

Floer cohomology with gerbes

Michael Thaddeus

Department of Mathematics, Columbia University, New York, N.Y. 10027

This is a written account of expository lectures delivered at the summer school on “Enumerative invariants in algebraic geometry and string theory” of the Centro Internazionale Matematico Estivo, held in Cetraro in June 2005. However, it differs considerably from the lectures as they were actually given. Three of the lectures in the series were devoted to the recent work of Donaldson-Thomas, Maulik-Nekrasov-Okounkov-Pandharipande, and Nakajima-Yoshioka. Since this is well documented in the literature, it seemed needless to write it up again. Instead, what follows is a greatly expanded version of the other lectures, which were a little more speculative and the least strictly confined to algebraic geometry. However, they should interest algebraic geometers who have been contemplating *orbifold cohomology* and its close relative, the so-called *Fantechi-Göttsche ring*, which are discussed in the final portion of these notes.

Indeed, we intend to argue that orbifold cohomology is essentially the same as a symplectic cohomology theory, namely *Floer cohomology*. More specifically, the quantum product structures on Floer cohomology and on the Fantechi-Göttsche ring should coincide. None of this should come as a surprise, since orbifold cohomology arose chiefly from the work of Chen-Ruan in the symplectic setting, and since the differentials in both theories involve the counting of holomorphic curves. Nevertheless, the links between the two theories are worth spelling out.

To illustrate this theme further, we will explain how both the Floer and orbifold theories can be enriched by introducing a *flat* $U(1)$ -gerbe. Such a gerbe on a manifold (or orbifold) induces flat line bundles on its *loop space* and on its *inertia stack*, leading to Floer and orbifold cohomology theories with local coefficients. We will again argue that these two theories correspond. To explain all of this properly, an extended digression on the basic definitions and properties of gerbes is needed; it comprises the second of the three lectures.

The plan of these notes is simple: the first lecture is a review of Floer cohomology; the second is a review of gerbes, as promised a moment ago; and the third introduces orbifold cohomology and its relatives, discusses how to add a gerbe, and interprets these constructions in terms of Floer theory. We conclude with some notes on the literature.

Since these are lecture notes, no attempt has been made to include rigorous proofs.

Supported by NSF grant DMS-0401128.

But many aspects of Floer cohomology, especially its product structures, are not well documented in the literature either, so the reader is cautioned to take what is said about Floer cohomology with a grain of salt. The same goes for the proposed identification between the quantum product structures. It is mildly speculative but presumably should not be impossible to prove by following what has been done for the case of the identity map. Anyhow, for the moment we content ourselves with a genial narrative of a heuristic nature, making no great demands upon the reader. It presents many more definitions than theorems, but it aspires to provide a framework in which theorems may be built.

In the third lecture, I assume some familiarity with the basic definitions and properties of quantum cohomology, as given for example in the Clay Institute volume *Mirror Symmetry* (see the notes on the literature for a reference).

Acknowledgements: I thank Kai Behrend, Barbara Fantechi, and Marco Manetti for their kind invitation to speak in Cetraro, and for their patience in awaiting these notes. I am also grateful to Jim Bryan and the University of British Columbia for their hospitality while the notes were written, and to Dan Abramovich, Behrang Noohi, and Hsian-Hua Tseng for very helpful conversations and advice.

Lecture 1: Floer cohomology

This is an optimist's account of the Floer cohomology of symplectic manifolds: its origins, its construction, the main theorems, and the algebraic structures into which it naturally fits. Let me emphasize that, as an optimist's account, it presents Floer cohomology as we would like it to be, not necessarily as it is. For example, Floer proved the Arnold conjecture only in the presence of some ugly technical hypotheses, which later mathematicians have labored tirelessly to eradicate. The present account pretends that they never existed.

Floer introduced his cohomology (in fact he used homology, but never mind) to prove the Arnold conjecture on the number of fixed points of an exact Hamiltonian flow. Like so much of symplectic geometry, this problem is rooted in classical mechanics.

Newton's second law

Suppose a particle is moving in a time-dependent force field $\vec{F}(t, \vec{q})$. Here we regard $F : \mathbf{R} \times \mathbf{R}^3 \rightarrow \mathbf{R}^3$ as a time-dependent vector field. Newton's second law says that the trajectory $q(t)$ satisfies $F = ma$, or, taking $m = 1$ for simplicity,

$$F(t, q(t)) = \frac{d^2}{dt^2} q(t).$$

This is a second-order differential equation for $q(t)$. It can be easier to solve, and perhaps even visualize, such equations by the standard trick of introducing a triple of extra variables \vec{p} and regarding the above as a first-order equation for $(p, q) \in \mathbf{R}^6$:

$$\begin{cases} F(t, q(t)) = \frac{d}{dt} p(t) \\ p(t) = \frac{d}{dt} q(t). \end{cases}$$

The solutions are flows along the time-dependent vector field on \mathbf{R}^6 whose value at (p, q) is $(F(t, q), p)$.

Notice that q describes the particle's position, and p describes its velocity or momentum. The space \mathbf{R}^6 of all (p, q) can therefore be regarded as the space of all initial conditions for the particle.

The Hamiltonian formalism

Hamiltonian mechanics takes off from here. The idea is to cast the construction above in terms of the symplectic form on $\mathbf{R}^6 = T^*\mathbf{R}^3$, and generalize it to an arbitrary symplectic manifold.

So let M be a symplectic manifold with symplectic form ω : call it the *phase space*. Let $H : \mathbf{R} \times M \rightarrow \mathbf{R}$ be any time-dependent smooth function on M : call it the *Hamiltonian*. The symplectic form induces an isomorphism $T^*M \cong TM$; use this to make the exterior derivative $dH \in \Gamma(T^*M)$ into a vector field $V_H \in \Gamma(TM)$. Since H can depend on the time $t \in \mathbf{R}$, V_H is really a time-dependent vector field $V_H(t)$.

The *exact Hamiltonian flow* of H is the 1-parameter family of symplectomorphisms of M

$$\Phi : \mathbf{R} \times M \rightarrow M$$

such that $d\Phi/dt = V_H(t)$ and $\Phi(0, x) = x$. Its existence and uniqueness are guaranteed by the standard theory of ODEs (at least for small t , or for M compact).

What makes this formalism so great is that it correctly describes the actual time evolution of a mechanical system when (1) the phase space M is the space of initial conditions of our system (i.e. possible positions and momenta) and (2) the Hamiltonian is the total energy, potential plus kinetic. The phase space will have a canonical symplectic form. Typically, it is of the form $M = T^*Q$ where Q parametrizes the possible configurations of the system. So one can choose (at least locally) position variables q_i and momentum variables p_i , the symplectic form is $\sum dp_i \wedge dq_i$, and kinetic energy (being essentially $\frac{1}{2}mv^2$) is some quadratic function of the momenta.

For example, from a time-independent quadratic potential on \mathbf{R}^3 , you should get the simple harmonic oscillator. Or for a force $F = \nabla\psi$ on \mathbf{R}^3 with time-dependent potential $\psi(q_1, q_2, q_3)$, you can recover the first-order equation of the previous section by taking $H = \psi(q_1, q_2, q_3) + \frac{1}{2}(p_1^2 + p_2^2 + p_3^2)$. Or, of course, you could look at n particles moving in \mathbf{R}^3 ; then the phase space will be \mathbf{R}^{6n} .

Parenthetically, let's clear up a confusing detail: why is the phase space M typically identified with a cotangent bundle T^*Q and not a tangent bundle TQ ? That is, why should momenta be considered cotangent vectors rather than tangent vectors? Of course, in physics we typically have a metric inducing an isomorphism between the two. Still, we may have muddied the waters somewhat by setting the mass equal to 1. The point is that momentum is a vector-valued quantity with units of g cm/sec; it should be regarded as pairing with velocity, a vector with units of cm/sec, to give energy, a scalar with units of g cm²/sec².

The Arnold conjecture

Arnold was interested in applications of Hamiltonian mechanics to real-life many-body problems, such as the long-term stability of the solar system. Then one is of course particularly interested in points in phase spaces that flow back to themselves, that is, $\Phi(t, x) = x$ for some $t > 0$, say $t = 1$. From now on let's write $\phi_t(x) = \Phi(t, x)$, so that $\phi_t : M \rightarrow M$ is a symplectomorphism.

Phase spaces in problems of physical interest are almost always noncompact, but Arnold realized that stronger statements might hold in the compact case. He conjectured the following:

If M is compact and ϕ_1 as above has nondegenerate fixed points, then the number of those fixed points is at least the sum of the Betti numbers of M .

Nondegeneracy of a fixed point x here means that $d\phi_1(x) - \text{id}$ is nonsingular. A more general version of the Arnold conjecture, which we omit, deals with the degenerate case.

In this situation the Lefschetz fixed-point formula implies that the number of fixed points is at least the Euler characteristic, that is, the alternating sum of the Betti numbers. Hence the Arnold conjecture gives a stronger lower bound in the exact Hamiltonian case, as long as some odd Betti number is nonzero. On the other hand, if we replace the exact form dH used to define an exact Hamiltonian flow by a general closed form (you might call this a *closed Hamiltonian flow*, but you can easily check that all 1-parameter families of symplectomorphisms starting at the identity are of this kind), then the Arnold

conjecture is false. Just consider the linear flow on a torus.

Floer's proof

Floer defines cohomology groups $HF^*(\phi)$ associated to any symplectomorphism, and shows

(1) that $HF^*(\text{id}) = H^*(M)$;

(2) that for ϕ with nondegenerate fixed points, $HF^*(\phi)$ can be calculated from a complex whose chains are formal linear combinations of fixed points;

(3) that $HF^*(\phi)$ is, in a suitable sense, invariant under composition with an exact Hamiltonian flow.

The Arnold conjecture is an immediate consequence, as the dimension of a chain complex must be at least the dimension of its cohomology.

The chain complex leading to this cohomology theory is an infinite-dimensional analogue of the Morse complex, so let's pause first to review the salient points about that.

Morse theory

Let X be a compact oriented manifold of finite dimension n . A *Morse function* $f : X \rightarrow \mathbf{R}$ is a smooth function with isolated critical points, at each of which the Hessian is nondegenerate. The *Hessian* is the matrix of second partials, but never mind: just recall instead that, according to the Morse lemma, this nondegeneracy is equivalent to the existence of local coordinates x_1, \dots, x_n in which

$$f(x_1, \dots, x_n) = -x_1^2 - x_2^2 - \dots - x_m^2 + x_{m+1}^2 + \dots + x_n^2.$$

The number of negative terms is called the *Morse index*.

Let C be the set of formal linear combinations of the critical points x_i (with, say, complex coefficients). This is a finite-dimensional vector space, and the Morse index $m(i)$ provides a grading. We can define a differential $d : C \rightarrow C$ by

$$d(x_i) = \sum_{j \mid m(j) - m(i) = 1} \#(i, j) x_j,$$

where $\#(i, j)$ denotes the number of gradient flow lines from x_i to x_j , counted with the appropriate signs. This means the following. Choose a Riemannian metric on X , so that the gradient ∇f is a vector field. The downward gradient flow from x_i and the upward gradient flow from x_j are submanifolds of dimension $m(i)$ and $n - m(j)$ respectively.

For a sufficiently general metric, they intersect transversely. The index difference being 1 then implies that they intersect in a finite number of flow lines. Choose an orientation of each downward flow; this induces an orientation of each upward flow. Each flow line from x_i to x_j then acquires a sign by comparing four orientations: those of X , the upward and downward flows, and the flow line itself.

It is easy to check that the choice of orientations makes no significant difference. A much harder fact is that $d^2 = 0$. One has to look at flows between critical points of index difference 2: instead of being parametrized by a finite set (= compact 0-manifold) as above, these are parametrized by a disjoint union of closed intervals (= compact 1-manifold), and the crucial point is that there are 0 points in the boundary, when they are counted with the appropriate signs.

So now we have a chain complex, and can take cohomology in the usual way. The amazing fact is that what we get is naturally isomorphic to the rational cohomology of the manifold X !

Notice that this immediately implies the Arnold conjecture in the time-independent case. For the nondegeneracy is then equivalent to the time-independent Hamiltonian $H : M \rightarrow M$ being a Morse function, and the fixed points of ϕ_1 are the critical points of H .

Bott-Morse theory

Morse functions always exist; in fact, they are dense among all smooth functions. Nevertheless, suppose fate has endowed us with some smooth $f : X \rightarrow \mathbf{R}$ which is not a Morse function. Can we still use it to determine the cohomology of X ? We could try perturbing f to get a Morse function. But often there is no choice of a perturbation which is practical for calculation.

There is one case where we still get some useful information. This is when f is a *Bott-Morse* function: that is, the critical points are a disjoint union of submanifolds, on whose normal bundles the Hessian is nondegenerate. In other words, near every critical point there exist local coordinates in which f can be expressed as before, except that some of the coordinates may be entirely absent. A good example is the pullback of a Morse function by the projection in a fiber bundle.

In the Bott-Morse case, there exists a spectral sequence whose E^2 term is the cohomology of the critical set, bigraded by the Morse index and the degree of the cohomology. It abuts to the cohomology of X . An easy exercise is to show that, in the original Morse case, this boils down to the cochain complex described before. A harder exercise is to

show that, in the example of the previous paragraph, it boils down to the Leray spectral sequence.

Morse theory on the loop space

Now let's return to our Floer set-up: a symplectomorphism $\phi : M \rightarrow M$. We might as well assume that M is connected. Let the *loop space* LM be the set of all smooth maps from the circle S^1 to M . In the case $\phi = \text{id}$, we will define Floer cohomology to be essentially the Morse cohomology of LM , with a "symplectic action function" F playing the role of the Morse function. The loop space is in some sense a manifold, but it is infinite-dimensional, and the upward and downward flows from the critical sets will both be infinite-dimensional as well, so it is lucky that we are optimists.

What is this function F ? Suppose first that $\pi_1(M) = 1$, so that LM is connected too. For any $\ell \in LM$, $\ell : S^1 \rightarrow M$, choose a map $\bar{\ell} : D^2 \rightarrow M$ extending ℓ , where D^2 is the disc, and let $F(\ell) = \int_{D^2} \bar{\ell}^* \omega$. This is only defined modulo the integrals of ω on spheres in M , but we can pass to the covering space $\tilde{L}M$ determined by the quotient $\pi_1(LM) \rightarrow \pi_2(M)$, and there F is defined without ambiguity. Indeed, $\tilde{L}M$ can be regarded as the space of loops plus homotopy classes of extensions $\bar{\ell}$.

If $\pi_1(M) \neq 1$, then LM has several components, and if we fix a loop in each, we can extend ℓ to a cylinder agreeing with the fixed loop on the other end, and proceed as before.

As a matter of fact, for general ϕ we can do something similar: let the *twisted loop space* be

$$L_\phi M = \{\ell : \mathbf{R} \rightarrow M \mid \ell(t+1) = \phi(\ell(t))\},$$

and fix a twisted loop in each connected component. A path from any twisted loop ℓ to the fixed one is a smooth map $\bar{\ell} : \mathbf{R} \times [0, 1] \rightarrow M$ satisfying the obvious periodicity and boundary properties, and we define $F(\ell) = \int_{[0,1] \times [0,1]} \bar{\ell}^* \omega$ as one would expect.

This function F is a very natural one. Indeed, we can define a symplectic form Ω on $L_\phi M$ as follows. The tangent space to $L_\phi M$ at ℓ consists of sections of $\ell^* TM$ which are periodic in a suitable sense. Define $\Omega(u, v) = \int_0^1 \omega(u, v) dt$. Then the Hamiltonian flow of F is exactly reparametrization of twisted loops by time translations.

Consequently, the critical points are exactly the constant loops: these must take values in the fixed-point set X^ϕ by the definition of the twisted loop space, so the critical set can be identified with X^ϕ . Paths in the twisted loop space are, of course, maps $\mathbf{R} \times \mathbf{R} \rightarrow M$ with the appropriate periodicity in the first factor. The gradient flow lines turn out to be exactly the *pseudo-holomorphic* maps, that is, maps whose derivatives

are linear over \mathbf{C} . (For brevity we refer to them henceforth as holomorphic.) Here the choice of an almost complex structure on M compatible with ω has induced a metric g on M and hence a metric G on $L_\phi M$.

Here is a sketch of why the gradient flows are exactly the holomorphic maps. Let $t+iu$ be coordinates on $\mathbf{R}^2 = \mathbf{C}$. A map $\ell : \mathbf{R}^2 \rightarrow M$ is a gradient flow if $\partial\ell/\partial u \in T_\ell L_\phi M$ is dual under the metric G to dF , that is, if for all $\nu \in T_\ell L_\phi M$,

$$G\left(\frac{\partial\ell}{\partial u}, \nu\right) = dF(\nu)$$

or

$$\int_0^1 g\left(\frac{\partial\ell}{\partial u}, \nu\right) dt = \int_0^1 \omega\left(\frac{\partial\ell}{\partial t}, \nu\right) dt.$$

Since $\omega(\mu, \nu) = g(i\mu, \nu)$, this is equivalent to

$$i \frac{\partial\ell}{\partial t} = \frac{\partial\ell}{\partial u},$$

which is the complex linearity of the derivative.

If everything is sufficiently generic, F is a Morse function. Then we can go ahead and define our Morse complex, where the differential d counts holomorphic maps. The key claims are that we can make things sufficiently generic by composing with some exact Hamiltonian flow, that $d^2 = 0$ as in the finite-dimensional case, and that the cohomology we get does not depend on the flow.

In many cases, F is not sufficiently generic, but it is still a Bott-Morse function: that is, the critical points are a union of submanifolds, and the Hessian is (in some infinite-dimensional sense!) nondegenerate on each normal bundle. Then we're going to get our spectral sequence. We presume that the Floer cohomology can be calculated from it: a highly nontrivial presumption, of course! This is not Floer's actual approach, but it is still a good way to think about it.

For example, if $\phi = \text{id}$ again, then there is just one critical submanifold, identified with M itself. Hence the E^2 term of the spectral sequence is supported in a single row, so we immediately conclude that $HF^*(\text{id}) = H^*(M)$, provided that our presumption is correct.

That sounds very nice, but only because we cheated. We neglected to pass to the cover $\tilde{L}M$. Up there, there are many critical submanifolds, all diffeomorphic to M but interchanged by deck transformations $\pi_2(M)$. If the Morse indices are different, the spectral sequence won't be supported in a single row, so we need another argument to

ensure that the differentials vanish. This is indeed true, but won't be justified, even heuristically, until we discuss the finite-order case a little later on.

So a more truthful statement is that $HF^*(\text{id})$ is a direct sum of many copies of $H^*(M, \mathbf{C})$, one for each element of $\pi_2(M)$. This is conveniently written by introducing $\Lambda = \mathbf{C}[\pi_2(M)]$, the group algebra of $\pi_2(M)$. For example, if $\pi_2(M) \cong \mathbf{Z}$, then $\Lambda \cong \mathbf{C}[q, q^{-1}]$. More generally, there will be variables q_1, q_2, \dots corresponding to generators β_1, β_2, \dots of $\pi_2(M)$. Then we have an isomorphism $HF^*(\text{id}) \cong H^*(M, \Lambda)$. We've glossed over the correct definition of the index, but suffice it to say that the correct grading of $q_i \in H^0(M, \Lambda)$ is $c_1(TM)[\beta_i]$. Here the almost complex structure on M has made the tangent bundle TM into a complex vector bundle.

Re-interpretation #1: sections of the symplectic mapping torus

If you don't like the periodicity condition on our holomorphic maps, here is another way to look at the flow lines. Let the integers act on $\mathbf{C} \times M$, on the first factor by translation by $\mathbf{Z} \subset \mathbf{C}$, on the second by iterating ϕ . This acts freely and symplectically, so the quotient M_ϕ is a symplectic manifold. It is a bundle over the cylinder whose fiber is M , and it admits a canonical flat connection whose monodromy is ϕ . For that reason we call it the *symplectic mapping torus*.

Fixed points of ϕ precisely correspond to flat sections of this bundle. Gradient flow lines of F correspond to holomorphic sections: indeed, both correspond to periodic maps $\ell : \mathbf{R}^2 \rightarrow M$ as in the previous section. And the convergence of a flow line to two given fixed points at its ends corresponds to the convergence of the holomorphic section to two given flat sections as we move toward the two ends of the cylinder.

The *periodic Floer homology* of Hutchings is a generalization of this in the case where M is a surface: one looks not only at fixed points, but at unordered k -tuples fixed by ϕ , and the differential consists of k -valued sections, possibly ramified. It is conjectured to be related to Seiberg-Witten Floer cohomology.

Re-interpretation #2: two Lagrangian submanifolds

Another flavor of Floer cohomology takes as its data a compact symplectic manifold N and two Lagrangian submanifolds $L_1, L_2 \subset N$. Act on one of them by an exact Hamiltonian flow until L_1 intersects L_2 transversely (exercise: this is possible). Then consider the Morse cohomology of the space of paths from L_1 to L_2 .

That is, define chains to be formal linear combinations of points $x_i \in L_1 \cap L_2$. And define a differential as before, but where $\#(i, j)$ now counts holomorphic maps from the

strip $[0, 1] \times \mathbf{R}$ to N such that the two ends of the strip converge to x_i and x_j . Once again, the grading is contrived in such a way that, if $m(i) - m(j) = 1$, we expect a finite number of such maps (modulo translations of the strip).

This flavor has to do with Floer's work on 3-manifold topology. For example, given a Heegaard decomposition of a 3-manifold, let N be the space of irreducible flat $SU(2)$ -connections on the bounding surface, and let L_1, L_2 be the connections that extend as flat connections over the two handlebodies. This satisfies the conditions of the previous paragraph except that N is not compact. Optimistically ignoring this technicality, we may state the *Atiyah-Floer conjecture* which claims that the symplectic Floer cohomology of this N agrees with the *instanton Floer cohomology* of the 3-manifold, also defined by Floer. We won't discuss it here except to say that it is roughly the Morse cohomology of the Chern-Simons function on the space of connections on the 3-manifold.

But we digress. Let's see how the previous flavor of Floer cohomology can be regarded as a special case of this one. Just take $N = M \times M$ with the symplectic form $\pi_1^*\omega - \pi_2^*\omega$ where π_1, π_2 are projections, and let L_1, L_2 be the diagonal and the graph of ϕ . The minus sign is chosen so that these will be Lagrangian. To see how the holomorphic curves in the two alternatives correspond, start with a section of the symplectic mapping torus, project the cylinder 2:1 onto a strip $[0, 1] \times \mathbf{R}$ (branched over the boundary components $0 \times \mathbf{R}$ and $1 \times \mathbf{R}$), trivialize the mapping torus in the natural way over the complement of $1 \times \mathbf{R}$, and define a map $[0, 1] \times \mathbf{R} \rightarrow M \times M$ taking a point on the strip to the values of the section at the two points of the cylinder above it, relative to this trivialization. An explicit formula is easy to write down, but why bother?

Product structures

Both of the alternatives above suggest a way to introduce a product structure on Floer cohomology. In fact, what we're going to define is a linear functional on

$$HF^*(\phi_1) \otimes HF^*(\phi_2) \otimes HF^*(\phi_3)$$

for any symplectomorphisms satisfying $\phi_1\phi_2\phi_3 = \text{id}$. (Technical detail: since infinitely many powers of q may appear in this element, we may have to pass to a slightly larger coefficient ring $\bar{\Lambda}$, the *Novikov ring*. For example, if $\Lambda = \mathbf{C}[q, q^{-1}]$, then $\bar{\Lambda} = \mathbf{C}[[q]][[q^{-1}]]$.)

Notice that the chains defining Floer cohomology for ϕ and ϕ^{-1} are formal linear combinations of the same fixed points. If one uses the Kronecker delta to define a nondegenerate pairing between these chains, this descends to a nondegenerate pairing $HF^*(\phi) \otimes HF^*(\phi^{-1}) \rightarrow \mathbf{C}$. The linear functional above can then be regarded as a linear

map

$$HF^*(\phi_1) \otimes HF^*(\phi_2) \longrightarrow HF^*(\phi_1\phi_2).$$

This ought to satisfy some kind of relation like associativity. In particular, for $\phi_1 = \phi_2 = \text{id}$, it ought to define an associative product on $HF^*(\text{id}) = H^*(M)$. For $\phi_1 = \text{id}$, it makes any $HF^*(\phi)$ into a module over $HF^*(\text{id})$. And so on.

In fact, it has been proved that the Floer product on $HF^*(\text{id})$ coincides with the quantum product coming from Gromov-Witten theory. So we can regard each $HF^*(\phi)$ as a module over the quantum cohomology ring.

Now, what is the linear functional we promised to define? In analogy with alternative #1, it's given by counting sections of a bundle over a sphere minus three points. (The cylinder was a sphere minus two points.) Call this surface S ; then $\pi_1(S)$ is free on two generators. Let $M_{\phi_1, \phi_2} = (\tilde{S} \times M)/\pi_1(S)$, where \tilde{S} is the universal cover and $\pi_1(S)$ acts on M via ϕ_1 and ϕ_2 . This is a symplectic bundle over S with fiber M . Now count holomorphic curves asymptotic to fixed points x_j, x_k, x_ℓ of ϕ_1, ϕ_2, ϕ_3 on the three ends.

One has to prove that this induces a homomorphism of complexes. The proof is supposed to be a gluing argument. So is the proof of associativity. The idea is to take a sphere minus three (resp. four) discs, and shrink a loop encircling one (resp. two) of the discs to a point. Then study the limiting behavior of holomorphic sections of the bundles with this base and fiber M as the loop shrinks.

By the way, how can all this be phrased in terms of alternative #2? It's easy to convince yourself that the product functional counts holomorphic triangles in $M \times M$ whose edges lie on the graphs of id , ϕ_1 , and $\phi_1\phi_2$. More generally, if $HF^*(L_1, L_2)$ denotes the Floer cohomology of two Lagrangian submanifolds, then there is supposed to be a product operation

$$HF^*(L_1, L_2) \otimes HF^*(L_2, L_3) \longrightarrow HF^*(L_1, L_3)$$

which counts holomorphic triangles with edges in L_1, L_2, L_3 .

From either point of view, it's clear that there is no reason to stop with three punctures. One can include any number, working with a sphere minus n points in alternative 1, or an n -gon in alternative 2, and they will induce $(n-1)$ -ary operations on the chain complexes which will descend to Massey products on the cohomology. The compatibility between these operations seems to be what Fukaya is describing in his definition of an A^∞ category. There's a substantial literature about complexes equipped with such operations, which it would be quite interesting to apply to Floer cohomology. E.g. the

Massey products on a compact Kähler manifold are known to vanish. Is this true of the Floer Massey products?

If you want to go even further, there's no need to insist that S be a punctured sphere: it could be a surface of any genus. Correspondingly, instead of n -gons, you could look at non-simply-connected domains.

The finite-order case

If ϕ has finite order, say $\phi^k = \text{id}$, then $L_\phi M$ can be regarded as a subspace of LM just by speeding up the path by a factor of k . The symplectic action function on $\tilde{L}M$ restricts to the one on $\tilde{L}_\phi M$, up to a scalar multiple. The Hamiltonian flow of F is reparametrization by time translations, but translations by integer values now act trivially, so the flow induces a circle action. In this situation — when the Hamiltonian flow of F induces a circle action — we say that F is the *moment map* for the action.

Now in finite dimensions, it is well known that moment maps for circle actions are *perfect* Bott-Morse functions, meaning that the differentials in the associated spectral sequence are all zero, or equivalently, that the associated Morse inequalities are equalities. Let's suppose that this remains true in our infinite-dimensional setting. If so, we conclude that *if ϕ has finite order, then*

$$HF^*(\phi) \cong H^*(M^\phi).$$

The author has been informed by Hutchings that, under some technical hypotheses, this result can be proved rigorously. It is, of course, a generalization of Floer's result that $HF^*(\text{id}) = H^*(M)$.

Givental's philosophy

Givental's philosophy is that Floer cohomology leads in a natural way to differential equations, and to solutions of those equations. These solutions are in some sense generating functions for numbers of rational curves on M ; for example, when M is the quintic threefold, we get the famous Picard-Fuchs equation predicted by mirror symmetry.

Givental considers *equivariant* Floer cohomology (even though this is hard to define rigorously): the circle S^1 acts on LM by rotating the loop. He denotes the generator of $H^*(BS^1)$ by \hbar . Every symplectic form ω on M induces an equivariantly closed 2-form ρ on $\tilde{L}M$. Indeed, with respect to the symplectic form Ω on LM defined earlier, the circle action given by reparametrization is Hamiltonian when we pass to the cover $\tilde{L}M$, and the moment map is exactly the action function F . It is part of the usual package

of ideas in equivariant cohomology that $p = \Omega + F$ can be regarded as an equivariantly closed 2-form, the *Duistermaat-Heckman form*.

Suppose for simplicity that M is simply connected. Then $\pi_2(M) = H_2(M, \mathbf{Z})$ by the Hurewicz theorem. If this has rank k , let $\omega_1, \dots, \omega_k$ be a basis consisting of integral symplectic forms, and let q_1, \dots, q_k be the deck transformations of $\tilde{L}M$ corresponding to the dual basis of $H_2(M, \mathbf{Z})$. We can act on the Floer cohomology $HF^*(\text{id})$ by multiplication by p_i , or by pullback by q_i . These operations all turn out to commute, except that

$$p_i q_i - q_i p_i = \hbar q_i.$$

The noncommutative algebra D over \mathbf{C} generated by p_i and q_i (and q_i^{-1} , since this is the inverse deck transformation), with these relations, is a familiar one. At any rate, it can be regarded as an algebra of differential operators if we set $q_i = e^{t_i}$ and $p_i = \hbar \partial / \partial t_i$.

So we should think of a D -module, such as $HF^*(\text{id})$, as a sheaf on a torus $(\mathbf{C}^\times)^k$ equipped with a connection (or rather, a 1-parameter family of connections parametrized by \hbar). As a $\mathbf{C}[q_i]$ -module, $HF^*(\text{id})$ is free, so the sheaf is a trivial bundle. Only the connection is nontrivial. What we want to know is encoded in the flat sections of the bundle, which are functions of the q_i (and \hbar) with values in $H^*(M)$.

Suppose we are in the good case where $H^*(M)$ is generated by H^2 . Then $HF^*(\text{id})$ is a principal D -module generated by $1 \in H^0(M)$: this is plausible, since p_i tends to the cup product with ω_i as $q_i \rightarrow 0$. So there is a canonical surjection of D -modules $D \rightarrow HF^*(\text{id})$. Its kernel K is generated by a finite number of differential operators, and setting these to zero gives the differential equations that determine what we want to know.

Indeed, knowing the flat sections is the same as knowing $\text{Hom}_D(HF^*(\text{id}), \mathcal{O})$, where \mathcal{O} is the sheaf of regular functions on the torus, for such homomorphisms are just the constant maps in terms of a basis of flat sections. On the other hand, such a thing is also the same as a module homomorphism $D \rightarrow \mathcal{O}$ which kills K . It is determined by its value at 1, and this consists of a function which satisfies all the differential equations in K .

This heuristic argument inspired Givental's approach to determining the Gromov-Witten invariants for the quintic threefold, and more generally for Calabi-Yau complete intersections in toric varieties. Instead of using the loop space, he uses spaces of stable maps, which he regards as finite-dimensional approximations to the loop space.

Lecture 2: Gerbes

And now for something completely different: the definition of a gerbe. The motivation for introducing them is quite simple. We want to consider Floer cohomology with local coefficients in a flat $U(1)$ -bundle over the loop space LM (and its twisted variants). This bundle should of course come from some kind of geometric structure on M , and a $U(1)$ -gerbe will be the best candidate.

Here is the first clue to what a gerbe should be. Isomorphism classes of flat line bundles on LM correspond to $H^1(LM, U(1))$. There is a natural *transgression map* $H^2(M, U(1)) \rightarrow H^1(LM, U(1))$ given by taking the Künneth component in $H^1 \otimes H^1$ of the pullback by the evaluation map $LM \times S^1 \rightarrow M$. So we might expect gerbes to be objects whose isomorphism classes correspond to $H^2(M, U(1))$.

The good news: such objects exist. They were created in the 1960s by Giraud, who was chiefly interested in nonabelian structure groups. Abelian gerbes were discussed in more detail by Brylinski in a book some 25 years later. The bad news: gerbes rely on the theory of stacks, which we now review in the briefest possible terms.

Definition of stacks

Let \mathcal{T} be the category of topological spaces (and continuous maps). The category \mathcal{S} of principal G -bundles (and bundle maps) has an obvious covariant functor to \mathcal{T} , namely passing to the base space. It enjoys the following properties.

- (a) **Inverses:** any bundle map over the identity $X \rightarrow X$ is an isomorphism.
- (b) **Pullbacks:** given any P over X and any continuous $f : Y \rightarrow X$, there exists Q over Y with a bundle map $Q \rightarrow P$, namely the pullback $Q = f^*P$. It is unique up to unique isomorphism, and it satisfies the obvious universal property for bundle maps over some $Z \rightarrow X$ factoring through Y .
- (c) **Gluing of bundles:** given an open cover U_α of X , bundles P_α over U_α , and isomorphisms $f_{\alpha\beta}$ on the double overlaps (with $f_{\alpha\alpha} := \text{id}$) satisfying $f_{\alpha\beta}f_{\beta\gamma}f_{\gamma\alpha} = \text{id}$ on the triple overlaps, there exists a bundle P over X with isomorphisms g_α over U_α to each P_α satisfying $f_{\alpha\beta}g_\beta = g_\alpha$.
- (d) **Gluing of bundle isomorphisms:** given two bundles P, P' over X , an open cover U_α , and isomorphisms from P to P' over each U_α agreeing on the double overlaps, there is a unique global isomorphism from P to P' agreeing with all the given ones. (Note this implies that gluing of bundles is unique up to isomorphism.)

A *stack* over \mathcal{T} is simply any category \mathcal{S} , equipped with a covariant functor to \mathcal{T} ,

that satisfies properties (a), (b), (c), (d). Here, of course, “bundle” should be replaced by “object” and “bundle map” by “morphism.” In this setting the properties have new, alarming names: (a) and (b) make \mathcal{S} a *category fibered in groupoids*; (c) says that *descent data are effective* and (d) says that *automorphisms are a sheaf*. Notice, by the way, that (c) and (d) implicitly use (b).

You don’t really need the base category to be that of topological spaces, of course. It can be any category where the objects are equipped with a Grothendieck topology, such as schemes with the étale topology, which allows us to make sense of open covers.

Examples of stacks

(1) The stack of principal G -bundles described above is called the *classifying stack* and denoted BG .

(2) The stack of *flat* principal G -bundles, that is, G -bundles equipped with an atlas whose transition functions are locally constant, with the obvious notion of flat bundle maps. In a flat bundle, nearby fibers (i.e. those in a contractible neighborhood) may be canonically identified.

(3) For a fixed space X , the category whose objects are continuous maps $Y \rightarrow X$ and whose morphisms are commutative triangles ending at X . The covariant functor takes a map to its domain. This is a stack, denoted $[X]$ or simply X . Note that in this case (d) becomes trivial, because isomorphism is just equality.

(4) For a topological group G acting on X , the category whose objects lying over Y are pairs consisting of principal G -bundles $P \rightarrow Y$ and G -equivariant maps $P \rightarrow X$. We leave it to the reader to figure out what the morphisms are. This is a stack, denoted $[X/G]$. Notice that this simultaneously generalizes (1), which is the case $[\cdot/G]$, and (3), which is the case $[X/\cdot]$, where \cdot denotes a point.

(5) A more exotic example: for a fixed line bundle $L \rightarrow X$ and a fixed integer n , the category whose objects are triples consisting of a map $f : Y \rightarrow X$, a line bundle $M \rightarrow Y$, and an isomorphism $M^{\otimes n} \cong f^*L$. This was studied by Cadman, who called it the *stack of n th roots*.

(6) For any two stacks, there is a *Cartesian product* stack whose objects are pairs of objects lying over the same space. For example, an object of $X \times BG$ is a map $Y \rightarrow X$ and a principal G -bundle $P \rightarrow Y$.

Morphisms and 2-morphisms

In the theory of categories, much mischief is caused by our inability to declare that two given objects are equal. In the category of finite-dimensional complex vector spaces, for example, we can't say that $V^{**} = V$. The only accurate statement is that they are naturally isomorphic. So if D is the functor taking a vector space to its dual, we can't say that $DD = \text{id}$. We can only say that there is a natural transformation of functors $DD \Rightarrow \text{id}$. We encounter similar mischief in the theory of stacks.

A *morphism of stacks* $\mathcal{S}' \rightarrow \mathcal{S}$ is a functor between categories compatible with the covariant functors to \mathcal{T} . A stack equipped with a morphism to \mathcal{S} is called a *stack over* \mathcal{S} .

But, if $F, F' : \mathcal{S}' \rightarrow \mathcal{S}$ are both morphisms of stacks, we also have the mind-expanding concept of a *2-morphism of morphisms* $\Theta : F \Rightarrow F'$, which is a natural transformation of the corresponding functors. Likewise, a *2-isomorphism of morphisms* is a natural isomorphism of the corresponding functors. For example, if $BGL(n)$ is the stack of (frame bundles of) rank n complex vector bundles, then taking the dual bundle defines a morphism $D : BGL(n) \rightarrow BGL(n)$ of stacks, and there exists a 2-isomorphism $DD \Rightarrow \text{id}$.

(Exercise: show that, for X a space and \mathcal{S} a stack, the category of stack morphisms $X \rightarrow \mathcal{S}$ and 2-morphisms is equivalent to the category of objects of \mathcal{S} lying over X and morphisms of \mathcal{S} lying over $\text{id} : X \rightarrow X$.)

(Another exercise: show that the category of automorphisms of BG is equivalent to the category of G -bitorsors, that is, G -bundles over a point equipped with a left G -action commuting with the usual right G -action. Hint: define a G -bundle over a stack and observe that any functor on G -bundles over spaces extends canonically to G -bundles over stacks; then consider the image of the tautological G -bundle over BG .)

As a consequence of the mischief, many of the familiar concepts we have in the category of spaces extend to stacks in a more convoluted fashion than one might expect. The basic point is that, instead of just requiring that two objects be equal, we have to choose an isomorphism. We give four key examples.

(1) **The fibered product.** If \mathcal{R} and \mathcal{R}' are stacks with morphisms F and F' to \mathcal{S} , the fibered product $\mathcal{R} \times_{\mathcal{S}} \mathcal{R}'$ consists of triples: an object R of \mathcal{R} , an object R' of \mathcal{R}' , and a choice of an isomorphism $F(R) \rightarrow F'(R')$. (Exercise: express Cadman's stack of n th roots as a fibered product. Another exercise: a 2-automorphism of F induces an automorphism of the fibered product.)

(2) **Commutative diagrams.** A diagram of stack morphisms isn't just commutative:

we have to make it so by choosing a 2-isomorphism. With a triangle of stacks, for instance, we write the symbol $\Downarrow \Theta$ inside the triangle:

$$\begin{array}{ccc}
 & \mathcal{R} & \\
 \nearrow & \Downarrow \Theta & \searrow \\
 \mathcal{Q} & \longrightarrow & \mathcal{S}
 \end{array}$$

to indicate that there is a 2-isomorphism Θ between the two stack morphisms $\mathcal{Q} \rightarrow \mathcal{S}$. When four such triangles with 2-isomorphisms fit together to form a tetrahedron, there is a natural compatibility condition between the 2-isomorphisms (which we leave to the alert reader to work out). If it is satisfied, the tetrahedron is said to be *commutative*.

(3) **Group actions on stacks.** Let Γ be a finite group. A Γ -action on a stack \mathcal{S} consists not only of morphisms $F_\gamma : \mathcal{S} \rightarrow \mathcal{S}$ for each $\gamma \in \Gamma$ (with $F_e := \text{id}$), but also of 2-isomorphisms $\Theta_{\gamma, \gamma'} : F_\gamma F_{\gamma'} \Rightarrow F_{\gamma\gamma'}$ such that, for any three $\gamma, \gamma', \gamma'' \in \Gamma$, the four 2-isomorphisms $\Theta_{\gamma, \gamma'}$, $\Theta_{\gamma', \gamma''}$, $\Theta_{\gamma\gamma', \gamma''}$, and $\Theta_{\gamma, \gamma'\gamma''}$ form a commutative tetrahedron in the sense alluded to above.

(Exercise: show that the category of Γ -actions on BG is equivalent to the category of extensions of Γ by G . Hint: use the previous exercise.)

(4) **Gluing of stacks.** Let $X = \bigcup X_\alpha$ be a space with an open cover. Recall that a collection of spaces $\pi_\alpha : S_\alpha \rightarrow X_\alpha$ may be glued along the open subsets $S_{\alpha\beta} = \pi_\alpha^{-1}(X_\beta)$ using homeomorphisms $f_{\alpha\beta} : S_{\beta\alpha} \rightarrow S_{\alpha\beta}$ (with $f_{\alpha\alpha} := \text{id}$), provided that they satisfy $f_{\alpha\beta} f_{\beta\gamma} f_{\gamma\alpha} = \text{id}$ on the triple overlaps. Here is the analogous statement for stacks. If \mathcal{S}_α are stacks over X_α , then they may be glued along $\mathcal{S}_{\alpha\beta} = \mathcal{S}_\alpha \times_X X_\beta$ using isomorphisms $F_{\alpha\beta} : \mathcal{S}_{\beta\alpha} \rightarrow \mathcal{S}_{\alpha\beta}$ with $F_{\alpha\alpha} = \text{id}$, provided that there exist 2-isomorphisms $\Theta_{\alpha\beta\gamma} : F_{\alpha\beta} F_{\beta\gamma} F_{\gamma\alpha} \Rightarrow \text{id}$ which in turn form a commutative tetrahedron over the quadruple overlaps.

In the last two examples, the choices of 2-isomorphisms had to satisfy a further condition, namely the commutativity of a tetrahedron. One might ask: why is this adequate? Why isn't some further choice of 3-isomorphisms necessary, and so on? The answer is that categories aren't the most abstract possible structure. In a category, the collection of morphisms between two fixed objects is assumed to be a set. Consequently, it is meaningful to speak of two given 2-isomorphisms as being equal (in contrast to 1-morphisms), since a 2-morphism $F \Rightarrow G$ consists of an element of the set of morphisms $F(C) \rightarrow G(C)$ for each object C .

One can, of course, define a more abstract entity, a 2-category, where even the morphisms between fixed objects merely comprise a category. Continuing recursively, one

can even define 3-categories, 4-categories, and so on, with their corresponding 2-stacks, 3-stacks, 4-stacks. . . Luckily, we will not have to enter this dizzying hall of mirrors.

Definition of gerbes

Let's return to the definition of stacks given a while back, and to the principal example BG . This stack actually satisfies two more properties, clearly analogous to (c) and (d):

(c'): **Local existence of bundles:** given any space Y , there is an open cover U_α of Y such that U_α is the base space of a G -bundle.

(d'): **Local existence of bundle isomorphisms:** given two bundles P and P' over Y , there is an open cover U_α such that $P|_{U_\alpha} \cong P'|_{U_\alpha}$.

Of course, (c') could not be more trivial for BG , since the trivial cover and the trivial bundle will do. However, the relative versions of both properties are interesting.

A stack \mathcal{S} over a space X is said to be a *gerbe* over X if:

(c') for any $f : Y \rightarrow X$ there is an open cover U_α of Y so that there exists an object of \mathcal{S} lying over each restriction $f|_{U_\alpha}$; and

(d') for any $f : Y \rightarrow X$ and any two objects P, P' of \mathcal{S} lying over f , there exists an open cover U_α of Y so that $P|_{U_\alpha} \cong P'|_{U_\alpha}$.

For example, BG is a gerbe over a point. We wish to exhibit nontrivial examples of gerbes over larger spaces.

The gerbe of liftings

To do this, recall first that for any homomorphism $\rho : G \rightarrow H$ of Lie groups, one defines the *extension of structure group* of a principal G -bundle $P \rightarrow X$ to be the twisted quotient $P_\rho = (P \times H)/G$, where G acts on H via ρ . It is a principal H -bundle over X . As our main example, let $\rho : \text{GL}(n) \rightarrow \text{PGL}(n)$ be the projection; then extension by ρ takes a vector bundle to its projectivization. (Here we have intentionally blurred the distinction between the equivalent categories of vector bundles and of frame bundles.)

Now let X be a topological space and P a principal H -bundle. Consider the category of triples consisting of (i) a map $f : Y \rightarrow X$; (ii) a principal G -bundle $Q \rightarrow Y$; (iii) an isomorphism $Q_\rho \rightarrow f^*P$. It is easily verified that this is a gerbe B over X ; call it the *gerbe of liftings* of P . In the main example, Q is a vector bundle over Y whose projectivization is identified with the pullback of a given projective bundle P .

We recognize these triples, don't we? They resemble the triples defining the fibered

product two sections back. Indeed, extension of structure group by ρ defines a natural transformation from G -bundles to H -bundles, hence a morphism $B\rho : BG \rightarrow BH$; on the other hand, P defines a morphism $X \rightarrow BH$, and our gerbe of liftings is nothing but $X \times_{BH} BG$.

It is easiest to understand this gerbe in the case where ρ is surjective, so that we have a short exact sequence

$$1 \longrightarrow A \xrightarrow{\sigma} G \xrightarrow{\rho} H \longrightarrow 1$$

with A normal. Consider first the case where P is trivialized. Then Q is a principal G -bundle with Q_ρ trivialized, and this precisely means that its structure group is reduced to A , that is, we get a bundle R with $R_\sigma = Q$. Conversely any such R gives Q with Q_ρ trivialized. Hence in this case the gerbe $B \cong X \times BA$.

On the other hand, if P is nontrivial, the gerbe may not be such a product. For example, if P is a projective bundle which does not lift to a vector bundle, then B has no global objects over the identity $X \rightarrow X$, but $X \times BA$ does.

In light of all this, a gerbe of liftings is a locally trivial bundle, in the category of stacks, with base X and fiber BA . At least morally speaking, one would like to say that H acts on BA , and that the gerbe is the associated BA -bundle to P . (Exercise: prove this when H is finite. Hint: use the previous exercise.) However, general group actions in the category of stacks turn out to be very slippery.

If the extension of groups is *central*, that is, $A \subset Z(G)$, then things become much simpler, at least at a conceptual level. To begin with, A must be abelian, and “an abelian group is a group in the category of groups,” that is, the group operations of multiplication and inversion are group homomorphisms $A \times A \rightarrow A$ and $A \rightarrow A$, respectively. Consequently, there are good notions of tensor product, and of dual, for A -bundles: namely, the extension of structure group by these homomorphisms. This in turn implies that there are natural morphisms $BA \times BA \rightarrow BA$ and $BA \rightarrow BA$ making BA into an abelian *group stack* in some sense. The central extension can be regarded as a principal A -bundle over H determining a morphism $H \rightarrow BA$, and this morphism is a homomorphism of group stacks. The gerbe of liftings is therefore a principal BA -bundle in this case. However, all this is not as easy to formulate rigorously as it seems, as the precise definition of a group stack is very confusing: associativity need not hold exactly, but only up to 2-isomorphisms which themselves must satisfy compatibility conditions. . .

The lien of a gerbe

Roughly speaking, an arbitrary gerbe may be described as in the previous section,

except that A , instead of being a fixed group, may be a sheaf of groups on X . This sheaf is called the *lien* or *band* of the gerbe. However, since nonabelian phenomena introduce some subtleties, we will discuss only the case analogous to the central extension of the last paragraph. This completely obscures the nonabelian motivation of the founders of the subject, but it is nevertheless sufficient for our purposes.

So let \mathcal{F} be a sheaf of abelian groups over X . An \mathcal{F} -torsor is a sheaf of sets over X equipped with an atlas of local isomorphisms to \mathcal{F} whose transition functions are given by multiplication by sections of \mathcal{F} . An \mathcal{F} -isomorphism of two \mathcal{F} -torsors is an isomorphism of sheaves locally given by multiplication by sections of \mathcal{F} .

Hence an \mathcal{F} -torsor is acted on by \mathcal{F} itself, and indeed is locally isomorphic to \mathcal{F} as an \mathcal{F} -sheaf, but without a choice of an identity element. For example, if \mathcal{F} is the sheaf of continuous functions with values in an abelian group A , then an \mathcal{F} -torsor is a principal A -bundle. Or if \mathcal{F} is the sheaf of locally constant functions with values in A , then an \mathcal{F} -torsor is a flat A -bundle.

There is a binary operation on \mathcal{F} -torsors taking L and L' to $(L \times_X L')/\mathcal{F}$ (with the antidiagonal action), which we denote $L \otimes L'$. For principal A -bundles, it agrees with the tensor product defined before.

Notice, if L, L' are fixed \mathcal{F} -torsors, that the sheaf of \mathcal{F} -isomorphisms $\text{Isom}(L, L')$ is itself an \mathcal{F} -torsor, and the sheaf of \mathcal{F} -automorphisms $\text{Aut } L = \text{Isom}(L, L)$ is canonically isomorphic to \mathcal{F} itself.

The collection of all \mathcal{F} -torsors forms a stack, indeed a gerbe, $B\mathcal{F}$ over X . More precisely, an object of $B\mathcal{F}$ consists of a map $g : Y \rightarrow X$ and a $g^*\mathcal{F}$ -torsor.

An \mathcal{F} -gerbe, then, is defined analogously to a torsor: it is a gerbe over X equipped with an atlas of local isomorphisms to $B\mathcal{F}$ whose transition functions are given by tensor product by sections of $B\mathcal{F}$, that is, torsors on the double overlaps.

For example, the gerbe of liftings of a central extension of H by A is an A -gerbe.

An \mathcal{F} -morphism of \mathcal{F} -gerbes is defined in the obvious way, as is an \mathcal{F} -2-morphism. To simplify notation, from now on *morphism* will always refer to an \mathcal{F} -morphism where \mathcal{F} -gerbes are concerned, and likewise for 2-morphisms. Note that this is a nontrivial restriction: for example, passage to the dual defines an automorphism of BA , but not an A -automorphism. (Exercise: in terms of the previous exercises, A -automorphisms correspond to the subcategory of bitorsors isomorphic to the trivial one.)

With this convention, an automorphism of an \mathcal{F} -gerbe, more or less by definition, is given by $L \otimes$ (that is, tensor product with L) for a fixed \mathcal{F} -torsor L . This induces

an equivalence of categories, so the 2-morphisms $L \otimes \Rightarrow L' \otimes$ (of gerbe automorphisms) correspond to morphisms $L \rightarrow L'$ (of torsors). In particular, the 2-automorphisms $L \otimes \Rightarrow L \otimes$ correspond naturally to sections of \mathcal{F} itself.

We can recover the lien from the gerbe. Suppose we are given a gerbe all of whose objects have abelian automorphism groups. Then the sheaves of automorphisms of any two objects are canonically isomorphic, so they glue together to give a globally defined sheaf \mathcal{F} of abelian groups. It is easy to show that the gerbe is then an \mathcal{F} -gerbe. However, if some automorphism groups are nonabelian, this gives rise to the complications ominously alluded to above.

Classification of gerbes

At last we are in the position to state a classification result. To avoid complications we confine ourselves to the abelian case, as before.

Theorem (Giraud). The group of isomorphism classes of \mathcal{F} -gerbes is isomorphic to $H^2(X, \mathcal{F})$.

Sketch of proof: Trivialize the gerbe on a cover by open sets X_α . The transition functors $F_{\alpha\beta}$ then correspond to \mathcal{F} -torsors $L_{\alpha\beta}$ on $X_{\alpha\beta} = X_\alpha \cap X_\beta$. After refining the cover if necessary, we may choose trivialisations of these torsors. But, on the triple overlaps $X_{\alpha\beta\gamma}$, we also have the trivialisations of $L_{\alpha\beta} \otimes L_{\beta\gamma} \otimes L_{\gamma\alpha}$ given by the 2-isomorphisms $F_{\alpha\beta} F_{\beta\gamma} F_{\gamma\alpha} \Rightarrow \text{id}$. These then determine sections of \mathcal{F} on $X_{\alpha\beta\gamma}$ which constitute a Čech 2-cochain. The tetrahedron condition on the 2-isomorphisms precisely implies that this is closed; and changing the trivialisations of the torsors $L_{\alpha\beta}$ adds an exact cocycle.

Allowing the base space to be a stack

A general philosophy is that everything that can be done for manifolds should also be attempted for orbifolds. More broadly, everything that can be done for spaces should also be attempted for stacks. In this spirit, we describe here what is meant by a sheaf, a torsor, or a gerbe whose base space is itself a stack. The definition resembles that of a characteristic class.

Let \mathcal{S} be a stack. A *sheaf over \mathcal{S}* is a functor \mathcal{F} , over the category of topological spaces, from \mathcal{S} to the category of sheaves. That is, it assigns to every object of \mathcal{S} over Y a sheaf F over Y , and to every morphism of objects over $g : Y \rightarrow Y'$ an isomorphism $F \cong g^* F'$. A *torsor* for a given sheaf is defined similarly.

However, we won't define a *gerbe over \mathcal{S}* in the same way, for gerbes (like all stacks)

don't just constitute a category, but rather a 2-category. Instead, a gerbe B over \mathcal{S} is a stack over \mathcal{S} such that for all objects of \mathcal{S} over Y , the fibered product $Y \times_{\mathcal{S}} B$ is a gerbe over Y . An \mathcal{F} -gerbe is defined similarly for a sheaf \mathcal{F} over \mathcal{S} .

(Exercise: a sheaf of abelian groups over BG corresponds naturally to an abelian group A with a G -action by group automorphisms. A gerbe over BG with lien A corresponds naturally to an extension of G by A so that the action of G on A in the extension is the given one.)

Definition of orbifolds

We want to conclude this lecture with a description of the Strominger-Yau-Zaslow proposal for mirror symmetry. To do so, we need two more definitions: of orbifolds and of twisted vector bundles.

First, orbifolds. Roughly speaking, these are stacks locally isomorphic to a quotient of a manifold by a finite group. Readers are cautioned that this definition may differ in a few respects from those in the literature.

Let S be a space. An *orbispace* \mathcal{S} with *coarse moduli space* S is a stack over S so that there exists an open cover $S = \cup S_{\alpha}$ satisfying $S_{\alpha} \times_S \mathcal{S} \cong [X_{\alpha}/\Gamma_{\alpha}]$, where Γ_{α} is a finite group, and the induced map $X_{\alpha}/\Gamma_{\alpha} \rightarrow S_{\alpha}$ of spaces is a homeomorphism. It is an *orbifold* if each X_{α} is a manifold.

A *smooth structure* on an orbifold is a choice of smooth structure on each X_{α} so that X_{α} and X_{β} induce the same smooth structure on the covering space $X_{\alpha} \times_S X_{\beta}$. (Exercise: this implies that each Γ_{α} acts smoothly.) A *complex structure* on an orbifold is defined similarly.

Twisted vector bundles

Let B be a gerbe over X with structure group $U(1)$. As we have seen, B is a fiber bundle over X with fiber $BU(1)$. A *twisted vector bundle* for B is a vector bundle over B whose restriction to each fiber is a representation of $U(1)$ (using the last exercise) of pure weight 1.

These are called "twisted" since they can be regarded as locally trivial on open sets $X_{\alpha} \subset X$, with transition functions $f_{\alpha\beta} : X_{\alpha\beta} \rightarrow GL(n)$. Instead of the usual cocycle condition, we require that $f_{\alpha\beta} f_{\beta\gamma} f_{\gamma\alpha} = b_{\alpha\beta\gamma} \text{id}$ where b is a cocycle representative for the isomorphism class of B in $H^2(X, U(1))$.

The same applies to flat gerbes and flat vector bundles.

Twisted vector bundles for a given gerbe clearly form an abelian category, so a *twisted K-theory* may be defined. If the gerbe is trivial, we recover ordinary K-theory. However, twisted K-theory for a fixed gerbe does not admit a tensor product: rather, we would have to sum over all gerbes, or at least all powers of the fixed one.

Strominger-Yau-Zaslow

The proposal of Strominger-Yau-Zaslow on mirror symmetry can be described in the language of gerbes and orbifolds. Their remarkable idea is that mirror partners should be Calabi-Yau orbifolds M and \hat{M} of complex dimension n which admit proper maps to the same orbifold Z of real dimension n :

$$\begin{array}{ccc} M & & \hat{M} \\ \pi \searrow & & \swarrow \hat{\pi} \\ & Z & \end{array}$$

so that, if z is a regular value of π and $\hat{\pi}$, the fibers $L_z = \pi^{-1}(z)$ and $\hat{L}_z = \hat{\pi}^{-1}(z)$ are special Lagrangian tori which are in some sense dual to each other. Here *Lagrangian* means Lagrangian with respect to the Kähler form, and *special* means that the imaginary part of the nonzero holomorphic n -form that exists on any Calabi-Yau vanishes on the torus.

The duality between the tori can be required in a strong sense originally envisioned by SYZ, or in a more general sense proposed by Hitchin and involving flat gerbes.

In the original formulation of SYZ, the maps π and $\hat{\pi}$ are assumed to have special Lagrangian sections, giving a basepoint for each L_z and \hat{L}_z . This canonically makes them into Lie groups, since a choice of a basis for T_z^*Z determines, via the Kähler form, n commuting vector fields on L_z and \hat{L}_z whose flows define a diffeomorphism to $(S^1)^n$. We then ask for isomorphisms of Lie groups (smoothly depending on z)

$$\hat{L}_z \cong \text{Hom}(\pi_1(L_z), \text{U}(1))$$

and vice versa. That is, the tori parametrize isomorphism classes of flat $\text{U}(1)$ -bundles on each other.

This formulation was generalized by Hitchin to the case of torus families without sections. It turns out that the absence of a section on M reflects the non-triviality of a gerbe on \hat{M} , and vice versa.

So suppose now that M (resp. \hat{M}) is equipped with a flat orbifold $\text{U}(1)$ -gerbe B (resp. \hat{B}) trivial on the fibers of π (resp. $\hat{\pi}$). We can now ask each torus to parametrize isomorphism classes of *twisted* flat $\text{U}(1)$ -bundles on the other torus. More than that,

we can ask $B|_{L_z}$ to be identified with the stack of twisted flat $U(1)$ -bundles on \hat{L}_z , and vice versa. Of course, we want this identification to depend smoothly on $z \in Z$, and we leave it to the reader to specify exactly what this means.

It is extremely difficult to find examples of special Lagrangian tori on Calabi-Yau manifolds. The consensus in the field seems to be that the requirements of SYZ as stated above are too stringent, and that perhaps they must only be satisfied in some limiting sense, say near the “large complex structure limit” in the moduli space of complex structures on the Calabi-Yau. However, the author has studied a few cases where for relatively straightforward reasons (because the metric is, say, hyperkähler or flat) the requirements of SYZ, in the gerbe sense, are seen to be satisfied precisely.

Lecture 3: Orbifold cohomology and its relatives

What kind of cohomology can be defined for orbifolds? The simplest answer is given in the first section below. Cohomology can be defined for any coefficient ring, or indeed, any sheaf on a stack, in such a way that, if \mathcal{M} is an orbifold with coarse moduli space M ,

$$H^*(\mathcal{M}, \mathbf{C}) = H^*(M, \mathbf{C}).$$

However, it has been known for a long time that, for the purposes of string theory, mirror symmetry, and so on, a more refined form of cohomology is preferable. This is the *orbifold cohomology* theory $H_{\text{orb}}^*(\mathcal{M}, \mathbf{C})$, which as a vector space is

$$H_{\text{orb}}^*(\mathcal{M}) = H^*(\mathcal{J}\mathcal{M}).$$

Here $\mathcal{J}\mathcal{M}$ is the so-called *inertia stack*, to be introduced shortly.

We did not specify what coefficient ring to take on the right-hand side, but suppose we choose the Novikov ring from Lecture 1, which is the coefficient ring for Floer cohomology. Then orbifold cohomology admits a quantum cup product whose associativity is a deep and significant fact. Indeed, this is the main reason for studying orbifold cohomology. However, we won't delve into the construction of the product or the proof of associativity. Rather, after defining orbifold cohomology, we will introduce some of its variants and relatives — the version with a flat $U(1)$ -gerbe, for example, and the *Fantechi-Göttsche ring* defined for a global quotient $[X/\Gamma]$ — and then explain how we expect all of these structures to be related to Floer theory.

Cohomology of sheaves on stacks

Just as a sheaf \mathcal{F} on a stack \mathcal{S} is a rule assigning to each object S of \mathcal{S} over Y

a sheaf F_S over Y , we can define a *cohomology class* for \mathcal{F} to be a rule assigning to each S an element of $H^*(Y, F_S)$ in a manner compatible with pullbacks. In more fancy categorical language, this is the limit of the functor $H^* \circ \mathcal{F}$ from \mathcal{S} to the category of abelian groups. It is clear that this is a group provided that it is a set! For reasonable sheaves and stacks, this will be true.

For example, if $[X/\Gamma]$ is an orbifold with a sheaf \mathcal{F} regarded as an equivariant sheaf on X , then clearly

$$H^*([X/\Gamma], \mathcal{F}) = H^*(X, \mathcal{F})^\Gamma,$$

where the superscript on the right-hand side denotes the invariant part. If K is a field of characteristic 0, then a theorem of Grothendieck gives a canonical isomorphism

$$H^*(X, K)^\Gamma \cong H^*(X/\Gamma, K),$$

so the cohomology of a global quotient (with coefficients in K) coincides with the cohomology of its coarse moduli space.

We can then conclude that the same is true for an arbitrary orbifold \mathcal{M} by using the Mayer-Vietoris spectral sequence. Use a countable atlas where every open set is a global quotient $[X_\alpha/\Gamma_\alpha]$; then the natural map $[X/\Gamma] \rightarrow X/\Gamma$ induces isomorphisms $H^*([X_\alpha/\Gamma_\alpha], K) \cong H^*(X_\alpha/\Gamma_\alpha, K)$, and similarly for double overlaps, triple overlaps, and so on. Hence it induces isomorphisms between the double complexes that appear in the Mayer-Vietoris spectral sequences for \mathcal{M} and its coarse moduli space M , and so we conclude that it induces an isomorphism

$$H^*(\mathcal{M}, K) \cong H^*(M, K)$$

when K is a field of characteristic 0.

(Exercise: show that for an arbitrary topological group G and coefficient ring R , there is a natural isomorphism $H^*([X/G], R) \cong H_G^*(X, R)$ where the right-hand side is equivariant cohomology.)

The inertia stack

Let \mathcal{S} be a stack. We can associate to it another stack, the *inertia stack* $\mathcal{I}\mathcal{S}$. This is defined to be the stack whose objects over Y are pairs consisting of an object of \mathcal{S} over Y and an automorphism of that object over the identity on Y , and whose morphisms are commutative squares.

If the stack is a quotient by a finite group, the inertia stack can be described explicitly.

Proposition. There is a natural isomorphism

$$\mathcal{J}[X/\Gamma] \cong \bigsqcup_{[\gamma]} [X^\gamma/C(\gamma)],$$

where the disjoint union runs over conjugacy classes in Γ , $X^\gamma = \{x \in X \mid \gamma x = x\}$ is the fixed-point set, and $C(\gamma)$ denotes the centralizer of $\gamma \in \Gamma$.

Sketch of proof. An object of $[X/\Gamma]$ consists of a principal Γ -bundle $P \rightarrow Y$ together with a Γ -equivariant map $P \rightarrow X$. Hence an object of $\mathcal{J}[X/\Gamma]$ consists of those two things plus an automorphism of P preserving the equivariant map. Since Γ is discrete, any automorphism is given by the right action of some $\gamma \in \Gamma$ commuting with the monodromy group, that is, the image of $\pi_1(Y) \rightarrow \Gamma$. Thus the structure group is reduced to $C(\gamma)$, so we get a principal $C(\gamma)$ -bundle and an equivariant map to X which, since it is preserved by γ , must have image in X^γ .

It follows directly that, if \mathcal{M} is an orbifold, then so is $\mathcal{J}\mathcal{M}$ (though with components of different dimensions).

(Exercise: prove that there is a natural isomorphism $\mathcal{J}\mathcal{S} \cong \mathcal{S} \times_{\mathcal{S} \times \mathcal{S}} \mathcal{S}$ for any stack \mathcal{S} .)

Orbifold cohomology

Henceforth, assume that our orbifold \mathcal{M} is *Kähler*, that is, locally $[X_\alpha/\Gamma_\alpha]$ with X_α a Kähler manifold so that X_α and X_β induce the same Kähler structure on the covering space $X_\alpha \times_{\mathcal{M}} X_\beta$. We may then define the *orbifold cohomology* of M to be

$$H_{\text{orb}}^*(\mathcal{M}, \mathbf{C}) = H^*(\mathcal{J}\mathcal{M}, \mathbf{C}).$$

To be more precise, the grading on the orbifold cohomology is not the usual one. Rather, the different connected components have the degrees of their cohomology shifted by different amounts. For a connected component of $[X^\gamma/C(\gamma)] \subset \mathcal{J}[X/\Gamma]$, the so-called *fermionic shift* is defined as follows. Since γ has finite order, it acts on the tangent space $T_x X$ at a point $x \in X^\gamma$ with weights $e^{2\pi i w_1}, \dots, e^{2\pi i w_n}$ for some rational numbers $w_1, \dots, w_n \in [0, 1)$. (This is why we need \mathcal{M} Kähler, or at least complex: so that the w_j will be well defined.) Then let $F(\gamma) = \sum_j w_j$. The notation suggests that $F(\gamma)$ is the same on all connected components of $X^\gamma/C(\gamma)$, which is true in most interesting cases. In any case, the grading of the cohomology of the component of $X^\gamma/C(\gamma)$ containing x should be increased by $2F(\gamma)$. For example, the correct grading for $H_{\text{orb}}^*[X/\Gamma]$ is

$$H_{\text{orb}}^k[X/\Gamma] = \bigoplus_{[\gamma]} H^{k-2F(\gamma)}(X^\gamma, \mathbf{C})^{C(\gamma)}.$$

Warning: the fermionic shift may not be an integer! But it will be in many interesting cases, like that of a global quotient $[X/\Gamma]$ provided that the canonical bundle of X has a nowhere vanishing section preserved by Γ (which we might call a *Calabi-Yau orbifold*).

(Exercise: prove that the orbifold Betti numbers of a compact complex orbifold satisfy Poincaré duality. If this is too hard, do it only for $[X/\Gamma]$.)

As we mentioned before, the main interest of orbifold cohomology is that $H_{\text{orb}}^*(\mathcal{M}, \bar{\Lambda}) = H_{\text{orb}}^*(\mathcal{M}, \mathbf{C}) \otimes_{\mathbf{C}} \bar{\Lambda}$ admits an associative quantum product, where $\bar{\Lambda}$ is the Novikov ring from Lecture 1. Indeed, stacks of stable maps to the orbifold \mathcal{M} have been constructed, as discussed in the notes of Abramovich in this volume, and their evaluation maps naturally take values in \mathcal{JM} . So Gromov-Witten invariants provide structure constants for a quantum cup product on $H^*(\mathcal{JM})$.

There are, of course, algebra homomorphisms $\mathbf{C} \rightarrow \bar{\Lambda} \rightarrow \mathbf{C}$ (the latter given by taking the constant term), and it is tempting to use these, together with the quantum product on $H_{\text{orb}}^*(\mathcal{M}, \bar{\Lambda})$, to define a product on $H_{\text{orb}}^*(\mathcal{M}, \mathbf{C})$. This is the so-called *orbifold product*, which in fact slightly predates the orbifold quantum product. It involves only the contributions of stable maps of degree 0. Nevertheless, it usually differs from the standard cup product, as there usually exist stable maps which have degree 0 (indeed, their images in the coarse moduli space are just points) but whose evaluations at different marked points lie in different components of \mathcal{JM} .

Twisted orbifold cohomology

Suppose, now, that we have a flat $U(1)$ -gerbe B on our orbifold \mathcal{M} . This immediately induces a flat $U(1)$ -torsor on \mathcal{JM} . Indeed, each object of \mathcal{JM} consists of an object of \mathcal{M} (say over Y) and an automorphism of that object (over $\text{id} : Y \rightarrow Y$), hence an automorphism of the $U(1)$ -gerbe $Y \times_{\mathcal{M}} B$ over Y , hence a $U(1)$ -torsor on Y .

Let LB be the flat complex line bundle over \mathcal{JM} associated to this torsor. Now define the *twisted orbifold cohomology* to be simply

$$H_{\text{orb}}^*(\mathcal{M}, B) = H^*(\mathcal{JM}, LB),$$

where the right-hand side refers to cohomology with local coefficients.

The degree should be again adjusted by the fermionic shift, which is the same as before. For a trivial gerbe, we recover the previous notion of orbifold cohomology.

Let's spell out what this is for a global quotient $\mathcal{M} = [X/\Gamma]$. The line bundle LB over \mathcal{JM} can be regarded as a collection, indexed by $\gamma \in \Gamma$, of $C(\gamma)$ -equivariant line

bundles $L_\gamma B$ over X^γ ; that is,

$$L_\gamma B = LB|_{[X^\gamma/C(\gamma)]}.$$

Then

$$H_{\text{orb}}^*(\mathcal{M}, B) = \bigoplus_{[\gamma]} H^*(X^\gamma, L_\gamma B)^{C(\gamma)}.$$

Again, there should be a notion of quantum product on this twisted orbifold cohomology after we tensor with the Novikov ring. What is needed is to show that the flat line bundles agree under the pullbacks to stable map spaces by the relevant evaluation maps.

The case of discrete torsion

One particularly attractive case has received the most attention in the literature: that of a global quotient $[X/\Gamma]$ with a flat $U(1)$ -gerbe pulled back Γ -equivariantly from a point, that is, a flat $U(1)$ -gerbe pulled back from $B\Gamma$. These are classified, as we saw, by $H^2(B\Gamma, U(1))$. This group is known in the physics literature as the *discrete torsion*, and in the mathematics literature as the *Schur multiplier*. It may be interpreted (and computed) as the group cohomology of Γ with coefficients in the trivial module $U(1)$. It can also be regarded as classifying central extensions

$$1 \longrightarrow U(1) \longrightarrow \tilde{\Gamma} \longrightarrow \Gamma \longrightarrow 1.$$

What makes such gerbes attractive is, firstly, that they are relatively plentiful: for example, $H^2(\mathbf{Z}_n \times \mathbf{Z}_n, U(1)) \cong \mathbf{Z}_n$. But also, the flat line bundles $L_\gamma B$ can be calculated over a point and then pulled back to X^γ . Consequently, the underlying line bundles are trivial; only the action of the centralizer $C(\gamma)$ is nontrivial. In the literature, this is sometimes called the *phase*: a homomorphism $C(\gamma) \rightarrow U(1)$.

One can easily show, if $\langle \cdot, \cdot \rangle : \Gamma \times \Gamma \rightarrow U(1)$ is a 2-cocycle representing an element B of discrete torsion in group cohomology, that the phase is given by

$$\delta \mapsto \frac{\langle \gamma, \delta \rangle}{\langle \delta, \gamma \rangle}.$$

Hence the summand $H^*(X^\gamma, L_\gamma B)^{C(\gamma)}$ that appears in the definition of $H_{\text{orb}}^*([X/\Gamma], B)$ is simply the isotypical summand of $H^*(X^\gamma, \mathbf{C})$, regarded as a representation of $C(\gamma)$, that transforms according to the inverse of the phase above.

The Fantechi-Göttsche ring

In fact, for a global quotient $[X/\Gamma]$ there is supposed to be a larger ring, equipped with a Γ -action, so that the orbifold cohomology can be recovered as the invariant part. This is the *Fantechi-Göttsche ring*.

Additively it is quite simple: just take

$$HFG^*(X, \Gamma) = \bigoplus_{\gamma \in \Gamma} H^*(X^\gamma, \bar{\Lambda}).$$

Notice that the sum runs over group elements, not just conjugacy classes.

As a representation of Γ it is also quite simple: for each $\delta \in \Gamma$, there is a natural isomorphism $X^\gamma \rightarrow X^{\delta\gamma\delta^{-1}}$, hence a pullback on the cohomology that induces an automorphism of $HFG^*(X, \Gamma)$. These fit together to give a Γ -action that acts on the Γ -grading by conjugation.

The nontrivial part is the quantum multiplication. The claim is that there are spaces, akin to those of stable maps, but somehow rigidified so that Γ acts nontrivially on them, and so that the evaluation map goes to $\bigsqcup_\gamma X^\gamma$ instead of just the inertia stack. One should then use these spaces, as in the usual definition of quantum cohomology, to define maps $H^*(X^\gamma, \bar{\Lambda}) \otimes H^*(X^{\gamma'}, \bar{\Lambda}) \rightarrow H^*(X^{\gamma\gamma'}, \bar{\Lambda})$.

These spaces, and their virtual classes, are constructed by Fantechi and Göttsche for stable maps of degree 0. As a result, they obtain a ring with degree 0 terms only, whose invariant part carries the orbifold product. But there is every reason to expect a quantum product in this setting.

Twisting the Fantechi-Göttsche ring with discrete torsion

As the reader may be suspecting, we would like a version of the Fantechi-Göttsche ring which involves a flat unitary gerbe. Let's first indicate how to do this for an element of discrete torsion.

As before, represent our element of discrete torsion by a 2-cocycle $\langle \cdot, \cdot \rangle : \Gamma \times \Gamma \rightarrow U(1)$. Being closed under the differential means that for all $f, g, h \in \Gamma$,

$$\frac{\langle f, g \rangle \langle fg, h \rangle}{\langle f, gh \rangle \langle g, h \rangle} = 1.$$

Now for any two elements $a_g \in H^*(X^g)$ and $b_h \in H^*(X^h)$, regarded as summands of $HFG^*(X, \Gamma)$, we have the usual quantum Fantechi-Göttsche product $a_g \cdot b_h \in H^*(X^{gh})$. Now define a new product by

$$a_g * b_h = \langle g, h \rangle a_g \cdot b_h.$$

This need not be commutative or even super-commutative, but it is associative: in fact closedness precisely guarantees this.

The action of Γ on $HFG^*(X, \Gamma)$ given above is no longer a ring homomorphism for the $*$ product. Instead, we need to twist the action as follows: the action of $h \in \Gamma$ takes $H^*(X^g)$ to $H^*(X^{hgh^{-1}})$ by the same map as before, but multiplied by the rather odd factor

$$\frac{\langle h, g \rangle \langle hg, h^{-1} \rangle}{\langle h, h^{-1} \rangle}.$$

The justification for this is that first of all, it now acts by ring homomorphisms for the $*$ product, and second of all, the part invariant under all $h \in \Gamma$ is now twisted orbifold cohomology in the sense defined above.

Twisting it with an arbitrary flat unitary gerbe

Next, let's see how the previous section is a special case of putting in an equivariant flat $U(1)$ -gerbe. So let B be such a gerbe on X , equivariant under Γ , or equivalently, a gerbe on $[X/\Gamma]$. As before we get a flat line bundle $L_g B$ over X^g , with a lifting of the $C(g)$ -action. Additively, we define

$$HFG^*(X, \Gamma; B) = \bigoplus_{g \in \Gamma} H^*(X^g, L_g B),$$

where the terms on the right are cohomology with local coefficients.

As before, to extend the quantum Fantechi-Göttsche product to this twisted case, one would have to show that the flat line bundles agree under the pullbacks, by the relevant evaluation maps, to the spaces akin to those of stable maps. (Exercise: carry this out for degree 0 maps. This amounts to showing that when restricted to $X^{g,h} = X^g \cap X^h$, there is a natural isomorphism $L_{gh} B \cong L_g B \otimes L_h B$.)

There is also, of course, a natural isomorphism induced by $h \in \Gamma$,

$$H^*(X^g, L_g B) \longrightarrow H^*(X^{hgh^{-1}}, L_{hgh^{-1}} B),$$

and so Γ acts on $HFG^*(X, \Gamma; B)$, and the invariant part is the twisted orbifold cohomology. Let's check that, in the case when B is discrete torsion, this isomorphism is simply the one induced by the identification $X^g \rightarrow X^{hgh^{-1}}$, times the rather odd factor.

Let $\pi : \tilde{\Gamma} \rightarrow \Gamma$ be the central extension determined by B . The automorphism of the category $BU(1)$ induced by multiplication by g is, of course, just tensorization by the $U(1)$ -torsor $\pi^{-1}(g)$: that is, there is a canonical isomorphism $L_g B \cong \pi^{-1}(g)$. Now, once a cocycle representative is chosen for B , we can identify $\tilde{\Gamma}$ with the product

$\Gamma \times U(1)$, and hence $\pi^{-1}(g)$ with $U(1)$, but this does not respect the group operation. Nevertheless, let's write $\{g, t\}$ for an element of $\Gamma \times U(1) = \tilde{\Gamma}$. In terms of this, the map $\pi^{-1}(g) \rightarrow \pi^{-1}(hgh^{-1})$ is given by conjugation by any element of $\pi^{-1}(h)$, so we might as well take $\{h, 1\}$. Then what we need to compute is

$$\begin{aligned} \{h, 1\} \{g, t\} \{h, 1\}^{-1} &= \{h, 1\} \{g, t\} \{h^{-1}, 1/\langle h, h^{-1} \rangle\} \\ &= \{hg, \langle h, g \rangle t\} \{h^{-1}, 1/\langle h, h^{-1} \rangle\} \\ &= \left\{ hgh^{-1}, \frac{\langle h, g \rangle \langle hg, h^{-1} \rangle}{\langle h, h^{-1} \rangle} t \right\} \end{aligned}$$

which shows that, in terms of the identification of $\tilde{\Gamma}$ as a product, the action of h multiplies $\pi^{-1}(g)$ by the rather odd factor.

The loop space of an orbifold

It is high time to explain what all of these rings are supposed to have to do with Floer cohomology. The claim is that each one can be realized as the Morse cohomology of a symplectic action function on an appropriate analogue of the loop space. Associated to each flat $U(1)$ -gerbe will be a flat line bundle on the loop space, and we should take Floer cohomology with local coefficients. The multiplications defined on orbifold and Floer cohomology should then coincide.

Let's begin with the untwisted Fantechi-Göttsche cohomology. Here, all the pieces are already in place. Observe that, since Γ is a finite group, every element acts on X as a finite-order symplectomorphism, so according to the conjecture from Lecture 1, additively

$$HFG^*(X, \Gamma) = \bigoplus_{\gamma \in \Gamma} H^*(X^\gamma, \bar{\Lambda}) = \bigoplus_{\gamma \in \Gamma} HF^*(\gamma).$$

But both sides also carry a product: on the right, this is thanks to the linear map

$$HF^*(\gamma) \otimes HF^*(\gamma') \longrightarrow HF^*(\gamma\gamma')$$

discussed in Lecture 1. The conjecture is that *this Floer product agrees with the quantum Fantechi-Göttsche product*.

To express this in terms of loop spaces, let $L_\Gamma X = \bigsqcup_{\gamma \in \Gamma} L_\gamma X$, where on the right-hand side γ is regarded as a symplectomorphism of X and $L_\gamma X$ is the twisted loop space as in Lecture 1. The group Γ acts on $L_\Gamma X$ by $\delta \cdot \ell = \delta\ell$; this takes $L_\gamma X$ to $L_{\delta\gamma\delta^{-1}} X$. But since $\gamma^m = \text{id}$ for $m = |\Gamma|$, there is also an action of $S^1 = \mathbf{R}/m\mathbf{Z}$ by translating the parameter. We refer to this as *rotating* the twisted loops. This action commutes with that of Γ , and its moment map is exactly the symplectic action function. The

fixed-point set is $\bigsqcup_{\gamma \in \Gamma} X^\gamma$. So the Fantechi-Göttsche ring is supposed to be the Morse cohomology of $L_\Gamma X$ with respect to the action function, which is a perfect Bott-Morse function.

One relatively tractable aspect of this conjecture should be the grading. We have explained how the Fantechi-Göttsche ring is graded: by the usual grading on cohomology corrected by the fermionic shift. On the other hand, Floer cohomology also carries a grading. Under the proposed isomorphism, these gradings presumably agree. Recall, though, that the fermionic shift can be fractional: this is already the case for the obvious action of \mathbf{Z}_n on the Riemann sphere. We artfully evaded discussing the Floer grading for nontrivial symplectomorphisms, but it evidently would have to take account of this.

Now, let's move on to consider orbifold cohomology. For an orbifold \mathcal{M} , the space of maps $S^1 \rightarrow \mathcal{M}$, in the sense of stacks, can be regarded as a stack $L\mathcal{M}$ in a natural way. An object of $L\mathcal{M}$ over Y is, of course, nothing but an object of \mathcal{M} over $Y \times S^1$. Indeed, we wish to regard $L\mathcal{M}$ as an infinite-dimensional symplectic orbifold, just as the loop space of a manifold is an infinite-dimensional symplectic manifold. We won't attempt to justify this beyond observing that, for a global quotient, we have $L[X/\Gamma] = (L_\Gamma X)/\Gamma$.

Once again a circle acts on $L\mathcal{M}$, and now the fixed-point stack of the circle action can be identified with the inertia stack $\mathcal{J}\mathcal{M}$. Again, these statements have to be considered imprecise since we haven't defined circle actions on stacks. But it is clear what we mean in the case of a global quotient, where we just have a circle action commuting with the Γ -action, and the inertial stack is

$$\mathcal{J}[X/\Gamma] = \left(\bigsqcup_{\gamma \in \Gamma} X^\gamma \right) / \Gamma.$$

So our claim is, once again, that we should regard $H_{\text{orb}}^*(\mathcal{M})$ as the Morse cohomology of $L\mathcal{M}$, and that the product structures in the Floer and orbifold settings should coincide. Here new technical obstacles would present themselves, for we are asking to do Morse theory on an orbifold, which is problematic even in finite dimensions.

Nevertheless, at a heuristic level, our claims are certainly very plausible. Both the Floer and orbifold cohomologies are defined by "counting" holomorphic maps from a thrice-punctured sphere to the orbifold. The difference lies in what we do to make this formal definition into a mathematically rigorous count. In algebraic geometry, one has the machinery of Gromov-Witten theory, with virtual classes and so on.

In symplectic geometry, on the other hand, one has to perturb the equations and their solutions. As we saw, with a single symplectomorphism ϕ , to define $HF^*(\phi)$ one should perturb with exact Hamiltonians until the fixed points are isolated. For three

symplectomorphisms satisfying $\phi_1\phi_2\phi_3 = \text{id}$, to define the map $HF^*(\phi_1) \otimes HF^*(\phi_2) \rightarrow HF^*(\phi_3^{-1})$ one should presumably perturb all three simultaneously so that their product remains trivial, but so that all three have isolated fixed points. Any map from a thrice-punctured sphere to a global quotient has monodromy of this form, so this indicates how to define the Floer product on a global quotient. On a general orbifold the situation is not so clear. However, Gromov-Witten invariants of orbifolds have been defined in the symplectic literature.

Addition of the gerbe

Now suppose that M , a compact Kähler orbifold, carries a flat $U(1)$ -gerbe B .

Consider a map $\ell : S^1 \rightarrow M$. This induces a flat $U(1)$ -gerbe ℓ^*B on S^1 . This in turn induces a flat $U(1)$ -gerbe on the universal cover \mathbf{R} , together with an automorphism covering the translation $t \mapsto t + 1$. But any flat $U(1)$ -gerbe on \mathbf{R} is trivial, and the trivialization determines another automorphism covering $t \mapsto t + 1$. Comparing the two gives a $U(1)$ -torsor over a point.

The same construction works in families, so any map $S^1 \times Y \rightarrow M$ determines a flat $U(1)$ -torsor over Y . In particular, there is a flat $U(1)$ -torsor LB over LM . The isomorphism class of LB is the image of the isomorphism class of B under the transgression map $H^2(M, U(1)) \rightarrow H^1(LM, U(1))$ defined at the beginning of Lecture 2.

Now, let T be a *trinion*, a sphere minus three disjoint disks, and consider a map $T \rightarrow M$. Again this induces a flat $U(1)$ -gerbe on the universal cover \tilde{T} , but now (since $\pi_1(T)$ has three generators whose product is 1) this leads to three automorphisms f_1, f_2, f_3 of the trivial gerbe on a point and a 2-isomorphism $f_1f_2f_3 \Rightarrow \text{id}$. The 2-isomorphism induces a trivialization of the tensor product $L_1 \otimes L_2 \otimes L_3$ of the three torsors coming from the boundary components.

Again this works in families, so if Y is any space of maps from the trinion to M , we get a trivialization of $ev_1^*LB \otimes ev_2^*LB \otimes ev_3^*LB$, where the *evaluation maps* $ev_i : Y \rightarrow LM$ are given by restriction to the boundary circles. This is why the quantum product makes sense with local coefficients in LB : when we pull back classes by ev_1 and ev_2 and cup them together, they push forward under ev_3 to a class with the appropriate local coefficients. (Note that reversing the orientation of a circle will dualize the relevant torsor.)

The non-orbifold case

Let's see how this plays out in the case where M is simply a compact Kähler manifold.

The isomorphism classes of gerbes then sit in the long exact sequence

$$H^2(M, \mathbf{Z}) \rightarrow H^2(M, \mathbf{R}) \rightarrow H^2(M, U(1)) \rightarrow H^3(M, \mathbf{Z}) \rightarrow H^3(M, \mathbf{R}).$$

The map from integral to real cohomology has as kernel the torsion classes and as image a full lattice, so this boils down to

$$0 \rightarrow \frac{H^2(M, \mathbf{R})}{H^2(M, \mathbf{Z})} \rightarrow H^2(M, U(1)) \rightarrow \text{Tors } H^3(M, \mathbf{Z}) \rightarrow 0,$$

which of course splits, though not canonically. Consider first what happens as the gerbe B ranges over the torus $H^2(M, \mathbf{R})/H^2(M, \mathbf{Z})$. In this case the following notation is convenient: for any $\beta \in H_2(M, \mathbf{Z})$, write $B^\beta = \exp 2\pi i B(\beta) \in U(1)$. The torsor LB restricted to the constant loops $M \subset LM$ is, of course, canonically trivial. But, if $F : T \rightarrow M$ is any map from the trinion to M taking the boundary circles to constant loops, the trivialization of $\text{ev}_1^*LB \otimes \text{ev}_2^*LB \otimes \text{ev}_3^*LB$ does not agree with the canonical one. Rather, as is easily checked, they differ by the scalar factor B^β , where $\beta = F_*[T]$ is the homology class of F (well defined since F is constant on boundary components).

This introduces an additional weighting factor of B^β in the contributions of degree β holomorphic maps $T \rightarrow M$ to the Floer product. Since these are already weighted by q^β , we conclude that *the Floer products parametrized by $B \in H^2(M, \mathbf{R})/H^2(M, \mathbf{Z})$ can be all be obtained from the usual one by the change of variables $q \mapsto Bq$.*

In fact, this story extends to the full group $H^2(M, U(1))$, including $\text{Tors } H^3(M, \mathbf{Z})$. For by the universal coefficient theorem $H^2(M, U(1)) = \text{Hom}(H_2(M, \mathbf{Z}), U(1))$, so any element whatsoever of $H^2(M, U(1))$ can be used to introduce a weighting factor on the homology classes of holomorphic maps. Nontrivial torsion in $H^3(M, \mathbf{Z})$ is equivalent to nontrivial torsion in $H_2(M, \mathbf{Z})$ and can be used to provide additional new weightings.

So in the non-equivariant case gerbes do not produce any real novelty. We just recover the usual family of weighting factors on homology classes of stable maps given us by quantum cohomology. This is not really surprising: the gerbe was supposed to produce local systems on LM , but then we passed to a cover $\tilde{L}M$ which trivialized those local systems. However, in the equivariant case we do get something new, namely the twisted quantum products.

The equivariant case

Much as before, if B is a $U(1)$ -gerbe on X , $\phi : X \rightarrow X$ a symplectomorphism, and an isomorphism $\phi^*B \cong B$ is given, then a $U(1)$ -torsor $L_\phi B$ is naturally induced on the twisted loop space $L_\phi X$. Now it is no longer true that the restriction of $L_\phi B$ to the constant loops $X^\phi \subset L_\phi X$ must be trivial.

On the loop space $L[X/\Gamma] = (L_\Gamma X)/\Gamma$ of a global quotient, then, we get a torsor LB extending the torsor on the inertia stack discussed before. The same thing is presumably true for an orbifold \mathcal{M} that is not a global quotient. For any space of maps from the trinomial to \mathcal{M} , there should be a trivialization of $ev_1^*LB \otimes ev_2^*LB \otimes ev_3^*LB$, and this should allow a twisted Floer product to be defined. At this point it should be clear: we conjecture that *this agrees with the twisted orbifold quantum product*.

An intriguing question: for the Lagrangian-intersection flavor $HF^*(L_1, L_2)$ of Floer cohomology, is there any analogous way to put in a gerbe?

A concluding puzzle

A basic theorem in K-theory asserts that, on a compact manifold X , the Chern character induces an isomorphism

$$K(X) \otimes \mathbf{C} \cong H^*(X, \mathbf{C}).$$

If a finite group Γ acts on X , then there is a similar theorem for the equivariant K-theory:

$$K_\Gamma(X) \otimes \mathbf{C} \cong \bigoplus_{[\gamma]} H^*(X^\gamma, \mathbf{C})^{C(\gamma)},$$

where the sum runs over conjugacy classes. The right-hand side is exactly what we have been calling $H_{\text{orb}}^*(X/\Gamma, \mathbf{C})$. This can also be made a ring isomorphism, provided that the product structure is appropriately defined on both sides. But it seems to be complicated: the usual product on K-theory goes over to the usual product on the cohomology of the inertia stack (not the orbifold cohomology), so to get a ring homomorphism to orbifold cohomology we have to adjust the operation on K-theory, which we might prefer not to do.

As we have discussed, both sides can be generalized by twisting with a Γ -equivariant gerbe B , so we might hope for something like

$$K_\Gamma(X, B) \otimes \mathbf{C} \cong \bigoplus_{[\gamma]} H^*(X^\gamma, L_\gamma B)^{C(\gamma)}.$$

But now the natural multiplicative structures on the two sides are of completely different types. The twisted K-theory on the left-hand side is a module over the untwisted K-theory $K_\Gamma(X)$, while the right-hand side is a ring in its own right. Can these two algebraic structures be related in any reasonable way?

Notes on the literature

Notes to Lecture 1

Although it is a textbook that does not purport to give all technical details, the best source for further reading on Floer homology is: D. McDuff and D.A. Salamon, *J-holomorphic curves and symplectic topology*, AMS, 2004, referenced hereinafter as McDuff-Salamon. This is a greatly expanded version of *J-holomorphic curves and quantum cohomology*, AMS, 1994. The same authors have also written a wider survey of symplectic geometry: *Introduction to symplectic topology*, Oxford, 1998.

The Hamiltonian formalism: See V.I. Arnold, *Mathematical methods of classical mechanics*, Grad. Texts in Math. 60, Springer, 1989, or V. Guillemin and S. Sternberg, *Symplectic techniques in physics*, Cambridge, 1984.

The Arnold conjecture: Floer's original papers are Morse theory for Lagrangian intersections, *J. Differential Geom.* 28 (1988) 513–547; The unregularized gradient flow of the symplectic action, *Comm. Pure Appl. Math.* 41 (1988) 775–813; Witten's complex and infinite-dimensional Morse theory, *J. Differential Geom.* 30 (1989) 207–221; Symplectic fixed points and holomorphic spheres, *Comm. Math. Phys.* 120 (1989) 575–611. The monotone hypothesis, a technical condition on the first Chern class of the tangent bundle, was removed by H. Hofer and D.A. Salamon, Floer homology and Novikov rings, *The Floer memorial volume*, Progr. Math. 133, Birkhäuser, 1995, and by G. Liu and G. Tian, Floer homology and Arnold conjecture, *J. Differential Geom.* 49 (1998) 1–74. For the Lefschetz fixed-point formula, see §11.26 of R. Bott and L.W. Tu, *Differential forms in algebraic topology*, Grad. Texts in Math. 82, Springer, 1982.

Morse theory: The classic reference is J. Milnor, *Morse theory*, Princeton, 1963. The point of view in which the differential counts flow lines did not become popular until the 1980s; for a winsome account from that era, see R. Bott, Morse theory indomitable, *Publ. Math. IHES* 68 (1988) 99–114.

Bott-Morse theory: The spectral sequence was introduced by Bott in An application of the Morse theory to the topology of Lie-groups, *Bull. Math. Soc. France* 84 (1956) 251–281. See the author's A perfect Morse function on the moduli space of flat connections, *Topology* 39 (2000) 773–787 for a concise account. A thorough discussion of Bott-Morse theory is in D.M. Austin and P.J. Braam, Morse-Bott theory and equivariant cohomology, *The Floer memorial volume*, Progr. Math. 133, Birkhäuser, 1995.

Morse theory on the loop space: See Floer's original papers. The Morse index in the Floer theory is called the *Conley-Zehnder* or *Maslov index*: see McDuff-Salamon,

§12.1.

Re-interpretations: An inspiring exposition on the various forms of Floer homology is by M.F. Atiyah, New invariants of 3- and 4-dimensional manifolds, *The mathematical heritage of Hermann Weyl (Durham, NC, 1987)*, Proc. Sympos. Pure Math. 48, AMS, 1988. Another is by J.-C. Sikorav, Homologie associée à une fonctionnelle (d'après A. Floer), *Astérisque* 201-203 (1991) 115–141. For the periodic Floer homology of Hutchings, see M. Hutchings, An index inequality for embedded pseudoholomorphic curves in symplectizations, *J. Eur. Math. Soc.* 4 (2002) 313–361, or M. Hutchings and M. Sullivan, The periodic Floer homology of a Dehn twist, *Algebr. Geom. Topol.* 5 (2005) 301–354.

Product structures: Proofs that the Floer product on $HF^*(id)$ coincides with the quantum product are given by S. Piunikhin, D. Salamon, and M. Schwarz, Symplectic Floer-Donaldson theory and quantum cohomology, *Contact and symplectic geometry (Cambridge, 1994)*, Cambridge, 1996, and by G. Liu and G. Tian, On the equivalence of multiplicative structures in Floer homology and quantum homology, *Acta Math. Sin. (Engl. Ser.)* 15 (1999) 53–80.

There are no details in the literature of the product structures for arbitrary symplectomorphisms. But there is a sketch in McDuff-Salamon, §12.6. And the case where M is a Riemann surface has been the subject of several papers, e.g. R. Gaiotto, Floer homology of algebraically finite mapping classes, *J. Symp. Geom.* 1 (2003) 715–765, and P. Seidel, The symplectic Floer homology of a Dehn twist, *Math. Res. Lett.* 3 (1996) 829–834. For the Novikov ring, see McDuff-Salamon §11.1. For the Fukaya category, see many of Fukaya's papers such as K. Fukaya, Floer homology and mirror symmetry I, *Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds*, AMS/IP Stud. Adv. Math. 23, AMS, 2001, or K. Fukaya and P. Seidel, Floer homology, A_∞ -categories and topological field theory, *Geometry and physics (Aarhus, 1995)*, Dekker, 1997.

The vanishing of the Massey products on a Kähler manifold is proved in P. Deligne, P. Griffiths, J. Morgan, and D. Sullivan, Real homotopy theory of Kähler manifolds, *Invent. Math.* 29 (1975) 245–274.

The finite-order case: On moment maps and perfect Bott-Morse functions, see F.C. Kirwan, *Cohomology of quotients in symplectic and algebraic geometry*, Princeton, 1984. On the finite-order case, a clearly relevant paper is that of A.B. Givental, Periodic mappings in symplectic topology, *Funct. Anal. Appl.* 23 (1989) 287–300.

Givental's philosophy is most fully laid out in Homological geometry and mirror symmetry, *Proceedings of the International Congress of Mathematicians (Zürich, 1994)*,

vol. 1, Birkhäuser, 1995. But see also his Equivariant Gromov-Witten invariants, *Internat. Math. Res. Notices* 1996 (1996) 613–663, as well as A.B. Givental and B. Kim, Quantum cohomology of flag manifolds and Toda lattices, *Comm. Math. Phys.* 168 (1995), 609–641.

For the “usual package of ideas in equivariant cohomology,” see the elegant exposition of M.F. Atiyah and R. Bott, The moment map and equivariant cohomology, *Topology* 23 (1984) 1–28.

Notes to Lecture 2

Much of the basic information on stacks is lifted from B. Fantechi, Stacks for everybody, *European Congress of Mathematics (Barcelona, 2000)*, vol. 1, Progr. Math. 201, Birkhäuser, 2001, and from W. Fulton, What is a stack? Lecture notes available from www.msri.org/publications/ln/msri/2002/introstacks/fulton/1/index.html. Some other readable sources are D. Edidin, B. Hassett, A. Kresch, and A. Vistoli, Brauer groups and quotient stacks, *Amer. J. Math.* 123 (2001) 761–777 and A. Vistoli’s appendix to Intersection theory on algebraic stacks and on their moduli spaces, *Invent. Math.* 97 (1989) 613–670. Much more formidable and comprehensive is the book of G. Laumon and L. Moret-Bailly, *Champs algébriques*, *Ergebnisse Math.* 39, Springer, 2000.

Examples of stacks: The stack of n th roots is discussed by C. Cadman, Using stacks to impose tangency conditions on curves, *Amer. J. Math.*, to appear.

Morphisms and 2-morphisms: A good basic reference on the relevant category theory is Appendix A of C. Weibel, *An introduction to homological algebra*, Cambridge, 1994. For bitorsors, the tetrahedron condition, and so on, see the book of L. Breen, On the classification of 2-gerbes and 2-stacks, *Astérisque* 225 (1994). Group actions on stacks are meticulously treated by M. Romagny, Group actions on stacks and applications, *Michigan Math. J.* 53 (2005) 209–236.

Definition of gerbes and the following 4 sections: The earliest and most comprehensive treatment of gerbes is in the book of J. Giraud, *Cohomologie non abélienne*, *Grund. Math. Wiss.* 179, Springer, 1971. Abelian gerbes are readably discussed by J.-L. Brylinski, *Loop spaces, characteristic classes and geometric quantization*, Progr. Math. 107, Birkhäuser, 1993. See also the book of Breen and the paper of Edidin et al. cited above.

Definition of orbifolds: A good general discussion, delivered with the author’s usual quirky charm, appears in §13 of the samizdat lecture notes of W. Thurston; for some reason this was not included in the version that appeared in book form, but it is available from

www.msri.org/communications/books/gt3m. Another approach to orbifolds, more closely related to stacks, is that via groupoids, due to Moerdijk and collaborators; see for example I. Moerdijk, Orbifolds as groupoids: an introduction, *Orbifolds in mathematics and physics (Madison, WI, 2001)* Contemp. Math. 310, AMS, 2002.

Twisted vector bundles: See, for example, E. Lupercio and B. Uribe, Gerbes over orbifolds and twisted K -theory, *Comm. Math. Phys.* 245 (2004) 449–489, or A. Adem and Y. Ruan, Twisted orbifold K -theory, *Comm. Math. Phys.* 237 (2003) 533–556.

Strominger-Yau-Zaslow: The original article by A. Strominger, E. Zaslow, and S.T. Yau, Mirror symmetry is T -duality, *Nuclear Phys. B* 479 (1996) 243–259, has spawned a vast literature; we mention only the addition of gerbes (a.k.a. “B-fields”) by N.J. Hitchin, Lectures on special Lagrangian submanifolds, *Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds*, AMS/IP Stud. Adv. Math. 23, AMS, 2001, and an appealing survey by R. Donagi and T. Pantev, Torus fibrations, gerbes, and duality, preprint. The author’s papers giving examples where SYZ is satisfied are M. Thaddeus, Mirror symmetry, Langlands duality, and commuting elements of Lie groups, *Internat. Math. Res. Notices* 2001 (2001) 1169–1193, and T. Hausel and M. Thaddeus, Mirror symmetry, Langlands duality, and the Hitchin system, *Invent. Math.* 153 (2003) 197–229.

Notes to Lecture 3

A good general reference on quantum cohomology and Gromov-Witten invariants (without orbifolds) is Part 4 of K. Hori et al., *Mirror symmetry*, AMS, 2003. This volume comprises the proceedings of a school run by the Clay Mathematics Institute.

Cohomology of sheaves on stacks: A convenient reference for Grothendieck’s theorem is I.G. Macdonald, Symmetric products of an algebraic curve, *Topology* 1 (1962) 319–343.

Orbifold cohomology: The orbifold product (where quantum parameters are set to zero) was introduced by W. Chen and Y. Ruan, A new cohomology theory of orbifold, *Comm. Math. Phys.* 248 (2004) 1–31. But the quantum product, though constructed later, appears to be more fundamental: for this see D. Abramovich, T. Graber, and A. Vistoli, Gromov-Witten theory of Deligne-Mumford stacks, preprint. See also Abramovich’s notes in this volume.

Twisted orbifold cohomology: Among the many interesting recent works on the subject, we mention only two by Y. Ruan: Discrete torsion and twisted orbifold cohomology, *J. Symplectic Geom.* 2 (2003) 1–24, and Stringy orbifolds, *Orbifolds in mathematics and physics (Madison, WI, 2001)*, Contemp. Math. 310, AMS, 2002.

The case of discrete torsion: The seminal physics paper is by C. Vafa and E. Witten, On orbifolds with discrete torsion, *J. Geom. Phys.* 15 (1995), 189–214. In fact a whole book had been written by a mathematician, G. Karpilovsky, *The Schur multiplier*, Oxford, 1987.

The Fantechi-Göttsche ring was introduced by B. Fantechi and L. Göttsche, Orbifold cohomology for global quotients, *Duke Math. J.* 117 (2003) 197–227. Since they set the quantum parameters to zero, the Γ -invariant part of their ring carries the orbifold product. Their product has not yet been fully extended to a quantum product, but there is some relevant discussion of the necessary rigidification in T. Jarvis, R. Kaufmann, and T. Kimura, Pointed admissible G -covers and G -equivariant cohomological field theories, *Compos. Math.* 141 (2005) 926–978, and in the 2006 Ph.D. thesis of Maciek Mizerski at the University of British Columbia.

The loop space of an orbifold: Gromov-Witten invariants for orbifolds are defined symplectically by W. Chen and Y. Ruan in Orbifold Gromov-Witten theory, *Orbifolds in mathematics and physics (Madison, WI, 2001)*, Contemp. Math. 310, AMS, 2002, and algebraically by D. Abramovich, T. Graber, and A. Vistoli, Gromov-Witten theory of Deligne-Mumford stacks, preprint.

A concluding puzzle: For the basic theorem in K-theory, see M.F. Atiyah, *K-theory*, Benjamin, 1967. The equivariant version of the theorem is usually attributed to M.F. Atiyah and G.B. Segal, On equivariant Euler characteristics, *J. Geom. Phys.* 6 (1989) 671–677. However, an alternative lineage for this result is traced by A. Adem and Y. Ruan, Twisted orbifold K -theory, *Comm. Math. Phys.* 237 (2003) 533–556. Adem and Ruan also give a ring isomorphism from equivariant K-theory to the cohomology of the inertia stack. The adjusted ring homomorphism going to orbifold cohomology is constructed by T. Jarvis, R. Kaufmann, and T. Kimura, Stringy K-theory and the Chern character, preprint. Another such construction, which extends to twisted K-theory, is given by A. Adem, Y. Ruan, and B. Zhang, A stringy product on twisted orbifold K-theory, preprint.