AN OBSTRUCTION BUNDLE RELATING GROMOV-WITTEN INVARIANTS OF CURVES AND KÄHLER SURFACES

By JUNHO LEE and THOMAS H. PARKER

Abstract. In a previous paper the authors defined symplectic "Local Gromov-Witten invariants" associated to spin curves and showed that the GW invariants of a Kähler surface X with $p_g > 0$ are a sum of such local GW invariants. This paper describes how the local GW invariants arise from an obstruction bundle (in the sense of Taubes) over the space of stable maps into curves. Together with the results of our earlier paper, this reduces the calculation of the GW invariants of elliptic and general-type complex surfaces to computations in the GW theory of curves with additional classes: the Euler classes of the (real) obstruction bundles.

On a compact Kähler surface X, a holomorphic 2-form α is a section of the canonical bundle whose zero locus is a canonical divisor D. Several years ago, the first author observed (see [Lee1]) that each such 2-form α naturally induces an almost complex structure J_α that satisfies a remarkable property:

Image Localization Property. *If a J_α-holomorphic map represents a nontrivial $(1,1)$ class, then its image lies in D.*

After further perturbing to a generic J near J_α, the images of all J-holomorphic maps cluster in ε-neighborhoods of the components of the canonical divisor D. This implies that the Gromov-Witten invariant of X is a sum

$$GW_{g,n}(X,A) = \sum_k GW^{\text{loc}}_{g,n}\big(D_k, d_k[D_k]\big)$$

over the connected components D_k of D of "local GW invariants" that count the contribution of maps whose image lies near D_k.

When D is smooth, the local invariants depend only on the normal bundle N to $D \subset X$. By the adjunction formula, N is a holomorphic square root of the canonical bundle K_D, that is, N is a theta-characteristic of the curve D and the pair (D,N) is a spin curve. The total space N_D of N has a tautological holomorphic 2-form α whose zero locus is the zero section $D \subset N_D$. For perturbations of the corresponding J_α, all J-holomorphic maps cluster around the zero section. These clusters define local GW invariants of the spin curve (D,N) which, the authors

Manuscript received September 20, 2009; revised July 26, 2011.
Research of the second author supported in part by the National Science Foundation.
American Journal of Mathematics 134 (2012), 453–506.

proved in [LP], depend only on the parity of N (i.e. on $h^0(N)$ mod 2). Altogether, we have

$$GW_{g,n}(X,A) = \sum_k \left(i_k\right)_* GW^{\mathrm{loc}}_{g,n}\left(N_{D_k}, d_k\right) \tag{0.1}$$

where $(i_k)_*$ is the induced map from the inclusion $D_k \subset X$. Thus the computation of the GW invariants of Kähler surfaces with $p_g > 0$ is reduced to the problem of calculating the local invariants of spin curves.

Recently, Kiem and Li [KL] defined the local invariants by algebraic geometry methods and proved the formulas for degree 1 and 2 local invariants conjectured by Maulik and Pandharipande [MP]. The first author [Lee2] reproved those formulas by adapting the symplectic sum formula of [IP2] to local GW invariants.

Because (0.1) applies to all GW invariants, not just those of the "embedded genus", one cannot apply Seiberg-Witten theory. Nor can the local invariants be computed by the usual methods of algebraic geometry, such as localization and Grothendieck-Riemann-Roch, because *the linearized J_α-holomorphic map equation is not complex linear*. In particular, when $\mathrm{genus}(D_k) > 0$ the local invariants in (0.1) are not the same as the "local GW invariants" used to study Calabi-Yau 3-folds [BP] or the "twisted GW invariants" defined by Givental.

While the local invariants are defined in terms of the GW theory of the (complex) *surface* N_D one would like, as a step toward computation, to recast them in terms of the much better-understood GW theory of *curves* (cf. [OP]). This paper uses geometric analysis arguments to prove that the local GW invariants of a spin curve (D,N) arise from a cycle in the space of stable maps into the curve D. The cycle is defined by constructing an "obstruction bundle". While the basic idea is clear and intuitive, the construction is difficult because of technical issues involving the construction of a complete space of maps.

The intuition goes like this. The tautological 2-form α on N_D determines an almost complex structure J_α. By the Image Localization Property, the space of stable J_α-holomorphic maps into N_D representing $d[D]$ is the same as the space of degree d stable maps into D:

$$\overline{\mathcal{M}}^{J_\alpha}_{g,n}\left(N_D, d[D]\right) = \overline{\mathcal{M}}_{g,n}(D,d). \tag{0.2}$$

Counting dimensions, one sees that the formal dimension of $\overline{\mathcal{M}}_{g,n}(D,d)$ is exactly twice the dimension of the virtual fundamental class that defines the local GW invariants of N_D (when $n=0$). The dimensions do not match because J_α is not generic. Perturbing J_α to a generic J effectively reduces the space of maps to a half-dimensional cycle in $\overline{\mathcal{M}}_{g,n}(D,d)$ that defines the local GW invariants of the spin curve (D,N). To understand this reduction, we use another remarkable property of the J_α-holomorphic maps:

Injectivity Property. *The linearization of the map $f \mapsto \overline{\partial}_{J_\alpha} f$, when restricted to the normal bundle, is an elliptic operator L_f whose kernel vanishes for every J_α-holomorphic map ($\alpha \neq 0$).*

The injectivity property implies that the vector spaces $\operatorname{coker} L_f$ have constant dimension. The construction of Section 8 shows that these cokernels form a locally trivial real vector bundle

$$\begin{array}{c} \mathcal{O}b \\ \big\downarrow \\ \overline{\mathcal{M}}_{g,n}(D,d) \end{array} \tag{0.3}$$

over the space (0.2) whose rank is the formal GW dimension for surface. This is an "obstruction' bundle" in the sense of Taubes: the Implicit Function Theorem shows that the space of perturbed holomorphic maps is diffeomorphic to the subset of $\overline{\mathcal{M}}_{g,n}(D,d)$ given by the zero set of a certain section of $\mathcal{O}b$. As always, we have the associated map

$$\widehat{ev} = st \times ev : \overline{\mathcal{M}}_{g,n}(D,d) \longrightarrow \overline{\mathcal{M}}_{g,n} \times D^n \tag{0.4}$$

whose first factor is the stabilization map and whose second factor records the images of the marked points. In this context, the GW invariants of the curve D are defined as the image of the virtual fundamental class $[\overline{\mathcal{M}}_{g,n}(D,d)]^{\mathrm{vir}}$ under $\widehat{ev}_*$. Our main result is that the local GW invariants are defined from these by capping with the Euler class of this obstruction bundle:

MAIN THEOREM. *There is an oriented real bundle $\mathcal{O}b$ over $\overline{\mathcal{M}}_{g,n}(D,d)$ whose isomorphism class depends only on the parity of the spin curve (D,N). When* $\operatorname{genus}(D) > 0$

$$GW^{loc}_{g,n}\big(N_D,d\big) = \widehat{ev}_*\big(\big[\overline{\mathcal{M}}_{g,n}(D,d)\big]^{vir} \cap e(\mathcal{O}b)\big), \tag{0.5}$$

where the virtual fundamental class $[\overline{\mathcal{M}}_{g,n}(D,d)]^{vir}$ defines the GW invariants of the curve D.

This formula describes how the local invariants evaluate on elements of the cohomology of $\overline{\mathcal{M}}_{g,n} \times D^n$. But in the context of this theorem, all descendant classes are pullbacks of classes in $H^*(\overline{\mathcal{M}}_{g,n} \times D^n)$ (see Section 1). Thus the equality (0.5) applies to descendant classes.

This theorem is a step toward computing the GW invariants of minimal Kähler surfaces with $p_g > 0$. For non-minimal surfaces, one would also need a version of the Main Theorem for the local GW invariants of an exceptional curve: when N is the bundle $\mathcal{O}(-1)$ over $D = \mathbb{P}^1$. This case is fundamentally different: it is simpler because the Image Localization and Injectivity Properties hold for the standard complex structure, but is more complicated because a lemma essential for the analysis—Lemma 1.2 below—does not apply when $D=\mathbb{P}^1$. As a result, the right-hand side of (0.5) is well-defined and the associated GW invariants are computable [FP], but our analysis does not show the equality in (0.5). This $D=\mathbb{P}^1$ case will be analyzed elsewhere.

The main theorem would not be difficult to prove if $\mathcal{O}b \to \overline{\mathcal{M}}_{g,n}(D,d)$ were a smooth vector bundle over a manifold, or if one could perturb to the smooth situation while retaining the Image Localization and Injectivity properties. Unfortunately, there do not currently exist theorems or techniques for smoothing the compactified moduli space. As a result, the proof of the main theorem must address two significant technical issues: (i) the lack of information about the structure of the moduli space near its boundary, and (ii) the local triviality of the obstruction bundle.

To deal with issue (i), we first perturb the J_α-holomorphic map equation in the manner of Ruan-Tian [RT2]. The resulting moduli space $\overline{\mathcal{M}}=\overline{\mathcal{M}}_{g,n}(D,d)$ of (J,ν)-holomorphic maps into D consists of a smooth top stratum $\mathcal{M}$ and boundary strata of codimension at least two. We fix a small neighborhood U of the boundary and consider its complement $\mathcal{M}_U=\mathcal{M}\setminus U$. Then, after smoothing and perturbing U if necessary, $\mathcal{M}_U$ is a compact smooth oriented manifold with boundary, so defines a *relative* homology class in $\mathcal{M}ap$. In Section 3 we prove that this "relative virtual fundamental class" defines the same Gromov-Witten invariants as other standard definitions, including the one used in algebraic geometry.

On the compact manifold $\mathcal{M}_U$, the (J,ν)-holomorphic map equation defines a section of the obstruction bundle, and the Euler class $e(\mathcal{O}b)$ is Poincaré dual to the zero set of any section that is transverse to zero. We achieve transversality by adding a second Ruan-Tian perturbation term. In Section 9 we prove a generalized Image Localization Theorem that shows that this second perturbation leaves $\mathcal{M}$ unchanged. The main theorem is proved in Section 10 by showing that the zero set of a section defines a rational homology class that is equal to the local GW invariants by cobordism, and on the other hand is equal to the right-hand side of (0.5) by Poincaré duality for $\mathcal{M}_U$.

The analysis aspects of the proof are aimed at issue (ii) above: proving that the obstruction bundle is locally trivial. The key difficulty is that the normal component of the linearization of the J_α-holomorphic map equation at a map f is an operator of the form $L_f=\overline{\partial}+A(df)$, and df is not pointwise bounded in the topology of the Gromov compactness theorem. Thus it is not clear whether L_f is continuous in f on the space of smooth maps—a fact we need in order to show that $\mathcal{O}b$, which is essentially $\operatorname{coker} L_f$, is locally trivial (in the literature, continuity is often implicitly assumed). For this purpose, we introduce a stronger topology on the space of maps in Section 2 and prove a strengthened version of the Gromov compactness theorem. Then in Section 4 we extend the operators L_f off the set of J_α-holomorphic maps as a family of "modified linearizations"; these are equally natural, but are more easily estimated, than the actual linearizations. Sections 5–7 develop the needed analysis results to show that L_f is continuous in f as an operator on appropriate weighted Sobolev spaces. The key to this analysis is the Poincaré inequalities proved in Section 6, which lead to estimates for the weighted norms that are uniform even as the domain degenerates to a nodal curve. These results are used in Section 8 to define the obstruction bundle and prove that it is locally trivial.

For computations, one would like to express the obstruction bundle $\mathcal{O}b$ in terms of algebraic geometry. At the Kähler structure J_0 on N_D, L_f is the $\overline{\partial}$-operator on the bundle f^*N, the fiber $\mathcal{O}b_f$ is $H^{0,1}(f^*N)$, and the Injectivity Property suggests that $h^0(f^*N) = 0$. This would imply that $\mathcal{O}b$ is the index bundle $\operatorname{ind}\overline{\partial}$. However, when D has genus $h > 1$ the Injectivity Property does not hold for J_0 and, as in Brill-Noether theory, $h^{0,1}(f^*N)$ can jump up at special maps. (Pandharipande and Maulik showed us a specific example where such a jump necessarily occurs in the moduli space, and a similar example appears in [KL].) Thus $\mathcal{O}b$ is not in general equal to the $\operatorname{ind}\overline{\partial}$. We clarify this in Section 11 by showing how the linearization $f \mapsto L_f$ defines a map from $\overline{\mathcal{M}}$ to a space $\mathcal{F}$ of *real* Fredholm operators. There are natural classes $\kappa_i \in H^*(\mathcal{F})$, first defined by Koschorke, that give obstructions to the bundle $\operatorname{ind}_{\mathbb{R}} L$ being an actual vector bundle rather than a virtual bundle. We prove that all Koschorke classes vanish for the family of maps defined by the moduli space (0.2), so $\mathcal{O}b = \operatorname{ind}_{\mathbb{R}} L$ is an actual bundle over the moduli space. This gives a homotopy-theoretic characterization of the obstruction bundle.

Zinger [Z] independently obtained (0.5) following the approach of Fukaya-Ono [FO] and Li-Tian [LT2]. He also describes some interesting generalizations of our main theorem.

Acknowledgment. We thank the referee for many helpful comments.

1. J_α-holomorphic maps with stabilized domains. We begin with a review of the setup for J_α-holomorphic maps; for details see [Lee1, LP]. Fix a Kähler surface X with complex structure J and geometric genus $p_g > 0$. Then the real vector space

$$\mathcal{H} = \operatorname{Re}\left(H^{2,0} \oplus H^{0,2}\right)$$

has dimension $2p_g > 0$. Using the Kähler metric compatible with J, each $\alpha \in \mathcal{H}$ defines an endomorphism K_α of TX by the equation

$$\langle u, K_\alpha v\rangle = \alpha(u,v). \tag{1.1}$$

Each K_α is skew-adjoint, anti-commutes with J, and satisfies $K_\alpha^2 = -|\alpha|^2 Id$. It follows that JK_α is skew-adjoint and $Id + JK_\alpha$ is invertible. Thus there is a family of almost complex structures

$$J_\alpha = \left(Id + JK_\alpha\right)^{-1} J \left(Id + JK_\alpha\right) \tag{1.2}$$

on X parameterized by $\alpha \in \mathcal{H}$. Note that while α is holomorphic, the corresponding almost complex structure J_α need not be integrable, and indeed, usually isn't integrable.

For each $\alpha \in \mathcal{H}$, we can consider the set of maps $f : C \to X$ from a connected complex curve with complex structure j into X that satisfy the J_α-holomorphic map equation

$$\overline{\partial}_{J_\alpha} f = 0 \tag{1.3}$$

where $\overline{\partial}_{J_\alpha} f = \frac{1}{2}(df + J_\alpha df j)$. It was shown in [Lee1] that these are exactly the maps that satisfy the "perturbed J-holomorphic map equation"

$$\overline{\partial}_J f - \nu_\alpha = 0 \quad \text{where } \nu_\alpha = K_\alpha\big(\partial_J f\big) j \tag{1.4}$$

(even though $\overline{\partial}_{J_\alpha} f$ is not equal to $\overline{\partial}_J f - \nu_\alpha$ for arbitrary maps). In this paper, we fix J and work with the J_α-holomorphic map equation in the form (1.4) rather than (1.3).

Because any holomorphic map into a Kähler surface represents a $(1,1)$ class, we can restrict attention to maps representing $(1,1)$ classes: the GW invariants vanish for all other classes. In this context, the first author observed that the following remarkable fact.

LEMMA 1.1. (Image Localization) *If $f : C \to X$ is a J_α-holomorphic map that represents a non-trivial $(1,1)$ class, then the image of f lies in one connected component D_k of the zero set of α.*

This fact leads to the general formula (0.1) expressing the GW invariant as a sum of local invariants associated with the components of the divisor of α. When such a divisor D is smooth with multiplicity one, the square of its holomorphic normal bundle N is the canonical bundle K_D, so (N,D) is a spin curve. In this case, it was shown in [LP] that the local GW invariants depend only on the genus and parity of the spin curve (D,N) (the parity is $h^0(N)$ mod 2).

Now consider a genus h spin curve (D,N), that is, a curve D with genus h and a holomorphic line bundle N on D with $N^2 = K_D$. The total space N_D of N has a complex structure J that makes the projection $\pi : N_D \to D$ holomorphic. We then have an exact sequence

$$0 \longrightarrow \pi^* N \longrightarrow TN_D \longrightarrow \pi^* TD \longrightarrow 0 \tag{1.5}$$

and hence the canonical bundle of N_D is $\pi^*(N^* \otimes K_D) = \pi^* N$. The tautological section of $\pi^* N$ is thus a holomorphic 2-form α on N that vanishes transversally along the zero section $D \subset N_D$. This 2-form α induces an almost complex structure J_α on N_D as in (1.2). The local GW invariants of D are the invariants associated with the space

$$\overline{\mathcal{M}}^{J_\alpha}_{g,n}(N_D, d) \tag{1.6}$$

of all stable J_α-holomorphic maps $f : C \to N_D$ whose domain has genus g and n marked points and whose image represents $d[D] \in H_2(N_D)$. By Proposition 1.1 this is the same as the space $\overline{\mathcal{M}}_{g,n}(D, d)$ of stable J-holomorphic maps into the zero section D with degree d. However, the linearizations of these equations differ in a way that will be crucial in later sections.

LEMMA 1.2. *If D has genus $h \geq 1$ and $2g+n \geq 3$, then the domain of every map in $\overline{\mathcal{M}}_{g,n}(D, d)$ is a stable curve.*

Proof. When D has genus $h \geq 1$ all rational components of the domain are mapped to points, so are stable curves (by the definition of stable map). Similarly, all components with genus one are mapped to points or have a node (so are stable), unless the domain is smooth and the map is etale, in which case the domain has at least one marked point since $2g+n \geq 3$. □

In particular, because all domains are stable, the relative cotangent bundles over $\overline{\mathcal{M}}_{g,n}(D, d)$ are pull-backs of the relative cotangent bundles over $\overline{\mathcal{M}}_{g,n}$ by the stabilization map. The descendent classes are thus pull-backs of cohomology classes via the map (0.4).

One usually takes $\overline{\mathcal{M}}_{g,n}$ to be a Deligne-Mumford space. However, it is more convenient to take it to be the moduli space described by Abramovich, Corti and Vistoli [ACV], building on the key work of Looijenga [Lo]. This is a finite branched cover of the compactified Deligne-Mumford space that is a fine moduli space and a smooth projective variety ([ACV, Theorem 7.6.4]); it solves the moduli problem for families of "G-twisted curves" that carry a principle G-bundle for a certain finite group G together with certain additional structure at their nodes and marked points (see [ACV] for details). As described in [IP1, Section 1], this modification has no substantial effect: one recovers the standard GW invariants by dividing by the degree of the cover. Accordingly, we will use the standard notation $\overline{\mathcal{M}}_{g,n}$ for compactified Deligne-Mumford space and leave the presence of twisted structures implicit. In this context, the space $\overline{\mathcal{M}}_{g,n}$ of G-twisted curves and the total space of its universal curve

$$\begin{array}{c} \overline{\mathcal{U}}_{g,n} \\ \big\downarrow{\scriptstyle \pi} \\ \overline{\mathcal{M}}_{g,n} \end{array} \tag{1.7}$$

are manifolds with Riemannian metrics.

Throughout, we will work with the moduli spaces of solutions of perturbations of (1.4), namely, solutions of the perturbed J_α-holomorphic map equation

$$\overline{\partial}_J f - K_\alpha\big(\partial_J f\big) j = \nu_f \tag{1.8}$$

for various perturbation terms ν. We will use perturbations of the type introduced by Ruan and Tian in [RT2], which can be described as follows. Fix an almost complex manifold (X, J) and let $\overline{\mathcal{U}}$ be the universal curve (1.7). Because G-twisted curves have no non-trivial automorphisms, the domain of a map $f : C \to X$ is uniquely identified (as a G-twisted curve) with a fiber of $\overline{\mathcal{U}}$, and the graph of f is a map $F : C \to \overline{\mathcal{U}} \times X$. Consider the bundle $T^*\overline{\mathcal{U}} \boxtimes TX = \mathrm{Hom}(\pi_1^* T\overline{\mathcal{U}}, \pi_2^* TX)$ over $\overline{\mathcal{U}} \times X$. A *Ruan-Tian perturbation* is an element ν of the space

$$\mathcal{P} = \Omega_J^{0,1}\big(T^*\overline{\mathcal{U}} \boxtimes TX\big) \tag{1.9}$$

of $(0,1)$ sections, that is, sections ν satisfying $\nu \circ j_u = -J \circ \nu$ where j_u is the complex structure on $\overline{\mathcal{U}}$. *Here, and everywhere below,* $(0,1)$ *components are determined by J, not* J_α. Restricting such a ν to the graph of f gives a form $\nu_f \in \Omega^{(0,1)}(f^*TX)$, defined by

$$v_f(x)(u) = v\big(x, f(x)\big)(u) \quad \forall x \in C,\ u \in T_x C, \tag{1.10}$$

that can be used in (1.8). Note that ν_f depends on f, but not on df.

We will routinely use the phrase "for generic ν" to mean "for ν in a Baire set in the space $\mathcal{P}$".

2. Convergence of maps with stable domains. The proof of the main theorem requires a careful definition of the space of maps as a topological space. The appropriate topology is *not* the one defined by Gromov compactness. Instead, we will use the stronger "λ-topology" defined by carefully-chosen weighted Sobolev norms. In this section, we set out the definitions and then prove that Gromov compactness holds in the λ-topology.

Let $\pi : \overline{\mathcal{U}} \to \overline{\mathcal{M}}_{g,n}$ be the universal G-twisted curve (1.7). There is a $\delta_0 > 0$ such that in each fiber $C_z = \pi^{-1}(z)$ the nodes are separated from each other and from the marked points by a distance of at least $4\delta_0$. We will scale the metric so that $\delta_0 = 1$ and choose a continuous function

$$\rho : \overline{\mathcal{U}} \longrightarrow \mathbb{R} \tag{2.1}$$

that is equal to the distance to the nodal variety on the set $\{\rho < 1\}$ and satisfies $1 \le \rho \le 2$ on the complementary set. For each $\delta < 1$ we will also write

$$C_z(\delta) = C_z \cap \{\rho \ge \delta\} \quad \text{and} \quad B_z(\delta) = C_z \cap \{\rho < \delta\}. \tag{2.2}$$

For small δ, $B_z(\delta)$ is a union of components, each being either a union of two disks with their center points identified (a node) or a thin annular neck (a near node).

Because the universal G-twisted curve (1.7) is a fine moduli space, the fibers have no non-trivial automorphisms and:

(i) Around each smooth fiber C_z there is a neighborhood V_z of z and smooth local trivialization

$$\phi : C_z \times V_z \longrightarrow U_z. \tag{2.3}$$

(ii) Around each fiber C_z with $B_z(\delta) \neq \emptyset$ there is a similar local trivialization

$$\phi : C_z(\delta) \times V_z \longrightarrow U_z \tag{2.4}$$

where $U_z = \pi^{-1}(V_z) \cap \{\rho > \delta\}$.

Now fix an isometric embedding of X into $\mathbb{R}^N$ for some N. The space of maps is defined using the following weighted Sobolev norms.

Definition 2.1. Fix a constant λ with $0 < \lambda < \frac{1}{6}$. For each $p \geq 2$ and each curve C, let $\mathcal{M}ap_\lambda(C,X)$ be the completion of the set of smooth maps $f : C \to X$ in the weighted norm

$$|||f|||_{1,p} = \left(\int_C \rho^{p-2-\lambda}\,|df|^p\right)^{\frac{1}{p}} + \left(\int_C \rho^{-\lambda}\,|df|^2\right)^{\frac{1}{2}} + \left(\int_C |f|^p\right)^{\frac{1}{p}} \tag{2.5}$$

where $|f|$ is defined by the embedding $X \subset \mathbb{R}^N$, and both $|df|$ and the measure on C are defined using the Riemannian metrics on C and X induced from the metrics on $\overline{\mathcal{U}}$ and $\mathbb{R}^N$.

If we replace C by $C_z(\delta)$ in the Definition 2.1 then the resulting norms are uniformly equivalent, for each δ, to the usual $L^{1,p}$ norm on the fibers of the local trivializations (2.3) and (2.4). After fixing p, we also define the *p-energy* $E_p(f)$ and *total energy* $E(f)$ of a map $f : C \to X$ to be

$$E(f) = E_p(f) + E_2(f) \quad \text{where } E_p(f) = \left(\int_C \rho^{p-2-\lambda}\,|df|^p\right)^{\frac{2}{p}}. \tag{2.6}$$

We will consider only the subset of $\mathcal{M}ap_\lambda(X)$ whose total energy is below a fixed level.

Definition 2.2. **(λ-topology)** For each number E and $p > 2$, set

$$\mathcal{M}ap^E_{g,n}(X) = \left\{(z,f) \mid z \in \overline{\mathcal{M}}_{g,n},\ f \in \mathcal{M}ap_\lambda\big(C_z, X\big) \text{ and } E(f) < E\right\}. \tag{2.7}$$

Give this space the λ-topology: a sequence (C_n, f_n) converges to (C,f) if (a) $C_n \to C$ in $\overline{\mathcal{M}}_{g,n}$, (b) $E_p(f_n) \to E_p(f)$ and $E_2(f_n) \to E_2(f)$, and (c) $f_n \to f$ in the norm (2.5) on $C_z(\delta)$ for every $\delta > 0$.

The convergence in (c) is defined using the trivialization (2.4). For simplicity, we will often denote the space (2.7) by $\mathcal{M}ap_\lambda(X)$.

When working with the norm (2.5) on curves C_z near a nodal fiber C_0 of the universal curve, it is helpful to use the definitions (2.2) to decompose C_z into the "thick" region $C_z(\delta)$ away from the nodes, and the "thin" regions $B_{z,i}(\delta) \subset C_z$ that lie within distance δ from one of the nodes n_i of C_0. The thin regions are of two types—nodes and necks—described as follows.

Near a node n_i of C_0, choose holomorphic coordinates $\{x_i\}$ on the universal curve such that

$$B_{z,i}(\delta) = \Big\{\big(x_1,\dots,x_n\big) \mid x_1x_2 = \mu,\ x_j = c_j\ \forall j>2, \text{ and } |x_1|^2 + |x_2|^2 < \delta\Big\} \tag{2.8}$$

for some constants $\mu, c_j \in \mathbb{C}$. For notational simplicity, we will often omit the subscript z in $C_z(\delta)$ and $B_{z,i}(\delta)$.

- **(Nodes)** When $\mu = 0$, $B_i(\delta)$ is diffeomorphic to two disks $D^1(\delta)$ and $D^2(\delta)$ with their center points identified. The punctured disk $D^1(\delta)\setminus\{0\}$ can be parameterized by the map

$$\phi : (t,\theta) \mapsto= \big(e^{-t-i\theta}, 0, c_3, \dots\big) \tag{2.9}$$

from the half-infinite cylinder $T^1 = (L,\infty)\times S^1$ where $L = -\ln\delta$. Composing ϕ with $x_1 \leftrightarrow x_2$ gives a parameterization of $D^2(\delta)\setminus\{0\}$. Together, $B_i(\delta)\setminus\{0\}$ is parameterized by a disjoint union $T(\delta) = T^1 \cup T^2$ of two cylinders.

- **(Necks)** When $\mu \neq 0$, $B_i(\delta)$ can be parameterized by the map

$$\phi : (t,\theta) \longmapsto \Big(x_1, \frac{\mu}{x_1}, c_3, \dots\Big) \quad \text{where } x_1 = \sqrt{|\mu|}\, e^{t+i\theta} \tag{2.10}$$

from the cylinder $T(\delta) = (-L, L)\times S^1$ for the unique positive L that satisfies $2|\mu|\cosh(2L) = \delta^2$. Note that $L \to \infty$ as the neck pinches.

To obtain uniform estimates, it is useful to express the norm (2.5) in the parameterizations (2.9) and (2.10). In the coordinate chart used in (2.8), the metric on $\overline{\mathcal{U}}$ is uniformly equivalent to the Euclidean metric g_0 in these coordinates, and the distance function is uniformly equivalent to the Euclidean distance. Furthermore, under the maps (2.9) and (2.10) g_0 is conformally related to the cylindrical metric $\hat{g} = dt^2 + d\theta^2$ by

$$\phi^* g_0 = \rho^2 \hat{g} \quad \text{where } \rho^2 = \begin{cases} e^{-2t} & \text{near a node} \\ 2|\mu|\cosh(2t) & \text{in a neck.} \end{cases} \tag{2.11}$$

Consequently, after identifying f and ρ with their ϕ-pullbacks we have, for both nodes and necks and for each $p \geq 2$,

$$\int_{B_i(\delta)} \rho^{p-2-\lambda}\, |df|^p \, dvol_g \cong \int_{T(\delta)} \rho^{-\lambda}\, |df|^p \, dvol_{\hat{g}} \tag{2.12}$$

where $\cong$ means that the ratio between the two sides is bounded above and below by positive constants that depend only on p and λ. It is useful to note that, because

ρ is essentially exponential in t, the integrals of its powers in the cylindrical metric satisfy

$$\int_{\rho\leq\delta} \rho^\gamma \, dt\, d\theta \leq c_\gamma\, \delta^\gamma \quad \text{for } \gamma > 0. \tag{2.13}$$

The following two lemmas give properties of maps in $\mathcal{M}ap^E_{g,n}(X)$. The first shows that the images of necks and nodal neighborhoods are uniformly small, and the second shows that when two maps are close in the λ-topology then the Hausdorff distance $\text{dist}_{\mathcal{H}}$ between their images is small.

LEMMA 2.3. *Fix g,n,E and $p>2$. Then there are constants $\delta_0>0$ and c such that for every $(C,f)\in\mathcal{M}ap^E_{g,n}(X)$ and every $\delta<\delta_0$, each component $B_i(\delta)$ of $B(\delta)\subset C$ satisfies*

$$\text{diam}\left(f\left(B_i(\delta)\right)\right) \leq c\,\delta^{\frac{\lambda}{p}}.$$

Proof. By the compactness of $\overline{\mathcal{U}}$, there is a $\delta_0>0$ so that whenever $\delta<\delta_0$ there is a nodal fiber C_0 of $\overline{\mathcal{U}}$ and holomorphic coordinates centered on a node $n_i\in C_0$ so that $B_i(\delta)$ is given by (2.8). Identifying $f: B_i(\delta)\to X$ with its ϕ-pullback $f: T\to X\subset\mathbb{R}^N$ as in (2.9) or (2.10), we have

$$\text{diam}\, f\left(B_i(\delta)\right) = \text{diam}\, f(T) = \underset{T}{\text{osc}}\, f.$$

In the node case, the Sobolev inequality $\text{osc}\, f\leq c_1\|df\|_p$ holds on each segment $T_n=[L+n,L+n+1]\times S^1$ of T, and (2.11) implies that $\rho\leq c_2\delta e^{-n}$ on T_n. Summing from $n=1$ to ∞ yields

$$\begin{aligned}\text{diam}\, f\left(B_i(\delta)\right) &\leq c_3\sum_n\left(\int_{T_n}\left(\delta e^{-n}\rho^{-1}\right)^\lambda |df|^p\right)^{\frac{1}{p}}\\ &\leq c_3\,\delta^{\frac{\lambda}{p}}\sum_n e^{-\frac{n\lambda}{p}}\left(\int_T \rho^{-\lambda}|df|^p\right)^{\frac{1}{p}} \leq c_4\sqrt{E}\,\delta^{\frac{\lambda}{p}}\end{aligned}$$

using (2.6) and (2.12). The same argument applies in the neck case, noting that (2.11) again implies that $\rho\leq c_5\delta e^{-n}$ on both $T_n=[L-n-1,L-n]\times S^1$ and on $T'_n=[-L+n,-L+n+1]\times S^1$ for $0\leq n<L$. $\square$

LEMMA 2.4. *Fix g,n,E and $p>2$. Given $(C_0,f)\in\mathcal{M}ap^E_{g,n}(X)$, $\varepsilon>0$ and a sufficiently small $\delta>0$, there are neighborhoods $\mathcal{N}_f(\varepsilon)$ and $\mathcal{N}_f(\varepsilon,\delta)$ of (C_0,f) in the λ-topology such that*

$$\text{dist}_{\mathcal{H}}\left(f\left(C_0\right),g(C)\right)\leq\varepsilon \quad \textit{for all } (C,g)\in\mathcal{N}_f(\varepsilon)$$

and

$$\|f-g\|_{\infty;C_0(\delta)}\leq\varepsilon \quad \textit{for all } (C,g)\in\mathcal{N}_f(\varepsilon,\delta).$$

Proof. By Lemma 2.3 we can choose $\delta_1 > 0$ so that $\mathrm{diam}(f(B_i(\delta_1))) \le \varepsilon/3$ for all $f \in \mathcal{M}ap^E_{g,n}(X)$. Fix δ with $2\delta \le \delta_1$. We can then choose a neighborhood $\mathcal{N}_f$ of (C_0, f) small enough so that the domains of all maps $g : C \to X$ in $\mathcal{N}_f$ lie in the uniform local trivialization (2.4) around C_0 for this δ. Decompose the domain C_0 of f into $C_0(\delta)$ and the union of neighborhoods $B_i(\delta)$ of its nodes. Then for each $(C, g) \in \mathcal{N}_f$ we have $\delta \le \rho \le 2$ on $C(\delta)$, so the norm (2.5) is uniformly equivalent to the (unweighted) $L^{1,p}$ norm on $C(\delta)$. Furthermore, the fibers of the local trivialization (2.4) have uniform geometry, so there is a uniform constant (depending on δ) for the usual Sobolev embedding $C^0 \subset L^{1,p}$. Consequently,

$$\|f - g\|_{\infty;C_0(\delta)} \le c(\delta)\ |||f - g|||_{1,p}.$$

We can then choose $\mathcal{N}_f(\varepsilon, \delta)$ to make the right-hand side less that $\varepsilon/3$ for all $g \in \mathcal{N}_f(\varepsilon, \delta)$. In particular, setting $\mathcal{N}_f(\varepsilon) = \mathcal{N}_f(\varepsilon, \delta_1/2)$, each $g \in \mathcal{N}_f(\varepsilon)$ satisfies

$$\begin{aligned} \mathrm{dist}_{\mathcal{H}}\left(f\left(C_0\right), g(C)\right) &\le \|f - g\|_{\infty;C_0(\delta)} + \sup_i \mathrm{diam}\, f\left(B_i(\delta)\right) + \sup_i \mathrm{diam}\, g\left(B_i(\delta)\right) \\ &\le \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3}. \end{aligned}$$

□

We next prove an enhanced version of the Gromov Compactness Theorem. It assumes that all maps have stable domains (cf. Lemma 1.2), and proves convergence in the λ-topology, which is stronger than the convergence in the standard formulations of Gromov Compactness (cf. [FO, IS2, RT2]).

THEOREM 2.5. (Compactness) *Suppose that all maps in $\overline{\mathcal{M}}^{J,\nu}_{g,n}(X, A)$ have stable domains. Then there is an $E = E(p, g, n, A)$ such that $\overline{\mathcal{M}}^{J,\nu}_{g,n}(X, A)$ is a compact subset of $\mathcal{M}ap^E(X)$ whenever* $\sup|\nu|$ *is small.*

Proof. Given a sequence of maps $f_k : C_k \to X$ in $\overline{\mathcal{M}}^{(J,\nu)}_{g,n}(X, A)$ we can consider their graphs $F_k : C_k \to \overline{\mathcal{U}} \times X$ as described before equation (1.9). Then $\{F_k\}$ is a sequence of J_ν-holomorphic maps with uniformly bounded energy. Moreover, as explained in Section 1, each C_k is a stable curve with no non-trivial automorphisms. The first factor of each F_k is therefore a diffeomorphism onto a fiber of the universal curve. The Gromov Compactness Theorem [IS2], applied to $\{F_k\}$ implies that there is a subsequence such that (i) the domains C_k converge in $\overline{\mathcal{U}}$ to a limit C_0, (ii) the maps f_k converge to a limit f_0 in Hausdorff distance and in $L^{1,p}$ on compact sets in the complement of the nodes of C_0, and (iii) the energy densities $|dF_k|\, dvol$ converge as measures.

Next fix $\delta \le 1$ small enough so that the maps (2.10) are defined on the 2δ-neighborhoods of the nodes. Because the energy densities converge as measures we can, given any $\varepsilon_0 > 0$, also assume that, for each node n_i of C_0, the energy in each neck $B_{k,i}(\delta)$ of each C_k is at most ε_0. But energy is conformally invariant so, as in (2.12) with $p=2$, the pullback maps $f_k : T(\delta) \to X$ satisfy the (J,ν)-holomorphic map equation and

$$\int_{T(\delta)} |df_k|^2 = \int_{B_{k,i}(\delta)} |df_k|^2 < \varepsilon_0. \tag{2.14}$$

Elliptic theory then gives a pointwise bound on $|df_k|$: by Lemma 5.1 of [IP2] there are constants c_1 and ε_0 so that whenever (2.14) holds we have $|df_k| < c_1\rho^{1/3}$ on $T(\delta)$; both c_1 and the maximum number of nodes depend only on g, n, A and $\sup|\nu|$. Integrating as in (2.12) and (2.13) shows that for any $p \geq 2$ and $\lambda < \frac{1}{6}$

$$\begin{aligned} \int_{B_{k,i}(\delta)} \rho^{p-2-\lambda} |df_k|^p &\leq c_2 \int_{T(\delta)} \rho^{-\lambda} |df_k|^p \leq c_2 c_1^p \int_{T(\delta)} \rho^{-\lambda+\frac{p}{3}} \\ &\leq c_3\, \delta^{\frac{p}{3}-\lambda} \leq c_3 \sqrt{\delta}. \end{aligned} \tag{2.15}$$

This, together with the convergence on each $C(\delta)$, implies that $E_p(f_k) \to E_p(f_0)$, that $E_2(f_k) \to E_2(f_0)$, and that there is a bound $E(f) \leq E$ for all (J,ν)-holomorphic maps f. Convergence in the λ-norm also follows: given $\varepsilon > 0$, choose δ small enough that the bound (2.15) is less than $(\varepsilon/4)^p$ and less than $(\varepsilon/4)^2$ when $p=2$. By the compactness of $X \subset \mathbb{R}^N$ the integral of $|f|^p$ over $B(\delta)$ is bounded by $c_4\,\mathrm{Area}(B(\delta)) \leq c_5\delta^2$; we can also assume this is less than $(\varepsilon/4)^p$. Finally, for this δ, the restriction of the norm (2.5) to $C(\delta)$ is uniformly equivalent to the unweighted $L^{1,p}$ norm, so we can choose K large enough so that, for all $k \geq K$, we have $|||f_k - f_0|||_{1,p;C(\delta)} < \varepsilon/4$, and hence $|||f_k - f_0|||_{1,p} < \varepsilon$. ☐

3. The virtual fundamental class. Gromov-Witten invariants of a closed symplectic manifold X have been defined in several different ways. When X is a smooth complex curve (the only case needed for this paper), all definitions apply and yield the same GW invariants. Each definition involves compact moduli spaces $\overline{\mathcal{M}}_{g,n}(X,A)$ of (formal) dimension $r = c_1(X)(A) + (\dim X - 3)(1-g) + n$ and the associated map

$$\widehat{ev} = st \times ev : \overline{\mathcal{M}}_{g,n}(X,A) \longrightarrow \overline{\mathcal{M}}_{g,n} \times X^n \tag{3.1}$$

defined as in (0.4). For our purposes, three different descriptions are relevant. One is algebraic and the other two are analytic:

(1) When X is projective, the space of stable maps $\overline{\mathcal{M}}^{\text{stable}}_{g,n}(X,A)$ is a projective variety and algebraic geometers (see [LT1, BF]) define the virtual fundamental class $[\overline{\mathcal{M}}_{g,n}(X,A)]^{virt}$ as an element of its Chow cohomology:

$$\left[\overline{\mathcal{M}}_{g,n}(X,A)\right]^{virt} \in A_r\big(\overline{\mathcal{M}}^{\text{stable}}_{g,n}(X,A)\big). \tag{3.2}$$

(2) When X is symplectic and semipositive, Ruan-Tian [RT2] showed that for generic (J,ν) the space $\overline{\mathcal{M}}^{(J,\nu)}_{g,n}(X,A)$ of all (J,ν)-holomorphic maps is the

union of (i) a $2r$-dimensional orbifold $\mathcal{M}$ consisting of maps from smooth domains and (ii) a stratified "boundary" whose image under the map (3.1) lies in a set of dimension at most $2r-2$. Consequently, the image of the moduli space represents a rational homology class

$$GW_{g,n}(X,A) \in H_{2r}\big(\overline{\mathcal{M}}_{g,n} \times X^n; \mathbb{Q}\big). \tag{3.3}$$

This Ruan-Tian GW class is independent of the generic (J,ν) and is a symplectic invariant of X. ("Semipositive" is a technical condition that is true whenever $\dim X \le 6$.)

(3) When X is symplectic, the construction of Li-Tian [LT2] defines a virtual fundamental class

$$\big[\overline{\mathcal{M}}_{g,n}(X,A)\big]^{\mathrm{vir}} \in H_{2r}\big(\mathcal{M}ap_{g,n}(X,A); \mathbb{Q}\big) \tag{3.4}$$

in the homology of the infinite-dimensional space of maps $\mathcal{M}ap_{g,n}(X,A)$. Because (3.1) extends to $\mathcal{M}ap_{g,n}(X,A)$ one can then evaluate the class (3.4) on classes in $H^*(\overline{\mathcal{M}}_{g,n} \times X^n)$.

Variations on construction 3 have been done by Fukaya-Ono [FO], Ruan [R] and Siebert [S].

When X is projective, Li and Tian proved in [LT3] that the virtual class (3.4) is the homology class underlying the Chow class (3.2) under the inclusion of the space of stable maps into $\mathcal{M}ap_{g,n}(X,A)$. When X is semipositive, one can show that the pushforward of the virtual class (3.4) by $\widehat{ev}_*$ is the Ruan-Tian class (3.8); we explicitly prove this for compact curves in Remark 10.2. Thus when X is a smooth compact curve D, all three definitions apply and, after pushing forwards by $\widehat{ev}_*$, define the same element in the homology of $\overline{\mathcal{M}}_{g,n} \times X^n$.

Later, to prove the Main Theorem stated in the introduction, we will work with the GW invariants defined by equation (3.3). In preparation, we now describe one version of the construction in the relevant case: when the domains of all (J,ν)-holomorphic maps have no non-trivial automorphisms (cf. Section 1). In this case, for generic (J,ν), each stratum in the moduli space

$$\overline{\mathcal{M}} = \overline{\mathcal{M}}^{(J,\nu)}_{g,n}(X,A) \tag{3.5}$$

(including the "top stratum" $\mathcal{M}$) is a smooth oriented manifold, the boundary $\partial\overline{\mathcal{M}} = \overline{\mathcal{M}} \setminus \mathcal{M}$ is compact, and $\widehat{ev}$ restricts to a smooth map to the compact oriented manifold $\overline{\mathcal{M}}_{g,n} \times X^n$ on each stratum. The Ruan-Tian invariant can be defined, using arguments of Kronheimer and Mrowka, by deleting a neighborhood of the boundary and considering the resulting relative homology class, as follows.

First, a simple compactness and transversality argument (see the first paragraph of the proof of Proposition 4.2 in [KM]) shows that there is an open neighborhood U of the image $\widehat{ev}(\partial\overline{\mathcal{M}})$ satisfying:

(a) $\overline{U}$ is a smooth manifold with boundary and $\widehat{ev}_{|\mathcal{M}}$ is transverse to ∂U, and

(b) there is a (finite) basis for $H_{\ell-2r}(\overline{\mathcal{M}}_{g,n}\times X^n;\mathbb{Q})$ represented by cycles disjoint from $\overline{U}$; here ℓ is the dimension of $\overline{\mathcal{M}}_{g,n}\times X^n$.

Definition 3.1. Given a moduli space (3.5), choose a set U as above and let $\overline{\mathcal{M}}_U\subset\mathcal{M}$ be the closed set

$$\overline{\mathcal{M}}_U=\overline{\mathcal{M}}\cap\widehat{ev}^{-1}\big(\overline{\mathcal{M}}_{g,n}\times X^n\setminus U\big)\subset\mathcal{M}. \tag{3.6}$$

It follows from (a) above that $\overline{\mathcal{M}}_U$ is a compact oriented manifold with boundary, and thus carries a fundamental class in relative homology that we denote by $[\overline{\mathcal{M}}_U]\in H_{2r}(\overline{\mathcal{M}}_U,\partial\overline{\mathcal{M}}_U;\mathbb{Q})$.

On the other hand, the inclusion of the pairs $k:(\overline{\mathcal{M}}_{g,n}\times X^n,\emptyset)\to(\overline{\mathcal{M}}_{g,n}\times X^n,\overline{U})$ induces a map in homology

$$k_*:H_{2r}\big(\overline{\mathcal{M}}_{g,n}\times X^n;\mathbb{Q}\big)\longrightarrow H_{2r}\big(\overline{\mathcal{M}}_{g,n}\times X^n,\overline{U};\mathbb{Q}\big).$$

Proposition 4.2 of [KM] then implies that there is a unique homology class $[\widehat{ev}(\overline{\mathcal{M}})]$ with

$$k_*\big[\widehat{ev}(\overline{\mathcal{M}})\big]=\widehat{ev}_*\big[\overline{\mathcal{M}}_U\big]. \tag{3.7}$$

In fact, [KM] showed that there is a rational singular smooth cycle B that represents $[\widehat{ev}(\overline{\mathcal{M}})]$ and agrees with $\widehat{ev}(\overline{\mathcal{M}}_U)$ outside of U. The uniqueness then follows from property (b) above because the inclusion

$$H_{2r}(\overline{U};\mathbb{Q})\longrightarrow H_{2r}\big(\overline{\mathcal{M}}_{g,n}\times X^n;\mathbb{Q}\big)$$

is trivial and hence k_* is injective. The class $[\widehat{ev}(\overline{\mathcal{M}})]=[B]$ is the GW class:

LEMMA 3.2. *With U and $\overline{\mathcal{M}}_U$ as above, the Ruan-Tian GW class is the unique rational homology class satisfying (3.7), namely*

$$GW_{g,n}(X,A)=\big[\widehat{ev}(\overline{\mathcal{M}})\big]\in H_{2r}\big(\overline{\mathcal{M}}_{g,n}\times X^n;\mathbb{Q}\big). \tag{3.8}$$

In particular, $[\widehat{ev}(\overline{\mathcal{M}})]$ is independent of the choice of U and the generic (J,ν).

Proof. For each $\gamma\in H^{2r}(\overline{\mathcal{M}}_{g,n}\times X^n;\mathbb{Q})$, Ruan and Tian choose a representative Γ of the Poincaré dual of γ with Γ transversal to $\widehat{ev}(\overline{\mathcal{M}})$ and disjoint from $\widehat{ev}(\partial\overline{\mathcal{M}})$. They defined $GW_{g,n}(X,A)(\gamma)$ as $\widehat{ev}(\mathcal{M})\cap\Gamma$ and showed these numbers are independent of the various choices made. But by property (b) above, we can assume that Γ is a linear combination of cycles that do not intersect $\overline{U}$. Then the geometric intersection $\widehat{ev}(\mathcal{M})\cap\Gamma$ is the same as $B\cap\Gamma$, which represents $[B](\gamma)=[\widehat{ev}(\overline{\mathcal{M}})](\gamma)$. ∎

In our case, $X = N_D$ is the total space of the spin curve (D,N). By the Image Localization Lemma 1.1, for small generic ν there is a ϵ-neighborhood $N_D(\varepsilon)$ of the zero section D in N_D such that every (J_α,ν)-holomorphic map has its image in that neighborhood. Although $N_D(\varepsilon)$ is not compact, it is a manifold with boundary and Lemma 3.2 still applies (again by Proposition 4.2 of [KM]). Thus the image of the moduli space represents a rational homology class

$$\text{(3.9)}\quad GW^{\text{loc}}_{g,n}\big(N_D,d\big) := \big[\widehat{ev}\big(\overline{\mathcal{M}}^{(J_\alpha,\nu)}_{g,n}\big(N_D,d[D]\big)\big)\big] \in H_{2\beta+2n}\big(\overline{\mathcal{M}}_{g,n}\times N_D^n;\mathbb{Q}\big).$$

Here the number β, which will occur frequently in our dimension counts, is given in terms of the degree d of the map and the genus h of D by

$$\text{(3.10)}\qquad \beta = d(1-h)+g-1.$$

The class (3.9) is the local GW class of the spin curve (D,N). As shown in [LP], for given g,n,d and h, it depends only on the parity $h^0(N)$.

4. The linearization operator. We now return to the specific situation in which the local GW invariants are defined. Thus, as described after Lemma 1.1, we consider J_α-holomorphic maps into the total space N_D of a spin bundle $N \to D$ over a curve D. This section shows how the special form of the J_α-holomorphic map equation implies vanishing theorems for the linearized operator. Along the way we introduce a family of operators that extends the relevant component of the linearization to general maps.

A neighborhood U of the zero section $D \subset N_D$ can be identified with a neighborhood of the zero section in the projectivization $\mathbb{P}(N\oplus\mathcal{O}_D)$ by a fiber-preserving biholomorphism. Since $\mathbb{P}(N\oplus\mathcal{O}_D)$ is Kähler we can use this identification to define a Kähler structure on U. With this Kähler structure, the exact sequence (1.5) of the underlying complex vector bundle splits as $TN_D = \pi^*TD\oplus\pi^*N$. We will often write this simply as

$$\text{(4.1)}\qquad TN_D = TD\oplus N,$$

leaving the pullbacks as understood; we will then call TD and N the "horizontal" and "vertical" subbundles, respectively. In this and the following sections ∇ will denote the covariant derivative on N induced by the Levi-Civita connection of the Kähler manifold N_D, that is, $\pi_N\circ\nabla^{LC}$ where π_N is the orthogonal projection onto the vertical component of (4.1).

Let $f: C\to N_D$ be a map from a smooth domain C. The linearization of the left-hand side of (1.4) defines an operator

$$D_f:\Gamma\big(f^*TN_D\big)\oplus H^{0,1}(C,TC)\longrightarrow\Gamma\big(\Lambda^{0,1}\big(f^*TN_D\big)\big)$$

given by $D_f(\xi,k) = \widehat{L}_f\xi + \frac{1}{2}(J-K_\alpha)df\,k$. The operator $\widehat{L}_f$ arises from the variation in the map with the complex structure on C held fixed, and the second term arises from the variation k of the complex structure on C. Under the splitting (4.1), $\widehat{L}_f$ decomposes as

$$\widehat{L}_f = \begin{pmatrix} \overline{\partial}_f^T & A \\ B & \widetilde{L}_f \end{pmatrix} : \Omega^0(f^*TD) \oplus \Omega^0(f^*N) \longrightarrow \Omega^{0,1}(f^*TD) \oplus \Omega^{0,1}(f^*N)$$

where A and B are certain bundle maps. For $\xi \in \Gamma(f^*N)$ and $v \in TC$, the component $\widetilde{L}_f$ is given by

$$\widetilde{L}_f\xi(v) = \tfrac{1}{2}\pi_N\big[\nabla_v\xi + J\nabla_{jv}\xi - \nabla_\xi K_\alpha\big(df(jv) + Jdf(v)\big) - K_\alpha\big(\nabla_{jv}\xi + J\nabla_v\xi\big)\big].$$

If the image of f lies in D then K_α vanishes and, by [LP, Lemma 8.2b], the ∇K_α term lies in f^*N. Thus this operator reduces to

$$\widetilde{L}_f = \overline{\partial}_f - \tfrac{1}{2}\nabla K_\alpha(df - Jdfj)j \tag{4.2}$$

where $\overline{\partial}_f$ is the $\overline{\partial}$-operator on f^*N.

Definition 4.1. To each map $f : C \to N_D$, we associate the "modified linearization operator" $L_f : \Omega^0(f^*N) \to \Omega^{0,1}(f^*N)$ given by

$$L_f = \overline{\partial}_f + R_\alpha \quad \text{where } R_\alpha = -\tfrac{1}{2}\pi_N\big[\nabla K_\alpha(df - Jdfj)j\big]. \tag{4.3}$$

If f is a J_α-holomorphic map—or is any other map whose image lies in the zero divisor—then L_f is the vertical-to-vertical component (4.2) of the linearized J_α-holomorphic map equation. Furthermore, if f is J_α-holomorphic, Lemma 8.2 of [LP] shows that $R_\alpha : f^*N \to T^*C \otimes f^*N$ is the complex anti-linear bundle map defined by $R_\alpha(\xi) = -\nabla_\xi K_\alpha \circ df \circ j$. Thus

> *$f \mapsto L_f$ is a family of operators parameterized by maps that agrees with $\widetilde{L}_f$ along the moduli space of J_α-holomorphic maps.*

Much of the rest of this paper will be devoted to understanding the analytic properties of this family of operators.

For a map $f : C \to N_D$ whose domain C is a connected nodal curve, we will regard L_f as an operator on a function space that can be described in terms of the normalization $\pi : \tilde{C} \to C$. The inverse image of each node $n_i \in C$ is a pair of points $p_i, q_i \in \tilde{C}$. For each component $\tilde{C}_k$ of the normalization, let $E_{k,f}$ be the space $\Omega^0(\tilde{C}_k, \pi^*f^*N)$ of smooth sections of π^*f^*N on $\tilde{C}_k$, and similarly let $F_{k,f} = \Omega^{0,1}(\tilde{C}_k, \pi^*f^*N)$. Combine these by setting

$$E_f = \Big\{\xi \in \bigoplus_k E_{k,f} \mid \xi(p_i) = \xi(q_i) \text{ for all } i\Big\} \quad \text{and} \quad F_f = \bigoplus_k F_{k,f}. \tag{4.4}$$

Formula (4.3) then defines an operator

$$L_f : E_f \longrightarrow F_f \tag{4.5}$$

whose restriction to each component $\tilde{C}_k$ gives operators $L_{k,f} : E_{k,f} \to F_{k,f}$ as in (4.3). Using Riemann-Roch (cf. [FO, Lemma 12.2]) one obtains

$$\operatorname{index} L_f = -2\beta \tag{4.6}$$

where β is given by (3.10).

In general, as f varies over the space of stable maps, one expects the dimensions of the kernels of these operators to jump, with compensating jumps in the dimensions of the cokernels. But the following theorem shows that this does not happen for the operators L_f as f varies over the space of J_α-holomorphic maps. This second remarkable fact about the J_α-holomorphic map equation plays a crucial role in our analysis. It implies, as we will show in Sections 5–8, that cokernels of the operators L_f form vector bundles over $\overline{\mathcal{M}}_{g,n}(D,d)$.

VANISHING THEOREM 4.2. *For each J_α-holomorphic map $f : C \to D$ in $\overline{\mathcal{M}}_{g,n}(D,d)$ with $d \neq 0$*

$$\ker L_f = 0 \quad \textit{and} \quad \dim \operatorname{coker} L_f = 2\beta.$$

Proof. Let d_k be the degree the restriction of f to one component $\tilde{C}_k$ of the normalization of C. If $d_k \neq 0$, $L_{k,f}$ is injective by [LP, Proposition 8.6]. On the other hand, when $d_k = 0$, the operator $L_{k,f}$ is the $\overline{\partial}$-operator on the trivial bundle whose kernel is the constant functions. Since any solution of $L_f \xi = 0$ restricts to a solution of $L_{k,f}\xi = 0$ on $\tilde{C}_k$, ξ vanishes on each component with $d_k \neq 0$ and is constant on each component with $d_k = 0$. But ξ is continuous at each node, so $\xi \equiv 0$. This shows L_f is injective and hence the dimension of its cokernel is the negative of its index (4.6). □

5. The bundles $\mathcal{E}$ and $\mathcal{F}$.

We now pass to the level of global analysis by completing the spaces E_f and F_f in (4.5) in appropriate Sobolev norms. As f varies, these define spaces $\mathcal{E}$ and $\mathcal{F}$, each with a projection to the space of map whose fibers are vector spaces and the modified linearization operator (4.5) defines a map $L : \mathcal{E} \to \mathcal{F}$ that is linear on fibers. However, it is not clear whether $\mathcal{E}$ and $\mathcal{F}$ are vector bundles—they may fail to be locally trivial over maps with nodal domains. We will return to the issue of local triviality in the Section 8. Here, in preparation, we define weighted Sobolev norms and show that L_f is a uniformly bounded operator in this context.

Throughout this and the next three sections, we will fix λ with $0 < \lambda < \frac{1}{6}$ and often omit it from the notation. All curves will be fibers of the universal G-twisted curve $\overline{\mathcal{U}}_{g,n}$, but we will suppress g and n from the notation. In Section 2, we fixed a Riemannian metric on $\overline{\mathcal{U}}_{g,n}$ and a defining function ρ for the nodal set. For each $p \geq 2$ define a norm on sections of a vector bundle over a curve C by

$$\|\eta\|_{p,\lambda}^p = \int_C \rho^{p-2-\lambda} |\eta|^p \, dvol_C. \tag{5.1}$$

The space $\mathcal{M}ap_\lambda(X)$ of Definition 2.1 carries two natural vector bundles $\mathcal{E}$ and $\mathcal{F}$; their fibers are described as follows.

• At a map (C, f), let $\mathcal{E}_f$ be the Banach space obtained as the completion of the space E_f in (4.4) with respect to the norm

$$|||\xi|||_{1,p} \; = \; \|\nabla\xi\|_{2,\lambda} + \|\nabla\xi\|_{p,\lambda} + \left(\int_C |\xi|^p\right)^{\frac{1}{p}}. \tag{5.2}$$

• Similarly, let $\mathcal{F}_f$ be the completion of the space F_f in (4.4) with respect to the weighted L^p norm

$$|||\eta|||_p = \|\eta\|_{2,\lambda} + \|\eta\|_{p,\lambda}. \tag{5.3}$$

On each component $B_i(\delta)$ of $B(\delta)$, we can choose coordinates and use the parameterizations (2.9) and (2.10) to rewrite each part of (5.2) as an integral over a cylinder $T(\delta)$. In particular, as in (2.12)

$$\int_{B_i(\delta)} \rho^{p-2-\lambda} |\nabla\xi|^p \cong \int_{T(\delta)} \rho^{-\lambda} |\nabla\xi|^p \quad \text{and} \quad \int_{B_i(\delta)} |\xi|^p \cong \int_{T(\delta)} \rho^{-2} |\xi|^p. \tag{5.4}$$

The first equivalence also holds when $\nabla\xi$ is replaced by any bundle-valued 1-form η.

On the other hand, on any $C(\delta) \subset C$ with $\delta > 0$, the function $\rho^{-\lambda}$ is bounded above and below by positive constants, so the norms (5.2) and (5.3) are equivalent to the standard unweighted $L^{1,p}$ and L^p norms respectively. This is true, in particular, on any smooth domain C (because then $C = C(\delta)$ for small δ). However, the norms in these equivalences depend on δ.

Remark 5.1. Our norms are closely related to those in [LT2], especially if one takes $\lambda = p - 2$ with $2 < p < 13/6$. They are chosen with two points in mind:

(i) When transferred to the cylinder by (5.4), the norm should include an exponential weight $\rho^{-\lambda} \approx e^{\lambda|t|}$ in order to obtain the Poincaré inequalities of Section 6. In terms of the original domain, this gives the awkward-looking weighting factor $\rho^{p-2-\lambda}$ in (2.5) and (5.2).

(ii) Weighted L^p norms give decay estimates like Lemma 2.3 and (5.7) below, while weighted L^2 norms are well-adapted for integration by parts arguments, as we will see in Section 6. The sum of weighted L^p and L^2 norms, as in (2.5), (5.2) and (5.3), has both advantages.

Definition 5.2. Fix $p > 2$, topologize the space

$$\mathcal{E} = \left\{ (C, f, \xi) \mid (C, f) \in \mathcal{M}ap_\lambda(X),\ \xi \in \mathcal{E}_f \right\}$$

by saying that a sequence (C_n, f_n, ξ_n) converges to (C, f, ξ) if

(a) $(C_n, f_n) \to (C, f)$ as in Definition 2.2,

(b) $|||\xi_n|||_{1,p} \to |||\xi|||_{1,p}$,

(c) $\xi_n \to \xi$ in the norm (5.2) on $C(\delta)$ for every $\delta > 0$ after identifying the domains by a trivialization (2.4).

The topological space $\mathcal{F}$ is defined similarly, using triples (C, f, η) with $\eta \in \mathcal{F}_f$ and using the norm (5.3) instead of (5.2).

Corollary 5.5 below shows that the matching condition (4.4) is preserved under convergence in this topology. In preparation, we prove two lemmas about the Sobolev norms (5.2) and (5.3). Recall that in dimension 2, there is a Sobolev embedding $L^{1,p} \subset L^\infty$ for $p > 2$. For the weighted norm (5.2) there is a similar embedding with a constant that is uniform even as necks pinch to become nodes:

LEMMA 5.3. *For each $p > 2$ there is a constant c, depending on p but uniform for domains C in the universal curve and the map $f \in \mathcal{M}ap^E_\lambda(X)$, such that*

$$\|\xi\|_\infty \le c\, |||\xi|||_{1,p} \quad \textit{for all } \xi \in \mathcal{E}_f. \tag{5.5}$$

In particular, convergence in the norm (5.2) implies convergence in C^0.

Proof. Set $\delta = \frac{1}{4}$. Then each fiber $C = C_z$ of the universal curve decomposes as $C(\delta) \cup B(\delta)$ as in (2.2) (we include the case with $B(\delta)$ empty and $C(\delta) = C$) and the diffeomorphisms (2.10) are defined on the 4δ-neighborhoods of the nodes. The domains $C(\delta)$ have uniform geometry, so the norm of the Sobolev embedding $L^{1,p} \hookrightarrow L^\infty$ on $C(\delta)$ is bounded independent of z. This, and the fact that $\frac{1}{4} \le \rho \le 2$ on $C(\delta)$, implies that $\|\xi\|_\infty \le c(p)\left(\|d|\xi|\|_{p,\lambda} + \|\xi\|_p\right)$ on $C(\delta)$. Also, on each component of $B(2\delta)$, the proof of Lemma 2.3 shows that $\operatorname{osc}|\xi| \le c\|d|\xi|\|_{p,\lambda}$. On both domains, we have $\|d|\xi|\|_{p,\lambda} \le c\|\nabla\xi\|_{p,\lambda}$ by "Kato's inequality": for each smooth section ξ and each connection ∇ compatible with the metric one has $|d|\xi|| \le |\nabla\xi|$ almost everywhere. The lemma follows because $\|\xi\|_\infty$ is bounded by the sup of $|\xi|$ on $C(\delta)$ plus the oscillation of $|\xi|$ on $B(2\delta)$. □

The next result improves Lemma 5.3 by bounding the oscillation of ξ—not just $|\xi|$—in the necks $B(\delta)$. This requires replacing Kato's inequality with bounds on the connection. Recall from Section 2 that we have fixed an isometric embedding of X into $\mathbb{R}^N$. Hence any $\xi \in \Gamma(f^*TX)$ can be written as an $\mathbb{R}^N$-valued function $\xi = (\xi^1, \ldots, \xi^N)$ on C and

$$\nabla\xi = d\xi + \left(f^*A\right)\xi \tag{5.6}$$

where ∇ is the pullback connection and f^*A is the pullback of the second fundamental form, which satisfies $|f^*A| \le |A|\,|df|$.

LEMMA 5.4. *For each $p>2$ there are constants c and δ_0 and a neighborhood $\mathcal{N}$ of the space of J_α-holomorphic maps in the λ-topology such that for each $f : C \to N_D$ in $\mathcal{N}$*

(a) the image $f(C)$ lies in the unit disk bundle $N_D(1)$,

(b) for each $\delta < \delta_0$ the oscillation of each $\xi \in \mathcal{E}_f$ on each component of $B(\delta)$ satisfies

$$\operatorname{osc}_{B(\delta)} \xi \le c\,\delta^{\frac{\lambda}{p}}\ |||\xi|||_{1,p}. \tag{5.7}$$

Proof. By Lemma 1.1 the image of each J_α-holomorphic map lies in D. Taking $\varepsilon = 1$ in Lemma 2.4 then shows that there is a neighborhood $\mathcal{N}_1$ of the space of J_α-holomorphic maps such that the images of all maps in $\mathcal{N}_1$ lie in $N_D(1)$.

Bounding the oscillation of each ξ exactly as in Lemma 2.3, we have

$$\operatorname{osc}_{B(\delta)} \xi \le c\,\delta^{\frac{\lambda}{p}} \left(\int_T \rho^{-\lambda}\, |d\xi|^p \right)^{\frac{1}{p}}.$$

We can use (5.6) to replace $d\xi$ by $\nabla\xi - A_f\, df\, \xi$. Part (a) above insures that the image of f lies in $N_D(1)$, and by compactness the second fundamental form A is uniformly bounded on $N_D(1)$. Also using Lemma 5.3, we therefore have

$$\|d\xi\|_{p,\lambda} \ \le\ \|\nabla\xi\|_{p,\lambda} + c\,\|\xi\|_\infty\, \|df\|_{p,\lambda} \le \left(1 + c\sqrt{E_p(f)}\right) \cdot |||\xi|||_{1,p}$$

where E_p is the p-energy (2.6). Finally, set $\mathcal{N} = \mathcal{N}_1 \cap \mathcal{M}ap^{2E}(N_D)$ where E is the constant in Theorem 2.5. Then $\mathcal{N}$ is a neighborhood of the space of J_α-holomorphic maps and each $f \in \mathcal{N}$ satisfies $E_p(f) \le 2E$. The lemma follows from the two displayed equations above. $\square$

One consequence of Lemma 5.4 is that each $\xi \in \mathcal{E}_f$ has a well-defined value $\xi(p_i)$ at each node p_i of the normalization of C, and these values satisfy the matching condition in the definition (4.4) of E_f. This matching condition is not mentioned in Definition 5.2. The following corollary shows that the convergence in Definition 5.2 preserves the matching condition.

COROLLARY 5.5. *If a sequence $\{(C_n, f_n, \xi_n)\}$ in $\mathcal{E}$ converges in the sense of Definition 5.2 to a limit (C, f, ξ), then the limit (C, f, ξ) satisfies the matching condition (4.4), so lies in $\mathcal{E}_f$.*

Proof. As in (4.4), the limit ξ is a section of the pullback of the bundle N to the normalization $\tilde{C}$. Fix a node n of C and let p and q be the two points of $\tilde{C}$

corresponding to the node n. By Lemma 5.4 for each $\varepsilon > 0$ we can fix a $\delta > 0$ such that

$$\mathrm{osc}_{D(p,\delta)}\, \xi \le \varepsilon \quad \mathrm{osc}_{D(q,\delta)}\, \xi \le \varepsilon \quad \mathrm{osc}_{B(\delta)}\, \xi_n \le \varepsilon \quad \forall n$$

where $D(p,\delta)$ and $D(p,\delta)$ are disks in $\tilde{C}$ centered at p and q. For this δ, we can identify $C_n(\delta)$ with $C(\delta)$ and then $\xi_n \to \xi$ on $C(\delta)$ as in Definition 5.2(c). Then $\partial B(\delta)$ consists of two disjoint circles, one identified with $\partial D(p,\delta)$ and the other with $\partial D(q,\delta)$. Choose points $x \in \partial D(p,\delta)$ and $y \in \partial D(q,\delta)$. Then by (5.5)

$$\big|\xi_n(x) - \xi(x)\big| \le c\, |||\xi_n - \xi|||_{1,p} \le \varepsilon$$

for all large n. The same inequality holds with x replaced by y. Altogether, $|\xi(p) - \xi(q)|$ is bounded by the sum of five terms

$$\big|\xi(p) - \xi(x)\big| + \big|\xi(x) - \xi_n(x)\big| + \big|\xi_n(x) - \xi_n(y)\big| + \big|\xi_n(y) - \xi(y)\big| + \big|\xi(y) - \xi(q)\big|$$

with each term bounded by ε. Thus $\big|\xi(p) - \xi(q)\big| \le 5\varepsilon$ for all $\varepsilon > 0$, so the limit ξ satisfies the matching condition $\xi(p) = \xi(q)$. $\square$

We next show that L_f is a bounded operator with the norms (5.2) and (5.3). For our purposes it is important to have bounds that are uniform for all maps f is some neighborhood of the space of J_α-holomorphic maps.

PROPOSITION 5.6. (Uniform boundedness of L_f) *For each $p > 2$ there is a constant $c = c(p)$ such that for each $f : C \to N_D$ in the neighborhood $\mathcal{N}$ of Lemma 5.4 and each $\xi \in \mathcal{E}_f$*

$$\big|\big|\big|L_f \xi\big|\big|\big|_p \le c\,(1 + \sqrt{E})\, |||\xi|||_{1,p}$$

where E is the constant of Theorem 2.5.

Proof. From (4.3) we have the pointwise bound $|L_f \xi| \le |\nabla \xi| + |\nabla \alpha||\xi||df|$. By Lemma 5.4 the image of f lies in a compact set in N_D where $|\nabla \alpha| < c_1$ for some c_1. Integrating using norms (5.2) and (5.3) and definition (2.6), we have

$$\big|\big|\big|L_f \xi\big|\big|\big|_p \le |||\xi|||_{1,p} + c_1 \|\xi\|_\infty \sqrt{E(f)}$$

with $\|\xi\|_\infty$ bounded as in (5.5) and $E(f) \le 2E$ as at the end of the proof of Lemma 5.4. The proposition follows. $\square$

6. Poincaré inequalities on weighted spaces. For each J_α-holomorphic map $f : C \to N_D$, the modified linearization operator (4.5) is an elliptic operator whose kernel is 0 by Theorem 4.2. Consequently, there is a Poincaré inequality of the form $\|\xi\| \le c\|L_f \xi\|$ for a variety of norms. The constant c in this inequality need not be uniform in a family of maps whose domains pinch to a nodal curve.

For the weighted Sobolev norms introduced in Section 5, however, there is such an inequality with uniform constant. This will be proven in Theorem 7.7 based on the results in this section.

Because L_f is the $\overline{\partial}$-operator plus lower order terms, the key step is establishing Poincaré inequalities for $\overline{\partial}$ on the cylinders that parameterize the thin part of the domain curves as described in Section 2. Thus we continue to use the norms (5.1)–(5.3), but now will work on the $T(\delta)$ side of (5.4) with the $p=2$ norms. These cylinders are of two types: half-infinite cylinders, which parameterize neighborhoods of nodes of C, and finite cylinders that parameterize necks in a nearly-nodal domain C. In both cases we will consider complex-valued functions $\phi : T \to \mathbb{C}$ satisfying

$$\|\overline{\partial}\phi\|_{2,\lambda}^2 = \int_T \rho^{-\lambda}\, |\overline{\partial}\phi|^2 \,<\, \infty \quad \text{and} \quad \int_T |d\phi|^2 < \infty \tag{6.1}$$

and seek Poincaré inequalities with constants independent of the length of the cylinder.

We begin by observing that if ϕ satisfies (6.1) then the weighted L^2 norm of $d\phi$ is finite. This is immediately true for smooth domains because $\rho^{-\lambda}$ is bounded on the domain. The following lemma shows that it is also true in the nodal case.

LEMMA 6.1. (Nodes) *Let* $T=[0,\infty)\times S^1$ *with coordinates* (t,θ). *If* $\phi : T\to\mathbb{C}$ *is a smooth function satisfying (6.1), then* $\|d\phi\|_{2,\lambda}<\infty$ *and for* $0<\lambda<\frac{1}{6}$

$$\int_{T'} e^{\lambda t}\,|d\phi|^2 \le 6\int_T e^{\lambda t}\,|\overline{\partial}\phi|^2 + 3\int_E e^{\lambda t}\,|d\phi|^2$$

where $T'=[\frac{1}{2},\infty)\times S^1$ *and* E *denotes the end* $[0,\frac{1}{2}]\times S^1$.

Proof. Writing $\phi = u+iv$, we have

$$|d\phi|^2\,dt\,d\theta = 4|\overline{\partial}\phi|^2\,dt\,d\theta \,+\, 2d(u\,dv).$$

Multiplying by $e^{\lambda t}$ and noting that $d\big(e^{\lambda t}u\,dv\big) = e^{\lambda t}d(u\,dv)+\lambda e^{\lambda t}u\,v_\theta\,dt\,d\theta$ gives

$$e^{\lambda t}|d\phi|^2\,dt\,d\theta \;=\; 4e^{\lambda t}|\overline{\partial}\phi|^2\,dt\,d\theta \,+\, 2d\big(e^{\lambda t}u\,dv\big) - \lambda e^{\lambda t}u\,v_\theta\,dt\,d\theta.$$

Integrating over any circle $S_t=\{t\}\times S^1$ and integrating by parts, one also sees that

$$\int_{S_t} u\,v_\theta = \int_{S_t} \tfrac{i}{2}\left(\phi\,\overline{\phi_\theta}\right) \le \tfrac{1}{2}\int_{S_t} |d\phi|^2$$

(this last inequality can be checked by Fourier expansions). Fixing $s\ge 0$ and integrating over the annulus $A_{st}=[s,t]\times S^1\subset T$ using the above equations and

Stokes' Theorem then gives

$$\int_{A_{st}} e^{\lambda t}|d\phi|^2 \leq 4\int_T e^{\lambda t}|\overline{\partial}\phi|^2 + 2\int_{S_t} e^{\lambda t}|d\phi|^2 + 2\int_{S_s} e^{\lambda s}|d\phi|^2 + \tfrac{\lambda}{2}\int_{A_{st}} e^{\lambda t}|d\phi|^2.$$

Setting

$$\Phi(t) = \int_{A_{st}} e^{\lambda t}|d\phi|^2 - 6\int_T e^{\lambda t}|\overline{\partial}\phi|^2 - 3\int_{S_s} e^{\lambda s}|d\phi|^2, \tag{6.2}$$

rearranging, and using the fact that $0 < \lambda < \frac{1}{6}$, we obtain the inequality

$$\Phi(t) \leq 3\Phi'(t). \tag{6.3}$$

Now suppose $\Phi(t_0) > 0$ for some $t_0 > s$. Then $\Phi(t) > 0$ for all $t > t_0$. Integrating (6.3) over the interval $[t, L]$ shows that

$$\Phi(t) \leq e^{\frac{t-L}{3}}\, E(L)$$

for $t_0 \leq t < L$. In particular,

$$0 < \Phi\big(t_0\big) \leq c\big(t_0\big)\, e^{\frac{-L}{3}} \int_0^{2\pi}\int_0^L e^{\lambda t}|d\phi|^2 \leq c\big(t_0\big)\, e^{(\lambda - \frac{1}{3})L}\, \|d\phi\|_2^2.$$

Since $\lambda < \frac{1}{6}$ we have a contradiction for large L. Therefore, $\Phi(t) \leq 0$ for all t and all s. With this, the lemma follows by integrating both sides of (6.2) over the s variable from 0 to $\frac{1}{2}$. □

LEMMA 6.2. (Nodes) *For each smooth function $\phi : T \to \mathbb{C}$ satisfying (6.1), there exists a unique constant $\overline{\phi} \in \mathbb{C}$ such that the integrals over the sets T', T and E of Lemma 6.1 satisfy*

$$\int_{T'} \rho^{-\lambda}\,|\phi - \overline{\phi}|^2 \leq c(\lambda)\int_T \rho^{-\lambda}\,|\overline{\partial}\phi|^2 + 8\int_E \rho^{-\lambda}\,|d\phi|^2 \tag{6.4}$$

where $c(\lambda) = 4 + 8\lambda^{-2}$.

Proof. First define a function $g(t)$ by averaging over circles $S_t = \{t\} \times S^1$

$$g(t) = \frac{1}{2\pi}\int_{S_t} \phi. \tag{6.5}$$

The oscillation of $g(t)$ on $[L,\infty)$ can be estimated by using the fundamental theorem of calculus and applying Holder's inequality to $g' = \rho^{\lambda/2} \cdot \rho^{-\lambda/2} g'$:

$$\begin{aligned} \operatorname{osc}_{[L,\infty)} g &\leq \int_L^\infty |g'| \leq \left(\int_L^\infty e^{-\lambda t}\, dt\right)^{\frac{1}{2}} \left(\int_L^\infty \rho^{-\lambda}|g'|^2\, dt\right)^{\frac{1}{2}} \\ &\leq \tfrac{1}{\sqrt{\lambda}}\, e^{-\frac{\lambda L}{2}}\, \|d\phi\|_{2,\lambda;[L,\infty)} \end{aligned} \tag{6.6}$$

(observe that $\|d\phi\|_{2,\lambda}$ is finite by Lemma 6.1). It follows that $g(t)$ approaches a constant $\overline{\phi} \in \mathbb{C}$ as $t \to \infty$, and $|g(t) - \overline{\phi}|$ is bounded for $t \geq L$ by the right-hand side of (6.6). Writing $\lambda e^{\lambda t}\, dt = d(e^{\lambda t})$ and integrating by parts gives

$$\lambda \int_0^L e^{\lambda t} |g - \overline{\phi}|^2\, dt \leq e^{\lambda t}\, |g - \overline{\phi}|^2 \Big|_0^L + \int_0^L e^{\lambda t} \left(\varepsilon |g - \overline{\phi}|^2 + \varepsilon^{-1} |g'|^2\right) dt$$

for any $\varepsilon > 0$. The lower boundary term is negative, so can be dropped. The upper boundary term is dominated by $e^{\lambda L}$ times the square of the right-hand side of (6.6), so vanishes in the limit $L \to \infty$. Taking $\varepsilon = \frac{1}{2}\lambda$ and rearranging gives

$$\int_0^\infty e^{\lambda t}\, |g - \overline{\phi}|^2\, dt \;\leq\; 4\lambda^{-2} \int_0^\infty e^{\lambda t}\, |g'|^2\, dt. \tag{6.7}$$

Next, write the Fourier expansion of ϕ in the form

$$\phi(t,\theta) - \overline{\phi} \;=\; f(t) \;+\; \sum_{n \neq 0} a_n(t)\, e^{in\theta} \tag{6.8}$$

with $f(t) = g(t) - \overline{\phi}$. Integrating over the circles S_t, we have

$$\frac{1}{2\pi} \int_{S_t} |\phi - \overline{\phi}|^2 \;=\; |f(t)|^2 + \sum_{n \neq 0} |a_n(t)|^2 \tag{6.9}$$

and

$$\frac{1}{\pi} \int_{S_t} |\overline{\partial}\phi|^2 = |f'|^2 + \sum_{n \neq 0} \big|b_n(t)\big|^2 \quad \text{where } b_n = a_n' - n a_n. \tag{6.10}$$

Let $\beta = \beta_L(t)$ be a smooth positive cutoff function equal to 1 on $[\frac{1}{2}, L]$, with support in $[0, L + \frac{1}{2}]$ and satisfying $\beta \leq 1$ and $|d\beta| \leq 4$. Integrating $\beta e^{\lambda t}\, |b_n|^2$ by parts and observing that $|b_n|^2 \geq -n(|a_n|^2)' + n^2 |a_n|^2$ yields

$$\int_0^\infty \beta\, e^{\lambda t} \sum_{n \neq 0} \big|b_n\big|^2 \geq \frac{1}{2} \int_0^\infty \beta\, e^{\lambda t} \sum_{n \neq 0} \big|a_n\big|^2 - \frac{4}{2\pi} \int_{E \cup T_L} e^{\lambda t}\, |d\phi|^2,$$

where $T_L = [L, \infty) \times S^1$ and where we have used the inequalities $n^2 - |n|\lambda \geq \frac{1}{2}$ and $n|a_n|^2 \leq n^2 |a_n|^2$ for $n \neq 0$. Now take $L \to \infty$. Because $\|d\phi\|_{2,\lambda}$ is finite by Lemma 6.1, the part of the last integral over T_L vanishes in the limit, leaving

$$2\pi \int_{\frac{1}{2}}^\infty e^{\lambda t} \sum_{n \neq 0} |a_n|^2 \leq 4\pi \int_0^\infty e^{\lambda t} \sum_{n \neq 0} |b_n(t)|^2 + 8 \int_E e^{\lambda t}\, |d\phi|^2.$$

On the other hand, multiplying (6.9) by $e^{\lambda t}$, integrating over t, and using (6.7) gives

$$\int_{T'} e^{\lambda t}\, |\phi - \overline{\phi}|^2 \leq \frac{8\pi}{\lambda^2} \int_0^\infty e^{\lambda t}\, |f'(t)|^2 + 2\pi \int_{\frac{1}{2}}^\infty e^{\lambda t} \sum |a_n|^2.$$

The desired inequality (6.4) follows by combining the last two displayed equations and observing that each term on the right-hand side of (6.10) is less than the left-hand side of (6.10).

Finally, to prove the uniqueness of $\overline{\phi}$, suppose that (6.4) holds for two constants $\overline{\phi}_1$ and $\overline{\phi}_2$. Then by the triangle inequality

$$\int_0^L \int_0^{2\pi} e^{\lambda t} |\overline{\phi}_1 - \overline{\phi}_2|^2 = 2\pi\lambda^{-1}\left(e^{\lambda L} - 1\right) |\overline{\phi}_1 - \overline{\phi}_2|^2$$

is bounded, uniformly in L. Taking $L \to \infty$ then shows that $\overline{\phi}_1 = \overline{\phi}_2$. □

LEMMA 6.3. (Necks) *Let $T = [-L, L] \times S^1$ with coordinates (t, θ) and $\rho^2 = 2|\mu|\cosh(2t)$. Then for each smooth function $\phi : T \to \mathbb{C}$, there exists a constant $\overline{\phi} \in \mathbb{C}$ such that*

$$\int_{T'} \rho^{-\lambda} |\phi - \overline{\phi}|^2 \le c(\lambda) \int_T \rho^{-\lambda} |\overline{\partial}\phi|^2 + 8 \int_E \rho^{-\lambda} |d\phi|^2 \tag{6.11}$$

where $T' = [-L + \frac{1}{2}, L - \frac{1}{2}] \times S^1$, E is the end $E = T \setminus T'$ and $c(\lambda) = 4 + 16\lambda^{-2}$.

Proof. Define $\overline{\phi} \in \mathbb{C}$ be the average value of ϕ over the center circle $S_0 = \{0\} \times S^1$:

$$\overline{\phi} = \frac{1}{2\pi} \int_{S_0} \phi. \tag{6.12}$$

Then $\phi - \overline{\phi}$ has the Fourier expansion (6.8) with $f(0) = 0$. For $t \ge 0$ write $\rho^{-\lambda} = \sigma'$ on $[-L, L]$ where $\sigma(t) = -\int_t^L \rho^{-\lambda}(r)\, dr$. Integrating by parts, we obtain

$$\int_0^L \rho^{-\lambda} |f|^2\, dt = \sigma |f|^2 \Big|_0^L - \int_0^L \sigma\, 2\,\mathrm{Re}\, \overline{f} df \le 0 + \int_0^L |\sigma| \left(\varepsilon |f|^2 + \varepsilon^{-1} |f'|^2\right) dt$$

for any $\varepsilon > 0$. Because $\rho^2 \le 2|\mu| e^{2t} \le 2\rho^2$ we also have $|\sigma| \le 2\rho^{-\lambda}/\lambda$. After again taking $\varepsilon = \frac{1}{4}\lambda$ and rearranging, we are left with the inequality

$$\int_0^L \rho^{-\lambda}\, |f|^2\, dt \le 16\lambda^{-2} \int_0^L \rho^{-\lambda} |f'|^2\, dt. \tag{6.13}$$

Because ρ is an even function, (6.13) also holds on the interval $[-L, 0]$ and hence on $[L, L]$.

Now repeat the end of the proof of Lemma 6.2. Integrating the Fourier expansion (6.8) over the circles S_t again gives (6.9) and (6.10). This time, relate f and f' by (6.13) and choose a smooth positive cutoff function β that is equal to 1 on $[-L + \frac{1}{2}, L - \frac{1}{2}]$, with support in $[-L, L]$ and satisfying $\beta \le 1$ and $|d\beta| \le 4$. Multiply the term linear in n by $\beta\rho^{-\lambda}$, integrate over t, and integrate by parts, noting that $|d \log \rho| = |\tanh(2t)| \le 1$ from (2.11). The lemma follows. □

Remark. While the constant $\overline{\phi}$ in Lemma 6.3 is not unique, the end of the proof of Lemma 6.2 shows that the range of possibilities for $\overline{\phi}$ shrinks exponentially as the neck T becomes longer.

Together, Lemmas 6.2 and 6.3 give a Poincaré inequality for functions on the thin part $B(\delta)$ of a curve C. As in Section 2, each component of the thin part is parameterized by a set $T(\delta)$ that is either a finite cylinder or two half-infinite cylinders. In both cases we use the weighted norm (2.12) on functions.

PROPOSITION 6.4. *There is a constant δ_0 such that for each $\delta < \delta_0$ and each $L^{1,2}_{loc}$ function $\phi : T(2\delta) \to \mathbb{C}$ satisfying (6.1), there exists a $\overline{\phi} \in \mathbb{C}$ with*

$$\int_{T(\delta)} \rho^{-\lambda}\,|\phi-\overline{\phi}|^2 \leq 32\lambda^{-2}\int_{T(2\delta)} \rho^{-\lambda}\,|\overline{\partial}\phi|^2 + 8\int_{B(2\delta)\setminus B(\delta)} \rho^{-\lambda}\,|d\phi|^2. \tag{6.14}$$

Proof. In the chart (2.8), the metric and the distance function on $B(\delta)$ induced by the metric on the universal curve are, for small δ, uniformly close to the metric and the distance function induced on the curve $xy = \mu$ by the Euclidean metric on $\mathbb{C}^2$. Thus it suffices to show (6.14) for the Euclidean metric, taking ρ^2 to be $|x|^2 + |y|^2$. Because smooth functions are dense in $L^{1,2}$ it also suffices to assume that ϕ is smooth.

Each component of $B(2\delta)$ corresponds, under the parameterization ϕ of (2.9) or (2.10), to a cylinder $T(2\delta)$ isometric to either the cylinder T in Lemma 6.2 or the T in Lemma 6.3. In both cases, the image $\phi(E)$ lies in the annulus $B(2\delta)\setminus B(\delta)$ (note $\sqrt{e} < 2$). Thus (6.14) follows from Lemmas 6.2 and 6.3. $\square$

7. Estimates on the linearization L_f. This section contains the essential elliptic estimates for the modified linearization operator L_f defined by (4.3). Such estimates are standard for unweighted Sobolev spaces, but we need them for the weighted spaces $\mathcal{E}_f$ and $\mathcal{F}_f$, and we need them to be locally uniform on the space of maps $\mathcal{M}ap_\lambda(X)$ of Definition 2.2. Care is needed because on weighted Sobolev spaces the Rellich compactness lemma may fail and because elliptic operators need not be Fredholm (cf. [L]).

7.1. The interior estimate. The first step is to parlay the Poincaré inequality (6.14) into an estimate on the thin region $B(\delta)$. This step is possible only because we are working with weighted norms.

LEMMA 7.1. (Interior estimate) *Fix $p > 2$ and a J_α-holomorphic map f. Then there is a neighborhood $\mathcal{N}_f$ of f in $\mathcal{M}ap_\lambda(N_D)$ and constants δ_1, c_1 and $c_2 = c_2(\delta)$ such that if $g \in \mathcal{N}_f$, $\delta < \delta_1$, $\xi \in L^{1,p}$ and $L_g\xi \in \mathcal{F}_g$, then $\xi \in \mathcal{E}_g$ and*

$$|||\xi|||_{1,p;B(\delta)} \leq c_1\,|||L_g\xi|||_{p;B(4\delta)} + c_2\,|||\xi|||_{1,p;B(4\delta)\setminus B(\delta)}\,. \tag{7.1}$$

Proof. By Lemma 1.1 the image of f lies in the set where $\alpha = 0$, and by Lemma 2.3 we can choose $\mathcal{N}_f$ and δ small enough that the image of each component of $B(4\delta)$ lies in a holomorphic coordinate chart for the bundle $N \to N_D$. In such a chart, we can regard $g \in \mathcal{N}_f$ and any $\xi \in \mathcal{E}_g$ as complex-valued functions on $B(2\delta)$ or, equivalently, on $T(2\delta)$ via (2.9) and (2.10). To obtain uniform estimates, we will work on $T(2\delta)$, and return to $B(2\delta)$ at the end of the proof.

Using formula (4.3) we can write L_g as the standard $\overline{\partial}$-operator on functions plus additional terms:

(7.2)
$$L_g\xi = \overline{\partial}\xi + (A_g\xi)^{0,1} \quad \text{where } A_g \text{ has the form } A_g\xi = g^*\Gamma\xi + \nabla_\xi K_\alpha \circ dg \circ j.$$

Here j is the complex structure on C, Γ is a term built from the Christoffel symbols satisfying $|g^*\Gamma| \le c_3 \operatorname{diam}(f(B(\delta)))|dg| \le c_4|dg|$, $|\nabla K_\alpha \circ dg| \le c_5\,|dg|$. Similarly, $\nabla\xi = d\xi + \Gamma\xi$ with the same bounds on Γ. Hence

$$\big||L_g\xi| - |\overline{\partial}\xi|\big| + \big||\nabla\xi| - |d\xi|\big| \le c_6\,|dg|\,|\xi|. \tag{7.3}$$

By the definition of $\mathcal{M}ap_\lambda(N_D)$, the weighted p-norm $|||dg|||_p$ is finite. Furthermore, $|\xi|$ is bounded because of the hypothesis $\xi \in L^{1,p}$ and the Sobolev embedding $L^{1,p} \subset C^0$. These facts insure that both the operator A_g in (7.2) and the right-hand side of (7.3) are in $\mathcal{F}_g$. We then see that both conditions in equation (6.1) hold:

- Using the hypothesis $L_g\xi \in \mathcal{F}_g$, (7.3) implies that $\overline{\partial}\xi \in \mathcal{F}_g$.
- By the hypothesis $\xi \in L^{1,p}$, the L^2 norm of $\nabla\xi$ on $B(\delta)$ is finite. But then $\nabla\xi \in L^2(T(\delta))$ (take $p = 2$ and $\lambda = 0$ in (5.4)), and hence (7.3) shows that $d\xi \in L^2(T(\delta))$.

Thus the lemmas of Section 6 apply to ξ.

Now restrict attention to one component of $B(\delta)$ and set $\zeta = \xi - \overline{\xi}$ where $\overline{\xi} \in \mathbb{C}$ is the constant provided by Proposition 6.4. For each integer n let $A(n) \subset T$ be the annulus with $n \le t \le n+1$ and let $A^+(n)$ be the annulus with $n-1 \le t \le n+2$. Since $\zeta \in L^{1,p}$ on A_n, standard elliptic theory (Lemma 1.2.2 of [IS1] or Lemma C.2.1 of [MS] and interpolation) shows that

$$\|\zeta\|_{1,p;A(n)} \le c_7 \left(\|\overline{\partial}\zeta\|_{p;A^+(n)} + \|\zeta\|_{2;A^+(n)} \right) \tag{7.4}$$

where these are unweighted L^p norms. The constant c_7 is independent of n because the $A^+(n)$ are isometric to one another. Holder's inequality shows that (7.4) also holds when the p on the left side is replaced by 2 (retaining the p on the right side).

Recall from (2.11) that ρ is equal to e^{-t} on the half-infinite cylinders and satisfies $|\mu|e^{2|t|} \le \rho^2 \le 2|\mu|e^{2|t|}$ on the finite cylinders. In both cases, the ratio of the maximum to the minimum value of $\rho^{-\lambda}$ across $A^+(n)$ is bounded uniformly in n. Therefore (7.4) remains valid (with a modified constant) when the norms are weighted by $\rho^{-\lambda}$; taking the p-power of both sides gives the factor of δ below. Note that the union of the $A^+(n)$ lies in $T(2\delta)$ and that each point on $T(\delta)$ lies in three of the $A^+(n)$. Summing (7.4) on n and using the inequality $\rho \leq 2\delta$ on $T(\delta)$, one therefore obtains

$$\int_{T(\delta)} \rho^{-\lambda} |d\zeta|^p \leq c_8 \int_{T(2\delta)} \rho^{-\lambda} |\overline{\partial}\zeta|^p + c_8 (2\delta)^{\frac{p-2}{p}\lambda} \left(\int_{T(2\delta)} \rho^{-\lambda} |\zeta|^2 \right)^{\frac{p}{2}}.$$

Here we have used the fact that $\overline{\partial}\zeta = \overline{\partial}\xi \in \mathcal{F}_g$ which, by Lemma 6.2, insures that the two integrals on the right-hand side—and therefore all three integrals—are finite. It also insures that the $|\zeta|^2$ term is bounded as in the Poincaré inequality (6.14) on $T(4\delta)$. Consequently,

$$\|d\zeta\|_{p,\lambda;T(\delta)} \leq c_9 \Big(\|\overline{\partial}\zeta\|_{p,\lambda;T(4\delta)} + \|\overline{\partial}\zeta\|_{2,\lambda;T(4\delta)} + \|d\zeta\|_{2,\lambda;T(4\delta)\setminus T(\delta)} \Big).$$

But $\zeta = \xi - \overline{\xi}$ satisfies $d\zeta = d\xi$ and $\overline{\partial}\zeta = \overline{\partial}\xi$. After applying (7.3) we obtain

$$\begin{aligned} \|\nabla\xi\|_{p,\lambda;T(\delta)} \leq c_{10} \Big(& \|L_g\xi\|_{p,\lambda;T} + \|L_g\xi\|_{2,\lambda;T} + \|dg\,\xi\|_{p,\lambda;T} \\ & + \|dg\,\xi\|_{2,\lambda;T} + \|\nabla\zeta\|_{2,\lambda;T\setminus T(\delta)} \Big) \end{aligned}$$

where T denotes $T(4\delta)$. This inequality holds in particular for $p = 2$. Adding the p and $p = 2$ cases yields

$$|||\nabla\xi|||_{p;T(\delta)} \leq c_{11} \Big(|||L_g\xi|||_{p;T} + \|\xi\|_{\infty,T}\, |||dg|||_{p;T} \Big) + c_{12}\, |||\xi|||_{1,p;T\setminus T(\delta)}. \tag{7.5}$$

Now Lemmas 7.2 and 7.3 below show that there is a constant c_{13}, independent of δ, such that

$$\|\xi\|_{\infty,T}\, |||dg|||_{p;T} \leq c_{13} \left(\delta^{\frac{\lambda}{p}}\, |||\nabla\xi|||_{p;T(\delta)} + \delta^{-\frac{2}{p}}\, |||\xi|||_{1,p;T\setminus T(\delta)} \right) \cdot \delta^{\frac{1}{3}-\frac{\lambda}{2}}.$$

Hence after taking δ sufficiently small, the middle term on the right-hand side of (7.5) can be absorbed partly on the other side of the equation and partly into the last term in (7.5). Returning to $B(\delta)$ and noting the equivalence at the beginning of the proof of Proposition 6.4, we then have

$$|||\nabla\xi|||_{p;B(\delta)} \leq c_{14}\, |||L_g\xi|||_{p;B(4\delta)} + c_{15}(\delta)\, |||\xi|||_{1,p;B(4\delta)\setminus B(\delta)}. \tag{7.6}$$

Finally, the last term in the norm (5.2) on $B(\delta)$ can be bounded using (7.7) below:

$$\left(\int_{B(\delta)} |\xi|^p \right)^{\frac{1}{p}} \leq \big[\mathrm{Area}(B(\delta))\big]^{\frac{1}{p}} \|\xi\|_\infty \leq c_{16}\, \delta^{\frac{2}{p}} \left(|||\nabla\xi|||_{p;B(\delta)} + |||\xi|||_{1,p;B(4\delta)\setminus B(\delta)} \right).$$

The inequality of the lemma follows by combining this with (7.6) and taking δ sufficiently small. Since we have shown these integrals are finite, we also conclude that $\xi \in \mathcal{E}_g$. $\square$

The above proof made use of the following two technical lemmas.

LEMMA 7.2. *Fix $p > 2$. There is a constant $c = c(p)$, independent of $\delta \le 1$, such that for each $g : T(4\delta) \to X$ with finite p-energy, each $\xi \in \mathcal{E}_g$ satisfies*

$$\|\xi\|_{\infty,T(4\delta)} \le c\left(\delta^{\frac{\lambda}{p}}\, |||\nabla\xi|||_{p;T(\delta)} + \delta^{-\frac{2}{p}}\, |||\xi|||_{1,p;T(4\delta)\setminus T(\delta)}\right). \tag{7.7}$$

Proof. As in Lemmas 6.1 and 6.3, let E be the end of $T = T(4\delta)$ isometric to one or two copies of $[0, \frac{1}{2}] \times S^1$. We can bound the oscillation of $|\xi|$ as in the proofs of Lemmas 2.3 and 5.3, and bound $|\xi|$ on E using the Sobolev embedding $L^{1,p} \subset L^\infty$:

$$\|\xi\|^p_{\infty,T} \;\le\; (\mathrm{osc}_T\,|\xi|)^p + \|\xi\|^p_{\infty;E} \le c_1\,\delta^\lambda\, |||\nabla\xi|||_{p;T} + \int_E |d|\xi||^p + |\xi|^p.$$

But $|d|\xi||^p \le \delta^\lambda \rho^{-\lambda} |\nabla\xi|^p$ by Kato's inequality, so we obtain

$$\|\xi\|^p_{\infty,T} \le c_1\,\delta^\lambda \left(|||\nabla\xi|||^p_{p;T(\delta)} + |||\nabla\xi|||^p_{p;T\setminus T(\delta)}\right) + \delta^\lambda \int_E \rho^{-\lambda}|\nabla\xi|^p \;+ \delta^{-2}\int_E \rho^2|\xi|^p.$$

Translating the norm (5.2) to T using (5.4), one sees that this gives (7.7) for $\delta \le 1$ because $E \subset T(4\delta)\setminus T(\delta)$ under the transformations (2.9) and (2.10). □

LEMMA 7.3. *Fix $p > 2$ and a J_α-holomorphic map f. There are positive constants δ_0 and c such that for each $\delta < \delta_0$ there is a neighborhood $\mathcal{N}_f(\delta)$ of f such that, for each $(C, g) \in \mathcal{N}_f(\delta)$, the function $|dg|$ satisfies $|||dg|||_p \le c\,\delta^{\frac{1}{3}-\frac{\lambda}{2}}$ on $B(\delta) \subset C$.*

Proof. The energy bound (2.15) on the J_α-holomorphic map f shows that $|||df|||_{p;B_0(\delta)} \le c_1\,\delta^{\frac{1}{3}-\frac{\lambda}{2}}$ where $B_0(\delta)$ is the thin part of the domain C_0 of f. Also, by the triangle inequality,

$$\begin{aligned} |||dg|||_{p;B(\delta)} &\le |||df|||_{p;B_0(\delta)} + \left|\sqrt{E_p(g,\delta)} - \sqrt{E_p(f,\delta)}\right| \\ &\quad + \left|\sqrt{E_2(g,\delta)} - \sqrt{E_2(f,\delta)}\right| \end{aligned} \tag{7.8}$$

where $E_p(g,\delta)$ is the p-energy integral (2.6) on the thin part $B(\delta)$ of the domain C of g. The last two terms in (7.8) vanish as $g \to f$ in the λ-topology because, by parts (b) and (c) of Definition 2.2,

$$E_2(g,\delta) = E_2(g) - \|dg\|^2_{2,\lambda;C(\delta)} \longrightarrow E_2(f) - \|df\|^2_{2,\lambda;C_0(\delta)} \;=\; E_2(f,\delta).$$

One similarly sees that $E_p(g,\delta) \to E_p(f,\delta)$. Thus $|||dg|||_p \le c\,\delta^{\frac{1}{3}-\frac{\lambda}{2}}$ whenever g is close to f. □

7.2. The exterior estimate. For a fixed map $f : C \to N_D$, standard theory gives a elliptic estimate on the thick part $C(\delta)$ of the form

$$\|\xi\|_{1,p,C(2\delta)} \leq c\left(\|L_f\xi\|_{p,C(\delta)} + \|\xi\|_{p,C(\delta)}\right) \tag{7.9}$$

where these are unweighted $L^{1,p}$ and L^p norms. The next lemma shows that the corresponding inequality holds for the weighted norms, uniformly for maps near f. In fact, because the weighting function ρ is bounded above and below on $C(\delta)$, the extension to weighted norms, with a constant depending on δ, is immediate. The issue, then, is whether the weighted version of (7.9) holds with a constant that is locally constant on the space of maps with the λ-topology of Definition 2.2.

LEMMA 7.4. (Uniform exterior estimate) *Fix $p > 2$ and a J_α-holomorphic map (C_0, f). Then for each δ, $0 < \delta < 1$, and each $\varepsilon > 0$, there is a neighborhood $\mathcal{N}_f$ of f in $\mathcal{M}ap_\lambda(N_D)$ and a constant $c = c(p,\delta)$ such that for all $g \in \mathcal{N}_f$, each $\xi \in L^{1,p}(C(\delta))$ satisfies*

$$|||\xi|||_{1,p;C(2\delta)} \leq c\left(|||L_g\xi|||_{1,p;C(\delta)} + |||\xi|||_{p;C(\delta)}\right) + \varepsilon\, |||\xi|||_{1,p;C(\delta)}\,. \tag{7.10}$$

Proof. First, by Definition 2.2(a), there is a neighborhood $\mathcal{N}_f \subset \mathcal{N}$ so that the domains of all maps (C,g) in $\mathcal{N}_f$ lie in a single chart in the universal curve of the form (2.3) or (2.4). This chart gives an identification $C(\delta) = C_0(\delta)$ that is C^1 close to an isometry. We can regard each map $g : C(\delta) \to N_D$ in $\mathcal{N}_f$ as a map $g : C_0(\delta) \to N_D$; by the proof of Lemma 2.4 these satisfy

$$\sup_{x\in C_0(\delta)} \operatorname{dist}(f(x), g(x)) \leq c_1\, |||f-g|||_{1,p;C_0(\delta)}\,. \tag{7.11}$$

Recall that the J_α-holomorphic map f is smooth and, by Lemma 1.1, has image in the zero section D of N_D. Thus (7.11) implies that the images of all $g \in \mathcal{N}_f$ lie in a neighborhood of the zero section where there is a uniform bound on $|\nabla K_\alpha|$.

Now fix $x \in C_0(2\delta)$ and choose a disk $D(x,2r) \subset C_0(\delta)$. From Definition 4.1, the operator L_f has the form $L_f = \overline{\partial} + R_\alpha$ where $|R_\alpha| \leq 2|\nabla K_\alpha||df|$ is bounded. Thus $R_\alpha \in L^p$. Elliptic theory for such operators, as in [MS, Appendix C.2], yields an estimate

$$\|\xi\|_{1,p,D(x,r)} \leq c_3\Big(\|L_f\xi\|_{p,D(x,2r)} + \|\xi\|_{p,D(x,2r)}\Big)$$

with c_3 depending on (C_0, f) and r. We can add $\|\xi\|_{1,2,D(x,r)}$ to the left-hand side by Holder's inequality. We can also insert the weighting function ρ into each of the integrals, noting that $\delta \leq \rho \leq 2$ on $D(x,2r) \subset C(\delta)$. The result is

$$|||\xi|||_{1,p;D(x,r)} \leq c_4\Big(\,\Big|\Big|\Big|L_f\xi\Big|\Big|\Big|_{p;D(x,2r)} + |||\xi|||_{p;D(x,2r)}\Big) \tag{7.12}$$

for all $\xi \in \mathcal{E}_f$.

We next estimate $|L_f - L_g|$, first on a small disk, then on $C(2\delta)$. Fix $x \in C(2\delta)$. After making r smaller, we may assume that $f(D(x,2r))$ lies in a holomorphic coordinate chart $U \subset N_D$ with coordinates $\{x^i\}$. In this chart, J is constant. In light of (7.11) there is a neighborhood $\mathcal{N}_f(x)$ of f such that the images $g(D(x,2r))$ for all $g \in \mathcal{N}_f(x)$ lie in U. As in the proof of Lemma 7.1 we can then regard $g \in \mathcal{N}_f(x)$ and any $\xi \in \mathcal{E}_g$ as complex-valued functions on $D(x,2r)$. Since the J_α-holomorphic map f has its image in D, we can further assume that the images of g for all $g \in \mathcal{N}_f(x)$ lies in the $\frac{1}{2}$-neighborhood of the D in N_D.

As in (7.2), write L_g as the standard $\overline{\partial}$-operator on functions plus additional terms:

$$L_g\xi = \overline{\partial}_C\xi + \pi^{0,1}_{C,g}\left(g^*\Gamma\xi + \nabla_\xi K_\alpha dgj\right) \tag{7.13}$$

where $\overline{\partial}_C\xi = \frac{1}{2}(d\xi + Jd\xi j)$ and the projection onto g^*N-valued $(0,1)$ forms is defined in terms of the complex structure j on C by

$$\pi^{0,1}_{C,g}(\eta)(v) = \frac{1}{2}\Big(\eta(v) + J\eta(jv)\Big).$$

Let j_0 be the complex structure on C_0. Note that (i) we may assume $|j|$ is bounded on $C_0(\delta)$ by making $\mathcal{N}_f$ smaller since $|j_0| = 1$ on C_0 and (ii) the smooth quantities Γ, J and ∇K_α are bounded. Thus (7.3) and (7.13) give the pointwise inequality:

$$\big|(L_g - L_f)\xi\big| \le c_5\big(|f-g| + |df - dg| + |j - j_0|\big)\,\|\xi\|_\infty + |j_0 - j|\,|\nabla\xi| \tag{7.14}$$

where c_5 depends only on (C_0, f) and δ.

Now, fix $\varepsilon > 0$. After making $\mathcal{N}_f$ smaller, we may assume that

$$\big\|j_0 - j\big\|_{\infty;C_0(\delta)} + |||f-g|||_{1,p;C_0(\delta)} < \varepsilon$$

for all $g \in \mathcal{N}_f(x)$ (cf. Definition 2.2). Consequently, (7.14) and Lemma 5.3 give

$$\Big|\Big|\Big|(L_g - L_f)\xi\Big|\Big|\Big|_{p;D(x,2r)} \le c_7\varepsilon\, |||\xi|||_{1,p;D(x,2r)}. \tag{7.15}$$

Combining (7.12) and (7.15) then yields

$$|||\xi|||_{1,p;D(x,r)} \le c_8\Big(|||L_g\xi|||_{p;D(x,2r)} + |||\xi|||_{p;D(x,2r)} + \varepsilon\, |||\xi|||_{1,p;D(x,2r)}\Big) \tag{7.16}$$

for all $g \in \mathcal{N}_f(x)$. Finally, the compact curve $C(\delta)$ is covered by disks $D = D(x,2r)$ on which (7.16) holds. Let $2r_0$ be the Lebesgue number of this cover. Then (7.16) holds on each disk with center in $C(2\delta)$ and radius less than $2r_0$. Choose a finite cover of $C(2\delta)$ by such disks $D(x_i, 2r_0)$ so that each point of $C(\delta)$ lies in at most 20 disks. Summing the integrals in the estimate (7.16) and intersecting the corresponding neighborhoods $\mathcal{N}_f(x_i)$, one sees that (7.16) holds with $D(x,r)$ replaced by $C(2\delta)$ and $D(x,2r)$ replaced by $C(\delta)$, as required. $\square$

7.3. Global Poincaré inequalities. We can now combine the interior and exterior estimates to obtain a global Poincaré inequality. We do this first for a fixed J_α-holomorphic map f, and then use the Compactness Theorem 2.5 to obtain a Poincaré equality valid for all maps in a neighborhood of the moduli space.

PROPOSITION 7.5. *Fix $p > 2$ and a J_α-holomorphic map (C, f). Then there is a neighborhood $\mathcal{N}_f$ of (C, f) in $\mathcal{M}ap_\lambda(N_D)$ and a constant $c = c(p, f)$ such that the following holds: for each $(C', g) \in \mathcal{N}_f$, every section ξ of g^*N with (i) finite (unweighted) $L^{1,p}$ norm and (ii) $L_g\xi \in \mathcal{F}_g$, satisfies $\xi \in \mathcal{E}_g$ and*

$$\|\|\xi\|\|_{1,p} \leq c\,\|\|L_g\xi\|\|_p. \tag{7.17}$$

Proof. If this statement is false, there is a sequence of maps $(C_n, g_n) \to (C, f)$ and $\xi_n \in L^{1,p}(g_n^*N)$ with $\|\|\xi_n\|\|_{1,p} = 1$ for all n and $\|\|L_{g_n}\xi_n\|\|_p \to 0$. Fix $\delta > 0$. As in the proof of Lemma 7.10, we can identify each $C_n(\delta)$ with $C(\delta)$. Under this identification the complex structures on $C_n(\delta)$ converge in C^1 to that of $C(\delta)$ as $n \to \infty$ and the norms (5.2) and (5.3) are uniformly equivalent to the standard Sobolev $L^{1,p} \cap L^{1,2}$ and $L^p \cap L^2$ norms.

By the compactness of the Sobolev embedding $L^{1,p} \subset C^0$ there is a subsequence, still denoted $\{\xi_n\}$, that converges weakly in $L^{1,p}$ and strongly in C^0 to some $\xi_0 \in \mathcal{E}_f\big|_{C(\delta)}$. Hence $L_f\xi_0 = 0$. After replacing δ by $\delta/2$, then $\delta/4$, etc., repeatedly taking subsequences, and passing to a diagonal subsequence we obtain a continuous solution of $L_f\xi_0 = 0$ on $C \setminus \{\text{nodes}\}$, and this ξ_0 is bounded because the numbers $\|\xi_n\|_\infty$ are uniformly bounded by our hypothesis and Lemma 5.3. Lemma 7.6 below then shows that ξ_0 extends across the nodes to a smooth global solution of $L_f\xi_0 = 0$ on C_0. But $\ker L_f = 0$ by Vanishing Theorem 4.2, so in fact $\xi_0 = 0$. Thus $\xi_n \to 0$ in C^0 and consequently for each $\delta > 0$ we have

$$\|\|\xi_n\|\|_{p;C(\delta)} \longrightarrow 0. \tag{7.18}$$

On the other hand, for a fixed small δ and all large n, the maps g_n lie in the neighborhoods $\mathcal{N}_f$ of Lemmas 7.1 and 7.4. Applying the interior estimate (7.1) on $B(2\delta)$, we have

$$\begin{aligned} 1 = \|\|\xi_n\|\|_{1,p} &\leq \|\|\xi_n\|\|_{1,p;B(2\delta)} + \|\|\xi_n\|\|_{1,p;C(2\delta)} \\ &\leq c_1\,\|\|L_{g_n}\xi_n\|\|_{p;B(8\delta)} + c_2\,\|\|\xi_n\|\|_{1,p;B(8\delta)\setminus B(2\delta)} + \|\|\xi_n\|\|_{1,p;C(2\delta)} \\ &\leq c_1\,\|\|L_{g_n}\xi_n\|\|_p + c_3\,\|\|\xi_n\|\|_{1,p;C(2\delta)}. \end{aligned}$$

After bounding the last term by the exterior estimate (7.10) with $\varepsilon = 1/2c_3$, this becomes

$$1 \leq c_4\,\|\|L_{g_n}\xi_n\|\|_p + c_5\,\|\|\xi_n\|\|_{p;C(\delta)} + \tfrac{1}{2}\,\|\|\xi_n\|\|_{1,p;C(\delta)}.$$

This inequality, together with (7.18) and the normalization $|||\xi_n|||_{1,p} = 1$, contradicts the assumption that $|||L_{g_n}\xi_n|||_p \to 0$. $\square$

LEMMA 7.6. *Let $f : D \to X$ be a smooth map from a disk D with Riemannian metric g. Then any bounded weak solution of $L_f\xi = 0$ on $D \setminus \{0\}$ extends to a smooth solution of $L_f\xi = 0$ on D.*

Proof. Fix a smooth non-increasing function $\beta(t)$ of $t \in \mathbb{R}$ with $\beta = 1$ for $t \le \frac{1}{2}$ and $\beta = 0$ for $t \ge 1$. For $\delta > 0$, set $\beta_\delta = \beta(r/\delta)$ using polar coordinates on D. Write L_f as L and ξ as $\beta_\delta\xi + (1-\beta_\delta)\xi$. Noting that $d\beta$ has support on the disk $D(\delta)$ and that $|L((1-\beta_\delta)\xi)| \le |d\beta_\delta||\xi| \le c\delta^{-1}|\xi|$, we have the following bounds on the L^2 inner product:

$$
\begin{aligned}
|\langle L^*L\eta, \xi\rangle| &= |\langle L^*L\eta, \beta_\delta\xi\rangle| + |\langle L\eta, L(1-\beta_\delta)\xi\rangle| \\
&\le \left(\|L^*L\eta\|_{2;D} + c\delta^{-1}\|L\eta\|_{2;D(\delta)} \right) \|\xi\|_\infty \sqrt{\mathrm{vol}(D(\delta))}.
\end{aligned}
$$

The right-hand side vanishes as $\delta \to 0$. Thus ξ is a bounded weak solution of $L^*L\xi = 0$ on D.

After a conformal change of metric, we can choose coordinates around the origin in which L has the form $\overline{\partial} + A$ where A is a zeroth-order operator and $L^*L = \Delta + B$ where B is a first order operator. Standard elliptic theory then implies that ξ extends across the origin to a smooth solution of $L^*L\xi = 0$ on D. Taking the inner product with $\beta_\delta\xi$ and integrating by parts then shows that $L\xi = 0$ on D. $\square$

THEOREM 7.7. *For each p with $2 < p \le 2+\lambda$, there is a constant $c(p)$ and a neighborhood $\mathcal{N}$ of the space of stable J_α-holomorphic maps in $\mathcal{M}ap_\lambda(N_D, d)$ such that, for every $f \in \mathcal{N}$,*

$$
L_f : \mathcal{E}_f \to \mathcal{F}_f \tag{7.19}
$$

is a uniformly bounded Fredholm map with $\operatorname{index} L_f = -2\beta$ *as in (3.10) and with*

$$
|||\xi|||_{1,p} \le c(p)\, \left|\left|\left| L_f\xi \right|\right|\right|_p \quad \forall \xi \in \mathcal{E}_f. \tag{7.20}
$$

Proof. By Proposition 7.5, each stable J_α-holomorphic map f has a neighborhood $\mathcal{N}_f$ and an associated constant $c(p, f)$ so that (7.17) holds for all $g \in N_f$. These sets $\{\mathcal{N}_f\}$ cover the space of stable maps, so (7.20) follows by the Compactness Theorem 2.5.

We know from Proposition 5.6 that L_f is uniformly bounded for $f \in \mathcal{N}$, and $\ker L_f = 0$ by (7.20). Inequality (7.20) also implies that the range of L_f is closed: if $L_f\xi_k \to \eta$ then applying (7.20) to $\xi_k - \xi_\ell$ shows that $\{\xi_k\}$ is Cauchy, so $\xi_k \to \xi_0$ in $\mathcal{E}_f$ with $L_f\xi_0 = \lim L_f\xi_k = \eta$, so $\eta \in \operatorname{Im} L_f$. The proof is completed by noting that $\dim\operatorname{coker} L_f = 2\beta$ by Lemma 7.8 below. $\square$

LEMMA 7.8. *For each p with $2 < p \leq 2+\lambda$ and each (C,f) in the neighborhood $\mathcal{N}$ of Theorem 7.7, we can choose a 2β-dimensional subspace $W_f \subset \mathcal{F}_f$, consisting of smooth forms that vanish in a neighborhood of the nodes, that is complementary to the image of $L_f : \mathcal{E}_f \to \mathcal{F}_f$. Hence* $\dim \operatorname{coker} L_f = 2\beta$.

Proof. Fix $f \in \mathcal{N}$ and let $L^{1,p}(E_f)$ and $L^p(F_f)$ be the completions of E_f and F_f defined by (4.4) in the usual, unweighted $L^{1,p}$ and L^p norms. By [FO, Lemma 12.2] (see also [MS, Theorem C.1.10]) the modified linearization (4.5) extends to a bounded Fredholm map

$$L_f : L^{1,p}(E_f) \longrightarrow L^p(F_f)$$

whose index is -2β and whose kernel vanishes by Theorem 4.2. Hence we can choose linearly independent $(0,1)$ forms $\{w_j \mid j = 1,\ldots,2\beta\}$ in $L^p(F_f)$ so that $W_f = \operatorname{span}\{w_j\}$ is complementary to $\operatorname{Im} L_f$. We can assume that each w_j is smooth and vanishes in a neighborhood of all nodes because such forms are dense in L^p. Then each $\eta \in L^p(F_f)$ can be uniquely written as

$$\eta = L_f \xi + w \tag{7.21}$$

for some $\xi \in L^{1,p}(E_f)$ and $w \in W_f$. Also observe that, because ρ is bounded and $p \leq 2+\lambda$, each $\eta \in \mathcal{F}_f$ satisfies

$$\int_C |\eta|^p \leq c \int_C \rho^{p-2-\lambda}\, |\eta|^p.$$

Thus $\mathcal{F}_f \subset L^p(F_f)$. The same inequality for the weighted norm (5.2) shows that $\mathcal{E}_f \subset L^{1,p}(E_f)$. We also have $W_f \subset \mathcal{F}_f$ because ρ^{-1} is bounded outside the set where all $w \in W_f$ vanish.

The lemma follows easily: each $\eta \in \mathcal{F}_f$ lies in $L^p(F_f)$, so can be written in the form (7.21) for some unique $\xi \in L^{1,p}(E_f)$ and $w \in W_f$. We can then apply Proposition 7.5 to $L_f \xi = \eta - w \in \mathcal{F}_f$ to conclude that $\xi \in \mathcal{E}_f$. Thus each $\eta \in \mathcal{F}_f$ can be written as the sum of an element in $L_f \mathcal{E}_f$ and a $w \in W_f$, and uniqueness holds because $\mathcal{E}_f \subset L^{1,p}(E_f)$. We conclude that $\dim \operatorname{coker} L_f = \dim W_f = 2\beta$. □

8. Obstruction bundle $\mathcal{O}b$. Abstractly, $\mathcal{O}b$ is the topological vector bundle whose fiber at a J_α-holomorphic map $f : C \to N_D$ is $\operatorname{coker} L_f$. Thus it is defined by the exact sequence

$$0 \longrightarrow \mathcal{E} \xrightarrow{L} \mathcal{F} \xrightarrow{\rho} \mathcal{O}b \longrightarrow 0$$

of topological vector bundles over $\mathcal{M}ap_\lambda(X)$. The goal of this section is to give a concrete realization of $\mathcal{O}b$ and show that it is a locally trivial bundle on the space of maps. For this, we will split the sequence

$$0 \longrightarrow \mathcal{E}_f \xrightarrow{L_f} \mathcal{F}_f \xrightarrow{\rho} \operatorname{coker} L_f \longrightarrow 0 \tag{8.1}$$

by regarding $\operatorname{coker} L_f$ as a subspace of $\mathcal{F}_f$ for each map f. The key issue is constructing a splitting that is locally trivial around maps with nodal domains.

Consider the space $\mathcal{M}ap_{g,n}(N_D,d)$ of maps into the open complex surface N_D that represent the class $d[D]$. The following theorem is the main application of the analysis done in previous sections.

THEOREM 8.1. *There is a locally trivial real vector bundle $\mathcal{O}b$ defined in a λ-topology neighborhood $\mathcal{N}$ of $\overline{\mathcal{M}}_{g,n}(D,d)$ in $\mathcal{M}ap_{g,n}(N_D,d)$ such that the fiber of $\mathcal{O}b$ at each $f \in \mathcal{N}$ lies in $\mathcal{F}_f$, is complementary to the image of L_f, and has the same dimension as* $\operatorname{coker} L_f$. *This bundle $\mathcal{O}b$ has a canonical orientation.*

The remainder of this section proves Theorem 8.1 in four steps. The first step adopts the method of Fukaya and Ono [FO]. Steps 2 and 3 show that the cokernel bundle is a locally trivial vector bundle; this is the key feature of our situation that does not occur in [FO]. Step 4 constructs an embedding of the cokernel bundle into the bundle $\mathcal{F}$. *In these steps and in the subsequent sections we assume that p lies in the range $2 < p \le 2+\lambda$ in which Theorem 7.7 and Lemma 7.8 are valid.*

Step 1. The fiber $\mathcal{O}b_f$ at a map $f : C \to N_D$. For a fixed $f \in \overline{\mathcal{M}}_{g,n}(D,d)$ the subspace W_f of Lemma 7.8 specifies a (non-canonical) splitting of (8.1). As in Lemma 7.8, write W_f as the space of linearly independent $(0,1)$ forms $\{w_j \mid j = 1,\dots,2\beta\}$ in $\mathcal{F}_f$. Also consider the subspace

$$\mathcal{P}_0 = \Omega_0^{0,1}(T^*\overline{\mathcal{U}} \boxtimes N) \tag{8.2}$$

of the space (1.9) of Ruan-Tian perturbations that have values in the subspace $N \subset TN_D$ and which vanish in some open set containing all nodes. For each $v \in \mathcal{P}_0$ the restriction of v to the graph of f, defined as in (1.10), is an element $v_f \in \Omega_0^{0,1}(C, f^*N)$ that vanishes near all nodes. The following lemma shows that, conversely, all such elements of $\Omega_0^{0,1}(C, f^*N)$ arise as restrictions of elements of $\mathcal{P}_0$.

LEMMA 8.2. *For each $w \in \Omega_0^{0,1}(C, f^*N)$, there is a $v \in \mathcal{P}_0$ with $v_f = w$.*

Proof. Fix δ small enough that the support of w lies in $C(2\delta)$. The graph $G_f = \{(x, f(x)) \,|\, x \in C(\delta)\}$ is a submanifold of $C(\delta) \times N_D$, so has a 2ε-tubular neighborhood $\mathcal{O}$ consisting of points (x,y) such that there is a unique minimal geodesic in N_D from y to $f(x)$. Parallel transporting the fibers of N along these geodesics trivializes N over $\mathcal{O}$. Fix a cutoff function $\beta \in C^\infty(C(\delta) \times N_D)$ with support on $\mathcal{O}$ and with $\beta \equiv 1$ on the ε-tubular neighborhood of G_f. Then, regarding w as a form on G_f, we can extend to $\mathcal{O}$ by parallel transport, multiply by β, and extend by zero to obtain a section βw of $T^*C(\delta) \boxtimes N$ over $C(\delta) \times N_D$.

Now fix a local trivialization $\phi : C(\delta) \times V \to U$ of the universal curve $\overline{\mathcal{U}}$ around C as in (2.4) and write $\phi^{-1} = (\tau_1, \tau_2)$. Choose a smooth cutoff function β_U on $\overline{\mathcal{U}}$ with support on $\pi^{-1}(U)$ and with $\beta_U \equiv 1$ on a smaller neighborhood $\pi^{-1}(U')$ of C. The section $\tilde{v} = \beta_U \tau_1^*(\beta w)$ then extends by zero to a section of $T^*\overline{\mathcal{U}} \boxtimes N$ over $\overline{\mathcal{U}} \times N_D$. The Jj-antilinear component $v = \frac{1}{2}(\tilde{v} + J \circ \tilde{v} \circ j)$ is the desired $(0,1)$ form. $\square$

Step 2. Charts for the cokernel bundle. Fix $(C, f) \in \overline{\mathcal{M}}_{g,n}(D, d)$ and $W_f = \text{span}\{w_j\}$ as in Step 1. In light of Lemma 8.2, we can extend the forms in W_f to obtain a vector space

$$V_f = \text{span}\{v_j\} \subset \mathcal{P}_0.$$

Given a neighborhood $\mathcal{N}_f$ of (C, f) in the space of maps, we can regard V_f as a trivial vector bundle $V_f \times \mathcal{N}_f$ over $\mathcal{N}_f$. For each map $g \in \mathcal{N}_f$ the composition of the restriction (1.10) and the projection ρ in (8.1) gives a linear map R defined by $R(v, g) = \rho(v_g)$ and a diagram

$$\begin{tikzcd} V_f \times \mathcal{N}_f \arrow[rr, "R"] \arrow[dr] & & \operatorname{coker} L \arrow[dl] \\ & \mathcal{N}_f & \end{tikzcd} \tag{8.3}$$

where $\operatorname{coker} L$ is the topological vector bundle whose fiber at g is $\operatorname{coker} L_g$.

LEMMA 8.3. *Each J_α-holomorphic map f has a neighborhood $\mathcal{N}_f$ in $\mathcal{M}ap_{g,n}(N_D, d)$ for which the map R in (8.3) is injective.*

Proof. By Theorem 7.7 there is a neighborhood $\mathcal{N}$ of $\overline{\mathcal{M}}_{g,n}(D, d)$ such that for every $g \in \mathcal{N}$ we have $\ker L_g = 0$ and $\dim \operatorname{coker} L_g = 2\beta$. In particular, for each $g \in \mathcal{N}$, $\operatorname{coker} L_g$ has the same dimension as V_f. It suffices to find a neighborhood of f on which R is injective on fibers.

Let S_f denote the unit sphere in the finite-dimensional vector space V_f defined by the condition $\|v\|_\infty = 1$. Because S_f is compact and $\operatorname{im} L_g$ is closed, the infimum

$$\varepsilon_g = \inf\left\{ |||L_g\xi - v_g|||_p : \xi \in \mathcal{E}_g \text{ and } v \in S_f \right\} \tag{8.4}$$

is realized for each g and is equal to 0 if and only if $R(\cdot, g)$ is not injective. Thus, by our choice V_f, we have $\varepsilon_f > 0$, and we must show that $\varepsilon_g > 0$ for all g near f.

Suppose not. Then there is a sequence (C_n, g_n) converging to (C, f) in the λ-topology and corresponding sections $\xi_n \in \mathcal{E}_{g_n}$ and $v_n \in S_f$ satisfying

$$L_{g_n}\xi_n = \big(v_n\big)_{g_n}. \tag{8.5}$$

Fix any $\delta > 0$ small enough that the support of each $(v_n)_{g_n}$ is contained in $C(\delta)$ and, as in the proof of Lemma 7.4, identify the domains $C_n(\delta)$ with $C(\delta)$ for all large n. We then have the following three facts.

(a) Convergence in the λ-topology implies C^0 convergence. Hence, after identifying the domains by a trivialization (2.4), $|||v_{g_n} - v_g|||_p \to 0$ because the weight ρ in the norm (5.3) is bounded on the union of the supports of the v_{g_n}. The triangle inequality then shows that $|||(v_n)_{g_n} - v_g|||_p \to 0$.

(b) Lemma 5.3, Theorem 7.7, and the compactness of S_f yield a sup bound on ξ_n:

$$\|\xi_n\|_\infty \;\le\; c_1\,|||\xi_n|||_{1,p} \le c_2\,|||L_{g_n}\xi_n|||_p = c_2\,|||(v_n)_{g_n}|||_p < c_3 \tag{8.6}$$

where the constants c_1, c_2 and c_3 are independent of g_n and these norms are on all of C_n.

(c) Covering $C(\delta)$ by disks on which (7.15) holds shows that there is a c_4 such that for each $\varepsilon > 0$ there is an N so that

$$\Big|\Big|\Big|L_f\xi_n - (v_n)_{g_n}\Big|\Big|\Big|_{p;C(\delta)} = \Big|\Big|\Big|(L_f - L_{g_n})\xi_n\Big|\Big|\Big|_{p;C(\delta)} \le c_4\varepsilon\;|||\xi_n|||_{1,p} \quad \forall n \ge N.$$

Using Theorem 7.7 and writing $L_f\xi_n = (L_f\xi_n - (v_n)_{g_n}) + (v_n)_{g_n}$, we see that

$$|||\xi_n - \xi_m|||_{1,p;C(\delta)} \;\le\; c_5\,\Big|\Big|\Big|L_f(\xi_n - \xi_m)\Big|\Big|\Big|_{p;C(\delta)}$$

is bounded by

$$c_5\bigg(\Big|\Big|\Big|L_f\xi_n - (v_n)_{g_n}\Big|\Big|\Big|_{p;C(\delta)} + \Big|\Big|\Big|L_f\xi_m - (v_m)_{g_m}\Big|\Big|\Big|_{p;C(\delta)} \\ + |||(v_n)_{g_n} - (v_m)_{g_m}|||_{p;C(\delta)}\bigg)$$

for all large m and n. This, together with facts (a)–(c), shows that $\{\xi_n\}$ is Cauchy in the norm (5.2) on $C(\delta)$ and converges on $C(\delta)$ to a limit ξ that satisfies $L_f\xi = v_f$.

Now replace δ by $\delta/2$ and repeat to obtain a further subsequence that converges on $C(\delta/2)$ to a limit that extends the previous ξ. Repeat with $\delta/4$, etc. Then the diagonal sequence converges on $C_f \setminus \{\text{nodes}\}$ to a limit ξ that satisfies $L_f\xi = v_f$. Because $v_f \equiv 0$ near each node, Lemma 7.6 shows that ξ extends to a solution of $L_f\xi = v_f$ on all of C. Then Proposition 7.5 shows that $|||\xi|||_{1,p}$ is finite, and Corollary 5.5 shows that ξ satisfies the matching condition (4.4) at the nodes. Thus $\xi \in \mathcal{E}_f$, contradicting the assumption that $\varepsilon_f > 0$. $\square$

Step 3. The transition maps for the cokernel bundle. In order to show $\operatorname{coker} L$ is a vector bundle, it suffices to show that for any two maps f_1 and f_2 with $\mathcal{N}_{f_1} \cap \mathcal{N}_{f_2} \neq \emptyset$ and for each fixed $v \in V_{f_1}$ the map

$$\mathcal{N}_{f_1} \cap \mathcal{N}_{f_2} \longrightarrow V_{f_2} \text{ defined by } g \longrightarrow R_2^{-1} \circ R_1(v,g) \tag{8.7}$$

is continuous where $R_i : V_{f_i} \times \mathcal{N}_{f_i} \to \operatorname{coker} L$ is a linear map defined in (8.3) and $\mathcal{N}_{f_1}$ and $\mathcal{N}_{f_2}$ are neighborhoods as in Lemma 8.3.

LEMMA 8.4. *The map (8.7) is continuous. Consequently,* $\operatorname{coker} L$ *is a locally trivial vector bundle.*

Proof. With the above notation, for each $g \in \mathcal{N}_{f_1} \cap \mathcal{N}_{f_2}$, Lemma 8.3 shows that $\{u_g | u \in V_{f_2}\}$ is complementary to the image of L_g. Hence, as in the proof of Lemma 7.8, for each v_g there exist unique $\xi \in \mathcal{E}_g$ and $u \in V_{f_2}$ with $v_g = L_g \xi + u_g$. Therefore

$$L_g \xi = v_g - u_g \quad \text{and} \quad R_2^{-1} \circ R_1(v,g) = u. \tag{8.8}$$

We claim that u depends continuously on g. To that end, consider sequences $g_n \to g$, $u_n \in V_{f_2}$ and ξ_n such that

$$L_{g_n}\big(\xi_n\big) = v_{g_n} - \big(u_n\big)_{g_n}. \tag{8.9}$$

First consider the case where there is a uniform bound on the norms $\|u_n\|_\infty$. Because V_{f_2} is finite-dimensional, we can assume, after passing to a subsequence, that u_n converges to an element $u_0 \in V_{f_2}$ in the C^0 norm. Then (8.9) is the same as (8.5) with v and v_n replaced by $z = v - u_0$ and $z_n = v - u_n$ respectively. As in statement (a) in the proof of Lemma 8.3, the convergence $g_n \to g$ implies that $|||(z_n)_{g_n} - z_g|||_p \to 0$. We can therefore repeat the argument of Lemma 8.3 to conclude that $\{\xi_n\}$ converge on compact sets away from the nodes of C to a limit ξ_0 that satisfies $L_g(\xi_0) = v_g - (u_0)_g$. The uniqueness of u in (8.8) then implies $u = u_0$. As a result, the sequence u_n—not just a subsequence—converges to u. Thus u depends continuously on g.

In the remaining case, u_n is not bounded. After passing to a subsequence, we can assume that $\|u_n\|_\infty > n$ for all n. Note that $|||(u_n)_{g_n}|||_p \leq c\|u_n\|_\infty$ for some uniform constant c. Now, dividing equation (8.9) by $\|u_n\|_\infty$, we have

$$L_{g_n}\big(\eta_n\big) = \alpha_{g_n} - \big(\beta_n\big)_{g_n}$$

with $|||\alpha_{g_n}|||_p \leq \frac{1}{n} \limsup |||v_{g_n}|||_p \to 0$ and with $\|\beta_n\|_\infty = 1$ and $|||(\beta_n)_{g_n}|||_p \leq c$ for all n. Again repeating the convergence argument of Lemma 8.3 we obtain, in the limit, a ζ that satisfies $L_g \zeta = -\beta_g$ where $\beta \neq 0$ is the limit of the sequence β_n in V_{f_2}. But this means that $R_2(\beta, g) = 0$, contradicting Lemma 8.3. Therefore, the map (8.7) is continuous. The lemma follows. $\square$

Step 4. Definition of $\mathcal{O}b$. The procedure of Step 2 can be applied to each J_α-holomorphic map f to obtain a pair $(V_f, \mathcal{N}_f)$. By compactness, we can choose a finite set $\{f_i\}$ of J_α-holomorphic maps so that $\{\mathcal{N}_{f_i}\}$ covers $\overline{\mathcal{M}}_{g,n}(D,d)$. The vector spaces V_{f_i} consist of $(0,1)$ forms and together define a single finite-dimensional vector space span$\{V_{f_i}\}$. Let $\mathbf{V}$ denote the trivial vector bundle$\mathbf{V} = \mathcal{N} \times \text{span}\{V_{f_i}\}$ over $\mathcal{N} = \cup \mathcal{N}_{f_i}$. We then have an exact sequence of locally trivial vector bundles over $\mathcal{N}$:

$$0 \longrightarrow \ker\sigma \longrightarrow \mathbf{V} \stackrel{\sigma}{\longrightarrow} \operatorname{coker} L \longrightarrow 0$$

where σ is the composition of the restriction map (1.10) with the projection to coker L. Now, using a metric on the vector bundle $\mathbf{V}$ induced from metrics on $\overline{\mathcal{U}}$ and N_D, we define the obstruction bundle $\mathcal{O}b$ over $\mathcal{N}$ to be the orthogonal complement of $\ker\sigma$ in $\mathbf{V}$. Then σ restricts to a vector bundle isomorphism $\mathcal{O}b \stackrel{\cong}{\to} \operatorname{coker} L$ over $\mathcal{N}$. Lastly, the fact $\mathcal{O}b$ has a canonical orientation is proved in Lemma 11.1 below. This completes the proof of Theorem 8.1.

9. A generalized Image Localization Lemma. In this section we generalize the Image Localization Lemma of Section 1, showing that it applies not just to solutions of (1.4) but, more generally, to maps $f : C \to N_D$ that satisfy the perturbed J_α-holomorphic map equation (1.8) provided that ν is small and its vertical component lies in the fiber of the obstruction bundle. The proof uses a renormalization argument similar to those in [IP1, Sections 6 and 7]. In the statement of the theorem, and throughout this section, we will use the decomposition (4.1) to write the perturbation ν as the sum $\nu^T + \nu^N$ of horizontal and vertical components. The conclusion is the same as in the original localization lemma: the images of the maps lie in the divisor of α, which is the zero section of the bundle $N \to D$.

THEOREM 9.1. (Image Localization with Perturbations) *Fix $E > 0$ and a neighborhood U of the zero section in N_D. Then there is a $\delta_0 > 0$ such that if $\nu = \nu^T + \nu^N$ with $\|\nu\|_{C^1} \le \delta_0$, then the image of every map $f : C \to N_D$ in $\mathcal{M}ap^E_\lambda(U)$ satisfying*

$$\overline{\partial}_J f - K_\alpha\big(\partial_J f\big) j = \nu^T_f + P_f\big(\nu^N_f\big) \tag{9.1}$$

lies in D where P_f is the L^2-orthogonal projection onto the fiber $\mathcal{O}b_f$.

Proof. If this statement is false, there is a sequence $\{\nu_n\}$ of the stated form with $\nu_n \to 0$ in C^1, and a sequence of maps $f_n : C_n \to N_D$ in $\mathcal{M}ap^E_\lambda(U)$ satisfying the equation (9.1), each with at least one point not mapped into D. By Gromov compactness and Theorem 2.5 we may assume, after passing to a subsequence, that the f_n converge, in $C^0 \cap L^{1,2}$ and in C^∞ away from the nodes, to a J_α-holomorphic map $f_0 : C_0 \to N_D$. By Lemma 1.1 the image of f_0 lies in $D \subset N_D$.

We can now renormalize. Let ϕ_n be the projection of f_n into D and let $R_n : N_D \to N_D$ be dilation in the fibers by a factor of $1/t_n$, where each t_n is chosen so that the image of renormalized map $F_n = R_n \circ f_n$ lies in the unit disk bundle, but not in any smaller disk bundle. Because the images of the f_n converge pointwise to the zero section we have $t_n \to 0$.

Next, write F_n as the graph (ϕ_n, ξ_n) in N_D where ξ_n is a section of $\phi_n^* N$. Then the original maps are given by $f_n = (\phi_n, t_n \xi_n)$ and $\phi_n \to f_0$ as $n \to \infty$. For each map f, write $\Phi(f) = \overline{\partial}_J f - K_\alpha \partial f j$ for the left-hand side of (9.1). After trivializing the pullback $\pi^* N$ by parallel transport T along the fibers of $\pi : N_D \to D$, we can compare the vertical components Φ^N of $\Phi(\phi_n)$ and $\Phi(f_n)$ by writing $f_n = \exp_{\phi_n}(t_n \xi_n)$ and applying the first variation formula on C_n:

$$T\Phi^N(f_n) = \Phi^N(\phi_n) + L_{\phi_n}(t_n \xi_n) + O\Big(|t_n \xi_n|^2\Big). \tag{9.2}$$

Here L_{ϕ_n} is given by equation (4.3) along the image of ϕ_n and $\Phi^N(\phi_n) = 0$ since the image of ϕ_n lies in D. We also have

$$TP_{f_n}(\nu_{f_n}^N) = P_{\phi_n}(\nu_{\phi_n}^N) + \nabla_{t_n \xi_n}(P_{\phi_n}(\nu_{\phi_n}^N)) + O\Big(|t_n \xi_n|^2\Big) \tag{9.3}$$

where we have simplified notation by writing $(\nu_n)_{\phi_n}$ as ν_{ϕ_n}. Since $\Phi^N(f_n) = P_{f_n}(\nu_{f_n}^N)$ by equation (9.1) and $\|\xi_n\|_\infty = 1$, it follows from (9.2) and (9.3) that

$$t_n^{-1} P_{\phi_n}(\nu_{\phi_n}^N) - L_{\phi_n}(\xi_n) = -(\nabla_{\xi_n} P_{\phi_n})(\nu_{\phi_n}^N) - P_{\phi_n}(\nabla_{\xi_n} \nu_{\phi_n}^N) + O(t_n). \tag{9.4}$$

By definition, P is the L^2 projection onto a subspace spanned by smooth forms with support away from the nodes, where $\phi_n \to f_0$ in C^∞. Writing that projection as an integral, one sees that P_{ϕ_n} and ∇P_{ϕ_n} are bounded, uniformly in n. The assumption that $\nu_n \to 0$ in C^1 then implies that the right-hand side of (9.4) vanishes in the limit $n \to \infty$

Recall that $\mathcal{O}b_{\phi_n}$ is a complementary subspace to the image of L_{ϕ_n}. In fact, the proof of Lemma 8.3 shows that the quantity ε_g defined by (8.4), which measures the angle between these subspaces, satisfies $\varepsilon_{\phi_n} > c$ for some positive constant c and all sufficiently large n. We therefore conclude that

$$\left|t_n^{-1} P_{\phi_n}(\nu_{\phi_n}^N)\right| \longrightarrow 0 \quad \text{and} \quad L_{\phi_n}(\xi_n) \longrightarrow 0$$

as $n \to \infty$. Dividing (9.3) by t_n and using these limits (noting that parallel transport preserves norms) yields the stronger statement that

$$\left|t_n^{-1} P_{f_n}(\nu_{f_n}^N)\right| \longrightarrow 0 \quad \text{and} \quad L_{\phi_n}(\xi_n) \longrightarrow 0 \quad \text{as } n \longrightarrow \infty. \tag{9.5}$$

Now return to the equation $F_n = R_n \circ f_n$. Note that $dF_n = (R_n)_* df_n$ and, since R_n is a holomorphic diffeomorphism of the total space of N_D, J commutes with $(R_n)_*$. Letting K_n denote the pullback endomorphism $(R_n^{-1})^* K_\alpha =$

$(R_n)_* K_\alpha (R_n^{-1})_*$, equation (9.1) can be rewritten as

$$\overline{\partial}_J F_n - K_n \partial_J F_n j = (R_n)_* \left[\nu^T_{f_n} + P_{f_n}(\nu^N_{f_n})\right]. \tag{9.6}$$

Consider the limit as $n \to \infty$ of each term in (9.6). Under the splitting (4.1), $(R_n)_*$ is bounded on horizontal vectors and is multiplication by $1/t_n$ on vertical vectors, so by (9.5)

$$\left|(R_n)_* \left[\nu^T_{f_n} + P_f(\nu^N_{f_n})\right]\right| \leq c|\nu^T_{f_n}| + \left|t_n^{-1} P_{f_n}(\nu^N_{f_n})\right| \longrightarrow 0. \tag{9.7}$$

Because K_α is smooth and vanishes along the zero section, there is an endomorphism K' depending smoothly on the projection of x into D such that $K_\alpha(tx) = tK'_\alpha + O(t^2)$ as $t \to 0$. Hence the pullback K_n satisfies

$$\begin{aligned} K_n &= (R_n)_* \left(t_n K'_\alpha + O(t_n^2)\right)(R_n^{-1})_* \\ &= t_n \begin{pmatrix} a & 0 \\ b & \dfrac{1}{t_n} \end{pmatrix} \begin{pmatrix} k'_{hh} & k'_{vh} \\ k'_{hv} & k'_{vv} \end{pmatrix} \begin{pmatrix} a^{-1} & 0 \\ -t_n b a^{-1} & t_n \end{pmatrix} + O(t_n) \\ &= \begin{pmatrix} 0 & 0 \\ k'_{hv} a^{-1} & 0 \end{pmatrix} + O(t_n) \end{aligned}$$

as $t_n \to 0$. Thus $\{K_n\}$ converges to an endomorphism $\widehat{K}$ whose only non-zero components take horizontal vectors to vertical vectors.

We can now reexpress (9.6) as a J_n-holomorphic map equation. As in (1.2), K_n anti-commutes with J and satisfies $K_n^2 = -|\alpha_n|^2 Id$ where α_n is the $(0,2)$ form whose value at x is $\alpha_{t_n x}$. It follows that $Id + JK_n$ has a bounded inverse, namely $(1+|\alpha_n|^2)^{-1}(Id - JK_n)$, that $J_n = (Id + JK_n)^{-1} J (Id + JK_n)$ is an almost complex structure, and that F_n satisfies the equation

$$\overline{\partial}_{J_n} F_n = \frac{1}{2}\left(dF_n + J_n dF_n j\right) = \left(Id + JK_n\right)^{-1} (R_n)_* \left(\nu^T_{f_n} + P_{f_n}\nu^N_{f_n}\right). \tag{9.8}$$

The convergence $K_n \to \widehat{K}$ then implies that $\{J_n\}$ converges in C^0 to the almost complex structure $\widehat{J} = (Id + J\widehat{K})^{-1} J (Id + J\widehat{K})$, and (9.7) implies that the right-hand side of (9.8) converges to 0 in C^0.

At this point we have established that the maps $F_n : C_n \to N_D$ satisfy the J_n-holomorphic equation with a perturbation term that goes to 0 in C^0. By Gromov compactness (cf. [IP1, Theorem 1.6]) there is a connected curve C_0 and smooth map $F_0 : C_0 \to N_D$ such that, after passing to a subsequence, F_n converges to F_0 in $C^0 \cap L^{1,2}$ and in C^∞ away from the nodes of C_0. It follows that $\pi \circ F_0 = f_0$ and hence F_0 defines a section ξ_0 of the pullback bundle $f_0^* N$ by the formula $F_0 = (f_0, \xi_0)$.

If C_0 is smooth, we have $\phi_n \to f_0$ and $\xi_n \to \xi_0$ in C^∞. Then (9.5) and the normalization $\|\xi_n\|_\infty = 1$ imply that ξ_0 is a non-trivial solution of

$$L_{f_0}(\xi_0) = 0.$$

This contradicts the Vanishing Theorem 4.2.

When C_0 is nodal we get a similar contradiction on some irreducible component of C_0, as follows. Note that, because F_0 is $\widehat{J}$-holomorphic, the image $F_0(C_i)$ of each irreducible component C_i is either a single point or represents a non-trivial homology class. If all components that carry a non-trivial homology class were mapped into D then, since C_0 is connected and F_0 is continuous, the entire image would lie in D, contradicting our choice of renormalization. Thus there exists a component C of C_0 with $F_0(C) \nsubseteq D$. On this component, ξ_0 is not zero. The argument in the proof of Proposition 7.5 then shows that $\xi_n \to \xi_0$ in C^∞ on each compact subset of $C \setminus \{\text{nodes}\}$ and that $L_{f_0}\xi_0 = 0$. This again contradicts Theorem 4.2. $\square$

10. The proof of the Main Theorem. We now turn to the proof of the Main Theorem stated in the introduction: the local GW invariants of a spin curve (D,N) are given by the cap product

$$GW^{\text{loc}}_{g,n}(N_D,d) = \widehat{ev}_*\left(\left[\overline{\mathcal{M}}_{g,n}(D,d)\right]^{\text{vir}} \cap e(\mathcal{O}b)\right) \tag{10.1}$$

where $\mathcal{O}b$ is the obstruction bundle defined in Section 8. The basic idea is to turn on a generic Ruan-Tian perturbation of the type described in Theorem 9.1. The perturbation defines a section of $\mathcal{O}b$ whose zero set represents the Euler class $e(\mathcal{O}b)$, and on the other hand is cobordant to a cycle representing the local GW invariant. The proof consists of the three steps done in this section, together with the discussion of orientations in Section 11.

To simplify notation, we set

$$\begin{aligned} \mathcal{M}ap_D &= \mathcal{M}ap_{g,n}(D,d), & Y_D &= \overline{\mathcal{M}}_{g,n} \times D^n \\ \mathcal{M}ap &= \mathcal{M}ap_{g,n}(N_D,d), & Y &= \overline{\mathcal{M}}_{g,n} \times N_D^n. \end{aligned}$$

Let $\hat{\mathcal{F}} \to \mathcal{M}ap$ be the topological vector bundle whose fiber at f is $\Omega^{0,1}(f^*TN_D)$. This bundle has a section Φ given by the J_α-holomorphic map equation $\Phi(f) = \overline{\partial}_J f - K_\alpha(\partial_J f)j$. For each perturbation ν in $\mathcal{P}$ defined in (1.9), one can perturb Φ to obtain a section

$$\Phi_\nu(f) = \Phi(f) - \nu_f \tag{10.2}$$

where ν_f is the restriction of ν to the graph of f as in (1.10). The zero set $\overline{\mathcal{M}}^\nu = \Phi_\nu^{-1}(0)$ is the moduli space of (J_α,ν)-holomorphic maps into N_D.

By Theorem 8.1, the bundle $\mathcal{O}b$ is defined on a neighborhood $\mathcal{N}$ of the moduli space of J_α-holomorphic maps inside $\mathcal{M}ap$. Following [RT2], one can decompose $\mathcal{N}$ into strata $\mathcal{N}_B$ indexed by a finite collection of sets B that specify: (i) a homeomorphism type of the domain as a curve with marked points and (ii) a degree of map associated to each component of the domain.

Step 1. As in Section 9, we can decompose ν according to the splitting of f^*TN_D into horizontal and vertical subspaces. Fix a C^1-small perturbation ν_T that is generic in the subspace $\mathbb{P}_T \subset \mathbb{P}$ of perturbations that lie in the horizontal subspace and consider the section Φ_{ν_T} defined by (10.2).

By the Image Localization Theorem 9.1 the images of all (J_α, ν_T)-holomorphic maps lie in D. The transversality arguments of Section 3 in [RT2], applied to maps into D, imply that the moduli space

$$\overline{\mathcal{M}} = \Phi_{\nu_T}^{-1}(0) \subset \mathcal{N} \cap \mathcal{M}ap_D$$

has a stratification such that each stratum $\mathcal{M}_B = \overline{\mathcal{M}} \cap \mathcal{N}_B$ is a smooth oriented manifold of dimension $4\beta + 2n - 2k_B$ where $\beta = d(1-h) + g - 1$ as in (3.10) and k_B is the number of nodes of the domains of maps in $\mathcal{N}_B$.

Let $\mathcal{M}$ be the top stratum consisting of maps with smooth domains. Now apply the construction described at the end of Section 3: fix a neighborhood U of the image $\widehat{ev}(\partial\overline{\mathcal{M}})$ of the boundary $\partial\overline{\mathcal{M}} = \overline{\mathcal{M}} \setminus \mathcal{M}$ and consider the manifold with boundary $\overline{\mathcal{M}}_U \subset \mathcal{M}$ as in Definition 3.1. The following lemma shows that space $\overline{\mathcal{M}}_U$ represents the Li-Tian virtual fundamental class (3.4) for $(X, A) = (D, d)$ modulo small neighborhoods of its boundary.

LEMMA 10.1. *Let $\hat{U} = \widehat{ev}^{-1}(\overline{U})$ be the preimage of $\overline{U}$ under the evaluation map (0.4) on $\mathcal{M}ap_D$. Then*

$$j_*\left[\overline{\mathcal{M}}_{g,n}(D,d)\right]^{vir} = i_*\left[\overline{\mathcal{M}}_U\right] \in H_{4\beta+2n}\big(\mathcal{M}ap_D, \hat{U}\big)$$

where $j : \mathcal{M}ap_D \to (\mathcal{M}ap_D, \hat{U})$ and $i : (\overline{\mathcal{M}}_U, \partial\overline{\mathcal{M}}_U) \to (\mathcal{M}ap_D, \hat{U})$ are inclusion maps.

Proof. Let $\hat{\mathcal{F}}_D \to \mathcal{M}ap_D$ be the bundle whose fiber at f is $\Omega^{0,1}(f^*TD)$. Since $K_\alpha \equiv 0$ on D, we can regard Φ_{ν_T} as a section of $\hat{\mathcal{F}}_D$. The arguments used to prove [LT2, Proposition 3.4] then show how the section Φ_{ν_T} gives rise to a cover of the moduli space $\overline{\mathcal{M}} = \Phi_{\nu_T}^{-1}(0)$ by finitely many *smooth approximations* $\{(W_k, F_k)\}$. This means that

- each W_k is open in $\mathcal{M}ap_D$ and $\overline{\mathcal{M}} \subset \bigcup W_k$, and
- each F_k is a subbundle of $\hat{\mathcal{F}}_D$ over W_k with finite rank such that $\Phi_{\nu_T}^{-1}(F_k) \subset W_k$ is smooth and F_k restricts to a smooth bundle over $\Phi_{\nu_T}^{-1}(F_k)$ with smooth section Φ_{ν_T}.

Observe that, because the top stratum $\mathcal{M}$ is already smooth, we can assume that

• W_1 consists of all maps with smooth domains (so $W_1 \cap \overline{\mathcal{M}} = \mathcal{M}$) and F_1 has rank zero.

The proofs of [LT2, Proposition 2.2 and Theorem 1.2] describe how one can perturb the moduli space $\overline{\mathcal{M}}$ inside the union of the smooth manifolds $\Phi^{-1}(F_k)$ to obtain a cycle that represents the virtual fundamental class $[\overline{\mathcal{M}}_{g,n}(D,d)]^{\mathrm{vir}}$.

Now choose a neighborhood V of the image $\widehat{ev}(\partial\overline{\mathcal{M}})$ in Y_D with $\overline{V} \subset U$. Set $\hat{V} = \widehat{ev}^{-1}(V)$ and replace each W_i with $i \geq 2$ with $W_i \cap \hat{V}$. Then the open set $W = \bigcup_{k>1} W_k$ lies in $\hat{V}$. Following [LT2] we can then perturb $\overline{\mathcal{M}} \cap W$ while keeping $\overline{\mathcal{M}}$ fixed outside of $\hat{V}$—hence on $\overline{\mathcal{M}}_U$—to produce a cycle representing $[\overline{\mathcal{M}}_{g,n}(D,A)]^{\mathrm{vir}}$. Passing to homology proves the statement of the lemma. $\square$

Remark 10.2. The virtual fundamental class that appears in Lemma 10.1 is the Li-Tian class (3.4). We can verify that it is also the Ruan-Tian class: let $k : Y_D \to (Y_D, \overline{U})$ be an inclusion map. Since $\widehat{ev} \circ j = k \circ \widehat{ev}$, Lemma 10.1 gives $k_* \circ \widehat{ev}_*[\overline{\mathcal{M}}_{g,n}(D,d)]^{\mathrm{vir}} = \widehat{ev}_*[\overline{\mathcal{M}}_U]$. The uniqueness statement in Lemma 3.2 then shows that

$$\widehat{ev}_*\left[\overline{\mathcal{M}}_{g,n}(D,d)\right]^{\mathrm{vir}} = GW_{g,n}(D,d) \in H_{4\beta+2n}\left(\overline{\mathcal{M}}_{g,n} \times D^n; \mathbb{Q}\right).$$

Step 2. We now further perturb the section Φ_{ν_T} by adding a section of the obstruction bundle $\mathcal{O}b$ induced from a Ruan-Tian perturbation μ. Each $\mu \in \mathcal{P}$ defines a section $\hat{s}_\mu$ of the obstruction bundle $\mathcal{O}b$ over $\mathcal{N}$ by

$$\hat{s}_\mu(f) = P_f\left(\mu_f\right) \tag{10.3}$$

where P_f is the L^2-orthogonal projection onto the fiber $\mathcal{O}b_f$. Let s_μ denote the restriction of $\hat{s}_\mu$ to the moduli space $\overline{\mathcal{M}}$. Whenever μ has small C^1 norm, the Image Localization Theorem 9.1 implies that

$$\left(\Phi_{\nu_T} - \hat{s}_\mu\right)^{-1}(0) = \overline{\mathcal{M}} \cap \hat{s}_\mu^{-1}(0) = s_\mu^{-1}(0). \tag{10.4}$$

LEMMA 10.3. *For generic μ in $\mathcal{P}$, the space (10.4) has a stratification indexed by B such that each stratum $\mathcal{M}_B \cap s_\mu^{-1}(0)$ is a smooth manifold of dimension $2\beta + 2n - 2k_B$.*

Proof. The proof is a transversality argument using the universal moduli space over the space of perturbations $\mathcal{P}$ (cf. [RT2, Theorem 3.1]). The smooth bundle

$$\mathcal{O}b \longrightarrow \mathcal{M}_B \times \mathcal{P}$$

has rank 2β and has a section σ_B defined by $\sigma_B(f,\mu) = s_\mu(f)$ whose zero set is the universal moduli space associated with B. The differential of σ_B at (f,μ) is

$$\left(D\sigma_B\right)_{(f,\mu)}(\xi,\chi) = \left(Ds_\mu\right)_f(\xi) - P_f\left(\chi_f\right) \quad \text{for } \xi \in T_f\mathcal{M}_B,\ \chi \in T_\mu\mathcal{P}$$

where Ds_μ is the differential of the restriction of s_μ to $\mathcal{M}_B$ and χ_f is the restriction of χ to the graph of f as in (1.10). Since the map $\chi \to P_f(\chi_f)$ is onto, so is the differential $(D\sigma_B)_{(f,\mu)}$. The universal moduli space $\mathcal{U} = \sigma_B^{-1}(0)$ is therefore smooth. Now, consider the projection

$$\pi_B : \mathcal{U} \longrightarrow \mathcal{P}$$

given by $\pi_B(f,\mu) = \mu$. Since π_B is Fredholm, the Sard-Smale Theorem implies that for generic μ the fiber $\pi_B^{-1}(\mu) = \mathcal{M}_B \cap s_\mu^{-1}(0)$ is a smooth manifold of dimension $\dim \mathcal{M}_B - \operatorname{rank} \mathcal{O}b$, which is $2\beta + 2n - 2k_B$. $\square$

Fix a generic $\mu \in \mathbb{P}$ as in Lemma 10.3 and set

$$\bar{Z} = s_\mu^{-1}(0) = \left(\Phi_{\nu_T} - \hat{s}_\mu\right)^{-1}(0). \tag{10.5}$$

By Lemma 10.3 and the discussion above Lemma 3.2, applied to some neighborhood of $\widehat{ev}(\partial \bar{Z})$ where $\partial \bar{Z} = \bar{Z} \cap \partial \overline{\mathcal{M}}$, the image $\widehat{ev}(\overline{Z})$ defines a homology class

$$\left[\widehat{ev}(\bar{Z})\right] \in H_{2\beta+2n}\left(Y_D;\mathbb{Q}\right) \cong H_{2\beta+2n}(Y;\mathbb{Q}). \tag{10.6}$$

On the other hand, the local GW invariant class is determined by the image of the moduli space $\overline{\mathcal{M}}^\nu = \Phi_\nu^{-1}(0)$ for a generic $\nu \in \mathcal{P}$ as in (3.9):

$$GW_{g,n}^{loc}\left(N_D, d\right) = \left[\widehat{ev}\left(\overline{\mathcal{M}}^\nu\right)\right] \in H_{2\beta+2n}(Y;\mathbb{Q}). \tag{10.7}$$

Thus the local invariant is defined in terms of the zero set of Φ_ν for a fixed generic $\nu \in \mathcal{P}$, while the class (10.6) is defined in terms of the zero set of $\Phi_{\nu_T} - \hat{s}_\mu$. To show these are equal we introduce, for each pair (ν, υ) in $\mathcal{P} \times \mathcal{P}$, a section of the bundle $\hat{\mathcal{F}} \to \mathcal{M}ap$ defined by

$$\Phi_{\nu,\upsilon}(f) = \Phi(f) - \nu_f - \hat{s}_\upsilon(f).$$

The transversality argument of Lemma 10.3, now applied to the universal moduli space over the parameter space $\mathcal{P} \times \mathcal{P}$, shows that, for generic small (ν, υ) in $\mathcal{P} \times \mathcal{P}$, $\Phi_{\nu,\upsilon}^{-1}(0)$ has a stratification indexed by B such that each stratum $\mathcal{N}_B \cap \Phi_{\nu,\upsilon}^{-1}(0)$ is a smooth manifold of dimension $2\beta + 2n - 2k_B$. Standard cobordism arguments (cf. [RT2, Theorem 3.3]) then give

$$GW_{g,n}^{\mathrm{loc}}\left(N_D, d\right) = \left[\widehat{ev}\left(\Phi_{\nu,0}^{-1}(0)\right)\right] = \left[\widehat{ev}\left(\Phi_{\nu_T,\mu}^{-1}(0)\right)\right] = \left[\widehat{ev}(\bar{Z})\right]. \tag{10.8}$$

Step 3. It remains to show that $[\widehat{ev}(\bar{Z})]$ equals to the right-hand side of (10.1). Again, to avoid issues of smoothness near $\partial \overline{\mathcal{M}}$, we work in relative homology. The following lemma is, in essence, the statement that the zero set of a generic section of a vector bundle is Poincaré dual (in relative homology) to the Euler class.

LEMMA 10.4. *Let* $\bar{Z}_U = \overline{\mathcal{M}}_U \cap \bar{Z}$. *Then*

$$j_*\big([\overline{\mathcal{M}}_{g,n}(D,d)]^{vir} \cap e(\mathcal{O}b)\big) = i_*\big[\bar{Z}_U\big] \in H_{2\beta+2n}\big(\mathcal{M}ap_D, \hat{U}\big) \tag{10.9}$$

where $j : \mathcal{M}ap_D \to (\mathcal{M}ap_D, \hat{U})$ *and* $i : (\overline{\mathcal{M}}_U, \partial\overline{\mathcal{M}}_U) \to (\mathcal{M}ap_D, \hat{U})$ *are inclusion maps.*

Proof. The obstruction bundle $\mathcal{O}b$ restricts to a smooth bundle $\mathcal{O}b_U$ over $\overline{\mathcal{M}}_U$. Let s_U be the restriction of the section $\hat{s}_\mu$ in (10.3) to $\overline{\mathcal{M}}_U$. Without loss of generality, we can assume s_U is transverse to the zero section, so $\bar{Z}_U = s_U^{-1}(0)$ is a compact smooth manifold with boundary.

Define a double $\tilde{\mathcal{M}}_U$ of $\overline{\mathcal{M}}_U$ by identifying two copies of $\overline{\mathcal{M}}_U$ along $\partial\overline{\mathcal{M}}_U$. Similarly, define a vector bundle $\tilde{\mathcal{O}b}_U \to \tilde{\mathcal{M}}_U$ and its section $\tilde{s}_U$ that are doubles of $\mathcal{O}b_U$ and s_U respectively. Since $\tilde{\mathcal{M}}_U$ is a closed manifold, we have

$$\big[\tilde{\mathcal{M}}_U\big] \cap e\big(\tilde{\mathcal{O}b}_U\big) = \big[\tilde{Z}_U\big] \tag{10.10}$$

where $\tilde{Z}_U = \tilde{s}_U^{-1}(0)$. Consider the following commutative diagram:

$$\begin{array}{ccccc}
H_{4\beta+2n}(\tilde{\mathcal{M}}_U) & \xrightarrow{\kappa_*} & H_{4\beta+2n}(\tilde{\mathcal{M}}_U, \tilde{\mathcal{M}}_U \setminus \mathcal{M}_U^\circ) & \xleftarrow[\simeq]{\iota_*} & H_{4\beta+2n}(\overline{\mathcal{M}}_U, \partial\overline{\mathcal{M}}_U) \\
\downarrow{\scriptstyle \cap e(\tilde{\mathcal{O}b}_U)} & & \downarrow{\scriptstyle \cap e(\tilde{\mathcal{O}b}_U)} & & \downarrow{\scriptstyle \cap e(\mathcal{O}b_U)} \\
H_{2\beta+2n}(\tilde{\mathcal{M}}_U) & \xrightarrow{\kappa_*} & H_{2\beta+2n}(\tilde{\mathcal{M}}_U, \tilde{\mathcal{M}}_U \setminus \mathcal{M}_U^\circ) & \xleftarrow[\simeq]{\iota_*} & H_{2\beta+2n}(\overline{\mathcal{M}}_U, \partial\overline{\mathcal{M}}_U)
\end{array} \tag{10.11}$$

where $\mathcal{M}_U^\circ = \overline{\mathcal{M}}_U \setminus \partial\overline{\mathcal{M}}_U$, κ and ι are inclusion maps, each rectangle commutes by the naturality of the cap product, and the horizontal isomorphisms follow by excision. Consequently, (10.10), (10.11) and the facts $\kappa_*[\tilde{\mathcal{M}}_U] = \iota_*[\overline{\mathcal{M}}_U]$ and $\kappa_*[\tilde{Z}_U] = \iota_*[\bar{Z}_U]$ show

$$\big[\overline{\mathcal{M}}_U\big] \cap e\big(\mathcal{O}b_U\big) = \big[\bar{Z}_U\big] \in H_{2\beta+2n}\big(\overline{\mathcal{M}}_U, \partial\overline{\mathcal{M}}_U\big).$$

Now, (10.9) follows from this equality, Lemma 10.1 and the naturality of cap product. □

Proof of Main Theorem. Recall that $\overline{\mathcal{M}}$ and $\bar{Z}$ are decomposed into smooth strata indexed by sets B such that $\dim \mathcal{M}_B = 4\beta + 2n - 2k_B$ and $\dim Z_B = 2\beta + 2n - 2k_B$ where $Z_B = \bar{Z} \cap \mathcal{M}_B$. Because Y_D is a compact manifold, any class in $H_*(Y_D;\mathbb{Q})$ can be represented by an embedded submanifold by Thom's Theorem [Th]. We can thus choose a basis for $H_{m-2\beta}(Y_D;\mathbb{Q})$, where $m + 2n = \dim Y_D$, represented by submanifolds $\mathcal{D}_i$ in general position with respect to all the restriction maps $\widehat{ev}_{|\mathcal{M}_B}$ and $\widehat{ev}_{|Z_B}$. Counting dimensions, one sees that

- each $\mathcal{D}_i$ is disjoint from the image $\widehat{ev}(\partial\bar{Z}) \subset \widehat{ev}(\partial\overline{\mathcal{M}})$, and
- we can choose a submanifold $\mathcal{D}$ representing the class $\widehat{ev}_*([\overline{\mathcal{M}}_{g,n}(D,d)]^{\mathrm{vir}} \cap e(\mathcal{O}b))$ that is disjoint from $\mathcal{D}_i \cap \widehat{ev}(\partial\overline{\mathcal{M}})$ for all i.

Hence, by shrinking the neighborhood U of $\widehat{ev}(\partial\overline{\mathcal{M}})$ if necessary, we can assume that

$$\overline{U}\cap\widehat{ev}(\bar{Z})\cap\mathcal{D}_i=\emptyset \quad\text{and}\quad \overline{U}\cap\mathcal{D}\cap\mathcal{D}_i=\emptyset \tag{10.12}$$

for all i. Now, consider the commutative diagram:

$$\begin{tikzcd} H_{2\beta+2n}(\mathcal{M}ap_D) \arrow[r, "j_*"] \arrow[d, "\widehat{ev}_*"] & H_{2\beta+2n}(\mathcal{M}ap_D,\hat{U}) \arrow[d, "\widehat{ev}_*"] & H_{2\beta+2n}(\overline{\mathcal{M}}_U,\partial\overline{\mathcal{M}}_U) \arrow[l, "i_*"'] \\ H_{2\beta+2n}(Y_D) \arrow[r, "k_*"] & H_{2\beta+2n}(Y_D,\bar{U}). \end{tikzcd} \tag{10.13}$$

Observe that (10.12) implies that for all i

(10.14)
$$(a)\ \big[\mathcal{D}\big]\cdot\big[\mathcal{D}_i\big]=k_*[\mathcal{D}]\cdot k_*\big[\mathcal{D}_i\big] \quad\text{and}\quad (b)\ \big[\widehat{ev}(\bar{Z})\big]\cdot\big[\mathcal{D}_i\big]=\widehat{ev}_*\circ i_*\big(\big[\bar{Z}_U\big]\big)\cdot k_*\big[\mathcal{D}_i\big]$$

where, in both of these equations, the dot on the right-hand side is the intersection pairing in $H_{2\beta+2n}(Y_D,\bar{U})\cong H_{2\beta+2n}(Y_D/U,\partial U)$. By the definition of $\mathcal{D}$, (10.14a) states that

$$\widehat{ev}_*\big(\big[\overline{\mathcal{M}}_{g,n}(D,d)\big]^{\mathrm{vir}}\cap e(\mathcal{O}b)\big)\cdot\big[\mathcal{D}_i\big]=k_*\circ\widehat{ev}_*\big(\big[\overline{\mathcal{M}}_{g,n}(D,d)\big]^{\mathrm{vir}}\cap e(\mathcal{O}b)\big)\cdot k_*\big[\mathcal{D}_i\big].$$

But the intersection on the right is, by the commutativity of the diagram, (10.9) and (10.14b),

$$\begin{aligned}\widehat{ev}_*\circ j_*\big(\big[\overline{\mathcal{M}}_{g,n}(D,d)\big]^{\mathrm{vir}}\cap e(\mathcal{O}b)\big)\cdot k_*\big[\mathcal{D}_i\big] &= \widehat{ev}_*\circ i_*\big(\big[\bar{Z}_U\big]\big)\cdot k_*\big[\mathcal{D}_i\big]\\ &= \big[\widehat{ev}(\bar{Z})\big]\cdot\big[\mathcal{D}_i\big].\end{aligned}$$

This shows $\widehat{ev}_*([\overline{\mathcal{M}}_{g,n}(D,d)]^{\mathrm{vir}}\cap e(\mathcal{O}b))=[\widehat{ev}(\bar{Z})]$ and hence, together with (10.8), completes the proof of the Main Theorem (10.1). $\square$

11. Secondary index invariants. This section puts the obstruction bundle in a general context and explains why Euler class $e(\mathcal{O}b)$ cannot in general be computed using Grothendieck-Riemann-Roch or the Families Index theorems.

Let E and F be Banach vector bundles over a compact parameter space X. We can consider the vector bundle

$$\begin{array}{c}\mathrm{Fred}_\ell(E,F)\\ \big\downarrow\\ X\end{array}$$

whose fiber over $x \in X$ is the space of real linear Fredholm maps from E_x to F_x with index $-\ell \leq 0$. A section L of $\mathrm{Fred}_\ell(E,F)$ then defines a family-index class

$$\mathrm{ind}\, L \in KR(X)$$

obtained by pulling back the class of the virtual bundle $[\ker L] - [\mathrm{coker}\, L]$ on $\mathrm{Fred}_\ell(E,F)$. The index theorem for families [AS] gives formulas for the Pontryagin classes $p_i(\mathrm{ind}\, L)$. But the Euler class does not factor through K-theory, and hence $e(\mathrm{ind}\, L)$ is not computable in the same way and in fact is not even defined in general.

In this context, for each $k > 0$, the set $A_{k,\ell}$ of all $L \in \mathrm{Fred}_\ell(E,F)$ with $\dim\ker L = k$ and $\dim\mathrm{coker}\, L = k+\ell$ is a submanifold of real codimension $k(k+\ell)$. As shown in [K], the closures $\overline{A}_{k,\ell}$ of these submanifolds define "Koschorke classes"

$$\kappa_{k,\ell} \in H^{k(k+\ell)}\big(\mathrm{Fred}_\ell(E,F)\big). \tag{11.1}$$

A section L of $\mathrm{Fred}_\ell(E,F)$ then defines cohomology classes $L^*\kappa_{k,\ell} \in H^{k(k+\ell)}(X)$. The classes $\{L^*\kappa_{k+i,\ell}\,|\, i>0\}$ are the obstructions to deforming L within its homotopy class to a family $\{L_x\}$ of operators with $\dim\ker L_x \leq k$ for all $x \in X$. In particular, when all Koschorke classes $\{L^*\kappa_{k,\ell}\}$ vanish, L can be deformed to a section $\tilde{L}$ with $\ker \tilde{L}_x = 0$ for all $x \in X$; $\mathrm{coker}\, \tilde{L}$ is then a rank ℓ vector bundle over X that represents $\mathrm{ind}\, L \in KR(X)$. In this sense, the Koschorke classes are the obstruction to realizing the family index—which is defined as a formal difference of bundles—as an *actual* vector bundle. Furthermore, when the Koschorke classes vanish this bundle is well-defined up to homotopy and hence, assuming $\mathrm{ind}\, L$ is an oriented bundle, the Euler class

$$e(\mathrm{ind}\, L) \in H^\ell(X)$$

is defined. This is a "secondary class" in the sense that it exists only for families with vanishing Koschorke classes.

Now consider the situation at hand, where $X = \overline{\mathcal{M}}_{g,n}(D,d) \subset \mathcal{M}ap_{g,n}(N_D,d)$, E and F are the bundles $\mathcal{E}$ and $\mathcal{F}$ defined in Section 5, and L is the linearization map $f \mapsto L_f$ with index $-\ell = -2\beta$. Then (except for the complication described below)

- $L^*\kappa_{k,2\beta} = 0$ for all $k > 0$ by the Vanishing Theorem 4.2.

This gives a global perspective on the role of Theorem 4.2: it ensures that all Koschorke classes vanish. In fact, it shows that after perturbing the Kähler structure J to J_α, the space of J_α-holomorphic maps is mapped by $\Psi(f) = L_f$ into a region in $\mathrm{Fred}_\ell(\mathcal{E},\mathcal{F})$ where the index bundle is an actual bundle. The pullback $\Psi^*(\mathrm{ind}\, L)$ is the obstruction bundle $\mathcal{O}b$ of the Main Theorem.

Now the complication: while the above paragraph provides valuable intuition, it is not rigorous until one prove that $\mathcal{E}$ and $\mathcal{F}$ are locally trivial bundles, or that the map $L:\mathcal{E}\to\mathcal{F}$ is homotopy equivalent to a map $L':\mathcal{E}'\to\mathcal{F}'$ between vector bundles with isomorphisms $\ker L\cong\ker L'$ and $\operatorname{coker} L\cong\operatorname{coker} L'$ for each map in the moduli space. It would be interesting to directly establish this general Koschorke picture.

Our proof of the Main Theorem in Section 10 used the fact that the obstruction bundle carries a canonical orientation. The proof, which we give now, fits into the above discussion of the space of Fredholm operators. Below, we write $\overline{\mathcal{M}}$ for $\overline{\mathcal{M}}_{g,n}(D,d)$.

LEMMA 11.1. *The bundle $\mathcal{O}b$ is orientable and has a canonical orientation.*

Proof. The linearization $f\mapsto L_f$ defines a map $L:\overline{\mathcal{M}}\to\text{Fred}$ whose image lies in the set of Fredholm operators with trivial kernel. Over this set, $\mathcal{O}b$—whose fiber is the cokernel—is isomorphic to the dual of the index bundle. The obstruction to orientability is therefore the pullback $w=L^*w_1(\det L)\in H^2(\overline{\mathcal{M}},\mathbb{Z}_2)$ of the first Stiefel-Whitney class of the real determinant of the index bundle. But the map L extends canonically to a homotopy $\overline{\mathcal{M}}\times[0,1]\to\text{Fred}$ by writing $L_f=\overline{\partial}+R_\alpha$ and setting $L_t(f)=\overline{\partial}+tR_\alpha$. Because $\det L$ is defined over the entire space of Fredholm operators, we have $w=L_0^*w_1(\det L)$. This is zero because the image of L_0 consists of complex Fredholm operators, whose kernel and cokernels have canonical complex orientations.

This shows, in fact, that the orientation bundle $\Lambda^{\text{top}}\mathcal{O}b$ is trivial over $\overline{\mathcal{M}}\times[0,1]$. There are therefore two orientation classes (nowhere-vanishing sections modulo multiplication by positive functions). The one that agrees with the complex orientation along $\overline{\mathcal{M}}\times\{0\}$ will be called the "canonical" orientation on $\mathcal{O}b$. □

Finally, we must specify the orientation on $\mathcal{O}b$ for which the Main Theorem holds. Let $\bar{Z}$ be the zero set of a transverse section $\mathcal{O}b\to\overline{\mathcal{M}}$. At each $(C,f)\in\bar{Z}$, the standard orientations of $\overline{\mathcal{M}}$ and $\bar{Z}$ used in GW theory are given by the determinant bundles $\det(\overline{\partial}^T\oplus Jdf)$ and $\det(\overline{\partial}^T\oplus\overline{\partial}\oplus Jdf)$ respectively. Thus in the equality (0.5) in the Main Theorem, the cycles representing the two sides are consistently oriented provided

$$\det\left(\overline{\partial}^T\oplus Jdf\right)=\det\left(\overline{\partial}^T\oplus\overline{\partial}\oplus Jdf\right)\otimes\Lambda^{\text{top}}\mathcal{O}b$$

This equality holds since the canonical orientation defined in the proof of Lemma 11.1 is

$$\Lambda^{\text{top}}\mathcal{O}b=\det(\overline{\partial})^*.$$

12. Remarks on calculating Euler classes. We conclude with some remarks on calculating the Euler class of the obstruction bundle. Algebraic geometers have a standard procedure for calculating the Euler class of the index bundles of families of $\overline{\partial}$ operators by using the Grothendieck-Riemann-Roch formula, often in conjunction with localization by a group action. This procedure is not, in general, applicable to finding the Euler class of the *real* bundle $\mathcal{O}b$. But it is worth noting that the GRR formula yields some information in the following two cases.

(1) Since the square of the Euler class is the top Pontryagin class, we have $e^2(\mathcal{O}b) = p_\beta(\mathcal{O}b) = (-1)^\beta c_{2\beta}(\mathcal{O}b \otimes_{\mathbb{R}} \mathbb{C})$. But $\mathcal{O}b \otimes_{\mathbb{R}} \mathbb{C}$ is the complex index bundle $\mathrm{ind}_{\mathbb{C}} L$ because $\ker L$ vanishes for all operators in the family. Thus the GRR formula is applicable for finding $e^2(\mathcal{O}b)$, but not for finding $e(\mathcal{O}b)$ itself.

(2) Recall Givental's notion of *twisted GW invariants*: each bundle E over X, determines a virtual bundle $\mathcal{E} \to \overline{\mathcal{M}}_{g,n}(X, A)$ over the space of stable maps whose fiber over a map $f : C \to X$ is $H^0(C, f^*E) \ominus H^1(C, f^*E)$. The corresponding twisted invariants are obtained by evaluating Chern classes of $\mathcal{E}$, together with τ classes, on the virtual fundamental class. In some especially simple situations, the Euler class of the obstruction bundle can be expressed in terms of twisted GW invariants, and these can be calculated.

LEMMA 12.1. *Over the space of J_α-holomorphic maps into N_D there is a locally trivial complex vector bundle Q that is isomorphic to the obstruction bundle $\mathcal{O}b$ as an* unoriented *real vector bundle, such that for each J_α-holomorphic map f the fiber Q_f of Q fits into the split exact sequence of complex vector spaces*

$$0 \longrightarrow \overline{\ker \overline{\partial}_f} \xrightarrow{R_\alpha} \operatorname{coker} \overline{\partial}_f \longrightarrow Q_f \longrightarrow 0 \tag{12.1}$$

where R_α is as in (4.3) and where $\overline{\partial}_f$ is an operator on f^*N between the unweighted Sobolev spaces $L^{1,2}(E_f)$ and $L^2(F_f)$ defined in Lemma 7.8.

Proof. We will repeatedly use the facts that R_α satisfies $JR_\alpha = -R_\alpha J$ and

$$\big\langle \overline{\partial}\xi_1, R_\alpha \xi_2 \big\rangle = -\big\langle R_\alpha \xi_1, \overline{\partial}\xi_2 \big\rangle \tag{12.2}$$

where $\langle \cdot, \cdot \rangle$ denote the standard, unweighted, L^2 inner product (cf. [LP, Section 8]). Fix a J_α-holomorphic map $f : C \to N_D$ and write $\overline{\partial}_f$ as $\overline{\partial}$. Regard $\operatorname{coker} \overline{\partial}$ as the L^2 orthogonal complement to the image of $\overline{\partial}$. First note that if $\xi \in \ker \overline{\partial}$ then by (12.2)

$$\langle \overline{\partial}\xi', R_\alpha \xi \rangle = -\langle R_\alpha \xi', \overline{\partial}\xi \rangle = 0$$

for all $\xi' \in \Omega^0(f^*N)$. Hence $R_\alpha \xi$ lies in $\operatorname{coker} \overline{\partial}$. Defining Q to be the L^2 perpendicular

$$Q = \big(R_\alpha(\ker \overline{\partial})\big)^\perp \subset \operatorname{coker} \overline{\partial}$$

gives the exact sequence (12.1) of complex vector spaces. Note that R_α, originally complex anti-linear, becomes complex linear when we reverse the complex structure on $\ker \overline{\partial}$ and that this definition of Q splits the sequence.

Next, regard $\operatorname{coker}(\overline{\partial}+R_\alpha)$ as the L^2 orthogonal complement to the image of $\overline{\partial}+R_\alpha$ and let $q:\operatorname{coker}(\overline{\partial}+R_\alpha)\to\operatorname{coker}\overline{\partial}$ be the L^2 orthogonal projection. Observe:

- q is injective: any $\eta\in\operatorname{coker}(\overline{\partial}+R_\alpha)$ with $q(\eta)=0$ has the form $\eta=\overline{\partial}\xi$ for some ξ, and hence vanishes because, by (12.2),

$$0=\big\langle\overline{\partial}\xi,(\overline{\partial}+R_\alpha)\xi\big\rangle=\|\overline{\partial}\xi\|^2=\|\eta\|^2.$$

- The image of q lies in Q: for any $\eta\in\operatorname{coker}(\overline{\partial}+R_\alpha)$ and $\xi\in\ker\overline{\partial}$ we have

$$0=\langle\eta,(\overline{\partial}+R_\alpha)\xi\big\rangle=\big\langle\eta,R_\alpha\xi\big\rangle=\big\langle q(\eta),R_\alpha\xi\big\rangle$$

(the last equality holds because $R_\alpha\xi\in\operatorname{coker}\overline{\partial}$ as above). Thus $q(\eta)$ is L^2 perpendicular to the image of R_α.

Now count dimensions. From (12.1) we have $\dim Q_f=-\operatorname{index}\overline{\partial}=-\operatorname{index}(\overline{\partial}+R_\alpha)=-2\beta$. But $\ker(\overline{\partial}+R_\alpha)=0$ by Theorem 4.2, so $\dim Q_f=\dim\operatorname{coker}(\overline{\partial}+R_\alpha)$. Thus q is an isomorphism between $\operatorname{coker}(\overline{\partial}+R_\alpha)$ and Q.

Finally, to relate Q to the (locally trivial) obstruction bundle, consider the L^2 perpendicular projection $\pi:\mathcal{O}b\to\operatorname{coker}(\overline{\partial}+R_\alpha)\cong Q$. Suppose that $v\in\mathcal{O}b$ satisfies $\pi(v)=0$. Then there is an $\xi\in L^{1,2}(E_f)$ with $(\overline{\partial}+R_\alpha)\xi=v$. Since v and f are smooth, elliptic regularity implies that ξ is smooth, so lies in the weighted space $\mathcal{E}_f$ for any $p\geq 2$. This contradicts Theorem 8.1 unless $v=0$. Thus π is injective, and therefore an isomorphism because both have dimension 2β. $\square$

As f varies across the space $\overline{\mathcal{M}}=\overline{\mathcal{M}}_{g,n}(D,d)$ of J_α-holomorphic maps, one obtains families

$$\ker\overline{\partial}\longrightarrow\overline{\mathcal{M}}\quad\text{and}\quad\operatorname{coker}\overline{\partial}\longrightarrow\overline{\mathcal{M}}$$

whose fibers are the complex vector spaces $\ker\overline{\partial}_f=H^0(f^*N)$ and $\operatorname{coker}\overline{\partial}_f=H^1(f^*N)$. In general, the dimensions of these fibers are not constant: the dimension of the kernel and the cokernel jumps up (by equal amounts) along a "jumping locus" in $\overline{\mathcal{M}}$. But away from the jumping locus Lemma 12.1 gives a formula, due to Kiem and Li [KL], for the Euler class of the obstruction bundle:

PROPOSITION 12.2. *Suppose* $\ker\overline{\partial}$ *and* $\operatorname{coker}\overline{\partial}$ *are locally trivial vector bundles over a set $Z\subset\overline{\mathcal{M}}$. Then there is an isomorphism of oriented real vector bundles*

$$\mathcal{O}b\simeq(-1)^{h_i}Q \tag{12.3}$$

over each component Z_i of Z where h_i is $h^0(f^*N)$ on Z_i. Consequently, $e(\mathcal{O}b)\in H^{2\beta}(Z)$ is

$$e(\mathcal{O}b)=\sum(-1)^{h_i}c_\beta\Big(\Big[\operatorname{coker}\overline{\partial}\big|_{Z_i}\Big]-\Big[\overline{\ker\overline{\partial}\big|_{Z_i}}\Big]\Big) \tag{12.4}$$

where $2\beta=-\operatorname{index}\overline{\partial}$ and the sum is over all connected components Z_i of Z.

Proof. Q has a complex orientation from (12.1) and the orientation of $\operatorname{coker}(\overline{\partial}+R_\alpha)=\mathcal{O}b$ is given by $\det(\overline{\partial})^*$ as in Lemma 11.1. These are related by

$$
\begin{aligned}
\det(\overline{\partial})^* &= \wedge^{\text{top}}\operatorname{coker}\overline{\partial}\otimes\wedge^{\text{top}}\ker\overline{\partial}\\
&= (-1)^{h^0}\wedge^{\text{top}}\operatorname{coker}\overline{\partial}\otimes\wedge^{\text{top}}\overline{\ker\overline{\partial}}\\
&= (-1)^{h^0}\det Q,
\end{aligned}
\tag{12.5}
$$

which gives (12.3). Taking Euler classes and noting that Q is a complex bundle, we have $e(\mathcal{O}b)=(-1)^{h_i}c_\beta(Q)$ and (12.4) follows from the exact sequence of Lemma 12.1. □

Example 12.3. Suppose that D is an elliptic curve with odd theta characteristic. Because the local GW invariants depend only on the parity of the theta characteristic, we can take N to be a trivial bundle. Then $\ker\overline{\partial}$ is the trivial line bundle $\mathbb{C}$ over $\overline{\mathcal{M}}$. In this case, we have

$$
e(\mathcal{O}b)=-c_\beta\big([\operatorname{coker}\overline{\partial}]-[\overline{\mathbb{C}}]\big)=-c_\beta(\operatorname{coker}\overline{\partial})=c_\beta\big(\operatorname{ind}_{\mathbb{C}}\overline{\partial}\big)
\tag{12.6}
$$

where $2\beta=-\operatorname{index}\overline{\partial}$ and where $\operatorname{ind}_{\mathbb{C}}\overline{\partial}$ is the complex index bundle.

The formula (12.6) shows that, in this case, the local GW invariants are special cases of Giventhal's twisted GW invariants of curves. These can be explicitly computed using the result of Proposition 2 in the paper of Faber and Pandharipande [FP].

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CENTRAL FLORIDA, ORLANDO, FL 32816
E-mail: Junho.Lee@ucf.edu

DEPARTMENT OF MATHEMATICS, MICHIGAN STATE UNIVERSITY, EAST LANSING, MI 48824
E-mail: parker@math.msu.edu

REFERENCES

[ACV] D. Abramovich, A. Corti, and A. Vistoli, Twisted bundles and admissible covers, *Comm. Algebra* **31** (2003), no. 8, 3547–3618.
[AS] M. F. Atiyah and I. M. Singer, The index of elliptic operators. IV, *Ann. of Math. (2)* **93** (1971), 119–138.
[BF] K. Behrend and B. Fantechi, The intrinsic normal cone, *Invent. Math.* **128** (1997), no. 1, 45–88.

[BP] J. Bryan and R. Pandharipande, The local Gromov-Witten theory of curves, *J. Amer. Math. Soc.* **21** (2008), no. 1, 101–136, With an appendix by J. Bryan, C. Faber, A. Okounkov, and R. Pandharipande.

[FP] C. Faber and R. Pandharipande, Hodge integrals and Gromov-Witten theory, *Invent. Math.* **139** (2000), no. 1, 173–199.

[FO] K. Fukaya and K. Ono, Arnold conjecture and Gromov-Witten invariant, *Topology* **38** (1999), no. 5, 933–1048.

[IP1] E.-N. Ionel and T. H. Parker, Relative Gromov-Witten invariants, *Ann. of Math. (2)* **157** (2003), no. 1, 45–96.

[IP2] ———, The symplectic sum formula for Gromov-Witten invariants, *Ann. of Math. (2)* **159** (2004), no. 3, 935–1025.

[IS1] S. Ivashkovich and V. Shevchishin, Structure of the moduli space in a neighborhood of a cusp-curve and meromorphic hulls, *Invent. Math.* **136** (1999), no. 3, 571–602.

[IS2] ———, Gromov compactness theorem for J-complex curves with boundary, *Internat. Math. Res. Notices* (2000), no. 22, 1167–1206.

[KL] Y. Kiem and J. Li, Gromov-Witten invariants of varieties with holomorphic 2-forms, preprint, arXiv:0707.2986.

[K] U. Koschorke, Infinite dimensional K-theory and characteristic classes of Fredholm bundle maps, *Global Analysis (Proc. Sympos. Pure Math., Vol. XV, Berkeley, Calif., 1968)*, American Mathematical Society, Providence, RI, 1970, pp. 95–133.

[KM] P. B. Kronheimer and T. S. Mrowka, Embedded surfaces and the structure of Donaldson's polynomial invariants, *J. Differential Geom.* **41** (1995), no. 3, 573–734.

[Lee1] J. Lee, Family Gromov-Witten invariants for Kähler surfaces, *Duke Math. J.* **123** (2004), no. 1, 209–233.

[Lee2] ———, Sum formulas for local Gromov-Witten invariants of spin curves, *Trans. Amer. Math. Soc.* (to appear).

[LP] J. Lee and T. H. Parker, A structure theorem for the Gromov-Witten invariants of Kähler surfaces, *J. Differential Geom.* **77** (2007), no. 3, 483–513.

[LT1] J. Li and G. Tian, Virtual moduli cycles and Gromov-Witten invariants of algebraic varieties, *J. Amer. Math. Soc.* **11** (1998), no. 1, 119–174.

[LT2] ———, Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds, *Topics in Symplectic 4-manifolds (Irvine, CA, 1996)*, *First Int. Press Lect. Ser., I*, Int. Press, Cambridge, MA, 1998, pp. 47–83.

[LT3] ———, Comparison of algebraic and symplectic Gromov-Witten invariants, *Asian J. Math.* **3** (1999), no. 3, 689–728.

[L] R. Lockhart, Fredholm, Hodge and Liouville theorems on noncompact manifolds, *Trans. Amer. Math. Soc.* **301** (1987), no. 1, 1–35.

[Lo] E. Looijenga, Smooth Deligne-Mumford compactifications by means of Prym level structures, *J. Algebraic Geom.* **3** (1994), no. 2, 283–293.

[MP] D. Maulik and R. Pandharipande, New calculations in Gromov-Witten theory, *Pure Appl. Math. Q.* **4** (2008), no. 2, Special Issue: In honor of Fedor Bogomolov. Part 1, 469–500.

[MS] D. McDuff and D. Salamon, *J-Holomorphic Curves and Symplectic Topology*, *American Mathematical Society Colloquium Publications*, vol. 52, American Mathematical Society, Providence, RI, 2004.

[OP] A. Okounkov and R. Pandharipande, Gromov-Witten theory, Hurwitz theory, and completed cycles, *Ann. of Math. (2)* **163** (2006), no. 2, 517–560.

[R] Y. Ruan, Virtual neighborhoods and pseudo-holomorphic curves, *Turkish J. Math.* **23** (1999), no. 1, 161–231, Proceedings of 6th Gökova Geometry-Topology Conference.

[RT1] Y. Ruan and G. Tian, A mathematical theory of quantum cohomology, *J. Differential Geom.* **42** (1995), no. 2, 259–367.

[RT2] ———, Higher genus symplectic invariants and sigma models coupled with gravity, *Invent. Math.* **130** (1997), no. 3, 455–516.

[S] B. Siebert, Gromov-Witten invariants for general symplectic manifolds, preprint, arXive:9608005.

[T] C. H. Taubes, Self-dual connections on 4-manifolds with indefinite intersection matrix, *J. Differential Geom.* **19** (1984), no. 2, 517–560.

[Th] R. Thom, Quelques propriétés globales des variétés différentiables, *Comment. Math. Helv.* **28** (1954), 17–86.

[Z] A. Zinger, A comparison theorem for Gromov-Witten invariants in the symplectic category, *Adv. Math.* **228** (2011), 535–574.