mathbb{Z} \times \mathbb{Z}$ on the set of preferred parametrizations. In this regard, note that the stabilizer of any given parametrization is the \mathbb{Z} subgroup in $\mathbb{Z} \times \mathbb{Z}$ of the integer multiples of Q .

As all parametrizations that arise henceforth for any given component of $C_0 - \Gamma$ are preferred parametrizations, the convention taken from here through the end of this article is that the word ‘parametrization’ refers in all cases to a preferred parametrization. The qualifier ‘preferred’ will not be written as its presence should be implicitly understood.

2.C Graphs such as $\underline{\Gamma}_o$ and preferred parametrizations

Let o denote a given multivalent vertex in T_C . This subsection describes how the graph $\underline{\Gamma}_o$ is used to define certain ‘canonical’ parametrizations for those components of $C_0 - \Gamma$ that are labeled by o ’s incident edges. This is preceded by a description of various properties of $\underline{\Gamma}_o$ that are used in subsequent parts of this article. The discussion here has seven parts.

Part 1 Suppose here that γ is an arc in $\underline{\Gamma}_o$ and let $e \in E_-$ and $e' \in E_+$ denote the incident edges to o that comprise its edge-pair label. Suppose that both K_e and $K_{e'}$ have been graced with parametrizations. View the interior of γ as an open arc in Γ_o and so an open arc in C_0 . As explained in the previous subsection, the parametrizations of both K_e and $K_{e'}$ extend to a neighborhood of $\text{int}(\gamma)$. As such, there is a ‘transition function’ that relates one of these extensions to the other. To describe this transition function, let \hat{v}_e denote a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued coordinate v on the parametrizing cylinder for K_e . Meanwhile, let $\hat{v}_{e'}$ denote a corresponding lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued coordinate on the parametrizing cylinder for $K_{e'}$. Then the coordinate transition function has the form

$$(2-14) \quad \alpha_{Q_e} \hat{v}_e = \alpha_{Q_{e'}} \hat{v}_{e'} - 2\pi \alpha_N,$$

where $N \equiv (n, n')$ is some ordered pair of integers and $\alpha_N = n'(1 - 3 \cos^2 \theta) - \sqrt{6n \cos \theta}$ is the $Q = N$ version of the function that appears in (2-2). In this regard, the pair N can be chosen so that the corresponding versions of (a, w) that appear in (2-5) are related via

$$(2-15) \quad \begin{aligned} a_e \left(\sigma, \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} \hat{v}_{e'} + 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)} \right) &= a_{e'}(\sigma, \hat{v}_{e'}), \\ w_e \left(\sigma, \frac{\alpha_{Q_{e'}}(\sigma)}{\alpha_{Q_e}(\sigma)} \hat{v}_{e'} + 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)} \right) &= w_{e'}(\sigma, \hat{v}_{e'}) \\ &+ \frac{1}{\alpha_{Q_e}(\sigma)} (q_{e'} q_{e'} - q_e q_{e'}) \hat{v}_{e'} - \frac{2\pi}{\alpha_{Q_e}(\sigma)} (q_e n' - q_{e'} n). \end{aligned}$$

Here, $(q_e, q_{e'})$ comprise the pair Q_e while $(q_{e'}, q_{e'})$ comprise $Q_{e'}$. With regards to these formulae, note that changing the lift \hat{v}_e by adding 2π has the effect of changing the integer pair from N to $N - Q_e$. Meanwhile, a change of the lift $\hat{v}_{e'}$ by the addition of 2π changes N to $N + Q_{e'}$.

Part 2 As is explained momentarily, a parametrization of K_e and a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued coordinate on the corresponding parametrizing cylinder determines a canonical ordered pair whose first component is a parametrization of $K_{e'}$ and whose second is a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued coordinate on the latter's parametrizing cylinder. Indeed, this canonical pair provides the $N = (0, 0)$ version of (2–14) and (2–15), and it is a consequence of (2–12) and (2–13) that there is a unique pair of parametrization and lift that makes both (2–14) and (2–15) hold with $N = (0, 0)$.

Keep in mind that this canonical parametrization and lift changes with a change in the initial parametrization for K_e and lift of its $\mathbb{R}/(2\pi\mathbb{Z})$ parameter. To make matters explicit, suppose that $N = (n, n')$ is an integer pair and that the latter changes the parametrization for K_e as depicted in (2–12) and (2–13). In addition, suppose that the value of the new \mathbb{R} -lift, \hat{v}_e^N , of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate on the parametrizing cylinder is related to the original at any given point by

$$(2-16) \quad \hat{v}_e^N = \hat{v}_e - 2\pi \frac{\alpha_N(\sigma)}{\alpha_{Q_e}(\sigma)}.$$

It now follows from (2–14) and (2–15) that the canonical e' pair of parametrization and lift are changed in the analogous manner by this same integer pair N . This is to say that the new parametrization of $K_{e'}$ is related to the original via the e' version of (2–12) and (2–13) as defined using the pair N . Meanwhile, the \mathbb{R} -valued lift, $\hat{v}_{e'}^N$, of the $\mathbb{R}/(2\pi\mathbb{Z})$ valued coordinate on the e' version of the parametrizing cylinder is obtained from the old at any given point by the version of (2–16) that substitutes e' for e .

With regards to \mathbb{R} -valued lifts of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate on a given parametrizing cylinder, note that any such lift is uniquely determined by the lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of any one point. In the applications below, a lift to \mathbb{R} is chosen for the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of a distinguished missing or singular point on the $\sigma = \theta_o$ boundary of the parametrizing cylinder. The latter lift is then used to define the lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate over the whole cylinder.

This last very straightforward observation is used below in the following context: Suppose that v_e is a distinguished missing or singular point on the $\sigma = \theta_o$ boundary circle of the parametrizing cylinder for e , and suppose that a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of v_e has been chosen. Now, let ν denote a path on this boundary circle that

begins at v_e and ends at some other missing or singular point. In this regard, the end point of ν may well be v_e . In any event, assume that ν can be written as a non-empty concatenation of segments that connect pairs of missing points, or connect pairs of singular points, or connect a missing point and a singular point. Note that any such segment can run either with or against the defined orientation of the $\sigma = \theta_o$ boundary circle.

Let γ' denote the image in $\underline{\Gamma}_o$ of the final segment on ν , and let e' denote the vertex that labels γ' with e . The lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of v_e then defines a lift to \mathbb{R} of the coordinate along the whole of ν and thus a lift over γ' . In this regard, note that the latter depends on the homotopy class rel end points of the path ν . In any event, this lift and the given parameterization of K_e determines a canonical parameterization of $K_{e'}$. Of course, it also determines a canonical lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the end point of ν .

Part 3 This and Part 4 of the subsection describe a generalization of these constructions. In this regard, the preceding definition suffices when T_C is a linear graph with no automorphisms, but more is needed for a more complicated graph.

This part of the story constitutes a digression whose purpose is to elaborate on some aspects of the graph $\underline{\Gamma}_o$. To start, let e denote an incident edge to o . Then e labels a certain circular graph, ℓ_{oe} , with oriented edges and labeled vertices that has an immersed image in $\underline{\Gamma}_o$. In order to describe ℓ_{oe} , it convenient to first choose a parametrization of the component, $K_e \subset C_0 - \Gamma$ that e labels. Such a parametrization identifies ℓ_{oe} with the $\sigma = \theta_o$ circle in the associated parametrizing cylinder. In this regard, the vertices of ℓ_{oe} are the set of missing and singular points on this circle, and its 'arcs' are then the arcs in this circle that connect these points. The orientation of an arc is defined by the restriction of dv . Meanwhile, the vertices are labeled by integers in the following manner: All singular points are labeled by the integer 0. Any given missing point has label $\pm m$ where m is a positive integer and where the $+$ sign is used if and only if the missing point corresponds to a concave side end of C . The integer m is the greatest common divisor of the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ over any given constant $|s|$ slice of the end.

The map from ℓ_{oe} to $\underline{\Gamma}_o$ is then induced by the extension of the parametrizing map. In this regard, recall that this extension maps the complement of the set of missing points on the $\sigma = \theta_o$ circle to $\Gamma_o \subset \underline{\Gamma}_o$. This map from ℓ_{oe} to $\underline{\Gamma}_o$ has the following properties: First, it sends vertices to vertices and arcs to arcs so as to preserve the vertex labels and the arc orientations. Second it is 1–1 on the complement of the vertices. Thus, the image of the map from ℓ_{oe} to $\underline{\Gamma}_o$ is a closed, oriented path in $\underline{\Gamma}_o$ that crosses no arc more than once. The arcs in the image of ℓ_{oe} are those whose edge pair label

contains e . In what follows, the homology class in $H_1(\underline{\Gamma}_o; \mathbb{Z})$ of the image of $\ell_{o,e}$ is denoted as $[\ell_{oe}]$.

Since the abstract circle ℓ_{oe} can be reconstituted (up to an automorphism) from its image as an oriented path in $\underline{\Gamma}_o$, the image is also denoted by ℓ_{oe} . Note as well that distinct parametrizations of K_e define isomorphic versions of ℓ_{oe} and that the associated maps to $\underline{\Gamma}_o$ are compatible with any such isomorphism.

Four properties of C_0 are reflected in the manner in which the collection $\{\ell_{oe}\}$ sit in $\underline{\Gamma}_o$. Their descriptions involve sets E_- and E_+ where E_- is the set of incident edges to o that connect to vertices with angles less than θ_o , while E_+ is the set of incident edges that connect to vertices with angles greater than θ_o .

Property 1 *Each arc in $\underline{\Gamma}_o$ is contained in precisely two versions of $\ell_{o(\cdot)}$, one labeled by an edge from E_+ and the other by an edge from E_- . These are the edges that comprise the label of the arc.*

Property 2 *The collection $\{[\ell_{oe}]\}_e$ is incident to o generates $H_1(\underline{\Gamma}_o; \mathbb{Z})$ subject to the following constraint:*

$$(2-17) \quad \sum_{e \in E_+} [\ell_{oe}] - \sum_{e \in E_-} [\ell_{oe}] = 0.$$

The first property is a direct consequence of the definitions. To explain the second property, first note that the closed parametrizing cylinders for the components of $C_0 - \Gamma$ that are labeled by o 's incident edges can be glued to $\underline{\Gamma}_o$ by identifying any given version of $\ell_{o(\cdot)} \subset \underline{\Gamma}_o$ with corresponding segment in the $\sigma = \theta_o$ boundary of the relevant parametrizing cylinder. The result of this gluing is a sphere with as many punctures as o has incident edges. By construction this multipunctured sphere deformation retracts onto $\underline{\Gamma}_o$. Meanwhile, its homology is generated by the collection $\{[\ell_{oe}]\}$ subject to the one constraint in (2-17).

The third property is actually a consequence of the first two, but can also be seen to follow from the fact that C_0 is a smooth, irreducible curve.

Property 3 *Let \hat{E} denote the set of o 's incident edges that appear in the pair labels of the incident half-arcs to a given vertex in $\underline{\Gamma}_o$. An equivalence relation on \hat{E} is generated by equating the two edges that come from the edge pair label of an incident half-arc to the given vertex. This relation defines just one equivalence class.*

The final property is a special case of the one just stated.

Property 4 Suppose that e and e' are incident edges to o , one in E_+ and the other in E_- , and suppose that $\gamma \subset \ell_{oe} \cap \ell_{oe'}$. If γ' follows γ in ℓ_{oe} , then γ' cannot follow γ in $\ell_{oe'}$ unless the vertex between them is bivalent.

Part 4 With the preliminary digression now over, this part of the subsection describes the advertised generalization of the definition in Part 2 of a canonical parametrization. The following definition is needed to set the stage:

Definition 2.2 A ‘concatenating path set’ is an ordered set, $\{v_1, v_2, \dots, v_N\}$, of labeled paths in $\underline{\Gamma}_o$ with the following properties:

- The label of each v_k consists of a specified direction of travel and a specified incident edge to o .
- If e denotes the edge label to a given v_k , then v_k is entirely contained in ℓ_{oe} . Moreover, v_k is the concatenation of a non-empty, ordered set of arcs in ℓ_{oe} that are crossed in their given order when the path is traversed from start to finish. Note that the arcs that comprise v_k need not be distinct and can be crossed in either direction.
- No two consecutive pairs v_k and v_{k+1} are labeled by the same incident edge to o .
- For each $1 \leq j < N$, the final arc on v_j is the starting arc on v_{j+1} .

The concatenating path set defines a directed path in $\underline{\Gamma}_o$ by sequentially traversing the paths that comprise the ordered set $\{v_1, \dots, v_N\}$ in their given order.

Granted this definition, here is the context for the discussion that follows: Suppose that one of o ’s incident edges has been designated as the ‘distinguished incident edge’. Let e denote the latter, and suppose that a parametrization is given for K_e , and that a vertex has been designated as the ‘distinguished vertex’ on the realization of ℓ_{oe} as the $\sigma = \theta_o$ boundary circle of the parametrizing cylinder. Let v_e denote this distinguished vertex, and let $v \in \underline{\Gamma}_o$ denote its image. Now, let e' label some incident edge to o with $e' = e$ allowed. Suppose, in addition, that a concatenating path set $v \equiv \{v_1, \dots, v_N\}$ has been chosen so that

- (2–18) • The edge label of v_1 is e , and v_1 starts at v .
 • The edge label of v_N is not e' but its final arc lies in $\ell_{oe'}$.

Properties 3 and 4 from Part 3 can be used to construct this sort of concatenating path set.

With this data set, a *parametrizing algorithm* is described next whose input is a pair consisting of a parametrization for e and a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of v_e in the $\sigma = \theta_o$ boundary of the parametrizing cylinder for K_e and whose output is a parametrization for $K_{e'}$ as well as a canonical \mathbb{R} -valued lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ parameter on the associated parametrizing cylinder. This output parameterization is the advertised ‘canonical’ parametrization for $K_{e'}$.

To describe this algorithm, note first that the observations from the final paragraph from Part 2 can be viewed as defining a parametrizing subroutine that takes as input the data:

- (2–19) • A parametrization of a component of $C_0 - \Gamma$ whose closure intersects Γ_o .
- A distinguished point on the $\sigma = \theta_o$ boundary of the parametrizing cylinder.
 - A lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of this distinguished point.
 - A non-trivial path in the $\sigma = \theta_o$ boundary of the parametrizing cylinder that starts at the distinguished point and consists of segments that connect missing and/or singular points on this boundary.

and gives as output:

- (2–20) • A parametrization of the as yet unparametrized component of $C_0 - \Gamma$ whose closure contains the image in Γ_o of the interior of the final segment on the chosen path.
- A lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate on the $\sigma = \theta_o$ boundary of the corresponding parametrizing cylinder for this second component of $C_0 - \Gamma$.

The parametrizing algorithm runs this subroutine N consecutive times. The first run uses as input the chosen parameterization for K_e , the point v_e , the chosen lift of its $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate, and the path v_1 . The second run of the subroutine uses as input the resulting parameterization from (2–20) of the component of $C_0 - \Gamma$ that shares the edge label with v_2 , the second to last vertex on v_1 with the lift of its $\mathbb{R}/(2\pi\mathbb{Z})$ parameter from (2–20), and the path v_2 . In general, the $(j + 1)$ 'st run of the subroutine uses the output from the j 'th run of the subroutine as it takes as input the parameterization in (2–20) for the component of $C_0 - \Gamma$ that shares the edge label with v_{j+1} , the second to last vertex on v_j , the lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate for v_j that is obtained from (2–20), and the path v_{j+1} . Because of (2–18), the N 'th run of the subroutine parametrizes $K_{e'}$ and gives a lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ parameter on its parametrizing cylinder. This parametrization of $K_{e'}$ is deemed ‘canonical’.

As just described, the canonical parameterization of a component of $C_0 - \Gamma$ whose closure intersects Γ_o depends on the following input data:

- (2–21) • A distinguished incident edge to o
- A parameterization of the corresponding component in $C_0 - \Gamma$.
 - A choice of a distinguished vertex in the $\sigma = \theta_o$ boundary of the parametrizing cylinder.
 - A lift to R of the $R/(2\pi\mathbb{Z})$ coordinate of the corresponding point in the $\sigma = \theta_o$ boundary of the parametrizing cylinder.
 - A choice of a concatenating path set that obeys the constraints in (2–18).

The dependence on this input data has a role in subsequent parts of this article and so warrants the discussion that occupies the remaining portions of this subsection.

Part 5 The question on the table now is that of the dependence of a canonical parametrization on the data in (2–21). What follows describes the situation for the various listed cases, starting at the top and descending. In these descriptions, the edge e denotes the original distinguished edge, and the edge e' denotes the edge that labels the component of $C_0 - \Gamma$ that is to receive the canonical parameterization. Use $\{v_1, \dots, v_N\}$ to denote the original concatenating path set.

Case 1 Suppose that a new distinguished edge, \hat{e} , has been selected. What follows describes input data for the parametrization algorithm, now run with the edge \hat{e} , that supplies a canonical parametrization for $K_{e'}$ that agrees with the original one. In short, the \hat{e} input data is obtained as follows: A concatenating path, \hat{v} , is chosen to first parametrize $K_{\hat{e}}$ given the original parametrization of K_e . This parametrization of $K_{\hat{e}}$ is then used as the input to the algorithm to obtain the \hat{e} version of the canonical parametrization of $K_{e'}$. The concatenating path that is used for this \hat{e} parametrization of $K_{e'}$ runs backward along \hat{v} to its starting vertex on ℓ_{oe} , and then forward along the concatenating path that is used by the parametrization algorithm to obtain the original canonical parametrization of $K_{e'}$.

Here are the details: Choose a concatenating path set, $\{\hat{v}_1, \dots, \hat{v}_{\hat{N}}\}$ whose end vertex lies on $K_{\hat{e}}$ and let \hat{v} denote its ending vertex. Let \hat{v} denote this path set, and use \hat{v} with the parametrizing algorithm to give $K_{\hat{e}}$ its canonical parameterization. Since the final arc from $\hat{v}_{\hat{N}}$ lies on $\ell_{o\hat{e}}$, the vertex \hat{v} has a canonical lift to the $\sigma = \theta_o$ circle of the parametrizing cylinder for $K_{\hat{e}}$. Use the latter for the new distinguished vertex, and use the value given by the parametrizing algorithm for the \mathbb{R} -valued lift of its $\mathbb{R}/(2\pi\mathbb{Z})$ parameter.

Now, take the following for the new concatenating path set: The first path in this set consists solely of the final arc in $\hat{v}_{\hat{N}}$, but traveled in the direction opposite to that

used by \widehat{v} . The second constituent path in the new concatenating path set is $\widehat{v}_{\widehat{N}}$, but traveled in the direction opposite to that by \widehat{v} . The third is \widehat{v}_{N-1} , also traveled in the reverse direction from end to start. Continue in this vein using the paths in \widehat{v} in reverse to return to the original distinguished vertex on ℓ_{oe} . However, instead of using the first constituent path, \widehat{v}_1 , of \widehat{v} in reverse here, use instead the concatenation that first takes \widehat{v}_1 in reverse and then adds to its end the first path, v_1 , in the concatenating path set that lead from the original distinguished vertex to $\ell_{oe'}$. This understood, take the subsequent constituent path in the new concatenating path set to be the second of the constituent paths, v_2 , from the original distinguished vertex to $\ell_{oe'}$. Then, complete the new path set by adding in their given order the constituent paths v_3, \dots, v_N from the original concatenating path set.

It is left as an exercise to verify that the output parameterization for $K_{e'}$ is not changed when this new data is used as input to the parametrizing algorithm.

Case 2 Suppose that the parameterization of K_e is changed by the action of a given integer pair N as depicted in (2–12) and (2–13). In addition, suppose that the new \mathbb{R} -lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate is related to the old via (2–16). Repeating what is said in Part 2 at each run of the parametrizing subroutine finds that the canonical parameterization of $K_{e'}$ is also changed by the action of the integer pair N via the e' version of (2–12) and (2–13).

Case 3 This is the story in the case that the distinguished vertex on the $\sigma = \theta_o$ boundary of the parametrizing cylinder is changed. Agree to use the same parameterization for K_e as the original. Also, agree to take the \mathbb{R} -valued lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate for the new vertex to be that given by the lift of the original distinguished vertex. Take the new concatenating path set as follows: The new version is $\{v_1', v_2, \dots, v_N\}$ where v_1' is the concatenation that travels from the new distinguished vertex to the original distinguished vertex in ℓ_{oe} , against the given orientation of ℓ_{oe} , and then proceeds outward from the old distinguished vertex along v_1 .

The parametrizing algorithm finds no difference between the new and old parameterizations of $K_{e'}$ when using this new input data.

Case 4 The change in just the \mathbb{R} -valued lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the distinguished point is given by the $N \in \mathbb{Z} \cdot Q_e$ version of Case 2. In particular, if $z \in \mathbb{Z}$ and if the new lift is related to the old via the $N = zQ_e$ version of (2–16), then the parameterization of $K_{e'}$ is changed by the action of $N = zQ_e$ via the e' version of (2–12) and (2–13).

Case 5 As is explained momentarily, the story for any given case from the final point in (2–21) can be obtained from that for the special case in which the new concatenating path set shares its respective starting and ending arcs with the original set. Meanwhile, the discussion for the latter case requires the introduction of some new notions and is deferred for this reason to the upcoming Part 7 of this subsection.

To begin the discussion here, suppose that the original concatenating path set is changed subject to (2–18) so as to obtain a new concatenating path set whose starting arc differs from the original. Let $\{v_1', v_2', \dots, v_{N'}'\}$ denote the new set. The resulting parametrization of $K_{e'}$ is not changed when the new concatenating path set has the form $\{v_1'', v_2', \dots, v_{N'}'\}$ where v_1'' is obtained from v_1' by adding two new arcs in ℓ_{oe} at its start, the first leading out along the starting arc of v_1 , and the second returning along this same arc to the distinguished vertex. The third arc in v_1'' is the first arc in v_1' , the fourth arc in v_1'' is the second arc in v_1' , and so on.

Suppose next that the final arc in the new concatenating path set $\{v_1', \dots, v_{N'}'\}$ differs from that in the original. In this case, the new concatenating path set can be modified by increasing the number of its constituent paths by 2 so as not to change the resulting parametrization of $K_{e'}$ but so as to have the desired final arc. To elaborate, let $\tau \in \ell_{oe'}$ denote the original final arc, viewed as a directed arc. Let e'' denote the incident edge to o that labels τ with e' . There are two cases to consider: In the first, direction along τ and direction along the final arc in $v_{N'}'$ give the same orientation to $\ell_{oe'}$. In this case, add a path, $v_{N'+1} \subset \ell_{oe'}$, that starts with the final arc in $v_{N'}'$, proceeds along it in the direction used by $v_{N'}'$, and continues until it first hits a vertex of τ . By assumption, this must be the starting vertex of τ . This understood, continue $v_{N'+1}$ by traversing τ to its end. The final path, $v_{N'+2}$, consists solely of τ , but viewed as in $\ell_{oe''}$.

In the second case, the arc τ and the final arc on $v_{N'}'$ define opposite orientations for $\ell_{oe'}$. In this case, $v_{N'+1}$ starts with the final arc in $v_{N'}'$ and continues along $\ell_{oe'}$ until the it first hits a vertex on τ . In this case, the vertex here is the end vertex. Continue $v_{N'+1}$ by traversing τ backwards to its start. Take $v_{N'+2}$ to be the path that starts at the end vertex of τ , traverses τ to its start and then reverses direction to retrace τ to its end again. Label this path with the edge e'' . A check of the parametrizing algorithm reveals that the canonical parameterization of $K_{e'}$ as defined using $\{v_1', \dots, v_{N'}', v_{N'+1}, v_{N'+2}\}$ is identical to that obtained using $\{v_1', \dots, v_{N'}'\}$.

Part 6 The remainder of the story for the final point in (2–21) requires the introduction of a certain ‘blow up’ of the graph $\underline{\Gamma}_o$. This blow up is a graph, $\underline{\Gamma}_o^*$, with labeled and oriented edges that comes equipped with a canonical ‘blow down’ map to $\underline{\Gamma}_o$ that maps certain edges to vertices of $\underline{\Gamma}_o$. Note that Property 4 from Part 3 is used implicitly in the the constructions that appear in the next few paragraphs.

To start the definition of $\underline{\Gamma}_o^*$, let v denote a vertex in $\underline{\Gamma}_o$ with non-zero integer assignment. Then v labels an oriented, circular subgraph $\ell^{*v} \subset \underline{\Gamma}_o^*$ whose vertices are in 1–1 correspondence with the incident half-arcs to v in $\underline{\Gamma}_o$. For the purpose of describing this correspondence, keep in mind that each vertex in $\underline{\Gamma}_o$ has an even number of incident half-arcs, half oriented towards the vertex and half oriented in the outward direction. The correspondence between the vertices of ℓ^{*v} and the incident half-arcs to v is defined so that the following is true: Let γ denote an incident half-arc to v with edge labels $e \in E_-$ and $e' \in E_+$. If γ is an inbound arc, then γ labels a vertex on ℓ^{*v} and the subsequent vertex in the oriented direction on ℓ^{*v} is labeled by the arc in ℓ_{oe} that follows γ . On the other hand, if γ is an outbound arc, then the subsequent vertex on ℓ^{*v} is labeled by the arc in $\ell_{oe'}$ that precedes γ .

This correspondence has the following consequence: The vertices that are met upon a circumnavigation of ℓ^{*v} correspond alternately to inbound and outbound incident half-arcs to the vertex v . Moreover, if a given vertex corresponds to an inbound arc in $\underline{\Gamma}_o$ with the edge pair label (e, e') , then the subsequent vertex is the subsequent arc in ℓ_{oe} and the previous vertex is the subsequent arc in $\ell_{oe'}$. On the other hand, if the vertex corresponds to an outbound arc with label (e, e') , then the subsequent vertex on ℓ^{*v} corresponds to the previous arc on $\ell_{oe'}$ and the previous vertex corresponds to the previous arc on ℓ_{oe} .

With the definition set, then each arc in ℓ^{*v} is labeled by an ordered pair whose first element is v and whose second is an incident edge to o . To elaborate, an incident edge $e \in E_-$ labels the arc when its starting vertex (as defined by its given orientation) corresponds to an inbound arc to v in $\underline{\Gamma}_o$ whose E_- edge label is e . An edge $e' \in E_+$ labels the arc when its starting vertex corresponds to an outbound arc in $\underline{\Gamma}_o$ whose E_+ label is e' .

The remaining arcs in $\underline{\Gamma}_o^*$ are labeled by pairs of incident edges to o , and this set enjoys a 1–1 correspondence with the arcs in $\underline{\Gamma}_o$ that respects their incident edge labels.

All of this is designed so that the map from $\underline{\Gamma}_o^*$ to $\underline{\Gamma}_o$ that collapses the circles $\{\ell^{*v}\}$ has the following properties: First, it is an orientation preserving map that sends vertices to vertices. Second, the map is 1–1 on neighborhoods of vertices with zero integer assignment. On the other hand, the inverse image of any given vertex in $\underline{\Gamma}_o$ with non-zero integer assignment is the circular subgraph that carries its label. Third, the inverse image of any arc in $\underline{\Gamma}_o$ is an arc that bears its same edge pair label. Finally, each ℓ_{oe} in $\underline{\Gamma}_o$ has a canonical inverse image as an embedded circular subgraph of $\underline{\Gamma}_o^*$. The latter is obtained from the inverse image of the arcs that comprise ℓ_{oe} by adding the arcs from the collection in $\cup_v \ell^{*v}$ whose label contains the edge e . The inverse image of ℓ_{oe} in $\underline{\Gamma}_o^*$ is denoted subsequently as ℓ_{oe}^* .

With their orientations as specified above, each edge labeled loop $\ell^*_{o(\cdot)}$ defines a homology class in $H_1(\underline{\Gamma}^*_o; \mathbb{Z})$ as does each vertex labeled $\ell^{*(\cdot)}$. In this regard, the collection of these classes, $\{[\ell^*_{oe}], [\ell^{*\nu}]\}$, generate the integral homology of $\underline{\Gamma}^*_o$ subject to the following single constraint:

$$(2-22) \quad \sum_{e \in E_+} [\ell^*_{oe}] - \sum_{e \in E_-} [\ell^*_{oe}] = \sum_{\nu} [\ell^{*\nu}].$$

This last identity provides a canonical class in the $\mathbb{Z} \times \mathbb{Z}$ valued cohomology of $\underline{\Gamma}^*_o$. Indeed, this class, ϕ_o , is defined so as to send any given $[\ell^*_{oe}]$ to Q_e . Meanwhile, ϕ_o sends the class $[\ell^{*\nu}]$ to the pair $m_{\nu}P_o$, where m_{ν} is the integer weight assigned to the vertex ν and P_o denotes here the relatively prime integer pair that defines θ_o via (1-8).

Here is one way to view $\underline{\Gamma}^*_o$: The result of attaching one boundary circle of a closed cylinder to each circle from the collection $\{\ell^*_{oe}\} \cup \{\ell^{*\nu}\}$ results in a space whose interior is a multipunctured sphere that is homeomorphic to Γ_o 's component of every small but positive δ version of the $|\theta - \theta_o| < \delta$ portion of C_0 .

Part 7 This part of the subsection describes how the canonical parametrization of $K_{e'}$ changes when the given concatenating path is changed with no change to the first traversed arc and no change either to the last arc or to the direction of traverse on the last arc.

To begin the story, note first that when $\{\nu_1, \dots, \nu_N\}$ is a concatenating path set, then each if its elements can be assigned a canonical inverse image in $\underline{\Gamma}^*_o$. The inverse image of a given ν_k is denoted here as ν^*_k . The latter path lies in the version of $\ell^*_{o(\cdot)}$ that shares ν_k 's incident edge label, it starts with the inverse image of ν_k 's starting arc and ends with the inverse image of ν_k 's ending arc. The next point to make is that the union of the collection $\{\nu^*_k\}_{1 \leq k \leq N}$ of such lifts defines a single directed path whose starting point is a vertex on the version of $\ell^{*(\cdot)}$ that projects to the starting vertex on ν_1 and whose end point is on the version of $\ell^{*(\cdot)}$ that projects to the ending vertex of ν_N . In this regard, the various subpaths that comprise this union are traversed in their given order in their given direction. Let ν^* denote the latter path.

Now, if ν and ν' are concatenating path sets that obey (2-18), and if they share both starting arcs and ending arcs, then the corresponding lifts, ν^* and ν'^* , share the same starting vertex and also share the same ending vertex. Thus, a closed loop, μ^* , is defined in $\underline{\Gamma}^*_o$ by traveling first on ν'^* from start to finish, and then traveling on ν^* in the 'wrong' direction, thus from its ending vertex back to its starting vertex.

The following lemma now summarizes the relation between the ν and ν' versions of the canonical parameterization for $K_{e'}$:

Lemma 2.3 *The action of the integer pair $-\phi_o([\mu^*]) \in \mathbb{Z} \times \mathbb{Z}$ on the canonical parameterization defined by ν gives the canonical parameterization defined by ν' .*

The remainder of this subsection is occupied with the

Proof of Lemma 2.3 The proof is given in six steps.

Step 1 Write $\nu' = \{\nu_1', \dots, \nu_{N'}'\}$ and let $\gamma' \subset \ell_{oe'}$ denote the directed path that first crosses the final arc on $\nu_{N'}'$ from start to finish in the direction used by ν' , and then recrosses this arc in the opposite direction. Label γ' with the incident edge e' . With γ' understood, introduce the concatenated path set given by the following ordered set of constituent paths:

$$(2-23) \quad \{\nu_1', \dots, \nu_{N'}', \gamma', \nu_N^{-1}, \dots, \nu_2^{-1}, \nu_1 \circ \nu_1^{-1}, \nu_2, \dots, \nu_N\}.$$

Here, $\nu_1 \circ \nu_1^{-1}$ denotes the path in ℓ_{oe} obtained by first traveling ν_1 in reverse from its end to its start, and then returning on ν_1 back to its end vertex. The canonical parameterization that the path set in (2-23) defines for $K_{e'}$ is the same as that assigned when using the path ν' .

Step 2 Let τ denote the first arc in ν_1 , viewed as a directed arc that starts at the distinguished vertex ν . Let \hat{e} denote the edge that labels the arc τ with e . Now, let γ denote the following directed path in $\ell_{o\hat{e}}$: Start at the ending vertex on τ and traverse this arc in reverse so as to return to the distinguished vertex. Then, reverse direction and retrace τ to return its end vertex. Thus, $\gamma = \tau \circ \tau^{-1}$. Label γ with the edge \hat{e} .

This labeled, directed path γ appears in the concatenating path set whose ordered constituent paths are

$$(2-24) \quad \{\nu_1', \dots, \nu_{N'}', \gamma', \nu_N^{-1}, \dots, \nu_2^{-1}, \nu_1^{-1}, \gamma\}.$$

Running the parametrizing algorithm on the set in (2-24) provides as output a new pair of parameterization of K_e and lift to \mathbb{R} of the $\mathbb{R}/2\pi\mathbb{Z}$ coordinate of the distinguished point on the $\sigma = \theta_o$ circle in the parametrizing cylinder. The latter pair is obtained from the former by the action of some integer pair as depicted in (2-12), (2-13) and (2-16).

This integer pair is relevant because the conclusions of Step 1 together with the Case 3 story for (2-21) from Part 5 imply that the canonical parametrization of $K_{e'}$ that is obtained using the concatenating path ν' is obtained from the canonical parametrization that is obtained using the concatenating path ν by the action of this same integer pair via the e' version of (2-12) and (2-13).

Step 3 Let \mathcal{C} denote the collection of concatenating path sets that obey the $e' = e$ version of (2–18) and have τ as both first and last arc. For example the set in (2–24) lies in \mathcal{C} . Note that each $\mu \in \mathcal{C}$ defines an integer pair as just described in Step 2, this denoted below by $N_*(\mu)$. To elaborate, any given $\mu \in \mathcal{C}$ can be used with the parametrizing algorithm to define a new parameterization for K_e and a corresponding lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the distinguished point on the $\sigma = \theta_o$ circle of the parametrizing domain. The latter pair is obtained from the original by the action of $N_*(\mu)$ as depicted in (2–12), (2–13) and (2–16).

The assignment $\mu \rightarrow N_*(\mu)$ defines a map from \mathcal{C} to $\mathbb{Z} \times \mathbb{Z}$ and the task is to identify this map. It is useful in this regard to introduce various structures of an algebraic nature to \mathcal{C} of which the first is the product operation, $\wp : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$, given by

$$(2-25) \quad \wp(\mu', \mu) = \{\mu'_1, \dots, \mu'_{N'}, \mu_1, \dots, \mu_N\}.$$

The conclusions from Step 2 imply that

$$(2-26) \quad N_*(\wp(\mu, \mu')) = N_*(\mu) + N_*(\mu').$$

The set $\mathcal{C} \rightarrow \mathcal{C}$ also admits a self map, t , that changes the sign of $N_*(\cdot)$. To elaborate, let γ_0 denote the directed path that starts at the distinguished vertex, runs along τ to its end, and then returns to the distinguished vertex by reversing direction on τ . Label γ_0 with the edge e . Now, let $\mu = \{\mu_1, \dots, \mu_N\} \in \mathcal{C}$ denote a given data set and set

$$(2-27) \quad t(\mu) = \{\gamma_0, \mu_N^{-1}, \mu_{N-1}^{-1}, \dots, \mu_1^{-1}, \gamma\},$$

The parametrizing algorithm finds that $N_*(t(\mu)) = -N_*(\mu)$.

By virtue of (2–26) and (2–27), the set \mathcal{C} contains elements with $N_*(\cdot) = (0, 0)$. Perhaps the simplest such element is the concatenating path set $\{\gamma, \gamma'\}$ with γ and γ' defined as in (2–27).

The next point to make is that the assignment $N_*: \mathcal{C} \rightarrow \mathbb{Z} \times \mathbb{Z}$ is insensitive to certain modifications of a given concatenating path set. To elaborate, suppose again that $\mu = \{\mu_1, \dots, \mu_N\} \in \mathcal{C}$. Suppose, in addition that $j < N$ has been specified. Let e_j denote the version of $\ell_{o(\cdot)}$ that contains μ_j and let τ_1 denote the final arc in μ_j viewed here as a directed arc. Use \hat{e}_j to denote the incident edge to o that is partnered with e_j in the label of τ_1 . Let $\{v_1, \dots, v_{N'}\}$ now denote a concatenating path set with the property that the first arc in v_1 has label \hat{e}_j and directed first arc equal to τ_1 . Let τ_N denote the final arc in v_N , also viewed as a directed arc. With the concatenations $\tau_N^{-1} \circ \tau_N$ and $\tau_1 \circ \tau_1^{-1}$ suitably labeled by edges, the set

$$(2-28) \quad \{\mu_1, \dots, \mu_j, v_1, \dots, v_N, (\tau_N^{-1} \circ \tau_N), v_N^{-1}, \dots, v_1^{-1}, (\tau_1 \circ \tau_1^{-1}), \mu_{j+1}, \dots, \mu_N\}.$$

is an element in \mathcal{C} . Running the parametrizing algorithm find that N_* has the same value on μ as it has on the set in (2–28).

Step 4 Let $\mu_e \subset \ell_{oe}$ denote the path that starts at the vertex v , runs out along τ and continues once around ℓ_{oe} so as to end after crossing τ a second time. The ordered pair $\{\mu_e, \gamma\}$ defines a concatenating path set in \mathcal{C} . The considerations in Case 4 of the preceding part of this subsection imply that N_* has value $\pm Q_e$ on this element, with the $-$ sign occurring if and only if travel along τ from v defines the oriented direction in ℓ_{oe} .

As will now be explained, every other version of $Q_{(\cdot)}$ whose label is an incident edge to o is a value of N_* in \mathcal{C} . To see how this comes about, let $e' \neq e$ denote a given incident edge to o and let $\{v_1, \dots, v_N\}$ denote a concatenating path set with the following properties: First, v_1 is labeled by e , it starts at v and its first arc is τ . Meanwhile, v_N is not labeled by e' but its last arc lies in $\ell_{oe'}$ and the direction along v_N on this last arc defines the given orientation of $\ell_{oe'}$. Let $\mu_{e'}$ denote the path in $\ell_{oe'}$ that starts with this last arc in v_N continues once around $\ell_{oe'}$ to the end of this last arc, and then reverses direction to retrace this last arc back to its starting vertex. Label $\mu_{e'}$ with the edge e' .

The ordered set $\{v_1, \dots, v_N, \mu_{e'}, v_N^{-1}, \dots, v_1^{-1}, \gamma\}$ is an element in \mathcal{C} , and a run of the parametrizing algorithm finds that N_* has the desired value $-Q_{e'}$ on this element.

Step 5 The results from the preceding steps suggest that the value of N_* on any given $\mu \in \mathcal{C}$ depends only on the homology class of the path in $\underline{\Gamma}_o$ that is defined by the union of μ 's constituent paths. However, this conclusion is wrong when $\underline{\Gamma}_o$ contains vertices that correspond to ends of C_0 .

That such is the case can be seen most readily when the vertex at the end of τ is a bivalent vertex. Let v' denote the latter vertex and let τ_1 denote the arc in ℓ_{oe} that shares v' with τ . Thus, both τ and τ_1 are labeled by the edge pair that consists of e and \hat{e} . This understood, then

$$(2-29) \quad \{\tau_1 \circ \tau, \tau \circ \tau^{-1} \circ \tau_1^{-1}\}$$

is a concatenating path set in \mathcal{C} if the first constituent path is labeled by e and the second by \hat{e} . As is explained next, a run of the parametrizing algorithm on this path finds that N_* assigns it either $m_{v'} P_o$ or $-m_{v'} P_o$. Here, $m_{v'}$ is the integer weight accorded v' as a vertex in $\underline{\Gamma}_o$, and P_o is the relatively prime integer pair that is defined by θ_o via (1–8). Meanwhile, the $+$ sign appears if and only if either $e \in E_-$ and travel from v along τ is in the oriented direction in ℓ_{oe} , or else $e \in E_+$ and travel from v along τ goes against the orientation from ℓ_{oe} .

To see why the parametrizing algorithm must give $\pm m_{v'} P_o$, it is necessary to go back to (2–5) for the cases of e and \hat{e} . Let R be some very large number, chosen so that $|s|$ has value R on the end $E \subset C_0$ that is labeled by the vertex v' . Write $P_o = (p_o, p_o')$ and keep in mind that the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about the $|s| = R$ slice of E are $|m_{v'}| p_o$ and $|m_{v'}| p_o$ granted that the latter are oriented by the pull-back of the 1-form $-\alpha$ with α depicted in (1–3).

As can be seen from (2–4), there are two points of the $|s| = R$ circle in E that lie on Γ_o , and one will lie in τ and the other in τ_1 when R is large. Label these points z and z' . Given that $K_{\hat{e}}$ is parametrized using the single element concatenating path set $\{\tau_1 \circ \tau\}$, it then follows that the respective values of the \mathbb{R} -valued functions \hat{v}_e and $\hat{v}_{e'}$ are very close at the respective points in the K_e and $K_{\hat{e}}$ parametrizing cylinders that map to z . They also have very similar values at the respective points that map to z' . The values of w_e and $w_{e'}$ at the respective points in the K_e and $K_{e'}$ parametrizing cylinders that map to z' are also nearly equal since these points are used in (2–14) and (2–15) to define the parameterization of $K_{e'}$.

However, given these near equalities, and given that the pair $(\frac{1}{2\pi} dt, \frac{1}{2\pi} d\varphi)$ has a non-zero integral in $\mathbb{Z} \times \mathbb{Z}$ around the $|s| = R$ circle, the form of (2–5) precludes nearly equal values for w_e and $w_{e'}$ at the respective points in their parametrizing cylinders that map to z . Indeed, the difference in their values is exactly accounted for by the $\mathbb{Z} \times \mathbb{Z}$ action of the appropriately signed version of $m_{v'} P_o$ on the given pair of K_e parameterization and lift to R of the $R/(2\pi\mathbb{Z})$ coordinate of v .

The analysis just done for this simple case can be repeated in the case that the vertex at the end of τ has some $2(k + 1)$ incident half-arcs. To elaborate, each such arc can be directed and then the set suitably labeled as $\{\tau, \tau_1, \dots, \tau_k\}$ so that

$$(2-30) \quad \{\tau_1 \circ \tau, \tau_2 \circ \tau_1^{-1}, \dots, \tau_k \circ \tau_{k-1}^{-1}, \tau \circ \tau^{-1} \circ \tau_k^{-1}\}$$

defines an element in \mathcal{C} . For example, τ_1 is the incident half-arc to v' in ℓ_{oe} that follows τ when traveling along τ in ℓ_{oe} . Meanwhile, τ_2 is the incident half-arc that follows τ_1^{-1} when traveling the latter in the indicated direction on the version of $\ell_{o(\cdot)}$ whose edge shares with e the label of τ_1 . The arc τ_3 is defined by the analogous rule when τ_2 is used in lieu of τ_1 ; in general, τ_k for $k > 3$ is defined by successive applications of this same rule using τ_{k-1} in lieu of τ_1 .

The parametrizing algorithm can be run using (2–30) and then (2–5) can be employed to prove that N_* again has value $\pm m_{v'} P_o$ with the $+$ sign appearing under the same circumstances as in case when v' is bivalent.

The story is much the same for a null-homotopic path in $\underline{\Gamma}_o$ of the following sort: Let v' denote any given arc in $\underline{\Gamma}_o$ and let $2(k + 1)$ denote the number of its incident

half-arcs. Direct and label these as $\{\tau', \tau_1, \dots, \tau_k\}$ by mimicking the scheme used in (2–30) with τ' replacing τ . Now, let $\{\nu_1, \dots, \nu_N\}$ denote a concatenating path set with the following properties: The path ν_1 starts with τ and is labeled by e . Meanwhile, the final arc on ν_N is τ' , but the edge label of ν_N is not that of both τ' and τ_1 . Given these properties, the ordered set

(2–31)

$$\{\nu_1, \dots, \nu_N, \tau_1 \circ \tau', \tau_2 \circ \tau_1^{-1}, \dots, \tau_k \circ \tau_{k-1}^{-1}, \tau^{-1} \circ \tau_k^{-1}, \nu_N^{-1}, \dots, \nu_1^{-1}, \gamma\}$$

defines an element in \mathcal{C} on which N_* has value $\pm m_\nu P_o$ with the \pm determined as before but for the replacement of τ' by τ .

Step 6 As explained previously, a concatenating path set defines a canonical path in the blow up graph $\underline{\Gamma}_o^*$. Indeed, each constituent subpath is specified as a concatenated union of arcs that all lie in a single version of $\ell_{o(\cdot)}$. Thus, each such path has a canonical lift to $\ell_e^* \subset \underline{\Gamma}_o^*$ to give a directed path of concatenated arcs. Moreover, consecutive paths from the given path set lift to paths in $\underline{\Gamma}_o^*$ so that their union defines a directed path with travel starting on the lift of first path and continuing on that of the second.

When $\mu \in \mathcal{C}$, let μ^* denote its lift. Let $\tau^* \subset \underline{\Gamma}_o^*$ denote the inverse image of the first arc in μ , and thus the first arc in μ^* . This is also the final arc in μ^* . Thus, travel on μ^* up to but not including the final appearance of τ^* defines a closed loop, $\hat{\mu}, \subset \underline{\Gamma}_o^*$, that starts and ends at the starting vertex of τ^* .

By virtue of what has been said in the previous steps, the value of N_* on μ is minus the value of the homomorphism ϕ_o on $\hat{\mu}$, and this last conclusion implies the assertion of Lemma 2.3. \square

3 The map from \mathcal{M} to \mathcal{O}_T

Suppose for the moment that \hat{A} is an asymptotic data set with $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$. Granted that \hat{A} also obeys the conditions in (1–16) that guarantee a non-empty version of $\mathcal{M}_{\hat{A}}$, then Theorem 1.2 asserts a diffeomorphism between $\mathcal{M}_{\hat{A}}$ and the product of \mathbb{R} and a space denoted by $\hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. Here, $\hat{O}^{\hat{A}}$ is the part of $O^{\hat{A}}$ from (1–21) where $\text{Aut}^{\hat{A}}$ acts freely. The purpose of this subsection is to describe one such map from $\mathcal{M}_{\hat{A}}$ to $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The map given here is the prototype for a suite of such maps that are used in subsequent sections to describe other components of the space moduli space of multiply punctured, pseudoholomorphic spheres.

3.A The more general context

If \hat{A} is an asymptotic data set, reintroduce from Subsection 1.A the set $\Lambda_{\hat{A}}$ of angles in $[0, \pi]$. Unless noted to the contrary, restrict attention in the remainder of this section to asymptotic data sets with the following properties:

- (3–1) *There is a unique element in \hat{A} that supplies the minimal angle in $\Lambda_{\hat{A}}$, and there is also a unique element in \hat{A} that supplies the maximal angle in $\Lambda_{\hat{A}}$.*

To make this explicit, note that \hat{A} has a unique element that gives $\Lambda_{\hat{A}}$'s minimal angle if and only if the one of following is true:

- (3–2) • *There are no $(1, \dots)$ elements in \hat{A} , the integer ζ_+ is zero, and there is a unique $(0, -\dots)$ element whose integer pair gives the minimal angle in $\Lambda_{\hat{A}}$ via (1–8).*
- *There are no $(1, \dots)$ elements in \hat{A} and $\zeta_+ > 0$.*
- *There is a unique $(1, \dots)$ element in \hat{A} and $\zeta_+ = 0$.*

An analogous set of conditions must hold when there is a unique element in \hat{A} that supplies the maximal angle to $\Lambda_{\hat{A}}$. Note that an asymptotic data set with $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ automatically obeys (3–1).

As is explained next, the data set \hat{A} can be used to define a linear graph, $T^{\hat{A}}$, much like that introduced in Subsection 1.B in the case $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$. As in Subsection 1.B, the graph $T^{\hat{A}}$ is viewed as a closed subinterval in $[0, \pi]$ whose vertices include its endpoints; the edges of $T^{\hat{A}}$ are the closed intervals that run from one vertex to the next. The vertices of $T^{\hat{A}}$ are in 1–1 correspondence with the angles in $\Lambda_{\hat{A}}$ and this correspondence is such that the angle of any vertex in $\Lambda_{\hat{A}}$ is also its angle in $[0, \pi]$.

The edges of $T^{\hat{A}}$ are labeled by integer pairs that are determined using the rules that follow. The notation is such that when e denotes an edge, then $Q_e = (q_e, q_e')$ denotes its integer pair label.

- (3–3) • *Let e denote the edge that contains the minimum angle vertex.*
- (a) *If this angle is positive, then Q_e is the integer pair from the $(0, -\dots)$ element in \hat{A} that supplies this minimal angle to $\Lambda_{\hat{A}}$.*
- (b) *If the minimal angle in $\Lambda_{\hat{A}}$ is zero and $\zeta_+ > 0$, then $Q_e = (0, -\zeta_+)$.*
- (c) *If the minimal angle in $\Lambda_{\hat{A}}$ is zero and $\zeta_+ = 0$, then $Q_e = -\varepsilon P$ given that the element from \hat{A} that supplies this angle to $\Lambda_{\hat{A}}$ has the form $(1, \varepsilon, P)$.*

- Let o denote a bivalent vertex in $T^{\hat{A}}$ and let e and e' denote its incident edges with e connecting o to a vertex with angle less than that of o . Then $Q_e = Q_{e'} + P_o$ where P_o is obtained by subtracting the sum of the integer pairs from the $(0, -, \dots)$ elements from \hat{A} that supply o 's angle to $\Lambda_{\hat{A}}$ from the some of the integer pairs from the $(0, +, \dots)$ elements from \hat{A} that supply o 's angle to $\Lambda_{\hat{A}}$.

In [15, Theorem 1.3] asserts that $\mathcal{M}_{\hat{A}}$ is non-empty if and only if the conditions in (1–16) hold for the current version of $T^{\hat{A}}$. Assume for the remainder of this subsection that \hat{A} obeys (3–1) and that the conditions of (1–16) hold for $T^{\hat{A}}$. What follows describes a generalization of the space $O^{\hat{A}}$ that appears in (1–21). This space is then used to parametrize a certain subspace in $\mathcal{M}_{\hat{A}}$.

To begin the description of $O^{\hat{A}}$, introduce $\hat{A}_* \subset \hat{A}$ to denote the subset of 4–tuples whose integer pairs define the non-extremal angles in $\Lambda_{\hat{A}}$. Associate to each $u \in \hat{A}_*$ a copy, \mathbb{R}_u of the affine line and let

$$(3-4) \quad \mathbb{R}^{\hat{A}} \subset \text{Maps}(\hat{A}_*; \mathbb{R})$$

denote the subspace of points where distinct u and u' from \hat{A}_* have distinct images in $\mathbb{R}/(2\pi\mathbb{Z})$ in the case that their integer pairs defined the same angle via (1–8). Let \mathbb{R}_- denote an auxiliary copy of \mathbb{R} .

As is explained next, there is an action of $\text{Maps}(\hat{A}_*; \mathbb{Z})$ on

$$(3-5) \quad \mathbb{R}_- \times \text{Maps}(\hat{A}_*; \mathbb{R})$$

that preserves $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. This action is trivial on \mathbb{R}_- . To describe the action on the second factor, let $u \in \hat{A}_*$ and let z_u denote the corresponding generator of $\text{Maps}(\hat{A}_*; \mathbb{Z})$. If $x \in \text{Maps}(\hat{A}_*; \mathbb{R})$, then $(z_u \cdot x)(\hat{u}) = x(\hat{u})$ if the integer pair from \hat{u} defines an angle via (1–8) that is less than that defined by the integer pair from u . Such is also the case when the two angles are equal except if $\hat{u} = u$. If $\hat{u} = u$, then $(z_u \cdot x)(u) = x(u) - 2\pi$. When the integer pair from \hat{u} defines an angle that is greater than that of the pair from u , then $(z_u \cdot x)(\hat{u})$ is obtained from $x(\hat{u})$ by adding

$$(3-6) \quad -2\pi \varepsilon_u \frac{p_{u'} p_{\hat{u}} - p_u p_{\hat{u}'}}{q_{\hat{e}}' p_{\hat{u}} - q_{\hat{e}} p_{\hat{u}'}}$$

where the notation is as follows: First, \hat{e} labels the edge that contains \hat{u} as its largest angle vertex and $(q_{\hat{e}}, q_{\hat{e}'})$ is the integer pair that is associated to \hat{e} . Meanwhile, $(p_u, p_{u'})$ is the integer pair from u and $(p_{\hat{u}}, p_{\hat{u}'})$ is that from \hat{u} . Finally, $\varepsilon_u \in \{\pm\}$ is the second entry in the 4–tuple u .

Also needed is the action of $\mathbb{Z} \times \mathbb{Z}$ on the space in (3–5) that is described just as in Step 2 of Part 1 in Subsection 1.A for the latter’s version of $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. This $\mathbb{Z} \times \mathbb{Z}$ action commutes with that of $\text{Maps}(\hat{A}_*; \mathbb{Z})$. This understood, define

$$(3-7) \quad O^{\hat{A}} \equiv [\mathbb{R}_- \times \mathbb{R}^{\hat{A}}]/[(\mathbb{Z} \times \mathbb{Z}) \times \text{Maps}(\hat{A}_*; \mathbb{Z})].$$

As is explained next, $O^{\hat{A}}$ is a smooth manifold.

To see that $O^{\hat{A}}$ is smooth, it is sufficient to verify that each point in (3–5) has the same stabilizer under the action of the group in (3–7). As is explained below, this stabilizer is a copy of \mathbb{Z} whose generator projects to $\text{Maps}(\hat{A}_+; \mathbb{Z})$ as $-\sum_u z_u$ and projects to $\mathbb{Z} \times \mathbb{Z}$ as the integer pair Q_e with e here denoting the edge in $T^{\hat{A}}$ that contains the smallest angle vertex. To prove such is the case, suppose that $\tau \in (\tau_-, x)$ is a point in the space depicted in (3–5), and that $g = (N, z)$ fixes τ . As τ_- is fixed, the pair N must have the form $r_- Q_e$ where Q_e here denotes the integer pair that is assigned to the edge with the smallest angle vertex and where r_- is a fraction whose denominator is the greatest common divisor of these same two integers.

To proceed, suppose next that $\hat{u} \in \hat{A}_*$ supplies an integer pair that gives the second smallest angle in $\Lambda_{\hat{A}}$. Since the pair $Q_{\hat{e}} = (q_{\hat{e}}, q_{\hat{e}}')$ in the relevant version of (1–19) is the same as the just defined Q_e , it follows that g fixes both τ_- and $x(\hat{u})$ if and only if $z(\hat{u}) = -r_-$. Note that this means that r_- is an integer since $z(\hat{u})$ is an integer.

Suppose next that $\hat{u} \in \hat{A}_*$ supplies an integer pair that gives the third smallest angle in $\Lambda_{\hat{A}}$. The corresponding $x(\hat{u})$ is then fixed if and only if

$$(3-8) \quad z(\hat{u}) + 2\pi \sum_u z(u) \varepsilon_u \frac{p_u' p_{\hat{u}} - p_u p_{\hat{u}}'}{q_{\hat{e}}' p_{\hat{u}} - q_{\hat{e}} p_{\hat{u}}'} + r_- \frac{q_e' p_{\hat{u}} - q_e p_{\hat{u}}'}{q_{\hat{e}}' p_{\hat{u}} - q_{\hat{e}} p_{\hat{u}}'} = 0,$$

where the sum is over the elements in \hat{A}_* whose integer pair defines the second smallest angle in $\Lambda_{\hat{A}}$. Here, e is the edge that has the smallest angle vertex in $T^{\hat{A}}$ and \hat{e} is the edge that contains the second and third smallest angle vertices. To make something of (3–8), note that each $z(u)$ that appears in the sum is $-r_-$. In addition, (3–3) identifies $\sum_u \varepsilon(p_u, p_u')$ with $Q_e - Q_{\hat{e}}$. These points understood, then (3–8) asserts that $x(\hat{u})$ is fixed by g if and only if $z(\hat{u}) = -r_-$.

One can now continue in this vein in an inductive fashion through the elements from \hat{A}_* with integer pairs that give successively larger angles in (1–8). In particular, an application of (3–3) at each step finds that $g = (N, z)$ fixes (τ_-, x) if and only if $N = r_- Q_e$ and z sends each element in \hat{A}_* to $-r_-$. This straightforward task is left to the reader.

Define the group $\text{Aut}^{\hat{A}}$ to be the group of 1–1 maps of \hat{A}_* to itself that only mix elements with identical 4–tuples. This group acts smoothly on $O^{\hat{A}}$. Set $\hat{O}^{\hat{A}} \subset O^{\hat{A}}$ to be the set of points where the action is free. Propositions 4.4 and 4.5 say more about $\hat{O}^{\hat{A}}$.

Theorem 1.2 in Subsection 1.B now has the following generalization:

Theorem 3.1 *Let \hat{A} denote an asymptotic data set that obeys (3–1) and whose graph $T^{\hat{A}}$ obeys the conditions in (1–16). The subspace of subvarieties in $\mathcal{M}_{\hat{A}}$ whose graph from Subsection 2.A is linear is a closed submanifold of $\mathcal{M}_{\hat{A}}$ that is diffeomorphic to $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. Moreover, there is a diffeomorphism that intertwines the \mathbb{R} action on $\mathcal{M}_{\hat{A}}$ as the group of constant translations along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$ with the \mathbb{R} action on $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ as the group of constant translations along the \mathbb{R} factor in the latter space.*

Let $\mathcal{M} \subset \mathcal{M}_{\hat{A}}$ denote the indicated subspace. The remaining subsections describe a map from \mathcal{M} to $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ that is seen in the next section to be a diffeomorphism between these two spaces.

3.B The map to \mathbb{R}

Fix $C \in \mathcal{M}$. The image of C in the \mathbb{R} factor of $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ is the simplest part of the story. Even so, its definition is different in each of the three cases of (3–2). These are treated in turn.

Case 1 Let $E \subset C$ denote the convex side end where the $|s| \rightarrow \infty$ limit of θ gives the minimal angle in $\Lambda_{\hat{A}}$. Associated to the latter is the real number b that appears in (2–4). In this regard, note that the integer n_E that appears here is zero and thus, b must be positive since θ_E is the infimum of θ on C . This understood, the map to \mathbb{R} sends the subvariety C to $-\zeta^{-1} \ln(b)$ where $\zeta \equiv \sqrt{6 \sin^2 \theta_E (1 + 3 \cos^2 \theta_E)} / (1 + 3 \cos^4 \theta_E)$.

Case 2 In this case, C intersects the $\theta = 0$ cylinder at a single point and C 's image in the \mathbb{R} factor of $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ is the s coordinate of this point.

Case 3 In this case, C has a single end whose constant $|s|$ slices limit to the $\theta = 0$ cylinder as $|s| \rightarrow \infty$. This end defines the positive constant \hat{c} that appears in (1–9). Note that the integers p and p' that appear in (1–9) comprise the pair from the $(1, \dots)$ element in \hat{A} . The image of C in \mathbb{R} is $-(\sqrt{\frac{3}{2}} + \frac{p'}{p})^{-1} \ln(\hat{c})$.

3.C The map to $O^{\hat{A}} / \text{Aut}^{\hat{A}}$

The definition of C 's assigned point in $O^{\hat{A}} / \text{Aut}^{\hat{A}}$ requires a preliminary digression to set the stage. To start, let o denote a bivalent vertex in $T^{\hat{A}}$ and let $\hat{A}_o \subset \hat{A}$ denote the subset of elements whose integer pair defines o 's angle via (1–8). Use n_o in what follows to denote the number of elements in \hat{A}_o .

Introduce S_o to denote the subspace of 1–1 maps from \hat{A}_o to $\mathbb{R}/(2\pi\mathbb{Z})$. The inverse image of S_o in $\text{Maps}(\hat{A}_o; \mathbb{R})$ is the part of the latter space that contributes to $\mathbb{R}^{\hat{A}}$. The components of S_o are in 1–1 correspondence with the set of cyclic orderings of the elements in \hat{A}_o . If $u_o \in \hat{A}_o$ is a chosen distinguished element, then any given component $S \subset S_o$ can be identified with

$$(3-9) \quad \mathbb{R}/(2\pi\mathbb{Z}) \times \Delta_o,$$

where $\Delta_o \subset \mathbb{R}^{n_o}$ is the open, $n_o - 1$ dimensional simplex in the positive quadrant where the sum of the coordinates is equal to 2π . Here, the identification in (3–9) is obtained as follows: Identify the cyclic ordering that defines S with the linear ordering that has u_o as the last element. This identification provides a 1–1 correspondence between \hat{A}_o and the set $\{1, \dots, n_o\}$. Granted this, the identification of S with the space in (3–9) is given by identifying a given $(\tau, (r_1, \dots)) \in \mathbb{R}/(2\pi\mathbb{Z}) \times \Delta_o$ with the map from \hat{A}_o to $\mathbb{R}/(2\pi\mathbb{Z})$ that sends the k 'th element in \hat{A}_o to the $\text{mod } (2\pi\mathbb{Z})$ reduction of $\tau + r_1 + \dots + r_k$.

Now suppose that F is a space with an action of $\text{Maps}(\hat{A}_o; \mathbb{Z})$. Let F_o denote the associated fiber bundle over S_o with fiber F . Thus,

$$(3-10) \quad F_o = (\text{Maps}(\hat{A}_o; \mathbb{R}) \times F) / \text{Maps}(\hat{A}_o; \mathbb{Z}).$$

The identification between S and the space in (3–9) is covered by one between $F_o|_S$ and

$$(3-11) \quad (\mathbb{R}_o \times F) / \mathbb{Z}_o \times \Delta_o,$$

where \mathbb{R}_o is a copy of \mathbb{R} while \mathbb{Z}_o is a copy of \mathbb{Z} whose generator acts on \mathbb{R}_o as the translation by -2π and on F as that of $\sum_u z_u \in \text{Maps}(\hat{A}_o; \mathbb{Z})$.

All of this has the following implications: Fix a distinguished element in each version of \hat{A}_o . This done, then any given component of $O^{\hat{A}}$ is identified with

$$(3-12) \quad (\times_o \Delta_o) \times [\mathbb{R}_- \times (\times_o \mathbb{R}_o)] / [(\mathbb{Z} \times \mathbb{Z}) \times (\times_o \mathbb{Z}_o)],$$

where the notation \times_o indicates a product that is labeled by the vertices in $T^{\hat{A}}$, and where the group actions are as follows: First, $N = (n, n') \in \mathbb{Z} \times \mathbb{Z}$ acts on \mathbb{R}_- as

before while acting on any given version of $\mathbb{R}_{\hat{\delta}}$ as translation by the element in (1–20). Meanwhile, any given \mathbb{Z}_o acts trivially on \mathbb{R}_- and also trivially on $\mathbb{R}_{\hat{\delta}}$ in the case that the angle of $\hat{\delta}$ is less than that of o . On the other hand, $1 \in \mathbb{Z}_o$ acts on \mathbb{R}_o as the translation by -2π and it acts on $\mathbb{R}_{\hat{\delta}}$ in the case that $\hat{\delta}$'s angle is greater than o 's angle as the translation by

$$(3-13) \quad -2\pi \frac{p_o' \hat{p}_{\hat{\delta}} - p_o \hat{p}'_{\hat{\delta}}}{q_{\hat{e}}' \hat{p}_{\hat{\delta}} - q_{\hat{e}} \hat{p}'_{\hat{\delta}}},$$

where the notation is as follows: First, $(\hat{p}_{\hat{\delta}}, \hat{p}'_{\hat{\delta}})$ is the relatively prime pair of integers that defines $\hat{\delta}$'s angle in (1–8). Second $P_o \equiv (p_o, p_o')$ is obtained by subtracting the sum of integer pairs from the $(0, -, \dots)$ elements in \hat{A}_o from the sum of those from the $(0, +, \dots)$ elements. Thus, P_o is the integer pair from o 's version of the second point in (3–3).

Granted what has just been said, if a distinguished element has been chosen in each version of \hat{A}_o , then a point for C in $O^{\hat{A}} / \text{Aut}^{\hat{A}}$ is obtained from an assigned point in

$$(3-14) \quad \mathbb{R}_- \times (\times_o(\mathbb{R}_o \times \Delta_o))$$

together with an assigned cyclic ordering for each version of \hat{A}_o . The story on these assignments appears below in four parts.

With the preceding understood, the digression is now ended.

Part 1 The assignment to C of a point in the \mathbb{R}_- factor of (3–14) requires the choice of a parameterization for the component of $C_0 - \Gamma$ whose labeling edge has the minimal angle vertex. With such a parameterization in hand, let w denote the associated function from (2–5). The \mathbb{R}_- assignment of C is the value of the expression

$$(3-15) \quad -\frac{1}{2\pi} \alpha_Q(\sigma) \int_{\mathbb{R}/(2\pi\mathbb{Z})} w(\sigma, v) dv$$

as computed using any σ that lies between the vertex angles on the edge in question. Here, Q is the integer pair that labels this same edge.

Part 2 As has most probably been noted, there is an evident ‘forgetful’ map from a graph T with a correspondence in (C_0, ϕ) to the graph $T^{\hat{A}}$ that is obtained by dropping the circular graph labels from the vertices in T . Thus, the map is an isomorphism of underlying graphs that respects vertex angles and the integer pair labels of the edges. This forgetful map is used implicitly in what follows to identify respective vertices and respective edges in the two graphs.

Let o denote a given bivalent vertex in T . Fix a 1–1 map from the vertices in the T version of $\underline{\Gamma}_o$ to \hat{A}_o with the following property: A vertex in $\underline{\Gamma}_o$ is assigned an element $u \in \hat{A}_o$ if and only if the integer label for the vertex is equal to $\varepsilon_u m_u$ where ε_u is the second entry to u and m_u is the greatest common divisor of the integer pair entry for u . Granted this 1–1 correspondence, define a cyclic ordering of \hat{A}_o using the ordering of the vertices on the circular graph $\underline{\Gamma}_o$ as they are met on an oriented circumnavigation.

Label the arcs in T 's version of $\underline{\Gamma}_o$ by integers consecutively from 1 to n_o so that the labeling gives the order in which arcs are crossed when circumnavigating $\underline{\Gamma}_o$ in its oriented direction when starting at the vertex that is mapped to the distinguished element in \hat{A}_o .

Having done the above, fix a correspondence of T in (C_0, ϕ) and thus define T_C .

Part 3 Let k denote an integer that labels a given arc on $\underline{\Gamma}_o$. To obtain the assignment for C in the k 'th coordinate factor of the simplex Δ_o in (3–16), first integrate the pull-back of $(1 - 3 \cos^2 \theta)d\varphi - \sqrt{6} \cos \theta dt$ over the k 'th arc in the corresponding C_0 version of the graph Γ_o . Then, divide the result by $2\pi\alpha_Q(\theta_o)$ where θ_o is o 's angle and where Q is the integer pair that labels the edge in $T^{\hat{A}}$ with largest angle θ_o .

Part 4 The assignment to C of a point in the $\times_o \mathbb{R}_o$ factors in (3–14) is made in an iterative fashion that starts with the minimal angle bivalent vertex in $T^{\hat{A}}$ and proceeds from vertex to vertex in their given order along $T^{\hat{A}}$. Here is the basic iteration step: Let o denote a bivalent vertex in $T^{\hat{A}}$, let e denote the edge in T_C that contains o as its largest angle vertex and let e' denote the edge that contains o as its smallest angle vertex. Suppose that e 's component $K_e \subset C_0 - \Gamma$ has been assigned a canonical parameterization. Let v denote the missing point on the $\sigma = \theta_o$ circle in the corresponding parametrizing cylinder that corresponds to the distinguished element in \hat{A}_o . Choose a lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of v . Use this lift as the value that is assigned to C in the factor \mathbb{R}_o that appears in (3–14). To initiate the next iteration round, use this same lift and the canonical parameterization of K_e in the manner described by Part 2 of Subsection 2.C to define the canonical parameterization of the component $K_{e'} \subset C_0 - \Gamma$.

The chosen parameterization from Part 1 above should be used for the canonical parameterization when starting the iteration at the smallest angled bivalent vertex.

3.D The invariance of the image in $O^{\hat{A}}$

The point just assigned to C in (3–14) required the following suite of choices:

- (3–16) • A distinguished element in each version of \hat{A}_o .
- A suitably constrained 1–1 correspondence from each $o \in T$ version of the vertex set in $\underline{\Gamma}_o$ to the corresponding \hat{A}_o .
 - A correspondence of T in (C_0, ϕ) .
 - A parameterization for the component of $C_0 - \Gamma$ whose corresponding edge contains the minimal angled vertex in T_C .
 - A lift to \mathbb{R} made for each multivalent vertex in T_C . In particular, let o denote such a vertex and let e denote the edge that contains o as its maximal angled vertex. The lift for o is that of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the missing point on the $\theta = \theta_o$ boundary of the parameterizing cylinder for the canonical parameterization of K_e that corresponds to the distinguished element in \hat{A}_o .

This subsection explains why the image in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ of the point given C in (3–14) is insensitive to the choices that are described in (3–16). These choices are considered below in the order 5, 4, 1, 2, 3. Note for future reference that the arguments given below proves that the image of C in $O^{\hat{A}}$ is already insensitive to changes that are described points 5, 4 and 1. In any event, the discussion that follows is in five parts, one for each of the points in (3–16).

Part 1 To start the explanation for the fifth point in (3–16), suppose that o is a multivalent vertex in T_C . Let θ_o denote o 's angle, let e denote the edge that has o as its largest angle vertex and let e' denote the edge that has o as its smallest angle vertex. Note that any parameterization of K_e has a ‘distinguished’ missing point on the $\sigma = \theta_o$ boundary circle, this the point that corresponds to the chosen distinguished point in \hat{A}_o .

Suppose that $\tau_o \in \mathbb{R}$ is the original lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the distinguished point on the $\sigma = \theta_o$ boundary of the parametrizing cylinder for K_e . Now change this lift to $\tau_o - 2\pi$. Such a change has no affect on the assignment of C in \mathbb{R}_- or in $\mathbb{R}_{\hat{\delta}}$ if $\hat{\delta}$ is a bivalent vertex with angle less than θ_o . Of course, it changes the assignment in \mathbb{R}_o from τ_o to $\tau_o - 2\pi$.

The affect of this change on the remaining $\mathbb{R}_{(c)}$ factors in (3–14) is examined vertex by vertex in order of increasing angle. To start this process, invoke the discussion in Case 4 in Part 5 of Subsection 2.C, to conclude that the change $\tau_o \rightarrow \tau_o - 2\pi$ changes the parameterization of $K_{e'}$ by the action of the integer pair Q_e . In this regard, note that the equality $Q_e = Q_{e'} + P_o$, implies that this change is identical to that obtained by the action of P_o . The latter view proves the more useful for what follows.

Let o' now denote the vertex with the largest angle on e' and suppose that o' is bivalent. Let τ^{old} denote the original assignment to C in $\mathbb{R}_{\hat{o}}$. As applied now to o' , the assertion of (2–16) and the conclusions of Case 2 of Part 5 in Subsection 2.C imply that there is a choice for the lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the distinguished point on the $\sigma = \theta_{o'}$ circle for the new parametrizing cylinder whose relation to the old is given by the $N = P_o$ version of

$$(3-17) \quad \tau^{\text{new}} = \tau^{\text{old}} - 2\pi \frac{\alpha_N(\theta_{o'})}{\alpha_{Q_{e'}}(\theta_{o'})} - 2\pi k_{o'}$$

with $k_{o'} \in \mathbb{Z}$.

To continue, let \hat{o} denote the vertex that shares an edge with o' but has greater angle. Suppose, for the sake of argument, that \hat{o} is bivalent. Let \hat{e} denote the edge between \hat{o} and o' . Apply Case 2 of Part 5 in Subsection 2.C here together with the identity $Q_{e'} = Q_{\hat{e}} + P_{o'}$ to conclude that the canonical parameterization for $K_{\hat{e}}$ is changed by the action of the integer pair $P_o + k_{o'} P_{o'}$. This then means that the old and new assignments for C in $\mathbb{R}_{\hat{o}}$ are related by the $\sigma = \theta_{\hat{o}}$ version of

$$(3-18) \quad \tau^{\text{new}} = \tau^{\text{old}} - 2\pi \frac{\alpha_{P_o}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} - 2\pi k_{o'} \frac{\alpha_{P_{o'}}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} - 2\pi k_{\hat{o}}$$

with $k_{\hat{o}} \in \mathbb{Z}$.

Continue in this vein vertex by vertex in order of increasing angle with applications of Case 2 in Part 5 of Subsection 2.C so as to obtain the generalization of (3–18) that is summarized as the next lemma.

Lemma 3.2 *The bivalent vertices of T with angles greater than θ_o label a collection, $\{k_{(\cdot)}\}$, of integers with the following significance: Let \hat{o} denote any bivalent vertex in T with angle greater than θ_o , and let \hat{e} denote the edge that has \hat{o} as its maximal angled vertex. Then, the change, $\tau_o \rightarrow \tau_o - 2\pi$ of the assignment to C in \mathbb{R}_o changes the assignment of C in $\mathbb{R}_{\hat{o}}$ by the rule*

$$(3-19) \quad \tau_{\hat{o}} \rightarrow \tau_{\hat{o}} - 2\pi \frac{\alpha_{P_o}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} - 2\pi \sum'_{o'} k_{o'} \frac{\alpha_{P_{o'}}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} - 2\pi k_{\hat{o}},$$

where the prime on the summation symbol is meant to indicate that the sum is over those bivalent vertices whose angles lie between θ_o and $\theta_{\hat{o}}$.

Together, (1–8) and (3–19) imply that the action of $\times_o \mathbb{Z}_o$ on (3–14) rectifies any change in the choices that are described in the fifth point of (3–16). Thus, the image of C in $O^{\hat{A}}$ is insensitive to any such change.

Part 2 Consider now the effect of a change as described by the fourth point in (3–16). The analysis of this change starts with its affect on the various \mathbb{R}_o factors. This can be analyzed vertex by vertex in order of increasing angle by successively invoking the observations from Case 2 in Part 5 of Subsection 2.C. In particular, with the help of (1–8), these observations dictate the following: Suppose that $N = (n, n') \in \mathbb{Z} \times \mathbb{Z}$ and that the chosen parameterization of the smallest angled component of $C_0 - \Gamma$ is changed by the action of N as described in (2–12) and (2–13). This changes C 's assignment in the $\times_o \mathbb{R}_o$ factor of (3–14) by the action on the original assignment in $\times_o \mathbb{R}_o$ of an element of the form $(N, \dots) \in (\mathbb{Z} \times \mathbb{Z}) \times (\times_o \mathbb{Z}_o)$.

As can be seen directly from (2–13) and (3–15), the change induced by N on the parameterization of the $C_0 - \Gamma$ component with the smallest θ values changes the assignment of C in \mathbb{R}_- by subtracting $2\pi(n'q_e - nq_e')$. Meanwhile, the integer pair N acts affinely on \mathbb{R}_- as described in the preceding subsection by subtracting $2\pi(n'q_e - nq_e')$.

Taken together, the conclusions of these last two paragraphs imply that the image of C in $O^{\hat{A}}$ is insensitive to any change that is described by the fourth point in (3–16).

Part 3 This part analyzes the affect on C 's assigned point in $O^{\hat{A}}$ of a change in the chosen distinguished element from any given version of $\hat{A}_{(\cdot)}$. The subsequent five steps prove that this assignment is not changed.

Step 1 The new and old version choices for distinguished elements result in respective new and old assignments of a point for C in $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. To be explicit, a choice of cyclic ordering and distinguished element in each version of $\hat{A}_{(\cdot)}$ identifies $\times_{\hat{\delta}}(\mathbb{R}_{\hat{\delta}} \times \Delta_{\hat{\delta}})$ with a component of $\mathbb{R}^{\hat{A}}$ as follows: The distinguished element and cyclic ordering of a given \hat{A}_o endow its elements in their cyclic order with a labeling by the integers in the set $\{1, \dots, n_o\}$ so that n_o labels the distinguished element. Granted this linear ordering, let u_k denote the k 'th element in some given \hat{A}_o . Then all maps in $\mathbb{R}^{\hat{A}}$ that arise from a point in $\times_{\hat{\delta}}(\mathbb{R}_{\hat{\delta}} \times \Delta_{\hat{\delta}})$ with a given $(\tau, (r_1, \dots)) \in \mathbb{R}_o \times \Delta_o$ send u_k to $\tau + r_1 + \dots + r_k - 2\pi \in \mathbb{R}$.

This identification of $\times_{\hat{\delta}}(\mathbb{R}_{\hat{\delta}} \times \Delta_{\hat{\delta}})$ with a component of $\mathbb{R}^{\hat{A}}$ results in the identification described earlier between the space in (3–12) and a component of $O_{\hat{A}}$.

Step 2 Let o denote a bivalent vertex in T_C , let $u \in \hat{A}_o$ the original choice for a distinguished element, and u' denote the new choice. As explained in Step 1, these choices result in respective points for C in $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. Such a change has no affect on

the assignment to \mathbb{R}_- . Let x and x' denote the respective original and new assignments of $\mathbb{R}^{\hat{A}}$.

Because the change from u to u' has no affect on C 's assignment in any $\mathbb{R}_{\hat{\delta}}$ or $\Delta_{\hat{\delta}}$ factor in (3–12) when $\theta_{\hat{\delta}} < \theta_o$, the maps x and x' send any given $\hat{u} \in \hat{A}_*$ to the same point in \mathbb{R} if the integer pair component of \hat{u} is less than θ_o .

Step 3 This step describes the image via x' of an element whose integer pair gives θ_o . This task requires an analysis of the change to C 's assignment to the Δ_o and \mathbb{R}_o factors in (3–14). Start this analysis by labeling the arcs in $\underline{\Gamma}_o$ from 1 through n_o with the first arc starting at u 's vertex in $\underline{\Gamma}_o$. Suppose that the arc that ends at the vertex that corresponds to u' is the k 'th arc. Let $r = (r_1, r_2, \dots) \in \Delta_o$ denote the original assignment for C . Then $r' = (r_{k+1}, r_{k+2}, \dots)$ gives the new assignment.

Consider next the change in the assignment to \mathbb{R}_o . The description here is simplest if it is agreed beforehand to keep the original parameterization of the component of $C_0 - \Gamma$ labeled by the smallest angle vertex in T_C and also to keep C 's assigned point in each $\mathbb{R}_{\hat{\delta}}$ in the case that $\theta_{\hat{\delta}} < \theta_o$. Granted this, let e now denote the edge that has o as its largest angle vertex. Then the canonical parametrization of the component $K_e \subset C_0 - \Gamma$ is unchanged. With this last point understood, take the lift to \mathbb{R} of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the missing point on the $\sigma = \theta_o$ circle that maps to u' to be that obtained from the original lift by adding $r_1 + \dots + r_k$.

The preceding conclusions have the following implications for the map x' : Let u_j denote the j 'th element in the original linear ordering of the set \hat{A}_o . Then $x(u_j) = x'(u_j)$ in the case that $j \in \{k + 1, \dots, n_o\}$ and $x'(u_j) = x(u_j) + 2\pi$ in the case that $j \in \{1, \dots, k\}$.

Step 4 This step describes the value of x' on the elements in \hat{A}_* whose integer pair defines an angle that is greater than θ_o . To start the analysis, suppose that \hat{o} is a bivalent vertex with angle greater than o 's angle. There is no change to the simplex $\Delta_{\hat{\delta}}$ with the change from u to u' . To consider the affect on the assignment in (3–12)'s factor $\mathbb{R}_{\hat{\delta}}$, let e' denote the edge in T_C with o as its smallest angle vertex. According to what is said in Case 3 of Part 5 and Lemma 2.3 of Part 6 from Subsection 2.C, the parameterization for $K_{e'}$ is changed by the action of the integer pair

$$(3-20) \quad N = \sum_{1 \leq j \leq k} P_{u_j}.$$

Given the discussion from Part 2 of Subsection 2.C and the fourth point in (3–16), this then has the following consequence: The new assignment for C in (3–14) can be made

consistent with what has been said so far and such that the new and old assignments to the $\theta_{\hat{o}} > \theta_o$ versions of $\mathbb{R}_{\hat{o}}$ change by the addition of

$$(3-21) \quad 2\pi \sum_{1 \leq j \leq k} \frac{p_{u_j}' \hat{p}_{\hat{o}} - p_{u_j} \hat{p}'_{\hat{o}}}{q_{\hat{e}}' \hat{p}_{\hat{o}} - q_{\hat{e}} \hat{p}'_{\hat{o}}}.$$

The implication for the map x' is as follows: Let $\hat{u} \in \hat{A}_*$ denote an element whose integer pair component defines via (1-8) an angle that is greater than θ_o . Then $x'(\hat{u})$ is obtained from $x(\hat{u})$ by acting on the former by adding the term in (3-21).

Step 5 The conclusions of the previous steps imply that the map x' is obtained from x by acting on the former with the element $-\sum_{1 \leq j \leq k} z(u_j)$ from $\text{Maps}(\hat{A}_*; \mathbb{Z})$.

Part 4 What follows in this part is an explanation of why the image of C in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ is insensitive to the change that is described by the second point in (3-16). To start, suppose that o is a multivalent vertex in T_C and the original correspondence between to the vertex set of $\underline{\Gamma}_o$ and \hat{A}_o is changed to some new assignment. The new assignment is thus obtained by composing the original with a 1-1 self map of \hat{A}_o that only permutes elements with identical 4-tuples. Let $\iota: \hat{A}_o \rightarrow \hat{A}_o$ denote this permutation.

As explained in Step 1 of the Part 3, the original point for C in $\mathbb{R}_- \times (\times_o \mathbb{R}_o)$ corresponds to an assigned point $(\tau_-, x) \in \mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. The change induced by ι in the identification between $\underline{\Gamma}_o$ and \hat{A}_o to C 's assigned point in $\mathbb{R}_- \times (\times_o \mathbb{R}_o)$ will change the assigned point in $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. Let (τ_-, x') denote this new point. The task ahead is to prove that (τ_-, x') and (τ_-, x) define the same point in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$.

To start this task, let \hat{e} denote the edge in T that contains T 's minimal angled vertex. Agree to keep the original parameterization of $K_{\hat{e}}$. This then means that $\tau_- = \tau_-'$. One can also arrange that $x'(\hat{u}) = x(\hat{u})$ if the integer pair from \hat{u} defines an angle via (1-8) that is less than θ_o . This is done by an iterative scheme that keeps the canonical parameterization unchanged on each component of $C_0 - \Gamma$ where $\theta < \theta_o$. Indeed, suppose that \hat{o} is a bivalent vertex, and suppose that the original parameterization is used for the component labeled by the edge with \hat{o} as its largest angle vertex. Then, the original parameterization will arise on the component labeled by the edge with \hat{o} as its smallest angle vertex if there is no change made to the lift at \hat{o} as described in the fourth point of (3-16).

To consider the behavior of x' on elements whose integer pair gives an angle as large as θ_o , remark that the new identification between $\underline{\Gamma}_o$ and \hat{A}_o has two affects. First, it changes the distinguished point on the $\sigma = \theta_o$ boundary of the parametrizing cylinder

for the component of $C_0 - \Gamma$ that is labeled by the edge with o as its largest angle vertex. It also changes the embedding of $\mathbb{R}_o \times \Delta_o$ into $\text{Maps}(\hat{A}_o; \mathbb{R})$.

Now, the change of the distinguished missing point can be undone if one first changes the original choice for the distinguished element in \hat{A}_o . As explained in Part 3, such a change modifies x to some $z_\iota \cdot x$ where $z_\iota \in \text{Maps}(\hat{A}_*, \mathbb{Z})$. Note that by virtue of what was said in the preceding paragraph, $z_\iota(\hat{u}) = 0$ in the case that the integer pair from \hat{u} defines an angle via (1–8) that is less than θ_o .

Meanwhile, the new embedding of $\mathbb{R}_o \times \Delta_o$ into $\text{Maps}(\hat{A}_o; \mathbb{R})$ only affects the value of $z_\iota \cdot x$ on elements u whose integer pair defines θ_o via (1–8). To see how x' differs from $z_\iota \cdot x$ on the latter set, note first that the new embedding of $\mathbb{R}_o \times \Delta_o$ to $\text{Maps}(\hat{A}_o; \mathbb{R})$ is obtained from the old by composing the latter with the action of the permutation ι . This then means that x' is obtained from $z_\iota \cdot x$ by the action of $\iota \in \text{Aut} \hat{A}$.

Part 5 A change in the choice for the correspondence of T in (C_0, ϕ) can be rectified by changing the choices for the first and second points in (3–16). This understood, then the image of C is insensitive to the choice of such a correspondence.

3.E The local structure of the map

The purpose here is to investigate the local structure around any given subvariety in \mathcal{M} of the map just defined to $\mathbb{R} \times O^{\hat{A}} / \text{Aut} \hat{A}$. However, there is one important point to establish first, this is summarized by:

Proposition 3.3 *The subspace $\mathcal{M} \subset \mathcal{M}_{\hat{A}}$ is a smooth submanifold whose dimension is $N_+ + N_- + 2$.*

Proof of Proposition 3.3 This is a special case in [15, Proposition 2.12]. □

The next proposition summarizes some of the salient local features of the map. In this regard, one should keep in mind that each $C \in \mathcal{M}$ has an open neighborhood with the following property: The graph $T_{(\cdot)}$ of any subvariety from the neighborhood is canonically isomorphic to T_C . Indeed, the ambiguity with the choice of an isomorphism between T_C and some $T_{C'}$ arises when there is no canonical pairing between the respective sets of ends in C and in C' that define identical elements in \hat{A} . However, if C' and C are close in \mathcal{M} , then each end of C' is very close at sufficiently large $|s|$ to a unique end of C , and vice versa. This geometric fact provides the canonical pairing of ends and thus the canonical identification between T_C and $T_{C'}$.

Having made this last point, it then follows that the map from \mathcal{M} to $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ has a canonical lift to $O^{\hat{A}}$ on a neighborhood of any given subvariety. The following proposition describes the nature of this lift.

Proposition 3.4 *Every subvariety in \mathcal{M} has an open neighborhood on which the map to $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ lifts as a smooth embedding onto an open set in $O^{\hat{A}}$.*

This proposition is also proved momentarily.

What follows is another way to view Proposition 3.4. To set things up, remark that all graphs in any given component of \mathcal{M} have isomorphic versions of $T_{(\cdot)}$. This said, fix a component and fix a graph, T , that is in the corresponding isomorphism class. Let $\mathcal{M}_{\hat{A},T}$ denote the corresponding component, and let $\mathcal{M}_{\hat{A},T}^{\Lambda}$ denote the set of pairs (C, T_C) where $C \in \mathcal{M}_{\hat{A},T}$ and T_C signifies a chosen correspondence of T in (C_0, ϕ) . The tautological projection from $\mathcal{M}_{\hat{A},T}$ to $\mathcal{M}_{\hat{A},T}^{\Lambda}$ defines the former as a covering space and principal $\text{Aut}(T)$ bundle over $\mathcal{M}_{\hat{A},T}^{\Lambda}$. Given a bivalent vertex $o \in T$, fix an admissible identification between the vertices of T 's version of $\underline{\Gamma}_o$ and the set \hat{A}_o . To elaborate, the identification is admissible if the second component of the assigned 4-tuple gives the sign of the integer label of the vertex, and if the greatest common divisor of the integer pair component of the 4-tuple is the absolute value of the integer label.

With these identification chosen, the map from $\mathcal{M}_{\hat{A},T}$ to $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ lifts as a map from $\mathcal{M}_{\hat{A},T}^{\Lambda}$ to $O^{\hat{A}}$ and Proposition 3.4 asserts that the lifted map is a local diffeomorphism. Note that $\mathcal{M}_{\hat{A},T}^{\Lambda}$ can be viewed as the moduli space of pairs that consist of a subvarieties with a distinct labeling of its ends.

As indicated by the discussion from Subsection 1.B in the paragraphs after Theorem 1.2, the various affine parameters that enter in the definition of $O^{\hat{A}}$ have direct geometric interpretations. To elaborate, suppose that T , $\mathcal{M}_{\hat{A},T}$ and $\mathcal{M}_{\hat{A},T}^{\Lambda}$ are as just defined. Let $o \in T$ denote a bivalent vertex and let e denote the edge of T that contains o as its largest angle vertex. Introduce (\hat{p}_o, \hat{p}'_o) to denote the relatively prime pair of integers that defines θ_o via (1–8). Now define a map from $\text{Maps}(\hat{A}_o; \mathbb{R})$ to $\text{Maps}(\hat{A}_o; \mathbb{R}/(2\pi\mathbb{Z}))$ as follows: First multiply any given $x \in \text{Maps}(\hat{A}; \mathbb{R})$ by $(q_{e'}\hat{p}_o - q_e\hat{p}'_o)$ and then take the $\text{mod } (2\pi\mathbb{Z})$ reduction of the result. Using the map for the relevant bivalent vertex, evaluation on any given $u \in \hat{A}_*$ defines a map from $\mathbb{R}^{\hat{A}}$ to $\mathbb{R}/(2\pi\mathbb{Z})$ that descends to give a map

$$(3-22) \quad \Psi_u: O^{\hat{A}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}.$$

The composition of such a map with the map from $\mathcal{M}_{\hat{A},T}^\Lambda$ has following geometric interpretation:

Proposition 3.5 *Let $(C, T_C) \in \mathcal{M}_{\hat{A},T}^\Lambda$, let o denote a bivalent vertex in T , let $u \in \hat{A}_o$, and let $E \subset C$ denote the end that corresponds via T_C to u . Then the restriction to the end E of $\hat{p}_o\varphi - \hat{p}'_o t$ has a unique $|s| \rightarrow \infty$ limit in $\mathbb{R}/(2\pi\mathbb{Z})$ and the latter is obtained by composing the map Ψ_u with the map from $\mathcal{M}_{\hat{A},T}^\Lambda$ to $O^{\hat{A}}$.*

The parameter on the line \mathbb{R}_- also has geometric interpretation. To describe the latter, first introduce m here to denote the greatest common divisor of the integer pair that is assigned to the edge in T with T 's smallest angle vertex. Then, the map from \mathbb{R}_- to $\mathbb{R}/(2\pi\mathbb{Z})$ that takes the $\text{mod } (2\pi\mathbb{Z})$ reduction of $\frac{1}{m}\tau_-$ descends to $O^{\hat{A}}/\text{Aut}^{\hat{A}}$. Let Ψ_- denote the latter map.

Proposition 3.6 *The composition of Ψ_- with the map from \mathcal{M} to $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ is the following:*

- *When the first point in (3–2) holds, let o denote the minimal angle vertex in T . Let $C \in \mathcal{M}$ and let $E \subset C$ denote the end where the $|s| \rightarrow \infty$ limit of θ is θ_o . Then this composition maps C to the $|s| \rightarrow \infty$ limit on E of $\hat{p}_o\varphi - \hat{p}'_o t$.*
- *When the second point in (3–2) holds, then this composition assigns to C the t -coordinate of its intersection point with the $\theta = 0$ locus.*
- *When the third point in (3–2) holds, write the $(1, \dots)$ element in \hat{A} as $(1, \varepsilon, m(\hat{p}_o, \hat{p}'_o))$. Let $E \subset C$ denote the end of C where the $|s| \rightarrow \infty$ limit of θ is 0. This composition then assigns to C the $|s| \rightarrow \infty$ limit of $-\varepsilon(\hat{p}_o\varphi - \hat{p}'_o t)$.*

Formal proofs of these last two propositions are omitted as both follow directly from (1–8) and (2–5) using the given definitions in Subsection 3.C above. The remarks that follow are meant to indicate what is going on. Consider first the assertion in Proposition 3.5. The point here is that if e is the edge that ends at o , and if (σ, v) parametrize K_e via (1–8), then (2–5) finds

$$(3-23) \quad \hat{p}_o\varphi - \hat{p}'_o t = (\hat{p}_o q' - \hat{p}'_o q)v + \hat{p}_o \sqrt{6} \cos(\sigma) - \hat{p}'_o (1 - 3 \cos^2 \sigma)w.$$

According to (1–8), the coefficient in front of w vanishes when σ is the angle that is assigned to the vertex o . As a consequence, the $|s| \rightarrow \infty$ limit of $\hat{p}_o\varphi - \hat{p}'_o t$ on the end that is associated to o is equal to $(\hat{p}_o q' - \hat{p}'_o q)v_o$ (modulo $2\pi\mathbb{Z}$), where v_o is the coordinate of the missing point on the boundary of the parametrizing cylinder that corresponds to the end in question.

The conclusions of Proposition 3.6 are derived using similar considerations. For example, in the first case of the proposition, the left most term in (3–23) is zero on the end in question and the right most term is $\alpha_Q(\sigma)w$. Thus, the claim follows directly using (3–15).

Proof of Proposition 3.4 The result follows using the previous two propositions in conjunction with [15, Proposition 2.13]. To elaborate, note that the differential structure on $\mathcal{M}_{\hat{A},T}^{\hat{A}}$ is that given by Theorem 1.1. [15, Proposition 2.13] guarantees that the angles from Propositions 3.5 and 3.6 provide local coordinates. \square

4 Proving diffeomorphisms

The previous section introduced the submanifold $\mathcal{M} \subset \mathcal{M}_{\hat{A}}$ of subvarieties with a corresponding $T_{(\cdot)}$ graph that is linear, and it described a map from \mathcal{M} to $\mathbb{R} \times \mathcal{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. This map is denoted in what follows by \mathfrak{B} . This section proves that the map \mathfrak{B} is a diffeomorphism onto $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The proof of this assertion verifies Theorem 3.1. Meanwhile, the proof introduces various techniques that are used in the subsequent sections to analyze the whole of $\mathcal{M}_{\hat{A}}^*$ and the other multiply punctured sphere moduli spaces.

The proof that \mathfrak{B} is a diffeomorphism has three parts. The first part establishes that the map is 1–1 onto its image. The second proves that the image is in $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The third proves that the map is proper onto $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$. Granted that the map is 1–1 and that its image is $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$, Proposition 3.4 establishes that the map is a local diffeomorphism. Granted that P is a proper map onto $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$, it diffeomorphically identifies \mathcal{M} with $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$.

The first subsection below introduces the machinery that is used to prove both that \mathfrak{B} is 1–1 and that its image is in $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The former conclusion is established in the subsequent subsection. Subsection 4.C constitutes a digression to describe $\mathcal{O}^{\hat{A}} - \hat{\mathcal{O}}^{\hat{A}}$, and then Subsection 4.D contains the proof that P maps \mathcal{M} onto $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The final subsection explains why \mathfrak{B} is a proper as a map onto $\mathbb{R} \times \hat{\mathcal{O}}^{\hat{A}} / \text{Aut}^{\hat{A}}$.

4.A Graphs and subvarieties

One can ask of any two elements from some version of $\mathcal{M}_{\hat{A}}^*$ whether one is obtained from the other by a constant translation along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. This section provides a sufficient condition for such to be the case. This condition is summarized

by Lemma 4.1, below. The next subsection considers the consequences of Lemma 4.1 in the case that both subvarieties come from Theorem 3.1's moduli space \mathcal{M} .

To set the stage for Lemma 4.1, suppose that (C_0, ϕ) and (C_0', ϕ') define points in some version of $\mathcal{M}^*_{\hat{A}}$. Then the same version of T must have correspondences in both (C_0, ϕ) and in (C_0', ϕ') . Choose respective correspondences and denote them as T_C and $T_{C'}$.

For each edge $e \subset T$, the parametrizing cylinder for the e -labeled component in $C_0 - \Gamma$ is the same as that used for $C_0' - \Gamma'$. This understood, say that respective parametrizations of the C_0 and C_0' versions of K_e are compatible when the following two conditions are met:

- (4-1)
- *When o is a multivalent vertex on e , then the respective extensions of the two parametrizations to the $\sigma = \theta_o$ circle in the parametrizing cylinder have identical sets of missing points, and also identical sets of singular points. Moreover, these identifications define identical versions of the circular graph ℓ_{oe} in the sense that a given point has the same integer label whether viewed in the C_0 or in the C_0' version.*
 - *This identification between the two versions of ℓ_{oe} is the same as that given by the chosen correspondences T_C and $T_{C'}$.*

Assume now that the respective parametrizations of the C_0 and C_0' versions of each $K_{(\cdot)}$ are compatible.

To continue, suppose again that e is an edge in T , and define functions (\hat{a}_e, \hat{w}_e) on any given C_0 version of K_e as follows: Let (a_e, w_e) denote the versions of the functions that parametrize C_0 's version of K_e via (2-5), and let (a_e', w_e') denote those that appear in the C_0' version. Both versions are pairs of functions on the parametrizing cylinder. Define $\hat{a}_e = a_e - a_e'$ and $\hat{w}_e = w_e - w_e'$.

Now make the following assumption:

- (4-2) *There exists a continuous function on the complement in C_0 of the critical point set of $\cos \theta$ whose pull-back from $C_0 - \Gamma$ via the various parametrizing maps gives the collection $\{\hat{w}_e\}$.*

Let \hat{w} denote this continuous function. As is explained below, such a function exists if and only if the identification between the $C_0 - \Gamma$ and $C_0' - \Gamma'$ via the compatible parametrizing maps extends to define a homeomorphism between C_0 and C_0' that identifies respective $\cos(\theta)$ critical point sets and is differentiable in their complements.

To finish the stage setting, introduce G to denote the portion of the $\hat{w} = 0$ locus that lies in the complement of the set of critical points of $\cos \theta$ on C_0 . The guarantee that ϕ' and ϕ are translates of each other involves G :

Lemma 4.1 *If $G \neq \emptyset$, then ϕ' is obtained from ϕ by composing the latter with a constant translation along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$.*

The proof of this lemma has three fundamental inputs. The first concerns the edge labeled set $\{\hat{a}_e \equiv a_e - a_{e'}\}$:

Lemma 4.2 *There is a continuous function on C_0 that pulls back from $C_0 - \Gamma$ via the parametrizing maps as the collection $\{\hat{a}_e\}$. Moreover, this function is smooth away from the critical points of the pull-back of θ .*

Let \hat{a} denote the function from Lemma 4.2.

The second input concerns the nature of G :

Lemma 4.3 *If G is neither empty nor all of the complement of the $\cos \theta$ critical point set, then it has the structure of an embedded, real analytic graph. In addition:*

- *Each edge is an embedded arc whose interior is oriented by the pull-back of $d\hat{a}$.*
- *The number of incident edges to any given vertex is even and at least four. Moreover, any circle about a vertex with sufficiently small and generic radius misses all vertices of G , and intersects the edges transversely. Furthermore, a circumnavigation of such a circle alternately meets inward pointing and outward pointing incident edges.*
- *Any sufficiently small and generic radius circle about a critical point of $\cos \theta$ misses all vertices of G and intersects the edges transversely. Furthermore, a circumnavigation of such a circle alternately meets inward pointing and outward pointing incident edges.*
- *If R is sufficiently large and generic, then the $|s| = R$ locus misses the vertices G and intersects the edges transversely. This locus intersects an even number of edges and a circumnavigation about any given component of the $|s| = R$ locus alternately meets edges where $|s|$ is respectively increasing and decreasing in the oriented direction.*
- *Let E denote an end of C_0 . Then, \hat{a} is bounded on $G \cap E$ and it has a unique $|s| \rightarrow \infty$ limit on $G \cap E$.*

The third input, and part of Lemma 4.2's proof, is a certain Cauchy–Riemann equation that is obeyed any given pair (\hat{a}_e, \hat{w}_e) :

$$(4-3) \quad \begin{aligned} \alpha_{Q_e} \hat{a}_{e\sigma} - \sqrt{6} \sin \sigma (1 + 3 \cos^2 \sigma) (w_e \hat{a}_{ev} + \hat{w}_e a_{e'v}) &= -\frac{1 + 3 \cos^4 \sigma}{\sin \sigma} \hat{w}_{ev} \\ (\alpha_{Q_e} \hat{w}_e)_\sigma - \sqrt{6} \sin \sigma (1 + 3 \cos^2 \sigma) (w_e \hat{w}_{ev} + \hat{w}_e w_{e'v}) &= \frac{1}{\sin \sigma} \hat{a}_{ev}, \end{aligned}$$

Indeed, this last equation appears when the $(a_{e'}, w_{e'})$ version of (2–6) is subtracted from the (a_e, w_e) version.

If Lemmas 4.2 and 4.3 are taken on faith for the moment, here is the proof of Lemma 4.1.

Proof of Lemma 4.1 If $G \neq \emptyset$, then there are two possibilities to consider, that where the closure of G is all of C_0 , and that where it is not. To analyze the first, appeal to (4–3) to conclude that each \hat{a}_e is constant when the closure of G is C_0 . According to Lemma 4.3, all of these constants are the same; ϕ' is obtained from ϕ by composing with a constant translation along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times g^2)$.

As is argued next, the case that G 's closure is not C_0 can not occur unless G is empty. Here is why: Since \hat{a} is bounded on G , it has some supremum. By virtue of the first point in Lemma 4.3, this supremum is not achieved in the interior of any edge. By virtue of the second point, it is not achieved at any vertex. Meanwhile, the third point implies that it is not achieved on the closure of G . The fourth and fifth points in Lemma 4.3 imply that the supremum is not an $|s| \rightarrow \infty$ limit of \hat{a} on G . This exhausts all possibilities and so $G = \emptyset$. □

To tie up the first of the loose ends, here is the proof of Lemma 4.2.

Proof of Lemma 4.2 Let $e \in T$ denote a given edge. A diffeomorphism is defined from C_0 's version of K_e to the C_0' version by pairing respective points that have the same inverse image in the parametrizing cylinder. Let ψ_e denote the latter map.

As is explained momentarily, there is a homeomorphism between C_0 and C_0' that is smooth away from the θ critical set and whose restriction to any given K_e is the corresponding ψ_e . This homeomorphism is denoted by ψ . Then $\hat{a} = \phi * s - \psi * \phi' * s$, and so \hat{a} is continuous and smooth away from the θ -critical set.

With the preceding understood, consider now the asserted existence of ψ . To start the analysis, focus attention on a multivalent vertex, $o \in T$. At issue here is whether the relevant versions of $\psi_{(\cdot)}$ fit together across the locus Γ_o in C_0 . For this purpose, let e denote one of o 's incident edges and let γ denote an arc in ℓ_{oe} with the latter viewed

as the $\sigma = \theta_o$ circle in the parametrizing cylinder for K_e . By virtue of the first point in (4–1), any point in the interior of γ simultaneously parametrizes a unique point in C_0 's version of Γ_o and also one in the C_0' version. This fact gives the map ψ_e a canonical extension as a homeomorphism from a neighborhood of the $\theta = \theta_o$ boundary in C_0 's version of the closure of K_e to a neighborhood of the $\theta = \theta_o$ boundary in the closure of the C_0' version.

To continue, let e' denote γ 's other edge label. Just as with ψ_e , the extension of $\psi_{e'}$ to γ defines a homeomorphism from its interior to that of an arc in the C_0' version of Γ_o . Together, (2–14) and the second point in (4–1) imply that these two homeomorphisms agree. As this conclusion holds for any arc of any version of Γ_o , it thus follows from (4–1) that the collection $\{\psi_e\}$ is the restriction of a homeomorphism from C_0 to C_0' .

Having established that ψ is at least continuous across Γ in C_0 , the next step is to verify that it is smooth across this locus. For this purpose, return to the arc γ in Γ_o , and let e and e' again denote its edge labels. Fix respective lifts, \hat{v}_e and $\hat{v}_{e'}$, for the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate functions on the e and e' versions of the parametrizing cylinder. Having done so, an integer pair arises in C_0 's version of (2–14) and (2–15), and a similar pair arises in the C_0' version. Let $L = (\ell, \ell')$ denote the difference between these two integer pairs. The assumptions in (4–2) force $\alpha_L(\theta_o)$ to vanish. This understood, then the assumption in (4–3) forces $(q_e\ell' - q_{e'}\ell)$ to vanish as well. These two vanishing conditions require L to vanish.

Granted that $L = 0$, fix a small radius disk about any given point in the interior of γ 's image in the version of Γ_o from C_0 . Let D denote this disk. If the radius of D is very small, then the parametrizing maps for both the C_0 versions of K_e and $K_{e'}$ extend to parametrize D . There is an analogous D' in C_0' . If the radius of D is small, both ψ_e and $\psi_{e'}$ extend as diffeomorphisms from D onto an open set in D' . Since ψ is continuous, these extensions agree on the $\theta = \theta_o$ locus in D . However, since $L = 0$, they agree on the whole of D . Thus, ψ is smooth on D . \square

Proof of Lemma 4.3 The analysis of G has four parts. The first studies G in $C_0 - \Gamma$, the second examines how G intersects Γ . Together, these prove that G has the structure of a graph. The third part establishes the third and fourth points of the lemma. The final part establishes the final point about \hat{a} .

Part 1 Fix an edge, $e \subset T$, so as to consider the part of the $\hat{w} = 0$ locus that lies in K_e . In this regard, assume that this locus is not the whole of K_e . It is a consequence of (4–3) that the pair (\hat{a}_e, \hat{w}_e) are real analytic on the interior of the parametrizing cylinder for K_e . Thus, the zero locus of \hat{w} in K_e is a 1–dimensional, real analytic variety. In particular, this gives it the structure of a graph whose vertices are the points

where both $d\hat{w}$ and \hat{w} are zero. An edge of G is the closure of a component of the complement in the $\hat{w} = 0$ locus of the set where $d\hat{w} = 0$. As a consequence of (4–3), the 1–form $d\hat{a}_e$ is zero at any given point on the $\hat{w} = 0$ locus if and only if $d\hat{w}_e$ is also zero there. It also follows from (4–3) that $d\hat{a}_e$ is not a multiple of $d\hat{w}_e$ at any point on the $\hat{w}_e = 0$ locus where both are non-zero. This implies that $d\hat{a}_e$ pulls-back to orient each edge of G . Moreover, the latter orientation is consistent with the one that comes by viewing the interior of any given edge as a portion of the boundary of the region where $\hat{w}_e \leq 0$.

To continue, remark that (4–3) also implies that G has no isolated points in K_e . To see why, first introduce the function

$$(4-4) \quad \sigma \rightarrow k(\sigma) = \left(\frac{1}{(1 + 3 \cos^4 \sigma) \alpha_{Q_e}(\sigma)} \right)^{1/2}.$$

Next, set $\eta \equiv \hat{a}_e - ik\hat{w}_e$. Now fix any point in the parametrizing cylinder, and standard arguments find a complex coordinate, z , centered on the point and a disk about that point that makes (4–3) equivalent to a single complex equation that has the form

$$(4-5) \quad \bar{\partial}\eta + u\hat{w}_e = 0.$$

Here, u is a smooth function on the disk. As η is real analytic, so it has a non-trivial Taylor’s expansion about any given point; and (4–5) implies that this expansion about any given $\hat{w}_e = 0$ point has the form

$$(4-6) \quad \eta = r + cz^p + o(|z|^{p+1}),$$

where r is real and p is a positive integer. This last observation precludes local extrema for \hat{w}_e where \hat{w}_e is zero. Therefore, G has no isolated points in K_e .

Equation (4–6) also implies that each vertex of G in K_e has at least four, and an even number of incident edges.

In any event, the second point in Lemma 4.3 for vertices in K_e follows using transversality theory and the fact that the $d\hat{a}$ orientation on any given edge is consistent with its orientation as part of the boundary of the region where $\hat{w} \leq 0$.

Part 2 Let $o \in T$ be a bivalent vertex, and let $\gamma \subset \Gamma_o$ be an arc. This part of the proof describes G ’s behavior in a neighborhood of γ . To start, focus attention on a small radius disk about a point in the interior of γ whose closure lies some positive distance from the vertices of γ . Let e and e' denote the edges that label γ . Thus, γ is in the closure of both K_e and $K_{e'}$. Because the pairs (a_e, w_e) and $(a_{e'}, w_{e'})$ extend over the $\sigma = \theta_o$ boundary of K_e ’s parametrizing domain, so the functions (\hat{a}_e, \hat{w}_e)

also extend. By taking the chosen disk to have small radius, the disk can be assumed to lie in the image of this extension of the parametrizing domain. Likewise, the pair $(\hat{a}_{e'}, \hat{w}_{e'})$ can be assumed to extend from the parametrizing cylinder of $K_{e'}$ to a subset in $[0, \pi] \times \mathbb{R}/(2\pi\mathbb{Z})$ that maps to the chosen disk via an extension of the e' version of (2–5).

A comparison of the respective extensions of (\hat{a}_e, \hat{w}_e) and $(\hat{a}_{e'}, \hat{w}_{e'})$ can be made using the corresponding C_0 and C_0' versions of (2–14) and (2–15). For this purpose, fix respective lifts, \hat{v}_e and $\hat{v}_{e'}$ chosen for the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate functions on the e and e' versions of the parametrizing cylinder. As remarked in the proof of Lemma 4.2, this can be done so that respective integer pairs that appear in the C_0 and C_0' versions of (2–14) and (2–15) are equal. Granted this, (2–15) implies that the coordinate transformation that is defined by (2–14) on the given disk identifies (\hat{a}_e, \hat{w}_e) with $(\hat{a}_{e'}, \hat{w}_{e'})$. As a consequence, the use on this disk of the extension of K_e 's parametrizing cylinder coordinates identifies $\hat{w}^{-1}(0)$ with the zero locus of \hat{w}_e .

With this last identification understood, then G 's intersection with the chosen disk must be a graph that has the structure described in Part 1 above for $G \cap K_e$. Moreover, any give edge of this intersection is either contained in the $\sigma = \theta_o$ locus, or it intersects this locus in a finite set. By the way, if G 's intersection with the chosen disk has an edge in the $\sigma = \theta_o$ locus, then it contains the whole of γ since this locus and G are real analytic.

Part 3 The transversality assertions in the third and fourth points follow from Sard's theorem. The assertion about the alternating orientation follows from the fact that the $d\hat{a}$ orientation on any given edge is consistent with its orientation as a part of the boundary of the region where $\hat{w} \leq 0$.

Part 4 The proof of the fifth point of the lemma requires some understanding of the behavior of \hat{a} on the ends of C_0 . For this purpose, suppose first that $E \subset C_0$ is an end with a $(\pm 1, \dots)$ label from \hat{A} . The identification $T_C = T_{C'}$ pairs E with a C_0' end, E' , that has the same label. This understood, a comparison between the E and E' versions of (1–9) proves that \hat{a}_e is bounded on E and has a unique limit on E as $\theta \rightarrow 0$.

Next, suppose that $E \subset C_0$ is an end with a $(0, -, \dots)$ label from \hat{A} where the constant n_E in (2–4) is zero. Let E' denote the corresponding end in C_0' as defined by the chosen identification of T_C and $T_{C'}$. Let b and b' denote the respective constants that appear in the E and E' versions of (2–14). It then follows from (2–14) that

$$(4-7) \quad \hat{a} = \frac{1}{r} \ln \left(\frac{b}{b'} \right) + o(1),$$

where r is the constant that appears in (2–4). Here, the term designated as $o(1)$ limits to zero as σ limits to the angle θ_E , thus as $|s| \rightarrow \infty$ on E .

Finally, suppose that $E \subset C_0$ is an end with a $(0, \dots)$ label from \hat{A} where the integer n_E in (2–4) is non-zero. It then follows from (2–4) that E corresponds to some multivalent vertex, $o \in T$. Now, equation (2–4) is one of a pair of equations that describe E at large $|s|$. To write this second equation, note that $(1 - 3 \cos^2 \theta_o)d\varphi - \sqrt{6} \cos \theta_o dt$ is exact on E . This understood, in [14, Equation (2.13)] with the analysis from of [14, Section 2] can be used to prove that any anti-derivative of this 1–form appears as follows: Let f denote the antiderivative, and let (ρ, τ) denote the coordinates that appear in (2–4). Then

$$(4-8) \quad f(\rho, \tau) = f_E + e^{-r\rho}(\kappa b \sin(n_E(\tau + \sigma)) + \hat{\delta}),$$

where n_E , r and b are the same constants that appear in E 's version (2–4). Meanwhile, κ is a constant that depends only n and on E 's label in \hat{A} . Finally, f_E is a constant and $\hat{\delta}$ is a term that limits to zero as $\rho \rightarrow \infty$.

To exploit (2–4) and (4–8) for the present purposes, keep in mind that the coordinate ρ is a constant, positive multiple of the function s . In this regard, it is important to realize that this multiplier depends only on E 's label in \hat{A} .

To see how (2–4) and (4–8) describe the behavior of \hat{a} on $G \cap E$, suppose that e is an edge of T that is incident to o , and suppose that the closure of K_e has non-compact intersection with E . This being the case, the end E then corresponds to a vertex on ℓ_{oe} and thus a missing point on the $\sigma = \theta_o$ circle in C_0 's version of the parametrizing cylinder for K_e . Fix an \mathbb{R} -valued lift, \hat{v} , of the coordinate v of this parametrizing cylinder that is defined in a contractible neighborhood of E 's missing point. To save on notation, assume that \hat{v} has value zero at this missing point. It then follows from (2–4) that

$$(4-9) \quad f \equiv \alpha_{Q_e}(\theta_o)\hat{v} + \sqrt{6} \left[(1 - 3 \cos^2 \theta_o) \cos \theta - \cos \theta_o(1 - 3 \cos^2 \theta) \right] w_e.$$

is a pull-back of an antiderivative of $(1 - 3 \cos^2 \theta_o)dt - \sqrt{6} \cos \theta_o d\varphi$ to a neighborhood of E 's missing point in the C_0 version of the parameter cylinder for K_e .

The respective identifications of T with both T_C and $T_{C'}$ pairs E with an end, $E' \subset C_0'$ that also corresponds to a vertex on $\underline{\Gamma}_o$. In this regard, the vertex is the same as the one given E , and this implies that the E' versions of (2–4) and (4–8) use the same integer n_E . Even so, the E' version can employ a different positive constant, b' .

There is also a primed version of (4–9), this denoted by f' . Note, in particular, that $f - f'$ is zero on the portion of $G \cap E$ in the closure of K_e . This being the case,

(2–4) and (4–8) imply that \hat{a} is bounded on this part of $G \cap E$, and that it has a unique $|s| \rightarrow \infty$ limit here which is given by (4–7). Since the pair b and b' are determined by E and E' without regard for the edge e , so the function \hat{a} has a unique $|s| \rightarrow \infty$ limit on $G \cap E$. \square

4.B Why the map from \mathcal{M} to $\mathbb{R} \times O^{\hat{A}} / \text{Aut}^{\hat{A}}$ is 1–1

The previous section defined a continuous map, P , from \mathcal{M} to $\mathbb{R} \times O^{\hat{A}} / \text{Aut}^{\hat{A}}$ and the purpose of this subsection is to explain why the latter map is 1–1. The explanation that follows is in three parts.

Part 1 To start, suppose that C and C' have the same image. As is explained in a moment, this requires that the same version of T have correspondences in both (C_0, ϕ) and in (C_0, ϕ') . Granted that such is the case, choose such correspondences. Such choices are implicit in Parts 2 and 3 that follow.

Let T_C denote a graph with a correspondence in (C_0, ϕ) and let $T_{C'}$ denote one with a correspondence in (C_0, ϕ') . To explain why T_C must be isomorphic to $T_{C'}$, remark first that the respective vertex sets enjoy a 1–1 correspondence that preserves angles, and that the latter correspondence induces one between the respective edge sets that preserves the integer pair labels. Thus, the issue here is whether the circular graphs that label the respective equal angle bivalent vertices in T_C and $T_{C'}$ are isomorphic.

To see that such is the case, let o denote a given bivalent vertex in T_C and let o' denote its equal angle partner in $T_{C'}$. Choose admissible identifications between $\underline{\Gamma}_o$ and \hat{A}_o , and likewise between $\underline{\Gamma}_{o'}$ and $\hat{A}_{o'}$. Here as before, an identification is deemed to be admissible when the second component of a 4–tuple gives the sign of the integer label of the corresponding vertex while the greatest common divisor of the integer pair component gives the absolute value of the integer label. These identifications induce respective cyclic orderings of \hat{A}_o . These cyclic orderings are relevant because $\underline{\Gamma}_o$ is isomorphic to $\underline{\Gamma}_{o'}$ as a labeled graph if and only if one cyclic ordering is obtained from the other by composing with a 1–1 self map of \hat{A}_o that only permutes identical 4–tuples.

To see that the cyclic orderings have the desired relation, take a lift of the common image point to $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. The image of \hat{A}_o by the $\mathbb{R}^{\hat{A}}$ component of this lift defines a set of n_o distinct points in $\mathbb{R}/(2\pi\mathbb{Z})$ and thus a cyclic ordering for \hat{A}_o . As dictated by the construction in Subsection 3.C, this cyclic ordering is obtained from those induced by the ordering of the vertices on the respective T_C and $T_{C'}$ versions of $\underline{\Gamma}_o$ by composing the latter with an appropriate permutation of \hat{A}_o that mixes only elements

with identical 4-tuples. Thus, the induced cyclic orderings on \widehat{A}_o from the T_C and $T_{C'}$ versions of $\underline{\Gamma}_o$ are related in the desired fashion.

Part 2 The respective images of C and C' in $O\widehat{A}/\text{Aut}\widehat{A}$ are defined by first assigning these subvarieties points in the space depicted in (3–14). As the images of C and C' agree in $O\widehat{A}/\text{Aut}\widehat{A}$, so their images agree in (3–12); and this means that the choices that are used to define the respective assignments in (3–14) can be made so that these assignments agree. This has the following consequences: First, the conditions in (4–1) are met at each multivalent vertex. To see why this is so, let $o \in T$ denote such a vertex and let e denote the incident edge with maximal angle θ_o . Since C and C' have the same image in Δ_o , the spacing of the missing points on the $\sigma = \theta_o$ circle in the parametrizing cylinder for K_e agree. Moreover, as the C and C' versions of the assigned point in \mathbb{R}_o agree, so the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinates of the respective distinguished missing points on the $\sigma = \theta_o$ circle must agree. This then implies that the respective C and C' versions of the missing point set on the $\sigma = \theta_o$ circle agree. Given the definitions of the respective T_C and $T_{C'}$ versions of Γ_o , this agreement of missing point sets induces the isomorphism between the two versions of ℓ_{oe} that comes from the given identification between T_C and $T_{C'}$ as the graph T .

The fact that C and C' define the same points in (3–14) has additional consequences. To describe the first, let e denote for the moment the edge of T that contains the minimal angle vertex of T . Since C and C' have the same image in \mathbb{R}_- , there are respective parametrizations of K_e that make their versions of (3–15) agree. Choose such parametrizations. Now, let o denote the vertex on e with the larger angle and, if o is bivalent, let e' denote the second of o 's incident edges. The point assigned both C and C' in the line \mathbb{R}_o endows $K_{e'}$ with its ‘canonical’ parameterization. This parameterization is such that the integer pair (n, n') that appears in (2–15) is zero. As such is the case for both the C and C' versions, so the functions \widehat{w}_e and $\widehat{w}_{e'}$ agree along the $\theta = \theta_o$ locus in C .

Part 3 Suppose that o is a bivalent vertex. Let e and e' denote the incident edges to o with e the edge where the maximum of θ is θ_o . Suppose that both the C and C' versions of K_e are given their ‘canonical’ parameterization as defined inductively using the following data: First, the already chosen parametrizations for the respective C and C' versions of the component of $C_0 - \Gamma$ whose labeling edge has the minimal angle vertex in T . Second, the assigned point in each $\mathbb{R}_{(\cdot)}$ factor from (3–14) whose label is a multivalent vertex with angle less than θ_o . The canonical parametrizations of the two versions of K_e and the assigned point in \mathbb{R}_o endows both the C and C' versions of $K_{e'}$ with a ‘canonical’ parameterization. The parametrizations of the C

and C' versions of K_e give the function \hat{w}_e on C 's version of K_e , while that of the two versions of $K_{e'}$ give $\hat{w}_{e'}$ on the C' version of $K_{e'}$. This understood, the argument used at the end of the previous paragraph finds that $\hat{w}_e = \hat{w}_{e'}$ along the $\theta = \theta_o$ locus in C .

Granted the preceding, it then follows that (4–2) holds when the components of the C and C' versions $C_0 - \Gamma$ are given the parametrizations just described. Since C and C' have the same image in the \mathbb{R} factor of $\mathbb{R} \times O_T / \text{Aut}(T)$, the desired identity $C = C'$ follows from Lemma 4.2 if the graph $G \subset C$ is non-empty. That such is the case follows from the identity between the assigned values of C and C' in (3–14)'s factor \mathbb{R}_- . Indeed, subtracting the C' version of (3–15) from the C version finds that \hat{w} has average 0 over any constant θ slice of the component of C 's version $C_0 - \Gamma$ whose labeling edge contains the minimal angle vertex in T . As such there must be a zero of \hat{w} on each such slice.

4.C The complement of $\hat{O}^{\hat{A}}$ in $O^{\hat{A}}$

Before proving that the map from \mathcal{M} to $O^{\hat{A}} / \text{Aut}^{\hat{A}}$ lands in $\hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$, it is worthwhile to describe the complement of $\hat{O}^{\hat{A}}$ in $O^{\hat{A}}$, this the subspace of points where $\text{Aut}^{\hat{A}}$ has non-trivial stabilizer.

To start, let o denote a multivalent vertex in $T^{\hat{A}}$ and let Cyc_o denote the set of cyclic orderings of \hat{A}_o . The components of $O^{\hat{A}}$ are in 1–1 correspondence with $\times_o \text{Cyc}_o$. When $v \in \times_o \text{Cyc}_o$, let $O^{\hat{A}}_v$ denote the corresponding component. If $x \in O^{\hat{A}}_v$ is fixed by some $\iota \in \text{Aut}^{\hat{A}}$, then ι must preserve the cyclic orderings that are determined by v .

To see the implications of this last observation, let o denote a bivalent vertex and let $\text{Aut}_{o,v}$ denote the group of permutations of \hat{A}_o that preserve the cyclic order defined by v while permuting only elements with identical 4–tuples. This is a cyclic group whose order is denoted in what follows by m_o . The group $\times_o \text{Aut}_{o,v}$ is the subgroup of $\text{Aut}^{\hat{A}}$ that preserves $O^{\hat{A}}_v$.

Set k_v to denote the greatest common divisor of the integers in the set that consists of m_- and the collection $\{m_o\}$. Thus, each $\text{Aut}_{o,v}$ has a unique $\mathbb{Z}/(k_v\mathbb{Z})$ subgroup. Note that in the case that $k_v > 1$, each such subgroup has a canonical generator. To explain, let ι_o denote the generator of $\text{Aut}_{o,v}$ that moves elements the minimal amount in the direction that increases the numerical order in the following sense: Fix a distinguished element in \hat{A}_o so as to turn v into a linear ordering with the distinguished element last. Then ι_o moves the distinguished element to the position numbered by n_o/m_o . The

canonical generator of the $\mathbb{Z}/(k_v\mathbb{Z})$ subgroup moves the distinguished element to the position numbered n_o/k_v .

Granted the preceding, a canonical $\mathbb{Z}/(k_v\mathbb{Z})$ subgroup of $\text{Aut}^{\hat{A}}$ is defined by the requirement that its generator project to any given $\text{Aut}_{o,v}$ so as to give the canonical generator of the latter's $\mathbb{Z}/(k_v\mathbb{Z})$ subgroup.

With this understood, consider:

Proposition 4.4 *The $\text{Aut}^{\hat{A}}$ stabilizer of any given point in $O^{\hat{A}}_v$ is a subgroup of the canonical $\mathbb{Z}/k_v\mathbb{Z}$ subgroup. Conversely, any subgroup of the latter has a non-empty set of fixed points in $O^{\hat{A}}_v$.*

This proposition is proved momentarily.

To view this fixed point set in a different light, fix a distinguished element in each $\hat{A}_{(\cdot)} \subset \hat{A}_*$ so as to use the description in (3–12) for $O^{\hat{A}}_v$. In particular, the space in (3–12) has the evident projection to $\times_o\Delta_o$ and this projection is equivariant with respect to the action of $\times_o\text{Aut}_{o,v}$ on $O^{\hat{A}}_v$ and an action on $\times_o\Delta_o$. In this regard, the action on $\times_o\Delta_o$ is the product of the action of each version of $\text{Aut}_{o,v}$ on the corresponding version of Δ_o that cyclically permutes the coordinates in the manner dictated by a given element in $\text{Aut}_{o,v}$. Here, the latter action is obtained by labeling the Euclidean coordinates of points in Δ_o by the elements in \hat{A}_o so that the j 'th coordinate corresponds to the j 'th element in \hat{A}_o using the linear ordering that gives v 's cyclic ordering and has the distinguished element last.

Granted all of this, consider:

Proposition 4.5 *Let G be a subgroup of the canonical $\mathbb{Z}/k_v\mathbb{Z}$ subgroup in $\text{Aut}^{\hat{A}}$. Then the set of points in (3–12) with $\text{Aut}^{\hat{A}}$ stabilizer G is the restriction of the fiber bundle projection from $O^{\hat{A}}$ to $\times_o\Delta_o$ over the set of points in $\times_o\Delta_o$ with stabilizer G .*

The remainder of this subsection contains the following proofs.

Proof of Propositions 4.4 and 4.5 The proof is given in six steps.

Step 1 Let $\hat{\text{Aut}}$ denote the semi direct product of the groups $\text{Aut}^{\hat{A}}$ and $(\mathbb{Z} \times \mathbb{Z}) \times \text{Maps}(\hat{A}_*; \mathbb{Z})$. This group acts on $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$. The stabilizers of points in $O^{\hat{A}}$ can be determined by studying the stabilizers in $\hat{\text{Aut}}$ of points in $\mathbb{R}_- \times \mathbb{R}^{\hat{A}}$ since the image in $O^{\hat{A}}$ of $(\tau_-, x) \in \mathbb{R}_- \times \mathbb{R}^{\hat{A}}$ is fixed by $g \in \text{Aut}^{\hat{A}}$ if and only if (τ_-, x) is

fixed by an element $g \in \widehat{\text{Aut}}$ that has the form of $(g, (N, z))$ with $N \in \mathbb{Z} \times \mathbb{Z}$ and $z \in \text{Maps}(\widehat{A}_*; \mathbb{Z})$. This understood, suppose now that the point (τ_-, x) projects to $O_{\widehat{A}_v}$ and that $g \in \times_o \text{Aut}_{o,v}$ fixes its image.

Step 2 To start the analysis of g , note that τ_- is fixed if and only if $N = r_- Q_e$ where e here denotes the edge in $T^{\widehat{A}}$ with the smallest angle vertex in $T^{\widehat{A}}$ and where r_- is a fraction with $m_- r_- \in \mathbb{Z}$.

Step 3 Let o now denote the vertex in $T^{\widehat{A}}$ with the second smallest angle. Assuming o is bivalent, g fixes both τ_- and the restriction of x to \widehat{A}_o if and only if

$$(4-10) \quad x(u) = x(g_o u) - 2\pi(z(u) + r_-)$$

for all $u \in \widehat{A}_o$. Here, g_o is the component of g in $\text{Aut}_{o,v}$. Let k_o denote the order of g_o . Using (4-10) some k_o times in succession finds that g fixes both τ_- and the restriction of x to \widehat{A}_o if and only if $m_- r_- \in \mathbb{Z}$ and

$$(4-11) \quad \sum_{0 \leq j < k_o} z(g_o^j u) + k_o r_- = 0$$

for all $u \in \widehat{A}_o$. Note that this last condition can be satisfied if and only if $k_o r_- \in \mathbb{Z}$. Note as well that the case $r_- \in \mathbb{Z}$ is precluded unless g_o is trivial since x is assumed to lie in $\mathbb{R}^{\widehat{A}}$.

Step 4 If $T^{\widehat{A}}$ has a second bivalent vertex, let \widehat{o} denote the one that shares an edge with o . Thus, \widehat{o} has the third smallest vertex angle. Then g fixes the restriction of x to $\widehat{A}_{\widehat{o}}$ if and only if

$$(4-12) \quad x(\widehat{u}) = x(g_{\widehat{o}} \widehat{u}) - 2\pi \left(z(\widehat{u}) + \sum_{u \in \widehat{A}_o} z(u) \varepsilon_u \frac{p_u' p_{\widehat{u}} - p_u p_{\widehat{u}}'}{q_{\widehat{e}}' p_{\widehat{u}} - q_{\widehat{e}} p_{\widehat{u}}'} + r_- \frac{q_{e'} p_{\widehat{u}} - q_e p_{\widehat{u}}'}{q_{\widehat{e}}' p_{\widehat{u}} - q_{\widehat{e}} p_{\widehat{u}}'} \right)$$

for all $\widehat{u} \in \widehat{A}_{\widehat{o}}$. Here, $g_{\widehat{o}}$ denotes g 's component in $\text{Aut}_{o,v}$ and \widehat{e} denotes the edge in $T^{\widehat{A}}$ that contains both o and \widehat{o} .

To make sense of this last condition, use (4-11) as applied to the various g_o orbits in \widehat{A}_o to identify the sum in (4-12) with

$$(4-13) \quad -r_- \frac{p_{o'} p_{\widehat{u}} - p_o p_{\widehat{u}}'}{q_{\widehat{e}}' p_{\widehat{u}} - q_{\widehat{e}} p_{\widehat{u}}'}$$

Then, use (3-3) to write P_o as $Q_e - Q_{\widehat{e}}$ and thus see that the equality in (4-12) is exactly the equation that results from (4-10) by replacing o with \widehat{o} and u with \widehat{u} .

This last point leads to the following conclusion: The element g fixes τ_- and the restriction of x to both \hat{A}_o and $\hat{A}_{\hat{o}}$ if and only if $m_-r_- \in \mathbb{Z}$ and both the (o, u) and (\hat{o}, \hat{u}) versions of (4–11) hold.

Step 5 With the help of (3–3), essentially the same analysis can be continued in an inductive fashion through the bivalent vertices with ever larger angle so as to draw the following conclusions:

- The element g fixes (τ_-, x) if and only if both $m_-r_- \in \mathbb{Z}$ and (4–10) holds for each bivalent vertex $o \in T^{\hat{A}}$.
- These conditions are satisfiable if and only if (4–11) holds for each bivalent vertex, and the latter can be satisfied if and only if $k_o r_- \in \mathbb{Z}$ for each bivalent vertex $o \in T^{\hat{A}}$.

Step 6 The ‘only if’ direction of the two conclusions from Step 5 have two relevant consequences, and here is the first: Because $x \in \mathbb{R}^{\hat{A}}$, the various versions of k_o must agree. Indeed, if not, let k denote the smallest and let o denote a vertex where $k_o = k$. Now, repeat the analysis with g replaced by g^k . The corresponding g^k versions of (4–11) has kr_- replacing r_- and o ’s version would require $kr_- \in \mathbb{Z}$. But then a version where $k_{(\cdot)} \neq k$ would need x to lie outside of $\mathbb{R}^{\hat{A}}$.

Here is the second consequence: Each k_o is a divisor of the corresponding m_o , this the order of $\text{Aut}_{o,v}$. As they are all equal to the same $k \in \mathbb{Z}$, so k divides each m_o and so k is a multiple of the greatest common divisor of the collection $\{m_o\}$. In addition as $kr_- \in \mathbb{Z}$, so k must also be a divisor of m_- . Thus, k is a multiple of k_v and g is an element in some subgroup of the canonical $\mathbb{Z}/k_v\mathbb{Z}$ subgroup of $\times_o \text{Aut}_{o,v}$.

Meanwhile, the ‘if’ directions of Step 5’s observations directly imply that any subgroup of the canonical $\mathbb{Z}/k_v\mathbb{Z}$ subgroup of $\times_o \text{Aut}_{o,v}$ has a non-empty set of fixed points, and that the fixed point set is described by Proposition 4.5. □

4.D Why the map from \mathcal{M} lands in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$

As indicated by the heading, the purpose of this subsection is to explain why the image of the map P from \mathcal{M} to $\mathbb{R} \times O^{\hat{A}} / \text{Aut}^{\hat{A}}$ lies in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. To this end, suppose that $C \in \mathcal{M}$. The image of C in $O^{\hat{A}} / \text{Aut}^{\hat{A}}$ has a lift to some component $O^{\hat{A}}_v \subset O^{\hat{A}}$, and suppose that this lift is fixed by some $g \in \times_o \text{Aut}_{o,v}$. This being the case, g is some multiple of the generator of the canonical $\mathbb{Z}/k_v\mathbb{Z}$ subgroup. Let $k \in \{0, \dots, k_v - 1\}$ denote this multiple. The proof that $k = 0$ has four parts.

Part 1 Fix a distinguished element in each version of \hat{A}_o so as to write $O^{\hat{A}_v}$ as in (3–12). Using this view of $O^{\hat{A}_v}$, lift C 's image as a point in (3–14) by choosing an admissible identification between the vertex set of each T_C version of $\underline{\Gamma}_o$ and the corresponding \hat{A}_o . Denote the resulting point in (3–14) as $(\tau_-, (\tau_{(\cdot)}, r_{(\cdot)}))$ where $\tau_- \in \mathbb{R}_-$, each τ_o is a point in the corresponding version of \mathbb{R}_o , and each r_o is a point in the corresponding version of Δ_o .

Let e denote for the moment the edge in $T^{\hat{A}}$ that contains the minimal angle vertex and choose a parameterization of K_e that makes (3–15) equal to τ_- . Use this parameterization with the data $(\tau_{(\cdot)})$ to inductively define canonical parameterizations for the remaining $C_0 - \Gamma$. This is done with the following induction step: Suppose that \hat{o} is a bivalent vertex and that the component of $C_0 - \Gamma$ whose label is the edge with largest angle $\theta_{\hat{o}}$ has its canonical parameterization. To obtain one for the component that has $\theta_{\hat{o}}$ as its smallest angle, first introduce the arc in $\underline{\Gamma}_{\hat{o}}$ starting at the vertex that corresponds to the distinguished element in $\hat{A}_{\hat{o}}$. This arc constitutes a one element concatenating path set. The latter with the lift $\tau_{\hat{o}}$ defines the canonical parameterization for the component of $C_0 - \Gamma$ whose label is the edge that has $\theta_{\hat{o}}$ as its smallest angle, this according to the rules laid out in Part 4 of Subsection 2.C.

The next step is to change each of these parameterization. To start, let e denote an edge of $T^{\hat{A}}$ and let ϕ_e denote the parametrizing map from the relevant parametrizing cylinder to K_e . The new parameterization of K_e is defined by composing ϕ_e with the diffeomorphism of $[0, \pi] \times \mathbb{R}/(2\pi\mathbb{Z})$ that pulls back the coordinate functions as

$$(4-14) \quad (\sigma, v) \rightarrow (\sigma, v - 2\pi k/k_v).$$

This new parameterization is of the sort that is described in Subsection 2.B because k_v evenly divides both integers from the pair associated to any edge of $T^{\hat{A}}$. This divisibility arises for the following reasons: First, k_v divides both integers from the edge with the smallest angle vertex. Second, it divides the order of $\text{Aut}_{o,v}$ for each bivalent vertex $o \in T^{\hat{A}}$ and the order of $\text{Aut}_{o,v}$ divides the corresponding version of P_o that appears in (3–3).

For reference below, note that if (a_e, w_e) denotes that original versions of the functions (a, w) that appear in (2–5), then the new versions, (a_e', w_e') are given by

$$(4-15) \quad a_e'(\sigma, v) = a_e(\sigma, v - 2\pi k/k_v) \quad \text{and} \quad w_e'(\sigma, v) = w_e(\sigma, v - 2\pi k/k_v).$$

Part 2 Agree to implicitly view g as an element in $\text{Aut}(T_C)$ via the chosen identification between any given T_C version of $\underline{\Gamma}_o$ and the corresponding \hat{A}_o . Granted this

view of g , let o now denote a bivalent vertex in $T^{\hat{A}}$. According to Proposition 4.5, g 's action on Δ_o must fix the point $r_o \in \Delta_o$ and this has the following consequence: Let v denote a given vertex in $\underline{\Gamma}_o$. Now sum the coordinates of r_o that correspond to the arcs in $\underline{\Gamma}_o$ that are met on the oriented path that starts at $g^{-1}v$ and ends at v . This sum is $2\pi k/k_v$.

This last observation implies that the diffeomorphism in (4-15) restricts to the $\sigma = \theta_o$ circle of the parametrizing cylinder so as to map missing points to missing points. Moreover, the labels granted these missing points as vertices in ℓ_{oe} are preserved by (4-14). Thus, both of the conditions in (4-1) hold for the new and original parametrizations using $C' = C$ and using g as the element in $\text{Aut}(T_C)$ for the isomorphism between T_C and itself.

Part 3 The condition in (4-2) also holds in this context. To see this, let o denote a bivalent vertex in $T^{\hat{A}}$, let e denote the edge that contains o as its maximal angle vertex, and let e' denote the edge that contains o as its minimal angle vertex. Let $\gamma \subset \underline{\Gamma}_o$ denote the arc that starts at the vertex that corresponds to the distinguished element in \hat{A}_o . This arc corresponds to respective arcs on the $\sigma = \theta_o$ boundary circle for the parametrizing domain of both K_e and $K_{e'}$. The $\mathbb{R}/(2\pi\mathbb{Z})$ coordinates on these arcs lift to \mathbb{R} so as to make the $N = 0$ version of (2-14) hold. These lifts identify the two boundary circle arcs. With this identification understood, the $N = 0$ version of (2-15) describes the relationship between w_e and $w_{e'}$ on the interior of γ when the latter is viewed as an arc in the $\theta = \theta_o$ locus in C 's model curve.

On this same arc, the pair w_e' and $w_{e'}$ are related by another version of (2-15); the version that applies to the arc that is mapped to γ by the inverse of the map in (4-14). However, using Lemma 2.3 and the fact that g acts on $\underline{\Gamma}_o$ as an isomorphism, the relationship between w_e' and $w_{e'}$ on γ 's image in C_0 is also given by the $N = 0$ version of (2-15). Indeed, Lemma 2.3 finds that

$$\begin{aligned}
 (4-16) \quad w_e(\sigma, \hat{v} - 2\pi k/k_v) &= w_{e'}(\sigma, \hat{v} - 2\pi k/k_v) \\
 &+ \frac{1}{\alpha_{Q_e}(\sigma)}(q_{e'}q_{e'} - q_e q_{e'}) (\hat{v} - 2\pi k/k_v) \\
 &+ \frac{2\pi}{\alpha_{Q_e}(\sigma)}(q_{e'}j - q_e j'),
 \end{aligned}$$

where (j, j') is proportional to the relatively prime integer pair that defines θ_o via (1-8). In particular, the proportionality factor here is minus the sum of the integer weights that are assigned to all vertices but v that lie on the oriented path in $\underline{\Gamma}_o$ from $g^{-1}(v)$ to v . Because g acts as an isometry, this sum must equal $(k/k_v)P_o$ with

P_o as in (3–3). Thus, (j, j') can be written as $(k/k_b)(Q_e - Q_{e'})$. With (j, j') so identified, the equality in (4–16) finds that $w_{e'}$ and $w_{e'}$ do indeed obey the predicted $N = 0$ version of (2–15) on γ .

As both $(w_e, w_{e'})$ and $(w_{e'}, w_{e'})$ obey the same $N = 0$ version of (2–15) on γ , it then follows that $\hat{w}_e = \hat{w}_{e'}$ along γ . The fact that this equality holds along the whole of Γ_o in C_0 can be seen with the help of (2–5). To see this, let γ' denote another arc in $\underline{\Gamma}_o$. The relationship between w_e and $w_{e'}$ on γ' can be obtained with the use of the e and e' versions of (2–5); it is given by the version of (2–15) where (j, j') is a certain multiple of the relatively prime pair that defines θ_o via (1–8). To be precise, this multiple is obtained by summing with an appropriate sign the weights of the vertices that lie on the oriented path in $\underline{\Gamma}_o$ that starts at the endpoint of γ and ends at the starting point of γ' . Since the same weights appear on the analogous path defined by $g^{-1}(\gamma)$ and $g^{-1}(\gamma')$, the relation between $w_{e'}$ and $w_{e'}$ on γ' is the same (j, j') version of (2–15). Thus, $\hat{w}_e = \hat{w}_{e'}$ along γ' .

Part 4 Let $G \subset C_0$ denote the part of the zero locus of \hat{w} that lies in the complement of the critical point set of the pull-back of $\cos(\theta)$. This set is described by Lemma 4.3. Moreover, $G \neq \emptyset$ since the w_e and $w_{e'}$ versions of the integral in (3–15) are identical in the case that e is any edge in $T^{\hat{A}}$.

To see the implications of Lemma 4.1 in this context, return momentarily to the proof of Lemma 4.2. The latter explains how the various versions of (4–14) fit together across $\Gamma \subset C_0$ so as to define a diffeomorphism, $\psi: C_0 \rightarrow C_0$. With ψ in hand, introduce the tautological map, $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$, onto C . Keep in mind that ϕ is almost everywhere 1–1. Let $\phi' \equiv \phi \circ \psi$. In the present context, Lemma 4.1 asserts that $\phi'(C_0)$ is obtained from C by a constant translation along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times S^2)$. However, this factor must be zero since the average of any given $a_{e'}$ around a constant σ circle is equal to that of a_e around the same circle. Thus, ϕ' also maps C_0 onto C . Granted this, then (4–14) implies that ψ is 1–1 if and only if k is zero.

4.E Why the map from \mathcal{M} is proper

The map P from $\mathcal{M} \rightarrow \mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ is proper if and only if all sequences in \mathcal{M} with convergent image in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ have convergent subsequences. The proof that such is the case is given here in three parts.

Part 1 Let $\{C_j\}_{j=1,2,\dots}$ denote a sequence with convergent image in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. The desired convergent subsequence in \mathcal{M} is found by first invoking [15, Proposition 3.7] to describe the $j \rightarrow \infty$ behavior of a subsequence from $\{C_j\}$ in terms of a

limiting data set, Ξ . Here, Ξ is a finite set of pairs where each has the form (S, n) with S being a pseudoholomorphic, multiply punctured sphere in $\mathbb{R} \times (S^1 \times S^2)$ and with n being a positive integer.

This first part of the subsection establishes that the conclusions of [15, Proposition 3.7] hold in the case that $K = \mathbb{R} \times (S^1 \times S^2)$. The precise version needed here is stated formally as follows.

Proposition 4.6 *Let $\{C_j\}_{j=1,2,\dots} \subset \mathcal{M}$ denote an infinite sequence with convergent image in $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$. There exists a subsequence, hence renumbered consecutively 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, pseudoholomorphic, multiply 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 .
- The following limit exists and is zero:

$$(4-17) \quad \lim_{j \rightarrow \infty} \left(\sup_{z \in C_j} \text{dist}(z, \cup_{\Xi} S) + \sup_{z \in \cup_{(S,n) \in \Xi} S} \text{dist}(C_j, z) \right)$$

The proof of this proposition appears momentarily. It is employed in the following manner to prove that the map P from \mathcal{M} to $\mathbb{R} \times \hat{O}^{\hat{A}} / \text{Aut}^{\hat{A}}$ is proper: Part 2 of this subsection uses Proposition 4.6 to prove that Ξ contains but a single element, this denoted as (S, n) . Part 3 of the subsection proves that $n = 1$ and that S is a subvariety from \mathcal{M} . Granted that such is the case, then (4-17) asserts that S is in fact the limit of $\{C_j\}$ with respect to the topology on \mathcal{M} as defined in (1-13). This last conclusion establishes that P is proper.

Proof of Proposition 4.6 Except for the second point in (4-17), this proposition restates conclusions from [15, Proposition 3.7]. The proof of the second point in (4-17), assumes it false so as to derive some nonsense. For this purpose, note that [15, Proposition 3.7] asserts that the second point in (4-17) holds if the supremums that appear are restricted to those points z that lie in any given compact subset of $\mathbb{R} \times (S^1 \times S^2)$.

The derivation of the required nonsense starts with the following lemma.

Lemma 4.7 *Assume that the second point in (4-17) does not hold for the given infinite sequence, $\{C_j\}$, of multiply punctured, pseudoholomorphic spheres. Then, there exists an \mathbb{R} -invariant, pseudoholomorphic cylinder, $S_* \subset \mathbb{R}$; an infinite subsequence of $\{C_j\}$ (hence renumbered consecutively from 1); and given $\varepsilon > 0$ but small, there is a real number s_0 , a sequence $\{s_{j-}\}_{j=1,2,\dots} \subset (-\infty, s_0]$ and a corresponding sequence $\{s_{j+}\}_{j=1,2,\dots} \subset [s_0, \infty)$; all with the following significance:*

- Both s_{j-} and s_{j+} are regular values of s on C_j .
- If either $\{s_{j-}\}$ and $\{s_{j+}\}$ is bounded, then it is convergent; but at least one of the two is unbounded.
- The $s \in [s_{j-}, s_{j+}]$ portion of C_j 's intersection with the radius ε tubular neighborhood of S_* has a connected component, C_{j*} , whose points have distance $\frac{1}{2}\varepsilon$ or less from S_* .
- Both the $s = s_{j-}$ and $s = s_{j+}$ slices of C_{j*} are non-empty and both contain points with distance $\frac{1}{2}\varepsilon$ from S_* .
- Let δ_j denote the maximum of the distances from the point on the $s = \frac{1}{2}(s_{j-} + s_{j+})$ locus in C_{j*} to S_* . The corresponding sequence $\{\delta_j\}$ is then decreasing with limit zero.

Lemma 4.7 is proved momentarily. To see where this lemma leads, suppose first that θ is neither 0 nor π on S_* . In this case, the conclusions in [15, Lemma 3.9] hold. But topological considerations find that [15, Lemma 3.9] and Lemma 4.7 lead to nonsense.

To elaborate, first let θ_* denote the constant value for θ on S_* and assume that θ_* is neither 0 nor π . Next, fix ε small and then j large. By virtue of what is said in [15, Lemma 3.9], there exists some very small and j -independent constants δ_{\pm} , either both positive and negative and with the following properties: First, neither $\theta_* + \delta_+$ or $\theta_* + \delta_-$ is an $|s| \rightarrow \infty$ limit of θ on C_j . Second, when j is large, $\theta_* + \delta_+$ is a value of θ on the $s = s_{j+}$ boundary of C_{j*} and $\theta_* + \delta_-$ is a value of θ on the $s = s_{j-}$ boundary. Let z_+ and z_- denote respective points on the $s = s_{j+}$ and $s = s_{j-}$ boundaries of C_{j*} where θ has the indicated value. Since both δ_- and δ_+ have small absolute value, the $\theta = \theta_* + \delta_-$ and $\theta = \theta_* + \delta_+$ loci are circles in the same component of C_j 's version of $C_0 - \Gamma$. Thus, there is a path in this component that lies in the set where $|\theta - \theta_*| \geq \min(|\delta_-|, |\delta_+|)$ and runs from z_- to z_+ . Denote the latter by γ . When j is large, the last point in Lemma 4.7 forbids an intersection between γ and the $s = \frac{1}{2}(s_{j-} + s_{j+})$ locus. On the other hand, there is a path in C_{j*} that runs from z_+ to z_- and does intersect this locus. The concatenation of the latter path with γ defines a closed loop in C_j that has non-zero intersection number with the constant s slices of C_{j*} . But no such loop exists since C_j has genus zero.

Suppose next that that $\theta = 0$ or π on S_* . In this case, the maximum value of θ on both the $s = s_{j-}$ and $s = s_{j+}$ slices of C_{j*} is bounded away from zero by a uniform multiple of ε^2 when ε is small. Meanwhile all values of θ on the $s = \frac{1}{2}(s_{j-} + s_{j+})$ slice of C_{j*} are bounded by $o(\delta_j^2)$. This being the case, the mountain pass lemma demands a non-extremal critical point of θ on C_j . Note in this regard that there are at most a finite

number of intersections between any two distinct, irreducible, pseudoholomorphic subvarieties.

Thus, in all cases, Lemma 4.7 leads to nonsense. Granted the lemma is correct, then its conclusions are false and so the second point in (4–17) must hold. \square

Proof of Lemma 4.7 Let Ξ denote the data set that is provided by [15, Proposition 3.7]. Since the second point in (4–17) is supposed to fail, there exists some subvariety, S , from Ξ ; an end $E \subset S$; a constant $R_0 \geq 1$; and, given $\varepsilon > 0$, a divergent sequence $\{R_j\} \subset [R_0, \infty)$; all with the following properties:

- (4–18) • Each sufficiently large j version of C_j intersects the $|s| \in [R_0, R_j]$ portion of the radius ε tubular neighborhood of E where the distance to E is no greater than $\frac{1}{2}\varepsilon$.
- Meanwhile, each such C_j has a point on the $|s| = R_j$ slice of this neighborhood where the distance to E is equal to $\frac{1}{2}\varepsilon$.

Of course, there may be more than one such S and more than one such end in S to which (4–18) applies. Lemma 4.7 holds if there exists a pair (S, E) as in (4–18) where S is not an \mathbb{R} invariant cylinder. Lemma 4.7 also holds if there exists a pair (S, E) as in (4–18) where S is an \mathbb{R} –invariant cylinder that intersects a subvariety from some other pair in Ξ . Indeed, such is the case because each C_j is irreducible. In fact, for this very reason, Lemma 4.7 holds if Ξ contains any pair whose subvariety is a not an \mathbb{R} –invariant cylinder.

To finish the argument for Lemma 4.7, consider now the case when all subvarieties from Ξ are \mathbb{R} –invariant cylinders. Even so, Lemma 4.7 must hold unless one of the following is true:

- (4–19) • Let S denote any cylinder from Ξ . Given $\varepsilon > 0$, there exists a divergent sequence $\{R_j\} \subset [0, \infty)$ such that each sufficiently large j version of C_j intersects the $s \in (-\infty, R_j]$ portion of the radius ε tubular neighborhood of S where the distance to S is no greater than $\frac{1}{2}\varepsilon$, and it intersects the $s = R_j$ slice at a point with distance to S equal to $\frac{1}{2}\varepsilon$
- Let S denote any cylinder from Ξ . Given $\varepsilon > 0$, there exists a divergent sequence $\{R_j\} \subset [0, \infty)$ such that each sufficiently large j version of C_j intersects the $s \in [-R_j, \infty)$ portion of the radius ε tubular neighborhood of S where the distance to S is no greater than $\frac{1}{2}\varepsilon$, and it intersects the $s = -R_j$ slice at a point with distance to S equal to $\frac{1}{2}\varepsilon$.

As is explained next, both possibilities violate the assumed convergence of the images of $\{C_j\}$ in the \mathbb{R} factor $\mathbb{R} \times \widehat{O} \widehat{A} / \text{Aut} \widehat{A}$.

In the case that the top point in (3–2) holds, this can be seen by making a new sequence whose j 'th subvariety is a j -dependent, constant translation of C_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. Were (4–19) to hold, a new sequence of this sort could be found whose limit data set in [15, Proposition 3.7] has an irreducible subvariety with the following incompatible properties: It is not an \mathbb{R} -invariant cylinder, it has the same θ infimum as each $\{C_j\}$, and there is no non-zero constant b that makes (2–4) hold for the corresponding end.

Were the second point in (3–2) to hold, then one of the limit cylinders from Ξ would be the $\theta = 0$ cylinder and then the convergence dictated by (4–19) would demand a $(1, \dots)$ element in \widehat{A} .

An argument much like that just used to rule out the first point of (3–2) rules out the third point as well. In this case, the derived nonsense from the translated sequence is an irreducible subvariety from the data set of the corresponding version of [15, Proposition 3.7] with the following mutually incompatible properties: First, it is not the $\theta = 0$ cylinder but contains an end where the $|s| \rightarrow \infty$ limit of θ is zero. Second, the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about the constant $|s|$ slices of this end have the form $\frac{1}{m} p$ and $\frac{1}{m} p'$ where (p, p') is the pair from the $(1, \dots)$ element in \widehat{A} and where $m \geq 1$ is some common divisor of this same pair. Finally, there would be no non-zero version of the constant \widehat{c} that would make (1–9) hold. Note that the argument for the second property uses the fact that there are no critical points of θ with θ value in $(0, \pi)$ on any C_j . Indeed, the lack of critical points implies that the component of the C_j version of $C_0 - \Gamma$ where the $|s| \rightarrow \infty$ limit of θ is zero is parametrized by a j -dependent map from a j -independent cylinder. Of course, this same cylinder parametrizes the corresponding component in any translation of C_j . Granted this, apply of [15, Proposition 3.7] to the translated sequence to find the asserted form for the integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$. \square

Part 2 This part of the argument proves that Ξ contains but a single element. The proof assumes the converse so as to derive an absurd conclusion. For this purpose, introduce $\mathcal{S} \subset \mathbb{R} \times (S^1 \times S^2)$ to denote the union of the subvarieties from S that are not \mathbb{R} -invariant cylinders. Then set Y to denote the subset of \mathcal{S} that contains the images of critical points of θ from the model curves of the irreducible components of \mathcal{S} , the singular points of $\cup_{(S,n) \in \Xi} S$, and the points on \mathcal{S} where $\theta \in \{0, \pi\}$. This set Y is finite.

The first point to make is that Ξ contains at least one element whose subvariety is not an \mathbb{R} -invariant cylinder. Indeed, this follows from an appeal to Proposition 4.6. However, Ξ cannot have two elements whose subvarieties are not \mathbb{R} -invariant. To explain, suppose first that there are two such subvarieties from Ξ and a value of θ that is taken simultaneously on both. Since Y is a finite set, this value can be taken so as to be disjoint from any value of θ on Y , and also disjoint from any value of θ from the set $\Lambda_{\hat{A}}$. Let θ_* denote such a value. Then, the $\theta = \theta_*$ locus in \mathcal{S} is at least two disjoint circles.

To see where this leads, introduce the notion a ‘ θ -preserving preimage’ in [15, Step 4 of Section 3.D]. Note here that a compact submanifold in $\mathcal{S} - Y$ has a well defined θ -preserving preimage even if C_j has nearby immersion points. This said, each circle from the $\theta = \theta_*$ locus in \mathcal{S} has θ -preserving preimages in each sufficiently large j version of C_j . The set of these preimages gives a set of disjoint, embedded circles on which $\theta = \theta_*$. However, there can be at most one such circle.

Thus, if Ξ has two subvarieties that are not \mathbb{R} -invariant, then the supremum of θ on one must be the infimum of θ on the other. To rule out this possibility, let S and S' denote the subvarieties that are involved, and suppose that θ_* is the supremum of θ on S and the infimum of θ on S' . Note that θ_* must be the value of θ on some bivalent vertex in $T\hat{A}$. Let e denote the edge in $T\hat{A}$ where the maximum of θ is θ_* and let e' denote that where the minimum of θ is θ_* . Let P denote the relatively prime integer pair that defines θ_* via (1-8). As is explained next, both Q_e and $Q_{e'}$ must be non-zero multiples of P , and this is impossible for the following reason: The $\theta = \theta_*$ locus in each C_j is a non-empty union of arcs on which both the $Q = Q_e$ and $Q = Q_{e'}$ versions of (2-2)'s function α_Q are positive. Thus, the $Q = P$ version of α_Q is positive at $\theta = \theta_*$, and this contradicts (1-8).

To see why Q_e is proportional to P , let $\delta > 0$ be very small. Proposition 4.6 guarantees that the integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ around the $\theta = \theta_o - \delta$ circle in any large j version of C_j is proportional to its integral around the analogous circle in S , and the latter integral is proportional to P_o . The analogous argument using the $\theta = \theta_* + \delta$ circles proves the claim about $Q_{e'}$.

The proof that there are no \mathbb{R} -invariant cylinders from Ξ is given next. For this purpose, assume to the contrary that one subvariety from Ξ is an \mathbb{R} -invariant cylinder. Let S_* denote this cylinder. The second point in (4-17) requires that θ 's value on S_* is its value at some bivalent vertex on T . Let o denote the vertex involved. In what follows, (p, p') are the relatively prime integers that define θ_o via (1-8).

Given $\varepsilon > 0$, let C_{j*} denote the subset of C_j whose points have distance less than ε from S_* . If j is large, the second point in (4-17) requires that C_{j*} have both a

concave side end and a convex side end of C_j . The next lemma asserts that there is a path in any large j version of C_{j*} from the convex side end in C_{j*} to the concave side one.

Lemma 4.8 *If $\varepsilon > 0$ is small, then all sufficiently large j versions of C_{j*} contain an embedding of \mathbb{R} on which s is neither bounded from above nor below. Moreover, at large $|s|$ this embedding sits in the $\theta = \theta_o$ locus in C_{j*} .*

Accept this lemma for the moment to see where it leads. For this purpose, note that the 1-form $pd\varphi - p'dt$ is zero on S_* and thus exact on the radius ε tubular neighborhood of S_* . It can therefore be written as df on this neighborhood where f is a smooth function that vanishes on S_* . Now let $\gamma_j \subset C_{j*}$ denote a given large j version of the line from Lemma 4.8. Because f vanishes on S_* , the line integral of $pd\varphi - p'dt$ along γ_j has small absolute value. In fact, if the j -version of this absolute value is denoted by ν_j , then the second point in (4-17) demands that $\lim_{j \rightarrow \infty} \nu_j = 0$. This last conclusion is nonsense as can be seen using (2-5) to reinterpret the integral of $pd\varphi - p'dt$ along γ_j as a j -independent multiple of the sum of one or more of the coordinates of the image of C_j in the simplex Δ_o . To explain, let e denote the incident edge in T that contains o as its largest angle vertex and fix any parameterization for K_e . Since the large $|s|$ part of γ_j lies in the $\theta = \theta_o$ locus, referral to (2-5) finds that ν_j is the absolute value of the integral of $(pq_e' - p'q_e)dv$ between two missing points on the $\sigma = \theta_o$ circle of parametrizing cylinder for K_e . According to Part 3 in Section 3c, the latter integral is a fixed multiple of a sum of one or more of the coordinates from C_j 's image in the Δ_o factor that is used in (3-12) to define C_j 's image in $\widehat{O}^{\widehat{A}}/\text{Aut}^{\widehat{A}}$. In particular, such a sum must have a positive, j -independent lower bound if the image of $\{C_j\}$ is to converge in $\widehat{O}^{\widehat{A}}/\text{Aut}^{\widehat{A}}$.

Proof of Lemma 4.8 There are three steps to the proof.

Step 1 Any given $(\rho, \chi) \in \mathbb{R} \times S^1$ defines a diffeomorphism of $\mathbb{R} \times (S^1 \times S^2)$ by sending (s, t, θ, φ) to $(s + \rho, t + p\chi, \theta, \varphi + p'\chi)$. All such diffeomorphisms act transitively on S_* with trivial stabilizer and so parametrize S_* once a point in S_* is chosen as the point where $\rho = 0$ and $\chi = 0$. Such a parameterization of S_* is used implicitly in what follows.

Fix a very small radius pseudoholomorphic disk in $\mathbb{R} \times (S^1 \times S^2)$ whose closure intersects S_* at a single point, this its center. For example, a disk of this sort can be found in one of the pseudoholomorphic cylinders where either t or φ is constant (see, eg [14, Subsection 4(a)].) Let D denote such a disk. The translates of D by the just

described $\mathbb{R} \times S^1$ group of diffeomorphisms defines an embedding of $S_* \times D$ into $\mathbb{R} \times (S^1 \times S^2)$ as a tubular neighborhood of S_* .

Step 2 Let S denote the subvariety from Ξ that is not \mathbb{R} -invariant. If S intersects S_* , it does so at a finite set of points. The subvariety S can also approach S_* at large $|s|$ if S has ends whose constant $|s|$ slices converge as $|s| \rightarrow \infty$ as a multiple cover of the Reeb orbit that defines S_* . In fact such ends must exist if $S \cap S_* = \emptyset$ since each C_j is connected and since (4–17) holds. By the same token, if there are no such ends, then S must intersect S_* .

Given $\varepsilon > 0$ and small, let $D_\varepsilon \subset D$ denote the radius ε subdisk about the origin. If ε is small enough, then the points where S intersects $S_* \times D_\varepsilon$ are of two sorts. The first lie in a small radius ball about the points where S intersects S_* . As $\varepsilon \rightarrow 0$, the radii of these balls can be taken to zero. The second sort are points where the absolute value of the parameter ρ on S_* is uniformly large. Of course, points of the latter sort exist if and only if S has an end whose constant $|s|$ slices converge as $|s| \rightarrow \infty$ to a multiple cover of the Reeb orbit that defines S_* . Assuming that such is the case, then each small ε has an assigned positive and large number, ρ_ε , this a lower bound for the absolute value of the parameter ρ on S_* at the points in $S \cap (S_* \times D_\varepsilon)$ with distance 1 or more from $S \cap S_\varepsilon$. In this regard, the assignment $\varepsilon \rightarrow \rho_\varepsilon$ can be made in a continuous fashion. It is perhaps needless to add that there is no finite $\varepsilon \rightarrow 0$ limit of the collection $\{\rho_\varepsilon\}_{\varepsilon > 0}$.

In the case where ρ_ε is defined, the analysis in [14, Section 2] can be used to find an even larger ρ'_ε together with non-negative integers n_+ and n_- with the following significance: There are precisely n_+ intersections between S and the disk fiber of $S_* \times D_\varepsilon$ over any point where $\rho \geq \rho'_\varepsilon$. Here, all have intersection number 1 and all occur in the subdisk $D_{\varepsilon/2}$. There are also precisely n_- intersections between S and the fiber of $S_* \times D_\varepsilon$ over any point where $\rho \leq -\rho'_\varepsilon$; all of these have intersection number 1 and lie in $D_{\varepsilon/2}$.

Step 3 Fix $\varepsilon > 0$ but very small, and let $C_{j\varepsilon} \subset C_j$ denote the intersection of C_j with $S_* \times D_\varepsilon$. Let n_* denote the integer that is paired with S_* in Ξ . If j is large, then $C_{j\varepsilon}$ intersects each fiber disk in $S_* \times D_\varepsilon$ where $|\rho| \leq \rho_\varepsilon$ in n_* points counted with multiplicity. In this regard, each such point has positive local intersection number. Moreover, these intersection points vary continuously with the chosen base point in S_* . This is to say that the assignment of the n_* intersection points to the base point defines a continuous map from the $|\rho| \leq \rho_\varepsilon$ portion of S_* to $\text{Sym}^{n_*}(D_\varepsilon)$. As a consequence, each component of the $\rho = \rho_\varepsilon$ slice of any large j version of $C_{j\varepsilon}$ is connected in the

$|\rho| \leq \rho_\varepsilon$ part of $C_{j\varepsilon} \cap (S_* \times D_\varepsilon)$ to at least one component in the $\rho = -\rho_\varepsilon$ slice, and vice versa.

Thus, if ε is small, then there is no obstruction to choosing a path between $\rho = \rho_\varepsilon$ and $\rho = -\rho_\varepsilon$ slices of any sufficiently large j version of $C_{j\varepsilon}$.

The argument just used can be repeated to extend the path just chosen, first as a path from the $\rho = \rho_\varepsilon'$ slice of $C_{j\varepsilon}$ to the latter's $\rho = -\rho_\varepsilon'$ slice, and then to arbitrarily large values of $|\rho|$. For example, to extend the path to where $\rho_\varepsilon = \rho \leq \rho_\varepsilon'$, note first that the number of intersections counted with multiplicity between $C_{j\varepsilon}$ and the fiber disks over this part of S_* may change. Even so, when j is very large, all such intersections lie either very close to S or very close to S_* . In particular, there are precisely n_* such intersections in each disk that lie very much closer to S_* than to S . Granted this, the preceding argument implies that any component of the $\rho = \rho_\varepsilon$ slice of $C_{j\varepsilon}$ is connected to one or more components of the $\rho = \rho_\varepsilon'$ slice of $C_{j\varepsilon}$.

To explain why this last extension can be continued to arbitrarily large values of ρ , note first that when j is large, then C_{j*} has precisely $n_* + n_+ n$ intersections with any given fiber disk over the $\rho \geq \rho_\varepsilon'$ portion of S_* when counted with multiplicity. Here, n is the integer that is paired with S in Ξ . Note that all intersections are again positive, and that these disk intersections now define a continuous map from the $\rho \geq \rho_\varepsilon'$ portion of S_* to $\text{Sym}^{n_* + n_+ n}(D)$. This then means that any given component of the $\rho = \rho_\varepsilon'$ slice of $C_{j\varepsilon}$ is part of a slice of a component of the $\rho \geq \rho_\varepsilon'$ part of $C_{j\varepsilon}$ where ρ has no finite upper bound. \square

Part 3 Let S denote the single element from Ξ and let S_0 denote the model curve for S . An argument from Part 2 can be used to prove that the pull-back of θ to S_0 has no non-extremal critical points. Indeed, were there such a point, then there would be an open interval of disconnected, compact θ level sets in S_0 . Most of these level sets would avoid the set Y , and any of the latter would have θ -preserving preimages in any sufficiently large j version of C_j . Of course, no such thing is possible since C_j lacks disconnected, compact θ level sets.

A very similar argument implies the following: Let $E \subset S$ denote any end where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π . Then E 's version of (2-4) has integer $n_E = 1$.

These last points have the following consequence: Let T_S denote a graph with a correspondence in a pair whose first component is the model curve for S . Then T_S is necessarily linear. Moreover, by virtue of Proposition 4.6, the vertex set of T_S enjoys an angle preserving, 1-1 correspondence with that of $T^{\hat{A}}$. Meanwhile, any two $C = C_j$ and $C = C_{j'}$ versions of the graph T_C must be isomorphic when j and j'

are large because the image of $\{C_j\}$ converges in $\widehat{O\hat{A}}/\text{Aut}\hat{A}$. Let T denote a graph in the isomorphism class. The remainder of this Part 3 explains why the graphs T_S and T are isomorphic. The explanation is given in eight steps.

Step 1 Each edge in T has its evident analog in T_S since the vertex sets share the same angle assignments. If $e \in T$ is a given edge, then the corresponding T and T_S versions of the integer pair Q_e are related as follows: The T version is n times the T_S version where n here is the integer that pairs with S in Ξ . This then means that $n = 1$ in the case that the integers that comprise the suite of T versions have 1 as their greatest common divisor. In particular if n is larger than 1, then it must divide both integers that comprise the version of $Q_{(\cdot)}$ whose labeling edge has the smallest angle vertex in T .

Step 2 Let o denote a bivalent vertex in T_S and let $E \subset S$ denote an end where the $|s| \rightarrow \infty$ limit of θ is θ_o . Since E 's version of (2–4) uses the integer $n_E = 1$, the $\theta = \theta_o$ locus in the model curve for S intersects the very large $|s|$ portion of E as a pair of open arcs, one oriented in the increasing $|s|$ direction and the other oriented with $|s|$ decreasing. Moreover, given a positive and very large number, R , there is an arc in the $\theta \leq \theta_o$ portion of E with the following properties:

- (4–20) • $|s| \geq R$ on the arc.
- The arc starts on one component of the $\theta = \theta_o$ locus in the $|s| \geq R$ part of E and ends on the other.
 - The function θ restricts to this arc so as to have but a single critical point, this its minimum.

Let ν denote such an arc. Note that ν has an analog that shares its endpoint and sits where $\theta \geq \theta_o$. The latter is denoted below by ν' .

Step 3 The arc ν has θ -preserving preimages in every large j version of C_j ; there are n such preimages, each with its endpoints on the $\theta = \theta_o$ locus in C_j . To see how these appear, let e denote the edge in T whose largest angle is θ_o . Fix a parameterization of the component K_e in C_j 's version of $C_0 - \Gamma$. Each θ preserving preimage of ν appears in the corresponding parametrizing cylinder as a closed arc with both endpoints on the $\sigma = \theta_o$ circle. Denote these arcs as $\{\nu_{j,k}\}_{1 \leq k \leq n}$.

To say more about these preimages, remember that the constant $|s|$ slices of E converge as $|s| \rightarrow \infty$ as a multiple cover of a $\theta = \theta_o$ closed Reeb orbit in $S^1 \times S^2$. The latter has a tubular neighborhood with a function, f , with the following two properties: First,

it vanishes on the Reeb orbit in question. Second, $df = \frac{1}{2\pi}(pd\varphi - p'dt)$ where p and p' are the relatively prime integers that define θ_o via (1–8). Because f has limit 0 as $|s| \rightarrow \infty$ on E , the integral over ν of df is small. Moreover, given $\varepsilon > 0$, there exists R_ε such that the integral of df over any $R \geq R_\varepsilon$ version of ν from (4–20) has absolute value less than ε .

Granted this, it then follows that the same is true for the integral of df over any $R \geq R_\varepsilon$ and sufficiently large j version of any of the θ -preserving preimages of ν . This understood, use the parametrization of K_e to identify the $\text{mod } (2\pi\mathbb{Z})$ image of the latter integral as that of $(pqe' - p'qe)dv$ between the two endpoints of the arc $\nu_{j,k}$. In particular, this means that when R is large and j is large, the two endpoints of $\nu_{j,k}$ on the $\sigma = \theta_o$ circle of the parametrizing cylinder are very close to each other. They divide the circle into one very short arc and one arc on which the integral of dv is almost 2π .

Step 4 Because the images of $\{C_j\}$ converge in $\widehat{O\hat{A}}/\text{Aut}\hat{A}$, their images converge in the Δ_o factor of (3–12). This fact and the final conclusion from Step 3 have the following consequence: The very short arc in the $\sigma = \theta_o$ parametrizing cylinder circle between the endpoints of any large R and correspondingly large j version of $\nu_{j,k}$ contains at most one missing point in its interior.

As is explained next, such a $\sigma = \theta_o$ arc must contain precisely one missing point. For this purpose, let θ_* denote the minimum of θ on ν . Were there no missing point in the indicated short arc, then one of the θ -preserving preimages of ν in C_j would be homotopic rel boundary in the $[\theta, \theta_*]$ portion of C_j to an arc lying entirely in the $\theta = \theta_o$ locus. Such a homotopy could then be projected back to S to give a homology rel boundary in the model curve for S between the arc ν and the disjoint union of an arc in the $\theta = \theta_o$ locus and a union of very small radius circles, each surrounding some point that maps to one of the immersion singularities in S . No such homology is possible because any constant $|s|$ slice of E generates a non-trivial homology class in S_0 .

Step 5 As a consequence of the result from Step 4, each end of S where $\lim_{|s| \rightarrow \infty} \theta$ is neither 0 nor π labels n ends of any sufficiently large j version of C_j . These n ends have the same $|s| \rightarrow \infty$ limit of θ as their namesake in S ; and by virtue of Proposition 4.6, they are all convex side if their namesake is a convex side end. Otherwise, they are all concave side ends. Here is one way to view this correspondence: Let $E \subset S$ denote an end as in the previous steps. Fix some large R , and the concatenation of the arc ν with its $\theta \geq \theta_o$ cousin ν' defines a closed loop in E that is homologous to a constant $|s|$ slice. Such a loop has θ -preserving preimages in each sufficiently large j version

of C_j . Any such preimage must be a union of one θ -preimage of ν and one of ν' . Indeed, it must, in any event, contain the same number of ν preimages as ν' preimages; and said number must be 1 because of the very small length of one of the arcs between the endpoints of each $\nu_{j,k}$ in the $\sigma = \theta_o$ circle of the relevant parametrizing cylinder.

Granted the preceding the concatenation of ν with ν' has n distinct θ preserving preimages in all large j versions of C_j . Since the aforementioned short arc between the endpoints of each $\nu_{j,k}$ contains one and only one missing point on the $\sigma = \theta_o$ circle, each preimage of the $\nu - \nu'$ concatenation is homologous in C_j to the constant $|s|$ slice in some end of C_j where the $|s| \rightarrow \infty$ limit of θ is θ_o . Moreover, these preimages account for n distinct ends in C_j . Finally, an appeal to Proposition 4.6 finds that distinct $\lim_{|s| \rightarrow \infty} \theta = \theta_o$ ends of S label disjoint, n -element subsets of such ends in C_j .

Step 6 An end $E \subset S$ corresponds to a vertex in the S version of the circular graph $\underline{\Gamma}_o$. As such, it comes with an integer weight. Meanwhile, the corresponding n ends in C_j correspond to n vertices on the C_j version of the graph $\underline{\Gamma}_o$. As is explained here, each of the latter vertices have the same integer weight as E 's vertex.

The conclusions of Step 5 guarantee that the $n + 1$ weights involved all have the same sign. Here is how to compute their magnitudes: The integral of the form $\frac{1}{2\pi}(pdt + p'd\varphi)$ over any constant $|s|$ slice of a relevant end has the form $m(p^2 + p'^2)$ where m is the desired magnitude.

Apply this last observation first to the concatenation of ν and its $\theta \geq \theta_o$ cousin ν' in a given end $E \subset S$. Since this concatenation is homologous to a constant $|s|$ slice, the form $\frac{1}{2\pi}(pdt + p'd\varphi)$ integrates over this concatenation to give $m_E(p^2 + p'^2)$ where m_E is the absolute value of the integer weight for E 's vertex in the S version of $\underline{\Gamma}_o$. Next, apply the observation to any one of the θ -preserving preimages of the $\nu - \nu'$ concatenation in each very large j version of C_j . As there are n of these, the integral over any one is $m_E(p^2 + p'^2)$ as well. The desired equality follows because one of these preimages is homologous to the constant $|s|$ slices any given E -labeled end in C_j .

Step 7 The results of the previous steps imply that T_S is isomorphic to T in the case that $n = 1$. Such an isomorphism is obtained via the correspondence given in Step 5 between the ends of S and those of C_j . In the hypothetical $n > 1$ case, it implies that each T version of the group Aut_o has a $\mathbb{Z}/(n\mathbb{Z})$ subgroup. More to the point, the following is true: Let $O^{\hat{A}}_\nu \subset O^{\hat{A}}$ denote a component whose $\text{Aut}^{\hat{A}}$ orbit contains the limit point of the image of $\{C_j\}$. Then $\text{Aut}^{\hat{A}}_\nu$ has a canonical $\mathbb{Z}/(n\mathbb{Z})$ subgroup

since n also divides the integer pair that is associated to the edge in T with the smallest angle vertex.

This step explains an observation that is used in the subsequent step in two ways: It is used to prove that the image of S in $\mathbb{R} \times \widehat{O}^{\widehat{A}} / \text{Aut}^{\widehat{A}}$ is the limit of the images of $\{C_j\}$ in the case that $n = 1$, and it is used to preclude the case that $n > 1$.

To start, let E and E' denote ends of S where $\lim_{|s| \rightarrow \infty} \theta = \theta_o$ and such that travel in the oriented direction along a component of the $\theta = \theta_o$ locus in the model curve of S proceeds from large $|s|$ on E to large $|s|$ on E' . The case that $E = E'$ is allowed here. In any event, let γ denote the component of the $\theta = \theta_o$ locus in question and let r_γ denote the integral along γ of the 1-form $(1 - 3 \cos^2 \theta_o) d\varphi - \sqrt{6} \cos \theta_o dt$.

Now, fix R very large and let z denote the endpoint on γ of E 's version of the arc ν . Meanwhile, let z' denote the endpoint on γ of the E' version of this arc in the case that $E' \neq E$. If $E' = E$, take z' to be the second of the two endpoints of the arc ν . Deform γ slightly if it passes through a point in Y so that the result misses Y , lives where $\theta \leq \theta_o$, and agrees with γ on E and E' . Let γ_R denote the portion of such a deformation that runs between z and z' . If R is large, then r_γ is very nearly the integral along γ_R of the pull-back of the 1-form $(1 - 3 \cos^2 \theta_o) d\varphi - \sqrt{6} \cos \theta_o dt$. Use $r_{\gamma,R}$ to denote the latter integral. Thus, the $R \rightarrow \infty$ limit of $\{r_{\gamma,R}\}$ is r_γ .

With R fixed and then j very large, the arc γ_R has n disjoint, θ preserving preimages in any parametrizing cylinder for the C_j version of the component K_e . Any such preimage runs from very close to one of E 's missing points on the $\sigma = \theta_o$ circle in the oriented direction to a point that is very close to the subsequent missing point, this one labeled by E' . This understood, it then follows from Proposition 4.6 that the integral of the $Q = Q_e$ version of $\alpha_Q(\theta_o) dv$ between these two missing point is very nearly r_γ .

Now, this arc between the two missing points labels one of the coordinates for C_j in the symplectic Δ_o that appears in (3–12), and it follows from what has just been said that the value of this coordinate is very nearly r_γ .

Step 8 In the case that $n = 1$, the results of the preceding step assert that the assigned point to any large j version of C_j in any given Δ_o is very near that assigned to S . The implication is that the $j \rightarrow \infty$ limit of these images is that of S .

In the hypothetical $n > 1$ case, the results from the preceding step imply that the image of any large j version of C_j in any given Δ_o is very close to the set of points that are fixed by the $\mathbb{Z}/(n\mathbb{Z})$ subgroup version of $\text{Aut}_{o,v}$. In fact, the results from the preceding step imply that the image of C_j in the space $\times_o \Delta_o$ is very near the fixed

set of the canonical $\mathbb{Z}/(n\mathbb{Z})$ subgroup of $\times_o \text{Aut}_{o,v} \subset \text{Aut}^{\hat{A}}$. Granted the observations from Proposition 4.6, this then implies that the image in $\hat{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$ of any large j version of C_j is very close in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ to $(O^{\hat{A}} - \hat{O}^{\hat{A}})/\text{Aut}^{\hat{A}}$.

This last conclusion rules out the $n > 1$ case because it is incompatible with the initial assumption of the convergence in $\hat{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$ of the image of $\{C_j\}$.

5 The first chapter of story when $N_- + \hat{N} + \zeta_- + \zeta_+$ is greater than 2

Introduce as in Subsection 1.C, the larger space $\mathcal{M}^*_{\hat{A}}$ whose elements consist of honest subvarieties in the sense of (1–5) along with ‘multiple covers’ of honest subvarieties. Subsection 1.C also introduced a stratification of $\mathcal{M}^*_{\hat{A}}$. The first subsection below summarizes results about the local structure of $\mathcal{M}^*_{\hat{A}}$ and its stratification. The remaining subsections contain the proofs of these results.

5.A The local structure of $\mathcal{M}^*_{\hat{A}}$ and its stratification

As outlined in the first section, the space $\mathcal{M}^*_{\hat{A}}$ consists of equivalence classes of pairs (C_0, ϕ) where C_0 is a complex curve homeomorphic to an $N_+ + N_- + \hat{N}$ times punctured sphere and ϕ is a proper, pseudoholomorphic map from C_0 into $\mathbb{R} \times (S^1 \times S^2)$ whose image is a subvariety as defined in (1–5). Moreover, the pair (C_0, ϕ) must be compatible with the data set \hat{A} in the following sense: First, there is a 1–1 correspondence between the ends of C_0 and the 4–tuples in \hat{A} ; and this correspondence must pair an end and a 4–tuple if and only if the 4–tuple comes from the end as described in Subsection 1.A. Second, the integers ζ_+ and ζ_- are the respective intersection numbers between C_0 and the $\theta = 0$ and $\theta = \pi$ cylinders. To be precise here, note that there is a finite number of $\theta = 0$ and $\theta = \pi$ points in C_0 . This understood, a sufficiently generic but compactly supported perturbation of ϕ gives an immersion of C_0 into $\mathbb{R} \times (S^1 \times S^2)$ that has transversal intersections with the $\theta = 0$ and $\theta = \pi$ loci. The latter has a well defined intersection number with both loci, and these are respectively ζ_+ and ζ_- . As noted in Subsection 1.C, pairs (C_0, ϕ) and (C_0, ϕ') define the same point in $\mathcal{M}^*_{\hat{A}}$ if ϕ' is obtained from ϕ by composing with a holomorphic diffeomorphism of C_0 .

A base for the topology on $\mathcal{M}^*_{\hat{A}}$ is depicted in (1–24). Theorem 1.3 asserts that $\mathcal{M}^*_{\hat{A}}$ is a smooth orbifold whose singular points consist of those pairs (C_0, ϕ) where there is a holomorphic diffeomorphism that fixes ϕ . Theorem 1.3 also asserts that the inclusion

of $\mathcal{M}_{\hat{A}}$ in $\mathcal{M}^*_{\hat{A}}$ is a smooth embedding onto an open subset. Theorem 1.3 is proved below so grant it for the time being.

As noted in Subsection 1.C, the strata of $\mathcal{M}^*_{\hat{A}}$ are indexed by ordered triples of the form (B, c, \mathfrak{d}) where $B \subset \hat{A}$ is a set of $(0, -, \dots)$ elements, c is a non-negative integer no greater than $N_+ + N_- + \hat{N} + \zeta_- + \zeta_+ - 2 - |B|$ and \mathfrak{d} is a partition of the integer $d \equiv N_+ + |B| + c$ as a sum of positive integers. By way of a reminder, the stratum $\mathcal{S}_{B,c,\mathfrak{d}} \subset \mathcal{M}^*_{\hat{A}}$ lies in the subset, $\mathcal{S}_{B,c}$ that consists of the equivalence classes of pairs (C_0, ϕ) that have two properties stated next. In these statements and subsequently, functions on $\mathbb{R} \times (S^1 \times S^2)$ and their pull-backs via ϕ are not distinguished by notation except in special circumstances. Here is the first property: The curve C_0 has precisely c critical points of θ where the value of this function is neither 0 nor π . Here is the second: The ends that correspond to elements in B are the sole convex side ends of C_0 where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π and whose version of (2–4) has a strictly positive integer n_E . To define $\mathcal{S}_{B,c,\mathfrak{d}}$, introduce the map I_d to denote the space of unordered d -tuples of points in $(0, \pi)$ and the map $f: \mathcal{S}_{B,c} \rightarrow I_d$ that sends a given (C_0, ϕ) to the d -tuple that consists of the critical values in $(0, \pi)$ of θ 's pull-back to C_0 and the $|s| \rightarrow \infty$ limits in $(0, \pi)$ of θ on the concave side ends of C_0 and on the ends that correspond to the 4-tuples in B . The stratum $\mathcal{S}_{B,c,\mathfrak{d}}$ is the inverse image via f of the stratum in I_d that is labeled by the partition \mathfrak{d} .

The next proposition describes the local structure of $\mathcal{S}_{B,c,\mathfrak{d}}$. The proposition speaks of a locally constant function on $\mathcal{S}_{B,c,\mathfrak{d}}$ whose value at a given point defined by some pair (C_0, ϕ) is the number of distinct critical values of θ in the subset of the critical values in $(0, \pi)$ that do not arise via (1–8) from an integer pair component of either a $(0, +, \dots)$ 4-tuple in \hat{A} or a 4-tuple in B . Let m denote this locally constant function.

Proposition 5.1 *If not empty, then the stratum $\mathcal{S}_{B,c,\mathfrak{d}}$ is a smooth orbifold in \mathcal{M}^* that intersects $\mathcal{M}^* - \mathbb{R}$ as a smooth manifold. In this regard, a given component has dimension $N_+ + |B| + c + m + 2$.*

A description of the components of any given version of $\mathcal{S}_{B,c,\mathfrak{d}}$ is provided in Sections 6 and 8.

The proof of Theorem 1.3 has five parts and these are presented the next two subsections. The final subsection contains the proof of Proposition 5.1.

5.B Parts 1–4 of the proof of Theorem 1.3

The proof is very much like that in [15, Proposition 2.9]. The five parts that follow focus on the points that differ.

Part 1 Suppose that (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A}}$. There is, in all cases, a fixed radius ball subbundle $B_1 \subset \phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ and an exponential map, e , that maps B_1 into $\mathbb{R} \times (S^1 \times S^2)$ so as to embed each fiber and send the zero section to C_0 . As noted in [15, Section 2.D] for pairs that map to $\mathcal{M}_{\hat{A}}$, the bundle $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ splits as a direct sum, $W \oplus N$, of complex line bundles such that the differential, ϕ_* , of ϕ maps $T_{1,0}C_0$ into W , and N restricts to the points where $\phi_* \neq 0$ as the pull-back normal bundle. In this regard, e can be chosen so as to embed the fibers of $B_1 \cap N$ and of $B_1 \cap W$ as pseudoholomorphic disks. Note that in case where ϕ is not almost everywhere 1–1, there is a complex curve, C_1 , with an attending, almost everywhere 1–1, pseudoholomorphic map, ϕ_1 , to $\mathbb{R} \times (S^1 \times S^2)$ whose image is $C \equiv \phi(C_0)$. In this case, ϕ factors as $\phi_1 \circ \psi$ where ψ is a holomorphic, branched covering map to C_1 . The $\phi_1^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ decomposition as $W \oplus N$ then pulls back by ψ to give the $W \oplus N$ decomposition for $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$. Because of this factoring property, the map e can be chosen to be invariant under the action on $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ of the group of holomorphic diffeomorphisms of C_0 that fix ϕ . Such a choice is assumed in what follows.

Part 2 [15, Section 2.D] described an operator, D_C , whose domain is a certain Hilbert space of sections of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ and whose range is a Hilbert space of sections of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2)) \otimes T^{0,1}C_0$. The discussion in [15, Section 2.D] involves only pairs (C_0, ϕ) that define points in $\mathcal{M}_{\hat{A}}$, but the story generalizes in an almost verbatim fashion to define D_C for any pair that defines a point in $\mathcal{M}^*_{\hat{A}}$.

By way of a reminder, D_C is defined from an operator, \underline{D} , whose kernel is the space of first order deformations of ϕ that result in maps that are pseudoholomorphic with respect to the given complex structure on C_0 and the almost complex structure J . The salient features of \underline{D} are as follows: First, \underline{D} is a first order differential operator that maps sections of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ to sections of $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ so as to map sections of W to those of $W \otimes T^{0,1}C_0$. In particular, if v is a section of $T_{1,0}C_0$, then ϕ_*v is a section of W and $\underline{D}\phi_*v = \phi_*(\bar{\partial}v)$. Second, composing \underline{D} with orthogonal projection onto the N summand in $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ defines an \mathbb{R} -linear operator that sends a section, η , of N to a section of the form

$$(5-1) \quad \bar{\partial}\eta + \nu\eta + \mu\bar{\eta};$$

here ν is a fixed section of $T^{1,0}C_0$ and μ is a fixed section of $N^2 \otimes T^{1,0}C_0$.

The operator \underline{D} has an extension as a Fredholm operator that maps a certain weighted Hilbert space completion of its range to that of its domain. The inner products that define the range and domain Hilbert spaces for sections of the N and $N \otimes T^{1,0}C_0$ summands are as depicted in [15, Equation (2.7)]. Similar weighted norms are used for

the respective W and $W \otimes T^{1,0}C_0$ summands, but the latter insure that all sections are square integrable. In this regard, some care must be taken when there are holomorphic diffeomorphisms of C_0 that preserve ϕ . To elaborate, note first that ϕ is finite to one, and as a consequence, the set of diffeomorphisms of C_0 that preserve ϕ defines a finite group. Let G_C denote the latter. The norms used for these Hilbert spaces can and should be taken to be G_C invariant.

Part 3 In the case that C_0 is a disk or cylinder, the operator D_C is \underline{D} in the just described Fredholm context. When C_0 has negative Euler characteristic, D_C is obtained from this Fredholm \underline{D} by composing with an orthogonal projection on the latter's range. The definition of this projection requires the choice of some $3(N_+ + N_- + \hat{N} - 1)$ dimensional, G_C -invariant vector space of sections of $T_{1,0}C_0 \otimes T^{1,0}C_0$ that projects isomorphically to the cokernel of $\bar{\partial}$. Let V denote the latter choice and let \prod denote the orthogonal projection onto ϕ_*V . Then $D_C = (1 - \prod)\underline{D}$.

The following proposition describes the important facts about the kernel and cokernel of D_C . The proof uses verbatim arguments from the proof of [15, Propositions 2.9] and so is omitted.

Proposition 5.2 *Suppose that (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A}}$. Then the operator D_C has index $N_+ + 2(N_- + \hat{N} + \zeta_+ + \zeta_- - 1)$. Moreover, this is its kernel dimension as its cokernel is trivial.*

Part 4 Here is the significance of D_C : A small ball in the vector space V parametrizes the complex structures on C_0 that are near to the given one. This understood, an element in the kernel of D_C gives a deformation of ϕ that is pseudoholomorphic to first order with respect to a complex structure that is parametrized by a point in V . More to the point, the implicit function theorem can be employed in a relatively standard manner to obtain the following description of a neighborhood of the point defined by (C_0, ϕ) in $\mathcal{M}^*_{\hat{A}}$: There is a ball, $B \subset \text{kernel}(D_C)$, a smooth function, f , from B to $\text{cokernel}(D_C)$ that vanishes with its differential at zero, and a homeomorphism between a neighborhood of (C_0, ϕ) 's point in $\mathcal{M}^*_{\hat{A}}$ and the quotient of $f^{-1}(0)$ by the action of the group G_C on the kernel of D_C . To elaborate, this homeomorphism comes from a G_C -equivariant map, F , from B to the domain space of D_C that maps the origin to 0 with differential at 0 the identity on $\text{kernel}(D_C)$. The homeomorphism is obtained by restricting the composition $e \circ F$ to $f^{-1}(0)$.

It is worth a moment now to say something about why these local charts map onto a neighborhood of (C_0, ϕ) as defined by (1–24). The point here is that if (C'_0, ϕ') is

close to (C_0, ϕ) in the sense of (1–24), then the map $\phi' \circ \psi$ can be obtained from ϕ by composing the exponential map from $\phi^*T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ with a small normed section. A change in the diffeomorphism changes the section, and [14, Proposition 2.2] can be used to find diffeomorphisms that give small normed section in the domain of D_C . Granted this, the implicit function theorem asserts that there is a unique such small normed section from the image of F .

Note that in the case that $\text{cokernel}(D_C) = \{0\}$, then $\text{kernel}(D_C)/G_C$ is a local Euclidean orbifold chart for a neighborhood of (C_0, ϕ) 's point in $\mathcal{M}^*_{\hat{A}}$. As is usually the case with implicit function theorem applications of the sort just described, these charts fit together to give a smooth orbifold structure to the set of points in $\mathcal{M}^*_{\hat{A}}$ that are defined by pairs (C_0, ϕ) with trivial D_C cokernel. Granted this, the assertion in Theorem 1.3 about the local structure of $\mathcal{M}^*_{\hat{A}}$ follows directly from Proposition 5.2.

5.C Part 5 of the proof of Theorem 1.3

This part explains why the set inclusion of $\mathcal{M}_{\hat{A}}$ into $\mathcal{M}^*_{\hat{A}}$ is a topological embedding. Note that the equivalence of the two topologies is, in fact, implied by the statement of [14, Proposition 3.2]. However, the arguments in [14] for this proposition focused for the most part on issues that are not present in the analogous compact symplectic manifold assertion and so left the proof of the equivalence to the reader. In its deference to [14, Proposition 3.2], the proof of [15, Proposition 2.9] does not address the implied equivalence between the two topologies on $\mathcal{M}_{\hat{A}}$. The explanation that follows has nine steps.

Step 1 The inclusion $\mathcal{M}_{\hat{A}} \rightarrow \mathcal{M}^*_{\hat{A}}$ is continuous since the condition for closeness given by (1–24) implies that given in (1–13). Thus, to prove it an embedding, it is enough to prove that the condition for closeness in (1–13) implies that in (1–24). This understood, the task is as follows: Fix $C \in \mathcal{M}_{\hat{A}}$ and suppose that some positive, but small κ is given. Find some positive κ' such that when $C' \in \mathcal{M}_{\hat{A}}$ obeys the κ' version of (1–13), then there is a diffeomorphism between C 's model curve and that of C' that makes the κ version of (1–24) hold. In what follows, C_0 and C_0' are the respective model curves for C and C' while ϕ and ϕ' are their respective pseudoholomorphic maps to $\mathbb{R} \times (S^1 \times S^2)$.

To construct the required diffeomorphism, let $\vartheta \subset C_0$ denote the set of points that are mapped by ϕ to singular points of C . The bundle N restricts to $C_0 - \vartheta$ as the normal bundle to the embedding. Fix $\varepsilon > 0$ but very much less than one, and $\delta \in (0, \varepsilon^4)$. Now let U denote the union of the radius δ disks about the points in ϑ and the $|s| \geq 1 + 1/\varepsilon$ portions of C_0 . In this regard, choose ε so that the $|s| \geq 1 + 1/\varepsilon$ part of U is far out on the ends of C_0 .

With ε and δ chosen, there is an exponential map that is defined on a small, constant radius disk bundle in N over $C_0 - U$ so as to embed this disk bundle as a tubular neighborhood of $\phi(C_0 - U)$ and to embed each fiber as a pseudoholomorphic disk. Note that there is quite a bit of freedom here with the choice for this exponential map and this freedom is used in what follows to fine tune things near the boundary of $C_0 - U$. In any event, suppose that the disk bundle and the exponential map have been fixed. Let $N_1 \subset N$ denote the disk bundle and $e: N_1 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ the exponential map.

Now if C' is very close to C in the sense of (1–13), then C' must intersect the image of each fiber of N_1 over $C_0 - U$ in precisely one point with multiplicity one. Indeed, because C and C' come from the same version of $\mathcal{M}_{\hat{A}}$, the net intersection number with any given fiber must be one. Meanwhile, all such intersection points count with positive weight by virtue of the fact that the fibers are embedded as pseudoholomorphic disks.

Because C' intersects the image of each fiber of N_1 over $C_0 - U$ just once, it can be written in the tubular neighborhood of $\phi(C_0 - U)$ as the image of $e \circ \eta$ where η is a very small normed section of N_1 .

The characterization of C' as the image of $e \circ \eta$ defines a diffeomorphism, ψ , between $C_0 - U$ and a part C_0' by demanding that $\phi' \circ \psi = e \circ \eta$. This diffeomorphism is such as to make $\text{dist}(\phi, \phi' \circ \psi)$ and the ratio $r(\psi)$ from (1–24) both very small on the whole of $C_0 - U$ in the case that C' is very close to C in the sense of (1–13). The argument as to why $r(\psi)$ is small is deferred to Step 3.

Step 2 This step constitutes a digression make four points about 4–dimensional pseudoholomorphic geometry. To set the stage, let X denote the 4–manifold and J an almost complex structure on X . The relevant case is that where $X = \mathbb{R} \times (S^1 \times S^2)$ and J is the almost complex structure that is described in Section 1. Let D denote a standard disk in \mathbb{C} , and suppose that an embedding of $D \times D$ into a X has been specified with the following specific property: The image of $D \times 0$ and the image of each $\{(z \times D)\}_{z \in D}$ disk is pseudoholomorphic.

Point 1 *There is a ball about the image of $(0, 0)$ in X with complex coordinates (x, y) that have three properties: First, $y = 0$ is in the image of the disk $D \times 0$. Second, each constant x disk lies in the image of some disk from the collection $\{z \times D\}_{z \in D}$. Finally, $T^{1,0}X$ is spanned over this coordinate chart by the 1–forms*

$$(5-2) \quad v \equiv dx + \sigma d\bar{x} \quad \text{and} \quad v' \equiv dy + \sigma' d\bar{x},$$

where σ and σ' vanish both at the origin and along the whole $y = 0$ locus.

The proof that such coordinates exist is straightforward and left to the reader.

To make the remaining points, let B denote the coordinate chart just described, let $\Omega \subset \mathbb{C}$ denote a disk and let $w: \Omega \rightarrow B$ denote a proper, pseudoholomorphic map.

Point 2 *The pull-back of x to Ω obeys $\bar{\partial}x + \sigma\bar{\partial}\bar{x} = 0$. Indeed, this follows by virtue of the fact that w^*v is a section of $T^{1,0}\Omega$. As a consequence, $|\bar{\partial}x| \ll |\partial x|$ when $|y|$ is small on the image of Ω . Note that this implies that the critical points of the pull-back of x are the zeros of ∂x .*

Point 3 *Let $z \in \Omega$ denote a zero of ∂x . There is holomorphic coordinate, w , for a neighborhood of z such that $\partial x = w^q + \mathcal{O}(|w|^{q+1})$ where q is a positive integer.*

Indeed, this can be seen from the following considerations: The holomorphic derivative of the equation from Point 2 gives one for ∂x that has the form $\bar{\partial}(\partial x) + \gamma\partial x + \hat{\gamma}\bar{\partial}\bar{x} = 0$, where γ and $\hat{\gamma}$ are smooth functions on Ω . This last equation implies that ∂x vanishes near z to leading order as a holomorphic function.

The fourth point is an immediate consequence of the latter two:

Point 4 *Viewed as mapping $w^{-1}(B)$ to \mathbb{C} , the function x looks locally like a ramified covering map onto its image.*

Step 3 This step explains why $r(\psi)$ is small at all points in $C_0 - U$ when C' obeys a sufficiently small κ' version of (1–13). To start, remark that by virtue of what is said in Step 2, any given point in $\phi(C_0 - U)$ has local complex coordinates (x, y) with the following three properties: First, each $x = \text{constant}$ disk is the image of a fiber of N_1 . Second, the disk where $y = 0$ is in C . Finally, $T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$ is spanned by the 1-forms v and v' as in (5–2).

Let B denote the domain in $\mathbb{R} \times (S^1 \times S^2)$ of these coordinates. The map ψ^{-1} on the ϕ' -inverse image of B is the composition of ϕ' with the projection to the $y = 0$ locus. This being the case, the fact that ϕ' is pseudoholomorphic implies that v must pull-back via ϕ' to C_0' as a form of type $(1, 0)$, and this implies that $r_z(\psi) = |\bar{\partial}x|/|\partial x|$ is the value of $|\sigma|$ at $(e \circ \eta)(z)$.

Granted the preceding, there exists (ε, δ) -dependent constants $\kappa_0 > 0$ and c_0 with the following significance: If $\kappa' < \kappa_0$ and if C' obeys the κ' version of (1–13), then ψ is well defined on $C_0 - U$. Moreover, both $\text{dist}(\phi, \phi' \circ \psi)$ and $r_{(\cdot)}(\psi)$ are bounded by an expression of the form $c_0\kappa'$ at all points in $C_0 - U$.

Step 4 To extend ψ , so as to make (1–24) hold for all z and very small κ , note first that topological considerations imply that the complement of $\psi(C_0 - U)$ in C_0' must be diffeomorphic to U , thus a union of some number of cylinders and some number of disks. Moreover, each cylinder must bound one of the $|s| \sim 1/\varepsilon$ circles in the boundary of $\psi(C_0 - U)$ and each disk must bound one of the radius $\mathcal{O}(\varepsilon)$ circles in the boundary.

The cylinder story is simpler than that for the disks, so it is treated first. For this purpose, let $S \subset U$ denote one of the cylinder components, and let γ denote its boundary circle. Let $\gamma' \subset C_0'$ denote $\psi(\gamma)$ and let $S' \subset C_0'$ denote the component of the complement of $\psi(C_0 - U)$ whose boundary circle is γ' . When ε is large, both S and S' are very close to an \mathbb{R} -invariant, pseudoholomorphic cylinder, S_0 . In this regard, take S_0 so that a multiple cover of its defining Reeb orbit is the $|s| \rightarrow \infty$ limit of the constant $|s|$ slices of S . Fix a point $x \in S_0$ and a pseudoholomorphic disk, D with center at x that is normal to S_0 . Since TS_0 is an orbit of the product of \mathbb{R} with a 1-parameter subgroup from the group \mathbb{T} generated by ∂_t and ∂_φ , the corresponding $\mathbb{R} \times S^1$ group can be used to translate a small subdisk in D centered at x to each point in S ; and these translates foliate a tubular neighborhood of S_0 in $\mathbb{R} \times (S^1 \times S^2)$ by pseudoholomorphic disks. The exponential map on N_1 from Step 1 can and should be chosen so as to map each fiber of N_1 near γ into one of these $\mathbb{R} \times S^1$ translates of D .

Near γ , both S and S' intersect each fiber the same number of times and in distinct points. Let m denote this number. Let π denote the projection from the tubular neighborhood of S_0 to S_0 that moves any given point to the center point of its particular translate of D . As will now be explained, the restriction of π to either S or S' defines a degree m , unramified covering of $\pi(S)$. To see why, use the first point in Step 2 to put coordinates (x, y) on a neighborhood of any given point in S_0 where the $x = \text{constant}$ slices are the translates of D , and where the $y = 0$ locus corresponds to S_0 . Moreover, the 1-forms ν and ν' from (5–2) span $T_{1,0}(\mathbb{R} \times (S^1 \times S^2))$ on this neighborhood.

As set up, the projection π on the parts of S and S' in this neighborhood is the function x . This understood, the fourth point of Step 2 implies that π restricts to either S or S' as a degree m , ramified cover over $\pi(S)$. As such, its critical points are isolated, and each counts positively to a ramification number. As both S and S' are cylinders, the ramification number must be zero and so π maps S and S' to $\pi(S)$ as honest degree m covers.

Now introduce the fibered product $S \times_\pi S'$, this the subspace in $S \times S'$ of pairs with the same image via π in S_0 . The latter is a smooth manifold with projections to S and to S' . In fact, because π is non-singular on both S and S' , these two projections are covering maps. Moreover, each is trivial because π has the same degree on S

as it has on S' . Thus, both such projections have sections. In particular, there is a unique section over S whose restriction to γ composes with the projection to S' as the map ψ . The latter section thus composes with projection to S' so as to extend ψ as a diffeomorphism from S to S' . This extension obeys any given small κ version of (1–24) over S if ε is small and then κ' very small.

Step 5 Suppose now that γ is a circle in ∂U that lies very close to some point $z \in \Xi$ and let γ' denote its ψ image in C_0' . Let D denote the disk in C_0 that γ bounds and let $D' \subset C_0'$ denote the disk that γ' bounds. Since the whole of D' is not in $\psi(C_0 - U)$, its ϕ' image must lie very close to $\phi(z)$ and thus very close to $\phi(D)$. In fact, the distance between any point of $\phi'(D')$ and any point of $\phi(D)$ will be $\mathcal{O}(\varepsilon)$ when ε is small. As is argued in the subsequent steps, there are diffeomorphism between D and D' that extend ψ with small $r(\psi)$.

To see why ψ can be extended to map from D to D' with small $r(\psi)$, remark that when ε is small, then the results from Step 2 can be used to find a holomorphic coordinate, u , that is defined on the radius 4ε disk centered at z , and complex coordinates (x, y) centered at $\phi(z)$ with the following five properties: First, ϕ on D has the form

$$(5-3) \quad \phi(u) = (u^{p+1}, 0) + \mathcal{O}(|u|^{p+2}),$$

where p is a non-negative integer. Moreover,

$$(5-4) \quad \phi^* dx = (p + 1)u^p du + \mathcal{O}(|u|^{p+1}) \quad \text{and} \quad \phi^* dy = \mathcal{O}(|u|^{p+1}).$$

Second, the constant x disks and the $y = 0$ disk are pseudoholomorphic. Third, the forms ν and ν' from (5-2) span $T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$ over the domain of these coordinates. Finally, the $x = \text{constant}$ disks where $|u| \geq 2\delta$ contain the image of the fiber disks in the bundle N_1 from Step 1. Use B in what follows to denote the domain of the (x, y) coordinates.

Consider first the pull-back of x to $\phi^{-1}(B)$. Take ε small, and granted that $x = u^{p+1} + \mathcal{O}(|u|^{p+2})$, if ε is small, then the restriction of ∂x to the u coordinate chart is non-zero away from the origin. Moreover, with $W \subset C_0$ denoting the inverse image via x of the radius ε^{p+1} disk about 0 in \mathbb{C} , the map x sends W to the radius ε^{p+1} disk in \mathbb{C} as a degree $p + 1$ ramified cover with a single ramification point where ∂x vanishes with degree p . Note that W is a disk.

Let $W' \subset C_0'$ denote the inverse image via x of the same radius ε^{p+1} disk in \mathbb{C} . When ε and κ' are small, then W' is also a disk since its boundary is a small radius, embedded circle in D' . Since x is pulled up from W near the boundary of W' , it has

degree $p + 1$ there as a map to the radius ε^{p+1} circle in \mathbb{C} . Moreover, an appeal to the second point in Step 2 finds that

$$(5-5) \quad |\bar{\partial}x| < c_\varepsilon \kappa' |\partial x|$$

on the whole of W' where c_ε is determine once and for all by ε . This last equation implies that all zeros count with positive multiplicity. Note that all occur in the complement of the ψ image of the $|u| > 2\delta$ portion of W .

Here is the final remark for this step: According to the fourth point, x maps W' to \mathbb{C} as a degree $p+1$, ramified cover. As it turns out, the sum of the orders of vanishing of ∂x at its zeros in W' is equal to p , but this fact is not proved directly.

Step 6 The argument for a small $r(\psi)$ extension of ψ over W' is simplest in the case that ∂x on W' is zero at a single point and this is also the only zero of $\partial\psi$ in W' . In the latter case, define first a C^1 extension as follows: The function x on both W and W' has a p 'th root, this denoted by $x^{1/p}$. On both W and W' this p 'th root provides a C^1 homeomorphism onto the radius ε disk in \mathbb{C} centered at 0. This function is smooth and, in both cases, maps the complement of $x^{-1}(0)$ diffeomorphically to the complement of 0 in the centered, radius ε disk. The composition of the map $x^{1/p}$ from W with its inverse to W' thus defines a C^1 homeomorphism between W and W' that is smooth except at a single point. Let ψ_0 denote the latter. Then $|\bar{\partial}\psi_0| \ll |\partial\psi_0|$ when ε is small by appeal to (5-5). Moreover, $|\partial\psi_0|$ is uniformly positive while $\bar{\partial}\psi_0$ is zero at the one non-smooth point. This understood, a suitable perturbation of ψ_0 then gives a smooth diffeomorphism, ψ , with small $r(\psi)$.

Step 7 To proceed with the general case, consider that x , when viewed as a map from W' to \mathbb{C} , pulls back the complex structure from \mathbb{C} on the complement of the points where ∂x vanishes. Indeed, this follows from (5-5). This pull-back complex structure extends over the zeros of ∂x to define a complex structure on W' that makes x a holomorphic map.

As will now be explained, there is a holomorphic coordinate for this new complex structure on W' that makes x out to be a polynomial. To find such a coordinate, remark that near the boundary of W' , x is pulled up from W . In particular, when ε is small, x has a $(p + 1)$ 'st root near the boundary of W' that maps the boundary of W' in a 1-1 fashion to the radius ε disk in \mathbb{C} . This understood, let D_∞ denote the complement in $\mathbb{C} \cup \infty$ of the radius ε disk about the origin, and let τ denote a holomorphic coordinate on D_∞ that vanishes at 0 and has constant norm $1/\varepsilon$ on the boundary of D_∞ . Now let M denote the complex curve obtained from the disjoint union of D_∞ and W' by identifying the boundary of D_∞ with the boundary of W' by

pairing points with $1/\tau = x^{1/(p+1)}$. Of course, M is $\mathbb{C}\mathbb{P}^1$ with strangely presented holomorphic coordinate patches. The point of this construction is that x extends to the whole of M as a degree $p + 1$ holomorphic map from $\mathbb{C}\mathbb{P}^1$ as M to $\mathbb{C}\mathbb{P}^1$ as $\mathbb{C} \cup \infty$. Moreover, this extension has the property that the inverse image by x of the point ∞ is a single point, this the origin in D_∞ . Thus, with the complement of the origin in D_∞ viewed as $\mathbb{C} \subset M$, this extension of x is a polynomial of degree $p + 1$ when written with the standard holomorphic coordinate on \mathbb{C} .

Step 8 Granted that ε and κ' are small, the next step provides a holomorphic coordinate, τ , on \mathbb{C} such that

$$(5-6) \quad |x(\tau) - \tau^{p+1}| \leq \varepsilon^{p+2} \quad \text{and} \quad |x'(\tau) - (p+1)\tau^p| \leq \varepsilon^{p+1}$$

at all points where $|x| \geq (\frac{1}{2}\varepsilon)^{p+1}$. Here, x' denotes the τ -derivative of x .

To see what this brings, identify W' with its image in \mathbb{C} via M . On the $|x| \geq (\frac{1}{2}\varepsilon)^{p+1}$ portion of W' , the function x is pulled up from W via ψ^{-1} . On W , $|x - u^{p+1}| \leq c \cdot \varepsilon^{p+2}$ where c is a fixed constant. Thus the fact that $|x - \tau^{p+1}| < \varepsilon^{p+2}$ where $|x| \geq (\frac{1}{2}\varepsilon)^{p+1}$ implies that τ can be chosen so that

$$(5-7) \quad |u - \psi^* \tau| \leq \varepsilon^2 \quad \text{and} \quad |du - \psi^* d\tau| \leq \varepsilon$$

where $|x| \geq (\frac{1}{2}\varepsilon)^{p+1}$ on W .

With this in mind, fix a favorite smooth function $\beta: [0, \infty) \rightarrow [0, 1]$ with value 1 on $[0, \frac{5}{8}]$, value 0 on $[\frac{7}{8}, \infty)$, and with $|d\beta| < 8$. With β chosen, define $\lambda: W \rightarrow W'$ by setting

$$(5-8) \quad \lambda * \tau = (1 - \beta(|u|/\varepsilon))\psi^* \tau + \beta(|u|/\varepsilon)u.$$

Thus, λ extends ψ . Moreover, if ε is small, then the inequalities in (5-7) guarantee that λ is a diffeomorphism.

It remains now to explain why $r(\lambda)$ is very small if both ε and κ' are small. For this purpose, note that $r(\lambda)$ is uniformly $\mathcal{O}(\varepsilon)$ where $|x| \geq (\frac{1}{2}\varepsilon)^{p+1}$ since the differentials of λ and ψ differ there by $\mathcal{O}(\varepsilon)$. On this rest of W , the function $r(\lambda)$ is the ratio of $|\bar{\partial}\tau|$ to $|\partial\tau|$ where $\bar{\partial}$ and ∂ are defined by the restriction to W' of the almost complex structure from $\mathbb{R} \times (S^1 \times S^2)$ and τ is considered here a function on W' . Indeed, such is the case since the pull-back via λ of τ is a holomorphic function on W . To compute this ratio, remember that x on W' is a holomorphic function of τ so $\bar{\partial}x = x'\bar{\partial}\tau$ and $\partial x = x'\partial\tau$. Thus, $|\bar{\partial}\tau|/|\partial\tau| = |\bar{\partial}x|/|\partial x|$ and this is very small if ε and κ' are small.

Step 9 The claim that (5-7) holds when ε is small is an immediate consequence of the following lemma:

Lemma 5.3 Fix an integer $p \geq 0$, and $\varepsilon, \varepsilon' > 0$. There exists $\rho \in (0, \varepsilon)$ with the following property: Let f denote a non-trivial polynomial of degree $p + 1$ on \mathbb{C} such that the locus where $|f| = \rho$ is a simple closed curve. Then there is a holomorphic coordinate τ on \mathbb{C} such that $|f - \tau^{p+1}| \leq \varepsilon' \varepsilon^{p+1}$ and $|f' - (p + 1)\tau^p| \leq \varepsilon' \varepsilon^p$ where $|f| > (\frac{1}{2}\varepsilon)^{p+1}$.

Proof of Lemma 5.3 A degree $p + 1$ polynomial determines, up to a $(p + 1)$ 'st root of unity, a holomorphic coordinate, τ , for \mathbb{C} and a set, Λ , of $p + 1$ not necessarily distinct complex numbers such that

$$(5-9) \quad f(\tau) = \prod_{b \in \Lambda} (\tau - b) \text{ and such that } \sum_{b \in \Lambda} b = 0.$$

It follows from this representation of f that there exists a constant, c_0 , such that no point in Λ has absolute value greater than $d \equiv c_0 \rho^{1/(p+1)}$ if the $|f| = \rho$ locus is connected. Here c_0 , and constants $\{c_j\}_{1 \leq j \leq 5}$ that follow are independent of f .

Meanwhile, $|f(\tau)| \leq c_1 d^{p+1}$ where $|\tau| \leq 2d$. Thus, the locus where $|f(\tau)| \geq (\frac{1}{2}\varepsilon)^{p+1}$ must occur where $|\tau| > 2d$ in the case that $\rho \leq c_2 \varepsilon^{p+1}$. However, $|f(\tau) - \tau^{p+1}| \leq c_3 d |\tau|^p$ at the points $|\tau| \geq 2d$. Thus, $|f - \tau^{p+1}| \leq c_4 \rho^{1/(p+1)} \varepsilon^p$ where $|f| \geq (\frac{1}{2}\varepsilon)^{p+1}$. This understood, there is a fifth f -independent constant, $c_5 \in (0, 1/c_4)$ such that $\rho \leq (c_5 \varepsilon' \varepsilon)^{p+1}$ makes the lemma's claim true. \square

5.D The proof of Proposition 5.1

The proof of this proposition is almost verbatim that of [15, Proposition 2.13]. The following three parts of the proof focus on the salient differences.

Part 1 Suppose that (C_0, ϕ) defines a point in $\mathcal{S}_{B,c,\delta}$. This first part of the proof describes what turns out to be the part of a complete set of local coordinates for a neighborhood of this point in $\mathcal{S}_{B,c,\delta}$. To start, let $\text{Crit}(C) \subset C_0$ denote the set of size c whose elements are the critical points of θ where $\theta \in (0, \pi)$. If $(C_0', \phi') \in \mathcal{S}_{B,c,\delta}$ defines a point near to that of (C_0, ϕ) , then (C_0', ϕ') is represented by a small normed section in the (C_0, ϕ) version of $\text{kernel}(D_C)$. Such a representation allows the c critical points of θ on C_0' to be partnered with the points in $\text{Crit}(C)$ so that a C_0' critical point maps very close to the image of its partner in $\mathbb{R} \times (S^1 \times S^2)$. This pairing of critical points also preserves the conditions that are defined by the partition d . In addition, the degree of vanishing of $d\theta$ at any given critical point in C_0' is the same as that of its partner in C_0 .

Keeping these facts in mind, let $z \in \text{Crit}(C)$. Fix a small ball, B , centered on z 's image in $\mathbb{R} \times (S^1 \times S^2)$ whose closure excludes the images of any other point in $\text{Crit}(C)$.

Introduce the function, r , on B as defined in (2–9). If z' is the critical point in C_0' that is paired with z , then the pair $(\theta(z'), r(z'))$ is well defined and in this way, some $2c$ functions, $\{(\theta_z, r_z)\}_{z \in \text{Crit}(C)}$, are defined on a neighborhood of (C_0, ϕ_0) 's point in $\mathcal{S}_{B,c,\delta}$, at least in the case where the group G_C is trivial. In the case $G_C \neq \{1\}$, then these functions are distinguishable only modulo the action of G_C .

Note that the collection $\{\theta_z\}_{z \in \text{Crit}(C)}$ defines at most m independent functions on a neighborhood of (C_0, ϕ) 's point in $\mathcal{S}_{B,c,\delta}$.

Part 2 This part of the proof supplies additional coordinates for a neighborhood in $\mathcal{S}_{B,c,\delta}$ of the point defined by (C_0, ϕ) . To start, let $E \subset C_0$ denote either one of the N_+ convex side ends where $\lim_{|s| \rightarrow \infty} \theta$ is neither 0 nor π , or one of the ends that correspond to a 4-tuple in B . In any case, let (p, p') denote the integer pair from the corresponding 4-tuple. Then the $\mathbb{R}/(2\pi\mathbb{Z})$ valued function $p\varphi - p't$ has an \mathbb{R} valued lift on E with a well defined $|s| \rightarrow \infty$ limit. Now, as in Part 1, if (C_0', ϕ') defines a point in $\mathcal{S}_{B,c,\delta}$ near to that of (C_0, ϕ) , then the ends of C_0' can be partnered with those of C_0 so that partners share the same \hat{A} 4-tuple and map very near each other in $\mathbb{R} \times (S^1 \times S^2)$. This understood, the function $p\varphi - p't$ has an \mathbb{R} -valued lift on E 's partner in C_0' with a well defined $|s| \rightarrow \infty$ limit that is very close to the corresponding limit on E .

The assignment of these limits to the points near the image of (C_0, ϕ) define a collection of $N_+ + |B|$ functions (modulo the action of G_C) on a neighborhood of (C_0, ϕ) 's point in $\mathcal{S}_{B,c,\delta}$. Let $\{\varpi_{+1}, \dots\}$ denote the N_+ functions so defined (modulo the G_C action) from the convex side ends, and let $\{\varpi_{-1}, \dots\}$ denote the corresponding collection of $|B|$ functions that come from the 4-tuples in B .

Part 3 Two more functions are defined here for a neighborhood of (C_0, ϕ) 's point in $\mathcal{S}_{B,c,\delta}$. For this purpose, choose either an $(1, \dots)$ element from \hat{A} , or a $(0, -, \dots)$ 4-tuple that is not from B , or a point $z \in C_0$ where θ is zero. In all three cases, a complex valued function, ϖ' , is defined modulo the G_C action on a neighborhood of (C_0, ϕ) 's point in $\mathcal{S}_{B,c,\delta}$ as in the statement of [15, Proposition 2.13].

To summarize the story, the function is defined on the point defined by some pair (C_0', ϕ') by first identifying the latter with an element near zero in $\text{kernel}(D_C)$. This done, then if ϖ' is defined from an end of C_0 , there is a partnered end, $E' \subset C_0'$. If, as before, (p, p') denotes the integer pair from the corresponding element in \hat{A} , then the phase of the complex number ϖ' is proportional to the $|s| \rightarrow \infty$ limit of $p'\varphi - pt$ on E' . The absolute value of the complex number is proportional to the logarithm of the constant b that appears in the E' version of (2–4) in the case that E' corresponds

to a $(0, -, \dots)$ element. In the case that E' corresponds to a $(1, \dots)$ element form \widehat{A} , the absolute value is proportional to the logarithm of the constant \widehat{c} that appears in the E' version of (1–9).

When ϖ' is defined from a $\theta = 0$ point $z \in C_0$, the chosen $\text{kernel}(D_C)$ element for (C_0, ϕ) partners the $\theta = 0$ points in C_0' with those in C_0 so that the partner of z is mapped very close to z 's image in $\mathbb{R} \times (S^1 \times S^2)$ and makes the same local contribution as does z to the count for ζ_+ . This understood, ϖ' is assigned the value of a holomorphic coordinate for a small disk in the $\theta = 0$ cylinder that is centered on z 's image.

Part 4 Granted the preceding, introduce the vector space $K^* \subset \text{kernel}(D_C)$ as defined in [15, (2.23)]. The asserted structure of $\mathcal{S}_{B,c,d}$ near the point defined by (C_0, ϕ) follows via the implicit function theorem with a proof that $K^* = \{0\}$. The argument for this conclusion, and the argument that proves $K^* = \{0\}$ are essentially verbatim copies of the arguments given for [15, Proposition 2.13]. Nothing new is needed to treat the case where ϕ is not almost everywhere 1–1.

Note that the proof that $K^* = \{0\}$ proves more, for it establishes the following:

Lemma 5.4 *If G_C is trivial, then the collection $\{\varpi_{+1}, \dots\}$, $\{\varpi_{-1}, \dots\}$, ϖ' , $\{r_z\}_{z \in \text{Crit}(C)}$ and a certain subset of m functions from the collection $\{\theta_z\}_{z \in \text{Crit}(C)}$ define coordinates for $\mathcal{S}_{B,c,d}$ near the point defined by (C_0, ϕ) . Here, the m functions from the collection $\{\theta_z\}_{z \in \text{Crit}(C)}$ are chosen so that the values of the chosen versions of θ_z at (C_0, ϕ) account for the critical values of θ that avoid the $|s| \rightarrow \infty$ limits of θ on the concave side ends in C_0 and on the ends that correspond to 4–tuples from B . In the case that $G_C \neq \{1\}$, then the analogous collection of $N_+ + |B| + c + m$ functions give smooth orbifold coordinates on a neighborhood in $\mathcal{S}_{B,c,d}$ of the point defined by (C_0, ϕ) .*

6 Slicing the strata

This section describes the structure of each component of any given stratum from Proposition 5.1. To explain the point of view here, let $\mathcal{S}_{b,c,d}$ denote a given stratum and let \mathcal{S} denote a given component. The component \mathcal{S} is then mapped to the m 'th symmetric product of $(0, \pi)$ using the critical values of θ that do not coincide with angles from $\Lambda_{\widehat{A}}$. As is explained in Section 8, this map fibers \mathcal{S} over a certain m -dimensional simplex and so the structure of \mathcal{S} is determined by that of a typical fiber. The subsections that follow focus on the structure of such a fiber. The principle results

are in Theorems 6.2 and 6.3 and in Propositions 6.4 and 6.7. Theorems 6.2 and 6.3 are proved in Section 7. Sections 8 and 9 use the results from this section to fully paint the picture of \mathcal{S} as a fiber bundle.

6.A Graphs for the stratification

The constructions from Subsection 2.A and Part 3 of Subsection 2.C associate a graph, $T_{(\cdot)}$, to each pair (C_0, ϕ) from $\mathcal{M}^*_{\hat{A}}$. This graph has labeled edges and vertices that reflect the structure of the level sets of the function θ on the subvariety. As it turns out, certain aspects of these graphs are constant on any given component of any given strata in $\mathcal{M}^*_{\hat{A}}$ and serve to classify these components. This subsection describes in more detail the graphs that are involved and the manner in which they classify components of the stratification.

To start, remember that a graph, T , of the sort under consideration is contractible and has labeled vertices and labeled edges. What follows summarizes what is involved.

The vertex labels Each vertex in T is labeled in part by an angle in $[0, \pi]$ subject to various constraints, the first of which are as follows:

- (6–1) • *The two vertices on any given edge have distinct angles.*
 • *No multivalent vertex angle is extremal in the set of the angles of the vertices on the union of its incident edges.*

The vertices have additional labels. To elaborate, a subset of the angle 0 vertices are labeled via a 1–1 correspondence with the set of $(1, \dots)$ elements in \hat{A} . The remaining angle 0 vertices are labeled by positive integers that sum to ζ_+ . Likewise, a subset of angle π vertices are labeled via a 1–1 correspondence with the set of $(-1, \dots)$ elements in \hat{A} ; and the remainder are labeled by negative integers that sum to $-\zeta_-$.

Each vertex with angle in $(0, \pi)$ is labeled jointly by a subset of the $(0, \dots)$ elements in \hat{A} and a certain sort of graph. To describe these labels, remark first that distinct vertices are assigned disjoint subsets, and that the union of these subsets is the whole set of $(0, \dots)$ elements in \hat{A} . The empty subset can only be assigned to vertices with three or more incident edges. Meanwhile, a monovalent vertex must get a singleton set with a $(0, -, \dots)$ element. Finally, the integer pair component of any element from an assigned subset defines the corresponding vertex angle via (1–8). The subset of \hat{A} that is assigned to a vertex o is denoted in what follows by \hat{A}_o .

The graph that is assigned to the vertex o is denoted by $\underline{\Gamma}_o$. The latter graph is connected, has labeled vertices and oriented, labeled arcs. Here, the edges of $\underline{\Gamma}_o$ are

called ‘arcs’ to avoid confusing them with the edges in T that are incident to o . The first Betti number of $\underline{\Gamma}_o$ must be one less than the number of incident edges in T to o . In particular, this means that $\underline{\Gamma}_o$ is a single point when o is monovalent. In addition, each vertex in $\underline{\Gamma}_o$ has an even number of incident half-arcs with half oriented so as to point towards the vertex and half are oriented to point away.

The labeling of the vertices in $\underline{\Gamma}_o$ is as follows: Each vertex is labeled with an integer, subject to two constraints. First, the number of elements in \hat{A}_o that are identical to a given 4-tuple is equal to the number of vertices in $\underline{\Gamma}_o$ where the sign of the label is the second component of the 4-tuple and where the absolute value of the label is the greatest common divisor of the integer pair from the 4-tuple. Here is the second constraint: Vertices with label 0 must have four or more incident half-arcs.

Each arc in $\underline{\Gamma}_o$ is labeled by a pair of o ’s incident edges subject to constraints that will now be described. For this purpose, partition the incident edge set to o as $E_- \cup E_+$ where E_- contains the edges on which o is the largest angle vertex and E_+ those on which o is the smallest angle vertex. The label of any given arc must contain one edge from E_- and one from E_+ . To say more, let e denote an incident edge to o . The collection of arcs whose label contains e concatenate to define an oriented, immersed loop in $\underline{\Gamma}_o$ such that a traverse of this loop crosses no arc more than once. Thus, this loop is the image of an abstract, oriented, circular graph, ℓ_{oe} , via an immersion that maps vertices to vertices and is 1–1 and orientation preserving on the edges. The collection $\{\ell_{oe}\}_{e \text{ is incident to } o}$ is then constrained by the rules in (2–17). Here, and in what follows, ℓ_{oe} is used to denote both the abstract circular graph and its image in $\underline{\Gamma}_o$ since the former can be recovered from the latter. For reference below: Properties 1–4 from Part 3 of Subsection 2.C are valid here; Properties 1 and 2 are assumed, while Properties 3 and 4 then follow automatically.

The label of any given vertex in T determines an ordered pair of integers. When o denotes a vertex, then P_o or (p_o, p_o') is used to denote the associated integer pair. Here are the rules for this assignment:

- (6–2)
- If o is a monovalent vertex with an assigned integer m , then $P_o = (0, -m)$.
 - If o is a monovalent vertex with label $(\pm 1, \dots)$ or $(0, -, \dots)$, then P_o is the label’s integer pair component.
 - If o is a multivalent vertex with angle in $(0, \pi)$, then P_o is obtained by subtracting the sum of the integer pair components of the $(0, -, \dots)$ elements in \hat{A}_o from the sum of the integer pair components of the $(0, +, \dots)$ elements in \hat{A}_o .

The edge labels Each edge in T is labeled by a non-zero ordered pair of integers. When e denotes an edge, its corresponding pair is denoted as Q_e or as $(q_e, q_{e'})$. These labels are determined by the branching of T and the data from \hat{A} according to the rules that follow. Note that these rules determine the labels via an induction that starts with the edges whose minimal angle vertices are monovalent. Note as well that this induction uses the fact that T is a tree. Here are the rules:

- (6-3) • If e is incident to a monovalent vertex, o , then $Q_e = \pm P_o$ where the $+$ sign is used in the following cases:
- (a) The vertex label consists of either a $(1, -, \dots)$ element from \hat{A} , or a positive integer, or a $(-1, +, \dots)$ element from \hat{A}
 - (b) The vertex angle is in $(0, \pi)$ and it is the lesser of the two vertex angles from e .
- If o is a multivalent vertex, then $\sum_{e \in E_-} Q_e - \sum_{e \in E_+} Q_e = P_o$.

There is one additional constraint:

- (6-4) Let e denote a given edge, and let $\theta_- < \theta_+$ denote the angles of the vertices on e . The function $\alpha_{Q_e}(\theta) = (1 - 3 \cos^2 \theta)q_{e'} - \sqrt{6} \cos \theta q_e$ is positive on (θ_-, θ_+) and zero at an endpoint if and only if the angle is in $(0, \pi)$ and the corresponding vertex on e is monovalent.

As done for example in [15, Section 5], this last condition can be rephrased as inequalities that reference only the integer pairs from the vertex angles.

Graph isomorphisms Let T and T' denote graphs of the sort just described. An isomorphism, ι , from T to T' consists of a 1-1 and onto map from T 's vertex set to the vertex set of T' along with a compatible map from T 's edge set to the set of T' edges. Both maps must preserve all labels. Thus, the respective integer pairs of an edge in T and its ι image in T' agree; the respective angles of a vertex in T and its ι image in T' agree; and the labels of a vertex in T and its ι image in T' are themselves isomorphic in the following sense: First, $P_{\iota(o)} = P_o$ for all vertices $o \in T$. Second, there is an associated collection of isomorphisms that are labeled by T 's multivalent vertices where the version, $\hat{\iota}_o$, labeled by vertex o is an isomorphism from the graph $\underline{\Gamma}_o$ to $\underline{\Gamma}_{\iota(o)}$. Thus, $\hat{\iota}_o$ consists of a 1-1 and onto map from the set of vertices in $\underline{\Gamma}_o$ to the set of vertices in $\underline{\Gamma}_{\iota(o)}$ along with a compatible 1-1 and onto map from the set of oriented arcs in $\underline{\Gamma}_o$ to the corresponding set in $\underline{\Gamma}_{\iota(o)}$. Both of these set maps must also preserve labels. Thus,

- the vertex set map must preserve the integer labels of the vertices,

- the arc set map sends an arc with label (e, e') to one with label $(\iota(e), \iota(e'))$.

(6–5)

An isomorphism from T to itself is deemed an automorphism of T . The set of such automorphisms is designated as $\text{Aut}(T)$. Composition makes $\text{Aut}(T)$ into a group and it is viewed as such in what follows.

Graph homotopy It proves useful to introduce an equivalence between graphs that is weaker than graph isomorphism. For this purpose, graphs T and T' are said to be ‘homotopic’ when there is a one parameter family, $\{T_\tau\}_{\tau \in [0,1]}$ of graphs with the following two properties: First, $T_0 = T$ and $T_1 = T'$. Second, the various T_τ differ one from another only in their vertex angles, and these angles change as continuous functions of τ with the constraint that the number of distinct, multivalent vertex angles is independent of τ . In any event, only the vertex angles can change; neither edge labeled integer pairs nor vertex labeled graphs are modified in any way.

Subsection 2.A and Part 3 of Subsection 2.C explain how a graph such as T can be assigned to any given pair (C_0, ϕ) from any given element in $\mathcal{M}^*_{\hat{A}}$. As remarked at the end of Subsection 2.A, the isomorphism type of the assigned graph depends only on the given element in $\mathcal{M}^*_{\hat{A}}$. The following is a consequence:

Proposition 6.1 *Two elements in any given component of any given stratum have homotopic graphs, and elements in either distinct strata or in distinct components of the same stratum have graphs that are not homotopic.*

This proposition is proved in Subsection 8.A.

The subspace $\mathcal{M}^*_{\hat{A},T}$ Suppose that T is a graph of the sort just described and let $\mathcal{M}^*_{\hat{A},T}$ denote the subset of elements in $\mathcal{M}^*_{\hat{A}}$ that are defined by a pair whose graph is isomorphic to T . Here is why such a space is relevant: Let \mathcal{S} denote a component of some stratum $\mathcal{S}_{B,c,\vartheta}$ from Proposition 5.1. A map from \mathcal{S} to the m ’th symmetric power of the interval $(0, \pi)$ is defined by taking the distinct critical values of θ in $(0, \pi)$ that do not arise via (1–8) from the integer pair of either a $(0, +, \dots)$ element from \hat{A} or an element from B . The typical fiber of this map is a version of $\mathcal{M}^*_{\hat{A},T}$ for a graph T .

The rest of this section contains a description of $\mathcal{M}^*_{\hat{A},T}$.

6.B The structure of $\mathcal{M}^*_{\hat{A},T}$

By way of a preview for what is to come, Theorem 6.2 below asserts that $\mathcal{M}^*_{\hat{A},T}$, when non-empty, is diffeomorphic as an orbifold to a space of the form $\mathbb{R} \times O_T / \text{Aut}(T)$ where O_T is a version of (3–12) where $\text{Aut}(T)$ acts. As O_T is connected, this implies that $\mathcal{M}^*_{\hat{A},T}$ is a connected component of a stratum of $\mathcal{M}^*_{\hat{A}}$.

The detailed description of O_T requires some preliminary stage setting, and this is done next in four parts. The length of this stage setting preamble is due for the most part to the subtle nature of the $\text{Aut}(T)$ action.

Part 1 This part of the digression describes the relevant version of (3–12) leaving aside the $\text{Aut}(T)$ action. To start the story, let o denote a multivalent vertex in T and let θ_o denote its assigned angle. Meanwhile, let E_o denote the set of incident edges to o , and let $\text{Arc}(\underline{\Gamma}_o)$ denote the set whose elements are the arcs in $\underline{\Gamma}_o$. Now define Δ_o to be the set of maps $r: \text{Arc}(\underline{\Gamma}_o) \rightarrow (0, \infty)$ that obey

$$(6-6) \quad \sum_{\gamma \in \ell_{o,e}} r(\gamma) = 2\pi \alpha_{Q_e}(\theta_o)$$

for all multivalent vertices $o \in T$ and edges $e \in E_o$. Note that Δ_o is a simplex whose dimension is one less than the number of vertices on $\underline{\Gamma}_o$.

Each multivalent vertex in T is also assigned a real line. If o designates such a vertex, then \mathbb{R}_o is used to denote its associated line. Let \mathbb{R}_- denote an auxiliary copy of \mathbb{R} . Then $\mathbb{R}_- \times_o (\mathbb{R}_o \times \Delta_o)$ appears below in the role that (3–14) plays in (3–12). Here, and in what follows, the symbol \times_o appears with no accompanying definition signifies a product that is indexed by the set of multivalent vertices in T .

To define the group actions in the case at hand, it is necessary to first choose an $\text{Aut}(T)$ invariant vertex in T . To obtain such a vertex, note that any connected, non-empty, $\text{Aut}(T)$ invariant subgraph with the least number of vertices amongst all subgraphs of this sort is necessarily a 1–vertex graph. Here is why: Were such a subgraph to have two or more vertices, there would be one that was both monovalent as a subgraph vertex and not $\text{Aut}(T)$ invariant. Removing this vertex, the interior of its incident edge, and their orbits under $\text{Aut}(T)$ would result in a non-empty, proper, $\text{Aut}(T)$ invariant subgraph of the original.

Fix a smallest angle, $\text{Aut}(T)$ –invariant vertex and denote it as \diamond . This done, it proves convenient to introduce \mathcal{V} to denote the set of multivalent vertices in $T - \diamond$. Each $o \in \mathcal{V}$ labels a copy of \mathbb{Z} and these groups are going to act in a mutually commuting fashion on $\mathbb{R}_- \times (\times_{\hat{o}} \mathbb{R}_{\hat{o}})$. To describe these actions, note first that when o is a multivalent

vertex in $T - \diamond$, then $T - o$ has a unique component that contains the vertex \diamond . Let T_o denote the closure in T of the complement of this component.

Granted this notation, then $1 \in \mathbb{Z}_o$ is defined to act trivially on \mathbb{R}_- and on $\mathbb{R}_{\hat{o}}$ in the case that $\hat{o} \notin T_o$. In the case that $\hat{o} \in T_o$, then $1 \in \mathbb{Z}_o$ acts on $\mathbb{R}_{\hat{o}}$ as the translation

$$(6-7) \quad -2\pi \frac{\alpha_{Q_e}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})},$$

where e and \hat{e} designate the respective edges that connect o to $T - T_o$ and \hat{o} to $T - T_{\hat{o}}$.

Meanwhile, define an action of the group $\mathbb{Z} \times \mathbb{Z}$ on $\times_{\hat{o} \in \mathcal{V}} \mathbb{R}_{\hat{o}}$ by having an integer pair $N = (n, n')$ act as the translation by

$$(6-8) \quad -2\pi \frac{\alpha_N(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})}.$$

This action commutes with that just defined $\times_{o \in \mathcal{V}} \mathbb{Z}_o$.

The extension of the $\mathbb{Z} \times \mathbb{Z}$ action to \mathbb{R}_- and \mathbb{R}_{\diamond} requires an additional choice, this the choice of a distinguished $\text{Aut}(T)$ orbit in the set of incident edges to \diamond . Let \hat{E} denote this distinguished orbit, let $m_{\hat{E}}$ denote the number of edges in \hat{E} , and let $Q_{\hat{E}} = (q_{\hat{E}}, q_{\hat{E}'})$ denote the integer pair that is associated to the edges that comprise \hat{E} . The action of N on \mathbb{R}_{\diamond} is then defined by the version of (6-8) that uses \diamond for \hat{o} and $Q_{\hat{E}}$ for $Q_{\hat{e}}$. Meanwhile, the action of N on \mathbb{R}_- is defined so that (n, n') acts as the translation by $-2\pi m_{\hat{E}}(n'q_{\hat{E}} - nq_{\hat{E}'})$.

Granted all of the above, set

$$(6-9) \quad O_T \equiv (\times_o \Delta_o) \times [\mathbb{R}_- \times (\times_o \mathbb{R}_o)] / [(Z \times Z) \times (\times_{\hat{o} \in \mathcal{V}} \mathbb{Z}_{\hat{o}})].$$

By way of reminder, \times_o here designates a product that is indexed by the full set of multivalent vertices in T , while $\times_{\hat{o} \in \mathcal{V}}$ designates a product that is indexed by the set of vertices in $T - \diamond$. It is left to the reader to verify that O_T is a smooth manifold. This O_T is the desired version of (3-12).

Part 2 This part of the story constitutes a digression to provide a sort of inductive description of $\text{Aut}(T)$ using as components a subgroup of automorphisms of Γ_{\diamond} and a collection of cyclic groups that are labeled by the multivalent vertices in $T - \diamond$. To start this description, let o denote a vertex in \mathcal{V} . Then $T - o$ has a one component that contains \diamond . Let T_o denote the closure of the remaining components. This is a connected, contractible subgraph of T . The vertex o also has a distinguished incident edge, this the edge $e \equiv e(o)$ that connects o to $T - T_o$. Let ℓ_o to denote the loop $\ell_{oe(o)}$.

Now, let $\text{Aut}(T_o) \subset \text{Aut}(T)$ denote the subgroup that acts trivially on the data from $T - T_o$. This group fixes o and so it has a canonical homomorphism in the group of automorphisms of the pair $\{\underline{\Gamma}_o, E_o\}$, where E_o here designates the set of incident edges to o . Let Aut_o denote the image of $\text{Aut}(T_o)$ in this last group. As Aut_o fixes the incident edge $e(o)$, so it preserves the loop ℓ_o and thus has an image in the cyclic group of automorphisms of the abstract version of ℓ_o . In fact, Aut_o is isomorphic to its image in this cyclic group. Such is the case by virtue of the following fact:

(6–10) *An automorphism of $\underline{\Gamma}_o$ that fixes any arc must be trivial.*

Indeed, suppose that γ is fixed and that e' is one of γ 's labeling edges. Then the automorphism must act trivially on $\ell_{oe'}$ and so fixes all of its arcs. It then follows from Property 3 of Part 3 in Subsection 2.C that the maximal collection of fixed arcs must constitute the whole of $\underline{\Gamma}_o$.

Let n_o denote the order of the cyclic group Aut_o . Since the automorphism group of ℓ_o has a canonical generator, so does Aut_o , this smallest power of the generator of $\text{Aut}(\ell_o)$ that resides in Aut_o . This generator provides a canonical isomorphism between Aut_o and $\mathbb{Z}/(n_o\mathbb{Z})$.

The preceding implies an exact sequence

$$(6–11) \quad 1 \rightarrow \times_{o'} \text{Aut}(T_{o'}) \rightarrow \text{Aut}(T_o) \rightarrow \mathbb{Z}/(n_o\mathbb{Z}) \rightarrow 1,$$

where the product in (6–11) is indexed by the vertices in T_o that share some edge with o . There is a similar exact sequence in the case that $o = \diamond$ but where $\mathbb{Z}/(n_o\mathbb{Z})$ is replaced by Aut_\diamond . Of course, if Aut_\diamond fixes an edge in E_\diamond , then $\text{Aut}_\diamond = \mathbb{Z}/(n_\diamond\mathbb{Z})$.

With regards to the various versions of Aut_o , keep in mind the following:

(6–12) *An element in Aut_o is trivial if it fixes more than two edges in E_o .*

To see why such is the case, first construct a closed, 2–dimensional cell complex by taking the 1–skeleton to be $\underline{\Gamma}_o$ and then attaching a disk to each $\ell_{o(\cdot)}$. An Euler class computation finds that this complex is a 2–sphere. The Aut_o action on $\underline{\Gamma}_o$ extends to the action on the sphere if it is agreed that its action on the 2–cells is obtained by a linear extension. This extended action is then a piecewise linear, orientation preserving action. In particular, if an edge is fixed by some element, the action on the corresponding disk is a rotation through a rational fraction of 2π and has the origin as its fixed point. Since $\text{Aut}(T)$ is finite, such an action can have at most two fixed points unless it is trivial.

The last point to make in this part of the subsection is that the various versions of (6–11) all split and these splitting can be used to write $\text{Aut}(T)$ as the iterated semi-direct

product

$$(6-13) \quad \text{Aut}(T) \approx \text{Aut}_\diamond \times \left[\times_o \left[\mathbb{Z}/(n_o\mathbb{Z}) \times \left[\times_{o'} \left[\mathbb{Z}/(n_{o'}\mathbb{Z}) \times \cdots \right] \cdots \right] \right] \right].$$

Here, the left most subscripted product is indexed by the vertices that share edges with \diamond ; meanwhile, any subscripted product, $\times_{\hat{o}}$, that appears as $\mathbb{Z}/(n_{\hat{o}}\mathbb{Z}) \times [\times_{\hat{o}'} \cdots]$ is indexed by the vertices in $T_{\hat{o}}$ that share an edge with \hat{o} . Moreover, such an isomorphism identifies any given $o \in T - \diamond$ version of $\text{Aut}(T_o)$ with the semi-direct product

$$(6-14) \quad \mathbb{Z}/(n_o\mathbb{Z}) \times \left[\times_{o'} \left[\mathbb{Z}/(n_{o'}\mathbb{Z}) \times \cdots \right] \cdots \right].$$

It is important to keep in mind that the isomorphism in (6-13) is not canonical. Rather, it requires the choice of a ‘distinguished’ vertex in each $o \neq \diamond$ version of ℓ_o subject to the following constraint: Let v_o denote the distinguished vertex for o and let $\iota \in \text{Aut}(T)$. Then v_o and $v_{\iota(o)}$ define the same $\text{Aut}(T)$ orbit in the set of vertices in $\cup_{\hat{o} \neq \diamond} \ell_{\hat{o}}$. Assume in what follows that such choices have been made.

To explain how (6-13) arises, note that the effect of an automorphism on any $o \neq \diamond$ version of Γ_o is determined completely by its affect on a single vertex in ℓ_o , thus on v_o . This follows by virtue of the fact that any given $\iota \in \text{Aut}(T)$ must map ℓ_o to $\ell_{\iota(o)}$ so as to preserve the cyclic order of the vertices. Granted this, a lift to $\text{Aut}(T)$ of any element in any given version of Aut_o can be defined in an iterated fashion along the following lines: Let ι denote the given element. The lift is defined to live in $\text{Aut}(T_o)$. The first step to defining this lift specifies its action on $\cup_{o'} \Gamma_{o'}$, where the union is labeled by the vertices in $T_o - o$ that share edges with o . Let o' denote such a vertex. Then $\iota(o')$ is determined apriori by the action of ι on o ’s incident edge set. The corresponding isomorphism from $\Gamma_{o'}$ to $\Gamma_{\iota(o')}$ is defined by requiring that $v_{o'}$ go to $v_{\iota(o')}$. Note that this now determines how ι acts on the set $\cup_{o'} E_{o'}$ whose elements are the incident edges to the various vertices in $T_o - o$ that share an edge with o . To continue, suppose that $o'' \in T_{o'} - o'$ shares an edge with o' . Then $\iota(o'')$ is determined apriori by the just defined action of ι on $\cup_{o'} E_{o'}$. The isomorphism between $\Gamma_{o''}$ and $\Gamma_{\iota(o'')}$ is again determined by the requirement that $\iota(v_{o''}) = v_{\iota(o'')}$. Continuing in this vein from a given vertex $\hat{o} \in T_o$ to those in $T_{\hat{o}} - \hat{o}$ that share edges with \hat{o} ends with an unambiguous definition for the lift of ι . In this regard, note that the lifts of elements ι and ι' from a given Aut_o must multiply to give the lift of $\iota \cdot \iota'$. Indeed, this is guaranteed by the afore-mentioned fact that the affect of an automorphism on any given $o \neq \diamond$ version of Γ_o is determined by its affect on v_o .

Part 3 The desired action of $\text{Aut}(T)$ on O_T is described by first introducing a certain central extension of $\text{Aut}(T)$ by $\mathbb{Z} \times \mathbb{Z}$ and describing an action of this extension on

$$(6-15) \quad (\times_o \Delta_o) \times (\mathbb{R}_- \times \mathbb{R}_\diamond) \times \left[\times_{\hat{o} \in \mathcal{V}} \mathbb{R}_{\hat{o}} \right] / \left[\times_{\hat{o}' \in \mathcal{V}} \mathbb{Z}_{\hat{o}'} \right].$$

Here, the symbol $\times_o \Delta_o$ designates the product of the simplices that are indexed by the multivalent vertices in T . The desired extension of $\text{Aut}(T)$ is denoted in what follows by $\widehat{\text{Aut}}(T)$. It is induced from a $\mathbb{Z} \times \mathbb{Z}$ central extension of Aut_\diamond in the following manner: Let $\widehat{\text{Aut}}_\diamond$ denote the extension of Aut_\diamond . Then $\widehat{\text{Aut}}(T)$ is the set of pairs $(\iota, g) \in \widehat{\text{Aut}}_\diamond \times \text{Aut}(T)$ that have the same image in Aut_\diamond . Note that the isomorphism given in (6–13) is covered by an analogous $\widehat{\text{Aut}}(T)$ version that has $\widehat{\text{Aut}}(T)$ on the left hand side and $\widehat{\text{Aut}}_\diamond$ instead of Aut_\diamond on the right.

To define $\widehat{\text{Aut}}_\diamond$, it is necessary to reintroduce the ‘blow up’ from Part 6 of Subsection 2.C of $\underline{\Gamma}_\diamond$. For this purpose, note that the any given $\underline{\Gamma}_o^*$ is available if Properties 1–4 from Part 3 of Subsection 2.C are satisfied. As observed in Subsection 6.A, these properties are present; thus any such $\underline{\Gamma}_o^*$ is well defined.

Granted the preceding, introduce $\underline{\Gamma}_\diamond^*$. Likewise, reintroduce the cohomology class $\phi_\diamond \in H^1(\underline{\Gamma}_\diamond^*; \mathbb{Z} \times \mathbb{Z})$ from this same Part 6 of Subsection 2.C. Use ϕ_\diamond to define a $\mathbb{Z} \times \mathbb{Z}$ covering space over $\underline{\Gamma}_\diamond^*$, this denoted in what follows by $\bar{\Gamma}^*$. The group of deck transformations of $\bar{\Gamma}^*$ is the group $\mathbb{Z} \times \mathbb{Z}$. This understood, the group Aut_\diamond has a central, $\mathbb{Z} \times \mathbb{Z}$ extension that acts as a group of automorphisms of the graph $\bar{\Gamma}^*$. The latter group is $\widehat{\text{Aut}}_\diamond$.

The action of $\widehat{\text{Aut}}(T)$ on the space in (6–15) is defined in the five steps that follow. In this regard, the first three steps reinterpret the various factors in (6–15).

Step 1 This first step reinterprets the factor $\mathbb{R}_\diamond \times \Delta_\diamond$. For this purpose, let $\text{Vert}_{\hat{E}}$ denote the set of vertices in $\bar{\Gamma}^*$ that project to some $e \in \hat{E}$ version of $\ell_{\diamond e}$. The plan is to interpret the product $\mathbb{R}_\diamond \times \Delta_\diamond$ as a fiber bundle over Δ_\diamond , this denoted by \mathbb{R}^Δ . In particular, \mathbb{R}^Δ is the linear subspace in $\text{Maps}(\text{Vert}_{\hat{E}}; \mathbb{R}) \times \Delta_\diamond$ where the pair (τ, r) obeys the following constraint: Let v and v' denote a pair of vertices from $\text{Vert}_{\hat{E}}$ and let $N \in \mathbb{Z} \times \mathbb{Z}$ denote any element that maps v to the component of $\bar{\Gamma}^*$ that contains v' . Then

$$(6-16) \quad \tau(v') = \tau(v) + 2\pi \frac{1}{\alpha_{Q_{\hat{E}}}(\theta_\diamond)} \left(\sum_{\gamma} \pm \hat{r}(\gamma) - 2\pi \alpha_N(\theta_\diamond) \right),$$

where the notation is as follows: First, the sum is over the arcs in $\bar{\Gamma}^*$ whose concatenated union defines a path that takes Nv to v' . Second, $\hat{r}(\gamma) = r(\gamma)$ if γ projects to an arc in $\underline{\Gamma}_\diamond$; if not, then $\hat{r}(\gamma) = 0$. Third, the $+$ sign appears when the arc is traversed in its oriented direction and the minus sign is used when the arc is traversed opposite to its orientation. Note that (6–6) guarantees that the expression on the right hand side of (6–16) is independent of the precise choice for N or for the path in question granted the given constraints on both.

Step 2 This step reinterprets the \mathbb{R}_- factor that appears in (6–15). For this purpose, suppose that e is an incident edge of \diamond , and then reintroduce the lift $\ell^*_{\diamond,e} \subset \Gamma^*_{\diamond}$ of the loop $\ell_{\diamond,e}$. Let L_e denote the set of components of the inverse image of $\ell^*_{\diamond,e}$ in $\bar{\Gamma}^*$. In this regard, each component of this inverse image is a linear subgraph of $\bar{\Gamma}^*$. Note that the deck transformation group $\mathbb{Z} \times \mathbb{Z}$ has a transitive action on L_e whereby the stabilizer of any given element is $\mathbb{Z} \cdot Q_e$. Let $\Lambda \equiv \times_{e \in \hat{E}} L_e$.

The plan is to identify \mathbb{R}_- as a certain subspace in $\text{Maps}(\Lambda; \mathbb{R})$. In particular, a map, τ_- , is in \mathbb{R}_- when the following is true: Given $L \in \Lambda$, and then $N = (n, n') \in \mathbb{Z} \times \mathbb{Z}$ and $e \in \hat{E}$, let L' denote the element in Λ that is obtained from L by acting on e 's entry by N . Then

$$(6-17) \quad \tau_-(L') = \tau_-(L) - 2\pi(n'q_{\hat{E}} - nq_{\hat{E}}').$$

The conditions depicted here define a 1–dimensional affine line in $\text{Maps}(\Lambda; \mathbb{R})$. The latter is denoted in what follows as \mathbb{R}^- .

Step 3 This step reinterprets the factor $[\times_{o \in \mathcal{V}} \mathbb{R}_o] / [\times_{o \in \mathcal{V}} \mathbb{Z}_o]$ in (6–15). For this purpose, suppose that e is incident to \diamond and introduce $\mathcal{V}(e)$ to denote the set of vertices in the component of $T - \diamond$ that contains $e - \diamond$. Set $W_e \equiv [\times_{\hat{o} \in \mathcal{V}(e)} \mathbb{R}_{\hat{o}}] / [\times_{\hat{o} \in \mathcal{V}(e)} \mathbb{Z}_{\hat{o}}]$ and so write

$$(6-18) \quad [\times_{o \in \mathcal{V}} \mathbb{R}_o] / [\times_{o \in \mathcal{V}} \mathbb{Z}_o] = \times_{e \in E_{\diamond}} W_e.$$

Now let $\hat{U}_e \subset \text{Maps}(L_e; W_e)$ denote the subspace of maps $x: L_e \rightarrow W_e$ that have the following property: If $\ell \in L_e$ and $N \in \mathbb{Z} \times \mathbb{Z}$, then $x(N \cdot \ell)$ and $x(\ell)$ have respective lifts to $\times_{\hat{o} \in \mathcal{V}(e)} \mathbb{R}_{\hat{o}}$ whose coordinates obey

$$(6-19) \quad x(N \cdot \ell)_{\hat{o}} = x(\ell)_{\hat{o}} - 2\pi \frac{\alpha_N(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})}$$

for each $\hat{o} \in \mathcal{V}(e)$.

Note that this condition is well defined even though the integer multiples of Q_e act trivially on L_e . Indeed, such is the case because $\mathbb{Z} \cdot Q_e$ also acts trivially on W_e . Since the action on L_e of $\mathbb{Z} \times \mathbb{Z}$ is transitive, the space \hat{U}_e is diffeomorphic to W_e .

Step 4 The space depicted in (6–15) is diffeomorphic to

$$(6-20) \quad \mathbb{R}^- \times \mathbb{R}^{\Delta} \times \left[\times_{e \in E_{\diamond}} (\hat{U}_e \times (\times_{\hat{o} \in \mathcal{V}(e)} \Delta_{\hat{o}})) \right].$$

The group $\hat{\text{Aut}}(T)$ will act on this version of (6–15). This step describes the action of the subgroups $\{\mathbb{Z}/(n_o\mathbb{Z})\}_{o \in \mathcal{V}}$ that appear in the semi-direct decomposition as depicted in the $\hat{\text{Aut}}(T)$ version of (6–13).

The definition of these actions requires an additional set of choices to be made for each vertex in \mathcal{V} : A ‘distinguished’ edge must be designated from each non-trivial $\mathbb{Z}/(n_o\mathbb{Z})$ orbit in E_o for each $o \in \mathcal{V}$. To explain, note first that (6–12) has the following consequence: The $\mathbb{Z}/(n_o\mathbb{Z})$ action on an orbit in E_o is either free or trivial. Thus, any non-trivial orbit has n_o elements and a canonical cyclic ordering. The choice of a distinguished element provides a compatible linear ordering with the distinguished element at the end. These distinguished edges should be chosen in a compatible fashion with the $\text{Aut}(T)$ action. This is to say that when $o \in \mathcal{V}$ and $e' \in E_o$ is the distinguished edge in its Aut_o orbit, then the following is true: Let $\iota \in \text{Aut}(T)$ denote an automorphism that sends the distinguished vertex in Γ_o to that in $\Gamma_{\iota(o)}$. Then $\iota(e')$ is the distinguished edge in its $\text{Aut}_{\iota(o)}$ orbit.

Let $\iota \in \widehat{\text{Aut}}(T)$ now denote the generator of the $\mathbb{Z}/(n_o\mathbb{Z})$ subgroup. This element acts trivially on \mathbb{R}^- , \mathbb{R}^Δ . Its action is also trivial on $\widehat{U}_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ unless $\mathcal{V}(e)$ contains o .

The affect of ι on the relevant version of $\widehat{U}_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ is defined from a certain action of ι on $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$. In this regard, note that ι_o has the action on $\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}}$ that sends a map $r: \text{Arc}(\Gamma_{\widehat{\delta}}) \rightarrow (0, \infty)$ to the map $\iota \cdot r: \text{Arc}(\Gamma_{\iota(\widehat{\delta})}) \rightarrow (0, \infty)$ that obeys

$$(6-21) \quad (\iota \cdot r)(\iota(\gamma)) = r(\gamma) \text{ for all } \gamma \in \text{Arc}(\Gamma_{\widehat{\delta}}).$$

The action that is described momentarily on $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ is intertwined by the projection to $\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}}$ with the action that is depicted by (6–21). In any event, an action of $\mathbb{Z}/(n_o\mathbb{Z})$ on $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ induces one on $\text{Maps}(L_e; W_e) \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ by composition, thus ι sends a pair (x, r) with $x: L_e \rightarrow W_e$ and $r \in (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ to the pair whose second component is $\iota \cdot r$ and whose first component is obtained from x by composing with ι 's affect on the W_e factor in $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$. As it turns out, this action of ι preserves the relation in (6–19) and so the induced action on $\text{Maps}(L_e; W_e) \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ induces the required action on $\widehat{U}_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$.

To define the action of ι on $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$ consider its affect on the image of some given point (τ, r) where $\tau \in \times_{\widehat{\delta} \in \mathcal{V}(e)} \mathbb{R}_{\widehat{\delta}}$. For this purpose, use $\tau_{\widehat{\delta}}$ to denote the $\mathbb{R}_{\widehat{\delta}}$ coordinate of τ , and use $r_{\widehat{\delta}}$ to denote the $\Delta_{\widehat{\delta}}$ coordinate. Also, use $[\tau, r]$ to denote the image point of (τ, r) in $W_e \times (\times_{\widehat{\delta} \in \mathcal{V}(e)} \Delta_{\widehat{\delta}})$. The point $\iota \cdot [\tau, r]$ has a lift, $(\tau', \iota \cdot r)$ with $\tau'_{\widehat{\delta}} = \tau_{\widehat{\delta}}$ unless $\widehat{\delta} \in T_o$. Meanwhile,

$$(6-22) \quad \tau'_o = \tau_o - 2\pi \frac{1}{\alpha_{Q_{e(o)}}(\theta_o)} \sum_{\gamma} r(\gamma),$$

where $e(o)$ here designates the edge that connects o to $T - T_o$ and where the sum is over the set of arcs on the oriented path in Γ_o between $\iota^{-1}(v_o)$ and v_o .

Consider next the case that $\hat{o} \in T_o - o$ and that \hat{o} 's component of $T_o - o$ is fixed by ι_o . This is to say that \hat{o} 's component is connected to o by an edge that is fixed by the $\mathbb{Z}/(n_o\mathbb{Z})$ action on o 's incident edge set. Let $e'(o)$ denote the latter edge. Then

$$(6-23) \quad \tau'_{\hat{o}} = \tau_{\hat{o}} - 2\pi \frac{1}{n_o} \frac{\alpha_{Q_{e(o)} - Q_{e'(o)}}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})},$$

where \hat{e} here denotes the edge that connects \hat{o} to $T - T_{\hat{o}}$.

Finally, suppose that \hat{o} is in a component of $T_o - o$ that is not fixed by ι . In this case,

$$(6-24) \quad \tau'_{\iota(\hat{o})} = \tau_{\hat{o}} - \varepsilon_{\hat{o}} 2\pi \frac{\alpha_{Q_{e(o)}}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})},$$

where $\varepsilon_{\hat{o}} = 0$ unless the edge that connects \hat{o} 's component of $T_o - o$ to o is a 'distinguished' edge in its $\mathbb{Z}/(n_o\mathbb{Z})$ orbit; in the latter case, $\varepsilon_{\hat{o}} = 1$.

It is left for the reader to verify that action just defined for ι on $W_e \times (\times_{\hat{o} \in \mathcal{V}(e)} \Delta_{\hat{o}})$ has ι^{n_o} acting as the identity. The reader is also asked to verify that the suite of these actions as o varies through \mathcal{V} defines compatible actions of the various versions of (6-14) on the space in (6-20).

Step 5 This step explains how the $\hat{\text{Aut}}_{\diamond}$ subgroup of $\hat{\text{Aut}}(T)$ acts on the space in (6-20). The explanation starts by describing the action on the factor \mathbb{R}^{Δ} . In this regard, keep in mind that \mathbb{R}^{Δ} is a real line bundle over Δ_{\diamond} and that there is an Aut_{\diamond} action on Δ_{\diamond} that has any given $\iota \in \text{Aut}_{\diamond}$ acting to send a map $r: \text{Arc}(\Gamma_{\diamond}) \rightarrow (0, \infty)$ to the map $\iota \cdot r$ that is given by the $\hat{o} = \diamond$ version of (6-21). The projection map from \mathbb{R}^{Δ} to Δ_{\diamond} will intertwine the desired action of $\hat{\text{Aut}}_{\diamond}$ with that of Aut_{\diamond} on Δ_{\diamond} . In any event, the action on \mathbb{R}^{Δ} is induced by the action on $\text{Maps}(\text{Vert}_{\hat{E}}; \mathbb{R}) \times \Delta_{\diamond}$ that has $\iota \in \hat{\text{Aut}}_{\diamond}$ sending a pair (τ, r) to the pair $(\iota \cdot \tau, \iota \cdot r)$ where $\iota \cdot \tau$ is defined by setting $(\iota \cdot \tau)(\iota(v)) \equiv \tau(v)$ for all $v \in \text{Vert}_{\hat{E}}$. This action preserves the relation in (6-16) so restricts to define an action of $\hat{\text{Aut}}_{\diamond}$ on \mathbb{R}^{Δ} .

To continue, consider next the action of $\hat{\text{Aut}}_{\diamond}$ on \mathbb{R}^{-} . The action in this case is induced from the action on $\text{Maps}(\Lambda, \mathbb{R})$ that has $\iota \in \hat{\text{Aut}}_{\diamond}$ sending a given map τ_{-} to the map $\iota \cdot \tau_{-}$ whose value on any given $\iota(L)$ is that of τ_{-} on L .

The final point is that of the $\hat{\text{Aut}}_{\diamond}$ action on the bracketed factor that appears in (6-20). To set the stage for the discussion, note that Aut_{\diamond} acts on the simplex product. Here, the affect of $\iota \in \text{Aut}_{\diamond}$ on any given $\hat{o} \in \mathcal{V}$ version of Δ_o sends $r \in \Delta_o$ to the map $\iota \cdot r \in \Delta_{\iota(o)}$ whose values are given by the rule in (6-21). The projection to $\times_{\hat{o} \in \mathcal{V}} \Delta_{\hat{o}}$ from the bracketed factor in (6-20) will intertwine the desired $\hat{\text{Aut}}_{\diamond}$ action with that

just described on $\times_{\hat{\delta} \in \mathcal{V}} \Delta_{\hat{\delta}}$. The action of $\widehat{\text{Aut}}_{\diamond}$ on the bracketed factor in (6–20) is induced by an action on

$$(6-25) \quad \times_{e \in E_{\diamond}} \left(\text{Maps}(L_e; W_e) \times \left(\times_{\hat{\delta} \in \mathcal{V}(e)} \Delta_{\hat{\delta}} \right) \right).$$

To describe the latter action, let w denote a point in (6–25), thus a tuple whose coordinates are indexed by \diamond 's incident edges. The coordinate with label $e \in E_{\diamond}$ consists of a map, x_e , from L_e to W_e together with a maps, $\{r_{\hat{\delta}}: \text{Arc}(\Gamma_{\hat{\delta}}) \rightarrow (0, \infty)\}_{\hat{\delta} \in \mathcal{V}(e)}$. Here, x_e has lifts as a map of L_e to $\times_{\hat{\delta} \in \mathcal{V}(e)} \mathbb{R}_{\hat{\delta}}$ that assigns a real number to each $(\ell, \hat{\delta}) \in L_e \times \mathcal{V}(e)$. Fix such a lift and use $\tau_{\hat{\delta}}(\ell) \in \mathbb{R}$ in what follows to denote the lift's assignment to a given $(\ell, \hat{\delta})$.

With the preceding understood, a given $\iota \in \widehat{\text{Aut}}_{\diamond}$ sends w to the tuple $\iota \cdot w$ whose coordinate with label e is denoted by $x'_{(\cdot)}$. The collection $\{x'_{(\cdot)}\}$ thus defines $\iota \cdot w$ given the intertwining requirement vis a vis the action of Aut_{\diamond} on $\times_{\hat{\delta}} \Delta_{\hat{\delta}}$. This understood, the various versions of $x'_{(\cdot)}$ are defined by requiring that any given version of $x'_{\iota(e)}$ lift so as to assign the real number $\tau_{\hat{\delta}}(\ell)$ to the pair $(\iota(\ell), \iota(\hat{\delta}))$ for all pairs $(\ell, \hat{\delta}) \in L_e \times \mathcal{V}(e)$.

It is left as an exercise with the definitions to verify that the rule just given specifies an action of $\widehat{\text{Aut}}_{\diamond}$ on (6–25) that preserves the subspace from the bracketed term in (6–20).

Granted that $\widehat{\text{Aut}}_{\diamond}$ acts as just described on the space in (6–20), it is a straightforward task to verify that the action is compatible vis-à-vis the $\widehat{\text{Aut}}(T)$ version of (6–13) with those defined in Step 4 of the various versions of (6–14). Such being the case, then these various actions define an action of $\widehat{\text{Aut}}(T)$ on the space in (6–20). This is the desired action.

Part 4 Introduce as notation O^*_T to denote the space depicted in (6–20). Then the space O_T is diffeomorphic to the quotient of O^*_T by the $\mathbb{Z} \times \mathbb{Z}$ subgroup in $\widehat{\text{Aut}}(T)$ that maps to the identity in $\text{Aut}(T)$. The group $\text{Aut}(T)$ now acts on O_T with the latter now viewed as $O^*_T/(\mathbb{Z} \times \mathbb{Z})$. Let $\widehat{O}_T \subset O_T$ denote the set of points where the $\text{Aut}(T)$ stabilizer is the identity. Thus, \widehat{O}_T is the $\mathbb{Z} \times \mathbb{Z}$ quotient of the points in (6–20) with trivial $\widehat{\text{Aut}}(T)$ stabilizer. In any event, $\widehat{O}_T/\text{Aut}(T)$ is a smooth manifold whose dimension is equal to one more than the number of vertices in $\cup_o \Gamma_o$; the union here is indexed by multivalent vertices in T . Meanwhile, $O_T/\text{Aut}(T)$ is a smooth orbifold.

With the preceding understood, consider:

Theorem 6.2 *If non-empty, $\mathcal{M}^*_{\widehat{A}, T}$ is diffeomorphic as an orbifold to $\mathbb{R} \times O_T/\text{Aut}(T)$. In particular the complement in $\mathcal{M}^*_{\widehat{A}, T}$ of \mathcal{R} is diffeomorphic to $\mathbb{R} \times \widehat{O}_T/\text{Aut}(T)$.*

There is an analogy in this case with the story told in Proposition 3.6 as there are orbifold diffeomorphisms from $\mathcal{M}^*_{\hat{A},T}$ that give a direct geometric interpretation to the various factors that enter the definition of $\mathbb{R} \times O_T / \text{Aut}(T)$. An elaboration requires a short digression for two new notions.

The first notion is that of the space $\mathcal{M}^*_{\hat{A},T}^\Lambda$ whose elements consist of equivalence classes of 4-tuples (C_0, ϕ, T_C) where (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A},T}$ while T_C is a fixed correspondence in (C_0, ϕ) . The equivalence relation equates $(C_0, \phi', T_{C'})$ with (C_0, ϕ, T_C) in the case that there is a holomorphic diffeomorphism $\psi: C_0 \rightarrow C_0$ such that $\phi' = \phi \circ \psi$, and such that $T_{C'}$ is obtained from T_C as follows: If T_C identifies a given component $K \subset C_0 - \Gamma$ with an edge $e \in T$, then $T_{C'}$ identifies $\psi^{-1}(K)$ with e ; and if T_C identifies any given arc $\gamma \subset \Gamma$ with an arc, γ_* , in a version of $\Gamma_{(\cdot)}$ from T , then $T_{C'}$ identifies $\psi^{-1}(\gamma)$ with γ_* .

Note that $\mathcal{M}^*_{\hat{A},T}^\Lambda$ is a smooth manifold since the group G_C of holomorphic diffeomorphisms of C_0 that fix ϕ acts freely on the set of correspondences of T in (C_0, ϕ) . Note in addition that the tautological map from $\mathcal{M}^*_{\hat{A},T}^\Lambda$ to $\mathcal{M}^*_{\hat{A},T}$ restricts over $\mathcal{M}_{\hat{A},T}$ as a covering space map with fiber $\text{Aut}(T)$.

To describe the second notion, fix a multivalent vertex $o \in T$ and a vertex $v \in \Gamma_o$ with non-zero integer label. In what follows, use (\hat{p}_o, \hat{p}'_o) to denote the relatively prime integer pair that defines the angle θ_o via (1–8). Define a map, $\Psi_v: O_T \rightarrow \mathbb{R}/(2\pi\mathbb{Z})$ as follows: Map $\mathbb{R}_o \times \Delta_o$ to $\mathbb{R}/2\pi\mathbb{Z}$ by the rule that sends $v \in \mathbb{R}_o$ and $r \in \Delta_o$ to

$$(6-26) \quad (\hat{p}_o q_{e'} - \hat{p}'_o q_e)v + \frac{(\hat{p}_o^2 + \hat{p}'_o'^2 \sin^2(\theta_o))^{1/2}}{(1 + 3 \cos^4(\theta_o))^{1/2}} \sum_{\gamma} \pm r(\gamma) \pmod{(2\pi\mathbb{Z})}.$$

Here, the notation is as follows: First, e is the distinguished edge for o . Second, the sum is indexed by the arcs in any ordered set of arcs from Γ_o that concatenate end to end so as to define a path that starts at Γ_o 's distinguished vertex and ends at v . Finally, the $+$ sign in this sum is taken if and only if the indicated arc is traversed in its oriented direction on the path to v . By virtue of (6–6), this map is insensitive to the choice for the concatenating set of arcs. As the map in (6–26) is also insensitive to the group actions that define O_T , so it descends as a map from O_T to $\mathbb{R}/2\pi\mathbb{Z}$. The latter is the map Ψ_v .

Theorem 6.3 *There are orbifold diffeomorphisms from $\mathcal{M}^*_{\hat{A},T}$ to $\mathbb{R} \times O_T / \text{Aut}(T)$ with the following properties:*

- *The projection to the \mathbb{R} factor intertwines \mathbb{R} 's action on $\mathcal{M}^*_{\hat{A},T}$ with its translation action on \mathbb{R} .*

- The orbifold diffeomorphism is covered by a diffeomorphism from $\mathcal{M}_{\hat{A},T}^\Lambda$ to $\mathbb{R} \times O_T$ with the following property: Let $o \in T$ be a multivalent vertex and let v be a vertex in $\underline{\Gamma}_o$ with weight $m_o \neq 0$. Let $(C_0, \phi, T_C) \in \mathcal{M}^*_{\hat{A},T}^\Lambda$, and let E denote the end in C_0 that corresponds via T_C to v . Then the $\mathbb{R}/(2\pi\mathbb{Z})$ valued function $\hat{p}_o\phi - \hat{p}'_o t$ on E has a unique limit as $|s| \rightarrow \infty$, and this limit is obtained by composing the map Ψ_u with the map from $\mathcal{M}^*_{\hat{A},T}^\Lambda$ to O_T .

As should be clear from the definitions in the next subsection, the other parameters that enter O_T 's definition can also be given direct geometric meaning.

Theorems 6.2 and 6.3 are proved in Section 7.

The following proposition says something about $O_T - \hat{O}_T$ and the corresponding stabilizers in $\text{Aut}(T)$.

Proposition 6.4 *The $\text{Aut}(T)$ stabilizer of any point in O_T projects isomorphically onto a subgroup of Aut_\diamond . Moreover, two stabilizer subgroups are conjugate in $\text{Aut}(T)$ if and only if they have conjugate images in Aut_\diamond .*

The next subsection defines a version of Theorem 6.3's map and the subsequent subsection proves that this map is continuous. The final subsection provides a simpler picture of the $\text{Aut}(T)$ action on O_T when Aut_\diamond fixes some incident edge to \diamond . For example, this is the case when T is a linear graph; and the results from this subsection can be used to deduce Theorem 3.1 from Theorem 6.2. The final subsection also contains the proof of Proposition 6.4.

6.C The map from $\mathcal{M}^*_{\hat{A},T}$ to $\mathbb{R} \times O_T / \text{Aut}(T)$

Suppose that (C_0, ϕ) gives rise to an element in $\mathcal{M}^*_{\hat{A},T}$. The image of this element in $\mathbb{R} \times O_T / \text{Aut}(T)$ is obtained from a point that is assigned to (C_0, ϕ) in $\mathbb{R}_- \times (\times_o(\mathbb{R}_o \times \Delta_o))$. That latter assignment requires extra choices, some made just once for all (C_0, ϕ) and others that are contingent on the particular pair. As is explained below, the contingent choices are not visible in $\mathbb{R} \times O_T / \text{Aut}(T)$. The story on the map to $\mathbb{R} \times O_T / \text{Aut}(T)$ is told in four parts.

Part 1 This part and Part 2 describe the choices that do not depend on the given point in $\mathcal{M}^*_{\hat{A},T}$ nor on its representative pair (C_0, ϕ) .

The first choice in this regard is an $\text{Aut}(T)$ orbit, \hat{E}' , of edges that end in vertices of T that are labeled by the smallest angle. This angle is denoted in what follows by θ_- .

The next series of choices involve the vertex \diamond . The first of these involves the already chosen $\text{Aut}(T)$ orbit \widehat{E} in the set of incident edges to \diamond . Choose a ‘distinguished’ edge in \widehat{E} , and a corresponding ‘distinguished’ vertex on the latter’s version of the loop $\ell_{\diamond(\cdot)}$. Let ℓ_{\diamond} denote the corresponding distinguished version of $\ell_{\diamond(\cdot)}$ and let $v_{\diamond} \in \ell_{\diamond}$ denote the distinguished vertex.

To continue this series of choices at \diamond , select a concatenating path set that connects the vertex v_{\diamond} to each version of $\ell_{\diamond(\cdot)} \neq \ell_{\diamond}$. The definition of such a path set is provided in Definition 2.2. However, the selection must be constrained by the two conditions in the version of (2–18) where the vertex is \diamond and e is the distinguished edge.

An analogous, but more constrained set of choices must be made for each multivalent vertex $o \in \mathcal{V}$. To elaborate, let e denote now the edge that connects o to $T - T_o$ and let $e' \neq e$ denote another incident edge to o that is fixed by the $\mathbb{Z}/(n_o\mathbb{Z})$ action. Choose a concatenating path set to connect v_o to $\ell_{oe'}$ subject to the two conditions in (2–18). Such a concatenating path set must also be chosen when e' is not fixed by $\mathbb{Z}/(n_o\mathbb{Z})$, but additional care must be taken. To describe what is involved here, let E denote a non-trivial $\mathbb{Z}/(n_o\mathbb{Z})$ orbit in o ’s incident edge set and let $e' \in E$ denote the distinguished edge. Choose a concatenating path set from v_o to $\ell_{oe'}$ subject to the constraints in (2–18). Let $\{v_1, \dots, v_N\}$ denote this chosen set. Now, let $\iota \in \mathbb{Z}/(n_o\mathbb{Z})$ and let $v^{\iota} \subset \ell_o$ denote the path that starts at v_o and proceeds opposite the oriented direction along ℓ_o to its end at $\iota(v_o)$. The set of paths $\{\iota(v_1) \circ v^{\iota}, \iota(v_2), \dots, \iota(v_N)\}$ constitutes a concatenating path set that starts at v_o and ends at $\ell_{o\iota(e')}$. Here, $\iota(v_1) \circ v^{\iota}$ denotes the concatenation of the two paths. This new set obeys the $\iota(e')$ version of (2–18). Use the various versions of this set for the required concatenating paths for the elements in $E - e'$.

The preceding choices of concatenating path sets at the vertices in \mathcal{V} must be made so as to be compatible with the $\text{Aut}(T)$ action in the following sense: Let $o \in \mathcal{V}$, let $e' \in E_o$, and let $\{v_1, \dots, v_N\}$ denote the chosen concatenating path set that starts at the vertex v_o and ends on $\ell_{oe'}$. If $\iota \in \text{Aut}(T)$ maps the distinguished vertex on ℓ_o to that on $\ell_{\iota(o)}$, then $\{\iota(v_1), \dots, \iota(v_N)\}$ is the chosen concatenating path set that starts at $v_{\iota(o)}$ and ends on $\ell_{\iota(o)\iota(e')}$.

Part 2 Choose a vertex, \widehat{v} , in $\overline{\Gamma}^*$ that projects to v_{\diamond} . The choice for this vertex has three consequences. First, it trivializes the fiber bundle \mathbb{R}^{Δ} in the following manner: Let $(\tau, r) \in \text{Maps}(\text{Vert}_{\widehat{E}}; \mathbb{R}) \times \Delta_{\diamond}$ denote any given pair in \mathbb{R}^{Δ} . The trivializing map then sends (τ, r) to $(\tau(\widehat{v}), r)$. With this trivialization understood, write \mathbb{R}^{Δ} as $\mathbb{R}_{\diamond} \times \Delta_{\diamond}$ where \mathbb{R}_{\diamond} is a copy of \mathbb{R} .

The selection of \widehat{v} identifies \mathbb{R}^{-} with a fixed copy of \mathbb{R} , this labeled as \mathbb{R}_{-} in what follows. To describe how the identification comes about, let e denote the chosen

distinguished edge in \widehat{E} . Then \widehat{v} sits on a unique line, ℓ , in L_e . Now, let e' denote some other edge in \widehat{E} . The e' version of the concatenating path set defines a path in Γ_\diamond that starts at v_o and ends on $\ell_{\diamond e'}$. This path has a canonical lift to Γ^* and the latter has a unique lift to a path in $\overline{\Gamma}^*$ that starts at \widehat{v} and ends on a particular line in $L_{e'}$. The collection consisting of ℓ and these other lines specifies a point $L \in \times_{\widehat{e} \in \widehat{E}} L_{\widehat{e}}$. As a map, $\tau_- \in \mathbb{R}^-$ is determined by its value on L , so the assignment $\tau_- \rightarrow \tau_-(L) \in \mathbb{R}$ identifies \mathbb{R}^- with a fixed line.

The selection of \widehat{v} also identifies each \widehat{U}_e with the corresponding space $W_e \equiv [\times_{\widehat{o} \in \mathcal{V}(e)} \mathbb{R}_{\widehat{o}}] / [\times_{\widehat{o} \in \mathcal{V}(e)} \mathbb{Z}_{\widehat{o}}]$ that appears in (6–18). Here is how: If e is the distinguished edge from \widehat{E} , then the lift \widehat{v} sits on a unique $\ell \in L_e$. Then, the assignment to a given $x \in \text{Maps}(L_e; W_e)$ of $x(\ell) \in W_e$ identifies \widehat{U}_e with W_e . If e' is any other edge from E_\diamond , the choice of \widehat{v} and the e' version of the concatenating path set determines a path from \widehat{v} to a line in $L_{e'}$. The values of $x \in \text{Maps}(L_{e'}, W_{e'})$ on the latter identify $\widehat{U}_{e'}$ with $W_{e'}$.

Part 3 Granted this identification between (6–20) and (6–15), a point will be assigned to the given pair (C_0, ϕ) in \mathbb{R} , in each copy of Δ_o , in each copy of \mathbb{R}_o , and in \mathbb{R}_- . These assignments all require additional choices that must be made separately for each pair. These choices are listed below:

Choice 1 A choice of a correspondence of T in (C_0, ϕ) .

This chosen correspondence is used implicitly in the subsequent choices.

Choice 2 A parameterization for the component in $C_0 - \Gamma$ that corresponds to the distinguished edge in \widehat{E} .

The next set of choices are labeled by the multivalent vertices in T . The choice at any given vertex o is contingent on an a priori specification of a canonical parameterization of the component of $C_0 - \Gamma$ that corresponds to a certain edge in E_o . In the case that the $o = \diamond$, the incident edge is the distinguished edge in \widehat{E} and the parameterization is that provided by Choice 2. In the case that $o \in \mathcal{V}$, the edge in question connects o to $T - T_o$.

Choice 3 (at o) A lift to \mathbb{R} for the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the point on the boundary of the associated parametrizing cylinder that corresponds to v_o .

The canonical parameterization that is used to make any given $o \in \mathcal{V}$ version of this last choice is determined in an inductive fashion by the choices for those $\widehat{o} \in T$ where $T_{\widehat{o}}$ contains o . This induction works as follows: Suppose that o is a multivalent vertex

in T and that a canonical parameterization has been given to the component of $C_0 - \Gamma$ whose labeling edge supplies the loop ℓ_o in $\underline{\Gamma}_o$. Let e' denote any other incident edge to o . Chosen already is a concatenating path set that obeys the corresponding version of the constraints in (2–18). The parametrizing algorithm from Part 4 of Subsection 2.C uses this path set with o 's version of Choice 3 to specify the canonical parameterization for the e' component of $C_0 - \Gamma$.

Part 4 There are three cases to distinguish in order to describe the data for the value of the \mathbb{R} factor in $\mathbb{R} \times O_T / \text{Aut}(T)$ on a point that comes from $\mathcal{M}^*_{\hat{A}, T}$.

Case 1 In this case $\theta_- \neq 0$. The orbit \hat{E}' corresponds to a set of convex side ends in C where the $|s| \rightarrow \infty$ limit of θ is θ_- . Use \hat{E}' to also denote this set of ends. Associated to each end $E \in \hat{E}'$ is the real number $b \equiv b(E)$ that appears in (2–3). Here, $b(E)$ must be positive since θ_- is the infimum of θ on C . The map to \mathbb{R} assigns to (C_0, ϕ) the real number $-\zeta^{-1} \sum_{E \in \hat{E}'} \ln(b(E))$ where $\zeta \equiv \sqrt{6 \sin^2 \theta_E (1 + 3 \cos^2 \theta_E) / (1 + 3 \cos^4 \theta_E)}$.

Case 2 In this case $\theta_- = 0$ and there are no $(1, \dots)$ elements in \hat{A} . Here, \hat{E}' corresponds to a set of disjoint disks in C_0 whose centers are $\theta = 0$ points. This understood, then the image of (C_0, ϕ) in the \mathbb{R} factor of $\mathbb{R} \times O_T / \text{Aut}(T)$ is the sum of the s -coordinates of the centers of these disks.

Case 3 In this case $\theta_- = 0$ and \hat{E}' corresponds to a set of ends in C_0 whose constant $|s|$ slices limit to the $\theta = 0$ cylinder as $|s| \rightarrow \infty$. Let \hat{E}' also denote this set of ends. Each $E \in \hat{E}'$ defines the positive constant $\hat{c} \equiv \hat{c}(E)$ that appears in (1–9). Note that the integer p and p' that appear in (1–9) comprise the pair from the corresponding $(1, \dots)$ element in \hat{A} . The image of C in \mathbb{R} is $-(\sqrt{\frac{3}{2}} + \frac{p'}{p})^{-1} \sum_{E \in \hat{E}'} \ln(\hat{c}(E))$.

The point assigned (C_0, ϕ) in any given version of Δ_o is defined as follows: As a point in Δ_o is a map from $\text{Arc}(\underline{\Gamma}_o)$ to $(0, \infty)$, it is sufficient to provide a positive number to any given arc subject to the constraints in (6–6). For this purpose, let $\gamma \subset \underline{\Gamma}_o$ denote an arc. Then γ corresponds via T_C to a component of the locus Γ in C_0 . The integral over this component of the pull-back of $(1 - 3 \cos^2 \theta)d\varphi - \sqrt{6} \cos \theta dt$ is the value on γ of C_0 's assigned point in Δ_o .

The point assigned C_0 in any given version of \mathbb{R}_o is the chosen lift from o 's version of Choice 3 above.

The point assigned to C_0 in \mathbb{R}_- is obtained as follows: Each $e' \in \hat{E}$ labels a component of $C_0 - \Gamma$; keep in mind that each such component has been given a parametrization.

Let $w_{e'}$ denote the function w that appears in the corresponding version of (2–5). Now let σ be any value that is taken by θ on each $e' \in \hat{E}$ labeled component of $C_0 - \Gamma$. With σ understood, then C_0 's assigned value in R_- is

$$(6-27) \quad -\frac{1}{2\pi} \alpha_{Q_{\hat{E}}}(\sigma) \sum_{e' \in \hat{E}} \int w_{e'}(\sigma, v) dv.$$

6.D Continuity for the map to $O_T / \text{Aut}(T)$

This subsection makes two key assertions. Here is the first assertion: The assignment to (C_0, ϕ) of its point in $\mathbb{R} \times O_T / \text{Aut}(T)$ is independent of the choices that are made in Part 3 of the preceding subsection.

The second assertion is that the image of (C_0, ϕ) in $\mathbb{R} \times O_T / \text{Aut}(T)$ factors through $\mathcal{M}^*_{\hat{A}, T}$. Granted that such is the case, the constructions in the preceding subsection define a map from $\mathcal{M}^*_{\hat{A}, T}$ to $\mathbb{R} \times O_T / \text{Aut}(T)$ and it follows from the first assertion using Lemma 5.4 that this is a continuous map.

The remainder of this subsection has nine parts that justify these two assertions.

Part 1 The first observation concerns Choice 3 from the preceding subsection. A referral to the conclusions of Cases 2 and 4 of Part 5 from Subsection 2.C finds that changes in any $o \in \mathcal{V}$ version of Choice 3 are already invisible in the space that is depicted in (6–15). A referral to these same cases in Part 5 of Subsection 2.C finds that a change in \diamond 's version of Choice 3 or a change in Choice 2 moves (C_0, ϕ) 's assigned point in (6–15) along its orbit under the $\mathbb{Z} \times \mathbb{Z}$ action that defines the quotient in (6–9). Note in this regard that any change in \diamond 's version of Choice 3 will move C 's point in (6–15) along the $\mathbb{Z} \cdot Q_{\hat{E}}$ subgroup in $\mathbb{Z} \times \mathbb{Z}$. In any event, a change in Choice 2 or any version of Choice 3 is invisible in O_T . The following lemma states this conclusion in a formal fashion:

Lemma 6.5 *The constructions in the preceding subsection defines a continuous map from $\mathcal{M}^*_{\hat{A}, T}^\Lambda$ to $\mathbb{R} \times O_T$.*

Proof of Lemma 6.5 The conclusions made just prior to the lemma assert that the constructions of the preceding subsection assign a unique point in $\mathbb{R} \times O_T$ to each triple of the form (C_0, ϕ, T_C) where (C_0, ϕ) defines an equivalence class in $\mathcal{M}^*_{\hat{A}, T}$ and T_C is a correspondence in (C_0, ϕ) for T . To prove the lemma, it is enough to prove that the point assigned this triple depends only on its equivalence class in $\mathcal{M}^*_{\hat{A}, T}^\Lambda$.

For this purpose, suppose that ψ is a holomorphic diffeomorphism of C_0 and let $\phi' = \phi \circ \psi$. Let T_C' denote the correspondence that is obtained from T_C as follows:

When $e \in T$ is an edge and K_e its corresponding component in $C_0 - \Gamma$ as given by T_C , then $T_{C'}$ makes e correspond to $\psi^{-1}(K_e)$. Likewise, if $\gamma \subset \underline{\Gamma}_o$ is an arc, then it corresponds via T_C to an arc, γ_C , in C_0 . The arc γ corresponds via $T_{C'}$ to $\psi^{-1}(\gamma_C)$. This is designed so that the ϕ image of any T_C labeled subset of C_0 is identical to the ϕ' image of the corresponding $T_{C'}$ labeled subset.

Granted the preceding conclusion, then (C_0, ϕ, T_C) and $(C_0, \phi', T_{C'})$ have the same image in the \mathbb{R} factor of $\mathbb{R} \times O_T / \text{Aut}(T)$ since the respective images in this factor are defined from the ϕ and ϕ' images of respective components with the same edge labels in the $\phi^*\theta$ and $\phi'^*\theta$ versions of $C_0 - \Gamma$. A similar line of reasoning as applied to the arcs in any given version $\underline{\Gamma}_o$ explains why the images of (C_0, ϕ, T_C) and $(C_0, \phi', T_{C'})$ agree in the corresponding Δ_o factor of O_T .

To continue, let e now denote the distinguished edge in the $\text{Aut}(T)$ orbit \hat{E} from \diamond 's incident edge set. The chosen canonical parametrization for K_e can be pulled back via ψ and this pull-back can be used for the canonical parametrization of $\psi^{-1}(K_e)$. If this is done, then the other choices from Part 3 of Subsection 6.C can be made so that the canonical parametrization for any given $\psi^{-1}(K_{\hat{e}})$ is the pull-back via ψ of that for $K_{\hat{e}}$. Indeed, since the ϕ image of $K_{\hat{e}}$ is the same as the ϕ' image of $\psi^{-1}(K_{\hat{e}})$, the lifts to \mathbb{R} that are used to inductively define these parametrizations can be taken to agree at each stage of the induction. In particular, doing so guarantees that the respective assignments to (C_0, ϕ, T_C) and $(C_0, \phi', T_{C'})$ in \mathbb{R}_- and in each \mathbb{R}_o factor of $\mathbb{R}_- \times (\times_o \mathbb{R}_o)$ also agree.

Part 2 Left yet to discuss is the affect in a change for the correspondence of T in (C_0, ϕ) . Now any change from the original correspondence can be viewed as the result of using the original correspondence after acting on T by an element $\iota \in \text{Aut}(T)$. Such is the view taken here. With this understood, consider: □

Lemma 6.6 *The change in the original correspondence by the action of $\iota \in \text{Aut}(T)$ changes the assigned point in $\mathbb{R} \times O_T$ by the $\text{Aut}(T)$ action of this same ι .*

Note that the first of the two assertion made at the outset of this subsection is an immediate consequence of Lemma 6.6. Meanwhile, the second of these assertions is a direct corollary to Lemmas 6.5 and 6.6.

The remainder of this part and the subsequent parts of this subsection contain the folowing proof.

Proof of Lemma 6.6 Consider first the case where o is a vertex in \mathcal{V} and the automorphism ι generates the corresponding $\mathbb{Z}/(n_o\mathbb{Z})$ subgroup. The composition of the

original correspondence with ι does not change the identification between (6–20) and (6–15). This said, the new correspondence does not change the assignment to (C_0, ϕ) in either \mathbb{R}_- nor in any $\hat{o} \in T - T_o$ version of $\mathbb{R}_{\hat{o}} \times \Delta_{\hat{o}}$.

To analyze the change of (C_0, ϕ) 's assigned point in the remaining factors $\mathbb{R}_{\hat{o}} \times \Delta_{\hat{o}}$, let \hat{o} denote a vertex in T_o and let γ denote an arc in $\underline{\Gamma}_{\hat{o}}$. The original correspondence identified (\hat{o}, γ) with a component of the complement of the critical point set in the $\theta = \theta_{\hat{o}}$ locus in C_0 . The new one identifies the latter component now with the arc $\iota(\gamma)$ in $\underline{\Gamma}_{\iota(\hat{o})}$. Thus, the new map for C_0 in $\Delta_{\iota(\hat{o})}$ gives the value on $\iota(\gamma)$ that the original map for C in $\Delta_{\hat{o}}$ gives γ . This is what (6–21) asserts. \square

Part 3 Consider now the change in the assignment to \mathbb{R}_o . Let e denote the edge that connects o to $T - T_o$. Choices 2 and 3 from Part 3 of the preceding subsection can be made for the new correspondence so that the original parametrization of K_e is unchanged. The original correspondence identifies the distinguished vertex v_o with a particular missing or singular point on the $\sigma = \theta_o$ circle of the parametrizing cylinder for K_e . Let v denote the latter point and let v' denote the missing or singular point on this same circle that is identified by the original correspondence with $\iota^{-1}(v_o)$. The new correspondence identifies v' with v_o . As a consequence, the o version of Choice 3 for the new correspondence can be taken to be a lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of v' . This is what (6–22) asserts.

Part 4 Turn next to the change in some $\mathbb{R}_{\hat{o}}$ in the case that \hat{o} is a vertex in $T_o - o$. Here, there are two cases to consider. The first case is that where \hat{o} is in a component of $T_o - o$ that is fixed by ι . Thus, the component connects to o through an ι -invariant edge in E_o . Let e' denote the latter. The change in the assignment to $\mathbb{R}_{\hat{o}}$ is determined up to Choice 3 modifications by the change in the canonical parametrization of $K_{e'}$. Indeed, if the parametrization changes due to the action of some integer pair N , then the conclusions in Part 5 of Subsection 2.C imply that the new assignments to the versions of $\mathbb{R}_{\hat{o}}$ that connect to o through e' can be made so that each differs from the original by the addition of the number depicted in (6–8).

Granted the preceding, let $\{v_1, \dots, v_N\}$ denote the concatenating path set that connects v_o to a vertex in $\ell_{oe'}$. The new parametrization can be viewed as one that is obtained via the parametrizing algorithm of Subsection 2.C by using the original parametrization of K_e , the original \mathbb{R} lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the point v on the $\sigma = \theta_o$ circle in the parametrizing domain, but a different concatenating path set from v_o to a vertex on $\ell_{oe'}$. To explain, let v^t denote the path in ℓ_o that starts at v_o and runs opposite the orientation to $\iota^{-1}(v_o)$. The new concatenating path set is $\{v^t, \iota^{-1}(v_1) \circ v^t, \iota^{-1}(v_2), \dots, \iota^{-1}(v_N)\}$. Let $\gamma_N \subset v_N$ denote the final arc. Thus, $\gamma_N \subset \ell_{oe'}$. Let

$v' \subset \ell_{oe'}$ denote the oriented path in $\ell_{oe'}$ that starts with $\iota^{-1}(\gamma_N)$ and then runs from its ending vertex in the oriented direction on $\ell_{oe'}$ to the end vertex of γ_N , then finishes by traversing γ_N against its orientation to its start vertex. The ordered set $\{\iota^{-1}(v_1) \circ v', \iota^{-1}(v_2), \dots, \iota^{-1}(v_N), v', v_N^{-1}, \dots, v_1^{-1}\}$ is a concatenating path set that starts at v_o and has final arc ending at v_o . Let μ denote the oriented loop in $\underline{\Gamma}_o$ that is obtained by traversing the constituent paths from left most to right most in this ordered set, taking into account that the the final arc of any one is the initial arc of the next. This loop has a canonical lift, μ^* , to a loop in $\underline{\Gamma}_o^*$. According to Lemma 2.3, the new parametrization of $K_{e'}$ is obtained from the old by the action of the integer pair $-\phi_o([\mu^*])$ where $[\mu^*]$ denotes the homology class of the oriented loop μ^* and ϕ_o denotes the canonically defined class in $H^1(\underline{\Gamma}_o^*; \mathbb{Z} \times \mathbb{Z})$.

The computation of $\phi_o([\mu_*])$ relies on the fact that ϕ_o is $\mathbb{Z}/(n_o\mathbb{Z})$ invariant. Thus, as ι is the generator,

$$(6-28) \quad \phi_o([\mu_*]) = \frac{1}{n_o} \phi_o([\mu_*] + [(\iota(\mu))_*] + \dots + [(\iota^{n_o-1}(\mu))_*]).$$

This understood, the sum of homology classes that appears here is the class that is obtained by traversing $\ell_{oe'}$ once in its oriented direction while traversing ℓ_{oe} once in the direction that is opposite to its orientation. Thus,

$$(6-29) \quad \phi_o([\mu_*]) = \frac{1}{n_o} (Q_{e'} - Q_e).$$

With the preceding understood, the resulting change in the assignment to (C_0, ϕ) in $\mathbb{R}_{\hat{o}}$ is that given by (6-23).

Part 5 This part considers the case that $\hat{o} \in T_o - o$ is a vertex that lies in a component that is not fixed by the $\mathbb{Z}/(n_o\mathbb{Z})$ action. Let $e' \in E_o$ again denote the edge that connects \hat{o} 's component of $T_o - o$ to o . Let $E \subset E_o$ denote the orbit of e' for the $\mathbb{Z}/(n_o\mathbb{Z})$ action. Let \hat{e} denote the edge in T that connects \hat{o} to $T - T_{\hat{o}}$. The component of $C_0 - \Gamma$ that is labeled by the new correspondence by $\iota(\hat{e})$ is the component that the old one assigned to \hat{e} . The parametrization of this component will change by the action of an integer pair, N . This understood, the new assignment in $\mathbb{R}_{\iota(\hat{o})}$ is thus obtained from the old assignment in $\mathbb{R}_{\hat{o}}$ by adding the term in (6-8).

The key point now is that the integer pair N is the same for all vertices in \hat{o} 's component of $T_o - o$. In particular, the integer pair N is the pair that changes the parameterization of the component that is labeled by e' using the old parameterization. Let K' denote the latter. To explore the parameterization change, let $\{v_1, \dots, v_N\}$ denote the chosen concatenating path set that is used originally to give the canonical parameterization to K' . Now, there are two cases to consider. In the first, e' is the distinguished edge in

its orbit. In this case, $\iota(e')$ is the first edge and the concatenating path that gives its new orientation is $\{\nu_1 \circ \nu, \nu_2, \dots, \nu_N\}$ where ν is the path that circumnavigates ℓ_o in the direction opposite to its orientation. According to the conclusions from Part 5 of Subsection 2.C, this then implies that $N = Q_e$. In the second case, e' is not the distinguished edge. In this case, the same concatenating path set as the original gives the new parametrization of K' and so $N = 0$. Thus, the affect of ι on the $\{\mathbb{R}_\diamond\}$ assignments are as depicted by (6–24).

Part 6 This and Parts 7–9 concern the effect on (C_0, ϕ) 's assigned point in $O_T / \text{Aut}(T)$ when the change in the correspondence is obtained using an element in (6–13)'s Aut_\diamond subgroup. Let ι denote such an element. The first point to consider is the affect of ι on the assignment to the simplex Δ_\diamond . Let r denote the original assignment and r' the new one. If the original correspondence identifies a given component of the T_C version of Γ_\diamond with a given arc γ in the T version of $\underline{\Gamma}_\diamond$, then the new correspondence identifies this same component with $\iota(\gamma)$. But this means that r' has the same value on $\iota(\gamma)$ as r does on γ . Thus, $r' = \iota(r)$.

This line of reasoning finds the analogous formula for the change in the assignment to $\times_{o \in \mathcal{V}} \Delta_o$. To be precise, use r_o to denote the original assignment to Δ_o and r'_o the new one. Then $r'_{\iota(o)}(\iota(\gamma)) = r_o(\gamma)$ for all $\gamma \in \underline{\Gamma}_o$.

Part 7 The story on the other aspects of (C_0, ϕ) 's assignment in $O_T / \text{Aut}(T)$ starts here by considering the new assignment for (C_0, ϕ) in the space depicted in (6–15). This assigned point has a partner in the space depicted in (6–20). The original assignment for (C_0, ϕ) also provides a point for (C_0, ϕ) in the space depicted in (6–20). These two points are compared in what follows.

The new assignment in (6–15) requires new versions of Choice 2 and Choice 3 from Part 3 of the preceding subsection. This is because ι can move the distinguished edge in \hat{E} . To make these new choices, let e denote the original distinguished edge in \hat{E} . The component of $C_0 - \Gamma$ that corresponds via the new correspondence to e is the component, denoted here by K , that corresponds via the original to $\iota^{-1}(e)$. Choice 2 requires a parametrization of K . A convenient choice is the canonical parametrization as defined by the original choices.

The \diamond version of Choice 3 requires a lift to \mathbb{R} of the point on the $\sigma = \theta_\diamond$ circle of K 's parametrizing cylinder that corresponds via the new correspondence to the distinguished vertex v_\diamond . This is the point that corresponds to $\iota^{-1}(v_\diamond)$ via the original correspondence. A convenient choice is obtained as follows: Let $\{\nu_1, \dots, \nu_N\}$ denote the concatenating path set that is used to define the original parametrization for K . Note that the ending vertex of ν_N corresponds via the original correspondence to a

certain point on the $\sigma = \theta_\diamond$ circle in the original parametrizing cylinder of K . Use z to denote this point. The parametrizing algorithm from Subsection 2.C provides an \mathbb{R} lift of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of z . Subtract from this the integral of dv in the oriented direction on the $\sigma = \theta_\diamond$ circle along the segment that starts at $\iota^{-1}(v_\diamond)$ and ends at z . Let τ_{new} denote the resulting real number. Use τ_{new} for the new version of \diamond 's Choice 3.

Part 8 This same τ_{new} is the new assignment for (C_0, ϕ) in the \mathbb{R}_\diamond factor in (6–15). The story on the new assignments to the other factors requires the preliminary digression that follows. To start the digression, let e be an incident edge to \diamond , and let $\{v^e_1, \dots, v^e_{N(e)}\}$ denote the chosen concatenating path set that connects v_\diamond to a vertex on ℓ_{oe} . Let \hat{e} denote the edge that is mapped by ι to the distinguished vertex in \hat{E} . The \hat{e} version of this concatenating path set is denoted at times as in Part 7, thus by $\{v_1, \dots, v_N\}$.

To continue, let v denote the path in $\ell_{o\hat{e}}$ that starts at the beginning of the final arc in v_N and, after reaching the end of this arc, then proceeds against the orientation to the point $\iota^{-1}(v_\diamond)$. Then both $\{v^e_1, \dots, v^e_{N(e)}\}$ and $\{v_1, \dots, v_N, \iota^{-1}(v^{\iota(e)}_1) \circ v, \iota^{-1}(v^{\iota(e)}_2), \dots, \iota^{-1}(v^{\iota(e)}_{N(\iota(e))})\}$ define concatenating path sets that run from v_\diamond to ℓ_{oe} . Let v' denote the path in $\ell_{\diamond e}$ that starts with the last arc in $\iota^{-1}(v^{\iota(e)}_{N(\iota(e))})$, then continues in the oriented direction from the latter's end to the start of the final arc in $v^e_{N(e)}$ and then traverses this last arc in reverse. So defined, the set

$$(6-30) \quad \{v_1, \dots, v_N, \iota^{-1}(v^{\iota(e)}_1) \circ v, \iota^{-1}(v^{\iota(e)}_2), \dots, \iota^{-1}(v^{\iota(e)}_{N(\iota(e))}), v', v^e_{N(e)}{}^{-1}, \dots, v^e_1{}^{-1}\}$$

concatenates by gluing the final arc of each constituent path to the first arc of the subsequent one so has to give a loop that starts and ends at v_\diamond . Let μ_e denote this loop.

As constructed, the loop μ_e as a canonical lift as a loop in Γ^*_\diamond , this denoted by μ_e^* . Let $N_e = (n_e, n_e')$ denote the value of $-\phi_\diamond$ on $[\mu_e^*]$. According to Lemma 2.3, the parameterization of the component of $C_0 - \Gamma$ that originally corresponded to e is changed with the new choices by the action of the integer N_e .

As a consequence of this last conclusion, the new assignment to (C_0, ϕ) in the \mathbb{R} -factor of (6–15) is obtained from the old by adding $-2\pi \sum_{e \in \hat{E}} (n_e' q_{\hat{E}} - n_e q_{\hat{E}}')$.

Here is a second consequence: The new assignment to (C_0, ϕ) in the $\times_{o \in \mathcal{V}} \mathbb{R}_o$ factor can be made so that if τ_o denotes the original assignment in \mathbb{R}_o and τ_o' the new one, then

$$(6-31) \quad \tau_{\iota(o)'} = \tau_o - 2\pi \frac{\alpha_{N_e}(\theta_{\hat{o}})}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})},$$

where e denotes the edge that connects o 's component of $T - \diamond$ to \diamond .

Note for the future that there is one more path in Γ_\diamond that plays an important role in the story, this the path that is constructed from the paths in the set $\{v_1, \dots, v_N, v\}$ by identifying the final arc in any v_j with the starting arc in the subsequent path. This path is denoted below as μ . It starts at v_\diamond and ends at $\iota^{-1}(v_o)$. The path μ has a canonical preimage in Γ_\diamond^* . The latter is denoted in what follows by μ^* .

Part 9 It remains now to compare the new and old assignments in (6–15) by viewing the latter as points in (6–20). There are three steps to this task.

Step 1 For this purpose, use $(\tau, r) \in \text{Maps}(\text{Vert}_{\hat{E}}) \times \Delta_\diamond$ to denote (C_0, ϕ) 's original assignment in \mathbb{R}^Δ . Use (τ', r') to denote the new assignment. As noted already, $r' = \iota(r)$. To see about τ' , let \hat{v} denote the chosen point in $\bar{\Gamma}^*$ over v_\diamond . The path μ^* has a lift to $\bar{\Gamma}^*$ that starts at \hat{v} and ends at some vertex \hat{v}' , a vertex that maps to $\iota^{-1}(v_\diamond)$. According to (6–16), the value of τ on \hat{v}' is τ_{new} . Meanwhile, τ_{new} is also equal to the value of τ' on \hat{v} . Indeed, such is the case because Part 7 found τ_{new} to be (C_0, ϕ) 's new \mathbb{R}_\diamond assignment in (6–15).

Now, ι has a unique lift in $\hat{\text{Aut}}_\diamond$ that sends \hat{v}' to \hat{v} . Use $\hat{\iota}$ to denote the latter. The conclusions of the last paragraph assert that $\tau'(\hat{v}) = \tau_{new} = \tau(\hat{\iota}^{-1}\hat{v})$. Thus, $\tau' = \hat{\iota}(\tau)$.

Step 2 To study the \mathbb{R}^- factors, let τ_- denote the original assignment and τ'_- the new one. The values for τ_- and τ'_- were defined by specifying them on a particular element in the set Λ . Denote this element by \mathcal{L} . To elaborate, the component of L with label any given $e \in \hat{E}$ was obtained as follows: Construct the path v^e in Γ_o from $\{v^e_1, \dots, v^e_{N(e)}\}$ by gluing the end arc of any $j < N(e)$ version to the starting arc of the subsequent version. This path has a canonical lift to Γ_o^* and thus a canonical lift to $\bar{\Gamma}^*$ as a path that starts at \hat{v} . The end vertex of this lift lies on an inverse image of $\ell^*_{\diamond e}$. This is the component of L that is labeled by e .

The conclusions in Part 7 about the \mathbb{R}_- factors can now be summarized by the relation $\tau'(L) = \tau(L) - 2\pi \sum_{e \in \hat{E}} (n_e' q_{\hat{E}} - n_e q_{\hat{E}}')$. As will now be explained, this is also the value of τ on $\hat{\iota}^{-1}(L)$. To see why, let $e \in \hat{E}$, let ℓ denote e 's component of L and let ℓ' denote the component for $\iota(e)$. Thus, $\hat{\iota}^{-1}(\ell')$ is in L_e and this is e 's component of $\hat{\iota}^{-1}(L)$. As such, it is obtained from ℓ by the action of some element in $\mathbb{Z} \times \mathbb{Z}$. To find this element, note that as ℓ was defined by the set $\{v^e_1, \dots, v^e_{N(e)}\}$, so $\hat{\iota}^{-1}(\ell')$ is defined by the ordered set $\{v_1, \dots, v_N, \iota^{-1}(v^{\iota(e)}_1) \circ v, \iota^{-1}(v^{\iota(e)}_2), \dots, \iota^{-1}(v^{\iota(e)}_{N(\iota(e))})\}$. This then implies that $\hat{\iota}^{-1}(\ell')$ can be obtained from ℓ by the action of N_e .

Granted that such is the case for all $e \in \widehat{E}$, then (6–17) gives $\tau(\widehat{\iota}^{-1}(L))$ the desired value.

Step 3 The final concern is that of the assignments to any given version of $\widehat{U}_{(\cdot)}$. For this purpose, suppose that e is an incident edge to \diamond . The old and new correspondences each assign (C_0, ϕ) a map from L_e to $W_e = (\times_{o \in \mathcal{V}(e)} \mathbb{R}_o) / (\times_{o \in \mathcal{V}(e)} \mathbb{Z}_o)$. These maps are defined by lifts to $\times_{o \in \mathcal{V}(e)} \mathbb{R}_o$. In what follows, $\tau_o(\ell)$ and $\tau'_o(\ell)$ are used to denote the respective old and new values on a given $\ell \in L_e$ for the factor labeled by a given $o \in \mathcal{V}(e)$.

Now, both maps to W_e are determined by their values on a particular element in L_e . The latter, ℓ^e , is determined as follows: As described previously, the paths that comprise the set $\{v^e_1, \dots, v^e_{N(e)}\}$ glue together in a sequential fashion so as to define a path from v_\diamond to a vertex on ℓ_{oe} . This path then lifts to a path that starts at \widehat{v} and ends at a vertex on a unique preimage of ℓ_{oe} in $\overline{\Gamma}^*$. This preimage is ℓ^e .

The components of Γ that are assigned by the new correspondence to vertices in $\mathcal{V}(\iota(e))$ are assigned to vertices in $\mathcal{V}(e)$ by the old correspondence. This understood, let $o \in \mathcal{V}(e)$. Equation (6–31) describes the relationship between $\tau'_{\iota(o)}(\ell^{\iota(e)})$ and $\tau_o(\ell)$.

As is explained next, (6–31) is also the formula for $\widehat{\iota}(x)$ on $\ell^{\iota(e)}$. Indeed, by definition, $\widehat{\iota}(x)(\widehat{\iota}(\ell^e)) = x(\ell^e)$. To see what this means, note that $\widehat{\iota}(\ell^e) \in L_{\iota(e)}$ and so $\ell^{\iota(e)}$ is obtained from $\widehat{\iota}(\ell^e)$ by the action of some integer pair. As argued at the end of Step 2 in the case when $e \in \widehat{E}$, this integer pair is N_e . This being the case, then (6–19) implies that x' equal $\widehat{\iota}(x)$.

6.E More about $\text{Aut}(T)$

The story told by Theorems 6.2 and 6.3 simplifies to some extent when Aut_\diamond fixes one of \diamond 's incident edges. As is explained below, there is no need in this case to use the space in (6–20) because the $\text{Aut}(T)$ action is readily visible on the space in (6–9).

To begin the story in this case, let e denote an incident edge to \diamond that is fixed by $\text{Aut}(T)$. Agree to use e for the distinguished incident edge orbit. Since Aut_\diamond fixes e , it must act as a subgroup of the group of automorphisms of the labeled graph $\ell_{\diamond e} \equiv \ell_\diamond$. Thus, Aut_\diamond is a cyclic group whose order is denoted by n_\diamond . This implies that $\text{Aut}(T)$ is isomorphic to the $o = \diamond$ version of (6–14). Granted these remarks, view $\text{Aut}_\diamond = \mathbb{Z}/(n_\diamond \mathbb{Z})$ as a subgroup of $\text{Aut}(T)$ using this same version of (6–14).

Define the action of Aut_\diamond on the space in (6–9) via an action of $\widehat{\text{Aut}}_\diamond$ on the space in (6–15). In this regard, note that the group $\widehat{\text{Aut}}_\diamond$ is isomorphic now to

$$(6-32) \quad \left[\left(\frac{1}{n_\diamond} \mathbb{Z} \right) \times (\mathbb{Z} \times \mathbb{Z}) \right] / \mathbb{Z},$$

where the notation is as follows: The first factor in the brackets arises as the \mathbb{Z} extension

$$(6-33) \quad 1 \rightarrow \mathbb{Z} \rightarrow \frac{1}{n_\diamond} \mathbb{Z} \rightarrow \mathbb{Z}/(n_\diamond \mathbb{Z}) \rightarrow 1,$$

where the \mathbb{Z} action is that of the subgroup that covers the identity in $\mathbb{Z}/(n_\diamond \mathbb{Z})$. Meanwhile, the \mathbb{Z} action on $\mathbb{Z} \times \mathbb{Z}$ is the subgroup $\mathbb{Z} \cdot Q_e$ whereby $1 \in \mathbb{Z}$ acts as $-Q_e$. This understood, the equivalence class in (6-32) of a pair $(z, N = (n, n'))$ with $z \cdot n_\diamond \in \mathbb{Z}$ and $N \in \mathbb{Z} \times \mathbb{Z}$ acts on the factor \mathbb{R}_- in (6-15) as the translation by $-2\pi(n'q_e - nq_e')$. Meanwhile, it acts on \mathbb{R}_\diamond as the translation by the $\hat{o} = \diamond$ version of

$$(6-34) \quad -2\pi \frac{\alpha_{N+zQ_e}(\theta_{\hat{o}})}{\alpha_{Q_e}(\theta_{\hat{o}})}.$$

The action of $\iota = [z, N]$ on the factors $\times_o \Delta_o$ is such that if $r \in \Delta_o$ and γ is an arc in Γ_o , then the resulting $\iota(r)(\iota(\gamma))$ equals $r(\gamma)$; here, $\iota(\gamma)$ is the image of γ under the induced map from the set of arcs in Γ_o to the set in $\Gamma_{\iota(o)}$.

The description of the action on $(\times_{o \in \mathcal{V}} \mathbb{R}_o)/(\times_{o \in \mathcal{V}} \mathbb{Z}_o)$ requires distinguishing two separate cases. For this purpose, keep in mind the following: When $E \subset E_\diamond$ is an Aut_\diamond orbit, then $\hat{\text{Aut}}_\diamond$ preserves the factor $\times_{e' \in E} [(\times_{o \in \mathcal{V}(e')} \mathbb{R}_o)/(\times_{o \in \mathcal{V}(e')} \mathbb{Z}_o)]$. Up for discussion first is the case where E is a single edge. In this regard, note that $\mathcal{V}(e) = \emptyset$ so there is at most one such E unless Aut_\diamond is trivial. Let e' denote the fixed edge. Then the action of $(z, N) \in \hat{\text{Aut}}_\diamond$ comes from the following action on $\times_{o \in \mathcal{V}(e')} \mathbb{R}_o$: Let τ denote a point in this product and let $\hat{o} \in \mathcal{V}(e')$. Then the value of $\iota(\tau)$ in the factor labeled by any given $\iota(\hat{o})$ is obtained by subtracting

$$(6-35) \quad \frac{2\pi}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} \left[\frac{1}{n_\diamond} \alpha_{Q_e - Q_{e'}}(\theta_{\hat{o}}) + \alpha_{N+zQ_e}(\theta_{\hat{o}}) \right]$$

from the value of τ in the factor labeled by \hat{o} . Here, \hat{e} is the edge that connects \hat{o} to $T - T_{\hat{o}}$.

If $E \subset E_\diamond$ is a non-trivial orbit of Aut_\diamond , then E has n_\diamond edges and they have a canonical cyclic ordering. The definition of the $\hat{\text{Aut}}_\diamond$ action on $\times_{e' \in E} [(\times_{o \in \mathcal{V}(e')} \mathbb{R}_o)/(\times_{o \in \mathcal{V}(e')} \mathbb{Z}_o)]$ requires the choice of a distinguished edge in E and the definition of a compatible linear ordering with the distinguished edge last. Granted such a choice, define the action of $(z, N) \in \hat{\text{Aut}}_\diamond$ from the following action on $\times_{o \in \mathcal{V}(e')} \mathbb{R}_o$: Let τ denote a point in this product and let $\hat{o} \in \mathcal{V}(e')$. Then the value of $\iota(\tau)$ in the factor labeled by any given $\iota(\hat{o})$ is obtained by subtracting

$$(6-36) \quad \frac{2\pi}{\alpha_{Q_{\hat{e}}}(\theta_{\hat{o}})} \left[\varepsilon_{\hat{o}} \alpha_{Q_e}(\theta_{\hat{o}}) + \alpha_{N+zQ_e}(\theta_{\hat{o}}) \right]$$

from the value of τ on \hat{o} . Here, $\varepsilon_{\hat{o}} = 0$ unless the edge that connects \hat{o} 's component of $T_{\diamond} - \diamond$ to \diamond is the distinguished edge in E . In this case, $\varepsilon_{\hat{o}} = 1$.

The identification given above between (6–15) and (6–20) intertwines these $\widehat{\text{Aut}}(T)$ actions if the various concatenating path sets are chosen in an appropriate fashion. This understood, the version given here of the $\text{Aut}(T)$ action on O_T can be used for $\widehat{O}_T / \text{Aut}(T)$ in the statements of Theorems 6.2 and 6.3.

With regards to $O_T - \widehat{O}_T$, more can be said here than is stated in Proposition 6.4. For this purpose, various notions must be introduced. The first is the integer, m_- , this the least common divisor of the integers that comprise the pair Q_e . Next, let $T^F \subset T$ denote the subgraph on which $\text{Aut}(T)$ acts trivially. Thus, the edge e is in T^F , but T^F may well be bigger. In any event, T^F is connected. Let k_T denote the greatest common divisor of the integers in the set that consists of m_- and the versions of n_o for vertices $o \in T^F$.

The next notion is that of a ‘canonical diagonal subgroup’ of $\text{Aut}(T)$. To define this notion, note that if $o \in T^F$, then $\mathbb{Z}/(n_o\mathbb{Z})$ has a unique $\mathbb{Z}/(k_T\mathbb{Z})$ subgroup and the latter has a canonical generator; this is the element that is the smallest multiple of the generator of $\mathbb{Z}/(n_o\mathbb{Z})$. Granted the preceding, a subgroup of $\text{Aut}(T)$ is a canonical diagonal subgroup when it has the following two properties: First, it is cyclic of order k_T . Second, it has a generator that maps to the canonical generator of each $\mathbb{Z}/(k_T\mathbb{Z})$ subgroup of each $o \in T^F$ version of $\mathbb{Z}/(n_o\mathbb{Z})$. According to Proposition 6.4, any two canonical diagonal subgroups are conjugate in $\text{Aut}(T)$.

With the introduction now over, consider:

Proposition 6.7 *The stabilizer in $\text{Aut}(T)$ of any given point in O_T is a subgroup of some canonical diagonal subgroup. Conversely, if G is a subgroup of a canonical diagonal subgroup of $\text{Aut}(T)$, then the fixed point set of G is the product of its corresponding fixed point set in $\times_o \Delta_o$ and a product of circles, one that corresponds to the factor \mathbb{R}_- and the rest labeled in a canonical fashion by the various orbits of G in T 's multivalent vertex set.*

The proof of this proposition uses arguments from the proof of Proposition 6.4 and so the latter proof is offered first.

Proof of Proposition 6.4 To prove the first assertion, suppose for argument's sake that an element $\iota \in \text{Aut}(T)$ maps to the identity in Aut_{\diamond} yet fixes a given point in O_T . This means that ι is determined by its components in the versions of $\text{Aut}(T_{(\cdot)})$ that are labeled by the vertices that share an edge with \diamond . Thus, it is sufficient to consider

the case where $\iota \in \text{Aut}(T_o)$ for some $o \in T - \diamond$. As such, it is permissible to view O_T as in (6–9). Granted this, then the arguments given in Subsection 3.C to prove Propositions 4.4 and 4.5 can be borrowed almost verbatim to prove that ι is the identity element.

The proof of the last part of the proposition uses an induction argument that moves from any given $o \in T$ to the vertices in $T_o - o$ that share its incident edges. The phrasing of the induction step uses the notion of the ‘generation number’ of a multivalent vertex in T . Here, \diamond is the only generation 0 vertex, and a vertex o has generation $k > 0$ when it shares an edge with a generation $k - 1$ vertex in $T - T_o$. Use $\text{Vert}(k)$ to denote the set of generation k vertices. \square

The lemma that follows facilitates the induction.

Lemma 6.8 *Suppose that $k > 0$ and that ι and ι' both stabilize points in O_T and have the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k)} \text{Aut}(T_o))$. Then, there exists some $j \in \times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o)$ such that $j\iota j^{-1}$ and ι' have the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o))$.*

This lemma is proved below. The actual induction step is made using the following generalization:

Lemma 6.9 *Suppose that $k > 0$ and suppose that G and G' are subgroups of $\text{Aut}(T)$ that stabilize points in O_T and have the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k)} \text{Aut}(T_o))$. Then, there exists some $j \in \times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o)$ such that jGj^{-1} and G' have the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o))$.*

Proof of Lemma 6.9 By appeal to Lemma 6.8 there are non-trivial subgroups in G and in G' that have the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o))$. Take $H \subset G$ and $H' \subset G'$ to be a maximal group of this sort in the following sense: There is no subgroup of G that properly contains H and has a partner in G' with the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o))$. The argument that follows deriving nonsense if $H \neq G$.

To start this derivation, suppose that $\iota \in G - H$ and let $\iota' \in G'$ denote the element that shares ι 's image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k)} \text{Aut}(T_o))$. A second appeal to Lemma 6.8 finds some $j \in \times_{o \in \text{Vert}(k)} \text{Aut}(T_o)$ such that $j\iota j^{-1} = \iota'$. Of concern is whether the commutators of j with the elements in H are all in $\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o)$. If such is the case, then the group generated by H and $j\iota j^{-1}$ contains H as a proper subgroup and has the same image in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o))$ as the group generated by

H and ι' . Of course, this contradicts the assumed maximality of H . In any event, j acts in $\mathbb{Z}/(n_o\mathbb{Z})$ as the translation by some integer, here denoted by j_o . Thus, $j_o \in \{0, \dots, n_o - 1\}$. Note that the commutators of j with the elements of H all lie in $\times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o)$ if and only if the assignment $o \rightarrow j_o$ is constant on H orbits in $\text{Vert}(k)$.

To see that such must be the case, note first that ι and ι' have identical actions on $\text{Vert}(k)$ since their actions here are determined by their actions on the set of edges that are incident to the vertices in $\text{Vert}(k-1)$. Thus if ι fixes $o \in \text{Vert}(k)$, then so does ι' . In this case, there is nothing lost by taking $j_o = 0$. Suppose next that ι does not fix o . Now, ι maps ℓ_o to $\ell_{\iota(o)}$ as does ι' , and the affect of the latter is that of $j_{\iota(o)}\iota'j_o^{-1}$.

Keep this last observation on hold for the moment for the following key observations: First, if $u \in \text{Aut}(T)$, then $n_o = n_{u(o)}$. Second, u induces a map from the vertex set on ℓ_o to that on $\ell_{u(o)}$ that respects the cyclic ordering. Third, this map intertwines the action of $i \in \mathbb{Z}/(n_o\mathbb{Z})$ with its corresponding action in $\mathbb{Z}/(n_{u(o)}\mathbb{Z})$. Written prosaically, $u \cdot i = i \cdot u$.

Here now are the salient implications of these last observations: Suppose that $o \in \text{Vert}(k)$ and that $h^{-1}\iota(o) = o$ for some $h \in H$. By appeal to Lemma 6.8, there exists an element $i_o \in \text{Aut}(T_o)$ such that $i_o(h^{-1}\iota)i_o^{-1} = h^{-1}\iota'$ in $\mathbb{Z}/(n_o\mathbb{Z})$ and so an appeal to the conclusions of the preceding paragraph finds that $h^{-1}\iota = h^{-1}\iota'$ in $\mathbb{Z}/(n_o\mathbb{Z})$. Meanwhile, the conclusions from the preceding two paragraphs imply that $h^{-1}\iota' = (j_{h(o)}j_o^{-1})h^{-1}\iota$ in $\mathbb{Z}/(n_o\mathbb{Z})$. Thus, $j_{h(o)} = j_o$ and so the assignment $o \rightarrow j_o$ is constant on the H -orbits in $\text{Vert}(k)$. This then means that $j_h j^{-1} h^{-1} \subset \times_{o \in \text{Vert}(k+1)} \text{Aut}(T_o)$ for all $h \in H$.

The proof of Lemma 6.8 requires some knowledge of the conditions that allow a given $\iota \in \text{Aut}(T)$ to have a fixed point in O_T . What follows is a digression in six parts to describe both necessary and sufficient conditions on ι . \square

Part 1 The automorphism ι has a lift in $\widehat{\text{Aut}}(T)$ with a fixed point in the space that is depicted in (6–20). Let ι also denote this lift and let b denote a fixed point for this lifted version of ι , thus a point in (6–20). To determine necessary and sufficient conditions for ι to fix b , it proves useful to make the choices that are described in Parts 1 and 2 of Subsection 6.C so as to view b as a point in the space that is depicted in (6–15). In this incarnation, b is a tuple $(\tau_-, (\tau_\diamond, r_\diamond), (r_o, \tau_o)_{o \in \mathcal{V}})$ where τ_- and τ_\diamond are real numbers, each r_o is in the corresponding simplex Δ_o and $(\tau_o)_{o \in \mathcal{V}} \in (\times_o \mathbb{R}_o)/(\times_o \mathbb{Z}_o)$.

What follows summarizes how the choices from Parts 1 and 2 of Subsection 6.C identify (6–20) with (6–15). To start, note that the \mathbb{R}^Δ factor of (6–20) is a pair whose second

component is r_\diamond and whose first is a map from $\text{Vert}_{\widehat{E}}$ to \mathbb{R} . The value of this map on the vertex \widehat{v} from Part 2 in Subsection 6.C gives the \mathbb{R}_\diamond factor in (6–15). The correspondence between the remaining factors in (6–20) and (6–15) uses an assigned component of the inverse image in $\overline{\Gamma}^*$ of each $\ell^*_{\diamond(\cdot)}$. Use ℓ^e to denote the component that is assigned to an edge $e \in E_\diamond$. The value on ℓ^e of the U_e factor in (6–20) gives the factor $\times_{o \in \mathcal{V}(e)} \mathbb{R}_o$ in (6–15). Meanwhile, those $\{\ell^e\}_{e \in \widehat{E}}$ that are indexed by the edges in \widehat{E} provide a canonical element in Λ and the values of (6–20)’s \mathbb{R}^- factor on the latter provides the \mathbb{R}_- factor in (6–15).

The definition of ℓ^e is that used in Part 8 of Subsection 6.D. In brief, the concatenating path from v_\diamond to $\ell_{\diamond e}$ provides a unique path in $\overline{\Gamma}^*$ from \widehat{v} to a vertex that projects to $\ell^*_{\diamond e}$. The latter vertex sits on ℓ^e .

Part 2 The most straightforward aspect of the fixed point condition concerns the collection $(r_o)_{o \in \mathcal{V} \cup \diamond}$. In particular, if ι fixes b , then

$$(6-37) \quad r_{\iota(o)}(\iota(\gamma)) = r_o(\gamma)$$

for each $o \in \mathcal{V} \cup \diamond$ and for each arc $\gamma \subset \underline{\Gamma}_o$. This is to say that ι fixes b ’s image in $\times_o \Delta_o$.

Part 3 Consider next the condition that $\iota(\tau_\diamond) = \tau_\diamond$. For this purpose, construct a path in $\underline{\Gamma}_\diamond$ that starts at v_\diamond and ends at $\iota^{-1}(v_\diamond)$ as follows: Let e denote the distinguished vertex in E_\diamond and let $\widehat{e} \equiv \iota^{-1}(e)$. Introduce $\{v_1, \dots, v_N\}$ to denote the chosen concatenating path set for the edge \widehat{e} . Thus, the last vertex on v_N lies on $\ell_{\diamond \widehat{e}}$. Let ν denote the path that starts with the final arc in v_N as it is traversed while traveling v_N ; after running along this arc, ν then proceeds in the oriented direction along $\ell_{\diamond \widehat{e}}$ to $\iota^{-1}(v_\diamond)$. The concatenating path set $\{v_1, \dots, v_N, \nu\}$ defines a unique path in $\overline{\Gamma}^*$ that starts at \widehat{v} and ends at a vertex on $\ell^{\widehat{e}}$ that projects to $\iota^{-1}(v_\diamond)$. Let \widehat{v}_ι denote this ending vertex. Meanwhile, let $\mu \subset \underline{\Gamma}_\diamond$ denote the path that is obtained from $\{v_1, \dots, v_N, \nu\}$ by identifying the final arc in each v_k with the initial arc in the subsequent path. Note that μ is the projection to $\underline{\Gamma}_\diamond$ of the path in $\overline{\Gamma}^*$ that was used to define \widehat{v}_ι .

The vertex $\iota^{-1}(\widehat{v})$ is some $\mathbb{Z} \times \mathbb{Z}$ translate of \widehat{v}_ι , thus, $N\widehat{v}_\iota$ with $N \in \mathbb{Z} \times \mathbb{Z}$. Granted all of this, the condition for ι to fix τ_\diamond is:

$$(6-38) \quad \sum_{\gamma \subset \mu} \pm r_\diamond(\gamma) = \alpha_N(\theta_\diamond),$$

where the sum is over the arcs in μ , and where the $+$ sign is used if and only if the arc is crossed in its oriented direction. Note that this condition concerns only r_\diamond and ι ’s image in $\widehat{\text{Aut}}_\diamond$. In particular, it says nothing about the value of τ_\diamond .

Part 4 The conditions that ι must satisfy to fix the rest of b involve an integer pair that is defined for each of \diamond 's incident edges. This integer pair is defined modulo $\mathbb{Z} \cdot Q_e$ as follows: If $e \in E_\diamond$, then $\iota(\ell^e)$ is a component of the inverse image of $\ell_{\diamond\iota(e)}^*$ and so is equal to a $\mathbb{Z} \times \mathbb{Z}$ translate of $\ell^{(e)}$. Let $R_e = (r_e, r'_e)$ denote such a pair. The condition $\iota(\tau_-) = \tau_-$ involves the collection $\{R_e\}_{e \in \hat{E}}$:

$$(6-39) \quad \sum_{e \in \hat{E}} (r'_e q_{\hat{E}} - r_e q_{\hat{E}'}) = 0.$$

Thus, this condition concerns only ι 's image in $\hat{\text{Aut}}_\diamond$; it says nothing about τ_- .

By the way, a particular version of R_e is $-(N_e + N)$, where N is the integer pair defined in Part 3 and where N_e is defined as in Part 8 of Subsection 6.D.

Part 5 Granted that ι fixes τ_- , τ_\diamond and (r_o) , consider next the conditions for ι 's fixing of $(\tau_o)_{o \in \mathcal{V}}$. In this regard, there are two cases to consider, the first where ι fixes a given vertex $\hat{o} \in \mathcal{V}$ and the second where $\iota(\hat{o}) \neq \hat{o}$. This part tells the story in the case that ι fixes \hat{o} . For this purpose, keep in mind the following feature of T : If \hat{o} is fixed by ι , then so is o if $\hat{o} \in T_o$. With the preceding, let $\iota_{\hat{o}} \in \{0, \dots, n_{\hat{o}} - 1\} = \mathbb{Z}/(n_o\mathbb{Z})$ denote the image of ι . Then \hat{o} 's version of (6-37) requires that ι acts so that

$$(6-40) \quad \tau_{\hat{o}} \rightarrow \tau_{\hat{o}} - \frac{2\pi}{\alpha_{Q_{e(\hat{o})}}(\theta_{\hat{o}})} \left(\frac{\iota_{\hat{o}}}{n_{\hat{o}}} \alpha_{Q_{e(\hat{o})}}(\theta_{\hat{o}}) + \sum_o \frac{\iota_o}{n_o} \alpha_{Q_{e(o)} - Q_{e'(o)}}(\theta_{\hat{o}}) + \alpha_{R_e}(\theta_{\hat{o}}) \right),$$

where the vertex labeled sum involves only vertices $o \in \mathcal{V} - \hat{o}$ where T_o contains \hat{o} . To explain the notation, $e(o)$ and $e'(o)$ are both incident edges to o , the former connecting o to $T - T_o$ and the latter connecting \hat{o} 's component of $T_o - 0$ to o . Finally, e is the incident edge to \diamond that connects \hat{o} 's component of $T - \diamond$ to \diamond .

There is a convenient way to rewrite (6-40) that uses the following observation about T : Each vertex in $T - \diamond$ is a monovalent vertex in a unique subgraph of T whose second monovalent vertex is \diamond . Moreover, if the given vertex is in generation k , then there are k vertices on this graph and their generation numbers increase by 1 as they are successively passed as the graph is traversed from \diamond . The interior vertices in the \hat{o} version of this graph are precisely the vertices that appear in the sum on the right hand side of (6-40). In this regard, if these interior vertices are labeled as $\{o_1, \dots, o_k = \hat{o}\}$ by their generation number, then $e(o_j) = e'(o_{j-1})$. This understood, then (6-40)

implies that ι acts on $\tau_{\hat{o}}$ by adding

$$(6-41) \quad -\frac{2\pi}{\alpha_{Q_e(\hat{o})}(\theta_{\hat{o}})} \left(\frac{\iota_{o_k}}{n_{o_k}} - \frac{\iota_{o_{k-1}}}{n_{o_{k-1}}} \right) \alpha_{Q_e(o_k)}(\theta_{\hat{o}}) + \dots \\ + \left(\frac{\iota_{o_2}}{n_{o_2}} - \frac{\iota_{o_1}}{n_{o_1}} \right) \alpha_{Q_e(o_2)}(\theta_{\hat{o}}) + \left(\frac{\iota_{o_1}}{n_{o_1}} \alpha_{Q_e}(\theta_{\hat{o}}) + \alpha_{R_e}(\theta_{\hat{o}}) \right).$$

This last formula has the following consequence when applied to \hat{o} and to each vertex from the collection $\{o_j\}_{1 \leq j < k}$: Each $o \in \{o_1, \dots, o_{k-1}, o_k = \hat{o}\}$ version of τ_o is fixed by ι if and only if there exists a collection, $\{c_1, \dots, c_k\}$, of integers such that

$$(6-42) \quad \left(\frac{\iota_{o_1}}{n_{o_1}} - c_1 \right) \alpha_{Q_e}(\theta_{o_1}) + \alpha_{R_e}(\theta_{o_1}) = 0, \\ \left(\frac{\iota_{o_2}}{n_{o_2}} - \frac{\iota_{o_1}}{n_{o_1}} - c_2 \right) \alpha_{Q_e(o_2)}(\theta_{o_2}) + \left(\frac{\iota_{o_1}}{n_{o_1}} - c_1 \right) \alpha_{Q_e}(\theta_{o_2}) + \alpha_{R_e}(\theta_{o_2}) = 0, \\ \text{and so on through} \\ \left(\frac{\iota_{o_k}}{n_{o_k}} - \frac{\iota_{o_{k-1}}}{n_{o_{k-1}}} - c_k \right) \alpha_{Q_e(o_k)}(\theta_{o_k}) + \dots + \left(\frac{\iota_{o_2}}{n_{o_2}} - \frac{\iota_{o_1}}{n_{o_1}} - c_2 \right) \alpha_{Q_e(o_2)}(\theta_{o_k}) + \\ \left(\frac{\iota_{o_1}}{n_{o_1}} - c_1 \right) \alpha_{Q_e}(\theta_{o_k}) + \alpha_{R_e}(\theta_{o_k}) = 0.$$

Part 6 The next case to consider is that where \hat{o} is not fixed by ι . As before, use k for \hat{o} 's generation number. In this case, the effect of ι is to map $\tau_{\hat{o}}$ to $\mathbb{R}_{\iota(\hat{o})}$ and this map has the schematic form

$$(6-43) \quad \tau_{\hat{o}} \rightarrow \tau_{\hat{o}} - 2\pi \frac{\alpha_{Z_{\hat{o}}}(\theta_{\hat{o}})}{\alpha_{Q_e(\hat{o})}(\theta_{\hat{o}})} - \frac{2\pi}{\alpha_{Q_e(\hat{o})}(\theta_{\hat{o}})} \sum_{\gamma} r_{\hat{o}}(\gamma) \pmod{2\pi\mathbb{Z}},$$

where $Z_{\hat{o}}$ is an ordered pair of rational numbers that is determined by the image of ι in $\widehat{\text{Aut}}(T)/(\times_{o \in \text{Vert}(k)} \text{Aut}(T_o))$. For example, when \hat{o} is a first generation vertex, then $Z_{\hat{o}} = R_e$ with e here denoting the edge that contains both \hat{o} and \diamond . Meanwhile, in all cases, the sum in (6-43) is indexed by the arcs in $\ell_{\hat{o}}$ that are crossed when traveling in the oriented direction from $\iota^{-1}(v_{\iota(\hat{o})})$ to $v_{\hat{o}}$.

Proof of Lemma 6.8 The automorphisms ι and ι' have identical actions on $\text{Vert}(k)$ since their actions on this set are determined by their image in $\widehat{\text{Aut}}(T)/(\times_{o \in \text{Vert}(k)} \text{Aut}(T_o))$. This understood, if ι fixes a given $\hat{o} \in \text{Vert}(k)$, then so does ι' and it follows from (6-42) that ι and ι' have identical images in $\mathbb{Z}/(n_{\hat{o}}\mathbb{Z})$.

Suppose next that $\hat{o} \in \text{Vert}(k)$ is not fixed by ι . To study this case, note first that any $j \in \times_{o \in \text{Vert}(k)} \text{Aut}(T_o)$ has an image in each $o \in \text{Vert}(k)$ version of $\mathbb{Z}/(n_o\mathbb{Z})$, and this image is denoted in what follows as j_o , a number from the set $\{0, \dots, n_o - 1\}$. Meanwhile, the $\widehat{\text{Aut}}(T)$ version of (6–13) assigns both ι and ι' integers in this same set, these are their respective factors in the $\mathbb{Z}/(n_o\mathbb{Z})$ summand. The latter are denoted here by ι_o and ι'_o . This understood, there exists $j \in \times_{o \in \text{Vert}(k)} \text{Aut}_o$ such that $j \iota j^{-1} = \iota'$ in $\text{Aut}(T)/(\times_{o \in \text{Vert}(k+1)} \text{Aut}_o)$ when there exists a collection $\{j_o \in \mathbb{Z}/(n_o\mathbb{Z})\}_{o \in \text{Vert}(k)}$ such that

$$(6-44) \quad \iota'_o - \iota_o = -(j_{\iota(o)} - j_o)$$

for each $o \in \text{Vert}(k)$. To see that this equation is solvable, let $\hat{o} \in \text{Vert}(k)$, let M denote \hat{o} 's orbit under the action of ι , and let m denote the number of elements in M . If M consists only of \hat{o} , then by virtue of what was said in the previous paragraph, $\iota_o = \iota'_o$ and one can take $j_o = 0$. In the case that $m > 1$, then (6–44) is solvable if and only if

$$(6-45) \quad \sum_{o \in M} (\iota_o - \iota'_o) = 0 \pmod{2\pi n_{\hat{o}}\mathbb{Z}}.$$

To prove (6–45), remark that ι^m and ι'^m both fix \hat{o} and so must be equal in $\mathbb{Z}/(n_{\hat{o}}\mathbb{Z})$. As ι^m acts in $\mathbb{Z}/(n_{\hat{o}}\mathbb{Z})$ as $\sum_{o \in M} \iota_o$ and ι'^m acts as $\sum_{o \in M} \iota'_o$, the equality in (6–45) follows. \square

Proof of Proposition 6.7 Suppose that $G \subset \text{Aut}(T)$ stabilizes some point. The first observation here stems from Proposition 6.4: Since G is isomorphic to its image in $\text{Aut}_{\diamond} = \mathbb{Z}/(n_{\diamond}\mathbb{Z})$, it must be a cyclic group with one generator. To see what this means, let e denote the distinguished incident edge to \diamond , now fixed by $\text{Aut}(T)$. Let $\iota \in G$ be the generator. As before, use ι to also denote the lift to $\widehat{\text{Aut}}(T)$ and let b denote the point in (6–15) that is fixed by this lift.

Let (z, N) denote ι 's image in (6–32). The assumption that ι fixes b has the following consequence: The factor in \mathbb{R}_- is fixed if and only if

$$(6-46) \quad N = \frac{k_-}{m_-} Q_e,$$

where $k_- \in \mathbb{Z}$ and where m_- is the greatest common divisor of the pairs that comprise Q_e . Meanwhile, b 's factor in \mathbb{R}_{\diamond} is then fixed if and only if

$$(6-47) \quad \frac{\iota_{\diamond}}{n_{\diamond}} + \frac{k_-}{m_-} + z = 0,$$

where $\iota_\diamond \in \{0, \dots, n_\diamond - 1\}$ is ι 's image in $\text{Aut}_\diamond = \mathbb{Z}/(n_\diamond\mathbb{Z})$. This requires that

$$(6-48) \quad \left(\frac{k_-}{m_-} + z \right) = -\frac{\iota_\diamond}{n_\diamond} \pmod{\mathbb{Z}}.$$

To proceed, suppose next that ι fixes some vertex $\hat{o} \in T - \diamond$ and that \hat{o} is in generation $k \geq 1$. Let $\{o_1, \dots, o_k = \hat{o}\}$ denote the vertices in the \hat{o} version of (6-42). By virtue of (6-35), these equations now imply that ι fixes $\tau_{\hat{o}}$ if and only if there are integers $\{c_1, \dots, c_k\}$ such that

$$(6-49) \quad \begin{aligned} & \left(\frac{\iota_{o_1}}{n_{o_1}} - \frac{\iota_\diamond}{n_\diamond} - c_1 \right) \alpha_{Q_{e(o_1)}}(\theta_{o_1}) + \left(\frac{\iota_\diamond}{n_\diamond} + \frac{k}{m_-} + z \right) \alpha_{Q_e}(\theta_{o_1}) = 0. \\ & \left(\frac{\iota_{o_2}}{n_{o_2}} - \frac{\iota_{o_1}}{n_{o_1}} - c_2 \right) \alpha_{Q_{e(o_2)}}(\theta_{o_2}) + \left(\frac{\iota_{o_1}}{n_{o_1}} - \frac{\iota_\diamond}{n_\diamond} - c_1 \right) \alpha_{Q_{e(o_1)}}(\theta_{o_2}) \\ & + \left(\frac{\iota_\diamond}{n_\diamond} + \frac{k_-}{m_-} + z \right) \alpha_{Q_e}(\theta_{o_2}) = 0, \text{ and so on through} \\ & \left(\frac{\iota_{o_k}}{n_{o_k}} - \frac{\iota_{o_{k-1}}}{n_{o_{k-1}}} - c_k \right) \alpha_{Q_{e(o_k)}}(\theta_{o_k}) + \dots + \left(\frac{\iota_{o_2}}{n_{o_2}} - \frac{\iota_{o_1}}{n_{o_1}} - c_2 \right) \alpha_{Q_{e(o_2)}}(\theta_{o_k}) \\ & + \left(\frac{\iota_{o_1}}{n_{o_1}} - \frac{\iota_\diamond}{n_\diamond} - c_1 \right) \alpha_{Q_{e(o_1)}}(\theta_{o_k}) + \left(\frac{\iota_\diamond}{n_\diamond} + \frac{k_-}{m_-} + z \right) \alpha_{Q_e}(\theta_{o_k}) = 0. \end{aligned}$$

Here, a given version of ι_o denotes ι 's image in $\mathbb{Z}/(n_o\mathbb{Z}) = \{0, \dots, n_o - 1\}$. Coupled with (6-47), this set of equations is satisfied if and only if each of the k versions of ι_o/n_o that appear here are identical. The preceding fact together with (6-48) implies that ι generates a canonical diagonal subgroup of $\text{Aut}(T)$.

The remaining assertions of Proposition 6.7 follow directly from the preceding analysis with the various versions of the formula in (6-43). The details here are straightforward and so left to the reader. □

7 Proof of Theorems 6.2 and 6.3

As the heading indicates, the purpose of this section is to supply the proofs of the main theorems from the previous section. The arguments to this end are much like those used to prove Theorem 3.1. In fact, there are many places where the arguments transfer in an almost verbatim form and, except for a comment to this effect, these parts are left to the reader. In any event, to start, use \mathfrak{X} in what follows to denote the map that is defined in Subsection 6.C. Theorem 6.2 is proved by establishing the following about \mathfrak{X} :

- (7–1)
- The map \mathfrak{X} lifts on some neighborhood of any given element in $\mathcal{M}^*_{\hat{A},T}$ as a local diffeomorphism from $\mathcal{M}^*_{\hat{A},T}^\Lambda$ onto an open set in $\mathbb{R} \times O_T$.
 - The map \mathfrak{X} is 1–1 into $\mathbb{R} \times O_T / \text{Aut}(T)$.
 - $\mathcal{R} \cap \mathcal{M}^*_{\hat{A},T}$ is sent to $\mathbb{R} \times (O_T - \hat{O}_T) / \text{Aut}(T)$ and its complement to $\mathbb{R} \times \hat{O}_T / \text{Aut}(T)$.
 - The map \mathfrak{X} is proper onto $\mathbb{R} \times O_T / \text{Aut}(T)$.

Here, $\mathcal{M}^*_{\hat{A},T}^\Lambda$ is defined as in Theorem 6.3 to be the space of equivalence class of triples (C_0, ϕ, T_C) where (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A},T}$ and T_C is a correspondence of T in (C_0, ϕ) .

The first subsection below proves Theorem 6.3 that given \mathfrak{X} supplies Theorem 6.2's diffeomorphism.

7.A The local structure of the map \mathfrak{X}

The arguments given in this subsection justify the first point in (7–1) and the assertions of Theorem 6.3. To start, note that Lemma 6.5 asserts that the map \mathfrak{X} lifts to define a continuous map from $\mathcal{M}^*_{\hat{A},T}^\Lambda$ to $\mathbb{R} \times O_T$. This understood, Proposition 5.1 and Lemma 5.4 imply that the lift is smooth and locally 1–1. The first point in (7–1) follows directly from this last conclusion.

Proof of Theorem 6.3 for the map \mathfrak{X} For the purposes of this proof, assume that the map \mathfrak{X} is one of Theorem 6.2's diffeomorphisms. The proof of Theorem 6.3 amounts to no more than verifying the asserted properties of \mathfrak{X} . This task is left to the reader with the following comment: Given the form of the parametrizations from Definition 2.1, all of these properties are direct consequences of the definitions given in Subsection 6.C. □

7.B Why the map \mathfrak{X} is 1–1 from $\mathcal{M}^*_{\hat{A},T}$ to $\mathbb{R} \times O_T / \text{Aut}(T)$

The proof that \mathfrak{X} is 1–1 applies Lemma 4.1 as in Subsection 4.B. To start, suppose that (C_0, ϕ) and (C'_0, ϕ') have the same image in $\mathbb{R} \times O_T / \text{Aut}(T)$. By assumption, T has respective correspondences, T_C and $T_{C'}$, in (C_0, ϕ) and in (C'_0, ϕ') and these can be fixed so that (C_0, ϕ) and (C'_0, ϕ') have the same image in $\mathbb{R} \times O_T$. This is because any correspondence of T in (C_0, ϕ) can be obtained from any other by composing the original with a suitable automorphism of T . And, noted in Lemma 6.6, such a change in the correspondence changes the assigned point in $\mathbb{R} \times O_T$ by the action of the relevant automorphism.

To say more, the choices in Parts 1 and 2 of Subsection 6.C should be used to identify O_T with the space in (6–9). With the identification understood, the assignments to (C_0, ϕ) and (C_0, ϕ') of their respective points in

$$(7-2) \quad \mathbb{R}_- \times (\mathbb{R}_\diamond \times \Delta_\diamond) \times (\times_{o \in \mathcal{V}} (\mathbb{R}_o \times \Delta_o))$$

can be made so as to agree. This is done in a sequential fashion by first arranging that the respective \mathbb{R}_- assignments agree. For the latter purpose, let e denote the distinguished edge in the distinguished Aut_\diamond orbit \hat{E} . The assignments to \mathbb{R}_- can be made equal by suitably choosing the respective parameterizations of e 's component in the C and C' versions of $C_0 - \Gamma$. With the \mathbb{R}_- assignments equal, the next step is to arrange so that the \mathbb{R}_\diamond assignments agree. Since the $\mathbb{Z} \cdot Q_{\hat{E}}$ subgroup of $\mathbb{Z} \times \mathbb{Z}$ acts trivially on \mathbb{R}_- and since its action on \mathbb{R}_\diamond changes the assignment by the action of $2\pi\mathbb{Z}$, the C' version of the lift to \mathbb{R} that defines its point in \mathbb{R}_\diamond can be chosen to make it agree with C 's assigned point. When applied inductively, versions of this last argument prove that any given $o \in \mathcal{V}$ version of the two assignments in \mathbb{R}_o can be made to agree. Here, the induction moves from any given $o \in \mathcal{V}$ to the vertices that share its edges in T_o .

Given that C and C' have the same assignments in (7–2), then minor modifications of the arguments used in Part 2 of Subsection 4.B prove that the conditions in (4–1) are met in this case.

The appeal to Lemma 4.1 also requires the condition in (4–2). To see why this condition holds, suppose that o is a multivalent vertex in T and that $\gamma \subset \underline{\Gamma}_o$ is an arc. Now let e denote one of the edge labels on γ and let e' denote the other. The choices made for the (C_0, ϕ) assignment to the space in (7–2) give parametrizations to both the e and e' components of $C_0 - \Gamma$. These define C_0 versions of the functions w_e and $w_{e'}$ whose difference along the interior of γ 's image in Γ is described by some $N = (n, n')$ version of (2–14) and (2–15). There are corresponding C_0' versions of w_e and $w_{e'}$, and the point here is that their difference is also described by this same $N = (n, n')$ version of (2–14) and (2–15). Indeed, such is the case because the integer N is obtained by using Lemma 2.3 to compare the respective canonical parametrizations of the e' component of either version of $C_0 - \Gamma$ and its C_0' counterpart with the parametrization that gives the $N = 0$ case of (2–14) and (2–15) across γ . Since the same integer pair appears in both the C_0 and C_0' versions of (2–14) and (2–15), so $\hat{w}_e = \hat{w}_{e'}$ along γ .

Given that the conditions in (4–1) and (4–2) have been met, the graph G is well defined. If $G \neq \emptyset$, then Lemma 4.1 asserts that $C_0' = C_0$ and that ϕ' is obtained from ϕ by a constant translation along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. This means that $\phi = \phi'$ since they share the same assignment in the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. To see that

$G \neq \emptyset$, return to the respective C_0 and C_0' versions of (6–27). Subtraction of the C_0' version from the C_0 version finds the value 0 for the integral of \hat{w} around a non-empty union of constant θ circles in $C_0 - \Gamma$. Either $\hat{w} = 0$ on some of these circles, in which case $G \neq \emptyset$, or else \hat{w} is negative on some and positive on others. But this last case also implies that $G \neq \emptyset$ since the complement in C_0 of the critical point set of $\cos \theta$ is path connected.

7.C The images of $\mathcal{M}^*_{\hat{A},T} \cap \mathcal{R}$ and $\mathcal{M}^*_{\hat{A},T} - \mathcal{R}$.

The first order of business here is to explain why $\mathcal{M}^*_{\hat{A},T} \cap \mathcal{R}$ is mapped to the image in $\mathcal{R} \times O_T / \text{Aut}(T)$ of the points where $\text{Aut}(T)$ acts with non-trivial stabilizer. For this purpose, suppose that (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A},T}$ and that there is a non-trivial group of holomorphic diffeomorphisms of C_0 that fix ϕ . Let ψ denote an element in this group. Let T_C denote a correspondence for T in (C_0, ϕ) . Then T_C defines a point for (C_0, ϕ) in $\mathbb{R} \times O_T$. Meanwhile, ψ defines a new correspondence for T in (C_0, ϕ) as follows: If e is an edge of T and $K_e \subset C_0 - \Gamma$ the component that originally corresponds to e , then $\psi^{-1}(K_e)$ gives the component for e from the new correspondence. Likewise, if o is a multivalent vertex and γ is an arc in $\underline{\Gamma}_o$, then the new correspondence assigns to γ the ψ -inverse image of what is assigned γ by the old correspondence. As noted in Lemma 6.6, the change of the correspondence changes the point in $\mathbb{R} \times O_T$ that is assigned to (C_0, ϕ) by the action of ι . Meanwhile, Lemma 6.5 asserts that the new assigned point is, after all, the same as the original. Thus the assigned point in $\mathbb{R} \times O_T$ is fixed by ι .

The argument as to why $\mathcal{M}^*_{\hat{A},T} - \mathcal{R}$ is mapped to $\mathcal{R} \times \hat{O}_T / \text{Aut}(T)$ is much like the argument in Subsection 4.D. To start, suppose that (C_0, ϕ) defines a point in $\mathcal{M}^*_{\hat{A},T} - \mathcal{R}$. Fix a correspondence, T_C , for T in (C_0, ϕ) . According to Lemma 6.5, this gives (C_0, ϕ) a point in O_T that defines the image of its equivalence class in $O_T / \text{Aut}(T)$. Now suppose that $\iota \in \text{Aut}(T)$ fixes this point in O_T . The assertion to prove is that ι is the identity element. The proof has three steps.

Step 1 The isomorphism ι can be used to change the correspondence T_C to a new correspondence for T in (C_0, ϕ) . The latter is denoted in what follows as T_C^ι and it is defined by the following two conditions: To state the first, let e denote an edge in T and use K_e to denote the component of $C_0 - \Gamma$ that is labeled by e using T_C . Meanwhile use K_e^ι to denote the component that is labeled by e using the new correspondence. Define the collection $\{K_e^\iota\}$ so that $K_e^\iota = \iota(K_e)$. To state the second condition, let γ denote an arc in some version of $\underline{\Gamma}_o$. Then the arc in Γ that corresponds via T_C to γ corresponds via T_C^ι to $\iota(\gamma)$.

Now, the assignment to C of a point in (7–2) using the correspondence T_C required parametrizations for each component of $C_0 - \Gamma$. Recall that these were assigned in a sequential manner starting with an arbitrary choice for the component that corresponds to the distinguished incident edge to \diamond . Such a choice and the assignment in \mathbb{R}_\diamond assigned parametrizations to all incident edges to \diamond . The sequential nature of the process appears when the parametrizations were assigned to the components of $C_0 - \Gamma$ that are labeled by the incident edges to a vertex $o \in \mathcal{V}$. In particular, the point in \mathbb{R}_o and the parametrization for the component labeled by the edge that connects o to $T - T_o$ give the parametrizations for the components that are labeled by the edges that connect o to $T_o - o$.

Make the choices that assign a point in (7–2) for C using the correspondence T_C . Let b denote this point in (7–2). Meanwhile, a second set of choices can be made so that b is also the assigned point in (7–2) for C using the correspondence T_C^t . In this way, each component of $C_0 - \Gamma$ receives two parametrizations; one in its incarnation as K_e and the other as K_{t_e} . This understood, define a map, $\psi_e: K_e \rightarrow K_{t_e}$ as follows: Let ϕ_e denote the parametrizing map from the relevant cylinder to K_e , and let $\phi_{t_e}^t$ denote the corresponding map to K_{t_e} . Note that the parametrizing cylinders are the same. Set $\psi_e \equiv \phi_{t_e}^t \circ (\phi_e)^{-1}$.

Step 2 This step proves that there is a continuous map $\psi: C_0 \rightarrow C_0$ whose restriction to any given K_e is the map ψ_e . The proof has two parts; the first verifies that any given ψ_e extends continuously to the closure of K_e . The second verifies that corresponding extensions agree where the closures overlap.

To start the first part, remark that ψ_e extends continuously to the closure of K_e when the following is true: Let $o \in e$ be a multivalent vertex and let v be any vertex in ℓ_{oe} . Then the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the point on the $\sigma = \theta_o$ boundary of the parametrizing cylinder that corresponds via ϕ_e to v is the same as the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the point that corresponds via $\phi_{t_e}^t$ to v . Here, one need check this condition at only one vertex since the condition in (6–37) makes the difference between the relevant $\mathbb{R}/(2\pi\mathbb{Z})$ values into a constant function on the vertex set of ℓ_{oe} .

To check this condition, consider first the case when either $o = \diamond$ and e is \diamond 's distinguished edge, or else $o \in \mathcal{V}$ and e is the edge that connects o to $T - T_o$. In these cases, ℓ_{oe} contains the distinguished vertex $v_o \in \underline{\Gamma}_o$. In particular, the value of the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of the point that corresponds via ϕ_e to v_o is the reduction modulo $2\pi\mathbb{Z}$ of the point, b , that C is assigned in (7–2) using the correspondence T_C . Meanwhile, the analogous $\phi_{t_e}^t$ version of this $\mathbb{R}/(2\pi\mathbb{Z})$ value is the reduction modulo $2\pi\mathbb{Z}$ of the point that C is assigned in (7–2) using the correspondence T_C^t . Since the latter point is also b , so the desired equality holds for the relevant $\mathbb{R}/(2\pi\mathbb{Z})$ values.

Consider next the case where either $o = \diamond$ and e is not the distinguished edge, or else $o \in \mathcal{V}$ and e connects o to $T_o - o$. To analyze this case, remember that a concatenating path set has been chosen whose first path starts at v_o and whose last path ends at a vertex on ℓ_{oe} . Let $\mu(e) \subset \underline{\Gamma}_o$ denote the path that is obtained from the concatenating path set by identifying the final arc in all but the last path with the initial arc in the subsequent path. Note that $\mu(e)$ is a concatenated union of arcs whose last vertex is on ℓ_{oe} . In particular, the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate that corresponds via ϕ_e to this last vertex is

$$(7-3) \quad \tau_o + \frac{2\pi}{\alpha_{Q_e}(\theta_o)} \sum_{\gamma \in \mu(e)} \pm r_o(\gamma) \pmod{(2\pi\mathbb{Z})},$$

where the notation is as follows: First, τ_o is b 's coordinate in the \mathbb{R}_o factor in (7-2) and r_o is b 's Δ_o factor. Meanwhile, the sum is over the arcs that are met sequentially in $\mu(e)$ as it is traversed from v_o to its end, and where the $+$ sign appears with an arc if and only if it is crossed in its oriented direction. In this regard, note that a given arc can appear more than once in (7-3), and with different signs in different appearances.

Of course, the analogous formula gives $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate that corresponds via ϕ^t_e to the final vertex on μ_e . Thus, the ϕ_e and ϕ^t_e versions of the relevant $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate agree.

To verify that the extensions of $\{\psi_e\}$ agree on the locus $\Gamma \subset C_0$, suppose that $o \in T$ is a multivalent vertex, that $\gamma \subset \underline{\Gamma}_o$ is an arc, and that e and e' are γ 's labeling edges. The correspondence T_C identifies γ with a component of the complement in Γ of θ 's critical points. Denote this component by γ^δ . Of course, γ also corresponds via T_C^t to some other component, $\gamma^{t\delta} = (\iota^{-1}(\gamma))^\delta$. The extension of ψ_e maps γ^δ to $\gamma^{t\delta}$ as a diffeomorphism because the $\mathbb{R}/(2\pi\mathbb{Z})$ coordinate of a point that corresponds via ϕ_e to either of its end vertices is the same as the corresponding ϕ^t_e coordinate. Likewise, $\psi_{e'}$ maps γ^δ diffeomorphically onto $\gamma^{t\delta}$. The agreement between these two diffeomorphisms then follows directly from the formula in Definition 2.1. In this regard, keep in mind that the 1-form $\sqrt{6 \cos \theta} d\varphi - (1 - 3 \cos^2 \theta) dt$ pulls back to the $\sigma = \theta_o$ circle in the parametrizing cylinder for K_e as the $Q = Q_e$ version of $\alpha_Q(\theta_o) dv$ while its pull-back to the circle in the $K_{e'}$ cylinder is the $Q = Q_{e'}$ version.

Step 3 Let $\phi: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ denote the tautological map onto C . With ψ now defined, the plan is to prove that $\phi \circ \psi = \phi$. Since $(C_0, \phi) \notin \mathcal{R}$, this then means that ψ is the identity map and so ι is trivial.

The argument for this employs Lemma 4.1. To conform with the notation used in Lemma 4.1, let C_0' denote C_0 and T_C^t as $T_{C'}$. This understood, the fact that ψ is

continuous implies that the respective parametrizations ϕ_e and ϕ'_e are compatible in the sense of the definition in (4–1). Next, let (a_e, w_e) denote the functions that appear in the ϕ_e version of (2–5) and let $(a_{e'}, w_{e'})$ denote the functions that appear in the ϕ'_e version of (2–5). Now set $\hat{a}_e \equiv a_e - a_{e'}$ and $\hat{w}_e \equiv w_e - w_{e'}$. As it turns out, there is a continuous function on the complement in C_0 of the $\cos \theta$ critical points whose pull-back from $C_0 - \Gamma$ via the maps $\{\phi_e\}$ gives the collection $\{\hat{w}_e\}$. This is to say that the assumption in (4–2) is satisfied. Indeed, the proof that the collection $\{\hat{w}_e\}$ comes from a continuous function is identical in all but cosmetics to the proof of the analogous assertion in the second to last paragraph of the preceding subsection.

Let \hat{w} denote the function on C_0 that gives the collection $\{\hat{w}_e\}$ and introduce G to denote the zero locus of \hat{w} in the complement of the $\cos \theta$ critical points. According to Lemma 4.1, either $G = \emptyset$ or the closure of G is all of C_0 . Now, $G \neq \emptyset$ since the respective assignments to C_0 using T_C and $T_{C'}$ in (7–2) agree, and so they agree in the \mathbb{R}_- factor in particular. Thus G 's closure is C_0 and so $\hat{w} \equiv 0$. Thus $\phi \circ \psi$ is obtained from ϕ via a constant translation along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. Moreover, this constant translation must be the identity since the two maps have the same image set. Hence $\phi \circ \psi = \phi$ and so ι must be the identity automorphism.

7.D Why the map to $\mathbb{R} \times \hat{O}_T / \text{Aut}(T)$ is proper

The proof that Subsection 6.C's map is proper is given here in three parts, each the analog of the corresponding part of Subsection 4.E.

Part 1 This part of the argument proves the $\mathcal{M}^*_{\hat{A}, T}$ version of Proposition 4.6. This is as follows:

Proposition 7.1 *Let $\{(C_{0j}, \phi_j)\}_{j=1,2,\dots}$ denote an infinite sequence of pairs that defines a sequence in $\mathcal{M}^*_{\hat{A}, T}$ with convergent image in $\mathbb{R} \times O_T / \text{Aut}(T)$. There exists a subsequence, hence renumbered consecutively 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, pseudoholomorphic, multiply punctured sphere; and these have the following properties with respect to the sequence of subsets $\{C_i \equiv \phi_i(C_{0i})\}$ in $\mathbb{R} \times (S^1 \times S^2)$:*

- $\lim_{j \rightarrow \infty} \int_{C_j} \varpi = \sum_{(S,n) \in \Xi} n \int_S \varpi$ for each compactly supported 2-form ϖ .
- The following limit exists and is zero:

$$(7-4) \quad \lim_{j \rightarrow \infty} \left(\sup_{z \in C_j} \text{dist}(z, \cup_{(S,n) \in \Xi} S) + \sup_{z \in \cup_{(S,n) \in \Xi} S} \text{dist}(C_j, z) \right).$$

Granted for the moment Proposition 7.1, the proof that Subsection 6.C's map is proper is obtained by the following line of reasoning: Start with a sequence $\{(C_{0i}, \phi_i)\}$ with convergent image in $\mathbb{R} \times O_T / \text{Aut}(T)$. The next part of this subsection proves that Proposition 7.1's set Ξ contains only one element. Let (S, n) denote this element. The third part of the subsection proves the following:

- (7–5) • If $n = 1$, then $S \in \mathcal{M}_{\hat{A}, T}$ and $\{C_j\}$ converges to S in $\mathcal{M}_{\hat{A}, T}$.
- If $n > 1$, then there exists a pair (S_n, ϕ) where ϕ sends S_n onto S as an n -fold, branched cover; and this pair defines a point in $\mathcal{M}^*_{\hat{A}, T}$ that is the limit point of the image in $\mathcal{M}^*_{\hat{A}, T}$ of $\{(C_{0j}, \phi_j)\}$.

Thus, a convergent subsequence exists for any sequence in $\mathcal{M}^*_{\hat{A}, T}$ with convergent image in $\mathbb{R} \times O_T / \text{Aut}(T)$ and this property characterizes a proper map.

It is assumed in the rest of this subsection that a correspondence has been chosen for T in each (C_{0j}, ϕ_j) .

Proof of Proposition 7.1 As with Proposition 4.6, all but (7–4) follows from [15, Proposition 3.7]. In this regard, (7–4) is the version of [15, Proposition 3.7] where the compact set involved, K , is replaced by $K = \mathbb{R} \times (S^1 \times S^2)$. In any event, assume that (7–4) does not hold so as to derive some patent nonsense. This derivation has six steps.

Step 1 This step proves that the conclusions in Lemma 4.7 hold if (7–4) does not hold. The argument here is almost the same as that used for the Subsection 4.E's version of the lemma. To start, note that an argument in Subsection 4.E for Lemma 4.7 works as well here to prove that its conclusions hold except possibly in the case that all subvarieties from Ξ are \mathbb{R} -invariant cylinders and that one of the two points in (4–19) hold.

To rule out the (4–19) cases, note first that the smallest vertex angle from T must be zero in the (4–19) cases. The argument is essentially that from Subsection 4.E: The convergence of the image of $\{(C_{0j}, \phi_j)\}$ in the \mathbb{R} factor of $\mathbb{R} \times O_T / \text{Aut}(T)$ precludes a non-zero smallest angle because there would otherwise be a sequence in $\mathcal{M}^*_{\hat{A}, T}$ with the following properties: First, its j 'th element is (C_0, ϕ_j') where ϕ_j' is obtained from ϕ_j by a j -dependent but constant translation along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$. Second, the limit data set for $\{\phi_j'(C_{0j})\}$ from [15, Proposition 3.7] has a subvariety that is not an \mathbb{R} -invariant cylinder, has the same θ infimum as each $C_j = \phi_j(C_{0j})$. Third, there is no non-zero constant b for one of the corresponding ends that makes (2–3) hold.

Granted the preceding, assume that the smallest vertex angle from T is equal to zero. Let \hat{E}' denote the $\text{Aut}(T)$ orbit of edges in T that is used in Subsection 6.C to define the image of $\{(C_{0j}, \phi_j)\}$ in the \mathbb{R} factor of $\mathbb{R} \times \hat{O}_T / \text{Aut}(T)$. Suppose first that \hat{E}' corresponds to a set of disjoint disks in each C_{0j} that intersect the $\theta = 0$ cylinder at their center points. Note that this case occurs only when \hat{A} lacks $(1, \dots)$ elements. To preclude this (4–19) case, first associate to each C_{0j} the set of its $\theta = 0$ points that correspond to the edges in \hat{E}' . Let t_j denote this set. Note that no sequence whose j 'th element is from t_j can remain bounded as $j \rightarrow \infty$. This is to say that the set of $|s|$ coordinates for such a sequence cannot be bounded. Indeed, a bounded sequence here would force \mathbb{E} to have a pair whose subvariety is the $\theta = 0$ cylinder. Granted (4–19), each C_{0j} would have an end where the $|s| \rightarrow \infty$ limit of θ is zero. But there are no such ends when \hat{A} lacks $(1, \dots)$ elements.

Meanwhile, note that as the image of $\{(C_{0j}, \phi_j)\}$ in the \mathbb{R} factor of $\mathbb{R} \times \hat{O}_T / \text{Aut}(T)$ converges, so the sequence whose j 'th component is the sum of the s values at the points in t_j also converges. Granted the conclusions of the preceding paragraph, this can occur only if there are two sequences whose j 'th element is in t_j , and these are such that s tends to ∞ on one and to $-\infty$ on the other. However, this conclusion leads afoul of (4–19) and the lack of $(1, \dots)$ elements in \hat{A} when one considers that there is a path in C_{0j} between the j 'th point in the one sequence and the j 'th point in the other.

To finish the story for the current version of Lemma 4.7, assume now that \hat{A} has some $(1, \dots)$ element and so the edges in \hat{E}' correspond to ends in any given C_{0j} where $\lim_{|s| \rightarrow \infty} \theta = 0$. In this case, the image of C_{0j} in the \mathbb{R} factor of $\mathbb{R} \times \hat{O}_T / \text{Aut}(T)$ is, up to a positive constant factor, minus the sum of the logarithms of the versions of the constant \hat{c} that appear in (1–9) for the ends in C_{0j} that correspond to the edges in \hat{E}' . Thus, the sequence of such sums converge. As a consequence, the sequence whose j 'th element is the minimal contribution to the j 'th sum can not diverge towards $+\infty$, nor can the sequence whose j 'th element is the maximal contribution to the j 'th sum diverge towards $-\infty$. This then gives the following nonsense: As in the analogous case from the original proof of Lemma 4.7, a new sequence can be constructed with the following mutually incompatible properties: First, its j 'th element is (C_{0j}, ϕ_j') where ϕ_j' is obtained from ϕ_j by a constant, but j -dependent translation along the \mathbb{R} factor in $\mathbb{R} \times (S^1 \times S^2)$. Second, the data set from the corresponding $\{\phi_j'(C_{0j})\}$ version of [15, Proposition 3.7] provides a subvariety that is not the $\theta = 0$ cylinder, but has an end where the $|s| \rightarrow \infty$ limit of θ is 0 and is such that the respective integrals of $\frac{1}{2\pi} dt$ and $\frac{1}{2\pi} d\varphi$ about its constant $|s|$ slices have the form $\frac{1}{m} p$ and $\frac{1}{m} p'$ where (p, p') is the integer pair that comes from the elements in \hat{A} that correspond to \hat{E}' , and where $m \geq 1$ is a common divisor of this pair. Finally, there is no non-zero \hat{c} that

for the resulting versions of (1–9). The argument here for the values of the integrals of $\frac{1}{2\pi}dt$ and $\frac{1}{2\pi}d\varphi$ is almost verbatim that used in the analogous part of Section 4's proof of Lemma 4.7. Indeed, the latter argument works because the same domain in $(0, \pi) \times \mathbb{R}/(2\pi\mathbb{Z})$ parametrizes all \widehat{E}' labeled components of all C_{0j} version of $C_0 - \Gamma$; thus this same domain parametrizes all of the \widehat{E}' labeled components of each version of $C_0 - \Gamma$ from the translated sequence.

Step 2 Lemma 4.7 provides the \mathbb{R} -invariant cylinder S_* , and given $\varepsilon > 0$, the real number s_0 , the sequences $\{s_{j-}\}$ and $\{s_{j+}\}$ and the component C_{j*} of the $s \in [s_{j-}, s_{j+}]$ part of each large j version of C_{0j} . The first point to make here is that θ is neither 0 nor π on S_* . Indeed, were this otherwise, then the mountain pass lemma would find a non-extremal critical point of θ on each large j version of C_{0j} in C_{j*} . In particular, the corresponding sequence of critical values would converge as $j \rightarrow \infty$ to either 0 or π , and this is nonsense since the critical values on any one C_{0j} are identical to those on any other.

Thus, θ on S_* must be some angle in $(0, \pi)$. Denote the latter by θ_* . [15, Lemma 3.9] can be used in the present context to draw the following conclusions: Given $\delta_0 > 0$, there exist j -independent constants $\delta \in (0, \delta_0)$ and $R_+, R_- \geq 0$ such that when j is large,

- (7–6) • θ has values both greater than $\theta_* + \delta$ and less than $\theta_* - \delta$ where $s = s_{j-} + R_-$ in C_{j*} .
- Either $|\theta - \theta_*|$ is strictly greater than δ where $s = s_{j+} - R_+$ in C_{j*} or else θ takes values both greater than $\theta_* + \delta$ and $\theta_* - \delta$ where $s = s_{j+} - R_+$ in C_{j*} .

As no generality is lost by choosing R_{\pm} so that both the $s = s_{j-} + R_-$ and $s = s_{j+} - R_+$ are regular values of s on each C_{j*} and that neither locus in C_{j*} contains a θ critical point. Such choices for R_{\pm} are assumed in what follows.

By contrast, the fourth point in Lemma 4.7 implies the existence of a sequence $\{\delta_j\}$ with limit zero such that $|\theta - \theta_*| < \delta_j$ where $s = \frac{1}{2}(s_{j-} + s_{j+})$. Granted all of this, the argument for Proposition 4.6 in Subsection 4.E works here to prove that θ_* is either a critical point of θ on C_{0j} or else the $|s| \rightarrow \infty$ value of θ on an end in C_{j0} whose version of (2–4) has integer $n_{(\cdot)} = 0$. In either case, θ_* is the angle of some multivalent vertex in T and C_{j*} must intersect some graph from the collection $\{\Gamma_o\}$. As C_{j*} is connected and θ has small variation on C_{j*} , it can intersect at most one such graph. Let $o \in T$ denote the labeling vertex.

Step 3 Let s_* denote a regular value of s on C_{j*} that is within 1 of $\frac{1}{2}(s_{j-} + s_{j+})$, has transversal intersection with the $\theta = \theta_*$ locus, and is such that $|\theta - \theta_*| < \delta_j$ on the $s = s_*$ slice of C_{j*} . By virtue of [15, Proposition 3.7], the $s = s_*$ slice of C_{j*} is not null-homologous in the radius ε tubular neighborhood of S_* . Thus this slice is not null-homologous in the complement of the $\theta = 0$ and $\theta = \pi$ loci in C_{0j} . Even so, the following is true: If j is large, then this $s = s_*$ slice is not homologous in the $\theta \notin \{0, \pi\}$ part of C_{0j} to a union of suitably oriented slices of ends of C_{0j} where the $|s| \rightarrow \infty$ limit of θ is θ_* . Indeed, were this not the case, then (7-5) demands the impossible: The union of the $s = s_*$ slice of C_{j*} and a collection of very large, but constant $|s|$ slices of C_{0j} is the boundary of a subset of C_j 's model curve where $\theta \notin \{0, \pi\}$ and whose interior has a local extreme point of θ . This is impossible because, as noted in Subsection 2.A, the only local minima or maxima of θ on a pseudoholomorphic subvariety occur where θ is respectively 0 or π .

By the way, a component of the $s = s_*$ slice of C_{j*} that is homologically non-trivial in the radius ε neighborhood of S_* must intersect the C_{0j} version of the locus Γ_o when j is large. To explain, introduce the 1-form $x = (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$. The latter is exact in the radius ε tubular neighborhood of S_* so has zero integral over any given component of the $s = s_*$ slice of C_{j*} . Now, let η denote a component of the $s = s_*$ slice of C_{j*} . If $|\theta - \theta_*| > 0$ on η , then η must sit as an embedded circle in some component of the C_{0j} version of $C_0 - \Gamma$. If η is homologically non-trivial in this component, then it is homotopic there to a constant θ slice. Thus, the integral of x over η is $\pm \alpha_Q(\theta_*)$ where Q is the integer pair that is assigned to the edge in T that labels η 's component of $C_0 - \Gamma$. Thus, $\alpha_Q(\theta_*) = 0$. However, this is nonsense since (7-5) requires that $\theta = \theta_*$ on the closure of this component.

Step 4 To say more about the $s = s_*$ slice of C_{j*} requires a digression to point out some features of the small but positive δ versions of the $|\theta - \theta_*| < \delta$ neighborhood in C_{0j} of the locus Γ_o .

The first point here is that such a neighborhood is homeomorphic to a multiply punctured sphere. In particular, the first homology of such a neighborhood is canonically isomorphic to the first homology of the graph $\underline{\Gamma}_o^*$. This is to say that a system of generators for the $|\theta - \theta_*| < \delta$ can be obtained as follows: First take a large $|s|$ slice of each end of C_{0j} that corresponds to a vertex on $\underline{\Gamma}_o$. These form a set, $\{\ell^{*v}\}$, where the label v runs through the vertices in $\underline{\Gamma}_o$ with non-zero integer label. Add to these the set, $\{\ell^*_{oe}\}$, where the label e runs through the incident edges to o and ℓ^*_{oe} denotes the $|\theta - \theta_*| = \frac{1}{2}\delta$ slice of e 's component of the C_{0j} version of $C_0 - \Gamma$. The collection $\{\ell^*_{oe}\} \cup \{\ell^{*v}\}$ then generates the homology of this neighborhood of $\underline{\Gamma}_o$ subject to the one constraint in (2-22).

The second point concerns the 1-form $x = (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$ where θ is near θ_* . As remarked previously, this form is exact near S_* . To elaborate on this, write S_* as $\mathbb{R} \times \gamma_*$ where $\gamma_* \subset S^1 \times S^2$ is the relevant $\theta = \theta_*$ Reeb orbit. Now, x is, of course, pulled back from $S^1 \times S^2$ and its namesake on $S^1 \times S^2$ can be written on the radius ε tubular neighborhood of γ_* as df where f is a smooth function that vanishes on γ_* . Over the whole of the $\theta = \theta_*$ locus in $S^1 \times S^2$, the form x can be written as df_* where f_* is a multivalued function that agrees with an \mathbb{R} -valued lift near γ_* that agrees with f . The function f_* is constant on each $\theta = \theta_*$ Reeb orbit and its values distinguish the various $\theta = \theta_*$ Reeb orbits. In this regard, the values of f_* are in $\mathbb{R}/(2\pi\kappa_*\mathbb{Z})$ where

$$(7-7) \quad \kappa_*^2 = \frac{(1 + \cos^4 \theta_*)}{(p^2 + p'^2 \sin^2 \theta_*)}.$$

Here (p, p') is the relatively prime integer pair that determines θ_* via (1-8).

Step 5 This step concerns the integral of the 1-form x over any given component of the $|\theta - \theta_*| > 0$ portion of the $s = s_*$ slice of C_{j*} . To say more, let η denote such a component. Since $x = df$ near η and η is close to S_* where f is zero, so the integral of x over η has absolute value no greater than some j and η independent multiple of δ_j .

To see the implications of this result, note that η sits in some component of the C_j version of $C_0 - \Gamma$ whose label is an incident edge to o . Fix a parametrization of this component, and the closure of η then corresponds to an embedded path in the closed parametrizing cylinder whose two endpoints are on the $\sigma = \theta_*$ boundary and whose interior lies where $0 < |\theta - \theta_*| < \delta_j$. As such, this version of η is homotopic rel its endpoints to an embedded path, $\mathfrak{p}(\eta)$, in the $\sigma = \theta_*$ circle of the parametrizing cylinder.

The path $\mathfrak{p}(\eta)$ may or may not pass through some missing points on the $\sigma = \theta_*$ circle. Let $\mathfrak{p}_0(\eta) \subset \mathfrak{p}(\eta)$ denote the complement of any such missing points. Thus, $\mathfrak{p}_0(\eta)$ corresponds to a properly embedded, disjoint set of paths in Γ_o with two endpoints in total, these the endpoints of the closure of η . By virtue of the second point in the preceding step, the integral of x over $\mathfrak{p}_0(\eta)$ is also bounded in absolute value by a j and η independent multiple of δ_j . Moreover, with a suitable orientation, the 1-form x is positive on $\mathfrak{p}_0(\eta)$. Thus, the integral of x over any subset of $\mathfrak{p}_0(\eta)$ is also bounded by the same multiple of δ_j .

Step 6 Let V now denote the set of components of the $|\theta - \theta_*| > 0$ part of the $s = s_*$ locus in C_{j*} . Let σ_j denote $\cup_{\eta \in V} \mathfrak{p}_0(\eta)$. This is the image in Γ_o via a proper, piecewise smooth map of a finite, disjoint set of copies of S^1 and \mathbb{R} . However, by

virtue of what was said at the end of the previous step, if an arc, γ , from Γ_o is contained in σ_j , then the integral of x over γ is bounded by a j -independent multiple of δ_j . As a consequence, no large j version of σ_j contains the whole of any arc in Γ_o . Here is why: The integral of x over an arc in Γ_o gives the value on the arc of (C_{0j}, ϕ_j) 's assigned point in the simplex Δ_o . Were γ in σ_j , then this assigned point would give an $\mathcal{O}(\delta_j)$ value to γ , and were such the case for an infinite set of large j versions of (C_{0j}, ϕ_j) , then the image of $\{(C_{0j}, \phi_j)\}$ in Δ_o could not converge.

In the case $\Gamma_o = \underline{\Gamma}_o$, then σ_j is compact. Since no arc is contained in σ_j , it defines the zero homology class. According to the what was said in Step 4, it therefore defines the zero homology class in the $|\theta - \theta_*| < \delta$ neighborhood of Γ_o . At the same time, σ_j is homologous in the $|\theta - \theta_*| < \delta$ neighborhood of Γ_o to the $s = s_*$ slice of C_{j*} and Step 3 found that latter's class is definitively not zero. This nonsense proves that Γ_o can not be compact.

Suppose now that Γ_o is not the whole of $\underline{\Gamma}_o$. In this case, the closure of σ_j in $\underline{\Gamma}_o^*$ is the image via the collapsing map of the image, σ_j^* , of a map from a finite set of circles into $\underline{\Gamma}_o^*$. Since this σ_j^* contains no arc inverse image from Γ_o , its homology class must be a multiple of sums of those generated by the vertex labeled loops from the set $\{\ell^{*v}\}$. At the same time, the discussion in Step 4 identifies $H_1(\underline{\Gamma}_o^*; \mathbb{Z})$ with the first homology of the $|\theta - \theta_*| < \delta$ neighborhood of Γ_o and with this understood any such σ_j^* is homologous modulo the classes from $\{\ell^{*v}\}$ to the $s = s_*$ slice of C_{j*} . This then means that the $s = s_*$ slice of C_{j*} is homologous to some union of constant $|s|$ slices of the $\theta = \theta_*$ ends of C_{0j} . However, as noted in Step 3, this is definitively not the case. This last bit of nonsense completes the proof of Proposition 7.1. \square

Part 2 This part of the subsection proves that Ξ has but a single element. The proof borrows much from the discussion in Part 2 of Subsection 4.E. In particular, it starts out just the same by assuming that Ξ has more than one element so as to derive some patent nonsense. In this regard, note that the conclusions from Proposition 7.1 require that Ξ have at least one element that is not an \mathbb{R} -invariant cylinder and the three steps that follow explain why there is at most one element in Ξ of this sort. Granted that only one subvariety from Ξ is not an \mathbb{R} -invariant cylinder the argument for precluding \mathbb{R} -invariant cylinders from Ξ can be taken verbatim from Part 2 of Subsection 4.E.

Step 1 Choose an infinite subsequence from $\{(C_{0j}, \phi_j)\}$ with the following property: Let o denote a multivalent vertex in T and v a vertex in Γ_o . For each j , let x_j denote the image in C_j of the critical point of θ that corresponds to v . Then either $\{x_j\}$ converges, or $\{s(x_j)\}$ is unbounded and strictly increasing or strictly decreasing. Here

is the second property: Agree henceforth to relabel this subsequence by the integers starting at 1.

Define as in Part 2 of Subsection 4.E the subvariety $\Sigma \subset \mathbb{R} \times (S^1 \times S^2)$ to be the union of the subvarieties from Ξ that are not \mathbb{R} -invariant cylinders. Let $Y \subset \Sigma$ denote the critical points of $\cos(\theta)$ on the irreducible components of Σ , the singular points in the subvariety $\cup_{(S,n) \in \Xi} S$, and the limit points of the sequences $\{x_j\}$ as defined in the previous paragraph. So defined, Y is a finite set.

Step 2 Suppose two subvarieties from Ξ are not \mathbb{R} -invariant cylinders. To obtain nonsense from this assumption, let S and S' denote distinct, subvarieties from Ξ that are not \mathbb{R} -invariant. The argument in the fourth paragraph of Part 2 in Subsection 4.E that ruled out two non- \mathbb{R} invariant subvarieties applies here to prove that S and S' can be chosen so that there is a value of θ that is taken simultaneously on S and S' . No generality is lost by choosing this angle to be in $(0, \pi)$, to be distinct from θ 's values on Y , and to be distinct from the angles that are assigned to the vertices in T . This being the case, then for each j , there are points z_j and z_j' in C_{0j} on which θ has this same value and such that the sequence $\{z_j\}$ converges to a point in S while $\{z_j'\}$ converges to one in S' .

As C_{0j} is irreducible, there is a path in C_{0j} that runs from z_j to z_j' . For the present purposes, some paths are better than others. In particular there exists $R > 1$ and for each j , such a path, γ_j , with the following properties:

- (7-8) • *The image of γ_j in $\mathbb{R} \times (S^1 \times S^2)$ avoids the radius $\frac{1}{R}$ balls about any point in Y .*
- $|s| \leq R$ on γ_j .

The existence of R and $\{\gamma_j\}$ is proved in the subsequent two steps. Of course, granted (7-8), then S and S' must coincide because Proposition 7.1 and [15, Proposition 3.7] put the whole of every large j version of γ_j very close to only one subvariety from Ξ .

Step 3 The proof of the first point in (7-8) starts with the assertion that there exists $\varepsilon > 0$ such that the following is true:

- (7-9) *Let o denote a multivalent vertex in T . If j is large, then no arc in the C_{0j} version of Γ_o lies entirely in the inverse image of a single radius ε ball in $\mathbb{R} \times (S^1 \times S^2)$.*

The first point of (7-8) follows directly from this assertion and [15, Lemma 3.10].

To prove the assertion, suppose that γ is an arc in Γ_o whose image in C_j is in a radius ρ ball. The integral of the 1-form $(1 - \cos^2 \theta_o)d\varphi - \sqrt{6} \cos \theta_o dt$ over γ is then bounded by a fixed multiple of ρ . Thus, if ρ is small, so (C_{0j}, ϕ_j) 's assigned point in the simplex Δ_o assigns a small value to γ . Since the sequence $\{(C_{0j}, \phi_j)\}$ has convergent image in Δ_o , all such values enjoy a positive, j -independent lower bound.

The proof of the second point in (7–8) requires the following assertion:

(7–10) *There exists $R > 1$ such that when j is large, then no arc in the C_j version of Γ_o lies entirely where $|s| \geq R$.*

To see why, note that by virtue of Proposition 7.1, if R is such that the $|s| \geq R$ part of Σ is contained in the ends of Σ , then any arc in Γ_o where $|s|$ is everywhere larger than R is very close to an \mathbb{R} -invariant cylinder where $\theta = \theta_o$. Now, as explained in Part 1, the closed 1-form $(1 - \cos^2 \theta_o)d\varphi - \sqrt{6} \cos \theta_o$ can be written as df on some fixed neighborhood of such a cylinder with f being the pull-back from a neighborhood of the corresponding Reeb orbit of a smooth function that vanishes on the Reeb orbit. This then means that the integral of this 1-form over any $|s| \geq R$ arc is very small if R is very large and, with R chosen, then j is sufficiently large. Thus, the point assigned to (C_{0j}, ϕ_j) in Δ_o gives a very small value to such an arc. Since the assigned values enjoy a j -independent, positive lower bound, there is an upper bound to the minimum value of $|s|$ on any arc from any large j version of Γ_o .

The assertion in (7–10) allows the large j versions of γ_j to cross the C_{0j} version of the locus $\Gamma \subset$ where $|s|$ enjoys a j -independent upper bound. Proposition 7.1 then implies that the large j versions of γ_j can be chosen so that $|s|$ also enjoys a j -independent upper bound on the portions of γ_j in the C_{0j} version of $C_0 - \Gamma$. To elaborate on this, keep in mind that Proposition 7.1 finds a lower bound to $|s|$ on any large j version of γ_j at angles that are uniformly bounded away from the $|s| \rightarrow \infty$ limits of θ on the ends of Σ . To choose γ_j with a j -independent upper bound for $|s|$ as θ nears such a limit, first set the stage by letting θ_* denote the angle in question and write $(1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$ as df_* where f_* is the multivalued function that is described in the fourth step in Part 1. As f_* is constant on the $\theta = \theta_*$ Reeb orbits, so the ends of Σ where $\lim_{|s| \rightarrow \infty} \theta = \theta_*$ account for only a finite set of values for f_* . Thus, $|s|$ on γ_j will enjoy a j -independent upper bound if γ_j is chosen so that it approaches and crosses the $\theta = \theta_*$ locus where f_* is uniformly far from its values on the ends of Σ where θ_* is the $|s| \rightarrow \infty$ limit of θ . Granted (7–10), such a version of γ_j can be chosen in a component of the C_{0j} version of $C_0 - \Gamma$ using a parametrization as depicted in (2–5). In this regard, the case where f_* approaches a constant as $\theta \rightarrow \theta_*$ on a given component can be avoided since it occurs if and only if

the part of the component where θ is nearly θ_* is entirely in the large $|s|$ part of some end of C_{0j} .

Part 3 Let (S, n) denote the single element in Ξ . This last part of the story explains why (7–5) is true. The story here starts by making the following point: Because the image of $\{(C_{0j}, \phi_j)\}$ converges in $O_T / \text{Aut}(T)$, the choices in Part 2 of Subsection 6.C can be made for each (C_{0j}, ϕ_j) so that the resulting sequence in $(\times_o \Delta_o) \times (\mathbb{R}_- \times \mathbb{R}_\diamond) \times [\times_{o \in \mathcal{V}} \mathbb{R}_o]$ also converges. These choices are assumed in what follows.

Let S_0 denote the model curve for S and let ϕ_0 denote the tautological, almost everywhere 1–1 map from S_0 onto S in $\mathbb{R} \times (S^1 \times S^2)$. The complement of the inverse image of Y in S_0 is embedded in $\mathbb{R} \times (S^1 \times S^2)$ and so has a well defined normal bundle, this denoted by N in what follows.

Now, fix $\varepsilon > 0$ but very small and with the following considerations: First, the various points in Y are pairwise separated by a distance that is much greater than ε . View the latter distance as $\mathcal{O}(1)$ relative to ε . Second, given a point $p \in Y$, there is only one value of θ on the set Y that is also a value of θ on the radius ε ball about p , this being $\theta(p)$. Finally, the $|s| \geq 1/\varepsilon$ portion of S_0 is far out in the ends of S_0 and is far from any point that maps to Y .

Given such ε , let S_ε denote the portion of the $|s| \leq 1/\varepsilon$ part of S_0 that is mapped by ϕ_0 to where the distance from Y is at least ε^4 . There is a subdisk bundle, $N_\varepsilon \subset N$, over a submanifold of S that has S_ε in its interior and an exponential map $e: N_\varepsilon \rightarrow \mathbb{R} \times (S^1 \times S^2)$ with the following properties: First, it embeds N_ε as a tubular neighborhood of this larger submanifold. Second, it embeds each fiber as a pseudoholomorphic disk. Finally, the function θ is constant on each such fiber disk. In what follows, N_ε is not distinguished notationally from its image in $\mathbb{R} \times (S^1 \times S^2)$ via e .

If ε is fixed in advance and j is large, then $\phi_j(C_{0j})$ intersects N_ε as an immersed submanifold such that the composition of ϕ_j followed by the projection to S_ε defines a degree n , unramified covering map from $\phi_j^{-1}(N_\varepsilon)$ to S_ε . Let π_j denote the latter map. Granted that π_j is a covering map, then (7–5) follows by analyzing the behavior of $\phi_j(C_{0j})$ where $|s| \geq 1/\varepsilon$ and also where the distance to Y is less than ε . This analysis is presented in fourteen steps. The first six describe the parts of the large j versions of C_{0j} that map to where the distance to Y is less than ε .

Step 1 Let z denote a point in S_0 that maps to Y_* and let $B \subset \mathbb{R} \times (S^1 \times S^2)$ denote the radius ε ball centered on z 's image. Let θ_* denote the value of θ at z and assume here that $\theta_* > 0$. If z is a critical point of θ , let m denote $\deg_z(d\theta)$, otherwise set

$m = 0$. Fix an embedded circle, $\nu \subset S_0$, around z so that its image lies in B , so that it intersects the $\theta = \theta_*$ locus transversally in $2m + 2$ points, and so that all points in ν are mapped by ϕ_0 to where the distance to z 's image is greater than ε^2 . Let δ denote the maximum of $|\theta - \theta_*|$ on ν . It follows from (2-11) that ν can be chosen so that $\delta \leq c(z) \cdot \varepsilon^2$ where $c(z)$ depends only on z . With ε small, the image of ν in S has θ -preserving preimages via π_j in every large j version of C_{0j} . Let ν' denote such a circle. Let k denote the degree of π_i on ν' . The loop ν' intersects the $\theta = \theta_*$ locus in C_{0j} transversely in $k(2m + 2)$ points. The upcoming parts of the story explain why this is possible when ε is small and j is large only if ν' bounds an embedded disk in $\phi_j^{-1}(B)$ that contains a single critical point of θ , one where $d\theta$ vanishes with degree m .

Step 2 A digression is need for some preliminary constructions. For this purpose, fix attention on a component, μ , of the $\theta - \theta_* > 0$ portion of ν' . The interior of μ lies in a component of C_{0j} 's version of $C_0 - \Gamma$. After parametrizing the latter, μ can be viewed as an embedded path in the corresponding parametrizing cylinder with both endpoints on the $\theta = \theta_*$ circle but otherwise disjoint from this circle. Remark next that μ is homotopic rel its endpoints in the parametrizing cylinder to a path, μ_* , on the $\theta = \theta_*$ circle. In fact, the concatenation of μ and μ_* bounds a topologically embedded disk in the parametrizing cylinder. Use D_μ to denote both the interior of this disk and its image in C_j . If $0 \leq \theta - \theta_*$ on μ , then $0 < \theta - \theta_* < \delta$ on D_μ , and if $0 \geq \theta - \theta_*$ on μ , then $0 > \theta - \theta_* > -\delta$ on D_μ .

As is explained next, the path μ_* must be contained in $\phi_j^{-1}(B)$ when ε is small and j is large. To see why, introduce the 1-form $x \equiv (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$. This form is positive on μ_* and this leads to nonsense if ε is small, j is large and μ_* exits B . To demonstrate the nonsense, let B_1 denote the ball centered at z 's image in $\mathbb{R} \times (S^1 \times S^2)$ whose radius is the minimum of 1 and half the distance between z 's image and any other point of Y . Introduce the function f on B_1 that obeys $df = x$ and that vanishes on the image of z . The integral of x over μ is the difference between the values of f on μ 's endpoints, and so is bounded in absolute value by a j -independent multiple of ε . Since μ_* and μ are homotopic real boundary in the parametrizing cylinder, the integral of x over μ_* is likewise bounded in absolute value by ε .

Keeping this in mind, note that x is positive on the $\theta = \theta_*$ locus in $S \cap (B_1 - B)$. In particular, when ε is small, there is a positive and ε -independent number that is less than the integral of x over each such component. When j is large, Proposition 7.1 requires that this same number serve as a lower bound to the integral of x in $C_{0j} \cap \phi^{-1}(B_1 - B)$. Now, if a large j version of μ_* leaves $\phi_j^{-1}(B)$, then it must

leave $\phi_j^{-1}(B_1)$ also since a large j version of μ_* in $\phi_j^{-1}(B_1 - B)$ is a θ -preserving preimage of a portion of the $\theta = \theta_*$ locus in S_ε . Thus, the $\mathcal{O}(\varepsilon)$ integral of x over the large j versions of μ_* precludes it leaving $\phi_j^{-1}(B)$.

There are $2k(m+1)$ versions of μ , μ_* and D_μ , one for each component of the part of v' where $|\theta - \theta_*| > 0$. Each such μ_* is in $\phi_j^{-1}(B)$.

Step 3 The fact that each μ_* is in $\phi_j^{-1}(B)$ implies the following: If $\varepsilon > 0$ is chosen sufficiently small, then any sufficiently large j versions of v' is homotopically trivial in C_{0j} . To elaborate, note first that any loop in the $|\theta - \theta_*| < \delta$ part of C_{0j} is homotopic to one that sits entirely in the $\theta = \theta_*$ locus except at points far out on ends of C_{0j} where the $|s| \rightarrow \infty$ limit of θ is θ_* . The fact that any given version of μ_* lies entirely in $\phi_j^{-1}(B)$ implies that v' is homotopic to the loop in the $\theta = \theta_*$ circle that is obtained by the evident concatenation of the $k(2m+2)$ versions of μ_* . This is a loop, v_* , that lives entirely in the $\theta = \theta_*$ locus.

If θ_* is not the angle of a multivalent vertex in T , then v_* lies in some component of $C_{0j} - \Gamma$. Let e denote its labeling edge from T . If v_* is not null-homotopic, then the integral of $x \equiv (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$ over v_* is a non-zero, integer multiple of the $Q = Q_e$ version of $\alpha_Q(\theta_*)$. However, as $\phi_j(v')$ sits in a small ball, the integral of x over v' is therefore tiny if ε is small, and so the integral of x over v_* must be zero.

If θ_* is a multivalent vertex angle from T and if v_* is not null homotopic, then it must contain at least one arc from some version of Γ_o . Such an arc thus sits in $\phi_j^{-1}(B)$. However, if an arc is mapped into B , then the integral of $x \equiv (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$ over the arc is smaller than a j -independent, constant multiple of ε . This means that the arc is assigned a very small number by (C_{0j}, ϕ_j) 's image in $\times_o \Delta_o$ if ε is small and j is large. However, such arc assignments enjoy a j -independent, positive lower bound, and so there is no arc in D' if ε is too small and j is large. Thus, v_* is null-homotopic and so is v' .

Step 4 As each version μ_* is in $\phi^{-1}(B)$, it follows that all version of D_μ approach v' from the same side, and therefore the union of the $2k(m+1)$ version of D_μ define an embedded disk in C_{0j} with boundary v' . Let D' denote this disk. As $|\theta - \theta_*| \leq \delta$, this disk must lie in $\phi^{-1}(B)$ when δ is small and j is large. Indeed, were D' to exit $\phi^{-1}(B)$, it would have to intersect a θ -preserving preimage in C_{0j} of an analog of v around some other point in S_0 that maps to the same point as z . Since D' is embedded, it would then contain this circle and thus contain a point where $|\theta - \theta_*| > \delta$.

As constructed, the complement in D' of the $\theta = \theta_*$ locus is the union of the $2k(m+1)$ versions of D_μ . Since all D_μ are disks, this implies that the $\theta = \theta_*$ locus in D' is

connected. Moreover, if either m or k is greater than 1, then there are at least 4 such disks and so there is a critical point of θ inside D' . If both m and k equal 1, then there is no critical point of θ in D' .

If ε is small and j is large, there can be at most one such critical point in D' . To see this, note first that any two θ critical points in D' are joined in D' by a constant θ path since the $\theta = \theta_*$ locus in D' is connected. A constant θ path between two θ critical points is a union of arcs in some C_{0j} version of a graph from the set $\{\Gamma(\cdot)\}$. However, as noted in the previous step, no such arc can be mapped by ϕ_j into B if ε is small.

Granted that there is but one θ critical point in D' , then the degree of vanishing of $d\theta$ there must be $k \cdot (m + 1) - 1$ since there are $2k(m + 1)$ disk components to the complement of the $\theta = \theta_*$ locus in D' . Indeed, this follows from (2–11).

Step 5 If $n \geq 2$, then the loop ν can have more than one θ -preserving preimage. Suppose ν' and ν'' are two distinct θ -preserving preimages in a large j version of C_{0j} . As is explained next, no path in C_{0j} from ν' to ν'' is contained entirely in $\phi_j^{-1}(B)$ if ε is small and j is taken very large. To see why this is the case, note that were such a path to exist, the variation of θ on it would be bounded by a j -independent multiple of ε . As a consequence, ν' and ν'' would have to intersect the same component of the $\theta = \theta_*$ locus in C_{0j} and so there would be a path on this locus between them. Now, the disk D' is contained in $\phi_j^{-1}(B)$ and ν'' cannot be in D' as the θ values on ν'' are identical to those on ν . Thus this hypothetical path from ν' to ν'' would have to travel from ν' on the portion of the $\theta = \theta_*$ locus that avoids D' . This part of the locus is in the tubular neighborhood N_ε and as it projects to the $\theta = \theta_*$ locus in S_ε , it can not hit a θ -preserving preimage of ν before it exits the larger ball B_1 .

Step 6 This step describes the large j versions of $\phi_j(C_{0j})$ near points where S intersects the $\theta = 0$ locus. A similar story can be told for the points near the $\theta = \pi$ locus. To start, suppose that $z \in S_0$ is a point where θ is 0. In this case, choose a loop ν in $\phi_0^{-1}(B)$ that bounds a disk in S_0 whose center is z . In this regard, choose ν so that the maximum value of θ on ν is much less than that of θ on the boundary of $S \cap B$. Take j large and let $\nu' \subset C_{0j}$ again denote a θ -preserving preimage of ν . As before, this is an embedded circle. As can be seen using Proposition 7.1 and the C_{0j} versions of the parametrizations from Definition 2.1, the circle ν' is embedded in a component of the C_{0j} version of $C_0 - \Gamma$ whose closure contains a point where $\theta = 0$. As a consequence, ν' bounds a disk in this closure on which the maximum of θ is achieved on ν' . Let D' denote this disk. Were D' to leave $\phi_j^{-1}(B)$, its ϕ_j -image would by necessity intersect the boundary of B very near some component of $S \cap B$.

But, a large j version of such a disk would then have interior points where θ was larger than its maximum on v' . Thus D' is in $\phi_j^{-1}(B)$.

Here is one more point: Let k denote the degree of the covering map from v' to v . If the local intersection number at z between S_0 and the $\theta = 0$ locus is denoted as q_z , then the local intersection number between D' and the $\theta = 0$ locus is $k \cdot q_z$.

Step 7 Let $E \subset S_0$ denote an end, let θ_* denote the $|s| \rightarrow \infty$ limit of θ on E , and let n_E denote the integer that appears in E 's version of (2–4). Assume until told otherwise that $\theta_* \in (0, \pi)$. Take $R_0 \gg 1$ so that $|s|$ takes the value R_0 on E and so that R_0 is much greater than the value of $|s|$ on any point from the set Y . As can be seen from (2–4), it is also possible to choose R_0 so that the $|s| \geq R_0$ portion of the $\theta = \theta_*$ locus consists of $2n_E$ properly embedded, half open arcs on which ds restricts without zeros. Take R_0 so that this last condition holds. In what follows, it is assumed that $1/\varepsilon$ is much greater than R_0 .

It follows from E 's version of (2–4) that when ε is sufficiently small, there exists an embedded loop, v , in the $|s| \in [\frac{1}{2\varepsilon} + 2, \frac{1}{\varepsilon} - 2]$ part of E with the following properties: First, v is homotopic to the $|s| = \frac{1}{2\varepsilon}$ slice of E . Second, v intersects the $\theta = \theta_*$ locus transversely, and in $2n_E$ points. Third, the maximum on v of $|\theta - \theta_*|$ is less than ε . Use δ in what follows to denote this maximum.

If j is sufficiently large, then v has θ -preserving preimages in C_j . Let v' denote such a preimage. Thus, v' also intersects the $\theta = \theta_*$ locus transversely, and the maximum of $|\theta - \theta_*|$ on v' is also δ . Finally, with k denoting the degree of the covering map π_j on v' , then v' intersects the $\theta = \theta_*$ locus $2kn_E$ times, each in a transversal fashion. Note that by virtue of Proposition 7.1, the loop v' is homotopically non-trivial in the complement in C_{0j} of the $\theta \in \{0, \pi\}$ locus.

Let μ denote the closure of a component of the $|\theta - \theta_*| > 0$ portion of v' . This μ can be viewed as an embedded path in a parametrizing cylinder for the component of the C_{0j} version of $C_0 - \Gamma$ that contains μ 's interior. Viewed in the parametrizing cylinder, μ lies in the interior save for its two boundary points. In the parametrizing cylinder, μ is homotopic rel boundary to a path, μ_* , that runs between the two endpoints of μ on the $\theta = \theta_*$ circle. This is to say that the concatenation of μ and μ_* bounds a topological disk in the parametrizing cylinder. Let D_μ denote the interior of this disk, and also its image in C_{0j} . As before, $0 < \theta - \theta_* < \delta$ on the disk D_μ when $0 \leq \theta - \theta_*$ on μ , and $0 > \theta - \theta_* > \delta$ on D_μ when $0 \geq \theta - \theta_*$ on μ .

The parametrizing cylinder that contains μ_* may or may not have missing points on its $\theta = \theta_*$ circle. In the case that some such points lie in μ_* , let $\mu_{*0} \subset \mu_*$ denote their complement. The latter has a corresponding image in C_{0j} . The concatenation of μ_{*0}

and μ is a piecewise smooth, properly embedded submanifold of C_{0j} that bounds the closure of the disk D_μ . In this regard, the complement of the θ -critical points in μ_{*0} is smoothly embedded, and a component of this complement that lacks an endpoint of μ is the whole of the interior of an arc in some graph from the C_{0j} version of the collection $\{\Gamma_{(\cdot)}\}$

Step 8 When ε is small, and with ε chosen, j is sufficiently large, the disk D_μ lies entirely in the $|s| > \frac{1}{2\varepsilon}$ portion of C_{0j} . To see why, note first that if R_0 is large, then the $|s| \geq R_0$ part of any end in S_0 where $\lim_{|s| \rightarrow \infty} \theta = \theta_*$ lies in a small radius tubular neighborhood of a $\theta = \theta_*$ pseudoholomorphic cylinder. In this tubular neighborhood, the 1-form $x \equiv (1 - 3 \cos^2 \theta_*)d\varphi - \sqrt{6} \cos \theta_* dt$ is the pull-back by the projection to $S^1 \times S^2$ of an exact form, df , where f is a function that vanishes on the Reeb orbit and is constant on any nearby $\theta = \theta_*$ Reeb orbit. This understood, the integral of x along any part of the $|s| \geq \frac{1}{2\varepsilon} > R_0 + 3$ portion of the $\theta = \theta_*$ locus in any given end of S_0 is very small in absolute value with the bound going to zero as $\varepsilon \rightarrow 0$. On the other hand, there exists $\kappa > 0$ with the following significance: If γ is a connected portion of the $\theta = \theta_*$ locus in an end of S_0 , and if γ runs from where $|s| = \frac{1}{2\varepsilon}$ to $|s| = R_0 + 1$, then the integral of x over γ is greater than κ .

If ε is sufficiently small, then the integral of x over μ will be less than κ since the integral is the difference between the values of f at the two endpoints. This then means that the integral over μ_* of the pull-back of x to the parametrizing cylinder is also less than κ . In this regard, note that x pulls back as $\alpha_Q(\theta_*)dv$ where Q is the integer pair that is associated to the edge label from μ 's component of $C_0 - \Gamma$.

Now, if j is very large and $|s| \leq \frac{1}{2\varepsilon}$ on μ_{*0} , then this part of μ_{*0} must be a θ -preserving preimage in C_{0j} of a component of the $\theta = \theta_*$ locus in S_0 in an end of S_0 where $\lim_{|s| \rightarrow \infty} \theta = \theta_*$. As all such components run from where $|s| > \frac{1}{2\varepsilon}$ to where $|s| = R_0$, so must μ_* . Granted the conclusions of the preceding paragraph, Proposition 7.1 demands that x have integral greater than κ on x_* . Hence, μ_{*0} can not enter the $|s| \leq \frac{1}{2\varepsilon}$ part of C_{0j} . This implies that the disk D_μ is also forbidden from the $|s| \leq \frac{1}{2\varepsilon}$ part of C_{0j} .

Step 9 There are $2k \cdot n_E$ versions each of μ , μ_* and D_μ . Since $|s| > \frac{1}{2\varepsilon}$ on all versions of μ_{*0} , they all leave v' from the same side of v' ; and this implies that the closure in C_{0j} of the union of the $2k \cdot n_E$ versions of D_μ is a submanifold of C_{0j} with a piecewise smooth boundary. One boundary component is v' , and were there more, they would sit inside the union of the $2k \cdot n_E$ versions of μ_{*0} . However, when $\frac{1}{2\varepsilon}$ is large and j also, then v' is the only boundary component. Indeed, the last remarks of Step 7 imply that any second component would necessarily contain

the whole interior of an arc in some graph from the C_{0j} version of the collection Γ_0 . By virtue of Proposition 7.1, this arc would sit very close to the end E when ε is small and when j is very large. As explained in the preceding step, taking ε small and j large makes the integral of the 1-form x over such an arc less than any given, positive number. But the integral of x over an arc from C_{0j} 's version of $\{\Gamma_{(\cdot)}\}$ enjoys a j -independent, positive lower bound since the image of the sequence $\{(C_{0j}, \phi_j)\}$ in $\times_o \Delta_o$ converges.

Let C' denote the closure of the union of these $2k \cdot n_E$ versions of D_μ . When ε is small and j is large, this smooth submanifold of C_{0j} is homeomorphic to the complement of the origin in the closed unit disk; thus a half open cylinder with boundary v' . Indeed, C' is not a disk because v' is homologically non-trivial in the complement of the $\theta = 0$ and $\theta = \pi$ loci. Here is why C' has but one puncture: Each puncture corresponds to an end of C_{0j} where the $|s| \rightarrow \infty$ limit of θ is θ_* . Thus, each corresponds to a vertex on a graph from C_{0j} 's version of $\{\Gamma_{(\cdot)}\}$. Since θ is nearly θ_* on C' and C' is connected, all such vertices are on the same graph. Let o denote the corresponding vertex. Since the complement of the $\theta = \theta_*$ locus in C' is a union of disks, the part of Γ_o in C' has connected closure in Γ_o . Thus, any two vertices in Γ_o that label ends in C' are joined by an arc in Γ_o whose interior lies entirely in C' . As argued in the preceding paragraph, there are no such arcs when ε and j are large.

The argument just used explains why C' has no critical points of θ on C_{0j} .

Step 10 The conclusions of the preceding step imply that the end $E \subset S_0$ where the $|s| \rightarrow \infty$ limit of θ is in $(0, \pi)$ corresponds to a set of ends of C_{0j} that are all very close to E in $\mathbb{R} \times (S^1 \times S^2)$. This collection is in 1-1 correspondence with the θ -preserving preimages of v in the sense that each pre-image bounds a properly embedded, half open cylinder in C_{0j} whose large $|s|$ part coincides with the large $|s|$ part of its corresponding end. If v' is a θ -preserving preimage of v , let $E' \subset C_{0j}$ denote the corresponding end. If k is the degree of the projection from v' to v , then the integer $n_{E'}$ in the E' version of (2-4) is $k \cdot n_E$.

Consider next the behavior of C_{0j} near an end $E \subset S_0$ where the $|s| \rightarrow \infty$ limit of θ is 0. The pair (p, p') are used in what follows to designate the integers that appear in E 's version of (1-9).

To start story in this case, take $R_0 \gg 1$ so that $|s|$ takes the value R_0 on E and so that R_0 is much greater than the value of $|s|$ on any point in Y and any $\theta = 0$ point in S . In particular, choose R_0 so that (1-9) describes the $|s| \geq R_0$ portion of E . It is assumed here that ε is such that $\frac{1}{2\varepsilon} \gg R_0$.

Now let ν denote a constant θ slice of E that lives where $|s| \in [\frac{1}{2\varepsilon} + 2, \frac{1}{\varepsilon} - 2]$. Let $\delta > 0$ denote the value of θ on ν . As before, $\delta < \varepsilon$ when ε is small. When j is very large, the circle ν has θ -preserving preimages in C_{0j} . Let ν' denote such a preimage, and let k denote the degree of the restriction of π_j as a map from ν' to ν .

As can be seen using Proposition 7.1 and the C_{0j} versions of the parametrizations from Definition 2.1, any small ε and large j version of ν' bounds a cylinder in C_{0j} on which θ is strictly positive but limits to zero as $|s| \rightarrow \infty$. As such, the large $|s|$ part of this cylinder is the large $|s|$ part of an end of C_{0j} whose associated integer pair is (kp, kp') . A cylinder in C_{0j} of the sort just described is denoted in what follows by C' .

The part of C_{0j} that maps near an end of S where the $|s| \rightarrow \infty$ limit of θ is π looks much the same as the description just given for the part near an end of S where the $|s| \rightarrow \infty$ limit of θ is 0. To summarize: Each end of S_0 where $\lim_{|s| \rightarrow \infty} \theta$ is either 0 or π corresponds to one or more ends of each large j version of C_{0j} . If the given end of S_0 is characterized by the 4-tuple $(\delta = \pm 1, \varepsilon = \pm, (p, p'))$, then each of the associated ends of C_{0j} is characterized by a 4-tuple with the same δ and ε , and with an integer pair that is some positive multiple of (p, p') . Moreover, these multiples from the associated ends to each S_0 end add up to the integer n .

Step 11 The step considers assertion in (7–5) for the case that the integer n is 1. To argue this case, let T_S denote the graph that is assigned to the pair (S_0, ϕ) . Granted that T_S is isomorphic to T , it follows that (S_0, ϕ) defines a point in $\mathcal{M}_{\hat{A}, T}$ and Proposition 7.1 asserts that $\{(C_{0j}, \phi_j)\}$ converges to (S_0, ϕ) in $\mathcal{M}_{\hat{A}, T}$. As noted in Theorem 1.3, this means that $\{(C_{0j}, \phi_j)\}$ converges to (S_0, ϕ) in $\mathcal{M}^*_{\hat{A}, T}$ which is the desired conclusion.

With the preceding understood, what follows explains why T_S is isomorphic to T . The explanation starts with a summary of results from the 10 steps just completed. The first point here is that the ends of S_0 and those of the large j versions of C_{0j} enjoy a 1–1 correspondence whereby corresponding pairs determine the same 4-tuple and map very near each other in $\mathbb{R} \times (S^1 \times S^2)$ when j is large. The second point is that the respective sets of $\theta = 0$ points in S_0 and in the large j versions of C_{0j} enjoy a 1–1 correspondence whereby corresponding pairs have identical local intersection numbers and map very near each other in this locus when j is large. A similar correspondence exists between the respective sets of $\theta = \pi$ points in S_0 and in C_{0j} . Granted these two points, it follows that $(S_0, \phi) \in \mathcal{M}_{\hat{A}}$.

In the case that the integer n is 1, the set of non-extremal critical points of θ on S_0 enjoy a 1–1 correspondence with the analogous set in any large j version of C_{0j} . This

correspondence is such that partnered critical points have the same critical value and the same degree of vanishing of $d\theta$. Moreover, corresponding critical points map very close to each other when j is large. Granted these conclusions and those of the first paragraph of this Step 11, it follows that the respective vertices in the graph T_S and in the graph T enjoy a 1–1 correspondence that partners pairs with equal angle.

To compare the edges of the graphs T_S and T , remark that by virtue of Proposition 7.1, the components of $S_0 - \Gamma$ enjoy a 1–1 correspondence between those of any large j version of $C_{0j} - \Gamma$ that pairs components that map very close to each other in $\mathbb{R} \times (S^1 \times S^2)$. Moreover, the respective ranges of θ on paired components are identical, and the respective integrals of $\frac{1}{2\pi} dt$ and of $\frac{1}{2\pi} d\varphi$ about the constant θ slices of paired components are equal. It follows from this that the edges of T_S and T enjoy a 1–1 correspondence that is consistent with the aforementioned vertex correspondence and preserves integer pair assignments.

It remains now to consider the respective T_S and T versions of the graphs $\{\underline{\Gamma}_{(\cdot)}\}$ that are assigned to paired, multivalent vertices. For this purpose, let o denote a multivalent vertex in T and also its partner in T_S . The correspondence between the respective sets of non-extremal critical points of θ in S_0 and in C_{0j} together with that between the respective sets of ends in S_0 and in C_{0j} defines a 1–1 correspondence between the vertices in the S_0 and C_{0j} versions of Γ_o . Moreover, the latter correspondence pairs vertices with the same labels and with the same number of incident half-arcs.

The conclusions of Step 4 imply that any compact portion of the interior of any arc in the S_0 version of Γ_o has a unique θ -preserving preimage in any sufficiently large j version of C_{0j} . This correspondence pairs the arcs in the S_0 and C_{0j} versions of Γ_o so that partnered arcs have the same edge pair labels. The conclusions of Steps 4 and 9 together imply that this arc correspondence is consistent with the just described pairing of the vertices of $\underline{\Gamma}_o$. As a result, the two versions of Γ_o are isomorphic via an isomorphism that respects the correspondences described previously between the respective edge and vertex sets of T_S and those of T .

Taken together, these correspondences describe the desired isomorphism between T_S and T .

Step 12 This step establishes the $n > 1$ cases of (7–5). To start the argument for this case, fix ε to be very small and note that each large j version of π_j defines $\phi_j^{-1}(N_\varepsilon)$ as a degree n , proper covering space over S_ε . The first task in this step is to explain why the sequence whose j 'th component is the supremum over $\phi_j^{-1}(N_\varepsilon)$ of the ratio $|\partial\pi_j|/|\partial\pi_j|$ limits to zero as $j \rightarrow \infty$. To see this, let $z \in S_\varepsilon$. As explained in the first point of Step 2 in Subsection 5.C, there is a neighborhood of $\phi(z)$ in $\mathbb{R} \times (S^1 \times S^2)$ with

complex valued coordinates (x, y) and a disk $D \subset S_\varepsilon$ with center z such that the $y = 0$ slice is the ϕ -image of D , the $x = \text{constant}$ slices are the fibers over D of the bundle N_ε , and the 1-forms in (5-2) span $T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$ over the coordinate patch. With D identified via ϕ with the $y = 0$ slice, the function x restricts as a holomorphic coordinate on D , and π_j on $\pi_j^{-1}(D)$ is the composition $x \circ \phi_j$. This understood, it follows that the ratio $|\bar{\partial}\pi_j|/|\partial\pi_j|$ on $\pi_j^{-1}(D)$ is the ratio $|\bar{\partial}x|/|\partial x| = |\sigma|$. As σ vanishes where $y = 0$, Proposition 7.1 implies that the sequence whose j 'th element is the supremum over $\pi_j^{-1}(D)$ of $|\sigma|$ converges to zero as $j \rightarrow \infty$.

The assertion that the sequence whose j 'th element is the supremum over $\pi_j^{-1}(S_\varepsilon)$ of $|\bar{\partial}\pi_j|/|\partial\pi_j|$ limits to zero as $j \rightarrow \infty$ follows from this analysis on disks by virtue of the fact that S_ε is compact.

The next task is to extend each large j version of π_j as a degree n , ramified cover of the whole of C_{0j} to S_0 so that

$$(7-11) \quad \lim_{j \rightarrow \infty} \sup_{C_{0j}} (|\bar{\partial}\pi_j|/|\partial\pi_j|) = 0.$$

To do so, remark that the complement in C_{0j} of $\phi_j^{-1}(N_\varepsilon)$ consists of a disjoint union of disks and cylinders. The set of disks is partitioned into subsets that are labeled by the disk components of $S_0 - S_\varepsilon$. Likewise, the set of cylinders is partitioned into subsets that are labeled by the cylindrical components of $S_0 - S_\varepsilon$.

What with the conclusions from Steps 9 and 10, the desired extension of each large j version of π_j over the cylindrical components of $C_{0j} - \pi_j^{-1}(S_\varepsilon)$ is obtained by copying in an almost verbatim fashion the construction that is described in Step 4 of Subsection 5.C. In this regard, note that this construction extends π_j as an unramified cover over the cylindrical components of $C_{0j} - \pi_j^{-1}(S_\varepsilon)$.

The map π_j is extended over the disk components of $C_{0j} - \phi_j(N_\varepsilon)$ momentarily. Granted that this has been done, here is how to complete the argument for (7-5). Remark first that by virtue of (7-11), the large j versions of π_j can be used to pull-back the complex structure from S_0 and so define a new complex structure on C_{0j} that makes π_j into a holomorphic map. Let S_{nj} denote C_{0j} with this new complex structure. Note that the pair $(S_{n,j}, \phi \circ \pi_j)$ defines an equivalence class in $\mathcal{M}^*_{\hat{A}, T}$. Furthermore, the points that are defined by the pairs $(S_{0j}, \phi \circ \pi_j)$ and (C_{0j}, ϕ_j) are very close in $\mathcal{M}^*_{\hat{A}, T}$. Indeed, to see this, take ψ in (1-24) to be the identity map. Proposition 7.1 asserts that $\phi \circ \pi_j$ and $\phi_j \circ \psi$ are very close when j is large. Meanwhile, $r(\psi) = |\bar{\partial}\pi_j|/|\partial\pi_j|$, and this is very small when j is large by virtue of (7-11).

To finish the argument, note that the ramification points for the map π_j converge in S_0 as $j \rightarrow \infty$ since none occur in the cylindrical components of $S_0 - S_\varepsilon$. Thus, the

sequence of complex curves $\{S_{nj}\}$ has a subsequence that converges to a complex curve. Let S_n denote the latter. The corresponding subsequence of holomorphic maps from the sequence $\{\pi_j\}$ likewise converges to a degree n , holomorphic, ramified covering, $\pi: S_n \rightarrow S_0$. This pair $(S_n, \phi \circ \pi)$ defines an equivalence class in $\mathcal{M}^*_{\hat{A}, T}$ that is a limit of the sequence that is defined by the pairs $\{(S_{0j}, \phi \circ \pi_j)\}$. Granted what was said in the preceding paragraph, the point defined by $(S_n, \phi \circ \pi)$ is necessarily the limit of the sequence $\{(C_{0j}, \phi_j)\}$.

The final two steps explain how π_j is extended as a branched cover over the disk components of each large j version of $C_{0j} - \pi_j^{-1}(S_\varepsilon)$ so as to satisfy (7–11).

Step 13 To start, let z denote the center point of a disk from $S_0 - S_\varepsilon$ and suppose first that $\theta(z)$ is neither 0 nor π . Let B denote the radius ε ball centered at z . Introduce the function r on B as defined in (2–9). As noted in the ensuing discussion,

$$(7-12) \quad dr = J \cdot d\theta + rd\theta$$

on B where $|r| \leq c \cdot \text{dist}(\cdot, \phi(z))$ with c being a constant.

Now, let $D \subset S_0$ denote a disk with center z whose boundary is in S_ε and whose ϕ image is in B . According to Property 5 from Subsection 2.B, if ε is small, there is a holomorphic coordinate on D such that the function $r + i\theta$ defined on B pulls back as indicated in (2–11) with θ_* set equal to $\theta(z)$ and with $m \geq 0$ denoting $\deg_z(d\theta)$. This understood, let $x \equiv \theta + ir - \theta(z)$. As a consequence of (2–11), $\tau \equiv x^{1/(m+1)}$ defines a class C^1 , complex valued coordinate on D when ε is small. Moreover, by virtue of (7–12),

$$(7-13) \quad |\bar{\partial}\tau| \leq c\rho|\partial\tau|$$

at the points whose image in B has distance ρ or less from $\phi(z)$. Here, c is independent of ρ when ρ is small.

To continue, let $D' \subset C_{0j}$ denote the disk that bounds a given component of $\pi_j^{-1}(\partial D)$ and let k denote the degree of π_j as a map from $\partial D'$ to ∂D . When j is large, D' contains a single disk component of $C_{0j} - \phi_j^{-1}(N_\varepsilon)$ and is mapped into B by ϕ_j . Also, D' contains inside it a single critical point of θ on C_{0j} . The aforementioned Property 5 in Subsection 2.B provides D' a holomorphic coordinate such that (2–11) holds with $k \cdot (m + 1) - 1$ replacing m . This understood, there is a complex valued constant, x_j , such that $x - x_j$ has a $k \cdot (m + 1)$ fold root on D' . This is to say that $\tau' \equiv (x - x_j)^{1/k(m+1)}$ defines a C^1 , complex valued function on D' when ε is small and j is large. In fact, it follows from (7–12) and the D' version of (2–11) that τ' defines a C^1 , complex coordinate on D' and that (7–13) holds with τ' replacing τ .

Let θ_j be such that the value of θ at the $\tau' = 0$ point in D' is $\theta(z) + \theta_j$. thus, $\theta_j = 0$ in the case that $m > 0$. Next, define $\psi: D' \rightarrow D$ by declaring that

$$(7-14) \quad \psi * \tau - \theta_j = \tau'^k.$$

This C^1 map realizes D' as a degree k branched cover over D with a single branch point. Moreover, this map restricts near $\partial D'$ so as to have the following property: The maps $\phi \circ \psi$ and ϕ_j send any given point to the same constant θ , pseudoholomorphic submanifold in B . Indeed, such is the case by virtue of the fact that r as well as θ are constant on any $\theta = \text{constant}$ submanifold in $\mathbb{R} \times (S^1 \times S^2)$. It follows from this last conclusion that τ' can be changed via multiplication by a k 'th root of unity so that it agrees with π_j near $\partial D'$. Thus, ψ extends π_j over D' as a C^1 , ramified cover.

Here is one last point about this extension of π_j : By virtue of (7-13) and its τ' analog, $|\bar{\partial}\psi| \ll |\partial\psi|$ over the whole of D' , and the sequence whose j 'th element is the supremum of $|\bar{\partial}\psi|/|\partial\psi|$ over the j 'th version of D' has limit zero as $j \rightarrow \infty$. Since ψ is differentiable and smooth save at $\psi^{-1}(z)$, it has a deformation that extends π_j over D' as a smooth map with $|\bar{\partial}\pi_j| \ll |\partial\pi_j|$ at all points. Moreover, these extensions can be made for each large j version of ψ so that the resulting sequence of supremums of $|\bar{\partial}\pi_j|/|\partial\pi_j|$ has limit zero as $j \rightarrow \infty$.

Step 14 To finish the story about π_j 's extension to the disk components of $C_{0j} - \pi_j^{-1}(S_\varepsilon)$, suppose now that z is a point in a disk component of $S_0 - S_\varepsilon$ where θ is zero. Let B denote the radius ε ball in $\mathbb{R} \times (S^1 \times S^2)$ centered at $\phi(z)$. The ball B has smooth complex coordinates (x, y) where

$$(7-15) \quad x = \sin \theta \exp\left(-i\left(\varphi - \frac{\sqrt{6} \cos \theta}{(1 - 3 \cos^2 \theta)} \hat{t}\right)\right) \text{ and } y = s - i\hat{t};$$

here \hat{t} is the \mathbb{R} -valued lift of the function t to B that vanishes at $\phi(z)$. These coordinates are such that each $x = \text{constant}$ disk is pseudoholomorphic and the $x = 0$ disk lies entirely in the $\theta = 0$ cylinder. In this regard, there is a complex valued function σ on B that vanishes at $x = 0$ and is such that $dx + \sigma d\bar{x}$ spans $T^{1,0}(\mathbb{R} \times (S^1 \times S^2))$ over the whole of B .

Now, let $D \subset S_0$ denote a disk centered at z with boundary in S_ε that is mapped by ϕ into B . Let q denote the intersection number between D and the $\theta = 0$ locus. Since ϕ is holomorphic, there is a holomorphic coordinate, u , on D such that x pulls back via ϕ as $u^q + \mathcal{O}(|u|^{q+1})$. Thus, $\tau \equiv x^{1/q}$ defines a C^1 , complex coordinate on D . Moreover, τ obeys (7-13) where the constant c is furnished by the expression for x in (7-15).

Now fix some very large j and let $D' \subset C_{0j}$ denote a disk that bounds a component of $\pi_j^{-1}(\partial D)$. Let k denote the degree of π_j 's restriction to $\partial D'$. Applying the argument just given to D' finds that $\tau' \equiv x^{1/kq}$ defines a C^1 , complex coordinate on D' that obeys its own version of (7–13) with the constant c used for the τ version. Granted all of this, define a map, $\psi: D' \rightarrow D$ by requiring that $\psi * \tau = \tau'^k$. Arguing as in the case where $\theta(z) \neq 0$ finds that such a map ψ can be modified by multiplication by a k 'th root of unity so as to provide a C^1 extension of π_j over D' . Moreover, this extension has a smooth perturbation as a degree k ramified cover with $|\bar{\partial}\pi_j| \ll |\partial\pi_j|$. Finally, these extensions can be made for each large j so that the sequence whose j 'th element is the supremum of $|\bar{\partial}\pi_j|/|\partial\pi_j|$ limits to zero as $j \rightarrow \infty$.

8 The strata and their classification

This section completes the story started in Section 5 by describing in more detail the strata of $\mathcal{M}^*_{\hat{A}}$. The following is a brief summary of the contents: The first subsection describes each stratum as a fiber bundle over a product of simplices whose typical fiber is some $\mathcal{M}^*_{\hat{A},T}$. This result is stated as Proposition 8.1. Theorem 6.2 and Proposition 8.1 thus give an explicit picture of any given component of any given stratum in $\mathcal{M}^*_{\hat{A}}$.

Meanwhile, Subsection 8.B describes necessary and sufficient conditions on a graph T that insure a non-empty $\mathcal{M}^*_{\hat{A},T}$. These are stated as Proposition 8.2. The final subsection proves Proposition 6.1, this the assertion that the homotopy type of a graph T arises from at most one component of at most one stratum of $\mathcal{M}^*_{\hat{A}}$. Together, Propositions 6.1 and 8.2 provide a complete classification of the components of the strata that comprise $\mathcal{M}^*_{\hat{A}}$.

8.A The structure of a stratum

As in Sections 5.A, let $\mathcal{S}_{B,c,\delta}$ denote a stratum of $\mathcal{M}^*_{\hat{A}}$ and let $\mathcal{S} \subset \mathcal{S}_{B,c,\delta}$ denote a connected component. A graph of the sort introduced in Subsection 6.A from any equivalence class in \mathcal{S} has some m distinct, multivalent vertex angles that do not arise via (1–8) from any $(0, +, \dots)$ element in \hat{A} nor from an element in B . This understood, a function, \mathfrak{p} , from \mathcal{S} to the m 'th symmetric product of $(0, \pi)$ is defined as follows: The value of \mathfrak{p} on the equivalence class defined by (C_0, ϕ) is the set of θ -values of the m compact but singular θ level sets in C_0 . If non-empty, then the inverse image via \mathfrak{p} of any given point in the m 'th symmetric product of $(0, \pi)$ is a version of $\mathcal{M}^*_{\hat{A},T}$ where T has precisely m distinct, multivalent vertex angles that do not come via (1–8) from an integer pair of any $(0, +, \dots)$ element in \hat{A} nor any element in B .

With the preceding understood, consider:

Proposition 8.1 *If non-empty, then \mathcal{S} is fibered by \mathfrak{p} over a product of open simplices.*

Proof of Proposition 8.1 It follows from Lemma 5.4 with (2–4), (2–11) and the implicit function theorem that the map \mathfrak{p} is a submersion. Granted this, then Theorem 6.2 implies that \mathfrak{p} fibers \mathcal{S} over its image in the m 'th symmetric product of $(0, \pi)$.

To picture the image of \mathcal{S} via \mathfrak{p} , return to the definition in Subsection 5.A of $\mathcal{S}_{B,c,\mathfrak{d}}$. The definition involved the subspace $\mathcal{S}_{B,c} \subset \mathcal{M}^*_{\hat{A}}$ whose elements are defined by pairs (C_0, ϕ) where the following two conditions are met: First, C_0 has precisely c critical points of θ where θ is neither 0 nor π . Second, the ends in C_0 that correspond to elements in B are the sole convex side ends of C_0 where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π and whose version of (2–4) has a strictly positive integer n_E . Letting $d = N_+ + |B| + c$ and I_d the d 'th symmetric product of $(0, \pi)$, then $\mathcal{S}_{B,c,\mathfrak{d}}$ consists of the inverse image of the stratum in I_d indexed by \mathfrak{d} via the map that assigns to a given (C_0, ϕ) the angles of the critical points in $(0, \pi)$ of θ as well as the angles in $(0, \pi)$ that are $|s| \rightarrow \infty$ limits of θ on the concave side ends in C_0 and the ends that correspond to 4-tuples from B . As in Subsection 5.A, denote the map from $\mathcal{S}_{b,c}$ to I_d as f .

Only m angles in the image of f can vary and their values define the map \mathfrak{p} . To say more, let $\Lambda_{+,B}$ denote the angles that are defined via (1–8) by the integer pairs from the $(0, +, \dots)$ elements in \hat{A} and from the elements in B . The complement of $\Lambda_{+,B}$ in $(0, \pi)$ is a union of open segments. The image of f consists of the set $\Lambda_{+,B}$ and then m angles that are distributed $(0, \pi) - \Lambda_{+,B}$. Each of the latter angles can vary as a function on \mathcal{S} , but only in a single component of $(0, \pi) - \Lambda_{+,B}$. However, keep in mind that two such angles in the same component can not coincide. Furthermore, the angles from f that lie in a given component of $(0, \pi) - \Lambda_{+,B}$ need not sweep out the whole component as functions on \mathcal{S} . The picture just drawn implies that the image of \mathfrak{p} is a product of simplices.

As is explained below, there are additional constraints on the range of variation of the angles from f . □

8.B An existence theorem

A component of any given stratum of $\mathcal{M}^*_{\hat{A}}$ has an associated homotopy type of graph. Though yet to be proved, Proposition 6.1 asserts that these homotopy types classify the components of the strata. The next proposition gives necessary and sufficient conditions for a given homotopy type of graph to arise as the label of a component of some stratum

of $\mathcal{M}^*_{\hat{A}}$. In this proposition and subsequently, the angle that is assigned to a given vertex o of a graph T is denoted by θ_o . As before, when e is an edge of T , then e 's assigned integer pair is denoted by Q_e or $(q_e, q_{e'})$.

Proposition 8.2 *Let T denote a graph of the sort that is described in Subsection 6.A. Then the space $\mathcal{M}^*_{\hat{A}, T}$ is non-empty if and only if the following conditions are met: Let e denote an edge of T and let o and o' denote the vertices at the ends of e where the convention taken has $\theta_{o'} > \theta_o$. Then the $Q = Q_e$ version of α_Q is positive on $(\theta_o, \theta_{o'})$, and it is zero at an endpoint if and only if the corresponding vertex is monovalent and labeled by a $(0, -, \dots)$ element from \hat{A} .*

To comment on these conditions, remark that if e is an edge, then its integer pair, Q_e , may or may not define an angle via (1–8). In any event, at least one of $\pm Q_e$ defines such an angle. In the case that Q_e defines an angle, denote the latter as θ_e , and in the case that $-Q_e$ defines an angle, denote the latter by θ_{-e} . These angles are the zeros of the function α_Q . Moreover, since the derivative of α_Q is positive at θ_e and negative at θ_{-e} , the following conditions are necessary and sufficient for α_Q to be positive on $(\theta_o, \theta_{o'})$:

- (8–1) • $\theta_o > \theta_e$ if only θ_e is defined.
 • $\theta_{o'} < \theta_{-e}$ if only θ_{-e} is defined.
 • $\theta_e < \theta_o < \theta_{o'} < \theta_{-e}$ if both θ_e and θ_{-e} are defined and $\theta_e > \theta_{-e}$.
 • Either $\theta_{o'} < \theta_{-e}$ or $\theta_o > \theta_e$ if both θ_e and θ_{-e} are defined and $\theta_{-e} < \theta_e$.

In this regard, note that θ_{-e} is not defined when $q_e > 0$ and $q_e/|q_{e'}| < \sqrt{\frac{3}{2}}$. On the otherhand, θ_e is not defined when $q_e < 0$ and $-q_e/|q_{e'}| < \sqrt{\frac{3}{2}}$. If both are defined, then $\theta_e < \theta_{-e}$ if $q_{e'}$ is positive, and $\theta_e > \theta_{-e}$ if $q_{e'}$ is negative.

Note that angle inequalities can be interpreted directly in terms of integer pairs. To elaborate, suppose that θ_Q is defined via (1–8) from an integer pair (q, q') and that θ_P is defined from another integer pair, (p, p') . Then

- (8–2) • If p is negative and q' is positive, then $\theta_P > \theta_Q$.
 • If p' and q' have the same sign, then $\theta_P > \theta_Q$ if $p'q - pq' < 0$.

With (8–2) in hand, the conditions in (8–1) can be written directly in terms of the data given in \hat{A} and T .

Proof of Proposition 8.2 The beginning of Subsection 2.A explains why the stated conditions are necessary. The strategy to prove that the stated conditions are sufficient is as follows: A given graph T is approximated by a sequence of graphs, all mutually homotopic, and chosen so that the corresponding versions of $\mathcal{M}^*_{\hat{A},(\cdot)}$ are nonempty. A sequence is constructed whose j 'th element is from the j 'th version of $\mathcal{M}^*_{\hat{A},(\cdot)}$. The latter sequence is then seen to converge to an element in $\mathcal{M}^*_{\hat{A},T}$. The four steps that follow give the details. In this regard, note that the convergence arguments are very much the same as those in Subsection 7.D and so only the novel points are noted.

Step 1 To start, say that a graph T of the sort that is described in Subsection 6.A is generic when it has the following properties: All multivalent vertices are either bivalent or trivalent, the trivalent vertex angles are pairwise distinct, and they are distinct from all bivalent vertex angles. [15, Theorem 1.3] asserts that $\mathcal{M}^*_{\hat{A},T}$ is nonempty when T is generic and obeys the stated conditions in Proposition 8.2.

Step 2 Let T denote a graph from Subsection 6.A and let $o \in T$ denote a bivalent vertex. Consider modifying T by replacing the graph $\underline{\Gamma}_o$ by a 'less valent' graph, $\underline{\Gamma}'_o$. This is done as follows: Suppose first that $\underline{\Gamma}_o$ has a vertex with valency greater than 4. Let v denote the latter. The new graph, $\underline{\Gamma}'_o$, is obtained from $\underline{\Gamma}_o - v$ by attaching the now dangling incident half-arcs that are incident to v to the vertices in a graph with two vertices and one arc between them. Three of the dangling half-arcs are attached to one of these two vertices and the remaining half-arcs are attached to the other. Now care must be taken here with the choice of the first three so as to insure that the arcs in the new graph have consistent labels by pairs of edges. In particular, this is done as follows: Take the first incident half-arc to be oriented as an incoming arc, and let (e, e') denote its edge pair label. The second incident half-arc must be the arc that follows the first in ℓ_{oe} . Note that it must also be distinct from the first arc. As a consequence, its pair label has the form (e, e'') . This second incident half-arc is outgoing. The third incident half-arc should be the arc that follows the first on $\ell_{oe'}$. It must also be distinct arc from the first arc. Thus, its edge label has the form (e''', e') . It too is outgoing. Note that the case $e''' = e$ is allowed. The new arc that runs between the two new vertices is oriented as an incoming half-arc and labeled by the edge pair (e''', e'') . One of the two new vertices should be labeled with the integer 0, the other with the integer that labels v . Note that there is a completely analogous construction that has all arc orientations reversed. The operation just described can be performed on any vertex with valency greater than four.

Suppose next that $\underline{\Gamma}_o$ has only 4-valent and bivalent vertices, but at least one 4-valent vertex with a non-zero integer label. Let v denote one of the latter vertices. In this

case, $\underline{\Gamma}'_o$ is obtained from $\underline{\Gamma}_o - v$ in the manner just described. One of the vertices in the new graph is bivalent and the other is 4-valent. The bivalent vertex should be given v 's integer label and the other should be labeled with 0. \square

Consider now:

Lemma 8.3 *Suppose that T' is obtained from T by modifying one version of $\underline{\Gamma}_{(\cdot)}$ as just described. Then $\mathcal{M}^*_{\hat{A},T}$ is nonempty if and only if $\mathcal{M}^*_{\hat{A},T'}$ is nonempty.*

Proof of Lemma 8.3 Suppose first that $\mathcal{M}^*_{\hat{A},T'}$ is nonempty. Let $o \in T$ denote the vertex involved and v the vertex in $\underline{\Gamma}_o$. Consider a sequence, $\{\lambda_j\}$, in the T' version of the space in (6–15) whose coordinates are j -independent but for the coordinate in the T' simplex Δ'_o . In the latter, the coordinate for the arc between the two vertices in $\underline{\Gamma}'_o$ that replace v should converge to zero. The remaining arc coordinates should converge as $j \rightarrow \infty$ so as to define a point in Δ_o . Use a j -independent value in \mathbb{R} and the image of $\{\lambda_j\}$ in $O_{T'}/\text{Aut}(T')$ to define a sequence of equivalence classes in $\mathcal{M}^*_{\hat{A},T'}$. The discussion in Subsection 7.D can be repeated now with a minor modification to prove that the sequence in $\mathcal{M}^*_{\hat{A},T'}$ converges to an element in $\mathcal{M}^*_{\hat{A},T}$. The salient modification replaces (7–9) and (7–10) with the assertion that only the one arc in $\underline{\Gamma}'_o$ that does not come from $\underline{\Gamma}_o - v$ can correspond to an arc in C_{0j} that either lies entirely in a radius ε ball or where $|s| \geq R$. \square

Suppose next that $\mathcal{M}^*_{\hat{A},T}$ is nonempty. The fact that $\mathcal{M}^*_{\hat{A},T'}$ is nonempty follows using Lemma 5.4 with (2–4), (2–11) and the implicit function theorem.

Step 3 Suppose now that each version from T of $\underline{\Gamma}_{(\cdot)}$ has only bivalent or 4-valent vertices and that all 4-valent vertices are labeled by zero. Thus, no modifications as described in Part 2 are possible. However, suppose that o is a multivalent vertex in T and that $\underline{\Gamma}_o$ contains two or more vertices with one being a 4-valent vertex. In this case, T is modified to produce a graph T' as follows: To start, let $v \in \underline{\Gamma}_o$ denote a 4-valent vertex in $\underline{\Gamma}_o$. The arc segments that are incident to v are labeled by pairs of edges, but there are either 3 or 4 edges in total involved. In any case, two connect o to respective vertices whose angles are either both smaller or both larger than θ_o . What follows assumes the former; the argument in the latter case is omitted since it differs only cosmetically.

Let e and e' denote the two edges that connect o to vertices with larger angle. The complement in T of the interiors of e and e' is disconnected, with three components, these denoted as T_e , $T_{e'}$ and T_- . The graph T_- contains o , while T_e contains the vertex opposite o on the edge e and $T_{e'}$ contains the vertex opposite o on the edge

e' . The graph T' is the union of T_e , $T_{e'}$, T_- and a trivalent graph, Y , with one vertex and three edges. The two vertices at the top of the Y are identified in T' with the respective e and e' vertices in T_e and $T_{e'}$. The vertex at the bottom of the Y is identified in T' with the vertex o in T_- . The vertex at the center of the Y is denoted by o' , and its angle, $\theta_{o'}$, is slightly larger than θ_o . The corresponding graph, $\underline{\Gamma}_{o'}$, is a figure 8 where the two small circles are the versions of $\ell_{o'(\cdot)}$ that are labeled by the two top edges of the Y graph. Meanwhile, the loop that traces the figure 8 is the version of $\ell_{o'(\cdot)}$ that is labeled by the bottom edge in the Y graph.

Let \hat{o} denote the vertex in T' that corresponds to $o \in T_-$. The vertex angle of \hat{o} is that of o , thus θ_o . The graph $\underline{\Gamma}_{\hat{o}}$ is obtained from $\underline{\Gamma}_o - v$ as follows: Attach the incoming arc segment to v with label e to the outgoing arc segment with label e' . Likewise, attach the incoming arc segment with label e' to that outgoing arc segment with label e . Then, replace e and e' in all arc labels by the label of the bottom edge in the figure Y .

As can be readily verified, a graph T' as just described satisfies the conditions in Proposition 8.2 if T does, and if the vertex angle $\theta_{o'}$ is sufficiently close to θ_o . This understood, consider:

Lemma 8.4 *Suppose that T' is a graph as just described, with $\theta_{o'}$ very close to θ_o . Then $\mathcal{M}^*_{\hat{A},T}$ is nonempty if and only if $\mathcal{M}^*_{\hat{A},T'}$ is nonempty.*

This lemma is proved momentarily.

Proposition 8.2 is a corollary to Lemmas 8.3 and 8.4 together with [15, Theorem 1.3]. The reason is that any given T can be sequentially modified as described first by Lemma 8.3 and then as in Lemma 8.4 so as to obtain a graph that is generic in the sense that is used by Step 1.

Step 4 This step contains the following proof.

Proof of Lemma 8.4 Suppose first that $\mathcal{M}^*_{\hat{A},T'}$ is nonempty when $\theta_{o'}$ is sufficiently close to θ_o . Take a sequence of angles $\{\theta_j\}_{j=1,2,\dots}$ that converges to θ_o from above with each very close to θ_o . Let $\{T_j\}_{j=1,2,\dots}$ denote a corresponding sequence of graphs where the j 'th version is T' with $\theta_{o'} = \theta_j$. The plan is to define a corresponding sequence in $\mathcal{M}^*_{\hat{A}}$ whose j 'th element is a point in the $T' = T_j$ version of $\mathcal{M}^*_{\hat{A},T'}$ by using Theorem 6.2 and a point in the T_j version of the space in (6–15). For this purpose, it is necessary to first make the choices that are described in Parts 1 and 2 of Subsection 6.C. Since the T_j 's are pairwise homotopic, these choices can be made for all at once. Granted that this is done, choose a corresponding sequence, $\{\lambda_j\}$, with

λ_j a point in the T_j version of the space in (6–15). This sequence should be chosen so that all of the factors are independent of the index j . Now define a corresponding sequence in $\mathcal{M}^*_{\hat{A}}$ whose j 'th element is in the $T' = T_j$ version of $\mathcal{M}^*_{\hat{A},T'}$ and is obtained from λ_j using the T_j version of Theorem 6.2 with some j -independent choice for the \mathbb{R} factor. Arguments that are much like those in Subsection 7.D can be employed to prove that such a sequence converges in $\mathcal{M}^*_{\hat{A}}$ and that the limit is in some $\mathcal{M}^*_{\hat{A},T''}$ where T'' is a graph of a rather special sort. In particular, if T'' is not isomorphic to T , then it has a vertex, \hat{o} , with angle θ_o such that the replacement of $\underline{\Gamma}_{\hat{o}}$ with $\underline{\Gamma}_o$ makes a graph that is isomorphic to T . As is explained next, a careful choice for the constant sequence $\{\lambda_j\}$ gives a version of T'' that is isomorphic to T .

Care must be taken only with the coordinates of λ_j in the $\mathbb{R}_{\hat{o}} \times \Delta_{\hat{o}}$ and $\mathbb{R}_{o'}$ factors in (6–15). To specify the latter, return to the construction of the graph $\underline{\Gamma}_{\hat{o}}$ in T' from $\underline{\Gamma}_o$. Let \hat{e} denote the bottom edge of the Y -graph portion of T' . There is a map from $\underline{\Gamma}_{\hat{o}}$ to $\underline{\Gamma}_o$ that is 1–1 except for two points on $\ell_{\hat{o}\hat{e}}$ that are both mapped to the vertex v . Let π denote the latter map. By assumption, there is another vertex besides v in $\ell_{oe} \cup \ell_{o'e}$. For the sake of argument, suppose ℓ_{oe} has a second vertex, then let v_0 denote the vertex on ℓ_{oe} that starts the arc in ℓ_{oe} that ends at v . Meanwhile, the abstract version of the loop $\ell_{o'\hat{e}}$ from the figure 8 graph $\underline{\Gamma}_{o'}$ has two vertices. Let v_1 and v_2 denote the latter where the convention has the arc that starts at v_1 and ends at v_2 mapping to $\ell_{o',e}$ in $\underline{\Gamma}_{o'}$.

The next step to choosing the coordinates of λ_j requires the introduction of the map from Δ_o to $\Delta_{\hat{o}}$ that is defined so that $r \in \Delta_o$ sends $\gamma \subset \underline{\Gamma}_{\hat{o}}$ to $\sum_{\gamma' \subset \pi(\gamma)} r(\gamma')$. Fix $r \in \Delta_o$ and use its image under this map for λ_j 's factor in $\Delta_{\hat{o}}$. Also, fix $\tau \in \mathbb{R}_{\hat{o}}$. To define the coordinate of λ_j in the \mathbb{R}_o factor, observe first that the concatenating path set for the loop $\ell_{\hat{o}\hat{e}}$ in $\underline{\Gamma}_{\hat{o}}$ can be used to assign a value in \mathbb{R} to the vertex $\pi^{-1}(v_0)$ from any given $\tau \in \mathbb{R}_{\hat{o}}$. This value is defined by starting with τ and adding or subtracting suitable multiples of the values given by the image of r in $\Delta_{\hat{o}}$ to the arcs on a certain path from the distinguished vertex in $\underline{\Gamma}_{\hat{o}}$ to $\pi^{-1}(v_0)$. In particular, the path uses the concatenating path set to get to $\ell_{\hat{o}\hat{e}}$ and then proceeds in the oriented direction on $\ell_{\hat{o}\hat{e}}$ to the vertex $\pi^{-1}(v_0)$. All of this is done so as to be compatible with the parametrizing algorithm as described in Section 2. Let τ_0 denote the \mathbb{R} value defined in this way from the pair (τ, r) . Now let τ_1 denote the result of adding to τ_0 the value that the image of \mathbb{R} assigns to the arc in $\ell_{\hat{o}\hat{e}}$ that starts at $\pi^{-1}(v_0)$. Use $\frac{1}{2}(\tau_0 + \tau_1)$ for λ_j 's coordinate in $\mathbb{R}_{o'}$.

Arguments that are much like those used in Subsection 7.D prove that the sequence $\{\lambda_j\}$ as just described defines a sequence in $\mathcal{M}^*_{\hat{A}}$ whose limit is in $\mathcal{M}^*_{\hat{A},T}$. The proof that $\mathcal{M}^*_{\hat{A},T'}$ is nonempty if $\mathcal{M}^*_{\hat{A},T}$ is obtained using Lemma 5.4 with (2–4), (2–11) and the implicit function theorem. \square

8.C Proof of Proposition 6.1

In order to prove Proposition 6.1, it is necessary to return to the milieu of Proposition 8.1 and obtain a more refined picture of the image of the map p . For this purpose, remember that there are m angles in the image of the map f that can vary on a given stratum component S . These angles are distributed amongst the various components of $(0, \pi) - \Lambda_{+,B}$. If T and T' are homotopic graphs, and if $o \in T$ and $o' \in T'$ are corresponding vertices, then either $\theta_o = \theta_{o'}$ and this angle is in $\Lambda_{+,B}$, or else θ_o and $\theta_{o'}$ are in the same component of $(0, \pi) - \Lambda_{+,B}$.

With this last point in mind, suppose that T_1 and T_2 are homotopic graphs with the following property: Let $\theta_1 \in (0, \pi) - \Lambda_{+,B}$ and suppose that there are multivalent vertices in T_1 with angle θ_1 . Let $V \subset T_1$ denote the latter set, and also the corresponding set in T_2 . Suppose that the vertices in the T_2 version of V are assigned angle $\theta_2 > \theta_1$. In addition, assume that both the $T = T_1$ and $T = T_2$ versions of $\mathcal{M}^*_{\hat{A},T}$ are nonempty. For each angle $\theta \in (\theta_1, \theta_2)$, let T_θ denote the version of T that is obtained from T_1 by assigning the angle θ to the vertices in V . Thus, each T_θ is homotopic to T .

Proposition 6.1 is now a consequence of the following:

Lemma 8.5 *If both the T_1 and T_2 versions of $\mathcal{M}^*_{\hat{A},T}$ are nonempty, then such is the case for each $T = T_\theta$ version in the case that $\theta \in [\theta_1, \theta_2]$.*

Proof of Lemma 8.5 Consider the following scenario:

Scenario 1 *An angle σ lies in θ_1 's component of $(0, \pi) - \Lambda_{+,B}$, the $T = T_\sigma$ version of $\mathcal{M}^*_{\hat{A},T}$ is empty, but all $T = T_\theta$ versions of $\mathcal{M}^*_{\hat{A},T}$ are nonempty when $\theta \in [\theta_1, \sigma)$.*

Note that Proposition 8.1 finds some such σ when the $T = T_1$ version of $\mathcal{M}^*_{\hat{A},T}$ is nonempty. As is explained momentarily, Scenario 1 occurs if and only if σ coincides with θ_{-e} in the case that e is an edge that connects a vertex from the T_1 version of V to a vertex with angle less than θ_1 .

Here is a second scenario:

Scenario 2 *An angle σ' lies in θ_1 's component of $(0, \pi) - \Lambda_{+,B}$, the $T = T_{\sigma'}$ version of $\mathcal{M}^*_{\hat{A},T}$ is empty, but all $T = T_\theta$ versions of $\mathcal{M}^*_{\hat{A},T}$ are nonempty when $\theta \in (\sigma', \theta_2]$.*

A cosmetic modification to the arguments given below to prove the assertion just made about Scenario 1 prove the following: Scenario 2 occurs if and only if σ' coincides with θ_e in the case that e is an edge that connects a vertex from the T_2 version of V to a vertex with angle greater than θ_2 .

Note that Scenario 1 precludes Scenario 2, and vice versa. Indeed, were both scenarios to occur, then both T_1 and T_2 would be in violation of the conditions in Proposition 8.2. Lemma 8.5 is a consequence of this fact that the two scenarios cannot both occur.

To explain the assertion about Scenario 1, note first that if $\sigma = \theta_{-e}$ with e as described, then the $T = T_\sigma$ version of $\mathcal{M}^*_{\hat{A},T}$ is empty because T_σ violates the conditions stated in Proposition 8.2. Suppose next that σ is some as yet undistinguished angle that gives Scenario 1. By virtue of the fact that the various $\theta \in [\theta_1, \sigma)$ versions of T_θ differ only in their vertex angle labels, the choices that are made in Parts 1 and 2 of Subsection 6.C can be made in a θ -independent fashion. This then identifies all $T = T_\theta$ versions of the space that is depicted in (6–15). Fix some element in this space, λ , and a real number, s_0 . Next, let $T = T_\theta$ and use (s_0, λ) in $\mathbb{R} \times \mathcal{O}_T / \text{Aut}(T)$ to define via Theorem 6.2 a point in this same $T = T_\theta$ version of $\mathcal{M}^*_{\hat{A},T}$. Let c_θ denote the latter sequence. As will now be explained, arguments much like those used in Subsection 7.D prove the following: The sequence $\{c_\theta\}$ converges as $\theta \rightarrow \sigma$ in $\mathcal{M}^*_{\hat{A}}$ to a point in the $T = T_\sigma$ version of $\mathcal{M}^*_{\hat{A},T}$ unless $\sigma = \theta_{-e}$ with e an edge that connects a vertex in V to a vertex with angle less than θ_1 . Indeed, all of the arguments in Subsection 7.D can be made in this situation except possibly those in Step 3 in Subsection 7.D from the proof of Proposition 7.1. Moreover, the arguments in Step 3 from the proof of Proposition 7.1 can also be made if the angle σ (this is the angle to use for θ_* in Step 3 of Subsection 7.D) is such that $\alpha_Q(\sigma) > 0$ when Q is the integer pair for any edge that is incident to a vertex in V . \square

9 Geometric limits

The purpose of this last section is to indicate how the various codimension 1 strata fit one against another to make the whole of $\mathcal{M}^*_{\hat{A}}$. The resulting picture of $\mathcal{M}^*_{\hat{A}}$ is by no means complete, and perhaps not very illuminating. However, what follows should indicate how tools from the previous sections can be used to add missing details.

This story starts with the codimension 0 strata, and so let \mathcal{S} denote a component of such a stratum in $\mathcal{M}^*_{\hat{A}}$. What follows summarizes some of results about \mathcal{S} from the previous sections. To start, introduce k to denote $N_- + \hat{N} + \zeta_- + \zeta_+ - 2$. The component \mathcal{S} lies in a stratum of the form $\mathcal{S}_{B,c,d}$ where $B = \emptyset$ and $c = k$, and d is the partition of $N_+ + k$ with the maximal number of elements. As a consequence,

Proposition 5.1's integer m is equal to k also. Thus, if T is a graph that arises from an element in \mathcal{S} , then T has k trivalent vertices, with no two angles identical and none an angle from an integer pair of any $(0, +, \dots)$ element in \hat{A} . If o is a trivalent vertex, then Γ_o is compact, a figure 8 with one vertex. The other multivalent vertices are bivalent. If o is a bivalent vertex in T , then Γ_o is a circular graph whose vertices correspond to the $(0, +, \dots)$ elements in \hat{A} with integer pairs that define o 's angle, θ_o , via (1–8).

This all translates into the following geometry: Suppose that (C_0, ϕ) defines an element in \mathcal{S} . Then the pull-back of θ to C_0 has non-degenerate critical points; the critical values are pairwise distinct, and none is an $|s| \rightarrow \infty$ limit of θ on a concave side end of C_0 where $\lim_{|s| \rightarrow \infty} \theta \in (0, \pi)$. If E is such a concave side end, then E 's version of (2–4) has integer $n_E = 1$. On the otherhand, if E is a convex side end where $\lim_{|s| \rightarrow \infty} \theta \in (0, \pi)$, then the integer n_E is zero. Finally, C_0 has transversal intersections with the $\theta = 0$ and $\theta = \pi$ cylinders.

To picture \mathcal{S} , note first that the image of Subsection 8.A's map \mathfrak{p} can be viewed as a k -dimensional product of simplices in $\times_k((0, \pi) - \Lambda_{+, \phi})$, this denoted in what follows as Δ^k . Then Subsection 6.C's map provides an orbifold diffeomorphism that identifies \mathcal{S} as $\mathbb{R} \times \mathbb{O}$, where \mathbb{O} is fibered by \mathfrak{p} over Δ^k . In addition, the typical fiber has the form $O_T / \text{Aut}(T)$ with T as just described. Because all such fibers have homotopic graphs, the fibration $\mathfrak{p}: \mathbb{O} \rightarrow \Delta^k$ can be trivialized. This is to say that there is an orbifold diffeomorphism from \mathcal{S} to

$$(9-1) \quad \mathbb{R} \times \mathcal{O} / \mathcal{A} \times \Delta^k,$$

where $\mathcal{O} = O_T$ and $\mathcal{A} = \text{Aut}(T)$ for a some fixed graph T .

The behavior of the codimension 0 strata near a given codimension 1 stratum component is rather benign by virtue of the fact that the codimension 1 strata are submanifolds where smooth and suborbifolds otherwise. To elaborate, remark that the \mathfrak{p} -image of a path in \mathcal{S} to a codimension 1 or larger stratum will limit to a boundary point of the closure of Δ^k in $\times_k((0, \pi) - \Lambda_{+, \phi})$. As explained below, each codimension 1 stratum component in the closure of \mathcal{S} correspond in a natural fashion to a codimension 1 face in this closure of Δ^k . In particular, if S_1 denotes a codimension 1 stratum component, then S_1 fibers over some $k - 1$ dimensional product of simplices, thus some Δ^{k-1} . In particular, S_1 is diffeomorphic as an orbifold to $\mathbb{R} \times \mathcal{O}_1 / \mathcal{A}^1 \times \Delta^{k-1}$ where \mathcal{O}_1 is some O_{T^1} and \mathcal{A}^1 the corresponding $\text{Aut}(T^1)$ with T^1 some fixed graph. In all cases, the group \mathcal{A}^1 has a representation in $\mathbb{Z}/2\mathbb{Z}$, and this understood, a neighborhood of S_1 in $\mathcal{M}^*_{\hat{A}, T}$ is diffeomorphic as an orbifold to

$$(9-2) \quad \mathbb{R} \times (\mathcal{O}_1 \times (-1, 1)) / \mathcal{A}^1 \times \Delta^{k-1},$$

where \mathcal{A}^1 acts on $(-1, 1)$ through the multiplicative action of $\mathbb{Z}/2\mathbb{Z}$ as ± 1 . In all cases, the stratum \mathcal{S}_1 appears in (9–2) as the locus where the coordinate in the $(-1, 1)$ factor is 0.

As it turns out, there is much more to say about how various codimension 0 and 1 strata fit around the codimension 2 strata in their closure. This aspect of the stratification is the focus of Subsection 9.A that follows.

Most of the rest of this section focuses on what can be viewed as the codimension 1 strata in a certain natural compactification of $\mathcal{M}^*_{\hat{A}}$. To elaborate for a moment on this point, remark that the identification given by Theorem 6.2 provides a natural compactification of any given $\mathcal{M}^*_{\hat{A}, T}$, this obtained from the compactification of \mathcal{O}_T that replaces each open simplex in (6–9) with the corresponding closed simplex. This and the replacement of Δ^k in (9–1) with its closure defines, up to an obvious factor of \mathbb{R} , a stratified space compactification of $\mathcal{M}^*_{\hat{A}}$. As it turns out, each added stratum in this compactification has a natural geometric interpretation in terms of multiply punctured sphere subvarieties. This interpretation is presented below in Sections 9.B for the additional codimension 1 strata in the compactifications of the various versions of $\mathcal{M}^*_{\hat{A}, T}$. Subsection 9.C goes on to describe Subsection 1.B’s compactification of the $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$ versions of $\mathcal{M}_{\hat{A}}$. The added codimension 1 strata for the compactification of $\mathcal{M}^*_{\hat{A}}$ when $N_- + \hat{N} + \zeta_- + \zeta_+ > 2$ are described in Sections 9.D and 9.E.

9.A The codimension one strata

The components of the codimension 1 strata in the case that $N_- + \hat{N} + \zeta_+ + \zeta_- \equiv k + 2$ are characterized in part by the structure of the graph that arises from a typical element. Here are the four possibilities for a component of a codimension 1 stratum:

- (9–3)
- *The graph has k trivalent vertices where precisely one pair have identical angles. Even so, no trivalent vertex angle comes via (1–8) from an integer pair of a $(0, +, \dots)$ element in \hat{A} .*
 - *The graph has $k - 2$ trivalent vertices and one 4–valent vertex. No two have the same angle and none comes via (1–8) from an integer pair of a $(0, +, \dots)$ element in \hat{A} . In addition, the 4–valent vertex is assigned a graph with two vertices.*
 - *The graph has $k - 1$ trivalent vertices and none comes via (1–8) from an integer pair of a $(0, +, \dots)$ element in \hat{A} . Meanwhile, there are $N_+ + 1$ bivalent vertices.*

- The graph has k trivalent vertices with pairwise distinct angles, and precisely one such angle comes via (1–8) from an integer pair of a $(0, +, \dots)$ element in \hat{A} . The latter has a graph with one vertex labeled with 0.

These four cases correspond to the following sorts of codimension 1 faces in the simplex Δ^k that appears in (9–1). The first and second points in (9–3) arise when two or more trivalent vertex angles lie in the same component of $(0, \pi) - \Lambda_{+, \emptyset}$. The first point can arise when there are no multivalent vertex angles between the angles of a pair of trivalent vertices from distinct edges in T . The second point can arise when no multivalent vertex angle lies between the angles of two trivalent vertices that share an edge.

The third point in (9–3) can arise when no multivalent vertex angle lies between the angles of a monovalent and trivalent vertex from a single edge. The fourth point in (9–3) can occur when no multivalent vertex angle lies between the angles of a bivalent and trivalent vertex from a single edge.

There is a corresponding geometric interpretation to the four strata in (9–3). To say more in this regard, suppose that (C_0, ϕ) defines an element in a codimension 1 stratum. If the component is characterized by the first or second points in (9–3), then there are two critical points of θ on C_0 with the same critical value in $(0, \pi)$. In the case of the first point, the corresponding θ level sets are disjoint; and they are not disjoint in the case of the second point. In the case of the third point, a convex side end version of (2–4) has $n_E = 1$. In the case of the fourth point, a critical value of θ coincides with the $|s| \rightarrow \infty$ limit of θ on some concave side end.

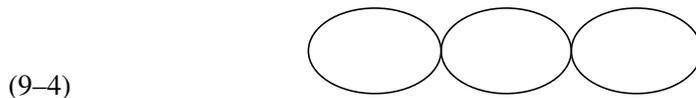
As is explained in the final subsection, sequences in a top dimensional stratum where a trivalent vertex angle limits to 0 or π can not converge to a codimension 1 stratum in $\mathcal{M}^*_{\hat{A}}$. Faces of the closure of Δ^k in $\times_k[0, \pi]$ that lie on the boundary of $\times_k[0, \pi]$ can give rise to codimension 2 strata.

The subsequent five parts of this subsection describe how the codimension 0 and the codimension 1 strata fit together inside $\mathcal{M}^*_{\hat{A}}$. The proofs of the assertions that are made below are omitted except for the comment that follows because the arguments in all cases are lengthy yet introduce no fundamentally new ideas. Here is the one comment: The proofs use the implicit function theorem to establish the pictures that are presented below of the relevant strata of $\mathcal{M}^*_{\hat{A}}$. In this regard, the arguments use the techniques that have already been introduced in Subsection 7.D, in much the same manner as they are used in Subsection 7.D, to prove that the picture that is provided by the implicit function theorem contains a full neighborhood of the stratum in question.

Part 1 This part describes a neighborhood of the strata whose elements have graphs that are described by the first point in (9–3). There are two cases to consider; the distinction is whether the group \mathcal{A}^1 that appears in (9–2) is larger than the automorphism group of the graphs that arise from elements in the nearby codimension 0 strata. To elaborate, the automorphism group of a graph from a codimension 0 stratum element must fix all trivalent vertices since these have distinct angles. However, the automorphism group of the codimension 1 stratum can, in principle, interchange the two trivalent vertices that share the same angle. Granted this, let \mathcal{A} denote the automorphism group of a graph from an element in a nearby codimension zero stratum and \mathcal{A}^1 denote that for a graph from an element in the codimension 1 stratum. The two cases under consideration here are those where $\mathcal{A}^1 \approx \mathcal{A}$ and where \mathcal{A}^1 is the semi-direct product of $\mathbb{Z}/2\mathbb{Z}$ with \mathcal{A} . In the former case, \mathcal{A}^1 acts trivially on the $(-1, 1)$ factor in (9–2). In the other case, \mathcal{A}^1 acts via its evident projection to $\mathbb{Z}/2\mathbb{Z}$.

To explain how this dichotomy of automorphism groups arises, let T^1 denote the graph from a typical element in the codimension 1 stratum. Thus, \mathcal{A}^1 is isomorphic to $\text{Aut}(T^1)$. The graph T^1 has an \mathcal{A}^1 -invariant, trivalent vertex o with the following property: Let T_1 , T_2 and T_3 denote the closures of the three components of $T^1 - o$. The labeling is such that T_1 and T_2 each contain one of the two trivalent vertices with equal angle. If T_1 is not isomorphic to T_2 , then $\mathcal{A}^1 \approx \mathcal{A}$. If T_1 is isomorphic to T_2 , then the distinguished $\mathbb{Z}/2\mathbb{Z}$ subgroup in \mathcal{A}^1 switches T_1 with T_2 .

Part 2 This part and the next part of the subsection consider the components of the codimension 1 stratum whose elements have graphs that are characterized by the second point in (9–3). There are also various cases to consider here. To elaborate, let T^1 denote such a graph, and let $o \in T^1$ denote the 4-valent vertex. Introduce E_- and E_+ to denote the sets of incident edges to o that respectively connect o to vertices with smaller angle and with larger angle. In the first case, either E_+ or E_- has a single edge. In the second case, both have two edges. Note that in the first case, the graph Γ_o associated to o has the form:



when E_+ and E_- each have two edges, then the associated graph Γ_o can be either the graph in (9–4) or the graph that follows.



This part of the subsection focuses on the case where E_- has three of o 's incident edges. The story when E_+ has three edges is identical but for notation to that told here. Part 3 of the subsection considers the case when both E_- and E_+ have two edges.

To start, label the three edges in E_- as $\{e_1, e_2, e_3\}$. The corresponding versions of $\ell_{o(\cdot)}$ label the three circles depicted in (9–4). Note that the central circle is distinguished, and each of the three edges here can have the central circle for its version of $\ell_{o(\cdot)}$. This is an important fact in what follows because it indicates that there can be three distinct codimension 1 stratum components involved.

Suppose that e_2 occupies the central circle. The image of $\mathcal{A}^1 = \text{Aut}(T^1)$ in Aut_o is either trivial or $\mathbb{Z}/(2\mathbb{Z})$. In this regard, let T_1, T_2 and T_3 denote the closures in T^1 of the respective components of $T^1 - o$ that contain the interiors of e_1, e_2 and e_3 . The case where the Aut_o image of \mathcal{A}^1 is $\mathbb{Z}/(2\mathbb{Z})$ arises when T_1 and T_3 are isomorphic.

Let \mathcal{S}_1 denote the corresponding codimension 1 stratum of $\mathcal{M}^*_{\hat{A}}$. In either case, a neighborhood of \mathcal{S}_1 in $\mathcal{M}^*_{\hat{A}}$ is diffeomorphic as an orbifold to the space depicted in (9–2). In this case, the group \mathcal{A}^1 acts on the $(-1, 1)$ factor via its image in Aut_o , either trivially or as the $\mathbb{Z}/(2\mathbb{Z})$ action as multiplication by ± 1 .

There is more to the picture just presented by virtue of the fact that any one of the three edges in E_- can label the middle circle in (9–4) and so there can be from 1 to 3 distinct codimension one strata involved here. What follows describes how these codimension 1 strata and their neighborhoods fit together in $\mathcal{M}^*_{\hat{A}}$.

To start, consider a graph, T , as described in the second point of (9–3) save that its 4-valent vertex, o , has a version of Γ_o with only one vertex. Thus, Γ_o is the union of three circles that intersect at a single point. These circles are the $e = e_1, e_2$ and e_3 versions of ℓ_{oe} . In this case, Aut_o is a subgroup of $\mathbb{Z}/(3\mathbb{Z})$, thus trivial if T_1, T_2 and T_3 are not mutually isomorphic. Note that there are two distinct versions of Γ_o in any case, these are distinguished as follows: Let \hat{e} denote the single edge in E_+ . Then $\ell_{o\hat{e}}$ has three vertices and three arcs. Each arc is labeled by e_1, e_2 and e_3 ; and the two versions of Γ_o are distinguished by the two possible cyclic orderings of the arcs that comprise $\ell_{o\hat{e}}$. Thus, there are, in fact, two possibilities for T . However, the two versions are isomorphic when two or more from the collection $\{T_j\}_{j=1,2,3}$ are isomorphic.

The homotopy type of a graph T as just described labels a codimension 2 stratum component, this diffeomorphic as an orbifold to $\mathbb{R} \times \mathcal{O}_2/\mathcal{A}^2 \times \Delta^{k-2}$. Here, \mathcal{O}_2 is diffeomorphic to \mathcal{O}_T and \mathcal{A}^2 isomorphic to $\text{Aut}(T)$.

Now introduce $Z \subset \mathbb{C}\mathbb{P}^1 = \mathbb{C} \cup \infty$ to denote the complement of the three cube roots of -1 . Thus, Z is a model for a standard ‘pair of pants’. Let $Z_1 \subset Z \cap \mathbb{C}\mathbb{P}^1$ denote

the three rays that go through 0 , ∞ , and the respective cube roots of 1 . The drawing in (1–27) depicts Z_1 in the finite part of Z .

There are three cases to consider. In the first, T_1 , T_2 and T_3 are pairwise non-isomorphic. In this case, there are three distinct codimension 1 strata components involved; and a neighborhood in $\mathcal{M}^*_{\hat{A}}$ of their union is diffeomorphic as an orbifold to

$$(9-6) \quad \mathbb{R} \times \mathcal{O}_2 / \mathcal{A}^2 \times Z \times \Delta^{k-2}.$$

Here, the three codimension 1 strata correspond to the loci $Z_1 - \{0, \infty\}$, and the two codimension 2 strata correspond to 0 and ∞ .

The second case occurs when two of the three graphs from $\{T_j\}_{j=1,2,3}$ are isomorphic. In this case, there can be as few as two distinct codimension 1 strata components involved. To describe a neighborhood of the union of these strata, note that the image of \mathcal{A}^2 in Aut_o in this case is $\mathbb{Z}/2\mathbb{Z}$. Thus, \mathcal{A}^2 acts on $\mathbb{C} \cup \infty$ through the action of $\mathbb{Z}/2\mathbb{Z}$ on $\mathbb{C} \cup \infty$ whose generator sends $z \rightarrow \bar{z}$. Note that this action commutes with the $\mathbb{Z}/2\mathbb{Z}$ action on $\mathbb{C} \cup \infty$ as $z \rightarrow -\bar{z}^{-1}$ and whose orbit space is $\mathbb{R}\mathbb{P}^2$. As a consequence \mathcal{A}^2 acts on $\mathbb{R}\mathbb{P}^2$ as well. With all of this understood, a neighborhood of their union in $\mathcal{M}^*_{\hat{A}}$ is diffeomorphic as an orbifold to

$$(9-7) \quad \mathbb{R} \times (\mathcal{O}_2 \times \bar{Z}) / \mathcal{A}^2 \times \Delta^{k-2},$$

where $\bar{Z} \subset \mathbb{R}\mathbb{P}^2$ is the image of Z . Here, the two codimension 1 strata correspond to the image in $\mathbb{R}\mathbb{P}^2$ of $Z_1 - \{0, \infty\}$ and the codimension 2 stratum to the image of $\{0, \infty\}$. In this regard, note that the rays of Z_1 through the non-trivial cube roots of 1 have the same image in $\mathbb{R}\mathbb{P}^2$, this distinct from the ray through 1 .

The final case occurs when T_1 , T_2 and T_3 are pairwise isomorphic. To picture this case, introduce the six element permutation group G of the set $\{1, 2, 3\}$. This group has the $\mathbb{Z}/3\mathbb{Z}$ subgroup of elements that preserves the cyclic order. The quotient of G by the latter group gives the parity homomorphism $G \rightarrow \mathbb{Z}/2\mathbb{Z}$. The group G acts on $\mathbb{C}\mathbb{P}^1$ as the group of complex automorphisms that permutes the cube roots of -1 . In the latter guise, G has generators $z \rightarrow 1/z$ and $z \rightarrow \lambda z$ where λ is a favorite, non-trivial cube root of 1 . Note that G also permutes the cube roots of 1 . In any event, G acts on $Z \subset \mathbb{C}\mathbb{P}^1$. This understood, a neighborhood of the codimension 2 stratum in $\mathcal{M}^*_{\hat{A}}$ is diffeomorphic as an orbifold to

$$(9-8) \quad \mathbb{R} \times [(\mathcal{O}_2 \times G) / \mathcal{A}^2 \times_G Z] \times \Delta^{k-2},$$

where \mathcal{A}^2 acts on G on its left side through $\text{Aut}_o = \mathbb{Z}/(3\mathbb{Z})$. Meanwhile, G acts on itself on its right side and on Z as noted above. The codimension 1 stratum appears in

(9–8) as the image of $Z_1 - \{0, \infty\}$. The codimension 2 strata appear as the image of 0 and ∞ from Z .

Part 3 What follows here is a description of a neighborhood in $\mathcal{M}^*_{\hat{A}}$ of the components of the codimension 1 strata whose elements have graphs that are characterized by the second point in (9–3) in the case that the 4–valent vertex has two edges that connect it to vertices with larger angle and two that connect it to vertices with smaller angle. To set the stage, let T^1 denote the graph in question and o the 4–valent vertex. Let e_- and e_-' denote the edges in E_- and let e_+ and e_+' denote those in E_+ . In what follows, α_- , α_-' , α_+ and α_+' denote $\alpha_Q(\theta_o)$ in the case that $Q = Q_e$ with e respectively the edges e_- , e_-' , e_+ and e_+' . Each of these functions is positive at θ_o , and by virtue of (2–17),

$$(9-9) \quad \alpha_- + \alpha_-' = \alpha_+ + \alpha_+'.$$

A distinction must now be made between the cases when the unordered sets $\{\alpha_-, \alpha_-'\}$ and $\{\alpha_+, \alpha_+'\}$ are distinct, and when they agree. Considered first as Case 1 is that when these two sets are distinct.

Case 1 Make the convention that $\alpha_- \geq \alpha_-'$ and that $\alpha_+ \geq \alpha_+'$. At least one of these is a strict inequality. Since $\alpha_- \neq \alpha_+$, one or the other is larger; and as the description for the $\alpha_- > \alpha_+$ case is identical to that when $\alpha_+ > \alpha_-$, the former is left to the reader. Thus, in what follows,

$$(9-10) \quad \alpha_+ > \alpha_- \geq \alpha_-' > \alpha_+'.$$

As noted briefly above, the graph Γ_o can be either as depicted in (9–4) or as in (9–5). In the case of (9–4), the assumption in (9–10) implies that there are only two consistent labelings of the arcs with pairs of incident edges. These are as follows:

- (9–11) • Both of the middle circle’s arcs are labeled by (e_-, e_+) , and the other two arcs are labeled by (e_-, e_+') and (e_-', e_+) .
- Both of the middle circle’s arcs are labeled by (e_-', e_+) , and the other two arcs are labeled by (e_-', e_+') and (e_-, e_+) .

Let T_- , T_-' , T_+ and T_+' denote the closures of the four components of $T^1 - o$, here labeled so as to indicate which contains which incident edge. By virtue of (9–10), the graphs T_+ and T_+' are not isomorphic. However, T_- and T_-' may be isomorphic. In the latter case, the two versions of T^1 that correspond to the two labelings in (9–11) are isomorphic. Otherwise, the two versions are not isomorphic. In any event, the image of $\text{Aut}(T^1)$ in Aut_o is trivial.

Granted what has just been said, a component of a stratum whose typical element has a graph such as either version of T^1 as just described is diffeomorphic as an orbifold to what is depicted in (9–2) with $\mathcal{A}^1 = \text{Aut}(T^1)$ acting trivially on $(-1, 1)$.

When the graph Γ_o is as depicted in (9–5), there is, up to isomorphism, only one way to label the arcs with pairs of edges. From right to left, the labeling is:

$$(9-12) \quad (e_-, e_+), \quad (e_-, e_+'), \quad (e_-' , e_+'), \quad (e_-' , e_+).$$

In the case that T_- is not isomorphic to T_-' , the image of $\text{Aut}(T^1)$ in Aut_o is trivial. In the case that these two graphs are isomorphic, the image is $\mathbb{Z}/2\mathbb{Z}$. In the latter case, the element -1 acts so as to rotate the diagram in (9–5) by π radians.

A neighborhood in $\mathcal{M}^*_\mathcal{A}$ of a component of a codimension 1 stratum that yields a the graph T^1 is diffeomorphic as an orbifold to the space in (9–2). In this regard, the action of $\mathcal{A}^1 = \text{Aut}(T^1)$ on $(-1, 1)$ is trivial if T_- and T_-' are not isomorphic. If they are isomorphic, then the action on $(-1, 1)$ is via its image in Aut_o as the multiplicative action of $\{\pm 1\} = \mathbb{Z}/2\mathbb{Z}$.

A particularly intriguing point here concerns how the codimension 0 and 1 strata fit around a codimension 2 stratum. The intrigue stems from the fact that the closures of the respective (9–4) and (9–5) cases for \mathcal{S}_1 intersect. In this regard, the relevant graph that describes the intersection obeys the second point of (9–3) save that the graph for the 4–valent vertex has but one vertex. Thus, the latter graph, Γ_o , consists of three circles that meet at a single point. Such a graph is obtained from (9–4) by shrinking either of the two arcs in the middle circle. The same graph is produced by shrinking either arc. However, the two labeled versions of (9–11) produce distinctly labeled versions of such a 1–vertex and 3–circle Γ_o . In particular, the respective top and bottom versions of (9–11) produce such a Γ_o whose arcs are labeled by

$$(9-13) \quad (e_-, e_+'), (e_-, e_+), (e_-' , e_+) \quad \text{and} \quad (e_-' , e_+'), (e_-' , e_+), (e_-, e_+).$$

Meanwhile, (9–5) yields a one vertex and three circle graph in two ways since either of the arcs labeled with e_+' can be shrunk. Note that the shrinking of an e_+ labeled arc is prohibited by the assumption in (9–10). There are two resulting versions of Γ_o , these distinguished by the arc labelings in (9–13).

Note that in the case that T_- and T_-' are isomorphic, the two versions of T that correspond to the two versions of Γ_o with the respective arc labelings in (9–13) are isomorphic. Otherwise, they are not isomorphic graphs.

Granted all of this, a picture of a neighborhood in $\mathcal{M}^*_\mathcal{A}$ of all of these strata is obtained as follows: Let T denote one or the other of the graphs that result from the two cases

in (9–13), and let \mathcal{O}_2 denote \mathcal{O}_T and \mathcal{A}^2 denote $\text{Aut}(T)$. Next, let $Z \subset \mathbb{C}$ denote the complement of 2 and -2 . Let $Z_1 \subset Z$ denote the union of the circles of radius 1 centered at 2 and -2 together with the arc $[-1, 1]$ along the real axis between them. The drawing in (1–28) depicts Z_1 in Z . In the case that T_- is not isomorphic to T_-' , a neighborhood of the codimension 1 and 2 strata just described is diffeomorphic as an orbifold to

$$(9-14) \quad \mathbb{R} \times \mathcal{O}_2 / \mathcal{A}^2 \times Z \times \Delta^{k-2}.$$

Here, the codimension 1 strata correspond to the complement in Z_1 of 1 and -1 , while the latter correspond to the codimension 2 strata. In this regard, subvarieties that map to the arc $(-1, 1) \subset Z_1$ have the 4-valent vertex graph in (9–5). Those on the two circular parts of Z_1 have graphs as in (9–4), and the two circles are distinguished by the two labelings in (9–11). Meanwhile, the subvarieties in the codimension 2 strata that map to $+1$ are distinguished from those that map to -1 by the two labelings in (9–13).

In the case that T_- and T_-' are isomorphic, then a neighborhood of the strata is diffeomorphic as an orbifold to

$$(9-15) \quad \mathbb{R} \times \mathcal{O}_2 / \mathcal{A}^2 \times \bar{Z} \times \Delta^{k-2},$$

where \bar{Z} is the quotient of Z via the action by multiplication of $\{\pm 1\}$ on \mathbb{C} .

Case 2 Consider now the case that the sets $\{\alpha_-, \alpha_-'\}$ and $\{\alpha_+, \alpha_+'\}$ are identical. Agree to label things so that $\alpha_- \geq \alpha_-'$ and $\alpha_+ \geq \alpha_+'$. Assume first that these are strict inequalities. In this case, the version of Γ_o in (9–4) has but one allowable arc labeling, this where the middle arcs are labeled by (e_-, e_+) and the outer two by (e_-, e_+') and (e_-', e_+) . Shrinking either of the two middle arcs yields a graph with three circles and one vertex. The two graphs so obtained have identical edge labels, (e_-, e_+) , (e_-, e_+') and (e_-', e_+) . The corresponding two versions of T are thus isomorphic.

In the case that the version of Γ_o is given by (9–5), there is, as before, only one way to label the graph in (9–5) by incident half-arcs. Note that maps in the corresponding Δ_o must assign a greater value to the (e_-, e_+) arc than to the (e_-', e_+') arc. Meanwhile, the arc labeled by (e_-, e_+') must be assigned the same value as the (e_-', e_+) arc. This being the case, the boundary of Δ_o can be reached either by shrinking the value of the (e_-', e_+') arc or by simultaneously shrinking the values of the (e_-, e_+') and (e_-', e_+) arcs. The face that corresponds to giving the (e_-', e_+') arc value zero corresponds to the codimension 2 stratum whose elements have a graph T as just described in the preceding paragraph. The other face of Δ_o does not correspond to an element in

$\mathcal{M}^*_{\hat{A}}$. Indeed, according to Lemma 9.6 to come, points on the latter face correspond to reducible subvarieties.

Granted this, here is a picture of a neighborhood in $\mathcal{M}^*_{\hat{A}}$ of these strata. Let $Z \subset \mathbb{C}$ denote the complement of the circle of radius 1 centered at 1, and let $Z_1 \subset Z$ denote the union of this circle with the negative real axis. Let T denote the graph that is referred to in the preceding paragraph. Set $\mathcal{O}_2 = \mathcal{O}_T$ and $\mathcal{A}^2 = \text{Aut}(T)$. Then the strata in question have a neighborhood in $\mathcal{M}^*_{\hat{A}}$ that is diffeomorphic as an orbifold to (9–14) with Z as just described. In this new version, the codimension 1 strata correspond to $Z_1 - 0$ and the codimension 2 stratum corresponds to 0. Here, the circular part of Z_1 corresponds to the component of the codimension 1 stratum whose elements have 4-valent vertex graphs as depicted in (9–4). The negative real axis corresponds to the component whose elements have 4-valent graphs as depicted in (9–5).

The final case to consider is that where $\alpha_- = \alpha'_- = \alpha_+ = \alpha'_+$. In this case, the 4-valent vertex in the codimension 1 stratum must have the form depicted in (9–5). Just one component is involved. The image of the corresponding version of $\mathcal{A}^1 = \text{Aut}(T^1)$ in Aut_o is either trivial, $\mathbb{Z}/2\mathbb{Z}$ or $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. The last case occurs when T_- is isomorphic to T'_- and when T_+ is isomorphic to T'_+ . The $\mathbb{Z}/2\mathbb{Z}$ case occurs when one or the other of these pairs consist of isomorphic graphs, but not both. The trivial case occurs when neither pair has isomorphic graphs. A neighborhood in $\mathcal{M}^*_{\hat{A}}$ is diffeomorphic as an orbifold to what is depicted in (9–2) where \mathcal{A}^1 acts on $(1, -1)$ via its image in Aut_o , thus via its image in the group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. Here, the action of the latter on $(-1, 1)$ has both the $\mathbb{Z}/2\mathbb{Z}$ generators multiplying by -1 . Note that in this case, neither boundary face of the Δ_o factor in \mathcal{O}_2 corresponds to a subvariety in $\mathcal{M}^*_{\hat{A}}$; both correspond to reducible subvarieties.

Part 4 This part describes the neighborhood of the codimension 1 strata in $\mathcal{M}^*_{\hat{A}}$ whose elements have graphs that are characterized by the third point in (9–3). If (C_0, ϕ) defines a point on such a stratum, then the extra bivalent vertex in T_C has graph Γ_o that is a single circle with one vertex. As noted previously, the latter vertex corresponds to a convex side end of C_0 where the $|s| \rightarrow \infty$ limit of θ is neither 0 nor π , and whose version of (2–4) has $n_E = 1$. Let o denote this vertex and θ_o its angle. Elements near this stratum have a trivalent vertex, \hat{o} , that is very nearly θ_o . If $\theta_{\hat{o}} > \theta_o$, the \hat{o} version of E_- has two of the three incident edges; if $\theta_{\hat{o}} < \theta_o$, then the corresponding E_+ has two of the three incident edges. The graphs from elements on the two sides of this stratum would otherwise be homotopic.

Suppose that $\{(C_{0j}, \phi_j)\}_{j=1,2,\dots}$ is a sequence in $\mathcal{M}^*_{\hat{A}}$ with limit (C_0, ϕ) and if this sequence is not on the stratum in question, then each C_{0j} has a unique θ critical point

that is very close to θ_o . Let z_j denote the latter point. Then $\{\theta(z_j)\}_{j=1,2,\dots}$ converges to θ_o but $\{s(z_j)\}_{j=1,2,\dots}$ is unbounded from below. To put this in a colloquial fashion, the position of this one critical point moves as $j \rightarrow \infty$ to more and more negative values of s .

Part 5 This part describes the neighborhoods of the codimension 1 strata whose elements have graphs that are characterized in the fourth point of (9–3). There are now two cases to consider. In the first case, the graph T for C_0 has N_+ bivalent vertices. This is to say that θ on C_0 has k non-degenerate critical points with one having θ value in Λ_+ . However, the constant θ level set through this critical point is compact like the others; thus still a figure 8. In this case, a neighborhood of the stratum is not very interesting and won't be discussed further except to say that the union of the two abutting codimension zero strata are described by a version of (9–1).

The more interesting case occurs when there are $N_+ - 1$ bivalent vertices. Assuming that this is the situation, let o denote the trivalent vertex in T whose angle is in $\Lambda_{+,\emptyset}$. The corresponding graph $\underline{\Gamma}_o$ is a figure 8 with one 4-valent vertex labeled with o and with all other vertices being bivalent and having positive integer labels. These bivalent vertices are in 1–1 correspondence with the $(0, +, \dots)$ elements in \hat{A} whose integer pair defines θ_o via (1–8).

The various versions of $\underline{\Gamma}_o$ are characterized by the manner in which these bivalent vertices are distributed on the two circles that comprise the figure 8. In this regard, the possible isomorphism types for $\underline{\Gamma}_o$ are described momentarily with the help of an ordered pair (\wp', \wp'') of linearly ordered sets that partition the set of $(0, +, \dots)$ elements in \hat{A} whose integer pairs define θ_o in (1–8). For the purposes of this description, label the incident edges to o as e , e' and e'' with the convention that e' and e'' connect o to vertices with both angles either greater than θ_o or both angles less than θ_o . The pair (\wp', \wp'') defines a version of $\underline{\Gamma}_o$ by using \wp' to label the bivalent vertices that are met when circumnavigating $\ell_{oe'}$ in the oriented direction starting at the 4-valent vertex. The elements in \wp'' have the analogous interpretation with regards to $\ell_{oe''}$. Note that this labeling of the versions of $\underline{\Gamma}_o$ can be redundant. In particular, such will be the case when \hat{A} has multiple copies of some $(0, +, \dots)$ whose integer pair gives θ_o via (1–8). The labeling is also redundant when Q_e and $Q_{e'}$ agree, for then the respective versions of $\underline{\Gamma}_o$ that are defined by (\wp', \wp'') and (\wp'', \wp') are identical.

The closures of these various strata components intersect in codimension 2 strata and it is intriguing to see how all of these strata fit together. Since the story in general can be complicated, attention is restricted in what follows to the case when θ_o is determined by the integer pair from a single $(0, +, \dots)$ element in \hat{A} . To describe this case, let T denote the graph that is described by the fourth point in (9–3) but where the version of

$\underline{\Gamma}_o$ has a single 4-valent vertex with positive label. Let T_e , $T_{e'}$ and $T_{e''}$ denote the closures of the components of $T-o$, these labeled so as to indicate which incident edge is in which component. Now, let $Z \subset \mathbb{C}$ denote the complement of $\{1, -1\}$ and let $Z_1 \subset Z$ denote the figure 8 locus where $|z^2 - 1| = 1$. This is depicted schematically in (1–26). In the case that $T_{e'}$ is not isomorphic to $T_{e''}$, a neighborhood in $\mathcal{M}_{\hat{A}}^*$ of the strata involved is diffeomorphic as an orbifold to the space depicted in (9–4) with Z as just described, with $\mathcal{O}_2 = \mathcal{O}_T$, and with $\mathcal{A}^2 = \text{Aut}(T)$. In the case at hand, the two codimension 1 strata are the two components of $Z_1 - 0$, and the codimension 2 stratum corresponds to 0.

In the case that $T_{e'} = T_{e''}$, there is but a single codimension 1 stratum component and one codimension 2 stratum component involved. To describe a neighborhood in $\mathcal{M}_{\hat{A}}^*$ of these strata, note that in this case $\text{Aut}(T)$ has image $\mathbb{Z}/2\mathbb{Z}$ in Aut_o . Thus, $\text{Aut}(T)$ acts on \mathbb{C} via this image and the $\mathbb{Z}/2\mathbb{Z}$ action is via multiplication by $\{\pm 1\}$ on \mathbb{C} . This understood, a neighborhood in $\mathcal{M}_{\hat{A}}^*$ of the codimension 1 and codimension 2 strata is diffeomorphic as an orbifold to $\mathbb{R} \times (\mathcal{O}_2 \times Z)/\mathcal{A}^2 \times \Delta^{k-2}$ where $\mathcal{O}_2 = \mathcal{O}_T$, $\mathcal{A}^2 = \text{Aut}(T)$ and Z is as described in the preceding paragraph. In this case, the codimension 1 stratum is the image of $Z_1 - 0$ and the codimension 2 stratum is the image of 0.

9.B Limits of sequence in $\mathcal{M}_{\hat{A}, T}^*$ that do not converge in $\mathcal{O}_T / \text{Aut}(T)$

The space $\mathcal{O}_T / \text{Aut}(T)$ has a natural compactification as a stratified space that is obtained by replacing each open simplex in (6–9) by the corresponding closed simplex. All of the added points in this compactification label subvarieties that are geometric limits of elements in $\mathcal{M}_{\hat{A}, T}^*$. The two parts of this subsection describe the various limits for sequences that converge to the added codimension 1 strata of the compactification. Analogous assertions for the codimension greater than 1 added strata when $N_- + \hat{N} + \zeta_- + \zeta_+ > 2$ are more complicated and so left to the more industrious readers.

The proofs of the lemmas that appear are omitted but for some sporadic comments. This is because the argument in each case would lengthen an already long exposition; in any event, but for straightforward modifications, each argument repeats the those that appear in Subsection 7.D.

Part 1 To set the scene for this part of the subsection, a multivalent vertex $o \in T$ has been fixed along with an arc $\gamma \subset \underline{\Gamma}_o$. In addition, a sequence $\{\lambda_j\}_{j=1,2,\dots}$ has been specified in T 's version of the space depicted in (6–9) that is constant but for the Δ_o factor. In this regard, the Δ_o part of λ_j maps γ to a number, r_j , such that the resulting

sequence $\{r_j\}_{j=1,2,\dots}$ converges to zero. Meanwhile, the analogous sequences that are defined by the remaining arcs in $\underline{\Gamma}_o$ converge as $j \rightarrow \infty$ to positive numbers. With $\{\lambda_j\}$ so chosen, fix $s_0 \in \mathbb{R}$ and then use the image of (s_0, λ_j) in $\mathbb{R} \times O_T / \text{Aut}(T)$ with the inverse of the map from Subsection 6.C to define a sequence in $\mathcal{M}^*_{\hat{A}, T}$. This scenario is assumed implicitly in the statements of Lemmas 9.1–9.4 that follow.

By the way, the arc γ can not start and end at the same vertex because each element in Δ_o assigns the same value to any given arc that starts and ends at the same vertex. To explain, note that such an arc defines a non-trivial class in $H_1(\underline{\Gamma}_o; \mathbb{Z})$, and the latter can be written as $\sum_e c_e [\ell_{oe}]$ with $\{c_e\}$ a collection of integers. Let $Q = \sum_e c_e Q_e$. As a consequence of (6–6), each element in Δ_o must assign $\alpha_Q(\theta_o)$ to γ .

Here is one comment with regards to proving Lemmas 9.1–9.4: Proposition 7.1 is still valid under the assumptions just made. In fact, what follows are assumptions that guarantee the validity of Proposition 7.1.

(9–16) *Let $\{\lambda_j\}_{j=1,2,\dots}$ denote a sequence from T 's version of the space in (6–15) that is constant but for the Δ_o factor. Meanwhile, let $\{r_j\}$ denote the corresponding sequence of Δ_o factors. This sequence of maps should converge so that the following is true: Let \mathcal{A} denote the set of arcs in $\underline{\Gamma}_o$ for which $\lim_{j \rightarrow \infty} r_j(\cdot)$ is zero. Then $\cup_{\gamma \in \mathcal{A}} \gamma$ is simply connected.*

As illustrated by what is discussed in Case 2 of Part 3 in Subsection 9.A, there may be codimension 1 faces of Δ_o whose maps assign 0 to two or more arcs from $\underline{\Gamma}_o$. As explained in Part 2 of this subsection, limits to these faces violate the preceding assumption. In any event, such faces are considered in Part 2.

In the various cases described below, a certain graph, T' , is defined from T by collapsing γ to a point and suitably interpreting the result. In all of the cases, T' is a graph that has the properties stated in Subsection 6.A. In what follows, \hat{A}' denotes the corresponding asymptotic data set.

In the meantime, a stratum in the compactified version of $O_T / \text{Aut}(T)$ arises as the image of the face of the compact version of Δ_o whose maps send γ to zero. This stratum has a natural interpretation as $O_{T'} / \text{Aut}(T')$. In each case that follows, the explanation of this interpretation is straightforward and left to the reader. Granted this interpretation, the sequence $\{\lambda_j\}$ in each of the cases that follow has a well defined limit point in $O_{T'} / \text{Aut}(T')$. This point is denoted by $[\lambda_0]$. The T' versions of Theorem 6.2 and the inverse of the map from Subsection 6.C assign a point in $\mathcal{M}^*_{\hat{A}}$ the pair $(s_0, [\lambda_0])$. Use (C_0, ϕ_0) to denote the latter point in $\mathcal{M}^*_{\hat{A}}$.

In the first lemma that follows, T' is identical to T but for the graph assigned to o . The latter, $\underline{\Gamma}'_o$, is obtained from $\underline{\Gamma}_o$ by removing the interior of the arc γ and replacing

its two vertices with a single vertex whose label is the sum of the integers that label the vertices on γ .

Lemma 9.1 *Assume that at least one vertex on γ is labeled by the integer 0. In this case, the sequence in $\mathcal{M}^*_{\hat{A}}$ converges in $\mathcal{M}^*_{\hat{A}}$ to the point defined by (C_0, ϕ_0) in $\mathcal{M}^*_{\hat{A}, T'}$.*

The next lemma, assumes that the two vertices on γ have either both positive or both negative integer assignments. Let m_1 and m_2 denote their respective integer assignments. The two vertices on γ correspond to respective 4-tuples a_1 and a_2 in \hat{A} of the form $(0, \varepsilon, |m_1|P)$ and $(0, \varepsilon, |m_2|P)$ where P is the relatively prime integer pair that defines θ_o via (1–8) and $\varepsilon = \text{sign}(m_1)$. Let \hat{A}' denote the asymptotic data set that is obtained from \hat{A} by first removing both a_1 and a_2 , and then adding $a' = (0, \varepsilon, |m_1 + m_2|P)$. Meanwhile, let T' denote the graph that is identical to T but for the graph assigned to o . This new assignment, $\underline{\Gamma}'_o$, is obtained from $\underline{\Gamma}_o$ by removing the interior of γ and replacing its two vertices by a single vertex whose integer assignment is $m_1 + m_2$.

Lemma 9.2 *With the circumstances as just indicated, the sequence converges to the element defined by (C_0, ϕ) in $\mathcal{M}^*_{\hat{A}', T'}$ in the following sense: The sequence in $\mathcal{M}^*_{\hat{A}, T}$ comes from a sequence $\{(C_{0j}, \phi_j)\}_{j=1,2,\dots}$ such that*

- $\lim_{j \rightarrow \infty} \int_{C_{0j}} \phi_j^* \varpi = \int_{C_0} \phi^* \varpi$ for each compactly supported 2-form ϖ .
- The following limit exists and is zero:

$$\lim_{j \rightarrow \infty} \left[\sup_{z \in C_{0j}} \text{dist}(\phi_j(z), \phi(C_0)) + \sup_{z \in C_0} \text{dist}(\phi_j(C_{0j}), \phi(z)) \right].$$

Moreover, there exists an end, $E \subset C_0$, that gives the 4-tuple a' and has the following significance: Given $R \gg 1$ and any sufficiently large j , there is a smooth, proper embedding, ψ_j , from the complement in C_0 of the $|s| > R$ portion of E into C_{0j} ; and these are such that

- The complement of the image of ψ_j is a properly embedded, thrice punctured sphere that contains two ends of C_{0j} , these supplying the elements a_1 and a_2 to \hat{A} .
- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi_j(z), \phi(z)) = 0$ for all z in the domain of ψ_j .
- $|\bar{\partial}\psi_j| \ll |\partial\psi_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi_j)} |\bar{\partial}\psi_j|/|\partial\psi_j| = 0$.

The final two lemmas in this Part 1 assume that both vertices on γ have nonzero integer assignments with differing signs. Let m_1 and $-m_2$ denote the two integers with the convention here that $m_1 > 0$ and $m_2 > 0$. The two vertices on γ correspond to 4-tuples a_1 and a_2 in \hat{A} where $a_1 = (0, +, m_1 P)$ and $a_2 = (0, -, m_2 P)$. Here, as before, P is the relatively prime pair that defines θ_o via (1–8).

The following lemma assumes that $m_1 \neq m_2$. Let $\varepsilon = \text{sign}(m_1 - m_2)$ and $m = |m_1 - m_2|$. For this case, let \hat{A}' denote the asymptotic data set that is obtained from \hat{A} by removing both a_1 and a_2 , and replacing them with $a' = (0, \varepsilon, mP)$. As in the previous cases, T' denotes the graph that is identical to T but for the graph assigned to o ; and the latter is obtained from Γ_o by removing the interior of γ and replacing its two vertices by a single vertex whose integer assignment is $m_1 + m_2$.

Lemma 9.3 *With the circumstances as just indicated, the sequence converges to the element defined by (C_0, ϕ) in $\mathcal{M}^*_{\hat{A}', T'}$, in the following sense: There is an \mathbb{R} -invariant cylinder, S , at angle θ_o and a sequence $\{(C_{0j}, \phi_j)\}$ that defines the given sequence, and these are such that*

- $\lim_{j \rightarrow \infty} \int_{C_{0j}} \phi_j^* \varpi = \int_{C_0} \phi^* \varpi + \min(m_1, m_2) \int_S \varpi$ when ϖ has compact support
- The following limit exists and is zero:

$$\lim_{j \rightarrow \infty} \left[\sup_{z \in C_{0j}} \text{dist}(\phi_j(z), \phi(C_0) \cup S) + \sup_{x \in \phi(C_0) \cup S} \text{dist}(\phi_j(C_{0j}), x) \right].$$

Moreover, there exists an end, $E \subset C_0$, that gives the 4-tuple a' and has the following significance: Given $R \gg 1$ and any sufficiently large j , there is a smooth, proper embedding, ψ_j , from the complement in C_0 of the $|s| > R$ portion of E into C_{0j} ; and these are such that

- The $|s| \rightarrow \infty$ limit of the constant $|s|$ slices of E converge as a degree m multiple cover of the Reeb orbit in $S^1 \times S^2$ that defines the cylinder S .
- The complement of the image of ψ_j is a properly embedded, thrice punctured sphere that contains two ends of C_{0j} , these supplying the elements a_1 and a_2 to \hat{A} .
- The constant $|s|$ slices of the latter ends converge as $|s| \rightarrow \infty$ as multiple covers of a $\theta = \theta_o$ Reeb orbit, v_j ; and the resulting sequence, $\{v_j\}$, of Reeb orbits converges as $j \rightarrow \infty$ to the Reeb orbit that defines S .
- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi_j(z), \phi(z)) = 0$ for all z in the domain of ψ_j .

- $|\bar{\partial}\psi_j| \ll |\partial\psi_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi_j)} |\bar{\partial}\psi_j|/|\partial\psi_j| = 0$.

Lemma 9.4 that follows assumes that the integers m_1 and m_2 are equal. In this case, \hat{A}' is obtained from \hat{A} by removing both a_1 and a_2 . Nothing is added. The versions of T' in this case depend on certain properties of $\underline{\Gamma}_o$. There are three cases. In the first case, the graph $\underline{\Gamma}_o$ has only two vertices and both are bivalent. In this case, the vertex o has just two incident edges and T' is obtained from T by removing the vertex o and concatenating these two edges as one. In the second case, $\underline{\Gamma}_o$ has more than two vertices, but the two vertices on γ are bivalent. In this case, T' differs from T only in the graph assigned to o . Here, the T' version is obtained from $\underline{\Gamma}_o$ by removing the whole of γ and concatenating the resulting free ends as a single arc in the new graph. This arc is labeled by the same pair of edges that label γ . In the third case, $\underline{\Gamma}_o$ may or may not have two vertices, but at least one of the vertices on γ has valency 4 or more. In this case, the interior of γ is removed from $\underline{\Gamma}_o$, and its two vertices are replaced by a single vertex whose integer assignment is zero. In any of these cases, write the sum of the valencies of the two vertices on γ as $k + 4$.

Lemma 9.4 *With the circumstances as just indicated, the sequence converges to the element defined by (C_0, ϕ) in $\mathcal{M}^*_{\hat{A}', T'}$ in the following sense: There is an \mathbb{R} -invariant cylinder, S , at angle θ_o and a sequence $\{(C_{0j}, \phi_j)\}$ that defines the given sequence, and these are such that*

- $\lim_{j \rightarrow \infty} \int_{C_{0j}} \phi_j^* \varpi = \int_{C_0} \phi^* \varpi + m_1 \int_S \varpi$ when ϖ has compact support
- The following limit exists and is zero:

$$\lim_{j \rightarrow \infty} \left[\sup_{z \in C_{0j}} \text{dist}(\phi_j(z), \phi(C_0) \cup S) + \sup_{x \in \phi(C_0) \cup S} \text{dist}(\phi_j(C_{0j}), x) \right].$$

Moreover, there exists a point $z \in C_0$ with the following significance: Given small but positive ε and any sufficiently large j , there is a smooth, proper embedding, ψ_j , from the complement in C_0 of the radius ε^2 disk about z in C_0 into C_{0j} ; and these are such that:

- If $k = 0$, then $d\theta|_z \neq 0$. If $k > 0$, then z is a critical point of θ and $\deg(d\theta|_z) = k$. In either case, $\phi(z) \in S$.
- The complement of the image of ψ_j is a properly embedded, thrice punctured sphere that contains two ends of C_{0j} , these supplying the elements a_1 and a_2 to \hat{A} .

- The constant $|s|$ slices of the latter ends converge as $|s| \rightarrow \infty$ as multiple covers of a $\theta = \theta_o$ Reeb orbit, v_j ; and the resulting sequence, $\{v_j\}$, of Reeb orbits converges as $j \rightarrow \infty$ to the Reeb orbit that defines S . The third end of this thrice punctured sphere is mapped by ϕ_j to the radius ε ball in $\mathbb{R} \times (S^1 \times S^2)$ about $\phi(z)$.
- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi_j(z), \phi(z)) = 0$ for all z in the domain of ψ_j .
- $|\bar{\partial}\psi_j| \ll |\partial\psi_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi_j)} |\bar{\partial}\psi_j|/|\partial\psi_j| = 0$.

Part 2 The preceding four lemmas require that $\{\lambda_j\}$ provide a convergent sequence with non-zero limit to each arc but one in $\underline{\Gamma}_o$. As seen in the previous subsection, there can arise situations where a codimension 1 face of Δ_o can not be reached by such a sequence. This can happen when the constraints in (6–6) require some collection of arcs in $\underline{\Gamma}_o$ to have identical assignments for each $r \in \Delta_o$.

To say more about this phenomena, note that when $r \in \Delta_o$ and a vertex $v \in \underline{\Gamma}_o$ are given first, and then a number with sufficiently small absolute value is chosen, a new point in Δ_o is obtained from r by adding the chosen number to r 's value on each outward pointing arc at v while subtracting the number from r 's value on each inward pointing arc. This observation has the following implication: Let γ and γ' denote arcs in $\underline{\Gamma}_o$. Then $r(\gamma) - r(\gamma')$ is independent of $r \in \Delta_o$ if and only if γ and γ' share the same starting vertex and also the same ending vertex.

To see what $r(\gamma) - r(\gamma')$ can be in this case, remark that the loop defined by traversing γ in its oriented direction and γ' in reverse defines a primitive homology class from $H_1(\underline{\Gamma}_o; \mathbb{Z})$ and so can be written as a linear combination $\sum_e c_e [\ell_{oe}]$ where $\{c_e\}$ is a collection of integers that are defined modulo adding a constant to each $e \in E_+$ version of c_e and subtracting the same constant from each $e \in E_-$ version. This implies that

$$(9-17) \quad r(\gamma) - r(\gamma') = 2\pi \sum_e c_e \alpha_{Q_e}(\theta_o).$$

Since the collection $\{c_e\}$ are integers, this last identity has the following implication:

$$(9-18) \quad \text{The values of } r(\gamma) \text{ and } r(\gamma') \text{ are equal for all } r \in \Delta_o \text{ for at most two values of } \theta_o \text{ unless } \sum_e c_e Q_e = 0.$$

Indeed, with $Q \equiv \sum_e c_e Q_e \neq 0$, then $r(\gamma) - r(\gamma') = 2\pi \alpha_Q(\theta_o)$; and this is zero when $Q \neq 0$ if and only if either Q or $-Q$ defines the angle θ_o via (1–8).

The following lemma about the set of integers $\{c_e\}$ plays a role in what follows.

Lemma 9.5 *Let γ and γ' denote arcs that share the same starting vertex and share the same ending vertex. Then the homology class of the loop that is obtained by traversing γ in its oriented direction and then γ' in reverse can be written as $\sum_e c_e [\ell_{oe}]$ where $c_e = 1$ or 0 when $e \in E_-$ and where $c_e = -1$ or 0 if $e \in E_+$. Moreover, $c_e \neq 0$ for at least one $e \in E_-$ and for at least one $e \in E_+$.*

Proof of Lemma 9.5 Let \mathbb{Z} denote the free \mathbb{Z} -module generated by the arcs in $\underline{\Gamma}_o$, and let η denote an arc. Introduce the homomorphism $f^\eta: \mathbb{Z} \rightarrow \mathbb{Z}$ that sends η to 1 and all other arcs to 0. Now, let (e, \hat{e}) denote the edges that label γ with the convention that $e \in E_-$ and $\hat{e} \in E_+$. As $f^\gamma(\gamma - \gamma') = 1$, it follows that $c_e + c_{\hat{e}} = 1$ since only ℓ_{oe} and $\ell_{o\hat{e}}$ contain γ . As a consequence, (2-17) can be used to choose the coefficients $\{c_{(\cdot)}\}$ so as to make $c_e = 1$ and $c_{\hat{e}} = 0$. Now let (e', \hat{e}') denote the edges that label γ' with $e' \in E_-$ and $\hat{e}' \in E_+$.

As is explained next, $c_{e'}$ can not equal -1 . To see why such is the case, introduce v to denote the vertex where γ and γ' start. An incident half-arc to v is called a 1_- half-arc when $c_{(\cdot)} = -1$ on the E_- component of its labeling pair. Meanwhile, a 1_+ half-arc has $c_{(\cdot)} = 1$ on the E_+ component of its labeling pair. Note that there are an even number of 1_- half-arcs and an even number of 1_+ half-arcs. Also, every 1_- half-arc except γ' in the case that $c_{e'} = -1$ is a 1_+ half-arc. On the other hand, a 1_+ half-arc is also a 1_- half-arc because $f^{\gamma'}(\gamma - \gamma') = -1$. Granted these observations, then parity considerations forbid $c_{e'}$ from having value -1 .

A very similar argument now proves that $c_{e'} = 0$ and $c_{\hat{e}'} = 1$. To see how this comes about, say that an incident half-arc to v is a 0_- half-arc if $c_{(\cdot)} = 0$ on the E_- component of its labeling pair. Meanwhile, a 0_+ half-arc has $c_{(\cdot)} = 0$ on the E_+ component of its labeling pair. There are an even number of 0_- half-arcs and an even number of 0_+ half-arcs. Also, every 0_+ half-arc except γ is a 0_- half-arc. On the other hand, every 0_- half-arc is a 0_+ half-arc except γ' if $c_{e'} = 0$. Granted these observations, then parity considerations require $c_{e'} = 0$ and then $c_{\hat{e}'} = -1$ because $f^{\gamma'}(\gamma - \gamma') = -1$.

Having established that $c_e = 1$, $c_{\hat{e}} = 0$, $c_{e'} = 0$ and $c_{\hat{e}'} = -1$, define a sequence of sets, $\mathcal{A}_1 \subset \mathcal{A}_2 \subset \dots \subset \mathcal{A}^{\gamma, \gamma'}$ of arcs in $\underline{\Gamma}_o - \{\gamma, \gamma'\}$ as follows: The arcs that comprise \mathcal{A}_1 are the arcs from $\ell_{oe} \cup \ell_{o\hat{e}'} - \{\gamma, \gamma'\}$. Those in any $k \geq 2$ version of \mathcal{A}_k are labeled by an edge from e that also labels an arc from \mathcal{A}_{k-1} . If e labels an arc in $\mathcal{A}^{\gamma, \gamma'}$, then $c_e = 1$ when $e \in E_-$, and $c_e = -1$ when $e \in E_+$. If e does not label an arc from $\mathcal{A}^{\gamma, \gamma'}$, then $c_e = 0$. As a parenthetical remark, note that the sets $\mathcal{A}^{\gamma, \gamma'}$ and $\mathcal{A}^{\gamma', \gamma}$ are complementary: An arc from $\underline{\Gamma}_o - \{\gamma, \gamma'\}$ in the former is not in the latter and vice versa. \square

To lay more groundwork for the assertions to come, suppose that an ordered pair of distinct vertices in $\underline{\Gamma}_o$ have been specified and let \mathbb{A} denote a set of arcs that start at the first vertex of the pair, end at the second and are all assigned equal value by every map in Δ_o . Other arcs can connect these vertices if there is a map in Δ_o that assigns them values that differ from its value on the arcs in \mathbb{A} . It is assumed in what follows that \mathbb{A} has two or more arcs. Now, define a new graph, $\underline{\Gamma}'_o$, from $\underline{\Gamma}_o$ by identifying the arcs in \mathbb{A} to a single vertex. This single vertex, v' , is labeled by the sum of the integers that label the two vertices on the arcs in \mathbb{A} . To obtain the other labels for $\underline{\Gamma}'_o$, declare that the collapsing map from $\underline{\Gamma}_o$ to $\underline{\Gamma}'_o$ induce an isomorphism between $\underline{\Gamma}_o - \cup_{\gamma \in \mathbb{A}} \gamma$ and $\underline{\Gamma}'_o - v'$ that preserves all integer labels to vertices and the edge label pairs of the arcs.

Since \mathbb{A} has more than one arc, the graph $\underline{\Gamma}'_o$ has the wrong Euler number for a graph that arises from a pair (C_0, ϕ) with C_0 a multiply punctured sphere. Indeed, let n denote the number of incident edges to o and let a denote the number of arcs in \mathbb{A} . The Euler number of $\underline{\Gamma}'_o$ is $1 - n + (a - 1)$ where as $1 - n$ is the correct number were $\underline{\Gamma}'_o$ to come from a multiply punctured sphere. As a consequence, the image in $H_1(\underline{\Gamma}'_o; \mathbb{Z})$ of the classes $\{\ell_{oe}\}$ must satisfy $a > 1$ constraints, thus more than the one that is depicted in (2–17). Of course, these extra constraints are generated by a set whose elements are labeled by distinct, unordered pairs of arcs from \mathbb{A} . In this regard, the constraint labeled by such a pair, $\{\gamma, \gamma'\}$, asserts that

$$(9-19) \quad \sum_e c_e[\ell_{oe}] = 0 \text{ in } H_1(\underline{\Gamma}'_o; \mathbb{Z}),$$

where the collection $\{c_e\}$ comes from the $\{\gamma, \gamma'\}$ version of Lemma 9.5.

The next task is to make some sense out of this. To start, fix an arc $\gamma \in \mathbb{A}$ and let e denote the edge from E_- that labels γ . Define sets $\mathbb{A}_1^\gamma \subset \mathbb{A}_2^\gamma \subset \dots \subset \mathbb{A}^\gamma$ of arcs in $\underline{\Gamma}'_o$ as follows: The arcs from ℓ_{oe} comprise \mathbb{A}_1^γ . Meanwhile, those in the $k \geq 2$ version of \mathbb{A}_k^γ are labeled by an edge in E_+ that also labels an arc in \mathbb{A}_{k-1}^γ . Note that $\mathbb{A}^\gamma \cap \mathbb{A}^{\gamma'} = \emptyset$ if $\gamma \neq \gamma'$. To see this, first return to the proof of Lemma 9.5 to see that $\mathbb{A}^\gamma \subset \mathbb{A}^{\gamma, \gamma'}$. Since $\mathbb{A}^{\gamma, \gamma'}$ is disjoint from $\mathbb{A}^{\gamma', \gamma}$, so \mathbb{A}^γ and $\mathbb{A}^{\gamma'}$ are also disjoint.

Let E^γ denote the set of o 's incident edges that label arcs in \mathbb{A}^γ . Note that E^γ is disjoint from $E^{\gamma'}$ in the case that $\gamma \neq \gamma'$. Indeed, such is the case since the former set have $c_{(\cdot)} = \pm 1$ in the $\{\gamma, \gamma'\}$ version of Lemma 9.5, while the latter set have $c_{(\cdot)} = 0$ in this same version. Granted this, let E_\pm^γ denote $E_\pm \cap E^\gamma$. It follows from Lemma 9.5 that the union of the arcs in \mathbb{A}^γ define a subgraph, $\underline{\Gamma}_o^\gamma \subset \underline{\Gamma}'_o$, whose \mathbb{Z} homology is generated by the collection $\{\ell_{o\hat{e}} : \hat{e} \in E^\gamma\}$ subject to the one constraint,

$$(9-20) \quad \sum_{\hat{e} \in E_-^\gamma} [\ell_{o\hat{e}}] - \sum_{\hat{e} \in E_+^\gamma} [\ell_{o\hat{e}}] = 0.$$

The various $\gamma \in \mathbb{A}$ versions of (9–20) generate the constraint in (2–17) and the full collection of the various $\{\gamma, \gamma'\}$ versions of the constraint in (9–19).

Return now to the graphs in the collection $\{\underline{\Gamma}_o^\gamma\}$. Remark first that any two such graphs with distinct arc labels intersect only at v' . On the other hand, $\cup_{\gamma \in \mathbb{A}} \underline{\Gamma}_o^\gamma$ is the whole of $\underline{\Gamma}_o'$. These observations follow from Property 3 from Part 3 of Subsection 2.C.

Each arc in $\underline{\Gamma}_o^\gamma$ already has an edge pair label, and all but one vertex has an integer label. The one as yet unlabeled vertex is the one that maps to v' in $\underline{\Gamma}_o'$. What follows explains how to deal with the latter so that the result, $\underline{\Gamma}_o^\gamma$, comes from the label of a vertex in a graph as described in a version of Subsection 6.A. To start, define

$$(9-21) \quad Q^\gamma \equiv \sum_{\hat{e} \in E^\gamma} Q_e - \sum_{\hat{e} \in E_+^\gamma} Q_e.$$

It follows from (9–18) that θ_o is defined via (1–8) by some relatively prime integer pair if Q^γ is non-zero; and then Q^γ is a multiple of this pair. Let \hat{m}^γ denote the latter multiple when Q^γ is non-zero, and set $\hat{m}^\gamma = 0$ otherwise. Meanwhile, let m^γ denote the sum of the integers that label the vertices in $\underline{\Gamma}_o^\gamma - v'$.

Now there are two cases to consider. If $\hat{m}^\gamma - m^\gamma$ is non-zero or if the vertex in $\underline{\Gamma}_o^\gamma$ that maps to v' has more than two incident edges, then the vertex that maps to v' is viewed in $\underline{\Gamma}_o^\gamma$ as an honest vertex with integer label $\hat{m}^\gamma - m^\gamma$. If the vertex in $\underline{\Gamma}_o^\gamma$ that maps to v' is bivalent and if $\hat{m}^\gamma - m^\gamma = 0$, then this vertex is invisible in $\underline{\Gamma}_o^\gamma$ and its incident arc are viewed as a single arc.

The corresponding Subsection 6.A graph is denoted in what follows as T^γ , and it is characterized as follows: There is a multivalent vertex, o^γ , in T^γ whose angle is θ_o and graph label is $\underline{\Gamma}_o^\gamma$. Meanwhile, $T^\gamma - o^\gamma$ is isomorphic to the union of the components of $T - o$ that contain the interiors of the edges from the set E^γ . The corresponding asymptotic data set for T^γ is denoted below by \hat{A}^γ .

All of this background is assumed in the upcoming Lemma 9.6. In addition, this lemma assumes that a sequence, $\{\lambda_j\}_{j=1,2,\dots}$, has been specified in T 's version of the space depicted in (6–15) with the following properties: The sequence is constant but for the Δ_o factor. Let $\{r_j\}$ denote the corresponding sequence in the Δ_o factor. Then $\{r_j\}$ converges and \mathbb{A} is the set of arcs in $\underline{\Gamma}_o$ where $\lim_{j \rightarrow \infty} r_j(\cdot)$ is zero. With $\{\lambda_j\}$ so chosen, fix a sequence $\{s_j\} \in \mathbb{R}$ and then use the image of (s_j, λ_j) in $\mathbb{R} \times O_T / \text{Aut}(T)$ with the inverse of the map from Subsection 6.C to define a sequence in $\mathcal{M}_{\hat{A}, T}^*$.

Lemma 9.6 *With the circumstances as just indicated, there exists the following: First, a pair, (S^γ, ϕ^γ) , that defines an element in each $\gamma \in \mathbb{A}$ version of the space $\mathcal{M}_{\hat{A}^\gamma, T^\gamma}^*$.*

Second, a subsequence of $\{\lambda_j\}$, hence renumbered consecutively from 1, and a sequence $\{(C_{j0}, \phi_j)\}$ that defines the corresponding subsequence in $\mathcal{M}^*_{\hat{A}, T}$. Third, a finite set, Ξ , of pairs of the form (S, ϕ) where each $(S, \phi) \in \Xi$ is one of the following

- An \mathbb{R} -invariant cylinder at an angle in $\Lambda_{\hat{A}}$ or at the angle θ_o .
- An element from the set $\{(S^\gamma, \phi^\gamma)\}_{\gamma \in \mathbb{A}}$.

Here is the significance: Given a compact set $K \subset \mathbb{R} \times (S^1 \times S^2)$, then

- $\lim_{j \rightarrow \infty} \int_{C_{j0}} \phi_j^* \varpi = \sum_{(S, \phi) \in \Xi} \int_S \phi^* \varpi$ for each 2-form ϖ with compact support in K .
- The following limit exists and is zero:

$$\lim_{j \rightarrow \infty} \left(\sup_{z \in \phi_j^{-1}(K)} \text{dist}(\phi_j(z), \cup_{(S, \phi) \in \Xi} \phi(S)) + \sup_{z \in \cup_{(S, \phi) \in \Xi} \phi^{-1}(K)} \text{dist}(\phi_j(C_{j0}), \phi(z)) \right).$$

Moreover, there is a subset $\mathbb{A}_0 \subset \mathbb{A}$ and two versions of Ξ that arises from suitable choices of $\{s_j\}$ such that the first contains only and all of the $\gamma \in \mathbb{A}_0$ versions of (S^γ, ϕ^γ) , and the second contains only and all of the $\gamma \in \mathbb{A} - \mathbb{A}_0$ versions.

A story can be told in the manner of Lemmas 9.1–9.4 that gives a much more detailed account of the convergence of the sequence in question. In particular, a part of this story describes a subvariety in $\times_{\gamma \in \mathbb{A}} O_{T^\gamma} / \text{Aut}(T^\gamma)$ that maps in a proper, finite to one fashion onto the given codimension 1 stratum of the compactification of $O_T / \text{Aut}(T)$. In any event, this more detailed account is left to the reader except for the following comment: As with Lemmas 9.1 and 9.4, various cases must be distinguished. In the present circumstances, these are characterized by

- the integers that are assigned the two vertices on the arcs in \mathbb{A} ;
- whether or not v' is represented by a vertex in the various versions of Γ_o^γ ;
- in the cases that v' is so represented, the integer label of the corresponding vertex.

9.C The compactification of $\mathcal{M}_{\hat{A}}$ in the case that $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$

The purpose of this subsection is to describe a compactification of the whole of $\mathcal{M}_{\hat{A}}$ in the case that \hat{A} is an asymptotic data set with $N_- + \hat{N} + \zeta_- + \zeta_+ = 2$.

The proofs of the various assertions that follow are omitted as they can be obtained in a straightforward manner using variations of arguments from Section 4 and Section 7. In this regard, (9–16) should be used to obtain the conclusions of Proposition 7.1.

To start the discussion, return to Theorem 1.2 where $\mathcal{M}_{\hat{A}}$ is described as $\hat{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$ with $\hat{O}^{\hat{A}}$ sitting in the space $O^{\hat{A}}$ in (1–21) as the subset where the $\text{Aut}^{\hat{A}}$ action is free. The compactification here is as described in Part 2 of Subsection 1.B, thus $\underline{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$ where

$$(9-22) \quad \underline{O}^{\hat{A}} \equiv [\mathbb{R}_- \times \text{Maps}(\hat{A}_+; \mathbb{R})] / [(\mathbb{Z} \times \mathbb{Z}) \times \text{Maps}(\hat{A}_+; \mathbb{Z})].$$

As indicated in Part 2 of Subsection 1.B, the points in $\underline{O}^{\hat{A}} - \hat{O}^{\hat{A}}$ describe bonafide subvarieties, but these lie in $\hat{A}' \neq \hat{A}$ versions of $\mathcal{M}_{\hat{A}'}$. This story is told in the two parts of the subsection that follow. The first part discusses the points in $O^{\hat{A}} - \hat{O}^{\hat{A}}$ while the second considers those in $\underline{O}^{\hat{A}} - O^{\hat{A}}$.

Part 1 The following summarizes most of the story on $O^{\hat{A}} - \hat{O}^{\hat{A}}$:

Proposition 9.7 *The map described in Subsection 3.C extends to define an orbifold diffeomorphism between $\mathbb{R} \times O^{\hat{A}}/\text{Aut}^{\hat{A}}$ and $\mathcal{M}^*_{\hat{A}}$.*

This proposition is a corollary of Theorem 6.2. To elaborate some on how Proposition 9.7 arises, suppose that o denotes a bivalent vertex in $T^{\hat{A}}$ and let \hat{A}_o denote the set of elements in \hat{A}_+ whose integer pair gives θ_o via (1–8). Let Cyc_o denote the set of cyclic orderings of \hat{A}_o . As noted in Subsection 4.C, the components of $O^{\hat{A}}$ are in 1–1 correspondence with the points in $\times_o \text{Cyc}_o$. Let v denote a point in the latter space and let $O^{\hat{A}}_v$ denote the corresponding component. Borrowing again from Subsection 4.C, let $\text{Aut}_{o,v}$ denote the group of permutations of \hat{A}_o that preserve the cyclic ordering from v while permuting only elements with identical 4–tuples, and set $\text{Aut}_v \equiv \times_o \text{Aut}_{o,v}$. This is the subgroup of $\text{Aut}^{\hat{A}}$ that preserves $O^{\hat{A}}_v$. As such, the image of $O^{\hat{A}}_v$ in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ is diffeomorphic as an orbifold to $O^{\hat{A}}_v/\text{Aut}_v$. The component $O^{\hat{A}}_v/\text{Aut}_v$ corresponds in the present context to one of Section 5’s top dimensional strata in $\mathcal{M}^*_{\hat{A}}$.

To say something more about the points in $O^{\hat{A}} - \hat{O}^{\hat{A}}$, suppose that v is as above and that Aut_v has a nontrivial, canonical $\mathbb{Z}/k\mathbb{Z}$ subgroup. In this case, $O^{\hat{A}}_v$ will have points where the $\text{Aut}^{\hat{A}}$ action is not free; and, of course, these give points in $O^{\hat{A}}/\text{Aut}^{\hat{A}}$ that are not in $\hat{O}^{\hat{A}}/\text{Aut}^{\hat{A}}$. Now, let G denote a nontrivial subgroup of some canonical $\mathbb{Z}/k\mathbb{Z}$ subgroup of Aut_v . Define \hat{A}'_+ to be the quotient of \hat{A}_+ by G . Thus, \hat{A}'_+ is obtained from \hat{A}_+ by replacing each G orbit in \hat{A}_+ by a single point. Note that the given cyclic ordering v for \hat{A}_+ induces one for \hat{A}'_+ ; the latter is denoted in what follows by v' .

Meanwhile, both ζ_- and ζ_+ must vanish when $\text{Aut}^{\hat{A}}$ has a nontrivial, canonical $\mathbb{Z}/k\mathbb{Z}$ subgroup. Moreover, $\hat{A} - \hat{A}_+$ must consist of two 4-tuples that are of the form $(0, -, Q)$ or $(\pm 1, \cdot, Q)$ where k evenly divides both components of Q . This understood, augment \hat{A}'_+ with the 4-tuples that are obtained from the latter two by replacing Q with $\frac{1}{|G|}Q$. Here, $|G|$ denotes the size of the group G . Use \hat{A}' to denote this augmented set. Note that \hat{A}' defines a graph, $T^{\hat{A}'}$, of the sort described in (1–15) that satisfies the condition stated in (1–16). In this regard, the vertices of $T^{\hat{A}'}$ enjoy an angle preserving 1–1 correspondence with those of $T^{\hat{A}}$. This correspondence induces one between the respective edges that has the following property: If Q is an integer pair from an edge in $T^{\hat{A}}$, then $\frac{1}{|G|}Q$ is the integer pair of its partner in $T^{\hat{A}'}$.

Now, let $O^{\hat{A}}_{v,G}$ denote the subset of points in $O^{\hat{A}}_v$ whose Aut_v stabilizer is G . Then the quotient of $O^{\hat{A}}_{v,G}$ by Aut_v is diffeomorphic to $\hat{O}^{\hat{A}'}_{v'}/\text{Aut}_{v'}$. Here, $\text{Aut}_{v'}$ is the subgroup of $\text{Aut}^{\hat{A}'}$ that preserves the cyclic order of \hat{A}'_+ that is inherited from v 's ordering of \hat{A}_+ . This diffeomorphism is induced by a map from $O^{\hat{A}}_v$ to $\hat{O}^{\hat{A}'}_{v'}$ that intertwines the action of Aut_v with that of $\text{Aut}_{v'}$. To describe the desired map, fix a point in each of the G orbits in \hat{A}_+ . Now, let λ denote a point in $\mathbb{R}_- \times \text{Maps}(\hat{A}_+; \mathbb{R})$ whose image lies in $O^{\hat{A}}_{v,G}$. Define from λ a point $\lambda' \in \mathbb{R}_- \times \text{Maps}(\hat{A}'_+; \mathbb{R})$ by taking its \mathbb{R}_- factor to be $\frac{1}{|G|}$ times that of λ while defining its value on any element in \hat{A}'_+ to be the value that λ assigns to the chosen inverse image in \hat{A}_+ . This assignment of λ' to λ defines a smooth map from $O^{\hat{A}}_{v,G}$ to $\hat{O}^{\hat{A}'}_{v'}$ that intertwines the Aut_v action with the $\text{Aut}_{v'}$ action. The image of this map is $\hat{O}^{\hat{A}'}_{v'}$ and the induced map, $\pi: O^{\hat{A}}_{v,G}/\text{Aut}_v \rightarrow \hat{O}^{\hat{A}'}_{v'}/\text{Aut}_{v'}$, is the desired diffeomorphism

With the preceding understood, consider now the following scenario: Let $\lambda_0 \in O^{\hat{A}}_{v,G}$ and let $\{\lambda_j\} \in \hat{O}^{\hat{A}'}_{v'}$ denote a sequence that converges to λ_0 . Fix $s_0 \in \mathbb{R}$ and let (C, ϕ) denote the point in $\mathcal{M}_{\hat{A}}$ that is defined by $(s_0, \pi(\lambda_0))$. Because C has genus zero, there is a unique holomorphic covering of C with the following properties: First, the covering space is a punctured sphere and the group of deck transformations is $\mathbb{Z}/(|G|\mathbb{Z})$. Second, the covering is trivial over each concave side end of C where the $|s| \rightarrow \infty$ limit of θ is in $(0, \pi)$. Third, the covering restricts over the other two ends of C as a connected covering space. Use C_0 to denote the covering space and let $\phi_0: C_0 \rightarrow \mathbb{R} \times (S^1 \times S^2)$ denote the composition of the map ϕ with the holomorphic covering map to C . Thus, (C_0, ϕ_0) defines a point in $\mathcal{M}^*_{\hat{A}}$. Meanwhile, use $\{(C_{j0}, \phi_j)\}$ to denote the sequence in $\mathcal{M}_{\hat{A}}$ that corresponds via the map from Section 3 to the sequence $\{(s_0, \lambda_j)\}$. Theorem 6.2 implies that $\{(C_{j0}, \phi_j)\}$ converges in $\mathcal{M}^*_{\hat{A}}$ to (C_0, ϕ_0) .

Part 2 This part of the subsection discusses the geometric significance of the points in $\underline{O}^{\hat{A}} - O^{\hat{A}}$. The story here starts with the definition of an $\text{Aut}^{\hat{A}}$ invariant stratification of $\underline{O}^{\hat{A}}$. In this regard, a given stratum of the stratification is labeled by a partition, \wp , of \hat{A}_+ of the following sort: Each partition subset consists of elements whose integer pair components define the same angle via (1–8). The stratum $\mathcal{O}_\wp \subset \underline{O}^{\hat{A}}$ comes via (9–22) from the subset maps from \hat{A}_+ to \mathbb{R} that have the following two properties: First, the map assigns the same value in $\mathbb{R}/(2\pi\mathbb{Z})$ to elements from the same partition subset. Second, the map assigns distinct values in $\mathbb{R}/(2\pi\mathbb{Z})$ to elements from distinct partition subsets when their associated integer pair components define the same angle via (1–8). The set $O^{\hat{A}}$ consists of the components of the strata where \wp consists solely of single element sets. In general, the stratum labeled by \wp has dimension $n_\wp + 2$, where n_\wp denotes the number of sets that comprise the partition \wp .

Associated to \wp is an asymptotic data set, \hat{A}' ; this is defined as follows: The integers (ζ_-, ζ_+) for \hat{A} are used for \hat{A}' as are any 4-tuples of the form $(\pm 1, \dots)$ or $(0, -, \dots)$. Meanwhile, the 4-tuples from \hat{A}'_+ are in 1–1 correspondence with the partition subsets of \wp . In particular, each is of the form $(0, +, \dots)$. The integer pair component for the 4-tuple labeled by a given partition subset is the sum of those that label its elements. Thus, all such integer pairs define the same angle via (1–8).

The next order of business is to define a certain map from \mathcal{O}_\wp to $O^{\hat{A}'}$. For this purpose, let λ denote a point in $\mathbb{R}_- \times \text{Maps}(\hat{A}_+; \mathbb{R})$. Define $\lambda' \in \mathbb{R}_- \times \text{Maps}(\hat{A}'_+; \mathbb{R})$ by using λ 's factor in \mathbb{R}_- for the \mathbb{R}_- factor of λ' , and by using λ 's value on any element from any given partition subset for the value of λ' on the corresponding 4-tuple in $\hat{A}'_{\wp+}$. This map induces a diffeomorphism between \mathcal{O}_\wp and $O^{\hat{A}'}$ and thus defines a map, π , from $\mathcal{O}_\wp / \text{Aut}^{\hat{A}}$ to $O^{\hat{A}'}/\text{Aut}^{\hat{A}'}$. Note here that $\text{Aut}^{\hat{A}}$ is, in all cases, a subgroup of $\text{Aut}^{\hat{A}'}$.

Points in $O^{\hat{A}'}$ where $\text{Aut}^{\hat{A}'}$ acts freely parametrize $\mathcal{M}_{\hat{A}'}$. As was just asserted in Proposition 9.7, the whole of $\mathbb{R} \times O^{\hat{A}'}/\text{Aut}^{\hat{A}'}$ parametrizes $\mathcal{M}^*_{\hat{A}'}$.

With all of this as background, consider now the following scenario: Fix $\lambda_0 \in \mathcal{O}_\wp$ and $s_0 \in \mathbb{R}$; and let (C_0, ϕ) define the point in $\mathcal{M}^*_{\hat{A}'}$ that is defined from $(s_0, \pi(\lambda_0))$ as given by the \hat{A}' version of Proposition 9.7. Use the asymptotic data from the ends of C_0 to define a 1–1 correspondence between this set of ends and the 4-tuples in \hat{A}' .

In the mean time, let $\{\lambda_j\} \in O^{\hat{A}}$ denote a sequence that converges to λ_0 ; and let $\{(C_{j0}, \phi_j)\}$ denote the sequence in $\mathcal{M}^*_{\hat{A}}$ that is defined from $\{(s_0, \lambda_j)\}$ as described in Proposition 9.7.

Lemma 9.8 *The following limit exists and is zero:*

$$\lim_{j \rightarrow \infty} \left[\sup_{z \in C_{0j}} \text{dist}(\phi_j(z), \phi(C_0)) + \sup_{z \in C_0} \text{dist}(\phi_j(C_{0j}), \phi(z)) \right].$$

Moreover, let R be such that the $|s| > R$ portion of C_0 lies in the ends of C_0 . Define $C_R \subset C_0$ to be the complement in C_0 of the $|s| > R$ portion of those ends whose associated \hat{A}' label corresponds to a partition subset with size greater than 1. Then, for all j sufficiently large, there exists an embedding, $\psi_j: C_R \rightarrow C_{j0}$, with the following properties:

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi_j(z), \phi(z)) = 0$ for all z in the domain of ψ_j .
- $|\bar{\partial}\psi_j| \ll |\partial\psi_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi_j)} |\bar{\partial}\psi_j|/|\partial\psi_j| = 0$.
- The complement in C_{0j} of the image of ψ_j consists of a disjoint union of multiply punctured spheres, each embedded in a tubular neighborhood of an \mathbb{R} -invariant cylinder so that the projection to the cylinder defines an orientation preserving, ramified covering map with finitely many ramification points.

9.D The codimension 1 strata in the compactification of $\mathcal{M}^*_{\hat{A}}$

This subsection describes the codimension 1 strata in the compactification of $\mathcal{M}^*_{\hat{A}}$ in the case that $N_- + \hat{N} + \zeta_- + \zeta_+ > 2$. To start, remark that there are two sorts of codimension 1 strata in these cases. Strata of the first sort arise from the codimension 1 strata in the compactifications of those $O_T/\text{Aut}(T)$ that appear in the various versions of (9–1). The geometric interpretation of these strata are described in Lemmas 9.2–9.4. In this regard, keep in mind that only circular versions of $\underline{\Gamma}_o$ are involved in the cases when the added strata are of codimension 1.

The second sort of stratum arises by adding certain dimension $k - 1$ faces to the simplex product Δ^k in (9–1). This section and the final section explore the geometry of the subvarieties in the codimension 0 strata that are near this second kind of added strata.

To start the discussion, let \mathcal{S} denote the codimension 0 stratum involved. The space Δ^k for \mathcal{S} is a k -dimensional product of simplices in $\times_k((0, \pi) - \Lambda_{+, \emptyset})$. The compactification of \mathcal{S} contains an added codimension 1 stratum when the following occurs: The homotopy type of graph for \mathcal{S} has an edge, e , with a trivalent vertex whose angle values on \mathcal{S} limit to the angle defined via (1–8) by Q_e , or to that defined by $-Q_e$. Denote the latter angle by θ_* ; it is the angle that defines the relevant codimension 1 face in the closure of Δ^k . The relevant trivalent vertex is denoted in what follows by o .

The upcoming Lemma 9.9 describes some of the basic features of sequence in \mathcal{S} whose image in Δ^k approach the face just described. To set the stage for the lemma, let T denote a graph whose homotopy type arises from the elements in \mathcal{S} . The other two of o 's incident edges are denoted by e' and \hat{e} with the one distinguished from the other in the following manner: The second vertices on e , e' and \hat{e} have angles that bracket θ_o in the interval $(0, \pi)$; and under the circumstances, the second angle on e is on the same side of θ_o as the second angle on one of the other two edges. The convention takes this other edge to be e' .

Let T_e denote the closure of the component of $T - o$ that contains the interior of e , and let T' denote the component of $T - \text{int}(e)$ that contains o . Thus, T_e and T' are closed subgraphs in T that intersect only at o . Both can be viewed as graphs that are described by versions of Subsection 6.A. To elaborate, the labeling of T_e is as follows: This graph has a monovalent vertex, o_e , at angle θ_* such that $T_e - o_e$ is isomorphic to the component of $T - o$ that contains the interior of e . Meanwhile, T' has a bivalent vertex, o' , at angle θ_* with the property that $T' - o'$ is isomorphic to the other two components of $T - o$. Let \hat{A}' denote the asymptotic data set that is defined by T' and let \hat{A}_e denote the corresponding set that comes from T_e .

Because the graphs that arise from the elements in \mathcal{S} are all homotopic to a single graph, the choices that are offered in Parts 1 and 2 of Subsection 6.C can be made once so as to hold for all fibers of the projection map in (9-1) to Δ^k . This done, consider a sequence in the space depicted in (9-1) whose factors are constant save for the coordinate in o 's factor of Δ^k . The latter angle for the j 'th element of the sequence is denoted θ_j , and the corresponding sequence of angles converges with limit θ_* . The inverse of the map in Subsection 6.C takes such a sequence and provides a sequence in \mathcal{S} with no limit in $\mathcal{M}^*_{\hat{A}}$. Let $\{(C_{0j}, \phi_j)\}_{j=1,2,\dots}$ denote a sequence of pairs that defines this sequence in \mathcal{S} .

With the preceding as background, consider:

Lemma 9.9 *Under the circumstances as just described, there exists the following:*

- Pairs (C_{0j}, ϕ_j) and (C_{ej}, ϕ_e) that define respective points in the (\hat{A}', T') and (\hat{A}_e, T_e) versions of $\mathcal{M}^*_{(\cdot, \cdot)}$.
- Sequences $\{r'_j\}_{j=1,2,\dots}$ and $\{r_{ej}\}_{j=1,2,\dots}$ of real numbers, the former increasing without limit and the latter decreasing without limit.
- Sequences $\{s'_j\}_{j=1,2,\dots}$ and $\{s_{ej}\}_{j=1,2,\dots}$ of real numbers, such that the former converges and the latter is decreasing and unbounded, or else the former is decreasing and unbounded while the latter converges.

- For all sufficiently large j , a proper embedding, ψ'_j , into C_{0j} from the complement in $C_{0'}$ of the $s > r'_j$ part of end in $C_{0'}$ that corresponds to the vertex $o' \in T'$.
- For all sufficiently large j , a proper embedding, ψ_{ej} , into C_{0j} from the complement in C_{e0} of the $s < -r_{ej}$ part of the end in C_{e0} that corresponds to the vertex $o_e \in T_e$.

Here is the significance: First, the image via ψ'_j of the $s = r'_j$ boundary of its domain coincides with the image via ψ_{ej} of the $s = r_{ej}$ boundary of the latter's domain. Otherwise, the images of the two maps are disjoint. Second,

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi'_j(z), \phi'_j(z)) = 0$ for any fixed $z \in C_{0'}$. Here, ϕ'_j is obtained from ϕ' by composing with the translation by s'_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$.
- $|\bar{\partial}\psi'_j| \ll |\partial\psi'_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi'_j)} |\bar{\partial}\psi'_j|/|\partial\psi'_j| = 0$.

Finally,

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi_{ej}(z), \phi_{ej}(z)) = 0$ for any fixed $z \in C_{e0}$. Here, ϕ_{ej} is obtained from ϕ_e by composing with the translation by s_{ej} along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$.
- $|\bar{\partial}\psi_{ej}| \ll |\partial\psi_{ej}|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi_{ej})} |\bar{\partial}\psi_{ej}|/|\partial\psi_{ej}| = 0$.

More can be said about the sorts of pairs $(C_{0'}, \phi')$ and (C_{0e}, ϕ_e) that appear here. For example, if the vertex $o_e \in T_e$ is fixed by $\text{Aut}(T_e)$, then the subset of such pairs in

$$(9-23) \quad \mathcal{M}^*_{\hat{A}', T'} \times \mathcal{M}^*_{\hat{A}_e, T_e}$$

is a smooth, codimension 2 suborbifold that can be defined as a fiber product using the maps that are depicted in [15, (2.19) and (2.20)]. The task of filling in the details is left to the reader.

The story told by Lemma 9.9 has analogues in higher codimensions. This can occur when some version of $\underline{\Gamma}_{(\cdot)}$ has an arc that begins and ends on the same vertex. By way of elaboration, suppose that o is a multivalent vertex in T and that $\gamma \subset \underline{\Gamma}_o$ is such an arc. As noted previously, the maps in Δ_o are identical on γ ; they send γ to $\alpha_Q(\theta_o)$ where $Q = \sum_e c_e Q_e$ with the constants $\{c_e\}$ defined by writing γ 's homology class as $\sum_e c_e [\ell_{oe}]$. Because the maps in Δ_o must send γ to a positive number, situations arise where the angle θ_o varies on graphs homotopic to T so that $\alpha_Q(\theta_o)$ limits to zero. Lemma 9.9 concerns the case that $\underline{\Gamma}_o$ is a figure 8 graph with a single 4-valent vertex. However, the event just described can occur in a more general context.

But for one comment, the general context is left for the reader to describe. Here is the comment: The techniques used in Steps 2 and 3 from the proof of Proposition 8.2 in Subsection 8.B can be used to prove that the event in question occurs only in the case that $\mathcal{M}^*_{\hat{A},T}$ lies in the closure of a codimension zero stratum of the sort that was just discussed. In particular, the angle where α_Q is zero defines a codimension 1 face of the space Δ^k that appears in the corresponding version of (9-1); this a face that provides one of the added strata in the closure of $\mathcal{M}^*_{\hat{A}}$.

9.E Faces of Δ^k where $\theta = 0$ or π

There are versions of $\mathcal{M}^*_{\hat{A}}$ that have top dimensional strata whose elements define a graph with an edge that has a trivalent vertex on one end and a $\theta = 0$ or $\theta = \pi$ vertex on the other. There can be strata in this case where the trivalent vertex angle limits to 0 or π as the case may be. This subsection discusses the behavior of the subvarieties where the graph has a very small trivalent vertex angle. The discussion for the case where the angle is close to π is not given since it is essentially identical to what follows.

Let \mathcal{S} denote the top dimensional stratum involved. In this case, one of the faces in Δ^k is characterized by a $\theta = 0$ assignment to a trivalent vertex. To start the story, note that as before, the choices in Parts 1 and 2 of Subsection 6.C can be made once so as to hold for all graphs that arise from elements in \mathcal{S} . This understood, let o denote the multivalent vertex involved and let e denote its edge. The assumption here is that there are no other vertex angles between θ_o and 0.

Fix a graph from an element in \mathcal{S} to call T , and then construct a sequence of homotopic graphs, $\{T_j\}_{j=1,2,\dots}$, where T_j varies from T only in the angle assigned to o . The latter angle is denoted now as θ_j and the assumption in what follows is that $\lim_{j \rightarrow \infty} \theta_j = 0$. The choices from Parts 1 and 2 of Subsection 6.C identify all $T = T_j$ versions of (6-15). This understood, fix a point, λ , in (6-15) and let (C_{0j}, ϕ_j) denote a pair that gives rise to the image of λ in the $T = T_j$ version of $\mathcal{M}^*_{\hat{A},T}$ via the inverse of the map from Subsection 6.C.

There are now two cases to consider. In the first, two of o 's incident edges connect o to respective vertices with larger angle. In the second case, e and another edge connect o to angle 0 vertices. The two parts to this subsection discuss these two cases.

Part 1 Let e' and e'' denote the two incident edges that connect o to vertices with larger angles. It must be assumed here that neither the $Q = Q_{e'}$ nor $Q = Q_{e''}$ version of α_Q has a zero between 0 and θ_o . Since $\alpha_Q(\theta_o) > 0$, this occurs when Q does not

define an angle via (1–8) or, if it does, then so does $-Q$, the former is greater than the latter and both are greater than θ_o .

The limiting behavior of $\{(C_{0j}, \phi_j)\}$ is described below in terms of a pair of graphs, T' and T'' , that are obtained from T in the following manner: Let $T_{e'}$ and $T_{e''}$ denote the closures in T of the respective components of $T - o$ that contain the interiors of e' and e'' . The graph T' is characterized as follows: It has an edge, \hat{e}' , with a $\theta = 0$ monovalent vertex such that the closure of $T' - \hat{e}'$ in T' is isomorphic to the closure of $T_{e'} - e'$ in T . In addition, $Q_{\hat{e}'} = Q_{e'}$. There is a similar characterization of T'' : It has an edge, \hat{e}'' , with a monovalent vertex such that the closure of $T'' - \hat{e}''$ in T'' is isomorphic to that of $T_{e''} - e''$ in T . Moreover, $Q_{\hat{e}''} = Q_{e''}$. Let \hat{A}' and \hat{A}'' denote the respective asymptotic data sets for the graphs T' and T'' . Note that these respective sets are determined from T and T' via the rules in (6–3).

Lemma 9.10 *With the scenario just described, there exist exist the following:*

- Pairs (C_0', ϕ') and (C_0'', ϕ'') that define respective elements in $\mathcal{M}^*_{\hat{A}', T'}$ and $\mathcal{M}^*_{\hat{A}'', T''}$.
- Sequences $\{s'_j\}_{j=1,2,\dots}$ and $\{s''_j\}_{j=1,2,\dots}$ of real numbers.
- For all sufficiently large j , a proper embedding, ψ'_j , into C_{0j} of the complement in C_0' of the $\theta < 2\theta_j$ part of the component of the C_0' version of $C_0 - \Gamma$ that corresponds to the edge $\hat{e}' \subset T'$.
- For all sufficiently large j , a proper embedding, ψ''_j , into C_{0j} of the complement in C_0'' of the $\theta < 2\theta_j$ part of the component of the C_0'' version of $C_0 - \Gamma$ of C_{e_0} that corresponds to the edge $\hat{e}'' \subset T''$.

Here is the significance: First, the images of ψ'_j and ψ''_j are disjoint and their complement in C_{0j} is the $\theta < 2\theta_j$ part of the closure of the components of the C_{0j} version of $C_0 - \Gamma$ that correspond to the edges e , e' and e'' . Second,

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi'_j(z), \phi'_j(z)) = 0$ for any fixed $z \in C_0'$. Here, ϕ'_j is obtained from ϕ' by composing with the translation by s'_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$.
- $|\bar{\partial}\psi'_j| \ll |\partial\psi'_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi'_j)} |\bar{\partial}\psi'_j|/|\partial\psi'_j| = 0$.

Finally,

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi''_j(z), \phi''_j(z)) = 0$ for any fixed $z \in C_0''$. Here, ϕ''_j is obtained from ϕ'' by composing with the translation by s''_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$

- $|\bar{\partial}\psi''_j| \ll |\partial\psi''_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi''_j)} |\bar{\partial}\psi''_j|/|\partial\psi''_j| = 0$.

As in the case with Lemma 9.9, the set of pairs in $\mathcal{M}^*_{\hat{A}', T'} \times \mathcal{M}^*_{\hat{A}'', T''}$ that arise can be described rather explicitly.

Part 2 In this case, the limit of $\{(C_{0j}, \phi_j)\}$ defines a point in $\mathcal{M}^*_{\hat{A}}$ only in the case that both integer pairs that connect o to $\theta = 0$ vertices have the form $(0, \cdot)$. The resulting stratum in $\mathcal{M}^*_{\hat{A}}$ has codimension 2 rather than codimension 1. Without this integer pair condition, the limit is described with the help of a point in some $\hat{A}' \neq \hat{A}$ moduli space $\mathcal{M}^*_{\hat{A}'}$.

In any case, to say more about the limit, introduce the graph T' that is characterized as follows: Let \hat{e} denote the edge in T that connects the vertex o to a larger angled vertex. The graph T' has an edge, \hat{e}' , with a monovalent vertex at angle 0 such that the closure in T' of $T' - \hat{e}'$ is isomorphic to the closure in T of the component of $T - o$ that contains the interior of $T - \hat{e}$. In addition, $Q_{\hat{e}'} = Q_{\hat{e}}$. Let \hat{A}' denote the asymptotic data set that is determined by T' .

Lemma 9.11 *With the scenario just described, there exist exist the following:*

- A pair (C_0', ϕ') that defines an element in $\mathcal{M}^*_{\hat{A}', T'}$.
- A sequence $\{s_j\}_{j=1,2,\dots}$ of real numbers.
- For all sufficiently large j , a proper embedding, ψ'_j , into C_{0j} of the complement in C_0' of the $\theta < 2\theta_j$ part of the component of the C_0' version of $C_0 - \Gamma$ that corresponds to the edge $\hat{e}' \subset T'$.

Here is the significance: First, the complement of the image of ψ'_j in C_{0j} is the $\theta < 2\theta_j$ part of the closure of the components of the C_{0j} version of $C_0 - \Gamma$ that correspond to the edges e, e' and e'' . Second,

- $\lim_{j \rightarrow \infty} \text{dist}(\phi_j \circ \psi'_j(z), \phi'_j(z)) = 0$ for any fixed $z \in C_0'$. Here, ϕ'_j is obtained from ϕ' by composing with the translation by s_j along the \mathbb{R} factor of $\mathbb{R} \times (S^1 \times S^2)$.
- $|\bar{\partial}\psi'_j| \ll |\partial\psi'_j|$ and $\lim_{j \rightarrow \infty} \sup_{\text{domain}(\psi'_j)} |\bar{\partial}\psi'_j|/|\partial\psi'_j| = 0$.

In the case that both Q_e and $Q_{e'}$ have first component zero, the sequence $\{s_j\}$ can be taken to be constant and the domain of ψ'_j can be taken to be the whole of C_0' . In this case, the range of ψ'_j is the whole of C_{0j} . In particular, the sequence in $\mathcal{M}^*_{\hat{A}}$ defined by $\{(C_{0j}, \phi_j)\}$ converges in $\mathcal{M}^*_{\hat{A}}$ to the point defined by (C_0', ϕ') .

Techniques of [15, Section 3] can be used to prove that every element in $\mathcal{M}^*_{\hat{A}', T'}$ arises as a limit in the sense just described of a sequence in the stratum \mathcal{S} .

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*, *GAF* 2000 (Tel Aviv, 1999), *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 Hofer**, *Dynamics, topology, and holomorphic curves*, from: “Proceedings of the International Congress of Mathematicians, Vol. I (Berlin, 1998)”, Extra Vol. I (1998) 255–280 MR1648034
- [6] **H Hofer**, *Holomorphic curves and dynamics in dimension three*, from: “Symplectic geometry and topology (Park City, UT, 1997)”, *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
- [9] **H Hofer, K Wysocki, E Zehnder**, *Properties of pseudoholomorphic curves in symplectizations. III. Fredholm theory*, from: “Topics in nonlinear analysis. The Herbert Amman Anniversary”, *Progr. Nonlinear Differential Equations Appl.* 35, Birkhäuser, Basel (1999) 381–475 MR1725579
- [10] **K Honda**, *An openness theorem for harmonic 2-forms on 4-manifolds*, *Illinois J. Math.* 44 (2000) 479–495 MR1772422
- [11] **D McDuff**, *The local behaviour of holomorphic curves in almost complex 4-manifolds*, *J. Differential Geom.* 34 (1991) 143–164 MR1114456
- [12] **C H Taubes**, *The geometry of the Seiberg-Witten invariants*, from: “Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998)”, Extra Vol. II (1998) 493–504 MR1648099
- [13] **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
- [14] **C H Taubes**, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$* , *Geom. Topol.* 6 (2002) 657–814 MR1943381

- [15] **CH Taubes**, *Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: properties and existence*, *Geom. Topol.* 10 (2006) 785–928 MR2240906

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

`chtaubes@math.harvard.edu`

Proposed: Rob Kirby

Received: 5 April 2005

Seconded: Jim Bryan, Ron Stern

Accepted: 25 October 2006