

CHAPTER 1

Introduction

The theory of J -holomorphic curves, introduced by Gromov in 1985, has profoundly influenced the study of symplectic geometry, and now permeates almost all its aspects. The methods are also of interest in the study of Kähler manifolds, since often when one studies properties of holomorphic curves in such manifolds it is useful to perturb the complex structure to be generic. The effect of this is to ensure that one is looking at persistent rather than accidental features of these curves. However, the perturbed structure may no longer be integrable, and so again one is led to the study of curves that are holomorphic with respect to some nonintegrable almost complex structure J .

These curves satisfy a nonlinear analogue of the Cauchy–Riemann equations. Before one can do anything useful with them, one must understand the elements of the theory of these equations; for example, know what conditions ensure that the solution spaces are finite dimensional manifolds and know ways of dealing with the fact that these solution spaces are usually noncompact. As explained in the preface, this chapter introduces all the basic concepts and outlines the ingredients needed to establish these results. Readers, specially those unfamiliar with the theory, should start by reading this chapter and then branch out into more detailed study of whichever aspects of the theory interests them.

We assume that the reader is familiar with the elements of symplectic geometry. There are several introductory books. One possible reference is McDuff–Salamon [277], but there are now more elementary treatments such as Berndt [35] and Cannas da Silva [59] as well as classics such as Arnold [15].

1.1. Symplectic manifolds

A symplectic structure on a smooth $2n$ -dimensional manifold M is a closed 2-form ω which is nondegenerate in the sense that the top-dimensional form ω^n does not vanish anywhere. By Darboux’s theorem, all symplectic forms are locally diffeomorphic to the standard symplectic form

$$\omega_0 = dx_1 \wedge dy_1 + \cdots + dx_n \wedge dy_n$$

on Euclidean space \mathbb{R}^{2n} with coordinates $(x_1, y_1, \dots, x_n, y_n)$. This makes it hard to get a handle on the global structure of symplectic manifolds. Variational techniques have been developed which allow one to investigate some global questions in Euclidean space and in manifolds such as cotangent bundles which have some linear structure: see [190, 277] and the references contained therein. But the method which applies to the widest variety of symplectic manifolds is that of J -holomorphic curves.

Here J is an almost complex structure on M which is tamed by ω . An almost complex structure is an automorphism J of the tangent bundle TM of M which

satisfies the identity $J^2 = -\mathbb{1}$. Thus J can be thought of as multiplication by i , and it makes TM into a complex vector bundle of dimension n . The form ω is said to **tame** J if

$$\omega(v, Jv) > 0$$

for all nonzero $v \in TM$. Geometrically, this means that ω restricts to a positive form on each complex line $L = \text{span}\{v, Jv\}$ in the tangent space T_xM . Given ω the set $\mathcal{J}_\tau(M, \omega)$ of almost complex structures tamed by ω is always nonempty and contractible. Note that it is very easy to construct and perturb tame almost complex structures, because they are defined by pointwise conditions. Note also that, because $\mathcal{J}_\tau(M, \omega)$ is path connected, different choices of $J \in \mathcal{J}_\tau(M, \omega)$ give rise to isomorphic complex vector bundles (TM, J) . Thus the Chern classes of these bundles are independent of the choice of J and will be denoted by $c_i(TM)$.¹

In what follows we shall only need to use the first Chern class, and what will be relevant is the value $c_1(A) := \langle c_1(TM), A \rangle$ which it takes on a homology class $A \in H_2(M)$. If A is represented by a smooth map $u : \Sigma \rightarrow M$, defined on a closed oriented 2-manifold Σ then $c_1(A) = c_1(E)$ is the first Chern number of the pullback tangent bundle $E := u^*TM \rightarrow \Sigma$. But every complex bundle E over a 2-manifold Σ decomposes as a sum of complex line bundles $E = L_1 \oplus \cdots \oplus L_n$. Correspondingly

$$c_1(E) = \sum_i c_1(L_i).$$

Since the first Chern number of a complex line bundle is just the same as its Euler number, it is often easy to calculate the $c_1(L_i)$ directly. For example, if A is the class of the sphere $S = \text{pt} \times S^2$ in $M = V \times S^2$ then it is easy to see that

$$TM|_S = TS \oplus L_2 \oplus \cdots \oplus L_n,$$

where the line bundles L_k are trivial. It follows that

$$c_1(A) = c_1(TM|_S) = c_1(TS) = \chi(S) = 2$$

where $\chi(S)$ is the Euler characteristic of S .

A smooth map $\phi : (M, J) \rightarrow (M', J')$ from one almost complex manifold to another is said to be (J, J') -**holomorphic** if and only if its derivative $d\phi(x) : T_xM \rightarrow T_{\phi(x)}M'$ is complex linear, that is

$$d\phi(x) \circ J(x) = J'(\phi(x)) \circ d\phi(x).$$

These are the Cauchy–Riemann equations, and, when (M, J) and (M', J') are both subsets of complex n -space \mathbb{C}^n , they are satisfied precisely by the holomorphic maps. An almost complex structure J is said to be **integrable** if it arises from an underlying complex structure on M . This is equivalent to saying that one can choose an atlas for M whose coordinate charts are (J, i) -holomorphic where i is the standard complex structure on \mathbb{C}^n . In this case the coordinate changes are holomorphic maps (in the usual sense) between open sets in \mathbb{C}^n . When M has dimension 2 a fundamental theorem says that all almost complex structures J on M are integrable: for a proof see Section E.4. However this is far from true in higher dimensions.

¹There is another space of almost complex structures naturally associated to (M, ω) , namely the set $\mathcal{J}(M, \omega)$ of ω -**compatible structures** defined in Section 3.1. For the present purposes one can use either space. In fact, to make our results as general as possible, we will often work with $\mathcal{J}(M, \omega)$ because this is very slightly harder: $\mathcal{J}_\tau(M, \omega)$ is open in the space of all almost complex structures, while $\mathcal{J}(M, \omega)$ is not.

The basic example of an almost complex symplectic manifold is standard Euclidean space $(\mathbb{R}^{2n}, \omega_0)$ with its standard almost complex structure J_0 obtained from the usual identification with \mathbb{C}^n via

$$z_j = x_j + iy_j$$

for $j = 1, \dots, n$. Thus J_0 maps the tangent vector $\partial/\partial x_j$ to $\partial/\partial y_j$ and maps $\partial/\partial y_j$ to $-\partial/\partial x_j$ in the standard basis of $\mathbb{R}^{2n} = T_z \mathbb{R}^{2n}$. In other words, the automorphism $J_0 : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$ is represented by the $2n \times 2n$ -matrix

$$J_0 := \begin{pmatrix} 0 & -1 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & -1 \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}.$$

Every Kähler manifold gives another example (but not every symplectic manifold admits a Kähler structure).

J -holomorphic curves. A J -holomorphic curve is a (j, J) -holomorphic map

$$u : \Sigma \rightarrow M$$

from a Riemann surface² (Σ, j) to an almost complex manifold (M, J) . Usually, we will take (Σ, j) to be the Riemann sphere $S^2 = \mathbb{C} \cup \{\infty\}$. In accordance with the terminology of complex geometry it is often convenient to think of the 2-sphere as the complex projective line $\mathbb{C}P^1$.

If u is an embedding (that is, an injective immersion) then its image C is a 2-dimensional submanifold of M whose tangent spaces $T_x C$ are J -invariant. Thus each tangent space is a complex line in TM . Further, by the taming condition, ω restricts to a positive form on each such line. Therefore C is a symplectic submanifold of M .³ Conversely, if $C \subset M$ is a 2-dimensional symplectic submanifold then there is an ω -tame almost complex structure J such that TC is J -invariant. (First define J on TC , then extend to the tangent spaces $T_x M$ for $x \in C$, and finally extend the section to the rest of M .) Since the restriction of J to TC is integrable, it follows that C is the image of a J -holomorphic curve. Thus embedded J -holomorphic curves are essentially the same as 2-dimensional symplectic submanifolds of M . This argument uses the 2-dimensionality of C in an essential way; it is not possible to construct a good theory of J -holomorphic maps from higher dimensional manifolds into (M, J) . (The symplectic vortex equations mentioned in Section 12.7 do not contradict this statement, since they have extra structure.)

Observe that according to this definition, a curve u is always parametrized. One should contrast this with the situation in complex geometry, where one often defines a curve by requiring it to be the common zero set of a certain number of holomorphic polynomials. Such an approach makes no sense in the nonintegrable, almost complex world, since when J is nonintegrable there usually are no holomorphic functions $(M, J) \rightarrow \mathbb{C}$.⁴

²A Riemann surface is a 1-dimensional complex manifold.

³A submanifold X of M is said to be symplectic if ω restricts to a nondegenerate form on X .

⁴However Donaldson has recently developed a powerful theory of *asymptotically holomorphic* functions $(M, J, \omega) \rightarrow \mathbb{C}$: see Donaldson [84, 85] and Auroux [26].

1.2. Moduli spaces: regularity and compactness

The crucial fact about J -holomorphic curves is that when J is generic they occur in finite dimensional families. These families make up finite dimensional manifolds

$$\mathcal{M}^*(A; J)$$

called **moduli spaces**, whose cobordism classes are independent of the particular J chosen in $\mathcal{J}_\tau(M, \omega)$. Here A is a homology class in $H_2(M, \mathbb{Z})$, and $\mathcal{M}^*(A; J)$ consists of essentially all J -holomorphic curves $u : \Sigma \rightarrow M$ that represent the class A . Although the manifold $\mathcal{M}^*(A; J)$ is almost never compact, it usually retains enough elements of compactness for one to be able to use it to define invariants.

Chapters 2–5 of this book are taken up with formulating and proving precise versions of the above statements. The results divide naturally into two groups, the first (discussed in Chapters 2 and 3) concerning issues of regularity and transversality and the second (discussed in Chapters 4 and 5) concerning the compactification of $\mathcal{M}^*(A; J)$. In the next paragraphs we pick out the most important concepts and theorems. For those new to the theory, these results together with their proofs are the ones to focus on first. The introductions to each chapter give further guidance on how to read the book.

Local properties. Chapter 2 concerns the local properties of J -holomorphic curves. The key results for future developments are the energy identity, which is crucial for later results on compactness, and Proposition 2.5.1, which gives a characterization of those curves which are not multiply covered. A curve $u : \Sigma \rightarrow M$ is said to be **multiply covered** if it is the composite of a holomorphic branched covering map $(\Sigma, j) \rightarrow (\Sigma', j')$ of degree greater than 1 with a J -holomorphic map $\Sigma' \rightarrow M$. It is called **simple** if it is not multiply covered. The multiply covered curves are often singular points in the moduli space $\mathcal{M}(A; J)$ of all J -holomorphic A -curves, and so one needs a workable criterion which guarantees that u is simple. We will say that a curve u is **somewhere injective** if there is a point $z \in \Sigma$ such that

$$du(z) \neq 0, \quad u^{-1}(u(z)) = \{z\}.$$

A point $z \in \Sigma$ with this property is called an **injective point** for u . Here is the statement of Proposition 2.5.1. It is an important ingredient of the proof of Theorem 3.1.6 which asserts that the moduli space $\mathcal{M}^*(A; J)$ of simple J -holomorphic A -curves is a smooth manifold for a generic ω -tame J .

PROPOSITION. *Every simple J -holomorphic curve u is somewhere injective. Moreover the set of injective points is open and dense in Σ .*

The proof is elementary, except for an appeal to a deep analytic theorem extending the unique continuation principle to J -holomorphic curves.

Fredholm theory. Fix a Riemann surface Σ of genus g and let $\mathcal{M}^*(A; J)$ denote the set of all *simple* J -holomorphic maps $u : \Sigma \rightarrow M$ which represent the class A .⁵ Here is an informal version of the main result of Chapter 3. Recall that a subset in a complete metric space is said to be residual (in the sense of Baire) if it

⁵In the first edition of this book we denoted this space simply by $\mathcal{M}(A; J)$. However, it is now customary to use the letter \mathcal{M} for full moduli spaces.

contains a countable intersection of open and dense sets. Baire's category theorem asserts that every such set is dense.

THEOREM A. *There is a residual subset $\mathcal{J}_{\text{reg}}(A) \subset \mathcal{J}_\tau(M, \omega)$ such that for each $J \in \mathcal{J}_{\text{reg}}(A)$ the space $\mathcal{M}^*(A; J)$ is a smooth manifold of dimension*

$$\dim \mathcal{M}^*(A; J) = n(2 - 2g) + 2c_1(A).$$

This manifold $\mathcal{M}^(A; J)$ carries a natural orientation.*

Another important theorem specifies the dependence of $\mathcal{M}^*(A; J)$ on the choice of J (Theorem 3.1.8). The basic reason why these theorems are valid is that the Cauchy–Riemann equation

$$du \circ j = J \circ du$$

is elliptic, and hence its linearization at a J -holomorphic curve is a Fredholm operator, denoted by D_u . The set $\mathcal{J}_{\text{reg}}(A)$ in Theorem A is the set of all those almost complex structures $J \in \mathcal{J}_\tau(M, \omega)$ such that the linearized operator D_u is onto for every $u \in \mathcal{M}^*(A; J)$. The elements of $\mathcal{J}_{\text{reg}}(A)$ are sometimes called **regular** almost complex structures. An interesting fact is that the taming condition on J is irrelevant here; the above results hold for *all* almost complex structures J on *any* compact manifold M .

A bounded linear operator $D : X \rightarrow Y$ between Banach spaces is said to be **Fredholm** if it has a finite dimensional kernel and a closed image of finite codimension in Y . The **index** of D is defined to be the difference in dimension between the kernel and cokernel of D :

$$\text{index } D = \dim \ker D - \dim \text{coker } D.$$

An important fact is that the set of Fredholm operators is open with respect to the norm topology and the Fredholm index is constant on each component of the set of Fredholm operators. Thus it does not change as D varies continuously, though of course the dimension of the kernel and cokernel can change.

As we show in Appendix A, Fredholm operators are essentially as well behaved as finite dimensional operators and they play an important role in infinite dimensional implicit function theorems. More precisely, assume that $\mathcal{F} : X \rightarrow Y$ is a C^∞ -smooth map between separable Banach spaces whose derivative $d\mathcal{F}(x) : X \rightarrow Y$ is Fredholm of index k at each point $x \in X$. If $y \in Y$ is a **regular value** of \mathcal{F} in the sense that $d\mathcal{F}(x)$ is surjective for all $x \in \mathcal{F}^{-1}(y)$, then, just as in the finite dimensional case, the inverse image

$$\mathcal{M} := \mathcal{F}^{-1}(y)$$

is a smooth manifold of dimension k . An infinite dimensional version of Sard's theorem says that almost all points of Y are regular values of \mathcal{F} . (Technically, the regular points form a residual set.) This theorem remains true if X and Y are separable Banach manifolds rather than Banach spaces. However it does not extend as stated to other kinds of infinite dimensional vector spaces, such as Fréchet spaces. Therefore, although the set \mathcal{J}_{reg} mentioned in Theorem A above does consist of the regular values of a Fredholm operator with target $\mathcal{J}_\tau(M, \omega)$, there are some additional technicalities in the proof because $\mathcal{J}_\tau(M, \omega)$ is a Fréchet rather than a Banach manifold.

J -holomorphic discs. Although most of the results in this book concern the properties of J -holomorphic spheres, we formulate and prove the foundational results in chapters 3 and 4 as well as appendices B and C for general compact Riemann surfaces with (possibly empty) boundary and thus include the case of J -holomorphic discs; we show in Chapter 3 that with Lagrangian boundary conditions the delbar operator is again Fredholm. Thus moduli spaces of J -holomorphic discs with boundary in a given Lagrangian submanifold again form finite dimensional manifolds for generic J . This fact is an important ingredient of the inductive proof of the Riemann–Roch theorem given in Appendix C. We give several applications of this theory in Section 9.2, notably the proof that there is no closed exact Lagrangian submanifold in \mathbb{R}^{2n} . The fact that Lagrangian boundary conditions are elliptic is also very important in Floer theory: see Chapter 12.

This is the first step in extending the theory to maps with more general domains. As a next step one could consider either closed domains of higher genus or noncompact domains of genus zero. Our work here is sufficient to understand closed domains (Σ, j) of higher genus with fixed complex structure j . Varying j does not significantly affect the Fredholm theory but does cause additional problems when discussing compactness. On the other hand, one needs considerably more sophisticated analysis in order to set up Fredholm theory for noncompact domains. (In this case one often studies perturbed equations: cf. the development of Floer theory in [116].) Both of these extensions are beyond the scope of the present book. They are also both needed to understand the full structure of Gromov–Witten invariants as displayed for example in Eliashberg, Givental and Hofer’s recent paper [101] on symplectic field theory.

Compactness. The next task is to develop an understanding of when the moduli spaces $\mathcal{M}^*(A; J)$ are compact. Here the taming condition plays an essential role. The symplectic form ω and an ω -tame almost complex structure J together determine a Riemannian metric

$$\langle v, w \rangle = \frac{1}{2}(\omega(v, Jw) + \omega(w, Jv))$$

on M and the **energy** of a J -holomorphic curve $u : \Sigma \rightarrow M$ with respect to this metric is given by the formula

$$E(u) = \frac{1}{2} \int_{\Sigma} |du(z)|_J^2 \, \text{dvol}_{\Sigma} = \int_{\Sigma} u^* \omega.$$

Thus the L^2 -norm of the derivative of a J -holomorphic curve satisfies a uniform bound which depends only on the homology class A represented by u . This in itself does not guarantee compactness because it is a borderline case for Sobolev estimates. (A uniform bound on the L^p -norms of du with $p > 2$ would guarantee compactness.)

Another manifestation of the failure of compactness can be observed in the fact that the energy $E(u)$ is invariant under conformal rescaling of the metric on Σ . This effect is particularly clear in the case where the domain Σ of our curves is the Riemann sphere $\mathbb{C}P^1$, since here there is a large group of global, rather than local, rescalings. Indeed, the noncompact group $G = \text{PSL}(2, \mathbb{C})$ acts on the Riemann sphere by conformal transformations

$$z \mapsto \frac{az + b}{cz + d}.$$

Thus each element $u \in \mathcal{M}^*(A; J)$ has a noncompact family of reparametrizations $u \circ \phi$, for $\phi \in G$, and so $\mathcal{M}^*(A; J)$ itself can never be both compact and nonempty (unless A is the zero class, in which case all the maps u are constant). However, the quotient space $\mathcal{M}^*(A; J)/G$ will sometimes be compact.

Recall that a homology class $B \in H_2(M)$ is called **spherical** if it is in the image of the Hurewicz homomorphism $\pi_2(M) \rightarrow H_2(M)$. Here is a statement of the first important theorem in Chapter 4. Observe that the hypothesis on A implies that no A -curve can be multiply covered, that is $\mathcal{M}^*(A; J) = \mathcal{M}(A; J)$.

THEOREM B. *Assume that there is no spherical homology class $B \in H_2(M)$ such that $0 < \omega(B) < \omega(A)$. Then the moduli space $\mathcal{M}^*(A; J)/G$ is compact.*

To prove this one shows that if u_ν is a sequence in $\mathcal{M}^*(A; J)$ which has no limit point in $\mathcal{M}^*(A; J)$ then there is a point $z \in \mathbb{C}P^1$ at which the derivatives $du_\nu(z)$ are unbounded. After passage to a subsequence one finds a decreasing sequence U_ν of neighbourhoods of z in $\mathbb{C}P^1$ whose images $u_\nu(U_\nu)$ converge in the limit to a nonconstant J -holomorphic sphere. If B is the homology class represented by this sphere, then either $\omega(B) = \omega(A)$, in which case the maps u_ν can be reparametrized so that they do converge in $\mathcal{M}^*(A; J)$, or $\omega(B)$ lies strictly between 0 and $\omega(A)$. This is the process of “bubbling”, which occurs in this context in a simple and geometrically clear way: see Section 4.2.

Theorem B implies that if $\omega(A)$ is already the smallest positive value assumed by ω on spheres then the moduli space $\mathcal{M}^*(A; J)/G$ is compact. As we shall see below, this is enough for some interesting applications. However, to cope with more general classes A we need to understand the complete limit of the sequence u_ν . The analysis required to solve this challenging problem is developed in the second half of Chapter 4. In particular, in order to see that the limiting curve is connected, we study the properties of long cylinders of small energy in Section 4.7.

The full structure of the limit is formulated in Theorem 5.3.1, a result known as Gromov’s compactness theorem. We give a more modern formulation than Gromov’s original statement in [160], using Kontsevich’s language of **stable maps**. Although not strictly necessary for the applications in this book, this language allows us to set up the Gromov–Witten invariants in their natural context and is the basis of many recent applications of the theory especially those involving quantum cohomology. Therefore, we develop it in considerable detail in Chapter 5 and Appendix D.

The best strategy for those new to the subject might be to try to understand the definition of Gromov convergence (Definition 5.2.1) and the statement of Theorem 5.3.1, leaving further details for later. Proposition 4.1.5 is also important; it says that for each J and each $c > 0$ there are only finitely many classes A with J -holomorphic representatives and energy $\omega(A)$ bounded by c .

1.3. Evaluation maps and pseudocycles

The Gromov–Witten invariants are built from the evaluation map

$$\mathcal{M}^*(A; J) \times \mathbb{C}P^1 \rightarrow M : (u, z) \mapsto u(z).$$

Note that this factors through the action of the reparametrization group G given by

$$\phi \cdot (u, z) = (u \circ \phi^{-1}, \phi(z)).$$

Hence we get a map defined on the quotient

$$\text{ev} = \text{ev}_J : \mathcal{M}^*(A; J) \times_{\mathbb{G}} \mathbb{C}P^1 \rightarrow M.$$

EXAMPLE. Let V be a closed symplectic manifold of dimension $2n - 2$ and

$$M := \mathbb{C}P^1 \times V$$

with the product symplectic form. Suppose that $\pi_2(V) = 0$. Then $A := [\mathbb{C}P^1 \times \text{pt}]$ generates the group of spherical 2-classes in M , and so $\omega(A)$ is necessarily the smallest value assumed by ω on the spherical classes. Theorems A and B therefore imply that the space $\mathcal{M}^*(A; J) \times_{\mathbb{G}} \mathbb{C}P^1$ is a compact manifold for generic J . Because $c_1(A) = 2$, in this case the dimension of $\mathcal{M}^*(A; J) \times_{\mathbb{G}} \mathbb{C}P^1$ is $2n$ and agrees with the dimension of M . Moreover, we will see in Chapter 3 (see Theorem 3.1.8) that different choices of J give rise to cobordant maps ev_J . Since cobordant maps have the same degree, this means that the degree of ev_J is independent of all choices. Now if $J = i \times J'$ is a product, where i denotes the standard complex structure on $\mathbb{C}P^1$, then it is easy to see that the elements of $\mathcal{M}^*(A; J)$ have the form

$$u(z) = (\phi(z), v_0)$$

where $v_0 \in V$ and $\phi \in \mathbb{G}$. It follows that the map ev_J has degree 1 for this choice of J and hence for every regular J .

Gromov used this fact in [160] to prove his celebrated nonsqueezing theorem.

THEOREM. *Let V be a closed symplectic manifold of dimension $2n - 2$ such that $\pi_2(V) = 0$. If ψ is a symplectic embedding of the ball $B^{2n}(r)$ of radius r into the cylinder $B^2(\lambda) \times V$ then $r \leq \lambda$.*

SKETCH OF PROOF. Embed the disc $B^2(\lambda)$ into a 2-sphere $\mathbb{C}P^1$ of area $\pi\lambda^2 + \varepsilon$, and let ω be the product symplectic structure on $\mathbb{C}P^1 \times V$. Let J' be an ω -tame almost complex structure on $\mathbb{C}P^1 \times V$ which, on the image of ψ , equals the push-forward by ψ of the standard structure J_0 of the ball $B^{2n}(r)$. Since the evaluation map $\text{ev}_{J'}$ has degree 1, there is a J' -holomorphic curve through every point of $\mathbb{C}P^1 \times V$. In particular, there is such a curve, C' say, through the image $\psi(0)$ of the center of the ball. This curve pulls back by ψ to a J_0 -holomorphic curve C through the center of the ball $B^{2n}(r)$. Since J_0 is standard, this curve C is holomorphic in the usual sense and so is a minimal surface in $B^{2n}(r)$. But it is a standard result in the theory of minimal surfaces that a properly embedded surface through the center of a ball in Euclidean space has area at least that of the flat disc, namely πr^2 . Further, because C is holomorphic, it is easy to check that its area is just given by the integral of the standard symplectic form ω_0 over it. Thus

$$\pi r^2 \leq \int_C \omega_0 = \int_{\psi^{-1}(C')} \psi^*(\omega) < \int_{C'} \omega = \omega(A) = \pi\lambda^2 + \varepsilon$$

where the middle inequality holds because $\psi(C)$ is only a part of C' . Since this is true for all $\varepsilon > 0$, the result follows. \square

More details of the above argument may be found in [160, 320]. In Section 9.3 we shall give full details of a slightly different proof that replaces the appeal to the theory of minimal surfaces by using the symplectic blow up. Chapter 9 also contains complete proofs of some of the other foundational results stated in Gromov's paper [160]. For example, in Section 9.2 we use the theory of J -holomorphic

discs developed in Chapter 3 to show that there is no closed exact Lagrangian submanifold in \mathbb{R}^{2n} when $n \geq 2$.

The Gromov–Witten pseudocycle. Often it is useful to evaluate the map u at more than one point. If one is evaluating at up to three points then, because G acts triply transitively on the sphere, it does not matter which points one chooses. However, if one is evaluating at more than three points (or if the domain is not a sphere) the choice of points does make a difference. One incorporates these points into the moduli space itself, hence making the basic object of study a tuple (u, z_1, \dots, z_k) consisting of an element $u \in \mathcal{M}(A; J)$ together with k pairwise distinct **marked points** $z_i \in S^2$. The space of all such tuples is denoted $\widetilde{\mathcal{M}}_{0,k}(A; J)$ and its quotient by G is denoted $\mathcal{M}_{0,k}(A; J)$.⁶ As before $\mathcal{M}_{0,k}^*(A; J)$ denotes the subset for which u is simple. It is a manifold of dimension

$$\dim \mathcal{M}_{0,k}^*(A; J) = 2n + 2c_1(A) + 2k - 6 =: \mu(A, k)$$

for generic J . Moreover, the obvious evaluation map

$$\widetilde{\mathcal{M}}_{0,k}(A; J) \rightarrow M^k : (u, z_1, \dots, z_k) \mapsto (u(z_1), \dots, u(z_k))$$

descends to

$$\text{ev} : \mathcal{M}_{0,k}(A; J) \rightarrow M^k.$$

If this map represents a cycle one can define the Gromov–Witten invariants as its intersection numbers with cycles in M^k .

Observe that even though $\mathcal{M}_{0,k}(A; J)$ consists of unparametrized maps it is never compact when $k \geq 2$ since the marked points are distinct by definition. Nevertheless, we would like to define the Gromov–Witten invariants using the d -dimensional homology class represented by the evaluation map

$$\text{ev} : \mathcal{M}_{0,k}^*(A; J) \rightarrow M^k,$$

where $d = \mu(A, k)$. Therefore, even if we know that $\mathcal{M}_{0,k}(A; J)$ is a manifold, we must work to show that it represents a cycle. The crucial point here is to understand its compactification. Kontsevich realised that the compactification of $\mathcal{M}_{0,k}(A; J)$ with the best structure is not the space of cusp-curves considered by Gromov but rather the space $\overline{\mathcal{M}}_{0,k}(A; J)$ of **stable maps**: see Definition 5.1.1. We formulate and prove Gromov’s compactness theorem in Chapter 5 using this language.

Even with a good understanding of the compactification $\overline{\mathcal{M}}_{0,k}(A; J)$ it is very hard to understand the sense in which ev represents a cycle for general manifolds (M, ω) , since multiply covered curves in lower strata can cause these to have too high a dimension.⁷ However, if (M, ω) is **semipositive** then we show in Chapter 6 that the evaluation map represents a **pseudocycle** of dimension $d := \mu(A, k)$. Intuitively speaking this means that its image can be compactified by adding a set of dimension at most $d - 2$.⁸ Thus its boundary has dimension at most $d - 2$ and so

⁶In this notation the 0 in the subscript denotes the genus of the domain. In this book we restrict to the genus 0 case, though much of what we say applies almost word for word in the case when (Σ, j) has arbitrary genus and fixed complex structure j .

⁷This problem has now been solved for general manifolds. By using delicate global perturbations one defines a d -dimensional *virtual moduli cycle*, which should be thought of as the fundamental class of $\mathcal{M}_{0,k}^*(A; J)$. Detailed references may be found in the introduction to Chapter 6. This is beyond the scope of the present book.

⁸A subset B of a smooth manifold X is said to be of **dimension at most** m if it is contained in the image of a smooth map $g : W \rightarrow X$ defined on a manifold W of dimension m .

is not seen from homological point of view. More formally we make the following definition.

DEFINITION. A d -dimensional **pseudocycle** in a manifold X is a smooth map

$$f : V \rightarrow X$$

defined on an oriented d -dimensional manifold V such that its image $f(V)$ has compact closure and its limit set Ω_f has dimension $\leq d - 2$. Here

$$\Omega_f = \bigcap_{\substack{K \subset V \\ K \text{ compact}}} \overline{f(V \setminus K)}$$

is the set of all limit points of sequences $f(x_\nu)$ where x_ν has no convergent subsequence in V .

This definition is explored in Section 6.5. The following theorem is proved in Section 6.6. Recall that \mathcal{J} denotes the space of of all ω -tame J .

THEOREM C. Let (M, ω) be a closed semipositive symplectic manifold. Then there is a residual set $\mathcal{J}_{\text{reg}}(M, \omega) \subset \mathcal{J}_\tau(M, \omega)$ with the following significance. If $A \in H_2(M; \mathbb{Z})$ satisfies $c_1(A) > 0$ and $J \in \mathcal{J}_{\text{reg}}(M, \omega)$ then the evaluation map

$$\text{ev} : \mathcal{M}_{0,k}^*(A; J) \rightarrow M^k$$

is a pseudocycle of dimension $\mu(A, k)$. Moreover, its homology class is independent of the choice of $J \in \mathcal{J}_{\text{reg}}(M, \omega)$.

The proof is sketched in the introduction to Chapter 6. It involves transversality results for evaluation maps that are carried out in Section 6.1-3. The reader is advised, in particular, to explore the precise definition of the set $\mathcal{J}_{\text{reg}}(M, \omega)$ in section 6.2 in terms of transversality for edge evaluation maps as without this precise formulation Theorem C cannot be fully understood.

There is an important variant of Theorem C that concerns the restriction of ev to the subspace $\mathcal{M}_{0,k}^*(A; \mathbf{w}, J)$ of $\widetilde{\mathcal{M}}_{0,k}^*(A; J)$ consisting of tuples (u, z_1, \dots, z_k) for which the marked points z_i associated to the indices $i \in I$ are fixed and set equal to w_i . (Here we assume that $\#I \geq 3$ and denote $\mathbf{w} := \{w_i\}_{i \in I}$.) To establish that the resulting map

$$\text{ev} : \mathcal{M}_{0,k}^*(A; \mathbf{w}, J) \rightarrow M^k$$

is a pseudocycle (this time of dimension $\mu(A, k) - 2(\#I - 3)$), one must introduce a wider class of perturbations of the Cauchy–Riemann equation, looking at solutions of a equation in which $J = \{J_z\}$ depends on $z \in S^2$. Geometrically, this is equivalent to looking at the graphs $z \mapsto (z, u(z))$ of suitable maps $u : S^2 \rightarrow M$. A similar approach is needed in order to prove Theorem C above in the case when $c_1(A) = 0$. To understand the extension of Theorem C to these cases one must redo the whole theory in a slightly more general setting. The details are explained in Section 6.7.

1.4. The Gromov–Witten invariants

The **Gromov–Witten invariants** are obtained by taking the intersection of the Gromov–Witten pseudocycle

$$\text{ev} : \mathcal{M}_{0,k}^*(A; J) \rightarrow M^k$$

with cycles of complementary dimension in M^k . More precisely, let $a_1, \dots, a_k \in H^*(M)$ be cohomology class of pure degrees such that

$$m := \sum_{j=1}^k \deg(a_j) = \mu(A, k).$$

Then we may choose a cycle $\alpha \subset M^k$ that is Poincaré dual to the cohomology class $a := \pi_1^* a_1 \smile \dots \smile \pi_k^* a_k$ and is **strongly transverse** to the Gromov–Witten pseudocycle $\text{ev} : \mathcal{M}_{0,k}^*(A; J) \rightarrow M^k$ in a sense made precise in Section 6.5. We then define the Gromov–Witten invariant as the oriented intersection number

$$\text{GW}_{A,k}^M(a_1, \dots, a_k) = \text{ev} \cdot \alpha.$$

We may choose the cycle α to be a product $\alpha_1 \times \dots \times \alpha_k$ where $\alpha_i \subset M$ is Poincaré dual to the class a_i . In this case the Gromov–Witten invariant is the oriented number of J -holomorphic curves $u : S^2 \rightarrow M$ in the homology class A which meet each of the cycles $\alpha_1, \dots, \alpha_k$.

The invariants $\text{GW}_{A,k}^{M,I}$, in which the marked points z_i are fixed for $i \in I$, are defined similarly by taking the intersection of a suitable evaluation map with α . As a first approximation, one can take this evaluation map to be the restriction of ev to the subspace $\mathcal{M}_{0,k}^*(A; \mathbf{w}, J)$ on which $z_i := w_i$ for $i \in I$. Note that when $k = 3$ the invariants agree:

$$\text{GW}_{A,3}^M(a_1, a_2, a_3) = \text{GW}_{A,3}^{M,\{1,2,3\}}(a_1, a_2, a_3).$$

However, when $\#I > 3$ the two invariants are different: see Section 7.3. In particular, it turns out that the invariants $\text{GW}_{A,k}^{M,I}$ have very interesting formal properties that reflect the operad structure on the homology $H_*(\overline{\mathcal{M}}_{0,k})$ of the genus zero Grothendieck–Knudsen spaces. (More on this below.)

An important point is that these invariants depend only on the deformation class of (M, ω) , i.e. on the component of ω in the space of symplectic forms on M . The proof of this statement uses the existence of the virtual moduli cycle. In this book we prove only enough to show that the invariants do not change if ω varies among semipositive forms. Since every symplectic form on a manifold of dimension at most 6 is semipositive, this implies deformation invariance in these dimensions.

Section 7.1 defines the invariants $\text{GW}_{A,k}^M$ and gives the first examples. This section contains many cross references to earlier results in order to make it accessible to readers who want to start there, referring back to the technical details as needed. The invariants $\text{GW}_{A,k}^{M,I}$ with some fixed marked points are defined in Section 7.3. The ideas are illustrated in Section 7.4 that shows how to count rational curves in projective spaces, while Section 7.5 discusses the so-called **full Gromov–Witten invariant**.

To explain this, observe that when the number k of marked points is at least 3 there is a forgetful map

$$\pi : \mathcal{M}_{0,k}(A; J) \rightarrow \overline{\mathcal{M}}_{0,k},$$

where $\overline{\mathcal{M}}_{0,k}$ is the Grothendieck–Knudsen moduli space⁹ of stable curves of genus 0 and k marked points and π takes a stable map $[\mathbf{u}, \mathbf{z}]$ to the underlying stable curve

⁹This is also known as the Deligne–Mumford space. However, the latter names are more properly attached to the corresponding spaces for higher genus domains.

[z]. Consider the map

$$\text{ev} \times \pi : \mathcal{M}_{0,k}(A; J) \rightarrow M^k \times \overline{\mathcal{M}}_{0,k}$$

This is a pseudocycle only under very special conditions (see Exercise 6.7.13). Nevertheless, one can get very useful insight into the expected properties of the Gromov–Witten invariants by looking at them from this point of view. In particular, as explained at the end of Chapter 6, for each homology class β in $\overline{\mathcal{M}}_{0,k}$ one should be able to define an invariant whose value

$$\text{GW}_{A,k}^M(a; \beta)$$

on the class $a \in H^*(M^k)$ is the intersection number of a cycle α representing the Poincaré dual to a with the restriction of ev to the preimage under π of a cycle in $\overline{\mathcal{M}}_{0,k}$ representing the class β . When β is the fundamental class $[\overline{\mathcal{M}}_{0,k}]$ we recover the invariant $\text{GW}_{A,k}^M$, while if $\beta = [\text{pt}] \in H_0(\overline{\mathcal{M}}_{0,k}, \mathbb{Q})$ we obtain the invariant $\text{GW}_{A,k}^{M, \{1, \dots, k\}}$ in which all marked points are fixed. The power of this approach is immediately evident. If we represent the class $[\text{pt}] \in H_0(\overline{\mathcal{M}}_{0,k})$ by a point \mathbf{w} in the top stratum $\mathcal{M}_{0,k}$ of $\overline{\mathcal{M}}_{0,k}$, then \mathbf{w} is a tuple of k distinct points in S^2 and the invariant $\text{GW}_{A,k}^M(a; \text{pt})$ is obviously just $\text{GW}_{A,k}^{M, \{1, \dots, k\}}$. However, if we take \mathbf{w} to lie in some other stratum then we are counting curves whose domain is modelled on some fixed tree. If the invariant $\text{GW}_{A,k}^M(a; \beta)$ is to be well defined then these two counts must be the same. This implies that the Gromov–Witten invariants should satisfy the so-called (*Splitting*) axiom.

In Section 7.5 we formulate the Kontsevich–Manin axioms for Gromov–Witten invariants, and interpret them in terms of our invariants $\text{GW}_{A,k}^{M,I}$. All the axioms are easy to establish except for the (*Splitting*) axiom. Since this requires the gluing theorem, its proof is deferred to Section 10.8. The most important special case of this axiom is the decomposition rule

$$\begin{aligned} \text{GW}_{A,4}^{M, \{1,2,3,4\}}(a_1, a_2, a_3, a_4) &= \sum_{B \in H_2(M)} \sum_{\nu\mu} \\ &\text{GW}_{A-B,3}^M(a_1, a_2, e_\nu) g^{\nu\mu} \text{GW}_{B,3}^M(e_\mu, a_3, a_4), \end{aligned}$$

where e_ν runs over a basis for the rational cohomology $H^*(M)$ and $g^{\nu\mu}$ is the inverse of the pairing matrix $g_{\nu\mu} := \int_M e_\nu \smile e_\mu$. This turns out to be a crucial ingredient in the proof of associativity of the quantum cup product.

Gluing. The last very technical chapter is Chapter 10 which develops a gluing principle for J -holomorphic curves. This should be thought of as the converse of Gromov’s compactness theorem. It asserts that if two (or more) J -holomorphic curves intersect and satisfy a suitable transversality condition then they can be approximated in the sense of Gromov convergence by a sequence of J -holomorphic curves $u^\alpha : S^2 \rightarrow M$ representing the sum of their homology classes. This theorem was first proved by Ruan–Tian [345]. We give a different argument that mimics constructions used in gauge theory and is based on a careful choice of cutoff functions. Our proof is essentially the same as that in the 1994 version of this book. However, it is given in much greater detail. Although it takes considerable work to establish all the properties of the gluing map, the underlying ideas are not too difficult to understand, and we have tried to write the chapter to make these accessible.

The techniques used in this chapter do not occur elsewhere in the book, and so one can postpone reading it without serious consequences. As one can see from Chapter 9, there are many interesting applications of J -holomorphic curves that do not use gluing. These applications are based on the existence of a J -holomorphic curve with suitable properties. However, as soon as one needs to count the number of curves, one almost always needs to use gluing in the guise of the (*Splitting*) axiom. But often one does not need the full force of this axiom. A special case, which is enough to establish quantum cohomology, is proved in Section 10.9.

1.5. Applications and further developments

There are two main chapters devoted to applications, Chapter 9, which proves many of the results first stated by Gromov in his seminal 1985 paper, and Chapter 11 about quantum cohomology. Both chapters contained detailed proofs of the main results, together with extended discussions of related questions, which are listed in the index under the individual topics and also under the general headings “comments” and “examples”. The most accessible applications are the proof of the nonsqueezing theorem in Section 9.3, the results on symplectic 4-manifolds and their groups of symplectomorphisms in Sections 9.4 and 9.5, and the results in Section 9.7 which use Gromov–Witten invariants to distinguish between symplectic structures. None of these sections requires any more than a basic knowledge of the Gromov–Witten invariant $\text{GW}_{A,k}^M$, though they do use the soft methods of symplectic geometry. In particular, the proofs use the statements of Theorems A,B,C above rather than the methods and ideas in their proofs.

Section 9.2 on obstructions to Lagrangian embeddings is also fairly straightforward. It involves studying the properties of a perturbed Cauchy–Riemann equation on a disc with Lagrangian boundary values, and hence needs a little more analytic preparation. On the other hand, the first theorem in this section is Gromov’s celebrated result that there are nonstandard symplectic structures on Euclidean space. The proof is a simple and direct argument due to Gromov that is based on properties of Lagrangian submanifolds in Euclidean space. In this situation there are no J -holomorphic spheres, which simplifies the discussion of compactness.

The other applications in Chapter 9 rely more heavily on Chapter 8. This is a preparatory chapter about Hamiltonian fibrations that introduces a geometric framework for studying perturbations of the Cauchy–Riemann equation. Solutions to these perturbed equations can be interpreted as sections of a trivial bundle $S^2 \times M \rightarrow S^2$ that are holomorphic for an appropriate almost complex structure \tilde{J}_H . (The difference between \tilde{J}_H and the product almost complex structure corresponds to the perturbation.) Although trivial bundles suffice for some of the applications (such as those in Section 9.1), the analysis is no harder if one considers general locally Hamiltonian fibrations over arbitrary Riemann surfaces. Just as when studying graphs in Section 6.7 one then needs to reprove the basic regularity and compactness results in this more general setup. The arguments are essentially the same, but as always there are a few tricky points where one needs to take care to get a sharp result. With this, one obtains a useful tool (the Seidel representation) for studying the fundamental group of the group of Hamiltonian symplectomorphisms. One also obtains a definition of higher genus Gromov–Witten invariants, in the case when the complex structure on the domain (Σ, j) is fixed.

Section 9.1 describes important applications of these ideas to Hamiltonian dynamics. The first main result is that every Hamiltonian flow on a semipositive symplectic manifold has at least one 1-periodic orbit. To prove this in full detail one does need most of Chapter 8; although one can restrict to the trivial bundle, one still needs to understand how energy is measured and how to deal with bubbling. The second set of results in this section prove the existence of two distinct 1-periodic orbits on symplectically aspherical manifolds. They are somewhat easier, since the hypothesis implies there is no bubbling. Section 9.6 gives a taste of the Hofer geometry on the Hamiltonian group, describing a simple consequence of the Seidel representation due to Polterovich and Seidel. This representation can really only be defined in the context of quantum or Floer cohomology. Therefore, we return to it in both Chapter 11 and Chapter 12.

Quantum cohomology. Chapter 11 on quantum cohomology is completely independent of Chapter 9, and uses Chapter 8 only in the section on the Seidel representation. Hence it can be read immediately after Chapter 7.

The basic idea in the definition of quantum cohomology is very easy to understand. Let us suppose for simplicity that (M, ω) is (spherically) monotone, that is the restriction of the symplectic class $[\omega]$ to the spherical homology classes in $H_2(M)$ is a positive multiple $\tau c_1(TM)$ of the first Chern class. Further, denote by $H^*(M)$ the integral cohomology of M modulo torsion. Thus one can think of $H^*(M)$ as the image of $H^*(M, \mathbb{Z})$ in $H^*(M, \mathbb{R})$. The advantage of neglecting torsion is that the group $H^k(M)$, for example, may be identified with the dual $\text{Hom}(H^{2n-k}(M), \mathbb{Z})$; in other words a k -dimensional cohomology class a may be described by specifying all the values of its pairings

$$\langle a, c \rangle := \int_M a \smile c$$

with the elements $c \in H^{2n-k}(M)$.

We define the quantum cup product $a * b$ of the classes $a \in H^k(M)$ and $b \in H^\ell(M)$ as the formal sum

$$(1.5.1) \quad a * b = \sum_{A \in H_2(M)} (a * b)_A q^{c_1(A)},$$

where q is an auxiliary variable, considered to be of degree 2, and the cohomology class $(a * b)_A \in H^{k+\ell-2c_1(A)}(M)$ is defined in terms of Gromov–Witten invariants by

$$\langle (a * b)_A, c \rangle = \text{GW}_{A,3}^M(a, b, c), \quad c \in H^*(M).$$

Note that the classes a, b, c satisfy the dimension condition

$$(1.5.2) \quad 2n + 2c_1(A) = \deg(a) + \deg(b) + \deg(c)$$

required for $\text{GW}_{A,3}^M$ to be nonzero. Thus $c_1(A) \leq 2n$ and, since $0 \leq c_1(A)$ by monotonicity, only finitely many powers of q occur in the formula (1.5.1). Moreover, the classes A which contribute to the coefficient of q^d satisfy $\omega(A) = \tau c_1(A) = \tau d$ and hence, by Proposition 4.1.5, only finitely many can be represented by J -holomorphic curves. This shows that the sum on the right hand side of (1.5.1) is finite. Since only nonnegative exponents of q occur in the sum (1.5.1) it follows that $a * b$ is an element of the group

$$\text{QH}^*(M; \mathbb{Z}[q]) = H^*(M) \otimes \mathbb{Z}[q],$$

where $\mathbb{Z}[q]$ is the polynomial ring in the variable q of degree 2. Extending by linearity, we therefore get a multiplication

$$\mathrm{QH}^*(M; \mathbb{Z}[q]) \otimes \mathrm{QH}^*(M; \mathbb{Z}[q]) \rightarrow \mathrm{QH}^*(M; \mathbb{Z}[q]).$$

It turns out that this multiplication is distributive over addition, skew-commutative and associative. The first two properties are obvious, but the last is much more subtle and depends on the (*Splitting*) axiom. If (M, ω) is not monotone, then the sum occurring in (1.5.1) may be infinite. To make sense of it, one must choose a suitable coefficient ring Λ . There are several possible choices for this quantum coefficient ring: see Example 11.1.4.

If $A = 0$, all J -holomorphic curves in the class A are constant. It follows that $\mathrm{GW}_{0,3}^M(a, b, c)$ is just the usual triple intersection $\int_M a \smile b \smile c$. Since $\omega(A) > 0$ for all other A with J -holomorphic representatives, the constant term in $a * b$ is just the usual cup product. Thus $a * b$ is a deformation of the usual cup product.

As an example, let M be complex projective n -space $\mathbb{C}P^n$ with its usual Kähler form. If p is the positive generator of $H^2(\mathbb{C}P^n)$, and if $L \in H_2(\mathbb{C}P^n)$ is the class represented by the line $\mathbb{C}P^1$, then the fact that there is a unique line through any two points is reflected in the identity

$$\langle (p * p^n)_L, p^n \rangle = \mathrm{GW}_{L,3}^{\mathbb{C}P^n}(p, p^n, p^n) = 1.$$

By equation (1.5.2), the other classes $(p * p^n)_A$ vanish for reasons of dimension. Thus $p * p^n = q1$, where $1 \in H^0(\mathbb{C}P^n)$ is the unit. Further one can show that $p^i * p^j = p^{i+j}$ for $i + j \leq n$. Hence the quantum cohomology of $\mathbb{C}P^n$ is given by

$$\mathrm{QH}^*(\mathbb{C}P^n; \mathbb{Z}[q]) = \frac{\mathbb{Z}[p, q]}{\langle p^{n+1} = q \rangle}.$$

Example 11.1.12 gives a direct proof. We also explain how to deduce this result from Batyrev’s formula for the quantum cohomology of toric manifolds (Theorem 11.3.4); cf. Exercise 11.3.11. The occurrence of the letters p, q is no accident here. In Section 11.3.2, we describe some recent work of Givental and Kim in which they interpret the quantum cohomology ring of flag manifolds as the ring of functions on a Lagrangian variety.

One of the aims of Chapter 11 is to demonstrate how much information is carried by the quantum cohomology. For example, in the case of Grassmannians and flag manifolds the quantum cohomology reflects structures such as the Verlinde algebra and the Toda lattice. These connections to other areas of mathematics are not entirely unexpected because quantum cohomology forms the “A-side” of the mirror symmetry conjecture and these other structures are related to the “B-side”. We shall not attempt to explain what this means, but we do discuss some of the important related concepts such as Frobenius manifolds and the Dubrovin connection. We also explain how the information contained in the Gromov–Witten invariants can be encoded by a function Φ called the Gromov–Witten potential. The (*Splitting*) axiom turns out to be equivalent to the fact that Φ satisfies a third order partial differential equation called the WDVV-equation.

Floer theory. The last chapter explains the basic definitions and structures in Floer theory without going into the analytic details. (Many of these are written up in the book Audin–Damian [24].) We outline a proof that the ring structure on Floer theory agrees with that in quantum cohomology, and indicate some of

the directions of current research. In particular, we define the spectral invariants for Hamiltonian symplectomorphisms due to Schwarz and Oh, and explain how they interact with the Seidel representation. Floer theory has many variants, some of which are outlined in Remark 12.5.5 and in Section 12.6. We end with a brief discussion of the symplectic vortex equations, that form a bridge between the Gromov–Witten invariants and gauge theory.

Concluding remarks. This book explains the fundamentals of the theory of J -holomorphic curves. There are several other quite different techniques used in symplectic topology — for example variational methods, generating functions, and Donaldson’s technique of almost-complex geometry — as well as many outgrowths of the kind of elliptic techniques discussed here such as Floer theory, symplectic homology, and symplectic field theory. It is natural to wonder where the theory of J -holomorphic spheres fits in this spectrum.

The purely variational methods and the method of generating functions seem to work only when the underlying space is something like a cotangent bundle and so has some linear structure. For example, as in Hofer [177], Viterbo [404] or Bialy–Polterovich [37], one can use these methods to define various versions of the Hofer norm on the group $\text{Ham}(\mathbb{R}^{2n}, \omega_0)$ of Hamiltonian symplectomorphisms of Euclidean space. Conley and Zehnder’s proof in [75] of the Arnold conjecture for the torus lifted the problem to Euclidean space and then used variational methods. More recently, Tamarkin [388] and Guillermou–Kashiwara–Shapira [166] have used sheaf-theoretic methods (a significant generalization of the idea of a generating function) to establish some fundamental nondisplaceability results. Another situation where these methods apply is that of a toric manifold, which as we explain in Chapter 11.3 is the quotient of an open subset of Euclidean space by a (complex) torus. For example, Givental used this approach in [145]. When considering general symplectic manifolds, it seems that one must use J -holomorphic curves or some version of Floer homology (which one can think of as a combination of variational with elliptic techniques). For example, Lalonde–McDuff’s proof in [223] that the Hofer norm is nondegenerate on completely arbitrary symplectic manifolds uses J -holomorphic spheres together with symplectic embedding techniques. Symplectic geometers have combined the available methods in many ingenious ways. We mention some of the possibilities at appropriate places in Chapters 9, 11 and 12.