

THE GOPAKUMAR-VAFA FORMULA FOR SYMPLECTIC MANIFOLDS

ELENY-NICOLETA IONEL AND THOMAS H. PARKER

ABSTRACT. The Gopakumar-Vafa conjecture predicts that the Gromov-Witten invariants of a Calabi-Yau 3-fold can be canonically expressed in terms of integer invariants called BPS numbers. Using the methods of symplectic Gromov-Witten theory, we prove that the Gopakumar-Vafa conjecture holds for any symplectic Calabi-Yau 6-manifold, and hence for Calabi-Yau 3-folds. The results extend to all symplectic 6-manifolds and to the genus zero GW invariants of semipositive manifolds.

The Gopakumar-Vafa conjecture [GV] predicts that the Gromov-Witten invariants $GW_{A,g}$ of a Calabi-Yau 3-fold can be expressed in terms of some other invariants $n_{A,h}$, called BPS numbers, by a transform between their generating functions:

$$\sum_{\substack{A \neq 0 \\ g}} GW_{A,g} t^{2g-2} q^A = \sum_{\substack{A \neq 0 \\ h}} n_{A,h} \sum_{k=1}^{\infty} \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{2h-2} q^{kA}. \quad (0.1)$$

The content of the conjecture is that, while the $GW_{A,g}$ are rational numbers, the BPS numbers $n_{A,h}$ are *integers*. (Gopakumar and Vafa also conjectured that for each $A \in H_2(X, \mathbb{Z})$, the coefficients of (0.1) satisfy $n_{A,h} = 0$ for large h ; we do not address this finiteness statement here.) It is natural to enlarge the context by regarding this as a conjecture about the Gromov-Witten invariants of any closed symplectic 6-manifold X that satisfies the topological Calabi-Yau condition $c_1(X) = 0$.

Formula (0.1) can be viewed as a statement about the structure of the space of solutions to the J -holomorphic map equation. For a generic almost complex structure J , each J -holomorphic map is the composition $f = \varphi \circ \rho$ of a multiple-cover ρ and an embedding φ . The embeddings are well-behaved: they have no nontrivial automorphisms, and the moduli space of J -holomorphic embeddings is a manifold. But multiply-covered maps cause severe analytical problems with transversality. In the symplectic construction of the GW invariants, these problems are avoided by lifting to a cover of the moduli space and turning on a lift-dependent perturbation ν of the equation; this destroys the multiple-cover structure and only shows that the numbers $GW_{A,g}$ are rational. But it also suggests an interpretation of the GV formula: the righthand side of (0.1) might be a sum over embeddings, with the sum over k counting the contributions of the multiple covers of each embedding.

This viewpoint is very similar to C. Taubes' work [T] relating Gromov invariants to the Seiberg-Witten invariants of 4-manifolds, and our approach has been fundamentally influenced by Taubes. It is also similar to the 4-dimensional situation described by Lee and Parker in [LP1] and [LP2]. In both cases, the set of J -holomorphic embeddings in each homology class is discrete and compact for generic J — a simplifying circumstance that does not appear to be true in the context of formula (0.1). Rather, for generic J and with a fixed bound E on area and genus, the moduli space $\mathcal{M}_{emb}(X)$ of embeddings is a countable set, possibly with accumulation points. With this picture in mind, our proof is based on three main ideas.

The first is the observation that, again for fixed J and E , the full moduli space $\overline{\mathcal{M}}(X)$ can be decomposed (in many ways) into finitely many “clusters” \mathcal{O}_j . Each cluster consists of all of the J -holomorphic maps, including multiple covers, whose image lies in the ε -tubular neighborhood of some smooth, embedded J -holomorphic “core curve” $C \subset X$. A cluster is an open and closed subset of the moduli space; it may have complicated internal structure, but there is a well-defined total contribution $GW(\mathcal{O})$ of all the maps in the cluster to the series (0.1). These contributions $GW(\mathcal{O})$ are local,

The research of EI was supported in part by NSF grant DMS-0905738 and that of TP by NSF grant DMS-1011793.

depending only on ε and J in the neighborhood, and it suffices to show that the GV conjecture holds for the contribution of each cluster.

The second observation is that there exist certain standard “elementary clusters” whose local invariants are explicitly computable. Results of Junho Lee [L] show that, for each embedded genus g curve C , there exists an almost complex structure J in an ε neighborhood U of C in X that makes C “super-rigid”, meaning that all J -holomorphic maps into U are in fact maps into C . For $g = 0$, one can take J to be the standard structure of the bundle $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$, but for higher genus J is a non-integrable almost complex structure. In Section 3 we compute the GW series $GW(\mathcal{O})$ of elementary clusters based on a calculation of Bryan and Pandharipande [BP2]. The resulting formula shows that the local version of the GV conjecture holds for elementary clusters.

The proof is completed by an isotopy argument in the spirit of Taubes’ work, and extending arguments in [IP1]. For a fixed cluster, we deform J in a neighborhood of the core curve to make it the J of an elementary cluster. During the isotopy, the cluster series $GW(\mathcal{O}_t)$ can change according to several types of wall-crossing formulas. For a generic isotopy, the core curve could disappear in a “creation-annihilation” singularity. To avoid this, we use a generic isotopy in which the restriction of J to the core curve is fixed; singularities then occur only when two core curves pass through one another momentarily. In Sections 5 and 6, we use Kuranishi models to show that the cluster series is invariant modulo contributions of finitely many clusters whose core curves have higher degree or genus. The Gopakumar-Vafa conjecture follows by induction.

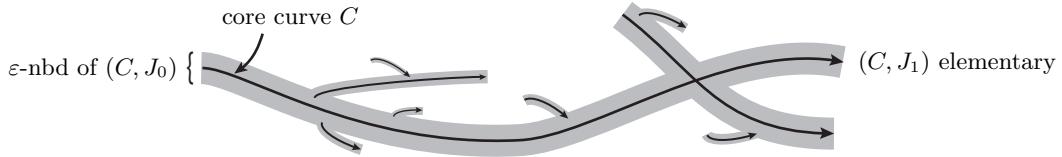


FIGURE 1. As J_t changes, embedded curves of higher genus or degree can emerge from, or sink into, an ε -tubular neighborhood of C , and the core curve C can pass through another embedded curve with the same genus and degree.

Our main results, Theorems 8.1 and 8.4, can be stated as a structure theorem.

Structure Theorem. *For any closed symplectic Calabi-Yau 6-manifold X , there exist unique integer invariants $e_{A,g}(X)$, indexed by non-zero classes $A \in H_2(X, \mathbb{Z})$ and genera $g \geq 0$, such that the series (0.1) has the form*

$$GW(X) = \sum_{A \neq 0} \sum_{g \geq 0} e_{A,g}(X) \cdot GW_g^{elem}(q^A, t) \quad (0.2)$$

where $GW_g^{elem}(q, t)$ is the universal power series (3.4), which depends on g , but not on X . Furthermore, all coefficients $n_{A,g}$ in (0.1) are integers.

There is an extensive literature revolving around this GV conjecture. J. Bryan and R. Pandharipande have a series of papers about it, including two ([BP1] and [BP2]) relevant to our approach. For algebraic 3-folds, several BPS-type integer invariants have been defined using holomorphic bundles, including the Pandharipande-Thomas [PT] and Donaldson-Thomas invariants, with conjectural GV-type correspondences GW/DT/PT between them. For toric 3-folds, Maulik, Oblomkov, Okounkov and Pandharipande proved the GW/DT correspondence by calculating both sides explicitly in a computational tour de force [MOOP]. Pandharipande and Pixton [PP] established the GW/PT correspondence for CY complete intersections in products of projective spaces. Other instances have been observed when a change of variables in the GW series produces integer invariants, including a formula for Fano classes ($A \in H_2(X)$ with $c_1(X)A > 0$) in symplectic 6-manifolds proved by A. Zinger [Z], and the computation of Klemm and Pandharipande for Calabi-Yau 4-folds [KP]. In Section 8 we combine our result on the Calabi-Yau classes first with Zinger’s to obtain a GV-type formula for all symplectic 6-manifolds, and then with

Klemm and Pandharipande's to obtain a GV-type formula for genus zero invariants of semipositive symplectic manifolds.

We thank the referees for their meticulous reviews and numerous insightful suggestions.

1. CURVES IN SYMPLECTIC CALABI-YAU 6-MANIFOLDS

The Gromov-Witten invariants of a closed symplectic manifold (X, ω) are constructed in two steps. One first forms the universal moduli space and its stabilization-evaluation map *se*

$$\begin{array}{ccc} \overline{\mathcal{M}}(X) & \xrightarrow{\text{se}} & \bigsqcup_{g,n} \overline{\mathcal{M}}_{g,n} \times X^n \\ \pi \downarrow & & \\ \mathcal{J} & & \end{array} \quad (1.1)$$

over the space \mathcal{J} of ω -tame almost complex structures. This moduli space consists of equivalence classes (up to reparametrizations of the domain) of pairs (f, J) , where $f : C \rightarrow (X, J)$ is a stable pseudo-holomorphic map whose domain C is a nodal marked Riemann surface; it has components $\overline{\mathcal{M}}_{A,g,n}(X)$ labelled by the genus g of C , the number n of marked points, and the homology class $A = f_*[C] \in H_2(X; \mathbb{Z})$ with $\omega(A) \geq 0$.

It is frequently convenient to define the *energy* of a triple (A, g, n) by

$$E(A, g, n) = \max\{\omega(A), g, n\} \geq 0$$

and to restrict attention to the subset $\overline{\mathcal{M}}^E(X)$ “below energy E ”, meaning the union of all components with $E(A, g, n) \leq E$, and the corresponding fibers $\overline{\mathcal{M}}^{J,E}(X)$ of (1.1). The restriction of π to each $\overline{\mathcal{M}}_{A,g,n}(X)$ is proper, and the fibers carry a d -dimensional virtual fundamental class

$$[\overline{\mathcal{M}}_{A,g,n}^J(X)]^{vir} \in \check{H}_d(\overline{\mathcal{M}}_{A,g,n}^J(X); \mathbb{Q}), \quad (1.2)$$

where \check{H}_* denotes Čech homology and

$$d = 2c_1(X)A + (\dim X - 6)(1 - g) + 2n \quad (1.3)$$

(cf. [P]). These are assembled in a formal power series

$$GW_*(X) = \sum_{\omega(A)>0} \sum_{g=0}^{\infty} [\overline{\mathcal{M}}_{A,g,n}^J(X)]^{vir} t^{2g-2} q^A \in \check{H}_*(\overline{\mathcal{M}}^J(X), \Lambda) \quad (1.4)$$

where Λ is the “rational Novikov ring” generated by t and $\{q^A\}$ with $tq^A = q^At$ and $q^{A+B} = q^A q^B$, with finitely many non-zero terms below each energy level. Note that the first sum, over all positive $A \in H_2(X; \mathbb{Z})$ omits any trivial contributions (i.e. from the class $A = 0$). For consistency, the term “ J -holomorphic map” will always mean a *non-trivial* map, and a “ J -holomorphic curve” in X is the image of such a map. When f is an embedding, we identify C with its image $f(C)$ in X .

Gromov-Witten invariants are especially simple if $c_1(X) = 0$ and $\dim X = 6$; such spaces are often called *symplectic Calabi-Yau 6-manifolds*. In this case all terms in (1.3) vanish when $n = 0$; the virtual fundamental class then has dimension $d = 0$ for all g and A , and we drop n from the notation. In this case, the Gromov-Witten numbers $GW_{A,g}(X) \in \mathbb{Q}$ are obtained by pairing the virtual fundamental class with 1 in $\check{H}^0(\overline{\mathcal{M}})$, and (1.4) becomes a Λ -valued series

$$GW(X) = \sum_{\omega(A)>0} \sum_{g=0}^{\infty} GW_{A,g}(X) t^{2g-2} q^A. \quad (1.5)$$

We will use the homology form (1.4) in Section 2, and use (1.5) thereafter.

Using the terminology of [MS, §2.5], a point $x \in C$ is called an *injective point* for a map $f : C \rightarrow X$ if $df(x) \neq 0$ and $f^{-1}(f(x)) = \{x\}$ when C is smooth; if C is nodal we also require that x is not a

node. A pseudo-holomorphic map $f : C \rightarrow X$ from a nodal curve is *simple* if it has an injective point on each irreducible component. The open subset of simple maps in a moduli space will be denoted by the subscript *simple*; for example

$$\overline{\mathcal{M}}^{J,E}(X)_{\text{simple}} \quad (1.6)$$

denotes the set of simple maps in $\overline{\mathcal{M}}^{J,E}(X)$, while $\mathcal{M}(X)_{\text{simple}}$ denotes the open subset of $\overline{\mathcal{M}}(X)$ consisting of simple maps with smooth domain. The Micallef-White Theorem, [MW] or [MS, Prop. 2.5.1], implies that the set of injective points of a simple J -holomorphic map is open and dense in C , and hence the image under f of this set is a submanifold of X .

A map $f \in \overline{\mathcal{M}}^{J,E}(X)$ is *regular* if the linearization $D_{f,J}$, given by (4.9), is onto; it is a *regular embedding* if it is both regular and an embedding. Since the index of $D_{f,J}$ is 0 by (1.3), each regular map f has a well-defined sign $(f, J) = \pm 1$, given by the mod 2 spectral flow from $D_{f,J}$ to any invertible complex operator. Finally, by a Baire subset of the parameter space we mean a countable intersection of open and dense sets.

Lemma 1.1. *Let (X, ω) be a symplectic Calabi-Yau 6-manifold. Then for each $E > 0$ there is a Baire subset \mathcal{J}_E^* of \mathcal{J} such that for each $J \in \mathcal{J}_E^*$*

- (a) *All simple J -holomorphic maps below energy E are regular embeddings with disjoint images.*
- (b) *The projection π in (1.1) is a local diffeomorphism around each regular embedding.*

Part (a) of Lemma 1.1 is proved in the Appendix, and part (b) is proved as Proposition 5.1b; these proofs use standard techniques. For our purposes, it is best to enlarge \mathcal{J}_E^* to a set $\mathcal{J}_{\text{isol}}^E$ that emphasizes slightly different properties.

Definition 1.2. *Denote by $\mathcal{J}_{\text{isol}}^E$ the set of all $J \in \mathcal{J}$ such that the moduli space (1.6), with the Gromov topology, consists of isolated points that are embeddings with disjoint images.*

Corollary 1.3. *$\mathcal{J}_{\text{isol}}^E$ is dense in \mathcal{J} .*

Proof. For each $J \in \mathcal{J}_E^*$, each simple J -holomorphic map f is an embedding, and is regular by part (a) of Lemma 1.1. By part (b), each such point (f, J) is an isolated point of the fiber $\mathcal{M}^{J,E}(X)_{\text{simple}}$ of π . \square

Unfortunately, the images in X of the pseudo-holomorphic maps that appear in Lemma 1.1 may accumulate. To focus on the images, consider the space $\text{Subsets}(X)$ of all non-empty compact subsets of X . Fix a background Riemannian metric on X with distance function d . Then $\text{Subsets}(X)$ is a metric space with the Hausdorff distance, defined by

$$d_H(A, B) = \sup_{a \in A} \inf_{b \in B} d(a, b) + \sup_{b \in B} \inf_{a \in A} d(a, b)$$

for $A, B \subseteq X$. Let c be the “underlying curve” map

$$c : \overline{\mathcal{M}}(X) \longrightarrow \text{Subsets}(X) \times \mathcal{J} \quad (1.7)$$

that associates to each (f, J) the pair $(f(C), J)$ with $f(C)$ regarded as subset of X . This map c is continuous, and Gromov compactness implies that its restriction to $\overline{\mathcal{M}}^E(X)$ is proper. Let $\mathcal{C}(X)$, \mathcal{C}^E and $\mathcal{C}^{J,E}$ denote, respectively, the images of $\overline{\mathcal{M}}(X)$, $\overline{\mathcal{M}}^E(X)$ and $\overline{\mathcal{M}}^{J,E}(X)$ under c . With this notation, there is a commutative diagram

$$\begin{array}{ccc} \overline{\mathcal{M}}^E(X) & & \\ \downarrow \pi & \searrow c & \\ \mathcal{J} & \swarrow & \mathcal{C}^E. \end{array} \quad (1.8)$$

Viewed differently, convergence in \mathcal{C}^E defines a topology on $\overline{\mathcal{M}}^E(X)$ that we will call the “rough topology”.

In general, non-trivial J -holomorphic maps $f : C \rightarrow X$ from nodal curves can be multiple covers, and simple maps can converge to multiply covered maps. Definition 1.2 constrains how limits are multiply covered, as the following lemma shows.

Lemma 1.4. *For $J \in \mathcal{J}_{isol}^E$,*

- (a) *Every J -holomorphic map $f : C \rightarrow X$ below energy E is a composition $\varphi \circ \rho$ of a holomorphic map $\rho : C \rightarrow C_{red}$ of complex curves and a J -holomorphic embedding $\varphi : C_{red} \rightarrow X$. This decomposition is unique up to reparametrizations of C_{red} . When f is simple, $\varphi = f$.*
- (b) *If $f : C \rightarrow X$ is a limit in the rough topology of a sequence $\{f_n\}$ in $\overline{\mathcal{M}}^{J,E}(X)$ with $d_H(f_n, f) \neq 0$ for all n , then the factorization $f = \varphi \circ \rho$ has either $\deg \rho > 1$ or $\text{genus}(C) > \text{genus}(C_{red})$.*

Proof. Let $\tilde{f} : \tilde{C} \rightarrow X$ be the lift of f to the normalization \tilde{C} of C . The arguments of [MS, §2.5] show that \tilde{f} factors as $\tilde{f} = \varphi \circ \tilde{\rho}$ where $\varphi : C_{red} \rightarrow X$ is a simple J -holomorphic map from a smooth, possibly disconnected domain C_{red} , and $\tilde{\rho} : \tilde{C} \rightarrow C_{red}$ is a map of complex curves; $\tilde{\rho}$ may take some components of \tilde{C} to points. But the assumption $J \in \mathcal{J}_{isol}^E$ implies that C_{red} is connected and φ is an embedding; $\tilde{\rho}$ then descends to a holomorphic map $\rho : C \rightarrow C_{red}$, unique up to reparametrization, with $f = \varphi \circ \rho$.

Part (b) immediately follows from Gromov compactness and the fact that simple maps are isolated for $J \in \mathcal{J}_{isol}^E$, and therefore the limit map $f = \varphi \circ \rho$ is not simple. \square

To each class $A \in H_2(X, \mathbb{Z})$ in the positive cone $\omega(A) > 0$ we associate a positive integer

$$d(A) = \text{lcm} \{ k \in \mathbb{N}_+ \mid A = kB \text{ where } B \in H_2(X, \mathbb{Z}) \}$$

called the *degree* of A , where lcm denotes the lowest common multiple. Let $\Omega(d)$ be the number of prime factors of d , counted with multiplicity. For any map f from a genus g curve representing a degree d class, we define its *level* to be

$$\ell(f) = \Omega(d) + g. \quad (1.9)$$

The components of the moduli space are filtered by the degree and genus, and therefore by their level; the level filtration will be used frequently in later sections. For each E , the sets

$$\overline{\mathcal{M}}_m^E(X) = \{(f, J) \in \overline{\mathcal{M}}^E(X) \mid \ell(f) \leq m\} \quad (1.10)$$

filter $\overline{\mathcal{M}}^E(X)$, and their images under (1.8)

$$\mathcal{C}_m^E = c(\overline{\mathcal{M}}_m^E(X))$$

filter the image $\mathcal{C}^E = c(\overline{\mathcal{M}}^E(X))$.

For each $J \in \mathcal{J}_{isol}^E$, the map c , when applied to multiply-covered maps, decreases the level but respects the filtration; for such J the fiber $\mathcal{C}_m^{J,E}$ of $\mathcal{C}_m^E \rightarrow \mathcal{J}$ is the collection of embedded J -holomorphic curves with level at most m . In this notation, $m = 0$ corresponds to genus zero curves representing primitive classes, and \mathcal{C}_m^E is the collection of embedded pseudo-holomorphic curves in X with $g + \Omega(d) \leq m$ and energy at most E .

Lemma 1.5. *For any fixed $J \in \mathcal{J}_{isol}^E$,*

- (a) $\mathcal{C}_m^{J,E} \subseteq \mathcal{C}_{m+1}^{J,E}$ is a filtration of $\mathcal{C}^{J,E}$ with $\mathcal{C}_m^{J,E} = \mathcal{C}^{J,E}$ for m large.
- (b) $\mathcal{C}_m^{J,E}$ and $\mathcal{C}^{J,E} = \bigcup \mathcal{C}_m^{J,E}$ are compact countable subsets of the metric space $\mathcal{C}(X)$.
- (c) *For any neighborhood U of $\mathcal{C}_{m-1}^{J,E}$, the set $\mathcal{C}_m^{J,E} \setminus U$ is a finite collection of embedded J -holomorphic curves.*

In particular, there are finitely many genus zero J -holomorphic curves with energy less than E representing primitive classes.

Proof. Condition (a) is immediate by definition and the fact that only finitely many homology classes are represented by J -holomorphic maps below energy E .

Next, each set $\mathcal{C}_m^{J,E}$ is compact because it is the image of a compact set, namely the fiber $\overline{\mathcal{M}}_m^{J,E}(X)$ of (1.10), under the continuous map (1.8). Then $\mathcal{C}_m^{J,E} \setminus U$ is a closed subset of the compact metric

space $\mathcal{C}_m^{J,E}$, so any infinite sequence $\{C_i\}$ has an accumulation point C_0 . Because only finitely many homology classes are represented by J -holomorphic maps below energy E we may assume, after passing to a subsequence, that they all have the same genus and same homology class $[C_i] = k\beta$ for the same primitive class β and the same k with $\Omega(k) + g \leq m$. By Lemma 1.4 the limit is a multiple cover of a curve of a strictly lower level, but that's impossible because U is open. Thus $\mathcal{C}_m^{J,E} \setminus U$ is finite. Finally, taking U_k to be the $1/k$ tubular neighborhood of $\mathcal{C}_{m-1}^{J,E}$, we conclude that $\mathcal{C}_m^{J,E} \setminus \mathcal{C}_{m-1}^{J,E} = \bigcup_k (\mathcal{C}_m^{J,E} \setminus U_k)$ is countable, and hence $\mathcal{C}_m^{J,E}$ and $\mathcal{C}^{J,E}$ are countable. \square

2. CLUSTERS IN SYMPLECTIC MANIFOLDS

For each subset $S \subset \mathcal{J}$, we can consider the moduli space $\overline{\mathcal{M}}^S(X) = \pi^{-1}(S)$ over S or, with an energy bound,

$$\overline{\mathcal{M}}^{S,E}(X) = \{(f, J) \mid f \in \overline{\mathcal{M}}^{J,E}(X), J \in S\}.$$

A *decomposition* of the moduli space over S is a way of writing it as a finite disjoint union $\bigsqcup_i \mathcal{O}_i$ of subsets \mathcal{O}_i that are both open and closed in the rough topology. Such a decomposition induces an isomorphism

$$\check{H}_* \left(\overline{\mathcal{M}}^{S,E}(X); \Lambda \right) \cong \bigoplus_i \check{H}_*(\mathcal{O}_i; \Lambda) \quad (2.1)$$

and a corresponding decomposition of the GW invariant for every $J \in S$, namely

$$GW_*^E(X) = \sum_i GW_*^E(\mathcal{O}_i), \quad (2.2)$$

where $GW_*^E(\mathcal{O}_i)$ is the component of $GW_*^E(X)$ in $\check{H}_*(\mathcal{O}_i; \Lambda)$, written as the series (1.4) under the isomorphism (2.1), and GW_*^E is the truncation of GW_* below energy E .

A simple way to obtain such sets \mathcal{O} is to choose an open set U in $\mathcal{C}(X)$ whose boundary does not intersect $\mathcal{C}^{J,E}$.

Lemma 2.1. *Each open subset U of $\mathcal{C}(X)$ with $\partial U \cap \mathcal{C}^{J,E} = \emptyset$ has a well-defined contribution $GW_*^E(U, J)$. The collection of J for which $\partial U \cap \mathcal{C}^{J,E} = \emptyset$ is open in the C^0 topology on \mathcal{J} , and the contribution $GW_*^E(U, J)$ is locally constant as a function of J .*

Proof. The assumption implies that the intersections of both U and its complement U^c with $\mathcal{C}^{J,E}$ are open subsets of the image $\mathcal{C}^{J,E}$ of c . Since c is continuous, $\mathcal{O} = c^{-1}(U)$ and $\mathcal{O}^c = c^{-1}(U^c)$ are open and closed subsets of $\overline{\mathcal{M}}^{J,E}(X)$. Define $GW_*^E(U, J) = c_* GW_*^E(\mathcal{O}, J)$ by formula (2.2) associated to the decomposition $\mathcal{O} \sqcup \mathcal{O}^c$.

Next note that the condition $\partial U \cap \mathcal{C}^{J,E} = \emptyset$ is an open condition on J . To see why, fix U and E . If this condition held for some J , but failed to hold for a sequence $J_k \rightarrow J$ in C^0 , there would be a sequence of J_k -holomorphic curves C_k in ∂U with energy $E(A, g, n)$ bounded by E . Applying Gromov compactness, one could find a subsequence converging to a J -holomorphic curve C . But ∂U and $\mathcal{C}^{J,E}$ are both closed in the Gromov topology, so the limit curve C would lie in $\partial U \cap \mathcal{C}^{J,E}$, contradicting the hypothesis.

Therefore there exists a C^0 ball V around J in \mathcal{J} such that $\partial U \cap \mathcal{C}^{E,V} = \emptyset$. This gives rise to the same decomposition $\mathcal{O} \sqcup \mathcal{O}^c$ but now of the moduli space $\overline{\mathcal{M}}^{V,E}(X)$ over the entire ball V , therefore the contribution of U is defined for each $J \in V$ and it is constant on V . \square

The geometric content of the contribution $GW_*^E(U, J)$ is most clearly seen by choosing U of the form $B(C, \varepsilon) \times \mathcal{J}$ for a ball $B(C, \varepsilon)$ of small radius ε in the Hausdorff distance centered at a J -holomorphic curve C . The subset of $\overline{\mathcal{M}}^J(X)$ that lies in $c^{-1}(U)$ is then a collection of J -holomorphic maps whose images are uniformly ε -close to C in X . We will call such a collection a cluster if the following properties hold.

Definition 2.2. A cluster $\mathcal{O} = (C, \varepsilon, J)$ in X below energy E (an “ E -cluster”) consists of an almost complex structure $J \in \mathcal{J}$, an embedded J -holomorphic curve C and a radius $\varepsilon > 0$ with the following properties:

- (a) All non-constant J -holomorphic maps in the ball $B(C, \varepsilon)$ with energy $\leq E$ represent $k[C]$ for some $k \geq 1$, and have genus $g \geq g(C)$;
- (b) C is the only J -holomorphic map in its degree and genus in the ball $B(C, \varepsilon)$.
- (c) There are no J -holomorphic curves with energy $\leq E$ at precisely ε Hausdorff distance from C .

The curve C is called the *core* of the cluster. Note that, by the definition of Hausdorff distance, curves in $B(C, \varepsilon)$ lie in the ε -tubular neighborhood of C in X .

The next lemma shows that small balls in the Hausdorff metric are often clusters. In fact, conditions (a) and (b) of Definition 2.2 are automatic for small ε when C is regular. Condition (c) is the important one: it implies that $\mathcal{O} = (C, \varepsilon, J)$ has a well-defined contribution

$$GW_*^E(\mathcal{O}) \in \check{H}_0(\mathcal{O}; \Lambda). \quad (2.3)$$

Lemma 2.3 (Cluster existence). *For each $J \in \mathcal{J}_{isol}^E(X)$ and each simple J -holomorphic curve C , the set S of $\varepsilon > 0$ for which the ball $B(C, \varepsilon)$ is an E -cluster is open and dense in a non-empty interval $[0, \varepsilon_C]$, and the contribution (2.3) is locally constant on S .*

Proof. By definition, for any $J \in \mathcal{J}_{isol}^E(X)$, all simple J -holomorphic curves are embedded and isolated in their degree and genus. Since $\mathcal{C}^{J,E}$ is compact, and an embedded curve C can appear as an accumulation point only of curves representing $k[C]$ and having genus at least that of C , Lemma 1.4 implies that there is an $\varepsilon_C > 0$ such that conditions (a) and (b) of Definition 2.2 hold for all $\varepsilon \leq \varepsilon_C$.

Finally, by Lemma 1.5, the image of $\mathcal{C}_m^{J,E}$ under the distance function $B(C, \varepsilon_C) \rightarrow [0, \varepsilon_C]$ is countable and compact, so its complement – which is the set of ε that satisfy condition (c) of Definition 2.2 – is open and dense. \square

Proposition 2.4 (Cluster decompositions). *Given E and $J \in \mathcal{J}_{isol}^E$, each open subset U of $\mathcal{C}(X)$ such that $\partial U \cap \mathcal{C}^{J,E} = \emptyset$ has a finite E -cluster decomposition $\{\mathcal{O}_i = (C_i, J, \varepsilon_i)\}$, and hence*

$$GW_*^E(U) = \sum_i GW_*^E(\mathcal{O}_i). \quad (2.4)$$

Proof. We will inductively construct cluster decompositions of the sets $U_m = U \cap \mathcal{C}_m^{J,E}$, beginning with the trivial case $U_{-1} = \emptyset$. This induction is finite because $\mathcal{C}_m^{J,E} = \mathcal{C}^{J,E}$ for m sufficiently large by Lemma 1.5a.

Suppose that $\{B_i = B(C_i, \varepsilon_i)\}$ is a cluster decomposition of U_m . This means that the balls B_i are disjoint, that the compact set U_m lies in $V = \bigsqcup B_i$, and there are no J -holomorphic curves on $\partial V = \bigsqcup \partial B_i$. Lemma 1.5 shows that $U_{m+1} \setminus V$ is a finite collection of curves $\{C_j\}$. None of these C_j lie on ∂V , so we can choose radii $\varepsilon_j > 0$ such that the balls $B'_j = B(C_j, \varepsilon_j)$ are clusters (by Lemma 2.3) and are disjoint from each other and from the balls B_i . These clusters B'_j , together with the original B_i are a cluster decomposition of U_{m+1} , completing the induction step. \square

Corollary 2.5 (Cluster refinement). *Fix any cluster $\mathcal{O} = (C, J, \varepsilon)$ with $J \in \mathcal{J}_{isol}^E$. For any $\varepsilon' \in (0, \varepsilon)$ for which $\mathcal{O}' = (C, J, \varepsilon')$ is a cluster, there exists finitely many higher level clusters $\{\mathcal{O}_i = (C_i, J, \varepsilon_i)\}$ such that*

$$GW_*^E(\mathcal{O}) = GW_*^E(\mathcal{O}') + \sum_i GW_*^E(\mathcal{O}_i). \quad (2.5)$$

Proof. Consider $U = B(C, \varepsilon) \setminus B(C, \varepsilon')$. Because both \mathcal{O} , \mathcal{O}' are clusters, (a) there are no curves in $\mathcal{C}^{J,E}$ on ∂U and (b) all the curves in $\mathcal{C}^{J,E} \cap U$ have strictly higher level compared to that of C (as C is the only bottom level curve in both clusters). By Proposition 2.4, there exists a cluster decomposition (2.4) of U , where the \mathcal{O}_i are strictly higher level. But condition (a) also implies that U and \mathcal{O}' give a decomposition of \mathcal{O} so $GW_*(\mathcal{O}) = GW_*(\mathcal{O}') + GW_*(U)$, which implies (2.5). \square

3. ELEMENTARY CLUSTERS AND THEIR CONTRIBUTIONS

The GW series can be explicitly calculated for one very special type of embedded curve. In this section we describe how this can be done by combining ideas already in the literature. We first construct such “elementary curves” using a remarkable non-integrable almost complex structure discovered by Junho Lee [L], and then point out that the GW series for these curves has been calculated by J. Bryan and R. Pandharipande [BP2].

Definition 3.1. *In a symplectic Calabi-Yau 6-manifold X , a cluster $\mathcal{O} = (C, \varepsilon, J)$ is called elementary if*

- (a) *The core curve C is balanced, meaning that its normal bundle splits as $N_C = L \oplus L$ in such a way that the normal operator D^N , given by (4.11) below, splits as $D' \oplus D'$.*
- (b) *The only non-trivial J -holomorphic maps into $B(C, \varepsilon)$ are multiple covers of the embedding $C \hookrightarrow X$.*
- (c) *For each such cover ρ , the pullback operator $\rho^* D'$ is injective (C is super-rigid).*

Property (c) implies that the core C of an elementary cluster is a regular J -holomorphic map, and one can show it also implies (b) for sufficiently small $\varepsilon > 0$ by the rescaling argument of [IP2]. Property (a) allows us to actually calculate the GW contribution.

When C is a rational curve, the unit disk bundle in $\mathcal{O}(-1) \oplus \mathcal{O}(-1)$ is an elementary cluster. The following proposition uses non-integrable almost complex structures to construct similar examples for any curve. The proof begins with the choice of a spin structure on C , i.e. a holomorphic line bundle $L \rightarrow C$ together with a (holomorphic) identification of L^2 with the canonical bundle K_C of C .

Proposition 3.2. *For every smooth complex curve C , there exists an elementary cluster whose core is C .*

Proof. Fix a curve C of genus g , a spin structure L , and a Kähler structure (J, g, ω) on the total space Y of $L \oplus L \rightarrow C$ compatible with its holomorphic structure. The canonical bundle $K_Y = \pi^*(L^{-2} \otimes K_C)$ of Y is then trivial, so Y is a Calabi-Yau 6-manifold. We will perturb J to obtain an almost complex structure on the unit disk bundle $U \subset Y$ that makes U an elementary cluster.

First consider the total space Z of $p : L \rightarrow C$. Its canonical bundle $K_Z = T^{2,0}Z$ is the pullback $p^*(L^{-1} \otimes K_C) = p^*(L)$, which has a canonical section β that vanishes transversally along the zero section. Pullback β , regarded as a 2-form, by the bundle projections $p_1, p_2 : Y \rightarrow Z$ onto the first and second copy of L and set $\alpha = p_1^*\beta + p_2^*\beta$. Then α is a closed $(2, 0)$ -form on Y that vanishes to first order along the zero section of $L \oplus L$, which is the core curve C of the disk bundle $U \subset Y$. Following Junho Lee, define a bundle map $K_\alpha : TY \rightarrow TY$ by $g(u, K_\alpha v) = (\alpha + \bar{\alpha})(u, v)$ for all $u, v \in TY$ and set

$$J_\alpha = (Id + JK_\alpha)^{-1}J(Id + JK_\alpha).$$

Then J_α is an ω -tame almost complex structure on U after replacing α by $t\alpha$ for small $t > 0$ (cf. Section 2 of [L]). In this context, Lee proved that *the image of every non-trivial J_α -holomorphic map into D is everywhere tangent to $\ker K_\alpha$* (cf. [L, (2.4)]). It is straightforward to check that at each point $p \in Y$ not on the zero section, $\ker K_\alpha$ is vertical. Consequently, any map whose image is tangent to $\ker K_\alpha$ lies in a fiber of $Y \rightarrow C$ or in the zero section. We conclude that every J_α -holomorphic map into U that represents $k[C]$ for some $k \neq 0$ is a map into the core curve C of the disk bundle U .

Along the zero section of Y and for $v \in TC$, $\nabla_v K_\alpha$ decomposes under the splitting $TC \oplus L \oplus L$ as $0 \oplus \nabla_v K_\beta \oplus \nabla_v K_\beta$. Correspondingly, the normal operator D^N , given by (4.11), splits as $D_\beta^N \oplus D_\beta^N$ as can be seen from [LP1, (8.4)]. As in Section 4 of [LP1], the normal projection of D_β is the sum $\bar{\partial} + R_\beta$ of the $\bar{\partial}$ -operator on L and a bundle map $R_\beta : L \rightarrow T^{0,1}C \otimes L$ that satisfies $JR_\beta = -R_\beta J$. The injectivity condition (c) of Definition 3.1 is then exactly the statement of Proposition 8.6 of [LP1]. \square

By Lemma 2.1 an elementary cluster has a well-defined invariant $GW_*(\mathcal{O}) \in \check{H}_*(\mathcal{O}; \Lambda)$, independent of ε , and a corresponding Λ -valued series $GW(\mathcal{O})$ as in (1.5). As usual, there is an associated *disconnected* invariant

$$Z(\mathcal{O}) = \exp(GW(\mathcal{O}))$$

obtained by exponentiating in the Novikov ring. It turns out that $Z(\mathcal{O})$ is the more easily calculated.

Proposition 3.3. *The disconnected GW invariant of an elementary cluster \mathcal{O} whose core C has genus g is given by*

$$Z(\mathcal{O}) = 1 + \sum_{d \geq 1} \sum_{\mu \vdash d} \prod_{\square \in \mu} \left(2 \sin \frac{h(\square)t}{2} \right)^{2g-2} q^{dC}, \quad (3.1)$$

where the second sum is over all partitions μ of d , the product is over the boxes in the Ferrers diagram of μ , and $h(\square)$ is the hooklength of $\square \in \mu$.

Proof. Because the linearization on covers of the core curve C is injective, the contribution to the GW invariant of its multiple covers can be calculated using the Euler class of Taubes obstruction bundle (this is the 6-dimensional version of the setup in [LP2], and a special case of Theorem 1.2 in [Z]).

Consider the moduli space $\overline{\mathcal{M}}_{d,\chi}^\circ(C)$ of degree d holomorphic maps ρ to C whose domain is possibly disconnected and has Euler characteristic χ , and where ρ is non-trivial on each connected component. This carries a virtual fundamental cycle $[\overline{\mathcal{M}}_{d,\chi}^\circ(C)]^{vir}$ of even complex dimension $b = d(2-2g) - \chi$. The operators D^N and D' of Definition 3.1 induce families of real operators \mathcal{D}^N and \mathcal{D}' over $\overline{\mathcal{M}}_{d,\chi}^\circ(C)$ whose fibers at ρ are the pullback operators $\rho^* D^N$ and $\rho^* D'$. By Definition 3.1a, the corresponding index bundles satisfy $\text{Ind } \mathcal{D}^N = \text{Ind } \mathcal{D}' \oplus \text{Ind } \mathcal{D}'$. *A priori*, these are real virtual bundles, but Definition 3.1c insures that the Taubes obstruction bundle $Ob = -\text{Ind}_{\mathbb{R}} \mathcal{D}^N$ is an actual vector bundle of rank b , equal to the direct sum of two copies of $Ob' = -\text{Ind}_{\mathbb{R}} \mathcal{D}'$. The bundles Ob and Ob' each come with a canonical orientation determined, on each connected component of the space of covers, by the spectral flow to an injective complex operator over one fixed cover ρ . Computing this spectral flow along a path of operators that respect the direct sum decomposition, one sees that $Ob = Ob' \oplus Ob'$ as oriented real bundles.

With this notation, the elementary contribution is equal to the integral of the Euler class

$$Z_{d,\chi}(\mathcal{O}) = \int_{[\overline{\mathcal{M}}_{d,\chi}^\circ(C)]^{vir}} e(Ob),$$

where

$$e(Ob) = e(Ob' \oplus Ob') = e(Ob') \cup e(Ob') = (-1)^{b/2} c_b(Ob' \otimes_{\mathbb{R}} \mathbb{C}).$$

This Chern class factors through K -theory (in general Euler classes do not). Because $D' = \bar{\partial}_L + R$ is a 0'th order deformation of the complex operator $\bar{\partial}_L$, the complexification of the index bundles of D' and $\bar{\partial}_L$ are equal in K -theory, so

$$c_b(Ob' \otimes_{\mathbb{R}} \mathbb{C}) = c_b(-\text{Ind } \bar{\partial}_L \otimes_{\mathbb{R}} \mathbb{C}) = c_b(-\text{Ind } \bar{\partial}_L \oplus (-\text{Ind } \bar{\partial}_L)^*).$$

Combining the last three displayed equations gives

$$Z_{d,\chi}(\mathcal{O}) = (-1)^{b/2} \int_{[\overline{\mathcal{M}}_{d,\chi}^\circ(C)]^{vir}} c_b(-\text{Ind } \bar{\partial}_L \oplus (-\text{Ind } \bar{\partial}_L)^*). \quad (3.2)$$

The right-hand side of (3.2) can be evaluated using equivariant techniques. The torus $T = \mathbb{C}^* \times \mathbb{C}^*$ acts on the total space Y of the holomorphic bundle $N_C = L \oplus L$. With the antidiagonal \mathbb{C}^* -action, Y is an equivariant local Calabi-Yau 3-fold. Bryan and Pandharipande defined a ‘residue’ generating function $Z^T(Y)$ whose coefficients

$$Z_{d,\chi}^T(Y) = \int_{[\overline{\mathcal{M}}_{d,\chi}^\circ(C)]^{vir}} c_b(-\text{Ind } \bar{\partial}_{L \oplus L})$$

are *equivariant* integrals (defined by localization). They proceeded to express them in terms of *ordinary* integrals:

$$Z_{d,\chi}^T(Y) = \sum_{b_1+b_2=b} \int_{[\overline{\mathcal{M}}_{d,\chi}^\circ(C)]^{vir}} (t_1/t_2)^{(b_2-b_1)/2} c_{b_1}(-\text{Ind } \bar{\partial}_L) c_{b_2}(-\text{Ind } \bar{\partial}_L) \quad (3.3)$$

where t_1, t_2 are the weights of the action (cf. page 105 of [BP2]). For the antidiagonal action $t_1 = -t_2$, (3.3) reduces to (3.2) after noting that $c_k(E^*) = (-1)^k c_k(E)$ for $E = -\text{Ind } \bar{\partial}_L$. On the other hand, for the antidiagonal action, Bryan and Pandharipande also explicitly calculated (3.3) to be the coefficient of the series appearing on the right hand side of (3.1) (Corollary 7.3 of [BP2]). This completes the proof. \square

The series (3.1) is a universal power series that depends only on the genus g of C . Thus we set

$$GW_g^{elem}(q, t) = \log Z_g^{elem}(q, t), \quad (3.4)$$

where

$$Z_g^{elem}(q, t) = 1 + \sum_{d \geq 1} \sum_{\mu \vdash d} \prod_{\square \in \mu} \left(2 \sin \frac{h(\square)t}{2} \right)^{2g-2} q^d. \quad (3.5)$$

In fact, taking log of (3.5) and separating the $d = 1$ term of the series,

$$GW_g^{elem}(q, t) = q \left(2 \sin \frac{t}{2} \right)^{2g-2} + \sum_{d \geq 2} \sum_{h \geq g} GW_{d,h}(g) q^d t^{2h-2}, \quad (3.6)$$

for some coefficients $GW_{d,h}(g) \in \mathbb{Q}$. Since the coefficient of the leading term qt^{2g-2} is $+1$, the core curve C of any elementary cluster has sign $C > 0$.

Now apply the ‘‘BPS transform’’, which takes an arbitrary element of the Novikov ring to another by

$$\sum_{A,g} N_{A,g} t^{2g-2} q^A = \sum_{A,g} n_{A,g} \sum_{k=1}^{\infty} \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{2g-2} q^{kA}.$$

This transform is well-defined and invertible (Proposition 2.1 of [BP1]). Thus for an elementary cluster \mathcal{O} whose core C has genus g we can write

$$GW(\mathcal{O}) = GW_g^{elem}(q^C, t) = \sum_{d \neq 0} \sum_h n_{d,h}(g) \sum_{k=1}^{\infty} \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{2h-2} q^{kdC} \quad (3.7)$$

for uniquely determined coefficients $n_{d,h}(g)$ that are, *a priori*, rational numbers. These coefficients have been explicitly calculated for low degree ($d \leq 2$) and for low genus ($g \leq 1$). A combinatorics argument now handles the case $g \geq 2$, yielding a basic fact:

Proposition 3.4. *The local Gopakumar-Vafa conjecture is true for elementary clusters. More specifically, the coefficients of the series (3.7) associated with a genus g elementary cluster \mathcal{O} satisfy*

- (a) (Integrality) $n_{d,h}(g) \in \mathbb{Z}$.
- (b) (Finiteness) for each d fixed, $n_{d,h}(g) = 0$ for $h < g$ or h large.
- (c) For $g = 0$, all $n_{d,h}(g)$ vanish except $n_{1,0}(g) = 1$.
- (d) For $g = 1$, all $n_{d,h}(g)$ vanish except $n_{d,1}(g) = 1$ for each $d \geq 1$.

Proof. When the core curve has genus zero and normal bundle $O(-1) \oplus O(-1)$, \mathcal{O} is an elementary cluster and its contribution to the GW invariant was first calculated by Faber and Pandharipande. Specifically, letting $c(h, d)$ denote the coefficient of $q^d t^{2h-2}$ in (3.6) with $g = 0$, equations (34), (35) and the middle displayed equation on page 192 of [FP] imply the formulas

$$c(h, d) = d^{2h-3} c(h, 1) \quad \sum c(h, 1) t^{2h-2} = \left(2 \sin \left(\frac{t}{2} \right) \right)^{-2}$$

(these are equations (1) and (2) in [P]). Consequently, the genus 0 elementary GW series is

$$GW_0^{elem}(q, t) = \sum_{h,k} c(h, k) t^{2h-2} q^k = \sum_k \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{-2} q^k.$$

Comparing with (3.7) gives (c).

For genus $g = 1$, (3.5) reduces to the generating function for the number $p(d)$ of unordered partitions of d :

$$Z_1^{elem}(q, t) = 1 + \sum_{d \geq 1} p(d) q^d = \prod_{d=1}^{\infty} \left(\frac{1}{1 - q^d} \right).$$

Hence

$$GW_1^{elem}(q, t) = \log Z_1^{elem}(q, t) = - \sum_{d=1}^{\infty} \log(1 - q^d) = \sum_{d=1}^{\infty} \sum_{k=1}^{\infty} \frac{q^{kd}}{k}.$$

Comparing with (3.7) gives (d).

In the higher genus case, both (a) and (b) are consequences of an algebraic fact about power series with integral coefficients that follows by combining several results in the paper [PT] of Pandharipande and Thomas. Making the change of variable $Q = e^{it}$, (3.7) becomes

$$\log Z_g^{elem} = \sum_{d \geq 1} \sum_h n_{d,h}(g) \sum_{k \geq 1} \frac{(-1)^{h-1}}{k} (Q^k + Q^{-k} - 2)^{h-1} q^{kd}.$$

On the other hand, for $g \geq 1$, (3.5) becomes

$$Z_g^{elem} = 1 + \sum_{d \geq 1} \sum_{\mu \vdash d} \prod_{\square \in \mu} (-1)^{g-1} \left(Q^{h(\square)} + Q^{-h(\square)} - 2 \right)^{g-1} q^d = \sum_{d=0}^{\infty} \sum_n A_{n,d} Q^n q^d \quad (3.8)$$

where, for each d , the inner sum is a Laurent polynomial in Q with integer coefficients $A_{n,d}$. These coefficients $A_{n,d}$ uniquely determine the numbers $n_{d,h}$. But by Theorem 3.20 of [PT], the integrality of the $A_{n,d}$ implies that all of the $n_{d,h}$ are also integers. Thus statement (a) holds.

For $g \geq 2$, the coefficient of q^d in (3.5) is a Taylor series in t^2 , $Z = 1 + t^{2g-2}q + O(t^{2g})$, so $\log Z = t^{2g-2}q + O(t^{2g})$. Comparing with (3.7) one sees that $n_{d,h}(g) = 0$ for all $h < g$, as in (b).

Finally, for genus $g \geq 2$, the inner sum in (3.8) is a Laurent polynomial in Q , symmetric in $Q \rightarrow Q^{-1}$, and with degree bounded by $(g-1) \sum h(\square) \leq d^2(g-1)$ (since for a partition of d , the hooklength $h(\square)$ of each box is at most d). This property is preserved under taking the log:

$$\log Z_g^{elem} = \log \sum_d \sum_n a_{d,n} Q^n q^d = \sum_{d \geq 1} \sum_{h \geq g} n_{d,h}(g) \sum_{k \geq 1} \frac{(-1)^{h-1}}{k} (Q^k + Q^{-k} - 2)^{h-1} q^{kd},$$

where $|n| \leq (g-1)d^2$. As in the proof of Lemma 3.12 of [PT], this implies the vanishing of $n_{d,h}(g)$ for large h . In fact, a proof by induction on d using the above bound implies that $n_{d,h}(g) = 0$ for $h-1 > d^2(g-1)$. \square

4. ANALYTIC PRELIMINARIES

Consider next the universal moduli space of simple maps (with smooth, connected domains)

$$\begin{array}{ccc} \mathcal{M}_{simple} & & \\ \downarrow \pi & & \\ \mathcal{J} & & \end{array} \quad (4.1)$$

with the projection $\pi(f, J) = J$. Note that \mathcal{M}_{simple} is an open subset of the universal moduli space $\overline{\mathcal{M}}(X)$ of (1.1), and simple maps have trivial automorphism groups (cf. [MS], Proposition 2.5.1). It contains the open subset \mathcal{M}_{emb} of maps that are embeddings.

To set up the analysis, we first work locally around a point $p = (f, J)$ in the moduli space consisting of a simple J -holomorphic map $f : C \rightarrow X$ whose domain is a smooth, connected marked complex curve $C = (\Sigma, x_1, \dots, x_n, j)$, and where J is a smooth almost complex structure and j is in the space $\mathcal{J}(\Sigma)$ of complex structures on Σ .

4.1. Slice and linearization. The moduli space $\mathcal{M}_{g,n}(X)$ is naturally a subset of the quotient of $Map(\Sigma, X) \times \mathcal{J}(\Sigma) \times \mathcal{J}$ by the action of the diffeomorphism group $Diff(\Sigma, \mathbf{x})$ of Σ that preserve each point in the set $\mathbf{x} = \{x_1, \dots, x_n\}$ of marked points. In practice, one chooses a local slice for the diffeomorphism action and regards the moduli space locally as a subset of the slice. For now, we assume that C has no automorphisms; this assumption will be removed at the end of this subsection. To define a slice, choose local holomorphic coordinates on a ball $B \subset \mathcal{M}_{g,n}$ centered at $[C] \in \mathcal{M}_{g,n}$. Then there is a universal deformation $\gamma : \mathcal{U} \rightarrow B$ with sections x_1, \dots, x_n . This means, in particular, that the central fiber $\gamma^{-1}(0)$ is identified with C as a marked complex curve, and every small deformation of C is equivalent under $Diff(\Sigma, \mathbf{x})$ to one and only one fiber of γ over B . Fix a smooth trivialization of \mathcal{U} in which the universal deformation is $\gamma : B \times C \rightarrow B$. This trivialization, regarded as a family of complex structures on (Σ, \mathbf{x}) , defines an embedding $B \rightarrow \mathcal{J}(\Sigma)$ whose image \mathcal{S}_C is a local slice for the action of the diffeomorphism group on $\mathcal{J}(\Sigma)$. The linearization of this embedding gives isomorphisms

$$T_C \mathcal{M}_{g,n} \cong T_0 B \xrightarrow{\cong} T_j \mathcal{S}_C. \quad (4.2)$$

Furthermore, the tangent space to the orbit \mathcal{O}_j of $Diff(\Sigma, \mathbf{x})$ on $\mathcal{J}(\Sigma)$ is the image of

$$\bar{\partial}_{TC} : \Omega_{\mathbf{x}}^0(TC) \rightarrow \Omega^{0,1}(TC),$$

where $\Omega_{\mathbf{x}}^0(TC)$ denotes the space of smooth sections of TC that vanish at the marked points.

The slice \mathcal{S}_C is transverse to this orbit at j , giving an isomorphism

$$T_j \mathcal{S}_C \cong T_j \mathcal{J}(\Sigma)/\mathcal{O}_j = \Omega^{0,1}(TC)/\text{im } \bar{\partial}_{TC} = H^{0,1}(TC), \quad (4.3)$$

where the last equality defines the vector space $H^{0,1}(TC)$. Consequently, the map

$$D_C : \Omega_{\mathbf{x}}^0(TC) \oplus T_j \mathcal{S}_C \rightarrow \Omega^{0,1}(TC) \quad (4.4)$$

defined by $D_C(\zeta, k) = \bar{\partial}_{TC}\zeta + jk$ is a complex-linear isomorphism.

To include maps, fix (f_0, J_0) , where f_0 is a J_0 -holomorphic map whose domain C_0 has $\text{Aut}(C_0) = 1$. Then

$$Slice_0 = Map(\Sigma, X) \times \mathcal{S}_{C_0} \times \mathcal{J} \quad (4.5)$$

is a local slice for the action of $Diff(\Sigma, \mathbf{x})$ on $Map(\Sigma, X) \times \mathcal{J}(\Sigma) \times \mathcal{J}$. This comes with a projection

$$\pi : Slice_0 \rightarrow \mathcal{J}$$

defined by $\pi(f, J) = J$, and a complex vector bundle $\mathcal{F} \rightarrow Slice_0$ whose fiber over $p = (f, J)$ is $\Omega^{0,1}(f^*TX)$. Near (f_0, J_0) , the moduli space, considered as a subset of $Slice_0$, is the zero set of the section Φ of \mathcal{F} defined by

$$\Phi(f, J) = \bar{\partial}_J f. \quad (4.6)$$

Under the isomorphism (4.2), the tangent bundle to the slice can be written as

$$T_{Slice_0} = \mathcal{E} \oplus T\mathcal{J}, \quad (4.7)$$

where $\mathcal{E} \rightarrow Slice_0$ is the complex bundle whose fiber at $p = (f, J)$ is $\mathcal{E}_p = \Omega^0(f^*TX) \oplus T_C \mathcal{M}_{g,n}$.

The linearization of the J -holomorphic map equation (4.6) at a solution p on the slice is the real operator $\mathcal{L}_p : \mathcal{E}_p \oplus T_J \mathcal{J} \rightarrow \mathcal{F}_p$ given by

$$\mathcal{L}_p(\xi, K) = D_p \xi + \frac{1}{2} K \circ f_* \circ j, \quad (4.8)$$

where $D_p : \mathcal{E}_p \rightarrow \mathcal{F}_p$ is the linearization under variations that fix J . Explicitly, D_p applied to $\xi = (\zeta, k) \in \Omega^0(f^*TX) \oplus T_C \mathcal{M}_{g,n}$ is

$$D_p(\xi)(w) = \frac{1}{2} [\nabla_w \zeta + J \nabla_{jw} \zeta + (\nabla_\zeta J)(f_*(jw)) + J f_* k(w)], \quad (4.9)$$

where ∇ is any torsion-free connection on TX ; at a solution p , D_p is independent of the connection (Lemma 2.1.2 [IS1] or Lemma 1.2.1 of [IS2]). Both \mathcal{L}_p and D_p depend on the J only through the 1-jet

of J along the image of f , and the spaces \mathcal{E}_p and \mathcal{F}_p depend only on the 0-jet of J along the image of f . Furthermore, when the variation $\zeta = f_*\zeta^T$ is tangent to C , (4.9) reduces to

$$D_p(f_*\zeta^T, k) = f_*D_C(\zeta^T, k), \quad (4.10)$$

where D_C is the isomorphism (4.4). Consequently, when $f : C \rightarrow X$ is a J -holomorphic immersion with normal bundle $N_C = f^*TX/TC$, the linearization (4.9) uniquely descends to a normal operator

$$D_p^N : \Gamma(N_C) \rightarrow \Omega^{0,1}(N_C). \quad (4.11)$$

Now consider a simple map $f_0 : C_0 \rightarrow X$ whose automorphism group $\text{Aut}(C_0)$ is non-trivial. As noted after (1.6), the injective points of f_0 are dense on each component. Hence we can choose ℓ additional injective points on C_0 so that the marked curve $\tilde{C}_0 = (C_0, x_1, \dots, x_{n+\ell})$ has $\text{Aut}(\tilde{C}_0) = 1$. We can then fix a trivialized local universal deformation of \tilde{C}_0 over a ball \tilde{B} as above. Set $\tilde{\Sigma} = (\Sigma, x_1, \dots, x_{n+\ell})$, and for each image point $y_i = f_0(x_i)$, $n < i \leq n + \ell$, choose a 4-ball V_i through y_i transverse to $f_0(C_0)$. Standard results (cf. Section 3.4 of [MS]) show that the space $\text{Map}_\ell(\tilde{\Sigma}, X)$ of maps satisfying $f(x_i) \in V_i$ for $n < i \leq n + \ell$ is locally a manifold near f_0 , and hence

$$\mathcal{Slice}_0 = \text{Map}_\ell(\tilde{\Sigma}, X) \times \tilde{B} \times \mathcal{J}$$

is a local slice for the action of $\text{Diff}(\tilde{\Sigma})$ on $\text{Map}_\ell(\tilde{\Sigma}, X) \times \mathcal{J}(\tilde{\Sigma}) \times \mathcal{J}$. Thus defined, each point in the slice is a pair $p = (f, J)$ where $f : \tilde{C} \rightarrow X$ is a map whose domain has no non-trivial automorphisms. The linearization

$$D_p : \mathcal{E}_p \rightarrow \mathcal{F}_p \quad (4.12)$$

is still given by (4.9), but where \mathcal{E} is now the bundle over the slice whose fiber over $p = (f, J)$ is

$$\mathcal{E}_p = \left\{ \zeta \in \Omega^0(f^*TX) \mid \zeta(x_i) \in TV_i \text{ for all } n+1 \leq i \leq n+\ell \right\} \oplus T_{\tilde{C}}\mathcal{M}_{g,n+\ell}. \quad (4.13)$$

For notational simplicity, we will henceforth write \tilde{C} as C , and always restrict to sections ζ with $\zeta(x_i) \in TV_i$. All the variations we construct in this and subsequent sections will be supported away from all marked points of \tilde{C} .

4.2. Completions, extensions and adjoints. Without changing notation, we now replace the space $\text{Map}_\ell(\tilde{\Sigma}, X)$ and the bundle \mathcal{E} by their completions in a Sobolev (k, p) norm with $p \geq 2$ and k large enough to ensure that elements in the completions are at least C^3 . Similarly, replace the bundle \mathcal{F} by its Sobolev $(k-1, p)$ completion, and the manifold \mathcal{J} by its completion in the C^m norm for some large m . As in Sections 3.1-3.2 of [MS], \mathcal{L}_p and D_p extend to bounded linear operators by the formulas (4.9).

Given a slice \mathcal{Slice}_0 , the linearizations (4.9) extend (non-canonically) to a family of operators parameterized by points $q = (f, J)$ in the slice as follows. Fix a Riemannian metric g_0 on X and let ∇^0 denote its Levi-Civita connection. Using the notation of (4.9), define D_q by

$$D_q(\zeta, k) = \frac{1}{2} (\nabla_{df}^0 \zeta + J \nabla_{dfj}^0 \zeta) + \frac{1}{4} (\nabla_\zeta^0 J)(dfj + Jdf) + \frac{1}{4} (Jdf + dfj)k. \quad (4.14)$$

This agrees with (4.9) for $(f, J) \in \mathcal{M}_{\text{simple}}$ because (4.9) is independent of the connection and $Jdf = dfj$. Again, (4.14) defines a family of bounded linear operators $D_q : \mathcal{E}_q \rightarrow \mathcal{F}_q$ between Sobolev completions.

Next fix a Riemannian metric g_S compatible with the complex structure on the local universal family of curves parameterized by \mathcal{S}_C . By restriction, g_S induces a Riemannian metric on each curve in the local family which, under the trivialization associated with the slice, gives rise to a family of metrics on Σ parameterized by the \mathcal{S}_C . These metrics, together with their associated volume forms and the fixed metric g_0 on X determine L^2 inner products $\langle \cdot, \cdot \rangle_{L^2}$ on $T_j\mathcal{S}_C$, \mathcal{E}_q and \mathcal{F}_q for each $q = (f, J) \in \mathcal{Slice}_0$.

Let D_q^* denote the formal L^2 adjoint of the operator D_q of (4.14), which is uniquely defined by

$$\langle D_q\xi, \eta \rangle_{L^2} = \langle \xi, D_q^*\eta \rangle_{L^2} \quad (4.15)$$

for all $\xi \in \Gamma(f^*TX) \oplus T_j\mathcal{S}_C$ and $\eta \in \Omega^{0,1}(f^*TX)$. The adjoint operators depend on the choice of metrics.

For an element $\xi = (\zeta, k)$ of \mathcal{E}_p at $p = (f, J)$ and an injective point x of f , we define $\xi^N(x)$ to be the component $\zeta^N(x)$ of $\zeta(x)$ normal to $f_*(T_x C)$ with respect to the metric g_0 on X .

We will repeatedly use the following simple consequence of elliptic theory.

Lemma 4.1. *Given $p = (f, J)$ in \mathcal{M}_{simple} , and nonzero $\kappa \in \ker D_p$ and $c \in \ker D_p^*$, there is an injective point $x \in C$ such that $c(x) \neq 0$ and $\kappa^N(x) \neq 0$.*

Proof. The fact that (4.4) is an isomorphism implies that κ cannot be everywhere tangent to C ; κ is also continuous because $\kappa \in \mathcal{E}_p$. Next, write $D_p(\zeta, k) = D_p^0\zeta + \frac{1}{2}Jf_*k$ where D_p^0 is the elliptic operator defined by the first three terms on the righthand side of (4.9). The equation $D_p^*c = 0$ means that the L^2 inner product $\langle D_p(\zeta, k), c \rangle_{L^2}$ is zero for all (ζ, k) and therefore, taking $k = 0$, $(D_p^0)^*c = 0$. Lemma 3.4.4 of [MS] then shows that c is in the Sobolev (k, p) space, so is continuous, and also shows that c cannot vanish identically on any open set in C . The lemma follows because the injective points are open and dense in C . \square

5. THE STRUCTURE OF THE MODULI SPACE

The moduli space \mathcal{M}_{simple} has a distinguished subset: the “wall” $\mathcal{W} \subset \mathcal{M}_{simple}$ defined as the stratified set

$$\mathcal{W} = \bigcup_{r \geq 1} \mathcal{W}^r \quad \mathcal{W}^r = \left\{ p = (f, J) \in \mathcal{M}_{simple} \mid \dim \ker D_p = r \right\}. \quad (5.1)$$

Observe that each embedded complex curve $\iota_C : C \hookrightarrow X$ determines a set $\mathcal{J}_C \subset \mathcal{J}$ consisting of all $J \in \mathcal{J}$ for which ι_C is J -holomorphic. The restriction of (4.1) over \mathcal{J}_C has a canonical section $J \mapsto (\iota_C, J)$ whose image is $\mathcal{M}_C = \{\iota_C\} \times \mathcal{J}_C$. The results of this section show that \mathcal{W}^1 is a codimension 1 submanifold with a distinguished submanifold $\mathcal{A} \subset \mathcal{W}^1$, that the other strata \mathcal{W}^r have higher codimension, and that the same is true for the subsets $\mathcal{W}^r \cap \mathcal{M}_C$ and $\mathcal{A} \cap \mathcal{M}_C$ of \mathcal{M}_C .

Proposition 5.1. *The universal moduli space in (4.1) has the following structure:*

- (a) \mathcal{M}_{simple} is a separable Banach manifold and its tangent space at $p = (f, J)$ is $\ker \mathcal{L}_p$.
- (b) \mathcal{W} is the set of critical points of the projection (4.1), and π is a local diffeomorphism on $\mathcal{M}_{simple} \setminus \mathcal{W}$.
- (c) \mathcal{W}^1 is a codimension 1 submanifold of \mathcal{M}_{simple} .
- (d) For each embedded curve C , $\mathcal{M}_C = \{\iota_C\} \times \mathcal{J}_C$ is transverse to \mathcal{W}^1 .

Proof. Parts (a) and (b) are well-known; we present their proofs for later use. By the Implicit Function Theorem, \mathcal{M}_{simple} is a manifold at those points p where \mathcal{L}_p is onto. Surjectivity fails at $p = (f, J)$ only if there is a non-zero $c \in \mathcal{F}_p$ that is L^2 orthogonal to $\mathcal{L}_p(\zeta, k, K)$ for all (ζ, k, K) . By (4.8) and (4.15), this implies that $D_p^*c = 0$ and

$$0 = \int_C \langle c, Kf_*j \rangle \quad (5.2)$$

for every variation K in J . By Lemma 4.1, there is an injective point $x \in C$ where $c(x) \neq 0$. One can then find a variation K_0 in J so that $K_0 f_* j = c$ at the point x (cf. [MS, Lemma 3.2.2]). Choose local coordinates $y = (y_1, y_2, \dots)$ on X centered at $f(x)$. Fix a non-negative bump function $\beta(y)$ supported in this coordinate chart and for each $\varepsilon > 0$ set $\beta_\varepsilon(y) = c(\varepsilon)\beta(y/\varepsilon)$, where $c(\varepsilon)$ is the constant determined by the normalization condition $\int_C f^* \beta_\varepsilon = 1$. Then for each continuous function φ on C we have

$$\lim_{\varepsilon \rightarrow 0} \int_C f^* \beta_\varepsilon \cdot \varphi = \varphi(x). \quad (5.3)$$

Substituting $K = \beta_\varepsilon K_0$ in (5.2) and taking the limit as $\varepsilon \rightarrow 0$ gives a contradiction.

(b) The projection $d\pi_p(\xi, K) = K$ of a non-zero element in $\ker \mathcal{L}_p$ is zero if and only if $\xi = (\zeta, k)$ is a non-zero element of $\ker D_p$. On the other hand, D_p has index 0, so is onto at each $p \notin \mathcal{W}$. At such p , for each K we can use (4.8) to obtain ξ with $\mathcal{L}_p(\xi, K) = 0$; then (ξ, K) is tangent to \mathcal{M} and $d\pi_p(\xi, K) = K$. Thus $d\pi$ is an isomorphism at each point not in \mathcal{W} , and \mathcal{W} is the collection of critical points of π .

(c) Let $\text{Fred} \rightarrow \mathcal{M}_{\text{simple}}$ be the bundle whose fiber over (f, J) is the space of index 0 Fredholm operators of type (4.9). Fred is the union of strata $\text{Fred}^r = \{D \in \text{Fred} \mid \dim \ker D = r\}$, where each Fred^r is a submanifold of codimension r^2 whose normal bundle at D is naturally identified with $\text{Hom}(\ker D, \text{coker } D)$ (cf. [K], §1.1b, c). It suffices to show that the section

$$\Psi(f, J) = D_{f, J}. \quad (5.4)$$

is transverse to Fred^1 .

For this purpose, we fix $p = (f, J) \in \mathcal{Slice}_0 \cap \mathcal{W}^1$ and consider a deformation $p_t = (f, J_t)$ in $\mathcal{M}_{\text{simple}}$ where both the map and the complex structure on the domain is fixed, while J_t is a path in \mathcal{J} whose restriction to $f(C)$ is fixed and which changes only in a small neighborhood U of the image $f(x)$ of an injective point x . Along the path p_t , \mathcal{E}_{p_t} is fixed, since it is independent of J , as is \mathcal{F}_{p_t} , which depends on J only through its restriction along $f(C)$. Thus we have a 1-parameter family of Fredholm maps $D_{p_t} : \mathcal{E}_p \rightarrow \mathcal{F}_p$ with fixed domain and target. The initial derivative \dot{p}_0 has the form $v = (0, K) \in \ker \mathcal{L}_p = T_p \mathcal{M}_{\text{simple}}$, where K vanishes along $f(C)$ and is supported in U . Hence the variation $(\delta_v D)_p = \frac{d}{dt}|_{t=0} D_{p_t}$ in the direction v is obtained by replacing J by J_t in (4.9) and taking the t -derivative at $t = 0$. Because the formula (4.9) is independent of the connection, we can take the variation with the connection fixed; the formula then shows that $(\delta_v D)_p$ is $\frac{1}{2}(\nabla K)f_*j$. Moreover, when applied to an element $\xi = (\zeta, k)$ of \mathcal{E}_p , this variation depends only on the normal component $\xi^N = (\zeta^N, 0)$ of ζ because $K \equiv 0$ on $f(C)$:

$$(\delta_v D)_p \xi = \frac{1}{2}(\nabla_{\xi^N} K)f_*j. \quad (5.5)$$

This formula is tensorial in ξ , does not depend on the connection, and is tensorial in $K \in T_J \mathcal{J}$ as long as $K \equiv 0$ along $f(C)$.

Now fix generators κ of $\ker D_p$ and c of $\ker D_p^*$. Since the normal space to Fred^1 at p is 1 dimensional, it is enough to construct a variation in p of the form $(0, K)$ such that the L^2 inner product $\langle c, (\delta_v D)_p \kappa \rangle_{L^2}$ is non-zero.

To find such a v , choose an injective point $x \in C$ of f with both $\kappa^N(x) \neq 0$ and $c(x) \neq 0$ as in Lemma 4.1. As in the proof of part (a), there is a $K_0 \in T_J \mathcal{J}$ supported near $f(x)$ such that $K_0 f_*j = c$ at x . Choose a local coordinate system $\{z, y_1, y_2, \dots\}$ centered at $f(x)$ with z a local complex coordinate on $f(C)$, and $\{y_i\}$ real coordinates with $\frac{\partial}{\partial y_1}|_{f(x)} = \kappa^N(x)$. Then $K = y_1 K_0$ vanishes along $f(C)$ and satisfies

$$(\nabla_{\kappa^N} K)f_*j = c \quad \text{at the single point } x. \quad (5.6)$$

Finally, set $v_\varepsilon = (0, 2\beta_\varepsilon K)$ with β_ε as in (5.3). Replacing K by $2\beta_\varepsilon K$ in (5.5) and using (5.6), one sees that

$$\lim_{\varepsilon \rightarrow 0} \int_C \langle c, (\delta_{v_\varepsilon} D)_p \kappa \rangle = |c(x)|^2 \neq 0. \quad (5.7)$$

In particular, there is a variation with $\langle c, (\delta_v D)_p \kappa \rangle_{L^2} \neq 0$, which proves statement (c).

Statement (d) holds because the transversality above was obtained using a variation $v_\varepsilon = (0, 2\beta_\varepsilon K)$ tangent to \mathcal{M}_C . \square

The proof of Proposition 5.1c extends to show that the part of \mathcal{W} not in \mathcal{W}^1 has codimension 3 in the following sense.

Definition 5.2. We say that a subset S of a manifold M has codimension k if it is contained in a countable union $\bigcup \rho_\ell(S^\ell)$ where each $\rho_\ell : S^\ell \rightarrow M$ is a Fredholm map of separable Banach manifolds with index $\rho_\ell \leq -k$.

Lemma 5.3. $\mathcal{W} \setminus \mathcal{W}^1$ has codimension 3 in \mathcal{M}_{emb} , and $(\mathcal{W} \setminus \mathcal{W}^1) \cap \mathcal{M}_C$ has codimension 3 in \mathcal{M}_C .

C. Taubes obtained a similar result in dimension 4 using analytic perturbation theory (Step 5 of the proof of Lemma 5.1 in [T]). A proof in the spirit of the above arguments is given in the appendix for Proposition A.4; Lemma 5.3 is a special case.

5.1. The structure of \mathcal{W}^1 . We next examine the portion of the top stratum \mathcal{W}^1 of the wall (5.1) that lies in the open subset \mathcal{M}_{emb} of \mathcal{M}_{simple} consisting of embedded maps. Our goal is to show that the projection

$$\pi_{\mathcal{W}} : \mathcal{W}^1 \rightarrow \mathcal{J} \quad (5.8)$$

obtained by restricting (4.1) to \mathcal{W}^1 is an immersion off a set of codimension 1. For this purpose, we first introduce locally defined functions ψ that vanish transversally along \mathcal{W}^1 .

Fix a point $p_0 \in \mathcal{W}^1$ and restrict to a slice $Slice_0$ containing p_0 . At every $p \in \mathcal{W}^1$, the operators D_p and D_p^* defined by (4.14) and (4.15) have 1-dimensional kernels. These kernels determine subbundles

$$\mathcal{E}^0 \subset \mathcal{E}, \quad \mathcal{F}^0 \subset \mathcal{F} \quad (5.9)$$

along \mathcal{W}^1 , and the projection $\pi_{cok} : \mathcal{F}^0 \rightarrow \mathcal{F}/\text{im } D$ onto the cokernel bundle along \mathcal{W}^1 is an isomorphism. By choosing sections of \mathcal{E}^0 and \mathcal{F}^0 along the submanifold \mathcal{W}^1 and extending, we can find smooth non-vanishing local sections κ of \mathcal{E} and c of \mathcal{F} , defined in a neighborhood U of p_0 in the slice, such that the restrictions to \mathcal{W}^1 are local sections of \mathcal{E}^0 and \mathcal{F}^0 respectively. Let $\psi : U \rightarrow \mathbb{R}$ be the function defined by

$$\psi(q) = \int_C \langle c_q, D_q \kappa_q \rangle, \quad (5.10)$$

using the same metrics and volume forms as in (4.15).

Clearly, ψ vanishes along $\mathcal{W}^1 \cap U$, where $D_q \kappa_q = 0$. Differentiating $\psi(q_t)$ for any path q_t in U with $q_0 = p \in \mathcal{W}^1 \cap U$ and initial velocity $q'_0 = u$ yields several terms, including the variation in the inner product and volume form. All except one vanish at $q_0 = p$ because $D_p \kappa_p = 0$ and $D_p^* c_p = 0$, showing that

$$(d\psi)_p(u) = \int_C \langle c_p, (\delta_u D)_p \kappa_p \rangle.$$

The proof of Proposition 5.1c produces variations showing that $(d\psi)_p \neq 0$ for all $p \in \mathcal{W}^1 \cap U$. Thus the restriction of ψ to $\mathcal{M}_{simple} \cap U$ vanishes transversally along $\mathcal{W}^1 \cap U$. In particular, we have

$$T_p \mathcal{W}^1 = (\ker d\psi)_p \quad \forall p \in \mathcal{W}^1 \cap U. \quad (5.11)$$

Lemma 5.4. Inside \mathcal{M}_{emb} , the subset \mathcal{A} of \mathcal{W}^1 where the projection (5.8) fails to be a immersion is a codimension 1 submanifold of \mathcal{W}^1 , and \mathcal{M}_C is transverse to \mathcal{A} .

Proof. Fix a point in \mathcal{A} and a slice containing it, and work locally in a neighborhood U on which ψ is defined by (5.10). Consider the vector field $v = (\kappa, 0)$ on U , where κ is the non-vanishing local section chosen above (5.10). The proof of Proposition 5.1b shows that at every $p \in \mathcal{W}^1$, $\ker(d\pi)_p$ is spanned by v_p , so $\pi_{\mathcal{W}}$ fails to be an immersion at p if and only if $v_p \in T_p \mathcal{W}^1$. Together with (5.11), this gives two local descriptions of \mathcal{A} :

$$\mathcal{A} = \{p \in \mathcal{W}^1 \mid (\kappa, 0)_p \in T_p \mathcal{W}^1\} = \{p \in \mathcal{W}^1 \mid (d\psi(v))_p = 0\}. \quad (5.12)$$

By the second description, it suffices to show that the restriction of the function $d\psi(v) : U \rightarrow \mathbb{R}$ to \mathcal{W}^1 vanishes transversally at each $p \in \mathcal{A}$.

Consider variations in $p = (f, J)$ of the form $w = (0, K)$, where K is an element of $T_J \mathcal{J}$ such that both $K = 0$ and $\nabla K = 0$ along the image $f(C)$. We will show that for every such variation

$$w = (0, K) \in T_p \mathcal{W}^1 \quad (5.13)$$

and

$$\nabla_w (d\psi(v))|_p = \frac{1}{2} \int_C \langle c, (\nabla_{\kappa^N} \nabla_{\kappa^N} K) df j \rangle. \quad (5.14)$$

The argument used to obtain (5.6) and (5.7) then produces a $K = \frac{1}{2}y_1^2\beta_\varepsilon K_0$ with $K = 0$ and $\nabla K = 0$ along $f(C)$ that makes the integral (5.14) non-zero, which shows transversality at p . Furthermore, if $p \in \mathcal{M}_C$, these variations $w = (0, K)$ are tangent to \mathcal{M}_C . Thus both parts of the lemma follow from (5.13) and (5.14).

We will prove (5.13) and (5.14) by constructing a 2-parameter family of deformations of p in U that is tangent to v and w at p . For clarity, write p as $p_0 = (f_0, J_0)$. Start by choosing a path J_t through J_0 with initial velocity $K \in T_{J_0}\mathcal{J}$ such that the restrictions of J_t and $\nabla^0 J_t$ to $f_0(C)$ are independent of t . Along the path $q_t = (f_0, J_t)$ in U , \mathcal{E}_{q_t} , \mathcal{F}_{q_t} and the operators D_{q_t} defined by (4.14) are constant, because all three depend only on the 1-jet of J_t along $f_0(C)$, which is fixed. Thus the path (f_0, J_t) is in \mathcal{W}^1 for all t , so its initial tangent w satisfies (5.13). We can assume that the local sections κ of \mathcal{E} and c of \mathcal{F} used to define ψ were chosen so that $\kappa = \kappa_0$ and $c = c_0$ along the path (f_0, J_t) .

Next, choose a smooth family of maps $f_s : (C, j_s) \rightarrow X$ with initial tangent vector κ_0 . Then

$$p_{s,t} = (f_s, J_t) \quad (5.15)$$

is a 2-parameter family in Slice_0 with $p_*\partial_s = (\kappa_0, 0) = v$ along $p_{0,t}$ and $p_*\partial_t = (0, K) = w$ at $s = t = 0$. Write the restrictions of the sections κ , c and the operators (4.14) as $\kappa_{s,t}$, $c_{s,t}$ and $D_{s,t}$ respectively, and set

$$\eta_{s,t} = D_{s,t}\kappa_{s,t}. \quad (5.16)$$

For $s = 0$, $\eta_{0,t} = D_{f_0, J_t}\kappa_0$ is 0 by construction; we also have $c_{0,t} = c_0$, $\kappa_{0,t} = \kappa_0$, and therefore

$$\eta_{0,0} = 0, \quad (\partial_t\eta)_{0,0} = 0, \quad (\partial_t c)_{0,0} = 0, \quad (\partial_t\kappa)_{0,0} = 0. \quad (5.17)$$

With this notation, the restriction of (5.10) to the family (5.15) is the function

$$\psi(s, t) = \psi(p_{s,t}) = \int_C \langle c_{s,t}, D_{s,t}\kappa_{s,t} \rangle_s = \int_C \langle c_{s,t}, \eta_{s,t} \rangle_s, \quad (5.18)$$

where the pointwise inner product and the area form depend on s but not on t . Along the path $p_{0,t}$, we have $d\psi(v) = d\psi(p_*\partial_s) = \partial_s\psi$. Differentiating in the w direction and noting that $w = p_*\partial_t$ at the origin then gives

$$\nabla_w(d\psi(v))_{p_0} = (\partial_t\partial_s\psi)_{0,0} = (\partial_s\partial_t\psi)_{0,0}.$$

To complete the proof, we will calculate $(\partial_s\partial_t\psi)_{0,0}$ by differentiating (5.18). For fixed s , $\kappa_{s,t}$ is a path of sections of the fixed bundle f_s^*TX , while both $c_{s,t}$ and $\eta_{s,t}$ are paths of 1-forms on C with values in the fixed bundle f_s^*TX (they are $(0,1)$ forms with respect to the pair (j_s, J_t)). Hence we have

$$(\partial_t\psi)_{s,0} = \int_C \langle (\partial_t c)_{s,0}, \eta_{s,0} \rangle_s + \int_C \langle c_{s,0}, (\partial_t\eta)_{s,0} \rangle_s. \quad (5.19)$$

Now differentiate (5.19) with respect to s and evaluate at $s = 0$ using (5.17). The contribution of the first integral vanishes because $(\partial_t c)_{0,0} = 0$ and $\eta_{0,0} = 0$, leaving

$$(\partial_s\partial_t\psi)_{0,0} = \partial_s|_{s=0} \int_C \langle c_{s,0}, (\partial_t\eta)_{s,0} \rangle_s. \quad (5.20)$$

By (5.16), $\eta_{s,t}$ is given by the operator (4.14) with $q = (f_s, J_t)$, applied to $\kappa_{s,t} = (\zeta_{s,t}, k_{s,t})$. In the resulting formula, fix s and differentiate with respect to t . Because f_s and j_s are independent of t and $(\partial_t J_t)_0 = K$, one sees that $(\partial_t\eta)_{s,0}$ has the general form

$$(\partial_t\eta)_{s,0} = D_{s,0}((\partial_t\kappa)_{s,0}) + \frac{1}{4}(\nabla_{\zeta_s}^0 K)(df_s j_s + J_0 df_s) + T_{s,t}(K), \quad (5.21)$$

where $T(K)$ is a sum of terms, each linear and tensorial in K . The contribution of the first term of (5.21) to (5.20) is

$$\partial_s|_{s=0} \int_C \langle c_{s,0}, D_{s,0}((\partial_t\kappa)_{s,0}) \rangle_s = \partial_s|_{s=0} \int_C \langle D_{s,0}^* c_{s,0}, (\partial_t\kappa)_{s,0} \rangle_s,$$

as in (4.15). Taking the s -derivative at $s = 0$ yields three terms, all of which vanish because $D_{0,0}^* c_{0,0} = 0$ and $(\partial_t\kappa)_{0,0} = 0$. Similarly inserting the remaining terms of (5.21) into (5.20) and differentiating yields

many terms; all but one vanish because K and $\nabla^0 K$ vanish along $f_0(C)$. After noting that f_0 is J_0 -holomorphic, one is left with

$$(\partial_s \partial_t \psi)_{0,0} = \frac{1}{2} \int_C \langle c_0, \nabla_{\zeta_0}^0 \nabla_{\zeta_0}^0 K d f_0 j_0 \rangle_0.$$

Because K and ∇K vanish along $f_0(C)$, this expression is independent of the connection, and its dependence on ζ_0 involves only the normal component $\zeta_0^N = \kappa_0^N$. Thus

$$\nabla_w(d\psi(v))|_{p_0} = (\partial_s \partial_t \psi)_{0,0} = \frac{1}{2} \int_C \langle c_0, \nabla_{\kappa_0^N} \nabla_{\kappa_0^N} K d f_0 j_0 \rangle_0.$$

This verifies (5.14) and completes the proof. \square

6. LOCAL MODELS FOR WALL CROSSINGS

We next study the local geometry of the moduli space around a point p on the wall $\mathcal{W}^1 \setminus \mathcal{A}$. We assume that p corresponds to a J -holomorphic embedding $f : C \rightarrow X$ of a smooth curve C , which we can regard as the inclusion $\iota_C : C \hookrightarrow X$ of its image. The goal is to show that the restriction $\pi_\gamma : \mathcal{M}_{emb}^\gamma \rightarrow \gamma$ of (4.1) over a generic path $\gamma \subset \mathcal{J}$ is a Morse function at p , and hence is locally described by a quadratic equation. Two types of smooth paths through J are relevant for our purposes:

TYPE A. γ_A is a path in \mathcal{J} such that the projection π is transverse to γ_A at p .

TYPE B. γ_B is a path in \mathcal{J}_C whose lift $\tilde{\gamma}_B = \{\iota_C\} \times \gamma_B$ to \mathcal{M}_C is transverse to \mathcal{W}^1 at p .

The lemmas at the end of this section show that both types of paths are generic. But first, we will use Kuranishi's method to construct a local model for the moduli space over these paths. It is convenient to study both types simultaneously by considering embedded parameterized disks

$$S = \{(t, s)\} \subset \mathcal{J} \tag{6.1}$$

whose t -axis is a path γ_B of Type B and whose s -axis is a path γ_A of Type A. We then restrict (4.1) over S to obtain

$$\pi_S : \mathcal{M}^S \rightarrow S, \tag{6.2}$$

where $\mathcal{M}^S = \pi^{-1}(S)$ is the moduli space over S .

Definition 6.1. A 3-ball $B \subset \text{Slice}_0$ with coordinates (x, y, z) centered at $p = (f, J) \in \mathcal{W}^1 \setminus \mathcal{A}$ is adapted to S at p if

- (a) $\pi : B \rightarrow S$ is given by $\pi(x, y, z) = (y, z)$.
- (b) $\gamma_B(t) = (t, 0)$ is a Type B path whose lift is $\tilde{\gamma}_B(t) = (0, t, 0)$.
- (c) $\gamma_A(s) = (0, s)$ is a Type A path.
- (d) In terms of the splitting (4.7), $T_p B$ is spanned by $\partial_x|_p = (\kappa, 0)$, $\partial_y|_p = (0, K_B)$, and $\partial_z|_p = (0, K_A)$, where κ generates $\ker D_p \cong \mathbb{R}$, $K_A = \dot{\gamma}_A(0)$ and $K_B = \dot{\gamma}_B(0)$.

With these assumptions, the transversality conditions in Type A and B are equivalent to

$$(a) \quad K_A \notin \text{im } d\pi_p \quad \text{and} \quad (b) \quad (0, K_B) \notin T_p \mathcal{W}^1 \tag{6.3}$$

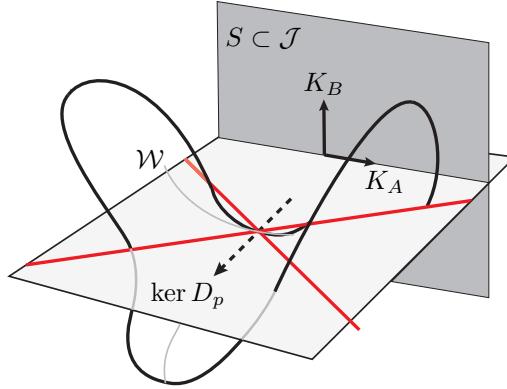
respectively. The requirement that $\gamma_B \subset \mathcal{J}_C$ also implies that K_B vanishes along $f(C)$, so $(0, K_B) \in \ker \mathcal{L}_p = T_p \mathcal{M}_{emb}$; because $d\pi_p(0, K) = K$ this also means that $K_B \in \text{im } d\pi_p$ and hence K_A and K_B are linearly independent.

Theorem 6.2 (Kuranishi model). *For each $p = (f, J)$ in $\mathcal{M}_{emb} \cap (\mathcal{W}^1 \setminus \mathcal{A})$ and each $S \subset \mathcal{J}$ as in (6.1) centered at J , there is a 3-ball B adapted to S at p such that \mathcal{M}^S is locally the 2-manifold*

$$V = \{(x, y, z) \in B \mid z = x(ax + by + r(x, y))\} \quad (6.4)$$

with $a, b \neq 0$, where $r(x, y) = O(x^2 + y^2)$ near the origin.

Theorem 6.2 shows that \mathcal{M}^S is locally a saddle surface in $B \subset \mathbb{R}^3$ given as the graph of a function $z = z(x, y)$ that has a non-degenerate critical point at the origin. Before giving the proof, we record the two cases that will be used in later sections.



Corollary 6.3. *Suppose that $p = (f, J) \in \mathcal{M}_{emb} \cap (\mathcal{W}^1 \setminus \mathcal{A})$ and γ is a path in \mathcal{J} through $\gamma(0) = J$.*

(a) *If γ is of Type A, then \mathcal{M}^γ is locally modeled at p by*

$$\{(x, t) \mid t = ax^2\} \quad \text{with } a \neq 0 \text{ and } \pi_\gamma(x, t) = t. \quad (6.5)$$

(b) *If γ is of Type B, then there is a disk $S \subset \mathcal{J}$ centered at p locally containing γ such that \mathcal{M}^S is locally*

$$\{(x, t, s) \in B \mid s = x(ax + bt + r(x, t))\} \quad \text{with } a, b \neq 0 \text{ and } \pi_S(x, t, s) = (t, s), \quad (6.6)$$

and the restriction to \mathcal{M}^γ is the restriction to the plane $s = 0$, namely

$$\{(x, t) \mid 0 = x(ax + bt + r(x, t))\}. \quad (6.7)$$

Proof. In the first case, take $\gamma_A = \gamma$, choose $\gamma_B \subset \mathcal{J}_C$ a path whose tangent vector K_B satisfies (6.3b), and choose a local embedded disk $S \subset \mathcal{J}$ containing both γ_A and γ_B locally near p . Then apply Theorem 6.2 and restrict to the plane $\{y = 0\}$ to get the local model for \mathcal{M}^γ of the form

$$\{(x, t) \mid t = x(ax + r(x))\} \quad \text{with } a \neq 0, \pi_\gamma(x, t) = t,$$

and with $r(x) = O(x^2)$ for small x . This becomes (6.5) after reparameterizing x .

Similarly, in the second case, take $\gamma_B = \gamma$, fix a direction $K_A \in T_J \mathcal{J}$ satisfying (6.3a), choose S with $T_J S = \text{span}(K_A, K_B)$ containing γ , and again apply Theorem 6.2. \square

The distinction between the local models (6.5) and (6.7) is crucial. By Lemma 6.5 below, Model (6.5) applies where a generic path in \mathcal{J} crosses a wall. This is precisely the local model for the creation (if $a > 0$) or annihilation (if $a < 0$) of a pair of curves in the moduli space. Similarly, Lemma 6.7 shows that Model (6.7) applies where a generic path in the subspace $\mathcal{J}_C \subset \mathcal{J}$ crosses a wall. But the model (6.7) is *not a manifold*: it is the union of two curves crossing transversely at the origin. It can be smoothed using the parameter s in (6.6), as will be done in Lemma 7.4.

Proof of Theorem 6.2. Specializing (4.5) and (4.7), one sees that the slice over S is $\text{Map}(\Sigma, X) \times \mathcal{S}_C \times S$ and that we can identify the first two factors with their tangent space at p , which is \mathcal{E}_p . We can also trivialize the bundle \mathcal{F} over a neighborhood of p . With these identifications, \mathcal{M}^S is the subset of $\mathcal{E}_p \times S$ that is the zero set of the \mathcal{F}_p -valued function F defined by (4.6). This has an expansion

$$F(\xi) = \mathcal{L}_p(\xi) + Q(\xi), \quad (6.8)$$

where $Q(\xi)$ vanishes to first order at $\xi = 0$. Next, fix a generator κ of $\ker D_p = \mathcal{E}^0$ and choose a decompositon $\mathcal{E}_p = \mathcal{E}^0 \oplus \mathcal{E}^+$. Then ξ can be written as $\xi = (x, y, z, \alpha)$, for coordinates (x, y, z) in a 3-ball $\mathcal{B} \subset \mathcal{E}^0 \times S$ around 0 and $\alpha \in \mathcal{E}^+$. Using (4.8), the linear term is

$$\begin{aligned} \mathcal{L}_p(\xi) &= xD_p\kappa + \frac{1}{2}(yK_B + zK_A)f_*j + D_p\alpha \\ &= zc_p + D_p\alpha, \end{aligned}$$

where we have set $c_p = \frac{1}{2}K_A f_* j$ and noted that $K_B = \dot{\gamma}_B(0)$ is tangent to a Type B path, so $K_B|_{f(C)} = 0$. Furthermore, the transversality assumption (6.3a) ensures that $c_p \notin \text{im } D_p$ as follows: if $c_p = D_p\mu$, then $(-\mu, K_A)$ is an element of $\ker \mathcal{L}_p = T_p \mathcal{M}_{\text{simple}}$ with $d\pi(-\mu, K_A) = K_A$, contradicting (6.3a). Thus \mathcal{F}_p decomposes as

$$\mathcal{F}_p = \widehat{\mathcal{F}}^0 \oplus \widehat{\mathcal{F}}^+$$

where $\widehat{\mathcal{F}}^0$ is the real span of c_p and $\widehat{\mathcal{F}}^+ = D_p \mathcal{E}^+$. Using this decomposition, write Q as $(Q_0 \cdot c_p, Q_1)$.

By Proposition A.4.1 of [MS] there is a bounded linear map $T : \widehat{\mathcal{F}}^+ \rightarrow \mathcal{E}^+$ that is a pseudo-inverse of D_p , which implies that $D_p T$ is the identity on $\widehat{\mathcal{F}}^+$. Define a map

$$\eta : \mathcal{E}^0 \times S \times \mathcal{E}^+ \rightarrow \mathcal{E}^0 \times S \times \mathcal{E}^+$$

by

$$\eta(x, y, z, \alpha) = (x, y, z, \eta_{x,y,z}(\alpha)), \quad \text{where} \quad \eta_{x,y,z}(\alpha) = \alpha + TQ_1(x, y, z, \alpha) \in \mathcal{E}^+.$$

By the Inverse Function Theorem, η is a local diffeomorphism near 0.

Using the above notation, (6.8) can be rewritten as the equation

$$F(x, y, z, \alpha) = ((z + Q_0)c_p, D_p(\eta_{x,y,z}(\alpha))) \in \widehat{\mathcal{F}}^0 \oplus \widehat{\mathcal{F}}^+.$$

This shows that $F(x, y, z, \alpha) = 0$ if and only if both $\eta_{x,y,z}(\alpha) = TD_p(\eta_{x,y,z}(\alpha)) = 0$, and $z = q$ where

$$q(x, y, z) = -Q_0 \circ \eta^{-1}(x, y, z, 0).$$

Thus there is a local diffeomorphism

$$\mathcal{M}^S \cong \{(x, y, z) \in B \mid z = q(x, y, z)\}.$$

The real-valued function q is smooth and vanishes to first order at the origin, so we can solve for z as a function of x and y to obtain

$$z = z(x, y) = ax^2 + bxy + cy^2 + r_2(x, y), \quad (6.9)$$

where the remainder r_2 vanishes to second order. In particular, locally near p , \mathcal{M}^S is a 2-manifold with coordinates (x, y) , in which the projection (6.2) is

$$\pi_S(x, y) = (y, z(x, y)),$$

so $d\pi_S = (dy, z_x dx + z_y dy)$.

Next observe that, since γ_B is a Type B path, (6.3b) implies that \mathcal{W}^1 is transverse to \mathcal{M}^S at p . Thus $\mathcal{W}^1 \cap \mathcal{M}^S$ is a 1-dimensional manifold near p . On the other hand, $\pi_S : \mathcal{M}^S \rightarrow S$ is a map between 2-manifolds, and the above formula shows that $\text{rank } d\pi_S \geq 1$. The proof of Proposition 5.1b then shows that the set of critical points of π_S is contained in $\mathcal{W}^1 \cap \mathcal{M}^S$.

In coordinates, the intersection $\mathcal{W}^1 \cap \mathcal{M}^S$ is locally the zero locus of the function

$$w(x, y) = z_x(x, y),$$

while $T_p(\mathcal{W}^1 \cap \mathcal{M}^S) = T_p \mathcal{W}^1 \cap T_p \mathcal{M}^S$ is the kernel of the 1-form on S

$$dw = z_{xx}dx + z_{yx}dy = 2adx + bdy$$

at the origin. Consequently, since $(\kappa, 0) = \partial_x$ and $(0, \dot{\gamma}_B(0)) = \partial_y$, we have

- $2a = dw(\kappa, 0) \neq 0$ by (5.12) because $p \notin \mathcal{A}$ so $(\kappa, 0) \notin T_p \mathcal{W}^1$.
- $b = dw(0, K_B) \neq 0$ by (6.3b) because $\tilde{\gamma}_B$ is transverse to \mathcal{W}^1 .
- $c = 0$ because, for small t , the path $\tilde{\gamma}_B(t) = (0, t, 0)$ lies in \mathcal{M}^S , so (6.9) becomes $0 = ct^2 + \mathcal{O}(t^3)$ for all small t .

In fact, the expansion in the third bullet point shows that $r_2(0, t) = 0$ for all small t , which implies that (6.9) has the form

$$z(x, y) = x(ax + by + r(x, y)),$$

where r vanishes to first order. \square

To apply Theorem 6.2 and its corollary, we will need several statements about generic paths in \mathcal{J} . These are consequences of the Sard-Smale Theorem ([S], [MS, Theorem A.5.1]) applied in the following manner. Suppose that $\pi : \mathcal{M} \rightarrow \mathcal{J}$ and $\rho : \mathcal{N} \rightarrow \mathcal{M}$ are Fredholm maps of separable C^ℓ Banach manifolds of index ι_π and ι_ρ respectively, and with $\ell \geq 2 + \max(0, \iota_\pi, \iota_\pi + \iota_\rho)$. Fix points $J_0, J_1 \in \mathcal{J}$, and let $\mathcal{P} = \mathcal{P}\mathcal{J}$ be the space of C^k , $k \gg 1$, paths $[0, 1] \rightarrow \mathcal{J}$ from J_0 to J_1 , which is a separable Banach manifold.

In this context, we have the following lemma.

Lemma 6.4. *Assume J_0, J_1 are regular values of π and $J_0, J_1 \notin (\pi \circ \rho)(\mathcal{N})$. Then there exists a Baire set \mathcal{P}^* of \mathcal{P} so that for each $\gamma \in \mathcal{P}^*$, $\mathcal{M}^\gamma = \pi^{-1}(\gamma)$ is an $\iota_\pi + 1$ dimensional submanifold of \mathcal{M} transverse to ρ . In particular, when $\iota_\rho + \iota_\pi \leq -2$, \mathcal{M}^γ is disjoint from $\rho(\mathcal{N})$. Consequently, for each subset $S \subset \mathcal{M}$ of codimension ≥ 2 , there is a Baire set \mathcal{P}_S^* so that for each $\gamma \in \mathcal{P}_S^*$, \mathcal{M}^γ is a manifold disjoint from S .*

Proof. The evaluation map $\text{ev} : \mathcal{P} \times I \rightarrow \mathcal{J}$ is a submersion of separable Banach manifolds, and the restriction of ev to the boundary $\mathcal{P} \times \partial I$ is also a submersion. It follows that the map

$$\varphi = \text{ev} \times \pi : \mathcal{P} \times I \times \mathcal{M} \rightarrow \mathcal{J} \times \mathcal{J}$$

is transverse to the diagonal $\Delta_{\mathcal{J}}$, so the fiber product

$$\widetilde{\mathcal{M}} = \varphi^{-1}(\Delta_{\mathcal{J}}) = \{(\gamma, t, f) \mid f \in \mathcal{M}^{\gamma(t)}\}$$

is a separable Banach manifold whose boundary is the fiber product of $\text{ev}|_{\mathcal{P} \times \partial I}$ and π . Furthermore, the projection $\tilde{\pi} : \widetilde{\mathcal{M}} \rightarrow \mathcal{P}$ is a Fredholm map of index $\iota_\pi + 1$, while the projection $p_3 : \widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ onto the third factor is a submersion:

$$\begin{array}{ccccccc} \widetilde{\mathcal{N}} & \xrightarrow{\tilde{\rho}} & \widetilde{\mathcal{M}} & \longrightarrow & \mathcal{P} \times I & \longrightarrow & \mathcal{P} \\ \downarrow & & \downarrow p_3 & & \downarrow \text{ev} & & \\ \mathcal{N} & \xrightarrow{\rho} & \mathcal{M} & \xrightarrow{\pi} & \mathcal{J}. & & \end{array} \quad (6.10)$$

In particular, the fiber product $\widetilde{\mathcal{N}}$ of ρ and p_3 is a manifold and the index of $\tilde{\rho}$ is the index of ρ .

By Sard-Smale Theorem applied to $\tilde{\pi} : \widetilde{\mathcal{M}} \rightarrow \mathcal{P}$, there is a Baire subset \mathcal{P}_1 of \mathcal{P} such that each point $\gamma \in \mathcal{P}_1$ is a regular value of $\tilde{\pi}$. It is straightforward to check that $\gamma \in \mathcal{P}$ is a regular value of $\tilde{\pi}$ if and only if the path γ is transverse to $\pi : \mathcal{M} \rightarrow \mathcal{J}$. When γ is a submanifold of \mathcal{J} , this latter transversality means that $\mathcal{M}^\gamma = \pi^{-1}(\gamma)$ is a submanifold of \mathcal{M} .

Similarly, again by the Sard-Smale Theorem, there is a Baire subset \mathcal{P}_2 of \mathcal{P} such that each point $\gamma \in \mathcal{P}_2$ is a regular point of the composition $\widetilde{\mathcal{N}} \rightarrow \mathcal{P}$ in the top row of (6.10), and again this occurs if and only if γ is transverse to $\pi \circ \rho$. (Note that J_0, J_1 are regular values of $\pi \circ \rho$ because they are not in the image of $\pi \circ \rho$.) When γ is a submanifold, this last transversality implies that $\mathcal{N}^\gamma = (\pi \circ \rho)^{-1}(\gamma) = \rho^{-1}(\mathcal{M}^\gamma)$ is a submanifold.

Finally, the collection \mathcal{P}_3 of embedded paths is open and dense in \mathcal{P} . Thus $\mathcal{P}^* = \mathcal{P}_1 \cap \mathcal{P}_2 \cap \mathcal{P}_3$ is a Baire set. For each $\gamma \in \mathcal{P}^*$, \mathcal{M}^γ is an $\iota_\pi + 1$ dimensional manifold, \mathcal{N}^γ is an $\iota_\pi + 1 + \iota_\rho$ dimensional

manifold, and is empty if $\iota_\pi + 1 + \iota_\rho < 0$. The last statement of the lemma follows from Definition 5.2 and the fact that every countable intersection of Baire sets is a Baire set. \square

We will apply this reasoning twice: first for paths in \mathcal{J} , then for paths in \mathcal{J}_C . Recall the notations \mathcal{J}_E^* and \mathcal{J}_{isol}^E from Section 1. For simplicity, we omit the X from the notation $\mathcal{M}(X)$ of the moduli spaces in Lemmas 6.5 and 6.7.

Lemma 6.5. *Any path in \mathcal{J} with endpoints in \mathcal{J}_E^* can be deformed, keeping its endpoints, to a path γ such that $\mathcal{M}_{simple}^{\gamma,E}$ is a 1-dimensional manifold, consisting of embeddings, and intersecting the wall \mathcal{W} transversely in isolated points, all in $\mathcal{W}^1 \setminus \mathcal{A}$. Moreover, any path γ with these properties is in \mathcal{J}_{isol}^E .*

Proof. By Proposition 5.1a, \mathcal{M}_{simple} is a manifold, and by Corollary A.3 the set \mathcal{NE} of all non-embedded simple maps is a codimension 2 subset. Next, restrict to \mathcal{M}_{emb} , noting that \mathcal{W}^1 is a submanifold of \mathcal{M}_{emb} by Proposition 5.1c, and that $\mathcal{A} \subset \mathcal{M}_{emb}$ is a codimension 2 submanifold by Lemma 5.4. Together with Lemma 5.3, this shows that $\mathcal{A} \cup \mathcal{NE} \cup (\mathcal{W} \setminus \mathcal{W}^1)$ is a set of codimension 2 in \mathcal{M}_{simple} .

Now apply Lemma 6.4 with $\mathcal{M} = \mathcal{M}_{simple}$ and $\mathcal{N} = \mathcal{A} \cup \mathcal{NE} \cup (\mathcal{W} \setminus \mathcal{W}^1)$, noting that the index $\iota_\pi = 0$ in this case. This gives a Baire subset \mathcal{P}_1 of \mathcal{P} over which $\pi^{-1}(\gamma)$ is a manifold of dimension 1 that does not intersect \mathcal{N} . Again apply Lemma 6.4, now with $\mathcal{M} = \mathcal{M}_{emb}$ and $\mathcal{N} = \mathcal{W}^1$, which is a codimension 1 manifold. Now we get a Baire subset \mathcal{P}_2 of \mathcal{P} over which $\pi^{-1}(\gamma)$ is a manifold of dimension 1 that is transverse to \mathcal{N} .

Consequently, for each γ in the Baire set $\mathcal{P}_1 \cap \mathcal{P}_2$, $\mathcal{M}_{simple}^{\gamma,E}$ is a 1-manifold (with boundary) transverse to \mathcal{W}^1 , intersecting the wall \mathcal{W} only along $\mathcal{W}^1 \setminus \mathcal{A}$, and consisting only of embedded curves. Note that the only critical points of $\pi_\gamma : \mathcal{M}_{simple}^{\gamma,E} \rightarrow [0, 1]$ are these wall-crossing points. The proof is completed by observing that the local model (6.5) shows that the wall crossing points are non-degenerate critical points of π and therefore (a) they are isolated points of $\mathcal{M}_{simple}^{\gamma,E}$ and (b) all points of the fiber $\mathcal{M}_{simple}^{\gamma(t),E}$ are isolated for each t . \square

Corollary 6.6. *\mathcal{J}_{isol}^E is a dense and path-connected subspace of \mathcal{J} .*

Proof. Density was shown in Corollary 1.3. Path-connectedness follows from Lemma 6.5 and the fact that \mathcal{J} is path-connected. \square

We conclude this section by proving a version of Lemma 6.5 for paths in the subspace \mathcal{J}_C of \mathcal{J} .

Lemma 6.7. *Any path in \mathcal{J}_C with endpoints in \mathcal{J}_E^* can be deformed, keeping its endpoints, to a path γ in $\mathcal{J}_C \cap \mathcal{J}_{isol}^E$ whose lift $\tilde{\gamma}$ intersects \mathcal{W} transversally at finitely many points, all in $\mathcal{W}^1 \setminus \mathcal{A}$.*

Proof. First consider the subset $\mathcal{M}_C = \{\iota_C\} \times \mathcal{J}_C$. This is a submanifold of \mathcal{M}_{emb} that is transverse to \mathcal{W}^1 by Proposition 5.1d, and to \mathcal{A} by Lemma 5.4. Furthermore, Lemma 5.3 shows that $(\mathcal{W} \setminus \mathcal{W}^1) \cap \mathcal{M}_C$ is a codimension 2 subset of \mathcal{M}_C . By Lemma 6.4, there is a Baire subset of paths γ in \mathcal{J}_C for which the lift $\tilde{\gamma} = (\iota_C, \gamma)$ to \mathcal{M}_{simple} intersects the wall only along $\mathcal{W}^1 \setminus \mathcal{A}$, and this intersection is a finite set of transverse points. For each intersection point $\tilde{\gamma}(t) \in \mathcal{W}^1 \setminus \mathcal{A}$, the local model (6.7) implies that the core curve $\tilde{\gamma}(t)$ is an isolated point of $\mathcal{M}^{\gamma(t)}$ (and is clearly embedded). The same conclusion is true for those t with $\tilde{\gamma}(t) \notin \mathcal{W}$ by Proposition 5.1b.

It remains to find another Baire subset of paths γ for which the points of $\mathcal{M}^{\gamma(t)} \setminus \mathcal{M}_C$ are embedded and isolated for each t . Denote by $\mathcal{M}_{simple}^* \rightarrow \mathcal{J}_C$ the collection of simple J -holomorphic maps f that have at least one point x_i on each component of their domain with $f(x_i) \in X \setminus C$. The results of Proposition 5.1, Lemma 5.3, Lemma 5.4 and Corollary A.3 all extend to the moduli space $\mathcal{M}_{simple}^* \rightarrow \mathcal{J}_C$ by using variations supported around the points $f(x_i)$, but vanishing along C ; such variations are tangent to \mathcal{J}_C . As in the proof of [MS, Lemma 3.4.3] a further variation, with support off C , can be used to ensure that all curves in \mathcal{M}_{simple}^* are transverse to C . Again, such variations are tangent to \mathcal{J}_C .

With this understood, the proof of Lemma 6.5 extends to give a Baire subset \mathcal{P}^* of the space $\mathcal{P}\mathcal{J}_C$ of paths in \mathcal{J}_C so that for each $\gamma \in \mathcal{P}^*$, the points of $\mathcal{M}_{simple}^{\gamma(t),E}$ are embedded and isolated for all t . \square

7. THE CLUSTER ISOTOPY THEOREM

For notational simplicity, given two clusters $\mathcal{O} = (C, \varepsilon, J)$ and $\mathcal{O}' = (C', \varepsilon', J')$ whose core curves C and C' have the same genus and homology class, write

$$GW^E(\mathcal{O}) \approx GW^E(\mathcal{O}')$$

to mean that the difference is a finite sum of terms of the form $\pm GW^E(C_i, \varepsilon_i, J_i)$ of strictly higher level (1.9) compared to that of C . With this notation, for example, the conclusion of the cluster refinement Corollary 2.5 simply says that for generic $0 < \varepsilon' < \varepsilon$

$$GW^E(C, \varepsilon, J) \approx GW^E(C, \varepsilon', J). \quad (7.1)$$

We now use the results of Sections 4 and 5 and an isotopy argument to prove that the GW series of every cluster is equivalent, in the above sense, to the series (3.4) of an elementary cluster. We call a cluster (C, J, ε) *regular* if the embedding $C \hookrightarrow X$ is a regular J -holomorphic map.

Theorem 7.1 (Cluster Isotopy). *For a regular cluster $\mathcal{O} = (C, J_0, \varepsilon_0)$ centered at an embedded genus g J_0 -holomorphic curve C ,*

$$GW^E(\mathcal{O}) \approx \text{sign}(C, J_0) GW_g^{elem, E}(q^C, t) \quad (7.2)$$

where the right-hand series is the truncation of (3.6) with $q = q^C$ below energy E .

Proof. The proof of Proposition 3.2 shows that there exist $J_1 \in \mathcal{J}_C$ and ε_1 so that $\mathcal{O}_{elem} = (C, \varepsilon_1, J_1)$ is an elementary cluster. In fact, we can assume that $J_1 \in \mathcal{J}_C \cap \mathcal{J}_E^*$ after a perturbation supported outside the $\varepsilon_1/2$ -neighborhood of C of the type constructed in the proof of [MS, Lemma 3.4.3]. Choose a path $\gamma(t) = J_t$ in \mathcal{J}_C from J_0 to J_1 (the proof of Theorem A.2 of [IP2] shows that \mathcal{J}_C is connected). By Lemma 6.7 we can assume, after a deformation, that γ is a path in $\mathcal{J}_C \cap \mathcal{J}_{isol}^E$ and there is a finite set $Sing = \{t_i\}$, not containing 0 or 1, such that $\tilde{\gamma} = (f_C, \gamma(t))$

- lies in $\mathcal{M}^\gamma \setminus \mathcal{W}$ for all $t \notin Sing$, and
- lies in a 2-dimensional surface V_i given by (6.6) for $t \in [t_i - \delta, t_i + \delta]$.

Choose $\delta > 0$ small enough so that the intervals $[t_i - \delta, t_i + \delta]$ do not overlap, and let their endpoints be $0 < \tau_1 < \dots < \tau_{2k} < 1$. For each i , fix a cluster $\mathcal{O}_i = (C, \varepsilon_i, J_{\tau_i})$. Then $\tilde{\gamma}$ can be regarded as the composition of paths $\tilde{\gamma}_i : [\tau_i, \tau_{i+1}] \rightarrow \mathcal{M}^\gamma$ of two types:

- (i) Paths in $\mathcal{M}_{emb} \setminus \mathcal{W}$. For these, Lemma 7.2 below shows that $GW^E(\mathcal{O}_i) \approx GW^E(\mathcal{O}_{i+1})$.
- (ii) Paths in $\mathcal{M}_{emb} \cap V_i$, crossing the wall transversally at a single point of $\mathcal{W}^1 \setminus \mathcal{A}$. For these, Lemma 7.4 below shows that $GW^E(\mathcal{O}_i) \approx -GW^E(\mathcal{O}_{i+1})$.

Altogether, we conclude that

$$GW^E(\mathcal{O}) \approx (-1)^\sigma GW^E(\mathcal{O}_{elem}),$$

where σ is the number of transverse wall-crossings, which is exactly the spectral flow of the operator D_p along the path $\tilde{\gamma}$. The path ends at an elementary cluster, which has positive sign by (3.6). Thus $(-1)^\sigma$ is exactly the sign of the initial curve (C, J_0) . \square

In the above proof, the assertion in Step (i) is a fact about isotopies with no wall-crossings. It can be stated as follows.

Lemma 7.2 (Simple Isotopy). *Fix $E > 0$. Then for any path (C_t, J_t) in $\mathcal{M}_{emb} \setminus \mathcal{W}$ with J_t in \mathcal{J}_{isol}^E and any $\varepsilon_0, \varepsilon_1$ such that $(C_0, \varepsilon_0, J_0)$ and $(C_1, \varepsilon_1, J_1)$ are clusters,*

$$GW^E(C_0, \varepsilon_0, J_0) \approx GW^E(C_1, \varepsilon_1, J_1). \quad (7.3)$$

Proof. It follows from Proposition 5.1b and the compactness of $[0, 1]$ that there is a $\delta > 0$ such that, for each $t \in [0, 1]$, C_t is the only J_t -holomorphic curve in its degree and genus in the ball $B(C_t, \delta)$ (in Hausdorff distance). By Lemma 2.3 we can choose, for each $0 \leq t \leq 1$, an $0 < \varepsilon_t < \delta$ such that $(C_t, \varepsilon_t, J_t)$ is a cluster. Then, by Lemma 2.1, $(C_s, \varepsilon_t, J_s)$ has a well-defined contribution $GW^E(C_s, \varepsilon_t, J_s)$ for all s in an open interval around t . These open intervals cover $[0, 1]$; take a finite subcover $\{I_k\}$. Then

$GW^E(C_s, \varepsilon_k, J_s)$ is constant for s in each I_k and C_s is the only J_s -holomorphic curve in its genus and homology class in that ball. Corollary 2.5 shows that on the intersection of two consecutive intervals the corresponding GW^E invariants differ by the contributions of higher-level clusters. The lemma follows. \square

By Lemma 6.5, each path in \mathcal{J} with endpoints in \mathcal{J}_E^* can be deformed, keeping its endpoints, to a path γ in \mathcal{J}_{isol}^E such that the projection

$$\pi_\gamma : \mathcal{M}_{emb}^\gamma \rightarrow \gamma \quad (7.4)$$

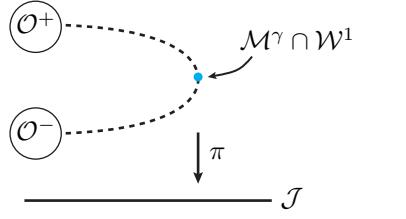
has only non-degenerate critical points, none an endpoint, each locally modeled by (6.5). If $a > 0$ in the local model, then γ can be parameterized so that $\gamma(t) = t$ and $\pi^{-1}(t) = \{x \mid t = ax^2\}$ is empty for $t < 0$ and is two distinct curves C_t^\pm for $0 < t < \delta$ (and vice versa if $a < 0$). A second isotopy lemma relates the GW invariants of clusters centered on these curves C_t^\pm .

Lemma 7.3 (Wall-crossing in \mathcal{J}). *Fix $E > 0$, a path γ in \mathcal{J}_{isol}^E and a non-degenerate critical point (C_0, J_0) of (7.4) for $J_0 = \gamma(0)$. Then there exists a $\delta > 0$ and a neighborhood U of (C_0, J_0) in \mathcal{M} such that if $0 \neq |t| < \delta$ and the sign of t is such that $\mathcal{M}^{\gamma(t)} \cap U = \{C_t^\pm\}$, then for two clusters $\mathcal{O}^+ = (C_t^+, \varepsilon, J_t)$, $\mathcal{O}^- = (C_t^-, \varepsilon', J_t)$ satisfy*

$$GW^E(\mathcal{O}^+) \approx -GW^E(\mathcal{O}^-).$$

Proof. The local model (6.5) at (C_0, J_0) implies that there is an $\varepsilon_1 > 0$ and a ball $U = B(C_0, \varepsilon_1)$ in $\mathcal{C}(X)$ that contains C_t^\pm and no other J_t holomorphic curves in the degree and genus of C_0 for all $|t| < \varepsilon_1$. Because $J_t \in \mathcal{J}_{isol}^E$, Lemma 2.3 ensures that ε_1 can be chosen so that $(C_0, \varepsilon_1, J_0)$ is a cluster. As J varies, the associated invariant $GW^E(U, J_s)$ is, by Lemma 2.1, well-defined and independent of s for small s .

The local model (6.5) shows that $U \cap \mathcal{C}^{J_s, E}$ is $\{C_t^\pm\}$ for $s = t$, and is empty for $s = -t$. Taking $s = t$ and applying Proposition 2.4, one sees that U decomposes into two clusters \mathcal{O}_t^\pm with $GW^E(\mathcal{O}^+) + GW^E(\mathcal{O}^-) \approx GW^E(U, J_s)$. Applying the same theorem with $s = -t$ shows that $GW^E(U, J_s) \approx 0$. The proof is completed by noting that the invariants $GW^E(\mathcal{O}^\pm)$ satisfy (7.1) as ε and ε' vary. \square



The core curve of a cluster does not persist through the wall-crossing described by Lemma 7.3. But the core curve remains if we fix the complex structure on the core curve C and cross the wall along a path γ in \mathcal{J}_C , as was done in the proof of Theorem 7.1. After a perturbation as in Lemma 6.7, the wall-crossing is locally modeled by (6.7). In this picture, for each $0 < t < \delta$, there are four curves to consider: the incoming core curve (C, J_{-t}) , the outgoing core curve (C, J_t) , and a second pair of curves (C'_{-t}, J_{-t}) and (C'_t, J_t) .

Lemma 7.4 (Wall-crossing in \mathcal{J}_C). *Fix $E > 0$ and a path γ in $\mathcal{J}_C \cap \mathcal{J}_{isol}^E$ so that $\tilde{\gamma}(t)$ crosses the wall transversally at $t = 0$ at a point (C, J_0) in $\mathcal{W}^1 \setminus \mathcal{A}$. Then there exists a $\delta > 0$ so that each incoming cluster $\mathcal{O}_{-\delta} = (C, \varepsilon, J_{-\delta})$ and each outgoing cluster $\mathcal{O}_\delta = (C, \varepsilon', J_\delta)$ satisfy*

$$GW^E(\mathcal{O}_{-\delta}) \approx -GW^E(\mathcal{O}_\delta).$$

Proof. Consider the local model $\mathcal{M}^S \rightarrow S$ given by (6.6). Its restriction over γ , given by (6.7), is two curves crossing at the origin. We will perturb this level set $\{z = 0\}$ in two opposite directions.

In the chart (6.6), with $\delta > 0$ fixed and small, $(C, J_{-\delta})$ has coordinates $(0, -\delta, 0)$, so can be perturbed to $(C_{-\delta, s}, J_{-\delta, s})$ with coordinates $(x(s), -\delta, s)$, where $x \approx -s/b\delta$ is the unique solution of $s = x(ax - b\delta + r(x, -\delta))$ with $x = O(s)$. The curves $(C'_{-\delta}, J_{-\delta})$ and (C, J_δ) can be similarly perturbed. By Lemma 2.3, the GW invariants of the corresponding clusters

$$A_s = (C_{-\delta, s}, \varepsilon, J_{-\delta, s}), \quad B_s = (C_{\delta, s}, \varepsilon, J_{\delta, s}), \quad D_s = (C'_{-\delta, s}, \varepsilon, J_{\delta, s})$$



FIGURE 4. These figures show the curve C (horizontal line) and curves C'_t (diagonal line) as t (the horizontal coordinate) varies. The circled labels refer to clusters at the ends of the dotted paths, with $s > 0$ in the first figure, $s < 0$ in the second, and $a, b > 0$ in both.

are locally constant in s : for sufficiently small s (and $\delta > 0$) we have

$$GW^E(A_s) = GW^E(\mathcal{O}_{-\delta}), \quad GW^E(B_s) = GW^E(\mathcal{O}_\delta), \quad GW^E(D_s) = GW^E(D_{-s}). \quad (7.5)$$

Assume $a > 0$ (else change $s \rightarrow -s$), and that $b > 0$ (else change $t \rightarrow -t$). The moduli space \mathcal{M}^S over S is locally near (C, J_0) the level set

$$\{(x, t, s) \mid s = x(ax + bt + r(x, t))\}.$$

For each fixed s, t small, this quadratic equation in x has either no solution or two solutions, except at a single point $x \approx -bt/2a$, where the tangent is in the kernel of the projection to \mathcal{J} , which means that this point lies on the wall and is a non-degenerate critical point of (7.4).

For a small positive s , the moduli space over $\gamma_s(t) = (t, s)$, $-\delta \leq t \leq \delta$ therefore contains a path in \mathcal{M}_{emb} from the core of cluster D_s to the core of B_s that does not cross the wall. After a small perturbation using Lemma 6.5, Lemma 7.2 applies to give:

$$GW^E(D_s) \approx GW^E(B_s). \quad (7.6)$$

For a small negative s the moduli space over $\gamma_s(t) = (t, s)$, $-\delta \leq t \leq 0$ is a path in \mathcal{M}_{emb} from the core of cluster A_s to the core of D_s , crossing now the wall transversally (at a point in $\mathcal{W}^1 \setminus \mathcal{A}$). Perturbing γ_s by Lemma 6.5 gives a path in \mathcal{J}_{isol} so Lemma 7.3 applies in this case to give

$$GW^E(A_s) \approx -GW^E(D_s). \quad (7.7)$$

The proof is completed by combining (7.5), (7.6) and (7.7). \square

8. STRUCTURE THEOREMS AND THE PROOF OF THE GV CONJECTURE

The isotopy results of the previous section lead quickly to a formula (8.1) that shows that the GW invariants have a remarkably simple structure. This formula is compatible with a simple geometric picture: if one could find a $J \in \mathcal{J}$ so that all simple J -holomorphic maps in X were elementary, then $GW(X)$ would have exactly the form (8.1), with $e_{A,g}(X)$ equal to the count of J -holomorphic curves with homology class A and genus g . However, it is far from clear whether any such J exists. Thus the coefficients $e_{A,g}(X)$ can be regarded as *virtual counts of elementary clusters in X* .

Theorem 8.1. *For any closed symplectic Calabi-Yau 6-manifold X , there exist unique integer invariants $e_{A,g}(X)$ such that*

$$GW(X) = \sum_{A \neq 0} \sum_{g \geq 0} e_{A,g}(X) \cdot GW_g^{elem}(q^A, t). \quad (8.1)$$

Proof. The uniqueness of the coefficients in (8.1) is easily shown because the collection of series

$$GW_g^{elem}(q^A, t) = t^{2g-2} q^A \left(1 + \text{higher order in } t \text{ and } q^A \right) \quad (8.2)$$

for $g \geq 0$ and $A \in H_2(X, \mathbb{Z})$ is linearly independent. To prove existence, fix E , choose any parameter $J \in \mathcal{J}_E^*$, and use Proposition 2.4 to write $GW^E(X)$ as a sum of finitely many cluster contributions. Formula (8.1) follows from the corresponding formula for each cluster, which is proved in Lemma 8.2 below, by taking $E \rightarrow \infty$. \square

Lemma 8.2. *For any regular E -cluster \mathcal{O} centered at a genus g curve C there exist unique integers $e_{d,h}(\mathcal{O})$, beginning with $e_{1,g}(\mathcal{O}) = \text{sign}(C)$, such that*

$$GW^E(\mathcal{O}) = \sum_{d \geq 1} \sum_{h \geq g} e_{d,h}(\mathcal{O}) GW_h^{elem,E}(q^{dC}, t), \quad (8.3)$$

where both sides are truncated below energy level E .

Proof. Because all J -holomorphic maps in \mathcal{O} represent $k[C]$ and have genus at least g , $GW^E(\mathcal{O})$ has the form

$$GW^E(\mathcal{O}) = \sum_{k \geq 1} \sum_{h \geq g} GW_{k,h}^E(\mathcal{O}) q^{kC} t^{2h-2} \quad (8.4)$$

with $k\omega(C) \leq E$ and $h \leq E$. Define the (C, g) -relative level of the monomial $t^{2h-2} q^{kC}$ to be $\Omega(k) + h - g$, and note that all terms in (8.4) have non-negative relative level. Let $GW_m^E(\mathcal{O})$ denote the sum of the terms in (8.4) with $\Omega(k) + h - g \leq m$. Substituting (8.2) into (8.3), one sees that $GW_m^E(\mathcal{O})$ involves only those $e_{d,h}(\mathcal{O})$ with $\Omega(d) + h - g \leq m$.

We will determine the weights in (8.3) by complete induction on m , starting with the trivial case $m = -1$. Thus we assume that for every regular cluster \mathcal{O} , whose core curve corresponds to any (A, g) , there is an expansion (8.3) truncated at (A, g) -relative level $m - 1$. Now by Theorem 7.1 we have

$$GW^E(\mathcal{O}) = \pm GW_g^{elem,E}(q^C, t) + \sum_{i \in I} \pm GW^E(\mathcal{O}_i),$$

where the \mathcal{O}_i are clusters whose core curves C_i have $[C_i] = k_i[C]$, genus $g_i \geq g$ and (C, g) -relative level $m_i = \Omega(k_i) + g_i - g > 0$. When $GW^E(\mathcal{O})$ is truncated at relative level m , each $GW^E(\mathcal{O}_i)$ is truncated at (C_i, g_i) -relative level $m - m_i < m$ so by induction,

$$GW^E(\mathcal{O}) = \pm GW_g^{elem,E}(q^C, t) + \sum_{i \in I} \pm \left(\sum_{d,h} e_{d,h}(\mathcal{O}_i) GW_h^{elem,E}(q^{dk_i C}, t) \right)$$

holds when truncated at (C, g) -relative level m . This completes the induction step. \square

In fact, we get the following result for *any* closed symplectic 6-dimensional manifold X as long as we restrict to the GW invariants coming only from classes $A \in H_2(X, \mathbb{Z})$ with vanishing Chern number $c_1(A) = c_1(X)A$.

Theorem 8.3. *Assume X is a closed symplectic 6-manifold. Then there exist unique integer invariants $e_{A,g}(X)$, defined for homology classes A with $c_1(A) = 0$, such that the GW invariant of X satisfies*

$$\sum_{\substack{A \neq 0 \\ c_1(A)=0}} \sum_{g \geq 0} GW_{A,g}(X) t^{2g-2} q^A = \sum_{\substack{A \neq 0 \\ c_1(A)=0}} \sum_{g \geq 0} e_{A,g}(X) \cdot GW_g^{elem}(q^A, t). \quad (8.5)$$

Proof. The dimension (1.3) is $2c_1(A)$, independent of the genus. It suffices to check that all the results in Sections 1-6 continue to hold as long as we replace everywhere $\overline{\mathcal{M}}(X)$ by the union of its zero dimensional pieces

$$\bigsqcup_{\substack{A \neq 0 \\ c_1(A)=0}} \overline{\mathcal{M}}_{A,g}(X). \quad (8.6)$$

A dimension count shows that for generic J the limit points of (8.6) in the rough topology (after restricting below fixed energy level E) can only be multiple covers of points of (8.6), and not of points with $c_1(A) \neq 0$. The rest is straightforward, and the details are left to the reader. \square

8.1. Proof of the GV conjecture. The GV conjecture follows easily from Theorem 8.1 and the explicit form of the GW invariant of an elementary cluster. For simplicity, set

$$\mathcal{E}_h(q, t) = \sum_{k=1}^{\infty} \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{2h-2} q^k.$$

With this notation, the GW invariant (3.7) of an elementary cluster whose core curve has genus g is

$$GW_g^{elem}(q, t) = \sum_{d \neq 0} \sum_{h \geq g} n_{d,h}(g) \mathcal{E}_h(q^d, t), \quad (8.7)$$

where the $n_{d,h}(g)$ are integers by Proposition 3.4a that vanish unless $h \geq g$ by Proposition 3.4b. The GV Conjecture then takes the following form.

Theorem 8.4. *Let X be a closed symplectic Calabi-Yau 6-manifold. Then there are unique integers $e_{A,h}(X)$ such that*

$$GW(X) = \sum_{A \neq 0} \sum_h e_{A,h}(X) \mathcal{E}_h(q^A, t). \quad (8.8)$$

In fact, these BPS numbers $n_{A,h}(X)$ are obtained from the virtual counts $e_{A,h}(X)$ of Theorem 8.1 by the universal formula involving the coefficients $n_{d,h}(g)$ in (8.7):

$$n_{A,h}(X) = \sum_{\substack{d, B \\ dB=A}} \sum_{g=0}^h e_{B,g}(X) \cdot n_{d,h}(g) \in \mathbb{Z}, \quad (8.9)$$

where the first sum is over all integers $d \geq 1$ and $B \in H_2(X, \mathbb{Z})$ such that $dB = A$ in $H_2(X, \mathbb{Z})$.

Proof. This follows immediately by combining (8.1) and (8.7) and rearranging the sums:

$$\begin{aligned} GW(X) &= \sum_{A \neq 0} \sum_{g \geq 0} e_{A,g}(X) \sum_{d \geq 1} \sum_{h \geq g} n_{d,h}(g) \mathcal{E}_h(q^{dA}, t) \\ &= \sum_{A \neq 0} \sum_{d \geq 1} \sum_h \left(\sum_{g \leq h} e_{A,g}(X) n_{d,h}(g) \right) \mathcal{E}_h(q^{dA}, t) \\ &= \sum_{A \neq 0} \sum_{\substack{d, B \\ dB=A}} \sum_h \left(\sum_{g \leq h} e_{B,g}(X) n_{d,h}(g) \right) \mathcal{E}_h(q^A, t). \end{aligned}$$

The rearrangements are justified by first working below an energy level $E(A, g) \leq E$, where all sums are finite. \square

9. EXTENSIONS OF THE GV STRUCTURE THEOREM

This section extends Theorem 8.4 in two different directions: to general symplectic 6-manifolds, and to the genus zero GW invariants of closed symplectic n -manifolds, $n \geq 6$, that are semipositive (as defined in [MS]), a class that includes symplectic Calabi-Yau manifolds. In fact, all transversality results were proved for simple maps in index zero moduli spaces. A version of the Cluster Decomposition Proposition 2.4 holds provided the underlying curve map (1.8) does not increase the dimension of such moduli spaces in the sense described below.

We restrict to the *primary GW invariants* of X , which are defined using the evaluation map (but not the stabilization map) in (1.1). For each collection $\{\gamma_i\} \subset H^*(X, \mathbb{Z})$ consider the generating function

$$GW^X(\gamma_1, \dots, \gamma_k) = \sum_{A \neq 0} \sum_{g \geq 0} \langle [\overline{\mathcal{M}}_{A,g,k}^J(X)]^{vir}, \text{ev}^*(\gamma_1 \times \dots \times \gamma_k) \rangle q^A t^{2g-2}. \quad (9.1)$$

The pairing is defined to be zero unless the formal dimension is zero, that is, unless $\iota = 0$, where

$$\iota = 2c_1(A) + (\dim X - 6)(1-g) + 2k - \sum_{i=1}^k \dim \gamma_i. \quad (9.2)$$

As usual, the pairing vanishes unless $\dim \gamma_i \geq 2$ for each i , so we henceforth assume this inequality. Throughout this section, we assume that $\dim X \geq 6$.

The coefficients in (9.1) are obtained by fixing pseudo-cycles $\beta_i : B_i \rightarrow X$ representing the Poincaré duals of γ_i (cf. [MS, Section 6.5]), and restricting to the index $\iota = 0$ constrained moduli space

$$\overline{\mathcal{M}}_{A,g,B}^J(X) = \overline{\mathcal{M}}_{A,g,k}^J(X) \times_{X^k} (B_1 \times \dots \times B_k) \quad (9.3)$$

(the fiber product of the evaluation map $\text{ev} : \overline{\mathcal{M}}_{A,g,k}^J(X) \rightarrow X^k$ and the map $B_1 \times \dots \times B_k \rightarrow X^k$). This gives rise in the usual way to the primary GW invariant that appears as the coefficients in (9.1) (cf. page 197 of [MS]).

Lemma 1.1 remains true for these index 0 constrained moduli spaces in the following form. Let \mathcal{D} be the (countable) set consisting of the indexing data (A, g, γ) appearing in (9.1). For each $\gamma = (\gamma_1, \dots, \gamma_k)$, choose a set of pseudo-cycle representatives β_1, \dots, β_k that are in general position. Standard transversality results show that, for each element of \mathcal{D} , each open stratum of the constrained universal moduli space

$$\overline{\mathcal{M}}_{A,g,B}(X)_{\text{simple}} \rightarrow \mathcal{J} \quad (9.4)$$

is a manifold. The Sard-Smale Theorem gives a Baire set of regular points in \mathcal{J} for the map (9.4) for each \mathcal{D} ; after intersecting over the elements of \mathcal{D} we can assume these are regular for all \mathcal{D} . Lemma A.1ab, together with the Sard-Smale Theorem, give two similar Baire sets. Another intersection produces a single Baire set \mathcal{J}^* of \mathcal{J} such that for each $J \in \mathcal{J}^*$ all index 0 moduli spaces (9.3) satisfy:

- (a) All simple J -holomorphic maps are regular, and are embeddings with pairwise disjoint images that are *B-regular*, meaning that for each $i = 1, \dots, k$, $f(x_i)$ is a regular value of β_i and $(\beta_i)_*(T_{b_i} B_i) \cap f_*(T_{x_i} C) = 0$ for each $b_i \in B_i$ with $\beta_i(b_i) = f(x_i)$.
- (b) The projection (9.4) is a local diffeomorphism around each map that is regular, *B-regular*, and an embedding.

Moreover, for each $J \in \mathcal{J}^*$, there are no simple J -holomorphic maps in the spaces (9.3) with $\iota < 0$. The universal moduli space constrained by B is

$$\overline{\mathcal{M}}_B(X) = \bigsqcup_{A,g} \overline{\mathcal{M}}_{A,g,B}(X), \quad (9.5)$$

where the disjoint union is only over those (A, g) for which ι in (9.2) is zero.

As in the proof of Lemma 1.4, any nontrivial J -holomorphic map $f : C \rightarrow X$ has an associated “reduced map” $\varphi : C_{\text{red}} \rightarrow X$, which is a simple J -holomorphic map with the same image as f . In this context, the underlying curve map (1.7) extends to a map

$$c : \overline{\mathcal{M}}_B(X) \longrightarrow \text{Subsets}(X) \times \mathcal{J} \times B_1 \times \dots \times B_k \quad (9.6)$$

defined by $(f, x_1, \dots, x_k, J, b_1, \dots, b_k) \mapsto (f(C), J, b_1, \dots, b_k)$. The examples below give structure theorems in cases where c does not increase the formal dimension (9.2), that is, where $\iota(\varphi) \leq \iota(f)$.

Under this assumption, we can replace $\overline{\mathcal{M}}(X)$ everywhere by (9.5) and all proofs in Sections 1-7, except those in Section 3, hold without change. In particular, there is a dense, path-connected set $\mathcal{J}_{\text{isol}}^E(B)$ corresponding to $\mathcal{J}_{\text{isol}}^E$ in Definition 1.2 but involving only maps in $\overline{\mathcal{M}}_B(X)$. Lemma 1.5 holds under the assumption above with $\mathcal{C}(X)$ replaced by the image of (9.6).

To finish, we must expand the definitions of “cluster” and “elementary cluster”. Define a B -constrained cluster exactly as in Definition 2.2, but using only elements $(f, \mathbf{x}, J, \mathbf{b})$ in $\overline{\mathcal{M}}_B(X)$. Thus the core C is a smooth embedded J -holomorphic curve $\iota_C : C \rightarrow X$ that we now assume is marked, B -regular as defined in (a) above, and decorated by a choice of $b_i \in \beta_i^{-1}(\iota_C(x_i))$ for each i . The contribution of a B -constrained cluster (C, J, ε) to $GW(\gamma)$ depends only on the restriction of J to the ε neighborhood of the core curve and the restriction of each β_i to the ε -ball in B_i centered at b_i , for $i = 1, \dots, k$. By a diffeomorphism, when ε is small, can identify the ε tubular neighborhood of C with an ε -disk bundle of the normal bundle $N_C \rightarrow C$, with C mapping to the zero section and each ε -ball around b_i in B_i mapping into a linear subspace of the fiber N_i over the points $p_i = \iota_C(x_i)$. One can then declare certain B -constrained clusters to be “elementary”. In both of the examples below there is a simple, natural way of doing this.

9.1. GV-formula for general symplectic 6-manifolds. For any closed symplectic 6-manifold X , the dimension (9.2) is

$$\iota = 2c_1(A) + \sum_{i=1}^k (2 - \dim \gamma_i),$$

independent of the genus. For the “Calabi-Yau classes” A with $c_1(A) = 0$ consider the GV-transform

$$\sum_{\substack{A \neq 0, g \\ c_1(A)=0}} GW_{A,g} q^A t^{2g-2} = \sum_{\substack{A \neq 0, g \\ c_1(A)=0}} n_{A,g} \sum_{k=1}^{\infty} \frac{1}{k} \left(2 \sin \frac{kt}{2} \right)^{2g-2} q^{kA}. \quad (9.7)$$

For “Fano classes” A with $c_1(A) > 0$ consider the following variation of the GV transform:

$$\sum_{\substack{A, g \\ c_1(A)>0}} GW_{A,g}(\gamma_1, \dots, \gamma_k) q^A t^{2g-2} = \sum_{\substack{A, g \\ c_1(A)>0}} n_{A,g}(\gamma_1, \dots, \gamma_k) \left(2 \sin \frac{t}{2} \right)^{c_1(A)+2g-2} t^{-c_1(A)} q^A \quad (9.8)$$

for each collection $\{\gamma_i\} \subset H^*(X, \mathbb{Z})$. The invariants $GW_{A,g}$ are zero for all classes A with $c_1(A) < 0$ (the moduli space without constraints is empty for $J \in \mathcal{J}^*$).

Theorem 9.1. *For a closed symplectic 6-dimensional manifold X , the coefficients of the primary GW series (9.7) and (9.8) have the following integrality properties:*

$$n_{A,g} \in \mathbb{Z} \quad \text{if } c_1(A) = 0, \quad \text{and} \quad n_{A,g}(\gamma_1, \dots, \gamma_k) \in \mathbb{Z} \quad \text{if } c_1(A) > 0$$

for all $\gamma_1, \dots, \gamma_k \in H^*(X, \mathbb{Z})$.

Proof. Fix $\gamma = \{\gamma_i\}$, corresponding constraints $B = \{B_i\}$, and a class $A \neq 0$ so that

$$\iota = 2c_1(A) + \sum (2 - \dim \gamma_i)$$

is zero. For each $J \in \mathcal{J}^*$, the resolution of each J -holomorphic map f , factors as $\varphi \circ \rho$ as described above, where $\varphi : C_{red} \rightarrow X$ is a simple J -holomorphic map. For each component Σ_i of C_{red} , let $A_i = \varphi_*[\Sigma_i] \in H_2(X, \mathbb{Z})$, and let $d_i \geq 1$ denote the degree of ρ over Σ_i , that is, the number of points in $\rho^{-1}(x)$ for a generic point $x \in \Sigma_i$. Then $A = \sum d_i A_i$.

Because $\varphi : C_{red} \rightarrow X$ is a simple J -holomorphic map, it cannot have any components with $c_1(A_i) < 0$ (such moduli spaces are empty for $J \in \mathcal{J}^*$). Moreover, the image of φ passes through all the constraints and represents $\sum A_i$. Hence

$$c_1(\varphi_*[C_{red}]) = \sum c_1(A_i) \leq \sum d_i c_1(A_i) = c_1(A);$$

in fact this must be an equality because otherwise the formal dimension of the constrained moduli space containing φ would be negative, contradiction. Thus for $J \in \mathcal{J}^*$, a multiple cover map represents a Calabi-Yau class if and only if its reduced curve also does, and represents a Fano class if and only if its reduction also does and $\deg \rho = 1$.

Consequently, the GW series separates into two independent contributions: a sum over the Calabi-Yau classes, where Theorem 8.3 applies, and a sum over the Fano classes that was studied by A. Zinger

[Z]. Theorem 8.3 combines with the proof of Theorem 8.4 to give the integrality of $n_{A,g}$ in (9.7). The Fano case is much simpler: there is no need to consider clusters because for $J \in \mathcal{J}^*$, every embedded J -holomorphic curve C with $c_1(A) > 0$ is isolated and super-rigid for the constrained moduli space, and the contribution of degree 1 maps from nodal curves to C is precisely

$$GW(C) = \left(2 \sin \frac{t}{2}\right)^{c_1(A)+2g-2} t^{-c_1(A)} q^C$$

(see (1.13) and (1.14) in [Z]). This completes the proof. \square

9.2. Genus zero invariants of semipositive manifolds. There is a similar structure theorem for the rational (genus zero) GW invariant of closed semipositive symplectic manifolds of dimension ≥ 6 . In this context, the appropriate GV transform has two parts:

- (1) For $c_1(A) = 0$, it is the Aspinwall-Morrison formula in the form given in equation (2) in [KP]:

$$\sum_{\substack{A \neq 0 \\ c_1(A)=0}} GW_{A,0}^X(\gamma_1, \dots, \gamma_k) q^A = \sum_{\substack{A \neq 0 \\ c_1(A)=0}} n_{A,0}^X(\gamma_1, \dots, \gamma_k) \sum_{d \geq 1} d^{k-3} q^{dA} \quad (9.9)$$

- (2) For $c_1(A) > 0$, it is (9.8) specialized to genus zero:

$$GW_{A,0}^X(\gamma_1, \dots, \gamma_k) = n_{A,0}^X(\gamma_1, \dots, \gamma_k). \quad (9.10)$$

As before, the invariants $GW_{A,g}$ are zero for all classes A with $c_1(A) < 0$ (since X is semipositive, there are no simple J -holomorphic spheres with $c_1(A) < 0$ for $J \in \mathcal{J}^*$).

Theorem 9.2. *For a closed semipositive symplectic manifold X of dimension at least 6, the coefficients (9.9) and (9.10) of the primary genus zero GW series have the following integrality property:*

$$n_{A,0}^X(\gamma_1, \dots, \gamma_k) \in \mathbb{Z}$$

for all $\gamma_1, \dots, \gamma_k \in H^*(X, \mathbb{Z})$.

Proof. Again fix γ , B , and A so that

$$\iota = 2c_1(A) + (\dim X - 6) + \sum (2 - \dim \gamma_j)$$

is zero. As before, assume f is a multiple cover with φ the reduced map, and A_i , d_i are degrees of its components. If the domain of φ has $r \geq 1$ components then its image has at least $r - 1$ self-intersection points (since the domain of f was connected). For $J \in \mathcal{J}^*$, these impose $(r - 1)(\dim X - 4)$ transversely cut conditions on simple maps, so the dimension of the moduli space containing φ is

$$\sum_{i=1}^r [2c_1(A_i) + (\dim X - 6)] - (r - 1)(\dim X - 4) + \sum_{j=1}^k (2 - \dim \gamma_j).$$

Since $A = \sum d_i A_i$ and $c_1(A_i) \geq 0$, this is less than or equal to $\iota - 2(r - 1) \leq \iota = 0$. But the moduli space is empty unless this is an equality, so we conclude that $r = 1$ and $d_1 = 1$ whenever $c_1(A) \neq 0$. Thus the GW series again separates into a sum of $c_1(A) = 0$ classes and a sum over $c_1(A) > 0$ classes.

The Fano case (9.10) is classical: dimension counts imply that for generic J the constrained moduli space consists only of simple maps, without any multiple covers thus the GW invariant is an integer.

For Calabi-Yau classes A in X^{2m} we declare a B -constrained cluster to be elementary if its core C is an embedded marked rational curve with normal bundle N biholomorphic to $\mathcal{O}(-1) \oplus \mathcal{O}(-1) \oplus (m-3)\mathcal{O}$, and the constraints B_i are linear subspaces of fibers $N_{f(x_i)}$ of N in general position in the sense that the only holomorphic section of N which intersects B_i for every $i = 1, \dots, k$ is the zero section. (Note that because $\iota = 0$, the sum of the codimensions of B_i in $N_{f(x_i)}$ is $2(m-3) = \dim_{\mathbb{R}} X - 6$.) The core curve C is then super-rigid (the constraints kill the kernel in the \mathcal{O} directions) and, as proved in [KP], the contribution of its multiple covers to the primary $GW(\gamma_1, \dots, \gamma_k)$ invariant is

$$\text{sign}_B C \cdot \sum_{d \geq 1} d^{k-3} q^d,$$

where $\text{sign}_B C = \pm 1$ is the sign of the core curve C as an element of the cutdown moduli space $\overline{\mathcal{M}}_{A,0,B}$. This sign can be explicitly calculated as the sign of the (transverse) intersection between the oriented linear subspace $B_1 \times \dots \times B_k$ and the image of the evaluation map $\text{ev}_x : H^0(C, f^*N) \rightarrow N_{f(x_1)} \times \dots \times N_{f(x_k)}$ on the space of holomorphic sections of f^*N . With this, Theorem 8.1 extends to give precisely (9.9). \square

APPENDIX

The proofs of Lemma 1.1a and Lemma 5.3 were deferred; here we give the details. The proofs are applications of transversality and the Sard-Smale Theorem. While we are primarily interested in Calabi-Yau 6-manifolds, Propositions A.2 and A.4 below apply to symplectic manifolds (X, ω) with $\dim X \geq 6$. Lemmas 1.1a and 5.3 are special cases of Corollary A.3 and Proposition A.4 respectively.

For any J -holomorphic map $f : C \rightarrow X$, its resolution $\tilde{f} : \tilde{C} \rightarrow X$ is J -holomorphic and has a smooth (but not necessarily connected) domain, so it suffices to work with such maps. For every integer $\ell \geq 0$, let

$$\mathcal{M}_{\ell, \text{simple}} \rightarrow \mathcal{J} \tag{A.1}$$

denote the universal moduli space consisting of tuples (f, C, J) , where f is a simple J -holomorphic map whose domain $C = (\Sigma, j, x_1, \dots, x_\ell)$ is a smooth and compact (but not necessarily connected) complex curve with ℓ marked points. Write $\mathcal{M}_{0, \text{simple}}$ as $\mathcal{M}_{\text{simple}}$, and let

$$\mathcal{NE} \subset \mathcal{M}_{\text{simple}}$$

be the subset of the universal moduli space consisting of simple maps that are not embedded. Maps in \mathcal{NE} either (a) are not one-to-one, or (b) are not immersions. Correspondingly, consider two types of subsets of the moduli spaces (A.1):

(a) For a pair of marked points, the inverse image of the diagonal under the evaluation map

$$\text{ev} : \mathcal{M}_{2, \text{simple}} \longrightarrow X \times X \tag{A.2}$$

is the subset

$$\text{ev}^{-1}(\Delta) \subset \mathcal{M}_{2, \text{simple}}$$

consisting of simple J -holomorphic maps f with smooth marked domain but whose image has a double point $f(x_1) = f(x_2)$.

(b) For a single marked point, there is a relative cotangent bundle \mathcal{L} over $\mathcal{M}_{1, \text{simple}}$ and an evaluation map $\text{ev} : \mathcal{M}_{1, \text{simple}} \rightarrow X$. The bundle

$$\mathcal{L} \otimes_{\mathbb{C}} \text{ev}^* TX \xleftarrow{\Phi} \mathcal{M}_{1, \text{simple}} \tag{A.3}$$

has a section Φ defined by $\Phi(f, x) = (df)_x \in T_x^* C \otimes_{\mathbb{C}} T_{f(x)} X$. The zero set $\Phi^{-1}(0)$ consists of simple, marked J -holomorphic maps f that are not immersions at the marked point x .

Lemma A.1. *For each $\ell \geq 0$, the moduli space (A.1) is a manifold. Furthermore,*

- (a) *the evaluation map (A.2) is transverse to the diagonal, and*
- (b) *the section (A.3) is transverse to the zero section.*

Proof. The first statement is proved the same way as Proposition 5.1a, regarding $\mathcal{M}_{\ell, \text{simple}}$ as locally contained in a slice. Statement (a) is true by Proposition 3.4.2 of [MS]. For (b), fix $p = (f, J)$ and recall that $T_p \mathcal{M}_{1, \text{simple}}$ is the kernel of the operator \mathcal{L}_p in (4.8) and (4.9). By Theorem E.1.2 of [MS], one can find a neighborhood U of x in C so that f is an embedding on $U^* = U \setminus \{x\}$ and U^* consists of injective points. We can further assume that U has a holomorphic coordinate z centered at x , and $f(U)$ lies in a coordinate ball B in X centered at $f(x)$. To prove (b), we pick $v \in T_{f(x)} X$, extend it to a smooth vector field on $f(U)$, and show that there is an element $(\zeta, 0, K)$ in $\ker \mathcal{L}_p$ such that (i) $\zeta(z) = zv + O(|z|^2)$, (ii) ζ is supported on a small ball around x in U , and (iii) K is supported on B .

To achieve (i)–(iii), one uses ideas from the proof of Lemma 3.4.3 of [MS]. Starting with a local solution $\xi_0 = (\zeta_0, 0)$ to the equation $D_p \xi_0 = 0$ that satisfies (i), one sets $\zeta = \beta \zeta_0$ where β is a bump function supported on U with $\beta = 1$ on a small ball around x , and then constructs a K satisfying (iii) and $\frac{1}{2}Kf_*j = -D_p(\zeta, 0)$ using the explicit formula in the proof of Lemma 3.2.2 of [MS]. Then $(\zeta, 0, K)$ is the desired element in $\ker \mathcal{L}_p$. \square

Next, note that for every $\ell \geq 0$ the map

$$\pi_\ell : \mathcal{M}_{\ell, simple} \rightarrow \mathcal{M}_{simple} \quad (\text{A.4})$$

that forgets the marked points is a submersion with index $\pi_\ell = 2\ell$.

Proposition A.2. *The set $\mathcal{NE} \subset \mathcal{M}_{simple}$ of simple maps that are not embeddings has codimension $\dim X - 4$ in the sense of Definition 5.2.*

Proof. As above, \mathcal{NE} is the union of $\pi_1(\iota_1 S^1)$ and $\pi_2(\iota_2 S^2)$, where $\iota_1 : S^1 = \Phi^{-1}(0) \hookrightarrow \mathcal{M}_{1,simple}$ and $\iota_2 : S^2 = \text{ev}^{-1}(\Delta) \hookrightarrow \mathcal{M}_{2,simple}$. Lemma A.1 implies that ι_1 and ι_2 are inclusions of submanifolds with index $\iota_\ell = -\text{codim } S^\ell$ for $\ell = 1, 2$. One sees from (A.2) and (A.3) that S^1 and S^2 both have codimension $\dim X$. Hence \mathcal{NE} is a set of codimension k where $k = -\text{index}(\pi_\ell \circ \iota_\ell) = -\text{index} \pi_\ell - \text{index} \iota_\ell = -2\ell + \dim X \geq \dim X - 4$. \square

Corollary A.3. *If X is a symplectic Calabi-Yau 6-manifold, there is a Baire set $\mathcal{J}^* \subset \mathcal{J}$ so that for each $J \in \mathcal{J}^*$ all simple J -holomorphic maps are regular, and are embeddings with pairwise disjoint images.*

Proof. The Sard-Smale Theorem implies that the regular values of the projection $\pi : \mathcal{M}_{simple} \rightarrow \mathcal{J}$ are a Baire set \mathcal{J}_1 in \mathcal{J} . When X is a Calabi-Yau 6-manifold, the index of π is 0, and Proposition A.2 shows that \mathcal{NE} is a codimension 2 subset of \mathcal{M}_{simple} . Applying Sard-Smale again, there is a Baire set \mathcal{J}_2 of \mathcal{J} such that for each $J \in \mathcal{J}^* = \mathcal{J}_1 \cap \mathcal{J}_2$, J is a regular value of π and $\pi^{-1}(J)$ is disjoint from \mathcal{NE} , which means that all simple J -holomorphic curves are embedded. \square

The remainder of the appendix is devoted to the proof of the following result, which immediately implies Lemma 5.3, and also generalizes parts (c) and (d) of Proposition 5.1.

Proposition A.4. *Suppose $\dim X \geq 6$ and \mathcal{N} is a component of \mathcal{M}_{simple} such that the projection $\pi : \mathcal{N} \rightarrow \mathcal{J}$ has index 0. Then*

(a) $\mathcal{W}^1 \cap \mathcal{N}$ is a codimension 1 submanifold of \mathcal{N} , and

(b) $(\mathcal{W} \setminus \mathcal{W}^1) \cap \mathcal{N}$ is a subset of \mathcal{N} of codimension ≥ 3 .

Furthermore, (a) and (b) hold with \mathcal{N} replaced by $\mathcal{N} \cap \mathcal{M}_C$.

The proof is based on another construction involving the spaces (A.1). For each $\ell \geq 0$, let

$$\mathcal{M}_{\ell, simple}^* \subseteq \mathcal{M}_{\ell, simple}$$

be the open set of all $q = (f, J, x_1, \dots, x_\ell)$ such that x_i is an injective point of f for $i = 1, \dots, \ell$. Pull back the bundles \mathcal{E} and \mathcal{F} to $\mathcal{M}_{\ell, simple}^*$ by the maps (A.4). Let $\mathcal{E}^\ell \rightarrow \mathcal{M}_{\ell, simple}^*$ be the subbundle of $\pi_\ell^* \mathcal{E}$ whose fiber at $q = (p, x_1, \dots, x_\ell)$ is the set

$$\mathcal{E}_q^\ell = \{ \xi \in \pi_\ell^* \mathcal{E}_p \mid \xi^N(x_i) = 0 \text{ for } i = 1, \dots, \ell \}$$

of elements of $\pi_\ell^* \mathcal{E}_p$ whose normal component vanishes at x_1, \dots, x_ℓ . For each $q = (p, \mathbf{x}) \in \mathcal{M}_{\ell, simple}^*$, \mathcal{E}_q^ℓ is a linear subspace of $\pi_\ell^* \mathcal{E}_p$ of codimension $\ell(\dim X - 2)$, and the linearization D_p at p , given in (4.9), restricts to a linear map

$$D_q^\ell : \mathcal{E}_q^\ell \rightarrow \pi_\ell^* \mathcal{F}_p. \quad (\text{A.5})$$

Regarding D_q^ℓ as the composition of the inclusion $\mathcal{E}_q^\ell \hookrightarrow \pi_\ell^* \mathcal{E}_p$ with D_p , one sees that D_q^ℓ is Fredholm with index

$$\iota_\ell = \text{index } D_q^\ell = \text{index } D_p - \ell(\dim X - 2) \quad (\text{A.6})$$

for all $q = (p, \mathbf{x}) \in \mathcal{M}_{\ell, simple}^*$. As in (5.4), the space of Fredholm operators of type (A.5) is a bundle $\text{Fred}_{\iota_\ell} \rightarrow \mathcal{M}_{\ell, simple}^*$ stratified by submanifolds $\text{Fred}_{\iota_\ell}^s$, and

$$\Psi^\ell(f, J) = D_q^\ell$$

defines a section of this bundle. Let

$$V^\ell \subseteq \mathcal{M}_{\ell, simple}^*$$

be the open set of all $q = (p, x_1, \dots, x_\ell)$ such that there exists an injective point y , distinct from the set $\{x_i\}$, such that $\ker D_{(q,y)}^{\ell+1} = 0$.

Lemma A.5. *The section Ψ^ℓ is transverse to $\text{Fred}_{\iota_\ell}^s$ along V^ℓ , as is its restriction to $\pi_\ell^{-1}(\mathcal{M}_C)$. Hence for $s \geq 1$ the sets*

$$S^{\ell,s} = V^\ell \cap (\Psi^\ell)^{-1}\text{Fred}_{\iota_\ell}^s = \{q \in V^\ell \mid \dim \ker D_q^\ell = s\}. \quad (\text{A.7})$$

and $S_C^{\ell,s} = S^{\ell,s} \cap \pi_\ell^{-1}(\mathcal{M}_C)$ are submanifolds of codimension $s(s - \iota_\ell)$, where ι_ℓ is given by (A.6).

Proof. To prove transversality at $q \in S^{\ell,s}$ we must show that the image of $(d\Psi^\ell)_q$ projects surjectively onto the normal space $\text{Hom}(\ker D_q^\ell, \text{coker } D_q^\ell)$ to $\text{Fred}_{\iota_\ell}^s$ at D_q^ℓ . Fix a slice containing q and identify $\text{coker } D_q^\ell$ with the kernel of the adjoint operator $(D_q^\ell)^*$ defined as in (4.15). By contradiction, assume there exists a non-zero element of the normal space, regarded as a linear map $A_q : \ker D_q^\ell \rightarrow \ker(D_q^\ell)^*$, such that $\langle A_q, (\delta_v D_q^\ell) \rangle_{L^2} = 0$ for every variation $v \in T_q \mathcal{M}_{\ell, simple}$. Since $A_q \neq 0$, there exists an L^2 -normalized $\kappa \in \ker D_q^\ell$ such that $c = A_q \kappa \in \ker(D_q^\ell)^*$ is nonzero. Fix an L^2 orthonormal basis $\{\kappa_i\}$ of $\ker D_q^\ell$ with $\kappa_1 = \kappa$. We then have

$$0 = \langle A_q, (\delta_v D_q^\ell) \rangle_{L^2} = \sum_i \int_C \langle A_q \kappa_i, (\delta_v D_q^\ell) \kappa_i \rangle \quad (\text{A.8})$$

for all $v \in T_q \mathcal{M}_{\ell, simple}$.

The assumption that $q \in V^\ell$ means, by definition, that there is an injective point $y \notin \{x_1, \dots, x_\ell\}$ such that the map

$$\text{ev}_x : \ker D_q^\ell \rightarrow N_{f(x)} \quad \text{given by } \kappa \mapsto \kappa^N(x) \quad (\text{A.9})$$

is injective for $x = y$, and hence for all x in a neighborhood of y . As in the proof of Lemma 4.1, there exists an injective point x in that neighborhood with $c(x) \neq 0$. For this x , the values $\{\kappa_i^N(x)\}$ are linearly independent because (A.9) is injective. Applying Lemma A.6 below with $\xi = \kappa_1^N(x) \neq 0$ and $V = \text{span}\{\kappa_i^N(x) \mid i \geq 2\}$ produces a K satisfying (A.11) below.

Now proceed as in the proof of Proposition 5.1c, taking $v_\varepsilon = (0, 2\beta_\varepsilon K)$ in (A.8). These variations do not affect the map f , the complex structure on the domain or the marked points, and hence do not change the domain and range of the operators (A.5). Because D_p and D_q^ℓ are differential operators with the same formula, the variation $(\delta_{v_\varepsilon} D_q^\ell)_q$ is again given by (5.5) with K replaced by $2\beta_\varepsilon K$.

After substituting and taking the limit $\varepsilon \rightarrow 0$ as in the proof of Proposition 5.1c, equation (A.8) implies that

$$0 = \sum_i \langle (A_q \kappa_i)(x), (\nabla_{\kappa_i^N(x)} K) f_* j \rangle = |c(x)|^2, \quad (\text{A.10})$$

where the last equality holds because $(\nabla_{\kappa_i^N(x)} K) f_* j = 0$ for all $i \geq 2$ by (A.11). This contradicts the fact that $c(x) \neq 0$, and hence establishes the transversality of Ψ^ℓ at $q \in S^{\ell,s}$. The restriction of Ψ to $\pi_\ell^{-1}(\mathcal{M}_C)$ is also transverse to $\text{Fred}_{\iota_\ell}^s$ because for each embedding $q \in S^{\ell,s}$ with image C , the variations v_ε above are tangent to $\pi_\ell^{-1}(\mathcal{M}_C)$. Thus $S^{\ell,s}$ and $S_C^{\ell,s}$ are manifolds whose codimension, in both cases, is the dimension of the normal space to $\text{Fred}_{-4\ell}^s$, which is $(\dim \ker D_q^\ell)(\dim \text{coker } D_q^\ell) = s(s - \iota_\ell)$. \square

Lemma A.6. Fix $p = (f, J) \in \mathcal{M}_{\text{simple}}$, an injective point $x \in C$, and a neighborhood U of $f(x)$. For any nonzero $\xi \in N_{f(x)}$, any subspace $V \subseteq N_{f(x)}$ not containing ξ , and any $c(x) \in T_x^*C \otimes_{\mathbb{C}} T_{f(x)}X$, there exists a $K \in T_J\mathcal{J}$, supported on U , vanishing along $f(C)$ such that, at the single point x :

$$(\nabla_\xi K)f_*j = c \quad \text{and} \quad (\nabla_w K)f_*j = 0 \quad \forall w \in V. \quad (\text{A.11})$$

Proof. As in Proposition 5.1a, there is a $K_0 \in T_J\mathcal{J}$ such that $K_0 f_*j = c$ at x . Choose a local coordinate system $\{z, y_1, y_2, \dots\}$ centered at $f(x)$ with z a local complex coordinate on $f(C)$, and $\{y_i\}$ real coordinates with $\frac{\partial}{\partial y_1}|_{f(x)} = \xi$ and $\frac{\partial}{\partial y_k}|_{f(x)} \in V$ for $2 \leq k \leq \dim V + 1$. Then $K = y_1 \beta K_0$ has the required properties where β is any smooth function supported in U with $\beta \equiv 1$ near the origin. \square

Proof of Proposition A.4. We begin by making a series of observations about the images of the sets V^ℓ and $S^{\ell,s}$ under the forgetful map (A.4).

- (i) *The images of the V^ℓ cover $\mathcal{M}_{\text{simple}}$.* If not, there would be a map $p \in \mathcal{M}_{\text{simple}}$ not in the image of any V^ℓ . Choose a dense sequence $\{x_1, x_2, \dots\}$ of distinct injective points in the domain C of p . Then for each ℓ , $q_\ell = (p, x_1, \dots, x_\ell) \notin V^\ell$, which implies that $\ker D_{q_{\ell+1}}^{\ell+1} \neq 0$. But then $\ker D_{q_{\ell+1}}^{\ell+1} \subseteq \ker D_{q_\ell}^\ell$ are nontrivial nested subspaces of the finite-dimensional vector space $\ker D_p$, so have a nonzero intersection. Hence there is a nonzero $\kappa \in \ker D_p$ whose normal component vanishes at all x_i , and therefore everywhere, contradicting Lemma 4.1.
- (ii) *The images of the $S^{\ell,s}$ with $s \geq 1$ cover \mathcal{W} .* Given $p \in \mathcal{W}$, we have $\ker D_p \neq 0$. As in (i), there is a sequence $\{x_i\} \subset C$ and an $m > 0$ such that $\ker D_{q_m}^m = 0$. Let ℓ be the largest k such that $\ker D_{q_k}^k \neq 0$. Then by (A.7), $q_\ell \in S^{\ell,s}$ for $s = \dim \ker D_{q_\ell}^\ell \geq 1$, and hence $p \in \pi_\ell(S^{\ell,s})$.
- (iii) *$S^{0,1} = \mathcal{W}^1$ is a submanifold of $\mathcal{M}_{\text{simple}}$.* Equation (A.7) shows that $S^{0,1} \subseteq \mathcal{W}^1$, while Lemma 4.1 implies $\mathcal{W}^1 \subseteq S^{0,1}$. Hence, by Lemma A.5, \mathcal{W}^1 is a submanifold of $\mathcal{M}_{\text{simple}}$ of codimension $1 - \iota_0$. In particular, for each component \mathcal{N} of $\mathcal{M}_{\text{simple}}$ with index $D_p = 0$, we have $\iota_0 = 0$ by (A.6), so the restriction $\mathcal{W}^1 \cap \mathcal{N}$ is codimension 1 submanifold of \mathcal{N} .
- (iv) *$\pi_\ell : S^{\ell,s} \rightarrow \mathcal{M}_{\text{simple}}$ is a Fredholm map of index $2\ell + s(\iota_\ell - s)$.* This map is the composition of the inclusion $S^{\ell,s} \rightarrow \mathcal{M}_{\ell,\text{simple}}$, which has index $s(\iota_\ell - s)$ by Lemma A.5, and the map (A.4) which has index 2ℓ .

By Facts (ii) and (iii), $\mathcal{W} \setminus \mathcal{W}^1$ is covered by the sets $\pi_\ell(S^{\ell,s})$ for $\ell \geq 0$, $s \geq 1$, and $(\ell, s) \neq (0, 1)$. By Fact (iv) and (A.6), the intersection of each of these sets with \mathcal{N} is the image of a Fredholm map of index $2\ell - \ell s(\dim X - 2) - s^2 \leq 2\ell - 4\ell s - s^2 \leq -3$. Thus $(\mathcal{W} \setminus \mathcal{W}^1) \cap \mathcal{N}$ is a set of codimension 3 in the sense of Definition 5.2.

The same proof applies if we restrict everything to \mathcal{M}_C instead of $\mathcal{M}_{\text{simple}}$. \square

REFERENCES

- [BP1] J. Bryan and R. Pandharipande, *BPS states of curves in Calabi-Yau 3-folds*, Geometry & Topology **5**(2001), 287–318.
- [BP2] _____, *The local Gromov-Witten theory of curves*, J. Amer. Math. Soc. **21**(2008), 101–136.
- [FP] K. Faber and R. Pandharipande, *Hodge integrals and Gromov-Witten theory*, Invent. Math. **139**(2000), 173–199.
- [GV] R. Gopakumar and C. Vafa, *M-Theory and Topological Strings I and II*, hep-th/9809187 and hep-th/9812127.
- [IP1] E.N. Ionel and T.H. Parker, *The Gromov invariants of Ruan-Tian and Taubes*, Math. Res. Lett. **4**(1997), 521–532.
- [IP2] _____, *Relative Gromov-Witten Invariants*, Annals of Math. **157**(2003), 45–96.
- [IS1] S. Ivashkovich and V. Shevchishin, *Pseudo-holomorphic curves and envelopes of meromorphy of two-spheres in \mathbf{CP}^2* , arXiv:math.CV/9804014.
- [IS2] S. Ivashkovich and V. Shevchishin, *Structure of the moduli space in a neighborhood of a cusp-curve and meromorphic hulls*, Invent. Math. **136**, (1999) 571–602.

- [KP] A. Klemm and R. Pandharipande, *Enumerative geometry of Calabi-Yau 4-folds*, Commun. Math. Phys. **281**(2008), 621-653.
- [K] U. Koschorke, *Infinite dimensional K-theory and characteristic classes of Fredholm bundle maps*, in Proceedings of the Symposia in Pure Mathematics, Vol. XV, pp 95-133. AMS, Providence, R.I. 1970.
- [L] J. Lee, *Holomorphic 2-forms and vanishing theorems for Gromov-Witten invariants*, Canad. Math. Bull. **52**(2009), 87-94.
- [LP1] J. Lee and T.H. Parker, *A Structure Theorem for the Gromov-Witten Invariants of Kähler Surfaces*. J. Diff. Geom. **77** (2007), 483-513.
- [LP2] J. Lee and T.H. Parker, *An obstruction bundle relating Gromov-Witten invariants of curves and Kähler surfaces*, American J. of Math. **134** (2012), 453-506.
- [MOOP] D. Maulik, A. Oblomkov, A. Okounkov, R. Pandharipande, *Gromov-Witten/Donaldson-Thomas correspondence for toric 3-folds*, Invent. Math. **186**(2011), 435-479.
- [MW] M. Micallef and B. White, *The structure of branch points in minimal surfaces and in pseudo holomorphic curves*, Ann. of Math. **141**(1995), 35-85.
- [MS] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology*, second ed., A.M.S. Colloquium Publications, **52**, AMS, Providence, RI, 2012.
- [P] R. Pandharipande, *Hodge integrals and degenerate contributions*, Commun. Math. Phys. **208**(1999), 489-506.
- [PP] R. Pandharipande and A. Pixton, *Gromov-Witten/Pairs correspondence for the quintic 3-fold*, arXiv:1206.5490.
- [PT] R. Pandharipande and R. P. Thomas, *Curve counting via stable pairs in the derived category*, Invent Math. **178**(2009), 407-447.
- [P] J. Pardon, *An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves*, preprint, arXiv:1508.03873, Geom. Topol. (to appear).
- [S] S. Smale, *An Infinite Dimensional Version of Sard's Theorem*, Amer. J. Math. **87**(1965), 861-866.
- [T] C. H. Taubes, *Counting pseudo-holomorphic curves in dimension four*, J. Diff. Geom. **44**(1996), 818-893.
- [Z] A. Zinger, *A comparison theorem for Gromov-Witten invariants in the symplectic category*, Adv. Math. **228**(2011), 535-574.

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY
E-mail address: ionel@math.stanford.edu

DEPARTMENT OF MATHEMATICS, MICHIGAN STATE UNIVERSITY
E-mail address: parker@math.msu.edu