

Relative Gromov-Witten invariants

By ELENY-NICOLETA IONEL and THOMAS H. PARKER*

Abstract

We define relative Gromov-Witten invariants of a symplectic manifold relative to a codimension-two symplectic submanifold. These invariants are the key ingredients in the symplectic sum formula of [IP4]. The main step is the construction of a compact space of ‘ V -stable’ maps. Simple special cases include the Hurwitz numbers for algebraic curves and the enumerative invariants of Caporaso and Harris.

Gromov-Witten invariants are invariants of a closed symplectic manifold (X, ω) . To define them, one introduces a compatible almost complex structure J and a perturbation term ν , and considers the maps $f : C \rightarrow X$ from a genus g complex curve C with n marked points which satisfy the pseudo-holomorphic map equation $\bar{\partial}f = \nu$ and represent a class $A = [f] \in H_2(X)$. The set of such maps, together with their limits, forms the compact space of stable maps $\overline{\mathcal{M}}_{g,n}(X, A)$. For each stable map, the domain determines a point in the Deligne-Mumford space $\overline{\mathcal{M}}_{g,n}$ of curves, and evaluation at each marked point determines a point in X . Thus there is a map

$$(0.1) \quad \overline{\mathcal{M}}_{g,n}(X, A) \rightarrow \overline{\mathcal{M}}_{g,n} \times X^n.$$

The Gromov-Witten invariant of (X, ω) is the homology class of the image for generic (J, ν) . It depends only on the isotopy class of the symplectic structure. By choosing bases of the cohomologies of $\overline{\mathcal{M}}_{g,n}$ and X^n , the GW invariant can be viewed as a collection of numbers that count the number of stable maps satisfying constraints. In important cases these numbers are equal to enumerative invariants defined by algebraic geometry.

In this article we construct Gromov-Witten invariants for a symplectic manifold (X, ω) relative to a codimension two symplectic submanifold V . These invariants are designed for use in formulas describing how GW invariants

*The research of both authors was partially supported by the N.S.F. The first author was also supported by a Sloan Research Fellowship.

behave under symplectic connect sums along V — an operation that removes V from X and replaces it with an open symplectic manifold Y with the symplectic structures matching on the overlap region. One expects the stable maps into the sum to be pairs of stable maps into the two sides which match in the middle. A sum formula thus requires a count of stable maps in X that keeps track of how the curves intersect V .

Of course, before speaking of stable maps one must extend J and ν to the connect sum. To ensure that there is such an extension we require that the pair (J, ν) be ‘ V -compatible’ as defined in Section 3. For such pairs, V is a J -holomorphic submanifold — something that is not true for generic (J, ν) . The relative invariant gives counts of stable maps for these special V -compatible pairs. These counts are different from those associated with the absolute GW invariants.

The restriction to V -compatible (J, ν) has repercussions. It means that pseudo-holomorphic maps $f : C \rightarrow V$ into V are automatically pseudo-holomorphic maps into X . Thus for V -compatible (J, ν) , stable maps may have domain components whose image lies entirely in V . This creates problems because such maps are not transverse to V . Worse, the moduli spaces of such maps can have dimension *larger* than the dimension of $\mathcal{M}_{g,n}(X, A)$. We circumvent these difficulties by restricting attention to the stable maps which *have no components mapped entirely into V* . Such ‘ V -regular’ maps intersect V in a finite set of points with multiplicity. After numbering these points, the space of V -regular maps separates into components labeled by vectors $s = (s_1, \dots, s_\ell)$, where ℓ is the number of intersection points and s_k is the multiplicity of the k^{th} intersection point. In Section 4 it is proved that each (irreducible) component $\mathcal{M}_{g,n,s}^V(X, A)$ of V -regular stable maps is an orbifold; its dimension depends on g, n, A and on the vector s .

The next step is to construct a space that records the points where a V -regular map intersects V and records the homology class of the map. There is an obvious map from $\mathcal{M}_{g,n,s}^V(X, A)$ to $H_2(X) \times V^\ell$ that would seem to serve this purpose. However, to be useful for a connect sum gluing theorem, the relative invariant should record the homology class of the curve in $X \setminus V$ rather than in X . These are additional data: two elements of $H_2(X \setminus V)$ represent the same element of $H_2(X)$ if they differ by an element of the set $\mathcal{R} \subset H_2(X \setminus V)$ of rim tori (the name refers to the fact that each such class can be represented by a torus embedded in the boundary of a tubular neighborhood of V). The subtlety is that this homology information is intertwined with the intersection data, and so the appropriate homology-intersection data form a covering space \mathcal{H}_X^V of $H_2(X) \times V^\ell$ with fiber \mathcal{R} . This is constructed in Section 5.

We then come to the key step of showing that the space \mathcal{M}^V of V -regular maps carries a fundamental homology class. For this we construct an orbifold compactification of \mathcal{M}^V — the space of V -stable maps. Since \mathcal{M}^V is a union

of open components of different dimensions the appropriate compactification is obtained by taking the closure of $\mathcal{M}_{g,n,s}^V(X, A)$ separately for each g, n, A and s . This is exactly the procedure one uses to decompose a reducible variety into its irreducible components. However, since we are not in the algebraic category, this closure must be defined via analysis.

The required analysis is carried out in Sections 6 and 7. There we study the sequences (f_n) of V -regular maps using an iterated renormalization procedure. We show that each such sequence limits to a stable map f with additional structure. The basic point is that some of the components of such limit maps have images lying in V , but along each component in V there is a section ξ of the normal bundle of V satisfying an elliptic equation $D^N \xi = 0$; this ξ ‘remembers’ the direction from which the image of that component came as it approached V . The components which carry these sections are partially ordered according to the rate at which they approach V as $f_n \rightarrow f$. We call the stable maps with this additional structure ‘ V -stable maps’. For each g, n, A and s the V -stable maps form a space $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$ which compactifies the space of V -regular maps by adding frontier strata of (real) codimension at least two.

This last point requires that (J, ν) be V -compatible. In Section 3 we show that for V -compatible (J, ν) the operator D^N commutes with J . Thus $\ker D^N$, when nonzero, has (real) dimension at least two. This ultimately leads to the proof in Section 7 that the frontier of the space of V -stable maps has codimension at least two. In contrast, for generic (J, ν) the space of V -stable maps is an orbifold with boundary and hence does not carry a fundamental homology class.

The endgame is then straightforward. The space of V -stable maps comes with a map

$$(0.2) \quad \overline{\mathcal{M}}_{g,n,s}^V(X, A) \rightarrow \overline{\mathcal{M}}_{g,n+\ell(s)} \times X^n \times \mathcal{H}_X^V$$

and relative invariants are defined in exactly the same way that the GW invariants are defined from (0.1). The new feature is the last factor, which allows us to control how the images of the maps intersect V . Thus the relative invariants give counts of V -stable maps with constraints on the complex structure of the domain, the images of the marked points, and the geometry of the intersection with V .

Section 1 describes the space of stable pseudo-holomorphic maps into a symplectic manifold, including some needed features that are not yet in the literature. These are used in Section 2 to define the GW invariants for symplectic manifolds and the associated invariants, which we call GT invariants, that count possible disconnected curves. We then bring in the symplectic submanifold V and develop the ideas described above. Sections 3 and 4 begin

with the definition of V -compatible pairs and proceed to a description of the structure of the space of V -regular maps. Section 5 introduces rim tori and the homology-intersection space \mathcal{H}_X^V .

For clarity, the construction of the space of V -stable maps is separated into two parts. Section 6 contains the analysis required for several special cases with increasingly complicated limit maps. The proofs of these cases establish all the analytic facts needed for the general case while avoiding the notational burden of delineating all ways that sequences of maps can degenerate. The key argument is that of Proposition 6.6, which is essentially a parametrized version of the original renormalization argument of [PW]. With this analysis in hand, we define general V -stable maps in Section 7, prove the needed transversality results and give the general dimension count showing that the frontier has sufficiently large codimension. In Section 8 the relative invariants are defined and shown to depend only on the isotopy class of the symplectic pair (X, V) . The final section presents three specific examples relating the relative invariants to some standard invariants of algebraic geometry and symplectic topology. Further applications are given in [IP4].

The results of this paper were announced in [IP3]. Related results are being developed by Eliashberg and Hofer [E] and Li and Ruan [LR]. Eliashberg and Hofer consider symplectic manifolds with contact boundary and assume that the Reeb vector field has finitely many simple closed orbits. When our case is viewed from that perspective, the contact manifold is the unit circle bundle of the normal bundle of V and all of its circle fibers – infinitely many – are closed orbits. In their first version, Li and Ruan also began with contact manifolds, but the approach in the most recent version of [LR] is similar to that of [IP3]. The relative invariants we define in this paper are more general than those of [LR] and appear, at least *a priori*, to give different gluing formulas.

Contents

1. Stable pseudo-holomorphic maps
 2. Symplectic invariants
 3. V -compatible perturbations
 4. Spaces of V -regular maps
 5. Intersection data and rim tori
 6. Limits of V -regular maps
 7. The space of V -stable maps
 8. Relative invariants
 9. Examples
- Appendix

1. Stable pseudo-holomorphic maps

The moduli space of (J, ν) -holomorphic maps from genus g curves with n marked points representing a class $A \in H_2(X)$ has a compactification $\overline{\mathcal{M}}_{g,n}(X, A)$. This comes with a map

$$(1.1) \quad \overline{\mathcal{M}}_{g,n}(X, A) \longrightarrow \overline{\mathcal{M}}_{g,n} \times X^n$$

where the first factor is the ‘‘stabilization’’ map st to the Deligne-Mumford moduli space (defined by collapsing all unstable components of the domain curve) and the second factor records the images of the marked points. The compactification carries a ‘virtual fundamental class’, which, together with the map (1.1), defines the Gromov-Witten invariants.

This picture is by now standard when X is a Kähler manifold. But in the general symplectic case, the construction of the compactification is scattered widely across the literature ([G], [PW], [P], [RT1], [RT2], [LT], [H] and [IS]) and some needed properties do not appear explicitly anywhere. Thus we devote this section to reviewing and augmenting the construction of the space of stable pseudo-holomorphic maps.

Families of algebraic curves are well-understood from the work of Mumford and others. A smooth genus g connected curve C with n marked points is *stable* if $2g + n \geq 3$, that is, if C is either a sphere with at least three marked points, a torus with at least one marked point, or has genus $g \geq 2$. The set of such curves, modulo diffeomorphisms, forms the Deligne-Mumford moduli space $\mathcal{M}_{g,n}$. This has a compactification $\overline{\mathcal{M}}_{g,n}$ that is a projective variety. Elements of $\overline{\mathcal{M}}_{g,n}$ are called ‘stable (g, n) -curves’; these are unions of smooth stable components C_i joined at d double points with a total of n marked points and Euler class $\chi(C) = 2 - 2g + d$. There is a universal curve

$$(1.2) \quad \overline{\mathcal{U}}_{g,n} = \overline{\mathcal{M}}_{g,n+1} \longrightarrow \overline{\mathcal{M}}_{g,n}$$

whose fiber over each point of $[j] \in \overline{\mathcal{M}}_{g,n}$ is a stable curve C in the equivalence class $[j]$ whenever $[j]$ has no automorphisms, and in general is a curve $C/\text{Aut}(C)$. To avoid these quotients we can lift to the moduli space of Prym structures as defined in [Lo]; this is a finite cover of the Deligne-Mumford compactification and is a manifold. The corresponding universal curve is a projective variety and is now a universal family, which we denote using the same notation (1.2). We also extend this construction to the unstable range by taking $\overline{\mathcal{M}}_{0,n} = \overline{\mathcal{M}}_{0,3}$ for $n \leq 2$ and $\overline{\mathcal{M}}_{1,0} = \overline{\mathcal{M}}_{1,1}$. We fix, once and for all, a holomorphic embedding of $\overline{\mathcal{U}}_{g,n}$ into some \mathbb{P}^N .

At this juncture one has a choice of either working throughout with curves with Prym structures, or working with ordinary curves and resolving the orbifold singularities in the Deligne-Mumford space whenever necessary by imposing Prym structures. Moving between the two viewpoints is straightforward;

see Section 2 of [RT2]. To keep the notation and discussion clear, we will consistently use ordinary curves, leaving it to the reader to introduce Prym structures when needed.

When one deals with maps $C \rightarrow X$ from a curve to another space one should use a different notion of stability. The next several definitions define ‘stable holomorphic maps’ and describe how they form a moduli space. We will use the term ‘special point’ to refer to a point that is either a marked point or a double point.

Definition 1.1. A *bubble domain* B of type (g, n) is a finite connected union of smooth oriented 2-manifolds B_i joined at double points together with n marked points, none of which are double points. The B_i , with their special points, are of two types:

- (a) stable components, and
- (b) unstable rational components, called ‘*unstable bubbles*’, which are spheres with a complex structure and one or two special points.

There must be at least one stable component. Collapsing the unstable components to points gives a *connected domain* $\text{st}(B)$ which is a stable genus g curve with n marked points.

Bubble domains can be constructed from a stable curve by replacing points by finite chains of 2-spheres. Alternatively, they can be obtained by pinching a set of nonintersecting embedded circles (possibly contractible) in a smooth 2-manifold. For our purposes, it is the latter viewpoint that is important. It can be formalized as follows.

Definition 1.2. A *resolution of a (g, n) bubble domain B with d double points* is a smooth oriented 2-manifold with genus g , d disjoint embedded circles γ_ℓ , and n marked points disjoint from the γ_ℓ , together with a map ‘resolution map’

$$r : \Sigma \rightarrow B$$

that respects orientation and marked points, takes each γ_ℓ to a double point of B , and restricts to a diffeomorphism from the complement of the γ_ℓ in B to the complement of the double points.

We can put a complex structure j on a bubble domain B by specifying an orientation-preserving map

$$(1.3) \quad \phi_0 : \text{st}(B) \rightarrow \overline{\mathcal{U}}_{g,n}$$

which is a diffeomorphism onto a fiber of $\overline{\mathcal{U}}_{g,n}$ and taking $j = j_\phi$ to be $\phi^* j_{\mathcal{U}}$ on the stable components of B and the standard complex structure on the unstable components. We will usually denote the complex curve (B, j) by the letter C .

We next define (J, ν) -holomorphic maps from bubble domains. These depend on the choice of an ω -compatible almost complex structure J (see (A.1) in the appendix), and on a ‘perturbation’ ν . This ν is chosen from the space of sections of the bundle $\text{Hom}(\pi_2^* T\mathbb{P}^N, \pi_1^* TX)$ over $X \times \mathbb{P}^N$ that are anti- J -linear:

$$\nu(j_P(v)) = -J(\nu(v)) \quad \forall v \in T\mathbb{P}^N$$

where j_P is the complex structure on \mathbb{P}^N . Let \mathcal{J} denote the space of such pairs (J, ν) , and fix one such pair.

Definition 1.3. A (J, ν) -holomorphic map from a bubble domain B is a map

$$(1.4) \quad (f, \phi) : B \longrightarrow X \times \overline{U}_{g,n} \subset X \times \mathbb{P}^N$$

with $\phi = \phi_0 \circ \text{st}$ as in (1.3) such that, on each component B_i of B , (f, ϕ) is a smooth solution of the inhomogeneous Cauchy-Riemann equation

$$(1.5) \quad \bar{\partial}_J f = (f, \phi)^* \nu$$

where $\bar{\partial}_J$ denotes the nonlinear elliptic operator $\frac{1}{2}(d + J_f \circ d \circ j_\phi)$. In particular, $\bar{\partial}_J f = 0$ on each unstable component.

Each map of the form (1.4) has degree (A, d) where $A = [f(B)] \in H_2(X; \mathbb{Z})$ and d is the degree of $\phi : \text{st}(B) \rightarrow \mathbb{P}^N$; $d \geq 0$ since ϕ preserves orientation and the fibers of \overline{U} are holomorphic. The ‘symplectic area’ of the image is the number

$$(1.6) \quad A(f, \phi) = \int_{(f, \phi)(B)} \omega \times \omega_{\mathbb{P}} = \int_B f^* \omega + \phi^* \omega_{\mathbb{P}} = \omega[A] + d$$

which depends only on the homology class of the map (f, ϕ) . Similarly, the energy of (f, ϕ) is

$$(1.7) \quad E(f, \phi) = \frac{1}{2} \int_B |d\phi|_\mu^2 + |df|_{J, \mu}^2 d\mu = d + \frac{1}{2} \int_B |df|_{J, \mu}^2 d\mu$$

where $|\cdot|_{J, \mu}$ is the norm defined by the metric on X determined by J and the metric μ on $\phi(B) \subset \mathbb{P}^N$. These integrands are conformally invariant, so the energy depends only on $[j_\phi]$. For $(J, 0)$ -holomorphic maps, the energy and the symplectic area are equal.

The following is the key definition for the entire theory.

Definition 1.4. A (J, ν) -holomorphic map (f, ϕ) is *stable* if each of its component maps $(f_i, \phi_i) = (f, \phi)|_{B_i}$ has positive energy.

This means that each component C_i of the domain is either a stable curve, or else the image of C_i carries a nontrivial homology class.

LEMMA 1.5. (a) *Every (J, ν) -holomorphic map has $E(f, \phi) \geq 1$.*

(b) *There is a constant $0 < \alpha_0 < 1$, depending only on (X, J) , such that every component (f_i, ϕ_i) of every stable (J, ν) -holomorphic map into X has $E(f_i, \phi_i) > \alpha_0$.*

(c) *Every (J, ν) -holomorphic map (f, ϕ) representing a homology class A satisfies*

$$E(f, \phi) \leq \omega(A) + C(3g - 3 + n)$$

where $C \geq 0$ is a constant which depends only on ν and the metric on $X \times \overline{\mathcal{U}}_{g,n}$ and which vanishes when $3g - 3 + n < 0$.

Proof. (a) If the component maps (f_i, ϕ_i) have degrees (d_i, A_i) then $E(f, \phi) = \sum E(f_i, \phi_i) \geq \sum d_i$ by (1.7). But $\sum d_i \geq 1$ because at least one component is stable.

(b) Siu and Yau [SY] showed that there is a constant α_0 , depending only on J , such that any smooth map $f : S^2 \rightarrow X$ that is nontrivial in homotopy satisfies

$$\frac{1}{2} \int_{S^2} |df|^2 > \alpha_0.$$

We may assume that $\alpha_0 < 1$. Then stable components have $E(f_i, \phi_i) \geq 1$ as above, and each unstable component either has $E(f_i, \phi_i) > \alpha_0$ or represents the trivial homology class. But in the latter case f_i is $(J, 0)$ -holomorphic, so $E(f_i, \phi_i) = A(f_i, \phi_i) = \omega[f_i] = 0$, contrary to the definition of stable map.

(c) This follows from straightforward estimates using (1.5) and (1.7), and the observation that curves in $\overline{\mathcal{M}}_{g,n}$ have at most $3g - 3 + n$ irreducible components. \square

Let $\mathcal{H}_{g,n}^{J,\nu}(X, A)$ denote the set of (J, ν) -holomorphic maps from a *smooth* oriented stable Riemann surface with genus g and n marked points to X with $[f] = A$ in $H_2(X; \mathbb{Z})$. Note that \mathcal{H} is invariant under the group $\text{Diff}(B)$ of diffeomorphisms of the domain that preserve orientation and marked points: if (f, ϕ) is (J, ν) -holomorphic then so is $(f \circ \psi, \phi \circ \psi)$ for any diffeomorphism ψ . Similarly, let $\overline{\mathcal{H}}_{g,n}^{J,\nu}(X, A)$ be the (larger) set of stable (J, ν) -holomorphic maps from a stable (g, n) bubble domain.

The main fact about (J, ν) -holomorphic maps — and the reason for introducing bubble domains — is the following convergence theorem. Roughly, it asserts that every sequence of (J, ν) -holomorphic maps from a smooth domain has a subsequence that converges modulo diffeomorphisms to a stable map. This result, first suggested by Gromov [G], is sometimes called the “Gromov Convergence Theorem”. The proof is the result of a series of papers dealing with progressively more general cases ([PW], [P], [RT1], [H], [IS]).

THEOREM 1.6 (Bubble Convergence). *Given any sequence (f_j, ϕ_j) of (J_i, ν_i) -holomorphic maps with n marked points, with $E(f_j, \phi_j) < E_0$ and $(J_i, \nu_i) \rightarrow (J, \nu)$ in C^k , $k \geq 0$, one can pass to a subsequence and find*

- (i) a (g, n) bubble domain B with resolution $r : \Sigma \rightarrow B$, and
- (ii) diffeomorphisms ψ_j of Σ preserving the orientation and the marked points,

so that the modified subsequence $(f_j \circ \psi_j, \phi_j \circ \psi_j)$ converges to a limit

$$\Sigma \xrightarrow{r} B \xrightarrow{(f, \phi)} X$$

where (f, ϕ) is a stable (J, ν) -holomorphic map. This convergence is in C^0 , in C^k on compact sets not intersecting the collapsing curves γ_ℓ of the resolution r , and the area and energy integrals (1.6) and (1.7) are preserved in the limit.

Under the convergence of Theorem 1.6, the image curves $(f_j, \phi_j)(B_j)$ in $X \times \mathbb{P}^N$ converge to $(f, \phi)(B)$ in the Hausdorff distance d_H , and the marked points and their images converge. Define a pseudo-distance on $\overline{\mathcal{H}}_{g,n}^{J,\nu}(X, A)$ by

$$(1.8) \quad d((f, \phi), (f', \phi')) = d_H(\phi(\Sigma), \phi'(\Sigma)) + d_H(f(\Sigma), f'(\Sigma)) \\ + \sum d_X(f(x_i), f'(x'_i))$$

where the sum is over all the marked points x_i . The space of stable maps, denoted

$$\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A) \quad \text{or} \quad \overline{\mathcal{M}}_{g,n}(X, A),$$

is the space of equivalence classes in $\overline{\mathcal{H}}_{g,n}^{J,\nu}(X, A)$, where two elements are equivalent if the distance (1.8) between them is zero. Thus orbits of the diffeomorphism group become single points in the quotient. We always assume the stability condition $2g + n \geq 3$.

The following structure theorem then follows from Theorem 1.6 above and the results of [RT1] and [RT2]. Its statement involves the canonical class K_X of (X, ω) and the following two terms.

Definition 1.7. (a) A symplectic manifold (X, ω) is called *semipositive* if there is no spherical homology class $A \in H_2(X)$ with $\omega(A) > 0$ and $0 < 2K_X[A] \leq \dim X - 6$.

(b) A stable map $F = (f, \phi)$ is *irreducible* if it is generically injective, i.e., if $F^{-1}(F(x)) = x$ for generic points x .

Let $\overline{\mathcal{M}}_{g,n}(X, A)^*$ be the moduli space of irreducible stable maps. Definition (1.7b) is equivalent to saying that the restriction of f to the union of the unstable components of its domain is generically injective (such maps are called *simple* in [MS]). Thus there are two types of reducible maps: maps whose re-

striction to some unstable rational component factors through a covering map $S^2 \rightarrow S^2$ of degree two or more, and maps with two or more unstable rational components with the same image.

THEOREM 1.8 (Stable Map Compactification). (a) $\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)$ is a compact metric space, and there are continuous maps

$$(1.9) \quad \mathcal{M}_{g,n}^{J,\nu}(X, A) \xrightarrow{\iota} \overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A) \xrightarrow{\text{st} \times \text{ev}} \overline{\mathcal{M}}_{g,n} \times X^n$$

where ι is an embedding, st is the stabilization map applied to the domain (B, j_ϕ) , and ev records the images of the marked points. The composition (1.9) is smooth.

(b) For generic (J, ν) , $\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)^*$ is an oriented orbifold of (real) dimension

$$(1.10) \quad -2K_X[A] + (\dim X - 6)(1 - g) + 2n.$$

Furthermore, each stratum $\mathcal{S}_k^* \subset \overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)^*$ consisting of maps whose domains have k double points is a suborbifold of (real) codimension $2k$.

(c) For generic (J, ν) , when X is semipositive or $\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)$ is irreducible, then the image of $\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)$ under $\text{st} \times \text{ev}$ carries a homology class.

The phrase ‘for generic (J, ν) ’ means that the statement holds for all (J, ν) in a second category subset of the space (3.3).

The manifold structure in (b) can be described as follows. Given a stable map (f, ϕ) with smooth domain B , choose a local trivialization $\mathcal{U}_{g,n} = \mathcal{M}_{g,n} \times B$ of the universal curve in a neighborhood U of $\phi(B)$. Then ϕ has the form $([j_\phi], \psi)$ for some diffeomorphism ψ of B , unique up to $\text{Aut}(B)$ (and unique when B has a Prym structure). Then

$$(1.11) \quad \mathcal{S}_\phi = \{(J, \nu)\text{-holomorphic } (f, \phi) \mid \phi = ([j_\phi], \text{id.})\}$$

is a slice for the action of the diffeomorphism group because any $(f', \phi') = (f', [j_{\phi'}], \psi)$ with $\phi'(B)$ in U is equivalent to $(f' \circ \psi^{-1}, [j_{\phi'}], \text{id.})$, uniquely as above. Thus the space of stable maps is locally modeled by the product of $\mathcal{M}_{g,n}$ and the set of (J, ν) -holomorphic maps from the fibers of the universal curve, which is a manifold as in [RT2].

The strata \mathcal{S}_k^* are orbifolds because with irreducible maps one can use variations in the pair (J, ν) to achieve the transversality needed to show that the moduli space is locally smooth and oriented for generic (J, ν) . This is proved in Lemma 4.9 in [RT1] and Theorem 3.11 in [RT2] (the proof also applies to irreducible maps with ghost bubbles, which are unnecessarily singled out in [RT1]). Moreover, the gluing theorem of Section 6 of [RT1] proves that \mathcal{S}_k^* has an orbifold tubular neighborhood in $\overline{\mathcal{M}}_{g,n}^{J,\nu}(X, A)^*$.

Theorem 1.8c was proved in [RT2] for semipositive (X, ω) by reducing the moduli space as follows. Every reducible stable map $f \in \overline{\mathcal{M}}_{g,n}(X, A)$ factors through an irreducible stable map $f_0 \in \overline{\mathcal{M}}_{g,n}(X, A_0)^*$ which has the same image as f , with the homology classes satisfying $\omega([f_0]) \leq \omega([f])$. Replacing each reducible f by f_0 yields a ‘reduced moduli space’ without reducible maps whose image under $\text{st} \times \text{ev}$ contains the image of the original moduli space. Semipositivity then implies that all boundary strata of the image of the reduced moduli space are of codimension at least 2.

Remark 1.9 (Stabilization). The semipositive assumption in Theorem 1.8c can be removed in several ways ([LT], [S], [FO], [R]), each leading to a moduli space which carries a “virtual fundamental class”, or at least whose image defines a homology class as in Theorem 1.8c. Unfortunately these approaches involve replacing the space of (J, ν) -holomorphic maps with a more complicated and abstract space. It is preferable, when possible, to work directly with (J, ν) -holomorphic maps where one can use the equation (1.5) to make specific geometric and P.D.E. arguments.

In a separate paper [IP5] we describe an alternative approach based on the idea of adding enough additional structure to insure that all stable (J, ν) -holomorphic maps are irreducible. More specifically, we develop a scheme for constructing a new moduli space $\tilde{\mathcal{M}}$ by consistently adding additional marked points to the domains and imposing constraints on them in such a way that (i) all maps in $\tilde{\mathcal{M}}$ are irreducible, and (ii) $\tilde{\mathcal{M}}$ is a finite (ramified) cover of the original moduli space. Theorem 1.8c then applies to $\tilde{\mathcal{M}}$ and hence $\tilde{\mathcal{M}}$, divided by the degree of the cover, defines a homology class.

2. Symplectic invariants

For generic (J, ν) the space of stable maps carries a fundamental homology class. For each g, n and A , the pushforward of that class under the evaluation map (1.1) or (1.9) is the ‘Gromov-Witten’ homology class

$$(2.1) \quad \left[\overline{\mathcal{M}}_{g,n}(X, A) \right] \in H_*(\overline{\mathcal{M}}_{g,n}; \mathbb{Q}) \otimes H_*(X^n; \mathbb{Q}).$$

A cobordism argument shows that this is independent of the choice of generic (J, ν) , and hence depends only on the symplectic manifold (X, ω) . Frequently, this Gromov-Witten invariant is thought of as the collection of numbers obtained by evaluating (2.1) on a basis of the dual cohomology group.

For our purposes it is convenient to assemble the GW invariants into power series in such a way that disjoint unions of maps correspond to products of the power series. We define those series in this section. Along the way we describe the geometric interpretation of the invariants.

Let $NH_2(X)$ denote the Novikov ring as in [MS]. The elements of $NH_2(X)$ are sums $\sum c_A t_A$ over $A \in H_2(X; \mathbb{Z})$ where $c_A \in \mathbb{Q}$, the t_A are variables satisfying $t_A t_B = t_{A+B}$, $c_A = 0$ if $\omega(A) < 0$, and where, for each $C > 0$ there are only finitely many nonzero coefficients c_A with energy $\omega(A) \leq C$. After summing on A and dualizing, (2.1) defines a map

$$(2.2) \quad \text{GW}_{g,n} : H^*(\overline{\mathcal{M}}_{g,n}) \otimes H^*(X^n) \rightarrow NH_2(X).$$

We can also sum over n and g by setting $\overline{\mathcal{M}} = \bigcup_{g,n} \overline{\mathcal{M}}_{g,n}$, letting $\mathbb{T}^*(X)$ denote the total (super)-tensor algebra $\mathbb{T}(H^*(X))$ on the rational cohomology of X , and introducing a variable λ to keep track of the Euler class. The total Gromov-Witten invariant of (X, ω) is then the map

$$(2.3) \quad \text{GW}_X : H^*(\overline{\mathcal{M}}) \otimes \mathbb{T}^*(X) \rightarrow NH_2(X)[\lambda].$$

defined by the Laurent series

$$(2.4) \quad \text{GW}_X = \sum_{A,g,n} \frac{1}{n!} \text{GW}_{X,A,g,n} t_A \lambda^{2g-2}.$$

The diagonal action of the symmetric group S_n on $\overline{\mathcal{M}}_{g,n} \times X^n$ leaves GW_X invariant up to sign, and if $\kappa \in H^*(\overline{\mathcal{M}}_{g,n})$ then $\text{GW}_X(\kappa, \alpha)$ vanishes unless α is a tensor of length n .

We can recover the familiar geometric interpretation of these invariants by evaluating on cohomology classes. Given $\kappa \in H^*(\overline{\mathcal{M}}; \mathbb{Q})$ and a vector $\alpha = (\alpha_1, \dots, \alpha_n)$ of rational cohomology classes in X of length $n = \ell(\alpha)$, fix a generic (J, ν) and generic geometric representatives K and A_i of the Poincaré duals of κ and of the α_i respectively. Then $\text{GW}_{X,A,g,n}(\kappa, \alpha)$ counts, with orientation, the number of genus g (J, ν) -holomorphic maps $f : C \rightarrow X$ with $C \in K$ and $f(x_i) \in A_i$ for each of the n marked points x_i . By the usual dimension counts, this vanishes unless

$$\deg \kappa + \sum \deg \alpha_i - 2\ell(\alpha) = (\dim X - 6)(1 - g) - 2K_X[A].$$

It is sometimes useful to incorporate the so-called ‘ ψ -classes’. There are canonically oriented real 2-plane bundles \mathcal{L}_i over $\overline{\mathcal{M}}_{g,n}(X, A)$ whose fiber at each map f is the cotangent space to the (unstabilized) domain curve at the i^{th} marked point. Let ψ_i be the Euler class of \mathcal{L}_i , and for each vector $D = (d_1, \dots, d_n)$ of nonnegative integers let $\psi_D = \psi_1^{d_1} \cup \dots \cup \psi_n^{d_n}$. Replacing the left-hand side of (2.1) by the pushforward of the cap product $\psi_D \cap [\overline{\mathcal{M}}_{g,n}(X, A)]$ and again dualizing gives invariants

$$(2.5) \quad \text{GW}_{X,g,n,D} : H^*(\overline{\mathcal{M}}_{g,n}) \otimes H^*(X^n) \rightarrow NH_2(X)$$

which agree with (2.2) when D is the zero vector. These invariants can be included in GW_X by adding variables in the series (2.4) which keep track of the vector D . To keep the notation manageable we will leave that embellishment to the reader.

The GW invariant (2.3) counts (J, ν) -holomorphic maps from *connected* domains. It is often more natural to work with maps whose domains are disjoint unions. Such J -holomorphic curves arose, for example, in Taubes' work on the Seiberg-Witten invariants ([T]). In fact, there is a simple and natural way of extending (2.3) to this more general case.

Let $\widetilde{\mathcal{M}}_{\chi, n}$ be the space of all compact Riemann surfaces of Euler characteristic χ with finitely many unordered components and with a total of n (ordered) marked points. For each such surface, after we fix an ordering of its components, the locations of the marked points define an ordered partition $\pi = (\pi_1, \dots, \pi_l) \in \mathcal{P}_n$. Hence

$$\widetilde{\mathcal{M}}_{\chi, n} = \bigsqcup_{\pi \in \mathcal{P}_n} \bigsqcup_{g_i} \left(\overline{\mathcal{M}}_{g_1, \pi_1} \times \dots \times \overline{\mathcal{M}}_{g_l, \pi_l} \right) / S_l$$

where \mathcal{P}_n is the set of all ordered partitions of the set $\{x_1, \dots, x_n\}$, $\overline{\mathcal{M}}_{g_i, \pi_i}$ is the space of stable curves with n_i marked points labeled by π_i , and where the second union is over all g_i with $\sum(2 - 2g_i) = \chi$. The symmetric group S_l acts by interchanging the components. Define the ‘‘Gromov-Taubes’’ invariant

$$(2.6) \quad \text{GT}_X : H^*(\widetilde{\mathcal{M}}) \otimes \mathbb{T}^*(X) \rightarrow NH_2(X)[\lambda]$$

by

$$(2.7) \quad \text{GT}_X = e^{\text{GW}_X}.$$

This exponential uses the ring structure on both sides of (2.6). Thus for $\alpha = \alpha_1 \otimes \dots \otimes \alpha_n$ and $\kappa = \kappa_1 \otimes \dots \otimes \kappa_l$,

$$\begin{aligned} & \text{GT}_{X, n}(\kappa, \alpha) \\ &= \sum_{\pi \in \mathcal{P}_n} \frac{\varepsilon(\pi)}{l!} \binom{n}{n_1, \dots, n_l} \text{GW}_{X, n_1}(\kappa_1, \alpha_{\pi_1}) \otimes \dots \otimes \text{GW}_{X, n_l}(\kappa_l, \alpha_{\pi_l}) \end{aligned}$$

where, for each partition $\pi = (\pi_1, \dots, \pi_l)$, α_{π_i} is the product of α_j for all $j \in \pi_i$ and $\varepsilon(\pi) = \pm 1$ depending on the sign of the permutation (π_1, \dots, π_l) and the degrees of α .

As before, when (2.7) is expanded as a Laurent series,

$$\text{GT}_X(\kappa, \alpha) = \sum_{A, \chi, n} \frac{1}{n!} \text{GT}_{X, A, \chi, n}(\kappa, \alpha) t_A \lambda^{-\chi},$$

the coefficients count the number of curves (not necessarily connected) with Euler characteristic χ representing A satisfying the constraints (κ, α) . Note that A and χ add when one takes disjoint unions, so that the variables t_A and λ multiply.

3. V -compatible perturbations

We now begin our main task: extending the symplectic invariants of Section 2 to invariants of (X, ω) relative to a codimension two symplectic submanifold V . Curves in X in general position will intersect such a submanifold V in a finite collection of points. Our relative invariants will still be a count of (J, ν) -holomorphic curves in X , but will also keep track of how those curves intersect V . But, instead of generic (J, ν) , they will count holomorphic curves for special (J, ν) : those ‘compatible’ to V in the sense of Definition 3.2 below.

Because (J, ν) is no longer generic, the construction of the space of stable maps must be thought through again and modified. That will be done over the next six sections. We begin in this section by developing some of the analytic tools that will be needed later.

The universal moduli space of stable maps $\overline{\mathcal{UM}}_{g,n}(X) \rightarrow \mathcal{J}$ is the set of all maps into X from *some* stable (g, n) curves which are (J, ν) -holomorphic for *some* $(J, \nu) \in \mathcal{J}$. If we fix a genus g two-manifold Σ , this is the set of (f, ϕ, J, ν) in $\text{Maps}(\Sigma, X \times \overline{\mathcal{U}}_{g,n}) \times \mathcal{J}$ with $\bar{\partial}_J f = \nu$. Equivalently, $\overline{\mathcal{UM}}_{g,n}(X)$ is the zero set of

$$(3.1) \quad \Phi(f, \phi, J, \nu) = \frac{1}{2} (df + J \circ df \circ j) - \nu$$

where j is the complex structure on the domain determined by ϕ . We will often abuse notation by writing j instead of ϕ .

In a neighborhood of (f, ϕ) the space of stable maps is modeled by the slice (1.11). Within that slice, the variation in ϕ lies in the tangent space to $\overline{\mathcal{M}}_{g,n}$, which is canonically identified with $H^{0,1}(TC)$ where C is the image of ϕ .

LEMMA 3.1. *The linearization of (3.1) at a point $(f, j, J, \nu) \in \mathcal{UM}_{g,n}$ is the elliptic operator*

$$D\Phi : \Gamma(f^*TX) \oplus H_j^{0,1}(TC) \oplus \text{End}(TX, J) \oplus \text{Hom}_J(T\mathbb{P}^N, TX) \rightarrow \Omega_j^{0,1}(f^*TX)$$

given by

$$D\Phi(\xi, k, K, \mu) = D_f(\xi, k) + \frac{1}{2}Kf_*j - \mu$$

where C is the domain of f and $D_f(\xi, k) = D\Phi(\xi, k, 0, 0)$ is defined by

$$(3.2) \quad D_f(\xi, k)(w) = \frac{1}{2} [\nabla_w \xi + J \nabla_{jw} \xi + (\nabla_\xi J)(f_*(jw)) + J f_* k(w)] \\ - (\nabla_\xi \nu)(w)$$

for each vector w tangent to the domain, where ∇ is the pullback connection on f^*TX .

Proof. The variations with respect to j, J and ν are obvious (cf. equation (3.9) in [RT2]), so we need only check the variation with respect to f . The calculation in [RT1, Lemma 6.3] gives

$$D_f(\xi, k)(w) = \frac{1}{2} \left[\nabla_w \xi + J \nabla_{jw} \xi + \frac{1}{2} (\nabla_\xi J)(f_*(jw) + J f_*(w)) + J f_* k(w) \right] \\ - (\nabla_\xi^J \nu)(w)$$

where $\nabla^J = \nabla + \frac{1}{2}(\nabla J)J$. By the equation $\Phi(f, j, J, \nu) = 0$, this agrees with (3.2). \square

As mentioned above, we will restrict attention to a subspace of \mathcal{J}^V consisting of pairs (J, ν) that are compatible with V in the following sense. Denote the orthogonal projection onto the normal bundle N_V by $\xi \mapsto \xi^N$; this uses the metric defined by ω and J and hence depends on J .

Definition 3.2. Let \mathcal{J}^V be the submanifold of \mathcal{J} consisting of pairs (J, ν) whose 1-jet along V satisfies the following three conditions:

$$(a) \quad J \text{ preserves } TV \text{ and } \nu^N|_V = 0,$$

and for all $\xi \in N_V, v \in TV$ and $w \in TC$

$$(3.3) \quad (b) \quad [(\nabla_\xi J + J \nabla_{J\xi} J)(v)]^N = [(\nabla_v J)\xi + J(\nabla_{Jv} J)\xi]^N, \\ (c) \quad [(\nabla_\xi \nu + J \nabla_{J\xi} \nu)(w)]^N = [(J \nabla_{\nu(w)} J)\xi]^N.$$

The first condition means that V is a J -holomorphic submanifold, and that (J, ν) -holomorphic curves in V are also (J, ν) -holomorphic in X . Conditions (b) and (c) relate to the variation of such maps; they are chosen to ensure that Lemma 3.3 below holds. Condition (b) is equivalent to the vanishing of some of the components of the Nijenhuis tensor N_J along V , namely that the normal component of $N_J(v, \xi)$ vanishes whenever v is tangent and ξ is normal to V . Thus (b) can be thought of as the ‘partial integrability’ of J along V .

For each (J, ν) -holomorphic map f whose image lies in V , we obtain an operator $D_f^N : \Gamma(f^*N_V) \rightarrow \Omega^{0,1}(f^*N_V)$ by restricting the linearization (3.2) to the normal bundle:

$$(3.4) \quad D_f^N(\xi) = [D_f(\xi, 0)]^N.$$

LEMMA 3.3. *Let $(J, \nu) \in \mathcal{J}^V$. Then for each (J, ν) -holomorphic map f whose image lies in V , D_f^N is a complex operator (that is, it commutes with J).*

Proof. Since J preserves the normal bundle, we must verify that $[D(J\xi) - JD(\xi)]^N = 0$ for each $\xi \in N_V$. By (3.2), the quantity $2J[D(J\xi) - JD(\xi)](w)$ is

$$(J\nabla_w J)\xi - (\nabla_{jw} J)\xi + \frac{1}{2}[\nabla_\xi J + J\nabla_{J\xi} J](f_*(jw) + Jf_*(w)) - 2[\nabla_\xi \nu + J\nabla_{J\xi} \nu](w).$$

After substituting $f_*(w) = 2\nu(w) - Jf_*(jw)$ into the first term and writing $v = f_*(jw)$, this becomes

$$\begin{aligned} & 2(J\nabla_{\nu(w)} J)\xi - (J\nabla_{Jv} J)\xi - (\nabla_v J)\xi + \nabla_\xi J(v) \\ & + J\nabla_{J\xi} J(v) - 2\nabla_\xi \nu(w) - 2J\nabla_{J\xi} \nu(w). \end{aligned}$$

Taking the normal component, we see that the sum of the second, third, fourth, and fifth terms vanishes by (3.3b), while the sum of the first, sixth, and seventh terms vanishes by (3.3c). \square

We conclude this section by giving a local normal form for holomorphic maps near the points where they intersect V . This will be used repeatedly later. The proof is adapted from McDuff [M].

Here is the context. Let V be a codimension two J -holomorphic submanifold of X and ν be a perturbation that vanishes in the normal direction to V as in (3.3). Fix a local holomorphic coordinate z on an open set \mathcal{O}_C in a Riemann surface C . Also fix local coordinates $\{v^i\}$ in an open set \mathcal{O}_V in V and extend these to local coordinates (v^i, x) for X with $x \equiv 0$ along V and so that $x = x^1 + ix^2$ along V with $J(\partial/\partial x^1) = \partial/\partial x^2$ and $J(\partial/\partial x^2) = -\partial/\partial x^1$.

LEMMA 3.4 (normal form). *Suppose that C is a smooth connected curve and $f : C \rightarrow X$ is a (J, ν) -holomorphic map that intersects V at a point $p = f(z_0) \in V$ with $z_0 \in \mathcal{O}_C$ and $p \in \mathcal{O}_V$. Then either (i) $f(C) \subset V$, or (ii) there is an integer $d > 0$ and a nonzero $a_0 \in \mathbb{C}$ so that in the above coordinates*

$$(3.5) \quad f(z, \bar{z}) = \left(p^i + O(|z|), a_0 z^d + O(|z|^{d+1}) \right)$$

where $O(|z|^k)$ denotes a function of z and \bar{z} that vanishes to order k at $z = 0$.

Proof. Let J_0 be the standard complex structure in the coordinates (v^i, x^α) . The components of the matrix of J then satisfy

$$(3.6) \quad (J - J_0)_j^i = O(|v| + |x|), \quad (J - J_0)_\beta^\alpha = O(|x|), \quad (J - J_0)_\alpha^i = O(|x|).$$

Set

$$A = (1 - J_0 J)^{-1}(1 + J_0 J) \quad \text{and} \quad \hat{\nu} = 2(1 - J J_0)^{-1} \nu.$$

With the usual definitions $\bar{\partial}f = \frac{1}{2}(df + J_0dfj)$ and $\partial f = \frac{1}{2}(df - J_0dfj)$, the (J, ν) -holomorphic map equation $\bar{\partial}_J f = \nu$ is equivalent to

$$(3.7) \quad \bar{\partial}f = A\partial f + \hat{\nu}.$$

Conditions (3.6) and the fact that the normal component of ν also vanishes along V give

$$A_j^i = O(|v| + |x|), \quad A_\beta^\alpha = O(|x|), \quad A_\alpha^i = O(|x|), \quad \nu^\alpha = O(|x|).$$

Now write $f = (v^i(z, \bar{z}), x^\alpha(z, \bar{z}))$. Because A_α^i vanishes along V and the functions $|dv^i|$ and $\partial A_\alpha^i / \partial x^\beta$ are bounded near z_0 , we obtain

$$\left| dA_\alpha^i \right| \leq \left| \frac{\partial A_\alpha^i}{\partial v^j} \cdot dv^j + \frac{\partial A_\alpha^i}{\partial x^\beta} \cdot dx^\beta \right| \leq c(|x| + |dx|).$$

Since ν^α also vanishes along V by Definition 3.3a, we get exactly the same bound on $|d\nu^\alpha|$. Returning to equation (3.7) and looking at the x components, we have

$$(3.8) \quad \bar{\partial}x^\alpha = A_i^\alpha \partial v^i + A_\beta^\alpha \partial x^\beta + \hat{\nu}^\alpha,$$

and hence

$$\partial \bar{\partial}x^\alpha = \partial A_i^\alpha \partial v^i + A_i^\alpha \partial^2 v^i + \partial A_\beta^\alpha \partial x^\beta + A_\beta^\alpha \partial^2 x^\beta + \partial \hat{\nu}^\alpha.$$

Because $\partial \bar{\partial}x^\alpha = 2\Delta x^\alpha$ and the derivatives of v and x are locally bounded this gives

$$|\Delta x^\alpha|^2 \leq c(|x|^2 + |\partial x|^2).$$

If x^α vanishes to infinite order at z_0 then Aronszajn's Unique Continuation theorem ([A, Remark 3]) implies that $x^\alpha \equiv 0$ in a neighborhood of z_0 , i.e. $f(C) \subset V$ locally. This statement is independent of coordinates. Consequently, the set of $z \in C$ where $f(z)$ contacts V to infinite order is both open and closed, so that $f(C) \subset V$. On the other hand, if the order of vanishing is finite, then $x^\alpha(z, \bar{z})$ has a Taylor expansion beginning with $\sum_{k=0}^d a_k \bar{z}^k z^{d-k}$ for some $0 < d < \infty$. Since $A_i^\alpha, A_\beta^\alpha$ and ν^α are all $O(|x|)$ and x is $O(|z|^d)$, (3.8) gives

$$\bar{\partial}x^\alpha = O(|x|) = O(|z|^d).$$

Differentiating, we conclude that the leading term is simply $a_0 z^d$. This gives (3.5). \square

4. Spaces of V -regular maps

We have chosen to work with holomorphic maps for (J, ν) compatible with V . For these special (J, ν) one can expect more holomorphic curves than are present for a completely general choice of (J, ν) . In particular, with our choice, any (J, ν) -holomorphic map into V is automatically holomorphic

as a map into X . Thus we have allowed stable holomorphic maps that are badly nontransverse to V — entire components can be mapped into V . We will exclude such maps and define the relative invariant using only ‘ V -regular’ maps.

Definition 4.1. A stable (J, ν) -holomorphic map into X is called *V -regular* if no component of its domain is mapped entirely into V and if none of the special points (i.e. marked or double points) on its domain are mapped into V .

The V -regular maps (including those with nodal domain) form an open subset of the space of stable maps, which we denote by $\mathcal{M}^V(X, A)$. In this section we will show how $\mathcal{M}^V(X, A)$ is a disjoint union of components, and how the irreducible part of each component is an orbifold for generic $(J, \nu) \in \mathcal{J}^V$.

Lemma 3.4 tells us that for each V -regular map f , the inverse image $f^{-1}(V)$ consists of isolated points p_i on the domain C distinct from the special points. It also shows that each p_i has a well-defined multiplicity s_i equal to the order of contact of the image of f with V at p_i . The list of multiplicities is a vector $s = (s_1, s_2, \dots, s_\ell)$ of integers $s_i \geq 1$. Let \mathcal{S} be the set of all such vectors and define the degree, length, and order of $s \in \mathcal{S}$ by

$$\deg s = \sum s_i, \quad \ell(s) = \ell, \quad |s| = s_1 s_2 \cdots s_\ell.$$

These vectors s label the components of $\mathcal{M}^V(X, A)$: associated to each s such that $\deg s = A \cdot V$ is the space

$$\mathcal{M}_{g,n,s}^V(X, A) \subset \mathcal{M}_{g,n+\ell(s)}(X, A)$$

of all V -regular maps f such that $f^{-1}(V)$ is exactly the marked points p_i , $1 \leq i \leq \ell(s)$, each with multiplicity s_i . Forgetting these last $\ell(s)$ points defines a projection

$$(4.1) \quad \begin{array}{c} \mathcal{M}_{g,n,s}^V(X, A) \\ \downarrow \\ \mathcal{M}_{g,n}^V(X, A) \end{array}$$

onto one component of $\mathcal{M}_{g,n}^V(X, A)$, which is the disjoint union of such components. Notice that for each s (4.1) is a covering space whose group of deck transformations is the group of renumberings of the last $\ell(s)$ marked points.

LEMMA 4.2. *For generic (J, ν) , the irreducible part of $\mathcal{M}_{g,n,s}^V(X, A)$ is an orbifold with*

$$(4.2) \quad \begin{aligned} \dim \mathcal{M}_{g,n,s}^V(X, A) = & -2K_X[A] + (\dim X - 6)(1 - g) \\ & + 2(n + \ell(s) - \deg s). \end{aligned}$$

Proof. We need only to show that the universal moduli space $\mathcal{UM}_{g,n,s}^*$ is a manifold (after passing to Prym covers); the Sard-Smale theorem then implies that for generic (J, ν) the moduli space $\mathcal{M}_{g,n,s}^V(X, A)^*$ is an orbifold of dimension equal to the (real) index of the linearization, which is precisely (4.2).

First, let $\mathcal{F}_{g,n}^V$ be the space of all data $(J, \nu, f, j, x_1, \dots, x_n)$ as in [RT2, Eq. (3.3)], but now taking f to be V -regular and $(J, \nu) \in \mathcal{J}^V$. Define Φ on $\mathcal{F}_{g,n}^V$ by $\Phi(J, \nu, f, j, \{x_i\}) = \bar{\partial}_{j,J} f - \nu$. The linearization $D\Phi$ is onto exactly as in equations (3.10) and (3.12) of [RT2], so that the universal moduli space $\mathcal{UM}_{g,n}^{V*} = \Phi^{-1}(0)$ is smooth and its dimension is given by (4.2) without the final ‘s’ terms.

It remains to show that the contact condition corresponding to each ordered sequence s is transverse; that will imply that $\mathcal{UM}_{g,n,s}^V(X, A)^*$ is a manifold. Consider the space $\text{Div}^d(C)$ of degree d effective divisors on C . This is a smooth manifold of complex dimension d . (Its differentiable structure is as described in [GH, p. 236]: given a divisor D_0 , choose local holomorphic coordinates z_k around the points of D_0 ; nearby divisors can be realized as the zeros of monic polynomials in these z_k and the coefficients of these polynomials provide a local chart on $\text{Div}^d(C)$.) Moreover, for each sequence s of degree d , let $\text{Div}_s(C) \subset \text{Div}^d(C)$ be the subset consisting of divisors of the form $\sum s_k y_k$. This is a smooth manifold of complex dimension $\ell(s)$.

For each sequence s of degree d define a map

$$\Psi_s : \mathcal{UM}_{g,n+\ell(s)}^V \longrightarrow \text{Div}^d(C) \times \text{Div}_s(C)$$

by

$$\Psi_s(J, \nu, f, j, \{x_i\}, \{y_k\}) = \left(f^{-1}(V), \sum_k s_k y_k \right)$$

where the y_k are the last $\ell(s)$ marked points. By Lemma 3.4, there are local coordinates z_l around the points $p_l \in C$ and $f(p_l) \in V \subset X$ such that the leading term of the normal component of f is $z_l^{d_l}$; hence

$$\Psi_s(J, \nu, f, j, \{x_i\}, \{y_k\}) = \left(\sum d_l p_l, \sum s_k y_k \right)$$

with $d_l \geq 1$, $\sum d_l = A \cdot V = d$. Let $\Delta \subset \text{Div}^d(C) \times \text{Div}_s(C)$ denote the diagonal of $\text{Div}_s(C) \times \text{Div}_s(C)$. Then

$$(4.3) \quad \mathcal{UM}_{g,n,s}^V = \Psi_s^{-1}(\Delta).$$

This is a manifold provided that Ψ_s is transverse to Δ . Thus it suffices to show that at each fixed $(J, \nu, f, j, \{x_i\}, \{y_k\}) \in \mathcal{UM}_{g,n,s}^V$ the differential $D\Psi$ is onto the tangent space of the first factor.

To verify that, we need only to construct a deformation

$$(J, \nu_t, f_t, j, \{x_i\}, \{y_k\})$$

that is tangent to $\mathcal{UM}_{g,n+\ell(s)}$ to first order in t , where the zeros of f_t^N are, to first order in t , the same as those of the polynomials $z_l^{d_l} + t\phi_l(z_l)$ where ϕ_l is an arbitrary polynomial in z_l of degree less than d_l defined near $z_l = 0$. In fact, by the linearity of $D\Psi$, it suffices to do this for $\phi_l(z_l) = z_l^k$ for each $0 \leq k < d_l$.

Choose smooth bump functions β_l supported in disjoint balls around the zeros of f with $\beta_l \equiv 1$ in a neighborhood of $z_l = 0$. For simplicity we fix l and omit it from the notation. We also fix local coordinates $\{v_j\}$ for V around $f(0)$, and extend these to coordinates (v_j, x) for X around $f(0)$, with V given locally as $x = x^1 + Jx^2 = 0$ as described before Lemma 3.4.

For any function $\eta(z)$ with $\eta(0) = 1$, we can construct maps

$$(4.4) \quad f_t = \left(f_0^T, f_0^N + t\beta z^k \eta \right).$$

It is easy to check that the zeros of the second factor have the form $z_t(1 + O(t))$ where the z_t are the zeros of $z^d + tz^k$. Then the variation \dot{f} at time $t = 0$ is $\xi = \beta z^k \eta e_N$, where e_N is a normal vector to V .

Keeping x, p, j, J fixed, we will show that we can choose η and a variation $\dot{\nu}$ in ν such that $(0, \dot{\nu}, \xi, 0, 0, 0)$ is tangent to $\mathcal{UM}_{g,n+\ell(s)}$. This requires two conditions on $(\xi, \dot{\nu})$.

(i) The variation in (J, ν) , which we are taking to be $(0, \dot{\nu})$, must be tangent to \mathcal{J}^V . Thus $\dot{\nu}$ must satisfy the linearization of equations (3.3), namely

$$\dot{\nu}^N = 0 \quad \text{and} \quad [\nabla_{e_N} \dot{\nu} + J \nabla_{J e_N} \dot{\nu}]^N(\cdot) = \left[(J \nabla_{\dot{\nu}(\cdot)} J) e_N \right]^N$$

along V , with e_N as above. This is true whenever $\dot{\nu}$, in the coordinates of Lemma 3.4, has an expansion off $x = 0$ of the form

$$(4.5) \quad \dot{\nu} = A(z, v) + B(z, v) \bar{x} + O(|z| |x|)$$

with $A^N = 0$ and $B^N = B^N(A) = \frac{1}{2} [J(\nabla_{A(\partial/\partial z)} J)(e_N)]^N$.

(ii) If $(0, \dot{\nu}, \xi, 0, 0, 0)$ is to be tangent to the universal moduli space it must be in the kernel of the linearized operator of Lemma 3.1, and so must satisfy

$$(4.6) \quad D\xi(z) - \dot{\nu}(z, f(z)) = 0$$

where D , which depends on f , is given in terms of the $\bar{\partial}$ operator of the pullback connection by

$$D\xi = \bar{\partial}_f \xi + \frac{1}{2} (\nabla_\xi J) df \circ j - \nabla_\xi \nu.$$

Near the origin in (z, v, x) coordinates, (4.5) is a condition on the 1-jet of $\dot{\nu}^N$ along the set where $x = 0$, and (4.6) is a condition along the graph $\{(z, v(z), z^d)\}$ of f_0 . Locally, these sets intersect only at the origin. Writing $\dot{\nu} = \dot{\nu}^V + \dot{\nu}^N$, we take

$$\dot{\nu}^V = [D\xi]^V$$

along the graph and extend it arbitrarily to a neighborhood of the origin. We

can then take $\dot{\nu}^N$ of the form (4.5) provided we can solve

$$(4.7) \quad D^N \xi(z) = \dot{\nu}^N(z, v(z), z^d) = B^N(D^V \xi(0)) \bar{z}^d + O(|z|^{d+1})$$

locally in a neighborhood of the origin with D^N as in (3.4).

Now, write $\xi = \alpha e_N + \beta J e_N$ where α and β are real, and identify this with $\xi = \zeta e_N$ where $\zeta = \alpha + i\beta$ is complex. Because D is an \mathbb{R} -linear first order operator, one finds that

$$(4.8) \quad D\xi = (\bar{\partial}\zeta) e_N + \zeta E + \bar{\zeta} F$$

where e_N is normal,

$$E = \frac{1}{2} [D(e_N) - JD(Je_N)] \quad \text{and} \quad F = \frac{1}{2} [D(e_N) + JD(Je_N)].$$

We need a solution of the form $\zeta = \beta z^k \eta$ near the origin. For this we can take $\beta \equiv 1$. The equation (4.7) we must solve has the form

$$(4.9) \quad -z^k \bar{\partial} \eta = z^k \eta E^N(z, \bar{z}) + \bar{z}^k \bar{\eta} F^N(z, \bar{z}) + B^N([D\xi(0)]^V) \bar{z}^d + O(|z|^{d+1}).$$

When $k = 0$, (4.9) has the form $\bar{\partial} \eta + a(z, \bar{z}) \eta + b(z, \bar{z}) \bar{\eta} = G(z, \bar{z})$, which can always be solved by power series. When $1 \leq k < d$, we have $\zeta(0) = 0$, so that $B^N([D\xi(0)]^V)$ vanishes by (4.8). Then using Lemma 4.3 below, (4.9) holds whenever η satisfies

$$-\bar{\partial} \eta = \eta E^N(z, \bar{z}) + a \bar{\eta} z^k z^{d-1-k} + O(|z|^{d+1-k}),$$

and this can also be solved by power series. \square

LEMMA 4.3. *Near the origin, $F^N = az^{d-1} + O(|z|^d)$ for some constant a .*

Proof. Fix a vector u tangent to the domain of f . Using the definition of F , equation (3.2), and the (J, ν) -holomorphic map equation $f_* u = 2\nu(u) - Jf_* ju$, one finds that $F^N(e_N)(u) = F^N(f_* u, u)$ where

$$(4.10) \quad \begin{aligned} 4F^N(U, u) &= J(\nabla_U J)e_N - (\nabla_{JU} J)e_N + (\nabla_{e_N} J)JU \\ &\quad + J(\nabla_{J e_N} J)JU + 2(\nabla_{J\nu(U)} J)e_N - 2(\nabla_{e_N} J)J\nu(u) \\ &\quad - 2(\nabla_{J e_N} J)\nu(u) - 2(\nabla_{e_N} \nu)u - 2J(\nabla_{J e_N} \nu)u. \end{aligned}$$

But the normal component of $U = f_* u$ is $dz^{d-1} \partial / \partial x$. Thus we can replace U in (4.10) by its component in the V direction; the difference has the form $z^{d-1} \Phi_1(z, \bar{z})$. In the resulting expression, the J is evaluated at the target point: $J := J(v(z), z^d)$. But

$$J(v(z), z^d) = J(v(z), 0) + O(|z|^d)$$

and similarly $\nabla J = (\nabla J)(v(z), 0) + O(|z|^d)$. Finally, with U tangent to V and J and ∇J replaced by their values at $(v(z), 0)$, one can check that (4.10) vanishes by (3.3). Lemma 4.3 follows. \square

5. Intersection data and rim tori

The images of two V -regular maps can be distinguished by (i) their intersection points with V , counted with multiplicity, and (ii) their homology classes $A \in H_2(X)$. One can go a bit further: if C_1 and C_2 are the images of two V -regular maps with the same data (i) and (ii), then the difference $[C_1 \# (-C_2)]$ represents a class in $H_2(X \setminus V)$. This section describes a space \mathcal{H}_X^V of data that include (i) and (ii) plus enough additional data to make this last distinction. Associating these data to a V -regular map then produces a continuous map

$$\mathcal{M}_{g,n}^V(X) \rightarrow \mathcal{H}_X^V.$$

It is this map, rather than the simpler map to the data (i) and (ii), that is needed for a gluing theorem for relative invariants ([IP4]).

We first need a space that records how V -regular maps intersect V . Recall that the domain of each $f \in \mathcal{M}_{g,n,s}^V$ has $n + \ell(s)$ marked points, the last $\ell(s)$ of which are mapped into V . Thus there is an *intersection map*

$$(5.1) \quad i_V : \mathcal{M}_{g,n,s}^V(X, A) \rightarrow V_s$$

that records the points and multiplicities where the image of f intersects V , namely

$$i_V(f, C, p_1 \dots, p_{n+\ell}) = ((f(p_{n+1}), s_1), \dots, (f(p_{n+\ell}), s_\ell)).$$

Here V_s is the space, diffeomorphic to $V^{\ell(s)}$, of all sets of pairs $((v_1, s_1), \dots, (v_\ell, s_\ell))$ with $v_i \in V$. This is, of course, simply the evaluation map at the last ℓ marked points, but cast in a form that keeps track of multiplicities.

To simplify notation, it is convenient to take the union over all sequences s to obtain the intersection map

$$(5.2) \quad i_V : \mathcal{M}_{g,n}^V(X, A) \longrightarrow \mathcal{S}V$$

where both

$$(5.3) \quad \mathcal{M}_{g,n}^V(X) = \coprod_A \coprod_s \mathcal{M}_{g,n,s}^V(X, A) \quad \text{and} \quad \mathcal{S}V = \coprod_s V_s$$

are given the topology of the disjoint union.

The next step is to augment $\mathcal{S}V$ with homology data to construct the space \mathcal{H}_X^V . The discussion in the first paragraph of this section might suggest taking \mathcal{H} to be $H_2(X \setminus V) \times \mathcal{S}V$. However, the above images C_1 and C_2 do not lie in $X \setminus V$ — only the difference does. In fact, the difference lies in

$$(5.4) \quad \mathcal{R} = \mathcal{R}_X^V = \ker [H_2(X \setminus V) \rightarrow H_2(X)].$$

Furthermore, there is a subtle twisting of these data, and \mathcal{H} turns out to be a nontrivial covering space over $H_2(X) \times \mathcal{S}V$ with \mathcal{R} acting as deck transforma-

tions — see (5.8) below. To clarify both these issues, we will compactify $X \setminus V$ and show how the images of V -regular maps determine cycles in a homology theory for the compactification.

Let $D(\varepsilon)$ be the ε -disk bundle in the normal bundle of V , identified with a tubular neighborhood of V . Choose a diffeomorphism of $X \setminus \overline{D(\varepsilon)}$ with $X \setminus V$ defined by the flow of a radial vector field and set $S = \partial \overline{D(\varepsilon)}$. Then

$$(5.5) \quad \widehat{X} = \left[X \setminus \overline{D(\varepsilon)} \right] \cup S$$

is a compact manifold with $\partial \widehat{X} = S$, and there is a projection $\pi : \widehat{X} \rightarrow X$ which is the projection $S \rightarrow V$ on the boundary and is a diffeomorphism in the interior.

The appropriate homology theory is built from chains which, like the images of V -regular maps, intersect V at finitely many points. Moreover, two cycles are homologous when they intersect V at the same points and their difference is trivial in $H_2(X \setminus V)$. We will give two equivalent descriptions of this homology theory.

For the first description, consider

- (i) the free abelian group C_k on k -dimensional simplices in \widehat{X} , and
- (ii) the subgroup D_k generated by the k -simplices that lie entirely in one circle fiber of $\partial \widehat{X}$.

Then $(C_*/D_*, \partial)$ is a chain complex over \mathbb{Z} . Let \mathcal{H} denote the 2-dimensional homology of this complex. Elements of $H_1(D_*)$ are linear combinations of the circle fibers of $\partial \widehat{X}$. Hence $H_1(D_*)$ can be identified with the space \mathcal{D} of divisors on V (a *divisor* is a finite set of points in V , each with sign and multiplicity). The long exact sequence of the pair (C_*, D_*) then becomes, in part,

$$(5.6) \quad 0 \longrightarrow H_2(\widehat{X}) \xrightarrow{\iota} H_2(C_*/D_*) \xrightarrow{\rho} \mathcal{D}.$$

For the second description we change the topology on X and \widehat{X} to separate cycles whose intersection with V is different. Let V^* be V with the discrete topology, and let S^* be S topologized as the disjoint union of its fiber circles. Then $\pi : S^* \rightarrow V^*$ and the inclusions $V^* \subset X$ and $S^* \subset \widehat{X}$ are continuous, and, when we use coefficients in \mathbb{Z} , $H_1(S^*)$ is identified with the space of divisors. The long exact sequence of the pair (\widehat{X}, S^*) again gives (5.6) with $H_2(\widehat{X}, S^*)$ in the middle. To fix notation we will use this second description.

The space in the middle of (5.6) is essentially the space of data we want. However, it is convenient to modify it in two ways. First, observe that projecting 2-cycles into X defines maps $\pi_* : H_2(\widehat{X}) \rightarrow H_2(X)$ and $\pi'_* : H_2(\widehat{X}, S^*) \rightarrow H_2(X)$ with $\pi_* = \pi'_* \circ \iota$. The kernel of π_* is exactly the space \mathcal{R} of (5.4), so that (5.6) can be rearranged to read

$$(5.7) \quad 0 \longrightarrow \mathcal{R} \xrightarrow{\iota} H_2(\widehat{X}, S^*) \xrightarrow{\rho} H_2(X) \times \mathcal{D}.$$

Second, in keeping with what we have done with V -regular maps, we can replace the space of divisors in (5.6) by the space $\mathcal{S}V$ of (5.3) which keeps track of the numbering of the intersection points, and whose topology separates strata with different multiplicity vectors s . There is a continuous covering map $\mathcal{S}V \rightarrow \mathcal{D}$ which replaces ordered points by unordered points. Pulling this covering back along the map $\rho : H_2(\widehat{X}, S^*) \rightarrow \mathcal{D}$ gives, at last, the desired space of data.

Definition 5.1. Let \mathcal{H}_X^V be the space $H_2(\widehat{X}, S^*) \times_{\mathcal{D}} \mathcal{S}V$.

With this definition, (5.7) lifts to a covering map

$$(5.8) \quad \begin{array}{ccc} \mathcal{R} & \longrightarrow & \mathcal{H}_X^V \\ & & \downarrow \varepsilon \\ & & H_2(X) \times \mathcal{S}V. \end{array}$$

where $H_2(X)$ has the discrete topology, $\mathcal{S}V$ is topologized as in (5.3), and $\varepsilon = (\pi'_*, \rho)$. This is also the right space for keeping track of the intersection-homology data: given a V -regular stable map f in X whose image is C , we can restrict to $X \setminus V$, lift to \widehat{X} , and take its closure, obtaining a curve \widehat{C} representing a class $[\widehat{C}]$ in $H_2(\widehat{X}, S^*)$. This is consistent with the intersection map (5.2) because $\rho[\widehat{C}] = \iota_V(f) \in \mathcal{D}$. Thus there is a well-defined map

$$(5.9) \quad h : \mathcal{M}_{g,n}^V(X) \longrightarrow \mathcal{H}_X^V$$

which lifts the intersection map (5.2) through (5.8). Of course, \mathcal{H}_X^V has components labeled by A and s , so this is a union of maps

$$(5.10) \quad h : \mathcal{M}_{g,n,s}^V(X, A) \longrightarrow \mathcal{H}_{X,A,s}^V$$

with $A \cdot V = \deg s$.

We conclude with a geometric description of elements of \mathcal{R} and of the twisting in the covering (5.8). Fix a small tubular neighborhood N of V in X and let π be the projection from the ‘rim’ ∂N to V . For each simple closed curve γ in V , $\pi^{-1}(\gamma)$ is a torus in ∂N ; such tori are called *rim tori*.

LEMMA 5.2. *Each element $R \in \mathcal{R}$ can be represented by a rim torus.*

Proof. Write X as the union of $X \setminus V$ and a neighborhood of V . Then the Mayer-Vietoris sequence

$$\longrightarrow H_2(\partial \widehat{X}) \xrightarrow{(\iota_*, \pi'_*)} H_2(X \setminus V) \oplus H_2(V) \longrightarrow H_2(X) \longrightarrow$$

shows that $(R, 0) = \iota_* \tau$ for some $\tau \in H_2(\partial \widehat{X})$ with $\pi_* \tau = 0$. The lemma then follows from the Gysin sequence for the oriented circle bundle $\pi : \partial \widehat{X} \rightarrow V$:

$$(5.11) \quad \longrightarrow H_3(V) \xrightarrow{\psi} H_1(V) \xrightarrow{\Delta} H_2(\partial\hat{X}) \xrightarrow{\pi_*} H_2(V) \longrightarrow$$

where ψ is given by the cap product with the first Chern class of the normal bundle to V in X . \square

Some rim tori are homologous to zero in $X \setminus V$ and hence do not contribute to \mathcal{R} . In fact, the proof of Lemma 5.2 shows that

$$\mathcal{R} = \text{image } [\iota_* \circ \Delta : H_1(V) \rightarrow H_2(X \setminus V)].$$

Now consider the image C of a V -regular map. Suppose for simplicity that C intersects V at a single point p with multiplicity one. Choose a loop $\gamma(t)$, $0 \leq t \leq 1$, in V with $\gamma(0) = \gamma(1) = p$, and let R be the rim torus $\pi^{-1}(\gamma)$. We can then modify C by removing the annulus of radius $\varepsilon/2 \leq r \leq \varepsilon$ around p in C and gluing in the rim torus R , tapered to have radius $\varepsilon(1 - t/2)$ over $\gamma(t)$. The resulting curve still intersects V only at p , but represents $[C] + [R]$. Thus this gluing acts as a deck transformation on $[C] \in \mathcal{H}_X^V$. Retracting the path γ , one also sees that each $\mathcal{H}_{X,A,s}^V$ is path connected.

Remark 5.3. There are no rim tori when $H_1(V) = 0$ or when the map $\iota_* \circ \Delta$ in (5.11) is zero. In that case \mathcal{H}_X^V is simply $H_2(X) \times \mathcal{S}V$. In practice, this makes the relative invariants significantly easier to deal with (see §9).

6. Limits of V -regular maps

In this and the next section we construct a compactification of each component of the space of V -regular maps. This compactification carries the “relative virtual class” that will enable us, in Section 8, to define the relative GW invariant.

One way to compactify $\mathcal{M}_{g,n,s}^V(X, A)$ is to take its closure

$$(6.1) \quad C\mathcal{M}_{g,n,s}^V(X, A)$$

in the space of stable maps $\overline{\mathcal{M}}_{g,n+\ell(s)}(X, A)$. Under the ‘bubble convergence’ of Theorem 1.6 the limits of the last $\ell(s)$ marked points are mapped into V . Thus the closure lies in the subset of $\overline{\mathcal{M}}_{g,n+\ell(s)}(X, A)$ consisting of stable maps whose last $\ell(s)$ marked points are mapped into V ; these still have associated multiplicities s_i , although the actual order of contact might be infinite.

The main step toward showing that this closure carries a fundamental homology class is to prove that the frontier $C\mathcal{M}^V \setminus \mathcal{M}^V$ is a subset of codimension at least two. For that, we examine the elements of $C\mathcal{M}^V$ and characterize those stable maps that are limits of V -regular maps. That characterization allows us to count the dimension of the frontier. The frontier is a subset of the space

of stable maps, so is stratified according to the type of bubble structure of the domain. Thus the goal of this section is to work towards a proof of the following statement about the structure of the closure CM^V .

PROPOSITION 6.1. *For generic $(J, \nu) \in \mathcal{J}^V$, each stratum of the irreducible part of*

$$CM_{g,n,s}^V(X, A) \setminus \mathcal{M}_{g,n,s}^V(X, A)$$

is an orbifold of dimension at least two less than the dimension (4.2) of $\mathcal{M}_{g,n,s}^V(X, A)$.

The closure CM^V contains strata corresponding to different types of limits. For clarity these will be treated in several separate steps:

Step 1: stable maps with no components or special points lying entirely in V ;

Step 2: a stable map with smooth domain which is mapped entirely into V ;

Step 3: maps with some components in V and some off V .

Step 1. For the strata consisting of stable maps with no components or special points in V the analysis is essentially standard (cf. [RT1]). Each stratum of this type is labeled by the genus and the number $d \geq 1$ of double points of their nodal domain curve B . Fix such a B . The corresponding stratum is the fiber of the universal space $\pi : \mathcal{UM}_{B,n,s}^V(X, A) \rightarrow \mathcal{J}^V$ of V -regular maps from B into X , and the irreducible part \mathcal{UM}^{V*} of \mathcal{UM}^V is an orbifold by the same transversality arguments as in [RT2].

LEMMA 6.2. *In this ‘Step 1’ case, for generic $(J, \nu) \in \mathcal{J}^V$, the irreducible part of the stratum $\mathcal{M}_{B,n,s}^V(X, A)$ of CM^V is an orbifold whose dimension is $2d$ less than the dimension (4.2) of $\mathcal{M}_{g,n,s}^V(X, A)$.*

Proof. Let $\tilde{B} \rightarrow B$ be the normalization of B . Then \tilde{B} is a (possibly disconnected) smooth curve with a pair of marked points for each double point of B . We will show that $\mathcal{UM}_{B,n,s}^V(X, A)^*$ is a suborbifold of $\mathcal{UM}_{g,n,s}^V(X, A)^*$ of codimension $2d$. Lemma 6.2 then follows by the Sard-Smale theorem.

Assume for simplicity that there is only one pair of such marked points (z_1, z_2) . Evaluation at z_1 and z_2 gives a map

$$\text{ev} : \mathcal{UM}_{B,n,s}^V(X, A)^* \rightarrow X \times X$$

and $\mathcal{UM}_{B,n,s}^V(X, A)^*$ is the inverse image of the diagonal Δ in $X \times X$. Since $\mathcal{UM}_{B,n,s}^V(X, A)^*$ is an orbifold, we need only check that this evaluation map is transversal to Δ .

To that end, fix $(f_0, J, \nu) \in \text{ev}^{-1}(\Delta)$. Choose local coordinates in X around $q = f_0(z_1) = f_0(z_2)$ and cutoff functions β_1 and β_2 supported in small

disks around z_1 and z_2 . Then, as in (4.4), we can modify f_0 locally around z_1 by $f_t = f_0 + t\beta_1 v$ and around z_2 by $f_t = f_0 - t\beta_2 v$, and modify ν to $\nu_t = \bar{\partial} f_t$ on the graph of f_t . The initial derivative of this path is a tangent vector w to $\mathcal{UM}_B^V(X)$ with $\text{ev}_*(w) = (v, -v)$. Thus ev is transversal to Δ . \square

Step 2. Consider the strata $C_1\mathcal{M}^V$ of \mathcal{CM}^V consists of all maps with smooth domain whose image is contained in V . Such maps lie in $\mathcal{M}_{g,n+\ell(s)}(V, A)$, and it might seem that we can focus on V and forget about X . But we are only examining the subset

$$C_1\mathcal{M}_{g,n,s}^V(X, A) \cap \mathcal{M}_{g,n+\ell(s)}(V, A)$$

that lies in the closure of $\mathcal{M}_{g,n,s}^V(X, A)$. The maps in this closure have a special property, stated as Lemma 6.3. This property involves the linearized operator.

For each $f \in \mathcal{M}_{g,n+\ell(s)}(V, A)$ denote by D^V the linearization of the equation $\bar{\partial} f = \nu$ at the map f . Note that the restriction map

$$\mathcal{J}^V \rightarrow \mathcal{J}(V)$$

that takes a compatible pair (J, ν) on X to its restriction to V is onto. Then by Theorem 4.2 for generic $(J, \nu) \in \mathcal{J}^V$ the irreducible part of the moduli space $\mathcal{M}_{g,n+\ell(s)}(V, A)$ is a smooth orbifold of (real) dimension

$$(6.2) \quad \text{index } D^V = -2K_V[A] + (\dim V - 6)(1 - g) + 2n + 2\ell(s).$$

There are several related operators associated with the maps f in this moduli space. First, there is the linearization D_s^X of the equation $\bar{\partial} f = \nu$; this acts on sections of f^*TX that have contact with V , described by the sequence s , and with index given by (4.2). Next, there is the operator D^N obtained by applying D^X to vector fields normal to V and then projecting back onto the subspace of normal vector fields. Completion in the Sobolev space with m derivatives in L^2 gives a bounded operator

$$(6.3) \quad D^N : L^{m,2}(f^*N_V) \rightarrow L^{m-1,2}(T^*C \otimes f^*N_V)$$

which is J -linear by Lemma 3.3. For $m > \deg s$ the sections that satisfy the linearization of the contact conditions specified by s form a closed J -invariant subspace $L_s^{m,2}(f^*N_V)$. Let D_s^N denote the restriction of D^N to that subspace $L_s^{m,2}$. The index of D_s^N is the index of D_s^X minus the index of D^V , so that

$$(6.4) \quad \text{index } D_s^N = 2(c_1(N_V)[A] + 1 - g - \deg s) = 2(1 - g)$$

since $\deg s = A \cdot V = c_1(N_V)[A]$.

LEMMA 6.3. *Each element of the closure $\mathcal{CM}_s^V(X)$ whose image is a single component that lies entirely in V is a map with $\ker D_s^N \neq 0$.*

Proof. This is seen by a renormalization argument similar to one in [T]. Suppose that a sequence $\{f_n\}$ in $\mathcal{M}_s^V(X)$ converges to $f \in \mathcal{M}(V)$; in the

present case there is no bubbling, so that $f_n \rightarrow f$ in C^∞ . For large n , the images of the maps f_n lie in a neighborhood of V , which we identify with a subset in the normal bundle N_V of V by the exponential map. Let ϕ_n be the projection of f_n to V along the fibers of N_V , so that $\phi_n \rightarrow f$ in C^∞ .

Next let $R_t : N_V \rightarrow N_V$ denote the dilation by a factor of $1/t$. Because the image of f_n is not contained in V there is, for each n , a unique $t = t_n$ for which the normal component of the pullback map $R_t(f_n)$ has C^1 norm equal to 1. These t_n are positive and $t_n \rightarrow 0$. Write R_{t_n} as R_n and consider the renormalized maps $F_n = R_n(f_n)$. These are holomorphic with respect to renormalized $(R_n^*J, R_n^*\nu)$; that is,

$$(6.5) \quad \bar{\partial}_{j, R_n^*J} F_n - R_n^* \nu = R_n^* (\bar{\partial}_{j, J} f_n - \nu) = 0.$$

By expanding in Taylor series one sees that $(R_n^*J, R_n^*\nu)$ converges in C^∞ to a limit (J_0, ν_0) ; this limit is dilation invariant and equal to the restriction of (J, ν) along V . The sequence $\{F_n\}$ is also bounded in C^1 . Therefore, after applying elliptic bootstrapping and passing to a subsequence, F_n converges in C^∞ to a limit F_0 which satisfies

$$\bar{\partial}_{j, J_0} F_0 - \nu_0 = 0.$$

We can also write F_n as $\exp_{\phi_n} \xi_n$ where $\xi_n \in \Gamma(\phi_n^* N_V)$ is the normal component of F_n , which has C^1 norm equal to 1. The above convergence implies that ξ_n converges in C^∞ to some nonzero $\xi \in \Gamma(f_0^* N_V)$. We claim that ξ is in the kernel of D_s^N along f_0 . In fact, since the f_n satisfy the contact constraints described by s and converge in C^∞ the limit ξ will have zeros described by s . Hence we need only show that $D_{f_0}^N \xi = 0$.

For fixed n , ϕ_n and $f_n = \exp_{\phi_n}(t_n \xi_n)$ are maps from the same domain so that by the definition of the linearization (for fixed J and ν)

$$P_n^{-1}(\bar{\partial}_J f_n - \nu_{f_n}) - (\bar{\partial}_J \phi_n - \nu_{\phi_n}) = D_{\phi_n}(t_n \xi_n, 0) + O(|t_n \xi_n|^2)$$

where P_n is the parallel transport along the curves $\exp_{\phi_n}(t \xi_n)$, $0 \leq t \leq t_n$. The first term in this equation vanishes because f_n is (J, ν) -holomorphic. Furthermore, because the image of ϕ_n lies in V , condition (3.3a) means that the normal component of $\bar{\partial}_J \phi_n - \nu_{\phi_n}$ vanishes. After dividing through by t_n and noting that $t_n^{-1} |t_n \xi_n|^2 \leq t_n$ we obtain

$$D_{f_0}^N \xi = \lim_{n \rightarrow \infty} D_{\phi_n}^N(\xi_n, 0) = 0. \quad \square$$

The operator D_s^N depends only on the 1-jet of $(J, \nu) \in \mathcal{J}^V$, so that we can consider the restriction map

$$(6.6) \quad \mathcal{J}^V \rightarrow \mathcal{J}^1$$

that takes a compatible pair (J, ν) on X to its 1-jet along V . This map is onto, and by Lemma 3.3, D_s^N is a complex operator for any $(J, \nu) \in \mathcal{J}^1$. Then D_s^N

defines a smooth section of

$$(6.7) \quad \begin{array}{c} \text{Fred} \\ \downarrow \\ \mathcal{UM}(V) \end{array}$$

where $\mathcal{UM}(V)$ is the universal moduli space of maps into V (which is a fiber bundle over \mathcal{J}^1), and where Fred is the bundle whose fiber at (f, j, J, ν) is the space of all complex linear Fredholm maps (6.3) of index $\iota \leq 0$. By a theorem of Koschorke [K], Fred is a disjoint union

$$\text{Fred} = \bigcup_k \text{Fred}_k$$

where Fred_k is the complex codimension $k(k - \iota)$ submanifold consisting of all the operators whose kernel is exactly k complex dimensional. In fact, the normal bundle to Fred_k in Fred, at an operator D , is $\text{Hom}(\ker D, \text{coker } D)$.

LEMMA 6.4. *The section D_s^N of (6.7) is transverse to each Fred_k .*

Proof. Fix $(f, j, J, \nu) \in \mathcal{UM}_s^V(X)$ such that the linearization D_s^N at (f, j, J, ν) lies on Fred_k . Let π^N be the projection onto the normal part, so that $D_s^N = \pi^N \circ D_s^X$. The lemma follows if we show that for any elements $\kappa \in \ker D_s^N$ and $c \in \ker (D_s^N)^*$ we can find a variation in (J, ν) such that

$$\langle c, (\delta D_s^N)\kappa \rangle \neq 0$$

(these brackets mean the L^2 inner product on the domain C and $(D_s^N)^*$ is the L^2 adjoint of D_s^N). But

$$(\delta D_s^N)\kappa = (\delta\pi^N)D_s^X\kappa + \pi^N(\delta D^X)\kappa + \pi^N D^X(\delta\pi^N)\kappa$$

with the linearization D^X is given by (3.2). We will take the variation with (f, j, J) fixed and ν varying as $\nu_t = \nu + t\mu$ with $\mu \equiv 0$ along V . Then π^N is fixed; i.e., it depends on J and f , but not on ν . Hence the above reduces to

$$(6.8) \quad \langle c, (\delta D^X)\kappa \rangle = -\langle c, \nabla_\kappa \mu \rangle.$$

This depends only on the 1-jet in the second variable of μ along V , where μ is the variation in $\nu(x, f(x))$.

Choose a point $x \in C$ such that $\kappa(x) \neq 0$. Let W be a neighborhood of x in \mathbb{P}^N and U a neighborhood of $f(x)$ in X such that κ has no zeros in U . To begin, c is defined only along the graph of f and is a $(0, 1)$ form with values in N_V . Extend c to a smooth section \tilde{c} of $\text{Hom}(T\mathbb{P}^N, TX)$ along $W \times U$ such that $\tilde{c}|_V$ is a section of $\text{Hom}(T\mathbb{P}^N, N_V)$. Multiply $\tilde{c}|_V$ by a smooth bump function β supported on $W \times U$ with $\beta \equiv 1$ on a slightly smaller open set.

Now construct the $(0, 1)$ form μ such that its 1-jet along V satisfies

$$\mu|_V = 0, \quad \nabla_{\kappa(y)}\mu(x, y) = (\beta\tilde{c})(x, y) \quad \text{and} \quad \nabla_{J\kappa(y)}\mu(x, y) = -J(\beta\tilde{c})(x, y).$$

The required compatibility conditions (3.3) are now satisfied because the right-hand side of (3.3c) vanishes since μ vanishes along V . Moreover,

$$\langle c, (\delta D_s^X)\kappa \rangle = - \int_C \langle c, \nabla_{\kappa}\mu \rangle = - \int_{C \cap U} \beta |c|^2.$$

But c satisfies the elliptic equation $(D_s^N)^*c = 0$, so by the unique continuation theorem for elliptic operators $|c|$ does not identically vanish on any open set. Thus we have found a nonzero variation. \square

PROPOSITION 6.5. *In this ‘Step 2’ case $C_1\mathcal{M}_{g,n,s}^V(X, A)$ is contained in the space*

$$(6.9) \quad \mathcal{M}'_{g,n,s} = \left\{ (f, j) \in \mathcal{M}_{g,n+\ell(s)}(V, A) \mid \dim \ker D_s^N \neq 0 \right\}.$$

Moreover, for generic $(J, \nu) \in \mathcal{J}^V$, the irreducible part of (6.9) is a suborbifold of $\mathcal{M}_{g,n+\ell(s)}(V, A)^*$ of dimension two less than (4.2).

Proof. The first statement follows from Lemma 6.3. Next, note that the dimension (6.2) of $\mathcal{M}_{g,n+\ell(s)}(V, A)$ differs from (4.2) by exactly the index (6.4) of D_s^N , so the second statement is trivially true if $\text{index } D_s^N > 0$. Thus we assume that $\iota = \text{index } D_s^N \leq 0$.

Lemma 6.4 implies that the set of pairs $(f, j, J, \nu) \in \mathcal{UM}_{g,n+\ell(s)}(V)$ for which D_s^N has a nontrivial kernel, namely

$$\mathcal{UM}' = D^{-1}(\text{Fred} \setminus \text{Fred}_0),$$

is a (real) codimension $2(1 - \iota)$ subset of $\mathcal{UM}_{g,n+\ell(s)}(V)$, and in fact a suborbifold off a set of codimension $4(2 - \iota)$. Since the projection $\pi : \mathcal{UM}' \rightarrow \mathcal{J}^1$ is Fredholm, the Sard-Smale theorem implies that for a second category set of $J \in \mathcal{J}^1$ the fiber $\pi^{-1}(J)$ — which is the space (6.9) — is an orbifold of (real) dimension

$$2 \text{index } D^V - 2(1 - \iota) = 2 \text{index } D^V + 2 \text{index } D_s^N - 2 = 2 \text{index } D_s^X - 2.$$

The inverse image of this second category set under (6.6) is a second category set in \mathcal{J}^V . Hence (6.9) is an orbifold for generic $(J, \nu) \in \mathcal{J}^V$, and has codimension at least two in $C\mathcal{M}_s^V(X)$. \square

Step 3. Next consider limit maps $f \in C\mathcal{M}_s^V(X)$ whose domain is the union $C = C_1 \cup C_2$ of bubble domains of genus g_1 and g_2 with f restricting to a V -regular map $f_1 : C_1 \rightarrow X$ and a map $f_2 : C_2 \rightarrow V$ into V . Limit maps f of this type arise, in particular, from sequences of maps in which either (a) two contact points collide in the domain or (b) one of the original n marked

points collides with a contact point because its image sinks into V . In either case the collision produces a ghost bubble map $f_2 : C_2 \rightarrow V$ which has energy at least α_V by Lemma 1.5.

In this Step 3 case, $f_1^{-1}(V)$ consists of the nodal points $C_1 \cap C_2$ and some of the last $\ell(s)$ marked points $p_k \in C$. The nodes are defined by identifying points $x_j \in C_1$ with $y_j \in C_2$. Since f_1 is V -regular and $f_1(x_j) \in V$, Lemma 3.4 associates a multiplicity s'_j to each x_j . Similarly, since f arises as a limit of V -regular maps the p_i , being limits of the contact points with V , have associated multiplicities. The set of p_i is split into the points $\{p_i^1\}$ on C_1 and $\{p_i^2\}$ on C_2 ; let $s^1 = (s_1^1, s_2^1, \dots)$ and $s^2 = (s_1^2, s_2^2, \dots)$ be the associated multiplicity vectors. Thus f is a pair

$$(6.10) \quad f = (f_1, f_2) \in \mathcal{M}_{g_1, n_1, s^1 \cup s'}^V(X, [f_1]) \times \overline{\mathcal{M}}_{g_2, n_2 + \ell(s^2) + \ell(s')} (V, [f_2])$$

with $n_1 + n_2 = n$, $[f_1] + [f_2] = A$, $\deg s^1 + \deg s' = [f_1] \cdot V$, and satisfying the matching conditions $f_1(x_j) = f_2(y_j)$.

PROPOSITION 6.6. *In this ‘Step 3’ case, the only elements (6.10) that lie in $\mathcal{CM}_s^V(X)$ are those for which there is a (singular) section $\xi \in \Gamma(f_2^* N_V)$ nontrivial on at least one component of C_2 with zeros of order s_i^2 at p_i^2 , poles of order s'_j at y_j (and nowhere else), and $D_{f_2}^N \xi = 0$ where $D_{f_2}^N$ is as in (3.4).*

The proof uses a renormalization argument similar to the one used in Lemma 6.3, but this time done in a compactification \mathbb{P}_V of the normal bundle $\pi : N_V \rightarrow V$. For clarity we describe \mathbb{P}_V before starting the proof.

Recall that N_V is a complex line bundle with an inner product and a compatible connection induced by the Riemannian connection on X . As a manifold \mathbb{P}_V is the fiberwise complex projectivization of the Whitney sum of N_V with the trivial complex line bundle

$$\pi_{\mathbb{P}} : \mathbb{P}_V = \mathbb{P}(N_V \oplus \mathbb{C}) \rightarrow V.$$

Note that the bundle map $\iota : N_V \hookrightarrow \mathbb{P}_V$ defined by $\iota(x) = [x, 1]$ on each fiber is an embedding onto the complement of the infinity section $V_{\infty} \subset \mathbb{P}_V$. The scalar multiplication map $R_t(\eta) = \eta/t$ on N_V defines a \mathbb{C}^* action on \mathbb{P}_V .

When V is a point we can identify \mathbb{P}_V with \mathbb{P}^1 and give it the Kähler structure $(\omega_{\varepsilon}, g_{\varepsilon}, j)$ of the 2-sphere of radius ε . Then $\iota : \mathbb{C} \rightarrow \mathbb{P}_V$ is a holomorphic map with $\iota^* g_{\varepsilon} = \phi_{\varepsilon}^2 [(dr)^2 + r^2(d\theta)^2]$ and $\iota^* \omega_{\varepsilon} = \phi_{\varepsilon}^2 r dr \wedge d\theta = d\psi_{\varepsilon} \wedge d\theta$ where

$$\phi_{\varepsilon}(r) = \frac{2\varepsilon}{1+r^2} \quad \text{and} \quad \psi_{\varepsilon}(r) = \frac{2\varepsilon^2 r^2}{1+r^2}.$$

This construction globalizes by interpreting r as the norm on the fibers of N_V , replacing $d\theta$ by the connection 1-form α on N_V and including the curvature

F_α of that connection. Thus

$$\iota^*\omega_\varepsilon = \pi^*\omega_V + \psi_\varepsilon\pi^*F_\alpha + d\psi_\varepsilon \wedge \alpha$$

is a closed form which is nondegenerate for small ε and whose restriction to each fiber of N_V agrees with the volume form on the 2-sphere of radius ε . Furthermore, at each point $p \in N_V$ the connection determines a horizontal subspace which identifies $T_p N_V$ with the fiber of $N_V \oplus TV$ at $\pi(p)$. But the fibers of N_V have a complex structure j_0 and a metric g_0 , and J and g on X restrict to V . One can then check that for small ε the form ω_ε ,

$$\tilde{J} = j_0 \oplus J|_V, \quad \text{and} \quad \tilde{g}_\varepsilon = (\phi_\varepsilon^2 g_0) \oplus g|_V$$

extend over V_∞ to define a tamed triple $(\omega_\varepsilon, \tilde{J}, \tilde{g}_\varepsilon)$ on \mathbb{P}_V . As in [PW], Lemma 1.5 holds for tamed structures, and so we can choose ε small enough that every J -holomorphic map f from S^2 onto a fiber of $\mathbb{P}_V \rightarrow V$ of degree $d \leq [f_1] \cdot V$ satisfies

$$(6.11) \quad \int_{S^2} |df|^2 \leq \frac{\alpha_V}{8}$$

where α_V is the constant associated with V by Lemma 1.5. We fix such an ε and write ω_ε as $\omega_{\mathbb{P}}$. Let V_0 denote the zero section of \mathbb{P}_V .

Now symplectically identify an $\varepsilon' < \varepsilon$ tubular neighborhood of V_0 in \mathbb{P}_V with a neighborhood of $V \subset X$ and pullback (J, g) from X to \mathbb{P}_V . Fix a bump function β supported on the ε' neighborhood of V_0 with $\beta = 1$ on the $\varepsilon'/2$ neighborhood. For each small $t > 0$ set $\beta_t = \beta \circ R_t$. Starting with the “background” metric $g' = \beta_t g + (1 - \beta_t)g_\varepsilon$, the procedure described in the appendix produces a compatible triple $(\omega_{\mathbb{P}}, J_t, g_t)$ on \mathbb{P}_V . Then as $t \rightarrow 0$ we have $J_t \rightarrow \tilde{J}$ in C^0 on \mathbb{P}_V and $g_t \rightarrow g_0$ on compact sets of $\mathbb{P}_V \setminus V_\infty$.

Proof of Proposition 6.6. Suppose that a sequence of V -regular maps $f_m : C_m \rightarrow X$ converges to $f = (f_1, f_2)$ as above. That means that the domains C_m converge to $C = C_1 \cup C_2$ and, as in Theorem 1.6, the f_m converge to $f : C \rightarrow X$ in C^0 and in energy, and C^∞ away from the nodes of C .

Around each node $x_j = y_j$ of $C_1 \cap C_2$ we have coordinates (z_j, w_j) in which C_m is locally the locus of $z_j w_j = \mu_{j,m}$ and C_1 is $\{z_j = 0\}$. Let $A_{j,m}$ be the annuli in the neck of C_m defined by $|\mu_{j,m}|/\delta \leq |z_j| \leq \delta$. We also let $C'_m \subset C_m$ denote the neck $A_m = \cup_j A_{j,m}$ together with everything on the C_2 side of A_m , f'_m be the restriction of f_m to C'_m , and let ϕ_m be the corresponding map into the universal curve as in (1.4).

The restrictions of f_m to $C'_m \setminus A_m$ converge to f_2 . Because the image of f_2 lies in V its energy is at least the constant α_V associated with V by Lemma 1.5. We can then fix δ small enough that the energy of $f = (f_1, f_2)$ inside the union of δ -balls around the nodes is at most $\alpha_V/32$. Then for large m

$$(6.12) \quad \int_{C'_m \setminus A_m} |d(\pi_V \circ f_m)|^2 + |d\phi_m|^2 \geq \alpha_V/2$$

and

$$\int_{A_m} |df_m|^2 + |d\phi_m|^2 \leq \alpha_V/16.$$

To renormalize, note that for large m the image of f'_m lies in a tubular neighborhood of V which is identified with a neighborhood of V_0 in \mathbb{P}_V . Hence f'_m gives rise to a one-parameter family of maps $\iota \circ R_t \circ f'_m$ into \mathbb{P}_V . We can consider the energy (1.7) of the corresponding map $(\iota \circ R_t \circ f'_m, \phi_m) : C'_m \rightarrow \mathbb{P}_V \times \bar{U}_{g,m}$ on the part of the domain which is mapped into the upper hemisphere \mathbb{P}_V^+ calculated using the metric \tilde{g}_ε on \mathbb{P}_V constructed above. That energy vanishes for large t and exceeds $\alpha_V/4$ for small t by (6.12). Therefore there is a unique $t = t_m$ such that the maps

$$g_m : C'_m \rightarrow \mathbb{P}_V \quad \text{by} \quad g_m = \iota \circ R_{t_m} \circ f_m$$

satisfy

$$(6.13) \quad \int_{g_m^{-1}(\mathbb{P}_V^+) \cup A_m} |dg_m|^2 + |d\phi_m|^2 = \alpha_V/4.$$

Note that $t_m \rightarrow 0$ because of (6.12) and the fact that $f_m(C'_m \setminus A_m) \rightarrow V$ pointwise.

Next consider the small annuli $B_{j,m}$ near $\partial C'_m$ defined by $\delta/2 \leq |w_j| \leq \delta$ and let $B_m = \cup_j B_{j,m}$. On each $B_{j,m}$ f_m converges in C^1 to $f_1 = a_j w_j^{s_j} + \dots$ and $f_m(B_{j,m})$ has small diameter. Hence, after possibly making δ smaller and passing to a subsequence, each $g_m(B_{j,m})$ lies in a coordinate neighborhood V_j centered at a point $q_j \in V_\infty$ with $\text{diam}^2(V_j) < \alpha_V/1000$. Fix a smooth bump function β on C_m which is supported on C'_m , satisfies $0 \leq \beta \leq 1$ and $\beta \equiv 1$ on $C'_m \setminus B_m$, and so that the integral of $|d\beta|^2$ over each $B_{j,m}$ is bounded by 100.

Now extend C'_m to a closed curve by smoothly attaching a disk D_j along the circle $\gamma_{j,m} = \{|w_j| = \delta\}$. Extend g_m to $\bar{g}_m : \bar{C}_m = C'_m \cup \{D_j\} \rightarrow \mathbb{P}_V$ by setting $\bar{g}_m(D_j) = q_j$ and coning off g_m on $B_{j,m}$ by the formula $\bar{g}_m = \beta \cdot g_m$ in the coordinates on V_j . The local expansion of f_1 shows that $f_m(\gamma_{j,m})$, oriented by the coordinate w_j , has winding number s_j around V_0 . The same is true of $g_m(\gamma_{j,m})$, so in homology $[\bar{g}_m]$ is $\iota_*[f_2] + sF$ where $s = \sum s_j$ and F is the fiber class of $\mathbb{P}_V \rightarrow V$.

By (6.13), the energy of \bar{g}_m on the region that is mapped into \mathbb{P}_V^+ is bounded by

$$(6.14) \quad \int_{g_m^{-1}(\mathbb{P}_V^+)} |dg_m|^2 + \sum_j \text{diam}^2(V_j) \int_{B_{j,m}} |d\beta|^2 \leq \frac{\alpha_V}{2}.$$

On the other hand, in the region mapped into \mathbb{P}_V^- , $\bar{g}_m = g_m$ is (J_m, ν_m) -holomorphic with $J_m \rightarrow \tilde{J}$ and $\nu_m \rightarrow \pi^* \nu_V$, so the energy in that region is

dominated by its symplectic area (1.6). Thus

$$\begin{aligned} E(\bar{g}_m) &\leq \frac{\alpha_V}{2} + c_1 \int_{g_m^{-1}(\mathbb{P}_V^-)} g_m^* \omega_{\mathbb{P}} \\ &\leq \frac{\alpha_V}{2} + c_1 \langle \omega_{\mathbb{P}}, [\bar{g}_m] \rangle + c_1 \int_{g_m^{-1}(\mathbb{P}_V^+)} |\bar{g}_m^* \omega_{\mathbb{P}}|. \end{aligned}$$

With (6.14) this gives a uniform energy bound of the form $E(\bar{g}_m) \leq c_1 \langle \omega, [f_2] + sF \rangle + c_2$.

This energy bound applies, *a fortiori*, to the restrictions g'_m of g_m to $C'_m \setminus B_m$. These g'_m are (J_m, ν_m) -holomorphic, so Theorem 1.6 provides a subsequence which converges to a $(\tilde{J}, \pi^* \nu_V)$ -holomorphic map whose domain is C_2 together with the disks $\{|w_j| \leq \delta/2\}$ in C_1 and possibly some bubble components.

After deleting those disks, the limit is a map $g_0 : \tilde{C}_2 \rightarrow \mathbb{P}_V$ with $g_0(y_j) \in V_\infty$ at marked points y_j . By construction, the projections $\pi \circ g'_m$ converge to f_2 , so the irreducible components of \tilde{C}_2 are of two types: (i) those biholomorphically identified with components of C_2 on which g_0 is a lift of f_2 to \mathbb{P}_V , and (ii) those mapped by g_0 into fibers of \mathbb{P}_V and also collapsed by the stabilization $\tilde{C}_2 \mapsto \text{st}(\tilde{C}_2)$. Then (6.13) implies that no type (i) component is mapped to V_∞ . The type (ii) components are $(J, 0)$ -holomorphic and on them $|d\phi|^2 \equiv 0$, so by (6.11) these components contribute a total of at most $\alpha_V/8$ to the integral (6.13). Thus (6.11) implies that at least one component of type (i) is not mapped into V_0 .

Lemma 3.4 shows each component of g_0 has a local expansion normal to V_∞ given by $b_j z_j^{d_j} + \dots$ at each y_j . To identify d_j we note that $\partial A_{j,m} = \gamma_{j,m} \cup \gamma'_{j,m}$ where $\gamma'_{j,m}$ is the circle $|z_j| = \delta$ oriented by z_j . The homology $g_m(A_{j,m}) \subset \mathbb{P}_V \setminus V_\infty$ then shows that d_j , which is the local winding number of $g_m(\gamma'_{j,m})$ with V_∞ , is equal to the local winding number of $g_m(\gamma_{j,m})$ with V_∞ , which is s_j .

The convergence $g'_m \rightarrow g_0$ on C_2 means that the sections $\xi_m = \iota^{-1} g_m$ of $f_2^* N_V$ converge to a nonzero $\xi = \iota^{-1} g_0$. Then $D^N \xi = 0$ as in the proof of Lemma 6.3, and our intersection number calculation shows that ξ has a pole of order s_j at each node y_j . Furthermore, the g_m have the same zeros, with multiplicity, as the f_m , so the zeros of ξ are exactly the last $\ell(s)$ marked points of the limit curve C_2 and the multiplicity vector associated with those zeros is the original s . Thus ξ is a nonzero element of $\ker D_{s,s'}^N$. \square

Proposition 6.6 shows that maps of the form (6.10) which are in the closure of $\mathcal{CM}_s^V(X)$ carry a special structure: a nonzero element ξ in the kernel of D_f^N with specified poles and zeros, defined on some component that is mapped into V . That adds constraints which enter the dimension counts needed to prove Proposition 6.1. In fact the proof shows that ξ vanishes only on those

components which sink into V_0 as the renormalized maps g_n converge. On those components we can renormalize again and proceed inductively. But instead of continuing down this road of special cases, we will define the special structure in the general case of maps with many components. Those maps form a space of ‘ V -stable maps’, and we will then do the dimension count once and for all in that context.

7. The space of V -stable maps

In the general case, the limit of a sequence of V -regular maps is a stable map whose components are of the types described in Steps 1–3 of Section 6. The components of the limit map are also partially ordered according to the rate at which they sink into V . In this section we introduce terminology which makes this precise, and then construct a compactification for the space of V -regular maps.

Let C be a stable curve. A *layer structure* on C is the assignment of an integer $\lambda_j = 0, 1, \dots$ to each irreducible component C_j of C . At least one component must have $\lambda_j = 0$ or 1. The union of all the components with $\lambda_j = k$ is the *layer k stable curve* $B_k \subset C$. Note that B_k might not be a connected curve.

Definition 7.1. A *marked layer structure* on $C \in \overline{\mathcal{M}}_{g,n+\ell}$ is a layer structure on C together with

- (i) a vector s giving the multiplicities of the last $\ell = \ell(s)$ marked points, and
- (ii) a vector t that assigns multiplicities to each double point of $B_k \cap B_l$, $k \neq l$.

Each layer B_k then has points $p_{k,i}$ of type (i) with multiplicity vector $s_k = (s_{k,i})$, and has double points with multiplicities. The double points separate into two types. We let t_k^+ be the vector derived from t that gives the multiplicities of the double points $y_{k,i}^+$ where B_k meets the higher layers, i.e. the points $B_k \cap C_j$ with $\lambda_j > k$. Let t_k^- be the similar vector of multiplicities of the double points $y_{k,i}^-$ where B_k meets the lower layers. Note that the double points within a layer are not assigned a multiplicity.

There are operators D_k^N akin to (6.3) defined on the layers B_k , $k \geq 1$, as follows. The marked points $y_{k,i}^-$ define $\ell(t_k^-)$ disjoint sections of the universal curve $\overline{\mathcal{U}}_{g,n+\ell} \rightarrow \overline{\mathcal{M}}_{g,n+\ell}$; in fact by compactness those sections have disjoint tubular neighborhoods. For each choice of $t = t_k^-$ and α , fix smooth weighting functions $W_{t,\alpha}$ whose restriction to each fiber of the universal curve has the

form $|z_j|^{\alpha+t_{k,j}^-}$ in some local coordinates z_j centered on $y_{k,j}^-$ and has no other zeros. Then given a stable map $f : B_k \rightarrow V$ let $L_{t,\delta}^m(f^*N_V)$ be the Hilbert space of all L_{loc}^m sections f^*N_V over $B_k \setminus \{y_{k,j}^-\}$ which are finite in the norm

$$\|\xi\|_{m,t,\delta}^2 = \sum_{l=0}^m \int_{B_k} |W_{t,l+\delta} \cdot \nabla^l \xi|^2.$$

For large m the elements ξ in this space have poles with $|\xi| \leq c|z_j|^{-t_j^- - \delta}$ at each $y_{k,j}^-$ and have $m-1$ continuous derivatives elsewhere on B_k . For such m let $L_{k,\delta}^m(f^*N_V)$ be the closed subspace of $L_{t_k,\delta}^m(f^*N_V)$ consisting of all sections that vanish to order $s_{k,i}$ at $p_{k,i}$ and order $t_{k,i}^+$ at $y_{k,i}^+$. By standard elliptic theory for weighted norms (cf. [L]) the operator D^N defines a bounded operator

$$(7.1) \quad D_k^N : L_{k,\delta}^m(f^*N_V) \rightarrow L_{k,\delta+1}^{m-1}(T^*C \otimes f^*N_V)$$

which, for generic $0 < \delta < 1$, is Fredholm with

$$\text{index}_{\mathbb{R}} D_k^N = 2c_1(N_V)A_k + \chi(B_k) + 2(\deg t_k^- - \deg s_k - \deg t_k^+) = \chi(B_k)$$

where $A_k = [f(B_k)]$ in $H_2(X)$. We used the fact that $c_1(N_V)A_k = \deg s_k + \deg t_k^+ - \deg t_k^-$ (since the Euler class of a line bundle can be computed from the zeros and poles of a section). Lemma 3.3 implies that the kernel of this operator is J -invariant, and so we can form the complex projective space $\mathbb{P}(\ker D_k^N)$.

Definition 7.2. A V -stable map is a stable map $(f, \phi) \in \overline{\mathcal{M}}_{g,n+\ell(s)}(X, A)$ together with

- (a) a marked layer structure on its domain C with $f|_{B_0}$ being V -regular, and
- (b) for each $k \geq 1$ an element $[\xi_k]$ of $\mathbb{P}(\ker D_k^N)$ defined on the layer B_k by a section ξ_k that is nontrivial on every irreducible component of B_k .

Let $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$ denote the set of all V -stable maps. This contains the set $\mathcal{M}_{g,n,s}^V(X, A)$ of V -regular maps as the open subset — the V -stable maps whose entire domain lies in layer 0. Forgetting the data $[\xi_k]$ defines a map

$$(7.2) \quad \overline{\mathcal{M}}_{g,n,s}^V(X, A) \xrightarrow{\beta} \overline{\mathcal{M}}_{g,n+\ell(s)}(X, A).$$

Each V -stable map $(f, \phi, [\xi_1], \dots, [\xi_r])$ determines an element of the space \mathcal{H}_X^V of Definition 5.1 as follows. For a very small ε , we can push the components in V off V by composing f with $\exp(\varepsilon^k \xi_k)$ and, for each k , smoothing the domain at the nodes $B_k \cap \left(\bigcup_{l>k} B_l \right)$ and smoothly joining the images where the zeros of $\varepsilon^k \xi_k$ on B_k approximate the poles of $\varepsilon^{k+1} \xi_{k+1}$. The resulting map

$$f_{\xi} = f|_{B_0} \# \exp(\varepsilon \xi_1) \# \dots \# \exp(\varepsilon^r \xi_r)$$

is V -regular, and so represents a homology class $h(f, \phi, [\xi_k]) = h(f_\xi) \in \mathcal{H}_X^V$ under (5.9). That class depends only on $[\xi]$: for different choices of the ε_k and of representatives of the $[\xi_k]$, the $\varepsilon^k \xi_k$ are homotopic through nonzero elements of the kernel with the same zeros and poles and hence represent the same element of \mathcal{H}_X^V . Thus there is a well-defined map

$$(7.3) \quad \overline{\mathcal{M}}_{g,n,s}^V(X, A) \xrightarrow{h} \mathcal{H}_{X,A,s}^V.$$

PROPOSITION 7.3. *There exists a topology on $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$ which makes it compact and for which the maps β of (7.2) and h of (7.3) are continuous and differentiable on each stratum.*

Proof. There are three steps to the proof. The first looks at sequences of V -regular maps (which are V -stable maps with trivial layer structure) and the second analyzes a general sequence of V -stable maps. The third step uses that analysis to define the topology on $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$.

Let $f_m : C_m \rightarrow X$ be a sequence of maps in $\mathcal{M}_{g,n,s}^V(X, A)$. By the bubble tree convergence Theorem 1.6, a subsequence, still called f_m , converges to a stable map $f : C \rightarrow X$. By successive renormalizations we will give the limit map f the structure of a V -stable map $(f, [\xi])$.

Since the last $\ell(s)$ marked points converge, the multiplicity vector s of f_m carries over to the limit, defining the vector s of Definition 7.1b. The rest of the layered structure is defined inductively. We assign $\lambda_j = 0$ to each component C_j unless $f(C_j) \subset V$, so that the layer B_0 consists of all components that are not mapped into V . Let $C(1)$ be the union of those components of C not in layer 0. Assign to each double point y of $B_0 \cap C(1)$ a multiplicity t_y equal to the order of contact of $f|_{B_0}$ with V at y .

Now apply the argument of Proposition 6.6. That produces renormalized maps $g_{m,1} = \exp \xi_{m,1}$ which converge to a nontrivial element of ξ_1 in $\ker D^N$ on $C(1)$. We assign $\lambda_j = 1$ to each component $C_j \subset C(1)$ on which ξ_1 is nonzero and denote the union of the remaining components by $C(2)$. Then ξ_1 is defined and nonzero on every component of B_1 . Moreover,

- (a) ξ_1 vanishes at the double points y where B_1 meets $C(2)$. We assign such y a multiplicity t_y equal to the order of vanishing of ξ_1 at that point.
- (b) As in the proof of Proposition 6.6, ξ_1 has a pole of order t_x at each $x \in B_0 \cap B_1$ and vanishes to order s_{1i} at the points p_{1i} in $B(1)$.
- (c) $\xi_{m,1} \rightarrow \xi_1$ and hence f_m and $f|_{B_0} \# \exp(\varepsilon \xi_1) \# g_{m,1}|_{C(2)}$ define the same element of \mathcal{H}_X^V for large m .

This defines $[\xi_1]$ and multiplicity vectors s and t on B_1 .

Next, repeat the renormalization on $C(2)$ and continue. This inductively defines a layer structure on C , multiplicities t for each point in $B_k \cap C(k+1)$, and determines a nontrivial element ξ_k of the kernel of D_k^N on each layer B_k . This process terminates because each $C(k)$ has fewer components than $C(k-1)$ while parts (a) and (c) of Lemma 1.5 give a uniform bound on the number of components. The end result is a nontrivial ξ_k on every component of B_k , $k \geq 1$. From (c) above we see that $\lim[f_m] = h(f, [\xi_k])$ in $\mathcal{H}_{X,A,s}^V$.

We next consider general sequences in $\overline{\mathcal{M}}_{g,n,s}^V(X,A)$. Given a sequence $F_m = (f_m, [\xi_{m,k}])$ of V -stable maps, we first form the maps $f_{m,0}$ obtained by restricting f_m to its bottom layer $B_m(0)$. The $f_{m,0}$ may represent different homology classes, but there is a uniform bound on their energy, so that the iterated renormalization argument of Lemma 7.3 produces a subsequence converging to a V -stable map F_0 . Similarly, the restrictions of f_m to $B(1)$ converge to a stable map f_1 into V and the renormalized maps $\exp(\xi_{m,1}) : B_m(1) \rightarrow \mathbb{P}_V$ have a subsequence converging to a limit which, on its bottom layer, has the form $\exp(\xi_{0,1})$ with $\xi_{0,1} \in \ker D^N$. Then $F_1 = (f_{m,1}, [\xi_{0,1}])$ is a V -stable map whose image lies in V and whose bottom layer fits with the top layer of F_0 to form a V -stable map $F_0 \cup F_1$. This process continues, and terminates because each layer carries energy at least α_V .

Finally, observe that this renormalization process can be read differently: it actually *defines* a notion of a convergence sequence of V -stable maps. Convergence in that sense defines a topology on $\overline{\mathcal{M}}_{g,n,s}^V(X,A)$, which we adopt as *the* topology on the space of V -stable maps. Reinterpreted, the above analysis shows that $\overline{\mathcal{M}}_{g,n,s}^V(X,A)$ is compact and h is continuous with that topology. \square

The next theorem is the key result needed to define the relative invariants; it implies and supersedes Proposition 6.1.

THEOREM 7.4. *The space of V -stable maps is compact and there is a continuous map*

$$(7.4) \quad \varepsilon_V : \overline{\mathcal{M}}_{g,n,s}^V(X,A) \xrightarrow{\text{st} \times \text{ev} \times \text{h}} \overline{\mathcal{M}}_{g,n+\ell(s)} \times X^n \times \mathcal{H}_{X,A,s}^V$$

obtained from (1.9) and (7.3). Furthermore, $\mathcal{M}_{g,n,s}^V(X,A)$ is oriented and the complement of $\mathcal{M}_{g,n,s}^V(X,A)$ in the irreducible part of $\overline{\mathcal{M}}_{g,n,s}^V(X,A)$ has codimension at least two.

Proof. To define the orientation, note that at each $f \in \mathcal{M}_s^V(X)$, the tangent space is the kernel of the linearized operator D_f . For generic (J, ν) the cokernel vanishes, so the tangent space is identified with the formal vector space $\ker D_f - \text{coker } D_f$. This is oriented by the mod 2 spectral flow of a path in the space of Fredholm operators that connects D_f to any operator that commutes with J , where the kernel and the cokernel are complex vector spaces

and hence are canonically oriented. This orients $\mathcal{M}_s^V(X)$, and the orientation extends to the compactification provided the frontier strata have codimension at least 2.

We established compactness above and the dimension statements are verified in the next two lemmas. \square

For the dimension counts we return to the notation of (7.1). The strata of the space of V -stable maps are labeled by curves $C = \bigcup B_k$ with a marked layer structure but no specified complex structure; we denote these strata by $\overline{\mathcal{M}}_{g,n,s}^V(C)$. For each layer B_k , $k \geq 1$, let

$$(7.5) \quad \overline{\mathcal{M}}_{B_k, n_k, s_k \cup t_k^+, t_k^-}^V(X, A_k)$$

denote the set of V -stable maps $(f, [\xi])$ where $f \in \mathcal{M}_{B_k, n_k + \ell(s_k \cup t_k^+ \cup t_k^-)}(V, A_k)$ is a map from B_k to V for which moreover the operator (7.1) has a nontrivial kernel on each irreducible component of B_k .

In this context consider the Hilbert bundle over the universal moduli space

$$L_k^M(V) \rightarrow \overline{\mathcal{UM}}_{B_k, n_k + \ell(s_k \cup t_k^+ \cup t_k^-)}(V, A_k)$$

whose fiber at $f : B_k \rightarrow V$ is the space $L_{k,\delta}^m(f^*N_V)$ of (7.1). It is straightforward to adapt the proof of Lemma 6.4 to show that D_k^N defines a section of

$$(7.6) \quad \begin{array}{c} \text{Fred} \left(L_{s,\delta}^m(V), L_{s,\delta+1}^{m-1}(V) \right) \\ \downarrow \\ \overline{\mathcal{UM}}_{B_k, n_k + \ell(s_k \cup t_k^+ \cup t_k^-)}(V, A_k) \end{array}$$

which is transverse to the subspaces Fred_r of operators with kernel of dimension $r \geq 1$.

LEMMA 7.5. *The irreducible part of the space (7.5) is an orbifold of “correct” dimension, which is at most*

$$(7.7) \quad d_k = 2 \left[-K_X[A_k] + \frac{1}{4}(\dim X - 6)\chi(B_k) + n_k + \ell(t_k^-) + \ell(s_k) + \ell(t_k^+) + \deg t_k^- - \deg t_k^+ - \deg s_k \right] - 2.$$

Proof. Consider an irreducible V -stable map $(f, [\xi])$ in the space (7.5). By Theorem 1.8b, generically $\mathcal{M}_{B_k, n_k + \ell(s_k \cup t_k^+ \cup t_k^-)}(V, A_k)^*$ is an orbifold of dimension

$$(7.8) \quad d'_k = 2 \left[-K_V[A_k] + \frac{1}{4}(\dim V - 6)\chi(B_k) + n_k + \ell(t_k^+) + \ell(t_k^-) + \ell(s_k) \right].$$

Comparing (7.7), (7.8) and the displayed formula following (7.1) we see that $d_k - d'_k = 2(\iota - 1)$ where $2\iota = \text{index } D_k^N$. The lemma follows immediately if $\iota > 0$.

When $\iota \leq 0$ we can use the element ξ_k and the transversality of D_k^N in (7.6) to conclude that the set of $f \in \mathcal{M}_{B_k}(V, A_k)^*$ with $\dim \ker D_k^N = 2r$ form a suborbifold of codimension $2r(-\iota + r)$. The lemma then follows because $r(r - \iota) \geq 1 - \iota$ for all $r \geq 1$. (Notice that this argument requires that ξ_k be nontrivial on every component of B_k as in Definition 7.2). \square

LEMMA 7.6. *Each irreducible stratum $\overline{\mathcal{M}}_{g,n,s}^V(C)$ is an orbifold whose dimension is $2(r + \sum \ell_k)$ less than that in (4.2), where r is the total number of nontrivial layers and ℓ_k is the number of double points of C in layer $k \geq 0$.*

Proof. The stratum $\overline{\mathcal{M}}_{g,n,s}^V(C)$ is the product of the spaces (7.5), one for each layer, constrained by the matching conditions $f(x) = f(y)$ at each of the $\ell(t)$ double points where B_k meets the other layers. A standard transversality argument [RT1] shows that the irreducible part of this space is an orbifold of the expected dimension. Thus

$$\dim \overline{\mathcal{M}}_{g,n,s}^V(C) \leq \sum_{k=0}^r d_k - \ell(t) \dim V.$$

Now substitute in (7.7) for d_k , $k \geq 1$ and sum, noting that (i) formula (7.7) is additive in A_k and s_k , (ii) each double point x between different layers contributes its multiplicity t_x to both t^+ and t^- so that $\sum \deg t_k^+ = \sum \deg t_k^-$, and (iii) the Euler characteristics add according to the formula

$$\sum \chi(B_k) = \chi(C) + 2\ell(t)$$

where $\ell(t) = \sum \ell(t_k^+) = \sum \ell(t_k^-)$. The result follows. \square

Remark 7.7. There is a different but equivalent viewpoint on what a V -stable map is. In Proposition 6.6 and Lemma 7.3 we inductively produced a limit map f , layers B_k , and limiting renormalized maps $g_k : B_k \rightarrow \mathbb{P}_V$ for $k \geq 1$. On each B_k we wrote g_k as $\exp_f(\xi_k)$ using the exponential map from $V_0 \subset P_V$; the information $\{\{\xi_k\}\}$ then defined a V -stable map as in Definition 7.2. Alternatively, we could have recorded the g_k themselves modulo the \mathbb{C}^* action on \mathbb{P}_V . From that perspective the limiting V -stable map is an equivalence class of continuous maps $f \cup g_1 \cup \cdots \cup g_r$ from $C = \bigcup_{k=0}^r B_k$ into the singular space

$$X \cup_{V=V_\infty} \mathbb{P}_V \cup_{V_0=V_\infty} \cdots \cup_{V_0=V_\infty} \mathbb{P}_V$$

with f mapping B_0 to X and each g_k , $k \geq 1$, mapping B_k to the k^{th} copy of \mathbb{P}_V with all maps V -regular along each intermediate copy of V and with two such maps $f \cup g_1 \cup \cdots \cup g_r$ equivalent if they lie in the same orbit of

the $(\mathbb{C}^*)^r = \mathbb{C}^* \times \cdots \times \mathbb{C}^*$ action. The renormalization procedure of Lemma 7.3 gives a uniform bound on the number of copies of \mathbb{P}_V for each homology class A .

The correspondence between these two viewpoints is clear. We found that the analytic technicalities were easiest using the description of Definition 7.2, but in general it is useful to keep both descriptions in mind.

8. Relative invariants

The relative Gromov-Witten invariant is the homology class obtained by pushing forward the compactified space of V -regular maps by the map (7.4):

$$\varepsilon_V : \overline{\mathcal{M}}_{g,n,s}^V(X, A) \longrightarrow \overline{\mathcal{M}}_{g,n+\ell(s)} \times X^n \times \mathcal{H}_X^V.$$

In this section we show that this yields a well-defined homology class. We then recast the relative invariants as Laurent series and explain their geometric interpretation.

THEOREM 8.1. *Assume X and V are semipositive or that the moduli space $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$ is generically irreducible. Then for generic $(J, \nu) \in \mathcal{J}^V$, the image of $\overline{\mathcal{M}}_{g,n,s}^V(X, A)$ under ε_V defines an element*

$$(8.1) \quad \text{GW}_{X,A,g,n,s}^V \in H_*(\overline{\mathcal{M}}_{g,n+\ell(s)} \times X^n \times \mathcal{H}_X^V; \mathbb{Q})$$

of dimension

$$(8.2) \quad -2K_X[A] + (\dim X - 6)(1 - g) + 2(n + \ell(s) - \deg s).$$

This homology class is independent of the generic $(J, \nu) \in \mathcal{J}^V$.

Proof. The spaces \mathcal{H}_X^V and $\overline{\mathcal{M}}_{g,n}$ are orbifolds, and therefore so is $Y = \overline{\mathcal{M}}_{g,n} \times X^n \times \mathcal{H}_X^V$. Fix a generic (J, ν) and consider the image of the smooth map $\varepsilon_V : \mathcal{M}_{g,n,s}^V \rightarrow Y$. Its frontier

$\text{Fr}(\varepsilon_V) = \{y \in Y \mid y = \lim \varepsilon_V(f_k) \text{ and } \{f_k\} \text{ has no convergent subsequence}\}$
is exactly the image

$$\left[\varepsilon_V \left(\overline{\mathcal{M}}_{g,n,s}^V(X, A) \setminus \mathcal{M}_{g,n,s}^V(X, A) \right) \right].$$

Then Theorem 7.4 (applied to the reduced moduli space when X, V are semipositive) implies that the frontier $\text{Fr}(\varepsilon_V)$ lies in a set of dimension two less than the dimension (8.2) of the image. Proposition 4.2 of [KM] then implies that the image $\varepsilon_V \left(\overline{\mathcal{M}}_{g,n,s}^V \right)$ carries a rational homology class of that dimension.

The last statement of the theorem follows by a cobordism argument. By Theorem A.2 of the appendix \mathcal{J}^V is path-connected; so any two generic pairs

(J_1, ν_1) and (J_2, ν_2) can be joined by a path $\gamma(t)$. The Sard-Smale theorem, applied to the space of such paths, shows that over the generic such $\gamma(t)$, the universal moduli space \mathcal{UM}^V over γ is an orbifold. Again the frontier of the image

$$\varepsilon_V \left(\pi^{-1}(\gamma) \right)$$

lies in a set of dimension two less than the dimension of this image. Proposition 4.4 of [KM] then implies that the homology classes (8.1) defined by (J_1, ν_1) and (J_2, ν_2) are the same. \square

Definition 8.2. Let (X, ω) be a closed symplectic manifold with a codimension two symplectic submanifold V . For each g, n , the *relative GW invariant* of (X, V, ω) is the homology class (8.1).

It is again convenient to assemble these invariants into a Laurent series. For that, we simply repeat the discussion leading from (2.1) to (2.3). Thus the full relative GW invariant is the map

$$(8.3) \quad \text{GW}_X^V : H^*(\overline{\mathcal{M}}) \otimes \mathbb{T}^*(X) \longrightarrow H_*(\mathcal{H}_X^V; \mathbb{Q}[\lambda])$$

where $\mathbb{T}^*(X)$ is the total tensor algebra $\mathbb{T}(H^*(X))$ on the rational cohomology.

As in (5.10), \mathcal{H}_X^V is a union of components labeled by A and s , so that

$$H_* \left(\mathcal{H}_X^V; \mathbb{Q}[\lambda] \right) = \bigoplus_{\substack{A, s \\ \deg s = A \cdot V}} H_* \left(\mathcal{H}_{X, A, s}^V \right) \otimes \mathbb{Q}[\lambda].$$

Thus there is an expansion

$$(8.4) \quad \text{GW}_X^V = \sum_{g, n} \frac{1}{n!} \sum_{\substack{A, s \\ \deg s = A \cdot V}} \frac{1}{\ell(s)!} \text{GW}_{X, A, g, n, s}^V t_A \lambda^{2g-2}$$

where the coefficients on the right lie in $H_* \left(\mathcal{H}_{X, A, s}^V \right)$.

Formula (2.7) extends this to a relative Gromov-Taubes invariant

$$(8.5) \quad \text{GT}_X^V = \exp(\text{GW}_X^V) : H^*(\widetilde{\mathcal{M}}) \otimes \mathbb{T}^*(X) \longrightarrow H_*(\mathcal{H}_X^V; \mathbb{Q}[\lambda]).$$

It is clear that these invariants are natural: if ϕ is a diffeomorphism of X then $V' = \phi^{-1}(V)$ is a symplectic submanifold of $(X', \phi^*\omega)$ and

$$\text{GW}_{X'}^{V'} = \text{GW}_X^V.$$

It is also clear that these invariants extend the GW invariants of Section 2. In fact, the entire construction carries through when V is the empty set. In that case \mathcal{H}_X^V is just $H_2(X)$ and the relative invariant takes values in $NH_2(X)$. The relative and absolute invariants are then equal:

$$\text{GW}_X^\emptyset = \text{GW}_X.$$

More importantly, the relative GW invariants are unchanged under symplectic isotopies, i.e. they are constant as we move along 1-parameter families (X, V_t, ω_t) consisting of a symplectic form ω_t on X and a codimension-two submanifold V_t which is symplectic for ω_t . More generally, we say (X, V, ω) is *deformation equivalent* to (X', V', ω') if there is a diffeomorphism $\phi: X' \rightarrow X$ such that $(X', \phi^{-1}(V), \phi^*\omega)$ is isotopic to (X', V', ω')

PROPOSITION 8.3. *The relative invariant GW_X^V depends only on the symplectic deformation class of (X, V, ω) .*

Proof. By naturality we need only verify invariance under symplectic isotopies. But that follows by essentially the same cobordism argument used in the last part of the proof of Theorem 8.2 (cf. [RT2, Lemma 4.9] and [LT, Prop. 2.3]). \square

The geometric meaning of the relative invariant is obtained by evaluating the homology classes (8.1) on dual cohomology classes, thereby re-expressing the invariant as a collection of numbers. To do that, choose $\kappa \in H^*(\overline{\mathcal{M}}_{g,n})$, a vector $\alpha = (\alpha_1, \dots, \alpha_n)$ of classes in $H^*(X)$, and a $\gamma \in H^*(\mathcal{H}_X^V)$, such that

$$(8.6) \quad \deg \kappa + \deg \alpha - 2\ell(\alpha) + \deg \gamma$$

is the dimension (8.2) of the homology class $\text{GW}_{X,A,g,n,s}^V$. Then the evaluation pairing gives numbers

$$(8.7) \quad \text{GW}_{X,A,g,n,s}^V(\kappa, \alpha, \gamma) = \left\langle \left[\text{GW}_{X,A,g,n,s}^V \right], \kappa \cup \alpha \cup \gamma \right\rangle.$$

It is these numbers that have a specific geometric interpretation.

PROPOSITION 8.4. *Fix generic geometric representatives $K \subset \overline{\mathcal{M}}_{g,n}$, $A_i \subset X$, and $\Gamma \subset \mathcal{H}_X^V$ of the Poincaré duals of κ , α and γ . Then the evaluation (8.7) counts the oriented number of genus g (J, ν) -holomorphic maps $f: C \rightarrow X$ with $C \in K$, $h(f) \in \Gamma$, and $f(x_i) \in A_i$ for each marked point x_i .*

Note that the condition $h(f) \in \Gamma$, defined by the map (5.9), constrains both the homology class A of the map and the boundary values of the curve. In the special case when there are no rim tori, these homology and the boundary value constraints can be fully separated as in the beginning of Section 9.

Also note that the invariant counts maps from a domain with $n + \ell(s)$ marked points, the last $\ell(s)$ of which are mapped into V . Two such maps with their last $\ell(s)$ marked points renumbered are considered different. This might seem to introduce an unnecessary redundancy, but the marking on the last set of points is needed to prove that two curves whose intersection with V is the same can be ‘glued together’ (see [IP4, Th. 5.6]).

It is useful to have a more general version of Proposition 8.4. If we drop the dimension restriction (8.6) then the set of (J, ν) -holomorphic maps

$$(8.8) \quad \mathcal{M}_{g,n,s}^V(X, A; \kappa, \alpha, \gamma)$$

satisfying the conditions of Proposition 8.4 will no longer be finite; its expected dimension is

$$(8.9) \quad \dim \mathcal{M}_{g,n,s}^V(X, A) - \deg \kappa - \deg \alpha + 2\ell(\alpha) - \deg \gamma$$

with the first term given by (8.2). Of course (8.8) depends on (J, ν) ; it is a fiber of a map

$$\mathcal{UM}_{g,n,s}^V(X, A; \kappa, \alpha, \gamma) \rightarrow \mathcal{J}^V$$

from the subset of the universal space $\mathcal{UM}_{g,n,s}^V(X, A)$ consisting of those maps satisfying the conditions of Proposition 8.4. Proposition 8.4 is then the 0-dimensional case of the following fact.

LEMMA 8.5. *For generic representatives K, A_i and Γ ,*

$$\mathcal{UM}_{g,n,s}^V(X, A; \kappa, \alpha, \gamma)^*$$

is an orbifold. Hence for generic (J, ν, K, A_i, Γ) , the irreducible part of $\mathcal{M}_{g,n,s}^V(X, A; \kappa, \alpha, \gamma)$ is an orbifold of dimension as in (8.9).

Proof. Consider the evaluation map $\text{st} \times \text{ev} : \mathcal{UM}_{g,n,s}^V(X, A) \rightarrow \mathcal{M}_{g,n} \times X^n \times V^{\ell(s)}$ by

$$\begin{aligned} & (f, j, x_1, \dots, x_n, (p_1, s_1) \dots, (p_\ell, s_\ell)) \\ & \mapsto ((j, x_1, \dots, x_n), f(x_1), \dots, f(x_n), f(p_1), \dots, f(p_\ell)). \end{aligned}$$

We can achieve transversality by moving the geometric representatives K, A_i, C_l of α, γ , exactly as in Section 4 of [RT1], only now keeping the representatives Γ in V . The dimension formula (8.9) follows because κ cuts down $\deg \kappa$ dimensions, and the α and γ constraints cut down by $\deg(\alpha)$ and $\deg(\gamma)$ dimensions, respectively. \square

Remark 8.6. So far, the relative invariant is defined by cutting down with geometric representatives of classes in $H_*(X)$. For the application to the gluing theorem in [IP4] it is useful to have a version of these invariants that allows constraints in $H_*(\hat{X}, S)$ (which are Poincaré dual to classes in $H^*(X \setminus V)$).

Let \hat{X} be the manifold (5.5) obtained from $X \setminus V$ by attaching as boundary a copy of the unit circle bundle $\pi : S \rightarrow V$ of the normal bundle of V in X . Suppose that Z is a symplectic sum obtained by gluing \hat{X} to a similar manifold \hat{Y} along S . We can then consider stable maps in Z constrained by classes B in $H_k(Z)$, i.e. the set of stable maps f with the image $f(x)$ of a marked point lying on a geometric representative ϕ of B . If we restrict ourselves to the \hat{X}

side, ϕ defines a class $[\phi] \in H_*(\hat{X}, S)$ and a subspace $\overline{\mathcal{M}}^V(\phi)$ of the space of V -stable maps.

We can repeat the above arguments to obtain GW invariants constrained by ϕ . However, in general $\overline{\mathcal{M}}^V(\phi)$ is an orbifold with boundary and consequently the invariant depends on the choice of the representative ϕ . In [IP4, §13] we will show that this dependence is rather mild and in practice can be handled as follows.

In the exact sequence

$$H_k(X \setminus V) = H_k(\hat{X}) \longrightarrow H_k(\hat{X}, S) \xrightarrow{\partial} H_{k-1}(S) \xrightarrow{\iota} H_{k-1}(\hat{X})$$

choose a splitting at ∂ over $K = \ker \iota$ and choose a geometric representative ϕ_B of each B in the image of this splitting. For each class in $H_k(\hat{X}, S)$ we can find a geometric representative ϕ with the same boundary as one of the chosen ϕ_B . Then $a = [\phi \# (-\phi_B)]$ lies in $H_k(X \setminus V)$ and

$$(8.10) \quad GW(\phi) = GW(\phi_B) + GW(a).$$

This defines GW on a complete set of representatives of $H_k(X; V)$ which combines with Poincaré duality to give a map

$$(8.11) \quad \text{GW} : H^*(\overline{\mathcal{M}}) \times \mathbb{T}^*(X \setminus V) \longrightarrow H_*(\mathcal{H}_X^V).$$

This map is neither canonical nor linear in the constraints, but is additive as in (8.10) and is determined by the invariants (8.3) and those for the chosen constraints ϕ_B . It provides a set of constraints that can be used in the gluing theorem to constrain by any class in $H^*(Z)$.

9. Examples

The relative GW invariants are designed to be used in ‘cutting and pasting’ arguments of symplectic topology. However, in several interesting cases they are identical to invariants from enumerative algebraic geometry. Here we give three such examples; in each case X, V are semipositive. Actual computations of the relative invariants in these cases are done in [IP4].

As noted in Remark 5.3, the description of the relative GW invariants is simplified considerably when there are no rim tori. This occurs whenever $H_1(V) = 0$ and more generally when every rim torus represents zero in $H_2(X \setminus V)$. In these cases there is no covering (5.8), and \mathcal{H}_X^V is the subset of $H_2(X) \times \mathcal{S}V$ consisting of pairs (A, s) with $\deg s = A \cdot V$. The homology of \mathcal{H}_X^V is the corresponding subalgebra of $NH_2(X) \otimes \mathbb{C}\mathbb{T}_*(V)$ where $\mathbb{C}\mathbb{T}_*(V)$ is the ‘contact tensor algebra’ of V :

$$\mathbb{C}\mathbb{T}_*(V) = \mathbb{T}(\mathbb{N} \times H_*(V)).$$

The relative invariants are then maps

$$H^*(\overline{\mathcal{M}}) \otimes \mathbb{T}^*(X) \rightarrow \mathbb{C}\mathbb{T}_*(V) \otimes NH_2(X)[\lambda]$$

and have Laurent expansions like (8.4) with coefficients in $\mathbb{C}\mathbb{T}_*(V)$, and $NH_2(X)$ is the Novikov ring (cf. Section 2). Fix dual bases γ_i of $H_*(V; \mathbb{Q})$ and γ^i of $H^*(V; \mathbb{Q})$. Then bases of the contact algebra and its dual are given by elements of the form

$$(9.1) \quad C_{s,\gamma} = C_{s_1,\gamma_1} \otimes \dots \otimes C_{s_\ell,\gamma_\ell} \quad \text{and} \quad C_{s,\gamma}^* = C_{s_\ell,\gamma_\ell} \otimes \dots \otimes C_{s_1,\gamma_1}$$

respectively, where $s_i \geq 1$ are integers. With κ and α as above, we can expand

$$(9.2) \quad \text{GW}^V(\kappa, \alpha) = \sum_{A,s,\gamma} \frac{1}{\ell(s)!} \text{GW}_{X,A,g}^V(\kappa, \alpha; C_{s,\gamma}^*) C_{s,\gamma} t_A \lambda^{2g-2}$$

where the coefficients count the oriented number of genus g (J, ν)-holomorphic, V -stable maps $f : C \rightarrow X$ with $C \in K$, $f(x_i) \in A_i$; these have a contact of order s_j with V along fixed representatives Γ_j of the Poincaré duals of the γ^j , where K and A_i are Poincaré duals of the κ and α_i .

With this background, we will describe two simple examples of relative invariants and their classical algebraic-geometry counterparts. While these are amongst the very simplest examples of relative invariants, each has a long history and has proved to be frustratingly difficult to compute by algebraic-geometric methods. However, in both examples recent progress has been made on calculating the invariants by using, in part, symplectic cut-and-paste arguments.

Example 9.1. The Hurwitz numbers are examples of GW invariants of \mathbb{P}^1 relative to several points in \mathbb{P}^1 .

The classical Hurwitz number $N_{g,d}$ counts the number of nonsingular, genus g curves expressible as d -sheeted covers of \mathbb{P}^1 with a *fixed* branch divisor in general position. They were first computed in [Hu] by combinatorial techniques. More generally, if α is an (unordered) partition of d then the Hurwitz number $N_{g,d}(\alpha)$ counts the number of smooth degree d maps from a genus g Riemann surface to \mathbb{P}^1 with the ramification above a fixed point p_0 as specified by the partition α , and simple branching at exactly $r(g, \alpha) = d + \ell(\alpha) + 2g - 2$ other fixed points in general position.

On the other hand, for any distinct fixed points p_0, \dots, p_r in $X = \mathbb{P}^1$, the set $V = \{p_0, \dots, p_r\}$ is a symplectic submanifold of X with no rim tori. The homology class A of the map is given by the degree d ; we take $\kappa = 1$ because we are imposing no constraints on the complex structure of the curves. Thus the relative GW invariant $\text{GW}_g^V(\mathbb{P}^1, d)$ has the form (9.2) with values in $\mathcal{S}V$. But $\mathcal{S}\{p_0\}$ is the disjoint union of copies of p_0 , one copy for each vector s with $\deg s = p_0 \cdot A = d$. Furthermore, the relative invariant is unchanged

when the marked points that are mapped into p_0 are permuted. We can then associate the generator of $H_0(V_s)$ coming from the point p_0 with the monomial $z^s = z_{s_1} \cdots z_{s_\ell}$. Thus we identify

$$H_*(\mathcal{S}\{p_0\}) = \bigoplus_s \mathbb{Z} = \mathbb{Z}[z_1, z_2, \dots]$$

where the last term is the polynomial ring on variables z_1, z_2, \dots . Then the Hurwitz number

$$N_{g,d}(\alpha) = \text{GW}_{\mathbb{P}^1, d, g}^V(z^s; b^r)$$

where $r = d + 2g - 2 + \ell(\alpha)$ and s is one – any one – of the ordered partitions obtained by ordering α . Geometrically, the variable z_i models a contact of order i at p_0 , r is the number of leftover simple branch points, and b^r denotes the condition that r simple branch points are mapped to r distinct fixed points $\{p_1, \dots, p_r\}$.

With this notation the Laurent series (8.4) of the relative invariant is:

$$\begin{aligned} (9.3) \quad \text{GW}_{\mathbb{P}^1}^p &= \sum_s \frac{1}{\ell(s)!} \text{GW}_{\mathbb{P}^1, d, g}^V(z^s; b^r) \zeta^s t^d \frac{u^r}{r!} \lambda^{2g-2} \\ &= \sum_\alpha N_{d,g}(\alpha) \zeta^\alpha t^d \frac{u^r}{r!} \lambda^{2g-2} \end{aligned}$$

where the monomial ζ^α is dual to z^α . (The $\ell!$ appears because our relative invariant orders the points in the inverse image of p , while the Hurwitz numbers do not.) This is a standard generating function for the Hurwitz numbers.

Example 9.2. The GT invariant of \mathbb{P}^2 relative to a line L is the collection of enumerative invariants introduced by Caporaso and Harris in [CH].

In [CH], Caporaso and Harris establish a recursion formula for the number of nodal curves in \mathbb{P}^2 . They separate the set of nodal curves into classes according to how the curves intersect a fixed line L . Specifically, for each pair of finite sequences $\alpha = (\alpha_1, \alpha_2, \dots)$ and $\beta = (\beta_1, \beta_2, \dots)$ they consider the number $N_{d,\delta}(\alpha, \beta)$ of degree d curves with δ double points, having a contact with L of order k at α_k fixed points and at β_k unspecified points of L for each $k = 1, 2, \dots$, and passing through $r = 2d + g - 1 + \ell(\beta)$ fixed points off L . Note that δ is determined by the adjunction formula $2g = (d - 1)(d - 2) - 2\delta$.

From our viewpoint $V = L$ is a symplectic submanifold of $X = \mathbb{P}^2$ with no rim tori. As in Example 9.1 the homology class A of the map is given by the degree d and we are imposing no constraints on the complex structure. This time $\mathcal{S}V$ is the disjoint union of products of copies of $V = \mathbb{P}^1$. Since V has only even-dimensional homology, the relative invariant is again unchanged under permutations of the marked points that are mapped into L . We can

then associate the generator of $H_0(V_s)$ with the monomial $y^s = y_{s_1} \cdots y_{s_\ell}$, the generator of $H_2(V_s)$ with $z^s = z_{s_1} \cdots z_{s_\ell}$. Thus

$$H_*(\mathcal{S}V) = \bigoplus_s (\mathbb{Z} \oplus \mathbb{Z}) = \mathbb{Z}[y_1, z_1, y_2, z_2, \dots].$$

Then

$$(9.4) \quad N_{d,\delta}(\alpha, \beta) = \text{GT}_{\mathbb{P}^2, dL, \chi, r}^L(p^r; y^\alpha, z^\beta) = \text{GT}_{\mathbb{P}^2, dL, \chi, r}^L(p^r; C_{(s, \gamma)})$$

where $\chi = 2 - 2g = -d(d-3) + 2\delta$ and where s is any one of the ordered sequences such that the basis element (9.1) satisfies

$$\alpha_k = \text{Card} \{i \mid (s_i, \gamma_i) = (k, [p])\},$$

and

$$\beta_k = \text{Card} \{i \mid (s_i, \gamma_i) = (k, [\mathbb{P}^1])\}.$$

After matching notation through (9.4), we see that the Caporaso-Harris recursion formula is a consequence of the gluing theorem for relative invariants; see [IP3].

Example 9.3. The GT invariant of the elliptic surfaces $E(n)$ relative to a fiber F .

This example appeared in [IP2]. Here $E(n) \rightarrow \mathbb{P}^1$ is the elliptic surface with a section of self-intersection $-n$, so that $E(0) = \mathbb{P}^1 \times T^2$, $E(1)$ is the rational elliptic surface, and $E(2) = K3$, each regarded as a symplectic manifold. We focus on counting the genus 1 (Euler characteristic 0) curves representing multiples of the fiber class. For generic (J, ν) that count is given by the numbers $\text{GT}_{E(n), mF, 0}$. As in [IP2] these agree with the Seiberg-Witten invariants and are determined by the generating function

$$(9.5) \quad \sum_m \text{GT}_{E(n), mF, 0} t_F^m = (1 - t_F)^{n-2}.$$

The geometric interpretation of this is given in [T], [IP2] and [IP3]. For generic J and $\nu = 0$ there are exactly $n - 2$ holomorphic fibers; these are of type $(0, -)$ for $n > 2$ and type $(0, +)$ for $n < 2$. The type determines the contribution to the GT invariant of the maps which multiply cover these fibers when we move from $(J, 0)$ to a generic (J, ν) . For type $(0, +)$ all covers contribute, giving the factor $(1 - t_F)^{-1}$, while curves of type $(0, -)$ contribute the opposite factor $(1 - t_F)$.

Now fix a generic fiber F_0 and restrict attention to F_0 -compatible $(J, 0)$. Then D^N is a complex operator by Lemma 3.3, so that F_0 is a holomorphic curve of type $(0, +)$ (cf. [T]). The relative GT invariant then, by the Definition 4.1, does not contain the contribution of F_0 and its multiple coverings. Thus

$$(9.6) \quad \sum_m \text{GT}_{E(n), mF, 0}^F t_F^m = (1 - t_F)^{n-1}.$$

In particular, the absolute and relative GT invariants are different. In this case there are no rim tori, F_0 -regular maps representing multiples of the fiber never intersect F_0 , and (9.6) agrees with the relative Seiberg-Witten invariants.

Similarly, the GT invariant of $E(0) = S^2 \times T^2$ relative to two copies of F is

$$\sum_m \text{GT}_{E(0), mF, 0}^{F, F} t_F^m = 1.$$

Appendix

The space \mathcal{J}^V of almost complex structures compatible with V was defined in Section 3. Here we show that \mathcal{J}^V is nonempty and path-connected. This fact was used in Section 7 to show that the relative GW invariants depend only on the symplectic structure.

An almost complex structure J on a symplectic manifold (X, ω) is *compatible* with ω if

$$(A.1) \quad g(X, Y) = \omega(X, JY)$$

defines a Riemannian metric; this implies that $g(JX, JY) = g(X, Y)$. Such a compatible J can always be constructed, as follows. After we fix a “background” metric g' , ω defines a skew-symmetric endomorphism A of T_p^*X at each point $p \in X$ by $\omega(X, Y) = g'(AX, Y)$. From linear algebra, any $A \in \text{GL}(n)$ can be *uniquely* expressed as $A = JS$ where J is orthogonal and S is positive definite and symmetric. Then $(-J^2)(J^t S J) = -J(JJ^t)S J = -AJ = A^t J = S J^t J = S$. Since $J^t S J$ is positive definite and symmetric, the uniqueness of the decomposition gives $J^2 = -\text{Id}$. Thus J is an almost-complex structure, and then $g(X, Y) = \omega(X, JY)$ is a J -compatible metric.

Given a symplectic submanifold $V \subset X$, let $N^g \subset TX$ and $N^\omega \subset TX$ denote the normal bundles to V defined by the metric g and the symplectic form ω respectively.

LEMMA A.1. *For compatible (ω, g, J) , V is J -invariant if and only if $N^\omega = N^g$.*

Proof. If $N^\omega = N^g$ then for any $X \in N^\omega$ and $v \in TV$, we have $g(X, Jv) = -\omega(X, v) = 0$, so that $Jv \in TV$; thus V is J -invariant. Conversely, if V is J -invariant, the equation $g(X, v) = \omega(X, Jv)$ implies that $N^\omega = N^g$. \square

THEOREM A.2. *The space \mathcal{J}^V of pairs (J, ν) satisfying (3.3) is nonempty and path-connected.*

Proof. Replacing ν by $t\nu$, $0 \leq t \leq 1$, gives a retraction of \mathcal{J}^V to the space \mathcal{J}_0^V of J satisfying (3.3a) and (3.3b). It therefore suffices to show that \mathcal{J}_0^V is nonempty and path-connected.

Given triples (ω, g_0, J_0) and (ω, g_1, J_1) in \mathcal{J}_0^V , we can apply the above construction to the path $h_t = (1-t)g_0 + tg_1$ to get a homotopy (ω, g_t, J_t) in which the decomposition $TV \oplus N^\omega$ is orthogonal under g_t and preserved by J_t . Thus each J_t satisfies (3.3a).

To finish the proof, we will modify the path (ω, g_t, J_t) to a path $(\omega, \tilde{g}_t, \tilde{J}_t)$ which also satisfies (3.3b). For notational simplicity we will omit the subscript t .

The Nijenhuis tensor, multiplied by $-J$, defines a linear map $L : N_V \rightarrow \text{Hom}(TV, N_V)$ by

$$\begin{aligned} L_\xi(v) &= (J[v, \xi] - [v, J\xi] - [Jv, \xi] - J[Jv, J\xi])^N \\ &= [(\nabla_\xi J)(v) + J(\nabla_{J\xi} J)(v) - (\nabla_v J)(\xi) - J(\nabla_{Jv} J)(\xi)]^N \end{aligned}$$

for $\xi \in N_V$. This is tensorial and J -anti-linear in v and ξ , and depends only on J (the second formula above holds for the Levi-Civita connection of *any* Riemannian metric). Extend L to a map $L : N_V \rightarrow \text{End}(TX)$ by setting $L_\xi(\eta) = 0$ for $\xi, \eta \in N_V$, and let $L^t : N_V \rightarrow \text{End}(TX)$ be its transpose. By extending L to a neighborhood of V , integrating for a short distance along the lines normal to V , and extending arbitrarily, we can find a $K \in \Gamma(\text{End}(TX))$ whose 1-jet along V satisfies

$$(A.2) \quad K|_V = 0 \quad \text{and} \quad \nabla_\xi K = -\frac{1}{2} (L_\xi + L_\xi^t) \quad \forall \xi \in N_V.$$

Then $KJ = -JK$ and K is self-adjoint with respect to g . Consequently, JK is self-adjoint, so that

$$g'(X, Y) = g(e^{KJ} X, Y)$$

defines a Riemannian metric, and it is straightforward to check that $J' := e^{JK} J = J e^{KJ}$ is orthogonal with respect to g' . With g' as background metric, the procedure described after (A.1) yields a compatible triple $(\omega, \tilde{g}, \tilde{J})$ where $A' = \tilde{J}\tilde{S}$ satisfies

$$g(JX, Y) = \omega(X, Y) = g'(A'X, Y) = g(e^{KJ} A'X, Y)$$

and therefore $A' = J'$. The uniqueness of the factorization $A' = \tilde{J}\tilde{S} = J' \cdot I$ then implies that $\tilde{J} = J' = J + K + \dots$ where the dots denote terms that vanish to second order along V . With that, we can evaluate

$$\tilde{L}_\xi(v) = [(\nabla_\xi \tilde{J})(v) + \tilde{J}(\nabla_{\tilde{J}\xi} \tilde{J})(v) - (\nabla_v \tilde{J})(\xi) - \tilde{J}(\nabla_{\tilde{J}v} \tilde{J})(\xi)]^N$$

along V . Using equation (A.2), and the facts that $L_\xi^t(v) = 0$ for $v \in TV$ and

$JL_{J\xi}(v) = L_\xi(v)$, we find that

$$\tilde{L}_\xi(v) = L_\xi(v) - \frac{1}{2}L_\xi(v) - \frac{1}{2}JL_{J\xi}(v) = 0.$$

Therefore $(\omega, \tilde{g}, \tilde{J})$ is a compatible triple satisfying (3.3a,b).

Applying this procedure to the path (ω, g_t, J_t) does not change g_t or J_t at $t = 0, 1$ (where $L_\xi(v)$ already vanishes) and hence gives the desired path $(\omega, \tilde{g}_t, \tilde{J}_t)$. \square

UNIVERSITY OF WISCONSIN-MADISON, MADISON, WI
E-mail address: ionel@math.wisc.edu

MICHIGAN STATE UNIVERSITY, EAST LANSING, MI
E-mail address: parker@math.msu.edu

REFERENCES

- [A] N. ARONSZAJN, A unique continuation theorem for solutions of elliptic partial differential equations or inequalities of second order, *J. Math Pures Appl.* **36** (1957), 235–249.
- [CH] L. CAPORASO and J. HARRIS, Counting plane curves of any genus, *Invent. Math.* **131** (1998), 345–392.
- [E] Y. ELIASHBERG, Invariants in contact topology, *Proc. Internat. Congress of Math.* **II** (Berlin, 1998), *Doc. Math.* (1998) *Extra Vol. II* (1998), 327–338.
- [FO] K. FUKAYA and K. ONO, Arnold conjecture and Gromov-Witten invariant, *Topology* **38** (1999), 933–1048.
- [G] M. GROMOV, Pseudo holomorphic curves in symplectic manifolds, *Invent. Math.* **82** (1985), 307–347.
- [GH] P. GRIFFITH and J. HARRIS, *Principles of Algebraic Geometry*, J. Wiley & Sons Inc., New York, 1978.
- [H] C. HUMMEL, *Gromov’s Compactness Theorem for Pseudo-Holomorphic Curves*, *Progr. in Math.* **151**, Birkhäuser Verlag, Basel, 1997.
- [Hu] A. HURWITZ, Über die Anzahl der Riemann’schen Flächen mit gegebenen Verzweigungspunkten, *Math. Ann.* **55** (1902), 53–66.
- [IP1] E. IONEL and T. PARKER, The Gromov invariants of Ruan-Tian and Taubes, *Math. Res. Lett.* **4** (1997), 521–532.
- [IP2] ———, Gromov invariants and symplectic maps, *Math. Ann.* **314** (1999), 127–158.
- [IP3] ———, Gromov-Witten invariants of symplectic sums, *Math. Res. Lett.* **5** (1998), 563i–576.
- [IP4] ———, The symplectic sum formula for Gromov-Witten invariants, *Ann. of Math.*, to appear; preprint, math.SG/0010217.
- [IP5] ———, The virtual fundamental class of the moduli space of pseudo-holomorphic maps, in preparation.
- [IS] S. IVASHKOVICH and V. SHEVCHISHIN, Complex curves in almost-complex manifolds and meromorphic hulls, preprint, CV/9912046.
- [K] U. KOSCHORKE, *Infinite Dimensional K-Theory and Characteristic Classes of Fredholm Bundle Maps*, in *Global Analysis, Proc. Symposia Pure Math.* **XV**, A. M. S., Providence, R.I., 1970, 95–133.
- [KM] P. KRONHEIMER and T. MROWKA, Embedded surfaces and the structure of Donaldson’s polynomial invariants, *J. Differential Geom.* **41** (1995), 573–734.
- [LR] A.-M. LI and Y. RUAN, Symplectic surgery and Gromov-Witten invariants of Calabi-Yau 3-folds I, *Invent. Math.* **145** (2001), 151–218; alg-geom/9803036.

- [LT] J. LI and G. TIAN, Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds, in *Topics in Symplectic 4-Manifolds* (Irvine, CA, 1996), 47–83, *First Internat. Press Lect. Ser. I*, International Press, Cambridge, MA, 1998.
- [L] R. LOCKHART, Fredholm, Hodge and Liouville theorems on noncompact manifolds, *Trans. A. M. S.* **301** (1987), 1–35.
- [Lo] E. LOOIJENGA, Smooth Deligne-Mumford compactifications by means of Prym level structures, *J. Algebraic Geom.* **3** (1994), 283–293.
- [M] D. MCDUFF, The local behaviour of holomorphic curves in almost complex 4-manifolds, *J. Differential Geom.* **34** (1991), 143–164.
- [MS] D. MCDUFF and D. SALAMON, *J-holomorphic Curves and Quantum Cohomology*, A.M.S., Providence, RI, *University Lecture Series* **6**, 1994.
- [P] T. PARKER, Compactified moduli spaces of pseudo-holomorphic curves, in *Mirror Symmetry, III* (Montreal, PQ, 1995), 77–113, *AMS/IP Stud. Adv. Math.* **10**, A. M. S., Providence, RI, 1999.
- [PW] T. PARKER and J. WOLFSON, Pseudo-holomorphic maps and bubble trees, *J. Geom. Anal.* **3** (1993) 63–98.
- [R] Y. RUAN, Virtual neighborhoods and pseudo-holomorphic curves, *Proc. of 6th Gokova Geometry-Topology Conference, Turkish J. Math.* **23** (1999), 161–231.
- [RT1] Y. RUAN and G. TIAN, A mathematical theory of quantum cohomology, *J. Differential Geom.* **42** (1995), 259–367.
- [RT2] ———, Higher genus symplectic invariants and sigma models coupled with gravity, *Invent. Math.* **130** (1997), 455–516.
- [S] B. SIEBERT, Symplectic Gromov-Witten invariants, in *New Trends in Algebraic Geometry* (Warwick, 1996), 375–424, *London Math. Soc. Lecture Note Ser.* **264**, Cambridge Univ. Press, Cambridge, 1999.
- [SY] Y.-T. SIU and S.-T. YAU, Compact Kähler manifolds of positive bisectional curvature, *Invent. Math.* **59** (1980), 189–204.
- [T] C. H. TAUBES, Counting pseudo-holomorphic curves in dimension 4, *J. Differential Geom.* **44** (1996), 818–893.

(Received July 25, 1999)