Quantization And Topological Field Theory

Edward Witten

Bonn

August 6, 2009

(with S. Gukov, arXiv:0809.0305)

According to textbooks, the passage from classical mechanics to quantum mechanics is made by replacing Poisson brackets by commutators.

Ideally, in other words, to quantize of a symplectic manifold M, one is supposed to define a Hilbert space H, and a map from functions on M to operators on H.

The map from functions on M to operators on H is supposed to take Poisson brackets to commutators. In other words, if O_f is the operator that corresponds to a function *f*, then we would like to have

$$[O_f , O_g] = O_{-i\{f,g\}}$$

where *{f,g}* is the Poisson bracket of *f* and *g*.

The only trouble is that this problem does not have a solution, even for a very simple classical phase space such as $M=R^2$ with its standard symplectic structure. If one picks on M an affine structure with linear coordinates x,p

such that $\{p,x\}=1$, and quantizes in the usual way with $p = -i\, d/dx$, so that $[p,x] = -i$,

then one finds the desired relation for quadratic polynomials in x and p, but not for polynomials of higher order (let alone functions that are not polynomial).

Concretely, a polynomial of higher order in x and p maps to a differential operator, say

x^2p^2 --> $-x^2\, d^2/dx^2$ (plus possible lower order terms).

When one commutes two such operators, one must make repeated use of the relation $[d/dx,x]=1$. Almost whenever this relation must be used more than once, the map from Poisson brackets to commutators fails.

The fact that it works for quadratic functions is an important gift of nature.

This gift has the following important consequence: The quadratic functions generate, by Poisson brackets,

the linear symplectomorphisms of R^2 , that is, the ones that preserve its `affine structure.' So the map from *linear* symplectomorphisms to quantum operators works out correctly; hence quantization does not depend on the choice of x and p, but only on the choice of affine structure, i.e. the choice of what we mean by linear coordinates.

We can quantize by letting x act by multiplication and p by – i d/dx, or by letting p act by multiplication and x by i d/dp. The two procedures are equivalent to each other … the equivalence is generated by the Fourier transform from a function f(x) to a function g(p).

In contrast to this, since the passage from nonquadratic functions on R^2 to operators on Hilbert space does not map Poisson brackets to commutators, symplectomorphisms that are not linear are not implemented quantum mechanically in a natural way.

Hence the standard quantization of R^2 does depend on the choice of affine structure.

Quantizations with different choices of affine structure are not equivalent – there is no natural map from one Hilbert space to the other.

For examples other than R^2, the problem is only worse.

Because quantization is ambiguous locally, there is also no general method to carry it out globally.

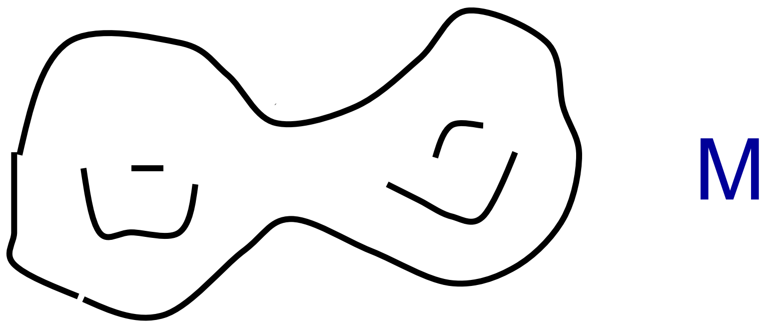

Ideally, we'd cover M with small open sets, quantize each one, and glue the results, but there are two basic reasons that this does not work. (Ambiguous locally, and no way to restrict to a small open set.)

There are three traditional responses to this phenomenon.

The traditional attitude of physicists is to ignore it, relying upon the fact that in a physics problem there is nearly always additional structure, beyond the symplectic structure of phase space, that makes it clear how one should be quantizing.

For example, for a nonrelativistic particle moving in order space R^3, the phase space is the cotangent bundle of R^3. R^3 has a natural affine structure given by the Euclidean metric; this also gives an affine structure to the cotangent bundle. We use this to quantize.

Even in more sophisticated examples, the approach of quantization is usually clear.

Quantum field theory, for example, is hard to understand mathematically, but usually not because of the sort of questions that I've mentioned.

One of the rare examples of a quantum field theory in which the subtleties of quantization actually are important is the topic of this conference – Chern-Simons topological field theory in three dimensions. (These subtleties are also important for representation theory more broadly.)

But before delving into the cases where the problem is important, I want to describe the two other traditional responses to the difficulty of quantization.

One approach is Deformation Quantization (Bayen et al (1978) …. Kontsevich (1997) Cattaneo-Felder (1999))

Here, given a symplectic manifold M, one does not aim to produce a Hilbert space H acted on by a noncommutative deformation of the ring of functions on M. Instead, one only aims to produce a noncommutative deformation of the ring of functions.

The passage from a function f to an element O_f of the ring is supposed to obey

$$O_f \, O_g \;= O_{fg} + hO_{\{f,g\}} + O(h^2)$$

There is a fairly satisfactory theory:

Over a ring of formal power series in a formal parameter *h* , such a deformation exists and is unique, up to automorphism.

Moreover, reasonable conditions are known under which one can take *h* to be a complex number, rather than a formal parameter

However, we get only an algebra and not a Hilbert space that the algebra acts on.

The difference is crucial, as we can see if we consider the case that the phase space is a two-dimensional sphere:

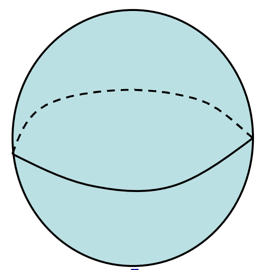

The algebra of functions is generated by x, y, z with the relation $x^2+y^2+z^2=1$ along with commutativity of x,y,z.

The deformed algebra of functions constructed in deformation quantization is obtained by replacing the statement of commutativity by a deformed version

$$xy - yx = hz$$

and cyclic permutations thereof.

For any value of h, this gives an associative algebra – in fact, a famous one, the universal envelopping algebra of SU(2).

But if we stop here, we completely miss one of the most basic facts in physics, which actually is the basis for the name \`quantization’:   Angular momentum is quantized, that is the angular momentum of a particle moving on the two sphere only takes integer or half integer values.

We only learn that angular momentum is quantized if we ask for the deformed algebra to act on a Hilbert space.

This happens at certain values of h where quantization is possible. (We return to this example later.)

In short, deformation quantization is a systematic theory of something, which moreover is a piece of the puzzle – in fact it will be part of the story that I will explain based on the A model -- but it is not quantization.

The main other attempt at a systematic theory is Geometric Quantization

(Bargmann, Kostant, Souriau. . .)

In Geometric Quantization, given a classical phase space (M, ω) the first step is to pick a prequantum line bundle, that is, a unitary line bundle $\mathcal{L} \longrightarrow M$ with a connection whose curvature is ω.

I believe that this is actually likely to be a first step in any approach to quantization, (including that based on the A-model).

Then one picks a polarization, that is, a maximal set of Poisson-commuting functions (roughly speaking, the x's as opposed to the p's).

For a real polarization, one quantizes by letting the x's act by multiplication and p as

d/dx, as textbooks would suggest.

There is an analog for polarizations of other types. For instance, for a complex polarization, one lets $z = x + i\,p$ act as multiplication while $x - ip$ acts as $h\, d/dz$.

Drawback:

The choice of a polarization is much more structure than just a symplectic structure.

If we really have to pick a polarization to quantize, then in passing to quantum mechanics we are losing almost all of the symmetry of an underlying classical symplectic manifold.

In general, geometric quantization is useful if there is a natural polarization, or in some special situations if there is a small family of natural polarizations and a good reason that they give equivalent results.

There is no systematic theory of what that good reason would be, only a few important examples that I will mention.

The most important example is certainly the quantization of R^{2n} with its standard symplectic structure. Generalizing what I said before for n=1, quantization depends only on a choice of affine structure, not on a further choice of polarization – that is, a separation into x's and p's.

So R^{2n} as an affine manifold has a natural quantization.

This has an important corollary: Let F be a group of linear symplectomorphisms of R^{2n} . Then we can define a new symplectic manifold M known as the symplectic quotient $R^{2n}//F$ … M is defined by setting to zero the Hamiltonian functions that generate the action of F and then dividing by F.

Let *H* be the Hilbert space obtained by quantizing R^{2n} . We would like to quantize

M, to get a new Hilbert space H_M.

Since F was a group of linear symplectomorphisms, it acts on H in a natural way, irrespective on what polarization was used to define H. It is reasonable to declare that H_M should be simply the F-invariant subspace of H.

Then F-invariant functions on R^{2n} act naturally on H_M , with the best properties one can hope for.

Here are two interesting cases:

- If R^{2n} admits an F-invariant polarization, then this descends to a polarization of M, and our procedure for quantization M could be obtained in Geometric Quantization by quantizing M with this polarization.

  Going farther, if there is one F-invariant polarization, there may be many of them.

  These will descend to a family of polarizations of M, and there are natural equivalences between the associated quantizations of M …

  this is interesting since there is no general theory of such equivalences.

2) Alternatively, suppose that there is no F-invariant polarization of R^{2n}. Then we have a well-motivated procedure for quantization of M – take the F-invariant part of *H* – but it cannot be interpreted (in any obvious way) in terms of Geometric Quantization.

As I have told you, the Chern-Simons topological field theory in 3 dimensions is one of the few real life examples in which physicists actually grapple with these questions. It is also a good – though infinite dimensional – example of the situation that I have just described.

Basic ingredients in this theory are a finite dimensional Lie group G and a Riemann surface C.

The theory is much easier if G is compact, but we also consider other cases.

We let M denote the moduli space of homomorphisms from the fundamental group of the Riemann surface C into G.

If one is given a nondegenerate quadratic form on g, the Lie algebra of G, then M becomes a symplectic manifold. This is what we would like to quantize. (We impose an integrality condition on the quadratic form to make this possible.)

This problem can be interpreted as an infinite-dimensional version of what I explained earlier for the following reason:

Let E be a G-bundle over C (topologically trivial, for example) and endow E with a connection A. Let Ω be the infinite dimensional group of bundle automorphisms – i.e. roughly the group of maps from C to the finite-dimensional group G. Let Ã be the space of all connections on E. Then Ã is actually an infinite dimensional affine space, which also has a natural symplectic structure.

Ã is an affine space because the difference between any two connections is a one-form, say α, with values in ad(E), and the space of such one-forms is a linear space.

Ã is symplectic because if β and α are two such one-forms, one can define a skew pairing

$$(\alpha, \beta)_{\tilde{A}} = \int_C (\alpha, \beta)$$

and this skew form on the tangent space to Ã actually endows it with a symplectic form.

It turns out that the space M that we want to quantize can be interpreted as the symplectic quotient of the infinite dimensional space Ã of all connections by the group Ω of gauge transformations.

(The Hamiltonian functions that generate the action of Ω are simply the curvature, so the symplectic quotient consists of connections of zero curvature modulo gauge transformations – this is the same as representations of the fundamental group.)

We are therefore in an infinite-dimensional version of a problem we know how to solve in finite dimensions: we want to quantize the symplectic quotient of an affine space. One might hope the finite dimensional analysis applies to this infinite dimensional case.

This has actually only been done successfully for the case of compact G

An important property of M in this example, for compact G, is that it has a natural family of complex polarizations. Indeed, if we pick a complex structure on C, turning it into a complex Riemann surface, then M becomes (by a theorem of Narasimhan and Seshadri) an algebraic variety, the moduli space of stable G-bundles over C. As such it acquires a complex polarization, and can be quantized by Geometric Quantization – the resulting Hilbert space is the space of `nonabelian theta functions.'

The only problem is that, since C admits many complex structures, the resulting quantization of M may not be unique – generically it would not be.

In this case, however, the polarizations on M descend from natural Ω – invariant complex polarizations on the infinite dimensional space Ã . Hence one can imitate what one would do if Ã were finite-dimensional: quantize Ã and restrict to the Ω – invariant subspace, using the linear nature of Ã to show that the choice of polarization does not matter.

Following this direction, one can actually construct a projectively flat connection (over the Teichmuller space of C) enabling one to eliminate the dependence of the quantization of M on the choice of a polarization.

(Axelrod, DellaPietra, and EW 1990)

The existence of this flat connection has also been understood using conformal field theory (Knizhnik-Zamolodchikov, Bernard…) and directly from algebraic geometry of M (Hitchin, Faltings,…)

By contrast, quantization for the case that G is a noncompact semisimple Lie group is not well understood (but for split real forms see work of Chekhov-Fock, Fock-Goncharov, and Chekhov-Penner and on complex groups Dimofte, Gukov, Lemells, and Zagier).

I won’t try to review all the arguments, but instead I will use this example as an illustration of an alternative approach to quantization...

This is the approach to quantization (Gukov and EW, 2008) via the two-dimensional A-model. (The A-model may be most familiar via the Floer-Fukaya category and mirror symmetry. Suggestions that the A-model is related to geometric quantization were first made in Bressler and Soibelman (2002), Kapustin (2005). Also – examples by Aldi & Zaslow (2005).)

In this approach to quantization…

We start with a classical phase space (M, ω) that we wish to quantize.

As in Geometric Quantization (but unlike Deformation Quantization, in which, as a result, angular momentum isn't quantized) we start by picking a prequantum line bundle $\mathcal{L} \longrightarrow M$ of curvature ω .

The next step is to pick a *complexification* Y of M. Y is a complex manifold, with a complex conjugation operation

$$\tau : Y \to Y$$

that has M as a fixed point set. Y is supposed to be a complex symplectic manifold with a holomorphic two-form Ω that is a complexification of ω in the sense that along Y, $\omega = \mathrm{Re}\,\Omega$

Also $\mathcal{L}$ must extend to a line bundle over Y with curvature $\mathrm{Re}\,\Omega$

The construction is most interesting if Y is an affine variety, which means informally that there are lots of holomorphic functions on Y … so many that the restrictions to M of holomorphic functions on Y give a good approximation to the space of all functions on M.

Finally, there is a very important technical condition – Y, when understood as a real symplectic manifold with symplectic form $\mathrm{Im}\,\Omega$, must have a good A model

Without this condition, we really wouldn’t get anywhere, since complexifications of real manifolds are a dime a dozen.

I'll say more about the A-model later, but for Y to have a good A-model in the sense of two-dimensional topological field theory should imply that deformation quantization of Y gives an actual deformation of the ring of functions on Y depending on a complex parameter, not just a formal deformation over a power series ring.)

(Kontsevich has recently given a criterion for deformation quantization of a complex manifold to give an actual deformation, not just a formal one. His criterion is similar to the sort of criterion that could lead to a good A-model.)

The complexification Y plays a role in our approach to quantization that is similar to the role played by a polarization in Geometric Quantization.

There is no free lunch, and just like a polarization, a suitable Y may not exist or be unique.

Just as Geometric Quantization is most useful when there is a natural choice of polarization (or a small family of choices with special properties), the approach to quantization via the A-model is most useful when there is a natural choice of Y or a small family of good choices.

Rather than start right away with details about the A-model, I am going to contrast the workings of Geometric Quantization and A-model based quantization in a standard example.

We will also see where Deformation Quantization comes in.

We pick $M = S^2$ to be a two-sphere, with a symplectic form such that $\int_M \omega = 2\pi n$ for an integer n

We expect to get a Hilbert space of dimension n . (In Deformation Quantization, one need not require n to be integral.)

In Geometric Quantization, one picks a polarization, that is a complex structure J on S^2. This determines a metric on S^2, namely $g = \omega J$, and any metric at all can arise in this way. A given choice breaks almost all of the classical symplectomorphism symmetry, leaving only a finite-dimensional subgroup that preserves the metric.

The largest this subgroup can be is SO(3), and it is convenient to pick the polarization so that this is the case.

If the metric has SO(3) symmetry, then we can embed the two-sphere S^2 in $\mathbf{R}^3$ in the standard fashion

$$x^2 + y^2 + z^2 = 1$$

Thus, there are distinguished coordinates x, y, z, which are unique up to an SO(3) rotation.

Once we picked the complex structure J, the pre-quantum line bundle $\mathcal{L}$ became a holomorphic line bundle, and we define the quantum Hilbert space to be

$$\mathcal{H} = H^0(S^2, \mathcal{L})$$

This procedure of quantization is SO(3)-invariant, so the group SO(3) – or its double cover SU(2) if n is odd – acts naturally on $\mathcal{H}$.

Classically, the functions x,y,z generate SO(3) via Poisson brackets.

Since SO(3) is a symmetry of the structure used for quantization, its generators act naturally on the quantum Hilbert space …

as the SU(2) generators.

It is then fairly natural to declare that symmetric traceless polynomials in x, y, z act quantum mechanically via the corresponding polynomials in the SU(2) generators.

So this is Geometric Quantization of the two-sphere.

(Deformation Quantization of M gives part of the same answer with more symmetry.)

Now let us discuss how we approach the same problem using the A-model.

Here to proceed, we need to pick a suitable complexification Y of the two-sphere.

There cannot be a completely natural choice, because the infinite-dimensional group of symplectomorphisms is not going to act on any complexification.

However, we can make a nice choice if we are given an embedding of the two-sphere in $\mathbf{R}^3$ in the standard fashion

$$x^2 + y^2 + z^2 = 1$$

We simply declare that Y is defined by the same equation, but now with x, y, z understood as complex variables.

The holomorphic symplectic form $\Omega = n dx \wedge dy/z$ has all the desired properties, so we are in business.

Now in quantization based on the A-model, the class of functions on M that we quantize are the ones that extend to holomorphic functions on Y.

In our example, polynomials in x, y, z have this property, and this is the good class of functions that we consider.

Let $\mathcal{R}$ be the ring of these polynomial functions.

The functions x, y, z are distinguished in this approach because they generate $\mathcal{R}$

A related fact is that they are the holomorphic functions on Y that grow most slowly at infinity … Recall that in geometric quantization, the same functions are distinguished because they generate the SO(3) symmetry of the complex structure J.

The A-model now does two things for us

(iii) It constructs a noncommutative deformation of the ring $\mathcal{R}$. One can think of this step as deformation quantization of Y (not M!). We call the deformed ring $\tilde{\mathcal{R}}$

(ii) It constructs a Hilbert space that $\tilde{\mathcal{R}}$ acts on.

One interesting point is that construction of $\tilde{\mathcal{R}}$ in step (1) depends only on Y and not on M, the space we are trying to quantize.

There can be many different M's whose complexification is the same Y, and then the same algebra $\tilde{\mathcal{R}}$ will act in the quantization of each of those M's.

This is important in representation theory of semi-simple Lie groups.

(note work of R. Brylinski 1997, 1998)

Now I will say a little bit about how this comes out of the A-model.

The A-model is a two-dimensional topological field theory that in its simplest version counts holomorphic curves (or pseudo-holomorphic curves a la Gromov) in M. It enters string theory because of its role in determining Yukawa couplings in heterotic string theory, and related observables.

A brane is a boundary condition in the two-dimensional field theory.

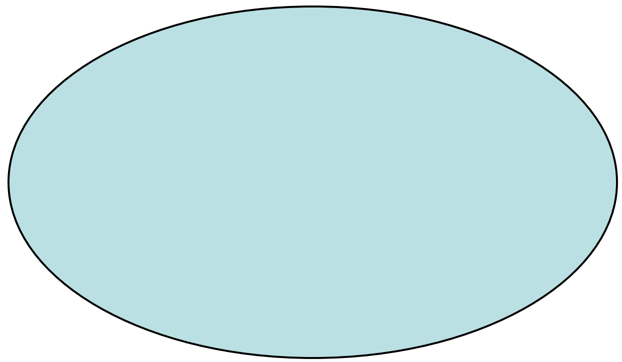

In the A-model, the most familiar branes are the Lagrangian A-branes, related to Floer theory among other things. Their support is middle-dimensional.

However (Kapustin and Orlov; Gualtieri) the A-model also has branes whose support is above the middle dimension.

The conditions for existence of such an A-brane are a little special. The simplest example actually arises for a complex symplectic manifold such as our Y if one endows it with the real symplectic form

$\mathrm{Im}\,\Omega$ so as to have an A-model.

In this case, Y admits a space-filling A-brane whose "Chan-Paton line bundle" is our friend $\mathcal{L}$, the extension to Y of the prequantum line bundle over M.

This brane is a basic example of a coisotropic A-brane and it was important in work of Kapustin and me on geometric Langlands.

We called it $\mathcal{B}_{cc}$, the canonical coisotropic A-brane

In general, if $\mathcal{B}$ is any brane, the space of $(\mathcal{B}, \mathcal{B})$ strings is an algebra:

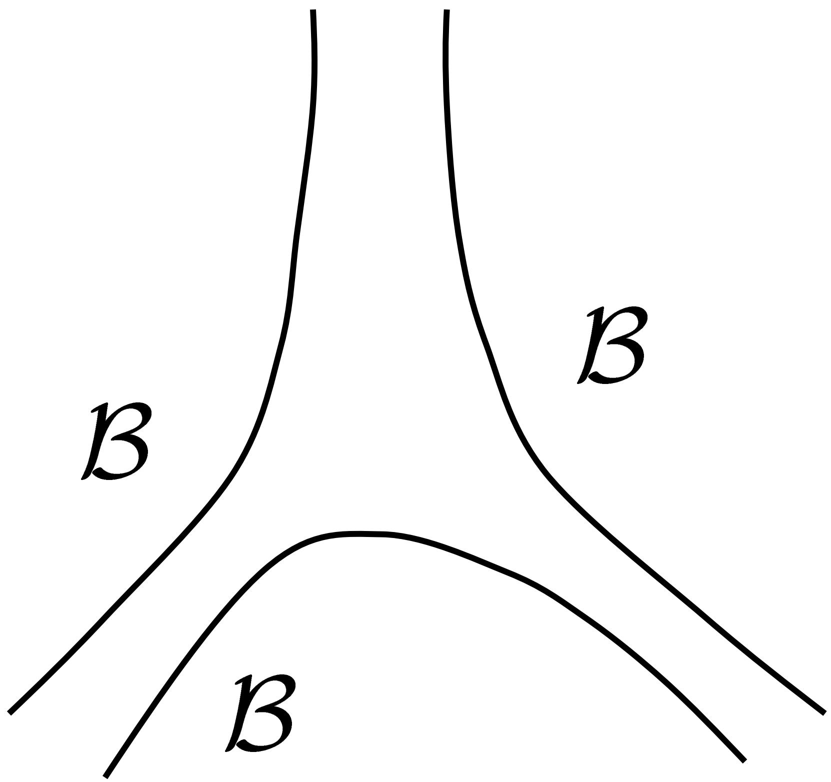

In the case of the canonical coisotropic brane, it turns out that the space of $(\mathcal{B}_{cc}, \mathcal{B}_{cc})$ strings is the algebra $\tilde{\mathcal{R}}$

that I told you about before.

Now we would like to find something that $\tilde{\mathcal{R}}$ can act on.

There is a simple way to do this: if $\mathcal{B}$ is any other A-brane, then $\tilde{\mathcal{R}}$ acts on the space of $(\mathcal{B}_{cc}, \mathcal{B})$ strings.

This is also explained by a simple picture:

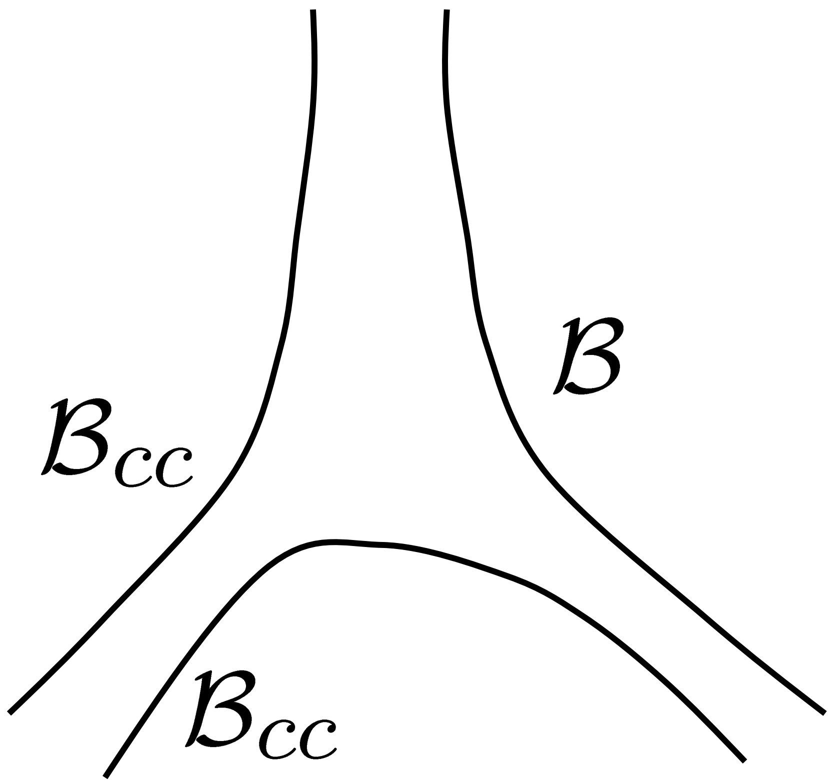

In the case at hand, we take $\mathcal{B}$ to be the ordinary Lagrangian A-brane supported on M … note that M, though (by definition) symplectic for $\mathrm{Re}\,\Omega$, is Lagrangian for $\mathrm{Im}\,\Omega$.

The desired Hilbert space for quantization of M is then simply the space of $(\mathcal{B}_{cc}, \mathcal{B})$ strings. (This is shown by a short formal calculation.)

I think I will simply conclude by revisiting some of our examples in this context.

Consider the quantization of $\mathbf{R}^{2n}$ with its standard symplectic form. Quantization definitely *does* depend on a choice of what one means by linear functions, i.e. a choice of coordinates $q^i,\ i = 1, \ldots, 2n$ whose Poisson brackets are constants.

However, quantization does *not* depend on a separation of the q's into positions and momenta … one makes such a choice in quantization, but the resulting Hilbert space does not depend on the choice.

This is usually proved by showing that the symplectic group $Sp(2n,\mathrm{R})$ (or rather its double cover) acts on the quantum Hilbert space.

From the point of view of the A-model, we simply observe that, once the linear functions q^i are picked, we can complexify $\mathbf{R}^{2n}$ by simply regarding the q^i as complex variables.

The complexification $Y = \mathbf{C}^{2n}$ certainly has a good A-model – the associated sigma model is a free field theory.

We didn't have to separate the q^i into coordinates and momenta in order to define the complexification Y of $\mathbf{R}^{2n}$.
So we conclude that quantization of $\mathbf{R}^{2n}$ does not depend on such a separation.

Finally, let us reconsider the example of quantizing M, the moduli space of representations of the fundamental group of a surface C in a Lie group G. First suppose that G is compact.

M has a natural complexification Y, the moduli space of representations in the complex Lie group G_C obtained by complexifying G.

Y is a complex symplectic manifold with all of the properties needed for the A-model.

Moreover, Y has a good A-model because it has a complete hyper-Kahler metric described by Hitchin. Indeed, for any choice of complex structure on C, Y acquires a complete hyper-Kahler metric.

Just as in the approach by geometric quantization, a complex structure on C must be picked to endow Y with a hyper-Kahler metric and enable us to define its A-model.

But now, instead of using gauge theory arguments and the linearity of the space of connections to argue that quantization does not depend on the complex structure of C, we argue the same thing using the topological invariance of the A-model … the fact that the A-model does not really depend on the choice of metric on Y.

What do we gain by looking at things in this way? One thing that I think we may gain is a better framework for thinking about the quantization problem for the noncompact real forms. Let G* be another real form of G. Let M* be the symplectic manifold of representations of the fundamental group of C into G*. Then just like M, M* is naturally embedded in Y, the space of representations in the complex Lie group.

This gives another brane in the same A-model.

Just as M is the support of a Lagrangian A-brane $\mathcal{B}$, by the same token M* is the support of a Lagrangian A-brane $\mathcal{B}^*$. By considering the space of ($\mathcal{B}_{cc}$, $\mathcal{B}^*$) strings rather than ($\mathcal{B}_{cc}$, $\mathcal{B}$) strings, we can formally use the A-model to quantize M*.

It must be independent of the choice of complex structure on C for the same A-model reason that this is so for M.

The same ring $\tilde{\mathcal{R}}$ of $(\mathcal{B}_{cc}, \mathcal{B}_{cc})$ strings that acts on quantization of M will also act on quantization of M*, since they are branes in the same A-model.

This ring is a deformation of the commutative ring of functions on the complex variety Y … which is generated by traces of holonomies. Its noncommutative deformation depends on a complex variable usually called q. It enters here as an A-model parameter. (Specific restrictions on q are necessary for the A-branes appropriate to different real forms to exist or to get an involution leading to a unitary structure. For instance, for compact G, q should be a root of unity.)

So the same deformed ring of holonomy functions acts in quantization of the moduli spaces for the different real forms.

(There is a slight shift in the parameters that was actually computed in another way in Bar-Natan and EW, 1990. This deformed ring has been studied various ways in algebraic geometry, and studied in differential geometry by Andersen,Gamnelgard,Lauritsen.)

I have to tell you that at the moment I do not know how to use this A-model reasoning to compute anything about the quantization of M*. But there is something that makes me think it is a good viewpoint.

It is the only quantum field theory approach I know in which it is natural to compare the different real forms. That is an aspect of reality that is taken for granted mathematically, but usually hard to incorporate in quantum field theory.